CONVERSE AND ACHIEVABILITY BOUNDS ON ERROR OF APPROXIMATE DIFFERENTIAL PRIVACY

Anonymous authors
Paper under double-blind review

ABSTRACT

Considerable effort has been devoted to designing additive noise mechanisms that achieve (ϵ, δ)-differential privacy for queries with ℓ_2 sensitivity Δ_2 while minimizing the resulting error. Let $\mathcal{V}_M(\Delta_2)$ denote the infimum of the mean squared error (MSE) over all M-dimensional additive noise distributions that are (ϵ, δ)-differentially private for every such query. We establish converse and achievability bounds on this optimal MSE. The converse (impossibility) shows that no feasible additive noise can achieve an MSE below $A_M(\Delta_2)$ by more than $\mathcal{O}(\Delta_2^2 M^{-2})$, while under a further condition on the noise shape, the achievability (possibility) shows that generalized gamma noise achieves an MSE of at most $A_M(\Delta_2) + \mathcal{O}(\Delta_2^2 M^{-2})$, where

$$A_M(\Delta_2) = M\Delta_2^2 v_0 - \Delta_2^2 C(\epsilon, \delta) + \frac{\Delta_2^2 J(\epsilon, \delta)}{M}.$$

Here, $\Delta_2^2 v_0$ is the exactly calibrated Gaussian variance per coordinate, and $C(\epsilon, \delta)$ and $J(\epsilon, \delta)$ are explicit functions of the privacy parameters alone. We derive the converse through the duality of error minimization under privacy constraints. To establish achievability, we develop an efficient noise calibration procedure that attains this upper bound. Our procedure determines the two shape parameters of the generalized gamma noise in closed form, avoiding the numerical shape search used by existing approaches. Relaxing the mild assumption on shape parameters required by our construction leaves the converse unchanged and increases the achievable upper bound by only $\mathcal{O}(\Delta_2^2 M^{-1})$. Numerical evaluations support the theoretical results.

1 INTRODUCTION

Adding Gaussian noise (Dwork et al., 2006; Dwork & Roth, 2014) is the standard way to release privacy-sensitive data, and it underpins privacy-preserving optimization at scale (Abadi et al., 2016). Its privacy profile can be calibrated exactly (Balle & Wang, 2018), and its rotational symmetry matches an ℓ_2 sensitivity constraint. However, calibration only optimizes the scale of Gaussian noise, and does not tell whether alternative noise achieves smaller error under the same guarantee.

Several alternatives have been proposed for ℓ_2 sensitivity under (ϵ, δ)-differential privacy, including the ℓ_2 mechanism (Joseph et al., 2025) and the rank-one singular Gaussian mechanism (Ji & Li, 2024). In essence, the noise distributions of these mechanisms belong to one broader family, namely spherically symmetric noise whose radius follows a generalized gamma distribution, which Liu et al. (2026b) and Wei et al. (2026) recently studied as a whole. By reducing any feasible noise to a spherically symmetric one through Haar averaging and testing it against a single spherical threshold event, Wei et al. (2026) show that Gaussian noise is first-order optimal as $M \to \infty$ among all data-independent additive mechanisms, while the error at finite dimension remains open. Two questions are therefore open at a given finite dimension M. **Q1:** How much error can a non-Gaussian mechanism save? **Q2:** Can the best non-Gaussian mechanism be written in closed form, or must it be found by numerical search (Wei et al., 2026; Yang et al., 2025)?

We answer both questions. For a probability distribution P on $\mathbb{R}^M$ and a vector $h \in \mathbb{R}^M$, let P_h denote the distribution of $Z + h$ with $Z \sim P$. For an M-dimensional query of ℓ_2 sensitivity at most Δ_2, we consider the primal optimization problem

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

$$\inf_{P} \mathbb{E}_{Z\sim P}\big[\|Z\|_2^2\big] \quad \text{subject to} \quad D_\epsilon(P\|P_h) \le \delta \quad \text{for all } h \in \mathbb{R}^M \text{ with } \|h\|_2 \le \Delta_2, \tag{1}$$

where D_ϵ is the hockey-stick divergence (Definition 3) and the infimum is taken over all probability distributions on $\mathbb{R}^M$. The constraint $D_\epsilon(P\|P_h) \le \delta$ for all $\|h\|_2 \le \Delta_2$ holds exactly when the additive mechanism with noise $Z \sim P$ is (ϵ, δ)-DP on every query with ℓ_2 sensitivity at most Δ_2 (Balle & Wang, 2018).

Definition 1 (Optimal MSE (OMSE))**.** *Let $\mathcal{V}_M(\Delta_2)$ be the infimum of the primal problem in eq. (1). We refer to $\mathcal{V}_M(\Delta_2)$ as the optimal mean squared error (OMSE) over all data-independent additive mechanisms that are (ϵ, δ)-DP for every M-dimensional query with ℓ_2 sensitivity at most Δ_2.*

Define $A_M(\Delta_2) := M\Delta_2^2 v_0 - \Delta_2^2 C(\epsilon,\delta) + \frac{\Delta_2^2 J(\epsilon,\delta)}{M}$, where $M\Delta_2^2 v_0$ is the MSE of the exactly calibrated Gaussian mechanism (Balle & Wang, 2018), whose variance per coordinate is $\Delta_2^2 v_0$, and $C(\epsilon,\delta) \ge 0$ and $J(\epsilon,\delta)$ are constants that depend only on (ϵ,δ) (Appendix B.4).

To answer **Q1**, we show that

- A **converse** bound, $\mathcal{V}_M(\Delta_2) \ge A_M(\Delta_2) - \mathcal{O}(\Delta_2^2 M^{-2})$, which expresses impossibility. No additive noise can achieve MSE that is smaller than $A_M(\Delta_2)$ by more than an order of $\frac{1}{M^2}$. The formal result is provided in Theorem 2.
- An **achievability** bound, $\mathcal{V}_M(\Delta_2) \le A_M(\Delta_2) + \mathcal{O}(\Delta_2^2 M^{-2})$, which expresses possibility. There exists additive noise that can achieve MSE exceeding $A_M(\Delta_2)$ by at most an order of $\frac{1}{M^2}$. The formal result is presented in Theorem 3 under a further condition on the noise shape, and Theorem 4 brackets the $\frac{1}{M}$ coefficient where that condition fails.

To answer **Q2**, we develop the closed-form generalized gamma mechanism (CGG, Algorithm 1). The radius of the noise has a generalized gamma distribution as in Wei et al. (2026), yet our procedure avoids an explicit numerical search and is two orders of magnitude faster than a shape search over the same family (Section 4). Utility-wise, the CGG mechanism attains the achievability bound, and the searched mechanism of Wei et al. (2026) improves on it by at most $\mathcal{O}(\Delta_2^2 M^{-2})$. Although CGG depends on a mild assumption on shape parameters, relaxing it leaves the converse unchanged and increases the achievable upper bound by only $\mathcal{O}(\Delta_2^2 M^{-1})$. Numerical evaluations support the theoretical results.

2 Preliminaries

2.1 Differential privacy

Fix a dimension $M \ge 2$, privacy parameters $\epsilon > 0$ and $0 < \delta < 1$, and a query f with values in $\mathbb{R}^M$ whose ℓ_2 sensitivity is at most $\Delta_2 > 0$, so that $\|f(\boldsymbol{x}) - f(\boldsymbol{x}')\|_2 \le \Delta_2$ for every pair of neighboring data sets $\boldsymbol{x}, \boldsymbol{x}'$, which differ in one record (Dwork & Roth, 2014). Throughout, we study additive mechanisms $\mathcal{M}(\boldsymbol{x}) = f(\boldsymbol{x}) + Z$ where the noise $Z \in \mathbb{R}^M$ is drawn independently of the data.

Definition 2 (Differential privacy (Dwork et al., 2006))**.** *A randomized mechanism $\mathcal{M}$ is (ϵ, δ)-differentially private, or (ϵ, δ)-DP, if for every pair of neighboring data sets $\boldsymbol{x}, \boldsymbol{x}'$ and every measurable output set $\mathcal{S}$, it satisfies $\Pr[\mathcal{M}(\boldsymbol{x}) \in \mathcal{S}] \le e^\epsilon \Pr[\mathcal{M}(\boldsymbol{x}') \in \mathcal{S}] + \delta$.*

Definition 3 (Hockey-stick divergence (Sason & Verdú, 2016))**.** *For probability measures P, Q on $\mathbb{R}^M$,*

$$D_\epsilon(P\|Q) = \sup_{\mathcal{S}} \big[P(\mathcal{S}) - e^\epsilon Q(\mathcal{S})\big] = \int \big[\mathrm{d}P - e^\epsilon \mathrm{d}Q\big]_+, \tag{2}$$

where $\mathcal{S}$ ranges over measurable sets and the integral is the total mass of the positive part of the signed measure $P - e^\epsilon Q$.

Theorem 1 (Analytic Gaussian mechanism (AGM) (Balle & Wang, 2018))**.** *The Gaussian mechanism $\mathcal{M}(\boldsymbol{x}) = f(\boldsymbol{x}) + Z$ with $Z \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_M)$ is (ϵ, δ)-DP for every query of ℓ_2 sensitivity at most Δ_2 if and only if $\Phi\left(\frac{\Delta_2}{2\sigma} - \frac{\epsilon\sigma}{\Delta_2}\right) - e^\epsilon \Phi\left(-\frac{\Delta_2}{2\sigma} - \frac{\epsilon\sigma}{\Delta_2}\right) \le \delta$, where Φ is the standard normal cumulative distribution function.*

2.2 SPHERICAL NOISE

Lemma 1 (Polar decomposition (Fang et al., 2013; Paolella, 2018)). *A spherically symmetric random vector $Z \in \mathbb{R}^M$ satisfies $Z \overset{d}{=} RU$, where $R \geq 0$ is a scalar random variable (the radius), U is uniformly distributed on the unit sphere $\mathbb{S}^{M-1}$, and R and U are independent.*

Definition 4 (Generalized gamma noise). *For $a, q, p > 0$, the generalized gamma density with scale a, shape q, and power p is (Stacy, 1962)*

$$f_R(r) = \frac{p}{a^q \Gamma\left(\frac{q}{p}\right)} r^{q-1} e^{-\left(\frac{r}{a}\right)^p}, \qquad r > 0. \tag{3}$$

Generalized gamma noise with these parameters is the vector $Z = RU$ where the radius R has density f_R and U is uniform on $\mathbb{S}^{M-1}$ and independent of R.

The density of Z on $\mathbb{R}^M$ is proportional to $\|z\|_2^{q-M} e^{-\left(\frac{\|z\|_2}{a}\right)^p}$ (Appendix B). The special case $q = M$ removes the power of $\|z\|_2$ and contains the ℓ_2 mechanism ($p = 1$) and the Gaussian mechanism ($p = 2$).

Throughout our analysis, we define the normalized shape and the effective shape

$$\rho = \frac{q}{M}, \qquad \tau = \rho p.$$

With $q = \rho M$, the exponent of $\|z\|_2$ becomes $(\rho - 1)M$, so the density of Z is nonincreasing in $\|z\|_2$ exactly when $\rho \leq 1$. For $\rho < 1$ this exponent is negative and the density is unbounded at the origin, with an integrable singularity for every $\rho > 0$. We also use $\mathrm{Gamma}(k, 1)$ to represent the gamma distribution with shape k and unit scale.

Lemma 2 (Haar symmetrization, Lemma 3.3 of Wei et al. (2026)). *Write $\mathcal{L}(Z)$ for the distribution of a random variable Z. Let $Z \in \mathbb{R}^M$ satisfy $D_\epsilon(\mathcal{L}(Z)\|\mathcal{L}(Z+h)) \leq \delta$ for every $\|h\|_2 \leq \Delta_2$ and let O be Haar-distributed on the orthogonal group of $\mathbb{R}^M$ and independent of Z. Then OZ is spherically symmetric, $\mathbb{E}[\|OZ\|_2^2] = \mathbb{E}[\|Z\|_2^2]$, and $D_\epsilon(\mathcal{L}(OZ)\|\mathcal{L}(OZ+h)) \leq \delta$ for every $\|h\|_2 \leq \Delta_2$.*

3 CONVERSE AND ACHIEVABILITY BOUNDS

Lemma 2 helps convert the analysis of an arbitrary noise distribution P to spherically symmetric ones. For every feasible P and $Z \sim P$, the Haar symmetrization OZ is feasible for eq. (1) with the same MSE, so the infimum in eq. (1) equals the infimum over spherically symmetric distributions. Hence it suffices to study the converse and achievability bounds for spherically symmetric noise, and in the following analysis we focus on $Z = RU$ as in Lemma 1. Given Z, we call $V = \frac{\|Z\|_2^2}{M\Delta_2^2}$ its normalized radial energy, so that $\mathbb{E}[\|Z\|_2^2] = M\Delta_2^2\mathbb{E}[V]$. Dividing Z and h by Δ_2 leaves every hockey-stick divergence unchanged and divides the MSE by Δ_2^2. Hence $\mathcal{V}_M(\Delta_2) = \Delta_2^2\mathcal{V}_M(1)$, and our numerical results all take $\Delta_2 = 1$.

Substituting the normalized variance $v = \frac{\sigma^2}{\Delta_2^2}$ into Theorem 1 gives the Gaussian privacy profile (Balle & Wang, 2018)

$$g(v) = \Phi\left(\frac{1}{2\sqrt{v}} - \epsilon\sqrt{v}\right) - e^\epsilon \Phi\left(-\frac{1}{2\sqrt{v}} - \epsilon\sqrt{v}\right), \qquad v > 0.$$

$g(v)$ decreases strictly from 1 to 0, so $g(v_0) = \delta$ has a unique root v_0, and $\Delta_2^2 v_0$ is the exactly calibrated Gaussian variance per coordinate (Balle & Wang, 2018). It is concave on $(0, v_c)$ and convex on (v_c, ∞) with a single inflection point v_c (Appendix B). We use g' and g'' to represent the first and second order derivatives of g, respectively. To help presentation, we denote

$$\tau_* := \frac{1}{\epsilon}\sqrt{\frac{-2g''(v_0)}{g'(v_0)}}, \qquad p_* := -\frac{c_1(v_0, \tau_*)}{\epsilon^2 v_0^2}, \qquad \rho_* := \frac{\tau_*}{p_*}, \tag{4}$$

where τ_* is real exactly when $v_0 > v_c$ and c_1 is an explicit rational function of (ϵ, v_0, τ_*) given in eq. (59) of Appendix E.4.

Furthermore, we consider the following two mild assumptions.

Assumption 1. *The privacy budget* (ϵ, δ) *satisfies* $v_0 > v_c$ *and* $\delta - v_0 g'(v_0) < 1$.

Assumption 2. *The constants satisfy* $p_* \geq \tau_*$, *equivalently* $0 < \rho_* \leq 1$.

Assumption 1 puts the tangent line of g at v_0 strictly below g elsewhere (Lemma 6). Assumption 2 makes the generalized gamma density with normalized shape ρ_* and effective shape τ_* nonincreasing, so that privacy at one calibrated shift covers the whole ball (Lemma 10). Section 3.4 provides converse and achievability bounds when Assumption 2 is relaxed.

3.1 Converse bound

Theorem 2 (Converse bound)**.** *Under Assumption 1, there are finite constants* $K_{\mathrm{cv}}, M_{\mathrm{cv}}$ *that depend only on* (ϵ, δ) *such that for every* $M \geq M_{\mathrm{cv}}$,

$$\mathcal{V}_M(\Delta_2) \geq A_M(\Delta_2) - \frac{\Delta_2^2 K_{\mathrm{cv}}}{M^2} = A_M(\Delta_2) - \mathcal{O}(\Delta_2^2 M^{-2}). \tag{5}$$

Proof sketch. Appendix D gives the full proof in six steps.

Step 1 (Weak duality, Lemma 5). For a threshold function θ on $[0, \infty)$ define $\bar{S}_{M,\theta} = \left\{ y \in \mathbb{R}^M : y_1 \leq \theta\left(\frac{\|y\|_2^2}{M}\right) \right\}$ and $h_M[\theta](v) = \Pr\left(\sqrt{Mv}U \in \bar{S}_{M,\theta}\right) - e^{\epsilon} \Pr\left(\sqrt{Mv}U + e_1 \in \bar{S}_{M,\theta}\right)$, where y_1 is the projection of y on e_1 and v is a realization of $V = \frac{\|Z\|_2^2}{M\Delta_2^2}$. By the supremum in eq. (2), privacy at the shift $\Delta_2 e_1$ implies $\mathbb{E}[h_M[\theta](V)] \leq \delta$ for every θ. Relaxed to this single constraint, the primal problem is the linear program that minimizes $\mathbb{E}[V]$ over the distribution of V. We prove its weak duality (Anderson & Nash, 1987): if $\lambda \geq 0$ and $v + \lambda h_M[\theta](v) \geq \kappa$ for every $v \geq 0$, then $\mathbb{E}[V] \geq \kappa - \lambda\delta$, and therefore $\mathcal{V}_M(\Delta_2) \geq \Delta_2^2 M(\kappa - \lambda\delta)$. We call $(\theta, \lambda, \kappa)$ a certificate and $\kappa - \lambda\delta$ its value.

Step 2 (Strict global tangent, Lemma 6). Consider the Gaussian threshold $\theta_G(v) = \frac{1}{2} - \epsilon v$, whose mass tends to the Gaussian profile $g(v)$ as $M \to \infty$, and the tangent ℓ of g at its root v_0. We show that under Assumption 1, $g \geq \ell$ on $[0, \infty)$ with equality only at v_0. With $\lambda = -\frac{1}{g'(v_0)}$, the limiting certificate has value v_0, which recovers the first-order optimality of Gaussian noise as $M \to \infty$ (Wei et al., 2026).

Step 3 (Critical branch, Lemma 7). Let $j(v, z)$ be the limit of $h_M[\theta](v)$ as $M \to \infty$ for any continuous threshold with $\theta(v) = z$. We construct an analytic threshold Θ near v_0 with $j(v, \Theta(v)) = \ell(v)$ on a whole interval and $\Theta(v_0) = \theta_G(v_0)$. We prove that its slope at v_0 exceeds that of θ_G by $\frac{\epsilon\tau_*}{2}$ with τ_* from eq. (4).

Step 4 (Uniform expansion, Lemmas 4 and 8). We use a two-term normal approximation to one coordinate of a uniform point on $\mathbb{S}^{M-1}$ with $\mathcal{O}(M^{-3})$ remainder, uniform on compact sets. We prove that for any smooth threshold $\theta_M = T + \frac{u}{M} + \frac{b}{M^2}$, the mass $h_M[\theta_M](v)$ equals $j(v, T)$ plus explicit $\frac{1}{M}$ and $\frac{1}{M^2}$ coefficients in T, u, b plus $\mathcal{O}(M^{-3})$, uniformly in v and with one v derivative.

Step 5 (Global support, Lemma 9). We use the threshold θ_M that equals Θ near v_0 and θ_G elsewhere, glued smoothly, plus corrections $\frac{u}{M} + \frac{b}{M^2}$. By **Step 4**, $h_M[\theta_M](v)$ expands through order M^{-2}. We choose u and b so that near v_0 the $\frac{1}{M}$ coefficient is affine in v and the $\frac{1}{M^2}$ coefficient is constant, and $\lambda_M = -\frac{1}{g'(v_0)} + \mathcal{O}(M^{-1})$ so that the affine part cancels. Away from v_0, the leading term keeps the positive margin of **Step 2**. We prove that $(\theta_M, \lambda_M, \kappa_M)$ is a certificate with value $v_0 - \frac{C(\epsilon,\delta)}{M} + \mathcal{O}(M^{-2})$, where $C(\epsilon, \delta) = \frac{1}{2}\epsilon^2 v_0^2 (2 - \tau_*)^2$ is the constant-order Gaussian excess and vanishes only at the Gaussian shape $\tau_* = 2$ (Appendix D.5).

Step 6 (The $\frac{1}{M^2}$ coefficient, Appendix D.5). We use the one parameter left free in **Step 5**, the value of u at v_0. The explicit $\frac{1}{M^2}$ term is a concave quadratic in it. We show that its maximum is $J(\epsilon, \delta)$. This is the same number as the smallest $\frac{1}{M}$ coefficient that generalized gamma noise with τ_* attains over p (eq. (15) in Appendix B.4). Substituting the maximizer into **Step 1** proves eq. (5). K_{cv} collects the $\mathcal{O}(M^{-3})$ remainders and M_{cv} is the dimension beyond which **Steps 4 and 5** hold. □

3.2 THE CLOSED-FORM GENERALIZED GAMMA MECHANISM

To prove the achievability bound, we construct a mechanism whose MSE exceeds $A_M(\Delta_2)$ by at most $\mathcal{O}(\Delta_2^2 M^{-2})$. We refer to it as the closed-form generalized gamma (CGG) mechanism. It adds generalized gamma noise with the shape parameters of eq. (4).

Definition 5 (Closed-form generalized gamma mechanism)**.** *Let Assumptions 1 and 2 hold. The closed-form generalized gamma (CGG) mechanism releases*

$$\mathcal{M}^{\mathrm{CGG}}(\boldsymbol{x}) = f(\boldsymbol{x}) + Z_M^{\mathrm{CGG}}, \qquad Z_M^{\mathrm{CGG}} = R_M U,$$

where Z_M^{CGG} is generalized gamma noise (Definition 4) with shape $q = \rho_* M$, power $p = p_*$, and scale a_M. The scale a_M is the smallest for which $D_\epsilon(\mathcal{L}(Z_M^{\mathrm{CGG}}) \| \mathcal{L}(Z_M^{\mathrm{CGG}} + \Delta_2 e_1)) \le \delta$. Its error is $\mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2) := \mathbb{E}\left[\|Z_M^{\mathrm{CGG}}\|_2^2\right]$.

The effective shape ρp equals τ_*, which by Step 5 of Theorem 2 carries the constant-order Gaussian excess $C(\epsilon, \delta) = \frac{1}{2}\epsilon^2 v_0^2 (2 - \tau_*)^2$. The power p equals p_*, which by Step 6 minimizes the $\frac{1}{M}$ coefficient. Since both τ_* and p_* are closed-form functions of (ϵ, δ), the mechanism needs no numerical search over the shape. Algorithm 1 generates the noise.

Algorithm 1 The CGG Mechanism

Require: $\epsilon > 0$, $\delta \in (0, 1)$ with Assumptions 1 and 2, $\Delta_2 > 0$, $M \ge 2$, query $f(\boldsymbol{x}) \in \mathbb{R}^M$
Ensure: Private output $\tilde{f}(\boldsymbol{x}) \in \mathbb{R}^M$
1: Compute v_0 from $g(v_0) = \delta$, and evaluate τ_*, p_*, ρ_* via eq. (4) ▷ Once per (ϵ, δ)
2: Set $q \leftarrow \rho_* M$ and $p \leftarrow p_*$
3: Find a_M with $D_\epsilon(\mathcal{L}(Z) \| \mathcal{L}(Z + \Delta_2 e_1)) = \delta$ by root-finding on the scale a of generalized gamma noise Z with shape q and power p ▷ Once per (M, Δ_2)
4: Sample $T \sim \mathrm{Gamma}(q/p, 1)$ and U uniform on $\mathbb{S}^{M-1}$ ▷ $U = N/\|N\|_2$, $N \sim \mathcal{N}(0, \mathbf{I}_M)$
5: **return** $\tilde{f}(\boldsymbol{x}) = f(\boldsymbol{x}) + a_M T^{1/p} U$

Steps 1 and 3 each solve one scalar equation, and Step 4 samples exactly because $T = (\frac{R}{a_M})^p \sim \mathrm{Gamma}(\frac{q}{p}, 1)$. The mechanism is (ϵ, δ)-DP for every $M \ge 2$. Under Assumption 2, its density $\propto \|z\|_2^{(\rho_* - 1)M} e^{-(\|z\|_2 / a_M)^{p_*}}$ is nonincreasing in $\|z\|_2$ and privacy at the single shift of Step 3 covers the whole ball (Lemma 10). $\mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2)$ is closed form in a_M and two gamma functions (Appendix E.6).

3.3 ACHIEVABILITY BOUND

Theorem 3 (Achievability bound)**.** *Under Assumptions 1 and 2, the noise Z_M^{CGG} of Algorithm 1 is feasible for eq. (1) at every $M \ge 2$, and there are finite constants $K_{\mathrm{ach}}, M_{\mathrm{ach}}$ that depend only on (ϵ, δ) such that for every $M \ge M_{\mathrm{ach}}$*

$$\mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2) \le A_M(\Delta_2) + \frac{\Delta_2^2 K_{\mathrm{ach}}}{M^2} = A_M(\Delta_2) + \mathcal{O}(\Delta_2^2 M^{-2}). \tag{6}$$

Proof sketch. Appendix E gives the full proof in three steps.

Step 1 (Feasibility, Lemma 10 and Corollary 3). Under Assumption 2, $q = \rho_* M \le M$, so the density of Z_M^{CGG} is nonincreasing in $\|z\|_2$ and privacy at $\Delta_2 e_1$ extends to every $\|h\|_2 \le \Delta_2$ (Lemma 10). The single-shift profile is continuous and nonincreasing in the scale with limits 1 and 0, so a_M exists (Corollary 3). Hence Z_M^{CGG} is a feasible solution of eq. (1) at every $M \ge 2$.

Step 2 (Uniform expansion, Lemmas 16 and 17). We expand the single-shift profile of generalized gamma noise with shape (ρ, p) to order M^{-2}, uniformly for (ρ, p) in compact sets that include $\rho = 1$, and invert the calibration equation. We prove that the calibrated mean energy is $v_0 - \frac{C_\tau(v_0)}{M} + \frac{J_{\tau,p}(v_0)}{M^2} + \mathcal{O}(M^{-3})$ with $\tau = \rho p$, where $C_\tau(v_0)$ is the constant-order correction of effective shape τ and $J_{\tau,p}(v_0)$ is a quadratic in p, both explicit rational functions of (ϵ, v_0, τ, p) (Proposition 2). Multiplying by $M\Delta_2^2$ gives the calibrated error of every (ρ, p) in the family (Proposition 3).

Step 3 (Choice of shape, Proposition 2). $C_\tau(v_0)$ is strictly concave in τ with maximizer τ_* and maximum $C(\epsilon,\delta)$. At $\tau=\tau_*$, the quadratic $J_{\tau_*,p}(v_0)$ has minimizer p_* and minimum $J(\epsilon,\delta)$. Setting $(q,p)=(\rho_* M, p_*)$ in Proposition 3 proves eq. (6). □

Corollary 1 (Sharp expansion of the optimal error)**.** *Under Assumptions 1 and 2, there are finite K, M_0 such that for every $M \geq M_0$*

$$A_M(\Delta_2) - \frac{\Delta_2^2 K}{M^2} \leq \mathcal{V}_M(\Delta_2) \leq \mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2) \leq A_M(\Delta_2) + \frac{\Delta_2^2 K}{M^2}. \tag{7}$$

Hence $\mathcal{V}_M(\Delta_2) = A_M(\Delta_2) + \mathcal{O}(\Delta_2^2 M^{-2})$ and $0 \leq \mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2) - \mathcal{V}_M(\Delta_2) \leq \frac{2\Delta_2^2 K}{M^2}$.

Proof sketch. Combine Theorems 2 and 3 with the feasibility of Z_M^{CGG}, taking $K = \max\{K_{\mathrm{cv}}, K_{\mathrm{ach}}\}$ and $M_0 = \max\{M_{\mathrm{cv}}, M_{\mathrm{ach}}\}$. The two $\frac{1}{M}$ coefficients coincide because the maximum in Step 6 of Theorem 2 and the minimum in Step 3 of Theorem 3 are the same number $J(\epsilon,\delta)$ (Appendix F.2). □

Corollary 2 (Gaussian excess)**.** *Let $\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) = M\Delta_2^2 v_0$ be the error of the exactly calibrated Gaussian mechanism (Balle & Wang, 2018) and call $\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) - \mathcal{V}_M(\Delta_2)$ the Gaussian excess. Under Assumptions 1 and 2,*

$$\frac{\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) - \mathcal{V}_M(\Delta_2)}{\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2)} = \frac{C(\epsilon,\delta)}{M v_0} - \frac{J(\epsilon,\delta)}{M^2 v_0} + \mathcal{O}(M^{-3}). \tag{8}$$

Under Assumption 1 alone the left side is at most $\frac{C(\epsilon,\delta)}{M v_0} + \frac{K}{M^2 v_0}$ for a finite K and all $M \geq M_{\mathrm{cv}}$.

Proof sketch. By eq. (7), $\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) - \mathcal{V}_M(\Delta_2) = \Delta_2^2 C(\epsilon,\delta) - \frac{\Delta_2^2 J(\epsilon,\delta)}{M} + \mathcal{O}(\Delta_2^2 M^{-2})$. Dividing by $\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) = M\Delta_2^2 v_0$ proves eq. (8). For the second statement we use only eq. (5), which bounds $\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) - \mathbb{E}_P\|Z\|_2^2$ above by $\Delta_2^2 C(\epsilon,\delta) + \mathcal{O}(\Delta_2^2 M^{-1})$ for every feasible P, and divide again (Appendix F.3). The case $C(\epsilon,\delta)=0$, that is $\tau_*=2$, is Corollary 5. □

The Gaussian mechanism is the member $(\rho,p)=(1,2)$ of the family, so $C_2(v_0)=0$ and $C(\epsilon,\delta) \geq 0$ is the advantage of the best effective shape over the Gaussian one. Both shape parameters are needed to reach it. The ℓ_2 mechanism $(\rho,p)=(1,1)$ misses the constant order. The power-radial subfamily $\rho=1$ with $p=\tau_*$ matches the constant order but misses the $\frac{1}{M}$ coefficient by $B_* = \frac{1}{2}\epsilon^2 v_0^2 (p_*-\tau_*)^2$, which is positive unless $p_*=\tau_*$ (Corollary 6). All constants depend on (ϵ,δ) only through ϵ and v_0. Table 5 lists them at the four privacy pairs of Section 4.

The spherical generalized gamma (SGG) mechanism of Wei et al. (2026) chooses (q,p) by minimizing the calibrated error at each (ϵ,δ,M). It repeats Step 3 of Algorithm 1 once per candidate (q,p) and returns no bound on its distance to the optimum. Its search domain contains the CGG parameters. By Corollary 1, the minimum over that domain lies below $\mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2)$ by at most $\frac{2\Delta_2^2 K}{M^2}$. Section 4.2 measures the difference.

3.4 Relaxing Assumption 2

When $p_* < \tau_*$, the generalized gamma noise that attains both $C(\epsilon,\delta)$ and $J(\epsilon,\delta)$ has $\rho > 1$ and a nonmonotone density. Without monotonicity, the divergence need not increase with the shift length, so privacy at the boundary shift does not cover the interior of the ball. Among choices with a monotone density and $\tau=\tau_*$, the smallest $\frac{1}{M}$ coefficient is attained at $p=\tau_*$, that is, at $\rho=1$. By eq. (16) it exceeds $J(\epsilon,\delta)$ by $B_{\mathrm{sh}} := \frac{1}{2}\epsilon^2 v_0^2 (\tau_*-p_*)_+^2$, which vanishes exactly under Assumption 2.

Theorem 4 (Bounds without Assumption 2)**.** *Under Assumption 1 alone, there are finite K, M_0 that depend only on (ϵ,δ) such that for every $M \geq M_0$,*

$$A_M(\Delta_2) - \frac{\Delta_2^2 K}{M^2} \leq \mathcal{V}_M(\Delta_2) \leq A_M(\Delta_2) + \frac{\Delta_2^2 B_{\mathrm{sh}}}{M} + \frac{\Delta_2^2 K}{M^2}. \tag{9}$$

Hence the constant order of $\mathcal{V}_M(\Delta_2)$ is sharp under Assumption 1 alone. Every limit point of the sequence $\frac{M}{\Delta_2^2}\{\mathcal{V}_M(\Delta_2) - M\Delta_2^2 v_0 + \Delta_2^2 C(\epsilon,\delta)\}$ lies in $[J(\epsilon,\delta), J(\epsilon,\delta)+B_{\mathrm{sh}}]$.

Appendix F.1 gives the proof. The lower bound is Theorem 2 and the upper bound is Proposition 3 at $(\rho, p) = (1, \tau_*)$. Appendix E.5 maps the region of the privacy plane where both assumptions hold, which we call the matching window. Assumption 2 fails for large ϵ at small δ (Table 4). Outside the window, Theorem 4 brackets the $\frac{1}{M}$ coefficient within a width B_{sh} that stays below 0.7% of $J(\epsilon, \delta)$ for $\delta = 10^{-5}$ and $\epsilon \leq 8$ (Remark 4).

4 Numerical results

4.1 Numerical converse and achievability bounds

We take $\epsilon \in \{0.5, 1\}$, $\delta \in \{10^{-5}, 10^{-3}\}$, and $\Delta_2 = 1$. Both assumptions hold strictly at all four pairs (Table 5). Table 1 reports numerical converse and achievability bounds $[\underline{\mathrm{MSE}}_M, \overline{\mathrm{MSE}}_M]$ on $\mathcal{V}_M(1)$, both checked by interval arithmetic (Johansson, 2017). The converse bound applies Lemma 5 with a threshold solved numerically at each dimension. The achievability bound is the error of the CGG mechanism at a slightly inflated scale whose privacy profile is checked below δ (Appendix H.5).

Table 1 confirms the three orders of eq. (7). The Gaussian excess $\mathrm{MSE}_M^{\mathrm{AGM}} - \overline{\mathrm{MSE}}_M$ converges to $C(\epsilon, \delta)$. At $(0.5, 10^{-5})$ it is 178.1, 232.6, 238.91, and 239.53 for $M = 10$ to 10^4, against $C(0.5, 10^{-5}) = 239.60$ in Table 5, and at $(1, 10^{-5})$ it is 55.43, 75.44, 77.85, and 78.08 against 78.11. Relative to $\mathrm{MSE}_M^{\mathrm{AGM}}$, the excess falls tenfold per decade of M, the $\frac{C(\epsilon,\delta)}{M v_0}$ rate of Corollary 2, from 39.8% at $M = 10$ to 5.42%, 0.559%, and 0.0561% at $(1, 10^{-5})$. The interval width bounds the gap of CGG to the optimum. At $(0.5, 10^{-5})$ it shrinks from 9.0 at $M = 10$ to 0.81, 7×10^{-4}, and 3×10^{-5}, so CGG leaves at most 5% of the excess at $M = 10$ and less than 10^{-5} of it at $M \geq 10^3$, the $\mathcal{O}(M^{-2})$ gap of Corollary 1.

M	δ	$\varepsilon = 0.5$		$\varepsilon = 1.0$	
		$\mathrm{MSE}_M^{\mathrm{AGM}}$	$[\underline{\mathrm{MSE}}_M, \overline{\mathrm{MSE}}_M]$	$\mathrm{MSE}_M^{\mathrm{AGM}}$	$[\underline{\mathrm{MSE}}_M, \overline{\mathrm{MSE}}_M]$
low-dimension converse (Appendix H.4)					
10^1	10^{-5}	494.466	[307.309, 316.324]	139.176	[80.808, 83.750]
	10^{-3}	212.533	[180.134, 182.540]	66.289	[53.104, 53.793]
10^2	10^{-5}	4944.659	[4711.254, 4712.060]	1391.761	[1316.053, 1316.322]
	10^{-3}	2125.328	[2093.343, 2093.776]	662.886	[648.946, 649.081]
high-dimension dual-event converse (Theorem 2, Appendix H.5)					
10^3	10^{-5}	49446.58640	[49207.67306, 49207.67376]	13917.61239	[13839.76612, 13839.76642]
	10^{-3}	21253.27972	[21221.52831, 21221.52840]	6628.85876	[6614.91053, 6614.91057]
10^4	10^{-5}	494465.86395	[494226.33297, 494226.33300]	139176.12395	[139098.03922, 139098.03923]
	10^{-3}	212532.79722	[212501.03491, 212501.03493]	66288.58764	[66274.62797, 66274.62798]

Table 1: Bounds on $\mathcal{V}_M(1)$ against $\mathrm{MSE}_M^{\mathrm{AGM}}$ at $\Delta_2 = 1$, rounded outward to three decimals in the upper block and five in the lower. Lower endpoints come from the dual construction named in each block header. Upper endpoints are the CGG mechanism at a rounded shape and a slightly inflated scale (Appendices H.4 and H.5).

extends the relative Gaussian excess from the four pairs to a 25×25 grid over $0.1 \leq \epsilon \leq 1$ and $10^{-6} \leq \delta \leq 10^{-2}$. The percentage of the Gaussian error that CGG removes grows as δ decreases and ϵ increases, where $C(\epsilon, \delta)$ is largest, and peaks at $(1, 10^{-6})$ with 48.0%, 7.20%, and 0.751% at $M = 10$, 10^2, and 10^3. At $(1, 10^{-5})$ the three panels read 39.8%, 5.42%, and 0.559%, as in Table 1. Averaged over the grid, the percentage falls by a factor of 7.9 from $M = 10$ to 10^2 and 9.8 from 10^2 to 10^3, and at $M = 10^3$ it matches $\frac{100 C(\epsilon,\delta)}{M v_0}$ within 0.01 percentage points at every grid point.

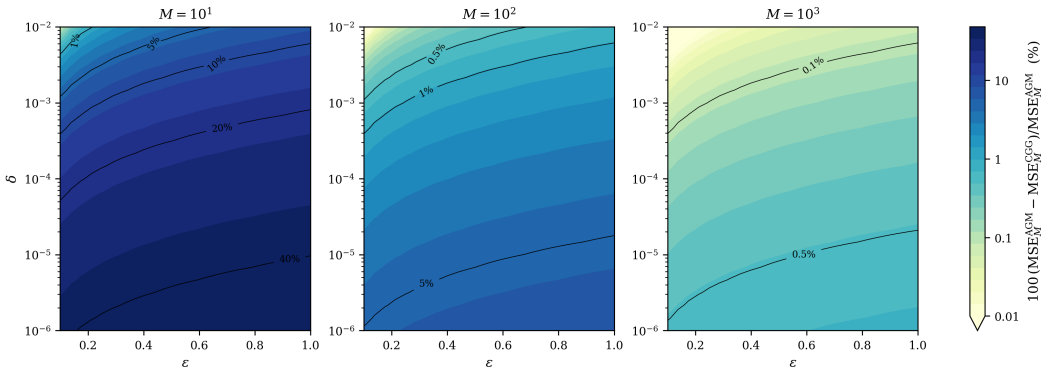


Figure 1: Percentage of the Gaussian error removed by CGG, $\frac{100(\mathrm{MSE}_M^{\mathrm{AGM}} - \mathrm{MSE}_M^{\mathrm{CGG}})}{\mathrm{MSE}_M^{\mathrm{AGM}}}$, over the privacy plane at $M \in \{10, 10^2, 10^3\}$.

4.2 The CGG mechanism against search and fixed shapes

Table 2 reports high-precision MSE values of four mechanisms at the same privacy pairs and $M \in \{10, 10^2, 10^3\}$. The analytic Gaussian mechanism (Balle & Wang, 2018) and the ℓ_2 mechanism (Joseph et al., 2025) are generalized gamma noise with $(q, p) = (M, 2)$ and $(M, 1)$. The other two are the SGG mechanism of Section 3.3 and the CGG mechanism. Resolving their $\mathcal{O}(M^{-2})$ differences needs the minimizer of the calibrated error to five decimals, so in place of the Bayesian optimization of Wei et al. (2026) we minimize the calibrated error directly over their box $0 < \rho \le 1$, $1 \le p \le 3$ by a 20×20 grid and bounded simplex refinement (Appendix H.2). Every scale is the root of the exact single-shift profile in Appendix E.2, computed by one routine for all four mechanisms.

ε	δ	M	MSE				Running Time (s)		
			AGM	ℓ_2	SGG	CGG	SGG	CGG	Ratio
0.5	10^{-5}	10^1	494.46586	317.41461	314.72680	316.31096	5.3	0.006	838
		10^2	4944.65864	4719.07075	4711.85539	4711.87138	8.8	0.013	671
		10^3	49446.58640	49215.61364	49207.67354	49207.67375	6.1	0.010	616
	10^{-3}	10^1	212.53280	187.17611	182.37533	182.53242	5.8	0.011	523
		10^2	2125.32797	2100.49103	2093.68987	2093.69153	6.6	0.010	681
		10^3	21253.27972	21228.55833	21221.52838	21221.52839	4.9	0.008	584
1.0	10^{-5}	10^1	139.17612	83.79809	83.23497	83.74623	5.8	0.010	603
		10^2	1391.76124	1318.04214	1316.26211	1316.26896	6.9	0.012	563
		10^3	13917.61239	13841.75463	13839.76634	13839.76641	5.9	0.010	610
	10^{-3}	10^1	66.28859	54.76165	53.70588	53.78993	5.9	0.010	563
		10^2	662.88588	650.77140	649.05342	649.05436	6.6	0.011	597
		10^3	6628.85876	6616.70970	6614.91056	6614.91057	5.4	0.009	593

Table 2: Exactly calibrated MSE and running time of four mechanisms at $\Delta_2 = 1$. The SGG column is the search run in Table 6, with 489 to 604 scale calibrations per cell versus one for CGG. Times are measured on one machine as one scale calibration, using the best of five repeats for AGM, ℓ_2, and CGG, and as one search run for SGG. The M-independent Step 1 of Algorithm 1 takes 13 to 17 milliseconds per privacy pair and is excluded. The CGG column is the uninflated error, hence the last-digit differences from Table 1 at $M = 10^3$ (Appendix H.5).

Table 2 separates the three levels of Corollary 6. The ℓ_2 mechanism misses the constant order. Its excess over CGG at $(0.5, 10^{-5})$ is 1.10, 7.20, and 7.94 for $M = 10$ to 10^3, approaching the limit $\epsilon^2 v_0^2 (\tau_* - 1)^2 = 8.02$, and at $(1, 10^{-3})$ it is 0.97, 1.72, and 1.80 against 1.81. The search improves on CGG by $\mathcal{O}(M^{-2})$ only. At $(0.5, 10^{-5})$ the SGG column lies below the CGG column by 1.58, 0.016,

and 0.00021, a hundredfold drop per decade, and at $(1, 10^{-3})$ by 0.084, 0.00094, and 0.0000095. Scaled by M^2, the twelve differences stay between 8 and 207, as Corollary 1 bounds them, and at $M = 10^3$ they are below a relative 10^{-8}. This improvement costs two orders of magnitude in setup time. CGG performs one scale calibration per cell in 6 to 13 milliseconds, while the search repeats it for 489 to 604 candidates in 4.9 to 8.8 seconds, the ratio of 523 to 838 in the last column, or 211 to 308 once the M-independent Step 1 of Algorithm 1 is counted on the CGG side. Appendix H.2 records the searched shapes, which approach (ρ_*, p_*) at rate M^{-1}, and a simplex warm-started at the CGG shape that removes the same $\mathcal{O}(M^{-2})$ amount with 95 to 239 calibrations.

5 Related work

Noise under different optimality. Several noise distributions have been proposed to achieve certain optimality, and the defined optimality differs by scenario. Under ϵ-DP, the staircase distribution, a piecewise-constant density whose steps decay geometrically by the factor $e^{-\epsilon}$, is optimal for a scalar query under certain symmetric costs (Geng & Viswanath, 2016b), and the same shape is optimal for vector queries under an arbitrary norm (Kulesza et al., 2025; Melbourne et al., 2026) and discrete queries (Joseph et al., 2026). Under $(0, \delta)$-DP, the optimal noise is uniform with an atom at the origin (Geng & Viswanath, 2016a; Geng et al., 2019). For a scalar query under (ϵ, δ)-DP, a truncated Laplace distribution is asymptotically optimal as $\epsilon \to 0$ (Geng et al., 2020), and Gaussian mixtures have been calibrated to the same constraint (Liu et al., 2026a). Optimal noise has also been studied under Rényi divergence (Gilani et al., 2025) and f-DP (Awan & Dong, 2022). In this work, we focus on (ϵ, δ)-DP for vector queries and the optimality defined by eq. (1).

Noise calibration under ℓ_2 sensitivity. Various works calibrate better noise under ℓ_2 sensitivity, such as the K-norm mechanisms (Hardt & Talwar, 2010; Awan & Slavković, 2021), coordinate-wise generalized Gaussian noise (Liu, 2019; Zhang et al., 2026; Rinberg et al., 2026), symmetric log-concave densities (Vinterbo, 2022), and noise derived from a target privacy loss random variable (Yang et al., 2025). The resulting noise is not necessarily spherically symmetric. Among spherically symmetric noise, recent developments include the rank-one singular Gaussian (Ji & Li, 2024), ℓ_2 (Joseph et al., 2025), and spherical generalized gamma (Wei et al., 2026) mechanisms, whose radii all follow a generalized gamma distribution (Liu et al., 2026b).

6 Discussion and limitations

Our setting is one data-independent additive release, the full ℓ_2 shift ball, the mean squared error of the raw noise, and a fixed privacy pair with $\epsilon, \delta > 0$. The optimum within it is determined through the constant and $\frac{1}{M}$ orders, so a constant-factor improvement must change one of these ingredients, for example a smaller or anisotropic sensitivity set, data-dependent adaptation, composition over many releases, a task-specific loss, or postprocessing. Carrying $A_M(\Delta_2)$ over to DP-SGD or to end-to-end learning requires a further theorem. The thresholds M_{cv} and M_{ach} are not explicit, since the proofs track the expansion remainders only up to constants. In practice, the numerical bounds of Section 4.1 take about two minutes at any given (ϵ, δ, M) and enclose $\mathcal{V}_M$ wherever Assumption 1 holds (Appendix H.5).

7 Conclusion

In this paper, we established converse and achievability bounds on the optimal MSE of data-independent additive mechanisms under (ϵ, δ)-differential privacy with ℓ_2 sensitivity Δ_2. Our converse shows that every feasible additive noise distribution incurs MSE at least $A_M(\Delta_2) - \mathcal{O}(\Delta_2^2 M^{-2})$. This provides a universal limit on how much noise design can reduce the error while preserving privacy. Our achievability result constructs noise with MSE at most $A_M(\Delta_2) + \mathcal{O}(\Delta_2^2 M^{-2})$. The noise calibration procedure determines the shape parameters in closed form, avoiding numerical shape search while attaining the optimal error to the stated accuracy. When the shape assumption required to generate the noise is relaxed, the achievable upper bound increases by only $\mathcal{O}(\Delta_2^2 M^{-1})$.

AI USE STATEMENT

In accordance with the ICLR 2027 AI policy, we disclose that we used large language models (such as Claude and ChatGPT) for two purposes. An LLM helped polish the prose, mainly in the introduction and the related work, by rephrasing sentences, correcting grammar, and making the exposition easier to follow. An LLM also helped locate related work, and we checked each reference it surfaced against the published source before citing it. The authors wrote the code that produces our numerical results and used an LLM only to check it.

The theory, proofs, experiments, and their interpretation are the authors' own work, and no LLM generated any reported numbers. We reviewed all AI-assisted text and take responsibility for the final content of this paper.

ETHICS STATEMENT

This work is theoretical. It involves no human subjects, personal or proprietary data, or data set release, and the reported numbers come from deterministic computations that depend only on the privacy parameters and the dimension. One application-side risk follows from the results. The mechanism of Definition 5 is (ϵ, δ)-DP because Assumption 2 makes its density nonincreasing in $\|z\|_2$, so a calibration at the single shift $\Delta_2 e_1$ covers the whole sensitivity ball (Lemma 10). Outside the region mapped in Appendix E.5, that argument fails, and an implementation calibrated to the same family at one shift there would release data without the stated guarantee. The authors declare no conflicts of interest.

REPRODUCIBILITY STATEMENT

Section 3 states the two assumptions the results rest on, and Appendix E.5 maps where each holds. Complete proofs are in Appendix D for the converse, Appendix E for the achievability bound, and Appendix F for Corollary 1 and the remaining corollaries. Appendix G states what the bounds at a given (ϵ, δ, M) consist of and what they establish, and Appendix H.5 records the data behind every entry of Table 1 together with the procedure that rechecks it at higher working precision. The supplementary material includes those records and the code that rebuilds and rechecks them, so every reported endpoint can be recomputed from the record alone. The same code regenerates every computed table and figure in the paper.

REFERENCES

Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pp. 308–318, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450341394. doi: 10.1145/2976749.2978318. URL https://doi.org/10.1145/2976749.2978318.

Edward J. Anderson and Peter Nash. *Linear Programming in Infinite-Dimensional Spaces: Theory and Applications*. John Wiley & Sons, Chichester, 1987.

Jordan Awan and Jinshuo Dong. Log-concave and multivariate canonical noise distributions for differential privacy. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 34229–34240. Curran Associates, Inc., 2022. doi: 10.52202/068431-2481. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/dd73933d99ccd7ffe2306adb95ec5d02-Paper-Conference.pdf.

Jordan Awan and Aleksandra Slavković. Structure and sensitivity in differential privacy: Comparing k-norm mechanisms. *Journal of the American Statistical Association*, 116(534):935–954, 2021. doi: 10.1080/01621459.2020.1773831. URL https://doi.org/10.1080/01621459.2020.1773831.

Borja Balle and Yu-Xiang Wang. Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In Jennifer Dy and Andreas Krause (eds.), *Pro-*

ceedings of the 35th International Conference on Machine Learning, volume 80 of *Proceedings of Machine Learning Research*, pp. 394–403. PMLR, 10–15 Jul 2018. URL https://proceedings.mlr.press/v80/balle18a.html.

Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in real algebraic geometry*. Springer, 2006.

Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 02 2013. ISBN 9780199535255. doi: 10.1093/acprof:oso/9780199535255.001.0001. URL https://doi.org/10.1093/acprof:oso/9780199535255.001.0001.

Persi Diaconis and David Freedman. A dozen de Finetti-style results in search of a theory. *Annales de l'I.H.P. Probabilités et statistiques*, 23(S2):397–423, 1987. URL https://www.numdam.org/item/AIHPB_1987__23_S2_397_0/.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4):211–407, August 2014. ISSN 1551-305X. doi: 10.1561/0400000042. URL https://doi.org/10.1561/0400000042.

Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In Serge Vaudenay (ed.), *Advances in Cryptology - EUROCRYPT 2006*, pp. 486–503, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. ISBN 978-3-540-34547-3.

K.T. Fang, S. Kotz, and K.W. Ng. *Symmetric Multivariate and Related Distributions*. Monographs on Statistics and Applied Probability. Springer US, 2013. ISBN 9781489929389. URL https://books.google.com/books?id=BU7-DAEACAAJ.

G.B. Folland. *Real Analysis: Modern Techniques and Their Applications*. Pure and Applied Mathematics: A Wiley Series of Texts, Monographs and Tracts. Wiley, 2013. ISBN 9781118626399. URL https://books.google.com/books?id=wI4fAwAAQBAJ.

Quan Geng and Pramod Viswanath. Optimal noise adding mechanisms for approximate differential privacy. *IEEE Transactions on Information Theory*, 62(2):952–969, 2016a. doi: 10.1109/TIT.2015.2504972.

Quan Geng and Pramod Viswanath. The optimal noise-adding mechanism in differential privacy. *IEEE Transactions on Information Theory*, 62(2):925–951, 2016b. doi: 10.1109/TIT.2015.2504967.

Quan Geng, Wei Ding, Ruiqi Guo, and Sanjiv Kumar. Optimal noise-adding mechanism in additive differential privacy. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 11–20. PMLR, 16–18 Apr 2019. URL https://proceedings.mlr.press/v89/geng19a.html.

Quan Geng, Wei Ding, Ruiqi Guo, and Sanjiv Kumar. Tight analysis of privacy and utility tradeoff in approximate differential privacy. In Silvia Chiappa and Roberto Calandra (eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 89–99. PMLR, 26–28 Aug 2020. URL https://proceedings.mlr.press/v108/geng20a.html.

Atefeh Gilani, Juan Felipe Gomez, Shahab Asoodeh, Flavio Calmon, Oliver Kosut, and Lalitha Sankar. Optimizing noise distributions for differential privacy. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 19505–19522. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/gilani25a.html.

Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, STOC '10, pp. 705–714, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781450300506. doi: 10.1145/1806689.1806786. URL https://doi.org/10.1145/1806689.1806786.

Tianxi Ji and Pan Li. Less is more: Revisiting the gaussian mechanism for differential privacy. In *33rd USENIX Security Symposium (USENIX Security 24)*, pp. 937–954, Philadelphia, PA, August 2024. USENIX Association. ISBN 978-1-939133-44-1. URL https://www.usenix.org/conference/usenixsecurity24/presentation/ji.

Fredrik Johansson. Arb: Efficient arbitrary-precision midpoint-radius interval arithmetic. *IEEE Transactions on Computers*, 66(8):1281–1292, 2017. doi: 10.1109/TC.2017.2690633.

Matthew Joseph, Alex Kulesza, and Alexander Yu. Approximate differential privacy of the ℓ_2 mechanism. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 28377–28392. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/joseph25a.html.

Matthew Joseph, Alex Kulesza, Yuyan Wang, and Alexander Yu. Ripple mechanisms for discrete and private statistics. In Steve Hanneke and Tor Lattimore (eds.), *Proceedings of Thirty Ninth Conference on Learning Theory*, volume 336 of *Proceedings of Machine Learning Research*, pp. 3856–3903. PMLR, 29 Jun–03 Jul 2026. URL https://proceedings.mlr.press/v336/joseph26a.html.

S.G. Krantz and H.R. Parks. *The Implicit Function Theorem: History, Theory, and Applications*. World Publishing Corporation, 2012. ISBN 9787510048036. URL https://books.google.com/books?id=ZW9z0gEACAAJ.

Alex Kulesza, Ananda Theertha Suresh, and Yuyan Wang. General staircase mechanisms for optimal differential privacy. In Yingzhen Li, Stephan Mandt, Shipra Agrawal, and Emtiyaz Khan (eds.), *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp. 4564–4572. PMLR, 03–05 May 2025. URL https://proceedings.mlr.press/v258/kulesza25a.html.

B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *The Annals of Statistics*, 28(5):1302–1338, 2000. ISSN 00905364, 21688966. URL http://www.jstor.org/stable/2674095.

Fang Liu. Generalized gaussian mechanism for differential privacy. *IEEE Transactions on Knowledge and Data Engineering*, 31(4):747–756, 2019. doi: 10.1109/TKDE.2018.2845388.

Huikang Liu, Aras Selvi, and Wolfram Wiesemann. Mind the gap: Mixtures of gaussians in approximate differential privacy. In *Forty-third International Conference on Machine Learning*, 2026a. URL https://openreview.net/forum?id=RdBsmhb9U7.

Shuainan Liu, Tianxi Ji, Zhongshuo Fang, Lu Wei, and Pan Li. Privacy loss of noise perturbation via concentration analysis of a product measure. *Proc. ACM Manag. Data*, 4(1), April 2026b. doi: 10.1145/3786680. URL https://doi.org/10.1145/3786680.

James Melbourne, Mario Diaz, and Shahab Asoodeh. Optimality of staircase mechanisms for vector queries under differential privacy, 2026. URL https://arxiv.org/abs/2601.14597.

Jerzy Neyman and Egon Sharpe Pearson. On the problem of the most efficient tests of statistical hypotheses. In *Breakthroughs in Statistics: Foundations and Basic Theory*, pp. 73–108. Springer, 1992.

Frank Olver, Daniel Lozier, Ronald Boisvert, and Charles Clark. *The NIST Handbook of Mathematical Functions*. Cambridge University Press, New York, NY, 2010-05-12 00:05:00 2010.

Marc S Paolella. *Linear models and time-series analysis: regression, ANOVA, ARMA and GARCH*. John Wiley & Sons, 2018.

Roy Rinberg, Ilia Shumailov, Vikrant Singhal, Rachel Cummings, and Nicolas Papernot. Beyond laplace and gaussian: Exploring the generalized gaussian mechanism for private machine learning, 2026. URL https://arxiv.org/abs/2506.12553.

Igal Sason and Sergio Verdú. f -divergence inequalities. *IEEE Transactions on Information Theory*, 62(11):5973–6006, 2016. doi: 10.1109/TIT.2016.2603151.

E. W. Stacy. A generalization of the gamma distribution. *The Annals of Mathematical Statistics*, 33(3):1187–1192, 1962. ISSN 00034851. URL http://www.jstor.org/stable/2237889.

Francesco Giacomo Tricomi, Arthur Erdélyi, et al. The asymptotic expansion of a ratio of gamma functions. *Pacific J. Math*, 1(1):133–142, 1951.

Staal A. Vinterbo. Differential privacy for symmetric log-concave mechanisms. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera (eds.), *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 6270–6291. PMLR, 28–30 Mar 2022. URL https://proceedings.mlr.press/v151/vinterbo22a.html.

Yu Wei, Alexander Bienstock, and Antigoni Polychroniadou. Asymptotic optimality of the high-dimensional gaussian mechanism and improved low-dimensional mechanisms for differential privacy. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=82Wosp2Iu1.

Qin Yang, Nicholas Stout, Meisam Mohammady, Han Wang, Ayesha Samreen, Christopher J. Quinn, Yan Yan, Ashish Kundu, and Yuan Hong. Plrv-o: Advancing differentially private deep learning via privacy loss random variable optimization. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security*, CCS '25, pp. 306–320, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400715259. doi: 10.1145/3719027.3765151. URL https://doi.org/10.1145/3719027.3765151.

Xiang Zhang, Mohamedou Ould Haye, and Yiqiang Q. Zhao. Scale analysis and shape selection for the generalized gaussian mechanism under approximate differential privacy, 2026. URL https://arxiv.org/abs/2608.31138.

APPENDIX

A NOTATION AND PROOF DEPENDENCY MAP

A.1 NOTATION

Table 3 collects the symbols used in the main text and the appendices.

Table 3: Symbols used in the main text and in this appendix.

Symbol	Meaning
Problem data	
M	dimension of the released vector
Δ_2	ℓ_2 sensitivity of the query f
ϵ, δ	privacy parameters
$\boldsymbol{x}, \boldsymbol{x}'$	neighboring data sets
$\mathcal{M}(\boldsymbol{x}) = f(\boldsymbol{x}) + Z$	additive mechanism
Noise and its scales	
$Z = RU$	noise vector, radius $R = \|Z\|_2$, direction U uniform on $\mathbb{S}^{M-1}$ and independent of R; U_1 its first coordinate
$\bar{R} = \frac{R}{\Delta_2}$	sensitivity-normalized radius
$V = \frac{\bar{R}^2}{M} = \frac{\|Z\|_2^2}{M\Delta_2^2}$	normalized radial energy
v	a value of the normalized radial energy; for calibrated generalized gamma noise, its normalized mean energy $\mathbb{E}[V]$
$\mathcal{L}(X)$	distribution of a random variable X
P, P_h	distribution of Z and of $Z + h$
e_1	first coordinate vector; the shift direction used in calibration
Error functionals	
$D_\epsilon(P\|Q)$	hockey-stick divergence, eq. (2)
$\mathcal{V}_M(\Delta_2)$	optimal mean squared error (OMSE) over additive noise, eq. (1)
$A_M(\Delta_2)$	three-term expansion $M\Delta_2^2 v_0 - \Delta_2^2 C(\epsilon,\delta) + \frac{\Delta_2^2 J(\epsilon,\delta)}{M}$, Section 1
$\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2)$	error of the analytic Gaussian mechanism (AGM), $M\Delta_2^2 v_0$
$\mathrm{MSE}_M^{\rho,p}(\Delta_2)$	error of exactly calibrated generalized gamma noise with shape parameters (ρ, p)
$\mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2)$, $\mathrm{MSE}_M^{+}(\Delta_2)$	the same at (ρ_*, p_*) and at the clipped shape (ρ_+, p_+), $p_+ = \max\{p_*, \tau_*\}$, $\rho_+ = \frac{\tau_*}{p_+}$ (Appendix F)
$\mathrm{MSE}_M^{\ell_2}$, $\mathrm{MSE}_M^{\mathrm{PR}}$, $\mathrm{MSE}_M^{\mathrm{SGG}}$	errors of the ℓ_2 mechanism, the power-radial subfamily at $(q,p) = (M,\tau_*)$, and the spherical generalized gamma (SGG) mechanism of Wei et al. (2026)
$\mathrm{MSE}_M^{\mathrm{cand}}$, $\overline{\mathrm{MSE}}_M$, $\underline{\mathrm{MSE}}_M$	exact error of the achievability candidate; numerical achievability and converse bounds
Gaussian profile	
$g(v)$	Gaussian privacy profile in the normalized variance $v = \frac{\sigma^2}{\Delta_2^2}$, Section 3
v_0	Gaussian root, $g(v_0) = \delta$
v_c	inflection point of g, $v_c = \frac{\sqrt{9+\epsilon^2}-3}{2\epsilon^2}$
$\ell(v)$	tangent to g at v_0
$v_{\tan}$, v_{shape}; $\delta_{\tan}$, δ_{shape}	roots of the tangent and shape boundaries in v, and their images $g(v_{\tan})$, $g(v_{\mathrm{shape}})$ (Appendix E.5)
Generalized gamma mechanisms and shape constants	
a, q, p	scale, shape, and power of the generalized gamma radius, eq. (3)
$\beta = a^{-p}$	rate form of the scale, used in the derivations

Symbol	Meaning
$\rho = \frac{q}{M}$, $\tau = \rho p$	normalized shape and effective shape; $\rho = 1$ is the power-radial subfamily $q = M$
$\nu = \frac{\rho M}{p}$, T, X	shape of the radial gamma variable $T = \bar{\beta}\bar{R}^p \sim \mathrm{Gamma}(\nu, 1)$, and the centered energy $X = \frac{\bar{R}^2}{Mv} - 1$ (Appendix E.3)
Z_M^{CGG}	noise vector of the CGG mechanism, $q = \rho_* M$, $p = p_*$, scale a_M (Definition 5, Algorithm 1)
$L_h(z)$	privacy loss at shift h, eq. (13)
$a_1(v)$, $a_2(v)$	$\frac{1}{M}$- and $\frac{1}{M^2}$-coefficients of the single-shift privacy profile at normalized mean energy v (Lemma 16)
$C_\tau(v)$, $J_{\tau,p}(v)$	constant-order and $\frac{1}{M}$-order correction functions; $J_{\tau,p}(v) = c_0 + c_1 p + c_2 p^2$ with c_0, c_1, c_2 rational in (v, ϵ, τ) (Proposition 2)
$E = \epsilon^2 v_0^2$, $Q = E + 3v_0 - \frac{1}{4}$	auxiliary constants; $Q = -\frac{2v_0^2 g''(v_0)}{g'(v_0)}$
$\tau_* = \sqrt{\frac{Q}{E}}$, $C(\epsilon, \delta)$	optimal effective shape and constant-order Gaussian excess, eq. (4) and Appendix B.4
p_*, $J(\epsilon, \delta)$, $\rho_* = \frac{\tau_*}{p_*}$	optimal power, $\frac{1}{M}$ coefficient, and induced shape ratio, eq. (4) and eq. (15)
$B_* = \frac{E}{2}(p_* - \tau_*)^2$	$\frac{1}{M}$-order penalty of the power-radial subfamily
$B_{\mathrm{sh}} = \frac{E}{2}(\tau_* - p_*)_+^2$	width of the $\frac{1}{M}$ bracket in Theorem 4
$K_{\mathrm{cv}}, M_{\mathrm{cv}}$; $K_{\mathrm{ach}}, M_{\mathrm{ach}}$; K, M_0	finite constants of Theorems 2 and 3 and of Corollary 1, depending only on (ϵ, δ)
Dual construction	
θ, $\bar{S}_{M,\theta}$	threshold function and the event it defines
$h_M[\theta](v)$	hockey-stick mass of that event at normalized radial energy v
λ, κ	dual multiplier and pointwise support level in Lemma 5; λ_M, κ_M their values in the construction of Lemma 9
$\theta_G(v) = \frac{1}{2} - \epsilon v$, $\theta_M = T + \frac{u}{M} + \frac{b}{M^2}$	Gaussian threshold; corrected threshold of Lemma 9
$j(v, z)$	$M \to \infty$ limit of the event mass at threshold z
Θ	critical analytic branch, Lemma 7
u_0, u_*; a_{11}, a_{20}	amplitude of the first-order threshold correction and its optimal value; expansion coefficients of the event mass at v_0, eq. (31)
$H_M(y)$	distribution function of $\sqrt{M}U_1$ (Lemma 4)
$v_{\max}, r_{\max}$	right endpoints of the recorded partitions on the converse and achievability sides (Appendix H.5)

The appendices also use the following symbols. Φ and ϕ are the standard normal distribution and density functions. Γ is the gamma function and $\mathrm{Gamma}(k, 1)$ is the gamma distribution with shape k and unit scale. Subscripts as in j_z and j_{zz} denote partial derivatives. $x_M \sim y_M$ means $\frac{x_M}{y_M} \to 1$. $K \Subset I$ means that K is compact and contained in the interior of I. $O_{C^1}(\eta^k)$ denotes a remainder whose C^1 norm on the stated set is $\mathcal{O}(\eta^k)$.

A.2 Roadmap

Appendix B records the Gaussian profile, the generalized gamma noise density, and the correction constants. Appendix C collects the two uniform estimates that both arguments use. Appendix D proves the converse (Theorem 2). Appendix E proves the achievability bound (Theorem 3), covers the computation of the CGG mechanism, and maps where the two assumptions hold (Appendix E.5). The two proofs can be read in either order. Appendix F matches the bounds and proves the corollaries and Theorem 4. Appendix G explains how to read the numerical bounds. Appendix H collects the remaining numerical results together with the data behind every entry of Table 1, and Appendix I records the boundary cases.

B EXACT REDUCTIONS AND REFERENCE PROFILES

This section records the Gaussian privacy profile, the generalized gamma noise density in the coordinates used later, and the correction constants.

B.1 THE GAUSSIAN PRIVACY PROFILE

Let ϕ and Φ be the standard normal density and distribution function. For $Z \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_M)$, the left side of the inequality in Theorem 1 is the exact single-shift profile $D_\epsilon(\mathcal{L}(Z) \| \mathcal{L}(Z + \Delta_2 e_1))$ (Balle & Wang, 2018). By rotation invariance and monotonicity in the shift length (Lemma 10), it is also the full-sensitivity profile. In the normalized variance $v = \frac{\sigma^2}{\Delta_2^2}$ it equals

$$g(v) = \Phi(A(v)) - e^\epsilon \Phi(B(v)), \qquad A(v) = \frac{1}{2\sqrt{v}} - \epsilon\sqrt{v}, \quad B(v) = -\frac{1}{2\sqrt{v}} - \epsilon\sqrt{v}.$$

Since $A(v)^2 - B(v)^2 = -2\epsilon$ we have $\phi(A(v)) = e^\epsilon \phi(B(v))$. With $A'(v) - B'(v) = -\frac{1}{2}v^{-\frac{3}{2}}$ and $\phi'(x) = -x\phi(x)$ this gives

$$g'(v) = \phi(A(v))\{A'(v) - B'(v)\} = -\frac{\phi(A(v))}{2v^{\frac{3}{2}}} < 0, \qquad g''(v) = \frac{\phi(A(v))}{4v^{\frac{7}{2}}}\left(\epsilon^2 v^2 + 3v - \frac{1}{4}\right). \tag{10}$$

Hence g is strictly decreasing with $g(0^+) = 1$ and $g(\infty) = 0$. The equation $g(v_0) = \delta$ therefore has a unique root for every $\epsilon > 0$ and $\delta \in (0, 1)$. The factor $\epsilon^2 v^2 + 3v - \frac{1}{4}$ in g'' has the single positive root $v_c = \frac{\sqrt{9+\epsilon^2}-3}{2\epsilon^2}$, so g is strictly concave on $(0, v_c)$ and strictly convex on (v_c, ∞).

B.2 THE GENERALIZED GAMMA RADIAL FAMILY

In the rate parameterization $\beta = a^{-p}$, the density eq. (3) reads

$$f_R(r) = \frac{p\beta^{\frac{q}{p}}}{\Gamma(\frac{q}{p})} r^{q-1} e^{-\beta r^p}, \qquad r > 0,$$

and $T = \beta R^p \sim \mathrm{Gamma}(\frac{q}{p}, 1)$. Hence

$$\mathbb{E}[R^s] = \beta^{-\frac{s}{p}} \frac{\Gamma(\frac{q+s}{p})}{\Gamma(\frac{q}{p})}, \qquad s > -q. \tag{11}$$

The normalized radius $\bar{R} = \frac{R}{\Delta_2}$ has the same form with rate $\bar{\beta} = \beta\Delta_2^p$.

B.3 THE POWER-RADIAL SUBFAMILY AND THE JACOBIAN TERM

Let $\omega_{M-1} = \frac{2\pi^{\frac{M}{2}}}{\Gamma(\frac{M}{2})}$ be the surface area of $\mathbb{S}^{M-1}$ and $r = \|z\|_2$. Since $Z = RU$ with U uniform and $\mathrm{d}z = r^{M-1}\mathrm{d}r\mathrm{d}S(u)$, the density of Z on $\mathbb{R}^M$ is

$$f_{q,p,\beta}(z) = \frac{f_R(r)}{\omega_{M-1} r^{M-1}} = \frac{p\beta^{\frac{q}{p}}}{\omega_{M-1}\Gamma(\frac{q}{p})} r^{q-M} e^{-\beta r^p}. \tag{12}$$

For a shift $h \in \mathbb{R}^M$ and $s = \|z - h\|_2$, the privacy loss at z with $r, s > 0$ is

$$L_h(z) = \log \frac{f_{q,p,\beta}(z)}{f_{q,p,\beta}(z-h)} = (q - M)\log\frac{r}{s} + \beta(s^p - r^p). \tag{13}$$

The first term comes from the power r^{q-M}, the radial density r^{q-1} divided by the Jacobian r^{M-1}. In the power-radial subfamily $q = M$ it vanishes for every M, z, and h, the density is $\propto e^{-\beta\|z\|_2^p}$, and the loss reduces to $\beta(\|z - h\|_2^p - \|z\|_2^p)$. The full family keeps this term through $q \neq M$, which is the second degree of freedom that matches the $\frac{1}{M}$ coefficient.

B.4 THE TWO CORRECTION CONSTANTS

Throughout the appendices $E := \epsilon^2 v_0^2$ and $Q := \epsilon^2 v_0^2 + 3v_0 - \frac{1}{4}$. By eq. (10) $Q = -\frac{2v_0^2 g''(v_0)}{g'(v_0)}$, so eq. (4) reads $\tau_* = \sqrt{\frac{Q}{E}}$ and $C(\epsilon, \delta) = \frac{E}{2}(2 - \tau_*)^2$. The constant Q is positive if and only if $v_0 > v_c$.

Expanding the privacy profile of generalized gamma noise around g (Lemma 16) produces two functions of the normalized mean energy v. The constant-order coefficient depends on $\tau = \rho p$ alone,

$$C_\tau(v) = \frac{5}{2}\epsilon^2 v^2 + \frac{3}{2}v - \frac{1}{8} - \epsilon^2 v^2 \tau - \frac{\epsilon^2 v^2 + 3v - \frac{1}{4}}{\tau}. \tag{14}$$

The $\frac{1}{M}$ coefficient $J_{\tau,p}(v)$ of eq. (58) depends on p as well. By Proposition 2, it is a quadratic in p with positive leading coefficient $c_2(v, \tau)$ and rational coefficients c_0, c_1, given in eq. (60) and eq. (59).

At $v = v_0$, the function eq. (14) is $C_\tau(v_0) = \frac{5}{2}E + \frac{3}{2}v_0 - \frac{1}{8} - E\tau - \frac{Q}{\tau}$ with $\partial_\tau^2 C_\tau(v_0) = -\frac{2Q}{\tau^3}$, which is negative under Assumption 1. Hence $\tau \mapsto C_\tau(v_0)$ is strictly concave with maximizer τ_* and maximum $C(\epsilon, \delta)$. At $\tau = \tau_*$ one has $c_2(v_0, \tau_*) = \frac{Q}{2\tau_*^2} = \frac{E}{2}$, so $J_{\tau_*,p}(v_0) = c_0(v_0, \tau_*) + c_1(v_0, \tau_*)p + \frac{E}{2}p^2$ and the constants p_*, ρ_* of eq. (4) together with $J(\epsilon, \delta)$ are

$$p_* = -\frac{c_1(v_0, \tau_*)}{E}, \qquad J(\epsilon, \delta) = c_0(v_0, \tau_*) - \frac{c_1(v_0, \tau_*)^2}{2E}, \qquad \rho_* = \frac{\tau_*}{p_*}. \tag{15}$$

The cost of a departure from either optimum is given by

$$C(\epsilon, \delta) - C_\tau(v_0) = \frac{E(\tau - \tau_*)^2}{\tau}, \qquad J_{\tau_*,p}(v_0) = J(\epsilon, \delta) + \frac{E}{2}(p - p_*)^2. \tag{16}$$

The first identity uses $Q = E\tau_*^2$. The second is a completed square in p at fixed $\tau = \tau_*$. The function $J_{\tau,p}(v_0)$ is stationary in p at p_* but not in τ, because c_0, c_1, c_2 all depend on τ. Two values of the second identity recur. The first is $B_* := \frac{E}{2}(p_* - \tau_*)^2$, the $\frac{1}{M}$ penalty of the power-radial subfamily $q = M$, where $\rho = 1$ forces $p = \tau_*$. The second is $B_{\mathrm{sh}} := \frac{E}{2}(\tau_* - p_*)_+^2$, the penalty of the clipped power $p_+ := \max\{p_*, \tau_*\}$, which differs from p_* only when $p_* < \tau_*$.

Remark 1 (The constants belong to the Gaussian profile)**.** *By eq. (4), eq. (15), and Proposition 2, τ_*, $C(\epsilon, \delta)$, p_*, and $J(\epsilon, \delta)$ are explicit functions of ϵ and v_0 alone. Theorem 2 therefore establishes the same constants without reference to any family of mechanisms.*

C UNIFORM ANALYTIC ESTIMATES

Both arguments use two dimension-uniform estimates that are independent of the mechanism. The first is a two-term expansion of a ratio of gamma functions. The second is a two-term normal approximation to the distribution function of one coordinate of a uniform point on the sphere. The normal limit of a sphere coordinate is classical (Diaconis & Freedman, 1987). Lemma 4 makes its second-order term and remainder explicit and uniform.

C.1 UNIFORM GAMMA-RATIO CONTROL

Throughout this section $B_n(\cdot)$ is the nth Bernoulli polynomial and $B_n = B_n(0)$ the nth Bernoulli number.

Lemma 3 (Uniform finite gamma-ratio expansion)**.** *Let $K_b \subset \mathbb{R}$ be compact and $N \geq 1$ an integer, and let a_0 be such that $a + b > 0$ for $a \geq a_0$, $b \in K_b$. Then, uniformly on this set,*

$$\log \frac{\Gamma(a+b)}{\Gamma(a)} = b \log a + \sum_{m=1}^{N-1} \frac{(-1)^{m+1}\{B_{m+1}(b) - B_{m+1}(0)\}}{m(m+1)a^m} + R_N(a, b), \tag{17}$$
$$|\partial_a^i \partial_b^j R_N(a, b)| \leq C_{N,K_b} a^{-N-i}, \ i, j \in \{0, 1\}.$$

In particular, for $a = \frac{\rho M}{p}$ with (ρ, p) in a compact subset of $(0, \infty)^2$, every finite gamma-ratio expansion below, together with one ρ or p derivative, has remainder $\mathcal{O}(M^{-N})$.

Proof. With $\psi = (\log\Gamma)'$ we have $\log\frac{\Gamma(a+b)}{\Gamma(a)} = \int_0^b \psi(a+t)\mathrm{d}t$. The Euler–Maclaurin expansion of ψ with remainder (Tricomi et al., 1951; Olver et al., 2010),

$$\psi(x) = \log x - \frac{1}{2x} - \sum_{r=1}^{L-1} \frac{B_{2r}}{2rx^{2r}} + \mathcal{R}_L(x), \qquad |\partial_x^i \mathcal{R}_L(x)| \le C_L x^{-2L-i}, \;\; i \in \{0,1\},$$

holds for x bounded away from zero. Take $2L \ge N$. Increase a_0 so that $a + t \ge \frac{a}{2}$ for $|t| \le \max_{K_b} |b|$. Taylor's formula with integral remainder gives, uniformly in such t,

$$(a+t)^{-r} = a^{-r} \sum_{k=0}^{N-1} \binom{-r}{k} \left(\frac{t}{a}\right)^k + O_{K_b,N}\left(a^{-r-N}\right),$$

and

$$\log(a+t) = \log a + \sum_{k=1}^{N-1} \frac{(-1)^{k+1}}{k} \left(\frac{t}{a}\right)^k + O_{K_b,N}\left(a^{-N}\right).$$

Substitute both into the digamma expansion and integrate each monomial t^k over $[0,b]$. The identity $B_m(b) - B_m(0) = m\int_0^b B_{m-1}(t)\mathrm{d}t$ collects the coefficient of a^{-m} into $\frac{(-1)^{m+1}\{B_{m+1}(b)-B_{m+1}(0)\}}{m(m+1)}$. The leading term integrates to $b\log a$, and all terms of order a^{-N} go into R_N. Differentiating the remainder factor $(a+t)^{-r}$ in a gives one factor of a^{-1}. Differentiating in b acts on the upper endpoint and on Bernoulli polynomials over the compact K_b. This gives the derivative bound. For $a = \frac{\rho M}{p}$ we have $a \asymp M$ and $\partial_\rho a, \partial_p a = \mathcal{O}(M)$, so the chain rule turns $\partial_a R_N = \mathcal{O}(M^{-N-1})$ into $\mathcal{O}(M^{-N})$. Finite exponentiation and finite binomial sums preserve the order. □

C.2 Uniform spherical-coordinate approximation

Lemma 4 (Sphere CDF and uniform tails)**.** *Let U be uniform on $\mathbb{S}^{M-1}$ and $H_M(y) = \Pr(\sqrt{M}U_1 \le y)$. For every $Y_0 < \infty$ and integer $k \ge 0$ there are $C_{Y_0,k}, M_{Y_0,k} < \infty$ such that, for $M \ge M_{Y_0,k}$,*

$$\sup_{|y|\le Y_0} \left| \partial_y^k \left[H_M(y) - \Phi(y) - \phi(y)\left\{ \frac{W_1(y)}{M} + \frac{W_2(y)}{M^2} \right\} \right] \right| \le C_{Y_0,k} M^{-3}, \tag{18}$$

with $W_1(y) = \frac{y^3-3y}{4}$ and $W_2(y) = \frac{-3y^7+31y^5-43y^3-21y}{96}$. The derivative identity $(\phi W_1)' = \phi P_1$ holds with $P_1(y) = \frac{-y^4+6y^2-3}{4}$. For all $M \ge 2$ and $x \ge 0$ we have $\Pr(\sqrt{M}U_1 \ge x) \le e^{-\frac{x^2}{2}}$, and likewise for the lower tail.

Proof. The density of $\sqrt{M}U_1$ is $H_M'(y) = \frac{\Gamma(\frac{M}{2})}{\sqrt{\pi M}\Gamma(\frac{M-1}{2})}\left(1-\frac{y^2}{M}\right)^{\frac{M-3}{2}}$ for $|y| < \sqrt{M}$ (Fang et al., 2013). Lemma 3 for the normalizing constant and $\log(1-\frac{y^2}{M}) = -\frac{y^2}{M} - \frac{y^4}{2M^2} - \frac{y^6}{3M^3} + \mathcal{O}(M^{-4})$ give, uniformly on bounded intervals,

$$\log\frac{H_M'(y)}{\phi(y)} = \frac{L_1(y)}{M} + \frac{L_2(y)}{M^2} + \mathcal{O}(M^{-3}), \qquad L_1 = -\frac{3}{4} + \frac{3y^2}{2} - \frac{y^4}{4}, \quad L_2 = -\frac{1}{2} + \frac{3y^4}{4} - \frac{y^6}{6},$$

so $\frac{H_M'}{\phi} = 1 + \frac{L_1}{M} + \frac{L_2 + \frac{1}{2}L_1^2}{M^2} + \mathcal{O}(M^{-3})$. Direct differentiation verifies $(\phi W_1)' = \phi L_1 = \phi P_1$ and $(\phi W_2)' = \phi\{L_2 + \frac{1}{2}L_1^2\}$. Integrating from 0 with $H_M(0) = \Phi(0)$ gives eq. (18) for $k = 0$.

For $k \ge 1$ we use a Cauchy estimate. Let $Y_1 = Y_0 + 2$ and $\mathcal{D} = \{y \in \mathbb{C} : |\operatorname{Re} y| \le Y_1, \, |\operatorname{Im} y| \le 1\}$. For $y \in \mathcal{D}$ and $M > 8(Y_1+1)^2$ we have $|\frac{y^2}{M}| < \frac{1}{2}$, so the principal branch of $(1-\frac{y^2}{M})^{\frac{M-3}{2}}$ is holomorphic on $\mathcal{D}$. Its primitive $\mathcal{H}_M(y) = \frac{1}{2} + \int_0^y \mathcal{H}_M'(\zeta)\mathrm{d}\zeta$ is holomorphic and agrees with H_M on $[-Y_1, Y_1]$. Since Φ, ϕ, W_1, W_2 are entire, $\mathcal{E}_M(y) = \mathcal{H}_M(y) - \Phi(y) - \phi(y)\{\frac{W_1(y)}{M} + \frac{W_2(y)}{M^2}\}$ is holomorphic on $\mathcal{D}$. The density expansion above holds there because $\log(1-\frac{y^2}{M}) = -\sum_{m\ge1}\frac{1}{m}(\frac{y^2}{M})^m$ converges uniformly on $|\frac{y^2}{M}| < \frac{1}{2}$. Hence $\sup_{\mathcal{D}}|\mathcal{E}_M| \le C_{Y_1}M^{-3}$. For real $|y| \le Y_0$, the unit disc about y lies in $\mathcal{D}$, so Cauchy's estimate gives $|\mathcal{E}_M^{(k)}(y)| \le k! C_{Y_1} M^{-3}$.

For the tails, odd moments vanish by symmetry and $\mathbb{E}[(\sqrt{M}U_1)^{2k}] = (2k-1)!!\frac{M^k}{M(M+2)\cdots(M+2k-2)} \le (2k-1)!!$. Hence $\mathbb{E}[e^{t\sqrt{M}U_1}] \le \sum_k \frac{t^{2k}(2k-1)!!}{(2k)!} = e^{\frac{t^2}{2}}$. Markov's inequality with $t = x$ gives $\Pr(\sqrt{M}U_1 \ge x) \le e^{-\frac{x^2}{2}}$. □

D THE CONVERSE FOR ARBITRARY ADDITIVE NOISE

By Lemma 2, it suffices to consider radial noise $Z = RU$, whose distribution is determined by that of the normalized radial energy $V = \frac{\|Z\|_2^2}{M\Delta_2^2}$ on $[0,\infty)$. We fix one distinguishing event, prove a support inequality pointwise in v, and integrate it against the distribution of V.

D.1 THE DISTINGUISHING EVENT

For a threshold function θ on $[0,\infty)$ define

$$\bar{S}_{M,\theta} = \left\{ y \in \mathbb{R}^M : y_1 \le \theta\left(\frac{\|y\|_2^2}{M}\right)\right\},$$
$$h_M[\theta](v) = \Pr\left(\sqrt{Mv}U \in \bar{S}_{M,\theta}\right) - e^{\epsilon}\Pr\left(\sqrt{Mv}U + e_1 \in \bar{S}_{M,\theta}\right).$$

For $Z = RU$ with $V = \frac{\|Z\|_2^2}{M\Delta_2^2}$, the event $\{Z \in \Delta_2\bar{S}_{M,\theta}\}$ has hockey-stick mass $\mathbb{E}[h_M[\theta](V)]$ at the shift $\Delta_2 e_1$, since conditionally on $V = v$, the normalized noise is $\sqrt{Mv}U$. The shift $\Delta_2 e_1$ lies in the sensitivity ball, so every feasible radial distribution satisfies $\mathbb{E}[h_M[\theta](V)] \le \delta$. The sensitivity scale has dropped out.

D.2 THE POINTWISE DUAL CERTIFICATE

Lemma 5 (Pointwise dual certificate). *If $\lambda \ge 0$ and $v + \lambda h_M[\theta](v) \ge \kappa$ for every $v \ge 0$, then $\mathcal{V}_M(\Delta_2) \ge \Delta_2^2 M(\kappa - \lambda\delta)$.*

Proof. Let $Z = RU$ be feasible. If $\mathbb{E}[\|Z\|_2^2] = \infty$ there is nothing to prove. The hypothesis holds at $v = V$ almost surely, so taking expectations gives $\mathbb{E}[V] + \lambda\mathbb{E}[h_M[\theta](V)] \ge \kappa$. Feasibility gives $\mathbb{E}[h_M[\theta](V)] \le \delta$ and $\lambda \ge 0$, hence $\mathbb{E}[V] \ge \kappa - \lambda\delta$ and $\mathbb{E}[\|Z\|_2^2] = M\Delta_2^2\mathbb{E}[V] \ge \Delta_2^2 M(\kappa - \lambda\delta)$. By Lemma 2, every feasible distribution has a feasible radial counterpart with the same error, so the infimum obeys the same bound. □

This is the weak direction of infinite-dimensional linear-programming duality (Anderson & Nash, 1987).

D.3 THE CRITICAL SUPPORTING EVENT

Let $j(v,z) = \Phi(\frac{z}{\sqrt{v}}) - e^{\epsilon}\Phi(\frac{z-1}{\sqrt{v}})$ be the $M \to \infty$ limit of the event mass at threshold z, and let $\ell(v) = \delta + g'(v_0)(v - v_0)$ be the tangent to g at v_0.

Lemma 6 (Strict global tangent). *Under Assumption 1, $g(v) \ge \ell(v)$ for all $v \ge 0$, with $g(0) := 1$ and equality only at v_0.*

Proof. By Appendix B.1 g is strictly concave on $(0, v_c)$ and strictly convex on (v_c,∞), and $v_0 > v_c$. On $[v_c,\infty)$ strict convexity gives

$$g(v) \ge g(v_0) + g'(v_0)(v - v_0) = \ell(v), \tag{19}$$

with equality only at v_0. On $[0, v_c]$, the function $g - \ell$ is concave. It is positive at 0, where $g(0) - \ell(0) = 1 - \delta + v_0 g'(v_0) > 0$ by the second inequality of Assumption 1, and positive at v_c by eq. (19), since $v_c \ne v_0$. A concave function that is positive at both endpoints is positive on the whole interval, so $g > \ell$ on $[0, v_c]$. □

For fixed v, the map $z \mapsto j(v,z)$ is maximized at the Gaussian threshold $\theta_G(v) = \frac{1}{2} - \epsilon v$ with maximum $g(v)$ (verified at the start of the proof of Lemma 7). Contact between $j(v,\theta_G(v)) = g(v)$ and ℓ therefore occurs at the single point v_0. The next lemma replaces it by contact on an interval.

Lemma 7 (Critical branch)**.** *There is an analytic threshold Θ on a neighborhood of v_0 with $j(v,\Theta(v)) = \ell(v)$, $\Theta(v_0) = \theta_G(v_0)$, and $\Theta'(v_0) = -\epsilon + \frac{\epsilon\tau_*}{2}$. The function $v \mapsto j_z(v,\Theta(v))$ has a simple zero at v_0.*

Proof. Differentiating in z,

$$j_z(v,z) = \frac{1}{\sqrt{v}}\left\{\phi\left(\frac{z}{\sqrt{v}}\right) - e^{\epsilon}\phi\left(\frac{z-1}{\sqrt{v}}\right)\right\}, \qquad j_{zz}(v,z) = \frac{-z\phi\left(\frac{z}{\sqrt{v}}\right) + e^{\epsilon}(z-1)\phi\left(\frac{z-1}{\sqrt{v}}\right)}{v^{\frac{3}{2}}}.$$

The ratio of the two densities is $e^{\frac{1-2z}{2v}-\epsilon}$, equal to 1 exactly at $z = \theta_G(v)$, so $j_z(v,\theta_G(v)) = 0$. There $\frac{z}{\sqrt{v}} = A(v)$ and $\frac{z-1}{\sqrt{v}} = B(v)$, and $\phi(A(v)) = e^{\epsilon}\phi(B(v))$ gives $j_{zz}(v,\theta_G(v)) = -\frac{\phi(A(v))}{v^{\frac{3}{2}}} < 0$. Hence $\theta_G(v)$ is the unique maximizer and $j(v,\theta_G(v)) = g(v)$.

Let $w = v - v_0$ and, for $w \neq 0$,

$$F(w,a) = \frac{j(v_0+w, \theta_G(v_0+w) + wa) - \ell(v_0+w)}{w^2}.$$

Since j_z vanishes along θ_G, Taylor expansion in the displacement wa gives $j(v_0+w, \theta_G(v_0+w) + wa) = g(v_0+w) + \frac{1}{2}j_{zz}(v_0+w, \theta_G(v_0+w))w^2a^2 + \mathcal{O}(w^3)$, and $g(v_0+w) - \ell(v_0+w) = \frac{1}{2}g''(v_0)w^2 + \mathcal{O}(w^3)$. Thus F extends analytically to $w = 0$ with

$$F(0,a) = \frac{1}{2}\left\{g''(v_0) + j_{zz}(v_0, \theta_G(v_0))a^2\right\}.$$

Its roots satisfy $a_c^2 = -\frac{g''(v_0)}{j_{zz}(v_0,\theta_G(v_0))} = \frac{\phi(A(v_0))Q}{4v_0^{\frac{7}{2}}} \cdot \frac{v_0^{\frac{3}{2}}}{\phi(A(v_0))} = \frac{Q}{4v_0^2}$ by eq. (10), so $a_c = \frac{\sqrt{Q}}{2v_0} = \frac{\epsilon\tau_*}{2}$ using $Q = E\tau_*^2$. We take the positive root. Remark 2 shows that the negative root yields the larger constant $\frac{E}{2}(2+\tau_*)^2$ instead of $C(\epsilon,\delta)$. Since $\partial_a F(0,a_c) = j_{zz}(v_0,\theta_G(v_0))a_c \neq 0$, the analytic implicit function theorem (Krantz & Parks, 2012) gives an analytic $a(v)$ near v_0 with $a(v_0) = a_c$ and $F(v-v_0, a(v)) = 0$. Set $\Theta(v) = \theta_G(v) + (v-v_0)a(v)$. Then $j(v,\Theta(v)) = \ell(v)$, $\Theta(v_0) = \theta_G(v_0)$, and $\Theta'(v_0) = \theta_G'(v_0) + a_c = -\epsilon + \frac{\epsilon\tau_*}{2}$.

Finally, $j_z(v,\theta_G(v)) \equiv 0$ gives $j_{zv} + j_{zz}\theta_G' = 0$ along θ_G, so

$$\frac{\mathrm{d}}{\mathrm{d}v} j_z(v,\Theta(v))\Big|_{v_0} = j_{zz}(v_0,\theta_G(v_0))\{\Theta'(v_0) - \theta_G'(v_0)\} = j_{zz}(v_0,\theta_G(v_0))a_c \neq 0,$$

also when $\tau_* = 2$, where $\Theta'(v_0) = 0$ but $a_c = \epsilon > 0$. □

D.4 Conditional event operator and its uniform expansion

All thresholds below equal $\theta_G(v) = \frac{1}{2} - \epsilon v$ outside a fixed compact subinterval of $(0,\infty)$, are extended by θ_G to negative arguments, and have bounded derivatives, so $\sup\theta' < \frac{M}{2}$ for large M. For $v > 0$, the map $x \mapsto x + 1 - \theta(v + \frac{2x+1}{M})$ is then strictly increasing and tends to $\pm\infty$ as $x \to \pm\infty$. Its unique zero $s_M(v)$ is defined by

$$s_M(v) + 1 = \theta\left(v + \frac{2s_M(v)+1}{M}\right). \tag{20}$$

On the sphere $y = \sqrt{Mv}U$ one has $\frac{\|y+e_1\|_2^2}{M} = v + \frac{2y_1+1}{M} \geq 0$, so the extension to negative arguments does not affect the event, and $y + e_1 \in \bar{S}_{M,\theta}$ exactly when $y_1 \leq s_M(v)$. With H_M, the sphere CDF of Lemma 4 clipped outside its support,

$$h_M[\theta](v) = H_M\left(\frac{\theta(v)}{\sqrt{v}}\right) - e^{\epsilon}H_M\left(\frac{s_M(v)}{\sqrt{v}}\right) \qquad (v > 0).$$

At $v = 0$, evaluate the two indicators directly.

Let $\eta = \frac{1}{M}$ and, for a general threshold θ, $A = \frac{\theta(v)}{\sqrt{v}}$ and $B = \frac{\theta(v)-1}{\sqrt{v}}$, which reduce to $A(v), B(v)$ of Appendix B.1 when $\theta = \theta_G$. Let

$$k_1 = (2\theta - 1)\theta', \qquad k_2 = 2(2\theta - 1)(\theta')^2 + \frac{1}{2}(2\theta - 1)^2\theta''. \tag{21}$$

Write $s_M = s_0 + \eta s_1 + \eta^2 s_2 + \mathcal{O}(\eta^3)$ and expand the right side of eq. (20) as $\theta\{v + \eta(2s_M + 1)\} = \theta + \eta(2s_0 + 1)\theta' + \eta^2\{2s_1\theta' + \frac{1}{2}(2s_0 + 1)^2\theta''\} + \mathcal{O}(\eta^3)$. Equating coefficients gives $s_0 = \theta - 1$, $s_1 = k_1$, $s_2 = k_2$, so $\frac{s_M}{\sqrt{v}} = B + \eta\frac{k_1}{\sqrt{v}} + \eta^2\frac{k_2}{\sqrt{v}} + \mathcal{O}(\eta^3)$. Lemma 4 at A, and at the expanded second argument together with $\Phi''(B) = -B\phi(B)$ and $(\phi W_1)' = \phi P_1$, gives

$$H_M\left(B + \eta\frac{k_1}{\sqrt{v}} + \eta^2\frac{k_2}{\sqrt{v}}\right) = \Phi(B) + \eta\phi(B)\left\{W_1(B) + \frac{k_1}{\sqrt{v}}\right\} + \eta^2\phi(B)\left\{W_2(B) + P_1(B)\frac{k_1}{\sqrt{v}} - \frac{B}{2}\left(\frac{k_1}{\sqrt{v}}\right)^2 + \frac{k_2}{\sqrt{v}}\right\} + \mathcal{O}(\eta^3).$$

Hence $h_M[\theta](v) = j(v, \theta(v)) + \eta H_1[\theta](v) + \eta^2 H_2[\theta](v) + \mathcal{O}(\eta^3)$ with

$$H_1[\theta] = \phi(A)W_1(A) - e^{\epsilon}\phi(B)\left[W_1(B) + \frac{k_1}{\sqrt{v}}\right], \tag{22}$$

$$H_2[\theta] = \phi(A)W_2(A) - e^{\epsilon}\phi(B)\left[W_2(B) + P_1(B)\frac{k_1}{\sqrt{v}} - \frac{B}{2}\left(\frac{k_1}{\sqrt{v}}\right)^2 + \frac{k_2}{\sqrt{v}}\right]. \tag{23}$$

Replacing θ by $\theta + \eta u$ changes A and B by $\eta\frac{u}{\sqrt{v}}$ and k_1 by $\eta\{2u\theta' + (2\theta - 1)u'\}$ to first order. The product rule then gives the derivative of H_1 in the direction u,

$$DH_1[\theta]u = [\phi(A)P_1(A) - e^{\epsilon}\phi(B)P_1(B)]\frac{u}{\sqrt{v}} - e^{\epsilon}\phi(B)\left[\frac{2u\theta' + (2\theta - 1)u'}{\sqrt{v}} - \frac{Buk_1}{v}\right]. \tag{24}$$

Lemma 8 (Uniform expansion for dimension-dependent thresholds)**.** *Let $K \Subset (0, \infty)$ and let K^+ be an open neighborhood of K with compact closure in $(0, \infty)$. Suppose $\theta_M = T + M^{-1}u + M^{-2}b$ on K^+ with $T, u, b \in C^5(K^+)$, their C^5 norms bounded independently of M, and each θ_M extended to a C^1 function on $\mathbb{R}$ such that, for some fixed $c < \frac{1}{2}$ and all large M,*

$$\sup_{x\in\mathbb{R}} \theta_M'(x) \le cM, \qquad \lim_{s\to\pm\infty}\left\{s + 1 - \theta_M\left(v + \frac{2s+1}{M}\right)\right\} = \pm\infty \quad (v \in K). \tag{25}$$

Then the root in eq. (20) is unique and there are C, M_0 such that, uniformly for $v \in K$ and $M \ge M_0$,

$$h_M[\theta_M](v) = j(v, T) + \frac{j_z(v, T)u + H_1[T]}{M} + \frac{j_z(v, T)b + \frac{1}{2}j_{zz}(v, T)u^2 + DH_1[T]u + H_2[T]}{M^2} + r_M(v), \tag{26}$$

with $\max_{m=0,1}\sup_{v\in K}|\partial_v^m r_M(v)| \le CM^{-3}$. In particular, the constant is uniform for the thresholds eq. (37) on every compact subset of $(0, \infty)$.

Proof. Let $F(s, v, \eta) = s + 1 - \theta_M(v + \eta(2s + 1))$. By eq. (25) $\partial_s F = 1 - 2\eta\theta_M' \ge 1 - 2c > 0$, so F is strictly increasing in s with limits of opposite sign and the root $s_M(v)$ is unique. With $s_0 = \theta_M(v) - 1$ the C^5 bound gives $F(s_0, v, \eta) = \theta_M(v) - \theta_M\{v + \eta(2\theta_M(v) - 1)\} = \mathcal{O}(\eta)$ uniformly on K. The mean-value theorem gives $|s_M - s_0| \le \frac{|F(s_0,v,\eta)|}{1-2c} = \mathcal{O}(\eta)$. For large M, the root and all shifted arguments lie in a fixed compact subset of K^+. The parameterized implicit function theorem and Taylor's theorem with integral remainder then give, for every threshold ϑ in this bounded C^5 family and uniformly on K,

$$s_M[\vartheta] = \vartheta - 1 + \eta k_1[\vartheta] + \eta^2 k_2[\vartheta] + \eta^3 r_{s,M}[\vartheta], \qquad \|r_{s,M}\|_{C^1(K)} \le C_1,$$

where $k_i[\vartheta]$ is eq. (21) with θ replaced by ϑ. Both CDF arguments stay in a fixed bounded interval, so the differentiated form of Lemma 4 and the expansions preceding the lemma yield

$$h_M[\vartheta] = j(v, \vartheta) + \eta H_1[\vartheta] + \eta^2 H_2[\vartheta] + O_{C^1(K)}(\eta^3) \tag{27}$$

over the whole family. Substitute $\vartheta = \theta_M = T + \eta u + \eta^2 b$ and expand in the threshold argument. This gives

$$j(v, \theta_M) = j(v, T) + \eta j_z(v, T)u + \eta^2 \left\{ j_z(v, T)b + \frac{1}{2} j_{zz}(v, T)u^2 \right\} + O_{C^1}(\eta^3),$$
$$H_1[\theta_M] = H_1[T] + \eta DH_1[T]u + O_{C^1}(\eta^2), \qquad H_2[\theta_M] = H_2[T] + O_{C^1}(\eta),$$

and collecting orders in eq. (27) gives eq. (26). The thresholds eq. (37) are θ_G plus fixed smooth compactly supported corrections. Their extensions therefore have bounded derivatives and the required endpoint behavior, and eq. (25) holds for all large M. □

D.5 THE UNIFORM GLOBAL SUPPORT

Lemma 9 (Uniform global support)**.** *Under Assumption 1 there are smooth thresholds θ_M and multipliers $\lambda_M = -\frac{1}{g'(v_0)} + \mathcal{O}(M^{-1}) > 0$ such that, for every $v \geq 0$ and all large M,*

$$v + \lambda_M h_M[\theta_M](v) \geq \kappa_M, \qquad \kappa_M - \lambda_M \delta = v_0 - \frac{C(\epsilon, \delta)}{M} + \frac{J(\epsilon, \delta)}{M^2} + \mathcal{O}(M^{-3}). \tag{28}$$

Proof. We prove the following quantitative form. For every $u_0 \in \mathbb{R}$ there are smooth thresholds θ_M, constants a_{11}, a_{20}, K_0, and multipliers

$$\lambda_M = -\left\{ g'(v_0) + \frac{a_{11}}{M} \right\}^{-1} > 0 \tag{29}$$

such that, for all $v \geq 0$ and large M,

$$v + \lambda_M h_M[\theta_M](v) \geq v_0 + \lambda_M \left\{ \delta + \frac{g'(v_0)C(\epsilon, \delta)}{M} + \frac{a_{20}}{M^2} \right\} - K_0 M^{-3}. \tag{30}$$

Expanding λ_M in eq. (30) gives $\kappa_M - \lambda_M \delta = v_0 - \frac{C(\epsilon,\delta)}{M} + \frac{C(\epsilon,\delta)a_{11} - a_{20}}{M^2 g'(v_0)} + \mathcal{O}(M^{-3})$. The last part of the proof shows

$$\frac{C(\epsilon, \delta)a_{11} - a_{20}}{g'(v_0)} = J(\epsilon, \delta) - (u_0 - u_*)^2, \qquad u_* = \frac{\epsilon}{2} \left\{ v_0(2 - p_*) + \epsilon v_0(\tau_* - 2) + \tau_* C(\epsilon, \delta) \right\}. \tag{31}$$

The lemma is the case $u_0 = u_*$.

Cutoffs and the glued threshold. Choose $0 < w_0 < w_1$ such that the branch Θ of Lemma 7 exists on an open set $\mathcal{O} \supset [v_0 - w_1, v_0 + w_1] \Subset (0, \infty)$. Let $I_0 = [v_0 - w_0, v_0 + w_0]$ and $I_1 = (v_0 - w_1, v_0 + w_1)$, and fix nested intervals $I_0 \Subset \Omega_0 \Subset \Omega \Subset I_1$. Let $\Psi(x) = \frac{e^{-1/x}}{e^{-1/x} + e^{-1/(1-x)}}$ for $0 < x < 1$, $\Psi = 0$ on $(-\infty, 0]$, and $\Psi = 1$ on $[1, \infty)$, so that $\Psi \in C^\infty(\mathbb{R})$ and $0 < \Psi < 1$ on $(0, 1)$. Set

$$\chi(v) = \Psi\left(\frac{w_1 - |v - v_0|}{w_1 - w_0} \right), \qquad \overline{\Theta}(v) = \theta_G(v) + \chi(v)\{\Theta(v) - \theta_G(v)\},$$

where $\chi\{\Theta - \theta_G\}$ is extended by zero outside $\mathcal{O}$. Since χ vanishes off $\overline{I}_1 \subset \mathcal{O}$, the extension is smooth, and $\chi \in C^\infty(\mathbb{R})$ because Ψ is constant where $|v - v_0| < w_0$. Thus

$$\chi = 1 \text{ on } I_0, \qquad 0 < \chi < 1 \text{ on } I_1 \setminus I_0, \qquad \chi = 0 \text{ outside } I_1. \tag{32}$$

Let $\zeta \in C_c^\infty(\Omega)$ with $0 \leq \zeta \leq 1$ and $\zeta = 1$ on a neighborhood of $\overline{\Omega}_0$. The same zero extension applies to ζu and ζb below. Since $I_0 \Subset \Omega_0$,

$$\{\chi = 1\} = I_0 \subset \{\zeta = 1\}. \tag{33}$$

Conditions eq. (32) and eq. (33) confine exact leading contact to the set where both corrections $\frac{u}{M}, \frac{b}{M^2}$ act at full strength. When $\zeta < 1$, the M^{-1} coefficient carries the residual $\{\zeta(v) - 1\}D(v)u(v)$, with D and u defined in the next paragraph, and there the leading margin $j(v, \overline{\Theta}(v)) - \ell(v)$ is strictly positive and absorbs it.

Leading contact set. We claim that $j(v, \overline{\Theta}(v)) \geq \ell(v)$ with equality exactly on I_0, that is,

$$\{v \geq 0 : j(v, \overline{\Theta}(v)) = \ell(v)\} = I_0. \tag{34}$$

On I_0 we have $\overline{\Theta} = \Theta$ and Lemma 7 gives equality. Outside I_1 we have $\overline{\Theta} = \theta_G$, so $j(v, \overline{\Theta}(v)) = g(v) > \ell(v)$ by Lemma 6, since $v \neq v_0$. On $I_1 \setminus I_0$ we have $0 < \chi < 1$. The density ratio in the proof of Lemma 7 shows $j_z(v, z) > 0$ for $z < \theta_G(v)$ and $j_z(v, z) < 0$ for $z > \theta_G(v)$, so $z \mapsto j(v, z)$ is strictly unimodal about $\theta_G(v)$. If $\Theta(v) \neq \theta_G(v)$ then $\overline{\Theta}(v)$ lies strictly between $\Theta(v)$ and $\theta_G(v)$ and $j(v, \overline{\Theta}(v)) > j(v, \Theta(v)) = \ell(v)$. If $\Theta(v) = \theta_G(v)$ then $j(v, \overline{\Theta}(v)) = g(v) > \ell(v)$.

First correction and the constant $C(\epsilon, \delta)$. Let $D(v) = j_z(v, \overline{\Theta}(v))$ and $H(v) = H_1[\overline{\Theta}](v)$. Since $\overline{\Theta} = \Theta$ on the neighborhood I_0 of v_0, all derivatives at v_0 may be computed on Θ. Lemma 7 gives $D(v_0) = 0$ and $D'(v_0) = j_{zz}(v_0, \theta_G(v_0))a_c \neq 0$. Shrinking w_1 if necessary, D has no other zero on Ω.

We compute $a_{10} := H(v_0)$. Let $A = A(v_0)$, $B = B(v_0)$, $\theta = \theta_G(v_0) = \frac{1}{2} - \epsilon v_0$, and $\theta' = \theta'_G(v_0) + \sigma a_c = -\epsilon + \sigma \frac{\epsilon \tau_*}{2}$, where $\sigma \in \{\pm 1\}$ is the root chosen in Lemma 7 ($\sigma = +1$ there). Then $2\theta - 1 = -2\epsilon v_0$ and $k_1 = (2\theta - 1)\theta' = \epsilon^2 v_0(2 - \sigma\tau_*)$. With $e^{\epsilon}\phi(B) = \phi(A)$, eq. (22) becomes $H(v_0) = \phi(A)\{W_1(A) - W_1(B) - \frac{k_1}{\sqrt{v_0}}\}$. From $A - B = \frac{1}{\sqrt{v_0}}$, $A + B = -2\epsilon\sqrt{v_0}$, and $AB = \epsilon^2 v_0 - \frac{1}{4v_0}$ we get $A^2 + AB + B^2 = 3\epsilon^2 v_0 + \frac{1}{4v_0}$ and $W_1(A) - W_1(B) = \frac{(A-B)(A^2+AB+B^2-3)}{4} = \frac{1}{4\sqrt{v_0}}(3\epsilon^2 v_0 + \frac{1}{4v_0} - 3)$. Hence $H(v_0) = \frac{\phi(A)}{\sqrt{v_0}}\Xi$ with $4v_0\Xi = 3E + \frac{1}{4} - 3v_0 - 4E(2 - \sigma\tau_*) = -E(2 - \sigma\tau_*)^2$, using $3v_0 - \frac{1}{4} = E(\tau_*^2 - 1)$ (that is, $Q = E\tau_*^2$) and $\sigma^2 = 1$. By eq. (10),

$$a_{10} = H(v_0) = -\frac{\phi(A)}{4v_0^{\frac{3}{2}}}E(2 - \sigma\tau_*)^2 = g'(v_0)\frac{E}{2}(2 - \sigma\tau_*)^2, \tag{35}$$

which for $\sigma = +1$ is $a_{10} = g'(v_0)C(\epsilon, \delta)$. Remark 2 uses eq. (35) for the other sign.

Fix $u_0 \in \mathbb{R}$ and set $a_{11} = H'(v_0) + D'(v_0)u_0$. The numerator $N_1(v) = a_{10} + a_{11}(v - v_0) - H(v)$ satisfies $N_1(v_0) = 0$ and $N_1'(v_0) = D'(v_0)u_0$. Since D has a simple zero at v_0 and no other zero on Ω, Hadamard's lemma makes $u = \frac{N_1}{D}$ smooth on Ω with $u(v_0) = u_0$, and

$$D(v)u(v) + H(v) = a_{10} + a_{11}(v - v_0) \qquad (v \in \Omega). \tag{36}$$

For the M^{-2} coefficient of eq. (26) let $F_2 = H_2[\overline{\Theta}] + DH_1[\overline{\Theta}]u + \frac{1}{2}j_{zz}(\cdot, \overline{\Theta})u^2$ and $a_{20} = F_2(v_0)$. The same argument makes $b = \frac{a_{20} - F_2}{D}$ smooth on Ω with $D(v)b(v) + F_2(v) = a_{20}$ on Ω and $b(v_0) = -\frac{F_2'(v_0)}{D'(v_0)}$. Set

$$\theta_M(v) = \overline{\Theta}(v) + \frac{\zeta(v)u(v)}{M} + \frac{\zeta(v)b(v)}{M^2}. \tag{37}$$

Each summand is smooth on $\mathbb{R}$, the corrections vanish outside Ω, and $\overline{\Theta} = \theta_G$ outside I_1, so θ_M satisfies the hypotheses of Lemma 8 for large M. On a neighborhood of $\overline{\Omega}_0$ we have $\zeta \equiv 1$, and eq. (26) with $T = \overline{\Theta}$ gives, uniformly on $\overline{\Omega}_0$,

$$h_M[\theta_M](v) = j(v, \overline{\Theta}(v)) + \frac{a_{10} + a_{11}(v - v_0)}{M} + \frac{a_{20}}{M^2} + \mathcal{O}(M^{-3}). \tag{38}$$

Support on four regions. Define λ_M by eq. (29). Since $g'(v_0) < 0$ we have $\lambda_M > 0$ for large M and $\lambda_M \to \lambda_\infty = -\frac{1}{g'(v_0)}$. Let

$$F_M(v) = v + \lambda_M h_M[\theta_M](v) - v_0 - \lambda_M \left\{\delta + \frac{a_{10}}{M} + \frac{a_{20}}{M^2}\right\}.$$

Contact region $\overline{\Omega}_0$. By eq. (38) and $j(v, \overline{\Theta}(v)) \geq \ell(v) = \delta + g'(v_0)(v - v_0)$,

$$F_M(v) \geq \left[1 + \lambda_M \left\{g'(v_0) + \frac{a_{11}}{M}\right\}\right](v - v_0) - K_1 M^{-3} = -K_1 M^{-3},$$

where the square bracket vanishes by eq. (29).

Small v. Fix $v_- < \inf I_1$. On $(0, v_-]$ we have $\theta_M = \theta_G$ and eq. (20) solves to $s_M(v) = -\frac{\frac{1}{2}+\epsilon v+\frac{\epsilon}{M}}{1+\frac{2\epsilon}{M}}$. Taking v_- small and M large gives $\theta_M(v) \geq \frac{1}{4}$ and $s_M(v) \leq -\frac{1}{4}$, so the tail bound of Lemma 4 yields $h_M[\theta_M](v) \geq 1 - (1+e^\epsilon)e^{-\frac{1}{32v}}$. At $v = 0$ we have $h_M[\theta_M](0) = 1$. The limiting margin is $\lambda_\infty\{1 - \ell(0)\} = \lambda_\infty\{1 - \delta + v_0 g'(v_0)\} > 0$ by Assumption 1, so $F_M > 0$ on $[0, v_-]$ for small v_- and large M.

Large v. Since $h_M[\theta_M] \geq -e^\epsilon$ and both λ_M and $v_0 + \lambda_M\{\delta + \frac{a_{10}}{M} + \frac{a_{20}}{M^2}\}$ are bounded, there is $v_+ > \sup I_1$ with $F_M > 0$ on $[v_+, \infty)$ for all large M.

Exterior $K = [v_-, v_+] \setminus \Omega_0$. As $K \cap I_0 = \emptyset$, eq. (34) gives $j(v, \overline{\Theta}(v)) > \ell(v)$ on K. The strict inequality $\chi < 1$ on $I_1 \setminus I_0$ is used here, because $K \supset \Omega \setminus \Omega_0 \neq \emptyset$ meets I_1, where Θ itself has exact contact. The continuous function $m(v) = \lambda_\infty\{j(v, \overline{\Theta}(v)) - \ell(v)\}$ has $m_0 = \min_K m > 0$. Lemma 8 with $\frac{1}{\lambda_\infty} = -g'(v_0)$ gives $F_M(v) = m(v) + \mathcal{O}(M^{-1})$ uniformly on K, so $F_M \geq \frac{m_0}{2}$ on K for large M.

The sets $[0, v_-]$, $\overline{\Omega}_0$, K, and $[v_+, \infty)$ cover $[0, \infty)$. With $K_0 \geq K_1$ this proves eq. (30) and hence eq. (28).

The identity eq. (31). On I_0 both cutoffs equal one, so $\theta_M = \Theta + \frac{u}{M} + \frac{b}{M^2}$, and we compute every derivative at v_0 on the analytic branch. Write $x = \frac{v-v_0}{v_0}$ and $f(x) = \frac{1}{2} - \epsilon v_0 + \epsilon v_0(\frac{\tau_*}{2} - 1)x + \gamma_2 x^2 + \gamma_3 x^3$ for the cubic jet of Θ, and $G(x; \gamma_2, \gamma_3) = j(v_0(1+x), f(x)) - \delta - g'(v_0)v_0 x$. Its constant and linear coefficients vanish by $j(v_0, \theta_G(v_0)) = \delta$ and the choice of ℓ. Its quadratic coefficient $\frac{v_0^2}{2}\{g''(v_0) + j_{zz}(v_0, \theta_G(v_0))(\frac{\epsilon\tau_*}{2})^2\}$ vanishes by $-\frac{g''}{j_{zz}} = \frac{Q}{4v_0^2} = \frac{\epsilon^2\tau_*^2}{4}$. With $[x^m]$ for Taylor coefficients, exact contact through order four fixes

$$\gamma_2 = -\frac{[x^3]G(x; 0, 0)}{[x^3]\{G(x; 1, 0) - G(x; 0, 0)\}}, \qquad \gamma_3 = -\frac{[x^4]G(x; \gamma_2, 0)}{[x^4]\{G(x; \gamma_2, 1) - G(x; \gamma_2, 0)\}}. \tag{39}$$

Both denominators equal $j_{zz}(v_0, \theta_G(v_0))v_0 a_c \neq 0$, because along the critical jet $j_z(v_0(1+x), f(x)) = j_{zz}(v_0, \theta_G(v_0))v_0 a_c x + \mathcal{O}(x^2)$ and perturbing γ_2 (resp. γ_3) by one changes f by x^2 (resp. x^3). In particular $\gamma_2 = \frac{\epsilon v_0 \tau_*}{8}(p_* + 2 - 2\tau_*)$.

Differentiating eq. (36) twice at v_0, where $D(v_0) = 0$ and $u(v_0) = u_0$, gives $a_{11} = H'(v_0) + D'(v_0)u_0$ and

$$u'(v_0) = -\frac{H''(v_0) + D''(v_0)u_0}{2D'(v_0)}. \tag{40}$$

In a_{11} and $a_{20} = H_2[\overline{\Theta}](v_0) + DH_1[\overline{\Theta}]u(v_0) + \frac{1}{2}j_{zz}(v_0, \theta_G(v_0))u_0^2$, evaluated through eq. (22), eq. (23), eq. (24), and eq. (40), u_0 enters linearly except for the term $\frac{1}{2}j_{zz}u_0^2$. Hence $\frac{C(\epsilon,\delta)a_{11}-a_{20}}{g'(v_0)}$ is a polynomial of degree two in u_0. We substitute the jet eq. (39), replace $e^\epsilon\phi(B(v_0))$ by $\phi(A(v_0))$, write $\Phi^{(m)}(z) = (-1)^{m-1}\,\mathrm{He}_{m-1}(z)\phi(z)$ with He_m, the probabilists' Hermite polynomials, reduce powers of τ_* by $E\tau_*^2 = Q$, and use $p_* = -\frac{c_1(v_0,\tau_*)}{E}$. The polynomial becomes

$$\frac{C(\epsilon, \delta)a_{11} - a_{20}}{g'(v_0)} = -u_0^2 + 2u_* u_0 + c_0(v_0, \tau_*) - \frac{c_1(v_0, \tau_*)^2}{2E} - u_*^2,$$

where each coefficient difference reduces to the zero rational function in (ϵ, v_0, τ_*) modulo $E\tau_*^2 - Q$. Since $c_0(v_0, \tau_*) - \frac{c_1(v_0,\tau_*)^2}{2E} = J(\epsilon, \delta)$ by eq. (15), the right side is $J(\epsilon, \delta) - (u_0 - u_*)^2$, which is eq. (31). □

Remark 2 (Why the positive root)**.** *Both roots $a_c = \pm\frac{\epsilon\tau_*}{2}$ of $F(0, a) = 0$ in Lemma 7 give an analytic branch with exact contact, and the proof of Lemma 9 goes through for either sign σ. By eq. (35), the resulting certificate yields the constant $\frac{E}{2}(2 - \sigma\tau_*)^2$, and $\frac{E}{2}(2 + \tau_*)^2 > \frac{E}{2}(2 - \tau_*)^2 = C(\epsilon, \delta)$ for $\tau_* > 0$. The negative root therefore gives the weaker bound $\mathcal{V}_M(\Delta_2) \geq M\Delta_2^2 v_0 - \frac{\Delta_2^2 E}{2}(2 + \tau_*)^2 + \mathcal{O}(\Delta_2^2 M^{-1})$. The positive root gives $C(\epsilon, \delta)$, which Corollary 1 shows to be sharp.*

D.6 Proof of Theorem 2

Proof. Take $u_0 = u_*$ in Lemma 9. Inequality eq. (30) is the hypothesis of Lemma 5 with $\lambda_M > 0$ from eq. (29) and $\kappa_M = v_0 + \lambda_M\{\delta + \frac{g'(v_0)C(\epsilon,\delta)}{M} + \frac{a_{20}}{M^2}\} - \frac{K_0}{M^3}$, so

$$\mathcal{V}_M(\Delta_2) \geq \Delta_2^2 M(\kappa_M - \lambda_M\delta) = M\Delta_2^2 v_0 + \Delta_2^2\lambda_M g'(v_0)C(\epsilon,\delta) + \frac{\Delta_2^2\lambda_M a_{20}}{M} - \frac{\Delta_2^2 K_0}{M^2}.$$

Expanding eq. (29) as $\lambda_M = -\frac{1}{g'(v_0)} + \frac{a_{11}}{g'(v_0)^2 M} + \mathcal{O}(M^{-2})$ gives

$$\mathcal{V}_M(\Delta_2) \geq M\Delta_2^2 v_0 - \Delta_2^2 C(\epsilon,\delta) + \frac{\Delta_2^2}{M}\cdot\frac{C(\epsilon,\delta)a_{11} - a_{20}}{g'(v_0)} - \frac{\Delta_2^2 K_{\mathrm{cv}}}{M^2},$$

and eq. (31) at $u_0 = u_*$ makes the $\frac{1}{M}$ coefficient $J(\epsilon,\delta)$. This is eq. (5). The constant K_{cv} dominates the collected $\mathcal{O}(M^{-2})$ terms and M_{cv} is the dimension beyond which Lemmas 8 and 9 apply. □

E Uniform expansion of the achievability mechanism

This section proves Theorem 3 in the order of dependence. Appendix E.2 gives the exact privacy profile of generalized gamma noise. Lemmas 11–14, Proposition 1, and Lemma 15 supply the uniform estimates behind its two-order expansion. Lemmas 16 and 17 give the expansion and the calibration inversion. Appendix E.4 gives closed forms for the correction coefficients, and Proposition 3 gives the calibrated error. Appendix E.5 maps the admissible parameter region. Appendix E.6 covers the computation of the CGG mechanism and Appendix E.7 proves the theorem.

E.1 Normalized density and sampling

For $\bar{Z} = \frac{Z}{\Delta_2} = \bar{R}U$ and $\bar{\beta} = \beta\Delta_2^p$, the density eq. (12) in sensitivity-normalized coordinates is

$$\bar{f}(r) = \frac{p\bar{\beta}^{\frac{q}{p}}}{\omega_{M-1}\Gamma(\frac{q}{p})} r^{q-M} e^{-\bar{\beta}r^p}. \tag{41}$$

The change of variables $t = \beta r^p$ in eq. (3) gives $f_T(t) = \frac{t^{\frac{q}{p}-1}e^{-t}}{\Gamma(\frac{q}{p})}$, which is the identity $T = \beta R^p \sim \mathrm{Gamma}(\frac{q}{p}, 1)$ used in eq. (11). It also gives Step 4 of Algorithm 1. The radius $R = \beta^{-\frac{1}{p}} T^{\frac{1}{p}} = aT^{\frac{1}{p}}$ is an exact draw, and $U = \frac{N}{\|N\|_2}$ with $N \sim \mathcal{N}(0, \mathbf{I}_M)$ independent of T is uniform on $\mathbb{S}^{M-1}$ and independent of R.

E.2 Exact privacy profile

For an absolutely continuous P with density f and privacy loss $L_t(z) = \log\frac{f(z)}{f(z-te_1)}$,

$$D_\epsilon(P\|P_{te_1}) = \int f(z)\left[1 - e^{\epsilon - L_t(z)}\right]_+ \mathrm{d}z = \mathbb{E}_P\left[1 - e^{\epsilon - L_t(Z)}\right]_+.$$

For $Z = RU$ with a density $f(z) = f_\circ(\|z\|_2)$ depending on the norm alone, conditioning on $R = r$ and on $U_1 = u$, where H_M is the sphere CDF of Lemma 4, gives

$$D_\epsilon(P\|P_{te_1}) = \int_0^\infty f_R(r)\int_{-1}^1 \left[1 - e^\epsilon \frac{f_\circ(\sqrt{r^2 + t^2 - 2rtu})}{f_\circ(r)}\right]_+ \mathrm{d}H_M(\sqrt{M}u)\mathrm{d}r.$$

Lemma 10 (Worst-case shift for radially decreasing noise, Proposition 4.1 and Lemma 4.2 of Wei et al. (2026))**.** *If $q \leq M$, the density of generalized gamma noise with shape q is nonincreasing in $\|z\|_2$. For any spherically symmetric P with nonincreasing density, $t \mapsto D_\epsilon(P\|P_{te_1})$ is nondecreasing on $[0,\infty)$. Hence $D_\epsilon(P\|P_{\Delta_2 e_1}) \leq \delta$ implies $D_\epsilon(P\|P_h) \leq \delta$ for every $\|h\|_2 \leq \Delta_2$.*

Corollary 3 (Scale monotonicity and endpoint limits)**.** *Let P have a nonincreasing spherical density and let P_s be the distribution of sZ with $Z \sim P$ and $s > 0$. Then*

$$\Psi_{\Delta_2}(s) = D_\epsilon(P_s\|(P_s)_{\Delta_2 e_1}) = D_\epsilon(P\|P_{\frac{\Delta_2}{s}e_1}) \tag{42}$$

is continuous and nonincreasing on $(0,\infty)$ *with* $\Psi_{\Delta_2}(0^+)=1$ *and* $\Psi_{\Delta_2}(\infty)=0$*. Hence for every* $\delta\in(0,1)$*, the feasible scales* $\{s:\Psi_{\Delta_2}(s)\le\delta\}$ *form a closed upper interval. If* $\Psi_{\Delta_2}(s_0)=\delta$ *with* Ψ_{Δ_2} *strictly decreasing near* s_0*, then* s_0 *is the smallest feasible scale.*

Proof. Applying the bijection $z\mapsto\frac{z}{s}$ to both arguments gives the equality in eq. (42), and Lemma 10 makes the right side nonincreasing in s (the shift monotonicity in the scale is Lemma C.1 of Wei et al. (2026)). Translation is continuous in L^1 and $|\min\{a,e^\epsilon b\}-\min\{a,e^\epsilon c\}|\le e^\epsilon|b-c|$, so $t\mapsto D_\epsilon(P\|P_{te_1})=1-\int\min\{f(z),e^\epsilon f(z-te_1)\}\mathrm{d}z$ is continuous, and so is Ψ_{Δ_2}. As $s\to\infty$, the shift $\frac{\Delta_2}{s}\to 0$ and $\int\min\{f,e^\epsilon f(\cdot-\frac{\Delta_2}{s}e_1)\}\to 1$. As $s\to 0$, the shift diverges, and approximating f in L^1 by a compactly supported function shows that the same integral tends to 0. Continuity makes the feasible set closed, and monotonicity makes it an upper interval. If $\Psi_{\Delta_2}(s_0)=\delta$ and some $s<s_0$ were feasible, monotonicity would force $\Psi_{\Delta_2}\equiv\delta$ on $[s,s_0]$, contradicting strict decrease near s_0. □

The following representation is Lemma 4.1 of Wei et al. (2026), written in the coordinates $(\bar R,\sqrt{M}U_1)$ and in terms of the radial potential. The roots $y_\pm$ appear explicitly in Appendix E.3. Let $\bar\varphi(r)=\bar\beta r^p+(M-q)\log r$, increasing for $q\le M$, and $y_\sigma=\bar\varphi^{-1}(\bar\varphi(r)+\sigma\epsilon)$ for $\sigma\in\{\pm1\}$. Then, with the integrand evaluated at $r=\bar R$,

$$D_\epsilon(P_Z\|(P_Z)_{\Delta_2 e_1})=\mathbb{E}_{\bar R}\left[H_M\left(\frac{\sqrt{M}(r^2+1-y_+^2)}{2r}\right)-e^\epsilon H_M\left(-\frac{\sqrt{M}(r^2+1-y_-^2)}{2r}\right)\right], \tag{43}$$

where the second term is zero when $\bar\varphi(r)-\epsilon$ lies below the range of $\bar\varphi$. Indeed $\bar f(r)\propto e^{-\bar\varphi(r)}$, so with $r=\frac{\|z\|_2}{\Delta_2}$ and $s=\frac{\|z-\Delta_2 e_1\|_2}{\Delta_2}$, the loss is $L_{\Delta_2}(z)=\bar\varphi(s)-\bar\varphi(r)$. The condition $L_{\Delta_2}>\epsilon$ is equivalent to $s>y_+$ and, since $s^2=r^2+1-2rU_1$, to $U_1<\frac{r^2+1-y_+^2}{2r}$, which gives the first term after conditioning on r. For the translated term, the change of variables $y\mapsto y+e_1$ turns the condition into $s<y_-$, whose angular form has the opposite sign. Weighting by e^ϵ and subtracting gives eq. (43). For $q>M$, the monotonicity reduction fails, so we must keep the supremum over shifts.

E.3 The attainable expansion

Parameterize generalized gamma noise by (ρ,p,v) with $q=\rho M$, $0<\rho\le 1$, $\tau=\rho p$, and the scale fixed by $\mathbb{E}[\bar R^2]=Mv$, so that the total error is $M\Delta_2^2 v$. By eq. (11),

$$\bar\beta_M(v)=\left[\frac{\Gamma(\frac{\rho M+2}{p})}{Mv\Gamma(\frac{\rho M}{p})}\right]^{\frac{p}{2}},\qquad \beta_M=\frac{\bar\beta_M}{\Delta_2^p}. \tag{44}$$

E.3.1 Uniform deviation, moment, and score bounds

All bounds below are uniform over

$$(\rho,p)\in\mathcal{K}=[\rho_{\min},1]\times[p_{\min},p_{\max}],\qquad 0<\rho_{\min}\le 1,\quad 0<p_{\min}\le p_{\max}<\infty.$$

Throughout Appendix E.3 $\nu=\frac{\rho M}{p}$ is the shape of the radial gamma variable $T=\bar\beta\bar R^p\sim\mathrm{Gamma}(\nu,1)$, $b=\frac{2}{p}$, $Y=\frac{T}{\nu}$, and

$$m=\mathbb{E}\left[T^b\right]=\frac{\Gamma(\nu+b)}{\Gamma(\nu)},\qquad 1+X=\frac{\bar R^2}{Mv}=\frac{T^b}{m},$$

where the second identity is the normalization $\mathbb{E}[\bar R^2]=Mv$. The distribution of X depends only on (ρ,p,M). The parameter v cancels, so we need no score in the v direction.

Lemma 11 (Gamma deviation (Laurent & Massart, 2000; Boucheron et al., 2013))**.** *Let* $T\sim\mathrm{Gamma}(\nu,1)$*. For* $s>0$ *and* $s\in(0,1)$ *respectively,*

$$\Pr\{T\ge\nu(1+s)\}\le e^{-\nu I_+(s)},\quad I_+(s)=s-\log(1+s)\ge\frac{s^2}{2(1+s)};$$

$$\Pr\{T\le\nu(1-s)\}\le e^{-\nu I_-(s)},\quad I_-(s)=-s-\log(1-s)\ge\frac{s^2}{2}.$$

Proof. Since $\mathbb{E}e^{\lambda T} = (1-\lambda)^{-\nu}$ for $\lambda < 1$, Markov's inequality gives $\Pr\{T \geq \nu(1+s)\} \leq \exp\{-\nu[\lambda(1+s) + \log(1-\lambda)]\}$, minimized at $\lambda = \frac{s}{1+s}$, which yields I_+. Likewise $\mathbb{E}e^{-\mu T} = (1+\mu)^{-\nu}$ with $\mu = \frac{s}{1-s}$ yields I_-. Both $I_\pm$ vanish to second order at 0 with $I_+'' = (1+s)^{-2}$ and $I_-'' = (1-s)^{-2} \geq 1$, giving the quadratic lower bounds. □

Lemma 12 (Uniform moment bounds)**.** *Fix integers $k, n \geq 0$ and $r > 0$. There are $C = C(k, n, r, \mathcal{K})$ and M_0 such that for $M \geq M_0$*

$$\mathbb{E}[|Y-1|^{2n}] \leq C\nu^{-n}, \qquad \mathbb{E}[|\log Y|^k] \leq C, \qquad \mathbb{E}[Y^r] + \mathbb{E}[Y^{-r}] \leq C, \qquad \mathbb{E}[X^{2n}] \leq C\nu^{-n}. \tag{45}$$

Proof. The jth cumulant of T is $\nu(j-1)!$, so $\mathbb{E}[(T-\nu)^{2n}]$ is a polynomial in ν of degree n, giving the first bound. Gautschi's inequality $x^{1-s} < \frac{\Gamma(x+1)}{\Gamma(x+s)} < (x+1)^{1-s}$, applied $\lceil r \rceil$ times to $\mathbb{E}[Y^{\pm r}] = \nu^{\mp r}\frac{\Gamma(\nu \pm r)}{\Gamma(\nu)}$, gives the third. Since $y \mapsto \frac{|\log y|^k}{y^{\frac{1}{2}} + y^{-\frac{1}{2}}}$ is bounded on $(0, \infty)$, the second follows from the third. For the fourth write $1 + X = c_M Y^b$ with $c_M = \frac{\nu^b}{m}$. Lemma 3 with $N = 1$ gives $|c_M - 1| \leq C\nu^{-1}$. The mean value theorem gives $|y^b - 1| \leq b|y-1|\max\{1, y^{b-1}\} \leq b|y-1|(1+y^{b-1})$ for every $b > 0$ and $y > 0$. By Cauchy–Schwarz, $(1+y^{b-1})^{4n} \leq 2^{4n-1}(1+y^{4n(b-1)})$, and the first and third bounds applied with $2n$ and with $r = 4n|b-1|$ give $\mathbb{E}[|Y^b - 1|^{2n}] \leq C\nu^{-n}$. Finally $X = (c_M - 1) + c_M(Y^b - 1)$ with $(x+y)^{2n} \leq 2^{2n-1}(x^{2n} + y^{2n})$ finishes. Constants depend on $\mathcal{K}$ only through $b = \frac{2}{p}$ and $\nu \asymp M$. □

Lemma 13 (Weighted radial tails)**.** *Let $W(T) = \sum_{i=1}^n \gamma_i T^{d_i}(\log T)^{k_i}$ with $|d_i| \leq d$, $0 \leq k_i \leq d$, $|\gamma_i| \leq \gamma_{\max}$, and let $\mathcal{A}_c = \{|X| > c\}$ for $c \in (0, \frac{1}{2}]$. There are $s = s(c, \mathcal{K}) > 0$, $C = C(d, n, c, \mathcal{K})$, and M_0 such that for $M \geq M_0$*

$$\int_{\mathcal{A}_c} |W(T)| \mathrm{d}\Pr \leq C\gamma_{\max}\nu^d(1+\log\nu)^d \exp\left\{-\frac{\nu}{2}\min\{I_+(s), I_-(s)\}\right\}.$$

In particular the left side is $\mathcal{O}(e^{-\xi M})$ for some $\xi(c, \mathcal{K}) > 0$, uniformly on $\mathcal{K}$.

Proof. From $1 + X = c_M Y^b$ with $|c_M - 1| \leq C\nu^{-1}$, for large M, the event $|X| > c$ forces $|Y^b - 1| \geq \frac{c}{2}$, hence $|Y - 1| \geq s$ for some $s(c, \mathcal{K}) > 0$ since $b \in [\frac{2}{p_{\max}}, \frac{2}{p_{\min}}]$. Lemma 11 gives $\Pr(\mathcal{A}_c) \leq 2e^{-\nu\min\{I_+(s), I_-(s)\}}$. By Cauchy–Schwarz, $\int_{\mathcal{A}_c} |W| \leq (\mathbb{E}W^2)^{\frac{1}{2}} \Pr(\mathcal{A}_c)^{\frac{1}{2}}$. Writing $T = \nu Y$ and $\log T = \log\nu + \log Y$, eq. (45) gives $\mathbb{E}W^2 \leq Cn\gamma_{\max}^2\nu^{2d}(1+\log\nu)^{2d}$. The prefactor is absorbed by halving the exponent, and $\nu \geq \frac{\rho_{\min} M}{p_{\max}}$ on $\mathcal{K}$. □

Lemma 14 (Explicit shape scores)**.** *Let $\pi_{\rho,p}$ be the density of X, $S_\alpha = \partial_\alpha \log \pi_{\rho,p}$ its score for $\alpha \in \{\rho, p\}$, $L = \log(1+X) = b\log T - \log m$, and $\psi = (\log\Gamma)'$. Then*

$$\begin{aligned}
S_\rho - \mathbb{E}[S_\rho] &= \frac{M}{2}(L - \mathbb{E}[L]) - \frac{p}{2}(\partial_\rho \log m)(T - \mathbb{E}[T]),\\
\partial_\rho \log m &= \frac{M}{p}\{\psi(\nu+b) - \psi(\nu)\} = \frac{2}{p\rho} + \mathcal{O}(M^{-1}),\\
S_p - \mathbb{E}[S_p] &= -\frac{1}{2}(TL - \mathbb{E}[TL]) - c_p(T - \mathbb{E}[T]),\\
c_p &= \frac{1}{2}\log m - \frac{\nu}{2}\{\psi(\nu+b) - \psi(\nu)\} - \frac{1}{p}\psi(\nu+b) = -\frac{1}{p} + \mathcal{O}(M^{-1}),
\end{aligned} \tag{46}$$

and there are C, M_0 such that for $M \geq M_0$ and $(\rho, p) \in \mathcal{K}$

$$\mathbb{E}[S_\rho^2] + \mathbb{E}[S_p^2] \leq CM. \tag{47}$$

Proof. With $t = m^{\frac{p}{2}}(1+x)^{\frac{p}{2}}$, the change of variables from $\mathrm{Gamma}(\nu, 1)$ gives, for $x > -1$,

$$\log \pi_{\rho,p}(x) = \left(\frac{\rho M}{2} - 1\right)\log(1+x) - t + \frac{\rho M}{2}\log m - \log\Gamma(\nu) + \log\frac{p}{2},$$

and at $x = X$ we have $t = T$ and $\log(1 + x) = L$. Differentiating in ρ with $\partial_\rho \nu = \frac{M}{p}$ and $\partial_\rho t = \frac{p}{2} t \partial_\rho \log m$ gives $S_\rho = \frac{M}{2} L - \frac{p}{2}(\partial_\rho \log m) T + \text{const}$. Differentiating in p with $\partial_p \nu = -\frac{\nu}{p}$ and $\partial_p t = t\{\frac{1}{2}(\log m + L) + \frac{p}{2}\partial_p \log m\}$ gives $S_p = -\frac{1}{2} TL - c_p T + \text{const}$. Centering yields eq. (46). The asymptotics use Lemma 3 with $N = 2$, namely $\log m = b \log \nu + \mathcal{O}(\nu^{-1})$, $\psi(\nu + b) = \log \nu + \mathcal{O}(\nu^{-1})$, and $\psi(\nu + b) - \psi(\nu) = \frac{b}{\nu} + \mathcal{O}(\nu^{-2})$. In c_p, the terms $\frac{b}{2} \log \nu$ and $-\frac{1}{p} \log \nu$ cancel exactly because $b = \frac{2}{p}$. This cancellation keeps c_p bounded and eq. (47) at order M.

For the variances, $\operatorname{Var} L = b^2 \psi'(\nu) \le C M^{-1}$, $\operatorname{Var} T = \nu \le CM$, and $TL = b\{\nu Y \log Y + (\log \nu - \frac{p}{2} \log m) T\}$ with $\log \nu - \frac{p}{2} \log m = \mathcal{O}(\nu^{-1})$. The second term has variance $\mathcal{O}(\nu^{-1})$, and $|y \log y| \le |y - 1|(1 + |\log y|)$ with eq. (45) gives $\operatorname{Var}(\nu Y \log Y) \le \nu^2 (\mathbb{E}(Y - 1)^4)^{\frac{1}{2}} (\mathbb{E}(1 + |\log Y|)^4)^{\frac{1}{2}} \le C\nu$. Hence $\operatorname{Var}(TL) \le CM$. Applying $\operatorname{Var}(W_1 + W_2) \le 2 \operatorname{Var} W_1 + 2 \operatorname{Var} W_2$ to eq. (46), whose coefficients are bounded, gives eq. (47). □

Corollary 4 (Differentiating an expectation under the shape)**.** *Let φ_α be measurable with $\mathbb{E}_\alpha[\varphi_\alpha(X)^2] < \infty$ and $\partial_\alpha \varphi_\alpha$ locally dominated. For $\alpha \in \{\rho, p\}$,*

$$\partial_\alpha \mathbb{E}_\alpha[\varphi_\alpha(X)] = \mathbb{E}_\alpha[\partial_\alpha \varphi_\alpha(X)] + \operatorname{Cov}_\alpha(\varphi_\alpha(X), S_\alpha(X)), \qquad |\operatorname{Cov}_\alpha| \le \sqrt{\operatorname{Var}_\alpha \varphi_\alpha(X)} \sqrt{CM}.$$

If $|\varphi_\alpha(x)| \le C_\varphi |x|^n$ near 0 and φ_α is bounded by a fixed polynomial-logarithmic weight elsewhere, then $\partial_\alpha \mathbb{E}_\alpha[\varphi_\alpha(X)] = \mathbb{E}_\alpha[\partial_\alpha \varphi_\alpha(X)] + \mathcal{O}(M^{\frac{1-n}{2}})$.

Proof. Differentiating $\int \varphi_\alpha \pi_{\rho,p}$ with $\partial_\alpha \pi_{\rho,p} = \pi_{\rho,p} S_\alpha$ and $\mathbb{E}_\alpha S_\alpha = \partial_\alpha \int \pi_{\rho,p} = 0$ gives the covariance form. Cauchy–Schwarz and eq. (47) give the bound. Splitting at $|X| \le c$, the central part has $\operatorname{Var} \varphi_\alpha \le C_\varphi^2 \mathbb{E} X^{2n} \le C C_\varphi^2 M^{-n}$ by eq. (45), and the complement is $\mathcal{O}(e^{-\xi M})$ by Lemma 13. Hence the covariance is $\mathcal{O}(M^{-\frac{n}{2}}) \mathcal{O}(M^{\frac{1}{2}})$. □

Corollary 4 fixes the Taylor orders in the proof of Lemma 16. A remainder of order n contributes $\mathcal{O}(M^{\frac{1-n}{2}})$ after one shape derivative. An $\mathcal{O}(M^{-3})$ derivative bound therefore requires $n \ge 7, 5, 3$ for the terms carrying $1, M^{-1}, M^{-2}$, that is, expansions through degrees six, four, and two.

E.3.2 Radial moments and scale

Proposition 1 (Normalized radial moments and scale expansion)**.** *With $X = \frac{\bar{R}^2}{Mv} - 1$, the scale eq. (44), $\nu = \frac{\rho M}{p}$, and $\tau = \rho p$, we have $\mathbb{E}X = 0$,*

$$\mathbb{E}[(1 + X)^j] = \frac{\Gamma(\nu + \frac{2j}{p}) \Gamma(\nu)^{j-1}}{\Gamma(\nu + \frac{2}{p})^j},$$

$$\begin{aligned} \mathbb{E}[X^2] &= \frac{4}{\tau M} + \frac{2(p-2)^2}{\tau^2 M^2} + \mathcal{O}(M^{-3}), \qquad \mathbb{E}[X^3] = \frac{8(6-p)}{\tau^2 M^2} + \mathcal{O}(M^{-3}), \\ \mathbb{E}[X^4] &= \frac{48}{\tau^2 M^2} + \mathcal{O}(M^{-3}), \qquad \mathbb{E}[X^5] = \mathcal{O}(M^{-3}), \qquad \mathbb{E}[|X|^6] = \mathcal{O}(M^{-3}), \end{aligned} \tag{48}$$

and

$$\bar{\beta}_M(v) = \frac{\rho M}{p (Mv)^{\frac{p}{2}}} \left[1 + \frac{2 - p}{2\rho M} + \frac{p^2 - 4}{24 \rho^2 M^2} + \mathcal{O}(M^{-3}) \right]. \tag{49}$$

Proof. By eq. (11), $\mathbb{E}[\bar{R}^{2j}] = \bar{\beta}^{-\frac{2j}{p}} \frac{\Gamma(\nu + \frac{2j}{p})}{\Gamma(\nu)}$. Dividing by the jth power of the normalization $Mv = \bar{\beta}^{-\frac{2}{p}} \frac{\Gamma(\nu + \frac{2}{p})}{\Gamma(\nu)}$ gives the moment formula, and $j = 1$ gives $\mathbb{E}[X] = 0$. Apply eq. (17) with $N = 4$ at $b = \frac{2j}{p}$ and $b = \frac{2}{p}$, subtract j copies of the latter, exponentiate, and use $\nu = \frac{\tau M}{p^2}$. With $\mu_j =$

$\mathbb{E}[(1+X)^j]$,

$$\mu_2 = 1 + \frac{4}{\tau M} + \frac{2(p-2)^2}{\tau^2 M^2} + \frac{2(p^4-12p^3+40p^2-48p+16)}{3\tau^3 M^3} + \mathcal{O}(M^{-4}),$$
$$\mu_3 = 1 + \frac{12}{\tau M} + \frac{6p^2-32p+72}{\tau^2 M^2} + \frac{2p^4-32p^3+176p^2-384p+288}{\tau^3 M^3} + \mathcal{O}(M^{-4}),$$
$$\mu_4 = 1 + \frac{24}{\tau M} + \frac{12p^2-80p+288}{\tau^2 M^2} + \frac{4p^4-80p^3+624p^2-1920p+2304}{\tau^3 M^3} + \mathcal{O}(M^{-4}),$$
$$\mu_5 = 1 + \frac{40}{\tau M} + \frac{20p^2-160p+800}{\tau^2 M^2} + \frac{20p^4-480p^3+4880p^2-19200p+32000}{3\tau^3 M^3} + \mathcal{O}(M^{-4}),$$
$$\mu_6 = 1 + \frac{60}{\tau M} + \frac{30p^2-280p+1800}{\tau^2 M^2} + \frac{10p^4-280p^3+3520p^2-16800p+36000}{\tau^3 M^3} + \mathcal{O}(M^{-4}).$$

Then $\mathbb{E}[X^k] = \sum_{j=0}^{k}(-1)^{k-j}\binom{k}{j}\mu_j$ gives

$$\mathbb{E}[X^2] = \frac{4}{\tau M} + \frac{2(p-2)^2}{\tau^2 M^2} + \mathcal{O}(M^{-3}),$$
$$\mathbb{E}[X^3] = \frac{8(6-p)}{\tau^2 M^2} - \frac{8(p^3-12p^2+36p-32)}{\tau^3 M^3} + \mathcal{O}(M^{-4}),$$
$$\mathbb{E}[X^4] = \frac{48}{\tau^2 M^2} + \frac{16(5p^2-36p+76)}{\tau^3 M^3} + \mathcal{O}(M^{-4}),$$
$$\mathbb{E}[X^5] = \frac{320(6-p)}{\tau^3 M^3} + \mathcal{O}(M^{-4}),$$
$$\mathbb{E}[X^6] = \frac{960}{\tau^3 M^3} + \mathcal{O}(M^{-4}),$$

which is eq. (48). For eq. (49) take logarithms in eq. (44), apply eq. (17) with $b = \frac{2}{p}$, and use $\nu = \frac{\rho M}{p}$. This gives $\log \frac{\bar{\beta}_M(v)p(Mv)^{\frac{p}{2}}}{\rho M} = \frac{2-p}{2\rho M} - \frac{(p-2)(p-4)}{12\rho^2 M^2} + \mathcal{O}(M^{-3})$. After exponentiating, the M^{-2} coefficient is $-\frac{(p-2)(p-4)}{12\rho^2} + \frac{(2-p)^2}{8\rho^2} = \frac{p^2-4}{24\rho^2}$. □

E.3.3 Expanding the exact radial thresholds

Let $w = v(1+x)$, $r = \sqrt{Mw}$, and $u_\sigma = \frac{y_\sigma^2}{r^2} - 1$. The roots in eq. (43) solve

$$\sigma\epsilon = M\left\{\frac{K}{p}\left[(1+u_\sigma)^{\frac{p}{2}} - 1\right] + \frac{1-\rho}{2}\log(1+u_\sigma)\right\}, \tag{50}$$

where by eq. (49) $K = \rho(1+x)^{\frac{p}{2}}\{1 + \frac{\Lambda_1}{M} + \frac{\Lambda_2}{M^2} + \mathcal{O}(M^{-3})\}$ with $\Lambda_1 = \frac{2-p}{2\rho}$ and $\Lambda_2 = \frac{p^2-4}{24\rho^2}$. Let $\eta = M^{-1}$, $u_\sigma = \eta t_{0\sigma} + \eta^2 t_{1\sigma} + \eta^3 t_{2\sigma} + \mathcal{O}(\eta^4)$, $K_0 = \rho(1+x)^{\frac{p}{2}}$, $K_i = K_0\Lambda_i$, $A_0 = 1-\rho+K_0$, and $\Lambda_0 = K_0(p-2) - 2(1-\rho)$. Expanding $(1+u)^{\frac{p}{2}} - 1 = \frac{p}{2}u + \frac{p(p-2)}{8}u^2 + \frac{p(p-2)(p-4)}{48}u^3 + \mathcal{O}(u^4)$ and $\log(1+u) = u - \frac{u^2}{2} + \frac{u^3}{3} + \mathcal{O}(u^4)$ in eq. (50), the orders η^0, η^1 give $\sigma\epsilon = \frac{A_0}{2}t_{0\sigma}$ and $0 = \frac{A_0}{2}t_{1\sigma} + \frac{K_1}{2}t_{0\sigma} + \frac{\Lambda_0}{8}t_{0\sigma}^2$, hence

$$t_{0\sigma} = \frac{2\sigma\epsilon}{A_0}, \qquad t_{1\sigma} = -\frac{2\sigma\epsilon K_1}{A_0^2} - \frac{\Lambda_0\epsilon^2}{A_0^3}.$$

The two CDF arguments in eq. (43) are exactly $a_{\sigma,M} = \frac{\sigma}{2\sqrt{w}} - \frac{\sigma\sqrt{w}}{2}\{t_{0\sigma} + \eta t_{1\sigma} + \eta^2 t_{2\sigma} + \mathcal{O}(\eta^3)\} = a_{0\sigma} + \eta a_{1\sigma} + \eta^2 a_{2\sigma} + \mathcal{O}(\eta^3)$ with

$$a_{0\sigma}(x) = \frac{\sigma}{2\sqrt{w}} - \frac{\epsilon\sqrt{w}}{A_0}, \qquad a_{1\sigma}(x) = \frac{\epsilon\sqrt{w}K_1}{A_0^2} + \frac{\sigma\epsilon^2\sqrt{w}\Lambda_0}{2A_0^3}. \tag{51}$$

At $x = 0$ we have $K_0 = \rho$, $A_0 = 1$, and $K_1 = \frac{2-p}{2}$. One further order of eq. (50) reads $\sigma\epsilon = \frac{1}{2}t + \eta\{\kappa_{11}t + \kappa_{20}t^2\} + \eta^2\{\kappa_{12}t + \kappa_{21}t^2 + \kappa_{30}t^3\} + \mathcal{O}(\eta^3)$ with $t = t_{0\sigma} + \eta t_{1\sigma} + \eta^2 t_{2\sigma}$ and

$$t_{0\sigma}(0) = 2\sigma\epsilon, \qquad t_{1\sigma}(0) = (p-2)\sigma\epsilon + (2-\tau)\epsilon^2,$$
$$\kappa_{11} = \frac{2-p}{4}, \qquad \kappa_{12} = \frac{p(p^2-4)}{48\tau}, \qquad \kappa_{20} = \frac{\tau-2}{8},$$
$$\kappa_{21} = -\frac{(p-2)^2}{16}, \qquad \kappa_{30} = \frac{\tau(p-6)+8}{48}.$$

Setting the η^2 coefficient to zero gives $t_{2\sigma}(0) = -2[\kappa_{11}t_{1\sigma} + \kappa_{12}t_{0\sigma} + 2\kappa_{20}t_{0\sigma}t_{1\sigma} + \kappa_{21}t_{0\sigma}^2 + \kappa_{30}t_{0\sigma}^3]$ and $a_{2\sigma}(0) = -\frac{\sigma\sqrt{v}}{2}t_{2\sigma}(0)$. The leading derivative of the root equation is $\frac{A_0}{2} > 0$, so the roots are locally unique. For $q = M$, the small-radius region where y_- does not exist lies outside a fixed neighborhood of $X = 0$ for large M and falls under the exponential tail bound.

E.3.4 Expectation and calibration inversion

For expressions indexed by σ write $\mathcal{D}F_\sigma = F_+ - e^\epsilon F_-$ and define

$$b_0(x) = \mathcal{D}\Phi(a_{0\sigma}(x)), \qquad b_1(x) = \mathcal{D}\{\phi(a_{0\sigma})[a_{1\sigma} + W_1(a_{0\sigma})]\},$$
$$b_2(0) = \mathcal{D}\left\{\phi(a_{0\sigma})\left[a_{2\sigma} - \frac{1}{2}a_{0\sigma}a_{1\sigma}^2 + P_1(a_{0\sigma})a_{1\sigma} + W_2(a_{0\sigma})\right]\right\}_{x=0}, \tag{52}$$

so that $b_0(0) = g(v)$. The function $b_2(x)$ is smooth near 0 by the parameterized implicit function theorem (Krantz & Parks, 2012), and we only need its value at 0. Taylor expansion in x with $\mathbb{E}X = 0$ and eq. (48) gives $\mathbb{E}[b_0(X)] = b_0(0) + \frac{2b_0''(0)}{\tau M} + \frac{1}{M^2}\{\frac{(p-2)^2}{\tau^2}b_0''(0) + \frac{4(6-p)}{3\tau^2}b_0'''(0) + \frac{2}{\tau^2}b_0^{(4)}(0)\} + \mathcal{O}(M^{-3})$, $\frac{1}{M}\mathbb{E}[b_1(X)] = \frac{b_1(0)}{M} + \frac{2b_1''(0)}{\tau M^2} + \mathcal{O}(M^{-3})$, and $\frac{1}{M^2}\mathbb{E}[b_2(X)] = \frac{b_2(0)}{M^2} + \mathcal{O}(M^{-3})$. Adding,

$$D_{\rho,p,M}(\epsilon; v) = g(v) + \frac{a_1(v)}{M} + \frac{a_2(v)}{M^2} + \mathcal{O}(M^{-3}),$$
$$a_1 = b_1(0) + \frac{2}{\tau}b_0''(0) = g'(v)C_\tau(v),$$
$$a_2 = b_2(0) + \frac{2}{\tau}b_1''(0) + \frac{(p-2)^2}{\tau^2}b_0''(0) + \frac{4(6-p)}{3\tau^2}b_0'''(0) + \frac{2}{\tau^2}b_0^{(4)}(0), \tag{53}$$

with all derivatives taken in x.

Lemma 15 (Uniform finite-order implicit expansion)**.** *Let ϑ range over a compact set, I a compact t interval, $\eta_M \to 0$, and*

$$F_M(t, \vartheta) = F_0(t, \vartheta) + \eta_M F_1(t, \vartheta) + \eta_M^2 F_2(t, \vartheta) + \eta_M^3 R_M(t, \vartheta) \tag{54}$$

on I, with F_i, R_M, $\partial_t R_M$, $\partial_\vartheta R_M$, $\partial_t\partial_\vartheta R_M$, $\partial_t^2 F_i$, and $\partial_t\partial_\vartheta F_i$ uniformly bounded. Suppose $\widehat{t}_M = t_0 + \eta_M t_1 + \eta_M^2 t_2$ satisfies $F_M(\widehat{t}_M, \vartheta) = O_{C^1(\vartheta)}(\eta_M^3)$, stays a fixed distance from ∂I, and $\partial_t F_M \geq c_F > 0$ on I. Then for large M there is a unique root t_M within $\mathcal{O}(\eta_M^3)$ of $\widehat{t}_M$, and $t_M = \widehat{t}_M + \eta_M^3 r_M(\vartheta)$ with $\|r_M\|_{C^1(\vartheta)} \leq C$.

Proof. If $|F_M(\widehat{t}_M, \vartheta)| \leq C_0\eta_M^3$ and $\omega > \frac{C_0}{c_F}$, then $F_M(\widehat{t}_M \pm \omega\eta_M^3, \vartheta)$ have opposite signs. Hence a root exists between them, unique by monotonicity, and $|t_M - \widehat{t}_M| \leq \frac{C_0\eta_M^3}{c_F}$ by the mean value theorem. For the parameter derivative, $\partial_\vartheta t_M - \partial_\vartheta\widehat{t}_M = -\frac{\partial_\vartheta F_M(t_M,\vartheta) + \partial_t F_M(t_M,\vartheta)\partial_\vartheta\widehat{t}_M}{\partial_t F_M(t_M,\vartheta)}$. At $t = \widehat{t}_M$, the numerator is $\partial_\vartheta[F_M(\widehat{t}_M, \vartheta)] = \mathcal{O}(\eta_M^3)$, and its change between $\widehat{t}_M$ and t_M is $\mathcal{O}(\eta_M^3)$ by the bounds on $\partial_t\partial_\vartheta F_M$, $\partial_t^2 F_M$, and $\partial_\vartheta\widehat{t}_M$. Dividing by c_F gives the claim. □

Lemma 16 (Uniform privacy profile expansion)**.** *Let $K_v, K_p \Subset (0, \infty)$ and $K_\rho = [\rho_{\min}, \rho_{\max}]$ with $0 < \rho_{\min} \leq \rho_{\max} \leq 1$. There are C and M_0 such that for $M \geq M_0$, $v \in K_v$, $\rho \in K_\rho$, $p \in K_p$,*

$$\max_{m=0,1}\left|\partial_v^m\left\{D_{\rho,p,M}(\epsilon; v) - g(v) - \frac{a_1(v)}{M} - \frac{a_2(v)}{M^2}\right\}\right| \leq \frac{C}{M^3}, \tag{55}$$

and the same holds after one derivative in ρ or p on the relative interior of its set, one-sided at $\rho = 1$.

Proof. Central region. Choose $c \in (0,1)$ so that $w = v(1+x)$ stays in a compact subset of $(0,\infty)$ for $v \in K_v$ and $|x| \le c$. By Lemma 3, $K = K_0 + \eta K_1 + \eta^2 K_2 + \eta^3 R_{K,M}$ with $R_{K,M}$ and its first v, ρ, p derivatives bounded. Substituting into eq. (50) with finite Taylor formulas for $(1+\eta t)^{\frac{p}{2}}$ and $\log(1+\eta t)$ gives a family of the form eq. (54) whose leading t-derivative is $\frac{A_0}{2} \ge \inf A_0 > 0$. The displayed $t_{0\sigma}, t_{1\sigma}, t_{2\sigma}$ cancel three orders, so Lemma 15 gives $t_\sigma = t_{0\sigma} + \eta t_{1\sigma} + \eta^2 t_{2\sigma} + \eta^3 r_{\sigma,M}$ with $r_{\sigma,M}$ bounded in C^1 of the parameters. Both sphere-CDF arguments then lie in a fixed bounded interval with uniform C^1 expansions, and Lemma 4 makes the conditional integrand $b_0(x) + \eta b_1(x) + \eta^2 b_2(x) + \eta^3 r_M(x)$ on $|x| \le c$, with the required parameter derivatives of r_M bounded. Expand b_0, b_1, b_2 through degrees six, four, and two with integral remainders. Uniform versions of eq. (48) make the expected remainders $\mathcal{O}(M^{-3})$ after the factors $1, M^{-1}, M^{-2}$. The polynomial terms are handled by the exact moment identities of Proposition 1 and eq. (45), each contributing a fixed power of M^{-1} times a bounded coefficient. For shape derivatives, Corollary 4 turns each remainder of order $n \in \{7,5,3\}$ into $\mathcal{O}(M^{\frac{1-n}{2}})$, hence $\mathcal{O}(M^{-3})$ after its prefactor, and the direct $\partial_\alpha \varphi_\alpha$ terms obey the same power bounds. The scores of Lemma 14 and the moments of Lemma 12 are uniform on $\mathcal{K}$, also after multiplication by a smooth central cutoff. A cutoff $\varkappa$ equal to one on $|X| \le \frac{c}{2}$ and zero on $|X| \ge c$ has derivatives supported on an exponentially unlikely annulus, so replacing cutoff moments by the full moments in eq. (48) changes every coefficient and derivative by $\mathcal{O}(e^{-c_1 M})$ (Lemmas 11 and 13).

Tail region. Differentiating in the coordinate X or in the roots $y_\pm$, which move with (v,ρ,p), produces terms such as $\partial_\alpha \log \|\bar{R}U - e_1\|_2$ that are singular at $z = e_1$ and not uniform as $\rho \uparrow 1$. We therefore differentiate in the fixed normalized Lebesgue coordinate z (the factors $\Delta_2^{\pm M}$ cancel). With $r = \|z\|_2$ and $s = \|z - e_1\|_2$, and with $\bar{f}_\alpha$ denoting the density eq. (41) as a function of the parameter $\alpha \in \{v,\rho,p\}$,

$$D_\epsilon(P_{\bar{Z}} \| (P_{\bar{Z}})_{e_1}) = \int_{\mathbb{R}^M} [\bar{f}_\alpha(r) - e^\epsilon \bar{f}_\alpha(s)]_+ \mathrm{d}z,$$

and since $x \mapsto x_+$ is one-Lipschitz,

$$|\partial_\alpha [\bar{f}_\alpha(r) - e^\epsilon \bar{f}_\alpha(s)]_+| \le |\partial_\alpha \bar{f}_\alpha(r)| + e^\epsilon |\partial_\alpha \bar{f}_\alpha(s)| \tag{56}$$

at differentiability points, with the same bound for difference quotients. No derivative of r or s appears. For $q \le M$, the map $z \mapsto \log \bar{f}_\alpha(r) - \log \bar{f}_\alpha(s) - \epsilon$ is real analytic and nonconstant off $\{0, e_1\}$, so its zero set is null and the difference quotients converge almost everywhere. Dominated convergence (Folland, 2013) then justifies differentiation. Next, $\partial_\alpha \bar{f}_\alpha = \bar{f}_\alpha \tilde{S}_\alpha$, where the density score $\tilde{S}_\alpha$ is a finite combination of 1, $\log r$, $\bar{\beta} r^p$, and $\bar{\beta} r^p \log r$ with coefficients polynomial in M and uniform on $\mathcal{K}$ by the digamma bounds in the proof of Lemma 3. Under $\bar{f}_\alpha$ we have $T = \bar{\beta} r^p \sim \mathrm{Gamma}(\nu, 1)$, and the normalization gives $\frac{r^2}{Mv} = (\frac{T}{\nu})^{\frac{2}{p}}\{1 + \mathcal{O}(M^{-1})\}$, so $|\frac{r^2}{Mv} - 1| > \frac{c}{4}$ forces $T < \lambda_- \nu$ or $T > \lambda_+ \nu$ for fixed $0 < \lambda_- < 1 < \lambda_+$. Lemma 13 with $W = \tilde{S}_\alpha$ yields

$$\int_{\{|\frac{\|z\|_2^2}{Mv} - 1| > \frac{c}{4}\}} |\partial_\alpha \bar{f}_\alpha(\|z\|_2)| \mathrm{d}z \le C_2 e^{-c_2 M}, \tag{57}$$

where the surface-area factor is part of the M-dimensional integral. Finally $|r - s| \le 1$ and v compact give, for large M,

$$\left| \frac{r^2}{Mv} - 1 \right| > \frac{c}{2} \implies \left| \frac{s^2}{Mv} - 1 \right| > \frac{c}{4},$$

so after $z \mapsto z + e_1$, the second term of eq. (56) is also bounded by eq. (57). For the v derivative, the cutoff contributes only on $\frac{c}{2} \le |X| \le c$ with bounded derivative, and for ρ, p, the cutoff is parameter independent. Hence the tail value and every stated derivative are $\mathcal{O}(e^{-c_3 M})$. This argument is uniform as $\rho \uparrow 1$. At $\rho = 1$, the event $\bar{\beta}\bar{R}^p \le \epsilon$ where y_- is absent has probability $\mathcal{O}(e^{-c_4 M})$, and the fixed-z proof needs no root there.

Adding the central expansion, with C^1 remainder CM^{-3}, and the exponentially small tail proves eq. (55). □

The second calibration correction is

$$J_{\tau,p}(v) = C_\tau(v) C'_\tau(v) + \frac{g''(v)}{2g'(v)} C_\tau(v)^2 - \frac{a_2(v)}{g'(v)}, \tag{58}$$

with τ held fixed in C'_τ.

Lemma 17 (Uniform calibration inversion)**.** *Under the conditions of Lemma 16 with $K_\rho \times K_p \supset \mathcal{K}$, there are $C_\mathcal{K}$ and $M_\mathcal{K}$ such that for $(\rho, p) \in \mathcal{K}$ and $M \geq M_\mathcal{K}$, the smallest feasible mean energy satisfies*

$$\left| v_M - v_0 + \frac{C_\tau(v_0)}{M} - \frac{J_{\tau,p}(v_0)}{M^2} \right| \leq \frac{C_\mathcal{K}}{M^3}, \qquad \tau = \rho p,$$

also along dimension-dependent sequences $(\rho_M, p_M) \in \mathcal{K}$, including $\rho_M \uparrow 1$.

Proof. With $\widehat{v}_M = v_0 + \frac{c}{M} + \frac{j}{M^2}$ and $g(v_0) = \delta$, expanding eq. (53) gives $D_{\rho,p,M}(\epsilon; \widehat{v}_M) - \delta = \frac{g'(v_0)c + a_1(v_0)}{M} + \frac{g'(v_0)j + \frac{1}{2}g''(v_0)c^2 + a_1'(v_0)c + a_2(v_0)}{M^2} + \mathcal{O}(M^{-3})$. The M^{-1} coefficient vanishes at $c = -C_\tau(v_0)$. With $a_1' = g''C_\tau + g'C_\tau'$ the M^{-2} coefficient vanishes at $j = J_{\tau,p}(v_0)$ by eq. (58). Continuity of $C_\tau, J_{\tau,p}$ on $\mathcal{K}$ keeps $\widehat{v}_M$ near v_0, and Lemma 16 gives $\sup_\mathcal{K} |D_{\rho,p,M}(\epsilon; \widehat{v}_M) - \delta| \leq C_0 M^{-3}$ and $\partial_v D_{\rho,p,M} = g'(v) + \mathcal{O}(M^{-1}) \leq -c_F < 0$ near v_0 for some $c_F > 0$. For $\omega > \frac{C_0}{c_F}$, the values at $\widehat{v}_M \pm \omega M^{-3}$ straddle δ, so a root exists and is unique nearby with $|v_M - \widehat{v}_M| \leq \omega M^{-3}$. Since $q = \rho M \leq M$, Corollary 3 identifies this local crossing with the global smallest feasible scale. □

E.4 Explicit correction coefficients

Proposition 2 (Explicit correction polynomial)**.** *For $v > 0$ let $E_v = \epsilon^2 v^2$ and $Q_v = E_v + 3v - \frac{1}{4}$. The coefficient defined by eq. (51)–eq. (58) is $J_{\tau,p}(v) = c_0(v,\tau) + c_1(v,\tau)p + c_2(v,\tau)p^2$ with $c_2(v,\tau) = \frac{Q_v}{2\tau^2}$,*

$$c_1(v,\tau) = \frac{1}{48v\tau^2}\Big\{16E_v^2(2\tau^3 - 3\tau^2 + 1) - 48E_v v(2\tau^3 - \tau^2) + 4E_v(3\tau^2 - 2) - 48v^2 - 16v + 1\Big\}, \tag{59}$$

$$\begin{aligned} c_0(v,\tau) = -\frac{1}{192v\tau^2}\Big\{ & E_v^2(192\tau^4 - 384\tau^3 + 160\tau^2) \\ & + E_v v(-576\tau^4 + 1152\tau^3 - 1992\tau^2 + 2208\tau - 672) \\ & + E_v(96\tau^3 - 100\tau^2 - 48\tau + 48) \\ & + v^2(-24\tau^2 + 864\tau - 864) + v(-2\tau^2 - 216\tau + 216) - \tau^2 + 12\tau - 12\Big\}. \end{aligned} \tag{60}$$

The Gaussian case gives $C_2(v) = 0$ and $J_{2,2}(v) = 0$.

Proof. Write $a_\sigma = a_{0\sigma}$ and $\ell_\sigma = a_{1\sigma} + W_1(a_\sigma)$. From $\Phi' = \phi$ and $\phi'(a) = -a\phi(a)$, for a smooth $a = a(x)$,

$$\begin{aligned} \frac{\mathrm{d}^2}{\mathrm{d}x^2}\Phi(a) &= \phi(a)\{a'' - a(a')^2\}, \\ \frac{\mathrm{d}^3}{\mathrm{d}x^3}\Phi(a) &= \phi(a)\{a''' - 3aa'a'' + (a^2 - 1)(a')^3\}, \\ \frac{\mathrm{d}^4}{\mathrm{d}x^4}\Phi(a) &= \phi(a)\{a^{(4)} - 4aa'a''' - 3a(a'')^2 + 6(a^2 - 1)(a')^2a'' + (3a - a^3)(a')^4\}, \\ \frac{\mathrm{d}^2}{\mathrm{d}x^2}\{\phi(a)\ell\} &= \phi(a)\left[\ell'' - 2aa'\ell' + \{(a^2 - 1)(a')^2 - aa''\}\ell\right], \end{aligned}$$

which are the only derivatives eq. (52) and eq. (53) require. At $x = 0$ we have $a_+(0) = A(v)$ and $a_-(0) = B(v)$, and $A^2 - B^2 = -2\epsilon$ gives $e^\epsilon\phi(B) = \phi(A)$. By eq. (10), $\frac{g''}{2g'} = -\frac{Q_v}{4v^2}$ and $-\frac{1}{g'} = \frac{2v^{\frac{3}{2}}}{\phi(A)}$, so the factor $\phi(A)$ cancels in eq. (58). The derivatives $a_\sigma^{(m)}(0)$, $\ell_\sigma^{(m)}(0)$ follow from eq. (51) by $\frac{\mathrm{d}^m}{\mathrm{d}x^m}(1 + x)^\gamma|_0 = \gamma(\gamma - 1)\cdots(\gamma - m + 1)$ and the reciprocal rule. Inserting them and $t_{2\sigma}(0)$ from the root expansion above makes every term a rational function of (v, ϵ, τ, p). The cubic powers of p from $t_{2\sigma}$ cancel between the two signs after $e^\epsilon\phi(B) = \phi(A)$, and sorting by $1, p, p^2$ gives $c_2 = \frac{Q_v}{2\tau^2}$ and the numerators eq. (59)–eq. (60), with no asymptotic replacement. Finally, $\tau = 2$ in eq. (14) gives $C_2(v) = \frac{1}{2}E_v + \frac{3}{2}v - \frac{1}{8} - \frac{1}{2}Q_v = 0$, and $(\tau, p) = (2, 2)$ in $c_0 + 2c_1 + 4c_2$ gives zero term by term. □

E.4.1 Error of the exactly calibrated distribution

Proposition 3 (Exactly calibrated generalized gamma error). *Fix $0 < \rho \le 1$, $p > 0$, $\tau = \rho p$, and the smallest exactly calibrated scale. Then*

$$\mathrm{MSE}_M^{\rho,p}(\Delta_2) = M\Delta_2^2 v_0 - \Delta_2^2 C_\tau(v_0) + \frac{\Delta_2^2 J_{\tau,p}(v_0)}{M} + \mathcal{O}(\Delta_2^2 M^{-2}), \tag{61}$$

uniformly for (ρ, p) in compact sets with ρ bounded away from zero, $\rho = 1$ included. The function $\tau \mapsto C_\tau(v_0)$ has maximum $C(\epsilon, \delta)$, attained only at τ_*, and $J_{\tau_*,p}(v_0) = J(\epsilon, \delta) + \frac{E}{2}(p - p_*)^2$.

Proof. By Lemma 17, $v_M = v_0 - \frac{C_\tau(v_0)}{M} + \frac{J_{\tau,p}(v_0)}{M^2} + \mathcal{O}(M^{-3})$ uniformly, and $\mathrm{MSE}_M^{\rho,p}(\Delta_2) = M\Delta_2^2 v_M$ gives eq. (61). The remaining claims are the concavity of $\tau \mapsto C_\tau(v_0)$, the value $C_{\tau_*}(v_0) = C(\epsilon, \delta)$, and the completed square for $J_{\tau_*,p}(v_0)$, all established in Appendix B.4 and recorded in eq. (16). □

E.5 Where the assumptions hold

Corollary 1 needs both assumptions, and each fails along one curve in the privacy plane. Assumption 1 fails at the root $v_{\tan}(\epsilon)$ of $1 - \{g(v) - vg'(v)\} = 0$ beyond v_c. Assumption 2 fails where $p_*(v) = \tau_*(v)$, at $v_{\mathrm{shape}}(\epsilon)$. Since g decreases, the two mark out the matching window $\delta \in (\delta_{\mathrm{shape}}, \delta_{\tan})$ with $\delta_{\tan} = g(v_{\tan})$ and $\delta_{\mathrm{shape}} = g(v_{\mathrm{shape}})$. Inside it, the converse and the achievability bound carry the same $\frac{1}{M}$ coefficient, which is the content of Corollary 1. The tangent boundary is the threshold $\delta^\star$ that Wei et al. (2026) compute for first-order optimality. The shape boundary is new and algebraic (Proposition 5) and is located by exact real-root isolation. This subsection locates both and maps the window.

Proposition 4 (Strong-privacy shape transition). *Fix $\epsilon > 0$ and let $\delta \downarrow 0$, equivalently $v_0 \to \infty$. Then*

$$\tau_* = 1 + \frac{3}{2\epsilon^2 v_0} - \frac{\epsilon^2 + 9}{8\epsilon^4 v_0^2} + \mathcal{O}(v_0^{-3}), \qquad p_* - \tau_* = \frac{3 - \epsilon^2}{12\epsilon^2 v_0} + \frac{\epsilon^2 + 45}{12\epsilon^4 v_0^2} + \mathcal{O}(v_0^{-3}). \tag{62}$$

Hence for all small δ we have $p_* > \tau_*$ if $\epsilon < \sqrt{3}$ and $p_* < \tau_*$ if $\epsilon > \sqrt{3}$. At $\epsilon = \sqrt{3}$, $p_* - \tau_* = \frac{4}{9v_0^2} + \mathcal{O}(v_0^{-3}) > 0$.

Proof. With $x = v_0^{-1}$ and $E = \epsilon^2 x^{-2}$ we have $\tau_*^2 = \frac{Q}{E} = 1 + \frac{3x}{\epsilon^2} - \frac{x^2}{4\epsilon^2}$, and $\sqrt{1+y} = 1 + \frac{y}{2} - \frac{y^2}{8} + \mathcal{O}(y^3)$ gives the first expansion. Then $\tau_*^3 = 1 + \frac{9x}{2\epsilon^2} + (\frac{27}{8\epsilon^4} - \frac{3}{8\epsilon^2})x^2 + \mathcal{O}(x^3)$ and $\tau_*^{-2} = 1 - \frac{3x}{\epsilon^2} + (\frac{9}{\epsilon^4} + \frac{1}{4\epsilon^2})x^2 + \mathcal{O}(x^3)$. Substituting these into $p_* = -\frac{c_1(v_0, \tau_*)}{E}$ from eq. (59), with $v_0 = x^{-1}$, gives power by power

$$p_* = 1 + \frac{21 - \epsilon^2}{12\epsilon^2} x + \frac{63 - \epsilon^2}{24\epsilon^4} x^2 + \mathcal{O}(x^3),$$

and subtracting τ_* gives the second expansion in eq. (62). The sign statements read off the x coefficient. At $\epsilon = \sqrt{3}$ the x^2 coefficient is $\frac{48}{108} = \frac{4}{9}$. □

The shape condition is algebraic at every privacy pair.

Proposition 5 (Algebraic form of the shape condition). *Fix $\epsilon > 0$ and $v > v_c$, so $Q = \epsilon^2 v^2 + 3v - \frac{1}{4} > 0$, and with $E = \epsilon^2 v^2$ let*

$$K = -2\left(16E^2 + 48Ev - 8E - 48v^2 - 4v + 1\right), \quad R = 16Q(2E - 3v),$$
$$P = \epsilon v K, \qquad \Delta = R^2 Q - P^2.$$

Then $p_*(v) - \tau_*(v) = -\frac{\tau_*(v)R + K}{48vQ}$, and $p_* \ge \tau_*$ if and only if $R\sqrt{Q} + P \le 0$, that is, if and only if one of

$$\text{(i) } R \le 0,\ P \le 0; \qquad \text{(ii) } R > 0,\ P < 0,\ \Delta \le 0; \qquad \text{(iii) } R < 0,\ P > 0,\ \Delta \ge 0 \tag{63}$$

holds. The sign of R is that of $2\epsilon^2 v - 3$, and K, R, P, Δ are polynomials in (v, ϵ) with integer coefficients, Δ of degree eight in v.

Proof. Write $c_1(v,\tau) = \frac{N_1(\tau)}{48v\tau^2}$, where N_1 is the brace in eq. (59). Reducing $\tau_*^2 = \frac{Q}{E}$ and $\tau_*^3 = \tau_* \frac{Q}{E}$ gives $N_1(\tau_*) = 32Q(E-3v)\tau_* + K'$ with $K' = 16E^2 - 48EQ + 48vQ + 12Q - 8E - 48v^2 - 16v + 1$, and substituting $Q = E + 3v - \frac{1}{4}$ yields $K' = K$. Since $E\tau_*^2 = Q$, we have $48vQ\frac{c_1(v,\tau_*)}{E} = N_1(\tau_*)$, hence $48vQ\{\frac{c_1(v,\tau_*)}{E} + \tau_*\} = \tau_*\{32Q(E-3v) + 48vQ\} + K = \tau_* R + K$, and $p_* = -\frac{c_1(v,\tau_*)}{E}$ gives the difference formula. As $\tau_* = \frac{\sqrt{Q}}{\epsilon v}$, the condition $p_* \geq \tau_*$ is equivalent to $\tau_* R + K \leq 0$ and, after multiplying by $\epsilon v > 0$, to $R\sqrt{Q} + P \leq 0$. It holds if $R, P \leq 0$ and fails if $R, P \geq 0$ except when $R = P = 0$. In the mixed cases squaring gives $R^2 Q \leq P^2$ for $R > 0 > P$ and $P^2 \leq R^2 Q$ for $R < 0 < P$, which is eq. (63). The identity $R = 16vQ(2\epsilon^2 v - 3)$ gives the sign of R. Both $R^2 Q$ and P^2 have degree ten in v with leading coefficient $256\epsilon^{10}$. Their degree-nine coefficients also agree, and the degree-eight coefficient does not vanish. □

Exact real-root isolation of Δ (Basu et al., 2006) together with eq. (63) produces the column v_{shape} of Table 4. At every budget with a shape root, the feasible set inside (v_c, ∞) is the single interval (v_c, v_{shape}). At $\epsilon = 3$, the shape root is the rational point $v_{\text{shape}} = \frac{1}{6}$.

Proposition 6 (Growth of the constant $C(\epsilon,\delta)$)**.** *Fix $\epsilon > 0$ and let $\delta \downarrow 0$. Then*

$$C(\epsilon,\delta) = \frac{\epsilon^2 v_0^2}{2} - \frac{3v_0}{2} + \frac{\epsilon^2 + 18}{8\epsilon^2} + \mathcal{O}(v_0^{-1}), \qquad \frac{C(\epsilon,\delta)}{M v_0} = \frac{\epsilon^2 v_0}{2M}\left\{1 - \frac{3}{\epsilon^2 v_0} + \mathcal{O}(v_0^{-2})\right\}, \tag{64}$$

and $\epsilon^2 v_0 = 2\log(\frac{1}{\delta})\{1 + \mathcal{O}(1)\}$.

Proof. From eq. (62) with $x = v_0^{-1}$ we have $(2-\tau_*)^2 = 1 - \frac{3x}{\epsilon^2} + \frac{\epsilon^2+18}{4\epsilon^4}x^2 + \mathcal{O}(x^3)$, and $C(\epsilon,\delta) = \frac{E}{2}(2-\tau_*)^2$ gives both displays. For the root, $A(v), B(v) \to -\infty$ with $|B| - |A| = v^{-\frac{1}{2}}$, $|A||B| = \epsilon^2 v - \frac{1}{4v}$, and $\phi(A) = e^{\epsilon}\phi(B)$. The Mills-ratio expansion $\Phi(-y) = \phi(y)\{y^{-1} - y^{-3} + \mathcal{O}(y^{-5})\}$ gives $g(v) = \frac{\phi(A(v))}{\epsilon^2 v^{\frac{3}{2}}}\{1 + \mathcal{O}(v^{-1})\}$. Taking logarithms at v_0 with $A(v_0)^2 = \epsilon^2 v_0 - \epsilon + \frac{1}{4v_0}$ yields $\log\frac{1}{\delta} = \frac{\epsilon^2 v_0}{2} + \frac{3}{2}\log v_0 + \mathcal{O}(1)$. □

Remark 3 (Reading the constant)**.** *By eq. (64) and Corollary 2, the leading coefficients behave like $\frac{2\log^2(\frac{1}{\delta})}{\epsilon^2}$ in absolute and $\frac{\log(\frac{1}{\delta})}{M}$ in relative terms as $\delta \downarrow 0$. Every theorem fixes (ϵ,δ) first and produces K, M_0 depending on that pair. This is a statement about coefficients, not a joint limit in (M,δ). The constant-order conclusion applies once M is large compared with $\epsilon^2 v_0$. For finite (M,ϵ,δ), the numerical bounds of Section 4 apply.*

Proposition 7 (Singular privacy axes)**.** *Fix $0 < \delta < 1$, $z_\delta = \Phi^{-1}(\frac{1+\delta}{2})$, and $v_{0,\delta} = \frac{1}{4z_\delta^2}$. As $\epsilon \downarrow 0$ we have $v_0 \to v_{0,\delta}$, and if $z_\delta < \sqrt{3}$ then $\epsilon\tau_* \to 2z_\delta\sqrt{3 - z_\delta^2} > 0$, so $\tau_* \to \infty$. On the other axis, for every $\epsilon < \infty$ and $v < \infty$, $\mathcal{N}(0, v\mathbf{I}_M)$ fails $(\epsilon, 0)$-DP under any nonzero shift.*

Proof. At $\epsilon = 0$, the profile $g(v) = 2\Phi(\frac{1}{2\sqrt{v}}) - 1$ is strictly decreasing, so $g(v) = \delta$ has the unique root $v_{0,\delta}$. Since $(\epsilon, v) \mapsto g_\epsilon(v)$ is continuous on $[0,\epsilon_0] \times [v_-, v_+]$ and $g_0(v_-) > \delta > g_0(v_+)$ for $v_- < v_{0,\delta} < v_+$, the same holds for small ϵ, so $v_0(\epsilon) \in (v_-, v_+)$ and every limit point solves $g_0 = \delta$. Hence $v_0(\epsilon) \to v_{0,\delta}$. Then $\epsilon\tau_* = \frac{\sqrt{\epsilon^2 v_0^2 + 3v_0 - \frac{1}{4}}}{v_0} \to \frac{\sqrt{3v_{0,\delta} - \frac{1}{4}}}{v_{0,\delta}} = 2z_\delta\sqrt{3 - z_\delta^2}$, positive when $z_\delta < \sqrt{3}$. For the pure axis, the Gaussian log-likelihood ratio at shift $h \neq 0$ is $\frac{\|h\|_2^2}{2\Delta_2^2 v} - \frac{z \cdot h}{\Delta_2^2 v}$, unbounded along $z = -th$ as $t \to \infty$, whereas $(\epsilon, 0)$-DP requires it to be at most ϵ almost everywhere. □

Remark 4 (Scope of the matching window)**.** *Table 4 evaluates the two boundaries on a rational budget grid. The patterns below describe the sampled rows.*

1. *At $\epsilon \in \{\frac{1}{4}, \frac{1}{2}, 1\}$ exact root isolation finds no shape root in (v_c, ∞), so Assumption 2 is never active and Corollary 1 applies for every $\delta < \delta_{\tan}$. At $\epsilon = 1$, $\delta_{\tan} \approx 0.65$.*
2. *From $\frac{7}{4}$ through 3, the window is nonempty but closes fast, with δ_{shape} rising from 7.1×10^{-171} to 0.3563 across those four budgets. At $\epsilon = \frac{5}{2}$, the window is $v_0 \in (0.141895, 1.689667)$, so $(\epsilon,\delta) = (\frac{5}{2}, 10^{-5})$ lies outside it. At $\epsilon = 3$ it is $v_0 \in (v_{\tan}, \frac{1}{6})$ with $v_{\tan} \approx 0.13254$.*

3. *At $\frac{13}{4}$, $\frac{7}{2}$, and 4, the window is empty because $v_{\mathrm{shape}} < v_{\tan}$. The shape boundary is algebraic, and the tangent boundary is transcendental. We verify only that the window closes between $\epsilon = 3$ and $\frac{13}{4}$ and report the crossing near 3.18 as a diagnostic.*

In the last two patterns, Theorems 2 and 4 remain valid, but exact $\frac{1}{M}$ matching is open. At $\delta = 10^{-5}$, the shape condition fails from $\epsilon = \frac{5}{2}$ onward on this grid. There B_{sh} is 4.4×10^{-4}, 5.1×10^{-3}, 2.4×10^{-2}, 8.7×10^{-2}, and 1.77×10^{-1} at $\epsilon = \frac{5}{2}, 3, 4, 6, 8$, against $C(\epsilon, \delta) = 18.7, 14.3, 9.4, 5.4, 3.8$ and $J(\epsilon, \delta) = 81.8, 66.0, 48.1, 32.3, 25.4$ at the same budgets. The undetermined share $\frac{B_{\mathrm{sh}}}{J(\epsilon,\delta)}$ of the $\frac{1}{M}$ coefficient therefore stays below 0.7% over that range. Assumption 1 holds throughout it. The margin $v_0 - v_c$ falls from 2.60 to 0.317 and $\delta - v_0 g'(v_0)$ stays between 1.1×10^{-4} and 1.4×10^{-4} against the bound 1. At budgets typical of private deep learning, $\epsilon \leq 8$ and $\delta = 10^{-5}$, the constant order is sharp. Proposition 4 explains the transition through the sign of $3 - \epsilon^2$ as $\delta \downarrow 0$. By Proposition 6, the constant $C(\epsilon, \delta)$ grows without bound as $\delta \downarrow 0$, so reading eq. (7) coefficient by coefficient requires M large compared with $\epsilon^2 v_0$ (Remark 3). By Proposition 7, both privacy axes are singular, so the interior expansion does not extend to the scalar endpoint optimizers. The four privacy pairs of Section 4 have $\epsilon \leq 1$ and sit in the first pattern. Closing the gap for large ϵ and small δ requires a mechanism outside the monotone-density range (Appendix I).

ϵ	v_c	$v_{\tan}$	v_{shape}	$\delta_{\tan}$	δ_{shape}	matching window
1/4	0.083189	0.175853	–	0.7367	–	nonempty
1/2	0.082763	0.174354	–	0.7070	–	nonempty
1	0.081139	0.168789	–	0.6492	–	nonempty
7/4	0.077243	0.156227	249.689228	0.5718	7.1×10^{-171}	nonempty
2	0.075694	0.151499	11.112624	0.5491	1.5×10^{-12}	nonempty
5/2	0.072410	0.141895	1.689667	0.5084	3.9×10^{-4}	nonempty
3	0.069036	0.132543	0.166667	0.4733	0.3563	nonempty
13/4	0.067359	0.128066	0.123275	0.4577	0.4784	empty
7/2	0.065705	0.123750	0.110681	0.4431	0.5059	empty
4	0.062500	0.115643	0.097276	0.4170	0.5198	empty

Table 4: Joint scope of Assumptions 1 and 2 on a rational budget grid. v_{shape} is obtained by exact real-root isolation of Δ in Proposition 5 with the sign conditions eq. (63). A dash means that the shape condition holds on all of (v_c, ∞). $v_{\tan}$ is an interval-arithmetic bound on the tangent boundary. $\delta_{\tan} = g(v_{\tan})$ coincides with the threshold $\delta^\star$ of Wei et al. (2026) (their Table 2) at the budgets they report, and $\delta_{\mathrm{shape}} = g(v_{\mathrm{shape}})$ has no counterpart there. The window $\delta \in (\delta_{\mathrm{shape}}, \delta_{\tan})$ is nonempty exactly when $v_{\mathrm{shape}} > v_{\tan}$.

Figure 2 classifies a grid in the privacy plane by solving, for each ϵ, the scalar equations of the tangent and shape boundaries and evaluating g. The solid orange curve is $\delta - v_0 g'(v_0) = 1$. The dashed orange curve $v_0 = v_c$ lies above it at every ϵ shown, so the condition $v_0 > v_c$ never restricts the region. The shape boundary has the dotted line $\epsilon = \sqrt{3}$ as asymptote (Proposition 4) and leaves the frame at the $\delta = 10^{-15}$ floor near $\epsilon = 1.96$. Below the blue curve, Theorem 4 still holds, but exact $\frac{1}{M}$ matching is open. A missing shape boundary means that $p_* \geq \tau_*$ persists as $\delta \downarrow 0$, consistent with Proposition 4.

E.6 Computing and sampling the CGG mechanism

This section supplies the details behind Algorithm 1.

Step 1. The profile g is strictly decreasing from 1 to 0 (Appendix B.1), so v_0 is found by bisection. The constants τ_*, p_*, ρ_* are then the closed forms eq. (4) with c_1 from eq. (59), where τ_* is real under Assumption 1 and $\rho_* \leq 1$ under Assumption 2. None of these numbers depends on M or Δ_2. Table 5 lists them at the privacy pairs of Section 4.

Step 3. The scale a enters eq. (3) as a pure scale, so $a \mapsto D_\epsilon(P_a \| (P_a)_{\Delta_2 e_1})$ is the map eq. (42) with $s = a$. Since $q \leq M$, Corollary 3 applies. The profile is continuous and nonincreasing in a with limits 1 and 0, so the level δ is attained, and the feasible scales form a closed upper interval.

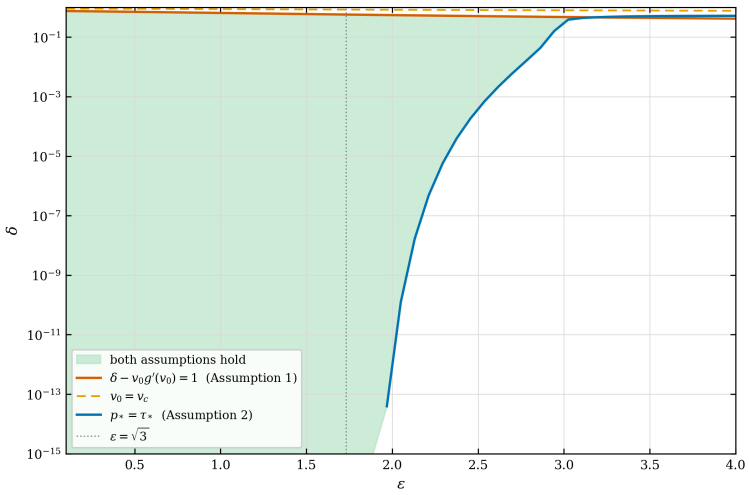


Figure 2: Scope of Assumptions 1 and 2 over $0.1 \le \epsilon \le 4$ and $10^{-15} \le \delta < 1$, with δ on a logarithmic scale. The solid orange curve is the tangent boundary of Assumption 1 and the blue curve is the shape boundary $p_* = \tau_*$ of Assumption 2. Both assumptions hold in the shaded region, which closes where the two curves cross near $\epsilon = 3.18$.

Its left endpoint a_M is the smallest feasible scale, and we can widen the bracket until it contains it. Each bisection iteration evaluates the exact profile eq. (43) once, a one-dimensional radial integral with H_M evaluated at the two radial thresholds. A few dozen iterations give double precision. By Lemma 10, privacy at $\Delta_2 e_1$ then holds for every $\|h\|_2 \le \Delta_2$. Since the calibration commutes with scaling, a_M at sensitivity Δ_2 is Δ_2 times the scale calibrated at $\Delta_2 = 1$.

Step 4. By Appendix E.1, $T = \left(\frac{R}{a}\right)^p \sim \mathrm{Gamma}(\frac{q}{p}, 1)$, so $R = aT^{\frac{1}{p}}$ is an exact draw from eq. (3), and $U = \frac{N}{\|N\|_2}$ is uniform and independent of R. No rejection step is involved. From eq. (11) with $s = 2$,

$$\mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2) = \mathbb{E}[R_M^2] = a_M^2 \frac{\Gamma\left(\frac{\rho_* M + 2}{p_*}\right)}{\Gamma\left(\frac{\rho_* M}{p_*}\right)},$$

the quantity reported in Tables 1 and 2.

Cost. Step 3 is the only M-dependent step and is one calibration, the operation the AGM performs for σ. A search over the shape parameters repeats it for every candidate visited. Section 4.2 reports the counts.

E.7 Proof of Theorem 3

Proof. Assumption 2 gives $p_* \ge \tau_* > 0$, so $\rho_* \in (0, 1]$ and $q = \rho_* M \le M$. The exponent $(\rho_* - 1)M \le 0$ of $\|z\|_2$ in eq. (12) makes the density nonincreasing for every M. By Corollary 3 a smallest scale a_M with $D_\epsilon(P \| P_{\Delta_2 e_1}) \le \delta$ exists, and Lemma 10 extends this to every $\|h\|_2 \le \Delta_2$, so Z_M^{CGG} is feasible for eq. (1). The pair (ρ_*, p_*) lies in a compact shape set with $\rho_* \le 1$, and a_M is the calibrated scale of Proposition 3, so eq. (61) applies with $C_{\tau_*}(v_0) = C(\epsilon, \delta)$ and $J_{\tau_*, p_*}(v_0) = J(\epsilon, \delta)$ by eq. (16). This gives

$$\mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2) \le M\Delta_2^2 v_0 - \Delta_2^2 C(\epsilon, \delta) + \frac{\Delta_2^2 J(\epsilon, \delta)}{M} + \frac{\Delta_2^2 K_{\mathrm{ach}}}{M^2} = A_M(\Delta_2) + \frac{\Delta_2^2 K_{\mathrm{ach}}}{M^2}$$

for $M \ge M_{\mathrm{ach}}$, which is eq. (6). Both constants depend only on (ϵ, δ), since $\rho_*, p_*, \tau_*, C(\epsilon, \delta), J(\epsilon, \delta)$ do and the remainder in Proposition 3 is uniform on a compact shape set containing (ρ_*, p_*). □

F MATCHING THE CONVERSE AND ACHIEVABILITY BOUNDS

This section combines Theorem 2 (Appendix D.6) with Proposition 3 (Appendix E.4.1). The identities eq. (16) express both bounds in terms of $A_M(\Delta_2)$. Recall from Appendix B.4 the clipped power $p_+ = \max\{p_*, \tau_*\}$ and $B_{\mathrm{sh}} = \frac{E}{2}(p_+ - p_*)^2$. Let $\rho_+ = \frac{\tau_*}{p_+} \le 1$ and let $\mathrm{MSE}_M^+(\Delta_2)$ be the exactly calibrated error of generalized gamma noise with $(q, p) = (\rho_+ M, p_+)$. When $p_* < \tau_*$ this is the power-radial distribution $(q, p) = (M, \tau_*)$.

F.1 PROOF OF THEOREM 4

Proof. The lower bound follows from Theorem 2, which uses only Assumption 1. For the upper bound, $p_+ \ge \tau_* > 0$ gives $0 < \rho_+ \le 1$ and $\rho_+ p_+ = \tau_*$, so (ρ_+, p_+) is admissible in Proposition 3. Writing its remainder as $\frac{\Delta_2^2 K_{\mathrm{U}}}{M^2}$ for $M \ge M_{\mathrm{U}}$ gives

$$\mathrm{MSE}_M^+(\Delta_2) \le M\Delta_2^2 v_0 - \Delta_2^2 C_{\tau_*}(v_0) + \frac{\Delta_2^2 J_{\tau_*,p_+}(v_0)}{M} + \frac{\Delta_2^2 K_{\mathrm{U}}}{M^2} = A_M(\Delta_2) + \frac{\Delta_2^2 B_{\mathrm{sh}}}{M} + \frac{\Delta_2^2 K_{\mathrm{U}}}{M^2},$$

by $C_{\tau_*}(v_0) = C(\epsilon, \delta)$ and $J_{\tau_*,p_+}(v_0) = J(\epsilon, \delta) + \frac{E}{2}(p_+ - p_*)^2$ from eq. (16). Since this mechanism is feasible, $\mathcal{V}_M(\Delta_2) \le \mathrm{MSE}_M^+(\Delta_2)$, which is eq. (9) with $K = \max\{K_{\mathrm{cv}}, K_{\mathrm{U}}\}$ and $M_0 = \max\{M_{\mathrm{cv}}, M_{\mathrm{U}}\}$. Subtracting $M\Delta_2^2 v_0 - \Delta_2^2 C(\epsilon, \delta)$ from eq. (9) and multiplying by $\frac{M}{\Delta_2^2}$ places the sequence $\frac{M}{\Delta_2^2}\{\mathcal{V}_M(\Delta_2) - M\Delta_2^2 v_0 + \Delta_2^2 C(\epsilon, \delta)\}$ in $[J(\epsilon, \delta) - \frac{K}{M}, J(\epsilon, \delta) + B_{\mathrm{sh}} + \frac{K}{M}]$, so every limit point lies in $[J(\epsilon, \delta), J(\epsilon, \delta) + B_{\mathrm{sh}}]$. Since $B_{\mathrm{sh}} \ge 0$ with equality exactly when $p_* \ge \tau_*$, the constant order is sharp wherever Assumption 1 holds. □

F.2 PROOF OF COROLLARY 1

Proof. Theorem 3 gives $\mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2) \le A_M(\Delta_2) + \frac{\Delta_2^2 K_{\mathrm{ach}}}{M^2}$ for $M \ge M_{\mathrm{ach}}$. Theorem 2 gives $\mathbb{E}\|Z\|_2^2 \ge A_M(\Delta_2) - \frac{\Delta_2^2 K_{\mathrm{cv}}}{M^2}$ for every feasible Z and $M \ge M_{\mathrm{cv}}$, hence for the infimum $\mathcal{V}_M(\Delta_2)$. Since Z_M^{CGG} is feasible, $\mathcal{V}_M(\Delta_2) \le \mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2)$. With $K = \max\{K_{\mathrm{cv}}, K_{\mathrm{ach}}\}$ and $M_0 = \max\{M_{\mathrm{cv}}, M_{\mathrm{ach}}\}$ this is eq. (7). The outer bounds give $\mathcal{V}_M(\Delta_2) = A_M(\Delta_2) + \mathcal{O}(\Delta_2^2 M^{-2})$. Subtracting the lower bound for $\mathcal{V}_M(\Delta_2)$ from the upper bound for $\mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2)$ gives $0 \le \mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2) - \mathcal{V}_M(\Delta_2) \le \frac{2\Delta_2^2 K}{M^2}$. The identity eq. (31) behind the agreement of the two $\frac{1}{M}$ coefficients is proved in Appendix D.5. □

F.3 PROOF OF COROLLARY 2

Proof. Write the outer bounds of eq. (7) as $\mathcal{V}_M(\Delta_2) = M\Delta_2^2 v_0 - \Delta_2^2 C(\epsilon, \delta) + \frac{\Delta_2^2 J(\epsilon,\delta)}{M} + \Delta_2^2 r_M$ with $|r_M| \le K M^{-2}$. Subtracting from $\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) = M\Delta_2^2 v_0$ gives

$$\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) - \mathcal{V}_M(\Delta_2) = \Delta_2^2 C(\epsilon, \delta) - \frac{\Delta_2^2 J(\epsilon, \delta)}{M} + \mathcal{O}(\Delta_2^2 M^{-2}), \tag{65}$$

and dividing by $M\Delta_2^2 v_0$ with v_0 fixed gives eq. (8).

For the second statement we use only the converse. Every feasible P has $\mathbb{E}_P\|Z\|_2^2 \ge \mathcal{V}_M(\Delta_2) \ge A_M(\Delta_2) - \frac{\Delta_2^2 K_{\mathrm{cv}}}{M^2}$ by Theorem 2, so

$$\frac{\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) - \mathbb{E}_P\|Z\|_2^2}{\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2)} \le \frac{C(\epsilon, \delta)}{M v_0} - \frac{J(\epsilon, \delta)}{M^2 v_0} + \frac{K_{\mathrm{cv}}}{M^3 v_0} \le \frac{C(\epsilon, \delta)}{M v_0} + \frac{K}{M^2 v_0}$$

with $K = |J(\epsilon, \delta)| + K_{\mathrm{cv}}$ and $M \ge \max\{M_{\mathrm{cv}}, 1\}$. □

Corollary 5 (Nondegenerate and degenerate Gaussian excess)**.** *Since $C(\epsilon, \delta) = \frac{E}{2}(2 - \tau_*)^2$, it is nonnegative and positive exactly when $\tau_* \ne 2$. Under Assumptions 1 and 2, if $\tau_* \ne 2$ then*

$$\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) - \mathcal{V}_M(\Delta_2) \sim \Delta_2^2 C(\epsilon, \delta), \qquad \frac{\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) - \mathcal{V}_M(\Delta_2)}{\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2)} \sim \frac{C(\epsilon, \delta)}{M v_0}. \tag{66}$$

If $\tau_* = 2$ then $C(\epsilon,\delta) = 0$, $J(\epsilon,\delta) = -\frac{E}{2}(2-p_*)^2 \le 0$, and

$$\begin{aligned}\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) - \mathcal{V}_M(\Delta_2) &= \frac{\Delta_2^2 E(2-p_*)^2}{2M} + \mathcal{O}(\Delta_2^2 M^{-2}),\\ \frac{\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2) - \mathcal{V}_M(\Delta_2)}{\mathrm{MSE}_M^{\mathrm{AGM}}(\Delta_2)} &= \frac{E(2-p_*)^2}{2v_0 M^2} + \mathcal{O}(M^{-3}).\end{aligned} \tag{67}$$

If also $p_* = 2$, both leading coefficients vanish and only the remainders $\mathcal{O}(\Delta_2^2 M^{-2})$ and $\mathcal{O}(M^{-3})$ remain.

Proof. We have $C(\epsilon,\delta) = \frac{E}{2}(2-\tau_*)^2 \ge 0$ with equality exactly at $\tau_* = 2$. If $\tau_* \ne 2$, eq. (65) and eq. (8) give ratios $1 + \mathcal{O}(M^{-1})$ to the stated leading terms, which is eq. (66). If $\tau_* = 2$, the member $(\rho, p) = (1, 2)$, for which $\tau = 2$, is Gaussian with exact error $M\Delta_2^2 v_0$ for every M, so $J_{2,2}(v_0) = 0$, and eq. (16) at $p = 2$ gives $J(\epsilon,\delta) = -\frac{E}{2}(2-p_*)^2$. Substitution into eq. (65) and eq. (8) gives eq. (67). □

Corollary 6 (Three approximation levels)**.** *Under Assumptions 1 and 2, let $\mathrm{MSE}_M^{\ell_2}(\Delta_2)$ and $\mathrm{MSE}_M^{\mathrm{PR}}(\Delta_2)$ be the exactly calibrated errors of the ℓ_2 mechanism $(q,p) = (M,1)$ and of power-radial noise $(q,p) = (M,\tau_*)$, and recall $B_* = \frac{E}{2}(p_* - \tau_*)^2$ from Appendix B.4. Then*

$$\begin{aligned}&\mathrm{MSE}_M^{\ell_2}(\Delta_2) - \mathcal{V}_M(\Delta_2) = \Delta_2^2 E(\tau_* - 1)^2 + \mathcal{O}(\Delta_2^2 M^{-1}),\\ &\mathrm{MSE}_M^{\mathrm{PR}}(\Delta_2) - \mathcal{V}_M(\Delta_2) = \frac{\Delta_2^2 B_*}{M} + \mathcal{O}(\Delta_2^2 M^{-2}), \qquad \mathrm{MSE}_M^{\mathrm{CGG}}(\Delta_2) - \mathcal{V}_M(\Delta_2) = \mathcal{O}(\Delta_2^2 M^{-2}).\end{aligned} \tag{68}$$

The Gaussian mechanism matches the $\mathcal{O}(M)$ term. The ℓ_2 mechanism misses the constant. Power-radial noise with $\tau = \tau_*$ matches the constant but misses the $\frac{1}{M}$ coefficient by B_*. The CGG mechanism matches all three orders. If $p_* = \tau_*$ then $\rho_* = 1$, the CGG and power-radial distributions coincide, and $B_* = 0$.

Proof. Proposition 3 at $(\rho, p) = (1, 1)$, so $\tau = 1$, and subtraction of eq. (7) give $\Delta_2^2\{C(\epsilon,\delta) - C_1(v_0)\} + \mathcal{O}(\Delta_2^2 M^{-1})$, and the first identity in eq. (16) at $\tau = 1$ evaluates the brace to $E(\tau_* - 1)^2$. For power-radial noise, $\rho = 1$ and $p = \tau_*$ give $\tau = \tau_*$, so the constant term cancels and the second identity in eq. (16) at $p = \tau_*$ leaves $\frac{\Delta_2^2}{M}\{J_{\tau_*,\tau_*}(v_0) - J(\epsilon,\delta)\} = \frac{\Delta_2^2 B_*}{M}$. The last line is Corollary 1. □

G Reading the numerical bounds

The bounds in Section 4 use $\Delta_2 = 1$. For a different sensitivity, multiply every MSE endpoint and absolute gap by Δ_2^2 and leave relative quantities unchanged. The bounds at one (ϵ, δ, M) consist of a feasible mechanism with exact error $\mathrm{MSE}_M^{\mathrm{cand}}$, an upper bound $\overline{\mathrm{MSE}}_M \ge \mathrm{MSE}_M^{\mathrm{cand}}$, and a dual certificate $(\theta, \lambda, \kappa)$ whose support inequality $v + \lambda h_M[\theta](v) \ge \kappa$ holds for all $v \ge 0$, both checked by interval arithmetic, giving $\underline{\mathrm{MSE}}_M = M(\kappa - \lambda\delta)$ by Lemma 5. Hence

$$\underline{\mathrm{MSE}}_M \le \mathcal{V}_M(1) \le \mathrm{MSE}_M^{\mathrm{cand}} \le \overline{\mathrm{MSE}}_M. \tag{69}$$

Lemma 5 holds at every $M \ge 2$, so the bounds hold at their own dimension. Theorem 2 separately states that a threshold with a prescribed expansion exists once $M \ge M_{\mathrm{cv}}$. The thresholds of Appendix H.5 use the expansion of Lemma 9 as a starting point, then solve for the knot values numerically at each M, maximizing the dual level subject to the support inequality. The interval check covers the resulting spline. Verifying the unrefined expansion at $M = 10^3$ would give a converse three digits weaker. Writing $\mathrm{MSE}_M^{\mathrm{cand}} = M\mathbb{E}[V]$, the unrounded gap decomposes as $\mathrm{MSE}_M^{\mathrm{cand}} - \underline{\mathrm{MSE}}_M = M\mathbb{E}[V + \lambda h_M[\theta](V) - \kappa] + M\lambda\{\delta - \mathbb{E}[h_M[\theta](V)]\}$, the slack in the support inequality plus the privacy slack, both nonnegative. The reported intervals contain all rounding and quadrature errors. Appendix H.5 records the data behind each entry and the recheck procedure.

Proposition 8 (Interpreting a certificate)**.** *Suppose $\mathrm{MSE}_M^{\mathrm{AGM}}(1) > \overline{\mathrm{MSE}}_M \ge \underline{\mathrm{MSE}}_M > 0$. Then*

$$\begin{aligned}\mathrm{MSE}_M^{\mathrm{AGM}} - \overline{\mathrm{MSE}}_M \le \mathrm{MSE}_M^{\mathrm{AGM}} - \mathcal{V}_M(1) \le \mathrm{MSE}_M^{\mathrm{AGM}} - \underline{\mathrm{MSE}}_M,\\ 0 \le \mathrm{MSE}_M^{\mathrm{cand}} - \mathcal{V}_M(1) \le \overline{\mathrm{MSE}}_M - \underline{\mathrm{MSE}}_M,\end{aligned}$$

and

$$\frac{\mathrm{MSE}_M^{\mathrm{AGM}} - \overline{\mathrm{MSE}}_M}{\mathrm{MSE}_M^{\mathrm{AGM}} - \underline{\mathrm{MSE}}_M} \le \frac{\mathrm{MSE}_M^{\mathrm{AGM}} - \mathrm{MSE}_M^{\mathrm{cand}}}{\mathrm{MSE}_M^{\mathrm{AGM}} - \mathcal{V}_M(1)}. \tag{70}$$

Proof. Subtracting eq. (69) from $\mathrm{MSE}_M^{\mathrm{AGM}}$ reverses the order and gives the first chain. The second follows from $\mathcal{V}_M(1) \le \mathrm{MSE}_M^{\mathrm{cand}} \le \overline{\mathrm{MSE}}_M$ and $\mathcal{V}_M(1) \ge \underline{\mathrm{MSE}}_M$. For the third, the numerator is at least $\mathrm{MSE}_M^{\mathrm{AGM}} - \overline{\mathrm{MSE}}_M > 0$ and the denominator lies in $(0, \mathrm{MSE}_M^{\mathrm{AGM}} - \underline{\mathrm{MSE}}_M]$. □

The three quantities are the Gaussian excess, the candidate's gap to the optimum, and the fraction of the excess it recovers. Section 4.1 reads them off the endpoints of Table 1, where the candidate is the CGG mechanism. Table 5 lists v_0 and the constants entering those bounds at the four privacy pairs. At $\Delta_2 = 1$ we have $\mathrm{MSE}_M^{\mathrm{AGM}} = M v_0$.

ϵ	δ	v_0	τ_*	p_*	ρ_*	$C(\epsilon,\delta)$	$J(\epsilon,\delta)$
0.5	10^{-5}	49.446586	1.114575	1.151758	0.967716	239.599740	687.726863
	10^{-3}	21.253280	1.249962	1.367634	0.913959	31.763512	12.123121
1.0	10^{-5}	13.917612	1.101936	1.129162	0.975888	78.111286	265.679925
	10^{-3}	6.628859	1.202862	1.280803	0.939147	13.960933	12.713967

Table 5: Constants of eq. (7) at the four privacy pairs of the experiments, rounded to six decimals. Assumption 1 holds with $v_0 - v_c$ at least 6.54 and $\delta - v_0 g'(v_0)$ at most 5.6×10^{-3}. Assumption 2 holds with $p_* - \tau_*$ at least 0.027, so $\rho_* \le 0.976$.

Those constants carry a matching guarantee only where both assumptions hold. Appendix E.5 locates the two boundaries that mark out that region, and Table 4, Figure 2, and Remark 4 map the window and record where each pattern holds.

H Additional numerical results

This section collects the figures behind Section 4 and six further items. These are the rates of the numerical bounds, the convergence and cost of the shape search, a direct numerical test of the constants in eq. (7) and eq. (68), bounds on $\mathcal{V}_M(1)$ at $M \in \{10, 10^2\}$, the data behind every entry of Table 1, and the calibrated error as a function of the shape parameters. All settings use $\Delta_2 = 1$ and the four privacy pairs of Table 5.

H.1 Rates of the numerical bounds

Figure 3 plots the sixteen intervals of Table 1 against M. The left panel gives the upper bound on the Gaussian excess and the error the CGG mechanism attains. The band between them is the saving any additive noise could still add, and it is narrower than the line width for $M \ge 10^2$. The ceiling falls by a factor between 9.7 and 10.1 over each of the last two decades, the $\frac{C(\epsilon,\delta)}{M v_0}$ rate of Corollary 2. Over the first decade, the factor is 7.7 to 10.1, with the $\mathcal{O}(M^{-2})$ remainder still contributing at $M = 10$. The right panel bounds the fraction of the excess the CGG mechanism leaves, at most 7.5% at $M = 10$ and at most $2.9 \times 10^{-5}\%$ at $M = 10^4$. Entries at $M \in \{10, 10^2\}$ use the finite-family converse of Appendix H.4 and are therefore weaker than Theorem 2.

H.2 Convergence and cost of the search

The CGG and SGG errors of Table 2 overlap when plotted against M, so Figure 4 plots their difference. It falls at rate M^{-2} for every privacy pair, and the difference scaled by M^2 remains bounded. Figure 5 shows the reason. The searched shape approaches the closed-form point (ρ_*, p_*) at rate M^{-1} in the (ρ, p) plane at three of the four privacy pairs. The exception is $(\epsilon, \delta) = (0.5, 10^{-5})$, where the $M = 10^2$ search stopped at the boundary $\rho = 1$, so its distance is not a minimizer's (Table 6). An $\mathcal{O}(M^{-1})$ displacement costs $\mathcal{O}(M^{-2})$ in the objective. By eq. (16), the constant coefficient is stationary in τ at τ_* and changes at second order, whereas the $\frac{1}{M}$ coefficient changes

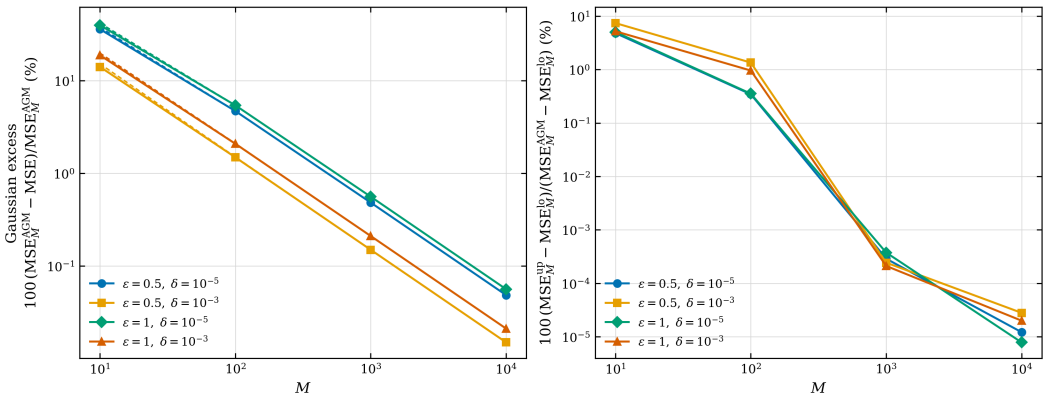


Figure 3: Left: Gaussian excess $\frac{100(\mathrm{MSE}_M^{\mathrm{AGM}} - \mathrm{MSE})}{\mathrm{MSE}_M^{\mathrm{AGM}}}$ against M, solid for the CGG mechanism and dashed for the numerical converse bound. Right: upper bound $\frac{100(\overline{\mathrm{MSE}}_M - \underline{\mathrm{MSE}}_M)}{\mathrm{MSE}_M^{\mathrm{AGM}} - \underline{\mathrm{MSE}}_M}$ on the fraction of the excess the CGG mechanism leaves, written $\mathrm{MSE}_M^{\mathrm{up}}$ and $\mathrm{MSE}_M^{\mathrm{lo}}$ on the axis. Entries at $M \in \{10^3, 10^4\}$ use the dual event of Lemma 5 and those at $M \in \{10, 10^2\}$ the finite-family converse of Appendix H.4. All sixteen are in Table 1.

at first order in τ but is already divided by M. Figure 6 plots the two halves of Table 2 on the same axes, with the setup times of Table 2. On this machine, the AGM is the cheapest at 0.04 to 0.14 milliseconds, since its profile is an explicit function of Φ and rooting it needs no quadrature. The ℓ_2 and CGG mechanisms both bisect against a radial quadrature and cost 5.7 to 14.1 milliseconds. The search costs 4.9 to 8.8 seconds, 523 to 838 times the CGG calibration, and recovers the $\mathcal{O}(M^{-2})$ amount above. The comparison separates a calibration at a closed-form shape from a search over shapes.

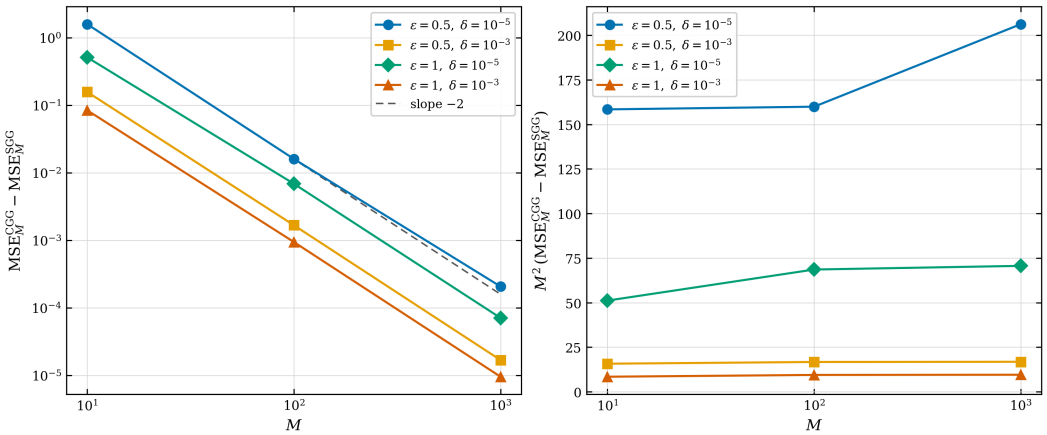


Figure 4: Left: $\mathrm{MSE}_M^{\mathrm{CGG}} - \mathrm{MSE}_M^{\mathrm{SGG}}$ against M on log–log axes with a reference line of slope -2. Right: the same difference scaled by M^2, which Corollary 1 bounds. Both panels come from the search run recorded in Table 6.

Budget, output, and precision of the search. We search the family with the density, parameterization, and box of Wei et al. (2026). Their outer bisection on an error budget and inner Bayesian optimization with early stopping return a feasible shape, whereas Table 6 needs the minimizer of the calibrated error to the fifth decimal. A grid and bounded simplex on the calibrated error itself reach that decimal, cost exactly the calibrations they perform, and run on the same exact profile

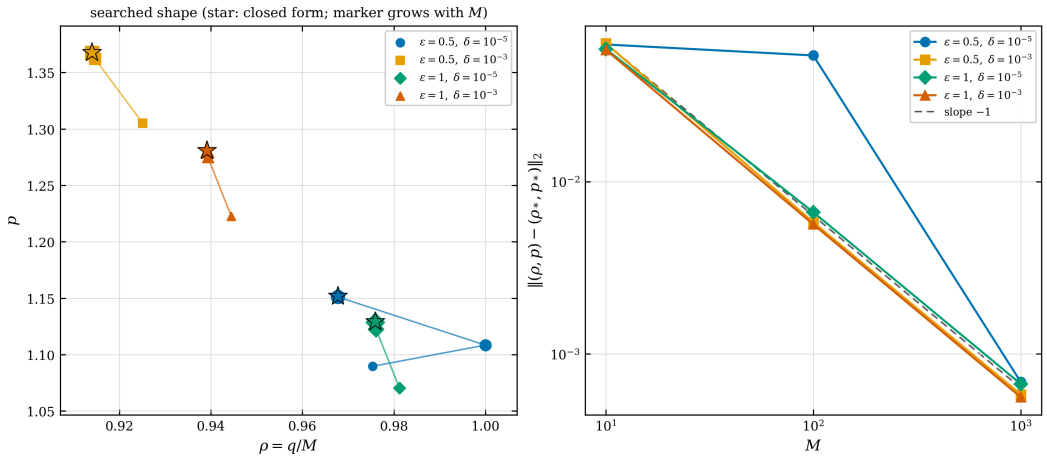


Figure 5: Left: searched shape $(\rho, p) = (\frac{q}{M}, p)$ at $M = 10, 10^2, 10^3$ for each privacy pair, the marker growing with M, and the closed-form point (ρ_*, p_*) (star). The axes show the corner of the search box $0 < \rho \leq 1$, $1 \leq p \leq 3$ of Wei et al. (2026) that contains every returned shape. Right: Euclidean distance from the searched shape to (ρ_*, p_*) against M with a reference line of slope -1.

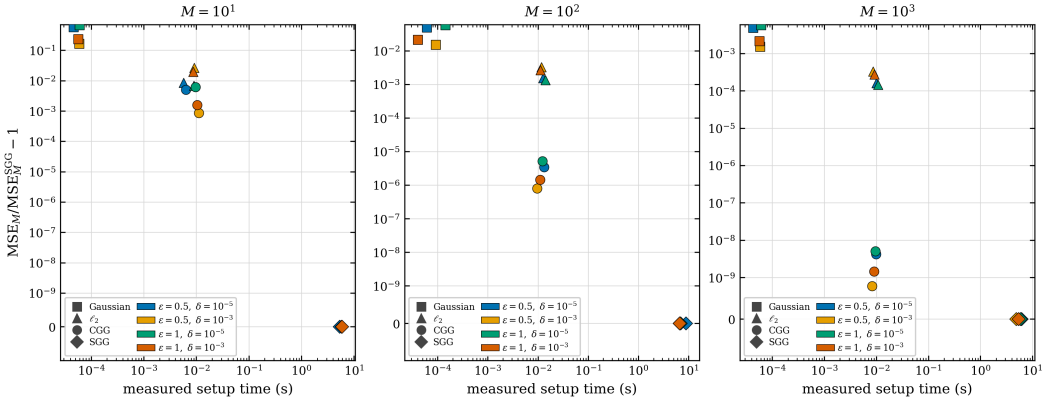


Figure 6: Relative error $\frac{\mathrm{MSE}_M}{\mathrm{MSE}_M^{\mathrm{SGG}}} - 1$ against measured setup time, one panel per M. SGG sits at 0 by construction. Marker shape indicates the mechanism and color indicates the privacy pair. For AGM, ℓ_2, and CGG each time is one scale calibration, best of five repeats started from scratch. For SGG it is one search run. The analytic Gaussian mechanism roots a closed-form profile and sits at the left edge. The ℓ_2 and CGG mechanisms root a radial quadrature and cost 5.7 to 14.1 milliseconds. The search costs seconds.

routine that calibrates the CGG mechanism. The implementation keeps the closed-form shape as a fallback candidate and returns it whenever the search ends above it, so the reported SGG error satisfies $\mathrm{MSE}_M^{\mathrm{SGG}} \leq \mathrm{MSE}_M^{\mathrm{CGG}}$ by construction. The fallback was not taken in any of the twelve cells, so every SGG entry of Table 2 and Table 6 is a search output.

The calibrated error is flat to about 10^{-9} along the ridge $\rho p = \tau_*$, so a direction-set method spends its budget resampling that ridge. A bounded simplex refinement reaches the fifth decimal in 90 to 200 calibrations. With the 20×20 grid that precedes it, the search performs 489 to 604 calibrations per cell, the range reported in Table 2. Table 2, Table 6, and Figures 4 to 6 come from a single run of this procedure. The reported quantity is the difference between the two mechanism columns. Recomputing it on the two returned shapes at a doubled quadrature order changes it by at most a relative 2.0×10^{-7}, which at $M = 10^3$ is 4×10^{-11}, five orders below the smallest difference in

Table 6. The size of the saving and its M^{-2} decay are stable across runs. The constant in front of M^{-2} is not. An earlier run of an equivalent search stopped in the interior rather than at the boundary $\rho = 1$ at $(\epsilon, \delta, M) = (0.5, 10^{-5}, 10^2)$ and reported a saving of 2.013×10^{-2} there instead of 1.599×10^{-2}.

ϵ	δ	M	ρ	p	Calibrations	$\mathrm{MSE}_M^{\mathrm{CGG}} - \mathrm{MSE}_M^{\mathrm{SGG}}$
0.5	10^{-5}	10^1	0.975358	1.089631	493	1.584×10^{0}
		10^2	1.000000	1.108389	604	1.599×10^{-2}
		10^3	0.967760	1.151072	522	2.062×10^{-4}
	10^{-3}	10^1	0.925080	1.305100	496	1.571×10^{-1}
		10^2	0.914669	1.361863	510	1.668×10^{-3}
		10^3	0.914030	1.367056	515	1.679×10^{-5}
1.0	10^{-5}	10^1	0.981216	1.070439	495	5.113×10^{-1}
		10^2	0.976108	1.122513	518	6.856×10^{-3}
		10^3	0.975905	1.128493	511	7.062×10^{-5}
	10^{-3}	10^1	0.944436	1.222674	489	8.404×10^{-2}
		10^2	0.939325	1.275151	503	9.409×10^{-4}
		10^3	0.939160	1.280241	523	9.512×10^{-6}

Table 6: The search run behind Table 2, giving the returned shape, the number of scale calibrations spent, and the resulting difference from the CGG mechanism. Compare (ρ_*, p_*) in Table 5. At $(\epsilon, \delta, M) = (0.5, 10^{-5}, 10^2)$, the simplex stopped at the boundary $\rho = 1$. The remaining cells stopped in the interior. Full working precision for every entry is in the supplementary material.

Warm start and a common timing basis. Table 7 starts the bounded simplex at the CGG shape, with the same tolerances and box and no grid. It uses 95 to 239 calibrations per cell, compared with 489 to 604 for grid plus simplex, and 0.75 to 3.1 seconds, compared with 4.9 to 8.8. It returns the same minima at eleven cells and, at $(\epsilon, \delta, M) = (0.5, 10^{-5}, 10^2)$, the interior minimum 2.013×10^{-2} below CGG that the grid run missed at the boundary $\rho = 1$. The error it removes below CGG is the $\mathcal{O}(M^{-2})$ amount of Corollary 1, at most a relative 6.1×10^{-3} at $M = 10$, 5.2×10^{-6} at $M = 10^2$, and 5.1×10^{-9} at $M = 10^3$. The last two columns put every setup cost on one basis, Step 1 plus one calibration for CGG, where Step 1 takes 13.8 to 17.2 milliseconds per privacy pair. On that basis, the grid search costs 211 to 308 times the CGG setup and 35 to 108 times the warm-started search.

ε	δ	M	Calibrations		Seconds		$\mathrm{MSE}_M^{\mathrm{CGG}} - \mathrm{MSE}_M^{\mathrm{warm}}$	Setup Cost over CGG	
			Grid	Warm	Grid	Warm		Grid	Warm
0.5	10^{-5}	10^1	493	95	5.3	0.75	1.584×10^{0}	240	35
		10^2	604	239	8.8	3.06	2.013×10^{-2}	308	108
		10^3	522	96	6.1	0.96	2.062×10^{-4}	235	38
	10^{-3}	10^1	496	96	5.8	1.18	1.571×10^{-1}	217	45
		10^2	510	218	6.6	2.49	1.668×10^{-3}	248	95
		10^3	515	110	4.9	0.92	1.679×10^{-5}	211	41
1.0	10^{-5}	10^1	495	103	5.8	1.11	5.113×10^{-1}	218	42
		10^2	518	105	6.9	1.22	6.856×10^{-3}	246	44
		10^3	511	95	5.9	0.97	7.062×10^{-5}	220	37
	10^{-3}	10^1	489	107	5.9	1.10	8.404×10^{-2}	243	46
		10^2	503	110	6.6	1.19	9.409×10^{-4}	290	53
		10^3	523	99	5.4	0.77	9.512×10^{-6}	258	38

Table 7: A simplex warm-started at the CGG shape against the grid-plus-simplex run of Table 6, at $\Delta_2 = 1$. "Calibrations" counts the scale calibrations each search performs and "Seconds" is its running time. The gain column is the error the warm start removes below the CGG mechanism. The last two columns divide each search's total setup time, Step 1 of Algorithm 1 plus one calibration plus the search, by the CGG setup time, Step 1 plus one calibration, on the same machine.

H.3 Numerical check of the expansion constants

The CGG error and the error of the power-radial subfamily $(q, p) = (M, \tau_*)$ are exact quadrature values, so we can test the constants in eq. (7) and eq. (68) directly. Figure 7 (left) plots $M^2\{\mathrm{MSE}_M^{\mathrm{CGG}}(1) - A_M(1)\}$ against M. Theorem 3 bounds it, and at every privacy pair it changes by less than 3% between $M = 10^2$ and $M = 10^4$. This confirms v_0, $C(\epsilon, \delta)$, and $J(\epsilon, \delta)$ together, since an error in v_0, $C(\epsilon, \delta)$, or $J(\epsilon, \delta)$ would make the curve grow like M^3, M^2, or M. The sign is negative at $\delta = 10^{-5}$ and positive at $\delta = 10^{-3}$, because A_M is an expansion center and the remainder takes either sign. The right panel plots $M\{\mathrm{MSE}_M^{\mathrm{PR}}(1) - \mathrm{MSE}_M^{\mathrm{CGG}}(1)\}/B_*$, which Corollary 6 predicts to converge to 1. At $M = 10^4$, the four ratios lie in $[1.0010, 1.0025]$. The negative value at $M = 10$ and $(\epsilon, \delta) = (1, 10^{-5})$ means that the power-radial subfamily is marginally better there, a finite-dimensional effect gone by $M = 10^2$. Figure 8 repeats the left panel for $0.1 \le \epsilon \le 1$. The curves for $M = 10^2$ and 10^3 agree to within 3.5% for $\epsilon \ge 0.3$. Over the whole sweep at these two dimensions the scaled remainder lies in $[-896, 752]$ at $\delta = 10^{-5}$ and in $[23, 225]$ at $\delta = 10^{-3}$, so $|\mathrm{MSE}_M^{\mathrm{CGG}}(1) - A_M(1)| \le \frac{900}{M^2}$ on this box, an empirical constant in place of K_{ach}. The $M = 10$ curve separates from them below $\epsilon \approx 0.6$ at $\delta = 10^{-5}$, where $\epsilon^2 v_0$ runs from 9.5 to 12.7, so $M = 10$ is not large compared with it while $M \ge 10^2$ is (Remark 3). At the same δ the $M = 10^2$ and 10^3 curves cross zero between $\epsilon = 0.1$ and $\epsilon = 0.125$, where the M^{-2} coefficient changes sign. The $M = 10$ curve stays negative across the whole sweep.

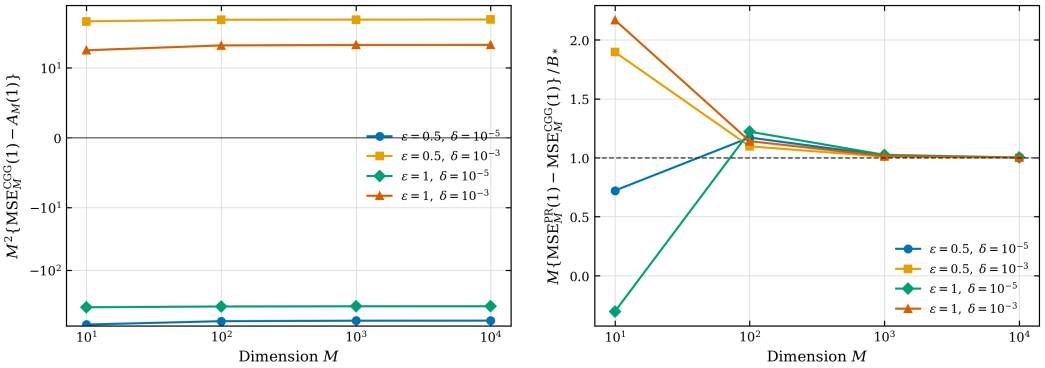


Figure 7: Left: $M^2\{\mathrm{MSE}_M^{\mathrm{CGG}}(1) - A_M(1)\}$ against M. Boundedness confirms the remainder order of Theorem 3 and the constants of A_M. Right: the loss of the power-radial subfamily over its predicted limit, $M\{\mathrm{MSE}_M^{\mathrm{PR}}(1) - \mathrm{MSE}_M^{\mathrm{CGG}}(1)\}/B_*$ against M, with the limit 1 of Corollary 6 dashed.

H.4 Bounds at small dimension

Lemma 5 holds for every $M \ge 2$. The dimension limit is the threshold in Lemma 9. The iteration that encloses the shifted root of eq. (20) contracts only while that threshold's slope stays below $\frac{M}{2}$, and at $(\epsilon, \delta) = (0.5, 10^{-5})$, the slope is bounded by 1.1×10^5 at $M = 10$ and by 1.2×10^3 at $M = 10^2$, against 26 at $M = 10^3$. The lower block of Table 1 therefore starts at $M = 10^3$. At $M \in \{10, 10^2\}$ we bound $\mathcal{V}_M(1)$ by a converse that uses no expansion in $\frac{1}{M}$. A finite family of distinguishing events (half-spaces, hyperboloids, balls, and the Neyman–Pearson events (Neyman & Pearson, 1992) of the calibrated Gaussian, ℓ_2, power-radial, and CGG candidates) is combined by a linear program into a dual certificate for eq. (1). We verify the support inequality over every radius using outward-rounded interval arithmetic (Johansson, 2017) and recheck it at higher working precision. The bound is weaker than Theorem 2 because the event family is finite. On the achievability side, the CGG scale is inflated by a factor of $1 + 2 \times 10^{-5}$ and its privacy profile is checked below δ by interval arithmetic, so its closed-form error is an upper bound. Over the four privacy pairs the two sides bracket $\mathcal{V}_{10}(1)$ within 9.02 and $\mathcal{V}_{10^2}(1)$ within 0.81 in absolute terms, the upper block of Table 1. As a check, at $(\epsilon, \delta) = (0.5, 10^{-5})$ and $M = 10$, the searched error 314.72680 of Table 2 lies inside the bracket $[307.30909, 316.32362]$. Figure 3 shows these eight brackets alongside the eight dual-event ones. The center $A_M(1)$ lies inside eleven of the sixteen brackets. The five ex-

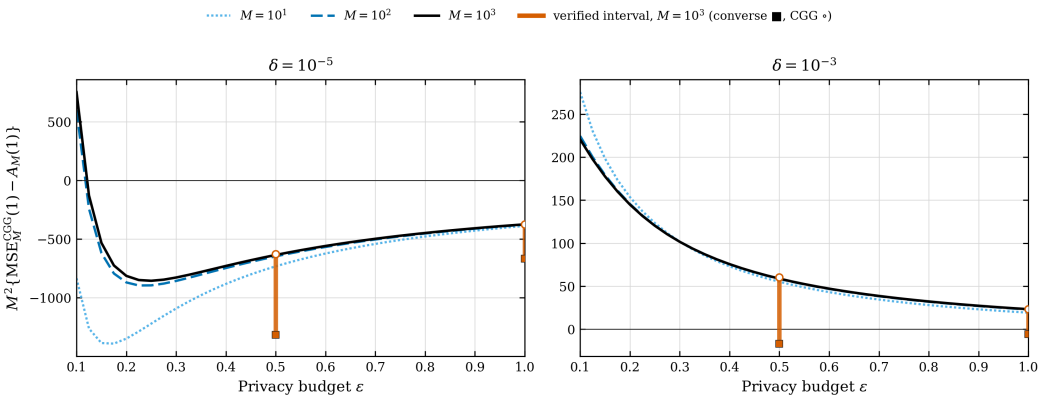


Figure 8: $M^2\{\mathrm{MSE}_M^{\mathrm{CGG}}(1) - A_M(1)\}$ against ϵ for $M \in \{10, 10^2, 10^3\}$ and $\delta \in \{10^{-5}, 10^{-3}\}$. Vertical bars at $M = 10^3$ and $\epsilon \in \{0.5, 1\}$ are the intervals of Table 1 on the same scaled axis (square for the converse, circle for CGG).

ceptions all have $\delta = 10^{-5}$, where $A_M(1)$ exceeds the upper bound, consistent with the negative remainder of those pairs in Figure 7.

H.5 Recorded data and independent recheck

Every entry of Table 1 rests on two records that can be rechecked without repeating the construction. Interval arithmetic returns, for every computed quantity, an interval that contains its exact value, and every step below carries one. The converse record described below is the dual-event one of the lower block. The finite-family records of the upper block store the event family and the linear-program multipliers instead and are rechecked in the same way.

The converse record. A converse record is a threshold θ, a multiplier $\lambda > 0$, a level κ, and a partition of $[0, v_{\max}]$. The threshold is a quintic Hermite spline through 161 knots, with exact rational values of $\theta, \theta', \theta''$ at each knot and extended off the knot window by the two endpoint tangents. It is therefore C^1 and piecewise polynomial, with a second derivative that is continuous except for a jump at each end of the knot window. The secant bound below needs only a Lipschitz first derivative, and on a cell containing a window endpoint, the interval bound $\overline{K}$ on $|F''|$ is taken over both branches. The knot window is $[\frac{1}{2}v_0, \frac{3}{2}v_0]$ and depends only on (ϵ, δ). Table 8 lists it. The knot values do depend on M. They solve a nonlinear program that maximizes κ subject to the support inequality at a fine set of collocation points in that window, starting from the expansion of Lemma 9 and holding the two endpoint slopes at $-\epsilon$. That solution is a proposal. The interval check below turns it into a bound.

M	ϵ	δ	knot window of θ	λ	κ	$v_{\max}$	cells
10^3	0.5	10^{-5}	$[24.7233, 74.1699]$	6.50799×10^5	55.715665	1723787	2999
10^3	0.5	10^{-3}	$[10.6266, 31.8799]$	5.46172×10^3	26.683250	14468	1858
10^3	1	10^{-5}	$[6.9588, 20.8764]$	1.65587×10^5	15.495633	615699	1723
10^3	1	10^{-3}	$[3.3144, 9.9433]$	1.44715×10^3	8.062063	5382	1880
10^4	0.5	10^{-5}	$[24.7233, 74.1699]$	6.57164×10^5	55.994276	1740646	10458
10^4	0.5	10^{-3}	$[10.6266, 31.8799]$	5.48080×10^3	26.730901	14519	6967
10^4	1	10^{-5}	$[6.9588, 20.8764]$	1.67437×10^5	15.584176	622580	7610
10^4	1	10^{-3}	$[3.3144, 9.9433]$	1.45390×10^3	8.081360	5407	6011

Table 8: Converse records behind Table 1, abridged. The full decimal expansions of λ and of the 161 knots are in the supplementary material and are read as exact rationals. The knot window of θ depends only on (ϵ, δ). "Cells" counts the recorded partition of $[0, v_{\max}]$ that the recheck re-verifies.

Covering the half-line. Write $F(v) = v + \lambda h_M[\theta](v)$ for the left side of Lemma 5. On $[0, v_{\max}]$, the record stores an adaptively refined partition into cells. Each cell $[l, r]$ is bounded below in two ways, and we keep the larger bound. The first is a direct interval bound on F over the cell. The second is the secant bound $\min\{F(l), F(r)\} - \frac{1}{8}\overline{K}(r-l)^2$ with $\overline{K}$ an interval bound on $|F''|$ there. The second preserves the cancellation between the two sphere-CDF terms of $h_M[\theta]$, which a direct interval bound loses. Inside each cell, the shifted threshold solving eq. (20) is bounded by an interval fixed-point iteration seeded by the spline's global slope bound, each step of which is itself a valid bound. The cell containing $v = 0$ is treated in closed form, since both thresholds then lie in the common left-tangent branch and $h_M[\theta]$ is an explicit difference of two values of H_M.

The partition stops at $v_{\max} = \lceil (1 + e^{\epsilon})\lambda \rceil + 1$. Beyond it, no computation is needed. Since $h_M[\theta] \geq -e^{\epsilon}$ pointwise, $F(v) \geq v - \lambda e^{\epsilon} \geq v_{\max} - \lambda e^{\epsilon} > \lambda$ for $v \geq v_{\max}$, and $\kappa < \lambda$ at every entry of Table 8. The recorded κ is the minimum of the cell bounds, reduced by a relative 10^{-13} and rounded down, and Lemma 5 gives $\underline{\mathrm{MSE}}_M = M(\kappa - \lambda\delta)$.

The event mass is evaluated exactly. The verification evaluates H_M exactly. Since $U_1^2 \sim \mathrm{Beta}(\frac{1}{2}, \frac{M-1}{2})$, the sphere distribution function of Lemma 4 is

$$H_M(y) = \frac{1}{2}\left\{1 + \mathrm{sgn}(y) I_{\frac{y^2}{M}}\left(\tfrac{1}{2}, \tfrac{M-1}{2}\right)\right\}, \qquad |y| \leq \sqrt{M}, \tag{71}$$

where I is the regularized incomplete beta function, and $H_M(y) \in \{0, 1\}$ outside $[-\sqrt{M}, \sqrt{M}]$. Every cell bound evaluates eq. (71) in interval arithmetic, and for $|y| \geq 12$ it uses the two-sided tail bound of Lemma 4,

$$H_M(y) \in \left[0, e^{-\frac{y^2}{2}}\right] \ \ (y \leq -12), \qquad H_M(y) \in \left[1 - e^{-\frac{y^2}{2}}, 1\right] \ \ (y \geq 12),$$

both consequences of $\Pr(\sqrt{M}U_1 \geq x) \leq e^{-\frac{x^2}{2}}$ and its reflection, valid for every $M \geq 2$. The expansion $H_M = \Phi + \phi\{\frac{W_1}{M} + \frac{W_2}{M^2}\} + \mathcal{O}(M^{-3})$ of Lemma 4 enters only the asymptotic proofs of Appendices D and E.

M	ϵ	δ	$\underline{\mathrm{MSE}}_M$	$\overline{\mathrm{MSE}}_M$	$\delta - \overline{D}_\epsilon$
10^3	0.5	10^{-5}	49207.673063944369065 $\pm 1.6 \times 10^{-16}$	49207.673753600053039 $\pm 3.4 \times 10^{-16}$	3.08×10^{-17}
10^3	0.5	10^{-3}	21221.528316234219018 $\pm 1.4 \times 10^{-17}$	21221.528393597940325 $\pm 1.9 \times 10^{-17}$	2.69×10^{-16}
10^3	1.0	10^{-5}	13839.766122041219512 $\pm 2.0 \times 10^{-16}$	13839.766414712629186 $\pm 4.2 \times 10^{-16}$	2.74×10^{-18}
10^3	1.0	10^{-3}	6614.9105383437834998 $\pm 1.6 \times 10^{-17}$	6614.9105678193820765 $\pm 3.2 \times 10^{-17}$	6.94×10^{-17}
10^4	0.5	10^{-5}	494226.33297080468023 $\pm 3.2 \times 10^{-15}$	494226.33299998053805 $\pm 3.0 \times 10^{-15}$	4.84×10^{-18}
10^4	0.5	10^{-3}	212501.03491708311540 $\pm 1.6 \times 10^{-15}$	212501.03492598559935 $\pm 4.2 \times 10^{-16}$	8.45×10^{-17}
10^4	1.0	10^{-5}	139098.03922210611015 $\pm 3.3 \times 10^{-15}$	139098.03922823122481 $\pm 4.2 \times 10^{-15}$	1.60×10^{-18}
10^4	1.0	10^{-3}	66274.627973170830131 $\pm 3.9 \times 10^{-16}$	66274.627975962804331 $\pm 3.0 \times 10^{-16}$	3.37×10^{-16}

Table 9: Unrounded endpoints behind Table 1, as intervals. Table 1 reports their outward roundings. The last column is the verified privacy margin $\delta - \overline{D}_\epsilon$, where $\overline{D}_\epsilon$ is the *upper* endpoint of the interval bound on $D_\epsilon(P_{Z_M} \| (P_{Z_M})_{\Delta_2 e_1})$ for the inflated candidate. Every margin is strictly positive, so feasibility holds for every value in that interval. The midpoint is printed to twenty significant digits, and the radius is the recorded one enlarged by that rounding and then rounded up, so each printed ball contains the recorded interval. The printed radii therefore lie between 10^{-17} and 5×10^{-15} while the recorded radii lie between 1.0×10^{-32} and 3.1×10^{-30}. The full-precision intervals are in the supplementary material.

The achievability record. An achievability record is a candidate (q, p, β), a radius $r_{\max}$, and a partition of $[0, r_{\max}]$. We compute the privacy profile at the shift $\Delta_2 e_1$ by integrating cell by cell using interval arithmetic. We choose $r_{\max}$ so the radial tail probability beyond it is below 10^{-15}, and add the incomplete-gamma tail to the interval bound on the privacy profile. That bound lies strictly below δ at every entry (Table 9), which requires inflating the reported scale slightly above the exactly calibrated one. The checked object is the inflated mechanism, so its error bounds $\mathcal{V}_M(1)$ from above. The recorded q and p are likewise finite decimals, read as exact rationals, and agree with $\rho_* M$ and p_* to a relative 10^{-15}. By Corollary 3, the calibrated scale at the rounded shape is at most the inflated scale, so the exactly calibrated error at that shape is at most $\overline{\mathrm{MSE}}_M$, which is what eq. (69) requires. Replacing the rounded shape by the exact one moves the calibrated error by less than the last printed digit. Since $\rho_* \le 1$ makes the radial density nonincreasing, Lemma 10 extends privacy to the whole sensitivity ball, and $\overline{\mathrm{MSE}}_M$ is the outward rounding of an interval bound on a ratio of two gamma functions. Across the eight entries, the inflation raises $\overline{\mathrm{MSE}}_M$ above the exactly calibrated CGG error by between 2.6×10^{-7} and 2.0×10^{-5}, at most two units in the fifth decimal. This, together with the outward rounding of $\overline{\mathrm{MSE}}_M$, is why the CGG column of Table 2, which reports the uninflated error, differs from $\overline{\mathrm{MSE}}_M$ in the last digit at three of the four pairs at $M = 10^3$.

One pair of bounds outside the matching window. At $(\epsilon, \delta) = (3, 10^{-5})$ Assumption 2 fails, with $p_* = 1.06192 < \tau_* = 1.07933$ and $B_{\mathrm{sh}} = 5.10 \times 10^{-3}$. The same construction at $M = 10^3$, with the achievability record at the clipped shape $(\rho, p) = (1, \tau_*)$ of Theorem 4, gives $\underline{\mathrm{MSE}}_M = 1919.55257$ and $\overline{\mathrm{MSE}}_M = 1919.55263$ against $\mathrm{MSE}_M^{\mathrm{AGM}} = 1933.75016$, a width of 4.84×10^{-5}. The candidate therefore recovers at least 99.9996% of the Gaussian excess. Both $A_M(1) = 1919.55278$ and $A_M(1) + \frac{B_{\mathrm{sh}}}{M} = 1919.55279$ lie 1.6×10^{-4} above the upper endpoint, and $\frac{B_{\mathrm{sh}}}{M} = 5.1 \times 10^{-6}$ is one order below the interval width. At this dimension, the $\mathcal{O}(M^{-2})$ remainder dominates the open $\frac{1}{M}$ term, so the bounds confirm the interval of Theorem 4 and leaves the coefficient itself open. The record is rechecked with the other sixteen.

Bounds at other parameters. The released builders take any privacy pair and dimension. The converse needs Assumption 1. The achievability candidate is the CGG mechanism inside the matching window and the candidate $(\rho, p) = (1, \tau_*)$ outside it, so it is feasible wherever Assumption 1 holds. The dual-event converse needs the spline slope below $\frac{M}{2}$, which holds at $M \ge 10^3$ for the pairs reported, and the finite-family converse of Appendix H.4 has no dimension threshold. One case takes about 90 seconds at $M = 10^3$ and about two minutes at $M = 10^4$ on one core, and the recheck runs at 224 bits.

Error propagation and recheck. The incomplete-beta and incomplete-gamma primitives carry guaranteed error bounds, the spline is evaluated on rational coefficients, and sums over cells accumulate as intervals, so every reported radius includes all rounding. The interval bound on the privacy profile of the final inflated candidate checks the scale, so the bisection that produced it is outside the trusted path. The recorded partitions are rechecked at higher working precision than they were built. The cells are shown to tile $[0, v_{\max}]$ or $[0, r_{\max}]$ without gap or overlap, each cell bound is recomputed, and each reported endpoint is rebuilt from the record alone, as $M(\kappa - \lambda\delta)$ on the converse side and as the second moment of the recorded candidate on the achievability side. The records are in the supplementary material, together with a check of the constants in eq. (15). One part confirms the identities eq. (16) symbolically from the closed forms of c_0, c_1, c_2. The other calibrates the CGG mechanism at $M \in \{2 \times 10^3, 4 \times 10^3, 8 \times 10^3\}$, forms $\widehat{J}_M = M^2\{v_M - v_0 + \frac{C(\epsilon,\delta)}{M}\}$ from the quadrature values, and recovers $J(\epsilon, \delta)$ to a relative 7×10^{-7} or better at all four privacy pairs from the Richardson combination $2\widehat{J}_{8000} - \widehat{J}_{4000}$, which removes the $\mathcal{O}(M^{-1})$ term of $\widehat{J}_M$.

H.6 The calibrated error as a function of the shape

The final calculation shows why the search in Section 4.2 recovers only $\mathcal{O}(M^{-2})$. At $M = 10^3$, for each privacy pair, we evaluate the exactly calibrated error of generalized gamma noise on a 21×21 grid of shapes (ρ, p) around (ρ_*, p_*) and record the relative increase over the CGG error (Figure 9). We clip the box at $\rho = 1$, beyond which the radial density is no longer nonincreasing, and Lemma 10 does not apply. That increase stays below 0.18% over the whole window. No shape

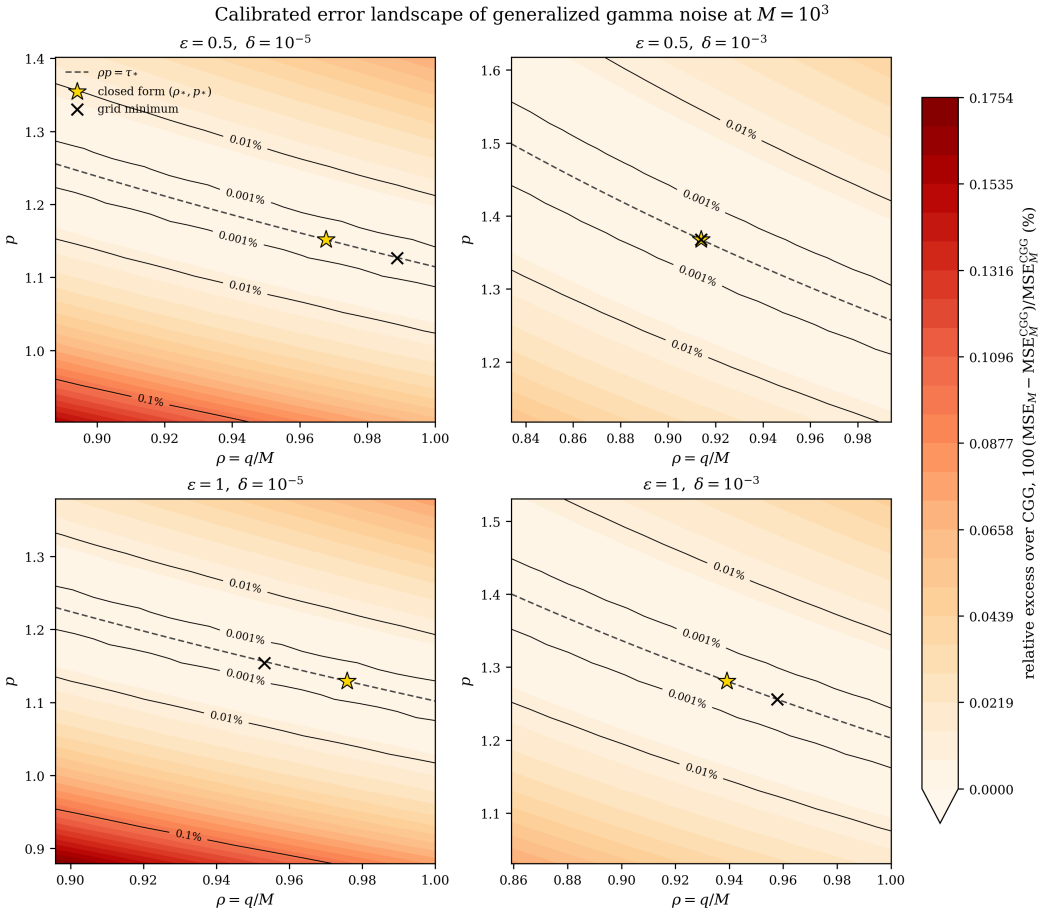


Figure 9: Relative increase of the exactly calibrated generalized gamma error over the CGG error at $M = 10^3$, on a 21×21 grid of shapes (ρ, p) around (ρ_*, p_*) (star), clipped at $\rho = 1$. One panel per privacy pair, each on its own axes since the grid is centered at that pair's (ρ_*, p_*). Level sets are elongated along the curve $\rho p = \tau_*$ of constant effective shape, on which the grid minimum (cross) also lies. That minimum improves on the CGG point only at $(\epsilon, \delta) = (0.5, 10^{-5})$, by $2.1 \times 10^{-8}\%$. At $(0.5, 10^{-3})$ it coincides with it, and at the other two pairs it is worse by at most $1.9 \times 10^{-6}\%$.

on the grid improves on the CGG mechanism by more than $2.1 \times 10^{-8}\%$, and at three of the four pairs none improves on it at all. The level sets are elongated along the curve $\rho p = \tau_*$ of constant effective shape. The error is sensitive to τ, which fixes the constant order through $C(\epsilon, \delta)$, and nearly flat along the curve, where only the $\frac{1}{M}$ coefficient moves. The grid minimum shows this directly. It sits at an effective shape within 2×10^{-3} of τ_* at every pair while its (ρ, p) is displaced by up to 0.034. A shape displaced by $\mathcal{O}(M^{-1})$ from (ρ_*, p_*) therefore changes the error by $\mathcal{O}(M^{-2})$, which is the behavior Figure 4 reports.

I Boundary cases and extensions

The sharp expansion concerns a fixed interior privacy pair, the full ℓ_2 shift ball, data-independent additive noise, and the mean squared error of the noise. This section records the linear program behind eq. (1), its relation to the scalar endpoint results, and the candidate mechanisms for the shape-infeasible region.

Measure-linear formulation. For a measurable $\mathcal{S} \subseteq \mathbb{R}^M$, a shift length $t \in [0, \Delta_2]$, and a radius $r \geq 0$, let $k_{t,\mathcal{S}}(r) = \Pr(rU \in \mathcal{S}) - e^\epsilon \Pr(rU + te_1 \in \mathcal{S})$. After the reduction of Lemma 2, problem eq. (1) is the infinite-dimensional linear program over radial distributions μ in the set $\mathcal{P}([0,\infty))$ of probability measures on $[0,\infty)$,

$$\mathcal{V}_M(\Delta_2) = \inf_{\mu \in \mathcal{P}([0,\infty))} \left\{ \int_0^\infty r^2 \mu(\mathrm{d}r) : \int_0^\infty k_{t,\mathcal{S}}(r)\mu(\mathrm{d}r) \leq \delta \ \ \forall t \in [0,\Delta_2],\ \forall \mathcal{S} \text{ measurable} \right\}. \tag{72}$$

Indeed, for $R \sim \mu$ independent of U, Tonelli's theorem gives $\int k_{t,\mathcal{S}} \mathrm{d}\mu = \Pr(RU \in \mathcal{S}) - e^\epsilon \Pr(RU + te_1 \in \mathcal{S})$. The supremum over $\mathcal{S}$ is $D_\epsilon(P \| P_{te_1})$. Rotation invariance lets the scalar shift length cover the ball, and $\|RU\|_2^2 = R^2$. The feasible set admits point masses, continuous radial densities, and mixtures. Lemma 5 selects one event from the constraint family and bounds it pointwise. This is the radial, high-dimensional counterpart of the linear-programming approach to optimal noise of Geng & Viswanath (2016b;a); Geng et al. (2019).

Scalar endpoints. At $\epsilon = 0$, the hockey-stick divergence is the total variation distance. In the scalar symmetric monotone class, Geng et al. (2019) reduce the constraint to a central-interval mass bound whose squared-loss optimizer in the high-privacy branch is uniform on $[-\frac{\Delta_2}{2\delta}, \frac{\Delta_2}{2\delta}]$, with error $\frac{\Delta_2^2}{12\delta^2}$. At $\delta = 0$, pure DP bounds the likelihood ratio almost everywhere, and the scalar optimizer is the staircase family of Geng & Viswanath (2016b). Proposition 7 shows that Gaussian feasibility and the fixed-shape expansion change singularly at both axes, so the interior expansion does not extend to these optimizers.

The shape-infeasible region. When $p_* < \tau_*$, the unconstrained minimizer in eq. (16) violates the monotone-density condition $p \geq \tau$. The constrained minimizer on the ridge is $p = \tau_*$ with penalty B_{sh}. Theorem 4 shows this penalty enters only the achievability side. The contact-set structure of the classical exact mechanisms suggests four radial candidate classes for closing the interval:

$f(r) \propto r^{q-M} e^{-\beta r^p} \mathbf{1}_{\{r \leq r_{\mathrm{cut}}\}}$	truncated generalized gamma noise,
$f(r) = \gamma_j$ for $r \in [r_j, r_{j+1})$	radial step distribution,
$\mu = w\mathrm{Dirac}_{r_\circ} + (1-w)\mu_{\mathrm{c}}$	point mass on a shell plus a continuous part,
$\mu = w\mu_1 + (1-w)\mu_2$	finite mixture of generalized gamma noise.

For nonmonotone candidates or candidates with point masses, Lemma 10 does not apply, so a valid construction must verify the constraint in eq. (72) at every $t \in [0, \Delta_2]$.