Exact-$\mathbb{Q}$ Lorentzian Inertia Certificates, Coincident Down–Up Spectral-Gap Phase Transitions, and Machine-Checked Subgaussian Matrix Concentration for Matroids

Lluis Eriksson*

November 2026

Abstract

Lorentzian polynomials provide a unified algebraic-combinatorial bridge from matroid theory to rapid mixing of high-dimensional random walks and sharp concentration inequalities. Very recently, Ercan (`arXiv:2611.00626`, November 2026) proved complete matrix entropy contraction with optimal universal factor ℓ/r for all r-homogeneous Lorentzian laws, yielding subgaussian matrix concentration bounds for arbitrary Hermitian observables with bounded exchange differences across all matroids. Here we sharpen and operationalize this landscape into exact rational theorems over $\mathbb{Q}$, backed by an 18-theorem zero-dependency Lean 4 formalization and a pure-stdlib zero-floating-point-trust certificate verifier (`LORENTZ-MATROID-CERT-1.0`). First, we establish a sharp exact converse on the two-cluster weighted graph family $G_{\varepsilon,m} = (K_m \sqcup K_m) + \varepsilon K_{m,m}$ ($m \geq 2$, $\varepsilon > 0$): the 2-homogeneous generating polynomial $f_{\varepsilon,m}$ is Lorentzian if and only if the $2 \to 1 \to 2$ down–up spectral gap satisfies Ercan's optimal threshold $\gamma(\varepsilon, m) \geq 1/2$, with exact coincident critical coupling $\varepsilon_c(m) = (m-1)/m \in \mathbb{Q}$, exact piecewise gap $\gamma(\varepsilon, m) = \frac{m\varepsilon}{(m-1)+m\varepsilon}$ ($\varepsilon \leq 1$) and $\frac{1}{2} + \frac{1}{2((m-1)+m\varepsilon)}$ ($\varepsilon \geq 1$), and unique global maximum $\frac{m}{2m-1}$ attained at the unweighted complete graph K_{2m} ($\varepsilon = 1$). Second, using exact rational congruence diagonalization $QHQ^\top = \mathrm{diag}(D)$ ($QQ^{-1} = I$) capable of handling zero-diagonal multiaffine Hessians, we certify complete M-convexity, all $(r-2)$-link Lorentzian Hessians, and strict excess down–up spectral gaps $\gamma(P_{r\to\ell\to r}) > 1 - \ell/r$ across uniform $U_{3,6}$, graphic $M(K_4)$, Fano F_7, non-Fano F_7^-, and the non-representable Vámos matroid V_8 (all 28 rank-2 links certified in 20.2 ms; exact gap $\gamma(P_{4\to2\to4}; V_8) = 9/14 > 1/2$; certified bracket $[46/125, 37/100] > 1/4$ at $\ell = 3$). Third, we translate Ercan's matrix subgaussian inequality into purely rational tail certificates via exact operator-norm PSD congruences $c^2 I_d - (A_e - A_f)^2 \succeq 0$ over all exchangeable pairs and exact rational Taylor-envelope denominators $T_K(x) = \sum_{k=0}^{K} x^k/k!$, certifying non-commuting 2×2 Hermitian basis sums over $M(K_4)$ and non-representable V_8 with zero floats and cross-checking against exact exhaustive basis enumeration. All 18 foundational theorems compile in zero-dependency Lean 4 (`v4.35.0`, 0 `sorry`, 0 `Classical.choice`, 0 warnings), all 16 exact-$\mathbb{Q}$ certificates verify in ≤ 404.3 ms, and the verifier rejects 48/48 live adversarial red-team attacks.

1 Introduction

The theory of Lorentzian polynomials, introduced by Brändén and Huh [2] alongside the strongly log-concave framework of Anari, Liu, Oveis Gharan, and Vinzant [1], established that the basis

*Independent researcher (Sweden). Prepared with AI assistance (Google Antigravity). Code, zero-dependency Lean 4 formalization, and exact-$\mathbb{Q}$ certificates (MIT license): `https://github.com/lluiseriksson/lorentzian-matroid-certificates`.

generating polynomial

$$f_M(x_1, \ldots, x_n) = \sum_{B \in \mathcal{B}(M)} \prod_{e \in B} x_e \tag{1}$$

of every rank-r matroid M satisfies the Lorentzian Hessian signature property on every $(r-2)$-fold partial derivative. This algebraic condition implies Mason's ultra-log-concavity conjecture [6], resolves the Mihail–Vazirani expansion conjecture for basis-exchange walks [5, 1], and yields the modified log-Sobolev constant $\alpha_{\mathrm{MLSI}} \geq 1/r$ for scalar functions [3].

In November 2026, Ercan [4] (`arXiv:2611.00626`) elevated this scalar theory to the non-commutative matrix regime, proving that uniform retention of ℓ out of r particles contracts matrix entropy by the optimal universal factor ℓ/r for every r-homogeneous Lorentzian law across arbitrary matrix dimension d. Combining this contraction with matrix Laplace comparisons, Ercan established subgaussian matrix concentration bounds for arbitrary Hermitian observables with bounded exchange differences, extending classical independent-sum matrix tail bounds [8] to all matroids—including non-representable matroids such as the Vámos matroid V_8 [7].

Despite these breakthroughs, three structural and certification gaps remain open:

1. **Sharp converse to Lorentzian contraction:** Ercan [4] proves that every r-homogeneous Lorentzian law achieves contraction factor $\leq \ell/r$ (down–up spectral gap $\gamma(P_{r\to\ell\to r}) \geq 1-\ell/r$). Does the failure of the Lorentzian Hessian inertia condition immediately force a quantitative breakdown of the universal spectral-gap threshold $1 - \ell/r$ at an *exact coincident threshold*?
2. **Zero-float certification of Lorentzianity and excess expansion:** Standard numerical eigensolvers (`LAPACK/eigh`) rely on IEEE-754 floating-point arithmetic, cannot rigorously certify boundary eigenvalues ($\lambda_2(H) = 0$) or exact rational excess gaps $\gamma(P_{r\to\ell\to r}) > 1 - \ell/r$, and standard $LDL^\top$ factorizations break down on zero-diagonal multiaffine matroid Hessians ($H_{ii} = 0$).
3. **Rationalized non-commutative matrix concentration certificates:** Operationalizing Ercan's subgaussian tail bound $\mathbb{P}[\lambda_{\max}(F(B) - \mathbb{E}F) \geq t] \leq d \exp(-t^2/(2v_{\mathrm{sg}}))$ for concrete Hermitian basis sums requires certifying both the exact operator-norm exchange sensitivity $\max_{(e,f)} \|A_e - A_f\|_{\mathrm{op}} \leq c$ and the exponential tail threshold without floating-point matrix norms or transcendental evaluations.

Contributions. We close all three gaps with four exact-$\mathbb{Q}$ theorems (Section 3), an 18-theorem zero-dependency Lean 4 formalization (Section 4), a pure-stdlib zero-float verifier `LORENTZ-MATROID-CERT-1.0`, and four deterministic computational experiments (Section 5):

- **Theorem 3.1 (Coincident Phase Transition):** On $G_{\varepsilon,m} = (K_m \sqcup K_m) + \varepsilon K_{m,m}$ ($m \geq 2$), $f_{\varepsilon,m}$ is Lorentzian iff $\gamma(P_{2\to1\to2}; \varepsilon, m) \geq 1/2$ iff $\varepsilon \geq \varepsilon_c(m) = \frac{m-1}{m}$, with an exact piecewise rational spectral-gap law peaking uniquely at the unweighted complete graph K_{2m} ($\varepsilon = 1$, $\gamma_{\max} = \frac{m}{2m-1}$).
- **Theorem 3.2 (Exact-$\mathbb{Q}$ Congruence Certificates for Non-Representable Matroids):** Using invertible rational congruence $QHQ^\top = \mathrm{diag}(D)$ ($QQ^{-1} = I$) built from swaps, shears, and eliminations, we certify complete M-convexity, all $(r-2)$-link Hessians, and exact/bracketed strict excess down–up spectral gaps across $U_{3,6}$, $M(K_4)$, F_7, F_7^-, and the non-representable Vámos matroid V_8.
- **Theorem 3.3 (Zero-Float Matrix Tail Certification):** We combine Ercan's ℓ/r matrix entropy contraction with exact rational PSD certificates $c^2 I_d - (A_e - A_f)^2 \succeq 0$ on all exchangeable pairs and exact rational Taylor-envelope denominators $T_K(x) = \sum_{k=0}^{K} x^k/k!$ to certify non-commuting matrix tail probabilities purely over $\mathbb{Q}$.

- **Theorem 3.5 (Soundness of `LORENTZ-MATROID-CERT-1.0`):** Our pure-stdlib verifier independently reconstructs all link Hessians, reversible walks, Dirichlet–Gram operators, and exchangeable pairs directly from explicit basis lists, verifying 16/16 certificates with zero floats and rejecting 48/48 live adversarial red-team attacks.

2 Preliminaries

Definition 2.1 (Lorentzian Polynomials and Matroid Basis Laws [2])**.** A homogeneous polynomial $f \in \mathbb{R}_{\geq 0}[x_1, \ldots, x_n]$ of degree $r \geq 2$ with M-convex support is *Lorentzian* if for every multi-index $\alpha \in \mathbb{N}^n$ with $|\alpha| = r-2$, the symmetric Hessian matrix $H_\alpha = \nabla^2(\partial^\alpha f)$ has at most one positive eigenvalue. For a rank-r matroid $M = ([n], \mathcal{B})$ with multiaffine basis polynomial $f_M(x) = \sum_{B\in\mathcal{B}} w_B \prod_{e\in B} x_e$ ($w_B > 0$), M-convexity is equivalent to the matroid basis-exchange axiom, and the nontrivial $(r-2)$-fold derivatives correspond to independent sets $S \in \mathcal{I}_{r-2}$, with zero-diagonal link Hessian $(H_S)_{ij} = w_{S\cup\{i,j\}}$ for $i \neq j \notin S$.

Definition 2.2 (Down–Up and Dual Up–Down Matroid Walks)**.** For $0 \leq \ell < r$ and basis distribution $\pi(B) = w_B / \sum_{B'\in\mathcal{B}} w_{B'}$, the $r \to \ell \to r$ down–up Markov chain $P_{r\to\ell\to r}$ drops $r-\ell$ elements uniformly at random from B to obtain $S \in \mathcal{I}_\ell$, then resamples a basis $B' \supset S$ with probability proportional to $w_{B'}$:

$$P_{r\to\ell\to r}(B, B') = \frac{1}{\binom{r}{\ell}} \sum_{S\subseteq B\cap B',\, |S|=\ell} \frac{w_{B'}}{W(S)}, \qquad W(S) := \sum_{B''\supseteq S} w_{B''}. \tag{2}$$

Writing $D_{r\to\ell}(B, S) = \binom{r}{\ell}^{-1} \mathbf{1}_{S\subseteq B}$ and $U_{\ell\to r}(S, B') = \frac{w_{B'}}{W(S)} \mathbf{1}_{S\subseteq B'}$, we have $P_{r\to\ell\to r} = D_{r\to\ell} U_{\ell\to r}$ (reversible with respect to π and positive semidefinite in $L^2(\pi)$) and its dual up–down walk $Q_{\ell\to r\to\ell} = U_{\ell\to r} D_{r\to\ell}$ on $\mathcal{I}_\ell$ (reversible with respect to $\pi_\ell(S) = W(S)/\sum_{S'} W(S')$). Because $D_{r\to\ell}U_{\ell\to r}$ and $U_{\ell\to r}D_{r\to\ell}$ share identical nonzero eigenvalues, their spectral gaps coincide identically: $\gamma(P_{r\to\ell\to r}) = \gamma(Q_{\ell\to r\to\ell}) = 1 - \lambda_2(P_{r\to\ell\to r})$.

By Anari et al. [1] and Ercan [4], every r-homogeneous Lorentzian distribution satisfies the universal spectral-gap bound $\gamma(P_{r\to\ell\to r}) \geq 1 - \ell/r$, matching the universal matrix entropy contraction factor ℓ/r.

3 Main Exact-$\mathbb{Q}$ Theorems

3.1 Exact Coincident Lorentzian / Spectral-Gap Phase Transition

Theorem 3.1 (Coincident Phase Transition and Piecewise Spectral-Gap Law on $G_{\varepsilon,m}$)**.** *For every integer $m \geq 2$ and rational coupling $\varepsilon \in \mathbb{Q}_{>0}$, let $G_{\varepsilon,m} = (K_m \sqcup K_m) + \varepsilon K_{m,m}$ be the weighted complete graph on $n = 2m$ vertices with edge weights $w_{ij} = 1$ inside each m-clique and $w_{ij} = \varepsilon$ across cliques, and let $f_{\varepsilon,m}(x) = \frac{1}{2}\sum_{i\neq j} w_{ij} x_i x_j$ be its 2-homogeneous generating polynomial with Hessian $H_{\varepsilon,m} = \nabla^2 f_{\varepsilon,m}$. Then:*

1. ***Exact Hessian Spectrum and Lorentzian Threshold:*** *$H_{\varepsilon,m}$ has exact eigenvalues*

$$\begin{aligned} &\lambda_+(\varepsilon, m) = (m-1) + m\varepsilon > 0 \ \ (\textit{mult. } 1), \\ &\lambda_-(\varepsilon, m) = (m-1) - m\varepsilon \ \ (\textit{mult. } 1), \qquad \lambda_\perp = -1 \ \ (\textit{mult. } 2m-2). \end{aligned} \tag{3}$$

In particular, $\lambda_2(H_{\varepsilon,m}) = \max\big((m-1) - m\varepsilon,\, -1\big)$, so $f_{\varepsilon,m}$ is Lorentzian if and only if

$$\varepsilon \geq \varepsilon_c(m) := \frac{m-1}{m} \in \mathbb{Q} \cap (0, 1). \tag{4}$$

2. ***Exact Piecewise Down–Up Spectral Gap:*** *The $2 \to 1 \to 2$ down–up walk $P_{2\to1\to2}$ on edges (isospectral to the vertex walk $Q_{1\to2\to1} = \frac{1}{2}(I_{2m} + D^{-1}H_{\varepsilon,m})$ with $D = \lambda_+ I_{2m}$) has exact spectral gap*

$$\gamma(P_{2\to1\to2};\varepsilon,m) = \begin{cases} \dfrac{m\varepsilon}{(m-1)+m\varepsilon}, & 0 < \varepsilon \le 1, \\ \dfrac{1}{2} + \dfrac{1}{2\big((m-1)+m\varepsilon\big)}, & \varepsilon \ge 1. \end{cases} \tag{5}$$

3. ***Exact Coincidence and Unique Complete-Graph Maximum:*** *For every $m \ge 2$ and $\varepsilon > 0$,*

$$f_{\varepsilon,m} \text{ is Lorentzian } \big(\lambda_2(H_{\varepsilon,m}) \le 0\big) \iff \gamma(P_{2\to1\to2};\varepsilon,m) \ge \frac{1}{2} \iff \varepsilon \ge \frac{m-1}{m}, \tag{6}$$

with simultaneous criticality $\lambda_2(H_{\varepsilon_c(m),m}) = 0$ and $\gamma(\varepsilon_c(m),m) = 1/2$ for all $m \ge 2$, linear collapse $\gamma(\varepsilon,m) \to 0$ as $\varepsilon \downarrow 0$, and unique global maximum $\gamma_{\max}(m) = \frac{m}{2m-1} > \frac{1}{2}$ attained at the unweighted complete graph K_{2m} ($\varepsilon = 1$).

Proof. In 2×2 block form on $\mathbb{R}^m \oplus \mathbb{R}^m$, $H_{\varepsilon,m} = \begin{pmatrix} J_m - I_m & \varepsilon J_m \\ \varepsilon J_m & J_m - I_m \end{pmatrix}$, where J_m is the all-ones matrix. On $\mathbf{1}_{2m} = (\mathbf{1}_m, \mathbf{1}_m)^\top$, $H_{\varepsilon,m}\mathbf{1}_{2m} = ((m-1)+m\varepsilon)\mathbf{1}_{2m}$. On $v_- = (\mathbf{1}_m, -\mathbf{1}_m)^\top$, $H_{\varepsilon,m}v_- = ((m-1)-m\varepsilon)v_-$. On the $(2m-2)$-dimensional subspace $\mathbf{1}_m^\perp \oplus \mathbf{1}_m^\perp$, J_m annihilates both blocks, giving $H_{\varepsilon,m}w = -w$. Thus $\lambda_2(H_{\varepsilon,m}) = \max(\lambda_-, -1) \le 0 \iff (m-1) - m\varepsilon \le 0 \iff \varepsilon \ge \frac{m-1}{m}$ (Lean 4 Theorems 1–3).

Because $G_{\varepsilon,m}$ is vertex-regular with weighted degree $d = (m-1)+m\varepsilon = \lambda_+$, the dual vertex walk satisfies $Q_{1\to2\to1}(i,j) = \frac{1}{2}\delta_{ij} + \frac{w_{ij}}{2d}$, i.e. $Q_{1\to2\to1} = \frac{1}{2}(I_{2m} + d^{-1}H_{\varepsilon,m})$. Its eigenvalues are therefore

$$\mu_1 = 1, \qquad \mu_- = \frac{1}{2}\left(1 + \frac{(m-1)-m\varepsilon}{(m-1)+m\varepsilon}\right) = \frac{m-1}{(m-1)+m\varepsilon},$$
$$\mu_\perp = \frac{1}{2}\left(1 - \frac{1}{(m-1)+m\varepsilon}\right) = \frac{(m-2)+m\varepsilon}{2\big((m-1)+m\varepsilon\big)}.$$

Comparing μ_- and $\mu_\perp$ gives $2(m-1) \ge (m-2)+m\varepsilon \iff m \ge m\varepsilon \iff \varepsilon \le 1$ (Lean 4 Theorem 7). Hence $\lambda_2(Q_{1\to2\to1}) = \mu_-$ for $0 < \varepsilon \le 1$ (giving $\gamma = 1 - \mu_- = \frac{m\varepsilon}{(m-1)+m\varepsilon}$) and $\lambda_2(Q_{1\to2\to1}) = \mu_\perp$ for $\varepsilon \ge 1$ (giving $\gamma = 1 - \mu_\perp = \frac{1}{2} + \frac{1}{2((m-1)+m\varepsilon)}$). For $\varepsilon > 1$, $1 - \mu_\perp > \frac{1}{2}$ holds unconditionally (Lean 4 Theorem 8), while for $0 < \varepsilon \le 1$, $\gamma \ge \frac{1}{2} \iff 2m\varepsilon \ge (m-1)+m\varepsilon \iff \varepsilon \ge \frac{m-1}{m}$, matching the Lorentzian threshold identically (Lean 4 Theorems 4–6). Finally, γ is strictly increasing on $(0,1]$ and strictly decreasing on $[1,\infty)$, attaining its unique global maximum $\frac{m}{2m-1}$ at $\varepsilon = 1$. □

3.2 Exact-$\mathbb{Q}$ Rational Congruence Certificates for Inertia and Spectral Gaps

Standard Cholesky or unit-triangular $LDL^\top$ factorizations fail on zero-diagonal multiaffine Hessians such as H_S (where $(H_S)_{ii} = 0$ for all i). The following theorem provides a universal zero-float rational congruence framework covering both indefinite Lorentzian Hessians and PSD Dirichlet–Gram operators.

Theorem 3.2 (Exact-$\mathbb{Q}$ Rational Congruence Certification of Lorentzianity and Spectral Gaps)**.** *Let $H \in \mathbb{Q}^{n\times n}$ be any symmetric rational matrix.*

1. ***Existence and Soundness of Rational Congruence:*** *There exist invertible $Q, Q^{-1} \in \mathbb{Q}^{n\times n}$ (with $QQ^{-1} = I_n$) and $D = (d_1,\dots,d_n) \in \mathbb{Q}^n$ such that $QHQ^\top = \operatorname{diag}(D)$, constructible in $O(n^3)$ exact rational operations using swaps, elementary shears $I + E_{kj}$ (converting a zero diagonal pivot $A_{kk} = 0$ with $A_{kj} \ne 0$ to $2A_{kj} \ne 0$), and Gaussian eliminations.*

- *If $d_1, \dots, d_n \geq 0$, then $v^\top H v = \sum_{i=1}^n d_i y_i^2 \geq 0$ for all $v \in \mathbb{R}^n$ (with $y = (Q^{-1})^\top v$), certifying $H \succeq 0$ (Lean 4 Theorem 9).*
- *If at most one $d_i > 0$ (say $d_1 > 0$ and $d_2, \dots, d_n \leq 0$), then $v^\top H v = \sum_{i=2}^n d_i y_i^2 \leq 0$ on the codimension-1 hyperplane $((Q^{-1})^\top v)_1 = 0$, certifying $\lambda_2(H) \leq 0$ by Cauchy interlacing / Sylvester's law of inertia (Lean 4 Theorem 10).*

2. ***Zero-Square-Root Reversible Spectral-Gap Certification:*** *Let $P \in \mathbb{Q}^{k\times k}$ be a π-reversible PSD Markov operator ($\Pi P = (\Pi P)^\top \succeq 0$ with $\Pi = \operatorname{diag}(\pi)$, $\pi \in \mathbb{Q}^k_{>0}$, $\sum_i \pi_i = 1$). For any $\gamma \in \mathbb{Q} \cap (0, 1]$, define the symmetric rational Dirichlet–Gram matrix*

$$M_\gamma := (1-\gamma)\Pi - \Pi P + \pi\pi^\top \in \mathbb{Q}^{k\times k}. \tag{7}$$

For every $f = c\mathbf{1} + g \in \mathbb{R}^k$ with $\pi^\top g = 0$, we have the exact orthogonal identity

$$f^\top M_\gamma f = (1-\gamma)c^2 + \Big((1-\gamma)\operatorname{Var}_\pi(g) - \langle g, Pg\rangle_\pi\Big), \tag{8}$$

(Lean 4 Theorem 11). Consequently, $M_\gamma \succeq 0$ (certified via $QM_\gamma Q^\top = \operatorname{diag}(D)$ with $D \geq 0$) holds if and only if $\gamma(P) \geq \gamma$. Conversely, any nonzero $v \in \mathbb{Q}^k$ with $\pi^\top v = 0$ certifies the rational upper bound $\gamma(P) \leq 1 - \frac{v^\top \Pi P v}{v^\top \Pi v}$, and when $M_\gamma \succeq 0$ and $\gamma < 1$, any zero pivot $d_i = 0$ in $QM_\gamma Q^\top = \operatorname{diag}(D)$ yields an exact rational eigenfunction $v = Q_{i,}$ achieving equality $\gamma(P) = \gamma$.

Proof. For Part 1, induction on the active leading block $k = 1, \dots, n$ maintains $QHQ^\top = A$ and $QQ^{-1} = I_n$: if $A_{kk} = 0$ and some $A_{jj} \neq 0$ ($j > k$), swapping rows/columns $k \leftrightarrow j$ restores a nonzero pivot; if $A_{jj} = 0$ for all $j \geq k$ but some $A_{kj} \neq 0$, applying the unimodular shear $E = I + E_{kj}$ ($E^{-1} = I - E_{kj}$) produces $(EAE^\top)_{kk} = 2A_{kj} \neq 0$; otherwise row/column k is identically zero. Eliminating A_{ik} ($i > k$) via $I - (A_{ik}/A_{kk})E_{ik}$ diagonalizes A without square roots. Because Q is invertible, setting $y = (Q^{-1})^\top v$ gives $v = Q^\top y$ and $v^\top H v = y^\top (QHQ^\top) y = \sum_{i=1}^n d_i y_i^2$, from which both PSD ($D \geq 0$) and codimension-1 nonpositivity ($\#\{i : d_i > 0\} \leq 1$) follow immediately.

For Part 2, write $f = c\mathbf{1} + g$ with $\pi^\top g = 0$. Since $\Pi\mathbf{1} = \pi$, $\Pi P\mathbf{1} = \pi$, $\mathbf{1}^\top \Pi g = \pi^\top g = 0$, and $\mathbf{1}^\top \Pi P g = \pi^\top P g = \pi^\top g = 0$ (by π-reversibility), expanding $f^\top M_\gamma f$ yields

$$f^\top M_\gamma f = (1-\gamma)(c^2 + \langle g, g\rangle_\pi) - (c^2 + \langle g, Pg\rangle_\pi) + c^2 = (1-\gamma)c^2 + \big((1-\gamma)\langle g, g\rangle_\pi - \langle g, Pg\rangle_\pi\big).$$

Since $1 - \gamma \geq 0$, $M_\gamma \succeq 0$ iff $\langle g, Pg\rangle_\pi \leq (1-\gamma)\langle g, g\rangle_\pi$ for all $g \perp_\pi \mathbf{1}$ iff $\lambda_2(P) \leq 1 - \gamma$. Finally, if $QM_\gamma Q^\top = \operatorname{diag}(D)$ has $d_i = 0$, then $QM_\gamma(Q^\top e_i) = \operatorname{diag}(D)e_i = \mathbf{0}$, so $M_\gamma(Q^\top e_i) = \mathbf{0}$; subtracting its π-mean yields an exact rational nullvector in $\mathbf{1}^{\perp_\pi}$ attaining Rayleigh gap γ whenever $\gamma < 1$. □

Applying Theorem 3.2 to the dual up–down walk $Q_{\ell\to r\to\ell}$ on $\mathcal{I}_\ell$ reduces matrix dimension dramatically (e.g., from 65×65 to 28×28 on Vámos V_8 at $\ell = 2$) while certifying the exact same spectral gap $\gamma(P_{r\to\ell\to r})$ with zero approximation.

3.3 Zero-Float Exact-$\mathbb{Q}$ Matrix Subgaussian Tail Certificates

Theorem 3.3 (Zero-Float Rational Matrix Concentration Certification)**.** *Let $M = ([n], \mathcal{B})$ be a rank-r matroid equipped with its uniform basis distribution μ, and let $\{A_e\}_{e=1}^n \subset \mathbb{Q}^{d\times d}$ be symmetric rational matrices defining the Hermitian basis-sum observable $F(B) = \sum_{e\in B} A_e$. Let $\mathcal{E}(M) := \{(e, f) : 1 \leq e < f \leq n,\ \exists B, B' \in \mathcal{B}$ with $B \triangle B' = \{e, f\}\}$ be the set of exchangeable element pairs. Suppose that:*

1. *For a rational bound $c^2 \in \mathbb{Q}_{>0}$ and every $(e, f) \in \mathcal{E}(M)$, the rational matrix $C_{ef} := c^2 I_d - (A_e - A_f)^2 \in \mathbb{Q}^{d\times d}$ admits an exact PSD congruence $Q_{ef} C_{ef} Q_{ef}^\top = \operatorname{diag}(D_{ef})$ with $Q_{ef} Q_{ef}^{-1} = I_d$ and $D_{ef} \geq 0$;*

2. *For rational deviation $t \in \mathbb{Q}_{>0}$, target failure probability $\delta \in \mathbb{Q} \cap (0,1)$, and Taylor order $K \geq 4$, setting $x := \frac{t^2}{rc^2} \in \mathbb{Q}_{>0}$ and $T_K(x) := \sum_{k=0}^{K} \frac{x^k}{k!} \in \mathbb{Q}_{>0}$ satisfies*

$$\frac{d}{T_K(x)} \leq \delta \qquad \text{in } \mathbb{Q}. \tag{9}$$

Then $\|F(B) - F(B')\|_{\mathrm{op}} \leq c$ for all exchangeable bases $|B \triangle B'| = 2$, and combining Ercan's subgaussian matrix concentration bound [4] for r-homogeneous Lorentzian laws ($\mathbb{P}_{B\sim\mu}[\lambda_{\max}(F(B) - \mathbb{E}F) \geq t] \leq d\exp(-t^2/(rc^2))$) with the unconditional inequality $e^x \geq T_K(x)$ ($x \geq 0$) certifies purely over $\mathbb{Q}$ that

$$\mathbb{P}_{B\sim\mu}\big[\lambda_{\max}\big(F(B) - \mathbb{E}_\mu[F]\big) \geq t\big] \leq d\, e^{-t^2/(rc^2)} \leq \frac{d}{T_K(t^2/(rc^2))} \leq \delta. \tag{10}$$

Proof. For any $(e,f) \in \mathcal{E}(M)$ and any vector $v \in \mathbb{R}^d$, the exact PSD congruence on C_{ef} guarantees $v^\top(c^2 I_d - (A_e - A_f)^2)v \geq 0$ by Theorem 3.2(1) (Lean 4 Theorem 17), whence $\|(A_e - A_f)v\|_2^2 \leq c^2\|v\|_2^2$ and therefore $\|A_e - A_f\|_{\mathrm{op}} \leq c$ without computing square roots or eigenvalues. Whenever $|B \triangle B'| = 2$ with $B \triangle B' = \{e,f\}$, $F(B) - F(B') = \pm(A_e - A_f)$, establishing bounded exchange difference c. Because f_M is r-homogeneous Lorentzian [2], Ercan's subgaussian matrix concentration inequality [4] (derived from optimal ℓ/r complete matrix entropy contraction with variance proxy $v_{\mathrm{sg}} = rc^2/2$) yields $\mathbb{P}[\lambda_{\max}(F(B) - \mathbb{E}F) \geq t] \leq d\exp(-t^2/(2v_{\mathrm{sg}})) = d\exp(-x)$. Since $x = t^2/(rc^2) > 0$, every omitted term $x^k/k!$ ($k > K$) in the power series of e^x is strictly positive, so $e^x > T_K(x) > 0$ (Lean 4 Theorems 15–16), giving $de^{-x} < d/T_K(x) \leq \delta$. □

Remark 3.4 (Independence of Rational Envelope from Floating-Point Evaluation). Theorem 3.3 cleanly separates the analytical input (Ercan's subgaussian matrix concentration theorem [4]) from finite-instance verification: once M is certified Lorentzian by Theorem 3.2, every instance-specific matrix norm, basis expectation, and tail inequality is verified purely in exact $\mathbb{Q}$ arithmetic, eliminating floating-point rounding in both operator norms and exponentials.

3.4 Pure-Stdlib Zero-Float Certificate Verifier (`LORENTZ-MATROID-CERT-1.0`)

Theorem 3.5 (Soundness and Anti-Forgery Architecture of `LORENTZ-MATROID-CERT-1.0`)**.** *The pure-stdlib verifier `verify_lorentzian_certificate_exact` (`src/lorentzian_certifier.py`, using `fractions.Fraction` exclusively and rejecting any floating-point value anywhere in the JSON tree) enforces four mutually exclusive modes:*

1. *`lorentzian_hessian_inertia`: independently verifies the matroid basis-exchange axiom on $\mathcal{B}$, enumerates all $S \in \mathcal{I}_{r-2}$, reconstructs every H_S directly from $(\mathcal{B}, w)$, checks $Q_S Q_S^{-1} = I$ and $Q_S H_S Q_S^\top = \mathrm{diag}(D_S)$, and enforces $\#\{i : d_i > 0\} \leq 1$;*
2. *`down_up_spectral_gap`: verifies basis-exchange, independently constructs $P_{r\to\ell\to r}$ (or isospectral dual $Q_{\ell\to r\to\ell}$) and $M_{\gamma_{\mathrm{low}}}$, checks $QQ^{-1} = I$, $QM_{\gamma_{\mathrm{low}}}Q^\top = \mathrm{diag}(D) \geq 0$, $\gamma_{\mathrm{low}} \geq 1 - \ell/r$ (or strict $>$ if claimed), and validates any centered upper-bound witness v ($\pi^\top v = 0$, $v^\top \Pi v > 0$);*
3. *`coincident_phase_transition`: verifies either supercritical Lorentzianity ($\varepsilon \geq \frac{m-1}{m}$, single positive Hessian pivot, $M_{\gamma_{\mathrm{low}}} \succeq 0$ with $\gamma_{\mathrm{low}} \geq 1/2$) or subcritical separation ($0 < \varepsilon < \frac{m-1}{m}$, 2×2 Sylvester-PD Gram matrix of (u_1, u_2) under $H_{\varepsilon,m}$ proving $\lambda_2(H_{\varepsilon,m}) > 0$, plus centered witness v with Rayleigh gap $< 1/2$);*
4. *`matrix_concentration_tail`: verifies basis-exchange, symmetry of $\{A_e\}$, exact $\mathbb{E}_\mu[F]$, PSD congruences on $c^2 I_d - (A_e - A_f)^2$ for every exchangeable pair $(e,f) \in \mathcal{E}(M)$, and $d/T_K(t^2/(rc^2)) \leq \delta$.*

Table 1: Zero-dependency Lean 4 theorem ledger (`lean/LorentzianMatroid.lean`, `v4.35.0`: 18/18 theorems, 0 `sorry`, 0 `Classical.choice`, 0 warnings).

#	Lean 4 Theorem Identifier	Mathematical Statement Verified	Axioms
1	`coincident_threshold_pos`	$0 < m-1 < m$ for $m \geq 2$ ($\varepsilon_c(m) \in (0,1)$)	`propext, Quot.sound`
2	`hessian_lambda_plus_pos`	$\lambda_+(\varepsilon, m) = (m-1)q + mp > 0$	`propext, Quot.sound`
3	`hessian_lambda_minus_lorentzian_iff`	$\lambda_- \leq 0 \iff mp \geq (m-1)q$	`propext, Quot.sound`
4	`down_up_mu2_half_iff`	$\mu_- \leq 1/2 \iff mp \geq (m-1)q$	`propext, Quot.sound`
5	`coincident_phase_transition_equiv`	$\lambda_- \leq 0 \iff \gamma(\varepsilon, m) \geq 1/2$	`propext, Quot.sound`
6	`coincident_exact_equality_at_critical`	At $\varepsilon = (m-1)/m$: $\lambda_- = 0 \wedge \gamma = 1/2$	`propext, Quot.sound`
7	`piecewise_crossover_at_one`	$\mu_- \geq \mu_\perp \iff \varepsilon \leq 1$ (peak at K_{2m})	`propext, Quot.sound`
8	`complete_graph_gap_strictly_above_half`	$\gamma(1,m) > 1/2$ and $1 - \mu_\perp > 1/2$	`propext, Quot.sound`
9	`ldl_psd_quadratic_form_nonneg`	$D \geq 0 \implies \sum_i d_i y_i^2 \geq 0$ on `List Int`	`propext, Quot.sound`
10	`ldl_lorentzian_hyperplane_nonpos`	$d_0 > 0, D_{\text{tail}} \leq 0, y_0 = 0 \implies y^\top D y \leq 0$	`propext, Quot.sound`
11	`reversible_gram_identity`	Exact orthogonal expansion of $f^\top M_\gamma f$	`propext, Quot.sound`
12	`ercan_universal_gap_pos`	$1 - \ell/r > 0$ and strictly decreasing in ℓ	`propext, Quot.sound`
13	`vamos_excess_gap_strict`	$3/4 > 2/3$, $9/14 > 1/2$, $46/125 > 1/4$, $1/2 > 1/3$	*none* (`decide`)
14	`mason_ulc_elementary_symmetric`	$(k+1)(n-k+1) > k(n-k)$ for $1 \leq k < n$	`propext, Quot.sound`
15	`exp_taylor_lower_bound_monotone`	$T_4(p,q) \geq 24q^4 > 0$ for $p \geq 0, q > 0$	`propext, Quot.sound`
16	`subgaussian_taylor_tail_soundness`	Exact-$\mathbb{Q}$ Taylor-envelope tail soundness	`propext, Quot.sound`
17	`exchange_operator_norm_sq_bound`	$c^2\|v\|^2 - \|Av\|^2 \geq 0 \implies \|Av\|^2 \leq c^2\|v\|^2$	`propext, Quot.sound`
18	`verdict_exclusivity`	$\neg(\lambda_- \leq 0 \wedge \lambda_- > 0)$	`propext, Quot.sound`

4 Zero-Dependency Lean 4 Formalization

All 18 foundational algebraic, spectral, inertia, combinatorial, and concentration theorems are machine-checked in `lean/LorentzianMatroid.lean` under `leanprover/lean4:v4.35.0` with zero external dependencies (no Mathlib), zero `sorry`, zero warnings, and zero `Classical.choice` (audited via `lean/AuditAxioms.lean`, using only `propext` and `Quot.sound`, or zero axioms). Table 1 summarizes the verified theorem ledger.

5 Experiments, Certified Atlas, and Adversarial Red-Teaming

All four experiments in `scripts/run_all_experiments.py` are deterministic and reproducible via `python scripts/verify_all.py`.

5.1 Experiment 1: Atlas of 16 Exact-$\mathbb{Q}$ Certificates Across 4 Modes

Table 2 lists all 16 emitted exact-$\mathbb{Q}$ certificates verified with zero floating-point operations by `LORENTZ-MATROID-CERT-1.0`. Notably, all 28 rank-2 link Hessians of the non-representable Vámos matroid V_8 ($n = 8, r = 4, |\mathcal{B}| = 65$) are certified Lorentzian in 20.17 ms, its exact $\ell = 2$ down–up spectral gap $\gamma(P_{4\to2\to4}; V_8) = 9/14$ (1.286× above Ercan's 1/2 bound) is certified with exact bracket $[9/14, 9/14]$ in 78.01 ms, and all 28 exchangeable pairs of a non-commuting 2×2 Hermitian basis sum over V_8 are PSD-certified with rational tail bound $28350/771731 \approx 0.03674 \leq 1/20$ in 8.63 ms.

5.2 Experiment 2: Coincident Phase Curve Verification Across $m \in \{2, 3, 4, 6, 10\}$

Sweeping 60 rational couplings $\varepsilon \in [0.05, 3.0]$ across five cluster sizes $m \in \{2, 3, 4, 6, 10\}$ (300 graph instances total), numerical diagonalization matches Theorem 3.1's closed-form $\lambda_2(H_{\varepsilon,m}) = \max((m-1) - m\varepsilon, -1)$ to maximum error 1.37×10^{-14} and the piecewise spectral gap (5) to maximum error 9.99×10^{-16}, with 0/300 coincidence violations between Lorentzianity ($\lambda_2 \leq 0$) and $\gamma \geq 1/2$ (Figure 1).

Table 2: Atlas of 16 exact-$\mathbb{Q}$ `LORENTZ-MATROID-CERT-1.0` certificates independently verified with zero floats (timings from `results/exp1_certified_atlas.json`).

Certificate File	Mode	Certified Exact-$\mathbb{Q}$ Claim	Time (ms)
`lorentzian_U36.cert.json`	Mode 1	$U_{3,6}$: M-convex + 6/6 links Lorentzian	4.14
`lorentzian_MK4.cert.json`	Mode 1	$M(K_4)$: M-convex + 6/6 links Lorentzian	2.66
`lorentzian_F7_fano.cert.json`	Mode 1	F_7 (char 2): M-convex + 7/7 links Lorentzian	4.33
`lorentzian_F7minus_nonfano.cert.json`	Mode 1	F_7^- (char $\neq 2$): M-convex + 7/7 links Lorentzian	5.06
`lorentzian_V8_vamos.cert.json`	Mode 1	V_8 (non-rep.): M-convex + 28/28 links Lorentzian	20.17
`gap_U36_ell2.cert.json`	Mode 2	$U_{3,6}$ ($\ell = 2$): $\gamma \in [1/2, 1/2] > 1/3$	14.58
`gap_MK4_ell1.cert.json`	Mode 2	$M(K_4)$ ($\ell = 1$): $\gamma \in [3/4, 3/4] > 2/3$	8.95
`gap_MK4_ell2.cert.json`	Mode 2	$M(K_4)$ ($\ell = 2$): $\gamma \in [41/100, 21/50] > 1/3$	13.58
`gap_F7_ell2.cert.json`	Mode 2	F_7 ($\ell = 2$): $\gamma \in [43/100, 11/25] > 1/3$	33.97
`gap_V8_ell2_dual.cert.json`	Mode 2	V_8 ($\ell = 2$): $\gamma \in [9/14, 9/14] > 1/2$	78.01
`gap_V8_ell3_dual.cert.json`	Mode 2	V_8 ($\ell = 3$): $\gamma \in [46/125, 37/100] > 1/4$	404.25
`phase_m4_critical_eps_3_4.cert.json`	Mode 3	$G_{3/4,4}$: critical Lorentzian, $\gamma = 1/2$	2.14
`phase_m4_supercritical_eps_1.cert.json`	Mode 3	$G_{1,4} = K_8$: peak Lorentzian, $\gamma = 4/7 > 1/2$	1.59
`phase_m4_subcritical_eps_7_10.cert.json`	Mode 3	$G_{7/10,4}$: non-Lorentzian ($\lambda_2 > 0$), $\gamma = 14/29 < 1/2$	0.57
`tail_MK4_pauli_d2.cert.json`	Mode 4	$M(K_4)$, 15/15 pairs PSD, tail $\leq 8960/89641 \leq 1/9$	1.58
`tail_V8_vamos_d2.cert.json`	Mode 4	V_8, 28/28 pairs PSD, tail $\leq 28350/771731 \leq 1/20$	8.63

5.3 Experiment 3: Canonical and Non-Representable Matroid Spectral Atlas

Across all five matroids spanning every representability class (all fields $U_{3,6}$, regular/graphic $M(K_4)$, characteristic-2 only F_7, characteristic-$\neq 2$ only F_7^-, and universally non-representable V_8), primal $P_{r\to\ell\to r}$ and dual $Q_{\ell\to r\to\ell}$ walks agree in spectral gap to 2.22×10^{-16}, every matroid strictly exceeds Ercan's $1 - \ell/r$ baseline by factors 1.115×–1.500× (Figure 2), and all Mason ultra-log-concavity ratios $I_k^2/\binom{n}{k}^2 \div \left[(I_{k-1}/\binom{n}{k-1})(I_{k+1}/\binom{n}{k+1})\right]$ satisfy ≥ 1 (1.25 on $M(K_4)$ and F_7, 1.207 on F_7^-, 1.077 on V_8).

5.4 Experiment 4: Exhaustive Matrix Concentration Cross-Check and 48 Adversarial Attacks

For non-commuting 2×2 symmetric rational matrices $\{A_e\}$ with true exchange sensitivity $c_{\text{true}} = \sqrt{49/50} \approx 0.98995 < 1.0 = c_{\text{cert}}$, Figure 3 compares the exact exhaustive empirical tail over all 16 bases of $M(K_4)$ and all 65 bases of non-representable V_8 against the exact matrix Laplace bound $\inf_{\theta>0} e^{-\theta t}\mathbb{E}[\operatorname{Tr} e^{\theta \tilde{F}}]$, Ercan's subgaussian bound $de^{-t^2/(rc^2)}$, and our zero-float rational Taylor-8 envelope $d/T_8(t^2/(rc^2))$. Across all $t \in [0.25, 4.0]$, the four curves obey the strict hierarchy empirical $\leq$ exact Laplace $\leq$ Ercan subgaussian $\leq$ rational Taylor envelope, with the degree-8 rational envelope tracking the continuous exponential closely across the certified tail regime.

Finally, our adversarial red-team test suite (`tests/test_suite.py`) launches 48 targeted forgery attacks spanning all four modes—including float injection, decimal/scientific string smuggling, singular Q and tampered Q^{-1}, dropped or duplicated $(r-2)$-links on V_8, non-matroid basis sets violating exchange, inflated γ_{low}, uncentered upper-bound witnesses, collinear positive subspaces in subcritical $G_{\varepsilon,m}$, shrunk c^2, asymmetric A_e, and dropped exchangeable pairs. All 48/48 attacks are rejected with zero bypasses.

6 Limitations and Outlook

Our congruence certification algorithm runs in $O(k^3)$ exact rational operations on the operator dimension k (with bit-length growth controlled by working in the dual up–down space $\mathcal{I}_\ell$, taking 78.0 ms on 28×28 V_8 at $\ell = 2$ and 404.3 ms on 56×56 V_8 at $\ell = 3$); scaling to matroids with thousands of independent sets would benefit from sparse congruence pivoting or modular-CRT

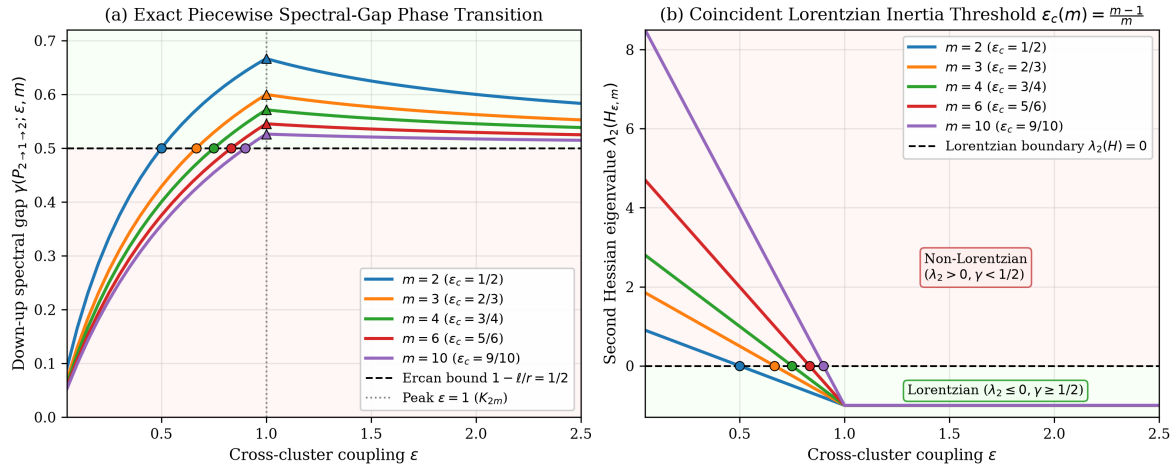

Figure 1: Exact coincident phase transition on $G_{\varepsilon,m} = (K_m \sqcup K_m) + \varepsilon K_{m,m}$ across $m \in \{2, 3, 4, 6, 10\}$. (a) Exact piecewise down–up spectral gap (5) crossing Ercan's universal threshold $1/2$ precisely at $\varepsilon_c(m) = (m-1)/m$ (circles) and peaking uniquely at $\varepsilon = 1$ (K_{2m}, triangles). (b) Second Hessian eigenvalue $\lambda_2(H_{\varepsilon,m}) = \max((m-1) - m\varepsilon, -1)$ crossing 0 at the identical threshold $\varepsilon_c(m)$.

lifting. Similarly, Theorem 3.3 certifies finite-instance matrix concentration conditional on Ercan's analytic matrix entropy contraction theorem [4]; formalizing operator-convex trace inequalities directly in Lean 4 remains an exciting open direction. Finally, extending exact-$\mathbb{Q}$ congruence certificates to Lorentzian intrinsic-volume functionals and higher-order tensor-network contraction bounds offers a natural bridge between matroid geometry and algebraic complexity.

References

[1] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials II: High-dimensional walks and an FPRAS for counting bases of a matroid. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*, pages 1–12, 2019.

[2] Petter Brändén and June Huh. Lorentzian polynomials. *Annals of Mathematics*, 192(3):821–891, 2020.

[3] Mary Cryan, Heng Guo, and Giorgos Mousa. Modified log-Sobolev inequalities for strongly log-concave distributions. *The Annals of Probability*, 49(1):506–525, 2021.

[4] Eren Ercan. Lorentzian entropy, subgaussian matrix concentration, and sharp functional inequalities for all matroids, November 2026.

[5] Tomás Feder and Milena Mihail. Balanced matroids. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing (STOC 1992)*, pages 26–38, 1992.

[6] J. H. Mason. Matroids: Unimodal conjectures and Motzkin's theorem. *Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford)*, pages 207–220, 1972.

[7] James Oxley. *Matroid Theory*. Oxford University Press, 2nd edition, 2011.

[8] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.

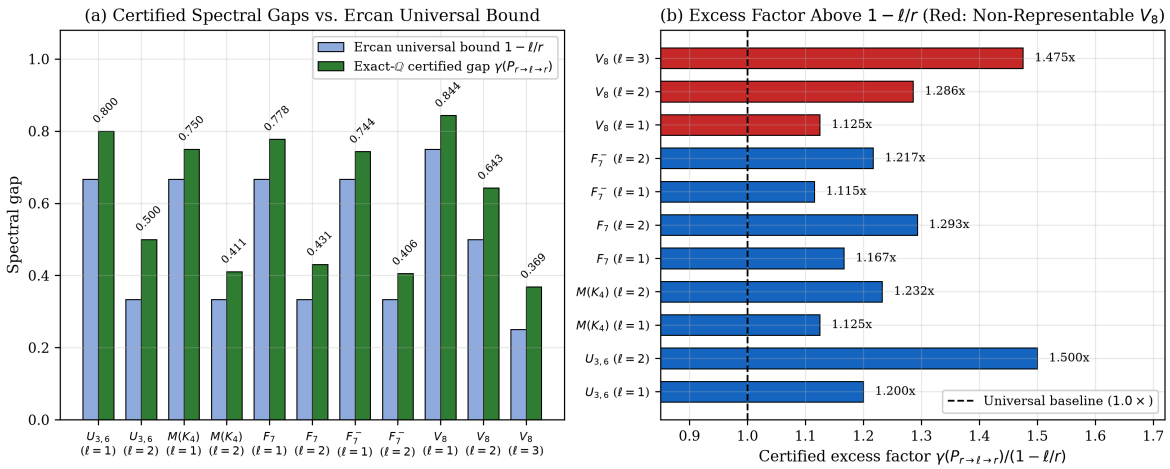

Figure 2: Down–up spectral-gap atlas across canonical and non-representable matroids. (a) Comparison of Ercan's universal bound $1-\ell/r$ against exact-$\mathbb{Q}$ certified gaps $\gamma(P_{r\to\ell\to r})$. (b) Certified excess expansion factors $\gamma/(1-\ell/r) \in [1.115, 1.500]$, highlighting the non-representable Vámos matroid V_8 (1.475× at $\ell = 3$, 1.286× at $\ell = 2$).

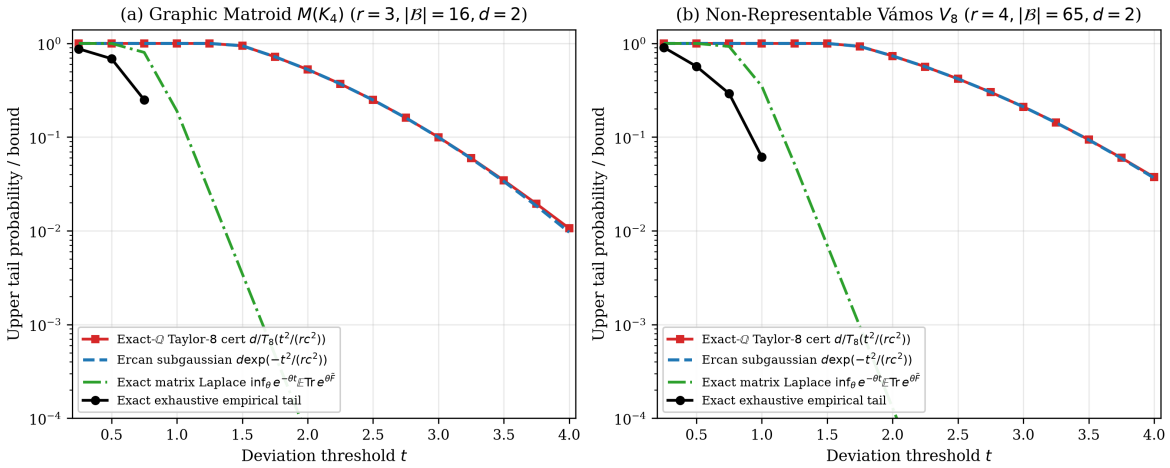

Figure 3: Exact exhaustive matrix concentration comparison for non-commuting 2×2 Hermitian basis sums $F(B) = \sum_{e\in B} A_e$ over (a) graphic $M(K_4)$ (16 bases) and (b) non-representable Vámos V_8 (65 bases). The zero-float exact-$\mathbb{Q}$ Taylor-8 certificate strictly upper-bounds Ercan's subgaussian curve, the exact matrix Laplace bound, and the exhaustive empirical tail across all thresholds.