WHEN DOES A REMAINING PRIVACY BUDGET PRESERVE EVERY ADMISSION DECISION?

Anonymous authors
Paper under double-blind review

ABSTRACT

A reusable privacy budget should preserve the choice of the next release. We ask when one experiment represents exactly the requests admitted by the original spent–cap comparison after independent composition. Every cap with a nonempty feasible set has a universal residual if and only if the spend has at most one finite privacy-loss value, with complete reveal as the degenerate case. The key is how composition acts on request profiles. One value rescales the whole profile; several values average different rescalings and can conceal a crossing. We realize this obstruction with valid, bounded-likelihood-ratio experiments for every finite-support multi-atom spend, with ordered compositions and a cap equal to one of those products. Weakly continuous composition semigroups with finite privacy-loss support at each time have the one-value structure, making parameter subtraction exact for arbitrary requests. A concrete Gaussian spent–cap pair has no universal residual and admits a randomized-response channel that its Gaussian-family residual excludes. An exact finite example visualizes the obstruction, and a model-selection illustration separates the additional channel accuracy from its downstream utility. The results identify when a single remaining-budget experiment preserves the choice of the next mechanism.

1 INTRODUCTION

Analysts sharing a private dataset may choose later queries only after seeing earlier answers (Rogers et al., 2016, Sec. 1). We study a one-step question: whether a remaining budget can preserve every feasible choice of the next release. An analyst receiving a remaining-budget experiment could test a proposed release against that experiment's full profile, without carrying the original spend–cap pair. Exact representation requires both safety and completeness: every accepted request must fit under the original cap, and every request that fits must remain available. Safety alone can still force extra noise or exclude a feasible release.

Differential privacy compares output distributions on neighboring datasets (Dwork & Roth, 2014). Concentrated and Rényi formulations provide additive divergence bounds (Bun & Steinke, 2016; Mironov, 2017); optimal and numerical composition methods evaluate specified sequences of releases (Kairouz et al., 2015; Murtagh & Vadhan, 2018; Gopi et al., 2021). Privacy profiles retain the full curve of one-sided privacy guarantees (Balle et al., 2020), and the hypothesis testing formulation supports comparisons between different mechanisms (Dong et al., 2022). These tools answer a forward question: does the spent experiment composed with a proposed request stay below a given cap? Representing what remains asks for an inverse operation before that request is known. The inverse must preserve all of these comparisons at once.

We formulate this requirement in the Blackwell order of binary experiments (Blackwell, 1953). A binary experiment is a pair of output laws under two hypotheses. Write $R \preceq_{\mathrm{B}} E$ when data-independent post-processing of E produces R, and write $C \otimes R$ for independent composition. The experiments C, B, and R bound the privacy of past releases, the total disclosure, and the proposed request, respectively. A *universal residual* is a feasible experiment $R^\star$ satisfying

$$C \otimes R \preceq_{\mathrm{B}} B \quad \Longleftrightarrow \quad R \preceq_{\mathrm{B}} R^\star \qquad \text{for every request } R.$$

Equivalently, it is the greatest element of the feasible request set. The question is therefore whether that set can be represented exactly by one experiment, even when the experiment may lie outside the spend's family.

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.

Gaussian differential privacy (GDP) makes the distinction concrete. The Gaussian experiment G_μ has an additive squared parameter, so a G_1 spend under a $G_{\sqrt{2}}$ cap leaves G_1 for another Gaussian request (Dong et al., 2022). We show that the same spent–cap pair also admits randomized response (RR) with truthful probability 0.70, whereas the Gaussian residual permits at most $\Phi(1/2) \approx 0.6915$ within symmetric RR; Φ is the standard-normal cumulative distribution function. No alternative experiment represents all feasible requests at this pair: the obstruction is the absence of a greatest feasible experiment.

Our main result classifies the spends for which a universal residual exists at every cap with a nonempty feasible set (Theorem 1). Apart from complete reveal, the finite privacy loss must be conditionally constant. The key is how composition acts on request profiles: one finite value applies a common rescaling, whereas several values average different rescalings and can conceal a crossing. The classification covers arbitrary privacy-loss laws.

For every finite-support spend with multiple finite values, we strengthen this obstruction to bounded requests whose compositions are already ordered. Perturbing likelihood-ratio densities while preserving normalization and the mean-one constraint realizes the concealed crossing with valid experiments. The larger product is a cap with no greatest residual, even though equivalence cancellation still holds; an exact rational example makes the loss of ordering visible. Weakly continuous composition semigroups with finite privacy-loss support at each time have the one-value structure, making parameter subtraction exact for arbitrary requests. Finally, the Gaussian example and a model-selection illustration connect the additional admitted choice to channel accuracy and downstream utility.

2 Related Work

Privacy filters govern interactions in which privacy parameters may be chosen after earlier outputs. Rogers et al. (2016) distinguish budget-enforcing filters from privacy-loss odometers; Tran et al. (2026) characterize when filtering by realized trade-off curves remains valid under fully adaptive composition. For a fixed spend and cap, Regehr et al. (2026, Thm. 1) give the residue-update safety condition and obtain the largest safe GDP-family update (Regehr et al., 2026, Thm. 5). Their supremum construction and composition criterion (Regehr et al., 2026, Prop. 11, Thm. 7) motivate the question here: when does one *feasible* residual represent every one-step request, including requests outside the spend's family?

Classical product cancellation recovers a request up to Blackwell equivalence from an equivalent product with a common regular factor (Torgersen, 1977). An exact residual needs the stronger ability to reflect one-sided comparisons and preserve feasible suprema. Related large-sample and catalytic results describe how additional composition can create comparisons: for generic bounded binary experiments, Mu et al. (2021, Thm. 1, Prop. 4) characterize tensor-power and independent-auxiliary comparisons by Rényi divergences. Farooq et al. (2024); Verhagen et al. (2025) extend such comparisons to tuples and varying support restrictions. In quantum settings, Perry et al. (2022) study asymptotic relative submajorization through a preordered semiring; Rethinasamy & Wilde (2020, Thm. 2) treat approximate classical transformations with full-support references and a catalyst whose marginal is returned. Our fixed-spend question has no chosen auxiliary or large-sample limit: it asks whether every cap has a greatest feasible next request.

Composition semigroups provide natural remaining-budget candidates because their parameters add. GDP composes exactly within its Gaussian family (Dong et al., 2022), while Pandey et al. (2025, Thm. 6) connect repeated-composition limits to infinitely divisible privacy losses. We classify weakly continuous semigroups whose privacy-loss support is finite at every time, without a uniform size bound. In this class, the fixed-spend structure makes parameter subtraction exact for arbitrary requests (Theorem 4).

3 Feasible Requests and Universal Residuals

A binary experiment is a pair $E = (P, Q)$ of probability measures on a common standard Borel output space. For a randomized mechanism and an ordered neighboring pair (x, x'), these are the output laws under x and x'. The Blackwell order $E_1 \preceq_{\mathrm{B}} E_2$ means that a single Markov kernel, independent of the hypothesis, maps the laws of E_2 to those of E_1. Thus E_2 is at least as informative

about the neighboring dataset. Mutual domination is written $E_1 \simeq_{\mathrm{B}} E_2$; statements about greatest elements are understood up to this equivalence. Canonical likelihood-ratio representatives retain singular mass; Appendix D gives the reduction used below.

An experiment $\bar{E}$ gives a global privacy bound for a declared mechanism class and adjacency relation when every covered experiment satisfies $E_{x,x'} \preceq_{\mathrm{B}} \bar{E}$. The bounds for past releases, the cap, and the request must cover the same privacy unit and mechanism scope. Completeness below is relative to these specified bounds. Independent composition is $E_1 \otimes E_2 = (P_1 \otimes P_2, Q_1 \otimes Q_2)$ and is monotone in each factor. The identity experiment $\mathrm{I} = (P, P)$ reveals nothing, while $\top$ has disjoint laws and reveals the hypothesis completely.

For current privacy bound C and cap B, *direct admission* accepts a request bound R precisely when $C \otimes R \preceq_{\mathrm{B}} B$. Define

$$\mathcal{R}_C(B) = \{R : C \otimes R \preceq_{\mathrm{B}} B\}.$$

If the actual experiments obey $C_{\mathrm{act}} \preceq_{\mathrm{B}} C$ and $R_{\mathrm{act}} \preceq_{\mathrm{B}} R$, independence and monotonicity give $C_{\mathrm{act}} \otimes R_{\mathrm{act}} \preceq_{\mathrm{B}} C \otimes R \preceq_{\mathrm{B}} B$. The updated privacy bound is $C \otimes R$, with B unchanged. The feasible set is nonempty exactly when $C \preceq_{\mathrm{B}} B$: the identity request is then feasible, and every composed request dominates C.

Every member of $\mathcal{R}_C(B)$ is a *safe residual*, since requests below it are also feasible. A *family-restricted residual* is a greatest member of $\mathcal{R}_C(B) \cap \mathcal{F}$ for a declared family $\mathcal{F}$, when one exists. A *universal residual*, or greatest residual, is a greatest member $R^\star$ of the unrestricted set:

$$C \otimes R \preceq_{\mathrm{B}} B \quad \Longleftrightarrow \quad R \preceq_{\mathrm{B}} R^\star \qquad \text{for every binary experiment } R.$$

Feasibility supplies safety; greatestness supplies completeness. A maximal member need only have no strictly larger feasible member and can still leave other feasible requests incomparable with it. Universal representation requires domination of every feasible member.

3.1 Privacy profiles and composition

Write $P = P_{\mathrm{ac}} + P_\perp$ for the absolutely continuous and singular parts of P relative to Q. The P-side privacy loss, or log-likelihood ratio (LLR), is

$$L_E = \log \frac{dP_{\mathrm{ac}}}{dQ} \quad \text{on common support,} \qquad L_E = +\infty \quad P_\perp\text{-a.s.}$$

Its law under P is the privacy-loss distribution (PLD). With $e^{-\infty} = 0$, the *hockey-stick profile* is

$$h_E(x) = \int (dP - x\, dQ)_+ = \mathbb{E}_P[(1 - xe^{-L_E})_+], \qquad x > 0.$$

At $x = e^\varepsilon$, this is the smallest corresponding one-sided δ bound. Binary-experiment comparison is equivalent to comparison of the entire profiles (Blackwell, 1953; Regehr et al., 2026):

$$E_1 \preceq_{\mathrm{B}} E_2 \quad \Longleftrightarrow \quad h_{E_1}(x) \le h_{E_2}(x) \quad \text{for all } x > 0.$$

The full threshold domain matters: a comparison at selected (ε, δ) values need not establish Blackwell domination.

Hypothesis reversal sends (P, Q) to (Q, P); an experiment is *symmetric* when it is equivalent to its reversal. The symmetric GDP experiment G_μ has P-side privacy loss $N(\mu^2/2, \mu^2)$ and satisfies $G_\mu \otimes G_\nu \simeq_{\mathrm{B}} G_{\sqrt{\mu^2+\nu^2}}$. Let Φ and ϕ be the standard-normal cumulative distribution function and density. Randomized response (Warner, 1965) R_ε reports its binary input truthfully with probability $p = e^\varepsilon/(1 + e^\varepsilon)$, otherwise flipping it. Its P-side privacy loss takes values $+\varepsilon$ and $-\varepsilon$ with probabilities p and $1 - p$. We also write this channel as $R(p)$.

An experiment is *regular* when its hypotheses have positive common mass, equivalently $\Pr_P(L_E < +\infty) > 0$. Torgersen's product theorem gives *equivalence cancellation* for such a common factor (Torgersen, 1977):

$$C \otimes R_1 \simeq_{\mathrm{B}} C \otimes R_2 \quad \Longrightarrow \quad R_1 \simeq_{\mathrm{B}} R_2.$$

Recovering an inequality rather than an equivalence is a stronger property. Section 5 separates these properties.

4 EXISTENCE OF UNIVERSAL PRIVACY RESIDUALS

Fix the spend C and let the cap vary over $B \succeq_{\mathrm{B}} C$. We say that C has the *all-cap existence property* if every cap B with a nonempty feasible set admits a universal residual $R^\star(C, B)$, which may be different for different caps. At each cap, a universal residual must combine all feasible choices without making that combination infeasible. Composition with one finite privacy-loss value applies a common rescaling to every request profile. Several values instead average different rescalings and can conceal a crossing. This distinction gives the general classification.

Theorem 1 (All-cap residual existence)**.** *For any binary experiment C, every nonempty $\mathcal{R}_C(B)$ has a greatest element if and only if either $C \simeq_{\mathrm{B}} \top$ or the conditional law of L_C given $L_C < +\infty$ is a point mass.*

Conditional constancy means that, whenever the spend does not reveal the hypothesis completely, its P-side privacy loss has the same finite value. Revealing mass is allowed. The cap quantifier is essential: a spend outside this class fails at some nonempty cap, not necessarily every cap. For complete reveal, the only nonempty feasible set is at a cap equivalent to $\top$, and its greatest residual is $\top$.

To see the mechanism, put $\theta = \Pr_P(L_C < +\infty) > 0$ and let W have the conditional law of e^{-L_C} given finite loss. Independent composition adds privacy losses, so conditioning on the spend gives

$$h_{C\otimes R}(x) = 1 - \theta + \theta\, \mathbb{E}[h_R(xW)], \qquad x > 0.$$

The revealing event contributes the same constant for every request; the finite loss chooses the threshold rescaling $x \mapsto xW$. When $W = e^{-c}$ is constant, write $p = \theta$. The spend has the three-output representative

$$P = (p, 1-p, 0), \qquad Q = (pe^{-c}, 0, 1 - pe^{-c}), \quad 0 < p \le 1, \quad pe^{-c} \le 1.$$

The common output has privacy loss c; the P-only output has loss $+\infty$; the Q-only output has zero P-probability. Its action is $h_{C\otimes R}(x) = 1 - p + ph_R(xe^{-c})$. At $c = 0$, this is a binary erasure experiment with common-output probability p. Nonzero c changes the odds on that output but retains a single threshold rescaling.

For any nonempty family $\mathcal{A}$, its *supremum* $\bigvee \mathcal{A}$ is the least experiment dominating every member. Its profile is $\sup_{R\in\mathcal{A}} h_R$ (Regehr et al., 2026, Prop. 11); for two requests this is their pointwise maximum. Thus $R^\star = \bigvee \mathcal{R}_C(B)$ is a valid experiment and is a universal residual precisely when it remains feasible. A single rescaling ensures feasibility:

$$\begin{aligned} h_{C\otimes R^\star}(x) &= 1 - p + p \sup_{R\in\mathcal{R}_C(B)} h_R(xe^{-c}) \\ &= \sup_{R\in\mathcal{R}_C(B)} h_{C\otimes R}(x) \le h_B(x). \end{aligned}$$

Write $T_C(R) = C \otimes R$. The following fixed-spend form of the residue criterion of Regehr et al. (2026, Thm. 7) shows that preservation of the supremum is also necessary for all-cap existence.

Proposition 2. *For a fixed factor C, the following are equivalent: every nonempty feasible set $\mathcal{R}_C(B)$ has a greatest element; and T_C preserves every nonempty supremum, meaning*

$$T_C\Big(\bigvee \mathcal{A}\Big) \simeq_{\mathrm{B}} \bigvee_{R\in\mathcal{A}} T_C(R) \quad \textit{for every nonempty } \mathcal{A}.$$

Equivalently, the first condition need only be required for $B \succeq_{\mathrm{B}} C$.

Appendix D.2 proves the equivalence by applying all-cap existence to the supremum of the composed requests; Appendix C gives the associated residual representation.

A nonconstant W breaks this preservation. Choose two request profiles that cross at a threshold with positive probability on both sides under the law of W. Taking their maximum before averaging retains the larger value on each side; averaging first loses that advantage. For the requests A, D constructed in Appendix D.3,

$$B = (C \otimes A) \vee (C \otimes D), \qquad h_{C\otimes(A\vee D)}(1) > h_B(1).$$

Both requests are feasible at B, but their least common upper bound is not. No greatest feasible residual can therefore exist. This proves the negative direction of Theorem 1 for arbitrary nonconstant finite-loss laws, using the supremum and gap mechanism of Regehr et al. (2026, Props. 11, 13). For finite-support spends, the next section strengthens the obstruction to bounded requests and a product cap.

5 A Stronger Obstruction for Finite-Support Spends

For a spend with several finite privacy-loss values, we construct bounded request likelihood ratios whose compositions are already ordered. The larger product itself is then a cap with no universal residual, rather than a supremum of products as in the general proof.

Let δ_c denote a point mass at c. We say that T_C *reflects order* if

$$C \otimes R_1 \preceq_{\mathrm{B}} C \otimes R_2 \quad \Longrightarrow \quad R_1 \preceq_{\mathrm{B}} R_2.$$

Unlike equivalence cancellation, this property recovers the direction of a one-sided comparison. The finite-support classification separates the two properties from all-cap residual existence.

Theorem 3 (Finite-support classification)**.** *Let the P-side privacy-loss distribution of C be*

$$\kappa\delta_{+\infty} + \sum_{i=1}^{m} p_i \delta_{c_i}, \qquad \kappa \geq 0, \quad p_i > 0, \quad c_1 < \cdots < c_m, \quad \kappa + \sum_i p_i = 1,$$

where m may be zero. Equivalence cancellation holds if and only if $m \geq 1$. Blackwell-order reflection holds if and only if $m = 1$. Every nonempty feasible residual set $\mathcal{R}_C(B)$ has a greatest element if and only if $m \leq 1$.

When $m = 0$, $C \simeq_{\mathrm{B}} \top$ and residual existence is degenerate. For a regular finite-support spend, order reflection and all-cap residual existence hold exactly when $m = 1$. A multi-atom spend still cancels equivalent products, even though it can lose the ordering of requests. Some caps remain representable: at $B = C$, feasibility and equivalence cancellation force the identity residual I; at $B = \top$, every request is feasible. The construction below identifies an intermediate product cap where two feasible requests have no feasible common upper bound.

5.1 Crossing profiles from valid requests

For the P-side PLD $\kappa\delta_{+\infty} + \sum_{i=1}^{m} p_i \delta_{c_i}$, independent composition adds privacy losses, giving

$$h_{C\otimes R}(x) = \kappa + \sum_{i=1}^{m} p_i h_R(x e^{-c_i}). \tag{1}$$

Thus each finite atom rescales the threshold of the request profile; revealing mass contributes the same constant to every request. Appendix D.1 derives the identity, including singular mass.

For several atoms, the weighted sum in (1) can conceal a crossing: composition offsets a positive profile difference at one threshold by negative differences at others, as Figure 1 illustrates. Our construction realizes this crossing with valid requests for every $m \geq 2$, with likelihood ratios bounded above and away from zero (Theorem 5). We seek a smooth function d that changes sign but becomes nonpositive after composition; realizing a scaled d as a difference of request profiles will then give the desired experiments. On the log-threshold axis, write $f(t) = d(e^t)$; composition acts on such differences by

$$(\mathcal{T}_C f)(t) = \sum_{i=1}^{m} p_i f(t - c_i).$$

Set $a = c_2 - c_1 > 0$ and choose $0 < \eta < p_2/p_1$. For a nonnegative, smooth bump φ with sufficiently narrow compact support around zero, define $b_j = a + c_j - c_1$ for $j = 2, \ldots, m$ and set

$$f(t) = -\varphi(t) + \eta\varphi(t - a) - \eta \sum_{j=2}^{m} \frac{p_j}{p_1} \varphi(t - b_j).$$

The input centers $0, a, b_2, \ldots, b_m$ are distinct. Taking $\varphi(0) > 0$ gives $f(0) < 0$ and $f(a) > 0$, so the input difference crosses. Only the bump at a is positive. Its translate by c_1 combines with the negative base bump at c_2, giving coefficient $\eta p_1 - p_2 < 0$. Each remaining positive translate is canceled by the compensator at b_j; all other terms are nonpositive. Hence $\mathcal{T}_C f \leq 0$ everywhere. This works for arbitrary atom locations and weights, including coincident output centers; Appendix D.6 checks the coefficients.

The sign pattern alone does not yet define requests; it must respect the mass and mean constraints of a likelihood-ratio law. To realize it as a profile difference, put $d(x) = f(\log x)$. Choose $0 < r < 1 < s$ so that (r, s) contains the supports of d and d'', and a smooth likelihood-ratio density q_0 on $[r, s]$ with mass and mean both one, positive throughout (r, s). Such a density is obtained by exponentially tilting a smooth positive density on (r, s): its mean varies continuously between the endpoints, so one tilt has mean one. Compact support gives

$$\int d''(z)\, dz = 0, \qquad \int z d''(z)\, dz = 0.$$

For sufficiently small $\rho > 0$, the densities $q_\pm = q_0 \pm (\rho/2)d''$ are nonnegative and still have mass and mean one. The mass condition makes $q_\pm$ probability laws under Q; the mean-one condition makes their likelihood-ratio-weighted versions probability laws under P. Therefore $Q_\pm(dz) = q_\pm(z)\, dz$ and $P_\pm(dz) = z q_\pm(z)\, dz$ define valid experiments, and integration by parts gives

$$h_{(P_+,Q_+)}(x) - h_{(P_-,Q_-)}(x) = \rho d(x).$$

Set $R_1 = (P_+, Q_+)$ and $R_2 = (P_-, Q_-)$. Their profiles cross, while their compositions are ordered. Taking $B = C \otimes R_2$ makes both requests feasible. Any greatest feasible $R^\star$ would dominate R_2, giving

$$C \otimes R_2 \preceq_{\mathrm{B}} C \otimes R^\star \preceq_{\mathrm{B}} B = C \otimes R_2.$$

Equivalence cancellation would force $R^\star \simeq_{\mathrm{B}} R_2$, contradicting $R_1 \not\preceq_{\mathrm{B}} R_2$. The realized crossing therefore rules out every candidate universal residual at this cap. Cancellation of equivalent products proves nonexistence here; it does not recover the inequality lost in composition. Appendices D.6 and D.7 give the complete proofs.

5.2 An exact finite obstruction

The following finite example makes the construction visible without approximation. Represent an experiment by the Q-side law of its likelihood ratio dP/dQ. Each pair lists a likelihood ratio and its Q-probability:

$$\begin{aligned} C &: \{(1/2, 2/3), (2, 1/3)\}, \\ R_1 &: \{(2/5, 5/7), (2, 1/7), (3, 1/7)\}, \\ R_2 &: \{(1/5, 10/21), (1, 3/7), (5, 2/21)\}. \end{aligned}$$

Each law has unit mass and likelihood-ratio mean one. The request profiles cross, since

$$h_{R_1}(2/5) - h_{R_2}(2/5) = -\frac{2}{21}, \qquad h_{R_1}(1) - h_{R_2}(1) = \frac{1}{21}.$$

To apply (1), convert the spend's Q-probabilities to P-probabilities by multiplying each mass by its likelihood ratio. The two P-side masses are $1/3$ and $2/3$, at privacy losses $-\log 2$ and $\log 2$. Thus, for any request in this example,

$$h_{C \otimes R}(x) = \frac{1}{3} h_R(2x) + \frac{2}{3} h_R(x/2).$$

The two different threshold rescalings expose different parts of the original crossing. Their weighted average can be nonpositive at every threshold even though neither request profile lies below the other. Their product profiles satisfy $C \otimes R_1 \preceq_{\mathrm{B}} C \otimes R_2$. Finite likelihood-ratio profiles are piecewise linear, so the complete breakpoint calculation in Appendix B.1 proves the inequality on the whole domain, rather than on a grid.

Figure 1 also explains why simply combining the feasible requests into a larger experiment does not repair the residual. The *supremum* $R_1 \vee R_2$ is the least experiment dominating both; its profile

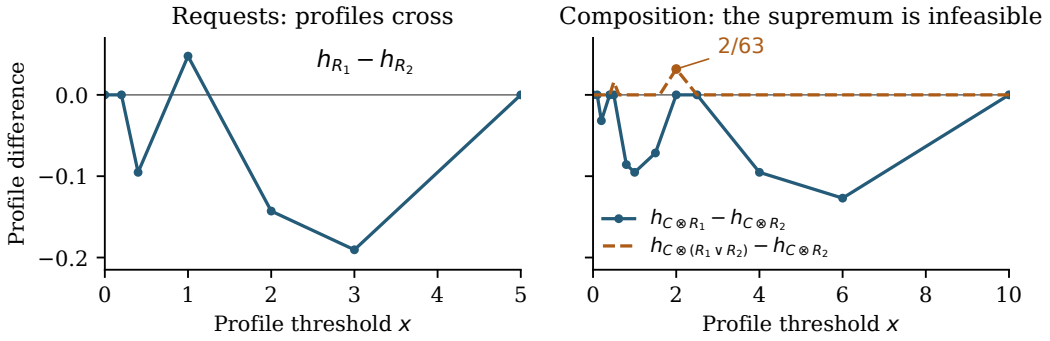

Figure 1: Profile differences for the exact rational example. The request profiles cross (left). After composition (right), $C \otimes R_1 \preceq_{\mathrm{B}} C \otimes R_2$ (solid), but $C \otimes (R_1 \vee R_2)$ exceeds the cap $C \otimes R_2$ (dashed) by $2/63$ at $x = 2$. Here $R_1 \vee R_2$ is the least experiment dominating both requests. Curves connect all breakpoints.

is $\max(h_{R_1}, h_{R_2})$. At threshold 2, composition weights the original profile differences at 1 and 4. Direct computation gives

$$h_{C\otimes R_1}(2) - h_{C\otimes R_2}(2) = 0.$$

Taking the maximum first retains the positive difference at 1 and removes the negative difference at 4. Consequently,

$$h_{C\otimes(R_1\vee R_2)}(2) - h_{C\otimes R_2}(2) = \frac{2}{63} > 0.$$

The least experiment containing both feasible choices is itself infeasible. Averaging and taking the pointwise maximum do not commute.

6 Universal Parameter Subtraction

A *tensor semigroup* is a family $\{C_t\}_{t\geq 0}$ with $C_0 \simeq_{\mathrm{B}} \mathrm{I}$ and $C_s \otimes C_t \simeq_{\mathrm{B}} C_{s+t}$. An additive parameter identifies a natural candidate remaining budget: $C_s \otimes C_{t-s} \simeq_{\mathrm{B}} C_t$ for $0 \leq s \leq t$. For it to represent arbitrary requests, the stronger equivalence must hold:

$$C_s \otimes R \preceq_{\mathrm{B}} C_t \quad \Longleftrightarrow \quad R \preceq_{\mathrm{B}} C_{t-s} \qquad \text{for every } 0 \leq s \leq t \text{ and every } R. \tag{2}$$

Order reflection of C_s supplies this step. It turns a family composition identity into a statement that covers requests outside the family as well.

The semigroup is *self-dual* when every C_t is equivalent to its hypothesis reversal. It is *weakly continuous* when $t_n \to t$ gives weak convergence of the P-side PLDs on $\mathbb{R} \cup \{+\infty\}$, including convergence of the revealing-event mass. Weak continuity together with finite support at every time forces the structure needed for universal subtraction.

Theorem 4 (Finite-support semigroups)**.** *Let $\{C_t\}_{t\geq 0}$ be a weakly continuous tensor semigroup with $C_0 = \mathrm{I}$, and suppose every P-side privacy-loss distribution has finite support. Then there are $\lambda \geq 0$ and $b \in \mathbb{R}$ with $\lambda + b \geq 0$ such that*

$$L_{C_t} = \begin{cases} bt, & \text{with probability } e^{-\lambda t}, \\ +\infty, & \text{with probability } 1 - e^{-\lambda t}. \end{cases}$$

Every positive-time factor therefore has one finite atom and possible revealing mass, and parameter subtraction satisfies (2) *for every request. If the semigroup is self-dual, then $b = 0$, leaving the erasure semigroup.*

The finite-support hypothesis has a specific role. Let p_t be the probability of finite privacy loss. The semigroup law gives $p_{s+t} = p_s p_t$, and continuity at the identity gives $p_t = e^{-\lambda t} > 0$. Conditional on finite loss, the law at time t is the n-fold convolution of the law at time t/n. If the smaller-time law had two points $u < v$, the convolution would contain the $n + 1$ distinct values

$$nu, (n-1)u + v, \ldots, u + (n-1)v, nv.$$

Choosing n larger than the finite support size at time t rules this out. The conditional law must be a point mass. Its location adds under composition, and continuity makes it bt. This argument needs finite support at each time but no uniform bound on the support size across time. Appendix D.8 supplies the proof and the constraints on the revealing masses.

Here λ is the rate at which P-side common mass disappears, and b is the finite privacy-loss drift per unit parameter. The Q-side common mass is $e^{-(\lambda+b)t}$, explaining $\lambda + b \geq 0$. Self-duality makes the two revealing masses equal, forcing $b = 0$. In that case C_t outputs an erasure symbol with probability $e^{-\lambda t}$ and otherwise reveals the hypothesis. Spending C_s under cap C_t leaves an erasure experiment with common-output probability $e^{-\lambda(t-s)}$, which represents arbitrary binary requests by order reflection.

The Gaussian family also has an additive parameter $t = \mu^2$, but its continuous privacy-loss support lies outside this theorem. The next section shows that its family residual can fail to preserve feasible requests.

7 Gaussian Admission and Channel Utility

Take spend $C = G_1$ and cap $B = G_{\sqrt{2}}$. Subtracting the squared Gaussian parameters gives G_1, the exact family-restricted residual. It safely admits every post-processing of G_1, but the following feasible randomized-response request lies outside that set.

Example 1. Let $\varepsilon = \log(7/3)$. Then

$$G_1 \otimes R_\varepsilon \preceq_{\mathrm{B}} G_{\sqrt{2}}, \qquad R_\varepsilon \not\preceq_{\mathrm{B}} G_1.$$

Moreover, the feasible residual set $\{R : G_1 \otimes R \preceq_{\mathrm{B}} G_{\sqrt{2}}\}$ has no greatest element.

The admission inequality is a comparison over all profile thresholds, proved by a single-crossing argument for explicit Gaussian-mixture privacy-loss densities in Appendix D.5. Non-domination is visible at one threshold:

$$h_{R(0.70)}(1) = \frac{2}{5} > 2\Phi(1/2) - 1 = h_{G_1}(1).$$

The product-cap argument from Section 5 also applies here. Suppose a greatest feasible $R^\star$ existed. Since G_1 itself is feasible, it would satisfy

$$G_1 \otimes G_1 \preceq_{\mathrm{B}} G_1 \otimes R^\star \preceq_{\mathrm{B}} G_{\sqrt{2}} \simeq_{\mathrm{B}} G_1 \otimes G_1.$$

Equivalence cancellation of the regular factor G_1 then gives $R^\star \simeq_{\mathrm{B}} G_1$, contradicting the feasibility of $R(0.70)$. The obstruction persists even when the candidate residual is unrestricted.

Table 1: Separate next-request choices under spend G_1 and cap $G_{\sqrt{2}}$. Each residual is safe, but neither preserves both choices.

Request R	Direct admission $G_1 \otimes R \preceq_{\mathrm{B}} G_{\sqrt{2}}$	Gaussian residual $R \preceq_{\mathrm{B}} G_1$	RR residual $R \preceq_{\mathrm{B}} R(0.70)$
Gaussian G_1	accept	accept	not admitted
RR $R(0.70)$	accept	not admitted	accept

Exporting $R(0.70)$ instead of G_1 reverses which request is lost in Table 1: its profile vanishes at $x = 7/3$, whereas $h_{G_1}(7/3) > 0$ because Gaussian privacy loss has an unbounded positive tail. The rows are separate possible next releases, not a sequence. Example 1 rules out every candidate universal residual, not just these two.

7.1 Channel accuracy and model-selection utility

Thresholding the G_1 privacy-loss statistic at zero produces $R(p_\star)$ with $p_\star = \Phi(1/2)$. Since $h_E(1)$ is total variation and cannot increase under post-processing, $|2p - 1| \leq 2\Phi(1/2) - 1$ is necessary for $R(p) \preceq_{\mathrm{B}} G_1$. Thus $p_\star$ is the largest truthful probability within symmetric RR allowed by the

Gaussian residual. Direct admission permits the more accurate $R(0.70)$ without changing the current privacy bound or cap.

Both policies below use the same selector and spent–cap pair; they differ only in the channel that releases the selected bit.

Table 2: Release of the same private bit under spend G_1 and cap $G_{\sqrt{2}}$. Lower equal-prior bit error is better. Values follow from the channel probabilities, not from model-fitting repetitions.

Admission rule	Executed channel	Bit error (%)
Direct spent–cap comparison	$R(0.70)$	**30.000**
Gaussian-residual policy	$R(p_\star)$	30.854

Table 2 reports the equal-prior bit error $1 - p$. The residual channel is obtained from $R(0.70)$ by an independent output flip with probability

$$q = \frac{0.70 - p_\star}{2(0.70) - 1} \approx 2.134\%.$$

Thus the loss imposed by this residual has an explicit post-processing interpretation. Lemma 8 gives the general formula.

For model selection, the candidate models are fitted using public data; the privacy guarantee covers the validation data used to choose between them. Let $S(D) \in \{0, 1\}$ be any possibly randomized selector on the private validation table D. Releasing $S(D)$ through RR with fresh randomness has privacy bound $R(p)$, independently of how the selector was constructed. The kernel in Appendix A.2 explicitly obtains the adjacent-dataset output laws by post-processing $R(p)$. The selected public model can subsequently be used without another access to the validation data.

Task utility also depends on what the input bit represents. In repetition i, let $S_i \in \{0, 1\}$ be a private selector's choice and let $U_i(j)$ be option j's evaluation utility, with larger values better. Averaging only RR randomness gives $\bar{U}_i(p) = pU_i(S_i) + (1 - p)U_i(1 - S_i)$, hence

$$\bar{U}_i(0.70) - \bar{U}_i(p_\star) = (0.70 - p_\star)\big[U_i(S_i) - U_i(1 - S_i)\big]. \tag{3}$$

Privacy admission fixes the coefficient; the selector and task determine the signed margin. A more truthful channel releases better choices more faithfully, but also worse ones. Task utility can therefore increase or decrease.

A constructed cost-sensitive Wine case uses two thresholds fitted on public data and fixed before private validation. Across 1000 prespecified split–context repetitions, the direct RR(0.70) policy gains 0.0215 percentage points over residual RR$(p_\star)$, with a 95% Monte Carlo interval of $[0.0181, 0.0249]$ on the fixed dataset. A public fixed-option rule reaches 97.2029% weighted accuracy, versus 97.0101% for direct RR: more accurate release of the selected bit does not by itself make private selection better than that rule. Appendix A gives the protocol, privacy proof, and policy utilities.

8 Conclusion

A universal residual is one greatest feasible experiment representing every next request. It exists at every feasible cap precisely when the spend has one finite privacy-loss value, with complete reveal as the degenerate case. For every finite-support multi-atom spend, bounded requests can have ordered compositions under a cap with no greatest residual. The semigroup result identifies families where parameter subtraction is exact for arbitrary requests. The Gaussian example distinguishes such a universal guarantee from a safe family residual; the selection identity separates channel accuracy from task utility.

A proposed request can still be checked by comparing its composition with the spend against the cap, even when no universal residual represents all future requests. Extending the bounded-request construction to continuous multi-value privacy-loss laws calls for controlling translated profiles while preserving the likelihood-ratio moment constraints.

Reproducibility. Appendix A specifies the model-selection protocol and its privacy guarantee. Appendix B gives the complete rational breakpoint calculation; Appendix D contains the proofs of the main results and supporting lemmas.

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

Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy profiles and amplification by subsampling. *Journal of Privacy and Confidentiality*, 10(1), 2020. doi: 10.29012/jpc.726. URL https://doi.org/10.29012/jpc.726.

David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953. doi: 10.1214/aoms/1177729032. URL https://doi.org/10.1214/aoms/1177729032.

Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In Martin Hirt and Adam Smith (eds.), *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. URL https://doi.org/10.1007/978-3-662-53641-4_24.

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. URL https://doi.org/10.1111/rssb.12454.

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. URL https://doi.org/10.1561/0400000042.

Muhammad Usman Farooq, Tobias Fritz, Erkka Haapasalo, and Marco Tomamichel. Matrix majorization in large samples. *IEEE Transactions on Information Theory*, 70(5):3118–3144, 2024. doi: 10.1109/TIT.2024.3352088. URL https://doi.org/10.1109/TIT.2024.3352088.

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. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper/2021/hash/6097d8f3714205740f30debe1166744e-Abstract.html.

Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1376–1385. PMLR, 2015. URL https://proceedings.mlr.press/v37/kairouz15.html.

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

Xiaosheng Mu, Luciano Pomatto, Philipp Strack, and Omer Tamuz. From Blackwell dominance in large samples to Rényi divergences and back again. *Econometrica*, 89(1):475–506, 2021. doi: 10.3982/ECTA17548. URL https://doi.org/10.3982/ECTA17548.

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. URL https://theoryofcomputing.org/articles/v014a008/.

Aaradhya Pandey, Arian Maleki, and Sanjeev Kulkarni. Infinitely divisible privacy and beyond I: Resolution of the $s^2 = 2k$ conjecture. *arXiv preprint arXiv:2512.00734*, 2025. URL https://arxiv.org/abs/2512.00734.

Christopher Perry, Péter Vrana, and Albert H. Werner. The semiring of dichotomies and asymptotic relative submajorization. *IEEE Transactions on Information Theory*, 68(1):311–321, 2022. doi: 10.1109/TIT.2021.3117440. URL https://doi.org/10.1109/TIT.2021.3117440.

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. *arXiv preprint arXiv:2602.15815*, 2026. URL https://arxiv.org/abs/2602.15815.

Soorya Rethinasamy and Mark M. Wilde. Relative entropy and catalytic relative majorization. *Physical Review Research*, 2(3):033455, 2020. doi: 10.1103/PhysRevResearch.2.033455. URL https://doi.org/10.1103/PhysRevResearch.2.033455.

Ryan M. Rogers, Aaron Roth, Jonathan Ullman, and Salil Vadhan. Privacy odometers and filters: Pay-as-You-Go composition. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper/2016/hash/58c54802a9fb9526cd0923353a34a7ae-Abstract.html.

Erik N. Torgersen. Mixtures and products of dominated experiments. *The Annals of Statistics*, 5 (1):44–64, 1977. doi: 10.1214/aos/1176343739. URL https://doi.org/10.1214/aos/1176343739.

Long Tran, Antti Koskela, Ossi Räisä, and Antti Honkela. f-differential privacy filters: Validity and approximate solutions. *arXiv preprint arXiv:2602.06756*, 2026. URL https://arxiv.org/abs/2602.06756.

Frits Verhagen, Marco Tomamichel, and Erkka Haapasalo. Matrix majorization in large samples with varying support restrictions. *IEEE Transactions on Information Theory*, 71(9):6517–6545, 2025. doi: 10.1109/TIT.2025.3585062. URL https://doi.org/10.1109/TIT.2025.3585062.

Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965. doi: 10.1080/01621459.1965.10480775. URL https://doi.org/10.1080/01621459.1965.10480775.

Appendix A gives the selection protocol and privacy guarantee, and Appendix B the exact finite calculation. Appendix C states additional structural results; Appendix D collects their proofs and the proofs of the main results.

A Selection Protocol and Request Privacy Bounds

The constructed Wine case uses the current privacy bound $C = G_1$, cap $B = G_{\sqrt{2}}$, and the same validation selector under both policies. Direct admission releases the selected bit through $R(0.70)$ and updates the current privacy bound to $G_1 \otimes R(0.70)$. The Gaussian-residual policy uses $R(p_\star)$, with $p_\star = \Phi(1/2)$, and updates it to $G_1 \otimes R(p_\star)$. Example 1 establishes the first admission. Thresholding G_1 at zero gives the second channel, so tensor monotonicity establishes its admission as well.

A.1 Data, options, and utility

The case uses scikit-learn 1.9.0 Wine classes 1 and 2, remapped to a 119-row, 13-feature binary table. In each of 1000 repetitions, a stratified per-class permutation assigns one third of each class, rounded down, to public fitting and divides the remainder evenly between validation and evaluation. This yields 39 public-fitting, 40 private-validation, and 40 evaluation rows. Standardized logistic regression with inverse regularization strength 1 is fitted on the public rows; feature means and scales are fitted on those same rows. The two options use probability thresholds 0.25 and 0.75, fixed before private validation; they share the same fitted model and trade the two class errors differently.

Each repetition draws uniformly between two mirrored cost contexts, with class weights $0.85/0.15$ and $0.15/0.85$. Utility is the weighted mean of the two class accuracies, with higher utility better. The selector chooses the threshold with greater weighted validation accuracy, breaking ties in favor of option 0. The same utility on the disjoint evaluation split defines $U_i(0)$ and $U_i(1)$. Both policies share the split, fitted model, thresholds, cost context, and selector. The context is available within the protected validation boundary but not to the downstream deployment rule; evaluation is used only for offline measurement. The thresholds, contexts, split rule, and metric were specified before fitting.

RR randomness is marginalized analytically using (3). The interval reported in Section 7.1 is a Student-t interval for the mean paired difference over the specified split/context repetitions, with the dataset fixed. It measures Monte Carlo variability of this protocol, not population uncertainty from a newly sampled dataset. A fixed-option rule uses no private validation information or hidden context; its higher mean utility here results from release randomization exceeding the selector's positive mean gain. The unrandomized selected option averages 97.7644%, compared with 97.2029% for fixed option 1 and 97.0101% after RR(0.70). The unrandomized value uses private validation information and is a diagnostic of selection quality.

Table 3: Mean weighted class accuracy on the fixed Wine dataset across 1000 repetitions. The two RR policies share the validation selector; the fixed rule always chooses the second threshold. Higher is better.

Rule	Mean utility (%)
Direct RR(0.70)	97.0101
Residual RR$(p_\star)$	96.9886
Fixed option 1	97.2029

A.2 Privacy guarantee for the selection release

Let S be a possibly randomized selector, and write $a = \Pr[S(D) = 1]$ and $b = \Pr[S(D') = 1]$ for adjacent validation datasets. The mechanism releases only Y, obtained by applying $R(p)$ to S with fresh randomness. Starting from $R(p) = (P_1, P_0)$, post-process its output with the kernel

$$T(1 \mid 1) = pa + (1-p)(1-b), \qquad T(1 \mid 0) = (1-p)(1-a) + pb.$$

Both entries are probabilities, and multiplication gives $P_1 T = \mathcal{L}(Y \mid D)$ and $P_0 T = \mathcal{L}(Y \mid D')$. Thus the selection release is Blackwell-dominated by $R(p)$ for every covered adjacent pair. The publicly fitted options and subsequent use of the released choice are post-processing. This supplies $R_{\mathrm{act}} \preceq_{\mathrm{B}} R(p)$ in the comparison of privacy bounds in Section 3.

B Exact Calculation for the Finite Obstruction

For an experiment represented by Q-side likelihood-ratio atoms (z_j, q_j), $h_E(x) = \sum_j q_j(z_j - x)_+$. Independent composition multiplies likelihood ratios and their Q-probabilities. This gives an exact finite calculation for the example in Section 5.2.

B.1 Breakpoint check for the rational obstruction

For the laws in Section 5.2, define $\Delta_C(x) = h_{C\otimes R_1}(x) - h_{C\otimes R_2}(x)$. The union of all product breakpoints and the exact differences there are

x	$\Delta_C(x)$	x	$\Delta_C(x)$
0	0	$3/2$	$-1/14$
$1/10$	0	2	0
$1/5$	$-2/63$	$5/2$	0
$2/5$	0	4	$-2/21$
$1/2$	0	6	$-8/63$
$4/5$	$-3/35$	10	0
1	$-2/21$		

Every entry is nonpositive. Linear interpolation between consecutive kinks proves $C\otimes R_1 \preceq_{\mathrm{B}} C\otimes R_2$ over the full profile domain.

C Additional Residual Structure

For $B \succeq_{\mathrm{B}} C$, denote a greatest residual by $B \oslash C$, when it exists. These caps form the upper cone above C. The equivalence

$$C \otimes R \preceq_{\mathrm{B}} B \quad \Longleftrightarrow \quad R \preceq_{\mathrm{B}} B \oslash C$$

says that the residual map is a partial right adjoint of composition on this cone. Proposition 2 characterizes its existence by preservation of nonempty suprema. This is a representation of the full feasible set by a single experiment.

The construction underlying Theorem 3 gives the following stronger description of the requests that establish failure.

Theorem 5. *Let the P-side privacy-loss distribution of C be*

$$\kappa\delta_{+\infty} + \sum_{i=1}^{m} p_i\delta_{c_i}, \qquad \kappa \geq 0, \quad p_i > 0, \quad c_1 < \cdots < c_m, \quad \kappa + \sum_i p_i = 1,$$

where m may be zero. The map $R \mapsto C \otimes R$ reflects Blackwell order if and only if $m = 1$. If $m \geq 2$, there are mutually absolutely continuous request experiments R_1, R_2 whose likelihood ratios are bounded above and away from zero, such that

$$R_1 \not\preceq_{\mathrm{B}} R_2, \qquad R_2 \not\preceq_{\mathrm{B}} R_1, \qquad C \otimes R_1 \preceq_{\mathrm{B}} C \otimes R_2.$$

For $m = 0$, $C \simeq_{\mathrm{B}} \top$ and its tensor map is constant.

The likelihood-ratio bounds show that failure does not require singular or unbounded request experiments. Singular mass in the spend contributes a common constant and does not affect the crossing construction.

The one-atom formula collects the positive consequences in a form that also applies to the semigroup calculation.

Theorem 6. *Suppose the P-side privacy-loss distribution of C is $p\delta_c + (1-p)\delta_{+\infty}$ with $p > 0$. Then, for every request R,*

$$h_{C\otimes R}(x) = (1-p) + p h_R(xe^{-c}).$$

Consequently T_C reflects Blackwell order, preserves every nonempty supremum, cancels equivalence, and gives a greatest residual for every cap whose feasible residual set is nonempty. If factors of this form constitute a tensor semigroup, parameter subtraction has exact admission semantics for every request experiment.

For the semigroup conclusion, $C_t \simeq_{\mathrm{B}} C_s \otimes C_{t-s}$ supplies the candidate residual; order reflection of C_s makes its comparison necessary as well as sufficient.

D Complete Proofs

Canonical experiments. Every binary pair is dominated by $M = (P+Q)/2$. Put $Z = dP/d(P+Q) \in [0,1]$ and let ν be its law under M. The induced pair is $P_Z(dz) = 2z\nu(dz)$ and $Q_Z(dz) = 2(1-z)\nu(dz)$. On a standard Borel space, the conditional law of the output given Z under M is a Markov kernel that reconstructs both P and Q, proving Blackwell equivalence. Atoms at 1 and 0 retain the P-exclusive and Q-exclusive masses. Thus dominated-experiment results apply without mutual absolute continuity: Torgersen's regularity condition requires only positive common mass. The nonempty profile suprema of Regehr et al. (2026, Prop. 11) also have canonical representatives, so taking a supremum stays in this experiment class.

Throughout, $E_1 \preceq_{\mathrm{B}} E_2$ means $h_{E_1} \le h_{E_2}$ pointwise, and equality of experiments is understood up to Blackwell equivalence. Suprema have pointwise supremum profiles.

D.1 Composition and output randomization

Lemma 7. *Suppose the P-side PLD of C is*

$$\kappa\delta_{+\infty} + \sum_{i=1}^{m} p_i \delta_{c_i}.$$

Then every request experiment R satisfies

$$h_{C\otimes R}(x) = \kappa + \sum_{i=1}^{m} p_i h_R(xe^{-c_i}). \tag{D.1}$$

Consequently, for any R_1, R_2, the revealing mass cancels:

$$h_{C\otimes R_1}(x) - h_{C\otimes R_2}(x) = \sum_{i=1}^{m} p_i \{h_{R_1}(xe^{-c_i}) - h_{R_2}(xe^{-c_i})\}. \tag{D.2}$$

Proof. Independent composition adds P-side privacy losses. On the event $L_C = +\infty$, the hockey-stick integrand equals one. Conditional on $L_C = c_i$, its expectation is $h_R(xe^{-c_i})$. Averaging gives (D.1), and subtracting two instances of (D.1) gives (D.2). □

The next elementary garbling supplies the channel comparison in Section 7.1.

Lemma 8. *If $1/2 \le p_1 \le p_2 < 1$, then $R(p_1) \preceq_{\mathrm{B}} R(p_2)$.*

Proof. If $p_2 = 1/2$, then $p_1 = p_2$ and the experiments are identical. Otherwise, set $q = (p_2 - p_1)/(2p_2 - 1) \in [0, 1/2]$ and independently flip the output of $R(p_2)$ with probability q. Its truthful probability becomes $p_2(1-q) + (1-p_2)q = p_1$, giving exactly $R(p_1)$. □

D.2 Proof of Proposition 2

Proof. The feasible set $\mathcal{R}_C(B)$ is nonempty exactly when $C \preceq_{\mathrm{B}} B$. Indeed, $\mathrm{I} \preceq_{\mathrm{B}} R$ for every R, so feasibility implies $C = C \otimes \mathrm{I} \preceq_{\mathrm{B}} C \otimes R \preceq_{\mathrm{B}} B$; conversely, I is feasible whenever $C \preceq_{\mathrm{B}} B$.

First suppose every nonempty feasible set has a greatest element, denoted $G_C(B)$. Greatestness and monotonicity give the adjunction

$$C \otimes R \preceq_{\mathrm{B}} B \quad \Longleftrightarrow \quad R \preceq_{\mathrm{B}} G_C(B), \qquad B \succeq_{\mathrm{B}} C. \tag{D.3}$$

Let $\mathcal{A}$ be a nonempty family and put $B_{\mathcal{A}} = \bigvee_{R\in\mathcal{A}} T_C(R)$. Every $R \in \mathcal{A}$ satisfies $R \preceq_{\mathrm{B}} G_C(B_{\mathcal{A}})$, so

$$T_C\Big(\bigvee \mathcal{A}\Big) \preceq_{\mathrm{B}} T_C(G_C(B_{\mathcal{A}})) \preceq_{\mathrm{B}} B_{\mathcal{A}}.$$

The reverse inequality follows from monotonicity of T_C. Thus T_C preserves the nonempty supremum.

Conversely, suppose T_C preserves every nonempty supremum. For $B \succeq_{\mathrm{B}} C$, define

$$G_C(B) = \bigvee \mathcal{R}_C(B).$$

The family is nonempty, and preservation gives

$$T_C(G_C(B)) = \bigvee_{R \in \mathcal{R}_C(B)} T_C(R) \preceq_{\mathrm{B}} B.$$

Hence $G_C(B)$ is feasible and, by construction, dominates every feasible request. It is the greatest residual $B \oslash C$, and (D.3) proves the stated partial-right-adjoint relation. □

D.3 Proof of Theorem 1

Proof. Put $\theta = \mathrm{Pr}_P(L_C < +\infty)$. If $\theta = 0$, then $C \simeq_{\mathrm{B}} \top$. The only nonempty feasible set has cap equivalent to $\top$ and contains every request, so its greatest element is $\top$.

Suppose $\theta > 0$, and let W have the conditional law of e^{-L_C} given $L_C < +\infty$. Then $0 < W < \infty$ almost surely. Independent composition and conditioning on L_C give, for every request R,

$$h_{C \otimes R}(x) = 1 - \theta + \theta\, \mathbb{E}[h_R(xW)], \qquad x > 0. \tag{D.4}$$

Indeed, on the P-exclusive event $L_C = +\infty$ the hockey-stick integrand equals one. On the finite event, averaging $(1 - xe^{-L_C} e^{-L_R})_+$ over the independent request gives $h_R(xW)$. All integrands lie in $[0, 1]$, so this conditioning needs no moment assumption. Q-exclusive mass has zero P-probability.

If $W = w$ almost surely, with $0 < w < \infty$, then (D.4) is the positive rescaling $h_R(x) \mapsto 1 - \theta + \theta h_R(xw)$. It commutes with every nonempty pointwise supremum. Proposition 2 therefore gives a greatest residual at every nonempty cap.

Now suppose W is nonconstant. Choose $a > 0$ such that $\Pr(W < a) > 0$ and $\Pr(W > a) > 0$; two distinct points of its essential support give such a separating value. Take $0 < \beta < \min\{1, 1/a\}$ and $\alpha = a\beta$, so $0 < \alpha, \beta < 1$. Define two binary-output request experiments by

$$A = ((1, 0), (\beta, 1 - \beta)), \qquad D = ((\alpha, 1 - \alpha), (1, 0)).$$

They have valid probability laws and profiles

$$h_A(z) = (1 - \beta z)_+, \qquad h_D(z) = 1 - \alpha + (\alpha - z)_+.$$

Since $\alpha < a < 1/\beta$, their difference is

$$h_A(z) - h_D(z) = \begin{cases} (1 - \beta)z, & 0 < z < \alpha, \\ \alpha - \beta z, & \alpha \le z < 1/\beta, \\ -(1 - \alpha), & z \ge 1/\beta. \end{cases}$$

It is strictly positive for $0 < z < a$ and strictly negative for $z > a$. The join $A \vee D$ is realized explicitly by $P = (\alpha, 1 - \alpha, 0)$ and $Q = (\beta, 0, 1 - \beta)$: its profile $1 - \alpha + (\alpha - \beta z)_+$ equals $\max\{h_A(z), h_D(z)\}$. This also agrees with the general supremum representation of Regehr et al. (2026, Prop. 11).

Set $X = h_A(W)$ and $Y = h_D(W)$. Both lie in $[0, 1]$, and each of the events $X > Y$ and $Y > X$ has positive probability. Consequently

$$\mathbb{E} \max(X, Y) - \mathbb{E}X = \mathbb{E}(Y - X)_+ > 0, \qquad \mathbb{E} \max(X, Y) - \mathbb{E}Y = \mathbb{E}(X - Y)_+ > 0.$$

Let $B = (C \otimes A) \vee (C \otimes D)$, a valid experiment by the same supremum representation. Both A and D are feasible at this cap. Applying (D.4) at $x = 1$ yields

$$h_{C \otimes (A \vee D)}(1) = 1 - \theta + \theta\, \mathbb{E} \max(X, Y) > 1 - \theta + \theta \max\{\mathbb{E}X, \mathbb{E}Y\} = h_B(1).$$

Any greatest feasible residual would dominate both A and D, hence their join; composition monotonicity would then make $C \otimes (A \vee D)$ feasible, contradicting the strict inequality. Thus a nonconstant finite-loss law always supplies a nonempty cap with no greatest residual. Together with the two positive cases, this proves the equivalence. □

D.4 Proof of Theorem 6

Proof. Lemma 7 with one finite atom gives

$$h_{C\otimes R}(x) = (1-p) + p h_R(xe^{-c}).$$

The condition $pe^{-c} \le 1$ needed for a valid experiment is automatic because the displayed law is the P-side PLD of the given factor C.

For two requests, the common revealing mass cancels and

$$h_{C\otimes R_1}(x) - h_{C\otimes R_2}(x) = p\{h_{R_1}(xe^{-c}) - h_{R_2}(xe^{-c})\}. \tag{D.5}$$

Since $p > 0$ and $x \mapsto xe^{-c}$ maps $(0,\infty)$ onto itself, (D.5) shows both directions of

$$C \otimes R_1 \preceq_{\mathrm{B}} C \otimes R_2 \quad \Longleftrightarrow \quad R_1 \preceq_{\mathrm{B}} R_2.$$

Thus Blackwell order is reflected, and equivalence cancellation follows by applying the same equivalence in both order directions.

Let $\mathcal{A}$ be nonempty. The pointwise profile representation of the binary-experiment supremum yields

$$\begin{aligned} h_{C\otimes\bigvee\mathcal{A}}(x) &= 1 - p + p \sup_{R\in\mathcal{A}} h_R(xe^{-c}) \\ &= \sup_{R\in\mathcal{A}} \{1 - p + p h_R(xe^{-c})\} \\ &= h_{\bigvee_{R\in\mathcal{A}}(C\otimes R)}(x). \end{aligned}$$

Hence T_C preserves every nonempty supremum, and Proposition 2 gives a greatest residual for every cap whose feasible residual set is nonempty. Finally, if such factors form a semigroup, then

$$C_s \otimes R \preceq_{\mathrm{B}} C_t = C_s \otimes C_{t-s}$$

is equivalent by order reflection to $R \preceq_{\mathrm{B}} C_{t-s}$, giving universal semigroup-subtraction semantics. □

D.5 Proof of Example 1

Under the P hypothesis, the P-side privacy loss of $G_1 \otimes R_\varepsilon$ has law

$$S \sim \frac{7}{10} N(1/2+\varepsilon, 1) + \frac{3}{10} N(1/2-\varepsilon, 1), \tag{D.6}$$

whereas that of $G_{\sqrt{2}}$ has law $T \sim N(1,2)$. Write their densities as f_S, f_T. Direct Gaussian algebra, using $e^\varepsilon = 7/3$, gives the even density ratio

$$\frac{f_S(a)}{f_T(a)} = \sqrt{2}\exp\left(-\frac{a^2}{4} + \frac{1}{8} - \frac{\varepsilon^2}{2}\right)\frac{\cosh(\varepsilon a)}{\cosh(\varepsilon/2)}. \tag{D.7}$$

We first record the one-sided comparison used to check the Blackwell-order direction without appealing to a numerical profile grid.

Lemma 9. *Let E and F be symmetric binary experiments, and let L_E, L_F denote their continuous P-side privacy losses with CDFs F_E, F_F. If*

$$F_F(a) \le F_E(a) \qquad \textit{for every } a \ge 0,$$

then $E \preceq_{\mathrm{B}} F$.

Proof. For $x = e^t \ge 1$, set $\psi_t(\ell) = (1 - e^{t-\ell})_+$. Tail integration gives

$$h_E(e^t) = \mathbb{E}[\psi_t(L_E)] = \int_t^\infty e^{t-a}\{1 - F_E(a)\}\, da,$$

and likewise for F. The assumed CDF inequality therefore gives $h_E(x) \le h_F(x)$ for $x \ge 1$. For any experiment,

$$h_{(P,Q)}(x) = 1 - x + x h_{(Q,P)}(1/x).$$

Symmetry identifies the reversed and original profiles, so the inequality for $1/x \ge 1$ also gives it for $0 < x < 1$. Thus $E \preceq_{\mathrm{B}} F$. □

Proof of Example 1. Let $\varepsilon = \log(7/3)$. The privacy-loss laws and density ratio are given in Equations (D.6) and (D.7). Let ℓ be the logarithm of the density ratio. For $a > 0$,

$$\ell'(a) = -\frac{a}{2} + \varepsilon \tanh(\varepsilon a).$$

The function $a \mapsto \tanh(\varepsilon a)/a$ is strictly decreasing: with $u = \varepsilon a$, the derivative has the sign of $u \operatorname{sech}^2 u - \tanh u < 0$, because

$$\sinh u \cosh u - u > 0, \qquad \frac{d}{du}(\sinh u \cosh u - u) = 2 \sinh^2 u > 0.$$

Moreover $\varepsilon^2 > 1/2$, so ℓ' is positive immediately to the right of zero and eventually negative. Hence ℓ first increases and then decreases.

The positive constants needed here can be bounded without a floating-point grid. The positive-term atanh series for $\log(7/3) = 2\sum_{n\ge 0}(2/5)^{2n+1}/(2n+1)$, together with alternating Taylor bounds for the normal integral, gives

$$\varepsilon^2 - \frac{1}{2} > 0.2179, \qquad \ell(0) > 0.0254.$$

Thus (D.7) starts above one and crosses one exactly once on $(0, \infty)$.

For an explicit finite calculation, write

$$A(y) = 2\sum_{j=0}^{11} \frac{y^{2j+1}}{2j+1}, \qquad A(y) \le \log\frac{1+y}{1-y} \le A(y) + \frac{2y^{25}}{25(1-y^2)}.$$

Use $y = 2/5$ for ε and $y = 17/67$ for $\log(8(7/10)(3/10))$; then $\ell(0) = \frac{1}{2}\log(8(7/10)(3/10)) + \frac{1}{8} - \frac{1}{2}\varepsilon^2$. For $J(x) = \int_0^x e^{-u^2/2}\,du$ and $0 \le x \le 7/5$, put

$$J_n(x) = \sum_{j=0}^{n} \frac{(-1)^j x^{2j+1}}{2^j j!(2j+1)}, \qquad J_{11}(x) \le J(x) \le J_{12}(x).$$

These alternating bounds and $70710678/10^8 < 1/\sqrt{2} < 70710679/10^8$ give $\sqrt{2\pi}D(0) > 0.03322$, where $D(a) = F_S(a) - F_T(a)$. Using $1/\sqrt{2\pi} > 0.3988$ (from $\pi < 22/7$) yields

$$\begin{aligned} D(0) &= \frac{7}{10}\Phi(-1/2-\varepsilon) + \frac{3}{10}\Phi(-1/2+\varepsilon) - \Phi(-1/\sqrt{2}) \\ &> 0.0132. \end{aligned}$$

Because $D' = f_S - f_T$, the single density crossing makes D first increase and then decrease. Since $D(0) > 0$ and $D(a) \to 0$ as $a \to \infty$, we have $F_T(a) \le F_S(a)$ for all $a \ge 0$. Both experiments in (D.6) are symmetric, so Lemma 9 yields

$$G_1 \otimes R_\varepsilon \preceq_{\mathrm{B}} G_{\sqrt{2}}. \tag{D.8}$$

At profile threshold $x = 1$,

$$h_{R_\varepsilon}(1) = \frac{e^\varepsilon - 1}{e^\varepsilon + 1} = \frac{2}{5},$$

whereas

$$h_{G_1}(1) = 2\Phi(1/2) - 1 = 2\int_0^{1/2} \phi(t)\,dt < \phi(0) = \frac{1}{\sqrt{2\pi}} < \frac{2}{5}.$$

The last inequality follows from $\pi > 25/8$. Therefore $R_\varepsilon \not\preceq_{\mathrm{B}} G_1$. This proves that the GDP-family-restricted residual G_1 is not a greatest residual in the full binary-experiment order.

It remains to prove the stronger nonexistence conclusion. Gaussian composition gives $G_1 \otimes G_1 \simeq_{\mathrm{B}} G_{\sqrt{2}}$, and (D.8) makes both G_1 and R_ε feasible. The Gaussian experiment G_1 is regular because its two hypotheses are mutually absolutely continuous; hence Torgersen's equivalence-cancellation theorem applies. If the feasible set had a greatest element $R^\star$, then $G_1 \preceq_{\mathrm{B}} R^\star$, and monotonicity plus feasibility would give

$$G_1 \otimes G_1 \preceq_{\mathrm{B}} G_1 \otimes R^\star \preceq_{\mathrm{B}} G_{\sqrt{2}} \simeq_{\mathrm{B}} G_1 \otimes G_1.$$

Product equivalence and regularity therefore allow cancellation of G_1, yielding $R^\star \simeq_{\mathrm{B}} G_1$. Greatestness would then force $R_\varepsilon \preceq_{\mathrm{B}} G_1$, a contradiction. The feasible residual set is therefore nonempty and has no greatest element. □

D.6 Proof of Theorem 5

Proof. For $m = 1$, Theorem 6 proves order reflection. For $m = 0$, the factor is complete reveal and $C \otimes R \simeq_{\mathrm{B}} \top$ for every R, so order reflection fails. It remains to treat $m \geq 2$.

Reindex the finite atoms as $c_0 < \cdots < c_k$ with $k = m - 1 \geq 1$. If $d(x) = h_{R_1}(x) - h_{R_2}(x)$ and $f(t) = d(e^t)$, (D.2) becomes

$$h_{C\otimes R_1}(e^t) - h_{C\otimes R_2}(e^t) = (\mathcal{T}_C f)(t), \qquad (\mathcal{T}_C f)(t) = \sum_{i=0}^{k} p_i f(t - c_i). \tag{D.9}$$

Set $a = c_1 - c_0 > 0$. Choose a nonnegative C^∞ bump φ with compact support and $\varphi(0) > 0$, and choose $0 < \eta < p_1/p_0$. For $j = 1, \ldots, k$, let $b_j = a + c_j - c_0$ and define

$$f(t) = -\varphi(t) + \eta\varphi(t - a) - \eta \sum_{j=1}^{k} \frac{p_j}{p_0} \varphi(t - b_j).$$

The only positive input bump is centered at a. After translation by kernel atom c_0, its output center $a + c_0 = c_1$ coincides with the negative base bump translated by c_1; their combined coefficient is $\eta p_0 - p_1 < 0$. For each $j \geq 1$, the positive output center $a + c_j$ coincides with the b_j compensator translated by c_0, whose coefficient $-\eta(p_j/p_0)p_0 = -\eta p_j$ exactly cancels $+\eta p_j$. Every remaining output term comes from a negative bump. Coincident centers can only add nonpositive coefficients. Hence

$$\mathcal{T}_C f \leq 0 \quad \text{everywhere.} \tag{D.10}$$

The input centers $0, a, b_1, \ldots, b_k$ are distinct. Taking the bump support sufficiently narrow gives $f(0) < 0$ and $f(a) > 0$.

We now realize this crossing by valid experiments. Put $d(x) = f(\log x)$. Then $d \in C_c^\infty((0, \infty))$. Choose $0 < r < 1 < s < \infty$ whose interior contains the supports of d and d''. There exists a smooth density q_0, positive on (r, s), supported on $[r, s]$, and satisfying

$$\int q_0(z)\, dz = 1, \qquad \int z q_0(z)\, dz = 1.$$

For example, exponentially tilt any smooth positive base density on (r, s); its mean varies continuously from r to s, so one tilt has mean one.

Compact support gives

$$\int d''(z)\, dz = 0, \qquad \int z d''(z)\, dz = 0. \tag{D.11}$$

The second identity follows by integration by parts, with all boundary terms zero. Since q_0 has a positive lower bound on $\mathrm{supp}(d'')$, choose $\rho > 0$ small enough that

$$q_\pm(z) = q_0(z) \pm \frac{\rho}{2} d''(z)$$

are strictly positive wherever q_0 is positive. Equation (D.11) shows that both $q_\pm$ have mass one and mean one. Thus

$$Q_\pm(dz) = q_\pm(z)\, dz, \qquad P_\pm(dz) = z q_\pm(z)\, dz$$

define binary experiments with likelihood ratio $z \in [r, s]$. They are bounded, and the two P laws and the two Q laws are mutually absolutely continuous.

Let $R_1 = (P_+, Q_+)$ and $R_2 = (P_-, Q_-)$. Twice integrating by parts gives, for every $\gamma > 0$,

$$h_{R_1}(\gamma) - h_{R_2}(\gamma) = \rho \int_\gamma^\infty (z - \gamma) d''(z)\, dz = \rho d(\gamma). \tag{D.12}$$

Because d has both signs, neither request dominates the other. Combining (D.9), (D.10), and (D.12) gives

$$h_{C\otimes R_1}(e^t) - h_{C\otimes R_2}(e^t) = \rho(\mathcal{T}_C f)(t) \leq 0$$

for every t, hence $C \otimes R_1 \preceq_{\mathrm{B}} C \otimes R_2$. This proves failure of order reflection for every $m \geq 2$ and completes the classification. □

D.7 Proof of Theorem 3

Proof. We separate the three possible atom counts.

$m = 0$**.** The hypotheses of C are mutually singular, so $C \simeq_{\mathrm{B}} \top$ and $\top \otimes R \simeq_{\mathrm{B}} \top$ for every R. Therefore equivalence cancellation and order reflection fail. A residual set is empty unless the budget is $\top$; at $\top$ every experiment is feasible and its greatest element is $\top$. Thus greatest-residual existence holds only in the stated degenerate sense.

$m = 1$**.** Theorem 6 gives equivalence cancellation, order reflection, and a greatest residual for every cap whose feasible residual set is nonempty.

$m \geq 2$**.** Positive finite P-side mass means that the two hypotheses of C have a nonzero common part, so C is regular. Torgersen's product-cancellation theorem therefore gives equivalence cancellation. Theorem 5 gives bounded, mutually absolutely continuous R_1, R_2 satisfying

$$R_1 \not\preceq_{\mathrm{B}} R_2, \qquad C \otimes R_1 \preceq_{\mathrm{B}} C \otimes R_2. \tag{D.13}$$

Thus order reflection fails. Put $B = C \otimes R_2$. Both requests are in $\mathcal{R}_C(B)$. If that set had a greatest element $R^\star$, then $R_2 \preceq_{\mathrm{B}} R^\star$ and

$$C \otimes R_2 \preceq_{\mathrm{B}} C \otimes R^\star \preceq_{\mathrm{B}} B = C \otimes R_2.$$

The two products are equivalent. Applying equivalence cancellation—not order reflection—gives $R^\star \simeq_{\mathrm{B}} R_2$. Greatestness would then force $R_1 \preceq_{\mathrm{B}} R_2$, contradicting (D.13). Hence no greatest residual exists for this cap.

The three cases prove equivalence cancellation exactly for $m \geq 1$, order reflection exactly for $m = 1$, and greatest-residual existence for every nonempty feasible set exactly for $m \leq 1$, together with the claimed no-greatest-residual budget for every $m \geq 2$. □

D.8 Proof of Theorem 4

Proof. Let

$$p_t = \Pr(L_{C_t} < +\infty).$$

The sum of two privacy losses is finite exactly when both summands are finite, so the semigroup law gives

$$p_{s+t} = p_s p_t, \qquad p_0 = 1.$$

Weak convergence to the identity PLD δ_0 and Portmanteau applied to the closed point $\{+\infty\}$ give $p_t \to 1$ as $t \downarrow 0$. Also $p_t > 0$ for finite t: otherwise $p_t = p_{t/n}^n = 0$ for every n, contradicting $p_{t/n} \to 1$. The continuous multiplicative Cauchy equation therefore has the form

$$p_t = e^{-\lambda t} \tag{D.14}$$

for some $\lambda \geq 0$.

Let μ_t be the finite subprobability part of the P-side PLD and normalize $\nu_t = \mu_t / p_t$. Finite parts convolve under composition, so

$$\nu_{s+t} = \nu_s * \nu_t. \tag{D.15}$$

Every ν_t has finite support. Fix $t > 0$, let $M = |\operatorname{supp}(\nu_t)|$, and choose an integer $n > M$. From

$$\nu_t = \nu_{t/n}^{*n}, \tag{D.16}$$

if $\nu_{t/n}$ had two support points $u < v$, the support on the right would contain the $n + 1$ distinct points

$$nu, (n-1)u + v, \ldots, u + (n-1)v, nv,$$

contradicting its size M. Hence $\nu_{t/n}$ is a point mass, and (D.16) makes ν_t a point mass as well. Write $\nu_t = \delta_{c_t}$.

Equation (D.15) gives $c_{s+t} = c_s + c_t$. Weak continuity at zero, together with $p_t \to 1$, forces $c_t \to 0$ as $t \downarrow 0$; otherwise the full PLD could not converge to δ_0. The continuous additive Cauchy equation

therefore gives $c_t = bt$ for some $b \in \mathbb{R}$. Combining this with (D.14) proves the displayed privacy-loss law.

The common finite event has Q-mass

$$p_t e^{-c_t} = e^{-(\lambda+b)t}.$$

Because it cannot exceed one, $\lambda + b \geq 0$. Each positive-time factor has positive finite mass and one finite atom, so Theorem 6 gives equivalence cancellation, order reflection, and greatest residuals. Cancelling C_s by order reflection in $C_s \otimes R \preceq_{\mathrm{B}} C_s \otimes C_{t-s}$ gives universal parameter-subtraction semantics.

Finally suppose the semigroup is self-dual. The original experiment has P-exclusive mass $1 - p_t$, while its reversal has P-exclusive mass equal to the original Q-exclusive mass $1 - p_t e^{-c_t}$. Blackwell equivalence preserves this mass, which is $\lim_{x\to\infty} h(x)$. Thus $p_t = p_t e^{-c_t}$. Since $p_t > 0$, $c_t = 0$ for every t, so $b = 0$ and only the erasure semigroup remains. □