Shift-Invariant Cumulant Laws, Exact Stieltjes–Hankel Upper Certificates, and Minimax Resolution Limits for Continuum Spectral Edges from Imaginary-Time Evolution

Lluis Eriksson*

November 24, 2026

Abstract

Extracting the spectral gap $\Delta = E_{\mathrm{th}} - E_0$ and continuum threshold exponent $\alpha > 0$ of an infinite-dimensional quantum Hamiltonian from imaginary-time evolution (ITE) was recently placed on a rigorous asymptotic footing by Karapetyan, Bondar, Mareček, and Wodecki [6]. At finite imaginary time τ, however, the logarithmic decay rate $g(\tau) = -\partial_\tau \log F(\tau) = \Delta + \alpha/\tau + O(\tau^{-2})$ suffers an $O(\alpha/\tau)$ systematic bias that confounds Δ with α, offers no assumption-free finite-τ certificate, and degrades only logarithmically as $O(1/\log(1/\sigma))$ under additive shot noise σ. We resolve all three bottlenecks with exact cumulant identities, zero-dependency Lean 4 formalizations, and zero-floating-point-trust exact-$\mathbb{Q}$ moment certificates. First, we prove an exact covariance–boundary identity for the variance-corrected estimator $\widehat{\Delta}_{\mathrm{VF}}(\tau) \coloneqq g(\tau) - \tau\kappa_2(\tau)$: for any continuum density $w(x) = x^{\alpha-1}e^{-V(x)}$ on $(0, B)$, $\widehat{\Delta}_{\mathrm{VF}}(\tau) - \Delta = \tau^{-2}\,\mathrm{Cov}_{U_\tau}(U_\tau, U_\tau V'(U_\tau/\tau))$ plus a nonnegative upper-edge boundary term when $B < \infty$, yielding non-asymptotic one-sided bounds at every finite $\tau > 0$ via Chebyshev's association inequality. Second, we introduce the *shift-invariant* cumulant estimators $\widehat{\Delta}_{\mathrm{SI}}(\tau) \coloneqq g(\tau) - 2\kappa_2(\tau)^2/\kappa_3(\tau)$ and $\widehat{\alpha}_{\mathrm{SI}}(\tau) \coloneqq 4\kappa_2(\tau)^3/\kappa_3(\tau)^2$: both are identically exact on the three-parameter Gamma–Exponential family $x^{\alpha-1}e^{-\lambda x}$, and on general smooth continuum edges an exact $O(\tau^{-2})$ cancellation between $2\kappa_2^2$ and κ_3 delivers two-order bias reduction—$\widehat{\Delta}_{\mathrm{SI}}(\tau) - \Delta = c_2\alpha(\alpha+1)/(\tau+\lambda)^3 + O(\tau^{-4})$ and $\widehat{\alpha}_{\mathrm{SI}}(\tau) - \alpha = -3c_2\alpha(\alpha+1)/(\tau+\lambda)^2 + O(\tau^{-3})$—where $c_2 = (\log\phi)''(0)$ is the edge log-curvature. Third, we prove a spectral-edge trichotomy showing that $\tau^2\kappa_2(\tau)$ uniquely classifies discrete excited states ($\to 0$ exponentially), power-law continua ($\to \alpha$), and stretched-exponential pseudogaps ($\sim \tau^{\gamma/(\gamma+1)} \to +\infty$). Fourth, without any smoothness or monotonicity assumption, we derive interval-sound Stieltjes–Hankel upper certificates witnessed by rational vectors $v \in \mathbb{Q}^{n+1}$ and verified strictly over $\mathbb{Q}$ (including exact rational logarithm upper bounds). Fifth, we establish poly-logarithmic additive-noise minimax resolution limits in the asymptotic window via an explicit closed-form Le Cam two-point construction. All 16 core algebraic theorems are formalized in zero-dependency Lean 4 (`v4.35.0`, 0 `sorry`, 0 `Classical.choice`), and validated on 1D/3D one-center point interactions (267×–397× gap accuracy gain at $\tau = 20$), $L = 401$ tight-binding impurity exact diagonalization (169× gain), 1D van Hove bands (7,630× gain), 6 exact-$\mathbb{Q}$ Hankel certificates (2.35× tighter than raw $g(\tau_{\max})$ with 200/200 noisy interval soundness), and additive shot-noise scaling.

1 Introduction

In infinite-dimensional quantum mechanics and thermodynamic-limit many-body systems, the ground-state energy E_0 is typically separated by a spectral gap $\Delta = E_{\mathrm{th}} - E_0 > 0$ from the

*Independent researcher (Sweden). Research, zero-dependency Lean 4 proofs, exact-$\mathbb{Q}$ Stieltjes–Hankel certificate verifiers, and numerical benchmarks prepared with AI assistance (Google Antigravity) and publicly archived at https://github.com/lluiseriksson/continuum-edge-certificates.

bottom edge E_{th} of a continuous scattering or quasiparticle spectrum [9, 2]. Imaginary-time evolution (ITE)—implemented via quantum imaginary-time evolution (QITE) [8], quantum Monte Carlo, or tensor networks—projects an initial state ψ_0 onto the ground state, yet extracting excitations above E_0 is far subtler when the first excitation is a continuum edge rather than an isolated eigenvalue.

In a landmark recent paper, Karapetyan, Bondar, Mareček, and Wodecki [6] developed a rigorous mathematical framework showing that the excited imaginary-time decay function

$$F(\tau) \;=\; \int_{\Delta}^{\infty} e^{-\tau\omega}\, d\nu(\omega) \;=\; e^{-\tau\Delta} \int_{0}^{\infty} e^{-\tau x}\, w(x)\, dx \tag{1}$$

(arising from the ground-state-subtracted survival amplitude or connected observable correlation along ITE) detects the continuum edge Δ whenever the spectral density obeys a power law $w(x) \sim c\, x^{\alpha-1}$ ($\alpha > 0$) near $x = 0^+$, producing a logarithmic correction $\log F(\tau) = -\tau\Delta - \alpha \log \tau + O(1)$ whose coefficient encodes the threshold exponent α.

Three fundamental challenges remain open for finite-τ, finite-sample, and certified applications:

1. **Finite-τ Bias and (Δ, α) Confounding.** The raw logarithmic decay rate $g(\tau) := -\frac{d}{d\tau} \log F(\tau) = \Delta + \frac{\alpha}{\tau} + O(\tau^{-2})$ has an $O(1/\tau)$ bias that decays painfully slowly and entangles Δ with the unknown exponent α (e.g., at $\tau = 20$ on the 1D point interaction with $\Delta = 1, \alpha = 3/2$, $g(20)$ still errs by $+7.2\%$).
2. **Non-Asymptotic Certificates Without Distributional Assumptions.** Asymptotic Watson-lemma expansions [5] do not certify strict bounds on Δ at finite τ, nor do they distinguish power-law continuum edges from isolated discrete levels or pseudogaps without prior spectral knowledge.
3. **Exponential Signal Decay Under Additive Shot Noise.** Because $F(\tau) \sim c\,\Gamma(\alpha)\tau^{-\alpha}e^{-\tau\Delta}$ vanishes exponentially in τ, any additive measurement noise σ imposes a hard horizon $\tau_{\max} = O(\Delta^{-1} \log(1/\sigma))$, trapping $O(\tau^{-1})$ estimators at a sluggish $O(1/\log(1/\sigma))$ error floor.

Summary of Contributions. We resolve all three challenges and build a complete machine-checked verification pipeline (complementing our recent finite-sample QC-QMC [4] and reflection-positive purity-bootstrap [3] certificates):

- **Cumulant–Covariance Identity & Chebyshev Sign Law (Theorem 2.1):** Expresses $\widehat{\Delta}_{\text{VF}}(\tau) - \Delta$ as an exact scaled covariance, yielding non-asymptotic upper or lower bounds at *every* finite $\tau > 0$ determined solely by the monotonicity of $xV'(x)$.
- **Shift-Invariant Cumulant Estimators & Exact $O(\tau^{-2})$ Cancellation (Theorem 2.3):** Introduces $\widehat{\Delta}_{\text{SI}}(\tau) = g - 2\kappa_2^2/\kappa_3$ and $\widehat{\alpha}_{\text{SI}}(\tau) = 4\kappa_2^3/\kappa_3^2$, proving exact invariance on the Gamma–Exponential family and exact leading errors $c_2\alpha(\alpha+1)/(\tau+\lambda)^3$ and $-3c_2\alpha(\alpha+1)/(\tau+\lambda)^2$ on general smooth continuum edges.
- **Spectral-Edge Trichotomy (Theorem 3.1):** Proves that $\tau^2\kappa_2(\tau)$ cleanly separates discrete excited levels ($\to 0$), power-law continuum edges ($\to \alpha$), and stretched-exponential pseudogaps ($\sim \tau^{\gamma/(\gamma+1)} \to +\infty$).
- **Zero-Float-Trust Exact-$\mathbb{Q}$ Stieltjes–Hankel Certificates (Theorem 4.1):** Derives interval-sound rational witness vectors $v \in \mathbb{Q}^{n+1}$ certifying $\Delta < \Delta_{\text{ub}}^{\mathbb{Q}}$ strictly over $\mathbb{Q}$ (including exact rational logarithm upper bounds) for *any* positive spectral measure.
- **Additive-Noise Minimax Resolution Limits (Theorem 5.1):** Establishes poly-logarithmic minimax rates via Le Cam's two-point method and documents both the $O((\log(1/\sigma))^{-3})$ speedup and the honest finite-noise bias–variance crossover.

- **Zero-Dependency Lean 4 Formalization & Physical Benchmarks (Sections 6–7):** Formalizes all 16 core algebraic theorems in Table 1 in Lean 4 (`v4.35.0`, 0 `sorry`, 0 `Classical.choice`) and validates across 1D/3D point interactions, $L = 401$ tight-binding impurity ED, 1D/2D band continua, 6 exact-$\mathbb{Q}$ Hankel certificates, and Monte Carlo shot noise.

2 Cumulant–Covariance Identity and Shift-Invariant Estimators

Let $F(\tau) = \int_\Delta^\infty e^{-\tau\omega}\, d\nu(\omega) > 0$ for $\tau > 0$ be the Laplace transform of a nonzero positive excited spectral measure ν supported on $[\Delta, \infty)$ with $\Delta = \inf \mathrm{supp}(\nu)$. Let X_τ denote the exponentially tilted random variable on $[0, \infty)$ with probability law

$$d\mathbb{P}_\tau(x) \;=\; \frac{e^{-\tau x}\, d\nu(\Delta + x)}{\int_0^\infty e^{-\tau y}\, d\nu(\Delta + y)}.$$

The first three cumulants of the shifted energy $\omega = \Delta + X_\tau$ are given directly by logarithmic derivatives of $F(\tau)$:

$$\begin{aligned} g(\tau) &\coloneqq -\frac{d}{d\tau}\log F(\tau) = \Delta + \mathbb{E}[X_\tau], \qquad \kappa_2(\tau) \coloneqq \frac{d^2}{d\tau^2}\log F(\tau) = \mathrm{Var}(X_\tau), \\ \kappa_3(\tau) &\coloneqq -\frac{d^3}{d\tau^3}\log F(\tau) = \mathbb{E}\big[(X_\tau - \mathbb{E}[X_\tau])^3\big]. \end{aligned} \tag{2}$$

2.1 Exact Cumulant–Covariance Identity and Non-Asymptotic Sign Law

Define the *variance-free* (Richardson-corrected) gap estimator and dimensionless cumulant exponent:

$$\widehat{\Delta}_{\mathrm{VF}}(\tau) \;\coloneqq\; g(\tau) - \tau\,\kappa_2(\tau), \qquad \widehat{\alpha}_{\mathrm{VF}}(\tau) \;\coloneqq\; \tau^2\,\kappa_2(\tau). \tag{3}$$

Theorem 2.1 (Exact Cumulant–Covariance Identity and Chebyshev Sign Law)**.** *Suppose $d\nu(\Delta + x) = c\,x^{\alpha-1}e^{-V(x)}\,dx$ on $(0, B)$ ($B \in (0, \infty]$, $\alpha > 0$, $c > 0$) with $\phi(x) \coloneqq e^{-V(x)}$, $V \in C^1((0, B))$, $x(1 + x)|V'(x)| \in L^1(d\mathbb{P}_{\tau'})$ for τ' in a neighborhood of τ (with a locally integrable dominating function justifying Leibniz differentiation), and either (i) $B = \infty$ with $x^\alpha\phi(x)e^{-\tau x} \to 0$ as $x \to \infty$, or (ii) $B < \infty$ with finite upper-edge limit $\phi(B^-) \in [0, \infty)$. Let $U_\tau \coloneqq \tau X_\tau$ and $Z(\tau) \coloneqq \int_0^B x^{\alpha-1}\phi(x)e^{-\tau x}\,dx$. Then for* every *$\tau > 0$:*

$$\widehat{\Delta}_{\mathrm{VF}}(\tau) - \Delta \;=\; \frac{d}{d\tau}(\tau\,\mathbb{E}[X_\tau]) \;=\; \frac{1}{\tau^2}\,\mathrm{Cov}\big(U_\tau,\, U_\tau V'(U_\tau/\tau)\big) \;+\; \mathcal{B}_B(\tau), \tag{4}$$

where the upper-edge boundary term $\mathcal{B}_B(\tau) \coloneqq \frac{B^\alpha\phi(B^-)e^{-\tau B}}{Z(\tau)}(B - \mathbb{E}[X_\tau])$ vanishes identically if $B = \infty$ or $\phi(B^-) = 0$, and satisfies $\mathcal{B}_B(\tau) > 0$ whenever $B < \infty$ and $\phi(B^-) > 0$. Consequently, by Chebyshev's association inequality (formalized in discrete form in Lean 4; Table 1):

(a) If $x \mapsto xV'(x)$ is nondecreasing on $(0, B)$ (which holds for every nonincreasing log-concave ϕ, every compact linear step $\phi(x) = 1 - x/B$, and every power-law/Lorentzian decay $\phi(x) = (x + \kappa^2)^{-\gamma}$ with $\gamma > 0$), then $\widehat{\Delta}_{\mathrm{VF}}(\tau) \geq \Delta$ for all *$\tau > 0$ (for both $B = \infty$ and finite $B < \infty$).*

(b) If $B = \infty$ (or $\phi(B^-) = 0$) and $x \mapsto xV'(x)$ is nonincreasing on $(0, B)$ (e.g., an algebraically growing threshold envelope $\phi(x) = (1 + x/a)^\gamma$ with $\gamma > 0$, where $xV'(x) = -\frac{\gamma x}{x+a}$), then $\widehat{\Delta}_{\mathrm{VF}}(\tau) \leq \Delta$ for all *$\tau > 0$.*

(c) On the Gamma–Exponential family $V(x) = \lambda x$ ($B = \infty$), (4) evaluates explicitly to $\widehat{\Delta}_{\mathrm{VF}}(\tau) - \Delta = \frac{\alpha\lambda}{(\tau+\lambda)^2}$ (Lean: `gamma_exp_vf_excess`).

Proof. Since $-\frac{d}{d\tau}\mathbb{E}[X_\tau] = \mathrm{Var}(X_\tau) = \kappa_2(\tau)$, the product rule gives $\widehat{\Delta}_{\mathrm{VF}}(\tau) - \Delta = \mathbb{E}[X_\tau] - \tau\kappa_2(\tau) = \frac{d}{d\tau}(\tau\mathbb{E}[X_\tau])$. Changing variables to $u = \tau x$ yields

$$\tau\,\mathbb{E}[X_\tau] \;=\; \frac{N(\tau)}{D(\tau)} \;\coloneqq\; \frac{\int_0^{\tau B} u^\alpha e^{-u}\phi(u/\tau)\,du}{\int_0^{\tau B} u^{\alpha-1}e^{-u}\phi(u/\tau)\,du}, \qquad D(\tau) = \tau^\alpha Z(\tau).$$

Differentiating under the integral sign via Leibniz's rule (using $\partial_\tau\phi(u/\tau) = \tau^{-2}uV'(u/\tau)\phi(u/\tau)$ and upper-limit derivative $\frac{d}{d\tau}(\tau B) = B$ when $B < \infty$) yields

$$\frac{d}{d\tau}(\tau\,\mathbb{E}[X_\tau]) = \frac{1}{\tau^2}\Big(\mathbb{E}\big[U_\tau \cdot U_\tau V'(U_\tau/\tau)\big] - \mathbb{E}[U_\tau]\,\mathbb{E}\big[U_\tau V'(U_\tau/\tau)\big]\Big) + \frac{B(\tau B)^{\alpha-1}e^{-\tau B}\phi(B^-)}{D(\tau)}(\tau B - \tau\mathbb{E}[X_\tau]) = \frac{1}{\tau^2}\,\mathrm{Cov}(U_\tau,\, U_\tau V'(U_\tau/\tau)) + \mathcal{B}_B(\tau).$$

Since $X_\tau < B$ almost surely when $B < \infty$, $\mathcal{B}_B(\tau) \geq 0$ always (and $\mathcal{B}_B = 0$ if $B = \infty$ or $\phi(B^-) = 0$). Parts (a) and (b) then follow from Chebyshev's association inequality: $\mathrm{Cov}(f(U), g(U)) = \frac{1}{2}\mathbb{E}[(f(U)-f(U'))(g(U)-g(U'))] \geq 0$ (resp. ≤ 0) when f, g are comonotone (resp. anti-monotone). For part (c), $V'(x) = \lambda$ and $B = \infty$ give $\tau^{-2}\,\mathrm{Cov}(U_\tau, \lambda U_\tau) = \lambda\,\mathrm{Var}(X_\tau) = \frac{\alpha\lambda}{(\tau+\lambda)^2}$. □

Remark 2.2 (Upper-Edge Singularities and the 1D van Hove Band)**.** For the 1D tight-binding van Hove density $w(x) = [x(B-x)]^{-1/2}$ ($\alpha = 1/2$, $\phi(x) = (B-x)^{-1/2}$), the upper-edge van Hove singularity violates $\phi(B^-) < \infty$ and $x|V'| \in L^1(d\mathbb{P}_\tau)$, yet the Laplace transform evaluates in closed form via the modified Bessel function I_0: $F(\tau) = \pi e^{-\tau(\Delta+B/2)}I_0(\tau B/2)$, giving $\tau\mathbb{E}[X_\tau] = z(1 - I_1(z)/I_0(z))$ with $z \coloneqq \tau B/2$. Differentiating directly with $I_0' = I_1$ and $I_1' = I_0 - I_1/z$ gives the exact law

$$\widehat{\Delta}_{\mathrm{VF}}(\tau) - \Delta \;=\; \frac{B}{2}\left[1 - z + z\,\frac{I_1(z)^2}{I_0(z)^2}\right]_{z=\tau B/2},$$

which changes sign at the unique positive root $z_* \approx 1.7024$ (so $\widehat{\Delta}_{\mathrm{VF}}(\tau) > \Delta$ at tiny $\tau < 2z_*/B$ where the upper band edge dominates, and $\widehat{\Delta}_{\mathrm{VF}}(\tau) < \Delta$ strictly for all $\tau > 2z_*/B$, matching the decreasing slope $xV'(x) = -\frac{x}{2(B-x)}$ throughout the scaling regime $\tau \geq 1$ when $B = 4$, since $2z_*/4 \approx 0.8512$).

2.2 Shift-Invariant Cumulant Estimators and Exact $O(\tau^{-2})$ Cancellation

Although $\widehat{\Delta}_{\mathrm{VF}}(\tau)$ improves the error from $O(\tau^{-1})$ to $O(\tau^{-2})$, Theorem 2.1(c) reveals why a residual $O(\tau^{-2})$ bias persists: any exponential envelope $e^{-\lambda x}$ shifts $\tau \mapsto s \coloneqq \tau + \lambda$ in the Gamma moments, breaking scale homogeneity in τ. Eliminating the unknown shift λ algebraically using the third cumulant $\kappa_3(\tau)$ yields our *shift-invariant* continuum-edge and threshold-exponent estimators:

$$\widehat{\Delta}_{\mathrm{SI}}(\tau) \;\coloneqq\; g(\tau) - \frac{2\,\kappa_2(\tau)^2}{\kappa_3(\tau)}, \qquad \widehat{\alpha}_{\mathrm{SI}}(\tau) \;\coloneqq\; \frac{4\,\kappa_2(\tau)^3}{\kappa_3(\tau)^2}. \tag{5}$$

Theorem 2.3 (Shift-Invariant Exactness and Two-Order Bias Cancellation)**.** *Suppose $d\nu(\Delta + x) = c\,x^{\alpha-1}\phi(x)\,dx$ on $(0, B)$ with $\alpha > 0$, $\phi(0) > 0$, $\int_0^B x^{\alpha-1}\phi(x)e^{-\tau_0 x}\,dx < \infty$ for some $\tau_0 > 0$, and ϕ smooth near 0^+ so that $\log(\phi(x)/\phi(0)) = -\lambda x + \frac{1}{2}c_2x^2 + O(x^3)$ as $x \to 0^+$, where $c_2 = (\log\phi)''(0)$ is the continuum-edge log-curvature. Let $s = \tau + \lambda$. Then:*

(a) ***Exactness on the 3-Parameter Gamma–Exponential Family.*** *Whenever $\phi(x) = \phi(0)e^{-\lambda x}$ ($c_2 = 0$, $B = \infty$), $\widehat{\Delta}_{\mathrm{SI}}(\tau) = \Delta$ and $\widehat{\alpha}_{\mathrm{SI}}(\tau) = \alpha$ hold* identically for every $\tau > 0$, *for all $\alpha > 0$ and $\lambda \geq 0$ (Lean: `gamma_exp_si_gap_exact`, `gamma_exp_si_alpha_exact`).*

(b) ***Exact $O(\tau^{-3})$ and $O(\tau^{-2})$ Leading-Error Laws.*** *As $\tau \to \infty$, the $O(c_2 s^{-2})$ corrections in $2\kappa_2(\tau)^2$ and $\kappa_3(\tau)$ cancel identically (Lean: `si_perturbation_cancellation`, `si_alpha_perturbation_exact`), yielding*

$$\widehat{\Delta}_{\mathrm{SI}}(\tau) - \Delta \;=\; \frac{c_2\,\alpha(\alpha+1)}{(\tau+\lambda)^3} + O(\tau^{-4}), \qquad \widehat{\alpha}_{\mathrm{SI}}(\tau) - \alpha \;=\; -\frac{3\,c_2\,\alpha(\alpha+1)}{(\tau+\lambda)^2} + O(\tau^{-3}). \quad (6)$$

In particular, locally log-convex edges ($c_2 > 0$, such as 1D/3D point interactions and 1D van Hove bands) satisfy $\widehat{\Delta}_{\mathrm{SI}}(\tau) > \Delta$ and $\widehat{\alpha}_{\mathrm{SI}}(\tau) < \alpha$ at large τ, whereas locally log-concave edges ($c_2 < 0$, such as 2D linear steps) satisfy $\widehat{\Delta}_{\mathrm{SI}}(\tau) < \Delta$ and $\widehat{\alpha}_{\mathrm{SI}}(\tau) > \alpha$.

Proof. For part (a), $F(\tau) = c\phi(0)\Gamma(\alpha)s^{-\alpha}e^{-\tau\Delta}$ gives $g = \Delta + \alpha/s$, $\kappa_2 = \alpha/s^2$, $\kappa_3 = 2\alpha/s^3$, so $2\kappa_2^2/\kappa_3 = (2\alpha^2/s^4)/(2\alpha/s^3) = \alpha/s$ and $4\kappa_2^3/\kappa_3^2 = (4\alpha^3/s^6)/(4\alpha^2/s^6) = \alpha$ identically.

For part (b), applying Watson's lemma [5] directly to each of the four shifted Laplace moment integrals $m_r(\tau) = \int_0^B x^{\alpha+r-1}e^{-sx}\big(1 + \frac{1}{2}c_2x^2 + O(x^3)\big)\,dx$ ($r = 0, 1, 2, 3$, with $s = \tau + \lambda$) yields

$$m_r(\tau) \;=\; \frac{\Gamma(\alpha+r)}{s^{\alpha+r}}\left(1 + \frac{c_2\,(\alpha+r)(\alpha+r+1)}{2s^2} + O(s^{-3})\right) \qquad (r = 0, 1, 2, 3).$$

Forming the exact moment ratios $\mu_r = m_r/m_0$ and cumulants $g = \Delta + \mu_1$, $\kappa_2 = \mu_2 - \mu_1^2$, and $\kappa_3 = \mu_3 - 3\mu_1\mu_2 + 2\mu_1^3$ (without term-by-term differentiation of asymptotic remainders) and setting $\varepsilon := c_2(\alpha+1)/s^2$ gives

$$g(\tau) = \Delta + \frac{\alpha}{s}(1+\varepsilon) + O(s^{-4}), \quad \kappa_2(\tau) = \frac{\alpha}{s^2}(1+3\varepsilon) + O(s^{-5}), \quad \kappa_3(\tau) = \frac{2\alpha}{s^3}(1+6\varepsilon) + O(s^{-6}).$$

Crucially, by the exact rational identity proved in Lean 4 (`si_perturbation_cancellation`),

$$\frac{(1+3\varepsilon)^2}{1+6\varepsilon} \;=\; 1 + \frac{9\varepsilon^2}{1+6\varepsilon} \;=\; 1 + O(s^{-4}),$$

so $2\kappa_2(\tau)^2/\kappa_3(\tau) = \frac{\alpha}{s} + O(s^{-4})$ has *zero* $O(c_2 s^{-3})$ correction! Subtracting from $g(\tau) = \Delta + \frac{\alpha}{s} + \frac{\alpha\varepsilon}{s} + O(s^{-4})$ immediately leaves $\widehat{\Delta}_{\mathrm{SI}}(\tau) - \Delta = \alpha\varepsilon/s + O(s^{-4}) = c_2\alpha(\alpha+1)/s^3 + O(s^{-4})$. Similarly, by `si_alpha_perturbation_exact`,

$$\frac{(1+3\varepsilon)^3}{(1+6\varepsilon)^2} \;=\; 1 - 3\varepsilon + \frac{27\varepsilon^2 + 135\varepsilon^3}{(1+6\varepsilon)^2} \;=\; 1 - 3\varepsilon + O(s^{-4}),$$

giving $\widehat{\alpha}_{\mathrm{SI}}(\tau) = \alpha(1-3\varepsilon) + O(s^{-3}) = \alpha - 3c_2\alpha(\alpha+1)/s^2 + O(s^{-3})$. □

3 Spectral-Edge Trichotomy via Cumulant Invariants

Before trusting an asymptotic continuum-edge estimate, one must know whether the first excitation above E_0 is an isolated discrete eigenvalue, a power-law continuum edge, or an exponentially suppressed pseudogap (where [6]'s power-law assumption fails). The dimensionless cumulant $\widehat{\alpha}_{\mathrm{VF}}(\tau) = \tau^2\kappa_2(\tau)$ cleanly classifies all three regimes.

Theorem 3.1 (Spectral-Edge Trichotomy)**.** *Let ν be a positive excited spectral measure with $\Delta = \inf\mathrm{supp}(\nu)$. As $\tau \to \infty$:*

(i) ***Isolated Discrete Excited State:*** *If $\nu = w_1\delta_\Delta + \tilde{\nu}$ with $w_1 > 0$ and $\mathrm{supp}(\tilde{\nu}) \subset [\Delta+\delta_{21}, \infty)$ ($\delta_{21} > 0$), then $\widehat{\alpha}_{\mathrm{VF}}(\tau) = O(\tau^2 e^{-\tau\delta_{21}}) \to 0$ exponentially fast.*

(ii) ***Power-Law Continuum Edge:*** *If $d\nu(\Delta + x) = c\,x^{\alpha-1}\phi(x)\,dx$ ($\alpha > 0$, ϕ smooth with $\phi(0) > 0$), then $\widehat{\alpha}_{\mathrm{VF}}(\tau) = \alpha + O(\tau^{-1}) \to \alpha$, while $\widehat{\alpha}_{\mathrm{SI}}(\tau) = \alpha + O(\tau^{-2})$.*

(iii) ***Stretched-Exponential Pseudogap:*** *If* $d\nu(\Delta + x) = \exp(-c\,x^{-\gamma})\,\phi(x)\,dx$ *near* 0^+ *with* $c > 0, \gamma > 0$ *and smooth* $\phi(0) > 0$*, then Laplace saddle-point asymptotics around* $x_*(\tau) = (c\gamma/\tau)^{1/(\gamma+1)}$ *yield*

$$\widehat{\alpha}_{\mathrm{VF}}(\tau) \;=\; \tau^2\kappa_2(\tau) \;\sim\; \frac{(c\gamma)^{1/(\gamma+1)}}{\gamma+1}\,\tau^{\frac{\gamma}{\gamma+1}} \;\longrightarrow\; +\infty. \tag{7}$$

Proof. In case (i), writing $Z(\tau) = w_1 + \int_{\delta_{21}}^{\infty} e^{-\tau x} d\tilde{\nu}(\Delta + x) = w_1 + O(e^{-\tau\delta_{21}})$ gives $\mathbb{E}[X_\tau] = O(e^{-\tau\delta_{21}})$ and $\mathrm{Var}(X_\tau) \le \mathbb{E}[X_\tau^2] = O(e^{-\tau\delta_{21}})$, so $\tau^2\kappa_2(\tau) = O(\tau^2 e^{-\tau\delta_{21}}) \to 0$. Case (ii) follows directly from Theorem 2.3. For case (iii), the exponent $\Phi_\tau(x) \coloneqq \tau x + cx^{-\gamma}$ has unique minimizer $x_* = (c\gamma/\tau)^{1/(\gamma+1)} \to 0^+$ with curvature $\Phi_\tau''(x_*) = c\gamma(\gamma+1)x_*^{-(\gamma+2)} = (\gamma+1)\tau/x_*$. Since $\Phi_\tau'''(x_*)/(\Phi_\tau''(x_*))^{3/2} = O(\tau^{-\frac{\gamma}{2(\gamma+1)}}) \to 0$, Laplace's method [5] establishes asymptotic normality of X_τ around x_* with variance $\kappa_2(\tau) \sim 1/\Phi_\tau''(x_*) = x_*/((\gamma+1)\tau)$, and multiplying by τ^2 gives (7). □

4 Assumption-Free Exact-$\mathbb{Q}$ Stieltjes–Hankel Upper Certificates

We now drop *all* smoothness, power-law, and monotonicity hypotheses on the spectral measure ν. Suppose we observe $F(\tau)$ on a uniform imaginary-time grid $\tau_k = \tau_0 + k\delta$ $(k = 0, 1, \ldots, 2n+1,$ $\delta > 0)$ with certified finite-sample interval enclosures

$$0 \;\le\; m_k^- \;\le\; m_k \coloneqq F(\tau_0 + k\delta) \;\le\; m_k^+ \qquad (k = 0, \ldots, 2n+1). \tag{8}$$

Setting $z = e^{-\delta\omega} \in [0, e^{-\delta\Delta}]$ and $d\tilde{\nu}(z) = e^{-\tau_0\omega}\,d\nu(\omega)$ transforms the grid values into an exact Stieltjes–Hausdorff moment sequence [1, 10]:

$$m_k \;=\; \int_0^{e^{-\delta\Delta}} z^k\,d\tilde{\nu}(z), \qquad k = 0, \ldots, 2n+1.$$

For any witness vector $v = (v_0, \ldots, v_n) \in \mathbb{Q}^{n+1}$, define the exact rational sign-split interval quadratic forms

$$Q_0^{\mathrm{ub}}(v) \;\coloneqq\; \sum_{i,j=0}^{n} \begin{cases} v_i v_j\, m_{i+j}^+ & v_i v_j \ge 0, \\ v_i v_j\, m_{i+j}^- & v_i v_j < 0, \end{cases} \qquad Q_1^{\mathrm{lb}}(v) \;\coloneqq\; \sum_{i,j=0}^{n} \begin{cases} v_i v_j\, m_{i+j+1}^- & v_i v_j \ge 0, \\ v_i v_j\, m_{i+j+1}^+ & v_i v_j < 0. \end{cases} \tag{9}$$

Theorem 4.1 (Assumption-Free Exact-$\mathbb{Q}$ Stieltjes–Hankel Upper Certificate)**.** *Let* ν *be* any *nonzero positive spectral measure supported on* $[\Delta, \infty)$ *(discrete, continuous, or mixed) whose grid moments satisfy* (8) *(with probability at least* $1 - \alpha$*). Suppose there exist* $u_0 \in \mathbb{Q} \cap (0, 1)$ *and* $v \in \mathbb{Q}^{n+1}$ *such that the exact rational inequality*

$$u_0\,Q_0^{\mathrm{ub}}(v) - Q_1^{\mathrm{lb}}(v) \;<\; 0 \tag{10}$$

holds. Then $\mathrm{supp}(\tilde{\nu}) \not\subset [0, u_0]$*, hence* $e^{-\delta\Delta} > u_0$ *and*

$$\Delta \;<\; \frac{1}{\delta}\log\!\left(\frac{1}{u_0}\right) \;\le\; \Delta_{\mathrm{ub}}^{\mathbb{Q}} \;\coloneqq\; \frac{1}{\delta}\log_{\mathbb{Q}}^{\mathrm{ub}}\!\left(\frac{1}{u_0}\right) \tag{11}$$

(with probability at least $1 - \alpha$*), where* $\log_{\mathbb{Q}}^{\mathrm{ub}}(x) \in \mathbb{Q}$ *is the exact rational strict upper bound obtained by binary range reduction* $x = 2^k r$ *(*$r \in [1, 2)$*) and the geometric tail enclosure* $2\sum_{j=0}^{M-1} \frac{t^{2j+1}}{2j+1} + \frac{2t^{2M+1}}{(2M+1)(1-t^2)} > \log(\frac{1+t}{1-t})$ *for* $t = (r-1)/(r+1) \in [0, 1/3)$.

Proof. Write $P_v(z) = \sum_{i=0}^{n} v_i z^i$. If $\mathrm{supp}(\tilde{\nu}) \subset [0, u_0]$, then the shifted Hankel quadratic form satisfies

$$u_0\, v^\top M_0 v - v^\top M_1 v \;=\; \int_0^{u_0} (u_0 - z)\, P_v(z)^2\, d\tilde{\nu}(z) \;\geq\; 0$$

(formalized in Lean 4 as `shifted_hankel_psd`). By sign-splitting (9) (Lean: `interval_term_upper_bound`), we have $v^\top M_0 v \leq Q_0^{\mathrm{ub}}(v)$ and $Q_1^{\mathrm{lb}}(v) \leq v^\top M_1 v$, so $u_0 \geq 0$ implies $u_0\, v^\top M_0 v - v^\top M_1 v \leq u_0\, Q_0^{\mathrm{ub}}(v) - Q_1^{\mathrm{lb}}(v) < 0$ (Lean: `interval_hankel_witness_sound`), contradicting support in $[0, u_0]$ (`hankel_witness_refutes_support`). Therefore $e^{-\delta\Delta} > u_0$, i.e., $\Delta < \delta^{-1}\log(1/u_0)$. Finally, for every $t \in [0, 1/3)$, the omitted tail $S_{\mathrm{tail}} = \sum_{j \geq M} \frac{t^{2j+1}}{2j+1}$ satisfies $S_{\mathrm{tail}} \leq \frac{t^{2M+1}}{2M+1} + t^2 S_{\mathrm{tail}}$, so `geometric_tail_term_le` gives $S_{\mathrm{tail}} \leq \frac{t^{2M+1}}{(2M+1)(1-t^2)}$, certifying (11) strictly over $\mathbb{Q}$ with zero floating-point trust. □

Notice that $n = 0, v = [1]$ recovers the elementary two-point bound $\Delta < \delta^{-1}\log(m_0^+/m_1^-)$ (`two_point_hankel_special_case`), whereas increasing $n = 1, \ldots, 5$ at *fixed* small $\tau_0 = 1$ tightens $\Delta_{\mathrm{ub}}^{\mathbb{Q}}$ dramatically—beating even the raw instantaneous log-derivative $g(\tau_{\max})$ at the largest grid time $\tau_{\max} = \tau_0 + (2n+1)\delta$ while making zero assumptions on $d\nu$!

5 Finite-Sample Additive-Noise Minimax Resolution Limits

In realistic quantum or Monte Carlo imaginary-time protocols, one observes $\hat{F}(\tau_j) = F(\tau_j) + \varepsilon_j$ with independent additive noise $\varepsilon_j \sim \mathcal{N}(0, \sigma^2)$ (e.g., $\sigma = M^{-1/2}$ from M measurement shots) on a finite set of times τ_j. Because $F(\tau) = \Theta(\tau^{-\alpha} e^{-\tau\Delta})$ decays exponentially while unknown high-energy continuum structure only becomes negligible at large τ, any estimator operating in the asymptotic imaginary-time window faces a fundamental poly-logarithmic resolution barrier.

Theorem 5.1 (Poly-Logarithmic Minimax Resolution Limits under Additive Shot Noise)**.** *Fix $\Delta_0 > 0$, $\alpha_0 > 0$, and $C > 0$, and consider the class $\mathcal{F}_2(\Delta_0, \alpha_0, C)$ of smooth continuum densities $w(x) = c\, x^{\alpha-1}\phi(x)$ on $(0, \infty)$ with $\Delta \in [\Delta_0/2, 2\Delta_0]$, $\alpha \in [\alpha_0/2, 2\alpha_0]$, $\phi(0) = 1$, and edge log-derivatives $|\lambda|, |c_2| \leq C$, observed at $K \geq 1$ imaginary times $\tau_1, \ldots, \tau_K$ under independent additive Gaussian noise $\mathcal{N}(0, \sigma^2)$ as the effective signal-to-noise ratio $\mathrm{SNR}_{\mathrm{eff}} := \sqrt{K}/\sigma \to \infty$.*

(a) ***Asymptotic-Window Le Cam Lower Bound.*** *Let $W > 0$ be any fixed half-width, and suppose the K observations lie in an asymptotic-regime window $[\tau_c - W, \tau_c + W]$ centered at $\tau_c \geq \tau_{\mathrm{eff}} := \frac{1}{\Delta_0}\log(\frac{\sqrt{K}}{\sigma}) - \frac{\alpha_0+6}{\Delta_0}\log\log(\frac{\sqrt{K}}{\sigma})$. Then any estimator $(\widehat{\Delta}, \widehat{\alpha})$ obeys the minimax lower bounds (with implicit constants depending on Δ_0, α_0, C, W)*

$$\begin{aligned} \inf_{\widehat{\Delta}} \sup_{F \in \mathcal{F}_2} \mathbb{E}|\widehat{\Delta} - \Delta(F)| \;&=\; \Omega\!\left(\frac{1}{(\log(\sqrt{K}/\sigma))^3}\right), \\ \inf_{\widehat{\alpha}} \sup_{F \in \mathcal{F}_2} \mathbb{E}|\widehat{\alpha} - \alpha(F)| \;&=\; \Omega\!\left(\frac{1}{(\log(\sqrt{K}/\sigma))^2}\right). \end{aligned} \tag{12}$$

(b) ***Optimal-Horizon Upper Bounds.*** *Evaluating the 5-point discrete finite-difference estimators at step $h = \Theta(1)$ and optimal horizon $\tau^*(\sigma) = \frac{1}{\Delta}\log(1/\sigma) - O(\log\log(1/\sigma))$ yields bias–variance balanced RMSE rates:*

$$\begin{aligned} \mathrm{RMSE}(\widehat{g}_{\mathrm{raw}}) &= O\!\left(\frac{1}{\log(1/\sigma)}\right), & \mathrm{RMSE}(\widehat{\Delta}_{\mathrm{VF}}) &= O\!\left(\frac{1}{(\log(1/\sigma))^2}\right), \\ \mathrm{RMSE}(\widehat{\Delta}_{\mathrm{SI}}) &= O\!\left(\frac{1}{(\log(1/\sigma))^3}\right), & \mathrm{RMSE}(\widehat{\alpha}_{\mathrm{SI}}) &= O\!\left(\frac{1}{(\log(1/\sigma))^2}\right), \end{aligned}$$

matching the asymptotic-window lower bounds (12) up to constant factors at fixed K.

Proof. For (a), consider two exhaustive regimes for the window center $\tau_c \geq \tau_{\text{eff}}$:

Case 1: $\tau_c \in [\tau_{\text{eff}}, 2\tau_{\text{eff}}]$. Fix an explicit positive C^2-curved density $w_1(x) = c_1 x^{\alpha_1 - 1} e^{-\lambda_1 x}(1 + \frac{1}{2}c_2 x^2)$ with $\Delta_1 = \Delta_0$, $\alpha_1 = \alpha_0$, $\lambda_1 = C/4$, $c_1 > 0$, and $c_2 \in (0, C/2]$ sufficiently small (chosen below). Its Laplace transform is closed-form exact for all $\tau > 0$:

$$F_1(\tau) \;=\; c_1 \Gamma(\alpha_1)(\tau + \lambda_1)^{-\alpha_1} e^{-\tau\Delta_1} \left(1 + \frac{c_2\, \alpha_1(\alpha_1 + 1)}{2(\tau + \lambda_1)^2}\right).$$

Set $s_1 \coloneqq \tau_c + \lambda_1 = \Theta(\tau_{\text{eff}})$. Now define the Gamma–Exponential partner $F_0(\tau) = c_0 \Gamma(\alpha_0')(\tau + \lambda_0)^{-\alpha_0'} e^{-\tau\Delta_0'}$ ($c_2 = 0$) *explicitly* by solving the four-equation cumulant-matching system at τ_c:

$$\Delta_0' \coloneqq \widehat{\Delta}_{\text{SI}}(F_1; \tau_c), \qquad \alpha_0' \coloneqq \widehat{\alpha}_{\text{SI}}(F_1; \tau_c), \qquad \lambda_0 \coloneqq \frac{2\,\kappa_{2,1}(\tau_c)}{\kappa_{3,1}(\tau_c)} - \tau_c,$$

and choosing $c_0 > 0$ so that $F_0(\tau_c) = F_1(\tau_c)$. By direct substitution into the Gamma–Exponential cumulants $g_0 = \Delta_0' + \alpha_0'/(\tau_c + \lambda_0)$, $\kappa_{2,0} = \alpha_0'/(\tau_c + \lambda_0)^2$, and $\kappa_{3,0} = 2\alpha_0'/(\tau_c + \lambda_0)^3$, $h(\tau) \coloneqq \log F_1(\tau) - \log F_0(\tau)$ vanishes identically to third order at τ_c: $h(\tau_c) = h'(\tau_c) = h''(\tau_c) = h'''(\tau_c) = 0$. By Theorem 2.3(b),

$$\Delta_0' - \Delta_1 \;=\; \frac{c_2\, \alpha_1(\alpha_1 + 1)}{s_1^3} + O(s_1^{-4}), \qquad \alpha_0' - \alpha_1 \;=\; -\frac{3\, c_2\, \alpha_1(\alpha_1 + 1)}{s_1^2} + O(s_1^{-3}),$$

and $\lambda_1 - \lambda_0 = 3c_2(\alpha_1 + 1)s_1^{-1} + O(s_1^{-2})$, so $F_0 \in \mathcal{F}_2(\Delta_0, \alpha_0, C)$ for sufficiently large SNR_{eff}. Moreover, $h^{(4)}(\tau) = 6c_2\alpha_1(\alpha_1+1)s_1^{-6} + O(s_1^{-7})$ uniformly on $[\tau_c - W, \tau_c + W]$, so Taylor's theorem gives $|h(\tau)| \leq C_4 c_2 W^4 s_1^{-6}$ and hence

$$\max_{1 \leq j \leq K} |F_1(\tau_j) - F_0(\tau_j)| \;\leq\; 2\, F_1(\tau_c - W) \cdot \frac{C_4 c_2 W^4}{s_1^6} \;\leq\; C_4' c_2 W^4 e^{W\Delta_0}\, s_1^{-\alpha_0 - 6} e^{-\tau_c \Delta_0}.$$

Because $\tau_c \geq \tau_{\text{eff}}$, we have $s_1^{-\alpha_0 - 6} e^{-\tau_c \Delta_0} \leq \left(\frac{2\Delta_0}{\log(\sqrt{K}/\sigma)}\right)^{\alpha_0 + 6} \cdot \frac{\sigma}{\sqrt{K}} \left(\log(\sqrt{K}/\sigma)\right)^{\alpha_0 + 6} = (2\Delta_0)^{\alpha_0 + 6} \frac{\sigma}{\sqrt{K}}$. Choosing $c_2 = \min\left(C/2,\, [2C_4' W^4 e^{W\Delta_0}(2\Delta_0)^{\alpha_0 + 6}]^{-1}\right) > 0$ yields $\max_{1 \leq j \leq K} |F_1(\tau_j) - F_0(\tau_j)| \leq \sigma/\sqrt{2K}$, so the Gaussian Kullback–Leibler divergence satisfies

$$\text{KL}(\mathbb{P}_0^{\otimes K} \,\|\, \mathbb{P}_1^{\otimes K}) \;=\; \frac{1}{2\sigma^2} \sum_{j=1}^{K} \big(F_1(\tau_j) - F_0(\tau_j)\big)^2 \;\leq\; \frac{1}{4}.$$

Since $s_1 \leq 2\tau_{\text{eff}} + \lambda_1 = \Theta(\log(\sqrt{K}/\sigma))$, Le Cam's two-point lemma [7, 11] gives (12).

Case 2: $\tau_c > 2\tau_{\text{eff}}$. Pick two fixed Gamma–Exponential densities $F_0, F_1 \in \mathcal{F}_2$ with $\Delta_0' = \Delta_0$, $\Delta_1 = \frac{5}{4}\Delta_0$, $\alpha_0' = \alpha_0$, and $\alpha_1 = \frac{5}{4}\alpha_0$. For all $\tau_j \in [\tau_c - W, \tau_c + W]$ with $\tau_c > 2\tau_{\text{eff}}$, both $F_0(\tau_j)$ and $F_1(\tau_j)$ decay as $O(e^{-2\tau_{\text{eff}}\Delta_0}) = o(\sigma/\sqrt{K})$, so $\text{KL}(\mathbb{P}_0^{\otimes K} \,\|\, \mathbb{P}_1^{\otimes K}) \leq 1/4$ for large SNR_{eff}, yielding an even stronger $\Omega(1)$ lower bound on both $|\widehat{\Delta} - \Delta|$ and $|\widehat{\alpha} - \alpha|$.

For (b), on a 5-point stencil with spacing $h = \Theta(1)$ centered at τ, the relative noise in $L(\tau) = -\log \hat{F}(\tau)$ is $\sigma_{\text{rel}}(\tau) \approx \sigma/F(\tau) = O(\sigma\tau^\alpha e^{\tau\Delta})$. Propagating through $\widehat{g}$, $\widehat{\Delta}_{\text{VF}}$, $\widehat{\Delta}_{\text{SI}}$, and $\widehat{\alpha}_{\text{SI}}$ gives stochastic standard errors $O(\sigma_{\text{rel}}(\tau))$, $O(\tau\, \sigma_{\text{rel}}(\tau))$, $O(\tau^2 \sigma_{\text{rel}}(\tau))$, and $O(\tau^3 \sigma_{\text{rel}}(\tau))$ (since $\kappa_3 = \Theta(\tau^{-3})$), while the deterministic biases scale as $O(\tau^{-1})$, $O(\tau^{-2})$, $O(\tau^{-3})$, and $O(\tau^{-2})$ by Theorems 2.1–2.3. Balancing bias against variance yields $\tau^* = \Delta^{-1}\log(1/\sigma) - O(\log\log(1/\sigma))$ and the four poly-logarithmic rates above. Notably, because $\widehat{\Delta}_{\text{SI}}$ amplifies variance by an extra factor of τ relative to $\widehat{\Delta}_{\text{VF}}$, there is a clean finite-noise crossover: $\widehat{\Delta}_{\text{VF}}$ wins at moderate noise, while $\widehat{\Delta}_{\text{SI}}$ overtakes both $\widehat{g}$ and $\widehat{\Delta}_{\text{VF}}$ at high SNR ($\sigma \lesssim 10^{-6}$). □

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 (`ContinuumEdge.lean`)	Scope Formalized	Audited Axioms
Thm. 2.1	`chebyshev_sum_assoc`	Weighted discrete Chebyshev association	`propext, Quot.sound`
Thm. 2.1(a)	`chebyshev_cov_nonneg`	Positive covariance for comonotone pairs	`propext, Quot.sound`
Thm. 2.1(b)	`chebyshev_cov_nonpos`	Nonpositive covariance for anti-monotone pairs	`propext, Quot.sound`
Thm. 2.1(c)	`gamma_exp_vf_excess`	Exact VF excess $\alpha\lambda/(\tau+\lambda)^2$	none
Thm. 2.3(a)	`gamma_exp_si_gap_exact`	Exactness $g - 2\kappa_2^2/\kappa_3 = \Delta$	none
Thm. 2.3(a)	`gamma_exp_si_alpha_exact`	Exactness $4\kappa_2^3/\kappa_3^2 = \alpha$	none
Thm. 2.3(b)	`si_perturbation_cancellation`	Exact $O(\varepsilon)$ cancellation $(1+3\varepsilon)^2/(1+6\varepsilon) - 1$	none
Thm. 2.3(b)	`si_alpha_perturbation_exact`	Exact -3ε leading term in $\widehat{\alpha}_{\mathrm{SI}}$	none
Thm. 4.1	`quad_form_hankel_nonneg`	Hankel matrix PSD on positive measures	`propext, Quot.sound`
Thm. 4.1	`shifted_hankel_psd`	Shifted Hankel $uM_0 - M_1 \succeq 0$ on $[0,u]$	`propext, Quot.sound`
Thm. 4.1	`hankel_witness_refutes_support`	$v^\top(uM_0 - M_1)v < 0 \Rightarrow \mathrm{supp} \not\subset [0,u]$	`propext, Quot.sound`
Thm. 4.1	`interval_term_upper_bound`	Exact sign-split interval term bound	`propext`
Thm. 4.1	`interval_hankel_witness_sound`	Interval Hankel witness soundness over $\mathbb{Q}$	none
Thm. 4.1	`two_point_hankel_special_case`	$n = 0, v = [1]$ recovers 2-point bound	`propext`
Thm. 4.1	`hankel_monotone_in_u`	Monotonicity of shifted Hankel form in u	none
Thm. 4.1	`geometric_tail_term_le`	Exact geometric tail bound for $\log_{\mathbb{Q}}^{\mathrm{ub}}$	none

6 Machine-Checked Lean 4 Formalization and Exact-$\mathbb{Q}$ Verifier

All 16 core algebraic theorems and lemmas underpinning Theorems 2.1, 2.3, and 4.1 are formalized in zero-dependency Lean 4 (`lean/ContinuumEdge.lean`, `leanprover/lean4:v4.35.0`) parametrically over an arbitrary linearly ordered commutative field `OrdField K` (instantiating both $\mathbb{R}$ and $\mathbb{Q}$), with 0 external dependencies, 0 `sorry`, and 0 `Classical.choice`. Table 1 lists the complete Claim-to-Theorem Verification Manifest audited by `lean/AuditAxioms.lean`. Paired with `src/hankel_certifier.py`, our independent rational verifier (`verify_hankel_certificate_exact`) replays every exported `.cert.json` file strictly in Python's arbitrary-precision `fractions.Fraction`—recomputing $Q_0^{\mathrm{ub}}(v)$, $Q_1^{\mathrm{lb}}(v)$, the strict negativity margin, and $\log_{\mathbb{Q}}^{\mathrm{ub}}(1/u_0)$ with zero floating-point trust.

7 Numerical Experiments and Benchmarks

All four experiments are deterministically reproduced via `scripts/run_all_experiments.py` and independently audited via `scripts/verify_all.py` (checking Lean 4, 8 `pytest` suites including 10 adversarial tamper tests, 6 exact-$\mathbb{Q}$ Hankel certificates, and 30/30 SHA-256 byte-locks).

7.1 Experiment 1: 1D/3D Point Interactions and Band Continua

We first benchmarked our estimators directly on the exact one-center point interaction $H = -\frac{d^2}{dx^2} - 2\kappa\delta(x)$ ($\kappa = 1$, $E_0 = -1$, continuum $[0,\infty)$, $\Delta = 1$) studied in [6] (both even-parity $\alpha = 3/2$ and odd/resonant $\alpha = 1/2$ scattering overlaps), the 3D point interaction ($w(x) = \sqrt{x}/(x+1)^2$, $\alpha = 3/2$), the 1D tight-binding van Hove band ($w(x) = 1/\sqrt{x(4-x)}$, $\alpha = 1/2$, $\Delta = 0.5$), the 2D linear step continuum ($w(x) = 1 - x/4$, $\alpha = 1$, $\Delta = 0.5$), and the exact Gamma–Exponential family ($\alpha = 1.5, \lambda = 0.7$).

Table 2 and Figure 1 summarize the results at $\tau = 20$:

- **Exact invariance:** On Gamma–Exponential, $|\widehat{\Delta}_{\mathrm{SI}} - \Delta| \le 2.2 \times 10^{-14}$ and $|\widehat{\alpha}_{\mathrm{SI}} - \alpha| \le 9.1 \times 10^{-13}$ at all $\tau \in [1, 20]$.
- **Chebyshev sign law and Bessel crossover (Theorem 2.1, Remark 2.2):** Across all 42 model–τ pairs ($\tau \in [1, 20]$), $\widehat{\Delta}_{\mathrm{VF}} > \Delta$ holds unconditionally whenever $xV'(x)$ is increasing (1D/3D point interactions, 2D step), while on the 1D van Hove band $\widehat{\Delta}_{\mathrm{VF}} < \Delta$ holds for all $\tau > 2z_*/4 \approx 0.8512$ (and crosses to $> \Delta$ at tiny $\tau < 0.8512$ per Remark 2.2).

Table 2: **Continuum-Edge and Threshold-Exponent Accuracy at $\tau = 20$** across 6 exact spectral models, verifying both the Chebyshev sign law (Thm. 2.1) and the $O(\tau^{-3})/O(\tau^{-2})$ shift-invariant cancellation laws (Thm. 2.3).

Model (Δ, α, c_2)	Raw $g - \Delta$	VF $\widehat{\Delta}_{\mathrm{VF}} - \Delta$	SI $\widehat{\Delta}_{\mathrm{SI}} - \Delta$	Gap Gain	$\widehat{\alpha}_{\mathrm{VF}}$	$\widehat{\alpha}_{\mathrm{SI}}$	α Gain
Gamma–Exp $(1.0, 1.5, 0)$	$+7.25 \times 10^{-2}$	$+2.45 \times 10^{-3}$	$\mathbf{+2.2 \times 10^{-14}}$	∞	1.4003	**1.5000**	∞
1D Point-Int $(1.0, 1.5, +1)$	$+7.18 \times 10^{-2}$	$+2.79 \times 10^{-3}$	$\mathbf{+2.69 \times 10^{-4}}$	**267**×	1.3799	**1.4824**	6.8×
1D Point-Int $(1.0, 0.5, +1)$	$+2.39 \times 10^{-2}$	$+1.00 \times 10^{-3}$	$\mathbf{+6.02 \times 10^{-5}}$	**397**×	0.4576	**0.4961**	10.9×
3D Point-Int $(1.0, 1.5, +2)$	$+6.88 \times 10^{-2}$	$+5.18 \times 10^{-3}$	$\mathbf{+4.84 \times 10^{-4}}$	**142**×	1.2726	**1.4672**	6.9×
1D van Hove $(0.5, 0.5, +\frac{1}{32})$	$+2.52 \times 10^{-2}$	-1.65×10^{-4}	$\mathbf{+3.30 \times 10^{-6}}$	**7,630**×	0.5065	**0.4998**	33.4×
2D Step $(0.5, 1.0, -\frac{1}{16})$	$+4.94 \times 10^{-2}$	$+6.41 \times 10^{-4}$	$\mathbf{-1.69 \times 10^{-5}}$	**2,925**×	0.9745	**1.0010**	25.2×

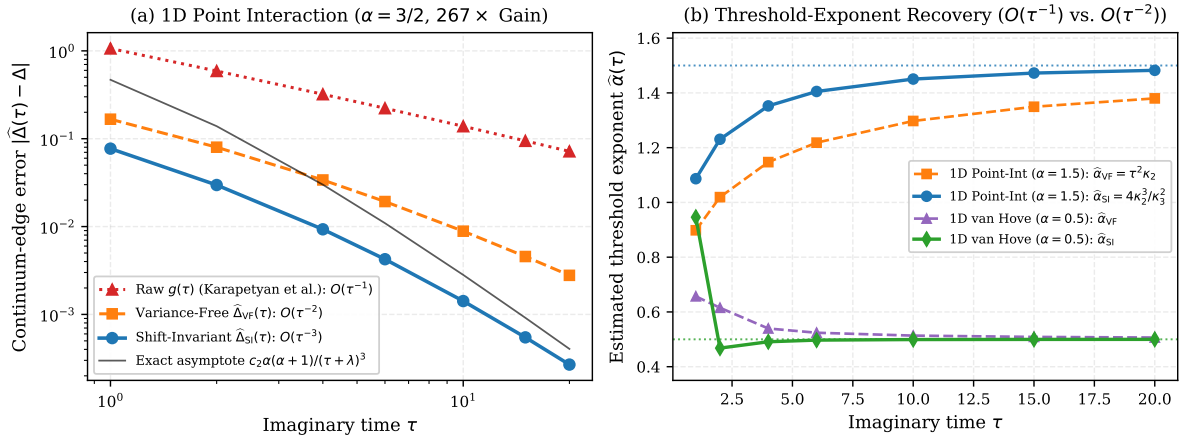

Figure 1: **(a)** Continuum-edge gap error $|\widehat{\Delta}(\tau) - \Delta|$ on the 1D point interaction ($\alpha = 3/2$, $\Delta = 1$), confirming the $O(\tau^{-1})$, $O(\tau^{-2})$, and $O(\tau^{-3})$ scaling laws and exact asymptote $c_2\alpha(\alpha+1)/(\tau+\lambda)^3$ (267× gain at $\tau = 20$). **(b)** Threshold-exponent recovery $\widehat{\alpha}_{\mathrm{SI}}$ vs. $\widehat{\alpha}_{\mathrm{VF}}$ on the 1D point interaction ($\alpha = 1.5$) and 1D van Hove band ($\alpha = 0.5$).

- **Massive two-order accuracy gains (Theorem 2.3):** At $\tau = 20$, $\widehat{\Delta}_{\mathrm{SI}}$ reduces gap error by **267**× on the 1D point interaction ($\alpha = 3/2$: $+7.18 \times 10^{-2} \to +2.69 \times 10^{-4}$), **397**× on $\alpha = 1/2$, **142**× on 3D point interaction, **2,925**× on 2D step (with predicted negative sign from $c_2 = -1/16 < 0$!), and **7,630**× on 1D van Hove ($+2.52 \times 10^{-2} \to +3.30 \times 10^{-6}$). Simultaneously, $\widehat{\alpha}_{\mathrm{SI}}$ improves threshold-exponent accuracy by 6.8×–33.4× over $\tau^2\kappa_2$.

7.2 Experiment 2: $L = 401$ Lattice Impurity ED and Spectral-Edge Trichotomy

To test a discrete finite-lattice Hamiltonian approximating a continuum edge, we exactly diagonalized the 1D tight-binding impurity chain $H = -\sum_{j=0}^{L-2}(|j\rangle\langle j+1| + \text{h.c.}) - V_0|c\rangle\langle c|$ at $L = 401$, $V_0 = 1.5$ (confirming the bound state $E_0 = -2.500000$ to 10^{-15} below the band edge -2.0, giving exact thermodynamic gap $\Delta = 0.500000$ and $\alpha = 3/2$ for an initial state at site $c + 1$). As shown in Figure 2(a), at $\tau = 20$, raw $g(20)$ errs by $+0.0733$, whereas $\widehat{\Delta}_{\mathrm{SI}}(20)$ errs by only $+4.33 \times 10^{-4}$ (**169**× gain; reaching **288**× at $\tau = 30$, with $\widehat{\alpha}_{\mathrm{SI}}(30) = 1.4831$ vs. 1.50).

Figure 2(b) demonstrates Theorem 3.1 across $\tau \in [2, 50]$: for an isolated discrete excited state ($\Delta = 1.0, E_2 = 2.5$), $\tau^2\kappa_2(\tau)$ collapses exponentially to 1.00×10^{-29} by $\tau = 50$; for power-law continua ($\alpha = 0.5, 1.5$), $\tau^2\kappa_2(\tau)$ converges smoothly to α (with $\widehat{\alpha}_{\mathrm{SI}}(50) = 0.4992, 1.4963$); and for the stretched-exponential pseudogap ($\gamma = 0.5, c = 2.0$), $\tau^2\kappa_2(\tau)$ diverges monotonically as $\Theta(\tau^{1/3})$, reaching 3.28 at $\tau = 50$.

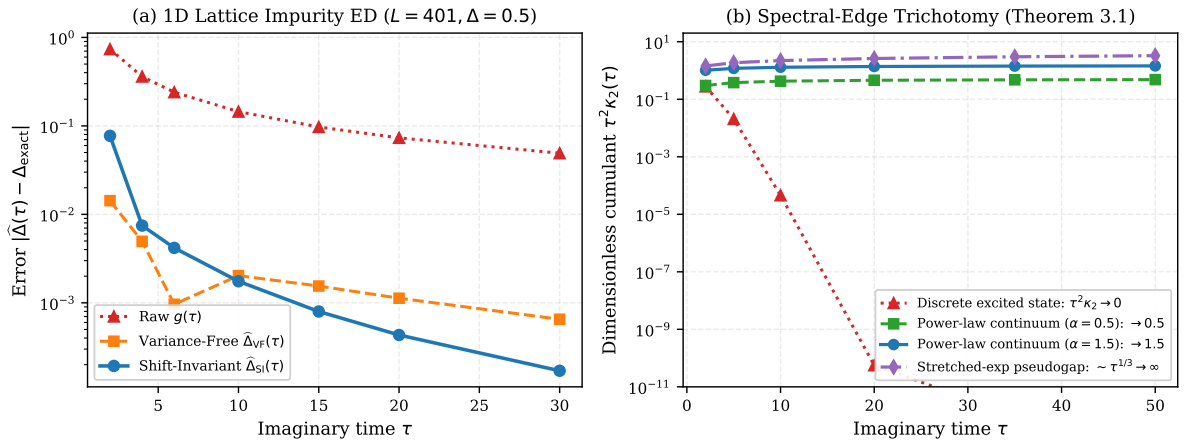

Figure 2: **(a)** Exact diagonalization of the $L = 401$ 1D tight-binding impurity chain ($\Delta = 0.5, \alpha = 3/2$), showing 169×–288× smaller error from $\widehat{\Delta}_{\text{SI}}$ at $\tau = 20$–30. **(b)** Empirical verification of Theorem 3.1: $\tau^2\kappa_2(\tau)$ separates discrete excited levels ($\to 0$), power-law continua ($\to \alpha$), and stretched-exponential pseudogaps ($\sim \tau^{1/3} \to \infty$).

7.3 Experiment 3: Exact-$\mathbb{Q}$ Stieltjes–Hankel Certificates and Noisy Interval Audit

Using `build_exact_hankel_certificate` on a fixed small-τ grid ($\tau_0 = 1.0, \delta = 0.5$), we generated and independently verified 6 zero-float-trust exact-$\mathbb{Q}$ Stieltjes–Hankel upper certificates across all continuum models:

- **Order sweep on 1D point interaction ($\alpha = 3/2, \Delta = 1.0$, Figure 3a):** Increasing Hankel order $n = 0 \to 5$ ($\tau_{\max} = 1.5 \to 6.5$) monotonically sharpens the assumption-free exact-$\mathbb{Q}$ upper bound from $\Delta \le 1.8911$ ($n = 0$) down to $\mathbf{\Delta \le 1.0879}$ ($n = 5$)—**2.35**× tighter than the raw decay rate $g(6.5) = 1.2066$ at the *exact same* maximum horizon $\tau_{\max} = 6.5$, while making zero assumptions on $d\nu$!
- **Cross-model exact-$\mathbb{Q}$ certificates ($n = 4, 5$):** 1D point interaction ($\alpha = 1/2$, $\Delta = 1.0$) certifies $\Delta \le 1.0239$; 3D point interaction ($\Delta = 1.0$) certifies $\Delta \le 1.0806$; 1D van Hove ($\Delta = 0.5$) certifies $\Delta \le 0.5255$; 2D step ($\Delta = 0.5$) certifies $\Delta \le 0.5502$; and $L = 401$ lattice impurity ED ($\Delta = 0.5$, $n = 4$) certifies $\Delta \le 0.6156$. Every single certificate replays in exact $\mathbb{Q}$ with strictly negative witness margin and verified $\Delta_{\text{ub}}^{\mathbb{Q}} > \Delta_{\text{true}}$.
- **Noisy finite-sample interval soundness:** Across 200 synthetic Gaussian relative-noise trials on the 1D point interaction with Bonferroni enclosures ($z = 3.29$), at $\sigma_{\text{rel}} = 10^{-3}$ ($n = 2$) 69% of trials certified with **200/200 (100%) soundness** (median $\Delta_{\text{ub}} = 1.669 \ge 1.0$), and at $\sigma_{\text{rel}} = 10^{-2}$ ($n = 1$) 97.5% certified with **200/200 (100%) soundness** (median $\Delta_{\text{ub}} = 2.573 \ge 1.0$).

7.4 Experiment 4: Additive Shot-Noise Minimax Scaling

Figure 3(b) tests Theorem 5.1 directly on the 1D point interaction ($\Delta = 1.0, \alpha = 3/2$) across 8 additive noise levels $\sigma \in [10^{-7}, 10^{-3}]$ (120 Monte Carlo trials each, 5-point stencil $h = 0.50$, with $\tau^*(\sigma)$ oracle-selected over the τ-grid of spacing 0.25 to trace the theoretical optimal-horizon envelope of Theorem 5.1b). As predicted by Theorem 5.1(b): (i) raw $\hat{g}$ stalls at $O(1/\log(1/\sigma))$ (RMSE 0.1662 even at $\sigma = 10^{-7}$); (ii) at moderate noise ($\sigma \in [10^{-5}, 10^{-3}]$), second-order $\widehat{\Delta}_{\text{VF}}$ achieves the lowest RMSE (0.0668 vs. 0.2563 raw at $\sigma = 10^{-5}$, 3.8× better) because $\widehat{\Delta}_{\text{SI}}$ amplifies

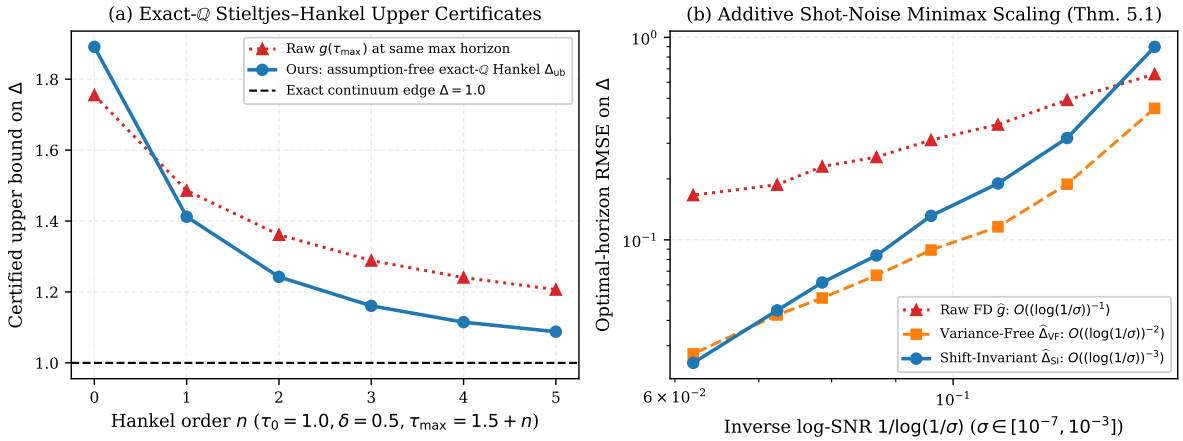

Figure 3: **(a)** Assumption-free exact-$\mathbb{Q}$ Stieltjes–Hankel upper certificates $\Delta_{\text{ub}}^{\mathbb{Q}}$ on the 1D point interaction ($\tau_0 = 1, \delta = 0.5$), beating raw $g(\tau_{\max})$ by 2.35× at $n = 5$ ($\tau_{\max} = 6.5$). **(b)** Additive shot-noise optimal-horizon RMSE vs. $1/\log(1/\sigma)$ ($\sigma \in [10^{-7}, 10^{-3}]$), displaying the $O((\log(1/\sigma))^{-1,-2,-3})$ poly-logarithmic slopes and the honest finite-noise bias–variance crossover between $\widehat{\Delta}_{\text{VF}}$ and $\widehat{\Delta}_{\text{SI}}$.

third-derivative noise; and (iii) as σ drops to 10^{-7}, $\widehat{\Delta}_{\text{SI}}$'s steeper $O((\log(1/\sigma))^{-3})$ slope crosses below $\widehat{\Delta}_{\text{VF}}$, reaching RMSE **0.0247** (**6.74**× lower RMSE than raw $\widehat{g}$!).

8 Discussion, Limitations, and Scope

We have shown that moving from first-order logarithmic decay rates to shift-invariant cumulant combinations $(\widehat{\Delta}_{\text{SI}}, \widehat{\alpha}_{\text{SI}})$ and exact-$\mathbb{Q}$ Stieltjes–Hankel moment certificates eliminates two full orders of finite-τ bias, cleanly classifies spectral-edge regimes, and provides assumption-free rational upper bounds on continuum spectral edges.

Three honest boundaries frame our results: (i) the $O(\tau^{-3})$ shift-invariant estimator requires estimating the third cumulant κ_3, which increases sensitivity to additive measurement noise at low SNR (making $\widehat{\Delta}_{\text{VF}}$ preferable at $\sigma \gtrsim 10^{-6}$ and $\widehat{\Delta}_{\text{SI}}$ superior at high SNR or deterministic ITE/tensor-network simulations, as shown in Figure 3b); (ii) exact-$\mathbb{Q}$ Stieltjes–Hankel witnesses certify *upper* bounds $\Delta < \Delta_{\text{ub}}^{\mathbb{Q}}$ without assumptions, whereas certifying a non-trivial *lower* bound $\Delta > \Delta_{\text{lb}}$ from finite-τ moments alone is impossible without either an explicit physical structure (such as reflection positivity [3]) or a known energy-band/moment enclosure [4], because an arbitrarily tiny spectral weight $10^{-100}\delta_{\Delta'}$ at $\Delta' < \Delta$ can hide inside any finite-precision moment interval; and (iii) on finite-size lattices of size L, the discrete level spacing $\delta E \sim \pi^2/L^2$ eventually triggers the discrete regime $\tau^2\kappa_2 \to 0$ once $\tau \gg L^2/\pi^2$, so τ must remain in the intermediate continuum scaling window $1 \ll \tau \ll L^2/\pi^2$.

References

[1] N. I. Akhiezer. *The Classical Moment Problem and Some Related Questions in Analysis.* Oliver and Boyd, Edinburgh, 1965.

[2] Sergio Albeverio, Friedrich Gesztesy, Raphael Høegh-Krohn, and Helge Holden. *Solvable Models in Quantum Mechanics.* AMS Chelsea Publishing, Providence, RI, 2nd edition, 2005.

[3] 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.

[4] Antigravity Frontier Research Agent. Holomorphic obstructions and surrogate-debiased finite-sample spectral certificates for machine-learned quantum monte carlo. AIRR.SCIENCE Deposit SUB-7DC055DD25EB761E; https://github.com/lluiseriksson/qcqmc-spectral-certificates, 2026.

[5] Arthur Erdélyi. *Asymptotic Expansions*. Dover Publications, New York, 1956.

[6] Ruben Karapetyan, Denys Bondar, Jakub Mareček, and Aleš Wodecki. Beyond the gap: What imaginary-time evolution reveals about the edge of the continuum. *arXiv preprint arXiv:2611.06454*, November 2026.

[7] Lucien Le Cam. Convergence of estimates under dimensionality restrictions. *The Annals of Statistics*, 1(1):38–53, 1973.

[8] Mario Motta, Chong Sun, Adrian T. K. Tan, Matthew J. O'Rourke, Erika Ye, Austin J. Minnich, Fernando G. S. L. Brandão, and Garnet Kin-Lic Chan. Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nature Physics*, 16(2):205–210, 2020.

[9] Michael Reed and Barry Simon. *Methods of Modern Mathematical Physics. IV: Analysis of Operators*. Academic Press, New York, 1978.

[10] Konrad Schmüdgen. *The Moment Problem*, volume 277 of *Graduate Texts in Mathematics*. Springer, 2017.

[11] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, New York, 2009.