Erciyes University Journal of Institue Of Science and Technology

Research Article
10.65520/erciyesfen.1968098

Imprint:
Volume: 42(3)
Year: 2026
Page: 51-67

Bekir Danış[a*]
Muhammed Hasdemir[b]

[a] Asst. Prof., Aydın Adnan Menderes University, bekir.danis@adu.edu.tr
[b] Asst. Prof., Aydın Adnan Menderes University, muhammed.hasdemir@adu.edu.tr

* Corresponding Author

Received: 6/10/2026
Accepted: 8/27/2026

Citation:
Bekir Danış, Muhammed Hasdemir (2026). An (ε, δ)-Differential Privacy Analysis of Adaptive Noise Calibration via Hockey-Stick Divergence. *Erciyes University Journal of Institue Of Science and Technology*, 42(3), 51-67. https://doi.org/10.65520/erciyesfen.*1968098*

Screened by iThenticate for Authors & Researchers

Except where otherwise noted, content in this article is licensed under a Creative Commons 4.0 International license. Icons by Font Awesome.

An (ε, δ)-Differential Privacy Analysis of Adaptive Noise Calibration via Hockey-Stick Divergence

Abstract

Standard differentially private mechanisms calibrate noise to the worst-case sensitivity of a query, which wastes utility whenever the actual data are well concentrated. This paper studies a two-stage adaptive mechanism that first obtains a privacy-preserving estimate of how spread out the data are, using the interquartile range (IQR) as the measure of dispersion, and then uses that estimate to shrink the noise added to the query answer whenever the data are concentrated. A tunable safety-floor parameter prevents the noise from being reduced below a prescribed minimum, so that the mechanism cannot become arbitrarily leaky even if the private dispersion estimate is itself inaccurate.

The paper makes four contributions, all established with complete proofs rather than heuristic argument. First, we derive the exact (tight) failure probability of the second stage using the hockey-stick divergence framework, and we show that a shortcut formula used in informal treatments of similar mechanisms systematically understates this failure probability and therefore does not constitute a valid privacy guarantee. Second, we prove a new "monotone-pinching" lemma showing that, under a mild regularity condition on how the data are spread across their range, the private IQR estimate obtained via the Exponential Mechanism cannot be off by more than one rank position, which yields an explicit and tight accuracy guarantee for the estimate. Third, we prove an adaptive composition theorem showing that combining the two stages incurs an unavoidable multiplicative penalty in the overall failure probability, and we give a precise, checkable condition under which the adaptive mechanism nonetheless achieves strictly lower estimation variance than the standard (non-adaptive) Laplace mechanism operating at the same total privacy budget. Fourth, we solve the resulting budget-allocation problem in closed form, characterising how the total privacy budget should be split between the two stages to minimise variance, together with an explicit condition under which a prescribed accuracy target is achievable at a given sample size. All theoretical bounds are illustrated numerically, and the accuracy guarantee for the private IQR estimate is validated against Monte Carlo simulation.

Keywords: Differential Privacy, Adaptive Noise Calibration, Hockey-stick Divergence, Exponential Mechanism, Interquartile Range, Utility Optimisation

Hockey-Stick Iraksaması Temelinde Uyarlamalı Gürültü Kalibrasyonunun (ε, δ)-Diferansiyel Gizlilik Açısından Analizi

Öz

Standart diferansiyel gizlilik mekanizmaları, gürültüyü sorgunun en kötü durum duyarlılığına göre kalibre eder; bu da veri yeterince az yayılmışken faydadan gereksiz yere ödün verilmesine yol açar. Bu çalışma, önce verinin ne

kadar yayıldığına dair gizliliği koruyan bir tahmin elde eden — yayılım ölçüsü olarak çeyrekler arası açıklığı (IQR) kullanan — ve ardından bu tahmini, veri yeterince yoğunsa sorgu yanıtına eklenen gürültüyü azaltmak için kullanan iki aşamalı uyarlamalı bir mekanizmayı incelemektedir. Ayarlanabilir bir güvenlik tabanı parametresi, özel yayılım tahmininin kendisi hatalı olsa bile gürültünün belirlenen bir asgari düzeyin altına inmesini engeller ve böylece mekanizmanın keyfi ölçüde gizlilik sızdırmasının önüne geçilir.

Çalışma, tamamı sezgisel değil eksiksiz ispatlarla desteklenen dört katkı sunmaktadır. İlk olarak, ikinci aşamanın tam (sıkı) başarısızlık olasılığı hockey-stick ıraksaması çerçevesi kullanılarak türetilmekte ve benzer mekanizmaların gayri resmi incelemelerinde kullanılan bir kısayol formülünün bu başarısızlık olasılığını sistematik biçimde olduğundan düşük gösterdiği, dolayısıyla geçerli bir gizlilik garantisi teşkil etmediği gösterilmektedir. İkinci olarak, verinin kendi aralığı içinde nasıl dağıldığına ilişkin hafif bir düzenlilik koşulu altında, Üstel Mekanizma aracılığıyla elde edilen özel IQR tahmininin bir sıra konumundan fazla sapamayacağını gösteren yeni bir "monoton sıkıştırma" (monotone-pinching) leması ispatlanmakta ve bu da tahmin için açık ve sıkı bir doğruluk garantisi sağlamaktadır. Üçüncü olarak, iki aşamayı birleştirmenin toplam başarısızlık olasılığında kaçınılmaz bir çarpımsal bedele yol açtığını gösteren bir uyarlamalı bileşim teoremi ispatlanmakta; ayrıca uyarlamalı mekanizmanın, aynı toplam gizlilik bütçesiyle çalışan standart (uyarlamalı olmayan) Laplace mekanizmasına kıyasla kesin biçimde daha düşük tahmin varyansı elde ettiği kesin ve denetlenebilir bir koşul verilmektedir. Dördüncü olarak, ortaya çıkan bütçe tahsisi problemi kapalı form olarak çözülmekte; toplam gizlilik bütçesinin varyansı en aza indirecek şekilde iki aşama arasında nasıl paylaştırılması gerektiği, belirli bir örneklem büyüklüğünde hedeflenen doğruluğun elde edilebilir olduğu açık bir koşulla birlikte karakterize edilmektedir. Tüm teorik sınırlar sayısal olarak örneklendirilmekte ve özel IQR tahminine ilişkin doğruluk garantisi Monte Carlo simülasyonuyla doğrulanmaktadır.

Anahtar kelimeler: Diferansiyel Gizlilik, Uyarlamalı Gürültü Kalibrasyonu, Hockey-stick Iraksaması, Üstel Mekanizma, Çeyrekler Arası Açıklık, Fayda Optimizasyonu

1. INTRODUCTION

Differential Privacy (DP) [3] provides a mathematically rigorous framework for releasing statistical summaries while protecting individual records. The Laplace mechanism [3] calibrates noise to the global sensitivity of the query, a worst-case quantity that leads to unnecessarily high noise when the dataset has low dispersion. Smooth sensitivity [8] was introduced to address this limitation, but enforcing pure ε-DP with local sensitivity typically requires heavy-tailed noise distributions with infinite variance. This motivates the (ε, δ)-DP relaxation, which permits lighter-tailed distributions at the cost of a small failure probability δ. Complementary perspectives on the privacy-utility trade-off are provided by instance-optimal mechanisms [1], Rényi DP [6], concentrated DP [5], the hypothesis-testing interpretation of privacy [2], a recent instance-optimal reformulation of the Laplace mechanism itself [9], Gaussian DP [10], the K-norm gradient mechanism [11], and geometry-aware adaptive Laplace mechanisms defined via directional operator norms on Riemannian manifolds [12].

Practical motivation. The tension between global-sensitivity noise and local data dispersion is not merely theoretical. In official statistics, a national statistical agency releasing income or expenditure summaries for many sub-populations (regions, age brackets, industry sectors) faces sub-populations whose dispersion can differ by orders of magnitude; calibrating every release to a single worst-case sensitivity bound wastes accuracy precisely on the low-dispersion sub-populations where a tighter estimate would be most valuable. In industry telemetry pipelines, where aggregate device or session statistics are published under a fixed privacy budget on a recurring (e.g. daily) basis, the same global-sensitivity bound is reused even when the underlying metric is known, from auxiliary or historical information, to be concentrated on a particular day; an adaptive scheme can recover utility on such days without weakening the guarantee on days when the data are genuinely dispersed. In both settings, the appeal of an adaptive mechanism is that it lets the data themselves — through a privately released summary statistic — decide how much noise is actually necessary, rather than committing in advance to a bound that must hold in the worst case over the entire data domain.

Problem statement. Adaptive noise calibration is, on its own, not new: local-sensitivity and smooth-sensitivity methods [8] already adjust noise to the data, and instance-optimal mechanisms [1]

CC: BY-NC-ND

pursue a related goal from a minimax perspective. What is missing from this literature, and what motivates the present paper, is a fully accountable two-stage design in which (a) the private statistic used to drive the noise calibration is itself an explicit, standard DP mechanism (the Exponential Mechanism) rather than a bespoke smooth-sensitivity construction, (b) the privacy cost of using that statistic to set a parameter of a second mechanism is tracked exactly, through the hockey-stick divergence, rather than bounded by a convenient but possibly invalid shortcut, and (c) the accuracy of the private statistic is controlled by an explicit, non-asymptotic, high-probability bound rather than an order-of-magnitude or asymptotic guarantee. Sections 3–5 show that closing these three gaps is analytically delicate: the naive failure-probability formula for Stage 2 is provably too small (Remark 4.2), the naive accuracy claim for Stage 1 requires a genuinely new combinatorial lemma to make rigorous (Lemma 3.4), and the resulting composed guarantee carries an factor that we show is unavoidable rather than an artifact of a loose proof (Remark 4.4).

The present paper proposes and analyses a two-stage adaptive mechanism that privately estimates the IQR of the dataset and uses this estimate to scale the Laplace noise injected into the primary query. Our analysis yields four main contributions. First, we compute the tight conditional failure probability $\delta(t)$ (Lemma 4.1) by integrating the Hockey-stick divergence over its true optimal set $(-\infty, y^*)$, and we show that restricting the integration to $(-\infty, 0]$ produces only a strict lower bound, which therefore constitutes an invalid privacy certificate (Remark 4.2). Second, under a boundary maximum-gap condition (Assumption 3.2), we establish through the monotone-pinching lemma (Lemma 3.4) that each order statistic $x_{(l)}$ belongs to the interval $[y_{l-1}, y_l]$, implying that $g(y_j) \in \{j, j+1\}$ and hence $d^* \leq 1$, which leads to the tight IQR accuracy bound $\kappa = 2(k^* + 2)h$. Third, we prove that the composed mechanism satisfies $\left(\varepsilon_{\text{total}}, e^{\varepsilon_1} \sup_{D \sim D'} E_{t \sim M_1(D')}[\delta(t)]\right)$-DP (Theorem 4.3), demonstrate that the factor e^{ε_1} is tight (Remark 4.4), and derive a sharp condition under which the adaptive mechanism attains a net variance reduction compared with the standard Laplace baseline (Remark 4.5). Finally, we characterise the optimal budget allocation $(\varepsilon_1^*, \varepsilon_2^*)$ (Proposition 5.2) together with an explicit feasibility condition, including the boundary case $\beta_0 = 1$.

Table 1. Summary of notation. For ease of reference, given the notational density of Sections 3–5, the symbols used throughout the paper are collected below in order of first appearance.

Symbol	**Meaning**	**First defined**
D, D′	neighbouring datasets (differ in one record)	Def. 2.1
ε, δ	privacy budget and failure probability	Def. 2.1
Δf	global sensitivity of the primary query f	Def. 2.2 / Assum. 3.9
Hγ(P‖Q)	hockey-stick divergence at level γ	Def. 2.3
Qp(D), IQR(D)	p-quantile and interquartile range of D	§2.1
ε1, ε2, εtotal	Stage-1, Stage-2, and total privacy budget	§3
τ	reference noise scale Δf/ε2	§3
n, h, Yn	sample size, grid spacing (B−A)/n, output grid	Assum. 3.1
x(i)	order statistics of D	Assum. 3.2
up(D,y), rankD(y)	Exponential-Mechanism utility and data rank	(3.2)
d*	optimal rank discrepancy	§3.1
g(yj)	grid-point rank function	Lemma 3.4
Q̃1, Q̃3, IQR̃	private quartile and IQR estimates	(3.3)
k*	high-probability rank-error budget	(3.6)
κ(n,ε1,β)	IQR accuracy bound	(3.4)
ϕ(t), b(t)	linear-ramp scaling function and adaptive noise scale	(3.8)
c	safety-floor parameter, c∈(0,1]	(3.8)
δ(t)	tight conditional Stage-2 failure probability	Lemma 4.1
δpartial(t)	invalid partial (shortcut) failure probability	Remark 4.2
δtotal	overall failure probability after composition	Thm. 4.3
δ̄(c)	worst-case expected failure probability over	Cor. 4.6

CC: BY-NC-ND

	Stage 1	
ρ*, ρ*(α)	dispersion ratio IQR(D)/τ, and its function of α	Thm. 5.1 / §5.1
α, β0	Stage-1 budget fraction ε1/εtotal, and IQR(D)εtotal/Δf	§5.1
α*	utility-optimal budget fraction	(5.5)

2. PRELIMINARIES

Definition 2.1. A randomised mechanism $M: \mathcal{D} \to \mathcal{Y}$ satisfies (ε, δ)-differential privacy if, for all neighbouring datasets $D \sim D'$ (differing in at most one record) and all measurable $S \subseteq \mathcal{Y}$,

$$\Pr[M(D) \in S] \leq e^{\varepsilon} \Pr[M(D') \in S] + \delta.$$

Definition 2.2. The ℓ_1-global sensitivity of a query $f: \mathcal{D} \to \mathbb{R}$ is

$$\Delta_f := \sup_{D \sim D'} |f(D) - f(D')|.$$

Definition 2.3. For distributions P, Q and $\gamma \geq 1$, the Hockey-stick divergence is

$$H_\gamma(P \| Q) = \sup_{S \in \mathcal{F}} \left(\Pr_{Y \sim P}[Y \in S] - \gamma \Pr_{Y \sim Q}[Y \in S] \right).$$

Lemma 2.4. Let P and Q be σ-finite distributions with densities p and q with respect to a common dominating measure μ. Then

$$H_\gamma(P \| Q) = \int_{\mathcal{Y}} \big(p(y) - \gamma q(y)\big)_+ \, d\mu(y), \qquad (2.1)$$

where $(x)_+ := \max\{0, x\}$. The supremum in Definition 2.3 is attained by $S^* = \{y : p(y) > \gamma q(y)\}$.

Proof. For any measurable S,

$$\Pr_P[S] - \gamma \Pr_Q[S] = \int_S (p - \gamma q) \, d\mu \leq \int (p - \gamma q)_+ \, d\mu,$$

establishing the upper bound. Attainment follows from

$$\Pr_P[S^*] - \gamma \Pr_Q[S^*] = \int_{S^*} (p - \gamma q) \, d\mu = \int (p - \gamma q)_+ \, d\mu,$$

since the integrand is positive on S^* and non-positive elsewhere. Measurability of S^* follows because p and q are μ-measurable densities, so $y \mapsto p(y) - \gamma q(y)$ is measurable and S^* is the preimage of $(0, \infty)$ under this map. On sets where $q(y) = 0$, the condition $p(y) > \gamma q(y)$ reduces to $p(y) > 0$, and the positive-part operation handles this without requiring division by q.

By Lemma 2.4, the minimum δ for which (D, D') satisfies (ε, δ)-DP under M is $H_{e^\varepsilon}\big(M(D) \| M(D')\big)$.

2.1. Quantile Definition and Interquartile Range

For a dataset $D = \{x_1, \ldots, x_n\} \subset [A, B]$ with order statistics $x_{(1)} \leq \cdots \leq x_{(n)}$, define the quantile at level $p \in (0,1)$ by

$$Q_p(D) := x_{(\lceil pn \rceil)}. \qquad (2.2)$$

The interquartile range is $\mathrm{IQR}(D) = Q_{0.75}(D) - Q_{0.25}(D)$.

3. TWO-STAGE MECHANISM

Let $\varepsilon_{\text{total}} = \varepsilon_1 + \varepsilon_2$ be the total privacy budget, and set $\tau := \Delta_f / \varepsilon_2$.

3.1. Stage 1: Private IQR Estimation

Assumption 3.1. Stage 1 uses the output space $\mathcal{Y}_n = \{A + jh : j = 0, \dots, n\}$ with spacing $h = (B - A)/n$.

Assumption 3.2. Setting $x_{(0)} := A$ and $x_{(n+1)} := B$, the dataset satisfies

$$\max_{j=0,1,\dots,n} \left(x_{(j+1)} - x_{(j)}\right) \le h. \tag{3.1}$$

Condition (3.1) bounds every consecutive gap, including boundary gaps, by at most $h = (B - A)/n$. We note that this is a joint condition on the data and the grid size n: as n increases, h decreases and the condition becomes more stringent. For data drawn from the uniform distribution on $[A, B]$, the expected maximum spacing among n i.i.d. points is of order $O(\log n/n)$, which coincides with $h = O(1/n)$ only up to a log$n$ factor; for distributions with density bounded away from zero on $[A, B]$, the same order holds up to a distribution-dependent constant. In practice, Assumption 3.2 is most naturally satisfied when measurements are recorded at precision h—i.e. the data themselves lie on the grid $\mathcal{Y}_n$—in which case every consecutive gap is a non-negative integer multiple of h and the condition reduces to requiring no gap exceeds h. When the condition fails, Remark 3.8 describes the degraded accuracy bound.

Remark 3.3. A minimum-gap condition is insufficient. For example, $n = 10$, $h = 0.1$, and

$$D = \{0, 0.1, 0.2, 0.3, 0.4, 0.6, 0.7, 0.8, 0.9, 1.0\}$$

satisfies $\min_j \left(x_{(j+1)} - x_{(j)}\right) = h$ but contains the interior gap $(0.4, 0.6) = 2h$, so a unit rank error produces a value error of $2h > h$. The boundary gaps $x_{(1)} - A = 0$ and $B - x_{(10)} = 0$ both satisfy the boundary extension of Assumption 3.2, but the interior violation breaks the sandwiching bound of Lemma 3.4.

For each target quantile $p \in \{0.25, 0.75\}$, set $r_p := \lceil pn \rceil$ and define

$$u_p(D, y) := -\left|\mathrm{rank}_D(y) - r_p\right|, \qquad \mathrm{rank}_D(y) := |\{x_i \in D : x_i \le y\}|, \qquad y \in \mathcal{Y}_n. \tag{3.2}$$

Adding or removing one record changes $\mathrm{rank}_D(y)$ by at most 1, so $\Delta_u = 1$.

Define the optimal utility set $\mathrm{OPT}_p(D) := \{y \in \mathcal{Y}_n : \mathrm{rank}_D(y) = r_p\}$ and the optimal rank discrepancy $d^* := \min_{y \in \mathcal{Y}_n} \left|\mathrm{rank}_D(y) - r_p\right|$. Note $d^* = 0$ if and only if $\mathrm{OPT}_p(D) \neq \varnothing$.

Lemma 3.4. Assume $n \ge 4$. Under Assumptions 3.1 and 3.2, write $y_j := A + jh$ for $j = 0, \dots, n$ and $g(y_j) := \mathrm{rank}_D(y_j)$. Then:

(i) $x_{(l)} \in [y_{l-1}, y_l]$ for every $l = 1, \dots, n$.

(ii) For $j = 0, \dots, n - 1$:

$$g(y_j) = j + \mathbf{1}[x_{(j+1)} = y_j] \in \{j, j + 1\}.$$

For $j = n$: $g(y_n) = n$.

(iii) For any target rank $r_p \in \{1, \dots, n - 1\}$, $d^* \in \{0, 1\}$.

Proof. Part (i): Sandwiching.

Upper bound $x_{(l)} \le y_l$ (ascending induction on l). Base case $l = 1$: Assumption 3.2 with $j = 0$ gives $x_{(1)} - A \le h$, so $x_{(1)} \le A + h = y_1$. Inductive step: assume $x_{(l-1)} \le y_{l-1}$. Assumption 3.2 with $j = l - 1$ gives $x_{(l)} - x_{(l-1)} \le h$, hence $x_{(l)} \le x_{(l-1)} + h \le y_{l-1} + h = y_l$.

Lower bound $x_{(l)} \ge y_{l-1}$ (descending induction on l). Base case $l = n$: Assumption 3.2 with $j = n$ gives $B - x_{(n)} \le h$, so $x_{(n)} \ge B - h = y_{n-1}$. Inductive step: assume $x_{(l+1)} \ge y_l$. Assumption 3.2 with $j = l$ gives $x_{(l+1)} - x_{(l)} \le h$, hence $x_{(l)} \ge x_{(l+1)} - h \ge y_l - h = y_{l-1}$.

Part (ii): Grid-rank formula. Fix $j \in \{0, \dots, n - 1\}$. By Part (i):

CC: BY-NC-ND

• For $l \le j$: $x_{(l)} \le y_l \le y_j$, contributing j to $g(y_j)$.

• For $l = j+1$: $x_{(j+1)} \in [y_j, y_{j+1}]$, so $x_{(j+1)} \le y_j$ iff $x_{(j+1)} = y_j$; contributes $\mathbf{1}[x_{(j+1)} = y_j] \in \{0,1\}$.

• For $l \ge j+2$: $x_{(l)} \ge y_{l-1} \ge y_{j+1} > y_j$; contributes 0.

Hence $g(y_j) = j + \mathbf{1}[x_{(j+1)} = y_j] \in \{j, j+1\}$. For $j = n$: all $l \in \{1, \dots, n\}$ satisfy $x_{(l)} \le B = y_n$, so $g(y_n) = n$. (The sentinel $x_{(n+1)} = B$ is not a data point and is not counted in rank_D.)

Part (iii): Discrepancy bound. For $r_p \in \{1, \dots, n-1\}$, both $r_p - 1$ and r_p lie in $\{0, \dots, n-1\}$, so Part (ii) gives:

$$g\left(y_{r_p-1}\right) \in \{r_p - 1, r_p\}, \qquad g\left(y_{r_p}\right) \in \{r_p, r_p + 1\}.$$

Hence

$$d^* \le \min\left(\left|g\left(y_{r_p-1}\right) - r_p\right|, \left|g\left(y_{r_p}\right) - r_p\right|\right) \le 1.$$

Since $d^* \ge 0$ is an integer, $d^* \in \{0,1\}$. The condition $r_p \in \{1, \dots, n-1\}$ holds for $p \in \{0.25, 0.75\}$ whenever $n \ge 4$: $r_{0.25} = \lceil n/4 \rceil \ge 1$ and $r_{0.75} = \lceil 3n/4 \rceil \le n-1$.

Remark 3.5. Part (ii) implies $g(y_{j+1}) - g(y_j) \in \{0,1,2\}$ for $j = 0, \dots, n-2$: a jump of 2 occurs when $x_{(j+1)} > y_j$ and $x_{(j+2)} = y_{j+1}$, and the skipped rank $j+1$ gives $d^* = 1$. Unattained rank values are always isolated, as a direct consequence of Part (i).

The Exponential Mechanism with effective budget $\varepsilon_1/2$ samples

$$\Pr[\tilde{Q}_p = y] \propto \exp\left(\frac{\varepsilon_1}{4} u_p(D, y)\right), \qquad y \in \mathcal{Y}_n, \tag{3.3}$$

where $\varepsilon_1/4 = (\varepsilon_1/2)/(2\Delta_u)$ [7]. After running (3.3) for $p = 0.25$ and $p = 0.75$ to obtain $\tilde{Q}_1$ and $\tilde{Q}_3$, set $\widetilde{\text{IQR}} := \max(0, \tilde{Q}_3 - \tilde{Q}_1)$.

Lemma 3.6. $M_1(D) := \widetilde{\text{IQR}}$ satisfies ε_1-DP.

Proof. Each run of (3.3) with budget $\varepsilon_1/2$ and $\Delta_u = 1$ is $(\varepsilon_1/2)$-DP [7]. Sequential composition [3] yields ε_1-DP for $(\tilde{Q}_1, \tilde{Q}_3)$, and post-processing immunity [4] extends this to $\widetilde{\text{IQR}}$.

Lemma 3.7. Assume $n \ge 4$. Under Assumptions 3.1 and 3.2, fix $\beta \in (0,1)$. With probability at least $1 - \beta$,

$$\left|\widetilde{\text{IQR}} - \text{IQR}(D)\right| \le \kappa(n, \varepsilon_1, \beta) := \frac{8(B-A)}{n\varepsilon_1}\left(\ln n + \ln\frac{4}{\beta}\right) + \frac{6(B-A)}{n}. \tag{3.4}$$

Proof. (a) Rank error bound. By Lemma 3.4(iii), $d^* \in \{0,1\}$. Applying [7, Theorem 6] with $\Delta_u = 1$, budget $\varepsilon_1/2$, $|\mathcal{Y}_n| = n+1$, and optimal score $-d^*$: for integer $k \ge 1$,

$$\Pr\left[\left|\text{rank}_D(\tilde{Q}_p) - r_p\right| \ge d^* + k\right] \le (n+1)\exp\left(-\frac{\varepsilon_1 k}{4}\right). \tag{3.5}$$

Since $n + 1 \le 2n$ for $n \ge 1$, set $2n\exp(-\varepsilon_1 k/4) = \beta/2$ to obtain

$$k_R^* := \frac{4}{\varepsilon_1}\ln\frac{4n}{\beta}. \tag{3.6}$$

Set $k^* := \lceil k_R^* \rceil$. A union bound over both quartile estimates gives: with probability at least $1 - \beta$,

$$\left|\text{rank}_D(\tilde{Q}_p) - r_p\right| \le d^* + k^* \le k^* + 1 \quad \text{for both } p \in \{0.25, 0.75\}. \tag{3.7}$$

CC: BY-NC-ND

(b) Rank-to-value conversion. Let $\ell := \mathrm{rank}_D(\tilde{Q}_p)$ with $|\ell - r_p| \le k^* + 1$.

Case $\ell = 0$: $\mathrm{rank}_D(\tilde{Q}_p) = 0$ means $\tilde{Q}_p < x_{(1)}$. By Lemma 3.4(i), $x_{(1)} \le y_1 = A + h$, so $\tilde{Q}_p < A + h$. Since $\tilde{Q}_p \in \mathcal{Y}_n$ and consecutive grid points are distance h apart, the only grid point in $[A, A + h)$ is $y_0 = A$. Hence $\tilde{Q}_p = A$. Then

$$\left|A - x_{(r_p)}\right| = x_{(r_p)} - A \le r_p h$$

by Lemma 3.4(i) (telescoping $\sum_{i=0}^{r_p-1}(x_{(i+1)} - x_{(i)}) \le r_p h$). Since $r_p = |\ell - r_p| \le k^* + 1$, $r_p h \le (k^* + 1)h < (k^* + 2)h$.

Case $\ell = n$: $\tilde{Q}_p \ge x_{(n)}$, so $\tilde{Q}_p - x_{(r_p)} \le B - x_{(r_p)}$. Telescoping: $B - x_{(r_p)} \le (n - r_p + 1)h$. Since $n - r_p \le k^* + 1$, $(n - r_p + 1)h \le (k^* + 2)h$.

Case $1 \le \ell \le n - 1$: $x_{(\ell)} \le \tilde{Q}_p < x_{(\ell+1)}$.

If $\ell \ge r_p$:

$$\tilde{Q}_p - x_{(r_p)} < x_{(\ell+1)} - x_{(r_p)} = \sum_{j=r_p}^{\ell}(x_{(j+1)} - x_{(j)}) \le (\ell - r_p + 1)h \le (k^* + 2)h.$$

If $\ell < r_p$:

$$x_{(r_p)} - \tilde{Q}_p \le x_{(r_p)} - x_{(\ell)} = \sum_{j=\ell}^{r_p-1}(x_{(j+1)} - x_{(j)}) \le (r_p - \ell)h \le (k^* + 1)h < (k^* + 2)h.$$

In all cases,

$$\left|\tilde{Q}_p - Q_p(D)\right| \le (k^* + 2)h.$$

(c) Propagation to the IQR. Write $A_3 := \tilde{Q}_3 - \tilde{Q}_1$, $B_3 := \mathrm{IQR}(D) \ge 0$. If $A_3 \ge 0$: triangle inequality gives

$$\left|\widetilde{\mathrm{IQR}} - \mathrm{IQR}(D)\right| = |A_3 - B_3| \le \left|\tilde{Q}_3 - Q_{0.75}\right| + \left|\tilde{Q}_1 - Q_{0.25}\right|.$$

If $A_3 < 0$: $\widetilde{\mathrm{IQR}} = 0$, and

$$\left|\widetilde{\mathrm{IQR}} - \mathrm{IQR}(D)\right| = B_3 \le |A_3 - B_3| \le \left|\tilde{Q}_3 - Q_{0.75}\right| + \left|\tilde{Q}_1 - Q_{0.25}\right|,$$

using $|A_3 - B_3| = B_3 - A_3 \ge B_3$ and the triangle inequality. In both cases,

$$\left|\widetilde{\mathrm{IQR}} - \mathrm{IQR}(D)\right| \le \left|\tilde{Q}_3 - Q_{0.75}\right| + \left|\tilde{Q}_1 - Q_{0.25}\right| \le 2(k^* + 2)h.$$

(d) Closed-form bound. Using $k^* \le k_R^* + 1$:

$$2(k^* + 2)h \le 2\left(\frac{4}{\varepsilon_1}\ln\frac{4n}{\beta} + 3\right)h = \frac{8(B - A)}{n\varepsilon_1}\left(\ln n + \ln\frac{4}{\beta}\right) + \frac{6(B - A)}{n},$$

which is (3.4). The formula uses $\ln(4/\beta)$ exactly; no further approximation is made.

CC: BY-NC-ND

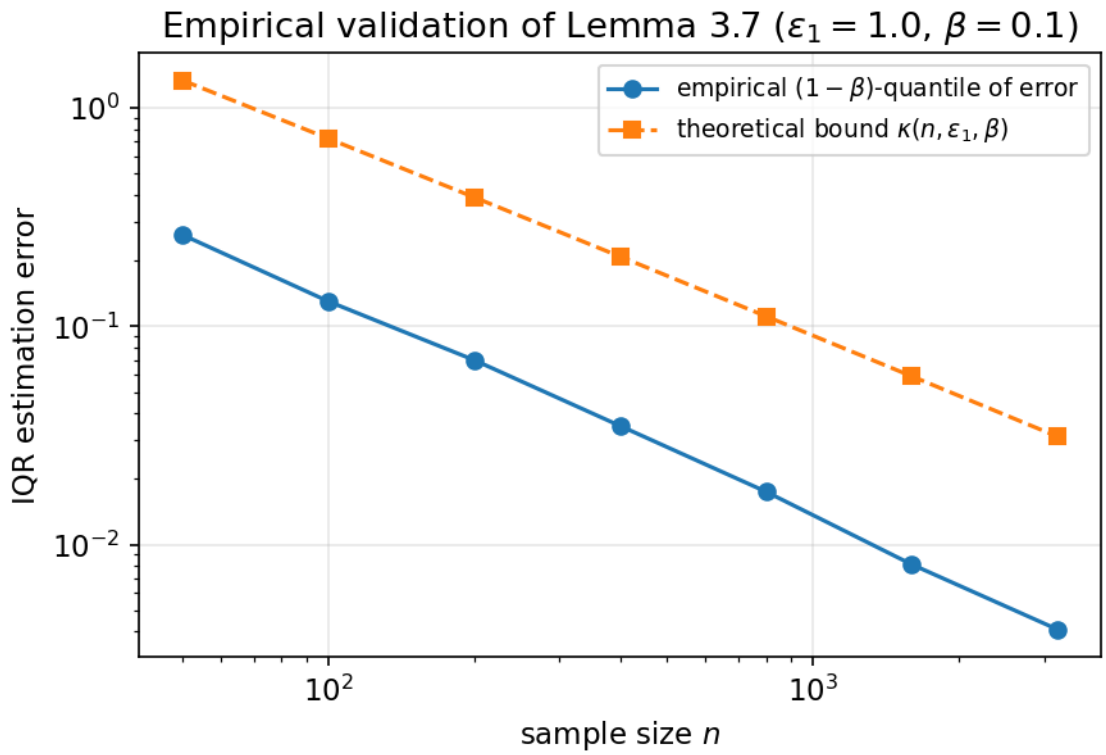

Figure 2. Monte Carlo validation of Lemma 3.7: the empirical (1−β)-quantile of |IQR̃−IQR(D)| over 400 trials per sample size (circles) against the theoretical bound κ(n,ε1,β) of (3.4) (squares), for ε1=1, β=0.1, data drawn uniformly on [0,1] and snapped to the grid Yn so that Assumption 3.2 holds exactly. The empirical error lies strictly below the theoretical bound at every sample size and both decay at the predicted O((log n)/n) rate, confirming that (3.4) is a valid, if conservative, non-asymptotic guarantee.

Remark 3.8. If Assumption 3.2 fails, the sandwiching of Lemma 3.4(i) breaks down and the only universally valid bound is

$$\left|\tilde{Q}_p - x_{(r_p)}\right| \le (k^* + 2)(B - A),$$

giving $\kappa = O\big((\log n)/\varepsilon_1\big)$ independent of n.

3.2. Stage 2: Adaptive Noise Injection

Assumption 3.9. The primary query $f: \mathcal{D} \to \mathbb{R}$ has known finite ℓ_1-global sensitivity Δ_f.

The linear-ramp scaling function $\phi(t) := \min(1, t/\tau)$ is non-decreasing and $(1/\tau)$-Lipschitz. The adaptive noise scale is

$$b(t) := \tau\big(c + (1-c)\phi(t)\big), \qquad c \in (0,1], \tag{3.8}$$

satisfying $c\tau \le b(t) \le \tau$. Stage 2 releases

$$M_2(D \mid \tilde{t}) = f(D) + Z_2, \qquad Z_2 \sim \mathrm{Lap}\big(b(\tilde{t})\big), \qquad \tilde{t} := \widetilde{\mathrm{IQR}}. \tag{3.9}$$

4. PRIVACY ANALYSIS

4.1. Tight Conditional Failure Probability

Lemma 4.1. Under Assumption 3.9, conditioned on $\tilde{t} = t$, $M_2(\cdot \mid t)$ satisfies $\big(\varepsilon_2, \delta(t)\big)$-DP with

CC: BY-NC-ND

$$\delta(t) = \begin{cases} 0, & \text{if } b(t) \geq \tau, \\ 1 - \exp\left(\frac{\varepsilon_2}{2} - \frac{\Delta_f}{2b(t)}\right), & \text{if } b(t) < \tau. \end{cases} \tag{4.1}$$

This equals $\sup_{D\sim D'} H_{e^{\varepsilon_2}}\big(M_2(D \mid t) \| M_2(D' \mid t)\big)$ and is the tightest possible δ.

Proof. Fix $D \sim D'$ with $\Delta := |f(D) - f(D')| \leq \Delta_f$. Set $f(D) = 0$, $f(D') = \Delta$, $b := b(t)$, $\gamma := e^{\varepsilon_2}$. Let

$$p(y) = \frac{1}{2b} e^{-|y|/b}, \qquad q(y) = \frac{1}{2b} e^{-|y-\Delta|/b}.$$

By Lemma 2.4, $H_\gamma(P\|Q) = \int (p - \gamma q)_+ \, dy$.

The log-likelihood ratio $L(y) := (|y - \Delta| - |y|)/b$ equals Δ/b for $y \leq 0$, $(\Delta - 2y)/b$ for $0 < y < \Delta$, and $-\Delta/b$ for $y \geq \Delta$.

Case $b \geq \tau$: For any $\Delta \leq \Delta_f$, $L(y) \leq \Delta/b \leq \Delta_f/b \leq \varepsilon_2$ for all y, so $S^* = \varnothing$ and $H_\gamma = 0$.

Case $b < \tau$: We determine $S^* = \{y : L(y) > \varepsilon_2\}$ for $\Delta \in (0, \Delta_f]$.

If $\Delta \leq \varepsilon_2 b$: $L(y) \leq \Delta/b \leq \varepsilon_2$ for all y, so $S^* = \varnothing$ and $H_\gamma = 0$.

If $\Delta > \varepsilon_2 b$: since $b < \tau = \Delta_f/\varepsilon_2$ implies $\Delta_f > \varepsilon_2 b$, this case applies at $\Delta = \Delta_f$. The condition $L(y) > \varepsilon_2$ holds on $(-\infty, 0]$ (where $L = \Delta/b > \varepsilon_2$) and on $(0, y^*)$ with $y^* := (\Delta - \varepsilon_2 b)/2 > 0$. On $[y^*, \Delta]$, $L \leq \varepsilon_2$; on $[\Delta, \infty)$, $L = -\Delta/b < 0$. Hence $S^* = (-\infty, y^*)$.

Setting $\alpha_0 := \varepsilon_2 - \Delta/b < 0$:

$$I_1 := \int_{-\infty}^{0} (p - \gamma q)\, dy = \frac{1}{2}(1 - e^{\alpha_0}),$$

$$I_2 := \int_{0}^{y^*} (p - \gamma q)\, dy = \frac{1}{2}\left[(1 + e^{\alpha_0}) - \left(e^{-y^*/b} + e^{\alpha_0 + y^*/b}\right)\right].$$

Substituting $y^*/b = -\alpha_0/2$ gives $e^{-y^*/b} = e^{\alpha_0 + y^*/b} = e^{\alpha_0/2}$, so $I_2 = \frac{1}{2}\left(1 - e^{\alpha_0/2}\right)^2 > 0$. With $x := e^{\alpha_0/2} \in (0,1)$:

$$H_\gamma = I_1 + I_2 = \frac{1}{2}(1 - x^2) + \frac{1}{2}(1 - x)^2 = 1 - x = 1 - e^{(\varepsilon_2 - \Delta/b)/2}.$$

The function $\Delta \mapsto H_\gamma$ equals 0 on $[0, \varepsilon_2 b]$ and $1 - e^{(\varepsilon_2 - \Delta/b)/2} > 0$ on $(\varepsilon_2 b, \infty)$, where it is strictly increasing. Since $b < \tau$ implies $\Delta_f > \varepsilon_2 b$, the supremum over $[0, \Delta_f]$ is attained at $\Delta = \Delta_f$, confirming (4.1). Continuity at $b = \tau$:

$$\lim_{b \uparrow \tau} \left(1 - e^{(\varepsilon_2 - \Delta_f/b)/2}\right) = 0.$$

Remark 4.2. Integrating only over $(-\infty, 0]$ yields the partial formula

$$\delta_{\text{partial}}(t) := I_1 = \frac{1}{2}\left(1 - e^{\varepsilon_2 - \Delta_f/b(t)}\right).$$

This formula arises from a common simplification in which the contribution of the interval $(0, y^*)$ to the Hockey-stick integral is dropped—an approximation implicit in informal privacy analyses of mechanisms with sub-sensitivity noise scales. Since $\delta(t) = \delta_{\text{partial}}(t) + I_2$ and

$$I_2 = \frac{1}{2}\left(1 - e^{(\varepsilon_2 - \Delta_f/b)/2}\right)^2 > 0$$

for all $b(t) < \tau$, δ_{partial} is a strict lower bound on the true Hockey-stick divergence and

CC: BY-NC-ND

constitutes an invalid privacy certificate. The gap I_2 grows as $b(t)$ decreases below τ, making the error most severe in the high-utility regime $b(t) \ll \tau$.

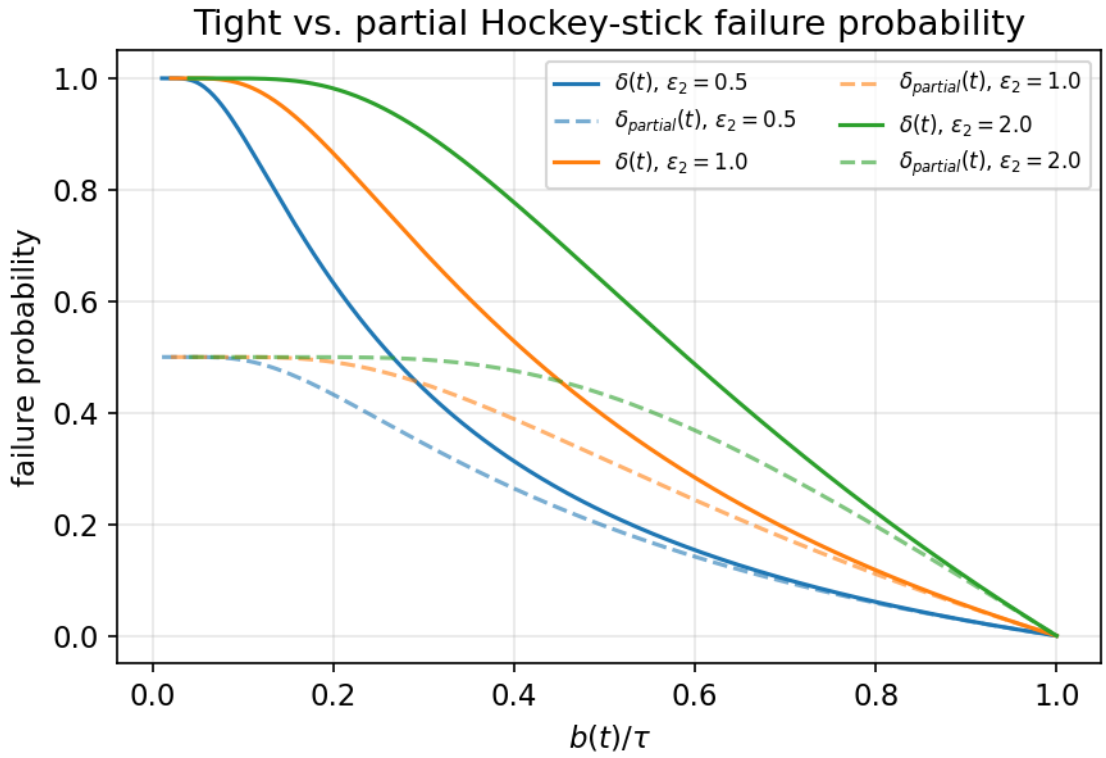

Figure 1. The tight failure probability δ(t) of Lemma 4.1 (solid) against the invalid partial formula δpartial(t) of Remark 4.2 (dashed), for three values of ε2, with Δf=1. The partial formula underestimates the true failure probability throughout the interior of (0,τ), and the gap is largest in the high-utility regime b(t)≪τ, exactly as predicted analytically.

4.2. Adaptive Composition

Theorem 4.3. Let $M(D) := \big(M_1(D), M_2\big(D \mid M_1(D)\big)\big)$. If M_1 is ε_1-DP and $M_2(\cdot \mid t)$ is $\big(\varepsilon_2, \delta(t)\big)$-DP for every t, then M satisfies $(\varepsilon_{\text{total}}, \delta_{\text{total}})$-DP with

$$\varepsilon_{\text{total}} = \varepsilon_1 + \varepsilon_2, \qquad \delta_{\text{total}} = e^{\varepsilon_1} \sup_{D \sim D'} E_{t \sim M_1(D')}[\delta(t)]. \qquad (4.2)$$

Proof. Fix $D \sim D'$ and measurable $S \subseteq \mathcal{Y}_1 \times \mathcal{Y}_2$. Write $S_t := \{y : (t, y) \in S\}$ and let $\mu_D, \mu_{D'}$ denote the distributions of $M_1(D), M_1(D')$ on $\mathcal{Y}_n$. The ε_1-DP of M_1 gives $\mu_D(\{t\}) \le e^{\varepsilon_1} \mu_{D'}(\{t\})$ for all t. Disintegrating and applying $\big(\varepsilon_2, \delta(t)\big)$-DP:

$$\begin{aligned}
\Pr[M(D) \in S] &= \sum_t \Pr\left[M_2(D \mid t) \in S_t\right] \mu_D(\{t\}) \\
&\le \sum_t \left(e^{\varepsilon_2} \Pr[M_2(D' \mid t) \in S_t] + \delta(t)\right) \mu_D(\{t\}) \\
&\le e^{\varepsilon_1 + \varepsilon_2} \sum_t \Pr\left[M_2(D' \mid t) \in S_t\right] \mu_{D'}(\{t\}) + e^{\varepsilon_1} \sum_t \delta(t) \mu_{D'}(\{t\}) \\
&= e^{\varepsilon_{\text{total}}} \Pr[M(D') \in S] + e^{\varepsilon_1} E_{t \sim M_1(D')}[\delta(t)],
\end{aligned}$$

where the third line applies $\mu_D(\{t\}) \le e^{\varepsilon_1} \mu_{D'}(\{t\})$ to both sums. Taking suprema over $D \sim D'$ and S yields (4.2).

CC: BY-NC-ND

Remark 4.4. Let $\mathcal{Y}_1 = \{0,1\}$, $\mu_D(\{0\}) = e^{\varepsilon_1}/(1+e^{\varepsilon_1})$, $\mu_{D'}(\{0\}) = 1/(1+e^{\varepsilon_1})$, $\delta(0) = \delta_0 > 0$, $\delta(1) = 0$. Then

$$\sum_t \delta(t)\mu_D(\{t\}) = e^{\varepsilon_1} E_{t\sim\mu_{D'}}[\delta(t)],$$

so the bound in (4.2) is achieved with equality.

Remark 4.5. The adaptive mechanism satisfies $(\varepsilon_{\text{total}}, \delta_{\text{total}})$-DP with $\delta_{\text{total}} > 0$, whereas the standard Laplace mechanism at the same total budget satisfies $(\varepsilon_{\text{total}}, 0)$-DP with noise scale $b_{\text{std}} = \Delta_f/\varepsilon_{\text{total}}$ and variance

$$\sigma_{\text{std}}^2 = 2\left(\frac{\Delta_f}{\varepsilon_{\text{total}}}\right)^2.$$

The adaptive mechanism is therefore always strictly weaker from the pure-DP perspective; its utility case rests entirely on variance reduction.

In the negligible-error regime ($\kappa \to 0$), the conditional variance is

$$2b(\tilde{t})^2 \approx 2\tau(\alpha)^2\big(c + (1-c)\rho^*(\alpha)\big)^2.$$

A net variance reduction below σ_{std}^2 requires

$$\tau(\alpha)\big(c + (1-c)\rho^*(\alpha)\big) < \frac{\Delta_f}{\varepsilon_{\text{total}}}. \qquad (4.3)$$

Since $\tau(\alpha) = \Delta_f/\big((1-\alpha)\varepsilon_{\text{total}}\big)$ and $\Delta_f/\varepsilon_{\text{total}} = \tau(\alpha)(1-\alpha)$, dividing both sides of (4.3) by $\tau(\alpha) > 0$ yields the equivalent condition

$$c + (1-c)\rho^*(\alpha) < 1 - \alpha, \qquad (4.4)$$

that is, the effective noise-scale reduction factor must be strictly less than the Stage-2 budget fraction. Combined with the variance reduction condition $\rho^* \le 1 - \kappa/\tau$ from Theorem 5.1, the full regime for a net utility gain over the standard Laplace baseline is

$$\rho^*(\alpha) < 1 - \frac{\kappa}{\tau} \quad \text{and} \quad c + (1-c)\rho^*(\alpha) < 1 - \alpha. \qquad (4.5)$$

When both conditions in (4.5) hold, the adaptive mechanism simultaneously achieves lower variance than the standard Laplace baseline at the cost of a positive δ_{total}.

CC: BY-NC-ND

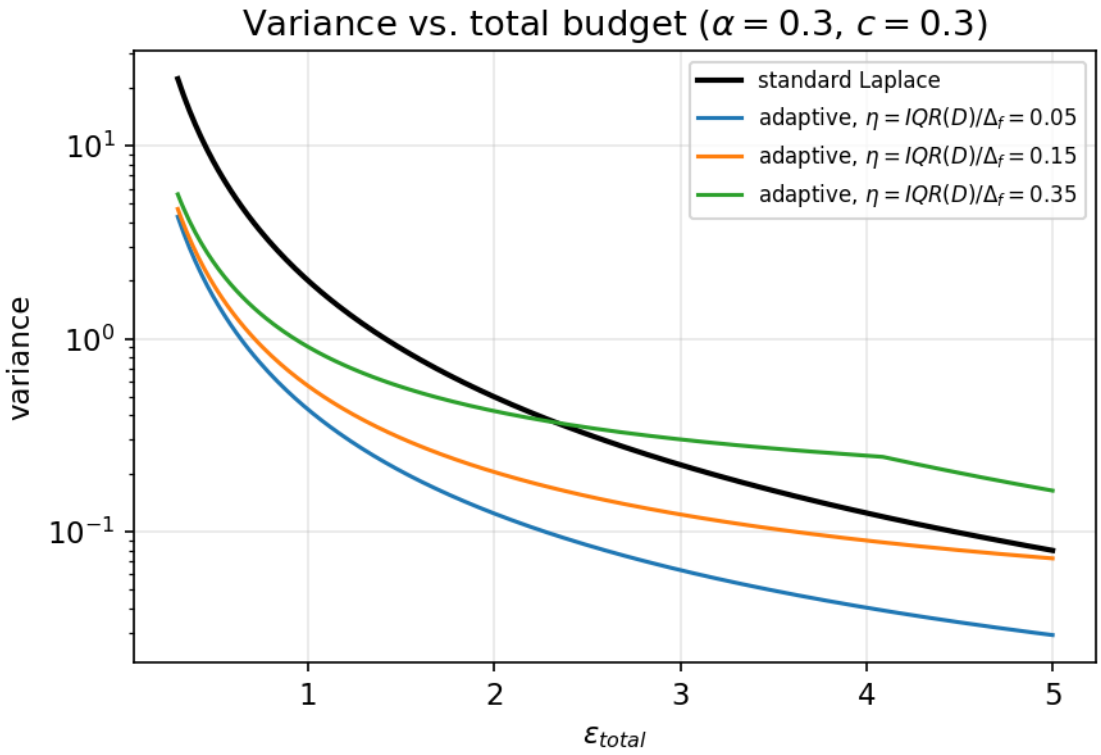

Figure 3. Estimation variance of the standard Laplace mechanism (black) against the adaptive mechanism (colored) as a function of the total privacy budget εtotal, for a fixed dataset with η:=IQR(D)/Δf ∈ {0.05, 0.15, 0.35}, with α=0.3, c=0.3 (log scale on the variance axis). Here ρ*(α)=min(1,(1−α)ηεtotal) is computed from (5.4) rather than held fixed, so it grows with the budget as prescribed by the model. For small η (concentrated data) the adaptive mechanism keeps a large advantage across the whole range; for larger η the advantage shrinks and eventually reverses once ρ*(α) saturates at 1, exactly the breakdown regime flagged by condition (4.4).

4.3. Safety-Floor Parameter and δ_{total}

Corollary 4.6. Define $\bar{\delta}(c) := \sup_{D\sim D'} E_{t\sim M_1(D')}[\delta(t)]$, so $\delta_{\text{total}} = e^{\varepsilon_1}\bar{\delta}(c)$. Then:

(i) $\bar{\delta}(c)$ is continuous and strictly decreasing on (0,1], with $\bar{\delta}(1) = 0$.

(ii) For every $c \in (0,1)$ and every finite $n \geq 1$:

$$\bar{\delta}(c) < \delta(0) = 1 - e^{(\varepsilon_2/2)(1-1/c)}.$$

When $c = 1$: $\bar{\delta}(1) = \delta(0) = 0$.

(iii) A sufficient condition for $\delta_{\text{total}} \leq \bar{\delta}_0$ is

$$c \geq c^* := \frac{\varepsilon_2}{\varepsilon_2 - 2\ln\left(1 - \bar{\delta}_0 e^{-\varepsilon_1}\right)}, \qquad (4.6)$$

valid whenever $\bar{\delta}_0 e^{-\varepsilon_1} < 1 - e^{-\varepsilon_2/2}$.

Proof. (i): Fix $c' < c$. Then $b'(t) < b(t)$ whenever $\phi(t) < 1$ (i.e. $\widetilde{\text{IQR}} < \tau$). The Exponential Mechanism assigns strictly positive probability to every $y \in \mathcal{Y}_n$; in particular $\Pr[\tilde{Q}_3 = \tilde{Q}_1 = A] > 0$, hence $\Pr[\widetilde{\text{IQR}} = 0] > 0 < \tau$, so $\Pr[\widetilde{\text{IQR}} < \tau] > 0$ under μ_D, for every D'. Since $\partial\delta/\partial b < 0$, the event $\{b'(t) < b(t)\}$ has positive μ_D,-measure and contributes strictly, giving $\bar{\delta}(c') > \bar{\delta}(c)$. Continuity follows from continuity of $\delta(t)$ in $b(t)$ and dominated convergence. When $c = 1$: $b(t) = \tau$ for all t, so

CC: BY-NC-ND

$\delta(t) = 0$ by Lemma 4.1, hence $\bar{\delta}(1) = 0$.

(ii): For $c \in (0,1)$: $b(t) \geq c\tau = b(0)$ implies $\delta(t) \leq \delta(0)$, so $\bar{\delta}(c) \leq \delta(0)$. For the strict inequality: the Exponential Mechanism assigns positive weight to every $y \in \mathcal{Y}_n$, so $\Pr[\widetilde{\mathrm{IQR}} > 0] > 0$. On $\{\widetilde{\mathrm{IQR}} > 0\}$, $\phi(\tilde{t}) > 0$ and $c < 1$, so $b(\tilde{t}) > c\tau = b(0)$, hence $\delta(\tilde{t}) < \delta(0)$. Therefore $E[\delta(\tilde{t})] < \delta(0)$ for every D', giving the strict inequality for all finite $n \geq 1$. When $c = 1$: $\bar{\delta}(1) = 0$ and $\delta(0)|_{c=1} = 1 - e^{(\varepsilon_2/2)(1-1)} = 0$, so $\bar{\delta}(1) = \delta(0) = 0$.

(iii): The denominator of (4.6) is positive iff $\bar{\delta}_0 e^{-\varepsilon_1} < 1 - e^{-\varepsilon_2/2}$, the stated validity condition. Under this condition $1 - \bar{\delta}_0 e^{-\varepsilon_1} \in (e^{-\varepsilon_2/2}, 1)$, so the logarithm is well-defined (argument in (0,1), hence negative), the denominator is positive, and c^* is real and finite. Moreover $c^* > 0$ (numerator $\varepsilon_2 > 0$ and denominator positive) and $c^* < 1$ (equivalent to $\ln(1 - \bar{\delta}_0 e^{-\varepsilon_1}) < 0$, which holds since $\bar{\delta}_0 > 0$). Rearranging

$$e^{\varepsilon_1}\left(1 - e^{(\varepsilon_2/2)(1-1/c)}\right) \leq \bar{\delta}_0$$

gives $c \geq c^*$.

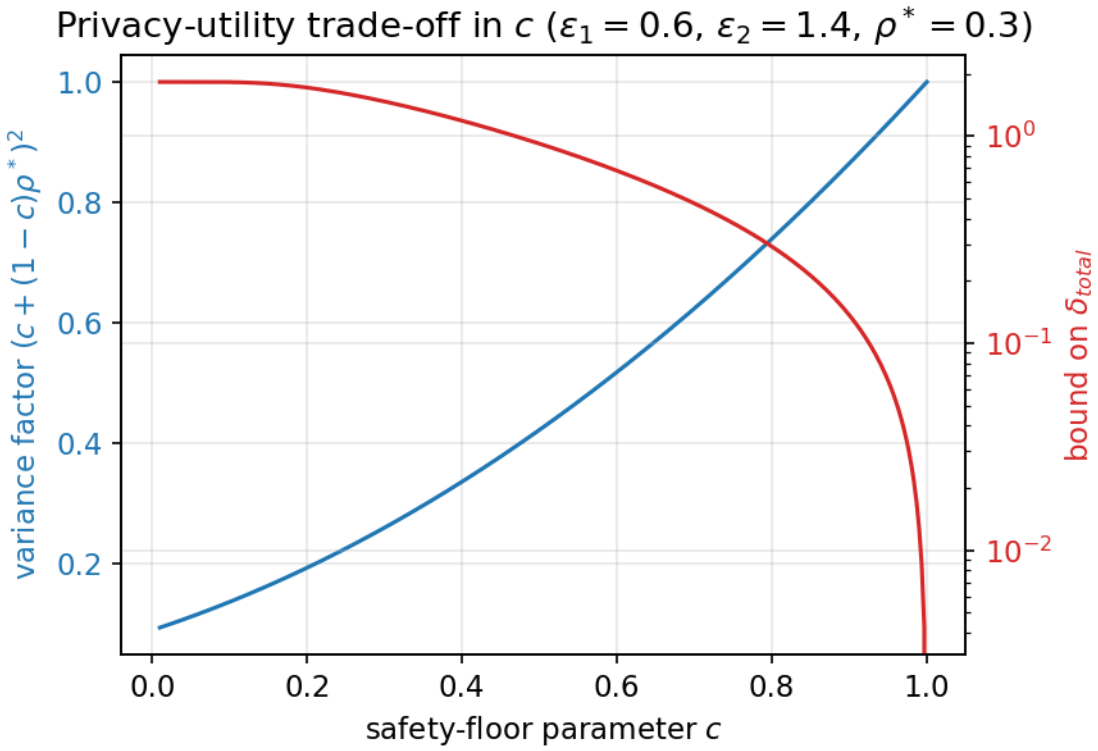

Figure 4. Effect of the safety-floor parameter c on the variance factor (c+(1−c)ρ*)² (blue, left axis) and on the bound e^ε1 δ̄(c) on δtotal (red, right axis, log scale), for ε1=0.6, ε2=1.4, ρ*=0.3. Increasing c monotonically reduces δtotal (Corollary 4.6(i)) at the cost of a monotonically increasing variance factor, illustrating the privacy-utility trade-off governed by c.

5. UTILITY ANALYSIS

Theorem 5.1. Assume $n \geq 4$. Let $\sigma^2_{\mathrm{Lap}} := 2\tau^2$, $\kappa := \kappa(n, \varepsilon_1, \beta)$ as in (3.4), and $\rho^* := \min(1, \mathrm{IQR}(D)/\tau)$. Under Assumptions 3.1 and 3.2, with probability at least $1 - \beta$:

$$2b(\tilde{t})^2 \leq 2\tau^2\left(c + (1-c)\min\left(1, \rho^* + \frac{\kappa}{\tau}\right)\right)^2 =: \sigma^2_{\mathrm{up}}, \qquad (5.1)$$

CC: BY-NC-ND

$$2b(\tilde{t})^2 \geq 2\tau^2\left(c + (1-c)\max\left(0, \rho^* - \frac{\kappa}{\tau}\right)\right)^2 =: \sigma_{\text{lo}}^2. \qquad (5.2)$$

If $\rho^* < 1 - \kappa/\tau$, then $\sigma_{\text{up}}^2 < \sigma_{\text{Lap}}^2 = 2\tau^2$, establishing a strict variance reduction relative to Stage 2 at full scale τ. Moreover,

$$2b(\tilde{t})^2 \xrightarrow{p} \sigma_{\text{Lap}}^2(c + (1-c)\rho^*)^2 < \sigma_{\text{Lap}}^2 \quad (n \to \infty, \rho^* < 1), \qquad (5.3)$$

at rate $\kappa/\tau = O\big((\log n)/(n\varepsilon_1)\big)$.

Proof. $\text{Var}(M_2 \mid \tilde{t} = t) = 2b(t)^2$ and ϕ is $(1/\tau)$-Lipschitz, so Lemma 3.7 (applicable since $n \geq 4$) gives $|\tilde{t} - \text{IQR}(D)| \leq \kappa$ with probability $\geq 1 - \beta$, implying $\phi(\tilde{t}) \leq \min(1, \rho^* + \kappa/\tau)$ and $\phi(\tilde{t}) \geq \max(0, \rho^* - \kappa/\tau)$. Bounds (5.1)-(5.2) follow from (3.8). For (5.3): $\kappa = O\big((\log n)/n\big) \to 0$, so for any $\varepsilon_c > 0$ and $\beta \in (0,1)$ there exists N_0 such that $\kappa < \varepsilon_c\tau$ for all $n \geq N_0$, giving $\Pr[|\phi(\tilde{t}) - \rho^*| \geq \varepsilon_c] \leq \beta$. Since β is arbitrary, $\phi(\tilde{t}) \xrightarrow{p} \rho^*$, and the continuous mapping theorem yields (5.3).

5.1. Optimal Budget Allocation

Let $E := \varepsilon_{\text{total}}$, $\varepsilon_1 = \alpha E$, $\varepsilon_2 = (1-\alpha)E$, $\tau(\alpha) := \Delta_f/\big((1-\alpha)E\big)$, $\beta_0 := \text{IQR}(D)E/\Delta_f$, and

$$\rho^*(\alpha) := \min(1, (1-\alpha)\beta_0). \qquad (5.4)$$

Proposition 5.2. Assume $n \geq 4$. In the negligible-error regime ($\kappa \to 0$), the variance $V(\alpha) := 2b(\alpha)^2$ with $b(\alpha) := \tau(\alpha)\big(c + (1-c)\rho^*(\alpha)\big)$ is strictly increasing in α in both regimes:

Case A ($\beta_0 > 1,\ \alpha < 1 - 1/\beta_0$): $\rho^*(\alpha) = 1,\ b(\alpha) = \tau(\alpha),\ V(\alpha) = 2\Delta_f^2/((1-\alpha)^2E^2)$; no IQR-adaptive gain.

Case B ($\big((1-\alpha)\beta_0 \leq 1\big)$, which covers all $\alpha > 0$ when $\beta_0 \leq 1$, all $\alpha \geq 1 - 1/\beta_0$ when $\beta_0 > 1$, and all $\alpha > 0$ when $\beta_0 = 1$):

$$b(\alpha) = \frac{c\Delta_f}{(1-\alpha)E} + (1-c)\text{IQR}(D) > 0.$$

In both cases $dV/d\alpha > 0$, so the utility-optimal allocation minimises α. Setting $\kappa(n, \alpha E, \beta) \leq \eta_0$ and solving from (3.4):

$$\alpha^* = \frac{8(B-A)\left(\ln n + \ln\frac{4}{\beta}\right)}{\big(n\eta_0 - 6(B-A)\big)E}, \qquad (5.5)$$

subject to the feasibility condition

$$n\eta_0 > 6(B-A). \qquad (5.6)$$

If (5.6) fails, the target accuracy η_0 is unachievable at sample size n.

Proof. In Case B: $\tau(\alpha)(1-\alpha)\beta_0 = \Delta_f\beta_0/E = \text{IQR}(D)$ gives the stated $b(\alpha)$. Case A: $dV/d\alpha = 4\Delta_f^2/((1-\alpha)^3E^2) > 0$. Case B:

$$\frac{dV}{d\alpha} = 4b(\alpha) \cdot \frac{c\Delta_f}{(1-\alpha)^2E} > 0.$$

Solving $\kappa(n, \alpha E, \beta) \leq \eta_0$ from (3.4) and rearranging yields (5.5), with the positivity of the denominator being equivalent to (5.6).

CC: BY-NC-ND

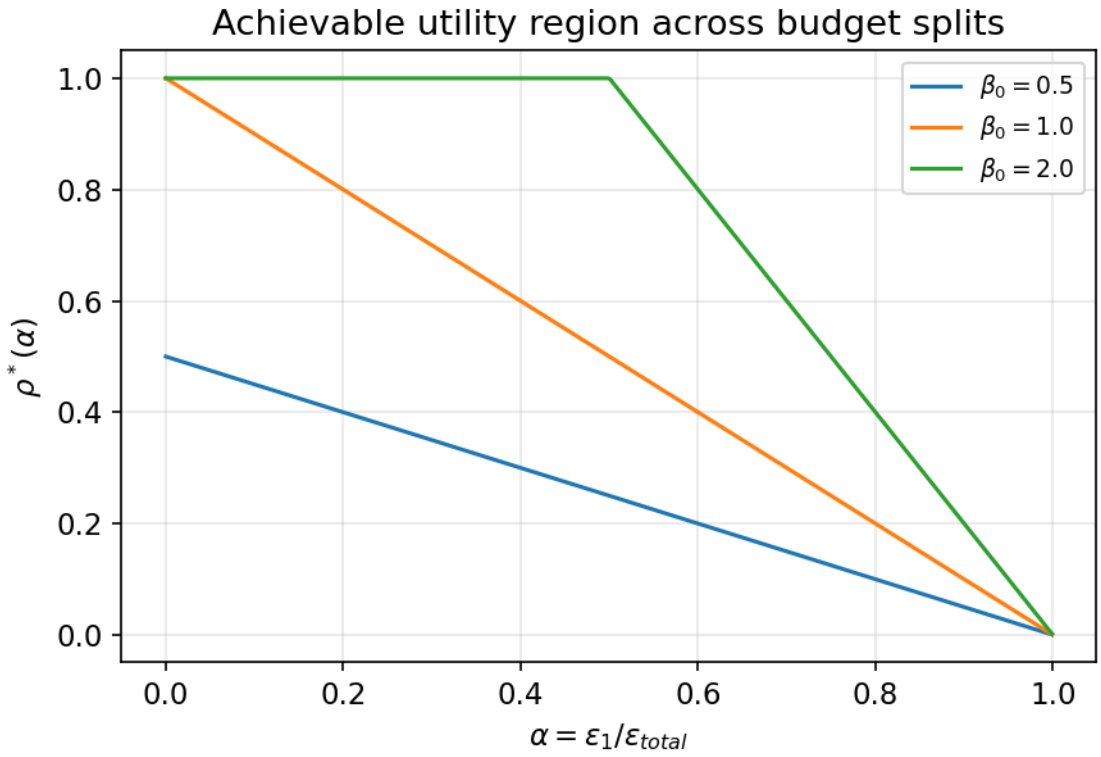

Figure 5. The achievable dispersion ratio ρ*(α)=min(1,(1−α)β0) of (5.4) as a function of the Stage-1 budget fraction α, for three values of β0=IQR(D)εtotal/Δf. Larger β0 (data concentrated relative to sensitivity and total budget) shifts the achievable-utility region upward and enlarges the range of α over which ρ*(α)=1 (Case A of Proposition 5.2, no further IQR-adaptive gain); smaller β0 confines the mechanism to Case B throughout, where variance is strictly increasing in α and the optimum is pushed toward α→0.

6. CONCLUSION

We have presented a rigorous analysis of a two-stage adaptive (ε, δ)-DP mechanism with four main contributions.

The tight conditional failure probability $\delta(t) = 1 - e^{\varepsilon_2/2 - \Delta_f/(2b(t))}$ is derived by integrating the Hockey-stick divergence over the full optimal set $(-\infty, y^*)$; restricting to $(-\infty, 0]$ drops the positive term I_2 and yields an invalid privacy certificate.

The monotone-pinching lemma establishes $x_{(l)} \in [y_{l-1}, y_l]$ under the boundary maximum-gap condition via induction, implying $g(y_j) \in \{j, j+1\}$ and $d^* \leq 1$, and yielding the IQR accuracy bound $\kappa = 2(k^* + 2)h$ with $k_R^* = \frac{4}{\varepsilon_1} \ln \frac{4n}{\beta}$.

The composition factor e^{ε_1} is tight in the abstract mechanism framework. The adaptive mechanism always provides strictly weaker pure-DP guarantees than the standard Laplace mechanism at the same budget; it achieves a net variance reduction over the standard Laplace baseline precisely when both conditions of (4.5) hold, namely $\rho^*(\alpha) < 1 - \kappa/\tau$ and $c + (1-c)\rho^*(\alpha) < 1 - \alpha$.

The utility-optimal budget allocation satisfies (5.5) subject to the feasibility condition $n\eta_0 > 6(B - A)$.

Several directions extend naturally from this analysis. First, the dispersion statistic used to drive the adaptive noise scale need not be the IQR: the same two-stage template — a privately released, high-probability-accurate summary statistic feeding a Lipschitz scaling function — applies to any robust

dispersion measure (e.g. the median absolute deviation), and it would be of interest to determine whether an analogue of the monotone-pinching lemma (Lemma 3.4) holds, and with what constants, for such alternatives. Second, the present composition analysis (Theorem 4.3) treats the two stages as a fixed, one-shot pipeline; extending the accountable-composition argument to a sequence of adaptively and repeatedly queried releases, in the spirit of the privacy-loss-accounting literature for iterative mechanisms, would broaden the applicability of the tight failure-probability formula of Lemma 4.1 beyond the two-stage setting studied here. Third, the safety-floor mechanism of Section 4.3 fixes c in advance; an adaptive or data-driven choice of c — subject to its own privacy accounting — could potentially improve on the variance/failure-probability trade-off of Corollary 4.6 without sacrificing the worst-case guarantee. Finally, while every bound in this paper is non-asymptotic and is validated numerically in Figures 1–5, an empirical evaluation on real-world datasets with heterogeneous, application-specific sensitivity structures (e.g. the official-statistics and telemetry settings discussed in the Introduction) would complement the theoretical contributions and is left to future work.

Peer-review: External, Independent.

Acknowledgements:
The authors would like to express their sincere thanks to the editor and the anonymous reviewers for their helpful comments and suggestions.

Declarations:

1. Statement of Originality:
This manuscript is the original work of the authors, has not been published previously, and is not under consideration for publication elsewhere, either in whole or in part.

2. Author Contributions:
Concept: BD,MH; ***Conceptualization:*** BD,MH; ***Literature Search:*** BD,MH; ***Data Collection:*** BD,MH; ***Data Processing:*** BD,MH; ***Analysis:*** BD,MH; ***Writing – original draft:*** BD,MH; ***Writing – review & editing:*** BD,MH.

3. Ethics approval:
Not applicable. This study is theoretical in nature and does not involve human participants, human data, or animal subjects.

4. Funding/Support:
This work has not received any funding or support.

5. Competing Interests:
The authors declare no competing interests.

6. GenAI Usage Statement:
No GenAI tools were used at any stage of the study.

7. Sustainable Development Goals:

☐ 1 NO POVERTY ☐ 2 ZERO HUNGER ☐ 3 GOOD HEALTH AND WELL-BEING ☐ 4 QUALITY EDUCATION ☐ 5 GENDER EQUALITY ☐ 6 CLEAN WATER AND SANITATION ☐ 7 AFFORDABLE AND CLEAN ENERGY
☐ 8 DECENT WORK AND ECONOMIC GROWTH ☒ 9 INDUSTRY, INNOVATION AND INFRASTRUCTURE ☐ 10 REDUCED INEQUALITIES ☐ 11 SUSTAINABLE CITIES AND COMMUNITIES ☐ 12 RESPONSIBLE CONSUMPTION AND PRODUCTION ☐ 13 CLIMATE ACTION ☐ 14 LIFE BELOW WATER
☐ 15 LIFE ON LAND ☐ 16 PEACE AND JUSTICE STRONG INSTITUTIONS ☐ 17 PARTNERSHIPS FOR THE GOALS THE GLOBAL GOALS For Sustainable Development

REFERENCES

[1] H. Asi and J. Duchi, Instance-optimality in differential privacy via approximate inverse sensitivity mechanisms, Advances in Neural Information Processing Systems, 33 (2020), 14106–14117.

[2] B. Balle, G. Barthe, M. Gaboardi, J. Hsu and T. Sato, Hypothesis testing interpretations and Rényi differential privacy, Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, Proc. Mach. Learn. Res., 108 (2020), 2496–2506.

[3] C. Dwork, F. McSherry, K. Nissim and A. Smith, Calibrating noise to sensitivity in private data analysis, Theory of Cryptography, 265–284, 2006.

[4] C. Dwork and A. Roth, The algorithmic foundations of differential privacy, Found. Trends Theor. Comput. Sci., 9(3–4) (2014), 211–407.

[5] C. Dwork and G. N. Rothblum, Concentrated differential privacy, arXiv preprint arXiv:1603.01887, 2016.

[6] I. Mironov, Rényi differential privacy, Proceedings of the 30th IEEE Computer Security Foundations Symposium, Santa Barbara, USA, 263–275, 2017.

[7] F. McSherry and K. Talwar, Mechanism design via differential privacy, Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science, Providence, USA, 94–103, 2007.

[8] K. Nissim, S. Raskhodnikova and A. Smith, Smooth sensitivity and sampling in private data analysis, Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, USA, 75–84, 2007.

[9] D. Durfee, Unifying Laplace mechanism with instance optimality in differential privacy, arXiv preprint arXiv:2505.02798, 2025.

[10] J. Dong, A. Roth and W. J. Su, Gaussian differential privacy, J. R. Stat. Soc. B, 84 (2022), 3–37.

[11] M. Reimherr and J. Awan, KNG: The K-norm gradient mechanism, Advances in Neural Information Processing Systems, 32 (2019).

[12] B. Danış, Operator normed adaptive Laplace mechanism: A manifold-based approach for differential privacy, Electronic Research Archive, 34(9) (2026), 6043–6071.

CC: BY-NC-ND