The Instance-Optimal Rate for Bounded Summation under Approximate Differential Privacy

September 1, 2026

Abstract

In the central replacement model, let $D = (x_1, \ldots, x_n) \in \{0, \ldots, U\}^n$, $S(D) = \sum_i x_i$, and $m(D) = \max\{1, x_1, \ldots, x_n\}$. Let $\rho^*_{\varepsilon,\delta,\beta}(n, U)$ be the infimum, over real-valued (ε, δ)-differentially private mechanisms, of the worst-case ratio $\varepsilon \operatorname{err}_\beta(M, D)/m(D)$, where err_β is pointwise $(1-\beta)$-quantile absolute error. There are universal constants $C, C_0 > 0$ such that, for integers $U \geq 2$ and $n \geq 1$, $0 < \varepsilon \leq 1$, $0 < \beta \leq 1/3$, $0 \leq \delta \leq n^{-2}$, and

$$n \geq \frac{C_0}{\varepsilon}\left(1 + \log\log U + \log\frac{1}{\beta}\right),$$

we prove

$$\frac{1}{200}\min\{A_{U,\beta}, B_{\varepsilon,\delta}\} \leq \rho^*_{\varepsilon,\delta,\beta}(n, U) \leq C\min\{A_{U,\beta}, B_{\varepsilon,\delta}\},$$

where $L_U = 1 + \lceil \log_2 U \rceil$, $A_{U,\beta} = 1 + \log(L_U/\beta)$, $B_{\varepsilon,\delta} = 1 + \log(1 + (e^\varepsilon - 1)/\delta)$ for $\delta > 0$, and $B_{\varepsilon,0} = +\infty$. The upper bound is attained by a central mechanism outputting a real estimate of $S(D)$. In the ideal exact-real and exact-sampling model, its pure-private branch runs in $O(n \log n + L_U)$ time; its approximate-private branch has deterministic error $O(m(D)B_{\varepsilon,\delta}/\varepsilon)$ and runs in time polynomial in n, $\log U$, $\log(1/\delta)$, and $\log(1/\varepsilon)$. The lower bound has no computational restriction.

Note. This paper was fully AI generated through a harness created by Richard Peng that discovers open problems and runs them through an automated research pipeline.

Contents

1 Introduction

Summation is among the most basic tasks in private data analysis. Consider n nonnegative contributions $x_1, \dots, x_n \in \{0, \dots, U\}$ and the query $S(D) = \sum_i x_i$. Under replacement adjacency, its global sensitivity is U, so the classical Laplace mechanism adds noise at scale U/ε [DMNS06]. This worst-case calibration treats a database whose largest entry is 1 in the same way as one whose largest entry is U. When U is only a conservative public bound, an error guarantee governed by the largest value actually present in the database is substantially more informative.

Clipping makes the issue concrete. For a threshold t, the clipped query

$$Q_t(D) = \sum_{i=1}^{n} \min\{x_i, t\}$$

has replacement sensitivity t, but it incurs the deterministic bias

$$S(D) - Q_t(D) = \sum_{i=1}^{n} (x_i - t)_+.$$

Thus a low threshold reduces privacy noise at the cost of discarding more of the sum, while a high threshold has the opposite effect. This bias–noise tradeoff also appears in contribution-bounding work, where the best cap can be characterized with the data held fixed [AKMV19]. In the present problem, however, the threshold must be chosen from the private data and one mechanism must give a simultaneous guarantee over every database.

We determine this uniform adaptation cost, up to universal constants, under central approximate differential privacy. The formal criterion is as follows.

Let $U \geq 2$ and $n \geq 1$ be integers, and let databases $D = (x_1, \dots, x_n)$ belong to $\{0, \dots, U\}^n$, with two databases adjacent when they differ in one coordinate. For a real-valued central mechanism M, define

$$S(D) = \sum_{i=1}^{n} x_i, \qquad m(D) = \max\{1, x_1, \dots, x_n\},$$

$$\mathrm{err}_\beta(M, D) = \inf\left\{a \geq 0 : \Pr\left[|M(D) - S(D)| \leq a\right] \geq 1 - \beta\right\},$$

and

$$\rho^*_{\varepsilon,\delta,\beta}(n, U) = \inf_M \sup_{D \in \{0,\dots,U\}^n} \frac{\varepsilon\, \mathrm{err}_\beta(M, D)}{m(D)},$$

where the infimum ranges over all (ε, δ)-differentially private mechanisms. The floor in $m(D)$ keeps the normalization nondegenerate on the all-zero database. Unlike a pointwise oracle benchmark, this definition requires the same private mechanism to attain its guarantee on all instances.

1.1 Our result

Put

$$L_U = 1 + \lceil \log_2 U \rceil, \qquad A_{U,\beta} = 1 + \log \frac{L_U}{\beta},$$

and, for $\delta > 0$,

$$B_{\varepsilon,\delta} = 1 + \log\left(1 + \frac{e^\varepsilon - 1}{\delta}\right), \qquad B_{\varepsilon,0} = +\infty.$$

Theorem 1.1 (Instance-optimal rate for bounded summation)**.** *There are universal constants $C, C_0 > 0$ such that the following holds for all integers $U \geq 2$ and $n \geq 1$. If*

$$0 < \varepsilon \leq 1, \qquad 0 < \beta \leq \frac{1}{3}, \qquad 0 \leq \delta \leq n^{-2},$$

and

$$n \geq \frac{C_0}{\varepsilon}\left(1 + \log\log U + \log\frac{1}{\beta}\right), \tag{1}$$

then

$$\frac{1}{200}\min\{A_{U,\beta}, B_{\varepsilon,\delta}\} \leq \rho^*_{\varepsilon,\delta,\beta}(n, U) \leq C\min\{A_{U,\beta}, B_{\varepsilon,\delta}\}. \tag{2}$$

The upper bound is attained by BranchClipSum*, which outputs a real estimate $\widehat{S}$ satisfying, for every D,*

$$\Pr\left[|\widehat{S} - S(D)| \leq \frac{Cm(D)}{\varepsilon}\min\{A_{U,\beta}, B_{\varepsilon,\delta}\}\right] \geq 1 - \beta.$$

Its pure-private branch uses $O(n \log n + L_U)$ exact-real operations. Its approximate-private branch has the displayed error on every execution and runs in time polynomial in n, $\log U$, $\log(1/\delta)$, and $\log(1/\varepsilon)$. The lower bound in Equation (2) applies to all mechanisms, without a computational restriction.

The two terms in Equation (2) describe different privacy regimes. For $0 < \varepsilon \leq 1$,

$$B_{\varepsilon,\delta} = \Theta\left(1 + \log\left(1 + \frac{\varepsilon}{\delta}\right)\right) \qquad (\delta > 0).$$

At $\delta = 0$, the second term is absent and the theorem gives the pure-private rate $\Theta(A_{U,\beta})$. For $\delta > 0$, the approximate branch is smaller precisely when

$$\delta > \frac{(e^\varepsilon - 1)\beta}{L_U - \beta}. \tag{3}$$

Crossing this threshold selects the better branch; obtaining a strict asymptotic-order improvement additionally requires $B_{\varepsilon,\delta} = o(A_{U,\beta})$.

For example, fix ε and β, take $n = \Theta(\log\log U)$ with a sufficiently large implicit constant to meet Equation (1), and set $\delta = n^{-2}$. Then

$$A_{U,\beta} = \Theta(\log\log U), \qquad B_{\varepsilon,\delta} = \Theta(\log\log\log U),$$

and hence

$$\rho^*_{\varepsilon,\delta,\beta}(n, U) = \Theta(\log\log\log U).$$

Thus approximate privacy gives a strict improvement over the pure-private rate in this regime, while the matching lower bound shows that the normalized ratio still diverges with U.

1.2 Relation to prior work

For fixed public n, estimating a one-dimensional empirical mean and estimating its sum differ only by multiplication by n. Huang, Liang, and Yi [HLY21] study this empirical statistic with radius- and diameter-adaptive guarantees under ρ-zCDP. Their privacy notion and their diameter benchmark differ from the standard (ε, δ) replacement-DP and origin-anchored maximum used here. For pure DP, Dong and Yi [DY23] give the direct maximum/radius-based clipping lineage

over integer domains and a finite-domain hard family yielding the double-logarithmic obstruction at constant failure probability. The $\delta = 0$ specialization of Theorem 1.1 has this double-logarithmic dependence and also records the dependence on β in our stated parameter regime.

The same bounded nonnegative sum and realized-maximum accuracy objective have also been studied under stronger or otherwise altered privacy models. Dong, Luo, Fanti, Shi, and Yi [DLF+24] obtain the $A_{U,\beta}$ accuracy scale with a one-round shuffle protocol; their bound does not use δ to improve that scale. Dong, Sun, and Yi [DSY23] give approximate-private max-contribution upper bounds for vector sums and relational aggregation; the one-dimensional case is a summation guarantee, but it uses add/remove user semantics and retains additional terms rather than giving the scalar replacement-DP transition in Equation (2). A general approximate-private wrapper for finite-range monotone functions due to Linder, Raskhodnikova, Smith, and Steinke [LRSS25] also implies a max-scaled sum bound under its add/remove and down-neighborhood formulation; this is an inferred specialization with a range overhead, rather than the present max-normalized minimax statement.

Several nearby lines optimize a different object. Contribution-bound selection for histograms competes with the best cap in a prescribed family under user-level privacy [LSZ+23], while subset-based private mean estimation uses an expected-loss benchmark defined through deletions from the observed dataset [DKSS23]. Fixed-sensitivity noise mechanisms, including truncated Laplace noise, assume that the sensitivity is already available as a public parameter [GDGK20]. These works respectively illuminate clipping, alternative notions of instance optimality, and the eventual noise release, but they do not privately identify the present instance scale and match it with a lower bound for the same central replacement-DP criterion.

1.3 Proof ideas

The upper bound begins by mapping each record to a dyadic scale and retaining the largest K scales. This top-K operation is stable under one replacement. Any interior point of the retained scales determines a threshold t satisfying $t < 2m(D)$ and clips fewer than K records; therefore the clipping bias is at most $Km(D)$. In the pure-private branch, the exponential mechanism [MT07] selects an interior scale with high probability and a Laplace release completes the estimate. In the approximate-private branch, an always-interior selector uses one level of the prefix reduction from the private interior-point literature [BNSV15], after which bounded truncated Laplace noise is added. The latter release is in the fixed-sensitivity tradition of Geng, Ding, Guo, and Kumar [GDGK20]. Publicly choosing the branch with the smaller of $A_{U,\beta}$ and $B_{\varepsilon,\delta}$ gives Branch-ClipSum.

For the lower bound, we compare the all-zero database with databases containing the same number of copies of geometrically separated values. Accuracy forces their output intervals to be disjoint. Exact approximate group privacy, with its additive δ loss retained, then yields both the $\log(L_U/\beta)$ obstruction and the $\log(1 + (e^{\varepsilon} - 1)/\delta)$ obstruction from a single packing. A separate three-point packing supplies the constant floor needed when these logarithms are small. Together the arguments give the uniform constant in the left side of Equation (2).

Organization. Section 2 fixes the privacy conventions and develops the ordered-multiset, selection, and prefix tools used in the proof. Section 3 gives a detailed overview of the two upper-bound branches and the lower-bound packing. The full construction and analysis appear in Section 4.

2 Preliminaries

Conventions and ordered multisets. All unqualified logarithms are natural, and $[n] = \{1, \dots, n\}$. Multisets retain multiplicities, and $\uplus$ denotes multiset union. Two same-size multisets differ by one replacement if they can be written as $C \uplus \{x\}$ and $C \uplus \{x'\}$. All mechanisms below that take multiset inputs are permutation invariant. Resource statements use ideal exact-real arithmetic and exact sampling from the displayed laws; the minimax statement is information theoretic.

In calls between displayed algorithms, public parameters fixed by the enclosing specification are ambient and may be suppressed. Every parameter whose value differs from its ambient value is displayed explicitly, while each algorithm's input line records all data and public parameters it uses.

For a finite multiset Y in a totally ordered set and $0 \le K \le |Y|$, let $\mathrm{Top}_K(Y)$ and $\mathrm{Bot}_K(Y)$ be its K largest and K smallest members, respectively, with multiplicity. An *interior point* of a nonempty multiset Y is an element of its ordered universe lying in $[\min Y, \max Y]$. For a point z in the same ordered universe, define its rank quality by

$$q_Y(z) = \min\big\{\#\{y \in Y : y \le z\}, \#\{y \in Y : y \ge z\}\big\}. \tag{4}$$

One replacement changes every value of q_Y by at most one. A median of Y has quality at least $|Y|/2$, whereas a point strictly outside $[\min Y, \max Y]$ has quality zero. For every integer $r \ge 1$,

$$q_Y(z) \ge r \quad \Longleftrightarrow \quad \#\{y \in Y : y \le z\} \ge r \ \text{ and } \ \#\{y \in Y : y \ge z\} \ge r. \tag{5}$$

We also write $c_z(Y) = \#\{y \in Y : y \ge z\}$. For $r \ge 1$, a point z is an *r-robust interior point* of Y if $\min Y \le z$ and $c_z(Y) \ge r$; these conditions imply that z is interior.

Lemma 2.1 (Stability of order truncation)**.** *If same-size multisets Y, Y' differ by one replacement, then, for every $0 \le K \le |Y|$, the multisets $\mathrm{Top}_K(Y), \mathrm{Top}_K(Y')$ differ by at most one replacement, as do $\mathrm{Bot}_K(Y), \mathrm{Bot}_K(Y')$. Consequently, applying a permutation-invariant private mechanism after either truncation preserves its privacy parameters with respect to the original input.*

Proof. Write $Y = C \uplus \{y\}$ and $Y' = C \uplus \{y'\}$. For $1 \le K < |Y|$, both top-K multisets contain $\mathrm{Top}_{K-1}(C)$ and one further member; the same statement holds for the bottom-K multisets. The cases $K = 0$ and $K = |Y|$ are immediate. □

Privacy tools. For $D, D' \in \{0, \dots, U\}^n$, write $D \sim D'$ when they differ in one coordinate; for multiset inputs, $D \sim D'$ denotes the one-replacement relation above. With either declared adjacency relation, a mechanism M is (ε, δ)-differentially private if, for every adjacent pair and every measurable set T,

$$\Pr[M(D) \in T] \le e^{\varepsilon} \Pr[M(D') \in T] + \delta. \tag{6}$$

The replacement ℓ_1-sensitivity of f taking values in $\mathbb{R}^k$ is $\sup_{D \sim D'} \|f(D) - f(D')\|_1$; for a scalar query this is its *replacement sensitivity*. We use postprocessing and adaptive basic composition: if, conditional on every earlier transcript, invocation j is $(\varepsilon_j, \delta_j)$-private, then the joint transcript is $(\sum_j \varepsilon_j, \sum_j \delta_j)$-private. We also use the exact group-privacy consequence of Equation (6): if D, D' are joined by a chain of k replacements, then

$$\Pr[M(D) \in T] \le e^{k\varepsilon} \Pr[M(D') \in T] + \delta \sum_{j=0}^{k-1} e^{j\varepsilon}. \tag{7}$$

For $s > 0$, $\mathrm{Lap}(s)$ denotes the law with density $(2s)^{-1}e^{-|w|/s}$. If $W \sim \mathrm{Lap}(s)$ and $u \geq 0$, then

$$\Pr[|W| > u] = e^{-u/s}, \qquad \Pr[W > u] = \Pr[W < -u] = \tfrac{1}{2}e^{-u/s}. \tag{8}$$

Thus adding independent $\mathrm{Lap}(\Delta/\alpha)$ noise to the coordinates of a query with ℓ_1-sensitivity at most Δ is $(\alpha, 0)$-private.

Suppose a quality function $q_D : \mathcal{R} \to \mathbb{R}$ on a finite range satisfies $|q_D(z) - q_{D'}(z)| \leq \Delta$ for every $D \sim D'$ and $z \in \mathcal{R}$, where $\Delta > 0$. The exponential mechanism with parameter α samples Z according to

$$\Pr[Z = z] \;\propto\; \exp\left(\frac{\alpha q_D(z)}{2\Delta}\right).$$

It is $(\alpha, 0)$-private. Moreover, for any $z_0 \in \mathcal{R}$ and $u \geq 0$,

$$\Pr[q_D(Z) \leq q_D(z_0) - u] \leq |\mathcal{R}| \exp\left(-\frac{\alpha u}{2\Delta}\right). \tag{9}$$

For probability laws P, Q, their total-variation distance is $d_{\mathrm{TV}}(P, Q) = \sup_T |P(T) - Q(T)|$; any coupling of random variables $X \sim P$ and $Y \sim Q$ gives $d_{\mathrm{TV}}(P, Q) \leq \Pr[X \neq Y]$. If a family P_D is (α, ζ)-private and $d_{\mathrm{TV}}(P_D, Q_D) \leq \tau$ for every D, then, for adjacent D, D' and every event T,

$$Q_D(T) \leq e^{\alpha} Q_{D'}(T) + \zeta + (1 + e^{\alpha})\tau. \tag{10}$$

For $\alpha, \Delta > 0$ and $0 < \zeta \leq 1/2$, define

$$R(\alpha, \zeta, \Delta) = \frac{\Delta}{\alpha} \log\left(1 + \frac{e^{\alpha} - 1}{2\zeta}\right).$$

The corresponding truncated-Laplace law has density proportional to $e^{-\alpha|w|/\Delta}$ on $[-R(\alpha, \zeta, \Delta), R(\alpha, \zeta, \Delta)]$ and density zero outside this interval.

Dyadic scales and binary prefixes. For $x \in \{0, \ldots, U\}$, define the dyadic scale and, for $t \geq 0$, the clipped sum by

$$b(x) = \lceil \log_2 \max\{1, x\} \rceil \in \{0, \ldots, L_U - 1\}, \tag{11}$$

$$Q_t(D) = \sum_{i=1}^{n} \min\{x_i, t\}. \tag{12}$$

Lemma 2.2 (Sensitivity of clipped summation)**.** *For every fixed $t \geq 0$, the replacement sensitivity of Q_t is at most t.*

Proof. A replacement changes one summand, and every summand lies in $[0, t]$. □

For the prefix construction, fix $d \geq 2$, put $\mathcal{X} = \{0, \ldots, d - 1\}$ and $\ell = \lceil \log_2 d \rceil$, and let $\mathrm{bin}_\ell(x) \in \{0, 1\}^\ell$ be the binary representation of $x \in \mathcal{X}$ padded with leading zeroes. If $x, x' \in \mathcal{X}$, their longest-common-prefix length is

$$\lambda(x, x') = \max\left\{s \in \{0, \ldots, \ell\} : \mathrm{bin}_\ell(x)_{1:s} = \mathrm{bin}_\ell(x')_{1:s}\right\}.$$

For $0 \leq s \leq \ell$ and $p \in \{0, 1\}^s$, let

$$C(p) = \{x \in \mathcal{X} : \mathrm{bin}_\ell(x)_{1:s} = p\}.$$

When $C(p)$ is nonempty, write $a(p) = \min C(p)$ and $b(p) = \max C(p)$; then $C(p) = [a(p), b(p)] \cap \mathcal{X}$, also when d is not a power of two. For a multiset Y over $\mathcal{X}$ and a fixed prefix length s, write

$$c_Y(p) = |Y \cap C(p)|, \qquad G_s(Y) = \{p \in \{0, 1\}^s : c_Y(p) > 0\}$$

for its prefix counts and its set of observed prefixes.

3 Overview

Put $T = \min\{A_{U,\beta}, B_{\varepsilon,\delta}\}$. The upper construction and two complementary packing arguments give the matching bounds. The main algorithm BranchClipSum receives the database and the public parameters, chooses publicly between a pure-private and an approximate-private clipping mechanism, and is (ε, δ)-private with pointwise $(1-\beta)$-quantile error $O(m(D)T/\varepsilon)$. The lower bound applies to every private mechanism, without a computational restriction. The sample hypothesis is used through the normalization $\varepsilon n \geq 64 A_{U,\beta}$, which closes the packing case requiring more than n records and, with $\delta \leq n^{-2}$, gives the parameter comparison needed by the approximate branch.

The common clipping reduction. Both upper branches first replace x_i by its dyadic scale $b(x_i)$ and form

$$Y = \mathrm{Top}_K\big(\{b(x_i) : i \in [n]\}\big).$$

This transformation serves two purposes. Dyadic scaling reduces threshold selection from the original universe to only $L_U = O(1 + \log U)$ ordered scales, while retaining the largest K scales makes it enough to locate an interior point rather than estimate the maximum itself. Moreover, top-K order truncation changes by at most one replacement when the input database does, so a private selector run on Y remains private for the original database.

The deterministic clipping lemma is the central reduction. If $z \in [\min Y, \max Y]$ and $t = 2^z$, then

$$t < 2m(D), \qquad \#\{i : x_i > t\} \leq K - 1, \qquad 0 \leq S(D) - Q_t(D) \leq K m(D). \tag{13}$$

Thus an interior scale simultaneously keeps the sensitivity of Q_t within a constant factor of the instance scale and limits the clipping bias. Once $Q_t(D)$ is privately released as $\widehat{S}$, the two contributions obey

$$|\widehat{S} - S(D)| \leq K m(D) + |\widehat{S} - Q_t(D)|.$$

If a prescribed value of K exceeds n, the corresponding branch instead returns zero; the elementary bound $\varepsilon|S(D)|/m(D) \leq \varepsilon n < \varepsilon K$ gives the same order of error.

The pure-private branch. PureClipSum takes $K_{\mathrm{pure}} = \left\lceil \frac{8}{\varepsilon} \log \frac{2L_U}{\beta} \right\rceil = O(A_{U,\beta}/\varepsilon)$. On the top scales it applies the exponential mechanism, with half the privacy budget, to the balanced rank quality q_Y. A median of Y has quality at least $K_{\mathrm{pure}}/2$, whereas every scale outside the sample range has quality zero, so the selected scale is interior with probability at least $1 - \beta/2$. Using the remaining privacy budget, it returns $Q_t(D) + Z$ for $Z \sim \mathrm{Lap}(2t/\varepsilon)$. Its standard tail bound and Equation (13) give normalized β-error $O(A_{U,\beta})$. The required sorting, rank-quality computation, and clipped summation take $O(n \log n + L_U)$ exact-real operations.

The approximate-private interior-point construction. To obtain the smaller $B_{\varepsilon,\delta}$ rate, the proof must avoid the pure branch's $\log(L_U/\beta)$ cost for making the threshold interior and must replace the unbounded Laplace tail by a deterministic noise bound. The main technical challenge is therefore an approximate-private selector on an ordered universe of size d whose output is interior on every execution and whose sample cost depends only doubly logarithmically on d.

The starting mechanism, OneLevelPrefix, takes a multiset $S \in \{0, \ldots, d-1\}^K$ and privacy and failure parameters (α, ζ, η). It is (α, ζ)-private and, except with probability η, returns a robust interior point using

$$K = O\left(\frac{1}{\alpha}\left(1 + \log\log(4d) + \log\frac{1}{\zeta} + \log\frac{1}{\eta} + \log\frac{1}{\alpha}\right)\right) \tag{14}$$

records. Its main intermediate objects preserve the rank information needed to lift a selected prefix length to a valid point of the original universe. The mechanism stably bottom-truncates S, randomly pairs the retained records, and forms the multiset of their longest-common-prefix lengths. An exponential mechanism chooses a prefix length with both lower and upper rank margins. The random-pairing separation lemma and the upper margin guarantee that some binary-prefix cylinder at the next level is well populated. This ensures that the SparsePrefix subroutine privately returns an observed cylinder rather than $\perp$ with high probability; the returned cylinder need not be the populous one. The subroutine perturbs only observed prefix counts. Approximate privacy pays for the possible singleton bin that can appear or disappear under one replacement, so there is no need to enumerate all prefixes. On the same rank events, for any returned observed cylinder the lower margin and the largest records omitted by bottom truncation make at least one endpoint robust; a noisy count selects a robust endpoint with high probability.

The base mechanism is only high-probability interior. The SureInterior conversion runs two independent base calls with reduced privacy parameters, returns the first output that actually lies in $[\min S, \max S]$, and returns $\min S$ only if both calls fail. This validation-and-fallback step is data dependent, but its law is within $O(\eta^2)$ in total variation of the base law conditioned on being interior. Taking η proportional to $\min\{\alpha, \zeta\}$ leaves enough privacy slack to absorb the correction. The result is an (α, ζ)-private, always-interior selector with the sample bound in Equation (14) with the $\log(1/\eta)$ term absorbed by $\log(1/\alpha) + \log(1/\zeta)$. It runs in time polynomial in K, $\log d$, $\log(1/\zeta)$, and $\log(1/\alpha)$.

The approximate-private branch. For $\delta > 0$, define

$$G = 1 + \log\log(4L_U) + \log\frac{1}{\delta} + \log\frac{1}{\varepsilon}, \qquad K_{\text{app}} = \left\lceil \frac{C_1 G}{\varepsilon} \right\rceil.$$

ApproxClipSum applies SureInterior with parameters $(\varepsilon/2, \delta/2)$ to the top K_{app} scales. This makes Equation (13) deterministic. It then invokes TruncatedLaplaceRelease on Q_t. For a sensitivity-Δ scalar query this subroutine uses a Laplace density truncated to radius $R(\alpha, \zeta, \Delta) = (\Delta/\alpha)\log(1 + (e^\alpha - 1)/(2\zeta))$. The radius is chosen so that neighboring densities have likelihood ratio at most e^α on their common support and boundary mass at most ζ; thus the release is approximately private with deterministically bounded noise. With $(\alpha, \zeta, \Delta) = (\varepsilon/2, \delta/2, t)$, it gives

$$|Z| \le \frac{2t}{\varepsilon}\log\left(1 + \frac{e^{\varepsilon/2} - 1}{\delta}\right) \le \frac{4m(D)}{\varepsilon} B_{\varepsilon,\delta}. \tag{15}$$

The assumptions $\delta \le n^{-2}$ and $\varepsilon n \ge 64 A_{U,\beta}$ imply the parameter comparison $G = O(B_{\varepsilon,\delta})$. Therefore the clipping bias $K_{\text{app}} m(D)$ and the bounded noise in Equation (15) together give deterministic error $O(m(D)B_{\varepsilon,\delta}/\varepsilon)$. This branch runs in time polynomial in n, $\log U$, $\log(1/\delta)$, and $\log(1/\varepsilon)$. The public wrapper runs the pure branch when $\delta = 0$ or $A_{U,\beta} \le B_{\varepsilon,\delta}$, and otherwise runs the approximate branch, proving the upper bound with exactly one invocation.

The matching lower bound. Fix any (ε, δ)-private mechanism M, let $R(M)$ be its worst normalized error, and choose $r > R(M)$. Set $a = r/\varepsilon$, $k = \lfloor 2a \rfloor + 1$, and $h = 1 + \lfloor \log_4 U \rfloor$. When $k \le n$, compare the all-zero database with the h databases having k copies of 4^j and zeros elsewhere. Accuracy places their outputs in pairwise disjoint intervals: the factor four separates consecutive scales, and $k > 2a$ separates all of them from the interval around zero. Retaining the additive loss in exact approximate group privacy and summing these disjoint events under the all-zero database

gives

$$e^{-k\varepsilon} \le \frac{\beta/h + \gamma}{1 - \beta + \gamma}, \qquad \gamma = \begin{cases} \delta/(e^{\varepsilon} - 1), & \delta > 0, \\ 0, & \delta = 0. \end{cases} \tag{16}$$

This single inequality yields the two alternatives $\log(h/\beta)$ and $\log((e^{\varepsilon} - 1)/\delta)$, which are, up to universal additive constants, $A_{U,\beta}$ and $B_{\varepsilon,\delta}$. If $k > n$, the normalized sample condition gives the desired lower bound directly.

The additive constants in this logarithmic packing matter when T is small. A separate argument assumes $R(M) < 1/20$, chooses its own $r \in (R(M), 1/20)$ and corresponding k, and packs databases with 0, k, and $2k$ ones to obtain a contradiction. This proves the uniform floor $R(M) \ge 1/20$. Combining that floor with Equation (16) gives $R(M) \ge T/200$. Taking the infimum over all mechanisms proves the stated information-theoretic lower bound and completes the match with BranchClipSum.

4 Proof of the Main Theorem

Unless a statement explicitly changes the scope, all parameters in this section satisfy the hypotheses of Theorem 1.1.

4.1 Scope and parameter normalization

The floor in $m(D)$ makes the normalized error meaningful at the all-zero database. The restriction on β is also substantive.

Lemma 4.1 (Necessity of a bounded failure probability)**.** *If the restriction $\beta \le 1/3$ is dropped and $1 - \beta \le 1/(nU + 1)$, then $\rho^*_{\varepsilon,\delta,\beta}(n, U) = 0$ for every $\varepsilon, \delta \ge 0$.*

Proof. Let M ignore its input and output a uniformly random member of $\{0, 1, \ldots, nU\}$. This mechanism is $(0, 0)$-differentially private. For every D, it outputs $S(D)$ with probability $1/(nU + 1) \ge 1 - \beta$, so $\mathrm{err}_\beta(M, D) = 0$ for every database. □

We next put the sample hypothesis into the form used by both the upper and lower bounds and relate the exact expression $B_{\varepsilon,\delta}$ to its more familiar scale.

Lemma 4.2 (Parameter normalization)**.** *There is a universal choice of C_0 such that Equation (1) implies*

$$\varepsilon n \ge 64 A_{U,\beta}. \tag{17}$$

Moreover, for $0 < \varepsilon \le 1$ and $\delta > 0$,

$$B_{\varepsilon,\delta} = \Theta\left(1 + \log\left(1 + \frac{\varepsilon}{\delta}\right)\right)$$

with universal constants.

Proof. For $U \ge 2$, $L_U \le 5 \log U$, and hence

$$A_{U,\beta} \le 1 + \log\log U + \log(1/\beta) + \log 5.$$

Because $\beta \le 1/3$, the expression in parentheses in Equation (1) is bounded below by a positive universal constant. Increasing C_0 therefore gives Equation (17). Finally,

$$\varepsilon \le e^{\varepsilon} - 1 \le (e - 1)\varepsilon \qquad (0 < \varepsilon \le 1),$$

and substitution in the definition of $B_{\varepsilon,\delta}$ proves the second assertion. □

4.2 Stable dyadic clipping

The following deterministic fact controls the clipping bias and the scale of the subsequent noisy release.

Lemma 4.3 (Bias from an interior dyadic scale)**.** *Let $y_{[1]} \geq \cdots \geq y_{[n]}$ be the sorted values of $b(x_i)$. If $1 \leq K \leq n$, $y_{[K]} \leq z \leq y_{[1]}$, and $t = 2^z$, then*

$$t < 2m(D), \qquad \#\{i : x_i > t\} \leq K - 1,$$

and therefore

$$0 \leq S(D) - Q_t(D) \leq Km(D). \tag{18}$$

Proof. Let $M_D = \max_i x_i$. Since $z \leq b(M_D)$,

$$t \leq 2^{b(M_D)} < 2\max\{1, M_D\} = 2m(D);$$

when $M_D \leq 1$, the sharper equality $t = 1$ holds. If $x_i > 2^z$, then $b(x_i) > z \geq y_{[K]}$. Fewer than K scales strictly exceed $y_{[K]}$, which proves the count bound. Hence

$$S(D) - Q_t(D) = \sum_i \max\{x_i - t, 0\} \leq \#\{i : x_i > t\}\, M_D \leq Km(D).$$

□

4.3 The pure-private branch

Set

$$K_{\text{pure}} = \left\lceil \frac{8}{\varepsilon} \log \frac{2L_U}{\beta} \right\rceil. \tag{19}$$

PureClipSum: Pure-private clipping and release.
Input: A database $D \in \{0, \ldots, U\}^n$; public integers $U \geq 2$, $n \geq 1$ and public parameters $0 < \varepsilon \leq 1$, $0 < \beta < 1$.
Output: A real estimate of $S(D)$.
Guarantee: $(\varepsilon, 0)$-privacy and normalized β-error $O(A_{U,\beta})$.
1. If $K_{\text{pure}} > n$, return 0.
2. Form $Y = \text{Top}_{K_{\text{pure}}}(\{b(x_i) : i \in [n]\})$ and sample $z \in \{0, \ldots, L_U - 1\}$ with probability proportional to $\exp(\varepsilon q_Y(z)/4)$.
3. Set $t = 2^z$, independently draw $Z \sim \text{Lap}(2t/\varepsilon)$, and return $Q_t(D) + Z$.

Lemma 4.4 (Pure-private clipping guarantee)**.** *For integers $U \geq 2$, $n \geq 1$ and parameters $0 < \varepsilon \leq 1$, $0 < \beta < 1$,* PureClipSum *is $(\varepsilon, 0)$-differentially private and, for every D,*

$$\Pr\left[|\textsc{PureClipSum}(D) - S(D)| \leq \frac{Cm(D)}{\varepsilon} A_{U,\beta}\right] \geq 1 - \beta$$

for a universal C. It runs in $O(n \log n + L_U)$ exact-real operations.

Proof. If $K_{\text{pure}} > n$, the output is data independent and

$$\frac{\varepsilon|S(D)|}{m(D)} \le \varepsilon n < \varepsilon K_{\text{pure}} = O(A_{U,\beta}).$$

Assume $K_{\text{pure}} \le n$. By Lemma 2.1 and the sensitivity-one quality Equation (4), the exponential-mechanism step is $(\varepsilon/2, 0)$-private. Conditional on every fixed preceding output t, Lemma 2.2 shows that Q_t has sensitivity at most t, so adding $\text{Lap}(2t/\varepsilon)$ is also $(\varepsilon/2, 0)$-private. Adaptive composition proves $(\varepsilon, 0)$-privacy.

Some median of Y has quality at least $K_{\text{pure}}/2$, while every point strictly outside $[\min Y, \max Y]$ has quality zero. Thus Equation (9) gives

$$\Pr[z \notin [\min Y, \max Y]] \le L_U e^{-\varepsilon K_{\text{pure}}/8} \le \frac{\beta}{2}.$$

The Laplace tail Equation (8) also gives

$$\Pr\left[|Z| > \frac{2t}{\varepsilon} \log \frac{2}{\beta}\right] \le \frac{\beta}{2}.$$

On the complementary events, Lemma 4.3 yields

$$|Q_t(D) + Z - S(D)| \le K_{\text{pure}} m(D) + \frac{4m(D)}{\varepsilon} \log \frac{2}{\beta} \le \frac{Cm(D)}{\varepsilon} A_{U,\beta}.$$

A union bound proves accuracy.

For the resource bound, compute and sort the n scales, fill an L_U-entry histogram of the selected scales, and use one prefix and one suffix pass to compute all qualities. These operations cost $O(n \log n + L_U)$. Sampling the resulting categorical law costs $O(L_U)$ exact-real operations, evaluating Q_t costs $O(n)$, and the Laplace draw costs $O(1)$. □

4.4 Bounded noise for approximate privacy

The approximate-private branch uses the truncated-Laplace law from the preliminaries. Writing $\lambda = \alpha/\Delta$, $R = R(\alpha, \zeta, \Delta)$, and $E = e^{-\lambda R}$, its magnitude is

$$-\lambda^{-1} \log\big(1 - (1 - E)V\big)$$

for V uniform on $[0, 1]$, and its sign is an independent fair sign.

TruncatedLaplaceRelease: Bounded-noise scalar release.
Input: A database D, a scalar query q, and public parameters $\Delta > 0$, $\alpha > 0$, and $0 < \zeta \le 1/2$, where q has replacement sensitivity at most Δ.
Output: A real release of $q(D)$.
Guarantee: (α, ζ)-privacy and deterministic error at most $R(\alpha, \zeta, \Delta)$.
1. Put $R = R(\alpha, \zeta, \Delta)$ and $\lambda = \alpha/\Delta$. Draw V uniformly from $[0, 1]$ and an independent fair sign $\sigma \in \{-1, 1\}$.
2. Set $W = -\sigma\lambda^{-1} \log(1 - (1 - e^{-\lambda R})V)$ and return $q(D) + W$.

Lemma 4.5 (Privacy and bounded error of truncated Laplace release)**.** *For every input in its specification,* TruncatedLaplaceRelease *is (α, ζ)-differentially private and has absolute error at most $R(\alpha, \zeta, \Delta)$ on every execution. If evaluating $q(D)$ takes time $T_q(D)$, its running time is $T_q(D) + O(1)$ in the exact-real model.*

Proof. Consider adjacent query values $q(D) \le q(D') = q(D) + s$ with $0 \le s \le \Delta$. Because $\zeta \le 1/2$, $R(\alpha, \zeta, \Delta) \ge \Delta$. On the overlap of the two shifted supports, the density ratio is at most $e^{\alpha s/\Delta} \le e^{\alpha}$. The part of either support outside the other has probability at most the mass of a boundary interval of width Δ. With $R = R(\alpha, \zeta, \Delta)$, this mass is

$$\frac{e^{-\alpha(R-\Delta)/\Delta} - e^{-\alpha R/\Delta}}{2(1 - e^{-\alpha R/\Delta})} = \frac{e^{\alpha} - 1}{2(e^{\alpha R/\Delta} - 1)} = \zeta.$$

Splitting an arbitrary event into its overlap and nonoverlap parts proves the privacy inequality in both directions. The support gives the deterministic error bound. The inverse-CDF representation in the algorithm uses a constant number of exact-real operations and random draws in addition to evaluating $q(D)$. □

4.5 A rank-preserving interior-point construction

We now construct an always-interior selector on the ordered universe $\mathcal{X} = \{0, \dots, d-1\}$ for $d \ge 2$, using the binary-prefix notation from the preliminaries. The next selector inspects only prefixes that occur in its input. Approximate privacy pays for the possible singleton prefix that appears or disappears under a replacement, avoiding enumeration of all 2^s prefixes.

For $0 < \rho, \zeta, \nu \le 1/4$, define

$$b_\rho = \frac{2}{\rho}, \qquad \tau_{\rho,\zeta} = 1 + \frac{2}{\rho} \log \frac{1 + e^{\rho}}{2\zeta}, \qquad h_{\rho,\zeta,\nu} = \left\lceil \tau_{\rho,\zeta} + \frac{2}{\rho} \log \frac{1}{2\nu} \right\rceil. \tag{20}$$

SparsePrefix: Private selection from observed prefixes.
Input: A multiset D over $\mathcal{X}$, a prefix length $s \in \{0, \dots, \ell\}$, and parameters $0 < \rho, \zeta, \nu \le 1/4$, in this order.
Output: An observed length-s prefix, or $\perp$.
Guarantee: (ρ, ζ)-privacy; if some prefix occurs at least $h_{\rho,\zeta,\nu}$ times, the output is an observed prefix with probability at least $1 - \nu$.
1. Form the observed-prefix histogram $\{(p, c_D(p)) : p \in G_s(D)\}$.
2. For each $p \in G_s(D)$, independently draw $E_p \sim \mathrm{Lap}(b_\rho)$ and declare p active when $c_D(p) + E_p \ge \tau_{\rho,\zeta}$.
3. Return the lexicographically first active prefix, or return $\perp$ if none is active.

Lemma 4.6 (Private sparse-prefix selection)**.** SparsePrefix *has the privacy and utility guarantees in its specification. It runs in $O(|D|\ell)$ bit operations and uses at most $|D|$ exact Laplace samples after the histogram is formed.*

Proof. Fix adjacent multisets D, D' and let $H = G_s(D) \cap G_s(D')$. Each of $G_s(D) \setminus H$ and $G_s(D') \setminus H$ has cardinality at most one, and a prefix in either difference has count one in the multiset where it occurs. Put $u = \zeta/(1 + e^{\rho})$. Such a singleton prefix is activated with probability

$$\Pr[1 + E_p \ge \tau_{\rho,\zeta}] = \frac{1}{2} \exp\left(-\frac{\tau_{\rho,\zeta} - 1}{2/\rho} \right) = \frac{\zeta}{1 + e^{\rho}} = u. \tag{21}$$

Let Q_D be the same procedure restricted to the common prefix set H, using the counts from D, and define $Q_{D'}$ analogously. Coupling noises on H shows that the output law of SparsePrefix(D)

is within total variation u of that of Q_D, and likewise for D'. On H, the two count vectors have ℓ_1-distance at most two. Adding independent $\mathrm{Lap}(2/\rho)$ noise therefore changes their joint density by a factor at most e^ρ. Thresholding and selecting the first active coordinate are postprocessing, so Q_D and $Q_{D'}$ satisfy pure ρ-privacy. Consequently, for every output event T,

$$\begin{aligned}\Pr[\textsc{SparsePrefix}(D) \in T] &\le \Pr[Q_D \in T] + u \\ &\le e^\rho \Pr[Q_{D'} \in T] + u \\ &\le e^\rho \Pr[\textsc{SparsePrefix}(D') \in T] + (1 + e^\rho)u,\end{aligned}$$

which is the (ρ, ζ)-privacy inequality. Interchanging D, D' gives the reverse direction.

If $c_D(p) \ge h_{\rho,\zeta,\nu}$ for some p, then this prefix is inactive only if

$$E_p \le -\frac{2}{\rho} \log \frac{1}{2\nu},$$

an event of probability at most ν. Whenever it is active, the algorithm returns some member of $G_s(D)$. Finally, a binary-trie scan constructs the at most $|D|$ observed bins and their counts in $O(|D|\ell)$ bit operations; the remaining resource claims follow directly from the algorithm. □

To turn a privately selected prefix length into a useful cylinder, we need a random pairing to contain a pair that is well separated in input rank.

Lemma 4.7 (Random-pairing separation)**.** *Let N be even, let $(\Pi_1, \dots, \Pi_N)$ be a uniformly random permutation of $\{1, \dots, N\}$, and pair consecutive positions. For every integer $r \ge 1$,*

$$\Pr\left[\#\left\{j : |\Pi_{2j-1} - \Pi_{2j}| \le r/12\right\} \ge r\right] \le 2^{-r}. \tag{22}$$

Proof. If $r > N/2$, the event is empty. Otherwise call a pair short when its labels differ by at most $r/12$. From a permutation with at least r short pairs, select canonically the first r of them. There are at most $\binom{N}{r}$ choices for the labels in their first positions, and each such label has at most $r/6$ possible partner labels. Compressing the selected directed pairs into blocks leaves $N - r$ distinct objects, which can be ordered in at most $(N - r)!$ ways. This gives an injective encoding after the canonical selection, and hence the probability is at most

$$\frac{\binom{N}{r}(r/6)^r (N-r)!}{N!} = \frac{(r/6)^r}{r!} \le \left(\frac{e}{6}\right)^r < 2^{-r},$$

where $r! \ge (r/e)^r$. □

We now choose public parameters that provide both lower and upper rank margins. For $0 < \alpha, \zeta, \eta \le 1/4$, put

$$\rho = \frac{\alpha}{3}, \qquad \nu = \frac{\eta}{4}, \qquad h = h_{\rho,\zeta,\nu}, \tag{23}$$

$$r = \left\lceil \max\left\{24h, \frac{2}{\rho}\log\frac{1}{2\nu}, \frac{\log(1/\nu)}{\log 2}, 12\right\}\right\rceil, \tag{24}$$

and

$$K_0 = 6r + \frac{8}{\rho}\log\frac{\ell+1}{\nu} + 2. \tag{25}$$

For an input size $K \ge K_0$, define

$$m = \left\lfloor \frac{K - 2r}{2} \right\rfloor, \qquad N = 2m. \tag{26}$$

These definitions ensure

$$K - N \geq 2r, \qquad m \geq 2r + \frac{4}{\rho} \log \frac{\ell + 1}{\nu}. \tag{27}$$

For the prefix-length multiset $W = (w_1, \ldots, w_m)$, we use the rank quality q_W and the two rank margins in Equations (4) and (5).

OneLevelPrefix: A robust one-level prefix lift.
Input: A multiset $S \in \{0, \ldots, d-1\}^K$ and parameters α, ζ, η satisfying $0 < \alpha, \zeta, \eta \leq 1/4$ and $K \geq K_0$, in this order.
Output: A point of $\{0, \ldots, d-1\}$.
Guarantee: The mechanism is permutation invariant and (α, ζ)-private; except with probability η, its output is an r-robust interior point of S.

1. Form $B = \mathrm{Bot}_N(S)$, uniformly order its labelled occurrences as $(y_1, \ldots, y_N)$, and set $w_j = \lambda(y_{2j-1}, y_{2j})$ for $j \in [m]$.
2. Sample $z \in \{0, \ldots, \ell\}$ with probability proportional to $\exp(\rho q_W(z)/2)$, where $W = (w_1, \ldots, w_m)$.
3. Set $s = \min\{z + 1, \ell\}$ and $p = \textsc{SparsePrefix}(B, s; \rho = \rho, \zeta = \zeta, \nu = \nu)$. If $p = \perp$, return 0; otherwise set $a = a(p)$ and $b = b(p)$.
4. If $z = \ell$, return a. Otherwise compute $c_b(S)$, draw $E \sim \mathrm{Lap}(1/\rho)$, and return b when $c_b(S) + E \geq 3r/2$, or a otherwise.

Lemma 4.8 (One-level robust prefix lift)**.** *For every input satisfying its specification,* OneLevelPrefix *is permutation invariant and (α, ζ)-differentially private. With probability at least $1 - \eta$, its output x is r-robust, namely*

$$\min S \leq x, \qquad \#\{y \in S : y \geq x\} \geq r,$$

and in particular is an interior point. Its running time is polynomial in K, $\log d$, $\log(1/\zeta)$, and $\log(1/\eta)$ in the exact-real model.

Proof. We begin with privacy. By Lemma 2.1, the bottom truncations of adjacent inputs differ by at most one replacement. Match their common labelled occurrences and their possible exceptional occurrences, and place matched occurrences in the same positions of the two uniform random orderings. This couples the orderings so that they differ in at most one position, and their paired LCP multisets therefore differ in at most one coordinate. Since Equation (4) has sensitivity one, the exponential mechanism selecting z is pure ρ-private.

Conditional on every selected z, the prefix length s is fixed, and Lemmas 2.1 and 4.6 make the sparse-prefix call (ρ, ζ)-private as a mechanism of the original input. Conditional on every transcript (z, p), the final count $c_b(S)$ has replacement sensitivity one; adding $\mathrm{Lap}(1/\rho)$ before thresholding is pure ρ-private. The remaining branches are postprocessing. Adaptive composition gives

$$(\rho, 0) + (\rho, \zeta) + (\rho, 0) = (\alpha, \zeta). \tag{28}$$

We turn to utility. A median of W has rank quality at least $m/2$. The exponential tail Equation (9) and Equation (27) imply, conditional on every W,

$$\Pr[q_W(z) < r \mid W] \leq (\ell + 1) \exp\left[-\frac{\rho}{2}\left(\frac{m}{2} - r\right)\right] \leq \nu. \tag{29}$$

Thus both rank margins in Equation (5) hold except with probability ν.

Label the occurrences of B by $1, \dots, N$ in nondecreasing value order, breaking ties arbitrarily. The random ordering induces a uniform permutation of these ranks. By Lemma 4.7 and the definition of r, except with probability at most

$$2^{-r} \le \nu, \tag{30}$$

fewer than r pairs have rank distance at most $r/12$. Suppose this event and $q_W(z) \ge r$ both hold. Among the at least r pairs with LCP length at least z, one pair has rank distance greater than $r/12$. More than $r/12$ records of B lie between its endpoints in value order, inclusively, and all share their first z bits. If $z < \ell$, at least half share one of the two possible next bits; if $z = \ell$, all have the full common binary representation. Hence some prefix of length $s = \min\{z+1, \ell\}$ occurs more than $r/24 \ge h$ times in B.

By Lemma 4.6, the call $\textsc{SparsePrefix}(B, s; \rho = \rho, \zeta = \zeta, \nu = \nu)$ returns an observed prefix p except with probability ν. Fix a witness $y^* \in B \cap C(p)$, so

$$a \le y^* \le b. \tag{31}$$

Every one of the at least $K - N \ge 2r$ omitted largest records is at least every member of B. If $z = \ell$, then $a = b = y^*$, so the returned value has a lower witness and at least $2r$ records at or above it.

Suppose $z < \ell$. The lower-rank margin in Equation (5) supplies a pair whose LCP length is at most z. The two endpoints cannot both lie in the length-$(z+1)$ cylinder $[a, b]$. Thus either its smaller endpoint is below a, or its larger endpoint is above b. In the former case, a has a lower witness and at least $2r$ inputs at or above it. In the latter case, b has the lower witness Equation (31), and the omitted records show that at least $2r$ inputs are at or above b. At least one endpoint is therefore robust.

It remains to analyze the noisy choice. The point a always has at least $2r$ inputs at or above it, and it is robust precisely when some input is at most a. The point b always has the witness $y^* \le b$, and it is robust whenever $c_b(S) \ge r$. If $c_b(S) < r$, choosing b requires $E \ge r/2$ and has probability at most

$$\Pr[E \ge r/2] \le \tfrac{1}{2} e^{-\rho r/2} \le \nu. \tag{32}$$

If a is not robust, the separating pair must instead have an endpoint above b, so $c_b(S) \ge 2r$; choosing a then requires $E < -r/2$ and has the same probability bound. If both endpoints are robust, either choice works.

The failures in Equation (29), Equation (30), the sparse-prefix call, and Equation (32) each have probability at most ν. Their union has probability at most $4\nu = \eta$, proving the claimed invariant.

Finally, sorting costs $O(K \log K)$ comparisons; computing the LCP values and the sparse observed-prefix histogram costs $O(K\ell)$ bit operations; histograms compute all $\ell + 1$ balanced qualities in $O(K + \ell)$ exact-real operations; and the endpoint count costs $O(K)$. Since $\ell = O(\log d)$ and all public thresholds use only the displayed logarithms, the running time has the asserted polynomial bound. Uniformly ordering labelled occurrences also shows permutation invariance. □

The preceding exact sample threshold has the asymptotic form needed by the summation algorithm.

Lemma 4.9 (One-level interior-point sample bound)**.** *Let*

$$\Xi = 1 + \log\log(4d) + \log\frac{1}{\zeta} + \log\frac{1}{\eta} + \log\frac{1}{\alpha}.$$

For $d \geq 2$ and $0 < \alpha, \zeta, \eta \leq 1/4$, the parameters defined in Equations (23) to (25) satisfy

$$K_0 \leq \frac{2256}{\alpha}\Xi.$$

Consequently, $K \geq 3000\Xi/\alpha$ implies the precondition of OneLevelPrefix.

Proof. Since $\rho = \alpha/3 \leq 1/12$ and $\nu = \eta/4$,

$$\log \frac{1+e^{\rho}}{2\zeta} \leq 1 + \log \frac{1}{\zeta}, \qquad \log \frac{1}{2\nu} = \log \frac{2}{\eta} \leq 1 + \log \frac{1}{\eta}.$$

Thus Equation (20) gives $h \leq 5\Xi/\rho$. Each of the other three entries in the maximum in Equation (24) is at most $120\Xi/\rho$, as is $24h$, and therefore

$$r \leq \frac{121\Xi}{\rho}. \tag{33}$$

Moreover,

$$\ell + 1 \leq \log_2(4d), \qquad \log \frac{\ell+1}{\nu} \leq 3\Xi.$$

Substituting these bounds into Equation (25), and using $2 \leq 2\Xi/\rho$, yields

$$K_0 \leq \frac{(6 \cdot 121 + 24 + 2)\Xi}{\rho} = \frac{752\Xi}{\rho} = \frac{2256\Xi}{\alpha}.$$

□

We next convert the high-probability interior point into an output that is interior on every execution. For $0 < \alpha, \zeta \leq 1/2$, define

$$\alpha_0 = \frac{\alpha}{4}, \qquad \zeta_0 = \frac{\zeta}{16}, \qquad \eta = \frac{\min\{\alpha, \zeta\}}{64}. \tag{34}$$

Let

$$K_{\text{sure}} = \frac{3000}{\alpha_0}\left(1 + \log\log(4d) + \log\frac{1}{\zeta_0} + \log\frac{1}{\eta} + \log\frac{1}{\alpha_0}\right). \tag{35}$$

By Lemmas 4.8 and 4.9, running OneLevelPrefix with parameters $(\alpha_0, \zeta_0, \eta)$ on at least K_{sure} records is (α_0, ζ_0)-private and returns an interior point with probability at least $1 - \eta$.

SureInterior: Conversion to an always-interior selector.
Input: A multiset $Y \in \{0, \ldots, d-1\}^K$ and parameters $0 < \alpha, \zeta \leq 1/2$, in this order, with $K \geq K_{\text{sure}}$.
Output: A point of $\{0, \ldots, d-1\}$.
Guarantee: (α, ζ)-privacy and output in $[\min Y, \max Y]$ on every execution.
1. Independently compute $Z_1 = \text{OneLevelPrefix}(Y; \alpha = \alpha_0, \zeta = \zeta_0, \eta = \eta)$ and $Z_2 = \text{OneLevelPrefix}(Y; \alpha = \alpha_0, \zeta = \zeta_0, \eta = \eta)$.
2. Return the first Z_j that belongs to $[\min Y, \max Y]$; if neither does, return $\min Y$.

Lemma 4.10 (Always-interior conversion)**.** *There is a universal C such that* SureInterior *is an (α, ζ)-differentially private, always-interior mechanism whenever*

$$K \geq \frac{C}{\alpha}\left(1 + \log\log(4d) + \log\frac{1}{\zeta} + \log\frac{1}{\alpha}\right). \tag{36}$$

If T_{base} is the running time of one base call with the parameters in Equation (34), then the running time is $2T_{\text{base}} + O(K)$. It is polynomial in K, $\log d$, $\log(1/\zeta)$, and $\log(1/\alpha)$.

Proof. The parameters in Equation (34) lie in the range required by Lemma 4.8, and

$$\log(1/\eta) \leq \log 64 + \log(1/\alpha) + \log(1/\zeta).$$

Thus a universal choice of C in Equation (36) makes its right-hand side at least K_{sure}.

Let P_Y be the output law of one base call, let $G_Y = [\min Y, \max Y]$, and let C_Y be P_Y conditioned on G_Y. For every measurable T, base privacy and base failure probability at most η imply

$$\begin{aligned} C_Y(T) &\leq \frac{P_Y(T)}{1-\eta} \\ &\leq \frac{e^{\alpha_0} P_{Y'}(T) + \zeta_0}{1-\eta} \\ &\leq \frac{e^{\alpha_0} C_{Y'}(T) + e^{\alpha_0}\eta + \zeta_0}{1-\eta}. \end{aligned}$$

Hence the family C_Y is (α_C, ζ_C)-private for

$$\alpha_C = \alpha_0 + \log\frac{1}{1-\eta}, \qquad \zeta_C = \frac{\zeta_0 + e^{\alpha_0}\eta}{1-\eta}. \tag{37}$$

For $x \in \mathcal{X}$, let δ_x denote the point-mass law at x. If $p_Y = P_Y(G_Y^c)$, the output law Q_Y of SureInterior is

$$Q_Y = (1 - p_Y^2) C_Y + p_Y^2\, \delta_{\min Y},$$

so $d_{\text{TV}}(Q_Y, C_Y) \leq \eta^2$. Apply Equation (10) with $\tau = \eta^2$ and the parameters in Equation (37). Since $\eta \leq 1/128$ and $\eta \leq \alpha/64$, we have $\log(1/(1-\eta)) \leq 2\eta$ and therefore

$$\alpha_C \leq \alpha/4 + 2\eta \leq \alpha/4 + \alpha/32 < \alpha.$$

Also $\eta \leq \zeta/64$, $e^{\alpha_0} \leq e^{1/8} < 8/7$, and therefore

$$\zeta_C \leq \frac{128}{127}\left(\frac{1}{16} + \frac{8/7}{64}\right)\zeta < \frac{\zeta}{12}.$$

Finally, $\eta^2 \leq \zeta/8192$ and $\alpha_C < 1/2$, so

$$\zeta_C + (1 + e^{\alpha_C})\eta^2 < \zeta.$$

This proves (α, ζ)-privacy. The output is interior by construction. One scan computes $\min Y$ and $\max Y$ in $O(K)$ time, and the algorithm makes exactly two base calls; the resource bounds now follow from Lemma 4.8. □

4.6 The approximate-private branch

It remains to compare the sample requirement of the interior-point construction with the target error rate. For $\delta > 0$, put

$$G = 1 + \log\log(4L_U) + \log\frac{1}{\delta} + \log\frac{1}{\varepsilon}. \tag{38}$$

Lemma 4.11 (Approximate-branch parameter comparison)**.** *Under Equation (1), if $0 < \delta \le n^{-2}$, then*

$$G \le CB_{\varepsilon,\delta} \tag{39}$$

for a universal C.

Proof. Put $h = 1 + \log L_U$. By Lemma 4.2, $\varepsilon n \ge 64A_{U,\beta} \ge 64h$. Write

$$u = \log\frac{\varepsilon}{\delta}, \qquad e_0 = \log\frac{1}{\varepsilon}.$$

Since $\delta \le n^{-2}$,

$$\frac{\varepsilon}{\delta} \ge \varepsilon n^2 = \frac{(\varepsilon n)^2}{\varepsilon} \ge \frac{(64h)^2}{\varepsilon}.$$

Consequently $u \ge e_0$, and $e^\varepsilon - 1 \ge \varepsilon$ gives $B_{\varepsilon,\delta} \ge 1 + u$. Furthermore,

$$\log\frac{1}{\delta} + \log\frac{1}{\varepsilon} = (u + e_0) + e_0 \le 3u \le 3B_{\varepsilon,\delta}.$$

Finally, $\log(4L_U) \le 2h$, whereas $u \ge 2\log(64h)$, and hence $1 + \log\log(4L_U) = O(1+u)$. Combining the estimates proves the claim. □

Fix a sufficiently large universal C_1 and define

$$K_{\text{app}} = \left\lceil \frac{C_1 G}{\varepsilon} \right\rceil. \tag{40}$$

The choice of C_1 ensures that Equation (36) holds with $d = L_U$, $\alpha = \varepsilon/2$, and $\zeta = \delta/2$ whenever $K_{\text{app}} \le n$.

ApproxClipSum: Always-accurate approximate-private clipping.
Input: A database $D \in \{0, \dots, U\}^n$; public parameters $U, n, \varepsilon, \delta, \beta$ satisfying the hypotheses of Theorem 1.1, with $\delta > 0$.
Output: A real estimate of $S(D)$.
Guarantee: (ε, δ)-privacy and deterministic normalized error $O(B_{\varepsilon,\delta})$.
1. If $K_{\text{app}} > n$, return 0.
2. Form $Y = \text{Top}_{K_{\text{app}}}(\{b(x_i) : i \in [n]\})$ and compute $z = \textsc{SureInterior}(Y, \varepsilon/2, \delta/2)$.
3. Put $t = 2^z$ and return $\textsc{TruncatedLaplaceRelease}(D, Q_t; \Delta = t, \alpha = \varepsilon/2, \zeta = \delta/2)$.

Lemma 4.12 (Approximate-private clipping guarantee)**.** *Under the hypotheses of Lemma 4.11,* ApproxClipSum *is (ε, δ)-differentially private and satisfies, for every D,*

$$|\textsc{ApproxClipSum}(D) - S(D)| \le \frac{Cm(D)}{\varepsilon} B_{\varepsilon,\delta}$$

on every execution. It runs in time polynomial in n, $\log U$, $\log(1/\delta)$, and $\log(1/\varepsilon)$ in the exact-real model.

Proof. If $K_{\mathrm{app}} > n$, the output is data independent and

$$\frac{\varepsilon|S(D)|}{m(D)} \le \varepsilon n < \varepsilon K_{\mathrm{app}} = O(G) = O(B_{\varepsilon,\delta})$$

by Lemma 4.11.

Suppose $K_{\mathrm{app}} \le n$. By Lemmas 2.1 and 4.10, the selector is $(\varepsilon/2, \delta/2)$-private. Conditional on every fixed selected threshold t, Lemmas 2.2 and 4.5 make the release $(\varepsilon/2, \delta/2)$-private. Adaptive composition proves overall (ε, δ)-privacy.

The selected z always lies between the minimum and maximum of Y. Thus Lemma 4.3 bounds the clipping bias by $K_{\mathrm{app}} m(D)$ and gives $t < 2m(D)$. The added noise has magnitude at most

$$\begin{aligned} R\left(\frac{\varepsilon}{2}, \frac{\delta}{2}, t\right) &= \frac{2t}{\varepsilon} \log\left(1 + \frac{e^{\varepsilon/2} - 1}{\delta}\right) \\ &\le \frac{4m(D)}{\varepsilon} B_{\varepsilon,\delta}. \end{aligned}$$

Consequently,

$$\frac{\varepsilon|\textsc{ApproxClipSum}(D) - S(D)|}{m(D)} \le \varepsilon K_{\mathrm{app}} + 4B_{\varepsilon,\delta} = O(G + B_{\varepsilon,\delta}) = O(B_{\varepsilon,\delta}),$$

deterministically.

For the resource bound, computing and sorting the scales costs $O(n \log n)$. The selector makes two calls to OneLevelPrefix, and evaluating Q_t costs $O(n)$; the bounded-noise release adds $O(1)$ exact-real operations. Here $K_{\mathrm{app}} \le n$, $d = L_U$, and the base parameters are

$$(\alpha_0, \zeta_0, \eta) = \left(\frac{\varepsilon}{8}, \frac{\delta}{32}, \frac{\min\{\varepsilon, \delta\}}{128}\right).$$

The claimed polynomial follows from Lemmas 4.8 and 4.10, since $L_U = O(1 + \log U)$ and $\log(1/\eta) = O(1 + \log(1/\varepsilon) + \log(1/\delta))$. □

4.7 The matching lower bound

For a fixed mechanism M, abbreviate its worst normalized error by

$$R(M) = \sup_{D \in \{0,\dots,U\}^n} \frac{\varepsilon \operatorname{err}_\beta(M, D)}{m(D)}.$$

The first packing places geometrically separated databases around the all-zero database and retains the exact additive loss from group privacy.

Lemma 4.13 (Geometric star-packing inequality)**.** *Let M be (ε, δ)-differentially private, let $r > R(M)$, put $a = r/\varepsilon$ and $k = \lfloor 2a \rfloor + 1$, and define*

$$h = 1 + \lfloor \log_4 U \rfloor, \qquad \gamma = \begin{cases} \delta/(e^\varepsilon - 1), & \delta > 0, \\ 0, & \delta = 0. \end{cases}$$

If $k \le n$, then

$$e^{-k\varepsilon} \le \frac{\beta/h + \gamma}{1 - \beta + \gamma}, \qquad k\varepsilon \ge \log \frac{1 - \beta + \gamma}{\beta/h + \gamma}. \tag{41}$$

Proof. For $j = 0, \dots, h-1$, let $v_j = 4^j \le U$. Let $D_\circ$ be the all-zero database, and let D_j contain k copies of v_j and $n-k$ zeroes. Since $r > R(M)$, the mechanism assigns probability at least $1-\beta$ to

$$I_\circ = [-a, a] \quad \text{under } D_\circ, \qquad I_j = [(k-a)v_j, (k+a)v_j] \quad \text{under } D_j.$$

Indeed, each displayed radius is strictly larger than $\mathrm{err}_\beta(M, D)$ for the corresponding database, and the set of radii attaining probability at least $1-\beta$ is upward closed. These intervals are pairwise disjoint. The inequality $k > 2a$ separates $I_\circ$ from I_0, and consecutive geometric intervals are disjoint because

$$(k+a)v_j < (k-a)v_{j+1} \quad \Longleftrightarrow \quad 5a < 3k,$$

which follows from $k > 2a$.

The replacement distance from $D_\circ$ to D_j is k. Applying Equation (7) from D_j toward $D_\circ$ to the event I_j and rearranging exactly yields

$$\Pr[M(D_\circ) \in I_j] \ge e^{-k\varepsilon}(1 - \beta + \gamma) - \gamma.$$

The interval $I_\circ$ already has probability at least $1-\beta$ under $D_\circ$, so the other disjoint intervals have total probability at most β. Summing the displayed lower bound on $\Pr[M(D_\circ) \in I_j]$ over j proves both forms in Equation (41). □

The logarithmic packing alone loses an additive constant. A three-point count packing supplies a uniform positive floor before those constants are absorbed.

Lemma 4.14 (Uniform constant error floor)**.** *If $\varepsilon n \ge 64 A_{U,\beta}$ and $\delta \le n^{-2}$, then every (ε, δ)-differentially private mechanism M satisfies $R(M) \ge 1/20$.*

Proof. Suppose that $R(M) < 1/20$, and choose $R(M) < r < 1/20$. Put

$$a = \frac{r}{\varepsilon}, \qquad k = \lfloor 2a \rfloor + 1, \qquad \gamma = \begin{cases} \delta/(e^\varepsilon - 1), & \delta > 0, \\ 0, & \delta = 0. \end{cases}$$

Then

$$k < \frac{0.1}{\varepsilon} + 1 \le \frac{1.1}{\varepsilon}.$$

Since $A_{U,\beta} \ge 1 + \log 6$, the sample hypothesis gives $n > 2.2/\varepsilon > 2k$.

For $q \in \{0, k, 2k\}$, let E_q contain q ones and $n-q$ zeroes, and set $J_q = [q-a, q+a]$. These intervals are disjoint because $k > 2a$. Moreover, $m(E_q) = 1$ and $a > \mathrm{err}_\beta(M, E_q)$; by upward closure of the admissible quantile radii, J_q has probability at least $1-\beta$ under E_q. Each of E_0 and E_{2k} is at replacement distance k from E_k. Group privacy in Equation (7), rearranged toward E_k, gives

$$\Pr[M(E_k) \in J_0],\ \Pr[M(E_k) \in J_{2k}] \ge e^{-k\varepsilon}(1-\beta) - \gamma.$$

Since J_k itself has mass at least $1-\beta$ under E_k, disjointness implies

$$\beta \ge 2e^{-k\varepsilon}(1-\beta) - 2\gamma. \tag{42}$$

Here $k\varepsilon < 1.1$. Also, using $e^\varepsilon - 1 \ge \varepsilon$, $\delta \le n^{-2}$, and the sample hypothesis,

$$\gamma \le \frac{1}{\varepsilon n^2} \le \frac{1}{(\varepsilon n)^2} < \frac{1}{18000}.$$

Because $1 - \beta \ge 2/3$, the right-hand side of Equation (42) is strictly greater than $(4/3)e^{-1.1} - 2/18000 > 1/3 \ge \beta$, a contradiction. □

Combining the exact star inequality with the constant floor yields the full rate. The case in which the packing would require too many records is handled directly.

Lemma 4.15 (Matching minimax lower bound)**.** *If $\varepsilon n \geq 64 A_{U,\beta}$ and $\delta \leq n^{-2}$, then*

$$\rho^*_{\varepsilon,\delta,\beta}(n, U) \geq \frac{1}{200} \min\{A_{U,\beta}, B_{\varepsilon,\delta}\}.$$

Proof. Fix an (ε, δ)-differentially private mechanism M. If $R(M) = +\infty$, there is nothing to prove. Otherwise choose $r > R(M)$ and put

$$a = \frac{r}{\varepsilon}, \qquad k = \lfloor 2a \rfloor + 1, \qquad T = \min\{A_{U,\beta}, B_{\varepsilon,\delta}\}, \qquad h = 1 + \lfloor \log_4 U \rfloor.$$

If $k > n$, then $\lfloor 2a \rfloor \geq n$, and hence

$$r \geq \frac{\varepsilon n}{2} \geq 32 A_{U,\beta} \geq 32T.$$

Assume now that $k \leq n$. From Lemma 4.13, weakening the numerator to $1 - \beta$, using $1 - \beta \geq 2/3$, and applying $x + y \leq 2 \max\{x, y\}$ gives

$$k\varepsilon \geq \min\left\{\log \frac{h}{\beta}, \log \frac{e^\varepsilon - 1}{\delta}\right\} - \log 3, \tag{43}$$

where the second logarithm is $+\infty$ when $\delta = 0$.

Writing $q = \lfloor (\log_2 U)/2 \rfloor$, we have $h = q + 1$ and $L_U \leq 2q + 3 \leq 3h$. Therefore

$$\log(h/\beta) \geq A_{U,\beta} - 1 - \log 3.$$

When $\delta > 0$, the standing hypotheses imply

$$\frac{e^\varepsilon - 1}{\delta} \geq \varepsilon n^2 = \frac{(\varepsilon n)^2}{\varepsilon} > 1,$$

and consequently

$$\log \frac{e^\varepsilon - 1}{\delta} \geq B_{\varepsilon,\delta} - 1 - \log 2.$$

Substitution in Equation (43) yields $k\varepsilon \geq T - 1 - 2\log 3$. On the other hand, $k\varepsilon \leq 2r + \varepsilon \leq 2r + 1$, so

$$r \geq \frac{T - 2 - 2\log 3}{2}. \tag{44}$$

By Lemma 4.14, $R(M) \geq 1/20$. If $T \geq 2(2 + 2\log 3)$, Equation (44) gives $r \geq T/4$. Otherwise $T < 8.4$ and $R(M) \geq 1/20 > T/200$. Letting $r \downarrow R(M)$ shows $R(M) \geq T/200$. Taking the infimum over all private mechanisms proves the lemma. □

4.8 Public branch selection and conclusion

The final mechanism chooses between the two upper constructions using only public parameters. Thus the branch decision consumes no privacy budget.

BranchClipSum: Instance-optimal bounded summation.
Input: A database $D \in \{0, \dots, U\}^n$; public parameters $U, n, \varepsilon, \delta, \beta$ satisfying the hypotheses of Theorem 1.1.
Output: A real estimate of $S(D)$.
Guarantee: (ε, δ)-privacy and normalized β-error $O(\min\{A_{U,\beta}, B_{\varepsilon,\delta}\})$.
1. Compute the public quantities $A_{U,\beta}$ and $B_{\varepsilon,\delta}$, using $B_{\varepsilon,0} = +\infty$.
2. If $\delta = 0$ or $A_{U,\beta} \le B_{\varepsilon,\delta}$, return PureClipSum$(D)$.
3. Otherwise return ApproxClipSum(D).

Lemma 4.16 (Best-branch privacy, accuracy, and resources)**.** *Under the hypotheses of Theorem 1.1,* BranchClipSum *is (ε, δ)-differentially private and, for every D,*

$$\Pr\left[|\text{BranchClipSum}(D) - S(D)| \le \frac{Cm(D)}{\varepsilon} \min\{A_{U,\beta}, B_{\varepsilon,\delta}\}\right] \ge 1 - \beta.$$

On the pure branch its running time is $O(n \log n + L_U)$; on the approximate branch it is polynomial in n, $\log U$, $\log(1/\delta)$, and $\log(1/\varepsilon)$.

Proof. The branch decision depends only on public parameters. On the pure branch, Lemma 4.4 gives $(\varepsilon, 0)$-privacy and normalized β-error $O(A_{U,\beta})$, and here $A_{U,\beta} = \min\{A_{U,\beta}, B_{\varepsilon,\delta}\}$. On the approximate branch, Lemma 4.12 gives (ε, δ)-privacy and deterministic normalized error $O(B_{\varepsilon,\delta})$, and here $B_{\varepsilon,\delta} < A_{U,\beta}$. Exactly one branch is invoked, so its resource bound is inherited unchanged. □

By Lemma 4.2, the sample-size condition in Theorem 1.1 implies the premise of Lemma 4.15, which proves the left inequality in Equation (2). The right inequality and the claimed output and resource guarantees follow from Lemma 4.16. This completes the proof of Theorem 1.1.

References

[AKMV19] Kareem Amin, Alex Kulesza, Andres Muñoz Medina, and Sergei Vassilvitskii. Bounding user contributions: A bias-variance trade-off in differential privacy. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 263–271. PMLR, 2019.

[BNSV15] Mark Bun, Kobbi Nissim, Uri Stemmer, and Salil P. Vadhan. Differentially private release and learning of threshold functions. In *56th Annual IEEE Symposium on Foundations of Computer Science*, pages 634–649. IEEE Computer Society, 2015.

[DKSS23] Travis Dick, Alex Kulesza, Ziteng Sun, and Ananda Theertha Suresh. Subset-based instance optimality in private estimation. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 7992–8014. PMLR, 2023.

[DLF+24] Wei Dong, Qiyao Luo, Giulia Fanti, Elaine Shi, and Ke Yi. Almost instance-optimal clipping for summation problems in the shuffle model of differential privacy. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security*, pages 1939–1953. Association for Computing Machinery, 2024.

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.

[DSY23] Wei Dong, Dajun Sun, and Ke Yi. Better than composition: How to answer multiple relational queries under differential privacy. *Proceedings of the ACM on Management of Data*, 1(2):123:1–123:26, 2023.

[DY23] Wei Dong and Ke Yi. Universal private estimators. In *Proceedings of the 42nd ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 195–206. Association for Computing Machinery, 2023.

[GDGK20] Quan Geng, Wei Ding, Ruiqi Guo, and Sanjiv Kumar. Tight analysis of privacy and utility tradeoff in approximate differential privacy. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 89–99. PMLR, 2020.

[HLY21] Ziyue Huang, Yuting Liang, and Ke Yi. Instance-optimal mean estimation under differential privacy. In *Advances in Neural Information Processing Systems*, volume 34, pages 25993–26004, 2021.

[LRSS25] Ephraim Linder, Sofya Raskhodnikova, Adam Smith, and Thomas Steinke. Privately evaluating untrusted black-box functions. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 2350–2361. Association for Computing Machinery, 2025.

[LSZ+23] Yuhan Liu, Ananda Theertha Suresh, Wennan Zhu, Peter Kairouz, and Marco Gruteser. Algorithms for bounding contribution for histogram estimation under user-level privacy. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 21969–21996. PMLR, 2023.

[MT07] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science*, pages 94–103. IEEE Computer Society, 2007.