Finite-spectral recovery of Weyl-blind twisting modes

Scott Schubr
Luminary Oddities

October 23, 2026

Abstract

Odd logarithmic twisting modes in a fixed polar-strip tube are invisible to the leading Weyl profile but detectable in the next heat profile. We give a quantitative continuation of this complete-spectrum result. An estimate on the actual curved normal cap yields heat remainder $O(T^{17/6})$, and hence normalized bias $O(T^{-1/6})$. A complete finite spectral prefix, certified spectral uncertainties, a counting envelope and a certified complex multiplier enclosure then give a deterministic recovery disk. With exact spectral inputs its fixed-mode cutoff rate is $O((\log\Lambda/\Lambda)^{1/12})$; the relative-noise balance is $O(\sigma^{1/7})$ before inversion of the heat multiplier. A long-straight-core construction proves why a uniform geometric prior is necessary. Two-sided multiplier bounds quantify exponential conditioning of the ideal lower heat profile. The geometric remainder constants exist but are not fully evaluated numerically. A concrete fundamental-mode interval execution is valid and noninformative.

Keywords: inverse spectrum; twisted tube; heat trace; finite spectral data; deterministic error bounds; geometric priors.

1 Model, predecessors and conclusions

To recover a fixed twisting mode from finite spectral data, we require a complete retained spectral prefix with multiplicities and certified uncertainties, a valid lower cutoff for every omitted level, a full counting or heat envelope, and a certified complex multiplier disk. The actual curved-cap estimate supplies the bias term. Its geometric constant and onset must be supplied, or replaced by another proved bias bound, for a numerical disk.

The resulting disk concerns one fixed mode and its geometric class. Converting a complex Fourier coefficient to a cosine amplitude doubles the radius. The long-core example explains the need for uniform geometric priors, and the reported noninformative execution is retained as a demonstration of the present numerical scale.

Let the connected polar section be

$$\omega = \operatorname{int}\left(\overline{(1,3)\times(0,1)} \cup \overline{(3,\sqrt{11})\times(0,2)}\right). \tag{1.1}$$

In particular, $r = 3$, $0 < q < 1$ is an interior aperture. Its Cartesian image is denoted by W. The Euclidean tube has parametrization

$$F(a,r,q) = (a, r\cos(q+\Theta(a)), r\sin(q+\Theta(a))), \qquad a > 0,$$

and Dirichlet boundary, including its cap at $a = 0$. Write H for its Dirichlet Laplacian and $\mathcal{H}(T) = \operatorname{Tr} e^{-H/T^2}$. On a sufficiently remote end $b = \Theta'$ is positive and increasing, with inverse

$$R(y) = yh(\log y), \quad L = 2\log 2, \quad \tau = \frac{\pi}{\log 2}, \quad h(u) = 1 + \sum_{k\in\mathcal{K}} h_k e^{ik\tau u}. \tag{1.2}$$

Here $\mathcal{K}$ is a finite set of nonzero odd integers, $h_{-k} = \overline{h_k}$, and $h > 0$, $h + h' > 0$. The initial part is a fixed smooth sweep of the same section. A fixed axial offset in the inverse law is allowed; it changes only a constant lower coefficient. Arbitrarily independently varied two ends are not considered. For two specified identical ends divide the trace and coefficients by their known multiplicity two.

The actual-tube heat theorem in [1], together with the simultaneous noncancellation theorem in [2], gives

$$\mathcal{H}(T) = A_0 T^4 + T^3 \left(c + \sum_{k \in \mathcal{K}} h_k C_k e^{ik\tau \log T} \right) + o(T^3), \qquad A_0 = \frac{1}{6\pi}. \tag{1.3}$$

The leading profile is independent of the odd coefficients. Our use of these two predecessor results is explicit: their full proofs are included as companion manuscripts in the publication package. The present paper supplies a rate and a finite-data error budget.

Theorem 1.1 (Actual-tube power remainder)**.** *For each fixed tube above there are finite $C_r, T_* \geq 1$ such that*

$$\left| \mathcal{H}(T) - A_0 T^4 - T^3 F_h(\log T) \right| \leq C_r T^{17/6}, \quad T \geq T_*, \qquad F_h(u) = c + \sum_{k \in \mathcal{K}} h_k C_k e^{ik\tau u}. \tag{1.4}$$

Common constants over a class require a common eventual-law onset, bounded core length and speed, positive common margins for h and $h + h'$, and uniform bounds for the finite derivatives used in the normal-form estimates below. This assertion is not uniform over uncontrolled cores. The constants in (1.4) *are not evaluated numerically in this paper.*

One explicit sufficient controlled class fixes the finite index set $\mathcal{K}$, a common bound on every $|h_k|$, positive common lower bounds on h and $h + h'$, a common speed-law onset in the absolute axial coordinate, a bounded fixed offset, and a common bound on the initial core's speed and length. The Fourier bounds then give uniform bounds on every derivative of h; inverse differentiation supplies all the required tail derivatives. This specification gives sufficient analytic uniformity, but does not assign numerical values to the resulting constants.

Sections 2–3 prove Theorem 1.1. Section 4 gives a finite recovery disk with every numerical input stated separately. Its exact-input rate for fixed odd k is $O_k((\log \Lambda/\Lambda)^{1/12})$, where all omitted eigenvalues are at least Λ. This is a sufficient rate, without an optimality claim. The cap argument uses standard Sobolev, Nash and localized heat tools. Its contribution is control of their dependence on the moving cap in this tube. General interior heat localization predates this work [6, 5]; an alternative proof using [4, 5] appears in Appendix A.

Inverse twisting results with richer observations must be distinguished from this eigenvalue-prefix problem. Choulli–Soccorsi [8] recover twisting information from a Dirichlet-to-Neumann map in a homogeneous cylindrical setting; their data and bounded-twist hypotheses differ from ours. Kavian–Kian–Soccorsi [9] study a periodic potential in a straight waveguide using Floquet spectral data that include boundary normal derivatives of eigenfunctions. Neither result supplies finite eigenvalue-only recovery of the present diverging tail. For practical heat fitting, the rectangular-box study [10] gives computational reconstruction demonstrations. Its fitted window and density-based tail approximation are a different numerical task from the deterministic complete-tail certificate here; our concrete benchmark does not achieve its practical accuracy.

2 The exact normal-cap core

In physical coordinates the Dirichlet form is

$$\int \left(|\partial_r u|^2 + r^{-2} |\partial_q u|^2 + |\partial_a u - b(a) \partial_q u|^2 \right) r \, \mathrm{d}r \, \mathrm{d}q \, \mathrm{d}a. \tag{2.1}$$

The remote normal coordinates of [1, Section 4] are

$$x = \Theta^{-1}(\Theta(a) + q), \qquad z = x - \frac{q}{b(x)(1 + b(x)^2 r^2)}. \tag{2.2}$$

Choose a large speed cutoff B and cut at $z = \ell_B = R(B)$, with the fixed offset included if present. The inverse clock is needed only near this remote cut, and need not exist through a possibly straight initial segment.

Lemma 2.1 (Flattening the actual cap)**.** *In untwisted coordinates the cap is a graph $a = g_B(r, q)$ with*

$$g_B = \ell_B - q/B + O(B^{-3}), \qquad |g_B - \ell_B| + |\nabla_W g_B| \le C/B.$$

The physical core Ω_B is the image of $(0, \ell_B) \times W$ under a map Ψ_B whose Jacobian is bounded above and below by fixed positive constants and whose derivative and inverse Lipschitz bound are at most CB.

Proof. The margins in (1.2) give $0 < c \le b' \le C$ and $b'' = O(b^{-1})$. At fixed (r, q) the second equation of (2.2), with $z = \ell_B$, gives $x = \ell_B + O(B^{-3})$ and $|\partial_r x| + |\partial_q x| = O(B^{-3})$ by implicit differentiation. Then $\Theta(g_B) = \Theta(x) - q$, giving the stated expansion and gradient bounds; for example $\partial_q g_B = (b(x)\partial_q x - 1)/b(g_B)$. Let χ_B be a fixed-shape cutoff on the last unit of $(0, \ell_B)$, zero below $\ell_B - 1$ and one near ℓ_B. The map

$$G_B(s, w) = (s + \chi_B(s)[g_B(w) - \ell_B], w)$$

has uniformly bounded derivative and inverse and Jacobian $1 + O(B^{-1})$. It fixes the lower cap and parameter lateral faces. Compose it with $F(a, w) = (a, \mathrm{Rot}_{\Theta(a)}\, w)$. On this core $\sup |b| \le CB$, including the fixed initial portion; F has Jacobian one and derivative bounded by CB on the bounded section. Its ambient inverse rotation on a disk containing W has the same Lipschitz bound. A fixed extension of g_B off W gives the corresponding bound for G_B^{-1}, also across the section's reflex corner. Thus $\Psi_B = F \circ G_B$ has all the claimed bounds. □

Let H_B be either the full Neumann realization on Ω_B or a mixed realization with Dirichlet physical lateral boundary and any prescribed Dirichlet/Neumann artificial caps. Its form is the exact Euclidean energy, not a separately diagonalized model operator. Coordinate and density transformations preserve this realization; they are not permission to change its cap condition.

Lemma 2.2 (Nash bound with cutoff dependence)**.** *For $0 < t \le 1$ the heat kernel of H_B satisfies*

$$0 \le k_B(t, X, Y) \le CB^3 t^{-3/2}, \qquad \int_{\Omega_B} k_B(t, X, Y)\, \mathrm{d}Y \le 1. \tag{2.3}$$

Proof. The fixed-section cylinders $V_\ell = (0, \ell) \times W$, $\ell \ge 2$, have the uniform inequality $\|v\|_6^2 \le C(\|\nabla v\|_2^2 + \|v\|_2^2)$. Extend transversely with a fixed Lipschitz-domain extension operator, reflect axially into collars of length one at each end, and cut off there. The disjoint collars and fixed transverse extension make all norms independent of ℓ; Euclidean Sobolev proves the claim. Pulling back once by Ψ_B gives

$$\|u\|_6^2 \le CB^2(\|\nabla u\|_2^2 + \|u\|_2^2). \tag{2.4}$$

Interpolation yields $\|u\|_2^{10/3} \le CB^2(q_B[u] + \|u\|_2^2)\|u\|_1^{4/3}$. For $y(t) = \|e^{-t(H_B+1)} f\|_2^2$, sub-Markov contractivity gives $y' \le -cB^{-2} y^{5/3}/\|f\|_1^{4/3}$. Integration and semigroup composition give the all-time $L^1 \to L^\infty$ bound $CB^3 t^{-3/2}$ for $H_B + 1$. Removing the shift costs at most e on $t \le 1$. The unshifted sub-Markov property gives the mass bound. □

Lemma 2.3 (Exact volume coefficient)**.** *For each of the actual core realizations above,*

$$\left|\operatorname{Tr} e^{-tH_B} - |\Omega_B|(4\pi t)^{-3/2}\right| \le CB^5 t^{-1}, \qquad 0 < t \le 1. \tag{2.5}$$

Proof. Fix an interior X with $d = \operatorname{dist}(X, \partial\Omega_B) \ge \sqrt{t}$. Choose a smooth cutoff supported in $B_d(X)$, equal to one near X, with derivatives bounded by C/d, C/d^2, supported on an annulus at distance comparable to d from X. Subtract the cutoff free heat kernel from k_B and use Duhamel. The forcing is $2\nabla\chi \cdot \nabla p_0 + (\Delta\chi)p_0$. For the time variable $s \le t/2$, its spatial L^1 norm is $C(d^{-2} + s^{-1})e^{-cd^2/s}$ and the other kernel is bounded by $CB^3 t^{-3/2}$. The time integral is bounded by $CB^3 t^{-3/2} e^{-c'd^2/t}$. For $s \ge t/2$ its pointwise norm is at most $Ct^{-3/2}(d^{-2} + t^{-1})e^{-cd^2/t}$. Use the mass bound in (2.3) and integrate in s. Since $t/d^2 \le 1$, the same estimate follows. Thus

$$|k_B(t, X, X) - (4\pi t)^{-3/2}| \le CB^3 t^{-3/2} e^{-c\operatorname{dist}(X,\partial\Omega_B)^2/t}. \tag{2.6}$$

On $d < \sqrt{t}$ this follows directly from (2.3), after enlarging C. The cutoff has interior support, so it satisfies every cap condition used here. The form semigroup, or interior smooth approximation, justifies Duhamel on the piecewise smooth domain. This argument follows the localization mechanism of [5, Proposition 4.1], with a split time integral exposing the polynomial dependence.

The product boundary layer obeys $|\{Y \in V_{\ell_B} : \operatorname{dist}(Y, \partial V_{\ell_B}) < a\}| \le C(\ell_B + 1)a$ for all $a > 0$. Lemma 2.1 and $\ell_B = O(B)$ therefore imply

$$|\{X \in \Omega_B : \operatorname{dist}(X, \partial\Omega_B) < a\}| \le CB^2 a.$$

Layer-cake integration gives an integral at most $CB^2\sqrt{t}$ for the exponential in (2.6). Its prefactor proves (2.5), with the exact free-volume coefficient. No Neumann domain monotonicity or planar substitute is used. □

3 Assembling the heat remainder

For clarity we state the quantitative inputs from the actual heat lift [1, Sections 3–4], including their cutoff dependence. On the remote end, clock, density removal and shear give weighted relative error $O(b^{-2})$, plus a bounded potential, against

$$D_r^2 + b(z)^2 D_q^2 + \frac{D_z^2}{b(z)^2 r^2}, \qquad D_j = -i\partial_j. \tag{3.1}$$

For example, before the shear the angular/axial energy is $r^{-2}|D_q + b^{-1}D_x|^2 + b^2|D_q|^2$. The shear in (2.2) cancels its leading mixed term; the Schur complement is $(1 + b^2r^2)^{-1}$, with relative $O(b^{-2})$ difference from $(b^2r^2)^{-1}$. The derivatives $b' = O(1)$, $b'' = O(b^{-1})$, and finite higher bounds follow by differentiating (1.2). Density derivatives and the polar potential give the stated bounded absolute term.

Partition $z \ge \ell_B$ into cells of length $\ell = T^{-1/2}$. Dirichlet/Neumann caps are inserted in these normal coordinates after the exact changes of variables. For frozen speed b_j, the trace and first normalized energy moment of a cell are bounded by

$$C(\ell T^3 + T^2/b_j)e^{-cb_j^2/T^2}. \tag{3.2}$$

Indeed radial Neumann relaxation into the two bands bounds the zero-axial-momentum trace by $CTZ_2(b_j^2/T^2) \le CT^2 b_j^{-1} e^{-cb_j^2/T^2}$, where $Z_a(t) = \sum_{n\ge1} e^{-t(\pi n/a)^2}$. The Gaussian axial sum gives (3.2); the energy moment follows from $xe^{-x} \le Ce^{-x/2}$. The cap sum versus momentum integral costs the zero-momentum trace, by monotonicity in momentum. The relative cell error

is $C(\ell/b_j + b_j^{-2})$. Keeping the exponential weight cell by cell, and using $\mathrm{d}z \asymp \mathrm{d}b$, bounds the remote comparison error by

$$C\{T^3/B + T^{5/2}\log(2+T) + T^2\}. \tag{3.3}$$

The two volume integrals are $CT^3\ell \int_B^\infty b^{-1}e^{-cb^2/T^2}\,\mathrm{d}b$ and $CT^3 \int_B^\infty b^{-2}e^{-cb^2/T^2}\,\mathrm{d}b$. The cap sum is bounded by $CT^2\ell^{-1}\int_B^\infty b^{-1}e^{-cb^2/T^2}\,\mathrm{d}b$. The remaining weighted terms are no larger than $C(T^2/B + T^{5/2}/B^2)$ and the bounded potential costs CT^2. Thus the constants in (3.3) are uniform for sufficiently large B; only the core still needs Lemma 2.3.

The frozen radial theorem in [1, Section 3] is

$$\begin{aligned} J_\eta(v) &= \frac{v}{4\pi\eta}\{4Z_1(v^2) + Z_2(v^2)\} + \frac{v}{2\sqrt{\pi}}G(v) + E_\eta(v), \\ G(v) &= 3D(v^2) - Z_1(v^2) - \frac{3+\sqrt{11}}{4}Z_2(v^2), \\ |E_\eta(v)| &\le C\eta v Z_2(v^2), \qquad \int_0^\infty vZ_2(v^2)\,\mathrm{d}v = \frac{1}{3}. \end{aligned} \tag{3.4}$$

Here J_η includes integration of axial momentum, and D is the connected-minus-cut relative heat trace of the step with widths one and two. It is a finite relative invariant, not a subtraction of infinite traces. To recall the normalization, integrating the axial Gaussian first in the killed bridge formula gives path weight $v(\int_0^1 \rho_s^{-2}\,\mathrm{d}s)^{-1/2}/(2\sqrt{\pi})$. For $\rho_s = r + \eta\beta_s$ its expansion is $v[r + \eta\int\beta_s\,\mathrm{d}s + O(\eta^2\sup|\beta|^2)]/(2\sqrt{\pi})$. The centered first variation vanishes in the bulk; each endpoint has bridge-hit weight $1/4$. Reconnecting at $r = 3$ yields $3D$; the weight-freezing error is $O(\eta Z_2)$ by the integrated Gaussian bridge supremum bound. This is the established radial estimate we use, not an isolated-width replacement of the aperture.

With $\eta = T^{-1}$, $b = Tv$ and momentum $p = T^2\xi$, (3.4) has total integrated error CT^2, since R' is bounded. Its bulk and boundary contributions are

$$\begin{aligned} &\frac{T^2}{4\pi}\int_B^\infty yR'(y)[4Z_1(y^2/T^2) + Z_2(y^2/T^2)]\,\mathrm{d}y, \\ &\frac{T^3}{2\sqrt{\pi}}\int_{B/T}^\infty vR'(Tv)G(v)\,\mathrm{d}v. \end{aligned} \tag{3.5}$$

For each fixed width a, the decreasing-sum integral bound gives $|Z_a(y^2/T^2) - aT/(2\sqrt{\pi}y)| \le C$ for all $y, T > 0$. Extending the bulk integral to zero therefore adds the planar core's exact free-volume term, with error CT^2B^2. The actual cap changes its volume by $O(B^{-1})$, costing CT^3/B. A fixed axial offset contributes a constant multiple of T^3, already in c. Because $vG(v)$ is bounded near zero and decays exponentially at infinity, the omitted boundary integral costs CT^2B. The Fourier evaluation of the full integrals is precisely the leading and lower profiles in (1.3).

Proof of Theorem 1.1. At $t = T^{-2}$ Lemma 2.3 contributes CT^2B^5. The preceding estimates give, uniformly for large B,

$$|\mathcal{H}(T) - A_0T^4 - T^3F_h(\log T)| \le C\{T^3/B + T^{5/2}\log(2+T) + T^2B^5 + T^2B^2 + T^2B + T^2\}.$$

Choose $B = T^{1/6}$ once this exceeds every geometric threshold. The first and third terms are $T^{17/6}$; all others are smaller. Each constant depends only on the stated fixed geometry and finite normal-form derivative bounds. The same proof with common bounds gives the controlled-class assertion. □

4 Finite spectral certification

Fix a nonzero odd k, put $\omega = k\tau$, and let $T = e^{u_0} \geq \max(1, T_*)$. Suppose the first M eigenvalues, with multiplicity and no missing low levels, have enclosures

$$\lambda_n \in [\ell_n, \mu_n + d_n], \quad \ell_n = \max(0, \mu_n - d_n), \quad \mu_n, d_n \geq 0.$$

Assume every omitted level is at least a certified $\Lambda \geq 0$, and the complete spectrum satisfies a known $N(x) \leq A_N(1+x)^2$ for $x \geq 0$. Define the finite functional

$$Q_M = \frac{1}{L}\int_{u_0}^{u_0+L} e^{-i\omega u}\left[e^{-3u}\sum_{n\leq M} e^{-\lambda_n e^{-2u}} - A_0 e^u\right] \mathrm{d}u. \tag{4.1}$$

For the full trace, Fourier orthogonality cancels c and all other modes. The bias is at most $C_r T^{-1/6}$. The leading integral is evaluated explicitly using the primitive $A_0 e^{(1-i\omega)u}/[L(1-i\omega)]$. Each spectral term has kernel

$$K_k(\lambda) = \frac{1}{2L}\int_{1/(16T^2)}^{1/T^2} t^{(1+i\omega)/2} e^{-\lambda t}\,\mathrm{d}t,$$

where the complex power uses the real logarithm of $t > 0$.

Theorem 4.1 (Deterministic recovery disk)**.** *Compute a numerical q_k from* (4.1) *using centers μ_n, with certified quadrature and arithmetic errors $E_{\mathrm{quad}}, E_{\mathrm{arith}}$. Set*

$$E_{\mathrm{tail}} = A_N T^{-3} e^{-\Lambda/(16T^2)}[(1+\Lambda)^2 + 32T^2(1+\Lambda) + 512T^4], \tag{4.2}$$

$$E_{\mathrm{noise}} = \frac{1}{L}\int_{u_0}^{u_0+L} e^{-5u}\sum_{n\leq M} d_n e^{-\ell_n e^{-2u}}\,\mathrm{d}u,$$

$$\Delta_k = C_r T^{-1/6} + E_{\mathrm{tail}} + E_{\mathrm{noise}} + E_{\mathrm{quad}} + E_{\mathrm{arith}}.$$

Then $|q_k - h_k C_k| \leq \Delta_k$. If $|C_k - c_k| \leq \eta_k < |c_k|$ is a separate certified complex multiplier disk and $|h_k| \leq H_k$ is a known prior, then

$$|q_k/c_k - h_k| \leq \frac{\Delta_k + H_k\eta_k}{|c_k|}. \tag{4.3}$$

For $h = 1 + \epsilon\cos(k\tau u)$ the amplitude is $2h_k$; double the disk radius and intersect with the real prior interval.

Proof. For the omitted counting measure, Tonelli gives

$$\sum_{n>M} e^{-t\lambda_n} \leq A_N t\int_\Lambda^\infty e^{-tx}(1+x)^2\,\mathrm{d}x = A_N e^{-t\Lambda}[(1+\Lambda)^2 + 2(1+\Lambda)/t + 2/t^2].$$

Its window maximum after multiplication by T^{-3} gives the stated tail bound. A sharper certificate may integrate this bound over the window. The scalar mean-value theorem gives $|e^{-t\lambda_n} - e^{-t\mu_n}| \leq t d_n e^{-t\ell_n}$, proving the noise term. Add these errors to the Fourier bias and validated evaluation errors. Finally $q_k - c_k h_k = (q_k - C_k h_k) + (C_k - c_k)h_k$ proves (4.3). □

An elementary quadrature certificate is available without special functions. With J equal midpoint panels of width $d = L/J$, the normalized error for the spectral part is at most $D_M d/4$, where

$$D_M = T^{-3}\sum_{n\leq M}[\sqrt{9+\omega^2} + 2\mu_n/T^2]e^{-\mu_n/(16T^2)}. \tag{4.4}$$

Differentiate the integrand; its derivative factor is $-3 - i\omega + 2\mu_n e^{-2u}$. Integrating $D_M|u - u_{\mathrm{mid}}|$ on each panel proves the formula. Higher-order validated integration can improve it. The full counting bound or a validated earlier-time heat upper bound is essential: a partial trace is a lower bound, and cannot serve as a complete-trace tail upper bound.

Corollary 4.2 (Exact-input cutoff rate)**.** *For exact retained eigenvalues and sufficiently accurate certified evaluation, each fixed mode admits forward error $O((\log\Lambda/\Lambda)^{1/12})$ as $\Lambda\to\infty$. With exact C_k this is also the coefficient rate, with a mode-dependent constant. A multiplier approximation must improve with the target tolerance, or its term in* (4.3) *remains.*

Proof. Choose $x = q\log\Lambda$, $q > 7/12$, and $T = \sqrt{\Lambda}/(4\sqrt{x})$. The bias is $4^{1/6}C_r(x/\Lambda)^{1/12}$. For $\Lambda \geq 1$ the tail is at most

$$64A_N\Lambda^{1/2-q}(4x^2+4x+2)/\sqrt{x},$$

which is lower order. Eventually $T \geq T_*$. □

Coverage through Λ needs no more than $A_N(1+\Lambda)^2$ levels. For fixed mode, controlled constants and exact multipliers, one conservative sufficient resource law is $M = O(\delta^{-24}\log^2(1/\delta))$ for forward tolerance δ. This is neither a recommendation nor a necessary budget.

5 Noise and the conditioning of the ideal heat profile

Suppose a validated global bound $\mathcal{H}(U) \leq K_H U^4$ is known for $U \geq 1$. Absolute errors bounded by η contribute pointwise normalized heat error at most

$$\min\{M\eta/T^5,\ \eta e^{\eta/T^2}K_H/T\}. \tag{5.1}$$

For relative errors $|\widetilde{\lambda}_n - \lambda_n| \leq \sigma\lambda_n$, $0 \leq \sigma < 1$, the corresponding bound is $8K_H\sigma T/[e(1-\sigma)^3]$. Indeed, with $x = \lambda_n/T^2$, each difference is bounded by $\sigma x e^{-(1-\sigma)x}$; use $xe^{-ax} \leq 2(ea)^{-1}e^{-ax/2}$ and sum at $U = T\sqrt{2/(1-\sigma)}$. On the full window T to $4T$ the constant is $32K_H/[e(1-\sigma)^3]$. Balancing this with $C_rT^{-1/6}$ gives $T \asymp \sigma^{-6/7}$ and normalized error $O(\sigma^{1/7})$, for small enough σ to exceed the threshold. Tail and numerical errors must also fit that budget. These are upper bounds for this estimator, not optimal raw-spectral noise laws.

We now specify the ideal multiplier conditioning used in the recovery disk. Let $D(t)$ be as in (3.4), $E(t) = D(t) - Z_1(t)/2$, and $M_E(z) = \frac{1}{2}\int_0^\infty t^{(z-1)/2}E(t)\,\mathrm{d}t$. Put $z_k = 1 + ik\tau$, $s_k = (z_k+1)/2$ and $A = 7/2 + \sqrt{11}$. The established odd-mode formula is

$$C_k = \frac{z_k}{2\sqrt{\pi}}\left[A\frac{\Gamma(s_k)\zeta(2s_k)}{2\pi^{2s_k}} + 3M_E(z_k)\right]. \tag{5.2}$$

The stronger positive-measure representation for the full D in [2] yields, with $b_k = k\pi/(2\log 2)$,

$$L_k < |C_k| < \frac{8767}{180}L_k < 49L_k, \qquad L_k = \frac{|z_k\Gamma(1+ib_k)|}{160\sqrt{\pi}}, \quad k > 0 \text{ odd}. \tag{5.3}$$

For completeness the upper estimate uses that representation in the form $C_k = 3z_k\Gamma(s_k)[Z_D(s_k) + ((2+\sqrt{11})/3)Z_1(s_k)]/(4\sqrt{\pi})$, where $|Z_D| < 1/12 + 21/800$ and $Z_1(s) = \pi^{-2s}\zeta(2s)$ has $|Z_1| \leq 1/6$. Together with $\sqrt{11} < 10/3$ the bracket is bounded by $8767/21600$, giving (5.3). The lower estimate is [2, Theorem 1.1]; it is not inferred merely from positivity of the corner remainder.

The gamma identity gives the exact scalar formula

$$L_k^2 = \frac{(1+4b_k^2)2b_k}{160^2(e^{\pi b_k} - e^{-\pi b_k})}, \qquad L_k \sim \frac{\sqrt{2}}{80}b_k^{3/2}e^{-\pi b_k/2}. \tag{5.4}$$

It decreases on positive odd k: the logarithmic derivative of $(1+4b^2)b/\sinh(\pi b)$ is less than $3/b - \pi < 0$ there. On the real odd Fourier subspace $|k| \leq K$, normalized Parseval therefore bounds the inverse norm of $h_k \mapsto C_kh_k$ between $1/(49L_K)$ and $1/L_K$. The lower test is a sufficiently small real cosine perturbation of $h = 1$, hence admissible. This exponential conditioning concerns the ideal lower heat profile. It is not a minimax lower bound for every method using raw eigenvalues. General transform instability has substantial predecessors, including [7].

6 Why a core prior is necessary

Proposition 6.1 (Uncontrolled-core obstruction)**.** *On the class allowing arbitrary bounded core length and tail onset, no class-uniform noisy finite-prefix recovery modulus tending to zero determines even one fixed nonzero odd cosine amplitude. This remains true for any proposed finite sample budget $M(\delta)$.*

Proof. Fix two distinct admissible profiles $h_0 = 1$ and $h_1 = 1 + \epsilon\cos(k\tau u)$, with $0 < \epsilon < (1 + k^2\tau^2)^{-1/2}$. Let $\beta_i = R_i^{-1}$. Choose a smooth nondecreasing cutoff χ, zero on $(-\infty, 0]$ and one on $[1, \infty)$, flat at both ends. Set

$$b_{i,a}(s) = 0 \ (s \le a), \qquad b_{i,a}(s) = \chi((s-a)/a)\beta_i(s) \ (s > a).$$

This is smooth and nonnegative, increasing on its eventual tail, and agrees exactly with $\beta_i(s)$ for $s \ge 2a$. In particular the eventual inverse is $R_i(y)$ in the original absolute coordinate, without a shifted Fourier phase. The section, rotated by angle -1, lies strictly on one side of the axis; the diverging-twist confinement theorem [3] gives compact resolvent, also with the one-end Dirichlet cap.

Let μ_1 be the first transverse Dirichlet eigenvalue on W. In the fixed cylinder the form is $\int(|\partial_s u - b\partial_\theta u|^2 + |\nabla_W u|^2)$. Transverse Poincare gives $H_{i,a} \ge \mu_1$. On the straight initial cylinder use the first j functions $\sin(n\pi s/a)\phi_1(w)$ and zero-extend. Their maximum Rayleigh quotient is $\mu_1 + (j\pi/a)^2$. Min–max therefore gives

$$\mu_1 \le \lambda_j(H_{i,a}) \le \mu_1 + (j\pi/a)^2.$$

For any $M, \delta > 0$, choose a so that $(M\pi/a)^2 \le 2\delta$. The two prefixes admit a common midpoint observation within δ of each, while their amplitudes differ by ϵ. Any estimator errs by at least $\epsilon/2$ on one input. Choose a after $M(\delta)$ for the final assertion. □

This proves neither exact equality of finite prefixes nor failure of convergence for one fixed known core. Its quantifiers explain why the priors in Theorem 1.1 cannot be discarded.

7 A concrete fundamental-mode certificate and its limits

To separate analytic rates from numerical usefulness, consider

$$R_\epsilon(y) = y[1 + \epsilon\cos(\tau\log y)], \quad |\epsilon| \le 1/20, \qquad b_\epsilon(s) = R_\epsilon^{-1}(s+1).$$

The cap is at $s = 0$ and there is no arbitrary transition. The inverse offset -1 contributes $-6/(4\pi)^{3/2}$ to the constant lower coefficient. Uniform bounds include $3/4 < R' < 5/4$, $4/5 < b' < 4/3$, $|b''| \le 3/b$, $|b'''| \le 40/b^2$. For $B \ge 64$, a direct cap flattening gives $\|u\|_6^2 \le 2^{15}B^2(q[u] + \|u\|_2^2)$ and $k_B(t, X, Y) \le 2^{26}B^3t^{-3/2}$ for $t \le 1$. These are intermediate bounds, not a value of C_r or T_*.

Here is a reproducible analytic global envelope, independent of the asymptotic remainder. Split at physical $s = R > 0$ with a Neumann interface. On the tail $b \ge 20(R+1)/21 = B_R$; every physical axial fiber has length at most $2/B_R$, with at most one Neumann end. Its energy is at least $\pi^2 B_R^2/16$. The core laterally zero-extends into the annular cylinder $0 < s < R$, $1 < r < \sqrt{11}$, all angles, with cap D/N. After multiplication by $\sqrt{r}$, comparison levels are

$$v_{nlm} = \frac{\pi^2(n+1/2)^2}{R^2} + \frac{\pi^2 l^2}{(\sqrt{11}-1)^2} + \frac{m^2}{11} - \frac{1}{4}, \quad n \ge 0,\ l \ge 1,\ m \in \mathbb{Z}.$$

This is a form embedding into a larger Dirichlet-lateral cylinder, not Neumann domain monotonicity. With $R = 2\sqrt{1+x}$ the tail gap exceeds x. Counting the three factors using $\pi > 3$ and $\sqrt{11} < 10/3$ gives bounds $(7/6)(1+x)$, $(7/9)\sqrt{1+x}$, $(23/3)\sqrt{1+x}$, respectively. Thus

$$N(x) \le 7(1+x)^2, \qquad \mathcal{H}(T) \le 35T^4 \quad (T \ge 1). \tag{7.1}$$

The heat bound follows by integrating the counting envelope.

Global physical-fiber Dirichlet Poincare and radial Poincare give for every level

$$\lambda_j \geq \frac{100\pi^2}{441} + \frac{\pi^2}{(\sqrt{11}-1)^2} - \frac{1}{4} > 3.8270183951.$$

Zero-extended sine trials on $0 < s < 1$, $1 < r < 3$, $0 < q < 1$ give

$$\lambda_j \leq \pi^2[j^2 + 5/4 + 1600/361].$$

These are actual whole-tube variational intervals; they are broad, overlap, and are uniform in ϵ. They are not measured point eigenvalues and cannot distinguish members of this family.

The monotonicity of E and the analytic mass bound in [2] give $|M_E(z_1)| < 21/(1600|s_1|)$. Validated complex gamma/zeta evaluation in (5.2) gives

$$|C_1 - (-0.03854139169433 + 0.01124277862658i)| < 0.020813, \qquad 0.01933 < |C_1| < 0.06097.$$

This disk includes uncertainty in the connected corner transform; its phase is not replaced by the lower bound (5.3). A 192-bit validated interval execution of (4.1), using the first 16 intervals, $T = 1$, and 4096 midpoint panels, has a conservative omitted-spectrum error below 3801.20, retained-level error below 12.597 and quadrature error below 0.001265. For this low-temperature execution a valid fallback bias bound is 141, rather than the unevaluated asymptotic $C_r T^{-1/6}$: (7.1), positivity, and $\|F_h\|_\infty < 1$ give $|T^{-3}\mathcal{H}(T) - A_0 T - F_h(\log T)| \leq 35T + 1$ on $1 \leq T \leq 4$. The resulting raw amplitude radius is about 197013. After intersection with its prior the answer remains exactly $[-0.05, 0.05]$. No amplitude, including the benchmark zero, is identified. These numerical diagnostics illustrate a failed informative recovery, not a successful inverse solve. The analytic theorems do not depend on these rounded numerical diagnostics.

8 Scope and next mathematical obligations

The power remainder and finite-cutoff theorem are quantitative analytic results with geometric constants. Fully numerical recovery still requires evaluation of every global comparison constant and threshold, useful complete spectral intervals, a complex multiplier disk, and a validated total error below the desired amplitude scale. The explicit Nash coefficient does not by itself evaluate the global remainder constant. Even $B \geq 64$ in the conservative choice $B = T^{1/6}$ requires $T \geq 64^6$; this is a limitation of the sufficient proof, not an information-theoretic lower bound. Sharper finite-energy comparison or direct validated forward-map inversion may be more useful than optimizing this asymptotic bound.

AI assistance. AI tools assisted with exploration, drafting and internal independent mathematical and arithmetic checks.

A Alternative Gaussian proof of the core estimate

Lemma 2.2 gives an all-time bound for the shifted full Neumann operator $H_N + 1$. Its Davies–Gaffney estimate follows from the exponential-weight energy identity in ambient distance, and is preserved by the positive shift. Coulhon–Sikora [4, Theorem 4.1] gives an off-diagonal bound with linear dependence on the diagonal coefficient CB^3:

$$k_N(t, X, Y) \leq C_\delta B^3 t^{-3/2} \exp[-|X - Y|^2/(4(1+\delta)t)], \quad 0 < t \leq 1.$$

The theorem's fixed-degree polynomial in $|X-Y|/\sqrt{t}$ is absorbed by the displayed fixed Gaussian slack. Unshifting costs at most e. Interior kernel continuity suffices on the open domain. The weighted Gaussian norm in [5, Proposition 4.1] is therefore at most CB^3. In dimension three its interior comparison with the free kernel has no logarithmic loss, and gives (2.6) for full Neumann boundary. Integrating the same CB^2a boundary layer proves (2.5) for H_N. For full Dirichlet boundary the killed-bridge exit bound gives the sharper CB^2t^{-1} error. Form order on this *same* domain places the mixed realization between full Dirichlet and full Neumann, proving (2.5) for it as well. No convex-domain two-term heat theorem is applied to this nonconvex core. This proof and the split-time proof establish one core lemma, not two independent spectral results.

References

[1] S. Schubr, *Heat-Trace Recovery of a Weyl-Blind Twisting Mode*, edition 1 (2026), companion manuscript supplied with this paper.

[2] S. Schubr, *Audibility of all odd blind modes in rapidly twisted polar-strip tubes*, edition 1 (2026), companion manuscript supplied with this paper.

[3] D. Krejčiřík, *Waveguides with asymptotically diverging twisting*, Applied Mathematics Letters **46** (2015), 7–10. arXiv:1407.4028.

[4] T. Coulhon and A. Sikora, *Gaussian heat kernel upper bounds via the Phragmén–Lindelöf theorem*, Proceedings of the London Mathematical Society (3)**96** (2008), 507–544. arXiv:math/0609429.

[5] R. L. Frank and S. Larson, *Uniform bounds for Neumann heat kernels and their traces in convex sets*, preprint (2026). arXiv:2601.07341v1.

[6] L. Li and A. Strohmaier, *Heat kernel estimates for general boundary problems*, Journal of Spectral Theory **6** (2016), 903–919. arXiv:1604.00784.

[7] R. R. Lederman and S. Steinerberger, *Stability estimates for truncated Fourier and Laplace transforms*, Integral Equations and Operator Theory **87** (2017), 529–543. arXiv:1605.03866.

[8] M. Choulli and E. Soccorsi, *An inverse anisotropic conductivity problem induced by twisting a homogeneous cylindrical domain*, Journal of Spectral Theory **5** (2015), 295–329. doi:10.4171/JST/99. Preprint title: *Recovering the twisting function in a twisted waveguide from the DN map*, arXiv:1209.5662v3.

[9] O. Kavian, Y. Kian and E. Soccorsi, *Uniqueness and stability results for an inverse spectral problem in a periodic waveguide*, preprint (2015). arXiv:1504.04267v1.

[10] E. Sultanow, A. Hatziiliou, C. May and N. Kratzke, *Hearing the Edges: Recovering a 3D Rectangular Box from Dirichlet Eigenvalues*, Axioms **15** (2026), 284. doi:10.3390/axioms15040284.