# Finite-Sample Spectral-Gap Falsification

## Exact weighted visibility minimax, hidden-atom LAN, and honest dependent tests

Lluis Eriksson

August 2026


### Abstract

Finite Euclidean correlation matrices are routinely converted into spectral gaps by generalized eigenvalue, Prony, or Lanczos procedures. We determine a finite-sample boundary for what such data can falsify and what they cannot certify by non-rejection. Let $T$ be a positive self-adjoint contraction and $B_n = V^*T^nV$ its transfer moments. The block Hankel pencil $H_N = [B_{i+j}]$, $G_N = [B_{i+j+1}]$ gives monotone Ritz lower bounds on the visible transfer edge. Exact rank stabilization is terminal: it makes the Krylov space invariant and recovers the complete visible spectrum from finitely many moments.

Away from terminality, an atom of weight $w$ at any $x_* = e^{-g}$ changes every moment by at most $w$ while forcing mass gap at most $g$. Taking $x_* < 1$ arbitrarily close to one proves non-identifiability even when an exactly degenerate vacuum excitation is forbidden. For Gaussian correlator sketches with nonsingular gapped covariance $K_0$, the hidden atom produces a rank-one spike after whitening. We compute its exact likelihood ratio and relative entropy. For $w_n = h/\sqrt{n}$ the experiment is locally asymptotically normal with information

$$I_{x_*} = \tfrac{1}{2}\{(\|K_0^{-1/2}\phi_N(x_*)\|^2 - 1)^2 + 2N + 1\},$$

giving the sharp local power envelope $1 - \Phi(z_{1-\alpha} - h\sqrt{I_{x_*}})$. In particular,

$$D(\mathsf{P}_{w,x_*}^{\otimes n}\|\mathsf{P}_0^{\otimes n}) \le \tfrac{1}{2}nw^2\|A_{x_*}\|_F^2, \qquad \mathrm{TV} \le \tfrac{1}{2}w\sqrt{n}\|A_{x_*}\|_F,$$

under an explicit local condition. Thus every level-$\alpha$ test has power at most $\alpha + O(w\sqrt{n})$ against alternatives whose positive gaps can tend to zero.

Quantitative visibility restores honest one-sided detection. We solve the relevant population optimization exactly. If at least mass $\gamma$ lies past $\theta + \delta$, the unique degree-$N$ weighted minimax filter is the normalized Chebyshev polynomial of the fourth kind,

$$q_N(x) = \frac{W_N(2x/\theta - 1)}{W_N(1 + 2\delta/\theta)}.$$

Its worst-case localizer is exactly $\theta(1-\gamma)/W_N(1+2\delta/\theta)^2 - \delta\gamma$; a two-atom measure attains equality. This gives a necessary-and-sufficient uniform sign threshold within the normalized polynomial-localizer class and strictly improves the usual first-kind stop-band envelope.

A simultaneous Wishart–Loewner confidence band yields finite-sample Gaussian level without a known covariance, asymptotic approximation, or sample split. We also prove a distribution-free companion within the bounded-iid model: paired differences remove an unknown mean, Hoeffding calibration gives finite-sample level, and a finite bank of filters may be selected without splitting at only a logarithmic penalty. A second theorem covers strictly stationary bounded $\beta$-mixing readouts. Sparse pairing, Berbee coupling, and an explicit lag-covariance correction give nonasymptotic level and power from one dependent trajectory whenever an external mixing envelope is available. The exact weighted margin gives a closed joint depth–sample tradeoff. We do not claim minimax optimality outside the declared fixed-window Gaussian experiment or normalized scalar polynomial-localizer class.

These tests falsify an overstated positive gap by detecting visible mass beyond $\theta$. Non-rejection alone is not a positive lower-confidence certificate for the gap. An interacting

ANNNI-chain experiment up to $2^{16}$ states exhibits the visibility issue: a parity-even probe is exactly blind to the lowest odd excitation, whereas a two-channel $(X, Z)$ family supplies a finite witness. Exact rational fixtures, an independent certificate, a Colab pipeline, and a Lean-checked hidden-atom core accompany the proofs.

## 1 The identifiability question

A spectral gap is a support exclusion, not merely a fitted exponential slope. In a transfer formulation, after removing the vacuum, it is the assertion

$$\operatorname{spec}(T) \subset [0, e^{-m}] \qquad (m > 0), \tag{1}$$

for a positive self-adjoint contraction $T$. Numerical and lattice calculations generally do not access $T$ directly. They access finitely many matrices

$$B_n = V^* T^n V, \qquad n = 0, 1, \ldots, M, \tag{2}$$

where the columns of $V : \mathbb{C}^d \to \mathcal{K}$ are states created by a chosen family of interpolating operators. These are Hausdorff moments of a positive matrix-valued spectral measure. This observation underlies generalized eigenvalue methods, Prony-type reconstruction, Rayleigh–Ritz and block Lanczos spectroscopy [1, 2, 3, 4, 5].

The inverse problem has two logically different targets:

(i) estimate spectral values that are visible to the chosen probes;

(ii) certify that no spectral value relevant to the claimed gap has been missed.

The first is an approximation problem. The second is an identifiability problem. A highly accurate estimate does not answer the second question if a small spectral component can hide below the error floor.

The direction of inference is crucial. Throughout, the statistical null is $\operatorname{supp} \nu \subseteq [0, \theta]$. Rejection detects visible edge mass beyond $\theta$, thereby falsifying an overstated positive gap and giving a lower bound on the visible transfer edge. Non-rejection alone does *not* exclude hidden mass and is not a positive lower-confidence bound on the gap.

This paper gives a six-branch answer.

**Exact terminal branch.** Rank stabilization of a block Hankel matrix is a finite, checkable condition under which the Krylov approximation terminates and the visible spectrum is recovered exactly.

**No-go branch.** For every positive componentwise error radius and every target $g > 0$, a spectral atom at $e^{-g} < 1$ can be hidden inside the data ball. Thus even a known unique vacuum does not permit a uniformly separated positive gap over the unconstrained model class.

**Statistical lower branch.** For nonsingular Gaussian correlator sketches, the near-critical atom has an exact rank-one-spike likelihood. Its $w \asymp n^{-1/2}$ local experiment is Gaussian with an explicit information and sharp power curve; below that scale all level tests are asymptotically powerless.

**Gaussian visibility branch.** A lower bound on the mass of any missed edge component restores a finite falsifier. A fourth-kind weighted-Chebyshev filter solves the population localizer minimax problem exactly and yields an explicit depth and noise margin; a Wishart–Loewner band converts it into an exact-level test and a closed depth–sample tradeoff.

**Bounded-readout branch.** For iid bounded readouts with unknown mean, paired differences and a one-sided Hoeffding band give a non-Gaussian finite-sample test. A finite filter bank may be searched on the same data with a disclosed multiplicity penalty.

**Dependent-readout branch.** For a strictly stationary bounded $\beta$-mixing trajectory, lagged pairing introduces a controlled covariance bias while separated pair blocks admit a Berbee coupling. This yields an explicit finite-sample test and power bound, including a closed geometric-mixing specialization.

The algebraic ingredients are classical: positive moment problems, Krylov termination, Chebyshev polynomials, and Gaussian LAN. The claim here is the assembled operational boundary—including the distribution-level lower and local-power laws, the exact weighted visibility minimax, and auditable finite-sample witnesses—not the invention of Rayleigh–Ritz or of the truncated moment problem.

## 2 Matrix moments and the visible spectral edge

Let $\mathcal{K}$ be a complex Hilbert space, let $0 \le T \le \mathbf{1}$ be self-adjoint, and let $V : \mathbb{C}^d \to \mathcal{K}$ be bounded. Write $E_T$ for the spectral resolution of $T$ and

$$\Sigma(A) = V^* E_T(A) V, \qquad B_n = \int_0^1 x^n \, d\Sigma(x) = V^* T^n V. \tag{3}$$

The visible cyclic subspace and its edge are

$$\mathcal{K}_V = \overline{\operatorname{span}}\{T^j \operatorname{ran} V : j \ge 0\}, \qquad \rho_V = \|T|_{\mathcal{K}_V}\|. \tag{4}$$

Only $T|_{\mathcal{K}_V}$ is identifiable from the moments. If $\mathcal{K}_V \neq \mathcal{K}$, a single probe family is blind to the orthogonal reducing complement.

For $N \ge 0$ define block Hankel matrices

$$H_N = [B_{i+j}]_{i,j=0}^N, \qquad G_N = [B_{i+j+1}]_{i,j=0}^N. \tag{5}$$

Given $c = (c_0, \dots, c_N) \in (\mathbb{C}^d)^{N+1}$, set

$$W_N c = \sum_{j=0}^N T^j V c_j, \qquad q_c(x) = \sum_{j=0}^N x^j c_j. \tag{6}$$

Then

$$\begin{aligned} c^* H_N c &= \|W_N c\|^2, \\ c^* G_N c &= \langle W_N c, T W_N c\rangle, \\ c^*(\theta H_N - G_N) c &= \int_0^1 (\theta - x)\, q_c(x)^* d\Sigma(x) q_c(x). \end{aligned} \tag{7}$$

**Proposition 2.1** (Ritz/localizer dictionary). *Let*

$$\rho_N = \sup_{c \notin \ker H_N} \frac{c^* G_N c}{c^* H_N c}. \tag{8}$$

*Then $\rho_N$ is the largest Ritz value of $T$ on $\mathcal{K}_N = \operatorname{ran} W_N$, and*

$$0 \le \rho_N \le \rho_{N+1} \le \rho_V, \qquad \rho_N \uparrow \rho_V. \tag{9}$$

*Moreover, $\theta H_N - G_N \not\succeq 0$ is a rigorous finite falsifier of $\rho_V \le \theta$, whereas $\theta H_N - G_N \succeq 0$ at one finite $N$ is not in general a certificate of that support bound.*

*Proof.* The first two identities in Eq. (7) identify the generalized Rayleigh quotient with that of $T$ on $\mathcal{K}_N$. The inclusions $\mathcal{K}_N \subseteq \mathcal{K}_{N+1}$ give monotonicity. Their union is dense in $\mathcal{K}_V$, so the variational suprema converge to $\|T|_{\mathcal{K}_V}\|$. If a localizer has a negative direction, Eq. (7) is incompatible with spectral support in $[0, \theta]$. Positivity on polynomials of bounded degree does not exclude spectral weight detectable only at larger degree. □

This proposition separates a common ambiguity. A Ritz value is a lower bound on the visible transfer edge and hence an *upper* bound on the visible mass $-\log \rho_V$. A negative localizer falsifies an overstated mass. Neither becomes a full-gap theorem until probe completeness is supplied.

## 3 The exact terminal branch

Finite exact recovery occurs when the moment problem becomes degenerate. This is the operator form of flat propagation in truncated moment theory and of exact termination in block Lanczos [9, 10, 11, 4].

**Theorem 3.1** (Flat Krylov termination)**.** *If*

$$\operatorname{rank} H_N = \operatorname{rank} H_{N+1} < \infty, \tag{10}$$

*then $\mathcal{K}_N = \mathcal{K}_{N+1}$, the space $\mathcal{K}_N$ is invariant under $T$, and*

$$\mathcal{K}_V = \mathcal{K}_N, \qquad \rho_V = \rho_N. \tag{11}$$

*The finite generalized pencil $(G_N, H_N)$ on the quotient by $\ker H_N$ recovers the complete visible spectrum, with multiplicities in the cyclic representation. If in addition $\mathcal{K}_V = \mathcal{K}$ and $T = e^{-aH}$ on the excitation space, the recovered edge certifies the full gap $m = -a^{-1} \log \rho_N$ (with the usual extended-value convention at $\rho_N = 0$).*

*Proof.* Since $H_N = W_N^* W_N$, $\operatorname{rank} H_N = \dim \mathcal{K}_N$. The nesting $\mathcal{K}_N \subseteq \mathcal{K}_{N+1}$ and Eq. (10) therefore imply equality. But $T\mathcal{K}_N \subseteq \mathcal{K}_{N+1}$, hence $T\mathcal{K}_N \subseteq \mathcal{K}_N$. All further Krylov spaces equal $\mathcal{K}_N$, so their closure is $\mathcal{K}_V = \mathcal{K}_N$. The compressed operator is now the exact restriction rather than an approximation, and the quotient pencil represents it in the nonorthogonal Krylov coordinates. □

**Example 3.2** (Exact rational terminal certificate)**.** Take

$$T = \operatorname{diag}(1/5, 1/2, 4/5), \qquad V = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{bmatrix}.$$

The tracked exact-arithmetic artifact constructs $B_0, \ldots, B_5$ and finds

$$\operatorname{rank} H_0 = 2, \qquad \operatorname{rank} H_1 = \operatorname{rank} H_2 = 3.$$

On a rational quotient basis the pencil determinant is

$$-\frac{9}{5000}(2\lambda - 1)(5\lambda - 4)(5\lambda - 1), \tag{12}$$

recovering $\{1/5, 1/2, 4/5\}$ exactly and hence the visible mass $-\log(4/5) = \log(5/4)$. No floating-point rank decision enters this example.

Flatness is powerful but singular. With noisy data, exact rank is destroyed; thresholding small singular values introduces a model choice. The next section shows that this is not merely a numerical inconvenience.

# 4 A finite-noise no-go theorem

The obstruction already appears for a scalar normalized spectral measure. Its matrix analogue follows by adding an identity-valued atom.

**Theorem 4.1** (Hidden-atom indistinguishability)**.** *Let $\mu$ be a probability measure on $[0,1]$ with moments $m_n = \int x^n \, d\mu(x)$. For $0 < w \le 1$ and $x_* \in [0,1]$ define*

$$\mu_{w,x_*} = (1-w)\mu + w\delta_{x_*}, \qquad m_n^{(w,x_*)} = (1-w)m_n + wx_*^n. \tag{13}$$

*Then, for every $n \ge 0$,*

$$|m_n^{(w,x_*)} - m_n| = w|x_*^n - m_n| \le w. \tag{14}$$

*If $x_* > \sup \operatorname{supp} \mu$, then $\mu_{w,x_*}$ has transfer edge $x_*$ and mass gap $-\log x_*$. In particular $x_* = 1$ gives zero gap.*

*If $\Sigma$ is an $I_d$-normalized positive matrix measure and $B_n = \int x^n d\Sigma(x)$, the same conclusion holds in operator norm for*

$$\Sigma_{w,x_*} = (1-w)\Sigma + wI_d\delta_{x_*}, \qquad B_n^{(w,x_*)} = (1-w)B_n + wx_*^n I_d, \tag{15}$$

*because $0 \preceq B_n \preceq I_d$ and $\|B_n^{(w,x_*)} - B_n\| \le w$.*

*Proof.* Since $0 \le x^n \le 1$ on $[0,1]$, one has $0 \le m_n \le 1$. Equation (14) follows by direct subtraction. If $x_*$ lies above the old support, it is the new support edge. The matrix statement is identical after replacing scalar order by Loewner order. □

**Corollary 4.2** (Near-critical obstruction with no zero mode)**.** *Let $\mu$ be supported in $[0,\theta]$ with $\theta < 1$. For every finite moment window, every componentwise radius $\varepsilon > 0$, and every*

$$0 < g < -\log\theta, \tag{16}$$

*the error ball contains a measure with transfer edge $e^{-g} < 1$ and strictly positive mass gap exactly $g$. Hence the admissible gaps in every such ball have infimum zero even after excluding an atom at one.*

*Proof.* Take $x_* = e^{-g} > \theta$ and $0 < w \le \min\{\varepsilon, 1\}$. Then Theorem 4.1 applies, the new edge is $x_* < 1$, and its mass gap is $-\log x_* = g$. Letting $g \downarrow 0$ proves the last statement. □

**Corollary 4.3** (Algorithm-independent impossibility)**.** *No estimator whose only input is a finite positive-moment window with a nonzero componentwise error bar can return a uniformly valid strictly positive gap lower bound over all positive measures whose nonvacuum transfer edge is strictly below one.*

This conclusion applies equally to nonlinear fits, Bayesian reconstructions, GEVP, Prony, Lanczos, Padé, semidefinite programs, and machine-learning estimators. It is a property of the information set, not of the solver. A denoising projection onto the positive Hankel cone can improve consistency and variance [8]; it cannot exclude the hidden measures in Theorems 4.1 and 4.2 without adding a visibility prior.

The theorem is deliberately worst-case. It does not say that small atoms are common or that gap estimation is useless. It says that a rigorous positive certificate must disclose what rules them out. The near-critical corollary is the physically relevant form when uniqueness of the vacuum is already known: the obstruction does not rely on an additional exactly zero-energy state.

## 5 A distribution-level Gaussian lower bound

Moment closeness alone does not quantify the difficulty of a statistical experiment. We now show that the hidden-atom mechanism persists at the level of the complete data law. Fix a finite degree $N$ and put

$$\phi_N(x) = (1, x, \dots, x^N, x^{1/2}, x^{3/2}, \dots, x^{N+1/2})^\mathsf{T}. \tag{17}$$

For a probability measure $\mu$ define

$$K_\mu = \int_0^1 \phi_N(x)\phi_N(x)^\mathsf{T}\, d\mu(x), \qquad z_1, \dots, z_n \overset{\text{iid}}{\sim} \mathcal{N}(0, K_\mu). \tag{18}$$

This is the finite-dimensional law obtained by applying a real isonormal Gaussian process to the vectors $T^jV$ and $T^{j+1/2}V$. Equivalently one may regard the Euclidean sampling step as half of the transfer step. No nonexistent "standard Gaussian vector" in an infinite-dimensional Hilbert space is assumed.

The two principal covariance blocks in (18) are precisely $H_N = [m_{i+j}]$ and $G_N = [m_{i+j+1}]$. Thus the sketch retains the localizer while providing a concrete finite-sample experiment.

**Theorem 5.1** (Near-critical Le Cam bound)**.** *Let $\mu_0$ be supported in $[0, \theta]$ and suppose $K_0 := K_{\mu_0} \succ 0$. Fix $x_* \in (\theta, 1]$, put $v_* := \phi_N(x_*)$, and*

$$\mu_{w,x_*} = (1-w)\mu_0 + w\delta_{x_*}, \qquad K_{w,x_*} = (1-w)K_0 + wv_*v_*^\mathsf{T}, \tag{19}$$

*and define*

$$A_{x_*} = K_0^{-1/2}(v_*v_*^\mathsf{T} - K_0)K_0^{-1/2}. \tag{20}$$

*If $w\|A_{x_*}\|_{\mathrm{op}} \le 1/2$, then, for $\mathsf{P}_0 = \mathcal{N}(0, K_0)$ and $\mathsf{P}_{w,x_*} = \mathcal{N}(0, K_{w,x_*})$,*

$$D(\mathsf{P}_{w,x_*}^{\otimes n} \| \mathsf{P}_0^{\otimes n}) \le \frac{nw^2}{2}\|A_{x_*}\|_F^2, \tag{21}$$

$$\|\mathsf{P}_{w,x_*}^{\otimes n} - \mathsf{P}_0^{\otimes n}\|_{\mathrm{TV}} \le \frac{w\sqrt{n}}{2}\|A_{x_*}\|_F. \tag{22}$$

*Consequently every test $\varphi \in [0,1]$ satisfying $\mathbb{E}_0\varphi \le \alpha$ obeys*

$$\mathbb{E}_{w,x_*}\varphi \le \alpha + \frac{w\sqrt{n}}{2}\|A_{x_*}\|_F. \tag{23}$$

*Here $\mu_0$ obeys the claimed positive edge bound, whereas $\mu_{w,x_*}$ has edge $x_*$ and mass gap $-\log x_*$. The case $x_* < 1$ introduces no exactly zero-energy excitation.*

*Proof.* Equation (19) gives $K_{w,x_*} = K_0^{1/2}(I + wA_{x_*})K_0^{1/2}$. If $t$ is an eigenvalue of $wA_{x_*}$, then $|t| \le 1/2$ and $t - \log(1+t) \le t^2$. The Gaussian relative-entropy formula therefore yields

$$D(\mathsf{P}_{w,x_*} \| \mathsf{P}_0) = \frac{1}{2}\sum_j \{t_j - \log(1+t_j)\} \le \frac{w^2}{2}\|A_{x_*}\|_F^2.$$

Relative entropy tensorizes, proving (21). Pinsker's inequality gives (22), and the variational definition of total variation gives (23); see, e.g., [20] for these standard testing reductions. □

The preceding inequality has an exact local counterpart. Put $d_N = 2N + 2$ and

$$a_x = K_0^{-1/2}\phi_N(x), \qquad \beta_x = \|a_x\|^2, \qquad I_x = \frac{1}{2}\big\{(\beta_x - 1)^2 + d_N - 1\big\}. \tag{24}$$

**Theorem 5.2** (Exact hidden-atom likelihood and local power)**.** *In the experiment of Theorem 5.1, the whitened perturbation is $A_x = a_x a_x^{\mathsf{T}} - I_{d_N}$. Its eigenvalues are $\beta_x - 1$ once and $-1$ with multiplicity $d_N - 1$. Consequently, with*

$$\tau_w = 1 + w(\beta_x - 1), \qquad \rho_w = 1 - w, \tag{25}$$

*the one-sample relative entropy is exactly*

$$D(\mathsf{P}_{w,x} \| \mathsf{P}_0) = \frac{1}{2}\left[(d_N - 1)\{-w - \log(1-w)\} + w(\beta_x - 1) - \log\{1 + w(\beta_x - 1)\}\right]. \tag{26}$$

*After whitening and rotating the spike direction, the log likelihood ratio for $n$ observations is*

$$\Lambda_n = -\frac{n}{2}\{(d_N - 1)\log\rho_w + \log\tau_w\} + \frac{1}{2}(1 - \tau_w^{-1})U + \frac{1}{2}(1 - \rho_w^{-1})V, \tag{27}$$

*where under the null $U \sim \chi^2_n$ and $V \sim \chi^2_{n(d_N-1)}$ are independent. Under the alternative, $U/\tau_w$ and $V/\rho_w$ have those same independent chi-square laws.*

*For $w_n = h/\sqrt{n}$ with fixed $h \geq 0$,*

$$\Lambda_n = h\Delta_n - \frac{1}{2}h^2 I_x + o_{\mathsf{P}_0}(1), \qquad \Delta_n \Longrightarrow \mathcal{N}(0, I_x). \tag{28}$$

*Thus the asymptotic power envelope among level-$\alpha$ tests of this fixed, known path is*

$$\pi_x(h) = 1 - \Phi\left(z_{1-\alpha} - h\sqrt{I_x}\right), \tag{29}$$

*and it is attained by the likelihood-ratio, equivalently local score, test. More generally, $w_n\sqrt{n} \to 0$ gives limiting power at most $\alpha$; $w_n\sqrt{n} \to h \in (0,\infty)$ gives (29); and $w_n\sqrt{n} \to \infty$, $w_n \to 0$, admits a consistent score test.*

*Proof.* The rank-one form of $A_x$ gives the stated spectrum, hence the eigenvalues (25) of $I + wA_x$. Substitution in the Gaussian entropy formula proves (26). In whitened coordinates, decompose each observation into its component parallel to $a_x$ and its orthogonal component. The sums of their squared norms are the independent chi-square variables in (27); rescaling gives their alternative laws.

For completeness, the central sequence is

$$\Delta_n = \frac{1}{2\sqrt{n}}\sum_{r=1}^{n}\{Z_r^{\mathsf{T}} A_x Z_r - \operatorname{tr} A_x\}, \qquad Z_r \overset{\text{iid}}{\sim} \mathcal{N}(0, I_{d_N}).$$

Its variance is $\operatorname{tr}(A_x^2)/2 = I_x$. A second-order Taylor expansion of (27) and the central limit theorem give (28). The Gaussian-shift Neyman–Pearson lemma then gives (29) and the three regimes; this is the standard finite-dimensional LAN argument [21]. □

The nonsingularity hypothesis is essential: adding an atom can open a new covariance direction and then be detectable immediately. It does not weaken the minimax conclusion. Measures on $2N+2$ distinct points in $(0,\theta)$ make $K_0$ positive definite because, after writing $y = \sqrt{x}$, the reordered features in (17) are the full Vandermonde family $1, y, \ldots, y^{2N+1}$. Moreover $x \mapsto A_x$ is continuous on every compact subinterval of $(\theta, 1]$. Hence for any $x_n \uparrow 1$ and $w_n = o(n^{-1/2})$, the local condition eventually holds and every level-$\alpha$ test has power at most $\alpha + o(1)$ against alternatives with strictly positive gaps $-\log x_n \downarrow 0$. This is an existential minimax obstruction for a fixed feature window and base covariance, not a pointwise or jointly uniform rate over all gapped covariances and growing degrees.

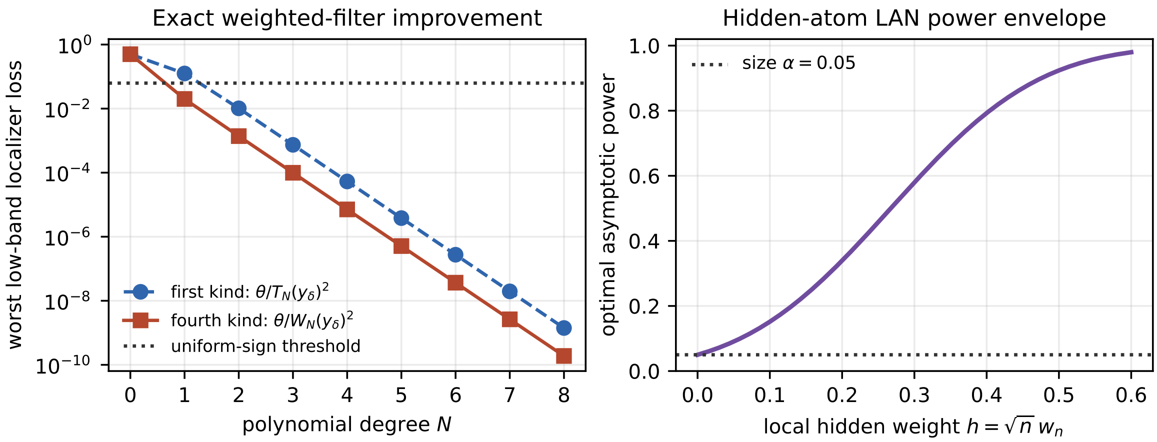


Figure 1: Two exact boundaries. Left: the fourth-kind weighted-localizer loss $\theta/\mathsf{W}_N(y_\delta)^2$ lies below the former first-kind uniform envelope and crosses the uniform-sign threshold one degree earlier in the displayed rational fixture. Right: the fixed-path LAN envelope (29); the local parameter is $h = \sqrt{n}\,w$, not the spectral gap $-\log x$.

**Corollary 5.3** (Weight, not the vanishing gap, sets the local scale)**.** *Let $x_n = e^{-g_n} \uparrow 1$ with $g_n \downarrow 0$ and keep $N, K_0$ fixed. Then $A_{x_n} \to A_1$ and $I_{x_n} \to I_1 > 0$. The three regimes of Theorem 5.2 remain valid with $I_x$ replaced by $I_1$. Hence, on a fixed feature window, an arbitrarily small positive gap does not by itself set the testing boundary: the hidden spectral weight is detectable at the $n^{-1/2}$ scale.*

This local statement is for a specified rank-one path. It is not a scan over an unknown atom location, a uniform theorem over ill-conditioned $K_0$, or a growing-degree minimax result. LAN and spiked-covariance power calculations are classical [22]; the contribution here is the exact reduction of the hidden spectral atom to this experiment and the resulting separation of gap size from visible weight.

## 6 Visibility restores finite falsification power

We now add a transparent quantitative visibility datum. Work first with a scalar probability measure $\nu$ on $[0, 1]$. Fix a claimed edge $0 < \theta < 1$, a resolution $0 < \delta \le 1 - \theta$, and a visibility floor $0 < \gamma \le 1$.

**Definition 6.1** (Visibility alternative)**.** The pair $(\delta, \gamma)$ is visible beyond $\theta$ if every admissible measure whose edge reaches $\theta + \delta$ satisfies

$$\nu([\theta + \delta, 1]) \ge \gamma. \tag{30}$$

The assumption is falsifiable and model-dependent. It can come from a frame bound for an operator family, a symmetry-resolved overlap estimate, a source theorem, or an experimentally calibrated minimum spectral weight. It cannot be manufactured from the same noisy moments without circularity, by Theorem 4.1.

Let $\mathsf{T}_N$ and $\mathsf{W}_N$ denote the Chebyshev polynomials of the first and fourth kinds [23]. The latter is characterized by

$$\mathsf{W}_N(\cos t) = \frac{\sin((N + \frac{1}{2})t)}{\sin(t/2)}, \qquad \mathsf{W}_N(\cosh \eta) = \frac{\sinh((N + \frac{1}{2})\eta)}{\sinh(\eta/2)}. \tag{31}$$

Set

$$y_\delta = 1 + \frac{2\delta}{\theta}, \qquad q_N(x) = \frac{\mathsf{W}_N(2x/\theta - 1)}{\mathsf{W}_N(y_\delta)}, \qquad e_N = \frac{\theta}{\mathsf{W}_N(y_\delta)^2}. \tag{32}$$

**Theorem 6.2** (Exact visibility-conditioned localizer boundary)**.** *For every probability measure satisfying* (30)*, define*

$$Q_N = \int_0^1 (\theta - x) q_N(x)^2 \, d\nu(x). \tag{33}$$

*Then*

$$Q_N \le e_N(1-\gamma) - \delta\gamma. \tag{34}$$

*The right-hand side is the exact supremum over this visibility class. Consequently $q_N$ uniformly falsifies the claimed support if and only if*

$$\mathsf{W}_N(y_\delta)^2 > \frac{\theta(1-\gamma)}{\delta\gamma}. \tag{35}$$

*Within this normalized degree-$N$ polynomial-localizer class the minimal depth is therefore exactly*

$$N_{\min} = \min\left\{ N \ge 0 : \mathsf{W}_N(y_\delta)^2 > \frac{\theta(1-\gamma)}{\delta\gamma} \right\}. \tag{36}$$

*Proof.* On $[\theta+\delta, 1]$, $q_N \ge 1$ because $\mathsf{W}_N$ is positive and increasing on $[1,\infty)$. On the low band, write $x = \theta(1+\cos t)/2$. Equation (31) gives the exact identity

$$(\theta - x) q_N(x)^2 = e_N \sin^2((N + \tfrac{1}{2})t) \le e_N. \tag{37}$$

If $l, s, h$ are the masses of $[0,\theta]$, $(\theta, \theta+\delta)$, and $[\theta+\delta, 1]$, then the three contributions to $Q_N$ are at most $e_N l$, 0, and $-\delta h$. Since $h \ge \gamma$ and $l \le 1-h$, (34) follows.

For $t_j = (2j+1)\pi/(2N+1)$, $j = 0, \dots, N$, put $x_j = \theta(1+\cos t_j)/2$. The two-atom measure

$$\nu_{*,j} = (1-\gamma)\delta_{x_j} + \gamma\delta_{\theta+\delta} \tag{38}$$

attains equality in (34) by (37) and $q_N(\theta+\delta) = 1$. This proves both sharpness and the equivalence (35). □

**Proposition 6.3** (Weighted polynomial minimax)**.** *Among all real polynomials $r$ of degree at most $N$ satisfying $r(\theta+\delta) = 1$,*

$$\inf_r \max_{0 \le x \le \theta} \sqrt{\theta - x}\, |r(x)| = \frac{\sqrt{\theta}}{\mathsf{W}_N(y_\delta)}. \tag{39}$$

*The filter $q_N$ attains the infimum and is the unique minimizer.*

*Proof.* For $q_N$, equality follows from (37). At the $N+1$ points $x_j$ used above, the weighted values alternate between the two extremes. If a normalized competitor had a strictly smaller maximum, the sign of $q_N - r$ would alternate at these points, giving at least $N$ roots in $[0,\theta]$. It also vanishes at $\theta+\delta$, which is impossible for a nonzero polynomial of degree at most $N$. The standard non-strict alternation argument gives uniqueness. This is the weighted Chebyshev alternation theorem specialized to the localizer weight $\sqrt{\theta - x}$. □

**Corollary 6.4** (Exact visibility minimax over the filter class)**.** *Let*

$$\mathcal{P}_N = \{ r \in \mathbb{R}[x] : \deg r \le N, \ r(\theta+\delta) = 1 \}$$

*and define the worst-case localizer risk*

$$\mathcal{R}_N(r) = \sup_{\substack{\nu \text{ probability on } [0,1] \\ \nu([\theta+\delta,1]) \ge \gamma}} \int (\theta - x) r(x)^2 \, d\nu(x).$$

*Then*

$$\inf_{r \in \mathcal{P}_N} \mathcal{R}_N(r) = e_N(1-\gamma) - \delta\gamma, \tag{40}$$

*and $q_N$ is the unique minimizer. Hence a uniformly negative localizer exists in $\mathcal{P}_N$ if and only if* (35) *holds.*

*Proof.* For any $r \in \mathcal{P}_N$, put $L(r) = \max_{[0,\theta]} \sqrt{\theta - x}|r(x)|$. A measure placing mass $1-\gamma$ at a maximizer of $L(r)$ and mass $\gamma$ at $\theta + \delta$ shows

$$\mathcal{R}_N(r) \geq (1-\gamma)L(r)^2 - \delta\gamma.$$

By Proposition 6.3, $L(r)^2 \geq e_N$, with equality only for $q_N$. The exact upper bound for $q_N$ is Theorem 6.2, proving (40) and uniqueness. □

The distinction from ordinary stop-band minimization matters. Writing $y_\delta = \cosh\eta$,

$$\mathsf{W}_N(\cosh\eta) = 1 + 2\sum_{k=1}^{N}\cosh(k\eta) > \mathsf{T}_N(\cosh\eta) \qquad (N \geq 1). \tag{41}$$

Thus the exact fourth-kind localizer loss $e_N$ is strictly smaller than the former first-kind uniform envelope $\theta/\mathsf{T}_N(y_\delta)^2$. The theorem does not assert optimality over nonlinear estimators, matrix-valued filters, or alternative observation models; it exactly solves the scalar normalized polynomial localizer problem generated by a finite moment window.

**Corollary 6.5** (Matrix-valued multichannel witness)**.** *Let $\Sigma([0,1]) = I_d$ and set $\nu = d^{-1}\operatorname{tr}\Sigma$. If $\nu([\theta+\delta,1]) \geq \gamma$ and Eq. (35) holds, then*

$$\theta H_N - G_N \not\succeq 0. \tag{42}$$

*Hence the matrix correlator window contains a finite direction that falsifies the claimed edge.*

*Proof.* Let $a_0,\dots,a_N$ be the coefficients of $q_N$ and, for each standard basis vector $e_\ell \in \mathbb{C}^d$, take the block coefficient vector $c^{(\ell)} = (a_0 e_\ell,\dots,a_N e_\ell)$. Then

$$\frac{1}{d}\sum_{\ell=1}^{d}(c^{(\ell)})^*(\theta H_N - G_N)c^{(\ell)} = Q_N < 0.$$

At least one summand is negative. □

The degree grows only logarithmically with $1/\gamma$ at fixed resolution. For small $\delta/\theta$, however, $\operatorname{arcosh}(1+2\delta/\theta) \sim 2\sqrt{\delta/\theta}$, exposing the expected edge-resolution cost. This is closely related to the Chebyshev mechanism in Kaniel–Paige–Saad convergence bounds, but here the output is a one-sided moment localizer with an explicit visibility premise, not merely an eigenvalue convergence estimate [24, 4].

## 6.1 Certified moment error

Write $q_N(x) = \sum_{j=0}^{N} a_j x^j$ and suppose measured moments satisfy $|\widetilde{m}_k - m_k| \leq \varepsilon$ for $0 \leq k \leq 2N+1$. Form the measured localizer $\widetilde{Q}_N$ by expanding Eq. (33) in those moments.

**Proposition 6.6** (Noise-safe sign test)**.** *The deterministic error obeys*

$$|\widetilde{Q}_N - Q_N| \leq (1+\theta)\varepsilon\left(\sum_{j=0}^{N}|a_j|\right)^2 =: R_N. \tag{43}$$

*Therefore*

$$\widetilde{Q}_N + R_N < 0 \tag{44}$$

*is a rigorous noisy-data falsifier of support in $[0,\theta]$. Under the visibility alternative, the sufficient worst-case condition*

$$e_N(1-\gamma) - \delta\gamma + 2R_N < 0 \tag{45}$$

*guarantees detection for every admissible error realization.*

*Proof.* If $q_N^2 = \sum_{k=0}^{2N} b_k x^k$, then $Q_N = \theta \sum b_k m_k - \sum b_k m_{k+1}$. Hence the error is at most $(1+\theta)\varepsilon \sum |b_k|$. Convolution gives $\sum |b_k| \le (\sum |a_j|)^2$. Under the visibility alternative, $\widetilde{Q}_N + R_N \le Q_N + 2R_N \le e_N(1-\gamma) - \delta\gamma + 2R_N$, which proves the stated worst-case condition. □

Power-basis coefficients can grow rapidly. Thus increasing $N$ is not monotonically beneficial at fixed componentwise precision: spectral localization improves while the rigorous error radius may worsen. The tracked certificates therefore evaluate the coefficient norm as well as the weighted population loss. Higher precision or a stable recurrence-based error model can extend the useful window, but must be declared explicitly.

# 7 An honest finite-sample Wishart test

The deterministic radius above ignores correlations. Under a declared Gaussian-sketch model, the full quadratic dependence can instead be handled exactly. Fix the fourth-kind polynomial $q_N$ before observing the data and set

$$u = q_N(T)V, \qquad v = Tq_N(T)V. \tag{46}$$

For clarity we state the scalar real case. Applying independent isonormal Gaussian processes $g_r$ to the finite span of $u, v$ gives

$$\begin{aligned} y_r &= (g_r(u), g_r(v))^\mathsf{T} \sim \mathcal{N}(0, K), \\ K &= \begin{bmatrix} P & X \\ X & R_2 \end{bmatrix}, \\ P &= \int q_N^2 d\nu, \qquad X = \int x q_N^2 d\nu, \qquad R_2 = \int x^2 q_N^2 d\nu. \end{aligned} \tag{47}$$

Put $S = n^{-1} \sum_r y_r y_r^\mathsf{T}$ and choose a significance level $0 < \alpha_{\mathrm{sig}} < 1$. Define

$$\eta = \frac{\sqrt{2} + \sqrt{2\log(2/\alpha_{\mathrm{sig}})}}{\sqrt{n}} < 1, \quad c_- = (1+\eta)^{-2}, \quad c_+ = (1-\eta)^{-2}, \tag{48}$$

and the simultaneous confidence interval

$$\mathcal{C}(S) = \{Q \in \mathbb{S}^2 : c_- S \preceq Q \preceq c_+ S\}. \tag{49}$$

We reject the claimed edge when no $Q \in \mathcal{C}(S)$ satisfies

$$\theta Q_{00} - Q_{01} \ge 0. \tag{50}$$

Rejection is a lower-confidence statement about the *visible transfer edge*, equivalently an upper-confidence statement about the visible mass. Non-rejection says only that the confidence band intersects the support-null constraint; it is not a positive lower-confidence certificate for the gap.

**Theorem 7.1** (Exact level and visibility-to-sample power)**.** *The test (50) has type-I error at most $\alpha_{\mathrm{sig}}$ for every Gaussian covariance generated by a measure supported in $[0, \theta]$, including singular covariances.*

*Assume now (30). Write*

$$\mu_N = \delta\gamma - e_N(1-\gamma). \tag{51}$$

*If $\mu_N > 0$, define*

$$r_N = \frac{\mu_N}{1+\theta}. \tag{52}$$

*Then the rejection probability is at least $1 - \alpha_{\mathrm{sig}}$ whenever*

$$n > \left(\sqrt{2} + \sqrt{2\log(2/\alpha_{\mathrm{sig}})}\right)^2 \frac{\left(\sqrt{1+r_N}+1\right)^4}{r_N^2}. \tag{53}$$

*Proof.* Let $Y$ be an $n \times 2$ standard Gaussian matrix. Writing the data matrix as $YK^{1/2}$, the Gaussian extreme-singular-value bound [19] gives, with probability at least $1 - \alpha_{\mathrm{sig}}$,

$$(1-\eta)^2 K \preceq S \preceq (1+\eta)^2 K.$$

Congruence proves the statement even when $K$ is singular. On this event the true $K$ belongs to (49); under the support null it also satisfies (50). Rejection is therefore impossible on the coverage event, proving the level claim. Filter selection used no data, so no sample split or post-selection correction is needed.

For power put

$$m = X - \theta P = \int (x-\theta) q_N(x)^2 d\nu(x), \quad C = \begin{bmatrix} \theta & -1/2 \\ -1/2 & 0 \end{bmatrix}, \quad w = \|K^{1/2} C K^{1/2}\|_*.$$

Let $a = K^{1/2} e_0$ and $b = K^{1/2} e_1$. The triangle inequality gives

$$w \le \theta \|a\|^2 + \left\| \frac{ab^\mathsf{T} + ba^\mathsf{T}}{2} \right\|_* \le \theta P + \sqrt{PR_2} \le (1+\theta)P, \tag{54}$$

where $R_2 \le P$ follows from $0 \le x \le 1$. Let $l, s, h$ be the masses of $[0, \theta]$, $(\theta, \theta+\delta)$, and $[\theta+\delta, 1]$. On $[0, \theta+\delta]$ one has $|q_N| \le 1$: on the low band use $|\mathsf{W}_N(y)| \le 2N+1 < \mathsf{W}_N(y_\delta)$, and on the middle band use monotonicity. On the high band $q_N \ge 1$. If $H$ denotes the positive contribution to $m$ from the high region and $L$ the magnitude of its negative low-region contribution, then

$$m \ge \delta h - e_N l \ge \mu_N, \qquad H \le m + e_N l.$$

The filtered norm has the regional bound

$$P \le l + s + \frac{H}{\delta} \le 1 - \gamma + \frac{m + e_N(1-\gamma)}{\delta}.$$

Combining this with (54), the ratio $m/w$ is bounded below by an increasing function of $m$. Substituting $m = \mu_N$ makes the bracket on the right exactly one, and hence $m/w \ge \mu_N/(1+\theta) = r_N$.

Every $Q \in \mathcal{C}(S)$ on the same coverage event obeys

$$\left(\frac{1-\eta}{1+\eta}\right)^2 K \preceq Q \preceq \left(\frac{1+\eta}{1-\eta}\right)^2 K.$$

Consequently $\mathrm{tr}(CQ) \le -m + \delta_+(\eta) w$, where $\delta_+(\eta) = ((1+\eta)/(1-\eta))^2 - 1$. If $\delta_+(\eta) < r_N$, every covariance in the band violates (50). Solving this inequality for $\eta$ and substituting (48) gives (53). □

**Corollary 7.2** (Closed depth–sample resource frontier)**.** *Fix a desired fractional margin $0 < \sigma < 1$ and choose the least $N$ such that*

$$\mathsf{W}_N(y_\delta)^2 \ge \frac{\theta(1-\gamma)}{(1-\sigma)\delta\gamma}. \tag{55}$$

*Then $\mu_N \ge \sigma\delta\gamma$. The precommitted two-coordinate Wishart test has rejection probability at least $1 - \alpha_{\mathrm{sig}}$ whenever*

$$n > \left(\sqrt{2} + \sqrt{2\log(2/\alpha_{\mathrm{sig}})}\right)^2 \frac{\left(\sqrt{1+r_\sigma}+1\right)^4}{r_\sigma^2}, \qquad r_\sigma = \frac{\sigma\delta\gamma}{1+\theta}. \tag{56}$$

*Thus filter depth is $O(\log(1/\gamma))$ at fixed resolution while the displayed Gaussian sample guarantee is $O(\gamma^{-2})$. This is an explicit achievable resource curve, not a joint minimax theorem when $N$, $K_0$, or the atom location vary with $n$.*

The dimension 2 in (53) is substantive: a precommitted polynomial compresses every raw sketch to the two filtered coordinates before forming the Wishart band. If the polynomial or degree is chosen adaptively from the same data, one may instead band the complete $(N+2)$-dimensional raw feature covariance of $(g(V), g(TV), \dots, g(T^{N+1}V))$. Its entries contain every moment through order $2N+2$, so $\theta H_N - G_N \succeq 0$ is a linear constraint on that covariance. The same coverage argument remains valid without splitting, but $\sqrt{2}$ in (53) is replaced by $\sqrt{N+2}$ and the power condition is witness-specific. The Gaussian, independence, and known-zero-mean hypotheses are essential; this is not an exact theorem for raw Markov-chain Monte Carlo averages.

# 8 A distribution-free bounded-readout test

Gaussianity and known centering are not required if the filtered readouts are independent and bounded. The following companion theorem also permits a finite filter bank to be searched on the same sample. Let $p^{(1)}, \dots, p^{(M)}$ be fixed before data collection and, for each $j$, let the raw two-coordinate readout $Y_s^{(j)} \in \mathbb{R}^2$ have arbitrary unknown mean and covariance

$$K_j = \begin{bmatrix} P_j & X_j \\ X_j & R_j \end{bmatrix}, \quad P_j = \int (p^{(j)})^2 d\nu, \quad X_j = \int x (p^{(j)})^2 d\nu. \tag{57}$$

The vectors $(Y_s^{(1)}, \dots, Y_s^{(M)})$ may be dependent across filters but are iid across $s$.

For a raw coordinate bound $B > 0$ define the exact range length

$$L_\theta(B) = B^2 \begin{cases} 4, & 0 < \theta \le \frac{1}{2}, \\ 2(1+\theta) + (2\theta)^{-1}, & \frac{1}{2} \le \theta < 1. \end{cases} \tag{58}$$

**Theorem 8.1** (Bounded unknown-mean filter bank). *Suppose $|Y_{s,k}^{(j)}| \le B_j$ almost surely for all $j$, $k \in \{0,1\}$, and $s = 1, \dots, 2n$. Form paired differences and localizer observations*

$$Z_r^{(j)} = \frac{Y_{2r}^{(j)} - Y_{2r-1}^{(j)}}{\sqrt{2}}, \quad W_r^{(j)} = \theta (Z_{r,0}^{(j)})^2 - Z_{r,0}^{(j)} Z_{r,1}^{(j)}, \quad \bar{W}_j = \frac{1}{n} \sum_{r=1}^{n} W_r^{(j)}. \tag{59}$$

*For $0 < \alpha < 1$, reject the support claim $\operatorname{supp} \nu \subseteq [0, \theta]$ when, for at least one $j$,*

$$\bar{W}_j + L_\theta(B_j) \sqrt{\frac{\log(M/\alpha)}{2n}} < 0. \tag{60}$$

*This test has type-I error at most $\alpha$ for every iid law satisfying the stated covariance and bounds, without Gaussianity or a known mean.*

*If for some $j_*$ the visible alternative gives $X_{j_*} - \theta P_{j_*} \ge \mu > 0$, then its rejection probability is at least $1 - \beta$ whenever*

$$n > \frac{L_\theta(B_{j_*})^2}{2\mu^2} \left( \sqrt{\log(M/\alpha)} + \sqrt{\log(1/\beta)} \right)^2. \tag{61}$$

*In particular the fourth-kind filter has $\mu = \mu_N = \delta\gamma - e_N(1-\gamma)$ whenever this number is positive.*

*Proof.* Pairing removes the unknown mean and preserves covariance:

$$\mathbb{E} Z_r^{(j)} = 0, \qquad \mathbb{E}[Z_r^{(j)} (Z_r^{(j)})^\mathsf{T}] = K_j.$$

Consequently

$$\mathbb{E} W_r^{(j)} = \theta P_j - X_j. \tag{62}$$

Under the support null this quantity is nonnegative for every polynomial filter. Since each coordinate of $Z_r^{(j)}$ is bounded by $M_j = \sqrt{2}B_j$, the quadratic function $q(a,b) = \theta a^2 - ab$ on $[-M_j, M_j]^2$ has maximum $M_j^2(1+\theta)$ and minimum

$$\min q = \begin{cases} M_j^2(\theta - 1), & \theta \le \frac{1}{2}, \\ -M_j^2/(4\theta), & \theta \ge \frac{1}{2}. \end{cases} \tag{63}$$

Their difference is exactly (58). Hoeffding's one-sided inequality [15] and a union bound over the $M$ filters prove the level statement; no independence across filters is used.

Under the alternative, (62) is at most $-\mu$ for $j_*$. A second one-sided Hoeffding bound shows that failure to reject has probability at most $\beta$ as soon as the population margin exceeds the sum of the $\alpha/M$ and $\beta$ deviations. Solving that inequality gives (61). The final statement is (34) with its sign reversed. □

**Example 8.2** (Exact rational calibration)**.** Take $\theta = 1/2$, $\delta = 1/4$, $\gamma = 1/5$, and $N = 2$. Then $\mathsf{T}_2(2) = 7$ while $\mathsf{W}_2(2) = 19$. The former first-kind uniform loss is $1/98$; the exact weighted loss is $e_N = 1/722$, and the sharp population margin is

$$\mu_N = \frac{1}{20} - \frac{4}{5}\frac{1}{722} = \frac{353}{7220}. \tag{64}$$

Indeed the fourth-kind filter already crosses the uniform-sign threshold at $N = 1$, whereas the first-kind envelope requires $N = 2$. At degree two, Eq. (53) requires $n > 265{,}338.8948$, hence 265,339 Gaussian sketches at $\alpha_{\text{sig}} = 0.05$. For $M = 3$, $B_j = 1$, and $\alpha = \beta = 0.05$, one has $L_{1/2}(1) = 4$ and (61) requires $n > 47{,}169.9371$, hence 47,170 paired samples or 94,340 raw observations. The standard-library certificate reproduces these numbers exactly up to the final transcendental evaluation. They are conservative worst-case guarantees, not a claim that bounded and Gaussian observations have equal information per raw sample.

The theorem is distribution-free only within its declared iid bounded-readout model. It does not turn correlated Markov-chain output into independent data. Its gain is orthogonal to the Wishart branch: it removes Gaussianity, known centering, and single-filter precommitment, while paying a range bound and two raw observations per paired sample.

# 9 A dependence-robust test from one trajectory

The iid hypothesis can be relaxed without pretending that correlated samples are independent. The price is an externally justified mixing envelope, sparse use of the trajectory, and an explicit lag-covariance bias. Let $\mathbf{Y}_s = (Y_s^{(1)}, \dots, Y_s^{(M)})$ be a strictly stationary process on a standard Borel space and put $\mathcal{F}_a^b = \sigma(\mathbf{Y}_s : a \le s \le b)$. We use the absolute regularity coefficient

$$\beta_{\text{mix}}(q) = \sup_t \mathbb{E}\left[ \sup_{A \in \mathcal{F}_{t+q}^{\infty}} \left| \mathbb{P}(A \mid \mathcal{F}_{-\infty}^t) - \mathbb{P}(A) \right| \right]. \tag{65}$$

With this convention, Berbee's coupling lemma supplies an independent copy of the future whose mismatch probability is at most $\beta_{\text{mix}}(q)$ [17].

For $q \ge 1$ and $r = 1, \dots, n$, define

$$t_r = 1 + 2q(r-1), \qquad Z_r^{(j)} = \frac{Y_{t_r+q}^{(j)} - Y_{t_r}^{(j)}}{\sqrt{2}}, \qquad W_r^{(j)} = \theta(Z_{r,0}^{(j)})^2 - Z_{r,0}^{(j)} Z_{r,1}^{(j)}. \tag{66}$$

Write $\bar{W}_j = n^{-1} \sum_r W_r^{(j)}$ and set

$$\varepsilon_{n,q} = (n-1)\beta_{\text{mix}}(q), \qquad b_j(q) = 2(1+\theta)B_j^2 \beta_{\text{mix}}(q). \tag{67}$$

**Theorem 9.1** (Bounded stationary $\beta$-mixing filter bank). *Suppose $|Y^{(j)}_{s,k}| \le B_j$ almost surely and that the contemporaneous covariance of $Y^{(j)}_s$ is the filtered matrix $K_j$ in (57). If $\varepsilon_{n,q} < \alpha$, reject the support claim $\operatorname{supp}\nu \subseteq [0,\theta]$ when, for at least one $j$,*

$$\bar{W}_j + b_j(q) + L_\theta(B_j)\sqrt{\frac{\log\bigl(M/(\alpha-\varepsilon_{n,q})\bigr)}{2n}} < 0. \tag{68}$$

*This test has type-I error at most $\alpha$ for every strictly stationary law satisfying the stated bounds and mixing coefficient.*

*Let $0 < \zeta < 1$ be a target miss probability. If for some $j_*$, $X_{j_*} - \theta P_{j_*} \ge \mu > 0$, if $\varepsilon_{n,q} < \min\{\alpha,\zeta\}$, and if $\mu > 2b_{j_*}(q)$, then the rejection probability is at least $1-\zeta$ whenever*

$$n > \frac{L_\theta(B_{j_*})^2}{2(\mu - 2b_{j_*}(q))^2}\left(\sqrt{\log\frac{M}{\alpha-\varepsilon_{n,q}}} + \sqrt{\log\frac{1}{\zeta-\varepsilon_{n,q}}}\right)^2. \tag{69}$$

*The test uses $2n$ retained readouts over a trajectory horizon $1+(2n-1)q$.*

*Proof.* Let $\Gamma_j(q) = \operatorname{Cov}(Y^{(j)}_{t_r}, Y^{(j)}_{t_r+q})$. A one-step Berbee coupling replaces $Y^{(j)}_{t_r+q}$ by an independent copy with the same law and mismatch probability at most $\beta_{\text{mix}}(q)$. Therefore every entry satisfies

$$|\Gamma_j(q)_{ab}| \le 2B_j^2\beta_{\text{mix}}(q). \tag{70}$$

Pairing still removes the stationary mean, but now

$$\mathbb{E}[Z^{(j)}_r(Z^{(j)}_r)^\intercal] = K_j - \tfrac{1}{2}\bigl(\Gamma_j(q) + \Gamma_j(q)^\intercal\bigr).$$

For $C = \left[\begin{smallmatrix} \theta & -1/2 \\ -1/2 & 0 \end{smallmatrix}\right]$, (70) gives

$$\left|\mathbb{E}W^{(j)}_r - \operatorname{tr}(CK_j)\right| \le b_j(q). \tag{71}$$

The pair blocks $U_r = (\mathbf{Y}_{t_r}, \mathbf{Y}_{t_r+q})$ are separated by $q$ time steps. Iterating Berbee's lemma constructs independent blocks $U'_1,\dots,U'_n$, each with the correct pair law, such that

$$\mathbb{P}\{(U_1,\dots,U_n) \ne (U'_1,\dots,U'_n)\} \le (n-1)\beta_{\text{mix}}(q) = \varepsilon_{n,q}. \tag{72}$$

Under the support null, $\operatorname{tr}(CK_j) = \theta P_j - X_j \ge 0$, so the coupled statistic has mean at least $-b_j(q)$. Its range length remains exactly $L_\theta(B_j)$. Hoeffding's inequality applied to the independent coupled blocks, followed by a union bound and (72), proves (68).

Under the alternative, the mean is at most $-\mu + b_{j_*}(q)$. The null-safe offset in (68) costs a second $b_{j_*}(q)$, leaving separation $\mu - 2b_{j_*}(q)$. A lower-tail Hoeffding bound with residual error budget $\zeta - \varepsilon_{n,q}$ yields (69). □

**Corollary 9.2** (Geometric absolute regularity). *Assume $\beta_{\text{mix}}(q) \le ce^{-q/\tau}$. Fix $0 < \varepsilon < \min\{\alpha,\zeta\}$ and $n \ge 2$, and take*

$$q_n = \max\left\{1, \left\lceil \tau\log\frac{c(n-1)}{\varepsilon}\right\rceil\right\}. \tag{73}$$

*Then $\varepsilon_{n,q_n} \le \varepsilon$ and*

$$b_j(q_n) \le \bar{b}_j(n) := \frac{2(1+\theta)B_j^2\varepsilon}{n-1}. \tag{74}$$

*Consequently Theorem 9.1 remains valid after replacing $\varepsilon_{n,q_n}$ by $\varepsilon$ and $b_j(q_n)$ by $\bar{b}_j(n)$. Thus a known geometric envelope gives a closed, checkable test from one stationary trajectory. Stationary geometrically ergodic Markov chains are an important source of geometrically $\beta$-mixing processes, but the constants must be justified for the chain at hand [18].*

Table 1: ANNNI scaling. $w_X, w_Z$ are squared overlaps with the first excited state. $\widehat{\rho}_6$ is the block-Krylov edge from moments through order 13. All rows refute the 35% overstated gap at degree one; the $X$-only family never refutes it through degree six.

| $L$ | $E_1 - E_0$ | $\rho = e^{-.2(E_1-E_0)}$ | $w_X$ | $w_Z$ | $\lvert\widehat{\rho}_6 - \rho\rvert$ |
|---|---|---|---|---|---|
| 6 | 2.0845863 | 0.6590755 | $3.4\,10^{-31}$ | 0.45160 | $1.9\,10^{-7}$ |
| 8 | 1.8574804 | 0.6897017 | $4.3\,10^{-32}$ | 0.39733 | $3.9\,10^{-6}$ |
| 10 | 1.7217194 | 0.7086852 | $9.8\,10^{-34}$ | 0.35390 | $2.1\,10^{-5}$ |
| 12 | 1.6335802 | 0.7212885 | $2.0\,10^{-32}$ | 0.31837 | $5.7\,10^{-5}$ |
| 14 | 1.5729200 | 0.7300925 | $2.8\,10^{-30}$ | 0.28900 | $8.6\,10^{-5}$ |
| 16 | 1.5293122 | 0.7364879 | $1.2\,10^{-31}$ | 0.26437 | $9.9\,10^{-5}$ |

**Example 9.3** (Dependent rational calibration)**.** Retain the rational margin $\mu = 353/7220$, $M = 3$, $B_j = 1$, and $\theta = 1/2$ from Eq. (64). Suppose $\beta_{\mathrm{mix}}(q) \le 2^{-q}$ and take $\alpha = \zeta = 0.05$, $\varepsilon = 0.01$. The conservative specialization of (69) first holds at $n = 50{,}177$ paired blocks. Equation (73) gives $q_n = 23$, so the test uses 100,354 retained readouts over 2,308,120 chain steps. The large horizon is not hidden in an effective-sample-size slogan: it is the explicit cost of the certified mixing envelope and worst-case bounded concentration.

This section is distribution-free only within the declared stationary, bounded, absolutely regular model. It does not estimate $\beta_{\mathrm{mix}}$ from the tested trajectory, does not cover burn-in, and does not assert that a generic MCMC chain satisfies the required envelope. Its content is a rigorous bridge for chains whose stationarity and mixing rate are established independently.

## 10 Interacting spin-chain test

Theorems about visibility should be tested in a model where blindness is exact and the transfer operator is not inserted by hand. We use the open quantum ANNNI chain

$$H_L = -J_1 \sum_{j=1}^{L-1} Z_j Z_{j+1} - J_2 \sum_{j=1}^{L-2} Z_j Z_{j+2} - h_x \sum_{j=1}^{L} X_j, \tag{75}$$

at $(J_1, J_2, h_x) = (1, 0.37, 2.2)$. This generic next-nearest-neighbour interacting instance is not treated through a free-fermion solution. It preserves the global parity $P = \prod_j X_j$. The computed ground state is even and the first excited state odd throughout $L = 6, 8, \ldots, 16$.

Let $|0\rangle$ be the independently computed ground state and, at the central site, define centered probes

$$|\psi_X\rangle = (X_c - \langle X_c\rangle)|0\rangle, \qquad |\psi_Z\rangle = (Z_c - \langle Z_c\rangle)|0\rangle. \tag{76}$$

Because $X_c$ is parity even, its overlap with the first odd state vanishes exactly. Because $Z_c$ is odd, it can see that state. With $\tau = 0.2$ we construct only the correlator matrices

$$B_n^{ab} = \langle \psi_a | e^{-n\tau(H_L - E_0)} | \psi_b \rangle, \qquad a, b \in \{X, Z\}. \tag{77}$$

The moment pipeline is then separated from the low-energy diagonalization used as ground truth.

Figures 2 and 3 show the two distinct outputs. The Ritz value estimates an edge; the negative localizer certifies that a proposed smaller edge is impossible for the measured spectral measure. The single even probe can look numerically stable while converging to the wrong sector.

At $L = 16$ the Hilbert-space dimension is 65,536. The independent gap is 1.5293121573, the true transfer edge is 0.7364879301, and the degree-six block value is 0.7363890221. Its absolute error is $9.89 \times 10^{-5}$. The first-state weights are $w_X = 1.17 \times 10^{-31}$ and $w_Z = 0.26437$. These numbers do not prove a thermodynamic mass gap; they validate the visibility mechanism and the correlator-only falsifier at nontrivial finite size.

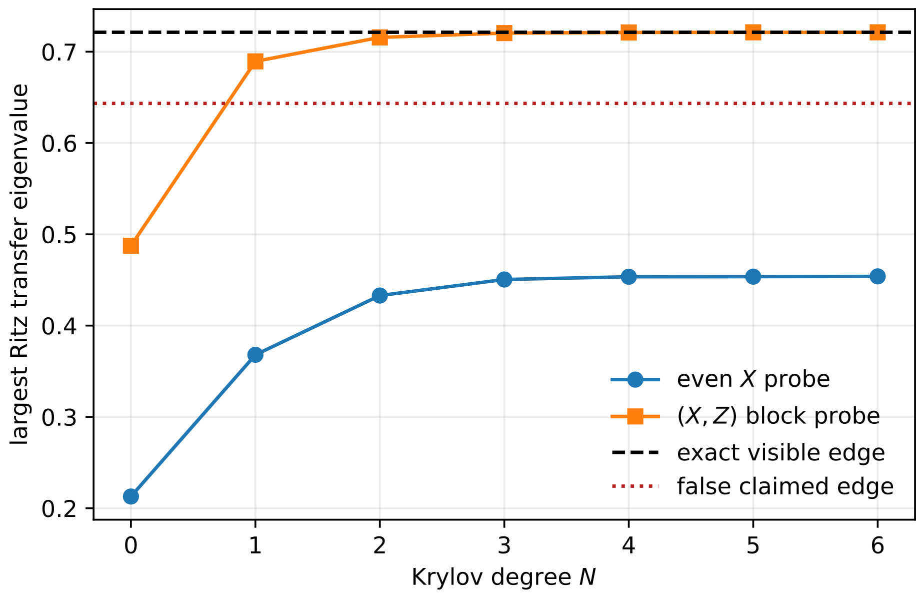


Figure 2: Largest Ritz transfer value for $L = 12$. The even $X$ probe converges to the lowest visible even excitation and misses the true edge. The $(X, Z)$ block family converges rapidly to the exact edge.

## 11 Reproducibility and formal boundary

All computational claims are generated from public, deterministic artifacts:

- exact rational flat-extension, hidden-atom, and Chebyshev artifacts;
- sub-second standard-library certificates for the near-critical atom, exact fourth-kind weighted minimax, Gaussian likelihood/LAN fixture, filter-first Wishart bound, bounded-readout Hoeffding calibration, and the complete $\beta$-mixing lag/coupling/horizon fixture;
- sparse ANNNI construction, independent low-energy benchmark, moment construction, block pencils, and localizers;
- tracked local JSON for $L = 6, 8, 10, 12$;
- a Colab summary for $L = 12, 14, 16$, including Python/NumPy/SciPy versions, source commit, and full-result SHA-256;
- scripts that regenerate all four figures.

Running the tracked `verify_all.py` from the repository root regenerates the exact artifact and audits every theorem-level invariant quoted from the local and Colab records; its README gives the one-line command. The independent certificate under `verification/statistical_gap_boundary` replays the finite-sample and revision fixtures using only the Python standard library; its README gives the commands. The public repository is `https://github.com/lluiseriksson/aqft-split-inclusion-series`. The exact/Krylov/ANNNI directory is frozen at commit `25f068c61905adb039f3b46f53ef23f5be9cc507`. The archived path is `verification/finite_window_gap_certificates`. The immutable Colab entry point is `colab_annni_scale.ipynb`.

The algebraic core of Theorem 4.1 is separately formalized in Lean 4. The immutable provenance is:

- commit: `1e41144d7a563c12f89f9b2ad34fd90aa52149d8`;

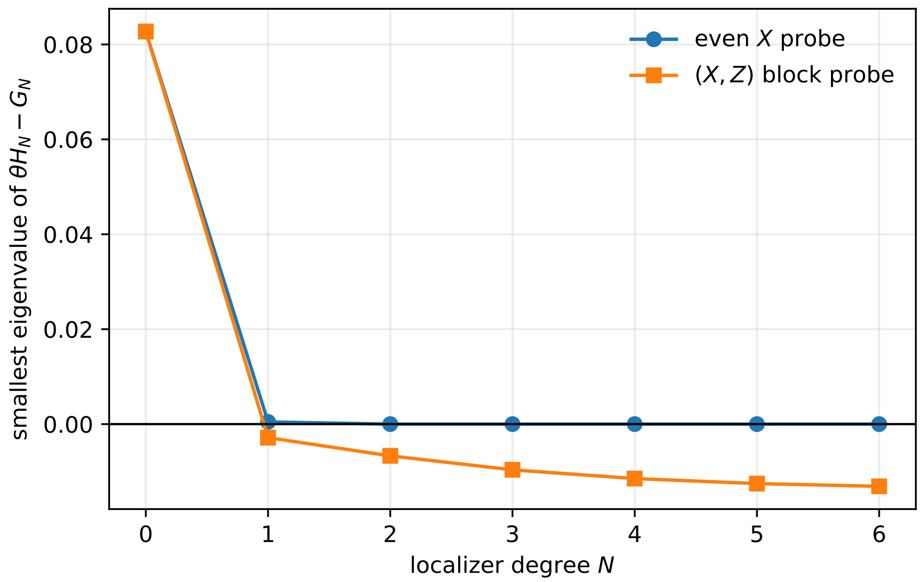


Figure 3: Smallest eigenvalue of the localizer for a gap overstated by 35%. The multichannel family gives a negative certificate from degree one; the symmetry-blind $X$ family never does.

- source: `LeanTransferMatrix/FiniteNoiseBlindness.lean`;
- axiom audit: `audit/FiniteNoiseBlindnessAxioms.lean`;
- review record: PR #25.

The build checks 8,158 targets with no `sorry` and no project axioms. The reported theorem dependencies are exactly {`propext`, `Classical.choice`, `Quot.sound`}.

The formalization proves the normalized endpoint hidden-atom mixture, componentwise distance bound, preservation of the interval constraints, existence inside every error ball $0 \le \varepsilon \le 1$, and $-\log 1 = 0$. It does not formalize the near-critical corollary, bounded-readout and $\beta$-mixing tests, spectral theorem, Chebyshev integral bound, or ANNNI numerics. The exact standard-library replay `certify_statistical_boundary.py` separately checks the near-critical fixture, mixing penalties, and all reported calibrations. These extensions remain paper proofs plus reproducible computations, not Lean theorems.

## 12 Relation to prior methods and novelty audit

This manuscript develops and replaces the unpublished technical draft *Noisy Euclidean Correlators Do Not Certify a Spectral Gap Without Visibility*; that draft was not submitted and is not intended as a separate paper. The exact terminal branch, endpoint no-go, Chebyshev construction, and ANNNI records are retained in revised form. The near-critical theorem, distribution-level Le Cam bound, exact hidden-atom likelihood and LAN phase, fourth-kind weighted minimax, Wishart visibility-to-sample theorem, and bounded unknown-mean filter-bank test are substantive extensions. The present version further adds a stationary $\beta$-mixing theorem with explicit coupling, bias, power, and trajectory-horizon costs.

The companion manuscript *Reflection Positivity and Exact Gap Certificates from Forgotten Quantum Order* [16] constructs a physical reflection-positive transfer operator and gives deterministic Hausdorff certificates. Its Gaussian confidence layer uses the same singular-value/Loewner device as Section 7. The present contribution is different: identifiability failure, near-critical and Le Cam lower bounds, quantitative visibility, explicit power conversion, and probe design. This cross-reference is included to make the shared statistical layer explicit.

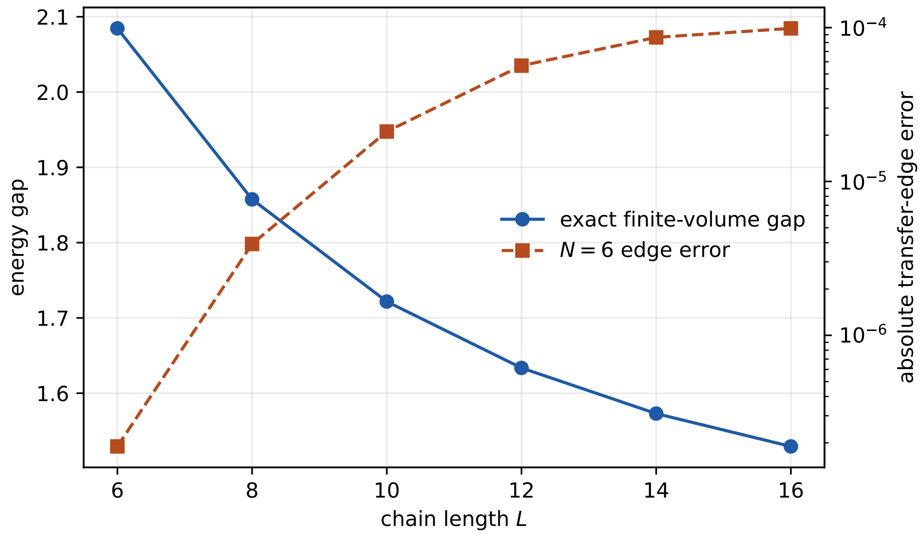


Figure 4: Independent finite-volume energy gap and degree-six transfer-edge error. The $L = 12, 14, 16$ run was executed in Colab Pro+ on the fixed source commit recorded in the artifact. This is a finite-size validation, not a thermodynamic-gap extrapolation.

Matrix Euclidean correlators as moment data, their relation to Rayleigh–Ritz/Lanczos, and rigorous bounds on smeared spectral functions are developed in [5]. The Prony–Ritz equivalence class and exact solution of finite-dimensional spectra from sufficient correlation data are analysed in [4]; matrix-polynomial combinations of GEVP and Prony appear in [2]. GEVP convergence with increasing operator bases is classical in lattice spectroscopy [1], and direct imaginary-time gap estimators have been tested in many-body models in [7, 6]. Recursive and flat matrix moment problems are studied in [9, 10, 11]; recent work also treats recovery of measures with atomic and continuous parts under separation hypotheses [12]. Sparse Hausdorff recovery from noisy moments assumes finite support, a minimum atom weight, and spectral separation [13]. These are complementary structured recovery regimes, not consequences of our visibility premise.

Reflection positivity, equations of motion, and KMS constraints have recently been assembled into semidefinite bootstrap bounds for continuous Euclidean two-point functions [14]. That programme constrains exact admissible correlators; it does not supply a sampling law or a simultaneous finite-sample confidence set for a random empirical covariance. Conversely, the Wishart–Loewner branch in Section 7 assumes independent Gaussian sketches and does not import equations of motion or KMS information. None of these classical or recent ingredients is claimed as new in isolation.

Table 2: Claim audit. "This paper" means the stated combination and specialization, not ownership of the classical ingredients.

| Item | Prior status | Contribution here |
|---|---|---|
| Matrix moments and block Ritz | Established in moment and lattice spectroscopy | Unified notation tied to a one-sided gap localizer |
| Exact finite recovery | Standard Krylov termination/flat moment phenomenon | Checkable terminal criterion stated directly for visible transfer gaps and verified by an exact rational artifact |
| Finite-noise hidden edge | Elementary mixture mechanism; small-component instability is familiar in inverse problems | Algorithm-independent arbitrarily small positive-gap construction inside every positive moment error ball, even when an extra zero mode is forbidden; endpoint core checked in Lean |
| Gaussian hidden edge | Gaussian relative entropy, Pinsker, and Le Cam testing are classical | Near-critical covariance family with exact rank-one-spike likelihood, fixed-window LAN information and power envelope, and positive gaps tending to zero; nonsingularity and fixed-window gates are stated |
| Quantitative visibility | Minimum weights and separation assumptions occur in sparse support recovery | One disclosed, model-dependent condition that restores a finite witness; no claim that it is the unique or weakest possible hypothesis |
| Chebyshev filtering | Classical in approximation and extremal eigenvalue convergence | Exact fourth-kind weighted-localizer minimax, an attained worst-case visibility envelope, a necessary-and-sufficient uniform sign threshold within the normalized scalar polynomial class, and a componentwise error radius |
| Wishart confidence | Gaussian singular-value concentration and Loewner bands are established tools | Exact-level covariance-feasibility test and an explicit visibility-to-margin-to-sample theorem after two-coordinate precommitted filtering |
| Bounded-readout confidence | Hoeffding calibration, paired differences, and Bonferroni control are classical | Distribution-free iid test with unknown mean, exact range constant, explicit power, and same-sample selection over a finite filter bank |
| Dependent-readout confidence | Absolute regularity, Berbee coupling, and Hoeffding concentration are classical | One-trajectory test with an explicit lag-covariance bias, coupling failure budget, power inequality, and geometric-mixing horizon; the mixing envelope is an external premise |
| Probe blindness | Symmetry selection rules are standard | Controlled interacting example separating a stable wrong-sector estimate from a multichannel finite certificate through $2^{16}$ states |

The main conceptual advance is therefore the assembled boundary theorem: finite noisy correlators support rigorous gap falsification, and under disclosed visibility they support finite detection, but they do not support an unconditional positive gap certificate. Exact flatness is the exceptional terminal regime. The Gaussian branches make both directions operational at finite sample size: an exact hidden-atom likelihood and LAN phase for a fixed nonsingular experiment, and an exact-level Wishart test with a closed depth–sample resource curve for a precommitted filtered experiment. The bounded-readout branch removes Gaussianity and known centering for iid data and permits finite-bank selection with a multiplicity penalty. The dependent branch gives a rigorous alternative for stationary bounded $\beta$-mixing output: it exposes rather than hides the price of decorrelation.

## 13 Consequences and limitations

**For lattice mass-gap programmes.** A finite correlator analysis can disprove an overstated transfer gap by a negative localizer. To prove a positive volume-uniform gap, one still needs a volume-uniform totality or visibility theorem for the chosen gauge-invariant operator family, controlled errors, and the reconstruction connecting the Euclidean measure to a positive transfer operator. The present results do not establish any of those model-dependent inputs for four-dimensional Yang–Mills.

**For operator design.** Adding channels is not merely a variance-reduction strategy. It changes the visible cyclic subspace. Symmetry-resolved operator families should be judged by quantitative frame or overlap bounds, not only by the apparent stability of their extracted levels.

**For noisy moment algorithms.** Projection to the positive Hankel cone is useful for restoring consistency, but a projected point estimate should not be reported as a support theorem. The visibility floor and the propagated sign margin in Eq. (44) are the appropriate certificate metadata.

**Limitations.** The fourth-kind filter solves the weighted mass-versus-localizer problem exactly only within normalized degree-$N$ scalar polynomial filters. Its sign condition is necessary and sufficient for uniform sign detection in that class, not a necessary identifiability threshold for arbitrary estimators. The componentwise error model ignores correlations and can be conservative. The Wishart theorem instead uses the full covariance, but requires independent exactly Gaussian, known-zero-mean sketches and a filter fixed before seeing the data. The bounded-iid theorem removes Gaussianity and known centering and allows a finite bank. The $\beta$-mixing theorem covers dependent stationary bounded output only when a valid mixing envelope is supplied externally; it does not infer that envelope, remove burn-in, or cover arbitrary MCMC output. The LAN boundary fixes the feature degree, base covariance, and atom path. It does not cover an unknown-location scan or a degree growing with $n$. When the filter changes with $\gamma$, the covariance and its conditioning also change, so the achievable depth–sample law is not claimed globally minimax. The ANNNI study is finite-volume and uses exact sparse propagation rather than Wishart, bounded-readout, or Monte Carlo data. No thermodynamic extrapolation is attempted. Data-driven mixing-rate certification, matrix-valued optimal filters, interval arithmetic, and formalization of the Chebyshev, Wishart, bounded-readout, and mixing theorems remain open extensions.

## 14 Conclusion

The distinction between estimating a visible level and certifying a spectral gap is structural. Block Hankel pencils give monotone visible-edge estimates and finite falsifiers. Exact rank stabilization closes the problem at finite depth. With any nonzero error and no visibility premise, however, an atom at $e^{-g} < 1$ with arbitrarily small $g > 0$ lies inside the data ball; this remains true when uniqueness of the vacuum is known. In a nonsingular Gaussian experiment its exact likelihood has an $n^{-1/2}$ local weight boundary and a closed Gaussian power envelope. A quantitative visibility floor is one explicit bridge back to a finite theorem: fourth-kind weighted-Chebyshev localization produces the exact worst-case signed margin in its polynomial class, and a two-coordinate Wishart–Loewner band turns that margin into an exact-level test with a closed depth–sample bound. For bounded iid readouts, paired differences and Hoeffding calibration provide a non-Gaussian, unknown-mean companion with finite-bank selection. For a stationary bounded $\beta$-mixing trajectory, sparse pairing and Berbee coupling add a certified dependence penalty, lag bias, and elapsed chain horizon. Rejection falsifies an overstated gap; non-rejection alone never becomes a positive gap certificate. The interacting spin-chain experiment shows the physical meaning of the boundary: a symmetry-blind probe can converge smoothly and still miss the gap-controlling state, while a multichannel family supplies an immediate negative witness. This is a finite-window falsification theorem, not a thermodynamic or Yang–Mills mass-gap proof.

**AI assistance disclosure.** AI tools assisted with literature discovery, adversarial proof checking, symbolic fixture verification, and manuscript editing. All theorem statements, scope decisions, computations, and the final text were reviewed by the author, who takes responsibility for the claims.

## References

[1] B. Blossier, M. Della Morte, G. von Hippel, T. Mendes, and R. Sommer, "On the generalized eigenvalue method for energies and matrix elements in lattice field theory," *JHEP* **04**, 094 (2009), arXiv:0902.1265.

[2] G. T. Fleming, "Beyond generalized eigenvalues in lattice quantum field theory" (2023), arXiv:2309.05111.

[3] M. L. Wagman, "Lanczos, the transfer matrix, and the signal-to-noise problem," *Phys. Rev. Lett.* **134**, 241901 (2025), arXiv:2406.20009.

[4] R. Abbott, D. C. Hackett, G. T. Fleming, D. A. Pefkou, and M. L. Wagman, "Filtered Rayleigh–Ritz is all you need" (2025), arXiv:2503.17357.

[5] R. Abbott, W. I. Jay, and P. R. Oare, "Moment problems and bounds for matrix-valued smeared spectral functions" (2025), arXiv:2508.01377.

[6] J. M. Leamer, A. B. Magann, G. McCaul, and D. I. Bondar, "Spectral gaps via imaginary time" (2026), arXiv:2303.02124.

[7] H. Suwa and S. Todo, "Generalized moment method for gap estimation and quantum Monte Carlo level spectroscopy," *Phys. Rev. Lett.* **115**, 080601 (2015), arXiv:1402.0847.

[8] Y. Yu, A. F. Kemper, C. Yang, and E. Gull, "Denoising of imaginary time response functions with Hankel projections," *Phys. Rev. Research* **6**, L032042 (2024), arXiv:2403.12349.

[9] R. Curto, A. Ech-charyfy, K. Idrissi, and E. H. Zerouali, "A recursive approach to the matrix moment problem" (2023), arXiv:2209.01630.

[10] R. E. Curto, L. A. Fialkow, and H. M. Möller, "The extremal truncated moment problem," arXiv:math/0610882.

[11] C. Mädler and K. Schmüdgen, "On the truncated matricial moment problem. I" (2023), arXiv:2310.00957.

[12] R. Karapetyan, S. Ma, A. Wodecki, and J. Mareček, "On moment-based recovery of measures with atomic and continuous parts" (2026), arXiv:2605.21644.

[13] S. Gordon, B. Mazaheri, L. J. Schulman, and Y. Rabani, "The sparse Hausdorff moment problem, with application to topic models" (2020), arXiv:2007.08101.

[14] M. Cho, B. Gabai, H. W. Lin, J. Yeh, and Z. Zheng, "Bootstrapping Euclidean two-point correlators" (2025), arXiv:2511.08560.

[15] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *J. Amer. Statist. Assoc.* **58**, 13–30 (1963), doi:10.1080/01621459.1963.10500830.

[16] L. Eriksson, "Reflection positivity and exact gap certificates from forgotten quantum order" (2026), companion manuscript, canonical PDF SHA-256 `d18422c9981769c4c36a7382e6831e200a94bcbb1fab447741ae240536e6b6d2`.

[17] H. C. P. Berbee, *Random Walks with Stationary Increments and Renewal Theory*, Mathematical Centre Tracts 112, Mathematisch Centrum, Amsterdam (1979), ISBN 978-90-6196-182-6, CWI publication 12694.

[18] M. Meitz and P. Saikkonen, "Subgeometric ergodicity and $\beta$-mixing," *J. Appl. Probab.* **58**, 594–608 (2021), doi:10.1017/jpr.2020.108, arXiv:1904.07103.

[19] M. Rudelson and R. Vershynin, "Non-asymptotic theory of random matrices: extreme singular values," in *Proceedings of the International Congress of Mathematicians 2010*, pp. 1576–1602, arXiv:1003.2990.

[20] A. B. Tsybakov, *Introduction to Nonparametric Estimation*, Springer, New York (2009).

[21] A. W. van der Vaart, *Asymptotic Statistics*, Cambridge University Press, Cambridge (1998), doi:10.1017/CBO9780511802256.

[22] A. Onatski, M. J. Moreira, and M. Hallin, "Asymptotic power of sphericity tests for high-dimensional data," *Ann. Statist.* **41**, 1204–1231 (2013), doi:10.1214/13-AOS1100, arXiv:1306.4867.

[23] J. C. Mason and D. C. Handscomb, *Chebyshev Polynomials*, Chapman & Hall/CRC, Boca Raton (2003), doi:10.1201/9781420036114.

[24] Y. Saad, "On the rates of convergence of the Lanczos and the block-Lanczos methods," *SIAM J. Numer. Anal.* **17**, 687–706 (1980).

[25] M. Lüscher, "Construction of a selfadjoint, strictly positive transfer matrix for Euclidean lattice gauge theories," *Commun. Math. Phys.* **54**, 283–292 (1977).