Small balls of flat-edge random Fourier paths: critical, supercritical and crossover regimes

Scott Schubr
Luminary Oddities

November 2026

Abstract

The limiting Kloosterman path of Kowalski and Sawin is a pinned random Fourier series whose coefficients are confined to $[-2,2]$. Its small-ball probabilities near a target path are governed by how expensive it is to place many coefficients close to the edge of their support. We study the pinned series with independent coefficients of density proportional to $\exp[-A(2-|x|)^{-a}]$ and determine the small-ball rates in all three regimes of the edge exponent a. At the critical exponent $a=2$ we prove that the worst-case negative log probability over half-Hölder centers is of exact order $\varepsilon^{-2}\log^2(1/\varepsilon)$ in every L^p, $1<p<\infty$, closing the logarithmic gap left in the uniform-norm estimates for this model. In L^2 we obtain the exact leading constant: it is $A\sigma^2/(4\pi^2)$, where σ is a logarithmic profile of the coefficient margins of the center; this gives explicit constants for exactly saturated centers and for margins closing like $|h|^{-b}$, namely $(4A/\pi^2)\big(2b/(2b+1)\big)^2$. For $a>2$, saturation at a fixed Fourier mode or at the endpoint gives the rate ε^{-a} with an explicit signed variational constant, while power-closing interior margins give the intermediate scale $\varepsilon^{-2(ab+1)/(2b+1)}$ with a closed-form beta-function constant. Finally, when $a\to 2$ with $(a-2)\log(1/\varepsilon)\to\tau$, a single logarithmic profile interpolates continuously between the critical and supercritical laws.

1 Introduction and results

The limiting Kloosterman path of Kowalski and Sawin [2, 3] is a pinned random Fourier series with independent Sato–Tate coefficients supported on $[-2,2]$. Its support is characterized by constraints on the normalized Fourier coefficients of the periodic part, and quantitative estimates for how much probability the path places near a given admissible curve are studied in [4, 6]. The key mechanism is a competition between two costs: approximating a target curve requires coefficients near prescribed values, and values near the edge ± 2 are improbable. To isolate this mechanism we keep the Fourier basis and the pin, and replace the Sato–Tate law by the flat-edge laws

$$f_{A,a}(x)=Z_{A,a}^{-1}\exp[-A(2-|x|)^{-a}]\mathbf{1}_{(-2,2)}(x),\qquad Z_{A,a}=\int_{-2}^{2}\exp[-A(2-|x|)^{-a}]\,dx,\tag{1}$$

with $A>0$ and $a>0$. The uniform-norm estimates in [6, Theorems 5.2 and 7.1] identify $a=2$ as the critical exponent for half-Hölder centers, with upper and lower bounds of orders $\varepsilon^{-2}\log^3(1/\varepsilon)$ and $\varepsilon^{-2}\log^2(1/\varepsilon)$. This paper determines the exact rates. At the critical exponent a square-function estimate combines the Fourier errors across dyadic scales and gives the sharp squared logarithm in every finite L^p (Theorem 1); in L^2 comparisons on the actual pinned path give the exact leading constant (Theorem 2). Above the critical exponent the rate becomes a power of ε^{-1},

with two distinct mechanisms (Theorems 4 and 5); near it, a logarithmic crossover connects the two descriptions (Theorem 6). Related scalar small-deviation results for weighted sums of positive variables are due to Rozovskii [5, 7]; here the linear drift and the complex constant fixed by the pin act on every Fourier mode and must be retained.

1.1 The pinned series

Write $e(t) = e^{2\pi i t}$, $v_0(t) = t$ and $v_h(t) = (e(ht) - 1)/(2\pi i h)$ for $h \neq 0$. Let $(X_h)_{h\in\mathbb{Z}}$ be independent with density (1), and put

$$K_{A,a}(t) = \sum_{h\in\mathbb{Z}} X_h v_h(t), \qquad 0 \le t \le 1, \tag{2}$$

using symmetric partial sums; these converge almost surely uniformly (Section 2). Norms refer to Lebesgue measure on $[0,1]$. A continuous $g : [0,1] \to \mathbb{C}$ is *admissible* if $g(0) = 0$, $\xi = g(1) \in [-2,2]$ is real, and the periodic part $u(t) = g(t) - \xi t$ satisfies

$$\widehat{u}(h) = \int_0^1 u(t)e(-ht)\,dt, \qquad \beta_h = 2\pi i h\,\widehat{u}(h) \in [-2,2] \quad (h \neq 0). \tag{3}$$

Write $m_h = 2 - |\beta_h|$ for the coefficient margins. For $L \ge 1$ let H_L be the set of admissible g with $|g(t) - g(s)| \le L|t - s|^{1/2}$. Throughout, $J = \log(1/\varepsilon)$.

1.2 The critical exponent

For $a = 2$ write $K_A = K_{A,2}$ and

$$B_p(g,\varepsilon) = -\log \mathbb{P}(\|K_A - g\|_p < \varepsilon).$$

Theorem 1. *Fix $A > 0$ and $L \ge 1$. For every fixed $1 < p < \infty$ there are $0 < c \le C < \infty$ such that*

$$c\varepsilon^{-2}J^2 \le \sup_{g\in H_L} B_p(g,\varepsilon) \le C\varepsilon^{-2}J^2.$$

The lower bound holds at the fixed center $g_q \in H_L$ of [6, Proposition 5.1], defined in Section 2. At $p = 1$ the same uniform upper bound holds and $B_1(g_q,\varepsilon) \ge c\varepsilon^{-2}J$.

For a fixed admissible g, write $m_h = 2 - |\beta_h|$ for $h \neq 0$. For $0 < \theta < 1$ define

$$S_\pm(\varepsilon,\theta) = \sum_{\substack{1\le|h|\le\lfloor\varepsilon^{-2}\rfloor \\ m_h\le\varepsilon^{1\pm\theta}\sqrt{|h|}}} \frac{1}{|h|}, \quad \sigma_{\mathrm{lo}} = \lim_{\theta\downarrow0}\liminf_{\varepsilon\downarrow0} \frac{S_+(\varepsilon,\theta)}{J}, \quad \sigma_{\mathrm{hi}} = \lim_{\theta\downarrow0}\limsup_{\varepsilon\downarrow0} \frac{S_-(\varepsilon,\theta)}{J}.$$

The outer limits exist by monotonicity. In particular $0 \le \sigma_{\mathrm{lo}} \le \sigma_{\mathrm{hi}} \le 4$.

Theorem 2. *For every fixed continuous admissible g, set $C_\varepsilon(g) = \varepsilon^2 B_2(g,\varepsilon)/J^2$. Then*

$$\frac{A\sigma_{\mathrm{lo}}^2}{4\pi^2} \le \liminf_{\varepsilon\downarrow0} C_\varepsilon(g) \le \limsup_{\varepsilon\downarrow0} C_\varepsilon(g) \le \frac{A\sigma_{\mathrm{hi}}^2}{4\pi^2}.$$

If the two profiles agree at σ, then $C_\varepsilon(g)$ converges to $A\sigma^2/(4\pi^2)$.

Corollary 3. *The following consequences concern the same normalization.*

1. *Suppose $m_h \in \{0\} \cup [m_0, 2]$ for some $m_0 > 0$, and the two-sign harmonic density of exact saturation exists:*
$$\kappa = \lim_{R\to\infty} \frac{1}{\log R} \sum_{\substack{1\le|h|\le R\\ m_h=0}} \frac{1}{|h|}.$$
Then $C_\varepsilon(g) \to A\kappa^2/\pi^2$. In particular $C_\varepsilon(g_q) \to 4A/(\pi^2 q^2)$.

2. *Suppose $m_h \asymp |h|^{-b}$ for $b > 0$ on both tails, with finitely many exceptional coefficients allowed. Then*
$$C_\varepsilon(g) \longrightarrow \frac{4A}{\pi^2}\left(\frac{2b}{2b+1}\right)^2.$$

3. *For every $L \ge 1$, one has $\limsup_{\varepsilon\downarrow 0} \sup_{g\in H_L} C_\varepsilon(g) \le 4A/\pi^2$. For $L \ge 80$, the limit of this supremum equals $4A/\pi^2$.*

1.3 Supercritical exponents

For general a write $B_{A,a}(g,\varepsilon) = -\log \mathbb{P}(\|K_{A,a} - g\|_2 < \varepsilon)$.

Let $S = \{h \neq 0 : m_h = 0\}$, $s_h = \operatorname{sgn}(\beta_h)$ for $h \in S$, and $w_h = (2\pi h)^{-2}$. For $z = (z_h)_{h\in S}$ with $z_h > 0$ and a scalar y, set
$$\mathcal{E}(y,z) = \frac{y^2}{4} + \sum_{h\in S} w_h(s_h z_h + y)^2, \qquad \mathcal{I}_a(y,z) = \sum_{h\in S} z_h^{-a} + P_a(y). \tag{4}$$
If the endpoint is interior, y ranges over $\mathbb{R}$ and $P_a(y) = 0$. If $|\xi| = 2$, set $s_0 = \operatorname{sgn}(\xi)$, require $-s_0 y > 0$, and put $P_a(y) = (-s_0 y)^{-a}$. Let $I_*(a)$ be the infimum of $\mathcal{I}_a$ over this domain under $\mathcal{E} \le 1$.

Theorem 4 (Exact saturation rate)**.** *For every fixed continuous admissible g and $a > 2$,*
$$\lim_{\varepsilon\downarrow 0} \varepsilon^a B_{A,a}(g,\varepsilon) = AI_*(a).$$
The infimum is attained and finite. It is positive precisely when the endpoint or at least one nonzero normalized coefficient is saturated.

For a wholly interior center the coefficient in Theorem 4 is zero. The next result determines the scale and constant when the margins close as powers on the two frequency tails.

Theorem 5 (Power-closing interior margins)**.** *Fix $a > 2$ and $b > 0$, and let g be a fixed continuous admissible center. Suppose $\xi \in (-2, 2)$, $m_h > 0$ for all $h \neq 0$, and*
$$m_h \sim c_+ h^{-b}, \qquad m_{-h} \sim c_- h^{-b} \quad (h \to +\infty), \qquad c_+, c_- > 0.$$
Set $q = 2(ab+1)/(2b+1)$, so $2 < q < a$. Then
$$\lim_{\varepsilon\downarrow 0} \varepsilon^q B_{A,a}(g,\varepsilon) = A\mathcal{C}.$$
Let $d = (a+2)b + 2$, $\rho = (ab+1)/d$, and $p_* = \rho/(1-\rho)$. The constant is
$$\mathcal{C} = \left[\frac{a}{ab+1}\,\mathrm{B}\!\left(1+\rho, \frac{a-2}{d}\right)\left(\frac{p_*}{2\pi^2 a}\right)^{\rho}\left(c_+^{-(a-2)/d} + c_-^{-(a-2)/d}\right)\right]^{1/(1-\rho)}, \tag{5}$$
where $\mathrm{B}(x,y) = \int_0^1 t^{x-1}(1-t)^{y-1}\,dt$.

1.4 A simultaneous approach to the critical exponent

Suppose $a_\varepsilon \to 2$ and $(a_\varepsilon - 2)\log(1/\varepsilon) \to \tau \in \mathbb{R}$. For $J = \log(1/\varepsilon)$ and $0 < \theta < 1$, define

$$S^\tau_\pm(\varepsilon,\theta) = \sum_{\substack{1\le|h|\le\lfloor\varepsilon^{-2}\rfloor \\ m_h\le\varepsilon^{1\pm\theta}\sqrt{|h|}}} \frac{\exp[-\tau\log|h|/(4J)]}{|h|}, \qquad \begin{aligned}\sigma^\tau_{\mathrm{lo}} &= \lim_{\theta\downarrow0}\liminf_{\varepsilon\downarrow0} S^\tau_+(\varepsilon,\theta)/J,\\ \sigma^\tau_{\mathrm{hi}} &= \lim_{\theta\downarrow0}\limsup_{\varepsilon\downarrow0} S^\tau_-(\varepsilon,\theta)/J.\end{aligned}$$

The outer limits exist by monotonicity, and the profiles are bounded for fixed τ.

Theorem 6 (Critical crossover)**.** *For a fixed continuous admissible g, set*

$$C_\varepsilon = \varepsilon^2 B_{A,a_\varepsilon}(g,\varepsilon)/J^2.$$

Then

$$\frac{Ae^\tau(\sigma^\tau_{\mathrm{lo}})^2}{4\pi^2} \le \liminf_{\varepsilon\downarrow0} C_\varepsilon \le \limsup_{\varepsilon\downarrow0} C_\varepsilon \le \frac{Ae^\tau(\sigma^\tau_{\mathrm{hi}})^2}{4\pi^2}.$$

Equality of the two profiles gives the corresponding exact limit. If $m_h \in \{0\}\cup[m_0,2]$ with $m_0 > 0$ and

$$\kappa = \lim_{R\to\infty}\frac{1}{\log R}\sum_{\substack{1\le|h|\le R\\ h\in S}}|h|^{-1}$$

exists, the limit is

$$\frac{4A\kappa^2}{\pi^2}\left(\frac{e^{\tau/2}-1}{\tau}\right)^2. \tag{6}$$

If $m_{\pm h} \asymp h^{-b}$ for $b > 0$, the limit is

$$\frac{16A}{\pi^2}\left(\frac{e^{\tau b/(2b+1)}-1}{\tau}\right)^2. \tag{7}$$

The expressions at $\tau = 0$ are given by continuity.

2 Preliminaries

The coefficients in (1) are centered and bounded by two. On a dyadic block $R < |h| \le 2R$, Hoeffding's exponential moment bound and a grid of $O(R)$ points give an expected sup norm $O(\sqrt{\log(eR)/R})$ for the largest symmetric partial polynomial in the block. Indeed, the centered unpinned polynomial has pointwise sub-Gaussian proxy $O(R^{-1})$. Take the union bound over the $O(R^2)$ partial polynomials and grid points, and apply the Bernstein inequality to pass from the grid to the circle. Pinning costs at most a factor two by subtraction of the value at zero. Summing these expected bounds over dyadic R gives the asserted almost sure uniform convergence, as in [6, Section 2]. The omitted tail beyond N is independent of the retained coefficients and has expected squared L^2 norm $O(N^{-1})$. This follows from $\|v_h\|_2^2 = 2/(2\pi h)^2$ and the identity

$$\mathbb{E}\Big\|\sum_h Y_h v_h\Big\|_2^2 = \sum_h \mathbb{E}|Y_h|^2\|v_h\|_2^2 \tag{8}$$

for independent centered coefficients, with passage to the L^2 limit for the tail. These bounds are uniform when a varies.

The periodic part of $K_{A,a}$ has purely imaginary mean, since each v_h with $h \neq 0$ does. Also $\widehat{u}(0)$ is purely imaginary: evaluate a smooth real even Fourier approximate identity at zero, where the nonzero coefficients give imaginary contributions, and pass to $u(0) = 0$. Thus with $\Delta = X_0 - \xi$,

$$\Re \widehat{K_{A,a} - g}(0) = \Delta/2, \qquad 2\pi i h \, \widehat{K_{A,a} - g}(h) = X_h - \beta_h - \Delta \quad (h \neq 0). \tag{9}$$

The second identity uses $\widehat{t}(h) = -1/(2\pi i h)$. On the ball, $|\Delta| \leq 2\varepsilon$. Consequently, if F is a finite set of frequencies with $\min_{h \in F} |h| \geq R$, the Euclidean norm of $(X_h - \beta_h)_{h \in F}$ with weights w_h is at most $\varepsilon(1 + O(R^{-1/2}))$. This follows by projecting away the drift on F and using $\sum_{|h| \geq R} w_h = O(R^{-1})$.

For $h \neq 0$, we have $|v_h(t)| \leq 1/(\pi|h|)$ and $\|v_h\|_2^2 = 2/(2\pi h)^2$. The flat-edge law is even, so its coefficients are centered. For a finite sum with independent centered real Y_h satisfying $|Y_h| \leq r_h$, Hoeffding's exponential moment inequality applied to its real and imaginary parts gives

$$\mathbb{E} \left\| \sum_h Y_h v_h \right\|_p^p \leq C_p \left(r_0^2 + \sum_{h \neq 0} r_h^2 / h^2 \right)^{p/2} \qquad (1 \leq p < \infty). \tag{10}$$

An absent frequency contributes zero. Indeed the two parts at each t are sub-Gaussian with the indicated common variance bound up to an absolute factor. Integrating their p-th moment bounds in t proves (10).

For the fixed center, take the Rudin–Shapiro signs r_k^{RS}, indexed from 1, satisfying

$$\sup_{|z|=1} \left| \sum_{k=1}^{m} r_k^{\mathrm{RS}} z^k \right| \leq C_{\mathrm{RS}} \sqrt{m}.$$

The complementary-polynomial identity gives this bound by decomposing a prefix into dyadic blocks, as in [6, Eq. (5.2)]. Define

$$g_1(t) = \frac{2}{\pi} \sum_{k \geq 1} \frac{r_k^{\mathrm{RS}}}{k} \sin(2\pi k t), \qquad g_q(t) = q^{-1} g_1(qt), \quad q \in \mathbb{N}.$$

Abel summation bounds the sine tail beyond m in sup norm by $O(m^{-1/2})$. The derivative of its degree-m partial sum has norm $O(\sqrt{m})$. Splitting increments at $m \asymp |t-s|^{-1}$ thus gives half-Hölder continuity. The calculation in [6, Proposition 5.1] gives the admissible unit-interval Hölder constant 80 for g_1. Its periodic extension has circular half-Hölder constant at most $80\sqrt{2}$, by splitting across the pin and applying $\sqrt{x} + \sqrt{y} \leq \sqrt{2(x+y)}$. Thus one may choose $q = \max(1, \lceil 2 \cdot 80^2 / L^2 \rceil)$ for arbitrary $L \geq 1$. The center g_q has zero mean and zero endpoints, and normalized coefficients equal to $2r_k^{\mathrm{RS}}$ at both signs of qk, with zero coefficients elsewhere.

Lemma 7 (Correction at the pin)**.** *For an admissible $g = \xi t + u$, let a finite set of real numbers b_h, $0 < |h| < N$, prescribe a Fourier approximation, and put*

$$E(t) = u(t) - \widehat{u}(0) - \sum_{0 < |h| < N} \frac{b_h}{2\pi i h} e(ht).$$

Then $E(0)$ is purely imaginary. On adding the pinned trigonometric polynomial $E(0)(P_N(t) - 1)$, with

$$P_N(t) = \frac{1}{N} \sum_{h=N}^{2N-1} \cos(2\pi h t),$$

to $\sum_{0<|h|<N} b_h v_h(t)$*, the error with respect to* u *becomes* $E(t) - E(0)P_N(t)$*. The extra normalized coefficients are real and have absolute value at most* $2\pi|E(0)|$*. Moreover,* $\|P_N\|_2 = (2N)^{-1/2}$*,* $\|P_N\|_p \le C_p N^{-1/p}$ *for* $1 < p < \infty$*, and* $\|P_N\|_1 \le C\log(eN)/N$*.*

Proof. Periodic smooth approximate identities converge at 0 to $u(0) = 0$. Their nonzero modes contribute purely imaginary values there, by (3) and real, even multipliers. Consequently $\widehat{u}(0)$ is purely imaginary, as is $E(0)$. The error of the original pinned approximation is $E(t) - E(0)$, by evaluating its unpinned counterpart at zero. The resulting extra normalized coefficient at $N \le |h| < 2N$ is $2\pi i h E(0)/(2N)$. Parseval gives the L^2 norm. The geometric sum gives $|P_N(t)| \le \min(1, C/(N\,\mathrm{dist}(t,\mathbb{Z})))$, and integration gives the other norm bounds. □

We use a smooth even multiplier supported on $|h| < N$, equal to one on $|h| \le N/2$, with values in $[0,1]$. Such cutoffs converge uniformly on a fixed continuous periodic u, since periodization of the rapidly decaying inverse transform gives a bounded approximate identity. Prescribing normalized coefficients bounded by two gives mean-zero L^2 error $O(N^{-1/2})$ from $|h| > N/2$, by (9) when the drift is bounded as well.

Finally, suppose that finitely many centers b_h are surrounded by intervals of half-width $r_h > 0$, with remaining margins $\ell_h = 2 - |b_h| - r_h > 0$. Relative to independent uniforms on these intervals, with the independent omitted tail retaining its law, the true probability of any event is at least the comparison probability times

$$\prod_h (2r_h/Z_{A,a}) \exp[-A\ell_h^{-a}]. \tag{11}$$

Conversely, for finitely many coordinates and any prescribed signs s_h, set $d_h = 2 - s_h X_h$. Then $2 - |X_h| \le d_h$. If an event implies $\sum d_h^{-a} \ge W$ on these coordinates, the true probability is at most $(4/Z_{A,a})^n e^{-AW}$, where n is their number. This uses the density bound on a box of volume 4^n.

3 Sharp order in finite L^p

3.1 A uniform lower bound on probability

Fix $p < \infty$, $p \ge 1$. Let η be smooth, with values in $[0,1]$, equal to one on $[0,1/2]$ and zero on $[3/4,\infty)$. For $0 < |h| < N$ put

$$w_h = (1 - \sqrt{|h|/N})\eta(|h|/N), \qquad b_h = w_h \beta_h,$$

and set $w_h = 0$ for $|h| \ge N$. Use the error E of Lemma 7. The periodic function u is half-Hölder on $\mathbb{T}$ with constant $O(L+2)$.

Let Δ_j be a smooth dyadic Fourier projection at scale 2^j, $j \ge 0$, with a partition of unity on the nonzero integer frequencies. We use the periodic Littlewood–Paley inequality [1, Chaps. 3 and 6], which for mean-zero F and $2 \le p < \infty$ gives

$$\|F\|_p \le C_p \left\| (\sum_j |\Delta_j F|^2)^{1/2} \right\|_p \le C_p (\sum_j \|\Delta_j F\|_p^2)^{1/2}.$$

For the error E, the rescaled shell multiplier of $\Delta_j E$ acting on u is an annular symbol with any fixed number of derivatives of size $O(\min(1, \sqrt{2^j/N}))$. On low shells this follows from $1 - w_h = \sqrt{|h|/N}$, and on the transition and high shells it follows from the smooth cutoff. On periodizing its inverse transform by Poisson summation, the kernel has zero integral and its absolute half-moment

with circular distance is $O(2^{-j/2}\min(1,\sqrt{2^j/N}))$. Indeed the inverse transform of the rescaled symbol has rapid decay with that same prefactor, and the scaling of the moment contributes $2^{-j/2}$. Convolving against $u(t-y)-u(t)$ now gives

$$\|\Delta_j E\|_\infty \le C_L 2^{-j/2}\min(1,\sqrt{2^j/N}). \tag{12}$$

The lowest nonzero frequencies are included by the initial annular projection. Summation of squares yields $\|E\|_p \le C_{p,L}\sqrt{\log(eN)/N}$ for $p \ge 2$. For $p < 2$ the same bound follows from L^2. Summing (12) directly also gives $|E(0)| \le C_L \log(eN)/\sqrt{N}$.

Use Lemma 7 to correct the pin. For every fixed $p < \infty$ the correction has L^p norm bounded by $C_{p,L}\sqrt{\log(eN)/N}$. Its coefficient centers tend uniformly to zero. Choose $N = \lceil D_{p,L}\varepsilon^{-2}\log(e/\varepsilon)\rceil$ for a sufficiently large fixed $D_{p,L}$. Shift ξ to $(1-c_p\varepsilon)\xi$ and use width $r_0 = c_p\varepsilon$ for a sufficiently small fixed $c_p > 0$. With these choices the deterministic error in the pinned approximation, including the drift, is at most $\varepsilon/2$.

Use widths $r_h = \min(1/2,\sqrt{|h|/N})$ for $0 < |h| < N$ and width $1/2$ on the correction band $N \le |h| < 2N$. For the low band the remaining distance from the edge is at least $\sqrt{|h|/N}$, since $|\beta_h| \le 2$ and $0 \le w_h \le 1 - \sqrt{|h|/N}$. For the correction band it is bounded below by a fixed positive constant. At $h = 0$ the remaining margin is at least $c_p\varepsilon$. Formula (11) then costs at most $C_{A,p}(N\log(eN) + \varepsilon^{-2})$ in negative logarithm. Here $\sum_{h<N} N/h = O(N\log(eN))$, and the logarithms of all widths have total magnitude $O(N\log(eN) + \log(1/\varepsilon))$.

For the finite interval perturbations, the variance proxy in (10) is $O(c_p^2\varepsilon^2 + \log(eN)/N)$. For the omitted tail it is $O(1/N)$, since $|X_h| \le 2$. By taking $D_{p,L}$ large and c_p small, the probability that each of these errors has norm at most $\varepsilon/8$ is at least $1/2$ by Markov's inequality. They are independent under the comparison. This leaves strict slack inside radius ε. The probability bound (11) proves the uniform upper bound on B_p in Theorem 1, including $p = 1$.

3.2 A lower bound on the negative log probability

Let $g = g_q$ and $D(t) = K_A(t) - X_0 t - g_q(t)$. Every nonzero v_h has purely imaginary integral, so on the small-ball event $|X_0| \le 2\varepsilon$ and $\|D\|_p \le 3\varepsilon$. Set $d_k = 2 - r_k^{\mathrm{RS}} X_{qk} > 0$ and

$$T_M(t) = \sum_{k=1}^{M} \frac{r_k^{\mathrm{RS}}}{\sqrt{k}} e(-qkt).$$

For $1 < p < \infty$ and conjugate exponent p', $\|T_M\|_{p'} \le C_{p'}\sqrt{\log(eM)}$. To prove it, first take $q = 1$, since the integer dilation preserves the norm. The Rudin–Shapiro prefix bound shows that partial sums of the signs times $e(-kt)$ on each dyadic shell are $O(2^{j/2})$. Abel summation with the smooth dyadic cutoff and $k^{-1/2}$, including the final truncation at M, gives a uniform $O(1)$ shell sup bound. The square-function inequality, or the L^2 bound if $p' < 2$, gives the desired estimate. For $p = 1$, summing the shell sup bounds gives $\|T_M\|_\infty \le C\log(eM)$.

In the pairing $\int_0^1 DT_M$, the mode qk contributes $id_k/(2\pi q k^{3/2})$, and no other modes contribute. Therefore $\sum_{k\le M} d_k/k^{3/2} \le C_{p,q}\varepsilon\sqrt{\log(eM)}$ for $p > 1$. For $p = 1$, replace the square root by $\log(eM)$. Hölder's inequality gives

$$\left(\sum_{k\le M}\frac{1}{k}\right)^3 \le \left(\sum_{k\le M} d_k^{-2}\right)\left(\sum_{k\le M}\frac{d_k}{k^{3/2}}\right)^2.$$

Thus on the event the sum of the inverse squares is at least $c_{p,q}\varepsilon^{-2}\log^2(eM)$ for $p>1$, and at least $c_q\varepsilon^{-2}\log(eM)$ for $p=1$. Since $2-|X_{qk}|\le d_k$, the joint density of these selected coefficients is at most $Z_A^{-M}\exp(-A\sum_{k\le M}d_k^{-2})$. The event projects into a set of volume at most 4^M. Choosing $M=\lfloor\varepsilon^{-2}\rfloor$ proves the claimed lower bounds on $B_p(g_q,\varepsilon)$.

4 The exact L^2 profile

4.1 A necessary spectral event

Let $g=\xi t+u$ be any fixed continuous admissible center. As in Lemma 7, $\Re\widehat{u}(0)=0$. Consequently $\Re\int_0^1(K_A-g)=(X_0-\xi)/2$, and the small-ball event implies $|X_0-a|\le 2\varepsilon$. For any finite set I of nonzero frequencies with $\min_{h\in I}|h|\ge R_0$, project K_A-g and the drift $(X_0-\xi)t$ onto I. Since $\widehat{t}(h)=-1/(2\pi ih)$, Parseval gives

$$\left(\sum_{h\in I}\frac{|X_h-\beta_h|^2}{(2\pi h)^2}\right)^{1/2}\le\varepsilon(1+O(R_0^{-1/2})). \tag{13}$$

When $\beta_h\neq 0$, put $d_h=2-\operatorname{sgn}(\beta_h)X_h$. Then $|d_h-m_h|=|X_h-\beta_h|$ and $2-|X_h|\le d_h$. If $R_0\to\infty$ and $\sum_{h\in I}m_h^2/h^2=o(\varepsilon^2)$, (13) and Cauchy–Schwarz give

$$\sum_{h\in I}d_h^{-2}\ge\frac{(\sum_{h\in I}|h|^{-1})^2}{4\pi^2(1+o(1))\varepsilon^2}. \tag{14}$$

Integrating the upper product-density bound as in the finite-p proof incurs at most $O_A(|I|)$ in the logarithm, which is negligible at scale $\varepsilon^{-2}J^2$ for $I\subset\{1\le|h|\le\varepsilon^{-2}\}$.

For fixed $\theta>0$ take the frequencies defining $S_+(\varepsilon,\theta)$, discarding those with $|h|<J^2$. The discarded harmonic sum is $O(\log J)$. On the chosen set $m_h\le\varepsilon^\theta$, so the normalizations d_h are defined for small ε. Moreover $\sum_I m_h^2/h^2\le O(\varepsilon^{2+2\theta}J)=o(\varepsilon^2)$. Applying (14), taking the lower limit in ε and then sending $\theta\downarrow 0$ proves the lower bound in Theorem 2. Empty selections contribute the trivial nonnegative lower bound.

4.2 Sufficient boxes and their cost

Let $M=\lfloor\varepsilon^{-2}J^{1/2}\rfloor$ and $\delta=J^{-1/16}$. Choose a set $B\subset\{1\le|h|\le M/2\}$ of coefficients to push inward. Put

$$S=\sum_{h\in B}|h|^{-1},\quad D_\varepsilon=S+J^{3/4},\quad \alpha=\frac{2\pi(1-4\delta)\varepsilon}{\sqrt{D_\varepsilon}},\quad a_h=\alpha\sqrt{|h|}\quad(h\neq 0).$$

Here $J^{3/4}\le D_\varepsilon=O(J)$, so $cJ^{-1/4}\le\alpha\sqrt{M}\le CJ^{-1/8}$. Take a smooth even multiplier $\eta_0\in[0,1]$, equal to one on $[-1/2,1/2]$ and zero outside $(-1,1)$. For $h\in B$ and for $M/2<|h|<M$, project the real coefficient $\eta_0(h/M)\beta_h$ onto the interval $[-2+a_h,2-a_h]$, giving b_h. For the other low frequencies set $b_h=\beta_h$. Above the cutoff there are initially no prescribed coefficients.

Let E be the mean-zero unpinned error of this approximation. At low frequencies a clipped coefficient changes by at most a_h, so the L^2 error from low frequencies is at most $\alpha\sqrt{S}/(2\pi)$. The error from $|h|>M/2$ is $O(M^{-1/2})=o(\delta\varepsilon)$ by the coefficient bound and Parseval. The operator with multiplier $\eta_0(h/M)$ is an approximate identity, by periodizing its rapidly decaying rescaled inverse transform. It converges uniformly to u, and the convergence is uniform in H_L for fixed L. The

additional clipping changes the unpinned sum at zero by at most $O(\sum_{1\le h\le M}\alpha/\sqrt{h}) = O(\alpha\sqrt{M})$. Therefore $E(0) = o(1)$, uniformly in H_L as well. Apply Lemma 7 on $M \le |h| < 2M$. The extra coefficients tend to zero, while the change in L^2 error is $O(|E(0)|M^{-1/2}) = o(\delta\varepsilon)$. Shift ξ to $(1-\delta^2\varepsilon)\xi$. The resulting deterministic error has norm at most $(1-4\delta)\varepsilon + o(\delta\varepsilon)$.

Give the nonzero prescribed coefficients the comparison half-width $r_h = \delta^4 a_h$ for $1 \le |h| < 2M$, and use $r_0 = \delta^2\varepsilon$ at zero. With U the centered finite perturbation, the variance identity gives

$$\mathbb{E}\|U\|_2^2 \lesssim \varepsilon^2\delta^4 + \delta^8\alpha^2\log(eM) \lesssim \varepsilon^2(\delta^4 + \delta^8 J^{1/4}) = O(\varepsilon^2 J^{-1/4}) = o(\delta^2\varepsilon^2).$$

The independent omitted tail has squared norm expectation $O(M^{-1}) = o(\delta^2\varepsilon^2)$. Markov's inequality puts both errors below $\delta\varepsilon$ with joint comparison probability $1 - o(1)$. This leaves a strict small-ball event.

On B and on the taper band the remaining coefficient margin is at least $(1-\delta^4)a_h$. The negative log density cost from B is therefore at most

$$\frac{ASD_\varepsilon}{4\pi^2(1-\delta^4)^2(1-4\delta)^2\varepsilon^2}. \tag{15}$$

The taper band costs $O_A(\alpha^{-2}) = O_A(\varepsilon^{-2}J)$ because its harmonic sum is bounded. The correction band costs $O_A(M)$. The shifted endpoint has remaining margin at least $\delta^2\varepsilon$ even at $\xi = \pm 2$, costing $O_A(\delta^{-4}\varepsilon^{-2})$. The logarithms of the nonzero widths sum in absolute value to $O(M\log J)$: use $\sum_{h\le 2M-1}\log h = \log((2M-1)!)$ and the bounds on $\alpha\sqrt{M}$. The width at zero contributes $O(J + \log J)$ and the normalization contributes $O_A(M)$. All of these costs are $o(\varepsilon^{-2}J^2)$, apart from (15) and the density costs of the unchanged low frequencies. As before the comparison is only over finitely many coefficients and the actual independent tail.

For fixed $\theta > 0$, choose B among the low modes to consist of those with $m_h \le \varepsilon^{1-\theta}\sqrt{|h|}$. The sum S differs from $S_-(\varepsilon,\theta)$ by at most $O(\log J)$, because only the range $\varepsilon^{-2} < |h| \le M/2$ can contribute a difference for small ε. At an unchanged low frequency,

$$\frac{r_h}{m_h} \le C\delta^4\frac{\varepsilon^\theta}{\sqrt{D_\varepsilon}} = o(1).$$

Thus these intervals are contained in the support and their total negative log density cost is bounded by

$$O_A\left(\sum_{\substack{1\le|h|\le M/2\\ h\notin B}} m_h^{-2}\right) = O_A(\varepsilon^{-2+2\theta}\log(eM)) = o(\varepsilon^{-2}J^2).$$

In (15), $SD_\varepsilon/J^2 = (S/J)^2 + o(1)$ because $S = O(J)$. The upper limit in ε, followed by $\theta \downarrow 0$, proves the upper bound in Theorem 2.

5 Evaluation of the profile

For exact saturation with a fixed interior margin m_0, all saturated frequencies contribute to both $S_\pm$. For fixed $\theta > 0$, the threshold of S_+ is bounded by ε^θ throughout its range, so eventually it admits only saturated frequencies. Their harmonic sum divided by J tends to 2κ. An unsaturated frequency counted by S_- must have $|h| \ge m_0^2\varepsilon^{-2+2\theta}$. The harmonic sum of both signs in that additional range is at most $4\theta J + O_{m_0}(1)$. Thus the two profiles agree at 2κ, including $\kappa = 0$. For g_q, the saturated frequencies are exactly $h = \pm qk$, giving $\kappa = 2/q$.

When $m_h \asymp |h|^{-b}$, $b > 0$, the logarithmic lower cutoffs for the two sums are

$$u_\pm = \frac{1 \pm \theta}{b + 1/2}.$$

Indeed the comparability constants place each cutoff between fixed positive multiples of $\varepsilon^{-u_\pm}$. For sufficiently small $\theta > 0$ we have $0 < u_\pm < 2$. The harmonic sums over both signs then satisfy $S_\pm(\varepsilon, \theta)/J \to 2(2 - u_\pm)$. The finite exceptions have bounded total harmonic weight. The common profile is $2(2 - 2/(2b+1))$, proving the second assertion of Corollary 3. For example, the sine series with normalized coefficients $\beta_{\pm h} = (2 - ch^{-b}) r_h^{\mathrm{RS}}$, $0 < c < 1$, is admissible and half-Hölder. Its correction from g_1 on a scale 2^j has sup norm $O(2^{-j(b+1/2)})$ by Abel summation, and derivative bound $O(2^{j(1/2-b)})$. Splitting increments at $2^j \asymp |t - s|^{-1}$ yields the half-Hölder bound, also when $b = 1/2$, where the low-frequency derivative bound grows only logarithmically.

For the uniform upper coefficient, the sufficient boxes may be used with every low mode in B. Then there are no unchanged modes to charge, $S/J \to 4$, and all deterministic and probabilistic error estimates are uniform over H_L. The fixed center $g_1 \in H_L$ for $L \geq 80$ attains the limiting coefficient by exact saturation. This proves the final assertion of Corollary 3.

6 The fixed saturation rate

6.1 Necessary cost and finite relaxations

On the ε-ball, take $y = \Delta/\varepsilon$ and $z_h = (2 - s_h X_h)/\varepsilon$ for $h \in S$. At a saturated endpoint $y = -s_0(2 - s_0 X_0)/\varepsilon$, as required. Equation (9) and Parseval give the constraint $\mathcal{E}(y, z) \leq 1$ when one keeps the real mean and the modes in S. For a finite $F \subset S$, let I_F be the infimum with only these modes in (4), retaining the endpoint condition. The finite coefficient cost is at least $\varepsilon^{-a} I_F$, where the endpoint is counted only if saturated. At an interior endpoint, simply allow any $y \in [-2, 2]$ in the finite necessary event. This condition is measurable by continuity of the minimum over that compact interval, so the bound following (11) applies. The number of retained coordinates is fixed. We obtain $\liminf_{\varepsilon \downarrow 0} \varepsilon^a B_{A,a}(g, \varepsilon) \geq A I_F$.

A feasible point of finite cost exists. One can take $z_h = |h|^{2/(a+2)}$, and $y = 0$ for an interior endpoint or $y = -s_0$ for a saturated endpoint, and then scale all coordinates by a common small factor. Both relevant tails converge because $2a/(a + 2) > 1$. The finite infima are nondecreasing along an exhaustion and bounded above by $I_*(a)$. In approximate finite minimizers with bounded cost, each retained z_h is bounded away from zero by its cost, and bounded above by $2 + 2\pi|h|$ by the energy and $|y| \leq 2$. A saturated y is also kept away from zero by the cost. Take a diagonal subsequence and use Fatou on the nonnegative series. The limit is feasible with cost at most the limiting finite infima. Thus these infima tend to $I_*(a)$, and the infimum is attained. Any saturation gives a positive cost from its individual upper bound on the deficit; with no saturation one can take $y = 0$ and the empty vector.

6.2 Sufficient boxes

Take a feasible near minimizer (y, z) and scale it by a factor $t < 1$, leaving fixed strict energy slack while multiplying its cost by t^{-a}. Denote the scaled values again by (y, z). Keep the designated saturated deficits only through a large fixed frequency R. For all larger frequencies use $e_h = |h|^{2/(a+2)}$. Both $\sum_{|h|>R} (e_h + |y|)^2/h^2$ and $\sum_{|h|>R} e_h^{-a}$ tend to zero as $R \to \infty$. Choose $N = \lceil \varepsilon^{-\gamma} \rceil$ with $2 < \gamma < (a + 2)/2$. Then

$$N^{-1/2} = o(\varepsilon), \qquad \varepsilon N^{2/(a+2)} = o(1), \qquad N \log(1/\varepsilon) = o(\varepsilon^{-a}).$$

Prescribe $b_0 = \xi + \varepsilon y$. For $0 < |h| \le R$, prescribe $\beta_h - \varepsilon s_h z_h$ when $h \in S$, and $\beta_h + \varepsilon y$ otherwise. For $R < |h| < N$, project $\beta_h + \varepsilon y$ onto $[-2 + \varepsilon e_h, 2 - \varepsilon e_h]$ and multiply by the smooth cutoff from Section 2. The projection displaces its input by at most $\varepsilon(e_h + |y|)$. Multiplication by the cutoff can only increase the margin.

The deterministic error from the real mean and modes through R has squared norm at most $\varepsilon^2 \mathcal{E}(y, z)$, since the drift is canceled on the low nonsaturated modes. Between R and $N/2$ the added squared norm is at most $C\varepsilon^2 \sum_{|h|>R} (e_h + |y|)^2 / h^2$. The mean-zero norm above $N/2$ is $O(N^{-1/2})$. The value of the unpinned error at zero tends to zero. The cutoff is a bounded approximate identity, and the total of the changes on the nonzero modes at zero is

$$O_R(\varepsilon) + O(\varepsilon N^{2/(a+2)}) + O(\varepsilon \log N) = o(1).$$

Use Lemma 7, leaving a correction of norm $o(1)N^{-1/2} = o(\varepsilon)$, with additional coefficients tending to zero. For sufficiently large R and small ε, the deterministic error lies strictly inside the target ball with slack at least a fixed positive multiple of ε.

Give the retained saturated modes widths $\delta\varepsilon z_h$, the saturated endpoint width $\delta\varepsilon|y|$ when applicable, and the prescribed $R < |h| < N$ modes widths $\delta\varepsilon e_h$. On the other modes below $2N$ use width $\delta\varepsilon$. For fixed small $\delta > 0$, all margins remain positive. By (8), the finite centered perturbation has expected squared norm $O(\delta^2\varepsilon^2)$, uniformly as N grows. Indeed $\sum_{h \in S} z_h^2 / h^2 < \infty$ by (4) and $|y| \le 2$, and $\sum e_h^2 / h^2 < \infty$. The independent tail has expected squared norm $O(N^{-1})$. For sufficiently small fixed δ the comparison probability of both errors lying within the slack is bounded below by a positive constant.

The respective remaining margins on the designated modes, at the saturated endpoint, and on the large projected band are at least $(1 - \delta)\varepsilon z_h$, $(1 - \delta)\varepsilon|y|$, and $(1 - \delta)\varepsilon e_h$. The other retained modes have positive fixed margins for ε small, and all interval log-widths contribute $O_R(N \log(1/\varepsilon))$. Formula (11) yields cost at most

$$A(1 - \delta)^{-a} \varepsilon^{-a} \left(\mathcal{I}_a(y, z) + \sum_{|h|>R} e_h^{-a} \right) + o(\varepsilon^{-a}).$$

Take the ε-limit first. Enlarge R, decrease δ, then remove the common scaling and approach the infimum. This gives the reverse bound in Theorem 4.

6.3 A scalar reduction and examples

Let $G_*(a) = \inf(\mathcal{E} + \mathcal{I}_a)$ under the same sign restrictions. Homogeneity gives

$$I_*(a) = \frac{2}{a} \left(\frac{a G_*(a)}{a + 2} \right)^{(a+2)/2}. \tag{16}$$

Indeed simultaneous scaling by $t > 0$ replaces the unconstrained objective by $t^2\mathcal{E} + t^{-a}\mathcal{I}_a$, whose minimum depends monotonically on $\mathcal{I}_a \mathcal{E}^{a/2}$. Normalizing $\mathcal{E}$ to one identifies the infimum of that product with $I_*(a)$ whenever a saturation is present. In that case the optimal balance for the sum is $2\mathcal{E} = a\mathcal{I}_a$, yielding (16). With no saturation both infima are zero.

For fixed y, the minimizing $z_h = r_h(s_h y)$ is given by the unique root

$$r_h(x)^{a+1}(r_h(x) + x) = \frac{a}{2w_h}, \qquad r_h(x) > \max(0, -x). \tag{17}$$

This is the stationarity equation of the strictly convex function $w_h(z+x)^2 + z^{-a}$. Implicit differentiation gives

$$r_h'(x) = -\frac{r_h(x)}{(a+1)(r_h(x)+x)+r_h(x)} \in (-1,0),$$

so $r_h(x)$ differs from $r_h(0) = (a/(2w_h))^{1/(a+2)}$ by at most $|x|$. The substituted series and its derivative converge uniformly on compact subsets of the scalar y-domain, by $\sum w_h^{a/(a+2)} < \infty$ and this bound on the roots. The substituted objective is convex by joint convexity before the coordinate minimization, and $y^2/4$ makes it strictly convex and coercive. At a saturated endpoint the barrier diverges as y tends to zero within its allowed half-line. Thus the unique scalar minimizer solves

$$\frac{y}{2} + 2\sum_{h\in S} w_h(s_h r_h(s_h y) + y) + P_a'(y) = 0. \tag{18}$$

For an interior endpoint, $P_a'(y) = 0$. For a saturated one, $P_a'(y) = as_0(-s_0 y)^{-a-1}$.

At an interior endpoint the minimum is at $y = 0$ precisely when $\sum_{h\in S} s_h w_h^{(a+1)/(a+2)} = 0$. In that balanced case the diagonal minimization gives

$$I_*(a) = \left(\sum_{h\in S} w_h^{a/(a+2)}\right)^{(a+2)/2},$$

by Hölder with equality at z_h proportional to $w_h^{-1/(a+2)}$. A nonzero signed sum makes the scalar minimizer have the opposite sign and makes the actual rate strictly smaller than the one obtained by requiring $y = 0$.

- With only endpoint saturation, $I_*(a) = 2^{-a}$.
- With an interior endpoint and one saturated mode h, minimizing over y gives energy $z_h^2/[4 + (2\pi h)^2]$. Thus $I_*(a) = [4 + (2\pi h)^2]^{-a/2}$, as for $g = 2v_h$.
- With two saturated modes $\pm k$ of equal sign and an interior endpoint, symmetry gives equal deficits z. Minimizing $y^2/4 + 2w_k(z+y)^2$ gives energy $2w_k z^2/(1+8w_k)$ for that common sign after replacing y by its signed version. The resulting coefficient is $I_*(a) = 2[4+(2\pi k)^2/2]^{-a/2}$. This includes the sine center $g(t) = 2\sin(2\pi kt)/(\pi k)$ used for the supercritical converse in [6, Theorem 7.1].
- With a saturated endpoint and one saturated mode h, put $u_* = s_h s_0$ and $k = z_h/(-s_0 y)$. The scalar objective after energy normalization is $(1 + k^{-a})[1/4 + w_h(k - u_*)^2]^{a/2}$. Its minimum is at the unique $k > 1$ solving

  $$(k - u_*)(k^{a+1} + u_*) = \frac{1}{4w_h}.$$

  The equation follows by differentiation; its left side is nonpositive on $0 < k \le 1$ and increases strictly from zero to infinity for $k > 1$, for either sign of u_*.

7 The rate for power-closing margins

7.1 The continuum minimization

For $\sigma \in \{+,-\}$, put $m_\sigma(u) = c_\sigma u^{-b}$, $u > 0$. Consider the minimum

$$\mathcal{C} = \inf \sum_\sigma \int_0^\infty z_\sigma(u)^{-a}\, du, \qquad z_\sigma \ge m_\sigma, \qquad \sum_\sigma \int_0^\infty \frac{(z_\sigma(u) - m_\sigma(u))^2}{4\pi^2 u^2}\, du \le 1, \tag{19}$$

over measurable profiles. With multiplier $\lambda > 0$ for the energy, the unique pointwise minimum is the solution

$$(z_\sigma - m_\sigma) z_\sigma^{a+1} = \frac{2\pi^2 a}{\lambda} u^2, \qquad z_\sigma > m_\sigma. \tag{20}$$

The derivative of the strictly convex objective $z^{-a} + \lambda(z - m_\sigma)^2/(4\pi^2 u^2)$ is negative at m_σ and positive for sufficiently large z. Its unconstrained minimum over positive z is consequently also the desired constrained one.

As $u \downarrow 0$, the excess $z_\sigma - m_\sigma$ is $O(u^{2+b(a+1)})$ from (20) and $z_\sigma \geq m_\sigma$. For large u, the equation implies $z_\sigma \asymp u^{2/(a+2)}$: first $z_\sigma \gtrsim u^{2/(a+2)}$, then $m_\sigma = o(z_\sigma)$. These estimates give finite cost and energy since $2a/(a+2) > 1$. The energy is strictly decreasing and continuous in λ by pointwise monotonicity and dominated convergence between fixed positive values of λ. As $\lambda \to \infty$, the energy tends to zero by dominated convergence. As $\lambda \downarrow 0$, it tends to infinity already on any compact positive interval. Hence a unique λ gives energy one. Integrating the pointwise minimization and subtracting λ proves that this profile attains (19), uniquely up to null sets. The same construction gives the unique profile at any positive energy, depending continuously on that energy in cost by dominated convergence.

7.2 Necessary cost on a growing frequency band

Take $n = \varepsilon^{-2/(2b+1)}$. Then $n^{-2b-1} = \varepsilon^2$ and $n^{ab+1} = \varepsilon^{-q}$. Fix $0 < \delta < R < \infty$ and retain h with $\delta \leq |h|/n \leq R$. For small ε their β_h are nonzero. Define

$$d_h = 2 - \operatorname{sgn}(\beta_h) X_h, \qquad z_h = n^b d_h, \qquad m_h^* = n^b m_h.$$

Using $|d_h - m_h| = |X_h - \beta_h|$, the Fourier projection bound following (9) gives on the ball

$$\frac{1}{n} \sum_{\delta \leq |h|/n \leq R} \frac{(z_h - m_h^*)^2}{4\pi^2 (|h|/n)^2} \leq 1 + o(1).$$

The normalized density cost on this band is $\varepsilon^q \sum d_h^{-a} = n^{-1} \sum z_h^{-a}$. Using the multiplier λ from (20), add λ times the displayed energy and minimize each term over $z_h > 0$. Subtracting at most $\lambda(1 + o(1))$ leaves a lower bound on the normalized cost. On each sign's compact band, m_h^* converges uniformly to $m_\sigma(|h|/n)$ by the assumed asymptotics. The minimized objective is continuous in baseline and frequency on compact subsets with both positive. Thus the sum of the minima converges to the integral of the pointwise minimum over $[\delta, R]$, summed over the signs.

Integrating the upper density bound incurs $O(n)$ in the logarithm for this fixed band, negligible at scale ε^{-q}. Let $\delta \downarrow 0$ and $R \to \infty$. The full integrated pointwise minimum is $\mathcal{C} + \lambda$, because the optimizing profile has energy one. This proves $\liminf_{\varepsilon \downarrow 0} \varepsilon^q B_{A,a}(g, \varepsilon) \geq A\mathcal{C}$.

7.3 Sufficient cost with the pinned path

Start from the optimizing continuum profile at energy $1 - \eta$, where $0 < \eta < 1$ is fixed. Its cost tends to $\mathcal{C}$ as $\eta \downarrow 0$. Fix $0 < \delta < R$, with R large enough to pay a small part of the energy slack using the tail estimate below. Choose $N = \lceil \varepsilon^{-\gamma} \rceil$ with

$$2 < \gamma < \min\left(q, \frac{(a+2)b + 2}{2b+1}\right).$$

Both upper endpoints exceed two. In particular

$$N^{-1/2} = o(\varepsilon), \quad N \log(1/\varepsilon) = o(\varepsilon^{-q}), \quad n^{-b}(N/n)^{2/(a+2)} = o(1).$$

The deterministic endpoint equals ξ. For $0 < |h| < \delta n$ keep β_h. On $\delta n \le |h| \le Rn$, prescribe normalized coefficient $\operatorname{sgn}(\beta_h)(2 - d_h)$, with $d_h = n^{-b} z_\sigma(|h|/n)$. For ε small these margins satisfy $d_h \ge m_h$ uniformly on the compact band, since (20) gives a strictly positive excess over the baseline there. On $Rn < |h| < N$ use the same sign and margin

$$d_h = \max\left(m_h, n^{-b}(|h|/n)^{2/(a+2)}\right).$$

Multiply the nonzero prescribed coefficients by the smooth cutoff of Section 2. Since it equals one below $N/2$, it preserves the chosen middle band for ε small. It only increases the margin elsewhere.

After division by ε^2, the middle band's squared Fourier error tends to the corresponding profile energy. From $Rn < |h| \le N/2$, the squared error is at most

$$C\varepsilon^2 \int_{R/2}^{\infty} u^{4/(a+2)-2}\, du.$$

This is the reciprocal-frequency square sum of the bound $n^{-b}(|h|/n)^{2/(a+2)}$ on the change in each normalized coefficient, converted to an integral. The modes above $N/2$ cost $O(N^{-1/2})$ in norm, including tapering and the omitted Fourier coefficients of the center. The unpinned periodic error at zero is $o(1)$: the cutoff itself approximates the continuous periodic part, the middle band changes its value by $O(n^{-b})$, and the upper band changes it by at most $O(n^{-b}(N/n)^{2/(a+2)})$. Use Lemma 7 on $N \le |h| < 2N$. The added normalized coefficients are $o(1)$ and the remaining norm cost is $o(\varepsilon)$. Thus with R large, the deterministic error is within radius ε with slack a fixed positive multiple of ε.

Use uniform comparison intervals of half-width $\nu\varepsilon$ around every prescribed coefficient through $|h| < 2N$. For fixed δ, R, every deterministic margin is $\gg n^{-b}$, using the finitely many fixed positive margins and the two tail asymptotics for the modes $|h| < \delta n$. Since $n^{-b}/\varepsilon \to \infty$, the relative loss of margin from these intervals is $o(1)$. The endpoint and correction band are uniformly interior. By (8), the finite perturbation has expected squared norm $O(\nu^2\varepsilon^2)$; the independent tail has expected squared norm $O(N^{-1})$. For a sufficiently small fixed ν, the probability of remaining within the slack under comparison is bounded below.

The low-band negative log density cost is at most $C\delta^{ab+1}\varepsilon^{-q} + O(1)$, by summing m_h^{-a}. The middle-band cost divided by ε^{-q} tends to the truncated profile cost, with the factor A from (1). The upper band contributes at most $C\varepsilon^{-q}\int_{R/2}^{\infty} u^{-2a/(a+2)}\, du$. The correction and endpoint cost $O(N)$ and the interval volumes cost $O(N\log(1/\varepsilon))$. Formula (11) therefore gives the reverse inequality after taking the ε-limit and sending $\delta \downarrow 0$, $R \to \infty$, and then $\eta \downarrow 0$.

7.4 Evaluation of the constant

Fix $\lambda > 0$ and let C_λ and E_λ be the cost and energy of the profiles in (20). For each sign write $z = c_\sigma u^{-b} t$ and $u = \ell v$, with $\ell = (c_\sigma^{a+2}\lambda/(2\pi^2 a))^{1/d}$. Then $t > 1$ satisfies $(t-1)t^{a+1} = v^d$, and the one-sign cost is

$$c_\sigma^{-a}\ell^{ab+1}\int_0^{\infty} t(v)^{-a} v^{ab}\, dv.$$

Integration by parts gives

$$\int_0^{\infty} t(v)^{-a} v^{ab}\, dv = \frac{a}{ab+1}\int_1^{\infty} t^{-a-1}[(t-1)t^{a+1}]^{\rho}\, dt = \frac{a}{ab+1}\,\mathrm{B}\left(1+\rho, \frac{a-2}{d}\right).$$

The boundary term vanishes at both ends because $ab + 1 > 0$ and $(a+2)\rho < a$. The second equality uses $x = 1 - 1/t$. The dilation $z_\sigma(u) \mapsto t^{-b} z_\sigma(u/t)$ multiplies cost by t^{ab+1} and energy by

$t^{-(2b+1)}$. Differentiating the minimizing objective along this dilation at $t = 1$ gives $\lambda E_\lambda = p_* C_\lambda$. At total energy one, $\lambda = p_* \mathcal{C}$. Summing the two costs and substituting this value of λ gives (5), using $(a+2)\rho - a = -(a-2)/d$.

Corollary 8 (Unequal closing powers)**.** *Under the interior assumptions of Theorem 5, suppose instead that $m_{\pm h} \sim c_\pm h^{-b_\pm}$ with $b_+ > b_- > 0$. The rate exponent is $q = 2(ab_+ + 1)/(2b_+ + 1)$. In (5) use $b = b_+$ and retain only the c_+ term. The roles of the signs reverse when $b_- > b_+$.*

Proof. For the necessary cost, retain only the positive scaled band and use the one-sign version of (19). For sufficiency use the two sign-specific scales, assigning energy $1-2\eta$ to the positive side and η to the negative side with $0 < \eta < 1/2$. A common cutoff exponent $\gamma > 2$ can satisfy both sets of strict upper bounds in the sufficient construction. The squared Fourier errors add by Parseval. Both pin displacements are $o(1)$, so one common reserved packet suffices. The probability estimates on the two bands add at their respective rates. The derivative $q'(b) = 2(a-2)/(2b+1)^2 > 0$ makes the negative-side cost negligible at the positive rate for fixed η. Sending $\eta \downarrow 0$ leaves the unit-energy one-sign constant as asserted. □

8 Proof of the critical crossover

Throughout this section $a = a_\varepsilon$. The normalizers satisfy $2e^{-A} \le Z_{A,a} \le 4$, by integrating over $[-1, 1]$ for the lower bound. Put $s = 2a/(a+2)$. For $1 \le |h| \le 2\varepsilon^{-2} J^{1/2}$,

$$|h|^{-s} = \frac{\exp[-\tau \log|h|/(4J)]}{|h|}(1 + o(1)) \tag{21}$$

uniformly, since $(s-1)J \to \tau/4$. Also $\varepsilon^{2-a} \to e^\tau$ and $J^{(a-2)/2} \to 1$.

8.1 Necessary bound

Fix $\theta > 0$ and retain the frequencies of S_+^τ with $|h| \ge J^2$. The omitted part of this harmonic sum is $O(\log J)$. The drift contribution on retained modes is $o(\varepsilon)$ by the projection bound following (9), and so is the weighted Euclidean norm of their baseline margins, since $\sum m_h^2/h^2 = O(\varepsilon^{2+2\theta} J)$. Their margins are at most ε^θ, so the corresponding β_h are nonzero for small ε. Writing $d_h = 2 - \operatorname{sgn}(\beta_h) X_h$ gives $\sum d_h^2/(2\pi h)^2 \le (1 + o(1))\varepsilon^2$ on the event. Hölder yields

$$\sum d_h^{-a} \;\ge\; (2\pi)^{-a} \varepsilon^{-a} (1 + o(1)) \left(\sum |h|^{-s} \right)^{(a+2)/2}.$$

Indeed $|h|^{-s} = (d_h^{-a})^{2/(a+2)} (d_h^2/h^2)^{a/(a+2)}$. The finite density bound after (11) incurs $O(\varepsilon^{-2})$ in the logarithm. Divide the remaining bound by $\varepsilon^{-2} J^2$ and use (21). Because $a \to 2$, the map $x \mapsto x^{(a+2)/2}$ tends uniformly to x^2 on the bounded range of $J^{-1} \sum |h|^{-s}$, including at zero. Take the lower limit and then $\theta \downarrow 0$, obtaining the claimed liminf. Empty selections give the nonnegative bound.

8.2 Sufficient bound

Let $N = \lfloor \varepsilon^{-2} J^{1/2} \rfloor$ and $\delta = J^{-1/16}$. For fixed $\theta > 0$, put

$$B = \{1 \le |h| \le N/2 : m_h \le \varepsilon^{1-\theta} \sqrt{|h|}\}, \qquad T = \sum_{h \in B} |h|^{-s}, \quad D = T + J^{3/4},$$

$$\alpha = \frac{2\pi(1-4\delta)\varepsilon}{\sqrt{D}}, \qquad v_h^* = \alpha|h|^{2/(a+2)}.$$

For $1 \le |h| < 2N$ these margins are comparable to $\alpha\sqrt{|h|}$ with fixed positive constants, by boundedness of $(a-2)\log(2N)$. Project the tapered coefficient $\eta_0(h/N)\beta_h$ onto $[-2+v_h^*, 2-v_h^*]$ on B and $N/2 < |h| < N$. Here η_0 is the smooth cutoff from Section 2. Leave the other low modes unchanged. The low mean-zero error has squared norm at most $\alpha^2 T/(4\pi^2)$. The upper band and omitted frequencies give norm $O(N^{-1/2}) = o(\delta\varepsilon)$. The unpinned error at zero tends to zero by the approximate identity and the bound $O(\alpha N^{2/(a+2)}) = o(1)$ on the additional change from projection. Use Lemma 7; the residual cost is $o(1)N^{-1/2} = o(\delta\varepsilon)$. Shift the endpoint to $(1-\delta^2\varepsilon)\xi$, leaving deterministic norm error at most $(1-4\delta)\varepsilon + o(\delta\varepsilon)$.

Use comparison widths $\delta^4 v_h^*$ on all nonzero modes below $2N$ and $\delta^2\varepsilon$ at zero. By (8) and the full weighted sum $O(J)$, the comparison perturbation has variance

$$O(\varepsilon^2\delta^4 + \delta^8\varepsilon^2 J/D) = o(\delta^2\varepsilon^2)$$

since $D \ge J^{3/4}$. The omitted independent tail has variance $O(N^{-1})$, also $o(\delta^2\varepsilon^2)$. Hence the probability of strict acceptance under comparison tends to one.

The main low-band cost is at most $A(1-\delta^4)^{-a}\alpha^{-a}T$. The taper band has weighted harmonic sum $O(1)$ and costs $O(\alpha^{-a})$. For unchanged low modes the relative width is $o(1)$, since $m_h > \varepsilon^{1-\theta}\sqrt{|h|}$. Their cost is $O(\varepsilon^{-a+a\theta}J)$ using $\sum_{1\le|h|<N}|h|^{-a/2} = O(J)$. The correction band costs $O(N)$, the endpoint costs $O(\delta^{-2a}\varepsilon^{-a})$, and the total logarithmic interval volumes cost $O(N\log J + J)$. Indeed the nonzero modes below $2N$ have widths comparable to $\delta^4\alpha\sqrt{|h|}$. Stirling's formula gives a total logarithmic cost of order $N[1+\log(1/(\delta^4\alpha\sqrt{N}))]$, and the endpoint contributes $O(J)$. All these extra terms are $o(\varepsilon^{-2}J^2)$. For the main term,

$$TD^{a/2}/J^2 = (T/J)^2 + o(1)$$

uniformly for $0 \le T = O(J)$, because $D - T = J^{3/4}$ and $(a-2)\log D = o(1)$. The difference between the cutoff $N/2$ and ε^{-2} has weighted harmonic size $O(\log J)$. Using (11), (21), and then sending $\theta \downarrow 0$ gives the limsup in Theorem 6.

8.3 Evaluation of the crossover profiles

The saturated modes define a harmonic measure on the scale $u = \log|h|/J$. After normalization by J, it converges weakly on $[0,2]$ to $\kappa\,du$. Its cumulative mass at $u > 0$ tends to κu by definition of κ. For fixed $\theta > 0$, the small threshold of S_+^τ selects only saturated frequencies for ε small. For S_-^τ any additional mode must have $|h| \ge m_0^2\varepsilon^{-2+2\theta}$. That range has weighted harmonic size $O_\tau(\theta J + 1)$. The common profile is $\kappa\int_0^2 e^{-\tau u/4}\,du$. Substitution gives (6).

For $m_{\pm h} \asymp h^{-b}$, the logarithmic lower cutoff of the two sums is $(1\pm\theta)/(b+1/2)$, up to $O(1/J)$. The two-sign harmonic measure restricted to the resulting range converges to $2\,du$ on that limiting interval. After $\theta \downarrow 0$, the common profile is $2\int_{2/(2b+1)}^2 e^{-\tau u/4}\,du$, which yields (7).

9 Limits of the fixed-exponent constants

The fixed-a constants also have limits as $a \downarrow 2$, which can be compared with the functions of τ above. If S has harmonic density κ as in Theorem 6, then

$$\lim_{a\downarrow 2}(a-2)^2 I_*(a) = 4\kappa^2/\pi^2. \tag{22}$$

To prove it put $T_a = \sum_{h\in S}(2\pi|h|)^{-2a/(a+2)}$ and $s = (a-2)/(a+2)$. If $F(R) = \sum_{h\in S,\ |h|\le R}|h|^{-1}$, integration by parts for the harmonic counting measure gives

$$s\sum_{h\in S}|h|^{-1-s} = s^2\int_0^\infty e^{-st}F(e^t)\,dt \longrightarrow \kappa.$$

After the change of variable $v = st$, this follows from $F(e^t)/t \to \kappa$ and $F(e^t) = O(1+t)$. Consequently $(a-2)T_a \to 2\kappa/\pi$. If $S \neq \varnothing$, then $T_a^{(a-2)/2} \to 1$, using a positive lower bound from one frequency and the upper bound $O((a-2)^{-1})$.

On the tail $|h| > R$ the weighted Euclidean norm of z_h for a feasible pair is at most $1 + 2(\sum_{|h|>R} w_h)^{1/2}$. Apply the diagonal Hölder bound to this tail, then let $a \downarrow 2$ and $R \to \infty$ to obtain the lower bound in (22). For an interior endpoint use $y = 0$ and the full diagonal minimizer to obtain $I_*(a) \le T_a^{(a+2)/2}$. For a saturated endpoint use $y = -s_0\sqrt{a-2}$ and first scale the diagonal deficits by $1 - |y|(1/4 + \sum_{h\in S} w_h)^{1/2}$. The energy is then at most one by the triangle inequality. The scaling tends to one and the endpoint cost $(a-2)^{-a/2}$ is negligible after multiplication by $(a-2)^2$. If S is empty, the result follows directly from the cases $I_*(a) = 0$ and $I_*(a) = 2^{-a}$.

For the interior power margins of Theorem 5,

$$\lim_{a\downarrow 2}(a-2)^2\mathcal{C} = 16/\pi^2. \tag{23}$$

Indeed $\rho \to 1/2$, $p_* \to 1$, and $x\,\mathrm{B}(1+\rho, x) \to 1$ as $x = (a-2)/d \downarrow 0$. The latter limit follows by splitting the beta integral at a fixed point near one and using continuity of t^ρ there, uniformly for ρ near $1/2$. Thus the bracket in (5) is asymptotic to $4/(\pi(a-2))$, and the exponent $1/(1-\rho) = 2 + O(a-2)$ gives (23).

For positive saturation density, (6) has large positive τ behavior $(4A\kappa^2/\pi^2)e^\tau/\tau^2$. For power-closing margins, (7) has behavior $(16A/\pi^2)e^{2\tau b/(2b+1)}/\tau^2$. These agree with the fixed-exponent scales and constants in (22) and (23) after identifying $\tau = (a-2)J$ and noting that $q - 2 = 2b(a-2)/(2b+1)$. The comparison concerns the separate limits of the coefficients and the crossover functions.

Acknowledgement. OpenAI Codex and Anthropic Claude assisted with proof development, independent checking, numerical verification and manuscript preparation. The author is responsible for the work.

References

[1] L. Grafakos, *Classical Fourier Analysis*, 3rd ed., Graduate Texts in Mathematics 249, Springer, 2014. doi:10.1007/978-1-4939-1194-3.

[2] E. Kowalski and W. Sawin, Kloosterman paths and the shape of exponential sums, *Compos. Math.* **152** (2016), 1489–1516. doi:10.1112/S0010437X16007351.

[3] E. Kowalski and W. Sawin, On the support of the Kloosterman paths, arXiv:1709.05192, 2017.

[4] D. Kyselov, Space-filling curves and small-ball probabilities for the Kloosterman process, arXiv:2610.22645v1, 2026.

[5] L. V. Rozovskii, Small deviation probabilities for a weighted sum of independent positive random variables with common distribution function that can decrease at zero fast enough, *Theory Probab. Appl.* **63** (2018), 155–163. doi:10.1137/S0040585X97T988976.

[6] S. Schubr, Roughness and boundary-density thresholds for small balls of the Kloosterman process, Luminary Research Archive, edition 1, 2026. doi:10.5281/zenodo.24409097.

[7] L. V. Rozovsky, Small deviation probabilities for sums of independent positive random variables, *Vestn. St. Petersbg. Univ., Math.* **53** (2020), 295–307. doi:10.1134/S1063454120030103.