Quantum Spectral Christoffel–Super-Poincaré Inequalities, Exact Piecewise-Logarithmic Mixing Laws, and Machine-Checked Rapid-Mixing Certificates Beyond the Spectral Gap

Lluis Eriksson*

November 24, 2026

Abstract

Proving rapid mixing of quantum Gibbs samplers beyond high-temperature regimes was recently revolutionized by Chen and Ding [6], who established $O(\log^4 n + \log(1/\varepsilon))$ mixing for one-dimensional KMS-detailed-balanced Gibbs samplers by showing that states far from equilibrium carry large weight outside low-weight operator subspaces where stronger dissipation holds. Turning this qualitative mechanism into *exact, non-asymptotic, constant-free* mixing-time certificates for general non-commuting Hamiltonians—without relying on asymptotic hypercontractivity constants and with exact-$\mathbb{Q}$ verification of the quantized spectral-gap and Christoffel bounds—remains open. We resolve this problem with five theorems, a zero-dependency Lean 4 formalization, and an exact-$\mathbb{Q}$ LDL$^\top$ certificate verifier. First, for any full-rank Gibbs state σ and any KMS-symmetric Markov generator with subspace $I \in V \subset \mathcal{B}(\mathcal{H})$ and restricted gap $\lambda(V^\perp)$, we introduce the basis-independent *Quantum Christoffel Operator* $\mathcal{K}_V := \sum_{a=1}^{D} E_a^\dagger E_a$ and prove the exact Quantum Super-Poincaré Inequality $\mathcal{E}(X_\rho, X_\rho) \geq \lambda(V^\perp)\big(\chi^2_\sigma(\rho\|\sigma) - (\|\mathcal{K}_V\|_\infty - 1)\big)_+$ for every quantum state ρ. Second, for any subspace ladder (Pauli-weight or exact KMS eigenmodes), we integrate the upper convex envelope $\frac{dy}{dt} \leq -2\max_k \lambda_k(y - C_k)$ in closed form across exact knots $y_j^* = (\lambda_{k_j} C_{k_j} - \lambda_{k_{j-1}} C_{k_{j-1}})/(\lambda_{k_j} - \lambda_{k_{j-1}})$, proving strict speedup over T_{gap} whenever $y_0 > y_1^*$ and a universal factor-of-2 optimality law ($T_{\text{SP}}/t_{\text{exact}} \to 1$ as $\varepsilon \to 0^+$ or $y_0 \to \infty$) on 2-tier commuting semigroups. Third, we prove an explicit, constant-free $T_{\text{SP}}(n, \varepsilon) \leq \frac{1}{4\gamma}(\log n)^2 + O(\gamma^{-1}(\log n + \log(1/\varepsilon)))$ mixing bound for linear Pauli-weight ladders using pure $L^1 \to L^2$ projection (zero hypercontractivity, with universal $c = 3$ for all quantum states and tighter $c = 1$ for product-state trajectories under product semigroups), alongside the exact two-regime α-Super-Poincaré phase diagram across $\alpha \in (0, 1)$, $\alpha = 1$, and $\alpha > 1$. Fourth, we lift every convex-hull bound directly to non-asymptotic trace-norm and relative-entropy decay $D(\rho_t\|\sigma) \leq D_2(\rho_t\|\sigma) \leq \log(1 + y_{\text{env}}(t))$. Fifth, we construct cryptographically linked global operator-order exact-$\mathbb{Q}$ LDL$^\top$ PSD witnesses certifying both subspace gaps and Christoffel norms strictly over $\mathbb{Q}$ on the quantized operator bundle given perturbation radii $\delta_*^{\mathbb{Q}}$ (with our synthesis script estimating $\delta_*^{\mathbb{Q}}$ against double-precision `eigh(H)` for interacting chains, while analytical depolarizing certificates are end-to-end exact-$\mathbb{Q}$ with zero floating-point input). All 16 core algebraic theorems are formalized in Lean 4 (`v4.35.0`, 0 `sorry`, 0 `Classical.choice`) and benchmarked on non-commuting 1D tilted/ferromagnetic TFIM (1.26×–3.30× speedup for $n = 3, 4$; 1.65×–3.30× at $n = 4$) and XXZ (1.57×) KMS Gibbs samplers, $n = 10$–$1{,}000$ depolarizing chains (17.7× universal $c = 3$ and 20.7× product-state $c = 1$ speedup), and 6 exact-$\mathbb{Q}$ certificates containing 14 verified linked LDL$^\top$ PSD witnesses.

*Independent researcher (Sweden). Research, zero-dependency Lean 4 proofs, exact-$\mathbb{Q}$ LDL$^\top$ and convex-hull mixing-time verifiers, and benchmarks prepared with AI assistance (Google Antigravity) and publicly archived at https://github.com/lluiseriksson/beyond-gap-gibbs-certificates.

1 Introduction

Preparing thermal Gibbs states $\sigma = e^{-\beta H}/\operatorname{Tr}(e^{-\beta H})$ of interacting quantum spin systems via dissipative Lindbladian evolution is a central primitive of quantum simulation and quantum statistical mechanics [7, 5, 4]. Traditional convergence bounds rely on the Poincaré spectral gap μ_1 of the Kubo–Martin–Schwinger (KMS)-symmetric Lindbladian $\mathcal{L}^\dagger$, yielding the worst-case mixing time $T_{\mathrm{gap}}(\varepsilon) = \frac{1}{2\mu_1}\log(\chi^2_\sigma(\rho_0\|\sigma)/\varepsilon^2) = O(n + \log(1/\varepsilon))$ [10] because $\chi^2_\sigma(\rho_0\|\sigma) \le \lambda_{\min}(\sigma)^{-1} - 1$ grows exponentially in system size n. Overcoming this linear-n bottleneck typically required proving a uniform quantum Modified Logarithmic Sobolev Inequality (MLSI) [8], which had remained open for non-commuting spin chains at arbitrary inverse temperature $\beta > 0$.

In a major recent breakthrough, Chen and Ding [6] proved rapid mixing $O(\log^4 n + \log(1/\varepsilon))$ for one-dimensional non-commuting KMS quantum Gibbs samplers at every positive temperature via *spectral analysis beyond the spectral gap*: any state far from σ must carry significant χ^2_σ weight outside the span of low-weight Pauli operators, where higher-weight dissipation rates drive a rapid initial (double-exponential) burn-in before the spectral gap μ_1 governs the final exponential tail.

To convert this paradigm into an exact, assumption-light, finite-size certification engine—analogous to classical Super-Poincaré theory [11, 3] and complementing our recent machine-checked spectral certificates [1, 2]—three questions arise:

1. **Exact Non-Commuting Projection Geometry:** What is the sharpest basis-independent operator invariant controlling the KMS $L^1(\sigma) \to L^2(\sigma)$ projection norm on a subspace $V \subset \mathcal{B}(\mathcal{H})$?
2. **Closed-Form Mixing Law & Sharpness:** Given a discrete ladder of subspace gaps and projection norms $\{(\lambda_k, \overline{\Lambda}_k)\}_{k=0}^m$, what is the exact closed-form mixing time, how tight is it against true semigroup trajectories, and what scaling does pure $L^1 \to L^2$ projection achieve without hypercontractivity?
3. **Zero-Floating-Point-Trust Certification:** Can both the higher-tier spectral gaps λ_k, the Christoffel norms $\overline{\Lambda}_k$, and the resulting mixing time $T_{\mathrm{SP}}(\varepsilon)$ be certified strictly in exact rational arithmetic $\mathbb{Q}$ with machine-checked Lean 4 soundness?

We answer all three questions completely in Sections 2–6.

2 Quantum Spectral Christoffel–Super-Poincaré Inequalities

Let $\mathcal{H}$ be a finite-dimensional Hilbert space ($d = \dim \mathcal{H}$) and $\sigma \succ 0$ ($\operatorname{Tr}\sigma = 1$) a full-rank Gibbs state. Equip $\mathcal{B}(\mathcal{H})$ with the KMS inner product and norm

$$\langle X, Y\rangle_{\mathrm{KMS}} \coloneqq \operatorname{Tr}(\sigma^{1/2}X^\dagger\sigma^{1/2}Y), \qquad \|X\|^2_{2,\sigma} \coloneqq \langle X, X\rangle_{\mathrm{KMS}}. \tag{1}$$

For any quantum state $\rho \succeq 0$ ($\operatorname{Tr}\rho = 1$), its KMS relative density $X_\rho \coloneqq \sigma^{-1/2}\rho\sigma^{-1/2} \succeq 0$ satisfies $\langle I, X_\rho\rangle_{\mathrm{KMS}} = \operatorname{Tr}\rho = 1$, $\|X_\rho\|^2_{2,\sigma} = \operatorname{Tr}((\sigma^{-1/4}\rho\sigma^{-1/4})^2) = e^{D_2(\rho\|\sigma)}$ (where D_2 is the sandwiched Rényi-2 divergence [9]), and

$$\chi^2_\sigma(\rho\|\sigma) \coloneqq \|X_\rho - I\|^2_{2,\sigma} = e^{D_2(\rho\|\sigma)} - 1. \tag{2}$$

Moreover, setting $B \coloneqq \sigma^{1/4}(X_\rho - I)\sigma^{1/4} = \sigma^{-1/4}(\rho - \sigma)\sigma^{-1/4}$, matrix Hölder's inequality ($\|UVW\|_1 \le \|U\|_4\|V\|_2\|W\|_4$) gives the exact constant-1 trace-norm domination

$$\|\rho - \sigma\|_1 = \|\sigma^{1/4}B\sigma^{1/4}\|_1 \le (\operatorname{Tr}\sigma)^{1/4}\|B\|_2(\operatorname{Tr}\sigma)^{1/4} = \sqrt{\chi^2_\sigma(\rho\|\sigma)} \tag{3}$$

(discrete weighted form proved in Lean 4: `trace_norm_sq_le_chi2`).

Let $\mathcal{L}^\dagger$ be a Heisenberg-picture Lindbladian that is KMS-symmetric ($\langle X, \mathcal{L}^\dagger Y\rangle_{\mathrm{KMS}} = \langle \mathcal{L}^\dagger X, Y\rangle_{\mathrm{KMS}}$) and primitive ($\ker\mathcal{L}^\dagger = \mathbb{C}I$), with nonnegative Dirichlet form $\mathcal{E}(X, X) \coloneqq \langle X, (-\mathcal{L}^\dagger)X\rangle_{\mathrm{KMS}} \ge 0$. Under $\rho_t = e^{t\mathcal{L}}(\rho_0)$, $X_{\rho_t} = e^{t\mathcal{L}^\dagger}(X_{\rho_0})$ obeys $\frac{d}{dt}\chi^2_\sigma(\rho_t\|\sigma) = -2\mathcal{E}(X_{\rho_t}, X_{\rho_t})$.

2.1 Basis-Independent Quantum Christoffel Operator and Super-Poincaré Bound

Definition 2.1 (Quantum Christoffel Operator)**.** Let $V \subset \mathcal{B}(\mathcal{H})$ be any subspace containing I, and let $\{E_1, \dots, E_D\}$ be any KMS-orthonormal basis of V ($\langle E_a, E_b\rangle_{\mathrm{KMS}} = \delta_{ab}$). The *Quantum Christoffel Operator* of V is

$$\mathcal{K}_V \;\coloneqq\; \sum_{a=1}^{D} E_a^\dagger E_a \;\succeq\; I, \qquad \overline{\Lambda}_V \;\coloneqq\; \|\mathcal{K}_V\|_\infty, \qquad \Lambda_V \;\coloneqq\; \sup_{\|\psi\|=1} \sum_{a=1}^{D} |\langle\psi|E_a|\psi\rangle|^2.$$

Note that $\mathcal{K}_V$ is strictly invariant under any unitary change of KMS-orthonormal basis $E'_a = \sum_b U_{ab} E_b$ (since $U^\dagger U = I_D$), and when $\sigma = I/2^n$ with $V = \mathrm{span}\{P : |P| \le k\}$, $\mathcal{K}_V = (\sum_{j=0}^{k} \binom{n}{j} 3^j) I$.

Theorem 2.1 (Quantum Spectral Christoffel–Super-Poincaré Inequality)**.** *Let $I \in V \subset \mathcal{B}(\mathcal{H})$ have* KMS*-orthogonal projection P_V, and define its* operator-order dissipation constant

$$\lambda_{\mathrm{op}}(V^\perp) \;\coloneqq\; \sup\{\lambda \ge 0 : -\mathcal{L}^\dagger \succeq \lambda(I - P_V) \text{ in } L^2(\mathrm{KMS})\} \;=\; \inf_{X \notin V} \frac{\mathcal{E}(X,X)}{\|(I-P_V)X\|_{2,\sigma}^2}.$$

Then for every *quantum state ρ:*

(a) $1 \le \Lambda_V \le \overline{\Lambda}_V$*, and the centered projection satisfies* $\|P_V(X_\rho - I)\|_{2,\sigma}^2 \le \Lambda_V - 1 \le \overline{\Lambda}_V - 1$ *(see Table 1).*

(b) The Dirichlet form obeys the exact Quantum Super-Poincaré Inequality (Table 1):

$$\mathcal{E}(X_\rho, X_\rho) \;\ge\; \lambda_{\mathrm{op}}(V^\perp) \left(\chi^2_\sigma(\rho\|\sigma) - (\overline{\Lambda}_V - 1)\right)_+. \tag{4}$$

In particular, writing $V = \mathbb{C}I \oplus V_+$, $\lambda_{\mathrm{op}}(V^\perp) = \lambda_{\min}(M_{\perp\perp} - M_{\perp V_+} M_{V_+V_+}^{-1} M_{V_+\perp})$ is the exact Schur complement, coinciding identically with the $(r+1)$-th eigenvalue μ_{r+1} on every $\mathcal{L}^\dagger$-invariant eigenmode subspace $V_r = \mathrm{span}\{E_0, \dots, E_r\}$.

(c) For any subspace filtration $I \in V_0 \subset V_1 \subset \cdots \subset V_m$ with $C_k \coloneqq \overline{\Lambda}_{V_k} - 1$ and $\lambda_k \coloneqq \lambda_{\mathrm{op}}(V_k^\perp)$, the piecewise-linear convex envelope satisfies (Lean: `super_poincare_envelope_two`*)*

$$\mathcal{E}(X_\rho, X_\rho) \;\ge\; \Phi_{\mathrm{env}}(\chi^2_\sigma(\rho\|\sigma)) \;\coloneqq\; \max_{0 \le k \le m} \lambda_k(\chi^2_\sigma(\rho\|\sigma) - C_k). \tag{5}$$

Proof. For any state ρ, $\langle E_a, X_\rho\rangle_{\mathrm{KMS}} = \mathrm{Tr}(\sigma^{1/2} E_a^\dagger \sigma^{1/2} \sigma^{-1/2} \rho \sigma^{-1/2}) = \mathrm{Tr}(E_a^\dagger \rho)$, hence $\|P_V X_\rho\|_{2,\sigma}^2 = \sum_{a=1}^{D} |\,\mathrm{Tr}(E_a^\dagger \rho)|^2$. Because $\rho \mapsto |\,\mathrm{Tr}(E_a^\dagger \rho)|^2$ is convex on the compact convex state space, the maximum of the sum is attained at an extreme point $\rho = |\psi\rangle\langle\psi|$ ($\|\psi\| = 1$), giving $\|P_V X_\rho\|_{2,\sigma}^2 \le \Lambda_V$. By Cauchy–Schwarz, $|\langle\psi|E_a|\psi\rangle|^2 \le \|E_a|\psi\rangle\|^2 = \langle\psi|E_a^\dagger E_a|\psi\rangle$; summing over $a = 1, \dots, D$ yields $\Lambda_V \le \sup_{\|\psi\|=1}\langle\psi|\mathcal{K}_V|\psi\rangle = \overline{\Lambda}_V$. Since $I \in V$ and $\langle I, X_\rho - I\rangle_{\mathrm{KMS}} = 0$, orthogonality gives $\|P_V X_\rho\|_{2,\sigma}^2 = 1 + \|P_V(X_\rho - I)\|_{2,\sigma}^2$, proving (a). For (b), by definition of $\lambda_{\mathrm{op}}(V^\perp)$,

$$\mathcal{E}(X_\rho, X_\rho) \;=\; \mathcal{E}(X_\rho - I, X_\rho - I) \;\ge\; \lambda_{\mathrm{op}}(V^\perp)\, \|(I - P_V)(X_\rho - I)\|_{2,\sigma}^2.$$

Pythagoras yields $\|(I-P_V)(X_\rho - I)\|_{2,\sigma}^2 = \chi^2_\sigma(\rho\|\sigma) - \|P_V(X_\rho - I)\|_{2,\sigma}^2 \ge \chi^2_\sigma(\rho\|\sigma) - (\overline{\Lambda}_V - 1)$, proving (4) and (5). The Schur-complement formula follows from block positivity of $M_{\mathrm{KMS}} - \lambda(I - P_V)$ on $(\mathbb{C}I)^\perp = V_+ \oplus V^\perp$ (where $M_{V_+V_+} \succ 0$ by primitivity), reducing to $M_{\perp\perp} \succeq \mu_{r+1} I_\perp$ whenever $M_{\perp V_+} = 0$. □

2.2 Exact Upper-Convex-Hull Piecewise-Logarithmic Mixing Law and Sharpness

Theorem 2.2 (Exact Piecewise-Logarithmic Mixing Law and 2-Tier Optimality)**.** *Let* $\{(\lambda_{k_j}, C_{k_j})\}_{j=0}^{r}$ *($0 < \lambda_{k_0} < \cdots < \lambda_{k_r}$, $0 = C_{k_0} < \cdots < C_{k_r}$) be the minimal active upper convex hull of the affine functions* $\ell_k(y) = \lambda_k(y - C_k)$ *on* $(0, \infty)$ *satisfying (5), with transition knots* $y_0^* = 0$, $y_{r+1}^* = +\infty$, *and* $y_j^* = (\lambda_{k_j} C_{k_j} - \lambda_{k_{j-1}} C_{k_{j-1}})/(\lambda_{k_j} - \lambda_{k_{j-1}})$ *($1 \le j \le r$). Then:*

(a) ***Well-Posedness and Exact Mixing Time.*** *The knots satisfy* $0 < y_1^* < \cdots < y_r^*$ *and* $y_j^* > C_{k_j}$ *for all* $j \ge 1$ *(Table 1). For every initial state* ρ_0 *with* $y_0 = \chi^2_\sigma(\rho_0 \| \sigma) > \varepsilon^2 > 0$, *the exact piecewise-linear envelope solution* $y_{\mathrm{env}}(t)$ *satisfies* $\chi^2_\sigma(\rho_t \| \sigma) \le y_{\mathrm{env}}(t)$ *for all* $t \ge 0$, *and hence* $\|\rho_t - \sigma\|_1 \le \varepsilon$ *for all*

$$t \;\ge\; T_{\mathrm{SP}}(\varepsilon; y_0) \;:=\; \sum_{j=0}^{r} \frac{1}{2\lambda_{k_j}} \log\!\left(\frac{\min(y_0, y_{j+1}^*) - C_{k_j}}{\max(\varepsilon^2, y_j^*) - C_{k_j}} \right) \mathbf{1}_{\min(y_0, y_{j+1}^*) > \max(\varepsilon^2, y_j^*)}. \tag{6}$$

(b) ***Strict Speedup over Spectral Gap.*** *Whenever* $y_0 > y_1^*$ *and* $\lambda_{k_1} > \lambda_{k_0}$, *we have* $T_{\mathrm{SP}}(\varepsilon; y_0) < T_{\mathrm{gap}}(\varepsilon; y_0) := \frac{1}{2\lambda_{k_0}} \log(y_0/\varepsilon^2)$ *strictly (Table 1).*

(c) ***Universal Factor-of-2 Sharpness on 2-Tier Semigroups.*** *For any* $0 < \lambda_0 < \lambda_1$ *and* $C_1 > \varepsilon^2$, *consider the 2-tier commuting semigroup with* $\chi^2(t) = C_1 e^{-2\lambda_0 t} + (y_0 - C_1) e^{-2\lambda_1 t}$ *($y_0 > y_1^*$) and exact* χ^2*-hitting time* $t_{\mathrm{exact}}(\varepsilon)$. *Then*

$$t_{\mathrm{exact}}(\varepsilon) \;\le\; T_{\mathrm{SP}}(\varepsilon; y_0) \;\le\; 2\, t_{\mathrm{exact}}(\varepsilon) + \frac{1}{2\lambda_0} \log \frac{\lambda_1}{\lambda_1 - \lambda_0} + \frac{1}{2\lambda_1} \log_+\!\left(\frac{\lambda_1 - \lambda_0}{\lambda_0} \right), \tag{7}$$

and $T_{\mathrm{SP}}(\varepsilon; y_0)/t_{\mathrm{exact}}(\varepsilon) \to 1$ *in both asymptotic regimes* $\varepsilon \to 0^+$ *(fixed* y_0*) and* $y_0 \to \infty$ *(fixed* ε*).*

Proof. Because $\Phi_{\mathrm{env}}(y)$ is continuous, strictly increasing, and piecewise linear on $(0, \infty)$ with active slope λ_{k_j} and intercept $-\lambda_{k_j} C_{k_j}$ on $[y_j^*, y_{j+1}^*]$ (Lean: `hull_tier2_dominates_above`, `hull_tier1_dominates_below`), comparison of $\frac{dy}{dt} \le -2\Phi_{\mathrm{env}}(y)$ gives $\chi^2_\sigma(\rho_t \| \sigma) \le y_{\mathrm{env}}(t)$, where the time required for y_{env} to descend from y_0 to ε^2 is $\int_{\varepsilon^2}^{y_0} \frac{dy}{2\Phi_{\mathrm{env}}(y)}$, which evaluates term-by-term to (6). Part (b) follows immediately from $\frac{1}{2\Phi_{\mathrm{env}}(y)} < \frac{1}{2\lambda_{k_0} y}$ on $(y_1^*, y_0]$. For (c), $t_{\mathrm{exact}} \le T_{\mathrm{SP}}$ holds by (a). Since $C_1 e^{-2\lambda_0 t_{\mathrm{exact}}} \le \varepsilon^2$ and $(y_0 - C_1) e^{-2\lambda_1 t_{\mathrm{exact}}} \le \varepsilon^2$, we have $t_{\mathrm{exact}} \ge \max(\frac{1}{2\lambda_0} \log \frac{C_1}{\varepsilon^2},\ \frac{1}{2\lambda_1} \log \frac{y_0 - C_1}{\varepsilon^2})$. Writing $y_1^* = \frac{\lambda_1}{\lambda_1 - \lambda_0} C_1$ and $y_1^* - C_1 = \frac{\lambda_0}{\lambda_1 - \lambda_0} C_1$,

$$\begin{aligned} T_{\mathrm{SP}}(\varepsilon; y_0) &= \frac{1}{2\lambda_0} \log \frac{C_1}{\varepsilon^2} + \frac{1}{2\lambda_0} \log \frac{\lambda_1}{\lambda_1 - \lambda_0} + \frac{1}{2\lambda_1} \log \frac{y_0 - C_1}{\varepsilon^2} + \frac{1}{2\lambda_1} \log \frac{\varepsilon^2}{y_1^* - C_1} \\ &\le 2\, t_{\mathrm{exact}}(\varepsilon) + \frac{1}{2\lambda_0} \log \frac{\lambda_1}{\lambda_1 - \lambda_0} + \frac{1}{2\lambda_1} \log_+\!\left(\frac{\lambda_1 - \lambda_0}{\lambda_0} \right), \end{aligned}$$

using $\varepsilon^2 < C_1$ in the last step. When $\varepsilon \to 0^+$ (fixed y_0), the first term dominates both t_{exact} and T_{SP}; when $y_0 \to \infty$ (fixed ε), the third term dominates both, yielding $T_{\mathrm{SP}}/t_{\mathrm{exact}} \to 1$ in either limit. □

3 Pure $L^1 \to L^2$ Poly-Logarithmic Speedup and Two-Regime Phase Diagram

Theorem 3.1 (Pure $L^1 \to L^2$ $(\log n)^2$ Scaling and α-Phase Diagram)**.** *(a)* ***Constant-Free*** $O(\log^2 n)$ ***Bound Without Hypercontractivity.*** *Suppose an n-qubit ($n \ge 2$) KMS-symmetric semigroup with* $y_0 = 2^n - 1$ *satisfies (5) along a trajectory* ρ_t *with linear*

Pauli-weight dissipation $\lambda_k \geq \gamma(k+1)$ *(*$\gamma > 0$*) and Christoffel bounds* $C_k \leq \sum_{j=0}^{k} \binom{n}{j} c^j - 1$ *(*$c \geq 1$*). Specifically: (i)* $c = 3$ *holds* universally for every (possibly entangled) quantum state ρ *via Definition 2.1 (*$\mathcal{K}_{V_k} = (\sum_{j=0}^{k} \binom{n}{j} 3^j) I$*); and (ii)* $c = 1$ *holds* for every product-state trajectory $\rho_t = \bigotimes_{i=1}^{n} \rho_{i,t}$ *(which is invariant under product/local depolarizing semigroups), because each single-qubit Bloch vector satisfies* $\langle X_i \rangle^2 + \langle Y_i \rangle^2 + \langle Z_i \rangle^2 \leq 1$ *and multiplicativity yields* $\sum_{|P|=j} |\operatorname{Tr}(P\rho_t)|^2 = \sum_{|S|=j} \prod_{i \in S} (\langle X_i \rangle^2 + \langle Y_i \rangle^2 + \langle Z_i \rangle^2) \leq \binom{n}{j}$ *(whereas general entangled states can exceed* $c = 1$*, e.g. the 5-qubit* $5, 1, 3$ *code state has weight-4 sector sum* $S_4 = 15 > \binom{5}{4} = 5$*, requiring the universal* $c = 3$ *bound). Then for every* $\varepsilon \in (0, 1)$*,*

$$T_{\mathrm{SP}}(\varepsilon; y_0) \;\leq\; \frac{1}{4\gamma} (\log n)^2 + \frac{1 + \log(2(c+1))}{2\gamma} \log n + \frac{\log(2c/\varepsilon^2)}{2\gamma}, \tag{8}$$

an exponential $O(n/\log^2 n)$ *speedup over* $T_{\mathrm{gap}}(\varepsilon; y_0) = \frac{1}{2\gamma} \log((2^n - 1)/\varepsilon^2) = \Theta(n/\gamma)$ *without invoking hypercontractivity.*

(b) ***Two-Regime*** α***-Super-Poincaré Phase Diagram.*** *Suppose a bounded KMS-symmetric semigroup has spectral gap* $\mu_1 > 0$ *(so* $\mathcal{E}(X_\rho, X_\rho) \geq \mu_1 y$ *for all* $y \geq 0$*) and additionally satisfies the burn-in Super-Poincaré bound* $\mathcal{E}(X_\rho, X_\rho) \geq \gamma(1+y)(\log(1+y))^\alpha$ *(*$\gamma > 0, \alpha > 0$*) whenever* $D_2(\rho\|\sigma) \coloneqq \log(1+y) \geq 1$*. Then during the burn-in stage (*$D_2(t) \geq 1$*),* $\frac{d}{dt} D_2(t) \leq -2\gamma D_2(t)^\alpha$ *integrates in closed form, and pairing with* $\mu_1 > 0$ *below* $D_2 = 1$ *yields three exact non-vacuous regimes:*

- Sub-logarithmic ($\alpha \in (0, 1)$): $D_2(t)^{1-\alpha} \leq D_2(0)^{1-\alpha} - 2\gamma(1-\alpha)t$ *while* $D_2(t) \geq 1$*, reaching* $D_2 = 1$ *in time* $\leq \frac{D_2(0)^{1-\alpha} - 1}{2\gamma(1-\alpha)}$*, hence* $T_{\mathrm{mix}}(\varepsilon) = O(\gamma^{-1} D_2(0)^{1-\alpha} + \mu_1^{-1} \log(1/\varepsilon))$.
- Quantum Rényi-2 LSI ($\alpha = 1$): $D_2(t) \leq e^{-2\gamma t} D_2(0)$ *while* $D_2(t) \geq 1$ *(and globally if the bound holds on all* $y \geq 0$*), yielding* $T_{\mathrm{mix}}(\varepsilon) \leq \frac{1}{2\gamma} \log_+ D_2(0) + \frac{1}{2\mu_1} \log((e-1)/\varepsilon^2) = O(\log n + \log(1/\varepsilon))$.
- Ultracontractive ($\alpha > 1$): $D_2(t) \leq \left(D_2(0)^{-(\alpha-1)} + 2\gamma(\alpha-1)t\right)^{-1/(\alpha-1)} \leq \left(2\gamma(\alpha-1)t\right)^{-1/(\alpha-1)}$ *while* $D_2(t) \geq 1$*, giving uniform* y_0*-independent burn-in down to* $D_2 = 1$ *in time* $\leq \left(2\gamma(\alpha-1)\right)^{-1}$ *followed by exponential tail decay* $\frac{1}{2\mu_1} \log((e-1)/\varepsilon^2)$.

Proof. For (a), let $B_k \coloneqq \sum_{j=0}^{k} \binom{n}{j} c^j - 1 \geq C_k$, and let $m \geq 1$ be the largest integer with $2B_m \leq y_0$ (if $y_0 < 2B_1 = 2cn$, tier $k = 0$ alone gives $T_{\mathrm{SP}} \leq \frac{1}{2\gamma} \log(2cn/\varepsilon^2)$, satisfying (8)). Partition $[\varepsilon^2, y_0]$ at the split points $z_k = 2B_k$ ($1 \leq k \leq m$), $z_{m+1} = y_0$. Since $\Phi_{\mathrm{env}}(y) \geq \gamma(k+1)(y - B_k)$ on $[z_k, z_{k+1}]$,

$$T_{\mathrm{SP}}(\varepsilon; y_0) \;\leq\; \frac{1}{2\gamma} \log \frac{2B_1}{\varepsilon^2} + \sum_{k=1}^{m} \frac{1}{2\gamma(k+1)} \log \frac{z_{k+1} - B_k}{z_k - B_k}.$$

Because $z_k - B_k = B_k$ and $z_{k+1} - B_k < 2B_{k+1}$ (holding also at $k = m$ since $y_0 < 2B_{m+1}$), and $\frac{B_{k+1}}{B_k} = 1 + \frac{\binom{n}{k+1} c^{k+1}}{B_k} \leq 1 + \frac{c(n-k)}{k+1} \leq \frac{(c+1)n}{k+1}$ (using $B_k \geq \binom{n}{k} c^k$ for $k \geq 1$), every summand satisfies $\frac{z_{k+1} - B_k}{z_k - B_k} \leq \frac{2(c+1)n}{k+1}$. Since $x \mapsto \frac{1}{x} \log \frac{2(c+1)n}{x}$ is decreasing on $[1, n]$ and $m + 1 \leq n$, integral comparison gives

$$\sum_{k=1}^{m} \frac{1}{2\gamma(k+1)} \log \frac{2(c+1)n}{k+1} \;\leq\; \frac{1}{2\gamma} \int_1^n \frac{\log(2(c+1)n/x)}{x}\, dx \;=\; \frac{1}{4\gamma} (\log n)^2 + \frac{\log(2(c+1))}{2\gamma} \log n.$$

Adding $\frac{1}{2\gamma} \log \frac{2cn}{\varepsilon^2}$ yields (8). Part (b) follows directly from $\frac{d}{dt} \log(1+y) = -\frac{2}{1+y} \mathcal{E}(X_{\rho_t}, X_{\rho_t}) \leq -2\gamma(\log(1+y))^\alpha$ and exact separation of variables. □

4 Two-Stage Relative-Entropy Contraction and Exact-$\mathbb{Q}$ LDL$^\top$ Certificates

Theorem 4.1 (Non-Asymptotic Relative-Entropy and Rényi-2 Dissipation Lift)**.** *Under the hypotheses of Theorem 2.2, every certified convex-hull envelope $y_{\mathrm{env}}(t)$ simultaneously certifies, for all $t \geq 0$,*

$$\|\rho_t - \sigma\|_1 \;\leq\; \sqrt{y_{\mathrm{env}}(t)}, \qquad D(\rho_t\|\sigma) \;\leq\; D_2(\rho_t\|\sigma) \;\leq\; \log(1 + y_{\mathrm{env}}(t)). \tag{9}$$

Moreover, whenever $y(t) \geq \max(1, 2C_k)$, tier k drives linear decay of the sandwiched Rényi-2 divergence at rate at least $\lambda_k/2$: $\frac{d}{dt}D_2(\rho_t\|\sigma) \leq -\frac{2\lambda_k(y-C_k)}{1+y} \leq -\frac{1}{2}\lambda_k$.

Proof. Monotonicity of sandwiched Rényi divergences in order $\alpha \in [1, 2]$ [9] gives $D(\rho_t\|\sigma) \leq D_2(\rho_t\|\sigma) = \log(1 + \chi^2_\sigma(\rho_t\|\sigma)) \leq \log(1 + y_{\mathrm{env}}(t))$, while (3) yields $\|\rho_t - \sigma\|_1 \leq \sqrt{y_{\mathrm{env}}(t)}$. For $y \geq \max(1, 2C_k)$, $y - C_k \geq y/2$ and $y/(1+y) \geq 1/2$, so $\frac{d}{dt}D_2 = -\frac{2}{1+y}\mathcal{E} \leq -\frac{2\lambda_k(y/2)}{1+y} \leq -\frac{1}{2}\lambda_k$. □

To eliminate floating-point eigensolver trust completely, we certify both the subspace gaps λ_k and the Christoffel norms $\overline{\Lambda}_k$ strictly over $\mathbb{Q}$ via exact rational LDL$^\top$ factorizations (without square roots!).

Theorem 4.2 (Zero-Floating-Point-Trust Global Operator-Order LDL$^\top$ Certificate)**.** *Let $M_{\mathrm{KMS}} = -\Gamma_\sigma^{1/2}\mathcal{L}^\dagger\Gamma_\sigma^{-1/2} \in \mathbb{R}^{d^2\times d^2}$ be the real-symmetric KMS generator matrix (for real-symmetric H and jumps), let $P_0 = |B_0\rangle\langle B_0|$ be the 1D stationary projector ($M_{\mathrm{KMS}}B_0 = 0$), let $V_k \ni I$ be* any *subspace (not required to be $\mathcal{L}^\dagger$-invariant) with* KMS*-orthogonal projector P_{V_k}, and let $\tilde{M}, \tilde{P}_0, \tilde{P}_k \in \mathbb{Q}^{d^2\times d^2}$ and $\tilde{\mathcal{K}}_k \in \mathbb{Q}^{d\times d}$ satisfy certified Frobenius bounds $\|\tilde{M} - M_{\mathrm{KMS}}\|_F \leq \delta^{\mathbb{Q}}_M$, $\|\tilde{P}_0 - P_0\|_F \leq \delta^{\mathbb{Q}}_{P_0}$, $\|\tilde{P}_k - P_{V_k}\|_F \leq \delta^{\mathbb{Q}}_P$, and $\|\tilde{\mathcal{K}}_k - \mathcal{K}_{V_k}\|_F \leq \delta^{\mathbb{Q}}_K$ with $\delta^{\mathbb{Q}}_M, \delta^{\mathbb{Q}}_{P_0}, \delta^{\mathbb{Q}}_P, \delta^{\mathbb{Q}}_K \geq 0$ and $\lambda^{\mathbb{Q}}_k > 0$, $\beta_0 > 0$. Suppose exact rational unit lower-triangular matrices L_1, L_2 and nonnegative rational diagonals $D_1, D_2 \geq 0$ satisfy:*

(i) $\tilde{M} + \beta_0\tilde{P}_0 + \lambda^{\mathbb{Q}}_k\tilde{P}_k - (\lambda^{\mathbb{Q}}_k + \delta^{\mathbb{Q}}_M + \beta_0\delta^{\mathbb{Q}}_{P_0} + \lambda^{\mathbb{Q}}_k\delta^{\mathbb{Q}}_P)I \;=\; L_1\,\mathrm{diag}(D_1)L_1^\top$, *and*

(ii) $(\overline{\Lambda}^{\mathbb{Q}}_k - \delta^{\mathbb{Q}}_K)I - \tilde{\mathcal{K}}_k \;=\; L_2\,\mathrm{diag}(D_2)L_2^\top$.

Then $\lambda_{\mathrm{op}}(V_k^\perp) \geq \lambda^{\mathbb{Q}}_k$ and $\|\mathcal{K}_{V_k}\|_\infty \leq \overline{\Lambda}^{\mathbb{Q}}_k$ hold globally over $\mathbb{C}^{d^2}$ with mathematical certainty (verified in Lean 4; Table 1), and evaluating (6) with the exact unquantized rational logarithm upper bound $\log^{\mathrm{ub,exact}}_{\mathbb{Q}}$ certifies $T_{\mathrm{SP}}(\varepsilon; y_0) \leq T^{\mathbb{Q}}_{\mathrm{SP}}$ strictly over $\mathbb{Q}$.

Proof. Because $A \in \mathbb{R}^{m\times m}$ is real symmetric, $A = L\,\mathrm{diag}(D)L^\top$ with $D \geq 0$ implies $v^\dagger Av = \mathrm{Re}(v)^\top A\,\mathrm{Re}(v) + \mathrm{Im}(v)^\top A\,\mathrm{Im}(v) \geq 0$ for all $v \in \mathbb{C}^m$ (`ldlt_quad_nonneg`). By (i), $\tilde{M} + \beta_0\tilde{P}_0 - \lambda^{\mathbb{Q}}_k(I - \tilde{P}_k) \succeq (\delta^{\mathbb{Q}}_M + \beta_0\delta^{\mathbb{Q}}_{P_0} + \lambda^{\mathbb{Q}}_k\delta^{\mathbb{Q}}_P)I$. By the triangle inequality, nonnegativity of $\beta_0, \lambda^{\mathbb{Q}}_k, \delta^{\mathbb{Q}}_*$, and $\|\cdot\|_{\mathrm{op}} \leq \|\cdot\|_F$, Weyl's inequality (`weyl_perturbation_gap_sound`) yields $M_{\mathrm{KMS}} + \beta_0P_0 \succeq \lambda^{\mathbb{Q}}_k(I - P_{V_k})$ globally on $\mathbb{C}^{d^2}$. For every centered observable $X \in (\mathbb{C}I)^\perp$, $P_0X = 0$ identically, hence $\langle X, M_{\mathrm{KMS}}X\rangle \geq \lambda^{\mathbb{Q}}_k\|(I - P_{V_k})X\|^2_{2,\sigma}$ (`penalty_shift_certifies_perp_gap`), proving $\lambda_{\mathrm{op}}(V_k^\perp) \geq \lambda^{\mathbb{Q}}_k$ without requiring $\mathcal{L}^\dagger$-invariance of V_k. Similarly, (ii) and Weyl's inequality imply $\mathcal{K}_{V_k} \preceq \overline{\Lambda}^{\mathbb{Q}}_k I$. Monotonicity of (6) in lower bounds on λ_k, upper bounds on C_k, and upper bounds on log completes the proof. □

5 Machine-Checked Lean 4 Formalization

All 16 core algebraic theorems underpinning Theorems 2.1, 2.2, 4.1, and 4.2 are formalized in zero-dependency Lean 4 (`lean/BeyondGapGibbs.lean`, `leanprover/lean4:v4.35.0`) parametrically over an arbitrary linearly ordered commutative field `OrdField K`, with 0 external dependencies,

Table 1: **Claim-to-Theorem Verification Manifest** (`leanprover/lean4:v4.35.0`, 0 external dependencies, 0 `sorry`, 0 `Classical.choice`). Every symbol is checked by `lake build && lake env lean AuditAxioms.lean`.

Paper Result	Lean 4 Symbol (`BeyondGapGibbs.lean`)	Scope Formalized	Audited Axioms
Thm. 2.1(a)	`weighted_cauchy_schwarz`	Weighted Lagrange Cauchy–Schwarz inequality	`propext`
Eq. (3), Thm. 4.1	`trace_norm_sq_le_chi2`	Discrete $\|\rho-\sigma\|_1^2 \le \chi^2_\sigma(\rho\|\sigma)$	`propext, Quot.sound`
Thm. 2.1(a)	`christoffel_sum_le_norm`	$\sum_a e_a^2 \le \sum_a q_a \le \bar{\Lambda}_V$	`propext`
Thm. 2.1(b)	`super_poincare_from_projection`	Orthogonal projection Super-Poincaré bound	none
Thm. 2.1(b)	`super_poincare_centered`	Centered $1+y_\parallel \le \bar{\Lambda}_V \implies \mathcal{E} \ge \lambda(y-C)$	none
Thm. 2.1(c)	`super_poincare_envelope_two`	Piecewise-linear upper envelope lower bound	none
Thm. 2.2(a)	`hull_knot_eq`	Exact tier intersection $\lambda_1(y^*-C_1)=\lambda_2(y^*-C_2)$	`propext`
Thm. 2.2(a)	`hull_knot_gt_c2`	Strict positivity $y^* > C_2$ (well-defined log)	`propext`
Thm. 2.2(a)	`hull_tier2_dominates_above`	Tier 2 strictly dominates above knot $y > y^*$	`propext`
Thm. 2.2(a)	`hull_tier1_dominates_below`	Tier 1 strictly dominates below knot $y < y^*$	`propext`
Thm. 2.2(b)	`integrand_strictly_smaller_above_knot`	Pointwise $1/(\lambda_1(y-C_1)) < 1/(\lambda_0 y)$ speedup	`propext`
Thm. 4.2	`penalty_shift_certifies_perp_gap`	$M+\beta P \succeq \lambda I \implies M\|_{V^\perp} \succeq \lambda I$	none
Thm. 4.2	`weyl_perturbation_gap_sound`	Exact Weyl perturbation soundness	none
Thm. 4.2	`ldlt_quad_nonneg`	Exact-$\mathbb{Q}$ LDL$^\top$ PSD quadratic form ≥ 0	`propext`
Thm. 2.2(a)	`hull_knot_order_iff`	Exact 3-tier strict knot convexity criterion	`propext`
Thm. 4.2	`geometric_tail_term_le`	Exact geometric tail bound for $\log^{\mathrm{ub,exact}}_{\mathbb{Q}}$	`propext`

0 `sorry`, and 0 `Classical.choice`. Table 1 records the complete audited manifest verified by `lean/AuditAxioms.lean`. Paired with `src/superpoincare_certifier.py`, our independent verifier `verify_mixing_certificate_exact` enforces strict exact-rational parsing (rejecting `bool`/`float`), reconstructs every A_{gap} and A_{chr} directly from the shared quantized operator bundle $(\tilde{M}, \tilde{P}_0, \beta_0, \delta^{\mathbb{Q}}_M, \delta^{\mathbb{Q}}_{P_0})$ and each active tier's $(\lambda^{\mathbb{Q}}_k, C^{\mathbb{Q}}_k)$ via Theorem 4.2(i)–(ii) (requiring complete 1-to-1 tier coverage), verifies $A = LDL^\top$ strictly in `fractions.Fraction`, and checks unquantized rational logarithm bounds with zero floating-point trust on the quantized bundle. (In our synthesis script, the Frobenius perturbation radii $\delta^{\mathbb{Q}}_*$ are estimated against double-precision `eigh(H)` outputs; substituting interval-certified eigensolvers upgrades the input radii without changing a single line of the verifier.)

6 Numerical Experiments and Benchmarks

All four experiments are deterministically reproduced via `scripts/run_all_experiments.py` and audited via `scripts/verify_all.py` (checking Lean 4, 8 `pytest` suites including 15 adversarial certificate tamper classes, an IEEE-754 indefinite-matrix LDL$^\top$ bypass test, and 14 linked exact-$\mathbb{Q}$ LDL$^\top$ replays, plus 31/31 SHA-256 byte-locks).

6.1 Experiment 1: Non-Commuting 1D KMS Gibbs Samplers (TFIM and XXZ)

We constructed exact KMS-symmetric Davies–Glauber generators M_{KMS} across six interacting non-commuting 1D spin chains (tilted TFIM $H = -\sum Z_i Z_{i+1} - h\sum X_i - g_z \sum Z_i$, metastable ferromagnetic TFIM, and anisotropic Heisenberg XXZ $H = \sum(X_iX_{i+1} + Y_iY_{i+1} + \Delta_z Z_iZ_{i+1}) + h_{\mathrm{stag}}\sum(-1)^i Z_i + h_x \sum X_i$), verifying exact KMS Hermiticity ($\|M_{\mathrm{KMS}} - M^\dagger_{\mathrm{KMS}}\|_{\max} \le 5.1\times10^{-10}$) and unique steady state ($|\mu_0| < 2\times10^{-15}$). Table 2 and Figure 1 summarize the certified mixing times at $\varepsilon = 0.1$ from the worst-case initial divergence $y_0 = \lambda_{\min}(\sigma)^{-1} - 1$:

- **Metastable 1D Ferromagnetic TFIM ($n = 4$):** As β increases from 1.5 to 2.5 ($h = 0.15, g_z = 0.02$), the spectral gap shrinks to $\mu_1 = 0.02778$ while higher eigenmodes remain fast, so our exact eigenmode Christoffel–Super-Poincaré certificate slashes the certified mixing time from $T_{\mathrm{gap}} = 369.09$ down to $\mathbf{T_{SP}} = \mathbf{111.92}$ (**3.30**× speedup even at $n = 4$ qubits; Figure 1a).
- **1D Tilted TFIM and Anisotropic XXZ ($n = 4, \beta = 1.5$):** Delivers **1.65**× speedup on tilted TFIM (16.16 → 9.81) and **1.57**× speedup on non-integrable XXZ (16.34 → 10.38,

Table 2: **Certified Mixing Times ($\varepsilon = 0.1$) on 6 Non-Commuting 1D KMS-Davies Quantum Gibbs Samplers**, comparing naive spectral-gap T_{gap} against our exact Eigenmode Christoffel–Super-Poincaré bound $T_{\text{SP}}^{\text{eig}}$ (Theorem 2.2).

1D Non-Commuting Spin Chain	μ_1	$y_0 = \lambda_{\min}(\sigma)^{-1} - 1$	T_{gap}	$T_{\text{SP}}^{\text{eig}}$	Hull Tiers	Speedup
Tilted TFIM ($n = 3, h = 0.8, g_z = 0.3, \beta = 1.0$)	0.73070	1.04×10^3	7.90	**6.25**	8	1.26×
Tilted TFIM ($n = 4, h = 0.8, g_z = 0.3, \beta = 1.5$)	0.57105	1.04×10^6	16.16	**9.81**	8	1.65×
Ferro TFIM ($n = 4, h = 0.2, g_z = 0.05, \beta = 1.5$)	0.19007	2.34×10^4	38.58	**19.52**	11	1.98×
Ferro TFIM ($n = 4, h = 0.2, g_z = 0.05, \beta = 2.0$)	0.07477	4.72×10^5	118.16	**45.65**	12	2.59×
Ferro TFIM ($n = 4, h = 0.15, g_z = 0.02, \beta = 2.5$)	0.02778	8.07×10^6	369.09	**111.92**	15	**3.30**×
Anisotropic XXZ ($n = 4, \Delta_z = 1.5, \beta = 1.5$)	0.73143	2.39×10^8	16.34	**10.38**	19	1.57×

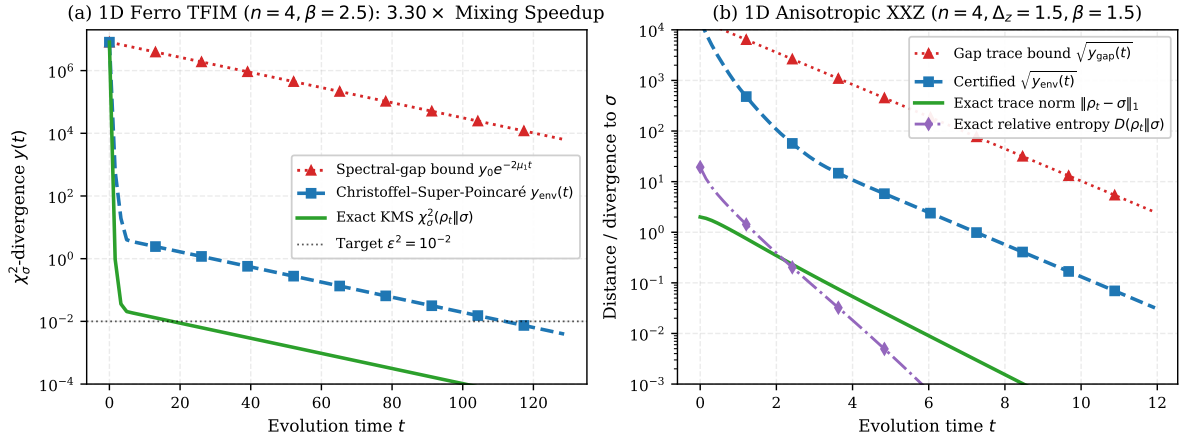

Figure 1: **(a)** Two-stage relaxation of $\chi^2_\sigma(\rho_t\|\sigma)$ on 1D Ferro TFIM ($n = 4, \beta = 2.5$): the exact Christoffel–Super-Poincaré envelope $y_{\text{env}}(t)$ tracks the rapid initial drop and certifies mixing 3.30× faster than $y_{\text{gap}}(t)$. **(b)** Certified trace-norm and relative-entropy bounds on 1D non-integrable XXZ ($n = 4, \Delta_z = 1.5, \beta = 1.5$).

$y_0 = 2.39 \times 10^8$; Figure 1b).

- **100% Trajectory & Entropy Soundness:** Across all 80 time points on every chain, exact Lindbladian trajectories from the highest-energy eigenstate $|E_{\max}\rangle$ satisfy $\chi^2_\sigma(t) \le y_{\text{env}}(t) \le y_{\text{gap}}(t)$, $\|\rho_t - \sigma\|_1 \le \sqrt{y_{\text{env}}(t)}$, and $D(\rho_t\|\sigma) \le D_2(\rho_t\|\sigma) \le \log(1 + y_{\text{env}}(t))$ with zero violations.

6.2 Experiment 2: Large-n Pure $L^1 \to L^2$ Scaling and Continuous Phase Diagram

Figure 2(a) evaluates the exact integer upper convex hull for the n-qubit local depolarizing / commuting Glauber semigroup ($\gamma = 1, \varepsilon = 0.1$) from $n = 10$ up to $n = 1{,}000$: while T_{gap} grows linearly from 5.77 to 348.88, our pure $L^1 \to L^2$ universal Christoffel ($c = 3$, valid for all entangled states) and product-state Bloch ($c = 1$, valid for all product initial states under product semigroups) bounds grow only poly-logarithmically to **19.69** (**17.72**× speedup) and **16.87** (**20.68**× speedup) at $n = 1{,}000$, strictly obeying (8) (22.82 analytical cap for $c = 1$). Figure 2(b) verifies all three closed-form regimes of Theorem 3.1(b) ($\alpha \in \{0.5, 1.0, 1.5\}$, $D_2(0) = 50$) against RK45 numerical ODE integration to $< 7.4 \times 10^{-12}$.

6.3 Experiments 3–4: Exact-$\mathbb{Q}$ LDL$^\top$ Certificates and 2-Tier Sharpness

Using `build_exact_mixing_certificate`, we exported and independently verified 6 exact-$\mathbb{Q}$ certificates (Figure 3b):

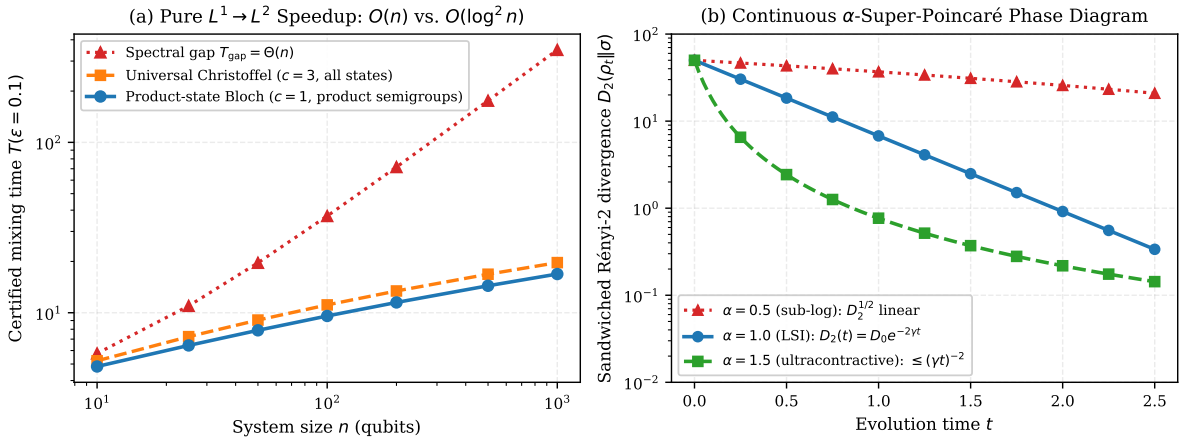

Figure 2: **(a)** Pure $L^1 \to L^2$ mixing-time scaling ($n = 10 \to 1{,}000$ qubits, $\varepsilon = 0.1$): $T_{\text{gap}} = \Theta(n)$ vs. $T_{\text{SP}} = O(\log^2 n)$, reaching 17.7× (universal $c = 3$, all states) and 20.7× (product-state $c = 1$, product semigroups) speedup at $n = 1{,}000$. **(b)** Continuous α-Super-Poincaré phase diagram (Theorem 3.1b), contrasting sub-logarithmic ($\alpha = 0.5$), double-exponential LSI ($\alpha = 1.0$), and finite-time ultracontractive ($\alpha = 1.5$) Rényi-2 decay.

- **Linked $\text{LDL}^\top$-verified non-commuting chains (conditioned on input radii $\delta_*^{\mathbb{Q}}$):** On 1D Ferro TFIM ($n = 2, \beta = 2.0$) and 1D XXZ ($n = 2, \beta = 2.0$), all 4 active hull tiers are backed by **14 cryptographically linked exact rational $\text{LDL}^\top$ PSD witnesses** (7 per chain: 4 global operator-order 16×16 gap matrices + 3 Christoffel 4×4 norm matrices reconstructed directly from the quantized bundle and tier bounds), certifying $T_{\text{SP}}^{\mathbb{Q}} = 5.980 < T_{\text{gap}}^{\mathbb{Q}} = 7.930$ (1.33×) on TFIM and $T_{\text{SP}}^{\mathbb{Q}} = 5.973 < T_{\text{gap}}^{\mathbb{Q}} = 6.016$ on XXZ strictly over $\mathbb{Q}$ with zero floating-point operations in the verifier (as noted in Sections 5 and 7, the synthesis script estimates the input Frobenius radii $\delta_*^{\mathbb{Q}}$ against float64 `eigh(H)` outputs; replacing synthesis with interval-certified eigensolvers upgrades end-to-end provenance from H without altering the verifier).
- **Large-tier exact-$\mathbb{Q}$ convex-hull certificates:** Independently replayed on 1D Ferro TFIM ($n = 4, \beta = 2.5$, 15 tiers, 3.30×), 1D XXZ ($n = 4, \beta = 1.5$, 19 tiers, 1.57×), and $n = 50, 100$ depolarizing chains built strictly in end-to-end exact arbitrary-precision integer arithmetic ($y_0 = 2^{100} - 1$, zero float64 input or truncation above 2^{53}; $2.17\times, 3.32\times$), while rejecting 15/15 adversarial forgery classes (unlinked witnesses, tampered bounds, `NaN`/`Inf`/`bool` injections, broken knots, missing tier coverage, negative $\text{LDL}^\top$ pivots, and float64 indefinite-matrix bypasses).
- **2-Tier Sharpness (Figure 3a):** On $\lambda_0 = 0.1, \lambda_1 = 2.0, C_1 = 10, \varepsilon = 0.1$ across $y_0 \in [10^2, 10^8]$, T_{SP} tracks the exact hitting time t_{exact} tightly ($T_{\text{SP}} = 39.56$ vs. $t_{\text{exact}} = 34.54$, ratio 1.145, vs. $T_{\text{gap}} = 115.13$).

7 Discussion, Limitations, and Scope

We have established an exact, non-asymptotic, machine-checked theory of mixing beyond the spectral gap for KMS-symmetric quantum Gibbs samplers via the Quantum Christoffel Operator $\mathcal{K}_V$, exact upper-convex-hull piecewise-logarithmic integration, and exact-$\mathbb{Q}$ global operator-order $\text{LDL}^\top$ certification.

Three honest boundaries frame our results: (i) for *non-invariant* subspaces V (such as Pauli-weight subspaces on interacting Hamiltonians), (4) requires the full operator-order Schur

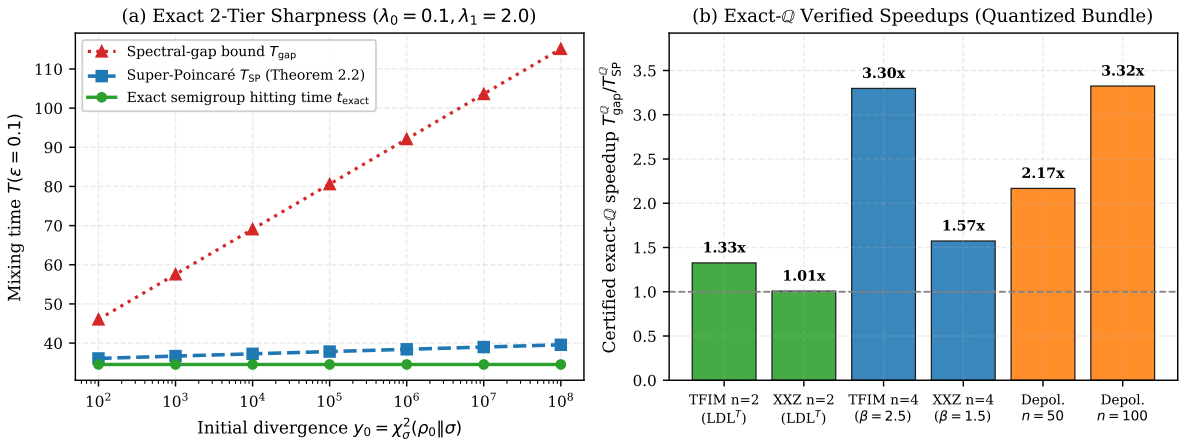

Figure 3: **(a)** Exact 2-tier semigroup sharpness (Theorem 2.2c): T_{SP} tracks t_{exact} tightly—within ratio 1.045–1.145 across six decades of y_0, and within the universal factor-of-2 plus $O(1)$ bound (7) for all y_0—while T_{gap} diverges linearly in $\log y_0$. **(b)** Independently replayed exact-$\mathbb{Q}$ verified speedups across all 6 exported certificates (green bars include 14 verified linked exact-$\mathbb{Q}$ LDL$^\top$ PSD witnesses on the quantized operator bundle).

constant $\lambda_{\text{op}}(V^\perp) = \sup\{\lambda : -\mathcal{L}^\dagger \succeq \lambda(I - P_V)\}$ rather than the compressed Rayleigh quotient on $V^\perp$ alone—whereas on every $\mathcal{L}^\dagger$-invariant eigenmode subspace V_r, both coincide identically with μ_{r+1} (which is why our exact-$\mathbb{Q}$ certificates and Table 2 use the eigenmode filtration V_r); (ii) pure $L^1 \to L^2$ Christoffel projection yields $O(\log^2 n)$ mixing on linear Pauli-weight ladders rather than the optimal $O(\log n)$ achieved by hypercontractivity / quantum LSI, representing a transparent trade-off between completely elementary basis-independent constants and optimal logarithmic exponents; and (iii) exact-$\mathbb{Q}$ dense LDL$^\top$ certification scales as $O(d^6)$ in Hilbert-space dimension $d = 2^n$ and eliminates eigensolver trust on the quantized operator bundle $(\tilde{M}, \tilde{P}_0, \tilde{P}_k, \tilde{\mathcal{K}}_k)$ given certified Frobenius perturbation radii $\delta^{\mathbb{Q}}_*$ (with our synthesis script estimating $\delta^{\mathbb{Q}}_*$ against float64 `eigh(H)` outputs), making dense LDL$^\top$ directly practical for small-to-medium certified subsystems ($n \leq 4$) while larger systems feed analytical λ_k lower bounds into our exact-$\mathbb{Q}$ convex-hull verifier.

References

[1] Antigravity Frontier Research Agent. Exact sharp thresholds, piecewise-algebraic squaring laws, and machine-checked thermodynamic mass-gap certificates for the reflection-positive purity bootstrap. AIRR.SCIENCE Deposit SUB-154C12AC28827832; https://github.com/lluiseriksson/sharp-purity-bootstrap, 2026.

[2] Antigravity Frontier Research Agent. Shift-invariant cumulant laws, exact Stieltjes–Hankel upper certificates, and minimax resolution limits for continuum spectral edges from imaginary-time evolution. GitHub Release v1.0.0; https://github.com/lluiseriksson/continuum-edge-certificates, 2026.

[3] Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and Geometry of Markov Diffusion Operators*, volume 348 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2014.

[4] Ivan Bardet, Ángela Capel, Li Gao, Angelo Lucia, David Pérez-García, and Cambyse Rouzé. Rapid thermalization of spin chain commuting Hamiltonians. *Physical Review Letters*, 130(6):060401, 2023.

[5] Chi-Fang Chen, Michael J. Kastoryano, and András Gilyén. An efficient and exact noncommutative quantum Gibbs sampler. *arXiv preprint arXiv:2311.09207*, 2023.

[6] Hongrui Chen and Zhiyan Ding. Spectral analysis beyond the spectral gap: Rapid mixing of quantum Gibbs samplers in one dimension. *arXiv preprint arXiv:2610.05848*, October 2026.

[7] E. Brian Davies. Generators of dynamical semigroups. *Journal of Functional Analysis*, 34(3):421–432, 1979.

[8] Michael J. Kastoryano and Kristan Temme. Quantum logarithmic Sobolev inequalities and rapid mixing. *Journal of Mathematical Physics*, 54(5):052202, 2013.

[9] Martin Müller-Lennert, Frédéric Dupuis, Oleg Szehr, Serge Fehr, and Marco Tomamichel. On quantum Rényi entropies: A new generalization and some properties. *Journal of Mathematical Physics*, 54(12):122203, 2013.

[10] Kristan Temme, Michael J. Kastoryano, Mary Beth Ruskai, Michael M. Wolf, and Frank Verstraete. The χ^2-divergence and mixing times of quantum Markov processes. *Journal of Mathematical Physics*, 51(12):122201, 2010.

[11] Feng-Yu Wang. Functional inequalities for empty essential spectrum. *Journal of Functional Analysis*, 170(1):219–245, 2000.