To Count or Not to Count: Practical DP Mean Estimation with Unknown Dataset Size

Marcel Neunhoeffer
Institute for Employment Research & Ludwig-Maximilians-Universität Munich
marcel.neunhoeffer@iab.de

Shlomi Hod
Weizenbaum Institute

Jörg Drechsler
Institute for Employment Research & Ludwig-Maximilians-Universität Munich

Abstract

Differentially private mean estimation with unknown dataset size n requires a choice: should the analyst spend part of the privacy budget estimating n, or avoid counting entirely? Several mechanisms exist for this problem, from OpenDP's resize mean [OpenDP, 2025] to the minimax-optimal 2D hourglass [Kulesza et al., 2024] and simplex augmentation [Fitzsimons et al., 2025], but no systematic comparison characterizes when each is preferred. We provide such a comparison. We derive closed-form MSE expressions for the presented strategies spanning a spectrum from *no counting* (fixed denominator) through *implicit counting* (simplex augmentation) to *explicit counting* (noisy count with budget splitting), and empirically evaluate them via simulation. Our main findings: (1) centering the sum at $(L+U)/2$, a step absent from current DP libraries, reduces variance by a factor up to 4; (2) when dataset size uncertainty is small, avoiding the count entirely is optimal; (3) when uncertainty is large, the simplex mechanism [Fitzsimons et al., 2025] dominates, achieving minimax-optimal worst-case MSE while adapting to the true mean, using only standard Laplace noise. These findings yield a practical two-level decision: *first decide whether to count, then decide how.*

1 Introduction

Differential privacy (DP) has become the standard for privacy-preserving data analysis [Dwork et al., 2006]. Mean estimation, i.e., releasing an estimate of $\bar{x} = \frac{1}{n}\sum_i x_i$ for $x_i \in [L, U]$, is a fundamental primitive. In the **add-remove model**, where neighboring datasets differ by one record, n is itself sensitive. This creates a tension: noise calibration depends on n, yet n is unknown.

Several mechanisms address this problem. Kulesza et al. [2024] (KSW) achieve the minimax-optimal MSE $= 2(U-L)^2/(n^2\varepsilon^2)$ uniformly via a 2D hourglass mechanism. Fitzsimons et al. [2025] show that simplex augmentation yields a "free" count estimate with no budget splitting. DP libraries such as OpenDP [OpenDP, 2025] offer resize mean pipelines based on the standard Laplace mechanism.

However, no prior work systematically compares these approaches or characterizes when each is preferred. The analyst faces two practical questions: *should I estimate n at all?* And if so, *how?*

Contributions. We provide an empirical and analytical comparison of several strategies for DP mean estimation with unknown n. (1) We observe that *centering* the sum at $c = (L+U)/2$ reduces sensitivity from $\max\{|U|, |L|\}$ to $(U-L)/2$ (Proposition 1), a prerequisite absent from current DP libraries, implying a factor 4 variance increase for non-negative data relative to the centered approach. (2) We derive closed-form MSE expressions for strategies spanning a spectrum from no counting to explicit counting, and characterize when each is preferred (Section 3). (3) We analyze the simplex mechanism [Fitzsimons et al., 2025] as a third paradigm, *implicit counting*, deriving its instance-dependent MSE and showing it Pareto-dominates the

minimax-optimal KSW mechanism (Section 3.3). (4) We demonstrate all results via simulation (Figure 1) and provide practical recommendations for DP mean estimation (Section 4).

2 Background and Centering

Setup. We observe $X = \{x_1, \ldots, x_n\}$ with $x_i \in [L, U]$ and $n \in [n_{\min}, n_{\max}]$. Under ε-DP in the add-remove model, we wish to estimate $\bar{x}$. All mechanisms use the Laplace mechanism: for a query f with global sensitivity Δ, release $f(X) + \mathrm{Lap}(\Delta/\varepsilon)$.
Current practice. OpenDP offers two pipelines for unknown n: (a) *resize mean*, padding/subsampling to fixed size m with Laplace noise (sensitivity $(U-L)/m$); and (b) *sum/count*, releasing noisy sum (sensitivity $\max\{|U|, |L|\}$) and noisy count separately. Neither approach centers the data.

Proposition 1 (Centering Reduces Sensitivity)**.** *Let $c = (L+U)/2$ and define the centered sum $S_c = \sum_{i=1}^{n}(x_i - c)$. In the add-remove model, the global sensitivity of S_c is $\Delta_c = (U-L)/2$, compared to $\Delta_S = \max\{|U|, |L|\}$ for the uncentered sum $S = \sum_i x_i$. The variance ratio of the Laplace mechanisms for these two sums is:*

$$\frac{\mathrm{Var}(\mathrm{Lap}(\Delta_S/\varepsilon))}{\mathrm{Var}(\mathrm{Lap}(\Delta_c/\varepsilon))} = \frac{4\max\{|U|, |L|\}^2}{(U-L)^2} \geq 1 \tag{1}$$

with equality when $L = -U$ (symmetric range) and ratio $= 4$ when $L = 0$.

Proof. Adding a record $x \in [L, U]$ changes S_c by $|x - c| \leq (U-L)/2$ (since c is the midpoint). Removing is symmetric. The variance ratio follows from $\mathrm{Var}(\mathrm{Lap}(b)) = 2b^2$. For non-negative data ($L = 0$), $\Delta_S = U$ vs. $\Delta_c = U/2$, a variance reduction by a factor of 4. For symmetric data ($L = -U$), the gains vanish. □

3 Estimation Strategies

We analyze strategies along a spectrum from explicit counting to no counting. All use centering at $c = (L+U)/2$ and $\Delta_c = (U-L)/2$. We compare these settings to OpenDP's resize mean in our simulations and in Appendix C.

3.1 Explicit Counting (Approach A)

Spend ε_1 on estimating n and $\varepsilon_2 = \varepsilon - \varepsilon_1$ on the centered sum:

$$\hat{n} = n + \mathrm{Lap}(1/\varepsilon_1), \qquad \mathrm{Var}(\hat{n}) = 2/\varepsilon_1^2 \tag{2}$$

$$\hat{S}_c = S_c + \mathrm{Lap}(\Delta_c/\varepsilon_2), \qquad \mathrm{Var}(\hat{S}_c) = (U-L)^2/(2\varepsilon_2^2) \tag{3}$$

The mean estimate is $\hat{\mu}_A = \hat{S}_c/\hat{n} + c$, which is ε-DP by sequential composition. A delta-method expansion (Appendix A) yields:

$$\boxed{\mathrm{MSE}_A \approx \frac{(U-L)^2}{2n^2\varepsilon_2^2} + \frac{2\mu_c^2}{n^2\varepsilon_1^2},} \tag{4}$$

where $\mu_c = \bar{x} - c$ is the centered mean. The first term is *sum noise*, irreducible for any Laplace-based approach. The second is *count noise*, proportional to μ_c^2. The count noise vanishes when $\bar{x} = c$.
Truncation and approximation quality. In practice, we clamp $\hat{n} \leftarrow \mathrm{clamp}(\hat{n}, n_{\min}, n_{\max})$, which is post-processing (hence still ε-DP). The delta-method used to obtain the MSE in Eq.(4)is a reasonable approximation when $n\varepsilon_1 \gg 1$. Figure 1(d) demonstrates this: the delta method tracks the empirical MSE closely for $n \geq 100$ ($n\varepsilon_1 \geq 15$) and overestimates modestly at $n = 50$ ($n\varepsilon_1 = 7.5$).
Adaptive budget split. Setting $\partial\,\mathrm{MSE}_A/\partial\varepsilon_1 = 0$ yields the optimal allocation:

$$\frac{\varepsilon_1}{\varepsilon_2} = \left(\frac{4\mu_c^2}{(U-L)^2}\right)^{1/3} \tag{5}$$

	MSE (leading constant $\times (U-L)^2/(n^2\varepsilon^2)$)			
$\bar{x}$ position	No count (App. B)	Explicit (App. A)	KSW 3 (2D, custom)	Simplex (Laplace)
$\bar{x} = c$	$\mathbf{1/2}^*$	$\mathbf{1/2}$	2	1
$\lvert\bar{x} - c\rvert = \frac{U-L}{4}$	bias*	≈ 1.2	2	$\mathbf{1.25}$
$\lvert\bar{x} - c\rvert = \frac{U-L}{2}$	bias*	4	$\mathbf{2}$	$\mathbf{2}$
Budget split?	No	Yes	No	No
Standard Laplace?	Yes	Yes	**No**	Yes

*MSE $= 1/2$ when $n = d^*$; otherwise bias $\beta^2\mu_c^2$ dominates, $\beta = (n_{\max} - n_{\min})/(n_{\max} + n_{\min})$.

Table 1: MSE comparison in units of $(U-L)^2/(n^2\varepsilon^2)$. KSW Alg. 3 [Kulesza et al., 2024] is minimax optimal but requires custom noise. The simplex [Fitzsimons et al., 2025] matches KSW 3 at the worst case and beats it elsewhere, using only standard Laplace noise. Approach A wins if the mean is close to the center of the range but has twice the optimal MSE at extreme values for the mean.

When $\mu_c \approx 0$: $\varepsilon_1 \to 0$, nearly all budget goes to the sum. When $|\mu_c| \approx (U-L)/2$ (extreme mean): $\varepsilon_1 \approx \varepsilon/2$, an even split. Since μ_c is unknown, conservative defaults are $\varepsilon_1 \approx 0.15\varepsilon$ or 0.5ε. In Appendix D we propose a three-phase protocol that auto-tunes the split, staying within $1.15\times$ the oracle MSE in our simulation.

Comparison to KSW. KSW Algorithm 2 [Kulesza et al., 2024] is Approach A with $\varepsilon_1 = \varepsilon/2$; our adaptive split generalizes it. KSW Algorithm 3 is minimax-optimal and avoids budget splitting via a custom 2D mechanism, achieving uniform MSE $= 2(U-L)^2/(n^2\varepsilon^2)$. We compare Algorithm A and KSW 3 in Columns 2 and 3 of Table 1 which summarizes the main findings of the paper by comparing all strategies. Approach A is $4\times$ better than the minimax optimum (KSW 3) if the mean is at the center of the data; at extreme means it is $2\times$ worse. This is an *instance-dependent* tradeoff: Approach A exploits dataset structure at the cost of worse worst-case guarantees. Figure 1(a) demonstrates these predictions. We also note that KSW 3 requires non-standard noise and is currently unavailable in DP libraries.

3.2 No Counting (Approach B)

When the uncertainty about the dataset size is moderate, the count query may not be worth its privacy cost. Approach B spends the full budget ε on the centered sum and divides by a fixed constant d: $\hat{\mu}_B = (S_c + \text{Lap}(\Delta_c/\varepsilon))/d + c$. This gives deterministic bias $(n/d - 1)\mu_c$ and noise η_S/d, where $\eta_S \sim \text{Lap}((U-L)/(2\varepsilon))$ yielding (Appendix A):

$$\boxed{\text{MSE}_B = \left(\frac{n}{d} - 1\right)^2 \mu_c^2 + \frac{(U-L)^2}{2d^2\varepsilon^2}} \tag{6}$$

Choosing d. Minimizing worst-case bias over $n \in \{n_{\min}, n_{\max}\}$ yields $d^* = (n_{\min} + n_{\max})/2$, with worst-case bias factor $\beta^2 = \left(\frac{n_{\max} - n_{\min}}{n_{\max} + n_{\min}}\right)^2$. Approach B wins when ε is small (splitting a tiny budget cripples both queries) or $n_{\max}/n_{\min}$ is small (bias is negligible). Note that Approach B is inconsistent: even as $\varepsilon \to \infty$, the bias $(n/d - 1)\mu_c$ persists whenever $n \neq d$.

3.3 Implicit Counting (Simplex)

Approaches A and B represent two extremes: count explicitly (budget splitting) or don't count (risk of bias). The simplex mechanism of Fitzsimons et al. [2025] offers a third option: *count implicitly*.

Augment each x_i to the pair $(x_i - L,\ U - x_i)$. Every row has L1 norm exactly $U - L$, so the 2D query $(s_1 = \sum(x_i - L),\ s_2 = \sum(U - x_i))$ has L1 sensitivity $U - L$ in the add-remove model. Releasing both column sums with independent $\text{Lap}((U-L)/\varepsilon)$ noise is ε-DP with *no budget splitting*. If we define $m_1 = s_1 + Z_1; Z_1 \sim \text{Lap}((U-L)/\varepsilon)$ and $m_2 = s_2 + Z_2; Z_2 \sim \text{Lap}((U-L)/\varepsilon)$, the count is recovered as $\hat{n} = (m_1 + m_2)/(U-L)$ and the mean as $\hat{\mu} = m_1/\hat{n} + L$, both via post-processing.

A delta-method expansion (following the same logic as Appendix A) yields:

$$\text{MSE}_{\text{simplex}} \approx \frac{(U-L)^2 + 4\mu_c^2}{n^2\varepsilon^2} \tag{7}$$

As summarized in Table 1, at $\bar{x} = c$, $MSE_{simplex} = (U-L)^2/(n^2\varepsilon^2)$, half of KSW Alg. 3 but 2× Approach A's oracle. At $|\bar{x} - c| = (U-L)/2$, $MSE_{simplex} = 2(U-L)^2/(n^2\varepsilon^2) = MSE_{KSW3} = 1/2MSE_A$. The simplex is minimax optimal (same worst case as KSW 3) yet better at non-worst case means, using only standard Laplace noise.

When to use the simplex. The simplex dominates all other approaches when uncertainty about n and $\bar{x}$ is large. It requires no tuning, no budget splitting, and achieves minimax-optimal worst-case MSE. Approach A is preferable only when the analyst has reason to believe $\bar{x} \approx c$, e.g., if it is prudent to assume that the distribution is approximately symmetric. The centering yields 2× lower MSE in this case. Figure 1(a) confirms this: the simplex (pink) remains below KSW 3 (dotted) over the full range of mean values and below Approach A at off-center means.

4 Discussion

Decision support. Here we offer some practical guidelines. First and foremost: Always center: compute $S_c = \sum(x_i - c)$ with $c = (L + U)/2$. Then make two decisions. *(1) To count or not to count?* When $\bar{x}$ is unknown, Approach B's worst-case bias $\beta^2\mu_c^2$ does not shrink with n or ε, so any counting mechanism dominates unless $n_{min} \approx n_{max}$. Only for this case is the no-count Approach B safe for all $\bar{x}$. If $\bar{x} \approx c$ is known a priori, Approach B remains competitive up to $n_{max} \leq 1.5n_{min}$. *(2) If counting, implicitly or explicitly?* Default to implicit counting via the simplex: it requires no tuning, no budget splitting, and is minimax optimal. Use explicit counting (Approach A) only when $\bar{x} \approx c$ is known, in which case the reduced sensitivity from centering justifies the budget-splitting cost; a three-phase protocol (Appendix D) auto-tunes the split.

Implications for DP libraries. OpenDP currently suffers from an up to factor four increase in variance by not centering and uses resize mean, which further suffers from an implicit $(U-L)$ vs. $(U-L)/2$ sensitivity for the sum (Appendix C). The simplex mechanism, already implementable with standard Laplace noise, always dominates resize mean and achieves KSW 3-level worst-case performance without custom noise distributions. These are practical improvements requiring no new theory.

Limitations. The delta-method MSE approximations lose accuracy when $n\varepsilon$ is small (Figure 1d). We assume known bounds $[L, U]$; extending to multivariate means and unknown bounds is future work.

References

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference (TCC)*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006. doi: 10.1007/11681878_14.

Jack Fitzsimons, James Honaker, Michael Shoemate, and Vikrant Singhal. Private means and the curious incident of the free lunch, 2025. URL `https://arxiv.org/abs/2408.10438`.

Alex Kulesza, Ananda Theertha Suresh, and Yuyan Wang. Mean estimation in the add-remove model of differential privacy. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

OpenDP. Aggregation: Mean, unknown dataset size. `https://docs.opendp.org/en/stable/api/user-guide/transformations/aggregation-mean.html`, 2025. Last accessed: 02/18/2026.

A MSE Derivations

A.1 Approach A: Centered Sum + Count

The mean estimate is $\hat{\mu}_A = \hat{S}_c/\hat{n} + c$ where $\hat{S}_c = S_c + \eta_S$ and $\hat{n} = n + \eta_n$ with $\eta_S \sim \text{Lap}((U-L)/(2\varepsilon_2))$ and $\eta_n \sim \text{Lap}(1/\varepsilon_1)$ independent.

Define $f(a,b) = a/b + c$. We estimate $\bar{x} = f(S_c, n) = S_c/n + c$. By the delta method (first-order Taylor-series expansion) applied at (S_c, n):

$$\begin{aligned}\hat{\mu}_A - \bar{x} &\approx \left.\frac{\partial f}{\partial a}\right|_{(S_c,n)} \eta_S + \left.\frac{\partial f}{\partial b}\right|_{(S_c,n)} \eta_n \\ &= \frac{1}{n}\eta_S - \frac{S_c}{n^2}\eta_n = \frac{1}{n}\eta_S - \frac{\bar{x}-c}{n}\eta_n\end{aligned}$$

Since $\eta_S \perp \eta_n$:

$$\begin{aligned}\text{MSE}_A &\approx \frac{1}{n^2}\,\text{Var}(\eta_S) + \frac{(\bar{x}-c)^2}{n^2}\,\text{Var}(\eta_n) \\ &= \frac{(U-L)^2}{2n^2\varepsilon_2^2} + \frac{2(\bar{x}-c)^2}{n^2\varepsilon_1^2}\end{aligned}$$

A.2 Optimal Budget Split for Approach A

Let $A = 2(\bar{x}-c)^2$ and $B = (U-L)^2/2$. We minimize $\text{MSE}_A(\varepsilon_1) = A/(n^2\varepsilon_1^2) + B/(n^2(\varepsilon-\varepsilon_1)^2)$ over ε_1:

$$\frac{d\,\text{MSE}_A}{d\varepsilon_1} = -\frac{2A}{n^2\varepsilon_1^3} + \frac{2B}{n^2(\varepsilon-\varepsilon_1)^3} = 0$$

This gives $A/\varepsilon_1^3 = B/(\varepsilon-\varepsilon_1)^3$, so:

$$\frac{\varepsilon_1}{\varepsilon-\varepsilon_1} = \left(\frac{A}{B}\right)^{1/3} = \left(\frac{4(\bar{x}-c)^2}{(U-L)^2}\right)^{1/3}$$

At $\bar{x} = c$: $\varepsilon_1 \to 0$. At $|\bar{x}-c| = (U-L)/2$: $\varepsilon_1/\varepsilon_2 = 1$, so $\varepsilon_1 = \varepsilon/2$.

Substituting the optimal split back: $\text{MSE}_A^* = (A^{1/3} + B^{1/3})^3/n^2\varepsilon^2$. At $\bar{x} = c$: $\text{MSE}_A^* = B/(n^2\varepsilon^2) = (U-L)^2/(2n^2\varepsilon^2)$, matching the oracle.

A.3 Approach B: Fixed Denominator

The estimate is $\hat{\mu}_B = (S_c + \eta_S)/d + c$ with $\eta_S \sim \text{Lap}((U-L)/(2\varepsilon))$.

$$\hat{\mu}_B - \bar{x} = \frac{n(\bar{x}-c) + \eta_S}{d} + c - \bar{x} = \left(\frac{n}{d} - 1\right)(\bar{x}-c) + \frac{\eta_S}{d}$$

The first term is deterministic bias; the second is random noise. Thus

$$\text{MSE}_B = \underbrace{\left(\frac{n}{d}-1\right)^2(\bar{x}-c)^2}_{\text{bias}^2} + \underbrace{\frac{(U-L)^2}{2d^2\varepsilon^2}}_{\text{variance}}$$

Minimax-optimal d. The squared bias term $(n/d-1)^2$ is maximized at $n = n_{\min}$ or $n = n_{\max}$. Setting $1 - n_{\min}/d = n_{\max}/d - 1$ yields $d^* = (n_{\min} + n_{\max})/2$, with worst-case bias factor $\left(\frac{n_{\max}-n_{\min}}{n_{\max}+n_{\min}}\right)^2$.

B The KSW Approach in Detail

We describe the two main algorithms of Kulesza et al. [2024] for ε-DP mean estimation in the add-remove model with unknown $n \in [n_{\min}, n_{\max}]$.

B.1 KSW Algorithm 2: Centered Sum + Count (Fixed Split)

Algorithm 2 uses the same centered sum + count structure as our Approach A, but with a fixed equal budget split:

1. **Center.** Set $c = (L + U)/2$.
2. **Noisy count.** Release $\hat{n} = n + \text{Lap}(2/\varepsilon)$, spending budget $\varepsilon/2$. Variance: $\text{Var}(\hat{n}) = 8/\varepsilon^2$.
3. **Noisy centered sum.** Release $\hat{S}_c = S_c + \text{Lap}((U - L)/\varepsilon)$, spending budget $\varepsilon/2$. Variance: $\text{Var}(\hat{S}_c) = 2(U - L)^2/\varepsilon^2$.
4. **Estimate.** Return $\hat{\mu} = \hat{S}_c/\hat{n} + c$.

By sequential composition, this is ε-DP. The MSE follows from our Approach A formula (Equation (4)) with $\varepsilon_1 = \varepsilon_2 = \varepsilon/2$:

$$\text{MSE}_{\text{KSW2}} = \frac{(U - L)^2}{2n^2(\varepsilon/2)^2} + \frac{2(\bar{x} - c)^2}{n^2(\varepsilon/2)^2} = \frac{2(U - L)^2}{n^2\varepsilon^2} + \frac{8(\bar{x} - c)^2}{n^2\varepsilon^2} \tag{8}$$

At the worst case $|\bar{x} - c| = (U - L)/2$, this gives $\text{MSE}_{\text{KSW2}} = 4(U - L)^2/(n^2\varepsilon^2)$.

Connection to Approach A. KSW Algorithm 2 is a special case of our Approach A with $\varepsilon_1 = \varepsilon/2$. Our adaptive budget split (Equation (5)) improves on this by allocating less budget to counting when $\bar{x} \approx c$, yielding MSE as low as $(U - L)^2/(2n^2\varepsilon^2)$, a $4\times$ improvement at centered means.

B.2 KSW Algorithm 3: 2D Hourglass Mechanism

Algorithm 3 avoids budget splitting entirely by using a single 2D query with a custom noise mechanism.

Instead of separately querying the sum and count, Algorithm 3 treats the pair (S_c, n) as a single 2D vector-valued query and adds 2D noise calibrated to the joint sensitivity.

Sensitivity geometry. In the add-remove model, adding or removing a record $x \in [L, U]$ changes the 2D query vector (S_c, n) by $(x - c, +1)$ or $(-(x - c), -1)$, respectively. The set of all possible changes is:

$$\mathcal{S} = \{(x - c, 1) : x \in [L, U]\} \cup \{(-(x - c), -1) : x \in [L, U]\}$$

This forms a set of line segments in $\mathbb{R}^2$ spanning from $(-(U - L)/2, +1)$ to $((U - L)/2, +1)$ and from $(-(U - L)/2, -1)$ to $((U - L)/2, -1)$, an "hourglass" shape.

Mechanism. The hourglass mechanism adds 2D noise (η_1, η_2) drawn from a distribution whose density is proportional to $\exp(-\varepsilon \cdot \|(\eta_1, \eta_2)\|_{\mathcal{K}})$, where $\|\cdot\|_{\mathcal{K}}$ is a norm whose unit ball $\mathcal{K}$ is the polar dual of the convex hull of $\mathcal{S}$. Intuitively, the noise is shaped to match the geometry of the sensitivity set, adding less noise in directions where the query is less sensitive. The mechanism releases:

$$(\hat{S}_c, \hat{n}) = (S_c, n) + (\eta_1, \eta_2)$$

and the mean estimate is $\hat{\mu} = \hat{S}_c/\hat{n} + c$.

MSE. Kulesza et al. [2024] show that this achieves:

$$\text{MSE}_{\text{KSW3}} = \frac{2(U - L)^2}{n^2\varepsilon^2} \tag{9}$$

uniformly over $\bar{x} \in [L, U]$. This is minimax optimal: no ε-DP mechanism can achieve lower worst-case MSE.

Why it avoids the budget-splitting tax. The crucial advantage is that the 2D mechanism uses the full budget ε for both the sum and count components simultaneously, rather than splitting ε into $\varepsilon_1 + \varepsilon_2$. In geometric terms, the hourglass noise distribution exploits the correlation structure between the sum and count sensitivity directions: when the added record x is near c (small sum change), the count change is still ± 1, and vice versa. A 1D Laplace mechanism cannot exploit this structure; it must treat sum and count as independent queries.

Comparison to our approaches. The hourglass mechanism outperforms our Approach A at extreme $\bar{x}$ (MSE $2\times$ vs. $4\times$ the oracle) and matches it at $\bar{x} \approx c$ up to the budget-splitting overhead. However, it requires implementing a custom 2D noise distribution, while our approaches use only the standard Laplace mechanism. Table 2 summarizes the full comparison across the counting spectrum.

Table 2: Complete comparison across the "how much to count" spectrum. MSE in units of $(U-L)^2/(n^2\varepsilon^2)$. Parameters: $d^* = (n_{\min}+n_{\max})/2$, $\beta = (n_{\max}-n_{\min})/(n_{\max}+n_{\min})$.

	MSE (units of $\frac{(U-L)^2}{n^2\varepsilon^2}$)		Budget	Requires
Approach	$\bar{x} = c$	$\lvert\bar{x}-c\rvert = \frac{U-L}{2}$	split?	custom mech.?
Resize Mean B (no count)	$2\left(\frac{n}{d^*}\right)^2$	$\beta^2 n^2\varepsilon^2 + 2\left(\frac{n}{d^*}\right)^2$	No	No
Resize Mean A (with count)	$\approx 2\left(\frac{n}{\hat{n}}\right)^2$	$\approx 2\left(\frac{n}{\hat{n}}\right)^2 + \frac{8(\bar{x}-c)^2}{n^2\varepsilon_1^2}$	Yes	No
Approach B (no count)	$\frac{1}{2}\left(\frac{n}{d^*}\right)^2$	$\beta^2 n^2\varepsilon^2 + \frac{1}{2}\left(\frac{n}{d^*}\right)^2$	No	No
Approach A (adaptive)	**1/2**	4	Yes	No
KSW Alg. 2 ($\varepsilon_1 = \varepsilon/2$)	2	4	Yes	No
KSW Alg. 3 (2D)	2	**2**	No	**Yes**

C OpenDP Resize Mean

OpenDP's *resize mean* pads (with constant w, typically $(L+U)/2$) or subsamples the dataset to a fixed size m, then releases the sample mean with Laplace noise calibrated to sensitivity $(U-L)/m$. Below we describe two variants, with and without counting, and compare them to the centered strategies of Section 3.

C.1 Resize Mean A (With Count)

Spend ε_1 on a noisy count via the geometric mechanism, clamp to $[n_{\min}, n_{\max}]$, then spend $\varepsilon_2 = \varepsilon - \varepsilon_1$ on the resize mean with $m = \hat{n}$:

$$\hat{n} = \text{clamp}\big(n + \text{Geo}(1/\varepsilon_1),\ n_{\min},\ n_{\max}\big), \tag{10}$$

$$\hat{\mu}_{\text{RM-A}} = \overline{X}_{\hat{n}} + \text{Lap}\big((U-L)/(\hat{n}\,\varepsilon_2)\big), \tag{11}$$

where $\overline{X}_m$ denotes the mean of the dataset padded or subsampled to size m. The sensitivity is $(U-L)/\hat{n}$. By sequential composition this is ε-DP.

When $\hat{n} \approx n$, the noise variance is $2(U-L)^2/(\hat{n}^2\varepsilon_2^2)$. Compared to Approach A, which has noise variance $(U-L)^2/(2n^2\varepsilon_2^2)$ from the centered sum, resize mean A incurs $4\times$ higher noise variance because it does not center: the uncentered sensitivity $(U-L)/m$ is twice the centered sensitivity $(U-L)/2$ when $m=1$, and for $L=0$ the effective uncentered sum sensitivity is U vs. $U/2$ centered.

C.2 Resize Mean B (No Count)

Resize to a fixed $d^* = (n_{\min}+n_{\max})/2$ and spend the full budget ε:

$$\hat{\mu}_{\text{RM-B}} = \overline{X}_{d^*} + \text{Lap}\big((U-L)/(d^*\varepsilon)\big). \tag{12}$$

The noise variance is $2(U-L)^2/(d^{*2}\varepsilon^2)$. Approach B, by contrast, divides the centered sum by d^* with noise variance $(U-L)^2/(2d^{*2}\varepsilon^2)$, a $4\times$ advantage from centering.

C.3 Why Centering Wins

The core difference is sensitivity. The centered sum has global sensitivity $\Delta_c = (U-L)/2$ regardless of n or m. The resize mean has sensitivity $(U-L)/m$, which depends on m and, crucially, applies to the *uncentered* data. For non-negative data ($L=0$), the uncentered sum sensitivity is $\max\{|U|,|L|\} = U$, while the centered sum sensitivity is $U/2$. The resize mean's per-record sensitivity $(U-L)/m = U/m$ is smaller than U only because it averages over m records, but the centered sum achieves the same averaging *and* has the centering advantage. Concretely, for $L=0$:

$$\frac{\text{Var(Resize Mean)}}{\text{Var(Centered Sum}/n)} = \frac{2(U/m)^2/\varepsilon^2}{2(U/2)^2/(n^2\varepsilon^2)} \cdot n^2 = \frac{4n^2}{m^2},$$

which equals 4 when $m = n$. The centered approaches thus dominate resize mean whenever $L \neq -U$.

C.4 Simulation Comparison

Figure 1 compares all eight strategies. At centered means ($\bar{x} = c$), Resize Mean A is 12.3× the oracle MSE vs. 1.08× for Approach A; at boundary means ($\bar{x}$ near L), it is 2.5× vs. 1.0×. Resize Mean B tracks Approach B closely; both suffer from denominator mismatch bias when $\bar{x} \neq c$, but with 4× higher noise variance. The KSW Alg. 3 minimax benchmark (MSE $= 2(U-L)^2/(n^2\varepsilon^2)$, uniform over $\bar{x}$) appears as a flat line in panel (a), confirming the advantage of instance-dependent approaches.

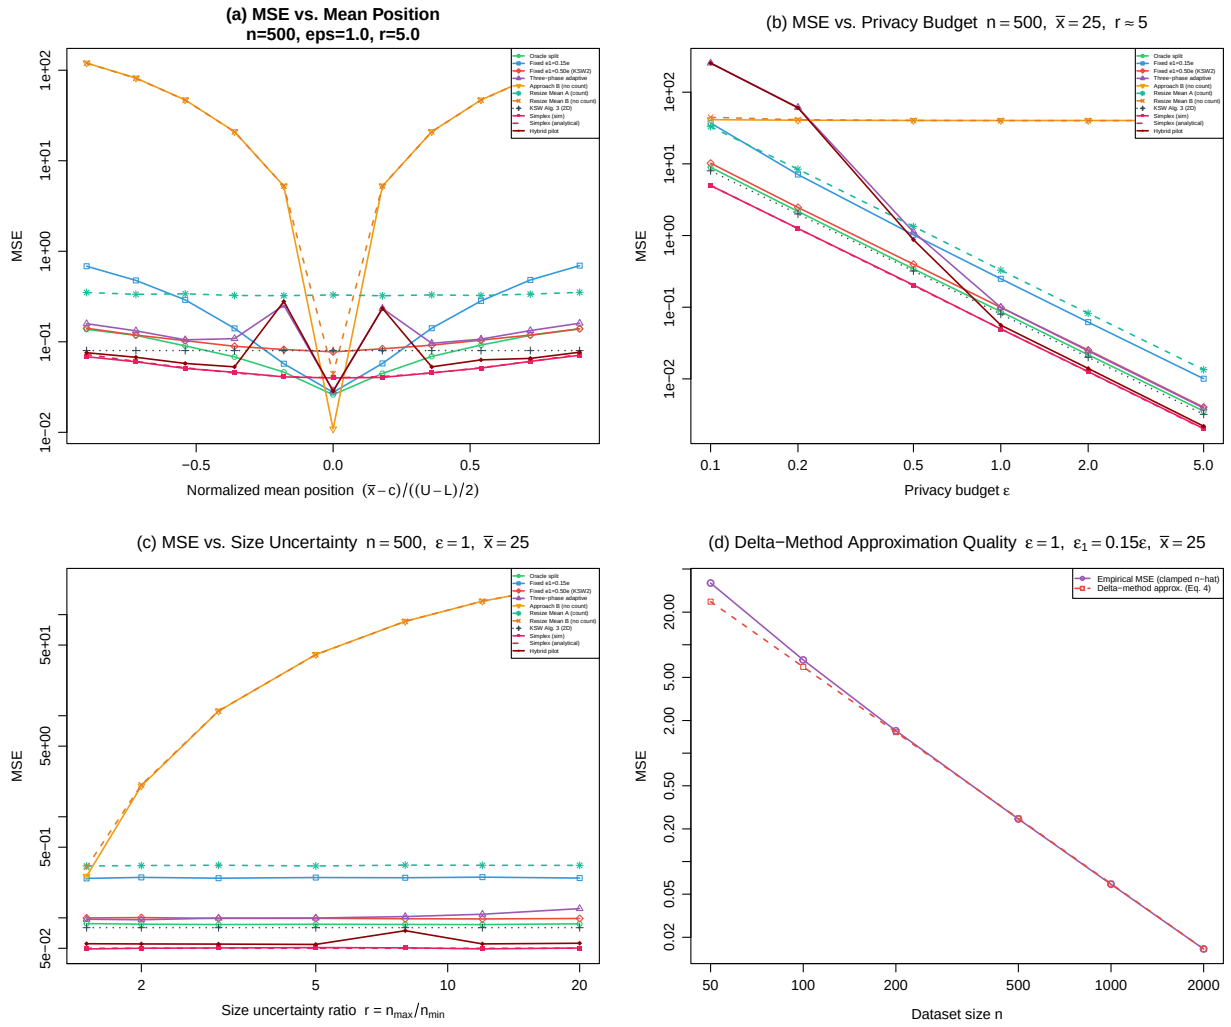

Figure 1: MSE comparison of eleven strategies ($n = 500$, $[L, U] = [0, 100]$, $r \approx 5$, 30,000 trials). (a) MSE vs. mean position at $\varepsilon = 1$: the hybrid pilot (dark red) tracks the oracle near $\bar{x} = c$ and the simplex at extreme means. (b) MSE vs. privacy budget at $\bar{x} = 25$. (c) MSE vs. size uncertainty ratio: simplex and hybrid are flat across r, while explicit counting approaches degrade. (d) Delta-method approximation quality: Eq. (4) tracks simulation closely for $n \geq 100$.

D Three-Phase Adaptive Protocol

The optimal budget split for Approach A (Equation (5)) depends on the unknown centered mean $\mu_c = \bar{x} - c$. We describe a end-to-end private protocol that estimates μ_c via a cheap pilot and sets the split automatically, requiring no user-specified tuning parameters beyond the standard DP inputs $(L, U, n_{\min}, n_{\max}, \varepsilon)$.

Algorithm 1 Three-Phase Adaptive DP Mean Estimation

Require: Data $X = \{x_1, \ldots, x_n\}$, bounds $[L, U]$, size range $[n_{\min}, n_{\max}]$, budget ε
Ensure: ε-DP estimate $\hat{\mu}$ of $\bar{x}$

— **Setup** —

$c \leftarrow (L+U)/2, \quad \Delta_c \leftarrow (U-L)/2, \quad d^* \leftarrow (n_{\min} + n_{\max})/2$

— **Phase 0: Pilot** (budget $\varepsilon_0 = 0.05\varepsilon$) —

$\tilde{\mu} \leftarrow \big(\sum_i (x_i - c) + \mathrm{Lap}(\Delta_c/\varepsilon_0)\big)/d^* + c$ ▷ Approach B with 5% budget

— **Phase 1: Adaptive split** —

$\sigma^2_{\text{pilot}} \leftarrow 2\Delta_c^2/(d^{*2}\varepsilon_0^2)$ ▷ Pilot noise variance

$\hat{\gamma} \leftarrow \mathrm{clamp}\big((\tilde{\mu} - c)^2 - \sigma^2_{\text{pilot}},\ 0,\ \Delta_c^2\big)$ ▷ Bias-corrected $(\bar{x} - c)^2$

$\rho \leftarrow (4\hat{\gamma}/(U-L)^2)^{1/3}$

$\varepsilon_1 \leftarrow \mathrm{clamp}\big(\varepsilon_{\text{rem}} \cdot \rho/(1+\rho),\ 0.01\varepsilon,\ \varepsilon_{\text{rem}}/2\big), \quad \varepsilon_2 \leftarrow \varepsilon_{\text{rem}} - \varepsilon_1$

— **Phase 2: Main estimation** (budget $\varepsilon_1 + \varepsilon_2 = \varepsilon_{\text{rem}}$) —

$\hat{n} \leftarrow \mathrm{clamp}\big(n + \mathrm{Lap}(1/\varepsilon_1),\ n_{\min},\ n_{\max}\big)$

$\hat{S}_c \leftarrow \sum_i (x_i - c) + \mathrm{Lap}(\Delta_c/\varepsilon_2)$

return $\hat{S}_c/\hat{n} + c$

Proposition 2 (Privacy)**.** *Algorithm 1 satisfies ε-differential privacy in the add-remove model.*

Proof. The algorithm accesses X through three Laplace queries: pilot sum (budget ε_0), count (budget ε_1), and main sum (budget ε_2). By sequential composition, the total cost is $\varepsilon_0 + \varepsilon_1 + \varepsilon_2 = \varepsilon$. The budget split in Phase 1 is post-processing of the pilot output; the clamping of $\hat{n}$ is post-processing of the noisy count. □

MSE analysis. Let $\alpha = \varepsilon_0/\varepsilon = 0.05$. The main estimation uses budget $\varepsilon_{\text{rem}} = (1-\alpha)\varepsilon$ instead of ε. Since MSE scales as $1/\varepsilon^2$, the pilot inflates MSE by at most $1/(1-\alpha)^2 = 1/0.95^2 \approx 1.11$. The remaining overhead comes from split misspecification: the pilot provides a noisy estimate $\hat{\gamma}$ of $(\bar{x} - c)^2$, but the MSE as a function of ε_1 is smooth and convex with a flat basin around its minimum, so moderate errors in ε_1 have small effects.

Simulation demonstration. Monte Carlo simulations (50,000 trials, $[L, U] = [0, 100]$, $n = 500$, $r \approx 5$, $\varepsilon = 1$) confirm that the three-phase protocol stays within 1.12–1.14× the oracle MSE across all mean positions: 1.12× at $\bar{x} = c$ (centered), 1.13× at $\bar{x} = (L+U)/4$ (quarter), and 1.14× at $\bar{x}$ near L (boundary). By contrast, the fixed split $\varepsilon_1 = 0.15\varepsilon$ is 5.01× worse at boundary means, and $\varepsilon_1 = 0.5\varepsilon$ is 3.11× worse at centered means. No fixed split performs uniformly well. Figure 1 shows the full comparison across all eight strategies.