PRIVACY ADMISSION FROM INCOMPLETE RÉNYI DIFFERENTIAL PRIVACY HISTORIES

Anonymous authors
Paper under double-blind review

ABSTRACT

A privacy accountant must decide whether a proposed release, together with past outputs, meets a total differential privacy budget. We study this decision when the historical record contains only upper bounds at finitely many Rényi differential privacy (RDP) orders, while the proposed request has a known privacy profile. The record defines a class of possible histories. For a request fixed independently of past private outputs and run with fresh randomness, we maximize its composed privacy value over that class while preserving the request as a factor. Retaining only multiplied RDP moments instead permits aggregate laws inconsistent with the specified request. We prove strict randomized-response and Gaussian-DP separations in which this information loss blocks safe admission. Moment duality gives checkable upper bounds with explicit allowances for historical tails and numerical errors. Controlled comparisons recover admissions that the aggregate summary cannot justify and distinguish this information gain from improvements in computed bounds. A second comparison holds the request profile fixed: direct optimization sharpens some admission bounds relative to a reusable historical profile, while reuse trades preparation cost for cheaper subsequent queries.

1 INTRODUCTION

Before a new data release is executed, a privacy accountant must verify that the combined output of past releases and the proposed request satisfies a public $(\varepsilon_B, \delta_B)$ differential privacy budget. This is a constraint on the joint disclosure, not a separate budget for the new release. An unnecessarily loose composition bound can therefore prevent a release even when the available information suffices to admit it. We ask how to use that information when the past and the proposed request are described at different levels of detail.

A finite Rényi differential privacy (RDP) record provides a compact interface between accountants. For example, heterogeneous randomized-response, Gaussian, and Laplace releases can be accounted for on a common list of RDP orders (Mironov, 2017, Section VII). A subsequent accountant can receive that vector of bounds without receiving the historical privacy-loss laws. Each entry bounds a likelihood-ratio moment, so the vector defines a class of compatible histories rather than a single law. Section 3 shows that even exact moments at finitely many orders generally leave uncertainty about future composition.

The proposed request can be described more precisely. Its mechanism and noise level may specify an entire privacy profile: its guarantee at every privacy threshold. We consider a request chosen independently of past private outputs and run with fresh randomness. Every possible joint law is then the product of a compatible history and this particular request. Multiplying historical and request moment bounds is valid, but keeping only the products drops the requirement that the joint law contain the specified request as a factor. Historical compression has already lost information; compressing the known request creates a second, avoidable loss.

Our accountant retains the request profile and maximizes its composed privacy value over all histories allowed by the record. A checked upper bound meeting the budget in every ordered neighboring-data direction authorizes execution without reconstructing the actual past. Section 4 proves that this distinction can change admission for randomized-response and Gaussian requests: every compatible

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.

product can meet a cap while an aggregate law satisfying the multiplied moments exceeds it. Exact optimization of the aggregate class therefore cannot remove the gap.

Section 5 turns the request-specific target into a moment optimization problem, with explicit tail and numerical margins. We also compare a request-preserving alternative: prepare a historical profile that dominates every compatible past, then compose it with later candidates. This separates the value of retaining information from the choice between a request-specific bound and a reusable one. Section 6 reports both admission decisions and complete accounting costs, distinguishing aggregate laws that rule out admission from computed brackets that remain inconclusive.

2 Related Work

RDP composes by adding logarithmic moment bounds (Mironov, 2017); moment accounting for private learning and concentrated privacy use related summaries (Abadi et al., 2016; Wang et al., 2019; Bun & Steinke, 2016). Classical and optimal composition characterize guarantees obtained from component (ε, δ) descriptions (Dwork & Roth, 2014; Kairouz et al., 2015; Murtagh & Vadhan, 2018). Pointwise-optimal conversion from an RDP curve to a trade-off guarantee characterizes what the RDP summary itself implies (Riess et al., 2026). The product constraint studied here supplies information beyond the multiplied moment summary, even when conversion of that summary is optimal.

Richer component descriptions support privacy-loss convolution, including fast Fourier transform methods (Koskela et al., 2020; Gopi et al., 2021), and characteristic-function accounting (Zhu et al., 2022). Pessimistic discretization gives safe finite approximations (Doroshenko et al., 2022), while multiple simultaneous DP constraints provide another component description (Cadir et al., 2026). Our historical input is a set of moment inequalities; the request retains an analytic or supplied profile. Gaussian DP provides an analytic example (Dong et al., 2022). Private-selection profiles (Koskela et al., 2024) also apply when the candidate set, parameters, and stopping rule are fixed independently of past private outputs.

Dominating-pair composition and profile suprema underlie the reusable historical-profile comparison (Zhu et al., 2022; Regehr et al., 2026). Privacy filters address adaptive spending and conditions under which composition commutes with profile suprema (Regehr et al., 2026). Here we optimize the independently fixed, fresh-randomness product using classical generalized-moment duality and finite-support representations (Winkler, 1988; Anderson & Nash, 1987). The supplement also uses Tchebycheff-system structure (Karlin & Studden, 1966) to characterize exact identification on finite-atomic classes. The next section specifies the information available to the privacy accountant and the admission guarantee it must establish.

3 Admission from an RDP History

The privacy accountant receives a historical record, a proposed request, and a public budget $(\varepsilon_B, \delta_B)$. Let $\mathcal{D}$ index ordered neighboring-data pairs, and let P_d, Q_d denote the corresponding laws of the combined output. The target is

$$P_d(A) \le e^{\varepsilon_B} Q_d(A) + \delta_B \quad \text{for every measurable output event } A \text{ and every } d \in \mathcal{D}.$$

Here e^{ε_B} is the permitted multiplicative probability ratio, and δ_B caps the remaining excess probability. This guarantee must hold after composition for every historical law consistent with the inputs.

Historical receipt. For $d \in \mathcal{D}$, let $P_{H,d}$ and $Q_{H,d}$ be the historical output laws and $E_{H,d} = (P_{H,d}, Q_{H,d})$ their binary experiment. We assume mutual absolute continuity and define the positive likelihood ratio under $Q_{H,d}$ by

$$X_d = \frac{dP_{H,d}}{dQ_{H,d}}, \qquad \mathbb{E}X_d = 1.$$

Its directional hockey-stick profile is $h_{E_{H,d}}(\gamma) = \mathbb{E}(X_d - \gamma)_+$, where $(z)_+ = \max\{z, 0\}$. At $\gamma = e^{\varepsilon}$, the bound $h_{E_{H,d}}(\gamma) \le \delta$ is equivalent to (ε, δ) privacy in direction d. Expectations of X_d use its $Q_{H,d}$ law.

At distinct RDP orders $\alpha_j > 1$, the historical receipt stores

$$M_d = \{M_{d,j}^+, M_{d,j}^-\}_{j=1}^m, \qquad \mathbb{E}X_d^{\alpha_j} \le M_{d,j}^+, \quad \mathbb{E}X_d^{1-\alpha_j} \le M_{d,j}^-.$$

Each bound is the exponential of $(\alpha_j - 1)$ times the corresponding RDP bound. The positive exponent represents forward RDP; the negative exponent represents reverse RDP because $\mathbb{E}_{Q_{H,d}} X_d^{1-\alpha_j} = \mathbb{E}_{P_{H,d}} X_d^{-\alpha_j}$. The signs distinguish exponent signs within one direction; the reversed neighboring pair has its own receipt. A saved RDP bound valid uniformly over ordered neighboring pairs can supply both signed bounds. Write $\mathcal{U}(M_d)$ for all strictly positive, mean-one likelihood-ratio laws satisfying these inequalities and any support or tail constraints supplied with the receipt. These are upper-bound constraints: a compatible law may have slack at any saved order, and admission must cover the entire class. Section 5 specifies the support and tail information used for numerical bounds.

Fixed request and admission rule. The request is a family of directional experiments $F_d = (P_{F,d}, Q_{F,d})$, selected independently of historical private outputs and run with fresh randomness on output space $\mathcal{Y}_F$. To include outputs possible only under $P_{F,d}$, decompose

$$P_{F,d} = P_{F,d}^{\mathrm{ac}} + P_{F,d}^{\perp}, \quad s_{F,d} = P_{F,d}^{\perp}(\mathcal{Y}_F), \quad Y_d = dP_{F,d}^{\mathrm{ac}}/dQ_{F,d}.$$

The singular mass $s_{F,d}$ enters the full request profile

$$h_{F_d}(t) = s_{F,d} + \mathbb{E}_{Q_{F,d}}(Y_d - t)_+.$$

Conditional on $X_d = x$, composition multiplies the request likelihood ratio by x. Its contribution at threshold γ is x times the request profile at γ/x. Thus

$$h_{E_{H,d} \otimes F_d}(\gamma) = \mathbb{E}[X_d h_{F_d}(\gamma/X_d)] = \mathbb{E}g_{\gamma,F_d}(X_d), \qquad g_{\gamma,F_d}(x) = x h_{F_d}(\gamma/x).$$

For every valid request profile, g_{γ,F_d} is convex, nondecreasing, and 1-Lipschitz in x. The same identity gives one-sided objective bounds from a profile enclosure. Even at one budget threshold, varying x uses the request profile across thresholds, retaining more information than its moments at the saved orders.

The request-specific target is the largest composed privacy value allowed by the historical record:

$$W_{F_d}(\gamma; M_d) = \sup_{\mathcal{L}(X_d) \in \mathcal{U}(M_d)} \mathbb{E}g_{\gamma,F_d}(X_d).$$

Only the historical law varies. A joint-moment accountant instead converts the request to bounds at the saved orders and multiplies them into the receipt. Let $N_{F_d,j}^+$ and $N_{F_d,j}^-$ bound the request moments, with $N_{F_d,j}^+ = \infty$ when $s_{F,d} > 0$ and $N_{F_d,j}^- = \infty$ when $Q_{F,d}(Y_d = 0) > 0$. When the required bounds are finite, define

$$\mathcal{K}_{F_d} = \{\mathcal{L}(X_d Y_d) : \mathcal{L}(X_d) \in \mathcal{U}(M_d),\ X_d \perp Y_d\},$$
$$\mathcal{A}_{\mathrm{agg},d} = \{\mathcal{L}(Z) : Z > 0,\ \mathbb{E}Z = 1,\ \mathbb{E}Z^{\alpha_j} \le M_{d,j}^+ N_{F_d,j}^+, \mathbb{E}Z^{1-\alpha_j} \le M_{d,j}^- N_{F_d,j}^-\ \forall j\}.$$

The product class $\mathcal{K}_{F_d}$ preserves the request; the aggregate class $\mathcal{A}_{\mathrm{agg},d}$ retains only multiplied moments. Independence gives $\mathcal{K}_{F_d} \subseteq \mathcal{A}_{\mathrm{agg},d}$ and hence

$$W_{F_d}(\gamma; M_d) = \sup_{\mathcal{L}(Z) \in \mathcal{K}_{F_d}} \mathbb{E}(Z - \gamma)_+ \le C_{\mathrm{agg},d}(\gamma; M_d) := \sup_{\mathcal{L}(Z) \in \mathcal{A}_{\mathrm{agg},d}} \mathbb{E}(Z - \gamma)_+.$$

The profile calculation handles singular request mass directly; the joint-moment calculation is inconclusive when a required request moment is infinite.

Set $\gamma_B = e^{\varepsilon_B}$. A checked upper bound on $W_{F_d}(\gamma_B; M_d)$ at most δ_B in every direction authorizes execution. A compatible history whose composed value exceeds the cap rules out this robust guarantee, without identifying the actual past. For the joint-moment comparison, an aggregate law above the cap rules out admission from that summary even when no compatible product crosses it. A computed lower–upper bracket that straddles the cap is inconclusive. Section 5 constructs these brackets from exact request profiles or outward enclosures.

What the receipt determines. A receipt is future-sufficient if it determines the composed privacy profile for every allowed request. The following result isolates historical compression by considering exact saved moments rather than upper bounds.

Proposition 1 (Information retained by finite receipts)**.** *No fixed finite set of RDP orders is universally future-sufficient over the full finite-output class: for every such set and every nondegenerate Gaussian factor, two bounded finite experiments can share all saved bidirectional moments yet have different future profiles. On the subclass whose likelihood-ratio law has at most k positive atoms, exactly $k-1$ distinct bidirectional orders are sufficient and, in general, necessary when moments are exact. On a fixed bounded interval, the integer orders $2, \dots, m+1$ give a constructive uniform ambiguity bound of order $O(m^{-1})$ for every fixed-factor objective.*

For the first statement with m saved orders, choose a likelihood-ratio law with strictly positive probabilities on $2m+3$ points straddling one. Each saved order contributes one positive and one negative moment constraint. Together with total mass and mean, these give $2m+2$ linear constraints. A small nonzero perturbation preserving them gives distinct histories with the same receipt. Gaussian composition adds Gaussian noise to the log likelihood ratio. This convolution is injective because the Gaussian characteristic function never vanishes. The composed profiles therefore differ at some threshold, and the accountant must cover both histories.

With at most k atoms per law, the difference of two histories instead has at most $2k$ support points. At $k-1$ distinct bidirectional orders, the moment functions, mass, and mean have full column rank on those points, forcing equal laws. Structural knowledge of the past therefore determines whether exact recovery from a finite record is possible. For example, the exact positive and negative order-two moments identify a history known to have at most two likelihood-ratio values. Without this restriction, matching those moments need not imply the same history. The number of saved orders is therefore meaningful together with the historical class to which the receipt applies.

On a bounded interval, more integer orders progressively reduce the uncertainty in future composition. To see why, let ν_1, ν_2 share the exact integer-order receipt on $[\ell, u]$. Their integrals agree for every polynomial p of degree at most $m+1$, so

$$\left| \int g_{\gamma,F}\, d\nu_1 - \int g_{\gamma,F}\, d\nu_2 \right| = \left| \int (g_{\gamma,F} - p)\, d(\nu_1 - \nu_2) \right| \le 2\|g_{\gamma,F} - p\|_\infty.$$

Here the norm is over $[\ell, u]$. Polynomial approximation of the 1-Lipschitz objective gives the stated $O(m^{-1})$ bound, uniformly over the request and threshold on this fixed interval. The supplement's information section proves all three statements. For the upper-bound receipts used by the accountant, the optimization retains all histories satisfying the bounds; the moment multipliers in Section 5 quantify how rounding those bounds affects the admission value.

4 When the Request Representation Changes Admission

Fix one ordered direction and suppress d and the receipt argument. A safe admission changes when a request-specific upper bound meets the cap but a feasible aggregate law exceeds it. We establish this separation using only the historical bidirectional order-two bounds $\mathbb{E}X^2, \mathbb{E}X^{-1} \le 1+\eta$, with $\eta \ge 0$; no historical support information is supplied.

The parameter η bounds deviation from the no-disclosure law $X = 1$. Since $\mathbb{E}X = 1$, the receipt implies $\mathbb{E}(X-1)^2 \le \eta$. The request-specific objective is 1-Lipschitz, so

$$W_F(\gamma) \le h_F(\gamma) + \sup_{\mathcal{U}(M)} \mathbb{E}|X-1| \le h_F(\gamma) + \sqrt{\eta}. \tag{1}$$

The standalone request value $h_F(\gamma)$ and the possible historical contribution $\sqrt{\eta}$ have separate roles. A cap at least their sum is safe for every compatible product. To separate the decisions, it remains to find a law satisfying the multiplied moments whose value exceeds that sum. The next theorem does so for two request families.

Let F_a be symmetric randomized response (RR) with $a > 1$: under Q, its likelihood ratio is a with probability $1/(a+1)$ and a^{-1} otherwise. Let G_μ denote Gaussian differential privacy (GDP), and let Φ be the standard-normal cumulative distribution function. The privacy loss of G_μ is $N(\mu^2/2, \mu^2)$ under P and $N(-\mu^2/2, \mu^2)$ under Q.

Theorem 2 (Strict separations)**.** *Suppose the historical bidirectional order-two receipt obeys* $\mathbb{E}X^2, \mathbb{E}X^{-1} \le 1 + \eta$.

Randomized response. *Let* F_a *be symmetric randomized response,* $q(a) = a^3 - 2a^2 - 2a + 2$*, and let* $a_0 \simeq 2.4812$ *be the largest real root of* q*. For* $a \ge a_0$*, put* $\gamma_a = 2a - 1$ *and* $p_a = (2a^2 + 3a - 1)^{-1}$*. Then*

$$W_{F_a}(\gamma_a) \le \sqrt{\eta}, \qquad C_{\mathrm{agg}}(\gamma_a) \ge p_a.$$

Hence the separation is strict when $\eta < p_a^2$.

Gaussian DP. *For* $F = G_\mu$ *with* $0 < \mu \le 4$*, define*

$$\Delta_G(\mu) = \tanh(\mu/2) - \{2\Phi(\mu/2) - 1\} > 0.$$

At $\gamma = 1$*,*

$$W_{G_\mu}(1) \le 2\Phi(\mu/2) - 1 + \sqrt{\eta}, \qquad C_{\mathrm{agg}}(1) \ge \tanh(\mu/2),$$

so strictness holds when $\eta < \Delta_G(\mu)^2$ *throughout the stated range* $0 < \mu \le 4$.

For RR, $\gamma_a = 2a - 1$ exceeds the largest request likelihood ratio a, so the request alone contributes zero at this threshold. Historical uncertainty contributes at most $\sqrt{\eta}$. The aggregate construction instead places mass p_a one unit above γ_a, giving value p_a. Every cap in $[\sqrt{\eta}, p_a)$ therefore admits the specified request but is crossed by a law allowed by its multiplied moments. The supplement's separation section gives the construction and an exact RR admission example.

For GDP, $\Delta_G(\mu)$ is the gap between the standalone value at $\gamma = 1$ and the value of a feasible aggregate distribution. When $\sqrt{\eta} < \Delta_G(\mu)$, every cap from $2\Phi(\mu/2) - 1 + \sqrt{\eta}$ up to, but excluding, $\tanh(\mu/2)$ separates the decisions. At $\eta = 0$, the receipt identifies the history as $X = 1$ almost surely, yet compressing the request still creates a gap; the supplement's Gaussian plot illustrates this case. For positive slack, useful admission bounds must account for the remaining historical uncertainty, which is the computational task addressed next.

5 A Checkable Fixed-Request Accountant

For the specified request, the quantity to compare with the public cap is the largest composed privacy value allowed by the historical receipt. We compute a bracket $[L, U]$ for this value, $W_F(\gamma_B; M)$, in each ordered direction, suppressing d below. A feasible historical law supplies L, and a pointwise majorant supplies U. An independent checker validates both before they enter the decision: $U \le \delta_B$ in every direction admits the request, while $L > \delta_B$ in one direction rules out the robust guarantee. The supplement's algorithm retains valid bounds during refinement and returns an inconclusive status with the checked brackets if its computation budget is exhausted before a decision.

5.1 Optimizing over a bounded history class

An upper bound must cover every likelihood-ratio law allowed by the receipt on a valid support interval $[\ell, u] \Subset (0, \infty)$. Write $W_F^{[\ell,u]}(\gamma; M)$ for the restricted value. Normalization, mean one, and the receipt bounds are linear constraints on the unknown law. The fixed request enters through $g_{\gamma,F}(x) = x h_F(\gamma/x)$. The primal searches for a compatible law; the dual searches for a combination of saved moment functions that majorizes this objective throughout the interval.

Theorem 3 (Fixed-factor moment optimization)**.** *Fix* $0 < \ell < 1 < u < \infty$ *and let* $\mathcal{M}$ *be the nonempty set of probability laws on* $[\ell, u]$ *satisfying* $\mathbb{E}X = 1$ *and* $\mathbb{E}X^{\alpha_j} \le M_j^+, \mathbb{E}X^{1-\alpha_j} \le M_j^-$ *for* $j = 1, \dots, m$*. For a fixed factor* F *with valid full profile,*

$$W_F^{[\ell,u]}(\gamma; M) = \max_{\nu \in \mathcal{M}} \int g_{\gamma,F}(x)\, d\nu(x).$$

The maximum is attained, and some optimizer has at most $2m + 3$ *atoms. If one law satisfies every moment inequality strictly, then the value equals the attained minimum*

$$\lambda_0 + \lambda_1 + \sum_{j=1}^{m} u_j M_j^+ + \sum_{j=1}^{m} v_j M_j^-$$

over $\lambda_0, \lambda_1 \in \mathbb{R}$ and $u_j, v_j \geq 0$, subject to $q(x) \geq g_{\gamma,F}(x)$ for every $x \in [\ell, u]$, where

$$q(x) = \lambda_0 + \lambda_1 x + \sum_{j=1}^{m} u_j x^{\alpha_j} + \sum_{j=1}^{m} v_j x^{1-\alpha_j}.$$

A checked finite-atomic primal law supplies a lower witness. Integrating $q \geq g_{\gamma,F}$ gives the dual upper bound: mass and mean one fix the affine terms, and the moment inequalities control the remaining terms because their multipliers are nonnegative. Every checked feasible majorant supplies a valid upper bound for admission. Strict feasibility additionally makes the primal and dual optima agree. These are classical generalized-moment tools (Winkler, 1988; Anderson & Nash, 1987), applied to an objective that retains the request's full profile.

To test admission with a low-cost upper bound, use convexity: the secant through the endpoint values of $g_{\gamma,F}$ lies above it on $[\ell, u]$. Taking expectations under a mean-one law yields

$$W_F^{[\ell,u]}(\gamma; M) \leq \frac{u-1}{u-\ell} g_{\gamma,F}(\ell) + \frac{1-\ell}{u-\ell} g_{\gamma,F}(u).$$

This affine choice of q has zero moment multipliers and needs only two outward endpoint evaluations. If the checked bound meets the cap, this direction is certified. Otherwise, optimizing the moment multipliers uses the rest of the receipt to seek a tighter bound.

5.2 From the moment problem to a checked decision

The support used for computation must follow from the input. Valid support bounds $[\ell, u] \Subset (0, \infty)$ give the compact problem directly. For an unbounded class, likelihood ratios outside a finite interval must also be covered. Positive and negative tail moments $\mathbb{E}X^\beta \leq B_+$ and $\mathbb{E}X^{-r} \leq B_-$, with $\beta > 1$ and $r > 0$, control this contribution. Saved orders can supply them through $\beta = \alpha_j$ and $r = \alpha_k - 1$; bounds at other exponents are additional historical information. We distinguish these tail constraints from the core orders used in the moment program.

Clipping an unbounded likelihood ratio and restoring its mean to one produces a compact law. The transformed support and relaxed moment bounds contain every such law, and a transfer margin covers the change in its composed profile.

The upper bound must hold between grid nodes as well as at them. On the transformed interval, the constructor proposes an endpoint secant or solves a nodewise linear program and rationalizes the dual coefficients. The checker certifies continuous majorization using a derivative bound and the objective's Lipschitz property. Alternatively, a tangent to the convex dual lies below it and a chord of the convex objective lies above it; a shift covering their difference certifies each cell. Request-profile enclosures and arithmetic errors enter outward.

The next theorem collects the allowances needed for a one-sided upper bound. Because the compact class is an outer relaxation, a lower-bound law must still be checked against all original receipt, support, and tail constraints.

Theorem 4 (One-sided admission bounds)**.** *Assume positive and negative tail moments $\mathbb{E}X^\beta \leq B_+$ and $\mathbb{E}X^{-r} \leq B_-$ for some $\beta > 1$ and $r > 0$. For $0 < \ell < 1 < u_{\max}$, set*

$$\delta_{\ell,u_{\max}} = B_+ u_{\max}^{1-\beta} + B_- \ell^{r+1} < 1.$$

Clipping and mean-normalizing X produces a compact outer moment problem with relaxed receipt bounds and uniform factor-composition error at most $2\delta_{\ell,u_{\max}}$. The transformed support is

$$\left[\frac{\ell}{1+\delta_{\ell,u_{\max}}}, \frac{u_{\max}}{1-\delta_{\ell,u_{\max}}}\right].$$

A dual majorant with derivative bound B is certified between historical-grid nodes of width δ_x by a node margin $(B+1)\delta_x/2$, plus the width of the outward $g_{\gamma,F}$ enclosure. For receipt sensitivity, fix the compact support interval and let $\mathbf{u} = (u_j)_{j=1}^m$ and $\mathbf{v} = (v_j)_{j=1}^m$ be the nonnegative moment multipliers of a dual solution within ε_d of the compact optimum. If the compact moment-bound vectors expand componentwise by (Δ^+, Δ^-), the compact robust value changes by at most

$$\varepsilon_d + \mathbf{u}\cdot\Delta^+ + \mathbf{v}\cdot\Delta^-.$$

Thus the bound remains valid under the stated outward perturbation, with the displayed change in compact robust value.

The compact problem's checked upper value, computed with relaxed moments on the transformed interval, bounds the original robust value after adding $2\delta_{\ell,u_{\max}}$. Mean normalization can move clipped ratios outside the raw interval $[\ell, u_{\max}]$; the expanded interval preserves coverage. The supplement gives the moment relaxations and full-profile brackets.

When a bracket remains inconclusive, the error terms distinguish numerical refinement from better historical information. A wider clipping interval, finer grid, or more accurate evaluation improves numerical control for the same receipt. Tighter saved bounds reduce the uncertainty being optimized. At a fixed compact interval, the dual multipliers quantify the effect of outward rounding at each order. An affine majorant has zero moment multipliers, so its bound depends on the endpoints and, for an unbounded history, the tail reduction. Every refined bound is checked again before use.

5.3 Direct optimization and profile reuse

When several fixed candidates share an unchanged receipt, a historical profile can be prepared once and composed with each request. This approach preserves the request while bounding the history before seeing its objective. Write $H(t) = \sup_{\nu \in \mathcal{U}(M)} h_\nu(t)$ for the largest compatible historical profile value at each threshold. With request singular mass s_F and $Y = dP_F^{\mathrm{ac}}/dQ_F$, composition gives

$$W_F(\gamma; M) = s_F + \sup_{\nu \in \mathcal{U}(M)} \mathbb{E}_{Q_F}[Y h_\nu(\gamma/Y)] \le s_F + \mathbb{E}_{Q_F}[Y H(\gamma/Y)],$$

with zero contribution when $Y = 0$. Direct optimization chooses one compatible history for the whole objective; the envelope permits different worst histories at different thresholds. A valid dominating profile can therefore represent a law that does not satisfy the receipt moments. The supplement constructs a checked upper curve for the compact outer class and adds the same tail margin as the direct calculation. Both methods retain the request: the computational choice is between tailoring a bound to the current objective and amortizing preparation across candidates sharing an unchanged receipt.

6 Admission Decisions and Computation

We evaluate the two comparisons developed above. The first tests whether preserving the specified request recovers admissions that multiplied moments cannot justify. The second holds the request profile fixed and compares the admission bounds and costs of direct optimization and reuse. Smaller checked upper values permit admission at tighter caps; a compatible law above the cap rules out the guarantee for its class. The workloads are controlled constructions. Exact history serves as a stronger-information reference, while receipt-based upper bounds cover every history allowed by their stated inputs.

6.1 What preserving the request recovers

The study draws 512 times with replacement from eight binary history deviations, seven RR request strengths, and seven caps, all at $\gamma = 5$; these cover 290 distinct triples. The generating history puts equal mass at $1 - \rho$ and $1 + \rho$. Both compressed-history methods receive only the shared bound $M = 1/(1 - \rho^2)$ on its positive and negative order-two moments; the historical mechanism and support are withheld. Only the exact-history reference uses the generating law. Each proposed request is supplied as the specified RR law.

All menu requests satisfy $1 < a < \gamma = 5$. A global quadratic majorant gives the factor-retained upper bound $c(a, \gamma)(M - 1)$ in both directions, where

$$c(a, \gamma) = \frac{a^2}{4(a+1)(\gamma - a)} + \frac{1}{4(a+1)(\gamma a - 1)}.$$

It improves on the square-root bound in Equation (1) throughout this menu. Both bounds use only the receipt and request, without numerical optimization or historical support. The joint-moment method

multiplies the same bounds by the request moments, then checks aggregate upper bounds and feasible laws above the cap. The supplement gives the majorant proof, parameter menus, and input-wise comparisons.

Table 1: Admission from a retained RR profile versus multiplied moments on 512 draws (290 distinct inputs). Joint-moment rejection is certified by an aggregate law above the cap. Only the exact-history reference receives the full generating law. Unique counts each observed triple once.

Factor retained	Joint moment	Exact-history reference	Draws	Unique
Accept	Inconclusive	Accept	164	89
Accept	Reject	Accept	329	191
Inconclusive	Reject	Accept	14	8
Inconclusive	Reject	Reject	5	2

Retaining the request permits 493 admissions. In 329, the aggregate class contains a law above the cap, establishing a strict information separation: even an exact optimizer for that class cannot justify admission. In the other 164, the implemented joint-moment bracket is inconclusive, so the result establishes a difference between computed bounds.

Counting each observed triple once gives 280 admissions among 290 inputs, with the same distinction shown in Table 1. The square-root bound admits 236 draws; the quadratic retains all of them and adds 257 using exactly the same information. Of the 19 cases still inconclusive under factor retention, exact history accepts 14 and rejects five. These brackets combine historical uncertainty with slack in the analytic upper bound and do not isolate the two effects.

A scheduled follow-up tests the same distinction after a shared update: both compressed representations admit a fixed Gaussian request and apply the same moment update, after which factor retention admits an RR request that an aggregate law rules out. The supplement's sequential-decisions section gives this setup, its sensitivity map, and a separate sequence where retaining the request recovers no admission.

6.2 Direct optimization and a reusable historical profile

We now compare the two request-preserving constructions in Section 5.3. The reusable method bounds the historical profile at 65 thresholds using the compact moment problem. The lower convex hull of the checked bounds dominates every compatible compact history; its slope changes give a finite mean-one law to compose with each request. Both methods add the same tail margin.

Both receive core orders 2–4 or 2–8, with tail exponents $(\beta, r) = (13, 11)$ in both panels. The core bounds are the moments of the equal-weight law on $\{3/4, 5/4\}$ expanded outward by 10^{-3}; the tail bounds use the corresponding moments without that expansion. Historical support is withheld, and both methods use raw clipping interval $[1/2, 2]$. Thus the panels vary core orders while holding tail information fixed. A feasible two-atom law supplies the common lower bound.

The menu contains RR strengths $a \in \{2, 3, 4\}$ and GDP strengths $\mu \in \{1/4, 1/2, 3/4\}$ at $\gamma \in \{1, 5\}$, and a supplied finite profile at $\gamma \in \{1, 4, 8\}$. These 15 request–threshold pairs are fixed candidates for one unchanged historical receipt. Each bracket is compared with all caps $\delta_B = j/100$, $j = 0, \ldots, 100$, giving 3030 paired directional cap decisions across the two panels. The counts measure budget locations at which a bound permits admission, rather than independent release trials. Both upper constructions use 256 historical cells and the continuous checks in Section 5.2, with an endpoint shortcut when the objective is affine. The supplement specifies the finite profile enclosure and reports every bracket.

Direct optimization admits 2432 cap cells and the reusable profile 2427. All five additional cells are GDP cases where reuse is inconclusive. For example, with core orders 2–4, the GDP request $\mu = 1/2$ at $\gamma = 1$ has common lower value 0.2215, direct upper value 0.2333, and reusable upper value 0.2428. At cap 0.24, the direct bound admits the request in this direction while the reusable bracket straddles the cap. These rounded values show the decision consequence of a tighter bound from identical information.

Table 2: Direct and reusable-profile decisions from identical inputs. Entries give accept, reject, and inconclusive counts over 101 directional caps per request–threshold pair: 606 cells per RR or GDP row and 303 per finite-profile row. Core orders vary; tail information is fixed.

Core orders	Request	Direct	Reusable profile
2–4	RR	448/144/14	448/144/14
2–4	GDP	527/73/6	524/73/9
2–4	Finite	240/13/50	240/13/50
2–8	RR	449/144/13	449/144/13
2–8	GDP	527/73/6	525/73/8
2–8	Finite	241/13/49	241/13/49

Both methods make the same RR and finite-profile decisions. Additional core orders recover two direct admissions and three reusable-profile admissions; doubling both grids leaves all paired decisions unchanged. The direct-only cells improve on an implemented reusable bound. In contrast, the first study's feasible aggregate laws establish an information loss that persists under exact aggregate optimization.

Table 3 separates preparation from subsequent queries. Prepared GDP queries are cheaper than fresh direct bounds, while the affine RR example favors direct endpoint checks. Including preparation, direct computation is faster for both complete 15-pair menus. The request objective and the number of candidates sharing a curve therefore determine the useful cost comparison.

Table 3: Measured median accounting costs with a supplied lower candidate. Direct calls construct and check a fresh bound; profile queries use a prepared curve. Complete passes cover all 15 request–threshold pairs and include preparation once for reuse.

Timed operation	Orders 2–4	Orders 2–8
Direct RR call, $a = 3$, $\gamma = 1$ (ms)	1.07	1.92
Reusable-profile query, same RR (ms)	3.81	5.99
Direct GDP call, $\mu = 1/2$, $\gamma = 1$ (ms)	407	539
Reusable-profile query, same GDP (ms)	59.8	75.6
Historical-profile preparation (s)	3.12	6.03
Complete direct 15-pair pass (s)	3.44	4.84
Complete reusable-profile 15-pair pass (s)	3.81	6.59

Measurements use one warmup and three repetitions. Direct calls include receipt preparation, bound construction and independent checking, profile evaluation, feasibility of the supplied lower candidate, and a directional cap comparison; candidate search is excluded. The supplement reports the hardware, software, ranges, and ordinary RDP update costs. An ordinary update multiplies moments without deciding the cap. Reuse applies to an unchanged receipt: after an executed release updates it, the curve requires revalidation or reconstruction.

7 Conclusion

Retaining an independently fixed request's profile preserves the product constraint that multiplied moments discard. The RR and GDP separations show that this changes safe admission even under exact aggregate optimization. Moment optimization turns the retained information into computable bounds with explicit tail and numerical margins. Direct bounds tailor computation to one request; historical-profile reuse shares preparation across candidates. Direct checks are faster for the tested complete menus, while prepared queries reduce repeated Gaussian accounting costs.

Future work can use receipt sensitivities to choose which orders and precisions to retain for anticipated requests, allocating more precision to moments that have a larger effect on the admission bound. Another direction is to update historical profiles efficiently as releases are executed. For transcript-dependent requests, the supplement's common-dominator bound provides a starting point under its measurability and garbling assumptions.

Reproducibility. The supplement gives complete proofs and experimental designs. The accompanying source README documents reproduction of the tables and timing measurements.

STATEMENT ON THE USE OF AI

Generative AI tools were used to revise wording and improve the clarity and readability of the manuscript; to identify potentially relevant literature, references, and search directions, and to assist with research execution, including drafting, debugging, and refining code. For mathematical proofs, AI was used primarily to review and check the authors' arguments for potential errors, inconsistencies, or gaps. AI-assisted manuscript revisions, literature-search suggestions, feedback on research ideas and proofs, and generated or modified code were reviewed and verified by the authors before incorporation into the final work. The authors independently assessed the accuracy and appropriateness of these outputs and retain full responsibility for the methodology, proofs, implementation, analysis, and conclusions presented in the paper.

REFERENCES

Martín 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*, pp. 308–318. ACM, 2016. doi: 10.1145/2976749.2978318.

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

Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography*, volume 9985 of *Lecture Notes in Computer Science*, pp. 635–658. Springer, 2016. doi: 10.1007/978-3-662-53641-4_24.

Cemre Cadir, Salim Najib, and Yanina Y. Shkel. Composition theorems for multiple differential privacy constraints, 2026. arXiv:2603.20968v2.

Jinshuo Dong, Aaron Roth, and Weijie J. Su. Gaussian differential privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):3–37, 2022. doi: 10.1111/rssb.12454.

Vadym Doroshenko, Badih Ghazi, Pritish Kamath, Ravi Kumar, and Pasin Manurangsi. Connect the dots: Tighter discrete approximations of privacy loss distributions. *Proceedings on Privacy Enhancing Technologies*, 2022(4):552–570, 2022. doi: 10.56553/popets-2022-0122.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014. doi: 10.1561/0400000042.

Sivakanth Gopi, Yin Tat Lee, and Lukas Wutschitz. Numerical composition of differential privacy. In *Advances in Neural Information Processing Systems*, volume 34, pp. 11631–11642, 2021.

Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pp. 1376–1385. PMLR, 2015.

Samuel Karlin and William J. Studden. *Tchebycheff Systems: With Applications in Analysis and Statistics*. Interscience, 1966.

Antti Koskela, Joonas Jälkö, and Antti Honkela. Computing tight differential privacy guarantees using FFT. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108, pp. 2560–2569. PMLR, 2020.

Antti Koskela, Rachel Emily Redberg, and Yu-Xiang Wang. Privacy profiles for private selection. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pp. 25313–25332. PMLR, 2024.

Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium*, pp. 263–275. IEEE, 2017. doi: 10.1109/CSF.2017.11.

Jack Murtagh and Salil Vadhan. The complexity of computing the optimal composition of differential privacy. *Theory of Computing*, 14(8):1–35, 2018. doi: 10.4086/toc.2018.v014a008.

Matthew Regehr, Bingshan Hu, Ethan Leeman, Pasin Manurangsi, Pierre Tholoniat, and Mathias Lécuyer. Privacy filters are captured by residues: A characterization of free natural filters and the cost of adaptivity, 2026. arXiv:2602.15815v2.

Anneliese Riess, Juan Felipe Gomez, Flavio du Pin Calmon, Julia Anne Schnabel, and Georgios Kaissis. Optimal conversion from Rényi differential privacy to f-differential privacy, 2026. arXiv:2602.04562v3.

Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. Subsampled Rényi differential privacy and analytical moments accountant. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89, pp. 1226–1235. PMLR, 2019.

Gerhard Winkler. Extreme points of moment sets. *Mathematics of Operations Research*, 13(4): 581–587, 1988. doi: 10.1287/moor.13.4.581.

Yuqing Zhu, Jinshuo Dong, and Yu-Xiang Wang. Optimal accounting of differential privacy via characteristic function. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, volume 151, pp. 4782–4817. PMLR, 2022.

A NOTATION AND RESULT LOCATIONS

The historical receipt contains moment upper bounds and any supplied support or tail constraints. The proposed request is fixed independently of historical private outputs and uses fresh randomness.

For a binary experiment $E = (P, Q)$ with $P \ll Q$, let $X = dP/dQ$ under Q, so $\mathbb{E}X = 1$. Its hockey-stick profile is

$$h_E(\gamma) = \mathbb{E}(X - \gamma)_+, \qquad \gamma \geq 0,$$

with an additive singular P-mass when needed. A bidirectional receipt at distinct orders $\alpha_j > 1$ stores

$$\mathbb{E}X^{\alpha_j} \leq M_j^+, \qquad \mathbb{E}X^{1-\alpha_j} \leq M_j^-, \qquad j = 1, \ldots, m.$$

The $+$ and $-$ labels refer to exponent signs within one ordered direction; admission requires a verified bound in both directions.

For receipt M, let $\mathcal{U}(M)$ contain all positive, mean-one likelihood-ratio laws satisfying these constraints. The fixed-request objective is $g_{\gamma,F}(x) = xh_F(\gamma/x)$, and $W_F(\gamma; M) = \sup_{\nu \in \mathcal{U}(M)} \mathbb{E}_\nu g_{\gamma,F}(X)$. The aggregate value $C_{\mathrm{agg}}(\gamma; M)$ maximizes the hockey-stick value over laws satisfying the multiplied historical and request moments. Table 4 locates the results supporting the main paper.

Table 4: Main-paper results and their supplement locations.

Main-paper item	Result proved here
Fixed-factor moment optimization	Theorem 12
One-sided admission bounds	Continuous margin, Proposition 17; compact reduction, Theorem 18; discretized and rounding bounds, Theorems 19–20
Strict separations	RR and GDP separations, Theorems 13 and 14
Exact RR instance	Example in Section D.3
Information retained by finite receipts	Finite-order-grid impossibility, finite-atomic identification, and bounded approximation, Theorems 5, 8, and 9

Figure 1 illustrates the Gaussian admission gap; its formulas and proof appear in Section D.3.

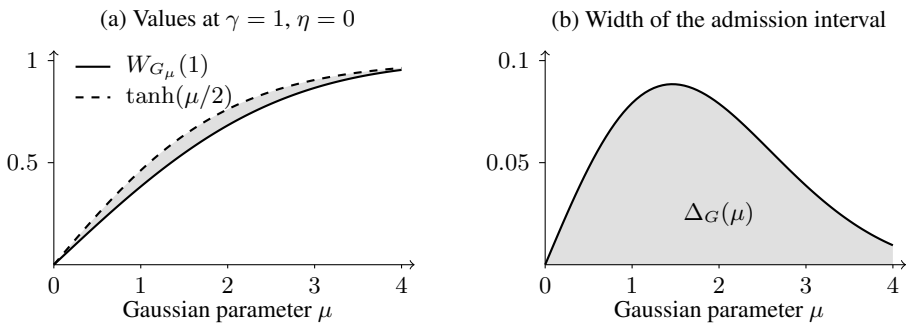

Figure 1: Gaussian separation at $\gamma = 1$ and $\eta = 0$ (Theorem 14). The history is $X = 1$ almost surely, so $W_{G_\mu}(1) = 2\Phi(\mu/2) - 1$. The dashed curve is the value $\tanh(\mu/2)$ of a feasible aggregate law. Caps in the shaded interval include its lower boundary and exclude its upper boundary: they admit the fixed Gaussian request but are exceeded by the aggregate law. The right panel shows the interval width $\Delta_G(\mu) > 0$ for $0 < \mu \leq 4$.

B COMPUTATION AND ADDITIONAL EXPERIMENTS

B.1 FROM A RECEIPT TO A DIRECTIONAL DECISION

Let $\mathcal{I}$ be a finite collection of directional inputs whose bounds cover every ordered neighboring-data pair. Each historical class is nonempty and comes with valid support or tail bounds. For a compact dual q with coefficients $(\lambda_0, \lambda_1, \mathbf{u}, \mathbf{v})$, write

$$D(q; \widetilde{M}) = \lambda_0 + \lambda_1 + \mathbf{u} \cdot \widetilde{M}^+ + \mathbf{v} \cdot \widetilde{M}^-.$$

In Algorithm 1, D^+ is an outward upper evaluation of this objective. CHECKMAJORANT checks nonnegative moment multipliers and continuous majorization using Proposition 17 or Lemma 21 with outward profile evaluations. FEASIBLE checks nonnegative atom masses, positive support, total mass and mean one, and every original receipt, support, and tail constraint. Failed or unavailable proposals leave the stored bounds unchanged. Bounds remain valid when refinement changes the clipping interval.

Algorithm 1 Checked admission of a fixed request from moment receipts

Require: Receipts M_d and profile enclosures for F_d, $d \in \mathcal{I}$; threshold γ_B, cap δ_B, computation budget
$(L_d, U_d, \nu_d^\star) \leftarrow (0, 1, \varnothing)$ for each $d \in \mathcal{I}$
while computation budget remains **do**
  for each d with $U_d > \delta_B$ **do**
    if valid support $[\ell_d, u_d]$ is supplied **then**
      $(I_d, \widetilde{M}_d, e_d) \leftarrow ([\ell_d, u_d], M_d, 0)$
    else
      Choose $0 < \ell < 1 < u$ with $\delta = B_+ u^{1-\beta} + B_- \ell^{r+1} < 1$
      $I_d \leftarrow [\ell/(1+\delta), u/(1-\delta)]$; $\widetilde{M}_d \leftarrow$ relaxed bounds in Theorem 18
      $e_d \leftarrow 2\delta$
    end if
    $q \leftarrow \mathrm{ProposeMajorant}(g_{\gamma_B, F_d}, I_d, \widetilde{M}_d)$
    if $\mathrm{CheckMajorant}(q, g_{\gamma_B, F_d}, I_d)$ **then**
      $U_d \leftarrow \min\{U_d, D^+(q; \widetilde{M}_d) + e_d\}$
    end if
    $\nu \leftarrow \mathrm{ProposeLaw}(M_d)$
    if $\mathrm{Feasible}(\nu; M_d)$ **then**
      $l \leftarrow \mathbb{E}_\nu g_{\gamma_B, F_d}(X)$
      if $l > L_d$ **then**
        $(L_d, \nu_d^\star) \leftarrow (l, \nu)$
      end if
    end if
    if $L_d > \delta_B$ **then**
      return (reject, $d, \nu_d^\star, L_d$)
    end if
  end for
  if $\max_{d \in \mathcal{I}} U_d \le \delta_B$ **then**
    return (admit, $(U_d)_{d \in \mathcal{I}}$)
  end if
  Refine clipping endpoints, grid, profile precision, or candidates
end while
return (inconclusive, $([L_d, U_d])_{d \in \mathcal{I}}$)

A reusable historical profile follows the same construction with hinge objectives at several historical thresholds; Section D.5 proves how the resulting curve is composed with subsequent requests.

B.2 REQUEST PROFILES

In the tables below, *Finite* denotes the fixed three-atom request whose Q-side likelihood ratio satisfies

$$\Pr(Y = 0) = \frac{1}{4}, \qquad \Pr(Y = 1) = \frac{5}{8}, \qquad \Pr(Y = 3) = \frac{1}{8}.$$

Its exact profile is $h_F(t) = \frac{5}{8}(1-t)_+ + \frac{1}{8}(3-t)_+$. For the supplied-profile comparisons, the accountant receives upper values at thresholds $0, 2, 3$ and the exact zero tail beyond 3. Chord

interpolation gives the conservative curve

$$\overline{h}_F(t) = \frac{5}{16}(2-t)_+ + \frac{1}{8}(3-t)_+.$$

This is the profile of a dominating law with masses $9/16$, $5/16$, and $1/8$ at 0, 2, and 3. Both upper constructions use this curve. The common lower evaluation uses the universal bound $\underline{h}_F(t) = (1-t)_+$, so $x\underline{h}_F(\gamma/x) = (x-\gamma)_+$. For the feasible equal-weight history on $\{3/4, 5/4\}$, this gives lower value $1/8$ at $\gamma = 1$ and zero at $\gamma = 4, 8$, as reported in the matched-menu tables. The exact request and history give $21/64$ at $\gamma = 1$; that finer information is reserved for the exact-law reference. RR and GDP are defined in Section D.3, which also gives the exact RR admission example.

B.3 Controlled finite-receipt workload

The workload draws 512 triples independently with replacement from the menus in Table 5. The generating history has equal Q-mass at $1-\rho$ and $1+\rho$. Both compressed-history methods receive only $M_2^+ = M_2^- = 1/(1-\rho^2)$; the exact-history reference also receives the law and support. The request is the specified RR law.

Factor retention uses the global quadratic bound below in both directions, with $\sqrt{1/(1-\rho^2)-1}$ as an outward-evaluated comparison bound. The joint-moment calculation uses the multiplied order-two bounds and checks aggregate upper bounds and feasible two-atom lower laws. Historical directions, exact product profiles, and aggregate laws above the cap are checked with rational arithmetic. An aggregate law above the cap certifies an information separation; an inconclusive aggregate bracket establishes only a difference between computed bounds.

Table 5: Controlled finite-receipt sampling design. Each draw selects one entry independently and uniformly with replacement from each ordered menu. All decimal values are exact.

Parameter	Values or specification
History deviation ρ	$0.002,\ 0.004,\ 0.008,\ 0.0125,$ $0.02,\ 0.03125,\ 0.05,\ 0.08$
RR likelihood ratio a	$2,\ 2.5,\ 3,\ 3.5,\ 4,\ 4.5,\ 4.75$
Privacy cap δ	$0.0025,\ 0.005,\ 0.01,\ 0.015,\ 0.02,\ 0.03,\ 0.05$
Threshold; sample size	$\gamma = 5$; 512 draws
Generating history	Equal Q-mass at $1-\rho$ and $1+\rho$
Supplied receipt	$M_2^+ = M_2^- = 1/(1-\rho^2)$

The 512 draws contain 290 distinct parameter triples and 222 repeated occurrences. Table 6 groups the same draws by historical slack, request strength, and cap; each block totals 512. These counts describe the fixed sampling design.

Table 6: Bound comparison on the same 512 RR draws. Square-root and quadratic columns use the same receipt; strict gap counts quadratic admissions with an aggregate distribution exceeding the cap. Each input block groups the same draws.

Input	Value	Draws	Square-root	Quadratic	Strict gap	Exact admit
History ρ	0.002	63	63	63	47	63
History ρ	0.004	56	45	56	34	56
History ρ	0.008	74	55	74	48	74
History ρ	0.0125	64	47	64	47	64
History ρ	0.02	65	17	65	43	65
History ρ	0.03125	59	9	59	44	59
History ρ	0.05	68	0	67	41	68
History ρ	0.08	63	0	45	25	58
Request a	2	88	42	88	0	88
Request a	2.5	69	35	69	0	69
Request a	3	67	21	67	60	67
Request a	3.5	71	34	70	70	71
Request a	4	58	32	55	55	58
Request a	4.5	84	38	80	80	84
Request a	4.75	75	34	64	64	70
Cap δ	0.0025	67	10	63	44	65
Cap δ	0.005	58	15	50	36	55
Cap δ	0.01	86	40	84	58	86
Cap δ	0.015	74	37	73	50	74
Cap δ	0.02	74	35	70	51	74
Cap δ	0.03	81	48	81	49	81
Cap δ	0.05	72	51	72	41	72

Among the 290 unique triples, the quadratic bound admits 280: 191 have an aggregate law above the cap and 89 have an inconclusive aggregate bracket. Of the ten remaining inputs, exact history accepts eight and rejects two.

The square-root bound admits 236 draws (133 unique inputs): 156 (91 unique) have an aggregate distribution exceeding the cap and 80 (42 unique) have an inconclusive aggregate bracket. Among its 276 inconclusive draws, 84 have an inconclusive aggregate bracket and exact-history admission; 187 have an aggregate distribution exceeding the cap and exact-history admission; five have an aggregate distribution exceeding the cap and exact-history rejection. Their unique counts are 47, 108, and two, respectively. The quadratic retains every square-root admission and adds 257 draws (147 unique inputs), including 173 (100 unique) with aggregate distributions exceeding the cap. Table 7 preserves the square-root counts by input.

Global quadratic bound. For $1 < a < \gamma$, the symmetric RR request has the composed objective

$$g_{\gamma,F_a}(x) = \frac{a}{a+1}(x-\gamma/a)_+ + \frac{1}{a+1}(x-\gamma a)_+.$$

For every $t > 1$ and $x > 0$, $(x-t)_+ \le (x-1)^2/[4(t-1)]$. When $x \le t$ the right side is nonnegative; when $x > t$ the difference equals $(x-(2t-1))^2/[4(t-1)]$. Applying this inequality to both hinges gives

$$g_{\gamma,F_a}(x) \le c(a,\gamma)(x-1)^2, \qquad c(a,\gamma) = \frac{a^2}{4(a+1)(\gamma-a)} + \frac{1}{4(a+1)(\gamma a-1)}.$$

Thus $\mathbb{E}X = 1$ and $\mathbb{E}X^2 \le M$ imply $W_{F_a}(\gamma; M) \le c(a,\gamma)(M-1)$ on the full positive half-line. The reverse history has likelihood ratio $1/X$ under P and second moment $\mathbb{E}_P X^{-2} = \mathbb{E}_Q X^{-1} \le M$; symmetric RR has the same request law after reversing direction. The bound therefore holds in both directions using only the saved receipt. Its dual coefficients are $(c, -2c, c, 0)$ for the constant, linear, positive second, and negative first powers. The independent checker reconstructs the three linear pieces of g from the request probabilities and minimizes the quadratic difference on each entire interval, including the unbounded tail. No clipping or numerical grid is needed. At $a = 3$ and $\gamma = 5$, $c = 2/7$: all 67 draws in this request group are admitted, compared with 21 under the square-root bound. The same global check covers all seven request strengths in Table 6.

Table 7: Square-root bound counts in the retained 512 draws. Each block groups the same draws by one input; values are displayed as decimals.

Input	Value	Draws	Admit	Strict gap	Exact admit
History ρ	0.002	63	63	47	63
History ρ	0.004	56	45	27	56
History ρ	0.008	74	55	34	74
History ρ	0.0125	64	47	36	64
History ρ	0.02	65	17	8	65
History ρ	0.03125	59	9	4	59
History ρ	0.05	68	0	0	68
History ρ	0.08	63	0	0	58
Request a	2	88	42	0	88
Request a	2.5	69	35	0	69
Request a	3	67	21	18	67
Request a	3.5	71	34	34	71
Request a	4	58	32	32	58
Request a	4.5	84	38	38	84
Request a	4.75	75	34	34	70
Cap δ	0.0025	67	10	8	65
Cap δ	0.005	58	15	12	55
Cap δ	0.01	86	40	24	86
Cap δ	0.015	74	37	27	74
Cap δ	0.02	74	35	25	74
Cap δ	0.03	81	48	30	81
Cap δ	0.05	72	51	30	72

Strict gap counts factor admissions with a joint-moment-compatible distribution exceeding the cap. Exact admit uses the complete historical law. The groups within each input sum to 512; the three input blocks are not additional samples.

B.4 Sequential decisions and sensitivity

The shared-update sequence starts from a binary history with deviation $\rho = 400/40001$. Both compressed representations admit a fixed Gaussian request with $\mu = 0.01$ at threshold $\gamma_B = 20$ and nominal cap $1/64$, then apply the same moment update. They next evaluate a scheduled RR request with $a = 4$ at threshold 7 and nominal cap 0.02. Factor retention admits this follow-up; a feasible aggregate law above the cap rules out joint-moment admission. Table 8 reports the executed-request counts and margins for the nominal setting and two cap perturbations.

Table 8: The shared GDP update followed by the RR decision at the nominal cap and two perturbations.

Cap case	GDP cap	RR cap	Accepted F/J	RR F/J	RR slack
Lower	$117/8000$	$19/1000$	$2/1$	yes/no	$6791/3200000$
Nominal	$1/64$	$1/50$	$2/1$	yes/no	$9991/3200000$
Upper	$133/8000$	$21/1000$	$2/1$	yes/no	$97/43000$

F/J denotes factor-retained versus joint-moment information. The cap columns correspond to the shared GDP request and RR follow-up, respectively; Accepted F/J gives the resulting request counts. Both apply the GDP update, while only F accepts the RR request. RR slack is its verified acceptance margin.

Table 9 varies the initial history deviation and both request strengths. Factor retention admits 35 of these settings and is inconclusive in ten; six of the latter have an aggregate law above the cap. Table 10 reports a separate five-position comparison: RR and GDP requests at threshold one surround a finite-profile request at threshold eight, with RR and GDP follow-ups. Both compressed representations reject the four RR/GDP requests using compatible laws above their caps and leave the finite request inconclusive. Exact history accepts that finite request. Thus this sequence records uncertainty left in the history even when the request profile is retained.

Table 9: Local decision map around the nominal setting $(\rho, \mu, a) = (400/40001, 1/100, 4)$. Rows vary history deviation ρ and shared GDP strength μ; columns vary RR follow-up strength a at fixed $(\gamma, \delta) = (7, 1/50)$. The remaining cells vary one or more coordinates around this setting.

ρ	μ	$7/2$	$15/4$	4	$17/4$	$9/2$
$1/140$	$1/200$	A/C	A/C	A/R	A/R	A/R
$1/140$	$1/100$	A/C	A/C	A/R	A/R	A/R
$1/140$	$3/200$	A/C	A/C	A/R	A/R	A/R
$400/40001$	$1/200$	A/C	A/C	A/R	A/R	A/R
$400/40001$	$1/100$	A/C	A/C	A/R	A/R	A/R
$400/40001$	$3/200$	A/C	A/C	A/R	A/R	A/R
$1/55$	$1/200$	A/C	A/C	A/R	A/R	A/R
$1/55$	$1/100$	C/C	C/C	C/R	C/R	C/R
$1/55$	$3/200$	C/C	C/C	C/R	C/R	C/R

Each entry is factor-retained/joint-moment: A is accept, R is reject with a distribution exceeding the cap in the corresponding compatibility class, and C is inconclusive. All 45 GDP decisions are A/A and apply the same signed update before the displayed RR decision.

Table 10: Decisions in the controlled five-request sequence.

Pos.	Request	δ_B	Factor retained	Joint moment	Exact history
1	RR, $\gamma = 1$	$1/100$	R	R	R
2	GDP, $\gamma = 1$	$1/100$	R	R	R
3	finite, $\gamma = 8$	$1/200$	C	C	A
4	RR, $\gamma = 1$	$1/100$	R	R	R
5	GDP follow-up, $\gamma = 1$	$1/100$	R	R	R

A: accept; R: reject with a distribution exceeding the cap in the corresponding compatibility class; C: inconclusive. The exact-history column uses the complete generating law for this sequence. The final row is the follow-up stage. Both compressed representations reject four requests and leave the finite request inconclusive.

B.5 Six-input bound and timing study

This study compares six complete receipt configurations and their computation costs. Both core orders and tail exponents change between panels; the matched-menu study in Section B.6 holds tail information fixed. The generating history has equal Q-mass at $3/4$ and $5/4$. The exact-history reference evaluates this law directly; the receipt-based upper bounds use only the moment and tail constraints below.

For the other methods, each saved order α supplies the bounds

$$M_\alpha^+ = \tfrac{1}{2}[(3/4)^\alpha + (5/4)^\alpha] + 10^{-3}, \qquad M_\alpha^- = \tfrac{1}{2}[(3/4)^{1-\alpha} + (5/4)^{1-\alpha}] + 10^{-3}.$$

The additive term is the receipt's outward expansion. Core orders 2–4 use tail exponents $(\beta, r) = (9, 7)$; orders 2–8 use $(13, 11)$. Their tail bounds are $B_+ = \frac{1}{2}[(3/4)^\beta + (5/4)^\beta]$ and $B_- = \frac{1}{2}[(3/4)^{-r} + (5/4)^{-r}]$. Both use raw clipping interval $[1/2, 2]$ before the mean-restoring transformation in Theorem 18. These tail constraints do not disclose the exact support to the receipt-based methods.

Request	Threshold γ	Cap, orders 2–4	Cap, orders 2–8
RR, $a = 3$	1	0.5107	0.5103
GDP, $\mu = 1/2$	1	0.2559	0.2559
Finite, three-atom profile above	8	0.0053	0.0006

The caps above are exact terminating decimals for the selected inputs. RR and the finite request use the affine secant through outward endpoint values of the objective on the transformed interval. Its moment multipliers are zero, so the tail constraints affect this bound through the transformed support and the clipping error. GDP uses 4096 equal subdivisions of that interval to optimize a dual with moment terms. Its checker evaluates all 4097 nodes and adds the continuous-domain margin. Profile evaluations use 160-bit precision with a 512-bit ceiling and outward rational endpoints. The nominal comparison tolerance is 10^{-4}; the reported intervals include the separate tail and numerical margins.

The reusable comparator here is a power-tail historical profile, distinct from the moment-based curve in the main matched-menu study. It joins a line starting at profile value one to a power-law tail determined by the highest saved positive moment. Compact endpoint bounds, the clipping error, and the original moment bound certify that this curve covers every receipt-compatible history. The construction then composes the experiment represented by this upper curve with the specified request, encloses the resulting integral, and bounds its profile by another curve of the same form. Both stages contribute to the reported upper bound. The power-tail construction is used in the six-input comparison and its associated timings.

Table 11 collects the admission outcomes for the exact and supplied-profile examples and six cases using receipts at orders 2–4 or 2–8.

Table 11: Admission outcomes for three exact or supplied-profile demonstrations and six certified cases using receipts at core orders 2–4 or 2–8.

Request and receipt orders	Factor retained	Joint moment	Exact history
Exact and supplied-profile demonstrations			
RR, exact example	A	R	–
GDP, supplied profile	A	–	–
Finite, supplied profile	A	–	–
Certified finite-receipt cases			
Finite, orders 2–4	C	C†	A
GDP, orders 2–4	C	R	A
RR, orders 2–4	C	R	A
Finite, orders 2–8	C	C†	A
GDP, orders 2–8	A	R	A
RR, orders 2–8	A	R	A

A: accept; C: inconclusive; R: reject with a distribution exceeding the cap in the corresponding compatibility class; C†: inconclusive because the reverse request moment is missing; –: not applicable. The exact-history column is a stronger-information reference.

The repeated measurement study uses RR, GDP, and a finite-output factor at orders 2–4 and 2–8. Each method has three warmups and 50 measurements with `time.perf_counter_ns` on Linux, Python 3.12.13, and an AMD Ryzen Threadripper 3990X. Table 12 isolates the cost of checking evidence that has already been constructed. Reported bytes measure the serialized evidence and decision, and the Factor calls column counts top-level request-profile evaluations. P95 is the 95th percentile of the measured times; one KiB is 1024 bytes. Ordinary signed-moment multiplication provides an update-cost reference.

Table 13 reports one construction and checking pass for all four methods together. Checking runs the independent verifier twice, the second time to validate its returned record; serialization to disk lies between the timed stages. These totals differ from per-method checks of preconstructed evidence and from the isolated fresh-construction timings below.

Isolated construction and checking times. The isolated study uses two warmups and ten measurements per operation on the same six inputs. Construction transforms the receipt, builds only the direct secant or moment dual, and encloses a supplied lower candidate's objective. Checking independently reconstructs the tail reduction, verifies continuous majorization, and reevaluates the objective endpoints. The ordinary update evaluates and multiplies request moments at the core orders; it reports unavailable reverse moments for the finite request. All evaluations within these operations are timed.

Input loading, receipt preparation, and lower-candidate feasibility occur before these timings. The equal-weight law on $\{3/4, 5/4\}$ is checked exactly for positive support, mass and mean one, every core moment bound, and both tail bounds. Each repetition constructs a fresh upper bound; the same feasible candidate supplies its lower endpoint.

Measurements use CPython 3.12.13, NumPy 2.5.2, SciPy 1.18.0, and gmpy2 2.3.1 in a Conda environment on the same Threadripper 3990X machine. The clock is `time.perf_counter_ns`; thread settings are inherited. Table 14 gives medians and observed ranges, including the larger wall-clock variation of GDP grid calculations.

Table 12: Bound-check and RDP-update measurements for all 18 configurations. Latencies are milliseconds over 50 measurements. Each timing covers the displayed computation; Record KiB is the bound-and-decision serialization.

Core orders	Request	Operation	State	Min	Median	P95	Max	Record KiB	Factor calls
2–4	RR	Factor upper check	C	0.697	0.716	0.739	0.751	1.4	4
2–4	RR	Joint rejection check	R	0.393	0.397	0.404	0.408	1.3	0
2–4	RR	RDP update	–	0.070	0.071	0.074	0.078	0.2	0
2–4	GDP	Factor upper check	C	2392.2	2420.7	2450.2	2468.5	7.4	4099
2–4	GDP	Joint rejection check	R	0.500	0.504	0.513	0.531	4.2	0
2–4	GDP	RDP update	–	0.093	0.095	0.096	0.104	1.1	0
2–4	Finite	Factor upper check	C	1.007	1.019	1.033	1.053	1.4	4
2–4	Finite	Joint rejection check	C†	0.263	0.265	0.273	0.296	0.9	0
2–4	Finite	RDP update	–	0.176	0.179	0.186	0.189	0.2	0
2–8	RR	Factor upper check	A	1.153	1.165	1.179	1.191	1.6	4
2–8	RR	Joint rejection check	R	0.681	0.687	0.697	0.706	1.8	0
2–8	RR	RDP update	–	0.163	0.164	0.168	0.180	0.4	0
2–8	GDP	Factor upper check	A	2700.0	2728.6	2762.3	2774.5	16.9	4099
2–8	GDP	Joint rejection check	R	0.896	0.899	0.911	0.918	8.1	0
2–8	GDP	RDP update	–	0.213	0.215	0.219	0.226	2.4	0
2–8	Finite	Factor upper check	C	1.487	1.504	1.526	1.541	1.6	4
2–8	Finite	Joint rejection check	C†	0.389	0.394	0.407	0.410	1.1	0
2–8	Finite	RDP update	–	0.241	0.243	0.248	0.262	0.2	0

A/R/C denote accept, rejection with a distribution exceeding the cap in the corresponding compatibility class, and inconclusive. C† marks a missing reverse request moment; – marks fields not produced by the RDP update. For factor and joint rows, Record KiB counts the serialized bound-and-decision record; Factor calls counts top-level profile evaluations. GDP uses a grid of 4096 subintervals, giving 4097 profile evaluations, plus evaluations at the two fixed exact-history atoms.

Table 13: Decisions and recorded construction/checking cost for the same six inputs. Times cover all four methods and validation of their verification outputs.

Request	Core orders	F/J/P/H	Build ms	Check ms	KiB
RR	2–4	?/R/?/A	6.7	65.9	40.0
RR	2–8	A/R/?/A	11.0	85.8	53.0
GDP	2–4	?/R/?/A	3045.0	5769.5	66.9
GDP	2–8	A/R/?/A	3342.8	6357.5	102.0
Finite	2–4	?/?†/?/A	7.3	64.1	38.2
Finite	2–8	?/?†/?/A	8.8	85.6	51.7

F: factor retained; J: joint moments; P: power-tail profile; H: exact-history reference. A: admit; R: distribution exceeding the cap; ?: inconclusive. †: a required reverse request moment is infinite. One recorded build and check per input; the checking stage repeats verification when validating its output.

Table 14: Isolated timing medians and observed ranges in milliseconds, over ten measurements after two warmups.

Request	Core orders	Construct	Check	Update
RR	2–4	0.741 [0.721, 0.754]	1.12 [1.1, 1.16]	0.11 [0.108, 0.117]
RR	2–8	0.873 [0.858, 0.997]	1.76 [1.16, 1.8]	0.251 [0.246, 0.254]
GDP	2–4	3060 [2490, 3710]	2690 [2460, 3710]	0.101 [0.0989, 0.106]
GDP	2–8	2940 [2810, 4080]	3790 [2840, 4410]	0.344 [0.341, 0.361]
Finite	2–4	0.745 [0.733, 0.763]	1.04 [1.03, 1.07]	0.182 [0.18, 0.193]
Finite	2–8	1.77 [1.76, 1.79]	2.34 [2.25, 2.34]	0.392 [0.386, 0.409]

Complete accounting time. Complete accounting time covers one continuous pass from an in-memory receipt and supplied two-atom candidate to a directional cap decision. It includes separate receipt preparation for construction and verification, exact candidate-feasibility checking, fresh upper-bound construction, independent verification, and rational cap comparison. Request-profile loading and evaluation within these operations are included. File input/output, candidate search, and execution of a private release are excluded. Table 15 reports directly measured totals for the same six inputs and software environment, with two warmups and ten measurements in a warm process.

Table 15: Complete accounting time for one directional cap decision with a supplied lower candidate on each of the six inputs.

Request	Core orders	Decision	Median ms	Min ms	Max ms
RR	2–4	?	2.18	2.15	2.23
GDP	2–4	?	5500	4900	7100
Finite	2–4	?	3.02	2.84	3.87
RR	2–8	A	2.96	2.86	3.43
GDP	2–8	A	5600	5500	8300
Finite	2–8	?	4.34	4.3	4.41

Median and observed range over ten measured repetitions after two warmups. Receipt preparation, candidate feasibility, upper-bound construction, independent verification, and the cap comparison are timed together. A: admit; ?: inconclusive.

B.6 Matched request menu with fixed tail information

The comparison uses core orders 2–4 and 2–8 with identical tail information. At every retained core order the historical bound is the corresponding moment of the equal-weight law on $\{3/4, 5/4\}$ plus 10^{-3}. Both panels retain the same bounds on $\mathbb{E}X^{13}$ and $\mathbb{E}X^{-11}$, and raw clipping interval $[1/2, 2]$. Theorem 18 expands the core bounds and rescales the interval; its common transfer margin $2\delta_{\ell,u}$ is approximately 0.01025. The generating law is used only for a feasible lower value. Upper bounds range over every history satisfying these receipt and tail bounds.

The moment-based profile construction uses 256 equal cells on the rescaled interval. Profile preparation uses 65 equally spaced thresholds, including both endpoints; the interior 63 thresholds each receive an independently checked dual. At an affine RR or finite-profile objective, the direct implementation uses an endpoint bound after checking that no objective kink lies in the interval interior. The finite request is the supplied upper profile from the earlier comparison, at the three thresholds reported below. Its positive moments come from the profile's dominating law with masses $9/16$, $5/16$, and $1/8$ at likelihood ratios 0, 2, and 3. Its reverse moments are infinite, which is retained in the ordinary-update reference.

The 15 requests are fixed candidates for one historical ledger. The 101 caps per request measure the budget locations of changes in directional admission; all 3030 request/receipt/cap cells are retained. Tables 16 and 17 report every bracket and separate distributions exceeding the cap from inconclusive outcomes. The five direct-only admissions occur at cap 0.16 for $\mu = 1/4$ and at 0.32 for $\mu = 3/4$ in both panels, and at 0.24 for $\mu = 1/2$ in the 2–4 panel, all at $\gamma = 1$. The refined construction uses 512 historical cells and 129 profile thresholds. Independently checking all refined bounds leaves every cap classification unchanged.

Table 16: All directional request brackets at core orders 2–4, with tail exponents $(13, 11)$. The lower value is common to both methods; U_D and U_P are the direct and moment-based profile upper bounds. Counts use exact bounds; displayed values are rounded.

Request	Strength	γ	Lower	U_D	U_P	D: A/R/?	P: A/R/?
RR	2	1	0.3333	0.3438	0.3438	66/34/1	66/34/1
RR	2	5	0.0000	0.0102	0.0102	99/0/2	99/0/2
RR	3	1	0.5000	0.5102	0.5102	49/50/2	49/50/2
RR	3	5	0.0000	0.0194	0.0194	99/0/2	99/0/2
RR	4	1	0.6000	0.6102	0.6102	39/60/2	39/60/2
RR	4	5	0.0000	0.0450	0.0450	96/0/5	96/0/5
GDP	$1/4$	1	0.1458	0.1588	0.1695	85/15/1	84/15/2
GDP	$1/4$	5	0.0000	0.0102	0.0102	99/1/1	99/1/1
GDP	$1/2$	1	0.2215	0.2333	0.2428	77/23/1	76/23/2
GDP	$1/2$	5	0.0005	0.0111	0.0115	99/1/1	99/1/1
GDP	$3/4$	1	0.3080	0.3192	0.3262	69/31/1	68/31/2
GDP	$3/4$	5	0.0125	0.0230	0.0255	98/2/1	98/2/1
Finite	–	1	0.1250	0.5729	0.5729	43/13/45	43/13/45
Finite	–	4	0.0000	0.0222	0.0224	98/0/3	98/0/3
Finite	–	8	0.0000	0.0102	0.0102	99/0/2	99/0/2

Table 17: All directional request brackets at core orders 2–8, with tail exponents $(13, 11)$. The lower value is common to both methods; U_D and U_P are the direct and moment-based profile upper bounds. Counts use exact bounds; displayed values are rounded.

Request	Strength	γ	Lower	U_D	U_P	D: A/R/?	P: A/R/?
RR	2	1	0.3333	0.3436	0.3437	66/34/1	66/34/1
RR	2	5	0.0000	0.0102	0.0102	99/0/2	99/0/2
RR	3	1	0.5000	0.5102	0.5102	49/50/2	49/50/2
RR	3	5	0.0000	0.0126	0.0126	99/0/2	99/0/2
RR	4	1	0.6000	0.6102	0.6102	39/60/2	39/60/2
RR	4	5	0.0000	0.0329	0.0330	97/0/4	97/0/4
GDP	$1/4$	1	0.1458	0.1580	0.1630	85/15/1	84/15/2
GDP	$1/4$	5	0.0000	0.0102	0.0102	99/1/1	99/1/1
GDP	$1/2$	1	0.2215	0.2328	0.2372	77/23/1	77/23/1
GDP	$1/2$	5	0.0005	0.0108	0.0110	99/1/1	99/1/1
GDP	$3/4$	1	0.3080	0.3189	0.3221	69/31/1	68/31/2
GDP	$3/4$	5	0.0125	0.0229	0.0239	98/2/1	98/2/1
Finite	–	1	0.1250	0.5728	0.5728	43/13/45	43/13/45
Finite	–	4	0.0000	0.0166	0.0167	99/0/2	99/0/2
Finite	–	8	0.0000	0.0102	0.0102	99/0/2	99/0/2

These cap cells compare the tightness of two implemented bounds from identical information. Each bracket is reused across caps; the cells are budget locations rather than independent release trials.

Table 18: Per-request times at core orders 2–4: median [minimum, maximum] milliseconds from three measurements after one warmup. Direct includes fresh construction and checking; the moment-based profile query reuses the checked historical law. Update is an ordinary moment update only.

Request	Strength	γ	Direct	Profile query	Update
RR	2	1	113 [113, 113]	4.69 [4.61, 6.15]	0.0829 [0.0812, 0.113]
RR	2	5	1.09 [1.03, 1.16]	2.08 [1.98, 2.14]	0.0761 [0.0757, 0.078]
RR	3	1	1.07 [1.06, 1.12]	3.81 [3.77, 3.83]	0.0763 [0.0762, 0.108]
RR	3	5	95.2 [94.1, 98.7]	2.41 [2.38, 2.67]	0.0826 [0.0819, 0.0831]
RR	4	1	1.17 [1.17, 1.22]	3.88 [3.85, 4.21]	0.0806 [0.0796, 0.0813]
RR	4	5	94.4 [94.2, 94.8]	2.94 [2.92, 3]	0.0807 [0.0795, 0.0826]
GDP	$1/4$	1	426 [423, 427]	62.6 [62.6, 63.1]	0.0987 [0.0954, 0.1]
GDP	$1/4$	5	692 [689, 702]	117 [115, 117]	0.0923 [0.0914, 0.115]
GDP	$1/2$	1	407 [406, 411]	59.8 [59.3, 60.2]	0.0928 [0.0909, 0.0941]
GDP	$1/2$	5	470 [468, 473]	73.2 [72.4, 75]	0.0972 [0.0959, 0.11]
GDP	$3/4$	1	404 [399, 407]	59 [58.8, 60.3]	0.111 [0.104, 0.123]
GDP	$3/4$	5	436 [430, 442]	64.6 [64, 65.5]	0.106 [0.0964, 0.138]
Finite	–	1	103 [103, 105]	5.54 [5.53, 5.54]	0.0347 [0.0343, 0.0351]
Finite	–	4	96.6 [96.6, 97.3]	3.3 [3.26, 3.34]	0.0345 [0.0331, 0.0352]
Finite	–	8	1.54 [1.52, 1.54]	2.22 [2.19, 2.31]	0.032 [0.0318, 0.0379]

Table 19: Per-request times at core orders 2–8: median [minimum, maximum] milliseconds from three measurements after one warmup. Direct includes fresh construction and checking; the moment-based profile query reuses the checked historical law. Update is an ordinary moment update only.

Request	Strength	γ	Direct	Profile query	Update
RR	2	1	240 [240, 249]	7.88 [7.61, 9.23]	0.184 [0.184, 0.187]
RR	2	5	1.72 [1.69, 1.74]	2.42 [2.42, 2.47]	0.176 [0.174, 0.186]
RR	3	1	1.92 [1.81, 1.97]	5.99 [5.95, 6.05]	0.192 [0.187, 0.201]
RR	3	5	194 [194, 202]	3.31 [3.26, 3.32]	0.189 [0.187, 0.195]
RR	4	1	1.79 [1.78, 1.79]	5.79 [5.75, 5.9]	0.184 [0.181, 0.209]
RR	4	5	203 [199, 204]	4.57 [4.48, 4.7]	0.188 [0.187, 0.189]
GDP	$1/4$	1	557 [554, 560]	78.7 [77.2, 82.1]	0.184 [0.184, 0.188]
GDP	$1/4$	5	806 [795, 863]	132 [132, 133]	0.207 [0.204, 0.229]
GDP	$1/2$	1	539 [535, 540]	75.6 [75.3, 77.6]	0.204 [0.191, 0.219]
GDP	$1/2$	5	621 [616, 632]	90.8 [90.2, 91.8]	0.187 [0.183, 0.201]
GDP	$3/4$	1	533 [529, 542]	74.3 [73.6, 74.5]	0.188 [0.185, 0.21]
GDP	$3/4$	5	567 [561, 589]	81.5 [81.4, 81.7]	0.241 [0.193, 0.274]
Finite	–	1	205 [205, 207]	8.96 [8.68, 9.14]	0.074 [0.0728, 0.0843]
Finite	–	4	201 [200, 202]	4.76 [4.58, 5.01]	0.0727 [0.0707, 0.0743]
Finite	–	8	2.09 [2.08, 2.15]	2.47 [2.43, 2.62]	0.0669 [0.0664, 0.0876]

Table 20: Moment-based profile preparation and separately measured complete 15-pair passes. Times are median [minimum, maximum] seconds. The profile pass includes fresh preparation once and all 15 queries.

Core orders	Atoms	Profile preparation	Direct menu	Profile menu
2–4	45	3.12 [3.06, 3.13]	3.44 [3.39, 3.5]	3.81 [3.8, 3.88]
2–8	49	6.03 [5.93, 6.04]	4.84 [4.75, 4.86]	6.59 [6.51, 6.61]

Timing uses the same machine and software environment, with one warmup and three measured repetitions. Direct calls start from an in-memory raw receipt and include preparation, construction, in-memory bound-record serialization, independent continuous verification, feasibility of the supplied equal-weight candidate on $\{3/4, 5/4\}$, and one rational cap comparison. Profile preparation includes all 63 dual constructions and checks, hull reconstruction, and exact verification of its atom law. A prepared query evaluates that law with a fresh request oracle, independently reevaluates the sum,

checks the common lower candidate, and classifies the cap. Request-oracle loading and evaluation are included; file input/output, candidate search, and private-release execution are excluded.

Timed decisions use cap $1/4$; the bound itself is independent of the cap. The full cap sweep is classified separately. Complete 15-pair passes in Table 20 are directly timed, including one fresh profile preparation for reuse. Ordinary RDP updates are an update-cost reference without a cap decision.

C Information retained by a finite receipt

The main accountant uses moment upper bounds and covers every compatible history. To isolate what finite historical compression itself loses, the results here also consider the stronger input of exact saved moments. A receipt is future-sufficient when it determines the composed profile for every allowed request. Even this stronger finite input is not universally future-sufficient; identification becomes possible under an explicit bound on the number of atoms. Finite-atomic identification uses classical Tchebycheff-system structure (Karlin & Studden, 1966).

Theorem 5 (No finite order grid is universally future-sufficient)**.** *For every finite set of orders $\{\alpha_1, \dots, \alpha_m\} \subset (1, \infty)$ and every $\mu > 0$, there exist two bounded finite-output binary experiments with identical exact bidirectional moments at all saved orders but different profiles after composition with G_μ.*

Proof. Use the $2m + 2$ functions

$$1, \quad x, \quad x^{\alpha_j}, \quad x^{1-\alpha_j}, \qquad j = 1, \dots, m.$$

Choose $2m + 3$ distinct positive nodes on both sides of one and strictly positive base weights w_i with total mass and mean one. The corresponding $(2m + 2) \times (2m + 3)$ evaluation matrix has a nonzero null vector c. Because the constant row is annihilated, c has both signs. For sufficiently small $\tau > 0$, the perturbed weights $w_i^{\pm} = w_i \pm \tau c_i$ remain positive. They define distinct finite likelihood-ratio laws with the same mass, mean, and all saved positive and negative moments.

Let $X_\pm$ have weights $w_i^{\pm}$ at these nodes. Let $U_\pm = \log X_\pm$ and let V be the Gaussian log-likelihood ratio of G_μ. If $U_+ + V$ and $U_- + V$ had the same law, their characteristic functions would be equal. The Gaussian characteristic function never vanishes, so the characteristic functions of U_+ and U_- would coincide, contradicting the distinct finite laws. The composed likelihood-ratio laws are therefore different. Since the stop-loss transform determines a nonnegative law, some threshold separates their profiles. □

Theorem 6 (Approximation duality)**.** *Fix $0 < \ell < 1 < u < \infty$, let $\mathcal{A} = \{\alpha_1, \dots, \alpha_m\}$, and define*

$$V_{\mathcal{A}} = \operatorname{span}\{1, x, x^{\alpha_j}, x^{1-\alpha_j} : j = 1, \dots, m\} \subset C([\ell, u]).$$

Here $C([\ell, u])$ is the space of continuous functions. For $g \in C([\ell, u])$, let $\Omega_{\mathcal{A}}(g)$ be the largest difference in $\int g \, d\nu$ between two mean-one probability laws on $[\ell, u]$ with identical moments on $V_{\mathcal{A}}$. Write $\operatorname{dist}_\infty(g, V_{\mathcal{A}}) = \inf_{v \in V_{\mathcal{A}}} \|g - v\|_\infty$ for the best uniform approximation error. With

$$\kappa_{\ell,u} = \frac{\min\{u - 1, 1 - \ell\}}{u - \ell},$$

$$2\kappa_{\ell,u} \operatorname{dist}_\infty(g, V_{\mathcal{A}}) \le \Omega_{\mathcal{A}}(g) \le 2 \operatorname{dist}_\infty(g, V_{\mathcal{A}}).$$

If the two laws may share any mean, their largest objective difference equals $2 \operatorname{dist}_\infty(g, V_{\mathcal{A}})$.

Proof. For any receipt-equivalent ν_1, ν_2 and any $v \in V_{\mathcal{A}}$,

$$\int g \, d(\nu_1 - \nu_2) = \int (g - v) \, d(\nu_1 - \nu_2),$$

so the absolute value is at most $2\|g - v\|_\infty$. Taking the infimum proves the upper bound.

For the lower bound, Hahn–Banach duality in $C([\ell, u])/V_{\mathcal{A}}$ gives a signed measure σ with total variation one, annihilating $V_{\mathcal{A}}$, and satisfying $\int g \, d\sigma = d$, where $d = \operatorname{dist}_\infty(g, V_{\mathcal{A}})$. Because 1 is in the span, the positive and negative Jordan parts each have mass $1/2$. Normalizing them gives two

probability laws with identical saved moments and a common mean r, and their objective gap is $2d$. If $r < 1$, mix both with the common atom δ_u using weight $(u-1)/(u-r)$. If $r > 1$, mix both with δ_ℓ using $(1-\ell)/(r-\ell)$. The resulting laws have mean one, preserve receipt equality, and retain an objective gap at least $2\kappa_{\ell,u}d$. If the common mean need not equal one, no mixing is necessary. □

Lemma 7 (Exponential generalized Vandermonde rank). *If $\lambda_1 < \cdots < \lambda_d$ and $t_1 < \cdots < t_n$ with $n \le d$, then the matrix $(e^{\lambda_r t_i})_{r,i}$ has full column rank. When $n = d$, its determinant is positive.*

Proof. A nonzero n-term exponential polynomial $\sum_{i=1}^{n} c_i e^{t_i \lambda}$ has at most $n-1$ distinct real zeros. This follows inductively by multiplying by $e^{-t_1 \lambda}$, differentiating, and applying Rolle's theorem to the resulting $(n-1)$-term exponential polynomial. Hence a nonzero column combination cannot vanish at $d \ge n$ distinct exponents.

For the square case the determinant never vanishes on the connected region of strictly ordered exponents and nodes, so its sign is constant. Set $t_i = \varepsilon s_i$ and let $\varepsilon \downarrow 0$. The leading Taylor term is a positive multiple of the product of the two Vandermonde determinants, so the sign is positive. Substituting $t_i = \log x_i$ gives the generalized Vandermonde statement for positive nodes. □

Theorem 8 (Exact finite-atomic threshold). *Let $\mathcal{C}_k$ be the class of likelihood-ratio laws with at most k positive atoms. Save exact bidirectional moments at m distinct orders greater than one, together with mass and mean one. Then $m \ge k-1$ uniquely determines every law in $\mathcal{C}_k$. If $m \le k-2$, for every such order grid there are two distinct laws in $\mathcal{C}_k$ with the same receipt and different future profiles after composition with any nondegenerate GDP factor.*

Proof. Order the $2m+2$ exponents as

$$1-\alpha_m, \ldots, 1-\alpha_1, 0, 1, \alpha_1, \ldots, \alpha_m.$$

For sufficiency, the signed difference of two k-atomic laws has at most $2k$ positive support points. If $m \ge k-1$, then $2k \le 2m+2$, and Lemma 7 gives full column rank for the evaluation matrix on the union support. A signed measure annihilating every saved function must therefore be zero.

For necessity, let $d = 2m+2$ and choose $d+1 = 2m+3$ distinct positive nodes. The $d \times (d+1)$ generalized Vandermonde matrix has a one-dimensional nullspace. The cofactor null vector has strictly alternating signs because all maximal minors have the same strict sign. Its positive and negative parts therefore define two distinct probability laws supported on at most $m+2$ and $m+1$ points, respectively, with identical saved moments and a common positive mean. Divide every node by that common mean to obtain mean-one laws; receipt equality is preserved because both sides of each moment equation acquire the same factor. When $m \le k-2$, both support sizes are at most k. The GDP convolution argument from Theorem 5 separates their future profiles. The case $k = 1, m = 0$ is included because mass and mean identify the sole atom as one. □

Theorem 9 (Constructive bounded-class upper bound). *Fix $[\ell, u] \Subset (0, \infty)$ and use the integer order grid $\{2, 3, \ldots, m+1\}$. For every fixed factor and threshold, the corresponding future-profile ambiguity is at most*

$$\frac{2C_J(u-\ell)}{m+1}$$

for a universal Jackson constant C_J. This is a constructive $O(m^{-1})$ upper bound for this integer-order construction.

Proof. The positive exponents contain $1, x, x^2, \ldots, x^{m+1}$, so the saved span contains every polynomial of degree at most $m+1$. Theorem 11 makes every fixed-factor objective 1-Lipschitz on $[\ell, u]$. Jackson approximation therefore gives a polynomial p of degree at most $m+1$ with

$$\|g_{\gamma,F} - p\|_\infty \le \frac{C_J(u-\ell)}{m+1}.$$

The upper half of Theorem 6 gives the result. □

D PROOFS AND NUMERICAL GUARANTEES

D.1 PRODUCT MODEL AND FIXED-REQUEST PROFILE IDENTITY

An experiment is Blackwell-dominated by another when a data-independent Markov kernel applied to the latter produces the former. Let Z be the historical transcript and write E_H for its binary experiment. When the future binary experiment is fixed independently of Z, write it as F.

Proposition 10 (Static product and common-dominator bound)**.** *1. If the future kernel is fixed independently of Z and the historical and future mechanisms use conditionally independent randomness given the dataset, the joint binary experiment is exactly $E_H \otimes F$.*

2. Suppose the transcript and output spaces are standard Borel. If the conditional future experiment is F_z and there exist one transcript-independent experiment $F^\star$ and a jointly measurable family of data-independent garbling kernels K_z such that $F_z = K_z \circ F^\star$ for every relevant z, then the actual sequential experiment is Blackwell-dominated by $E_H \otimes F^\star$.

The domination through $F^\star$ gives a conservative upper bound. It may be strict, and the construction does not identify the conditional optimum.

Proof. For the first part, conditional independence gives the product form under each of the two neighboring datasets, hence the product binary experiment. For the second part, start from (Z, U) generated by $E_H \otimes F^\star$. Keep Z and, conditional on the observed value $Z = z$, apply the data-independent kernel K_z to U. Joint measurability makes this a valid Markov kernel on the product output space. Its image is the actual sequential experiment, so the latter is a garbling of $E_H \otimes F^\star$ and is Blackwell-dominated by it. □

Theorem 11 (Fixed-request profile identity)**.** *Let the fixed factor $F = (P_F, Q_F)$ on output space Ω_F have Lebesgue decomposition*

$$P_F = P_F^{\mathrm{ac}} + P_F^{\perp}, \qquad s_F = P_F^{\perp}(\Omega_F), \qquad Y = \frac{dP_F^{\mathrm{ac}}}{dQ_F}.$$

Define the full profile

$$h_F(t) = s_F + \mathbb{E}_{Q_F}(Y - t)_+, \qquad t \geq 0.$$

Let $X > 0$ be the historical likelihood ratio of E_H, with $\mathbb{E}X = 1$, and assume independent randomness. Then, for every $\gamma \geq 0$,

$$h_{E_H \otimes F}(\gamma) = \mathbb{E}[X h_F(\gamma/X)]\,.$$

Consequently

$$g_{\gamma,F}(x) = x h_F(\gamma/x)$$

is convex, nondecreasing, and 1-Lipschitz on $(0, \infty)$. If $\underline{h}_F \leq h_F \leq \overline{h}_F$, then

$$x\underline{h}_F(\gamma/x) \leq g_{\gamma,F}(x) \leq x\overline{h}_F(\gamma/x)$$

provides one-sided enclosures of the fixed-request objective $g_{\gamma,F}$.

Proof. Condition on $X = x$. The absolutely continuous factor part contributes

$$\mathbb{E}_{Q_F}(xY - \gamma)_+ = x\mathbb{E}_{Q_F}(Y - \gamma/x)_+,$$

and the singular factor mass contributes xs_F. Averaging over the historical Q-law gives the identity. Equivalently,

$$g_{\gamma,F}(x) = xs_F + \mathbb{E}_{Q_F}(xY - \gamma)_+.$$

This is a nonnegative linear term plus an expectation of positive parts of affine functions, so it is convex and nondecreasing. Since $\mathbb{E}_{Q_F} Y = 1 - s_F$,

$$|g_{\gamma,F}(x) - g_{\gamma,F}(x')| \leq \{s_F + \mathbb{E}Y\}|x - x'| = |x - x'|.$$

The profile enclosure follows by multiplying by $x > 0$. □

Admission values and comparison class. For the receipt $M = \{M_j^+, M_j^-\}_{j=1}^m$, let $\mathcal{U}(M)$ be the laws of positive, mean-one likelihood ratios X that satisfy all of its moment bounds. The factor-retained robust value is

$$W_F(\gamma; M) = \sup_{\mathcal{L}(X)\in\mathcal{U}(M)} \mathbb{E}[Xh_F(\gamma/X)]\,.$$

At a public cap (γ_B, δ_B), a verified bound $W_F(\gamma_B; M) \le \delta_B$ certifies this direction; admission requires the same conclusion in every ordered direction.

For comparison, let the request-moment bounds satisfy

$$N_{F,j}^+ \ge \mathbb{E}_{Q_F} Y^{\alpha_j}, \qquad N_{F,j}^- \ge \mathbb{E}_{Q_F} Y^{1-\alpha_j}.$$

Set $N_{F,j}^+ = \infty$ when the request has singular P_F-mass and $N_{F,j}^- = \infty$ when $Q_F(Y = 0) > 0$. When the required bounds are finite, define

$$\begin{aligned}
\mathcal{K}_F &= \{\mathcal{L}(XY) : \mathcal{L}(X) \in \mathcal{U}(M),\ X \perp Y\},\\
\mathcal{A}_{\mathrm{agg}} &= \{\mathcal{L}(Z) : Z > 0,\ \mathbb{E}Z = 1,\\
&\qquad\qquad \mathbb{E}Z^{\alpha_j} \le M_j^+ N_{F,j}^+, \quad \mathbb{E}Z^{1-\alpha_j} \le M_j^- N_{F,j}^- \ \forall j\},\\
C_{\mathrm{agg}}(\gamma; M) &= \sup_{\mathcal{L}(Z)\in\mathcal{A}_{\mathrm{agg}}} \mathbb{E}(Z-\gamma)_+.
\end{aligned}$$

Thus $\mathcal{K}_F \subseteq \mathcal{A}_{\mathrm{agg}}$: the factor-retained value varies only the history, whereas the joint-moment class also permits unfactorized aggregate laws. A feasible law above the cap rules out robust admission for its compatibility class; a bracket straddling the cap is inconclusive. When a result below specifies a receipt, we suppress that fixed M and write $W_F(\gamma)$ and $C_{\mathrm{agg}}(\gamma)$.

D.2 Compact generalized-moment optimization

Theorem 12 (Fixed-factor moment optimization)**.** *Fix $0 < \ell < 1 < u < \infty$ and let $\mathcal{M}$ be the nonempty set of probability laws on $[\ell, u]$ satisfying*

$$\int x\, d\nu(x) = 1, \qquad \int x^{\alpha_j}\, d\nu(x) \le M_j^+, \qquad \int x^{1-\alpha_j}\, d\nu(x) \le M_j^-.$$

For a fixed factor with valid full profile,

$$W_F^{[\ell,u]}(\gamma; M) = \max_{\nu\in\mathcal{M}} \int g_{\gamma,F}(x)\, d\nu(x).$$

The maximum is attained, and some optimizer has at most $2m + 3$ atoms. If one law in $\mathcal{M}$ satisfies every moment inequality strictly, then the value equals the attained minimum

$$\lambda_0 + \lambda_1 + \sum_{j=1}^m u_j M_j^+ + \sum_{j=1}^m v_j M_j^-$$

over $\lambda_0, \lambda_1 \in \mathbb{R}$ and $u_j, v_j \ge 0$, subject to

$$\lambda_0 + \lambda_1 x + \sum_{j=1}^m u_j x^{\alpha_j} + \sum_{j=1}^m v_j x^{1-\alpha_j} \ge g_{\gamma,F}(x) \quad \textit{for every } x \in [\ell, u].$$

The stated atomic representation and dual formulation follow from classical generalized-moment machinery (Winkler, 1988; Anderson & Nash, 1987).

Proof. By Theorem 11, $g_{\gamma,F}$ is continuous on the compact interval. The space $\mathcal{P}([\ell, u])$ is weakly compact, and each displayed constraint is weakly closed because its integrand is continuous. Thus $\mathcal{M}$ is compact and the objective is weakly continuous, proving attainment.

Let $\nu^\star$ be optimal and define

$$T(x) = \big(x, x^{\alpha_1}, \ldots, x^{\alpha_m}, x^{1-\alpha_1}, \ldots, x^{1-\alpha_m}, g_{\gamma,F}(x)\big) \in \mathbb{R}^{2m+2}.$$

The vector $\int T\,d\nu^\star$ lies in the convex hull of $T([\ell,u])$. Carathéodory's theorem represents it by at most $2m+3$ points. The resulting atomic law preserves the mean, all saved moments, and the objective, so it is feasible and optimal.

Weak duality follows by integrating any feasible pointwise majorant. It remains to prove strong duality and dual attainment under strict slack. Enumerate the $2m$ moment functions by ϕ_i and their upper bounds by M_i. Define the perturbation value

$$V(r,s)=\sup\left\{\int g_{\gamma,F}\,d\nu:\nu\in\mathcal{P}([\ell,u]),\ \int x\,d\nu=1+r,\ \int\phi_i\,d\nu\le M_i+s_i\right\}.$$

Compactness makes V finite and upper semicontinuous on its effective domain. Let ν_0 satisfy every moment inequality with slack at least $\sigma>0$. For sufficiently small $r>0$, mix ν_0 with δ_u using weight $r/(u-1)$; for sufficiently small $r<0$, mix it with δ_ℓ using weight $(-r)/(1-\ell)$. The resulting law has mean $1+r$. Because there are only finitely many continuous moment functions, their changes are $O(|r|)$, so they retain at least $\sigma/2$ slack for small $|r|$. This remaining slack permits small positive or negative perturbations of every coordinate of s. Consequently the origin lies in the relative interior of the effective domain.

The finite concave function V therefore has a finite supergradient (a,b) at the origin:

$$V(r,s)\le V(0,0)+ar+b\cdot s.$$

Monotonicity under relaxation gives $b_i\ge 0$. Apply the inequality to an arbitrary probability law, and then to Dirac laws, to obtain a pointwise majorant of the required form. Setting $\lambda_1=a$ and choosing λ_0 so that the dual objective equals $V(0,0)$ proves zero gap and dual attainment. □

D.3 Strict separations and the exact admission decision

Both separations use only the historical bidirectional order-two bounds. The identity in Theorem 11 gives $W_F(\gamma)\le h_F(\gamma)+\sqrt{\eta}$; a feasible aggregate law gives a lower bound on $C_{\mathrm{agg}}(\gamma)$. A cap between these bounds changes robust admission. Figure 1 displays the Gaussian case.

Theorem 13 (Randomized-response strictness)**.** *Let $a>1$ and let F_a be the symmetric randomized-response factor with Q-side likelihood ratio*

$$Y_a=\begin{cases}a, & 1/(a+1),\\ a^{-1}, & a/(a+1).\end{cases}$$

Its two directional order-two moments equal $B(a)=a-1+a^{-1}$. Let

$$q(a)=a^3-2a^2-2a+2,$$

and let $a_0\simeq 2.4812$ be its largest real root. For $a\ge a_0$, set

$$\gamma_a=2a-1,\qquad p_a=(2a^2+3a-1)^{-1}.$$

If the historical receipt satisfies

$$\mathbb{E}X=1,\qquad \mathbb{E}X^2\le 1+\eta,\qquad \mathbb{E}X^{-1}\le 1+\eta,$$

then

$$W_{F_a}(\gamma_a)\le\sqrt{\eta},\qquad C_{\mathrm{agg}}(\gamma_a)\ge p_a.$$

The separation is strict when $\eta<p_a^2$.

Proof. Because $Y_a\le a<\gamma_a$, $(Y_a-\gamma_a)_+=0$. The positive-part map is 1-Lipschitz, so

$$(XY_a-\gamma_a)_+\le Y_a|X-1|.$$

Independence, $\mathbb{E}Y_a=1$, and Cauchy–Schwarz give

$$\mathbb{E}(XY_a-\gamma_a)_+\le\mathbb{E}|X-1|\le\sqrt{\mathbb{E}(X-1)^2}\le\sqrt{\eta}.$$

For the aggregate distribution define

$$Z_a=\begin{cases}2a, & p_a,\\ (a+1)/(a+2), & 1-p_a.\end{cases}$$

Direct calculation gives

$$\mathbb{E}Z_a = 1, \qquad \mathbb{E}Z_a^2 = \frac{3a+1}{a+2}, \qquad \mathbb{E}Z_a^{-1} = \frac{2a^2+4a-1}{2a(a+1)}.$$

Moreover,

$$B(a) - \mathbb{E}Z_a^2 = \frac{q(a)}{a(a+2)},$$

and

$$B(a) - \mathbb{E}Z_a^{-1} = \frac{2a^3-2a^2-4a+3}{2a(a+1)}.$$

The latter numerator is $q(a) + (a^3 - 2a + 1)$, with $a^3 - 2a + 1 > 0$ for $a > 1$. Thus both moments of this distribution are bounded by $B(a)$ when $q(a) \geq 0$, even before multiplication by the historical factor $1 + \eta$. Finally,

$$\mathbb{E}(Z_a - \gamma_a)_+ = p_a(2a - \gamma_a) = p_a.$$

□

For Gaussian DP, G_μ denotes the binary experiment whose log-likelihood ratio has law $N(\mu^2/2, \mu^2)$ under P and $N(-\mu^2/2, \mu^2)$ under Q. Write Φ and ϕ for the standard-normal distribution function and density.

Theorem 14 (Gaussian-DP strictness)**.** *Let the fixed factor be G_μ with $0 < \mu \leq 4$, and suppose the same bidirectional order-two historical receipt as in Theorem 13. Define*

$$\Delta_G(\mu) = \tanh(\mu/2) - \{2\Phi(\mu/2) - 1\}.$$

Then $\Delta_G(\mu) > 0$ on $(0, 4]$ and, at $\gamma = 1$,

$$W_{G_\mu}(1) \leq 2\Phi(\mu/2) - 1 + \sqrt{\eta}, \qquad C_{\mathrm{agg}}(1) \geq \tanh(\mu/2).$$

Hence the separation is strict when $\eta < \Delta_G(\mu)^2$ throughout the stated range $(0, 4]$.

Proof. The GDP hockey-stick value at threshold one is $2\Phi(\mu/2) - 1$. Theorem 11 and the historical second-moment bound give

$$W_{G_\mu}(1) \leq 2\Phi(\mu/2) - 1 + \mathbb{E}|X - 1| \leq 2\Phi(\mu/2) - 1 + \sqrt{\eta}.$$

For an aggregate distribution take symmetric randomized response with parameter $a = e^\mu$. Its two directional second moments equal $2\cosh\mu - 1$, and

$$2\cosh\mu - 1 = 1 + 2\sum_{k\geq 1} \frac{\mu^{2k}}{(2k)!} \leq 1 + \sum_{k\geq 1} \frac{\mu^{2k}}{k!} = e^{\mu^2},$$

because $2\,k! \leq (2k)!$ for $k \geq 1$. The GDP factor's directional order-two moments are e^{μ^2}, so this distribution belongs to the aggregate class. Its threshold-one value is

$$\frac{e^\mu - 1}{e^\mu + 1} = \tanh(\mu/2).$$

It remains to prove $\tanh t > \mathrm{erf}(t/\sqrt{2})$ for $0 < t \leq 2$, where $\mathrm{erf}(t/\sqrt{2}) = 2\Phi(t) - 1$. Put

$$d(t) = \tanh t - \mathrm{erf}(t/\sqrt{2}), \qquad d'(t) = \mathrm{sech}^2\, t - \sqrt{2/\pi}\, e^{-t^2/2}.$$

Define

$$R(t) = \log \frac{\mathrm{sech}^2\, t}{\sqrt{2/\pi}\, e^{-t^2/2}}.$$

The signs of d' and R agree, and $R'(t) = t - 2\tanh t$. The function $2\tanh t - t$ is strictly concave on $(0, \infty)$, starts at zero with positive derivative, and is negative at $t = 2$. Thus it has exactly one positive zero in $(0, 2)$, so R first decreases and then increases. Since $R(0) > 0$ and $R(2) < 0$, the derivative d' changes sign once and d first increases and then decreases.

The Mills lower bound

$$\Phi(-x) > \frac{x}{1+x^2}\phi(x)$$

follows by one integration by parts. At $x = 2$,

$$2\Phi(-2) > \frac{4e^{-2}}{5\sqrt{2\pi}}.$$

The inequality

$$\frac{4e^{-2}}{5\sqrt{2\pi}} > \frac{2}{e^4+1}$$

is equivalent to $2(e^2+e^{-2}) > 5\sqrt{2\pi}$. The left side exceeds 14 because $e^2 > 7$, while the right side is less than $5\sqrt{44/7} < 14$ because $\pi < 22/7$ and $25(44/7) < 14^2$. Therefore $2\Phi(-2) > 1 - \tanh 2$, so $d(2) > 0$. Since $d(0) = 0$ and d first increases and then decreases while remaining positive at the right endpoint, $d(t) > 0$ for all $t \in (0, 2]$. □

An exact RR admission example. Take symmetric RR with $a = 3$, which reports its binary input correctly with probability $3/4$. Use budget $(\varepsilon_B, \delta_B) = (\log 5, 1/40)$ and historical bounds $\mathbb{E}X^2, \mathbb{E}X^{-1} \le 1 + 1/2704$. The request's standalone profile is zero at threshold $\gamma_B = 5$, and the historical contribution is at most $\sqrt{\eta} = 1/52$. Theorem 13 therefore bounds $W_{F_3}(5)$ by $1/52$, below the public cap $1/40$. Symmetry and the reverse receipt give the same upper bound in the reverse direction, so the fixed request is admitted.

For the aggregate comparison, put probability $1/26$ at likelihood ratio 6 and $25/26$ at $4/5$. This law has mean one, positive second moment 2, and negative first moment $29/24$. Both moments are below the RR request's order-two value $7/3$, even before multiplication by the historical bound $1 + \eta$. It is therefore aggregate-compatible, yet its profile at threshold 5 is $1/26$, above the cap. The example instantiates the theorem's admission interval with exact rational checks; it rules out robust admission from the aggregate order-two summary.

D.4 Bounds for full profiles and continuous domains

Here a *stop-loss profile* is the hockey-stick curve $\gamma \mapsto \mathbb{E}(X-\gamma)_+$ defined above.

Proposition 15 (Profile-supremum validity)**.** *Let $\{h_i\}$ be mean-one stop-loss profiles with a common 1-Lipschitz bound and with a common tail control ensuring their pointwise supremum tends to zero. Then $h = \sup_i h_i$ is again a valid mean-one stop-loss profile.*

Proof. Each h_i is convex, nonincreasing, 1-Lipschitz, satisfies $h_i(0) = 1$, and tends to zero. These properties are preserved by the pointwise supremum under the common tail control. The right derivative $-h_+'$ is a right-continuous nonincreasing tail function taking values in $[0, 1]$, and

$$\int_0^\infty -h_+'(t)\,dt = h(0) = 1.$$

The corresponding mean-one nonnegative random variable has stop-loss transform h. This is the stop-loss-profile specialization of the privacy-loss distribution construction of the Blackwell lattice in Regehr et al. (2026, Proposition 11), used below in Theorem 16. □

The finite-node bracket below assumes an absolutely continuous request with a positive likelihood ratio Y satisfying $\mathbb{E}Y = 1$. Theorem 11 supplies the underlying profile identity when the request has singular mass.

Theorem 16 (Finite-node profile bracket)**.** *Let $\mathcal{M}$ be a compact historical likelihood-ratio class on $[\ell, u]$ with mean one, and let Y be the fixed-factor likelihood ratio with $\mathbb{E}Y = 1$. In every expectation below, Y is independent of the historical likelihood ratio X. Define*

$$W(\gamma) = \sup_{\nu \in \mathcal{M}} \mathbb{E}_{X\sim\nu, Y}(XY - \gamma)_+.$$

Let $0 = \gamma_0 < \cdots < \gamma_N = \Gamma$ have maximum gap Δ. At each node, let $\nu_i \in \mathcal{M}$ be a primal-feasible history law and define its composed profile by

$$h_i(t) := \mathbb{E}_{X\sim\nu_i, Y}(XY - t)_+, \qquad t \ge 0.$$

Suppose a dual cap $u_i \geq W(\gamma_i)$ satisfies $0 \leq u_i - h_i(\gamma_i) \leq \varepsilon$. Suppose also that, for some $\beta > 1$,

$$\sup_{\nu \in \mathcal{M}} \mathbb{E}X^\beta \leq M_\beta, \qquad \mathbb{E}Y^\beta \leq N_\beta,$$

and define

$$T_\beta(\gamma) = \frac{(\beta-1)^{\beta-1}}{\beta^\beta} M_\beta N_\beta \gamma^{1-\beta}.$$

Let $\underline{W} = \sup_i h_i$, and let $\overline{W}$ be the pointwise supremum of all mean-one stop-loss profiles satisfying the node caps and the tail cap $h(\gamma) \leq T_\beta(\gamma)$ for $\gamma \geq \Gamma$. Then

$$\underline{W} \leq W \leq \overline{W},$$

$\underline{W}$ and $\overline{W}$ are valid mean-one stop-loss profiles, and on $[0, \Gamma]$,

$$0 \leq W - \underline{W} \leq \varepsilon + \Delta, \qquad 0 \leq \overline{W} - W \leq \varepsilon + \Delta.$$

For $\gamma \geq \Gamma$ the bracket width is at most $T_\beta(\gamma)$. If $T_\beta(\Gamma) \leq \tau$, the global width is at most

$$\max\{2(\varepsilon + \Delta), \tau\}.$$

Proof. For every fixed history law, the composed stop-loss profile is convex, nonincreasing, and 1-Lipschitz. Their pointwise supremum W has the same properties, satisfies $W(0) = 1$, and tends to zero because $X \leq u$ and $\mathbb{E}Y = 1$. Proposition 15 therefore makes W a valid profile. Each h_i is feasible and hence lies below W, so $\underline{W} \leq W$ and $\underline{W}$ is valid.

For any $\gamma \in [0, \Gamma]$, choose a nearest node γ_i. Both W and h_i are 1-Lipschitz, so

$$W(\gamma) - h_i(\gamma) \leq |W(\gamma) - W(\gamma_i)| + \{W(\gamma_i) - h_i(\gamma_i)\} + |h_i(\gamma_i) - h_i(\gamma)| \leq \varepsilon + \Delta.$$

Taking the supremum over i gives the lower-envelope error.

For any $z, \gamma > 0$,

$$(z - \gamma)_+ \leq \frac{(\beta-1)^{\beta-1}}{\beta^\beta} \gamma^{1-\beta} z^\beta.$$

Thus $W \leq T_\beta$ on the tail, so W belongs to the upper class and $W \leq \overline{W}$. The common tail control and Proposition 15 make $\overline{W}$ valid. For any profile h in the upper class and a nearest node,

$$h(\gamma) - W(\gamma) \leq |h(\gamma) - h(\gamma_i)| + \{u_i - W(\gamma_i)\} + |W(\gamma_i) - W(\gamma)| \leq \varepsilon + \Delta.$$

Taking the supremum over h proves the upper-envelope error. The tail and global-width bounds follow directly. □

Proposition 17 (Continuous-domain dual margin)**.** *On $[\ell, u_{\max}]$, consider a dual majorant*

$$q(x) = \lambda_0 + \lambda_1 x + \sum_{j=1}^m u_j x^{\alpha_j} + \sum_{j=1}^m v_j x^{1-\alpha_j}, \qquad u_j, v_j \geq 0,$$

and set

$$L_q = |\lambda_1| + \sum_{j=1}^m u_j \alpha_j u_{\max}^{\alpha_j - 1} + \sum_{j=1}^m v_j (\alpha_j - 1) \ell^{-\alpha_j}.$$

Let a historical-support grid on $[\ell, u_{\max}]$ have maximum gap δ_x. If every node x_r satisfies

$$q(x_r) - g_{\gamma,F}(x_r) \geq \frac{L_q + 1}{2} \delta_x$$

with all values enclosed outward, then $q(x) \geq g_{\gamma,F}(x)$ for all $x \in [\ell, u_{\max}]$.

Proof. Termwise differentiation gives $\sup |q'| \leq L_q$, while Theorem 11 gives a 1-Lipschitz bound for $g_{\gamma,F}$. Therefore the residual $q - g_{\gamma,F}$ is $(L_q + 1)$-Lipschitz. Every point lies within $\delta_x/2$ of a node, so the displayed node margin preserves a nonnegative residual throughout the interval. □

Theorem 18 (Compact reduction for unbounded histories)**.** *Let $X > 0$ satisfy*

$$\mathbb{E}X = 1, \qquad \mathbb{E}X^{\beta} \le B_+, \qquad \mathbb{E}X^{-r} \le B_-$$

for some $\beta > 1$ and $r > 0$, together with the saved receipt moments. Choose $0 < \ell < 1 < u$ and set

$$\delta_{\ell,u} = B_+ u^{1-\beta} + B_- \ell^{r+1} < 1.$$

Define

$$\widetilde{X} = \min\{u, \max\{\ell, X\}\}, \qquad a = \mathbb{E}\widetilde{X}, \qquad X^c = \widetilde{X}/a.$$

Then $\mathbb{E}X^c = 1$,

$$X^c \in \left[\frac{\ell}{1+\delta_{\ell,u}}, \frac{u}{1-\delta_{\ell,u}}\right], \qquad \mathbb{E}|X - X^c| \le 2\delta_{\ell,u}.$$

For $s_j = \alpha_j - 1$,

$$\mathbb{E}(X^c)^{\alpha_j} \le \frac{M_j^+ + B_-\ell^{\alpha_j+r}}{(1-\delta_{\ell,u})^{\alpha_j}} =: \widetilde{M}_j^+,$$

and

$$\mathbb{E}(X^c)^{-s_j} \le (1+\delta_{\ell,u})^{s_j}\left(M_j^- + B_+ u^{-(\beta+s_j)}\right) =: \widetilde{M}_j^-.$$

For every independent valid factor and every $\gamma \ge 0$,

$$|\mathbb{E}g_{\gamma,F}(X) - \mathbb{E}g_{\gamma,F}(X^c)| \le 2\delta_{\ell,u}.$$

Consequently the original unbounded robust profile lies below the compact outer problem with the relaxed moments plus $2\delta_{\ell,u}$, and above the original class restricted to $[\ell, u]$.

Proof. For the upper tail,

$$\mathbb{E}(X-u)_+ \le \mathbb{E}[X\mathbf{1}\{X > u\}] \le u^{1-\beta}\mathbb{E}X^{\beta}.$$

For the lower tail,

$$\mathbb{E}(\ell - X)_+ \le \ell \Pr(X < \ell) \le B_-\ell^{r+1}.$$

Thus $\mathbb{E}|X - \widetilde{X}| \le \delta_{\ell,u}$ and $|a-1| \le \delta_{\ell,u}$. The support bound follows, and

$$\mathbb{E}|X - X^c| \le \mathbb{E}|X - \widetilde{X}| + \mathbb{E}|\widetilde{X} - \widetilde{X}/a| = \mathbb{E}|X - \widetilde{X}| + |a-1| \le 2\delta_{\ell,u}.$$

Upper clipping cannot increase positive moments. Lower clipping adds at most $\ell^{\alpha_j}\Pr(X < \ell) \le B_-\ell^{\alpha_j+r}$, and division by $a^{\alpha_j} \ge (1-\delta_{\ell,u})^{\alpha_j}$ gives the positive-moment relaxation. Lower clipping cannot increase negative moments. Upper clipping adds at most $u^{-s_j}\Pr(X > u) \le B_+u^{-(\beta+s_j)}$, and multiplication by $a^{s_j} \le (1+\delta_{\ell,u})^{s_j}$ gives the negative-moment relaxation. Finally, Theorem 11 gives

$$|\mathbb{E}g_{\gamma,F}(X) - \mathbb{E}g_{\gamma,F}(X^c)| \le \mathbb{E}|X - X^c| \le 2\delta_{\ell,u}.$$

Taking suprema proves the robust-profile bracket. □

For computation, discretize the compact history law on K historical-support nodes and evaluate the profile bracket at G threshold nodes. The resulting linear programs (LPs) are discretized versions of the compact primal and dual above.

Theorem 19 (Conditional discretization bound)**.** *Suppose the threshold grid has G points, the historical-support grid has K points, and the receipt has m bidirectional orders. At each threshold the discretized primal has K nonnegative weights, two equality constraints, and $2m$ moment inequalities. After splitting the two free affine coefficients, the discretized dual has $2m+4$ nonnegative variables and K node constraints. A full discretized bracket therefore uses $2G$ LP solves, at most GK cached request-profile evaluations, and $O(GKm)$ independent arithmetic checks.*

At a fixed threshold, suppose a semi-infinite dual has an ε_0-optimal feasible majorant q with $\sup_{[\ell,u]}|q'| \le B$. If the $g_{\gamma,F}$ upper enclosure has width at most ε_F at each node and the historical-support mesh is δ_x, then a constant shift of

$$\frac{B+1}{2}\delta_x + \varepsilon_F$$

produces a continuous-domain bound. Its value D_{δ_x} satisfies

$$W_F^{[\ell,u]}(\gamma; M) \le D_{\delta_x} \le W_F^{[\ell,u]}(\gamma; M) + \varepsilon_0 + \frac{B+1}{2}\delta_x + \varepsilon_F.$$

With nodewise primal-dual gaps at most ε_{PD} and threshold mesh Δ_γ, the compact profile bracket has width at most $2(\varepsilon_{PD} + \Delta_\gamma)$. These are conditional operation counts and convergence guarantees for the supplied discretized bounds.

Proof. The dimensions follow by direct counting. Evaluate the semi-infinite feasible q at the historical-support nodes. The outward enclosure costs at most ε_F. After the displayed constant shift, every node has residual margin $(B+1)\delta_x/2$. The derivative bound for q and the 1-Lipschitz bound for $g_{\gamma,F}$ cover each interval, proving feasibility. The objective rises by exactly the constant shift; weak duality gives the lower inequality. The profile-width statement is Theorem 16. □

Theorem 20 (Finite-precision receipt sensitivity)**.** *Fix a factor F, a threshold $\gamma \ge 0$, and a compact support interval $[\ell, u_{\max}]$. Let the compact problem on this interval use receipt upper bounds $M = (M^+, M^-)$, and let the outward-rounded bounds be*

$$\overline{M} = M + (\Delta^+, \Delta^-), \qquad \Delta^\pm \ge 0.$$

Let $\mathbf{u} = (u_j)_{j=1}^m$ and $\mathbf{v} = (v_j)_{j=1}^m$ be the nonnegative moment-multiplier vectors. Suppose $(\lambda_0, \lambda_1, \mathbf{u}, \mathbf{v})$ is dual feasible at M and has objective within ε_d of $W_F^{[\ell,u_{\max}]}(\gamma; M)$. Then

$$0 \le W_F^{[\ell,u_{\max}]}(\gamma; \overline{M}) - W_F^{[\ell,u_{\max}]}(\gamma; M) \le \varepsilon_d + \mathbf{u}\cdot\Delta^+ + \mathbf{v}\cdot\Delta^-.$$

Thus the displayed multiplier-weighted bound quantifies the effect of outward receipt rounding on the robust privacy value.

Proof. Relaxing moment upper bounds enlarges the primal class, giving the left inequality. The pointwise majorant does not depend on the numerical values of the moment bounds, so the same dual coefficients remain feasible. Their objective increases by exactly $\mathbf{u}\cdot\Delta^+ + \mathbf{v}\cdot\Delta^-$. Weak duality and the stated dual suboptimality give the upper bound. □

D.5 A REUSABLE HISTORICAL PROFILE FROM CHECKED MOMENT BOUNDS

The reusable-profile approach bounds the historical profile before the request is supplied. Write $\mathcal{U}(M)$ for the compatible historical laws and $H(t) = \sup_{\nu\in\mathcal{U}(M)} h_\nu(t)$. Let s_F be the request's singular P_F-mass and $Y = dP_{F,\mathrm{ac}}/dQ_F$. Independence and nonnegative integration give

$$W_F(\gamma; M) = s_F + \sup_{\nu\in\mathcal{U}(M)} \mathbb{E}_{Q_F}[Y h_\nu(\gamma/Y)] \le s_F + \mathbb{E}_{Q_F}[Y H(\gamma/Y)],$$

with zero contribution when $Y = 0$. The left side chooses one history for the complete composed objective; the right side takes a worst history separately at each threshold. Dominating pairs provide the composition interpretation (Zhu et al., 2022, Section 4.1); conditions for commutation with profile suprema appear in Regehr et al. (2026, Theorem 7(ii) and Lemma 3). The construction below gives a checked upper curve for the compact outer class used in the matched-menu comparison.

The direct accountant and the moment-based profile comparison use the same compact outer problem from Theorem 18. Write its support as $[L, U]$, where $L < 1 < U$, and retain the same relaxed moment bounds. For the direct calculation, the objective is $g_{\gamma,F}$. To construct a reusable historical profile, first solve the moment problem with objective $(x-t)_+$ at historical thresholds t, before supplying the request.

Lemma 21 (A continuous bound from convexity)**.** *Let q be the moment dual with nonnegative multipliers and let g be a convex objective on $[L, U]$. For each grid cell $[a, b]$, put $c = (a+b)/2$. If $G(a) \ge g(a)$ and $G(b) \ge g(b)$ are outward upper evaluations, then any nonnegative shift s satisfying*

$$s \ge \max_{\text{cells } [a,b]} \max_{z\in\{a,b\}} \{G(z) - q(c) - q'(c)(z-c)\}$$

makes $q + s$ a continuous majorant of g.

Proof. Both x^α and $x^{1-\alpha}$ are convex for $\alpha > 1$ and $x > 0$. Thus q is convex and its tangent at c is a lower bound throughout the cell. Convexity also places g below the chord joining its upper endpoint evaluations. The chord minus the tangent is affine, so its maximum on the cell is attained at an endpoint. Adding the displayed shift covers this maximum on every cell. Taking expectations under the saved moment bounds then certifies the shifted dual objective. □

At each threshold t_i, independently check a dual for $(x - t_i)_+$ using Lemma 21, and take the smaller of its bound and the endpoint-secant bound $(U - t_i)(1 - L)/(U - L)$. Include the exact points $(0, 1)$, $(L, 1 - L)$, and $(U, 0)$. The lower convex hull of these points and the checked upper bounds defines a piecewise-linear function $\bar{h}_H$, extended by zero beyond U. Its derivative jumps give masses w_i at positive likelihood ratios z_i. The checker reconstructs the hull with rational arithmetic and verifies nonnegative masses, total mass one, and $\sum_i w_i z_i = 1$.

Every feasible compact historical profile is convex and lies below all node bounds. Between any two hull vertices it therefore lies below their chord, so it is dominated by $\bar{h}_H$. The hull slopes lie in $[-1, 0]$ and increase to zero, making $\bar{h}_H$ itself a valid mean-one stop-loss profile. Stop-loss domination on a common mean implies domination of the expectation of every convex objective: a piecewise-linear convex objective is an affine function plus nonnegative hinge functions, and uniform approximation on $[L, U]$ extends the conclusion to continuous convex objectives. Consequently, for the original historical class,

$$W_F(\gamma; M) \le \sum_i w_i z_i h_F(\gamma/z_i) + 2\delta_{\ell,u}.$$

The sum is evaluated with the supplied request oracle. This comparison uses the request after constructing the historical profile and adds the same tail margin as the direct calculation. The finite historical law represented by the hull need not satisfy the receipt moments; its role is to dominate every compatible compact history.