Moment transfer and spectral bounds for Boolean Doob kernels

Lluis Eriksson

17 November 2026. Revision candidate

Abstract

For positive potentials on a Boolean cube, the stationary density of the noise Doob kernel and its normalized weighting density are different nonlinear functions of the potential. We give an admissible family for which the stationary critical moment is below two while the weighting critical moment is above two, despite a fixed positive stationary floor and exact marginals through order proportional to the square root of the dimension. This refutes transfer of that numerical cap, not the existence of a larger uniform bound. A nonlinear consequence of the stationary floor, combined with classical ordinary and reverse hypercontractivity, gives explicit inverse-potential moments and weighting-moment bounds below a computable finite threshold. Exact stationary matching then gives strong convergence of the two densities throughout this subcritical range. The finite bootstrap crosses the exponent needed for a dimension-free low-to-all operator bound. We obtain exponentially small error on degrees up to a fixed positive fraction of the stationary matching order. Non-Perron modes a fixed positive distance above the noise edge must concentrate on actual high-weighting sets of exponentially small stationary probability. A tensor example rules out an unrestricted high-degree hypercontractive repair; the high block for arbitrary potentials remains uncontrolled. For potentials supported on a binary linear-code Fourier algebra, we control every invariant sector of the full Doob kernel. Under a distance–cardinality condition its gap tends uniformly to the unweighted gap, for every positive quotient potential. Direct disjoint-block estimates cover a sharper parameter regime. A standard simplex code supplies positive potentials outside the disjoint-block class. Classical random linear coding extends this mechanism to positive-rate families at every sufficiently large length, with an explicit exponential gap estimate and normalized witnesses meeting any fixed stationary floor below one and the moment cap. Without marginal matching, the same stationary floor and critical cap also imply a dimension-free weak Poincare inequality and polynomial variance decay for bounded observables of the actual chain. At noise one half, an additional smoothing step gives the summable decay exponent $8/7$ and a dimension-free $O(N^{-1})$ mean-square bound for sample averages, including uniform initialization. These bounded-observable results do not imply a full L^2 gap. For arbitrary-valued positive potentials with a fixed positive minimum after energy normalization, the critical cap yields an explicit dimension-uniform full-chain gap, without matching or an upper bound on the potential. The original stationary floor supplies the required minimum for every two-level potential, recovering a full-gap class theorem. Any normalized family with collapsing gaps and the fixed critical cap must have potential minima tending to zero. A rare-set construction followed by exact Fourier correction and stable symmetric scaling shows that even all original stationary hypotheses do not force a positive normalized minimum. A critical hypercontractive sandwich and the proved stability of the reconstruction show that this same family has actual gap converging exponentially to the unweighted gap. A different exactly matched family has a noise denominator tending to zero and an unbounded pointwise potential-to-denominator ratio, yet its entire spectrum converges exponentially to Boolean noise. Thus neither a positive denominator floor nor a pointwise ratio ceiling follows from

the stationary assumptions. Noise smoothing of the actual potential preserves the stationary floor, but a four-point construction shows that every nontrivial smoothing can destroy exact stationary matching even while the critical cap survives. Smoothing the stationary density itself preserves the admissible constraints. For sufficiently strong fixed smoothing, a derived higher moment and finite inverse-moment crossing give uniform convergence of the reconstructed chain's entire spectrum to Boolean noise, exponentially in the matching degree. Fixed smoothing need not be a vanishing critical-norm perturbation of the original density. An exactly matched block-spike family disproves uniform stationary-variance upper comparison and two-sided Dirichlet-form equivalence between the original and regularized chains. Both full gaps in this family nevertheless tend to the noise gap. Direct spectral comparison with the original chain remains unresolved. The general stationary-gap question and critical weighting-moment bound remain open.

1 Setting and theorem map

Let μ be uniform probability on $\{-1,1\}^L$, $L \geq 1$, and let $0 < \rho < 1$ be fixed. Boolean product noise is denoted by $P = P_\rho$, with $P_\rho \chi_S = \rho^{|S|}\chi_S$ for Walsh characters χ_S. Unqualified expectations and norms use μ. For $h > 0$ set

$$g = Ph, \quad Z = \mathbb{E}hg, \quad \pi = hg/Z, \quad N = \mathbb{E}(h/g), \quad r = (h/g)/N.$$

Both π and r have mean one. The Markov kernel

$$K(x,y) = P(x,y)h(y)/g(x)$$

is reversible for $\pi\, d\mu$. Its eigenvalues are positive: conjugation by multiplication by $\sqrt{\pi}$ identifies it with $M_{\sqrt{h/g}} P M_{\sqrt{h/g}}$, a positive definite matrix. Its top eigenvalue is one and is simple because all entries are positive. Write $\lambda_2(K)$ for its second eigenvalue and $\mathrm{gap}(K) = 1 - \lambda_2(K)$. These are finite Markov-chain quantities.

The general question motivating this work fixes $b, C > 0$, puts

$$s = \frac{1+\rho}{1-\rho}, \qquad k_L = \min\{L, \lfloor C\sqrt{L} \rfloor\},$$

and imposes

$$\pi \geq b, \qquad \mathbb{E}\pi^s \leq 2, \qquad \mathbb{E}\pi\chi_S = 0 \quad (1 \leq |S| \leq k_L). \tag{1}$$

The last condition concerns the actual stationary density, on the original labelled coordinates. It is equivalent to uniform coordinate marginals of these orders. It does not assign exact marginals to r.

The results have distinct scopes. Section 2 gives one fixed ρ and floor for which the numerical implication $\mathbb{E}\pi^s \leq 2 \Rightarrow \mathbb{E}r^s \leq 2$ fails even under all of (1). Section 3 treats every fixed ρ and proves a dimension-free bound for $\|r\|_t$ for $1 \leq t < t_+(\rho) < s$; with matching it proves $r - \pi \to 0$ in this range. Section 4 treats every positive potential on specified code algebras, without a moment cap or floor for its spectral result. It controls the full chain, not only its quotient. Section 5 relates the sharper disjoint-block bounds and the simplex example to these conclusions. Section 6 proves positive-rate, all-large-length code existence with simultaneous column separation, specializes the full-chain estimate exponentially, and gives actual normalized stationary floor/cap witnesses. Sections 7–8 use the finite bootstrap to control the entire low-to-all operator through a fraction of the matching degree, localize actual non-Perron eigenfunctions,

and exclude a generic tail-space repair. Section 9 derives weak relaxation without matching and proves a summable rate and uniform-start estimation bound at $\rho = 1/2$. Section 10 proves a full gap from a fixed positive potential minimum and the critical cap, without matching or upper height restrictions, and derives the two-level class as a consequence of its stationary floor. Propositions 10.3 and 10.5 show that neither floor and cap alone nor the full exact-matching hypotheses supply that additional minimum, even under strong near-uniformity. Theorem 10.7 then proves that the exactly matched family retains the noise gap asymptotically, with exponential error. Section 11 constructs the opposite denominator defect, with $Ph \to 0$ and $h/Ph \to \infty$ at a point, while proving critical ratio convergence and full-spectrum stability for that new family. Section 12 separates potential smoothing from stationary-density smoothing: the former preserves the stationary floor but can break exact matching despite retaining the critical cap, whereas the latter preserves the constraints. Section 13 proves full-spectrum convergence to Boolean noise for the actual chain reconstructed after sufficiently strong fixed density smoothing, uniformly over all admissible inputs. It also refutes a vanishing critical-norm perturbation premise. Section 14 realizes a failure of separate variance and form upper comparisons on actual admissible chains, while controlling both full gaps. Direct comparison with the original spectrum remains open.

The manuscript groups analytical sources 349, 350, 351, corrected 353B and 355 with the separately audited 16 November low-to-full and dimension-free follow-ups A–D and the weak-relaxation and estimation notes E–G, together with the 17 November two-level and positive-background full-gap notes I and J and the vanishing-minimum notes K and L and their full-chain classification M, together with the exact-matching regularization obstruction N and the density-regularized full-spectrum theorem O, together with the actual separate-comparison obstruction P and the small-denominator family and spectral classification Q. The total-variation estimate from source 341 is proved again in Section 3. No spectral conclusion is inferred from moment convergence alone. No assertion about a physical Yang–Mills gap or a continuum limit is made.

2 Failure of exact critical-cap transfer

Theorem 2.1. *There are fixed $\rho_* \in (1/3, 1/2)$ and $b_* = 1/16$ such that, for every fixed $C > 0$, a sequence of strictly positive potentials satisfies* (1) *for every L and satisfies*

$$\mathbb{E}\pi_L^{s_*} < 2 < \mathbb{E}r_L^{s_*} \quad \textit{for all sufficiently large } L, \qquad s_* = (1+\rho_*)/(1-\rho_*).$$

Both displayed moments tend to two. This is not a divergent-moment or collapsing-gap example.

Proof. On a two-dimensional cube, in the order $(++, +-, -+, --)$, take

$$H = (29/16,\ 1/4,\ 1/8,\ 29/16) = 1 + cx - cy + dxy, \qquad c = 1/32, \quad d = 13/16.$$

Thus $\mathbb{E}H = 1$. Put $q = c(x - y)$, $v = 2c^2 = 1/512$, $w = d^2 = 169/256$. For an effective noise parameter $\theta \geq 0$,

$$G_\theta = P_\theta H = 1 + \theta q + \theta^2 dxy, \quad Z_\theta = \mathbb{E}HG_\theta = 1 + \theta v + \theta^2 w.$$

Take $h_\theta = H/\sqrt{Z_\theta}$, $\pi_\theta = HG_\theta/Z_\theta$, and $r_\theta = (H/G_\theta)/\mathbb{E}(H/G_\theta)$. These are actual densities from the same potential. At $\theta = 0$, both densities equal H, and differentiation gives

$$\pi_0' = H(q - v), \qquad r_0' = -H(q - v).$$

Indeed $\mathbb{E}Hq = v$ and $(\mathbb{E}H/G_\theta)'_{\theta=0} = -v$. For each fixed $t > 1$, let

$$C_t = \mathbb{E}H^t q - v\mathbb{E}H^t = \frac{31(1/4)^t - 33(1/8)^t - 2(29/16)^t}{2048}.$$

This is strictly negative. To see this, $f(t) = 31\,4^{-t} - 33\,8^{-t}$ has $f'(t) = \log 2\,8^{-t}(99 - 62\,2^t) < 0$ for $t \geq 1$, whereas $2(29/16)^t$ is increasing; at $t = 1$ both are $29/8$. Consequently

$$\frac{d}{d\theta}\mathbb{E}\pi_\theta^t\Big|_0 = tC_t < 0, \qquad \frac{d}{d\theta}\mathbb{E}r_\theta^t\Big|_0 = -tC_t > 0. \tag{2}$$

Let $M(t) = \mathbb{E}H^t$. Direct rational identities give $M(2) = 851/512 < 2$ and $M(3) = 24425/8192 > 2$. Strict convexity of $x \log x$ gives $M'(1) = \mathbb{E}H \log H > 0$, and $M''(t) = \mathbb{E}H^t(\log H)^2 > 0$. Thus there is a unique $t_* \in (2, 3)$ with $M(t_*) = 2$. Fix $\rho_* = (t_* - 1)/(t_* + 1)$, so its critical exponent is precisely t_*. Formula (2), with this exponent fixed, gives strict opposite inequalities around two for all sufficiently small positive θ. No varying exponent or numerical root is used.

For large L, let $m = \lfloor L/2 \rfloor$ and use two disjoint blocks B_1, B_2 of size m, with an unused coordinate when L is odd. Set $z_i = \chi_{B_i}$ and $\theta = \rho_*^m$. The variables z_1, z_2 are uniform and independent; restricted to their algebra, P_{ρ_*} is exactly the two-bit P_θ. Use $h_L = H(z_1, z_2)/\sqrt{Z_\theta}$. Products, positive inverses and expectations on the quotient agree with those above, so these are the required actual π_L, r_L.

Their Fourier support is contained in $\{\varnothing, B_1, B_2, B_1 \cup B_2\}$. In particular π_L has exact marginals through degree $m - 1$. Eventually $m > k_L$, $\theta \to 0$, and $\min \pi_L \to \min H = 1/8$. Choose a finite L_0 after which the floor $1/16$, the matching degree and the two strict moment inequalities all hold. For the remaining finitely many L, put $h_L = 1$. This defines a full admissible sequence. Continuity gives convergence of both moments to two. Section 5 shows that this two-block family in fact has gap tending to $1 - \rho_*$. □

The following estimates distinguish two simpler classes: two-level potentials cannot produce divergent critical weighting moments under a fixed stationary floor and cap, while product potentials obey the stronger constant-one moment comparison. Section 10 further proves a dimension-uniform full gap for the two-level class under the stationary floor and cap.

Proposition 2.2. *Normalize $\mathbb{E}hPh = 1$ and assume $\pi \geq b$. If h takes at most two values, then $g \geq b/2$ and $\mathbb{E}r^p \leq (4/b^3)^p \mathbb{E}\pi^p$ for every $p \geq 1$. If h is a product of positive one-coordinate factors, then $\mathbb{E}r^p \leq \mathbb{E}\pi^p$ for every $p \geq 1$, without a floor assumption. The latter also holds for products of factors on disjoint block parities, with their respective effective noise parameters.*

Proof. Write $a = \mathbb{E}h \leq 1$, the inequality following from the constant Fourier component of $\left\|P_{\sqrt{\rho}}h\right\|_2^2 = 1$. The floor implies $h \geq b/g$, so $a \geq b\mathbb{E}(1/g) \geq b/a$. Hence $a^2 \geq b$. Under $\pi\,d\mu$, $\mathbb{E}_\pi(1/g) = a$ and $\mathbb{E}_\pi(1/g^2) = N$, giving $N \geq a^2 \geq b$. For a nonconstant two-level h, let $h = m$ on a set of mass $q \in (0, 1)$ and $h = M > m$ on its complement. Integrating the floor on the first set gives $m\mathbb{E}(g1_A) \geq bq$, whence $m \geq bq$. Also $M(1 - q) \leq a \leq 1$. Since $g \geq m$ and $hg \geq b$ with $h \leq M$, we get $g \geq \max\{bq, b(1 - q)\} \geq b/2$. Thus $r = \pi/(Ng^2) \leq 4\pi/b^3$. The constant case is immediate.

For one bit, scaling removes an irrelevant constant, so take $H = 1 + tx$, $|t| < 1$. The normalized densities are $\pi = 1 + Ax$ and $r = 1 + Bx$, where

$$A = \frac{(1 + \rho)t}{1 + \rho t^2}, \qquad B = \frac{(1 - \rho)t}{1 - \rho t^2}.$$

Cross multiplication gives $|B| \le |A|$, since the difference of the two relevant positive numerators is $2\rho(1-t^2) > 0$. The function $[(1+c)^p + (1-c)^p]/2$ is even and nondecreasing for $c \in [0,1)$ and $p \ge 1$. This proves the one-bit assertion. Tensorization of the noise, the energy, N and the moments proves the product assertions. □

3 Inverse moments and subcritical transfer

In this section scale h so $Z = 1$, put $Q = P_{\sqrt{\rho}}$ and set $B = 2^{1/s}$. Assume $\pi = hg \ge b > 0$ and $\mathbb{E}\pi^s \le 2$; necessarily $b \le 1$. Define

$$R = 1 + \rho^{-1}, \qquad D_0 = 1 + \rho^{-3}, \qquad D_* = 2\rho^{-2} + \rho^{-3} - 1,$$
$$F(q) = 1 + \frac{sq/(s+q) - 1}{\rho^4}, \qquad q_\pm = \frac{1+\rho}{\rho^3}(1 + \rho^2 \pm \sqrt{1+\rho^2}).$$

Let $q_0 = D_*$, $q_{n+1} = F(q_n)$, and

$$A_n = b^{-1}(B/b)^n, \qquad t_n = \frac{sq_n}{q_n + 2s}, \qquad t_+ = \frac{sq_+}{q_+ + 2s}, \qquad K_n = Bb^{-3}(B/b)^{2n}.$$

Theorem 3.1. *For every L and $n \ge 0$,*

$$N \ge b, \quad gPg \ge b, \quad \left\|g^{-1}\right\|_{D_0} \le b^{-1}, \quad \left\|g^{-1}\right\|_{q_n} \le A_n, \quad \|r\|_{t_n} \le K_n.$$

Here $q_- < D_* < q_+$, $q_n \uparrow q_+ < \infty$, and $1 < t_0 < t_n < t_+ < s$ for $n > 0$. These conclusions need no matching. If additionally $\mathbb{E}\pi\chi_S = 0$ for $1 \le |S| \le k$, put

$$\eta_0 = \sqrt{\max\{\rho, (1-\rho)/2\}} < 1, \qquad e_k = B\eta_0^{k+1}.$$

For $1 < t < t_n$,

$$\|r - \pi\|_t \le (2e_k)^\alpha (K_n + B)^{1-\alpha}, \qquad \alpha = \frac{t_n - t}{t(t_n - 1)}. \tag{3}$$

Thus for every fixed $1 < t < t_+$, $\|r - \pi\|_t \to 0$ whenever $k \to \infty$, in particular for $k = k_L$.

We use the classical scalar Boolean hypercontractive inequalities [2]: for $1 < p \le q < \infty$, $\|P_\theta f\|_q \le \|f\|_p$ when $\theta^2 \le (p-1)/(q-1)$. For strictly positive f and finite $\beta < \alpha < 0$, reverse hypercontractivity gives $\|P_\theta f\|_\beta \ge \|f\|_\alpha$ when $\theta^2 \le (1-\alpha)/(1-\beta)$. Negative exponents denote generalized means $(\mathbb{E}f^a)^{1/a}$, not norms. In particular $\|f\|_{-a} = 1/\left\|f^{-1}\right\|_a$ for $a > 0$. No zero-exponent convention is needed in the uses below.

Proof of the inverse-moment and bootstrap assertions. The preceding proof gives $a = \mathbb{E}h \in [\sqrt{b}, 1]$ and $N \ge b$. Applying P to $h \ge b/g$ and using Jensen in each transition row yields

$$g \ge bP(1/g) \ge b/(Pg), \qquad g^{-1} \le (Pg)/b. \tag{4}$$

This is a consequence of the actual stationary floor, not an assumed bound on $\inf g$. Since $\|Qh\|_2 = 1$ and $Pg = P_{\rho^{3/2}}Qh$, ordinary hypercontractivity gives $\|Pg\|_{D_0} \le 1$, proving the first inverse bound. Likewise $g = P_{\sqrt{\rho}}Qh$ gives $\|g\|_R \le 1$. The floor yields $h^{-1} \le g/b$, so $\|h\|_{-R} \ge b$. Reverse hypercontractivity with $\alpha = -R$, $\beta = -D_*$ and $\theta = \rho$ applies because $D_* = (1+R)/\rho^2 - 1 > R$. It gives $\|g\|_{-D_*} \ge b$, or $\left\|g^{-1}\right\|_{D_*} \le b^{-1}$. Also $D_* - D_0 = 2(\rho^{-2} - 1) > 0$.

If $q > s/(s-1)$ and $\|g^{-1}\|_q \le A$, put $p = sq/(s+q) > 1$. Holder applied to $h = \pi/g$ gives $\|h\|_p \le BA$. As $Pg = P_{\rho^2}h$, hypercontractivity and (4) give

$$\left\|g^{-1}\right\|_{F(q)} \le b^{-1}\|Pg\|_{F(q)} \le (B/b)A.$$

The map has

$$F(q) - q = \frac{-\rho^4 q^2 + (s-1)(1-\rho^4)q - s(1-\rho^4)}{\rho^4(s+q)}, \qquad F'(q) = \frac{s^2}{\rho^4(s+q)^2} > 0.$$

The numerator has roots q_-, q_+: its discriminant equals $4\rho^2 s(1-\rho^4)$, and $s(1-\rho^4) = (1+\rho)^2(1+\rho^2)$. Direct bounds show

$$\rho^3(D_* - q_-) = \rho - \rho^2 - 2\rho^3 + (1+\rho)\sqrt{1+\rho^2} \ge (1+2\rho)(1-\rho^2) > 0,$$
$$\rho^3(q_+ - D_*) = -\rho + \rho^2 + 2\rho^3 + (1+\rho)\sqrt{1+\rho^2} \ge 1 + \rho^2 + 2\rho^3 > 0.$$

Since $D_* > R = 2s/(s-1)$, every step is valid. Monotonicity of F gives $q_n < q_{n+1} < q_+$, and continuity identifies the limit as q_+. Iteration gives A_n. Finally

$$r = \frac{\pi}{Ng^2}, \qquad \frac{1}{t_n} = \frac{1}{s} + \frac{2}{q_n}$$

and Holder give $\|r\|_{t_n} \le b^{-1}BA_n^2 = K_n$. The inequalities between the t_n follow from $D_* > 2s/(s-1)$ and monotonicity of $q \mapsto sq/(q+2s)$. □

For completeness we supply the stationary matching argument underlying (3); it does not require a pointwise floor.

Lemma 3.2. *Under energy normalization, if $\|\pi\|_p \le B_p$ for a finite $p > 1+\rho$ and π has exact marginals through degree k, set $\eta_p^2 = \rho\max\{1, 1/(p-1)\} < 1$ and $e = B_p\eta_p^{k+1}$. Then $a^2 \ge (1+e^2)^{-1}$, $a^2 \le N \le 1$, and $\|r - \pi\|_1 \le 2e$.*

Proof. Put $v = 1 - a^2 = \|Qh - a\|_2^2 \ge 0$, and let $\Pi_{>k}$ be Walsh projection above degree k. We have

$$\|\Pi_{>k}g\|_{p'} \le \eta_p^{k+1}\sqrt{v}, \qquad p' = p/(p-1).$$

For $p' \le 2$ this is the high-degree multiplier bound for Q, followed by norm monotonicity. For $p' > 2$, factor $Q = P_tP_\beta$, where $t = 1/\sqrt{p'-1}$ and $\beta = \sqrt{\rho}/t = \eta_p < 1$. Hypercontractivity sends $P_t : L^2 \to L^{p'}$, and the high-degree multiplier of P_β supplies the stated bound. The noise operators commute with Walsh projections; no norm bound for the projection on $L^{p'}$ is presumed. Matching gives $\mathbb{E}\pi\Pi_{\le k}g = a$. Hence $\mathbb{E}\pi g \le a + e\sqrt{v}$. Cauchy–Schwarz for the measure $h\,d\mu$ gives $1 = (\mathbb{E}hg)^2 \le a\mathbb{E}hg^2 = a\mathbb{E}\pi g$. Therefore $v/a \le e\sqrt{v}$. If $v > 0$, division yields $v \le a^2e^2$; the $v = 0$ case is immediate. Thus $a^2 \ge (1+e^2)^{-1}$ and $2(1-a) \le e^2$.

For the upper bound on N, put $u = h/g$ and $A = M_{\sqrt{u}}PM_{\sqrt{u}}$. The positive vector $\sqrt{hg}$ has A-eigenvalue one. Perron–Frobenius and positive semidefiniteness therefore give $\|A\| = 1$. With $X = M_{\sqrt{u}}Q$, the matrices $XX^* = A$ and $X^*X = QM_uQ$ have the same eigenvalues. Testing the latter at the unit vector 1 yields $N = \mathbb{E}u \le 1$. Cauchy–Schwarz under π gives $a^2 \le N$ as above. Since $r = \pi/(Ng^2)$,

$$\mathbb{E}\sqrt{r\pi} = a/\sqrt{N} \ge a, \qquad \|\sqrt{r} - \sqrt{\pi}\|_2^2 = 2(1 - a/\sqrt{N}) \le e^2.$$

Multiply the difference of the square roots by their sum and apply Cauchy–Schwarz; the latter sum has norm at most two. This gives $\|r-\pi\|_1 \le 2e$. Total variation is half this L^1 distance, not the entire distance. □

Completion of Theorem 3.1. Apply Lemma 3.2 with $p = s > 1+\rho$, $B_p = B$. Its parameter becomes η_0 in the theorem. Since $t_n < s$, $\|r-\pi\|_{t_n} \le K_n + B$. Interpolation between 1 and t_n, with $1/t = \alpha + (1-\alpha)/t_n$, gives (3). For each fixed $t < t_+$ choose one finite n with $t < t_n$ before taking $L \to \infty$. □

The ceiling is a limitation of this particular iteration. Its starting exponent lies strictly between the two fixed points and increases only to finite q_+. Moreover $A_n = b^{-1}(B/b)^n$ diverges as $n \to \infty$. We have neither a uniform estimate at q_+ nor one at t_+. Even a finite inverse bound at q_+ would leave $t_+ < s$. None of these facts disproves a different method for the critical moment.

4 Full-chain estimates for code-supported potentials

Identify Walsh indices with vectors in $\mathbb{F}_2^L$, and let $\mathcal{C}$ be a nonzero binary linear code of dimension n and minimum nonzero Hamming weight $D \ge 2$. Set $M = 2^n$ and

$$\delta = \sum_{c \in \mathcal{C}\setminus\{0\}} \rho^{|c|} < 1, \quad \eta = M\rho^{D/2}, \quad \alpha = \frac{1-\delta}{1+\delta}, \quad B_{\mathcal{C}} = \frac{\rho+\eta}{1-\delta}, \quad \varepsilon_{\mathrm{gap}} = \frac{(1-\rho^2)\delta}{\rho(1-\delta)}.$$

Theorem 4.1. *For every strictly positive h whose Fourier support is contained in $\mathcal{C}$, with no size restriction on its quotient values,*

$$\max\{0, \rho - \varepsilon_{\mathrm{gap}}\} \le \lambda_2(K) \le \min\{1, \max\{1-\alpha, B_{\mathcal{C}}\}\}. \tag{5}$$

Both π and r have exact marginals through degree $D-1$, and

$$\alpha^2 \le r/\pi \le \alpha^{-2}. \tag{6}$$

For any sequence of such codes with fixed ρ and $\eta \to 0$, $\mathrm{gap}(K) \to 1-\rho$, uniformly over all strictly positive quotient potentials. If additionally $\mathbb{E}\pi^s \le 2$, then $\mathbb{E}r^s \le 2\alpha^{-2s}$. The cap is needed for this absolute moment bound, not for (5) *or* (6).

Proof. Choose a basis $c_1, \dots, c_n$ and put $z_i(x) = \chi_{c_i}(x)$. This quotient map is surjective with equal fibers, and $h(x) = H(z(x))$ for a positive function H. Writing $c(I) = \sum_{i\in I} c_i$, the quotient noise density relative to uniform w is

$$b_0(z,w) = \sum_{I \subseteq [n]} \rho^{|c(I)|}\chi_I(z)\chi_I(w).$$

This is the actual positive stochastic quotient of product noise. In particular $1-\delta \le b_0 \le 1+\delta$. With $a = \mathbb{E}H$ and $G = PH$ on the quotient, $a(1-\delta) \le G \le a(1+\delta)$. The quotient Doob kernel has density $b_0(z,w)H(w)/G(z)$ and minorizes α times the probability density $H(w)/a$. Thus its oscillation contraction is at most $1-\alpha$. Every nonconstant eigenfunction has nonzero oscillation, so all its non-Perron eigenvalues have modulus at most $1-\alpha$.

It remains to control the other sectors, because the quotient spectrum alone is insufficient. For each coset $t+\mathcal{C}$, let $V_t = \chi_t\{F(z) : F \text{ any quotient function}\}$. These spaces are mutually

orthogonal under $\pi\,d\mu$: products of quotient functions with π remain in the $\mathcal{C}$ algebra, and a Walsh character outside $\mathcal{C}$ has zero expectation against that algebra. They span all 2^L dimensions. Multiplication by h and by $1/g$ preserves each sector, as does noise, so each is K-invariant. Under the identification $\chi_t F \leftrightarrow F$, its operator is

$$K_t F(z) = \mathbb{E}_w\, b_t(z,w)\frac{H(w)}{G(z)}F(w), \qquad b_t(z,w) = \sum_I \rho^{|t+c(I)|}\chi_I(z)\chi_I(w).$$

For $t \notin \mathcal{C}$ this is a signed kernel, not a Markov kernel. Choose a minimum-weight representative with $d = |t| \geq 1$. If $d \leq D/2$, all other vectors in the coset have weight at least $D - d \geq D/2$, so

$$\sum_{v \in t+\mathcal{C}} \rho^{|v|} \leq \rho^d + (M-1)\rho^{D/2} \leq \rho + \eta.$$

If $d > D/2$, all weights exceed $D/2$ and the sum is at most η. Consequently $|b_t| \leq \rho + \eta \leq B_{\mathcal{C}} b_0$, whence $|K_t F| \leq B_{\mathcal{C}} K_0 |F|$. Supremum norms bound every sector eigenvalue in modulus by $B_{\mathcal{C}}$. Together with the quotient bound and the global Markov bound one this proves the upper inequality in (5).

For the lower inequality, any singleton e_i is outside $\mathcal{C}$ because $D \geq 2$. In $b_{e_i} - \rho b_0$, a codeword not containing i contributes zero; a codeword containing i contributes coefficient $(1-\rho^2)\rho^{|c|-1}$. Hence

$$|b_{e_i} - \rho b_0| \leq (1-\rho^2)\delta/\rho \leq \varepsilon_{\mathrm{gap}} b_0.$$

It follows that $|K_{e_i} 1 - \rho| \leq \varepsilon_{\mathrm{gap}}$. The original character χ_{e_i} has mean zero under π and norm one, since π is in the $\mathcal{C}$ algebra. Its Rayleigh quotient is at least $\rho - \varepsilon_{\mathrm{gap}}$. The variational principle and positive definiteness yield the lower bound.

The algebra of quotient functions is closed under multiplication and positive inversion. Thus π and r are both $\mathcal{C}$-supported; their unit means imply the stated exact marginals. Let $\Pi = HG/Z$ be the quotient stationary density. Then

$$\mathbb{E}_\Pi G^{-2} = N/Z, \qquad r/\pi = \frac{G^{-2}}{\mathbb{E}_\Pi G^{-2}}.$$

The interval $a(1-\delta) \leq G \leq a(1+\delta)$ proves (6). Finally $\delta \leq (M-1)\rho^D \leq \eta\rho^{D/2} \leq \eta$. If $\eta \to 0$, then $\delta \to 0$, $\alpha \to 1$, $B_{\mathcal{C}} \to \rho$ and $\varepsilon_{\mathrm{gap}} \to 0$, uniformly in H. This proves the limiting and moment assertions. □

5 Disjoint blocks and a strictly larger potential class

5.1 The sharper direct block regime

Let $L = nm + d$ with $n \geq 1$, $m \geq 2$ and $d \geq 0$, and let $h = H(z_1, \ldots, z_n) > 0$, where z_i is the parity of the ith disjoint block of size m. Residual coordinates are unused. This is a special code algebra, but direct factorization yields better bounds than the coarse distance–cardinality estimate.

Proposition 5.1. *Put $\theta = \rho^m$, $\ell = (1-\theta)^n$, $u = (1+\theta)^n$, and*

$$\alpha_b = \ell/u, \qquad \beta_b = \frac{\rho + \rho^{m-1}}{1-\theta}, \qquad \varepsilon_{\mathrm{gap,b}} = \frac{(1-\rho^2)\rho^{m-1}}{1-\theta}.$$

For every positive H,

$$\max\{0, \rho - \varepsilon_{\mathrm{gap,b}}\} \le \lambda_2(K) \le \min\{1, \max\{1-\alpha_b, \beta_b\}\}, \qquad \alpha_b^2 \le r/\pi \le \alpha_b^{-2}.$$

Both densities have exact marginals through degree $m-1$. If $m \to \infty$ and $n\rho^m \to 0$, the gap tends uniformly in $H > 0$ to $1-\rho$.

Proof. The quotient noise density is $b_0(z,w) = \prod_i (1 + \theta z_i w_i)$, so $\ell \le b_0 \le u$ and $a\ell \le G \le au$. The same minorization and moment arguments as in Theorem 4.1 give the quotient and density bounds, now with α_b. For a sector representative, remove complete blocks from its support. Let $j_i \in \{0, \dots, m-1\}$ be its number of coordinates in block i, and e its residual weight. Factoring the Fourier sum gives

$$b_t(z,w) = \rho^e \prod_i (\rho^{j_i} + \rho^{m-j_i} z_i w_i).$$

For $j_i = 0$ the factor is that of b_0. If $1 \le j_i \le m-1$, its absolute ratio to $1 + \theta z_i w_i$ is bounded by β_b: the numerator is at most $\rho^{j_i} + \rho^{m-j_i} \le \rho + \rho^{m-1}$. This last inequality follows because the convex function $x \mapsto \rho^x + \rho^{m-x}$ attains its maximum on $[1, m-1]$ at an endpoint. If $\beta_b < 1$, a nontrivial sector with at least one interior factor therefore has $|b_t| \le \beta_b b_0$; if it has only residual support, its ratio is $\rho^e \le \rho < \beta_b$. This bounds those eigenvalues. If $\beta_b \ge 1$, the global Markov bound one already gives the stated upper bound; no invalid inequality $\beta_b^j \le \beta_b$ is used.

For a singleton in one block, its factor differs from $\rho(1 + \theta z_i w_i)$ by $(1-\rho^2)\rho^{m-1} z_i w_i$. The other factors are unchanged, yielding $|b_t - \rho b_0| \le \varepsilon_{\mathrm{gap,b}} b_0$. The singleton Rayleigh argument gives the lower bound. All nonzero vectors in the block code have weight at least m, giving matching. Finally $-\log \alpha_b = n \log[(1+\theta)/(1-\theta)] \le 4n\theta$ once $\theta \le 1/2$. Thus $\alpha_b \to 1$, $\beta_b \to \rho$ and $\varepsilon_{\mathrm{gap,b}} \to 0$ in the stated regime. □

Structural inclusion must not be confused with inclusion of the quantitative regimes. Theorem 4.1 applied to this code asks for $2^n \rho^{m/2} \to 0$, whereas Proposition 5.1 requires only $n\rho^m \to 0$ with $m \to \infty$. We retain the direct bounds instead of claiming the coarser condition subsumes them. In particular Proposition 5.1 with $n = 2$ applies to Theorem 2.1.

5.2 Simplex potentials outside disjoint-block representations

Proposition 5.2. *For each $j \ge 2$ put $L_j = 2^j - 1$ and index coordinates by nonzero $v \in \mathbb{F}_2^j$. The binary simplex code*

$$\mathcal{C}_j = \{c_a : a \in \mathbb{F}_2^j\}, \qquad c_a(v) = a \cdot v,$$

has dimension j and every nonzero word has weight $D_j = 2^{j-1}$. It satisfies the asymptotic condition in Theorem 4.1 for every fixed $0 < \rho < 1$. On this length subsequence its matching degree eventually exceeds $\lfloor C\sqrt{L_j} \rfloor$ for every fixed $C > 0$. There are strictly positive potentials supported on this code that cannot be represented as functions of disjoint equal-size block parities of size $m \ge 2$, even after coordinate permutations or enlargement of the disjoint-block algebra.

Proof. The map $a \mapsto c_a$ is injective by inspection at the unit-vector coordinates. A nonzero linear form takes the value one at exactly half the points of $\mathbb{F}_2^j$, and the excluded zero vector has value zero. Thus the dimension and weight assertions follow directly. Here

$$\delta_j = (2^j - 1)\rho^{2^{j-1}}, \qquad \eta_j = 2^j \rho^{2^{j-2}} \longrightarrow 0.$$

In particular $\delta_j < 1$ eventually. Also $D_j - 1 = (L_j - 1)/2$, proving the matching assertion on the stated subsequence.

Choose

$$h_j = 1 + \varepsilon_{\mathrm{amp}} \sum_{c \in \mathcal{C}_j \setminus \{0\}} \chi_c, \qquad 0 < \varepsilon_{\mathrm{amp}} < 1/(2^j - 1).$$

This is strictly positive, and its Fourier support is exactly $\mathcal{C}_j$. The amplitude $\varepsilon_{\mathrm{amp}}$ is not the spectral error $\varepsilon_{\mathrm{gap}}$. If h_j were a function of disjoint block parities, uniqueness of the Walsh expansion would put $\mathcal{C}_j$ inside the linear span of those block indicators. Every word in that span is constant within each block and vanishes on unused coordinates. But $\mathcal{C}_j$ separates every pair of distinct coordinates v, w: choose a with $a \cdot (v + w) = 1$. Also every coordinate is nonzero in some codeword. Consequently no block can contain two coordinates and none can be unused. This contradicts $m \geq 2$, and is unchanged by coordinate permutations or by allowing more disjoint blocks. □

The code construction is classical [4]; its elementary parameter proof is included to avoid treating a coding reference as a new theorem. The strict separation is an existence statement for the full-support positive potentials just exhibited. It does not hold for every potential in the code algebra, since constants also belong to disjoint-block algebras. Nor does this argument assert a prescribed stationary floor or moment cap for the displayed amplitude family. Its use here is to establish that the full-chain code theorem covers positive potentials beyond the disjoint-block class. Arbitrary invertible changes of the original Boolean coordinates are not permitted replacements for permutations: they generally change Hamming weights and the noise operator.

6 Positive-rate families at every sufficiently large length

The simplex example gives a transparent strict-class witness on a length subsequence. Classical random linear coding also realizes the spectral mechanism at positive rate for every sufficiently large length. The additional column conditions used below ensure strict separation from disjoint-block representations for the same codes, rather than for separately chosen examples.

Let $H_2(d) = -d \log_2 d - (1 - d) \log_2(1 - d)$ denote binary entropy.

Theorem 6.1. *Fix $0 < \rho < 1$, $0 < d < 1/2$ and*

$$0 < R_{\mathrm{code}} < \min\left\{1 - H_2(d),\ \frac{d}{2} \log_2(1/\rho)\right\}. \tag{7}$$

For every sufficiently large integer L there exists a binary linear code $\mathcal{C}_L \subset \mathbb{F}_2^L$ with dimension $n_L = \lfloor R_{\mathrm{code}} L \rfloor$, minimum distance $D_L \geq \lceil dL \rceil$, and a generator matrix with nonzero, pairwise distinct columns. Set

$$a_* = (d/2) \log(1/\rho) - R_{\mathrm{code}} \log 2 > 0, \qquad E_L = e^{-a_* L}, \qquad C_\rho = \max\{1 + \rho, (1 - \rho^2)/\rho\}.$$

For every strictly positive potential h supported on the Fourier algebra of this code, and all sufficiently large L,

$$|\mathrm{gap}(K_h) - (1 - \rho)| \leq C_\rho \frac{E_L}{1 - E_L}. \tag{8}$$

The bound is uniform over all such positive h, with no amplitude, stationary floor or moment assumption. Both π and r have exactly uniform labelled marginals through degree $\lceil dL \rceil - 1$.

Moreover, for every fixed $0 < b < 1$ and finite $s > 1$, there are potentials in these same algebras, at every sufficiently large length, with energy one, $\pi \geq b$ and $\mathbb{E}\pi^s \leq 2$, whose Fourier

support is the full code and which admit no disjoint equal-size block-parity representation with block size at least two. In particular these witnesses satisfy (1) *for its critical exponent and every fixed $C > 0$, eventually, and have the gap estimate* (8).

Proof. Both upper bounds in (7) are positive, so the allowed rate interval is nonempty. Put $n = \lfloor R_{\mathrm{code}} L \rfloor$, eventually at least one, and sample an n by L matrix G with independent fair binary entries. This is a probabilistic existence argument, not an executed random search. For each fixed nonzero $a \in \mathbb{F}_2^n$, the entries of aG are independent fair bits. The different vectors aG need not be independent. A union bound gives

$$\Pr\{\exists a \neq 0 : \ |aG| < dL\} \leq (2^n - 1)2^{-L} \sum_{0 \leq i < dL} \binom{L}{i} \leq 2^{n - L(1 - H_2(d))}.$$

For the last step, since $d < 1/2$ and $i \leq dL$,

$$d^i (1-d)^{L-i} \geq d^{dL} (1-d)^{(1-d)L} = 2^{-LH_2(d)}.$$

Multiply this lower bound by $\binom{L}{i}$, sum over the stated integer indices, and bound the corresponding binomial probability sum by one. This proves the entropy bound without assuming dL integral. The complement of the strict bad event gives weights at least $\lceil dL \rceil$. In particular a nonzero kernel vector is impossible, so G has full row rank and its row-image code has dimension n.

Each column is uniform in $\mathbb{F}_2^n$. A second union bound gives

$$\Pr\{\text{a zero column or an equal pair of columns}\} \leq \left[L + \frac{L(L-1)}{2} \right] 2^{-n}.$$

The sum of these two error bounds tends to zero. Indeed the first is at most $2^{-L(1 - H_2(d) - R_{\mathrm{code}})}$ and the second is at most $2[L + L(L-1)/2]2^{-R_{\mathrm{code}} L}$. For every sufficiently large integer L, their sum is less than one, so a single matrix has simultaneously the distance, rank and column properties. Select one for each such length. No efficient construction or nested family is asserted.

For the resulting code, $M_L = 2^{n_L}$ and $D_L \geq dL$ give

$$\eta_L = M_L \rho^{D_L/2} \leq E_L, \qquad \delta_L \leq M_L \rho^{D_L} \leq \eta_L \leq E_L.$$

Eventually $D_L \geq 2$ and $\delta_L < 1$, so Theorem 4.1 applies to every positive code potential. For $E_L \leq \min\{1/2, \rho/2\}$,

$$\begin{aligned} 1 - \alpha &= \frac{2\delta_L}{1 + \delta_L} \leq 2E_L \leq \rho, \\ B_{\mathcal{C}_L} &= \frac{\rho + \eta_L}{1 - \delta_L} \leq \rho + (1 + \rho) \frac{E_L}{1 - E_L}, \\ \varepsilon_{\mathrm{gap}} &\leq \frac{1 - \rho^2}{\rho} \frac{E_L}{1 - E_L}. \end{aligned}$$

The two sides of (5) therefore give (8); the clipping by zero and one only improves the estimates. The matching assertion follows from the same theorem and $D_L \geq \lceil dL \rceil$. In particular this matching degree eventually exceeds $\lfloor C\sqrt{L} \rfloor$ for every fixed $C > 0$.

It remains to construct the admissible strict-class witnesses. Nonzero columns ensure that every coordinate is active in some codeword; distinct columns ensure that every two coordinates

are separated by some codeword. For the latter, if the two columns are $v \neq w$, choose a linear functional detecting $v + w \neq 0$. For $0 < \tau < 1$, choose

$$h_0 = 1 + \frac{\tau}{M_L - 1} \sum_{c \in \mathcal{C}_L \setminus \{0\}} \chi_c.$$

Then $1 - \tau \leq h_0 \leq 1 + \tau$, and its Fourier support is exactly $\mathcal{C}_L$. If h_0 were a function of any number of disjoint equal-size block parities, its full support would force $\mathcal{C}_L$ into the span of their whole-block indicators. Every word in that span is constant within a block and zero on unused coordinates. Pair separation excludes any block of size at least two, and activity excludes residual coordinates. This rules out every representation allowed in Proposition 5.1, including larger quotient algebras and coordinate permutations.

For prescribed $0 < b < 1$ and finite $s > 1$, choose $\tau > 0$ sufficiently small that

$$A_\tau = \left(\frac{1+\tau}{1-\tau}\right)^2 \leq \min\{1/b, 2^{1/s}\}.$$

The right side exceeds one, so such a single positive τ, independent of L, exists. Positivity and the Markov property put both h_0 and Ph_0 in $[1-\tau, 1+\tau]$. Hence their product and its mean Z lie in $[(1-\tau)^2, (1+\tau)^2]$. Set $h = h_0/\sqrt{Z}$. This preserves its support and kernel, gives $\mathbb{E}hPh = 1$, and yields

$$A_\tau^{-1} \leq \pi = \frac{h_0 P h_0}{Z} \leq A_\tau, \qquad \pi \geq b, \qquad \mathbb{E}\pi^s \leq A_\tau^s \leq 2.$$

Thus all assertions concern the same normalized potential and its actual stationary density. The full-chain bound was already uniform in the potential and remains valid for these witnesses. □

The entropy-volume and random-linear-code ingredients are classical Gilbert–Varshamov arguments [5]; the proof is included with the precise rounding and simultaneous column requirements used here. The rate R_{code} is distinct from the inverse-moment exponent R of Section 3. The theorem supplies all-large-length positive-rate realizations of the code-restricted mechanism. It does not remove the code-algebra hypothesis, establish a general critical weighting cap, or provide a collapsing example. The strict-class conclusion is existential for the specified full-support potentials, not universal for all positive functions on the code. The prescribed-floor witness assertion excludes $b = 1$ and makes no uniform endpoint claim as ρ approaches zero or one.

7 Dimension-free spectral windows

Throughout this section, the noise parameter $0 < \rho < 1$ and stationary floor $b > 0$ are fixed. On the uniform Boolean cube, write $P = P_\rho$, $Q = P_{\sqrt{\rho}}$, $g = Ph$, $\pi = hg$, $u = h/g$, and $T = QM_uQ$. We retain $h > 0$, $\mathbb{E}hPh = 1$, $\pi \geq b$, $\mathbb{E}\pi^s \leq 2$, $s = (1+\rho)/(1-\rho)$, and exact stationary matching through degree k. These hypotheses are not removed below.

Lemma 7.1. *For every fixed $0 < \rho < 1$, a finite index n, independent of L, satisfies $t_n > 1 + \rho$ in Theorem 3.1.*

Proof. The inequality $t_n > 1 + \rho$ is equivalent to $q_n > q_c := 2(1+\rho)/\rho$. Direct substitution in the already defined bootstrap map gives

$$F(q_c) - q_c = \frac{(1-\rho^2)(3-\rho^2)}{\rho^3(2-\rho)} > 0.$$

The numerator of $F(q) - q$ is a negative-leading quadratic with roots q_-, q_+. Thus $q_- < q_c < q_+$. Since q_n increases to q_+, for some finite n, $q_n > q_c$. Fix such an index before taking dimension or matching-degree limits. □

Set $p_0 = 1 + \rho$, $t = t_n$ and $K_t = K_n$ at this fixed index. Then $p_0 < t < s$, and $\|u\|_t = N\|r\|_t \leq K_t$ by Lemma 3.2 and Theorem 3.1. With $B = 2^{1/s}$, $\eta_0 = \sqrt{\max\{\rho, (1-\rho)/2\}}$, $e = B\eta_0^{k+1}$, $p = (p_0 + t)/2$, and $\alpha = (1/p - 1/t)/(1 - 1/t)$, the discrepancy satisfies

$$\|u - \pi\|_p \leq \Delta_k := (2e)^\alpha (K_t + B)^{1-\alpha} + Be^2 \leq C_p \eta_0^{\alpha(k+1)}, \qquad C_p = (2B)^\alpha (K_t + B)^{1-\alpha} + B^3.$$

This uses $u - \pi = N(r - \pi) + (N - 1)\pi$, $1 - N \leq e^2$ from Lemma 3.2, and interpolation in Theorem 3.1. No endpoint estimate is used.

For a polynomial F of degree at most d and finite $r \geq 2$, classical Boolean hypercontractivity gives

$$\|QF\|_r \leq \max\{1, \sqrt{\rho(r-1)}\}^d \|F\|_2.$$

Indeed, multiply each degree-j Fourier coefficient by A^j, $A = \sqrt{\rho(r-1)}$, and then apply $P_{1/\sqrt{r-1}}$. The composition is Q and the first multiplier has L^2 norm at most $\max(1, A)^d$ on the polynomial subspace. The multiplier for $A > 1$ is algebraic, not a Markov kernel. This standard argument is the one used in O'Donnell, Theorem 9.21 [6].

Let R_d be the orthogonal projection onto degrees at most d, including constants. Set

$$r_1 = \frac{pp_0}{p - p_0}, \quad z = \frac{p_0 + \min(s, 2)}{2}, \quad r_2 = \frac{sz}{s - z}, \quad A_i = \sqrt{\rho(r_i - 1)}, \quad \beta = \sqrt{\frac{\rho}{z - 1}}.$$

Then $A_1, A_2 > 1$ and $0 < \beta < 1$.

Theorem 7.2. *For $0 \leq d \leq k$,*

$$\|(T - P)R_d\|_{2\to 2} = \|R_d(T - P)\|_{2\to 2} \leq \varepsilon_{k,d} := \Delta_k A_1^d + (B + 1)\beta^{k+1-d} A_2^d. \tag{9}$$

Proof. To prove this, write $(T - P)F = Q[(u - \pi)QF] + Q[(\pi - 1)QF]$ for $F = R_d F$. The first term is bounded by $\|u - \pi\|_p \|QF\|_{r_1}$, using $Q : L^{p_0} \to L^2$ and Hölder's inequality. The second product has no Fourier coefficients below degree $k + 1 - d$. Factor $Q = P_\beta P_{\sqrt{z-1}}$: the second factor contracts L^z to L^2, and the first damps that support by β^{k+1-d}. Hölder with s, r_2 and the polynomial bound complete the proof. The argument controls all low-degree inputs at once, without summing individual columns or introducing a binomial factor. Self-adjointness gives the row estimate. □

Choose

$$0 < c < \min\left\{1, \frac{\alpha \log(1/\eta_0)}{\log A_1}, \frac{\log(1/\beta)}{\log(A_2/\beta)}\right\}, \qquad d = \lfloor ck \rfloor.$$

Then the right side of (9) is bounded by $C_p \eta_0^\alpha e^{-\gamma_1 k} + (B + 1)\beta e^{-\gamma_2 k}$, where $\gamma_1 = \alpha \log(1/\eta_0) - c \log A_1 > 0$ and $\gamma_2 = \log(1/\beta) - c \log(A_2/\beta) > 0$. This dimension-free estimate requires only $k \to \infty$.

The actual spectral representation is important here. With $X = M_{\sqrt{u}} Q$, one has $XX^* = M_{\sqrt{u}} P M_{\sqrt{u}}$ and $X^* X = T$, so T has the actual spectrum of K and norm one. Its unit Perron vector is $w = Qh$, since $\|Qh\|_2^2 = 1$ and $TQh = QM_u Ph = Qh$. Lemma 3.2 gives $\|1 - w\|_2^2 = 2(1 - \mathbb{E}h) \leq e^2$. For a unit non-Perron eigenvector ψ of T with $\lambda \geq \rho + \zeta$, fixed $\zeta > 0$,

$$\|R_d \psi\|_2^2 \leq e^2 + (\varepsilon_{k,d}/\zeta)^2.$$

The nonconstant part follows by inverting $\lambda - P$ on the low subspace; the constant part follows from orthogonality to Qh and $\|1-Qh\|_2 \le e$. For $1 \le d \le \min(k, L-1)$, let $H_d = (I-R_d)T(I-R_d)$ on the high subspace. The block comparison yields

$$|\lambda_2(K) - \max\{\rho, \|H_d\|\}| \le 2\varepsilon_{k,d}.$$

Indeed, the low diagonal and off-diagonal errors each have norm at most $\varepsilon_{k,d}$; min-max applies to the resulting self-adjoint perturbation. This is a reduction to the actual high block, not a bound on its norm.

8 Necessary localization and a failed generic tail repair

Retain the setting of Section 7, with $d = \lfloor ck \rfloor$, fixed $\zeta > 0$, and all other parameters fixed. Let $Kf = \lambda f$ be an actual real eigenfunction with $\lambda \ge \rho + \zeta$, normalized by $\mathbb{E}\pi f = 0$ and $\mathbb{E}\pi f^2 = 1$, set $\psi = Q(hf)/\sqrt{\lambda}$. Then the exact pointwise identity is

$$u|Q\psi|^2 = \lambda \pi f^2.$$

Writing $\delta_k = e^2 + (\varepsilon_{k,d}/\zeta)^2$ and $U(M) = \{u > M\}$, with $\pi(A) := \mathbb{E}[\pi \mathbf{1}_A]$ for sets A, gives

$$\mathbb{E}[\pi f^2 \mathbf{1}_{U(M)^c}] \le \frac{M(\delta_k + \rho^{d+1})}{\lambda}, \qquad \pi(U(M)) \le B K_t^{t/s'} M^{-t/s'}, \quad s' = \frac{s}{s-1}.$$

Indeed $\|\psi\|_2 = 1$ and $\langle \psi, Qh \rangle = 0$ follow from $P(hf) = \lambda g f$ and the stationary normalization. The multiplier formula gives $\|Q\psi\|_2^2 \le \delta_k + \rho^{d+1}$; integrate the identity on $\{u \le M\}$. For the probability estimate, $\mu(U(M)) \le K_t^t M^{-t}$ and Hölder with $\|\pi\|_s \le B$ complete the proof. With $\gamma_0 = \min\{\log(1/\eta_0), \gamma_1, \gamma_2\}$, choose $0 < \theta < \min\{\log(1/\rho), 2\gamma_0/c\}$ and $M = e^{\theta d}$. As $k \to \infty$, the complement's stationary squared mass tends to zero, while the stationary probability $\pi(U(M))$ is at most a fixed constant times $e^{-\theta t d/s'}$. Both strict restrictions on θ are needed for this argument. At $k \sim C\sqrt{L}$, this is an exponential-in-$\sqrt{L}$ probability scale. It neither supplies a cut nor rules out these modes.

For completeness, this also forces divergence of each fixed higher stationary moment. For $2 < \ell < \infty$, set $a_k = M(\delta_k + \rho^{d+1})/\lambda \to 0$ and $C_t = B K_t^{t/s'}$. Hölder on the stationary probability space gives, for all sufficiently large k,

$$\|f\|_{L^\ell(\pi)} \ge \sqrt{1 - a_k}\, C_t^{-(1/2 - 1/\ell)} \exp\!\left(\frac{\theta t d}{s'} (1/2 - 1/\ell) \right).$$

Indeed, the squared mass on $U(M)$ is at least $1 - a_k > 0$, and is at most $\|f\|^2_{L^\ell(\pi)} \pi(U(M))^{1-2/\ell}$. This necessary growth is not a contradiction: no uniform upper bound on these eigenfunction norms has been established. A generic tail-space Hölder repair is unavailable. For fixed $1 + 1/\rho < q < \infty$ and nonnegative integer cutoffs $d_L = o(L)$, set $m = d_L + 1$, $\ell = L - m$, and

$$F_L(x) = \chi_{[m]}(x) \prod_{j=m+1}^{L} \frac{1 + \epsilon x_j}{\sqrt{1 + \epsilon^2}}.$$

For sufficiently large L, this has unit L^2 norm and no coefficients through degree d_L. Its smoothed norm is

$$\|QF_L\|_q = \rho^{m/2} \kappa^\ell, \quad \kappa = \frac{\left(((1 + \sqrt{\rho}\epsilon)^q + (1 - \sqrt{\rho}\epsilon)^q)/2\right)^{1/q}}{\sqrt{1 + \epsilon^2}}.$$

The expansion $\log\kappa = \frac{1}{2}(\rho(q-1)-1)\epsilon^2 + O(\epsilon^4)$ makes $\kappa > 1$ for a sufficiently small fixed $\epsilon \in (0,1)$. Thus $L^{-1}\log\|QF_L\|_q \to \log\kappa > 0$. This classical tensor sharpness mechanism refutes a dimension-uniform $L^2 \to L^q$ estimate on the entire high-degree space. It is not an admissible Doob counterexample, and does not contradict same-exponent $L^p \to L^p$ tail smoothing [7, 8]. For the available $t < s$, one has $2t' > 2s' = 1 + 1/\rho$, so this obstruction applies to the naive Hölder attempt to control the remaining high block using only $\|u\|_t$. It does not apply to the low-degree-input polynomial estimate used above.

9 Weak relaxation and estimation without matching

This section uses only the stationary floor and critical moment cap. It does not require exact coordinate marginals. Normalize $\mathbb{E}hPh = 1$, put $\nu = \pi\mu$ with $\pi = hg$, and retain the actual kernel K from Section 1. For real f set $O = \mathrm{osc}(f) = \max f - \min f$ and

$$\mathcal{D}_K(f) = \langle f, (I-K)f\rangle_\nu = \frac{1}{2}\mathbb{E}_{\mathfrak{m}}h(x)h(y)(f(x)-f(y))^2, \qquad \mathfrak{m}(x,y) = \mu(x)P(x,y).$$

Here $\mathfrak{m}$ is a probability measure on pairs. Write $\mathcal{D}_P$ for the same expression with the factor $h(x)h(y)$ removed. All constants below depend on the fixed noise and floor, and not on the dimension or the potential. Within this section, N denotes the integer sample length, not the weighting normalizer $\mathbb{E}(h/g)$ of Section 1; C and c below denote the local weak-Poincare constants, not earlier matching or window constants.

9.1 An inverse edge moment and a weak Poincare inequality

Theorem 9.1. *Assume $h > 0$, $\mathbb{E}hPh = 1$, $\pi \geq b > 0$, and $\mathbb{E}\pi^s \leq 2$, where $s = (1+\rho)/(1-\rho)$. Define*

$$s' = \frac{1+\rho}{2\rho}, \quad B = 2^{1/s}, \quad a = \frac{2\rho}{(1+\rho)^2}, \quad p = \frac{1-a}{a} = \frac{1+\rho^2}{2\rho},$$
$$C = B(1-\rho)^{-1/s'}b^{-2a}, \qquad c = C^{-1/a} = b^2(1-\rho)^{1+\rho}2^{-(1-\rho^2)/(2\rho)}.$$

Then every real f, every $\epsilon > 0$, and every integer $n \geq 0$ satisfy

$$\mathrm{Var}_\nu(f) \leq C\mathcal{D}_K(f)^a O^{2(1-a)}, \tag{10}$$
$$\mathrm{Var}_\nu(f) \leq c^{-1}\epsilon^{-p}\mathcal{D}_K(f) + \epsilon O^2, \tag{11}$$
$$\mathrm{Var}_\nu(K^n f) \leq O^2(1+pcn)^{-1/p}. \tag{12}$$

Proof. Let $Q = P_{\sqrt{\rho}}$. Energy normalization gives $\|Qh\|_2 = 1$. Ordinary Boolean hypercontractivity, in the form $P_\theta : L^u \to L^v$ for $1 < u \leq v < \infty$ when $\theta^2 \leq (u-1)/(v-1)$, gives $\|g\|_R \leq 1$ for $R = 1 + 1/\rho$, since $g = Q(Qh)$. The original floor gives $1/h \leq g/b$, hence $\|1/h\|_R \leq b^{-1}$. These inverse moments are derived, not assumed.

Set $r_0 = 1/\rho$, $w(x,y) = h(x)h(y)$, and $v = h^{-r_0}$. A second use of hypercontractivity, $Q : L^{1+\rho} \to L^2$, yields

$$\mathbb{E}_{\mathfrak{m}}w^{-r_0} = \langle v, Pv\rangle_\mu = \|Qv\|_2^2 \leq \|v\|_{1+\rho}^2 = \|1/h\|_{r_0(1+\rho)}^{2r_0} \leq b^{-2r_0}.$$

For $O > 0$, let $z = (f(x)-f(y))^2/O^2 \in [0,1]$ and $\alpha = r_0/(1+r_0) = 1/(1+\rho)$. Hölder's inequality gives

$$\mathbb{E}_{\mathfrak{m}}z \leq (\mathbb{E}_{\mathfrak{m}}wz)^\alpha(\mathbb{E}_{\mathfrak{m}}w^{-r_0}z)^{1-\alpha} \leq (2\mathcal{D}_K(f)/O^2)^\alpha b^{-2\alpha}.$$

Consequently

$$\mathcal{D}_P(f) \le 2^{\alpha-1} b^{-2\alpha} \mathcal{D}_K(f)^\alpha O^{2(1-\alpha)} \le b^{-2\alpha} \mathcal{D}_K(f)^\alpha O^{2(1-\alpha)}.$$

We deliberately discard a factor at most one. The Walsh spectrum gives $\mathrm{Var}_\mu(f) \le \mathcal{D}_P(f)/(1-\rho)$. For $m_f = \mathbb{E}_\mu f$, we have $|f - m_f| \le O$ and

$$\mathrm{Var}_\nu(f) \le \mathbb{E}_\mu \pi (f-m_f)^2 \le B (\mathbb{E}_\mu |f-m_f|^{2s'})^{1/s'} \le B O^{2/s} \mathrm{Var}_\mu(f)^{1/s'}.$$

Combining proves (10), since $a = \alpha/s'$. For $y = \mathrm{Var}_\nu(f)/O^2$ and $e = \mathcal{D}_K(f)/O^2$, it says $e \ge c y^{1+p}$. If $y \le \epsilon$, (11) is immediate; otherwise $y \le c^{-1}\epsilon^{-p} e$, proving the same assertion.

For completeness, the discrete-time decay is derived directly. The conjugate $M_{\sqrt{h/g}} P M_{\sqrt{h/g}}$ is positive semidefinite, so the self-adjoint Markov operator K has spectrum in $[0,1]$. For $F = K^n f$ this implies

$$\mathrm{Var}_\nu(F) - \mathrm{Var}_\nu(KF) = \langle F, (I-K^2)F\rangle_\nu \ge \mathcal{D}_K(F).$$

Also $\mathrm{osc}(K^n f) \le O$. Therefore $y_n = \mathrm{Var}_\nu(K^n f)/O^2$ obeys $y_{n+1} \le y_n - c y_n^{1+p}$. For positive consecutive terms, convexity of $t \mapsto t^{-p}$ gives

$$y_{n+1}^{-p} - y_n^{-p} \ge p(y_n - y_{n+1}) y_n^{-p-1} \ge pc.$$

Summing and using $y_0 \le 1$ proves (12). If a variance vanishes, all subsequent variances vanish; if $O = 0$, all statements are immediate. □

9.2 Summable relaxation at noise one half

The floor identity used in Section 3 supplies a stronger inverse moment at $\rho = 1/2$. One additional positive-moment smoothing step then makes the bounded-observable decay summable.

Theorem 9.2. *Assume $\rho = 1/2$, $h > 0$, $\mathbb{E}hPh = 1$, $\pi \ge b > 0$ and $\mathbb{E}\pi^3 \le 2$. Put*

$$c_* = b^4 2^{-61/24}, \qquad C_F = 3 + 32/c_*.$$

For all real f and integers $n \ge 0$,

$$\mathrm{Var}_\nu(K^n f) \le O^2 (1 + (7/8)c_* n)^{-8/7}.$$

For the stationary chain (X_i) with transition K and all integers $N \ge 1$,

$$\mathrm{Var}\left(\frac{1}{N}\sum_{i=0}^{N-1} f(X_i)\right) \le \frac{C_F O^2}{N}.$$

Proof. The floor gives $h \ge b/g$, hence $g \ge bP(1/g) \ge b/(Pg)$ by rowwise Jensen. Thus $1/g \le Pg/b$. Since $Pg = P_{\rho^{3/2}} Qh$ and $\|Qh\|_2 = 1$, hypercontractivity at $\rho = 1/2$ gives $\|Pg\|_9 \le 1$ and $\|1/g\|_9 \le b^{-1}$. With $B = 2^{1/3}$, Hölder and $h = \pi/g$ yield $\|h\|_{9/4} \le B/b$. Hypercontractivity $P_{1/2} : L^{9/4} \to L^6$ gives $\|g\|_6 \le B/b$. The original floor again implies

$$\|1/h\|_6 \le H := B/b^2.$$

In the notation of the preceding proof, now set $v = h^{-4}$. Then

$$\mathbb{E}_{\mathfrak{m}} w^{-4} = \|Qv\|_2^2 \le \|v\|_{3/2}^2 = \|1/h\|_6^8 \le H^8.$$

Repeating the edge comparison with $\alpha = 4/5$, and again discarding $2^{\alpha-1} \le 1$, gives

$$\mathcal{D}_P(f) \le H^{8/5}\mathcal{D}_K(f)^{4/5}O^{2/5}.$$

Use the noise gap $1/2$ and the variance transfer with $s' = 3/2$ to obtain

$$\mathrm{Var}_\nu(f) \le C_*\mathcal{D}_K(f)^{8/15}O^{14/15}, \qquad C_* = B\,2^{2/3}H^{16/15}.$$

Here $C_*^{-15/8} = b^4 2^{-61/24} = c_*$. The discrete recursion in the preceding proof applies with $a_* = 8/15$ and $p_* = (1-a_*)/a_* = 7/8$, proving the claimed decay.

Center f by its ν-mean and let $\gamma_j = \mathrm{Cov}(f(X_0), f(X_j)) = \langle f, K^j f\rangle_\nu$. Positivity and self-adjointness give

$$0 \le \gamma_{2m+1} \le \gamma_{2m} = \mathrm{Var}_\nu(K^m f).$$

With $A = (7/8)c_*$, the decreasing bound satisfies

$$\sum_{m\ge 1}(1+Am)^{-8/7} \le \int_0^\infty (1+Ax)^{-8/7}\,dx = 7/A = 8/c_*.$$

Consequently

$$\sum_{j\ge 1}\gamma_j \le \mathrm{Var}_\nu(f) + 2\sum_{m\ge 1}\mathrm{Var}_\nu(K^m f) \le O^2(1+16/c_*).$$

The exact finite-N covariance identity now gives

$$\mathrm{Var}\left(\frac{1}{N}\sum_{i=0}^{N-1} f(X_i)\right) = \frac{N\gamma_0 + 2\sum_{j=1}^{N-1}(N-j)\gamma_j}{N^2} \le \frac{\gamma_0 + 2\sum_{j\ge1}\gamma_j}{N} \le \frac{(3+32/c_*)O^2}{N}.$$

No conditional summation or continuous-time substitution is used. □

9.3 Uniform initialization and mean-square error

Corollary 9.3. *Under the assumptions of Theorem 9.2, start the actual K-chain in the uniform reference law μ. For every real f and integer $N \ge 1$,*

$$\mathbb{E}_{\mu\text{-}start}\left[\left(\frac{1}{N}\sum_{i=0}^{N-1} f(X_i) - \nu(f)\right)^2\right] \le \frac{C_F O^2}{bN}.$$

This is a mean-square error bound including any initialization bias.

Proof. For an initial law η, the finite-path likelihood ratio relative to the stationary path is exactly

$$\frac{d\mathbb{P}_\eta^{(N)}}{d\mathbb{P}_\nu^{(N)}}(x_0,\dots,x_{N-1}) = \frac{\eta(x_0)}{\nu(x_0)}.$$

The transition factors cancel. Hence $\eta \le M\nu$ implies $\mathbb{E}_\eta Z \le M\mathbb{E}_\nu Z$ for every nonnegative path functional Z. For $\eta = \mu$, the original floor supplies $\mu(x)/\nu(x) = 1/\pi(x) \le 1/b$. Apply this comparison to the squared error and use Theorem 9.2; stationarity makes its variance equal to its mean-square error about $\nu(f)$. □

The required initial-density domination is proved from the floor for this specific initialization. No dimension-free domination is asserted for a fixed starting vertex. The uniform-start estimator need not be unbiased or stationary. No concentration inequality or optimality claim is made.

These theorems concern oscillation-controlled observables. In particular, (10) applied to a nonconstant eigenfunction only gives

$$1-\lambda \geq c\left(\frac{\mathrm{Var}_\nu(f)}{\mathrm{osc}(f)^2}\right)^p.$$

The ratio may be arbitrarily small along a dimension-varying family. Thus these results do not establish a dimension-uniform $L^2(\nu)$ gap, do not bound the killed operator of Section 8, and do not remove matching from the stronger open spectral question.

10 Full gaps from a positive potential background

A fixed lower bound on the normalized potential allows arbitrary upper heights and arbitrarily many values. We first prove this general sufficient condition and then recover the two-level class from its stationary floor. The lower potential bound is not derived here for arbitrary potentials satisfying the original stationary hypotheses.

Theorem 10.1. *Fix $0<\rho<1$ and $0<a\leq 1$. Let $h\geq a$ on $\{-1,1\}^L$, $L\geq 1$, normalize $\mathbb{E}hP_\rho h=1$, and assume $\mathbb{E}\pi^s\leq 2$, where $\pi=hP_\rho h$ and $s=(1+\rho)/(1-\rho)$. Then the actual Doob kernel has*

$$\mathrm{gap}(K)\geq G_J(\rho,a)>0$$

uniformly in L and h. No matching or upper bound on h is needed. The following explicit constants suffice; all symbols are local:

$$s'=s/(s-1),\qquad B=2^{1/s},\qquad c=a/2,\qquad M=B/a,\qquad \eta=\frac{c}{c+M},$$

$$t=\min\left\{2^{-s},\left(\frac{\eta^2 c}{4M}\right)^s\right\},\qquad \epsilon=c^2 t,$$

$$\gamma=\frac{c^2(1-\rho)\epsilon^{s'-1}}{2^{s'}B^{2s'}},\qquad \Phi_*=\min\{\eta/2,\gamma\},\qquad G_J(\rho,a)=\Phi_*^2/4.$$

The constants are positive and deliberately conservative.

Proof. Put $P=P_\rho$, $Q=P_{\sqrt{\rho}}$, $g=Ph$, and $\nu=\pi\mu$. Since $g\geq a$, the cap gives $\|h\|_s\leq B/a=M$. Write $h=c+z$, so $z\geq c>0$, $\|z\|_s\leq M$, and $g=c+Pz$.

Critical multiplication and the actual denominator. Ordinary hypercontractivity [10] and $1/(1+\rho)=1/2+1/(2s)$ yield

$$\|QM_{\sqrt{z}}f\|_2\leq\|\sqrt{z}f\|_{1+\rho}\leq\|z\|_s^{1/2}\|f\|_2\leq\sqrt{M}\|f\|_2.$$

The same inequality for a function supported on C gives $\|Q1_C\|_{2\to 2}\leq\mu(C)^{1/(2s)}$. All operator norms in this proof, except those explicitly assigned to $L^2(\nu)$, use $L^2(\mu)$ or its orthogonal direct sum. Set $A=M_{\sqrt{z}}PM_{\sqrt{z}}$ and $R=M_{Pz}$. Then $0\preceq A\preceq MI$. Reversibility also gives the exact identity

$$\langle v,(R-A)v\rangle_\mu=\frac{1}{2}\sum_{x,y}\mu(x)P(x,y)z(x)z(y)\left(\frac{v(x)}{\sqrt{z(x)}}-\frac{v(y)}{\sqrt{z(y)}}\right)^2\geq 0.$$

Thus $A \preceq R$. With $\theta = c/(c+M) = \eta$ and $\lambda = M/(c+M)$, convex combination gives

$$A \preceq \theta M I + (1-\theta)R = \lambda(cI + R).$$

Congruence and $g = c + Pz$ prove

$$M_{1/\sqrt{g}} A M_{1/\sqrt{g}} \preceq (1-\eta)I. \tag{13}$$

This uses the nonconstant denominator exactly and assumes no upper bound on z.

An exact two-component representation. Let $S = M_{\sqrt{h/g}} P M_{\sqrt{h/g}}$, the conjugate of K by the isometry $f \mapsto \sqrt{\pi} f$ from $L^2(\nu)$ to $L^2(\mu)$. Define

$$Wf = \left(\sqrt{c/h}\, f, \sqrt{z/h}\, f\right), \qquad W^*W = I.$$

For $T = WSW^*$, direct multiplication gives

$$T_{ij} = M_{\sqrt{h_i/g}} P M_{\sqrt{h_j/g}}, \qquad h_0 = c, \quad h_1 = z.$$

In particular T_{11} is the operator in (13). Both W and W^* commute with the same set restriction on each component. Hence the compressions satisfy

$$S_{CC} = W_C^* T_{CC} W_C, \qquad \|S_{CC}\| \le \|T_{CC}\|. \tag{14}$$

The larger space supplies an exact identity for the original operator, not a new assumption or a substitute chain.

Small cuts. The multiplier $\sqrt{c/g}$ is at most one. Therefore

$$\|T_{00,CC}\| \le \|1_C P 1_C\| \le \mu(C)^{1/s}, \qquad \|T_{11,CC}\| \le 1 - \eta.$$

For the cross block factor $P = Q^2$. Bound the left diagonal multiplier by one and the far-right $g^{-1/2}$ by $c^{-1/2}$. The latter commutes with $\sqrt{z}$ and 1_C, without being moved through Q. Thus

$$\|T_{01,CC}\| \le c^{-1/2} \|1_C Q\| \, \|Q M_{\sqrt{z}} 1_C\| \le \sqrt{M/c}\, \mu(C)^{1/(2s)}.$$

If $0 < \nu(C) \le \epsilon$, then $\pi = hg \ge a^2 \ge c^2$ gives $\mu(C) \le t$. The diagonal blocks are bounded by $1/2$ and $1-\eta$, respectively, while the cross block is bounded by $\eta/2$. Since $0 < \eta < 1/2$, the block-diagonal part has norm at most $1-\eta$. The self-adjoint off-diagonal part has norm $\|T_{01,CC}\|$, so $\|T_{CC}\| \le 1 - \eta/2$.

Define the actual outward stationary flow

$$F_h(C) = \sum_{x \in C,\, y \notin C} \mu(x) P(x,y) h(x) h(y) = \mathcal{D}_K(1_C).$$

The vector $\sqrt{\pi}\, 1_C$ has squared norm $\nu(C)$. Using (14) and conjugacy yields

$$F_h(C) = \nu(C) - \langle 1_C, K 1_C \rangle_\nu \ge (\eta/2)\nu(C). \tag{15}$$

Every such cut is covered, regardless of the values of h within it.

Remaining cuts and the full gap. For $\epsilon \le \nu(C) \le 1/2$, the actual edge measure and $h \ge c$ give

$$F_h(C) \ge c^2 \mathcal{D}_P(1_C) \ge c^2(1-\rho)\mu(C)\mu(C^c).$$

The cap implies $\nu(E) \le B\mu(E)^{1/s'}$ for every E. Hence

$$\mu(C) \ge (\nu(C)/B)^{s'}, \qquad \mu(C^c) \ge (1/(2B))^{s'},$$

and consequently

$$\frac{F_h(C)}{\nu(C)} \ge \frac{c^2(1-\rho)\nu(C)^{s'-1}}{2^{s'}B^{2s'}} \ge \gamma.$$

Together with (15), these bounds imply $\Phi(K) \ge \Phi_*$ for conductance $\Phi(K) = \min_{0<\nu(C)\le 1/2} F_h(C)/\nu(C)$. The lazy reversible Cheeger theorem [13], applied to $\overline{K} = (I+K)/2$, gives

$$\mathrm{gap}(K)/2 \ge \tfrac{1}{2}(\Phi(K)/2)^2, \qquad \mathrm{gap}(K) \ge \Phi_*^2/4.$$

All transformations are valid also for constant h and $a = 1$, because $z \ge c > 0$. This proves the theorem. □

Corollary 10.2. *Fix $0 < \rho < 1$ and $0 < b < 1$. Every normalized positive potential taking at most two values and satisfying $\pi \ge b$ and $\mathbb{E}\pi^s \le 2$ obeys $\mathrm{gap}(K) \ge G_J(\rho, b/2) > 0$, without matching or any additional restriction on the level ratio or geometry.*

Proof. For a constant potential normalization gives $h = 1$. Otherwise let $h = M_0$ on a set of mass $\alpha \in (0,1)$ and $h = m_0 < M_0$ on its complement. The constant Fourier component gives $\mathbb{E}h \le 1$. Integrating the floor on the low set yields $m_0 \ge b(1-\alpha)$. Since $g \le M_0$, the floor there also gives $m_0 M_0 \ge b$, while $M_0\alpha \le \mathbb{E}h \le 1$ implies $m_0 \ge b\alpha$. Thus $h \ge m_0 \ge b/2$. Apply Theorem 10.1. □

This recovers the two-level conclusion, with possibly weaker constants than the separately preserved direct proof. Its geometric restriction has been replaced in the broader theorem by the explicit sufficient condition $h \ge a$; that condition is derived here only for the stated two-level class.

For any normalized family at fixed ρ with $\mathbb{E}\pi_L^s \le 2$, a collapsing gap necessarily implies $\min_x h_L(x) \to 0$. Otherwise a subsequence has a fixed positive minimum, contradicting Theorem 10.1. This is a necessary mechanism, not the construction of such a family. The original stationary floor and matching have not been shown to exclude it.

10.1 The stationary floor does not supply a potential minimum

The extra minimum in Theorem 10.1 cannot be derived from the stationary floor and cap alone, even for three-level potentials. The following construction does not impose exact matching.

Proposition 10.3. *Fix $0 < \rho < 1$, $0 < b < 1$ and any finite $s > 1$. There are energy-normalized positive potentials h_L such that $\pi_L = h_L P_\rho h_L \ge b$ and*

$$\|\pi_L - 1\|_s \longrightarrow 0, \qquad \|h_L - 1\|_{2s} \longrightarrow 0, \qquad \min h_L \longrightarrow 0.$$

For all sufficiently large L, the potential has exactly three values, $\mathbb{E}\pi_L^s \le 2$, and $(P_\rho h_L)(x_0) \to \infty$ at $x_0 = (1,\dots,1)$. These are actual Doob kernels. No exact stationary matching is imposed here; its critical-exponent spectral behavior is established below in Theorem 10.7.

Proof. Write $p = (1-\rho)/2$ and $q = (p+1/2)/2$, so $p < q < 1/2$. For Hamming distance D, set

$$A = \{y : 1 \le D(x_0, y) \le qL\}, \quad \alpha = \mu(A), \quad \delta = 2^{-L}, \quad r = P_\rho(x_0, A), \quad u = (1-p)^L.$$

The binomial distributions of D under μ and $P_\rho(x_0, \cdot)$, respectively, and Chebyshev's inequality give

$$\alpha \le \frac{1}{4L(1/2-q)^2} \to 0, \qquad 1 - r \le u + \frac{p(1-p)}{L(q-p)^2} \to 0.$$

For large L, A is nonempty and its complement contains points besides x_0. Also $P_\rho(y, x_0) \le u$ for every y, since $p < 1/2$. Put $M = \alpha^{-1/(4s)} > 1$ and $\beta = (1+b)/2$. There is a unique $0 < t < 1$ satisfying

$$t(ut + Mr + 1 - u - r) = \beta :$$

the left side is strictly increasing from zero, and at $t = 1$ it equals $1 + (M-1)r > 1 > \beta$. Moreover $t \le \beta/(Mr) \to 0$. Define $H = t$ at x_0, $H = M$ on A, and $H = 1$ elsewhere. With $P = P_\rho$, $Z = \mathbb{E}HPH$ and $h = H/\sqrt{Z}$, we have $H(x_0)(PH)(x_0) = \beta$ and $\mathbb{E}hPh = 1$ exactly. Positivity of P and $1 - 1_{\{x_0\}} \le H \le 1 + M1_A$ imply

$$1 - 2\delta + \delta u \le Z \le 1 + 2M\alpha + M^2\langle 1_A, P1_A\rangle \le 1 + 2M\alpha + M^2\alpha \longrightarrow 1.$$

Outside x_0, $H \ge 1$ and $PH \ge 1 - P(\cdot, x_0) \ge 1 - u$. Hence $\pi = HPH/Z \ge \min\{\beta, 1-u\}/Z \ge b$ eventually.

For $R = HPH$, the same upper bound on H gives

$$R \le 1 + M1_A + MP1_A + M^2 1_A P1_A.$$

Jensen's inequality and invariance of μ give $\|P1_A\|_s \le \alpha^{1/s}$; also $P1_A \le 1$. Consequently

$$\|(R-1)_+\|_s \le 2M\alpha^{1/s} + M^2\alpha^{1/s} = 2\alpha^{3/(4s)} + \alpha^{1/(2s)} \to 0.$$

The lower bound outside x_0 and $R(x_0) = \beta$ give $\|(1-R)_+\|_s \le u + (1-\beta)\delta^{1/s} \to 0$. Thus $\|R - 1\|_s \to 0$, and normalization yields $\|\pi - 1\|_s \to 0$ and $\mathbb{E}\pi^s \to 1$. Furthermore

$$\|H - 1\|_{2s}^{2s} = \alpha(M-1)^{2s} + \delta(1-t)^{2s} \le \alpha^{1/2} + \delta \to 0.$$

It follows that $\|h - 1\|_{2s} \to 0$ while $\min h = t/\sqrt{Z} \to 0$ and $(Ph)(x_0) = \beta/(t\sqrt{Z}) \to \infty$. Fill any remaining small dimensions with $h = 1$. The actual kernel $K(x,y) = P(x,y)h(y)/(Ph)(x)$ is stochastic, and $\mu(x)\pi(x)K(x,y) = \mu(x)P(x,y)h(x)h(y)$ proves reversibility. □

For $|f| \le 1$, the construction gives $|\mathbb{E}\pi_L f - \mathbb{E}f| \le \|\pi_L - 1\|_1 \to 0$, uniformly over such tests. This includes approximate marginal matching of every order, but no asymptotic error is set equal to zero. The original question requires exact labelled matching. Vanishing minima are a necessary condition for collapse supplied by Theorem 10.1, not a sufficient one; Proposition 10.3 supplies no gap-collapse example. The construction uses classical binomial estimates and elementary normalization. Its role here is to close this particular failed bridge; global priority and significance are not established.

10.2 Exact matching does not restore the potential minimum

The approximate matching in Proposition 10.3 can be corrected exactly while preserving its exceptional small potential value. We first establish the needed finite-dimensional stability estimate. Positive matrix scaling is classical; see the setup in Lee's lecture [1]. The arguments needed here are included in full.

Lemma 10.4. *Let P be a strictly positive symmetric stochastic matrix on n points with uniform probability μ. Every strictly positive mean-one density w has a unique positive solution of $hPh = w$. If P is also positive semidefinite, and $w_\tau = w_0 + \tau e \geq b_0 > 0$ for $0 \leq \tau \leq 1$ with $\mathbb{E}e = 0$, the corresponding solutions satisfy*

$$\|\log h_1 - \log h_0\|_\infty \leq \frac{\lceil \frac{1}{2}\log_2(n/b_0)\rceil + 2}{2b_0}\|e\|_\infty.$$

Proof. For $v \in \mathbb{R}^n$, consider

$$F_w(v) = \frac{1}{2}\sum_{x,y}\mu(x)P(x,y)e^{v_x+v_y} - \sum_x \mu(x)w(x)v_x.$$

The diagonal exponential terms make the Hessian positive definite; all other exponential terms are convex. Thus F_w is strictly convex. Writing $m = \max_x v_x$ and $d = \frac{1}{2}\min_x \mu(x)P(x,x) > 0$, we have $F_w(v) \geq de^{2m} - m$, because $\sum \mu w = 1$. Every sublevel set therefore has its maximum coordinate bounded above. If any coordinate tends to $-\infty$ while all coordinates are bounded above, its strictly positive linear weight forces $F_w \to \infty$. Thus sublevel sets are bounded, and a unique minimizer exists. Its stationary equation is $e^{v_x}(Pe^v)(x) = w(x)$, proving existence and uniqueness with $h = e^v$. The positive definite Hessian and the implicit function theorem give smooth dependence along the path w_τ; existence at every parameter and uniqueness join the local solutions.

For the actual kernel $K(x,y) = P(x,y)h(y)/(Ph)(x)$ with stationary $\nu = w\mu$, positive semidefiniteness of P implies that the spectrum of K lies in $[0,1]$. Let $\nu_* = \min_x \nu(x)$, $A = \nu_*^{-1/2}$, $m = \lceil \log_2 A \rceil$ and $T = (I-K)/2$. The absolute row sum of T is $1 - K(x,x) \leq 1$, whereas its $L^2(\nu)$ operator norm is at most $1/2$. Hence

$$\|T^j\|_{\infty\to\infty} \leq \min\{1, A2^{-j}\}.$$

Here the second estimate follows from $|f(x)| \leq \|f\|_{L^2(\nu)}/\sqrt{\nu_*}$. The convergent series and its norm bound give

$$(I+K)^{-1} = \frac{1}{2}\sum_{j\geq 0} T^j, \qquad \|(I+K)^{-1}\|_{\infty\to\infty} \leq \frac{1}{2}\left(m + \sum_{j=m}^{\infty} A2^{-j}\right) \leq \frac{m+2}{2}.$$

No Markov property of T or gap bound at the top of the spectrum is assumed. Differentiating the defining pointwise equation gives

$$\frac{d}{d\tau}\log w_\tau = (I + K_\tau)\frac{d}{d\tau}\log h_\tau.$$

Along the path $\nu_* \geq b_0/n$ and $\|d\log w_\tau/d\tau\|_\infty \leq \|e\|_\infty/b_0$. Integrating proves the displayed estimate. □

Proposition 10.5. *For every fixed $0 < \rho < 1$, $0 < b < 1$ and $C > 0$, there is a sequence satisfying all of (1), with energy normalization $\mathbb{E}h_L P_\rho h_L = 1$, such that, for $s = (1+\rho)/(1-\rho)$,*

$$\|\pi_L - 1\|_s \to 0, \qquad \|h_L - 1\|_{2s} \to 0, \qquad h_L(x_0) \to 0, \qquad (P_\rho h_L)(x_0) \to \infty.$$

In particular $\min h_L \to 0$. Theorem 10.7 classifies the actual gap of this family. Preservation of three potential values is not asserted.

Proof. Apply Proposition 10.3 with the larger stationary floor $b_1 = (1+b)/2$, using its notation $\pi, h, R, Z, \alpha, M, u, \delta$ and $\beta = (1+b_1)/2$. Its rare-set probability is exponentially small. Indeed, exponential Markov applied to $D \sim \mathrm{Bin}(L, 1/2)$ yields

$$\alpha \le \Pr(D \le qL) \le \left[e^{\lambda q} \frac{1 + e^{-\lambda}}{2} \right]^L.$$

The bracket is one at $\lambda = 0$ and its logarithmic derivative there is $q - 1/2 < 0$; choose a fixed small $\lambda > 0$ making it $a \in (0,1)$. The preceding proof also gives

$$\|R - 1\|_1 \le 2M\alpha + M^2\alpha + u + (1-\beta)\delta = 2\alpha^{1-1/(4s)} + \alpha^{1-1/(2s)} + u + (1-\beta)\delta.$$

Every term decays exponentially. Since $|Z - 1| \le \|R - 1\|_1$ and $Z \to 1$, there are fixed $A_1, c > 0$ with $\|\pi - 1\|_1 \le A_1 e^{-cL}$ for all sufficiently large L.

Put $k = k_L$ as in (1), and subtract the nonconstant low-degree Fourier projection:

$$e_L = -\sum_{1 \le |S| \le k} (\mathbb{E}\pi\chi_S)\chi_S, \qquad w_L = \pi + e_L.$$

Each coefficient has magnitude at most $\|\pi - 1\|_1$. The number of terms is at most kL^k eventually, whose logarithm is $O(\sqrt{L}\log L) = o(L)$. Consequently

$$\epsilon_L := \|e_L\|_\infty \le kL^k A_1 e^{-cL}, \qquad L\epsilon_L \to 0.$$

Orthogonality gives $\mathbb{E}w_L = 1$ and $\mathbb{E}w_L\chi_S = 0$ for every $1 \le |S| \le k$ exactly. Along $\pi + \tau e_L$, the floor is at least $b_1 - \epsilon_L \ge b$ eventually. Also $\|w_L - 1\|_s \le \|\pi - 1\|_s + \epsilon_L \to 0$, so its moment is eventually at most two.

The noise matrix is strictly positive and positive definite, with Walsh eigenvalues $\rho^{|S|} > 0$. Lemma 10.4 provides a positive $\widetilde{h}_L$ satisfying $\widetilde{h}_L P_\rho \widetilde{h}_L = w_L$, with

$$\eta_L := \|\log(\widetilde{h}_L / h_L)\|_\infty \le \frac{\lceil (L + \log_2(1/b))/2 \rceil + 2}{2b} \epsilon_L \to 0.$$

Thus $e^{-\eta_L} h_L \le \widetilde{h}_L \le e^{\eta_L} h_L$. The vanishing value at x_0 and the divergent value of $P_\rho h_L$ there persist. Moreover

$$\|\widetilde{h}_L - h_L\|_{2s} \le (e^{\eta_L} - 1)\|h_L\|_{2s} \to 0,$$

so the potential norm convergence persists as well. Its energy is $\mathbb{E}w_L = 1$, and its actual stationary density is exactly w_L. Use $\widetilde{h}_L = 1$ in the finitely many earlier dimensions and rename the resulting sequence h_L. □

This refutes the proposed implication from the full original stationary hypotheses to a uniform normalized potential minimum. It does not refute the spectral-gap question: the necessary condition for collapse from Theorem 10.1 has been attained, not shown sufficient. Further full-gap arguments must control the actual chain without deducing such a minimum. The finite reconstruction is an analytic existence proof, not an executed matrix-scaling certificate; no claim of global priority for its stability specialization is made.

10.3 The vanishing-minimum family retains the noise gap

The exact-matching construction above can be classified spectrally. Its small potential value does not produce a collapsing gap. The proof uses an estimate for the actual ratio $h/(Ph)$, rather than transferring stationary moment convergence directly to the spectrum.

Lemma 10.6. *Put $s = (1+\rho)/(1-\rho)$ and $Q = P_{\sqrt{\rho}}$ on the Boolean cube. For every real function v and every positive potential H, put $U = H/(PH)$. Then*

$$\|QM_vQ\|_{2\to 2} \le \|v\|_s, \qquad |\operatorname{gap}(K_H) - (1-\rho)| \le \|U-1\|_s.$$

More generally, every ordered eigenvalue of K_H differs from the corresponding ordered eigenvalue of P by at most $\|U-1\|_s$.

Proof. The conjugate exponent is $s' = (1+\rho)/(2\rho)$, so $2s' = 1 + 1/\rho$. Classical Boolean hypercontractivity gives $\|Qf\|_{2s'} \le \|f\|_2$: its noise condition is the equality $\sqrt{\rho} = \sqrt{(2-1)/(2s'-1)}$; see Chapter 10.1 of [6]. Holder's inequality with exponents $s, 2s', 2s'$ therefore gives

$$|\langle f, QM_vQg\rangle| \le \|v\|_s \|Qf\|_{2s'} \|Qg\|_{2s'} \le \|v\|_s \|f\|_2 \|g\|_2.$$

For the spectral assertion, the actual reversible conjugate of K_H is $S = M_{\sqrt{U}} P M_{\sqrt{U}}$. Set $X = M_{\sqrt{U}} Q$. Then $S = XX^*$ and $T = X^*X = QM_UQ$. Both factors of X are invertible: $U > 0$ and the Walsh eigenvalues of Q are positive. Thus $X^{-1}SX = T$ and T has the full spectrum of the actual chain. The first estimate yields

$$\|T - P\|_{2\to 2} = \|QM_{U-1}Q\|_{2\to 2} \le \|U-1\|_s.$$

The min–max principle compares every ordered eigenvalue of the two self-adjoint matrices. For $L \ge 1$, the second eigenvalue of P is ρ, proving the gap assertion, including its multiplicity. □

Theorem 10.7. *Fix $0 < \rho < 1$, $0 < b < 1$ and $C > 0$. The exactly matched family $\widetilde{h}_L$ constructed in the proof of Proposition 10.5 satisfies*

$$|\operatorname{gap}(K_{\widetilde{h}_L}) - (1-\rho)| \le C_1 e^{-c_1 L}$$

for all sufficiently large L, with constants $C_1 < \infty$ and $c_1 > 0$ depending only on the fixed parameters and construction choices. It satisfies all of (1)*, although its normalized minimum tends to zero and $(P\widetilde{h}_L)(x_0)$ tends to infinity.*

Proof. Use the precursor H from Proposition 10.3, with $s = (1+\rho)/(1-\rho)$, and retain its α, δ, u, M, t. Since $H \ge 1 - 1_{\{x_0\}}$ and $P(y, x_0) \le u$ for every y,

$$G := PH \ge 1 - u > 0$$

at every vertex, including x_0. No positive lower bound on t is needed. The disjoint exceptional supports and L^s contraction of P give, with $d_L = \alpha^{3/(4s)} + \delta^{1/s}$,

$$\|H-1\|_s \le \alpha^{1/s}(M-1) + \delta^{1/s}(1-t) \le d_L, \qquad \|H-G\|_s \le 2d_L.$$

Consequently the scale-invariant ratio satisfies

$$\|U-1\|_s = \|(H-G)/G\|_s \le \frac{2d_L}{1-u} =: \Delta_L.$$

Energy normalization does not change U. For the exact correction, Proposition 10.5 proves

$$\eta_L = \|\log(\widetilde{h}_L/h_L)\|_\infty \le \frac{\lceil (L+\log_2(1/b))/2\rceil + 2}{2b}\epsilon_L, \qquad \epsilon_L \le k_L L^{k_L} A_1 e^{-cL}.$$

Positivity of P yields the corresponding multiplicative bounds for the denominators, hence

$$e^{-2\eta_L} U \le \widetilde{U} := \frac{\widetilde{h}_L}{P\widetilde{h}_L} \le e^{2\eta_L} U.$$

Since $\|U\|_s \le 1 + \Delta_L$, it follows that

$$\|\widetilde{U} - 1\|_s \le \Delta_L + (e^{2\eta_L} - 1)(1 + \Delta_L) = e^{2\eta_L}\Delta_L + e^{2\eta_L} - 1. \tag{16}$$

The preceding construction proves $\alpha \le a^L$ for fixed $a < 1$; also $u = ((1+\rho)/2)^L$ and $\delta = 2^{-L}$. Finally $\log(k_L L^{k_L}) = o(L)$, so $\epsilon_L \le \exp(-cL + o(L))$. This is an upper bound, with no lower bound on ϵ_L asserted or needed. The factor multiplying ϵ_L in the estimate for η_L is $O(L)$. Thus the right side of (16) is bounded by $C_1 e^{-c_1 L}$ for a possibly smaller positive rate c_1. Lemma 10.6 now gives the asserted actual-gap estimate. The finitely many dimensions filled with $h = 1$ do not affect the limit. □

The classical sandwich estimate supplies no general bound on $\|h/(Ph) - 1\|_s$ from (1). Here the all-vertex denominator estimate is proved directly for the particular precursor, and the nonlinear reconstruction has proved multiplicative stability. This closes the constructed vanishing-minimum family as a candidate collapse mechanism while preserving its counterexample to deriving a potential minimum. The unrestricted stationary-gap question remains open.

11 Vanishing denominators without spectral collapse

The preceding family has small potential values but a uniformly positive noise denominator. We next rule out deriving a positive denominator floor from the original hypotheses. The same construction shows why excluding individual points with a large ratio is too strong a proposed remedy.

Theorem 11.1. *Fix $0 < \rho, b < 1$ and $C > 0$. There are energy-normalized positive potentials $\widetilde{h}_L$ satisfying all of* (1) *such that, at the all-plus vertex x_0, for a constant $\varepsilon > 0$,*

$$\widetilde{h}_L(x_0) \sim e^{\varepsilon L}, \qquad (P\widetilde{h}_L)(x_0) \sim e^{-\varepsilon L}, \qquad \widetilde{\pi}_L(x_0) \longrightarrow 1.$$

Thus $\widetilde{U}_L(x_0) \sim e^{2\varepsilon L}$, where $\widetilde{U}_L = \widetilde{h}_L/(P\widetilde{h}_L)$. Nevertheless there are constants $A < \infty$, $a > 0$ such that, for all sufficiently large L,

$$\|\widetilde{\pi}_L - 1\|_s + \left\|\widetilde{U}_L - 1\right\|_s \le Ae^{-aL}, \qquad \max_{1\le j\le 2^L} |\lambda_j(K_{\widetilde{h}_L}) - \lambda_j(P)| \le Ae^{-aL}.$$

Eigenvalues are ordered with multiplicities. In particular the full gap tends to $1 - \rho$. Constants and the eventual threshold may depend on the fixed parameters; no uniformity as ρ approaches its endpoints is claimed.

Proof. We use the classical Bernoulli concentration bound, including unequal parameters: if S sums L independent Bernoulli variables, then

$$\Pr(S - \mathbb{E}S \geq tL),\ \Pr(S - \mathbb{E}S \leq -tL) \leq e^{-2t^2 L} \qquad (t > 0).$$

Here is the needed proof. The log moment-generating function of a centered Bernoulli variable has value and first derivative zero at zero; its second derivative is the variance of a tilted Bernoulli law, at most $1/4$. Integrating twice bounds it by $\lambda^2/8$ for every real λ. Independence and exponential Markov at $\lambda = \pm 4t$ give the two bounds. No equal-parameter assumption is used.

Put

$$p = \frac{1-\rho}{2}, \quad q = p + \rho p = \frac{1-\rho^2}{2}, \quad a_0 = \frac{1+b}{2}, \quad \delta = \min\{\rho p/16, \rho^2/16\}, \quad r_0 = 2\delta/\rho.$$

Let $d(x)$ be the Hamming distance from x_0 and define

$$\begin{aligned} B_L &= \{x : |d(x)/L - p| \leq \delta\}, \\ A_L &= \{x : |d(x)/L - q| \leq 2\delta\}, \\ D_L &= \{x : (P1_{B_L})(x) > 1 - a_0\}. \end{aligned}$$

In this proof D_L is a set. We have $r_0 \leq p/8$, $q - p \geq 16\delta$ and $1/2 - q \geq 8\delta$. Thus A_L and B_L are disjoint. Conditional on x, the next noise step has independent coordinate flips and

$$\mathbb{E}_x d(Y) = pL + \rho d(x).$$

For $x \in B_L$ this mean lies between $(q - \rho\delta)L$ and $(q + \rho\delta)L$, so

$$P(x, A_L) \geq 1 - 2e^{-2\delta^2 L}, \qquad P(x_0, B_L) \geq 1 - 2e^{-2\delta^2 L}.$$

If $d(x)/L \geq r_0$, the mean is at least $(p+2\delta)L$ and the upper boundary of B_L is $(p+\delta)L$. Therefore $P(x, B_L) \leq e^{-2\delta^2 L}$. For all sufficiently large L,

$$x_0 \in D_L \subset \{x : d(x)/L < r_0\}.$$

The last ball is disjoint from both shells, since $r_0 < p - \delta$.

Choose the strictly positive constant

$$c = \min \left\{ \begin{matrix} 2\delta^2,\ 2(p - r_0)^2,\ 2(q - 2\delta - p)^2, \\ 2(1/2 - q - 2\delta)^2,\ 2(1/2 - p - \delta)^2,\ 2(1/2 - r_0)^2 \end{matrix} \right\}.$$

Under uniform measure $d(x)$ is binomial with mean $L/2$, whereas under $P(x_0, \cdot)$ it is binomial with mean pL. The concentration bounds and the containment of D_L consequently give

$$\begin{aligned} &\mu(A_L), \mu(B_L), \mu(D_L) \leq e^{-cL}, \\ &P(x_0, A_L), P(x_0, D_L) \leq e^{-cL}, \\ &\qquad P(x_0, B_L^c) \leq 2e^{-cL}, \qquad \inf_{x \in B_L} P(x, A_L) \geq 1 - 2e^{-cL}. \end{aligned} \tag{17}$$

These estimates include shell boundaries and the return to x_0.

Set $\varepsilon = c/(16s)$, $M = e^{\varepsilon L}$, and $E_L = A_L \cup D_L$. Define the actual potential and its product by

$$H = \begin{cases} M & \text{on } E_L, \\ M^{-1} & \text{on } B_L, \\ 1 & \text{elsewhere,} \end{cases} \qquad G = PH, \quad R = HG, \quad Z = \mathbb{E}R, \quad h = H/\sqrt{Z}.$$

Always $M^{-1} \le H, G \le M$. On E_L, $R \ge 1$. On B_L, $G \ge MP1_{A_L}$, hence $R \ge 1 - 2e^{-cL} \ge a_0$ eventually. Outside $E_L \cup B_L$, one has $H = 1$ and $x \notin D_L$, whence $G(x) \ge 1 - P(x, B_L) \ge a_0$. Thus $R \ge a_0$ at every vertex.

At $x_0 \in D_L$, splitting the noise expectation over the three regions gives

$$G(x_0) = M^{-1}P(x_0, B_L) + P(x_0, (B_L \cup E_L)^c) + MP(x_0, E_L).$$

Consequently

$$P(x_0, B_L) \le MG(x_0) \le 1 + 2Me^{-cL} + 2M^2e^{-cL} \longrightarrow 1,$$

because $2\varepsilon < c$. Hence $G(x_0) \sim M^{-1}$, $R(x_0) \to 1$, and $H(x_0)/G(x_0) \sim M^2$.

For finite $r \ge 1$ write $d_{r,L} = \|H - 1\|_r$. The support bounds in (17) imply

$$d_{r,L} \le M(2e^{-cL})^{1/r} + e^{-cL/r}, \qquad \|G - 1\|_r \le d_{r,L}.$$

The second inequality is noise contraction. Hölder therefore gives

$$\|R - 1\|_s \le d_{2s,L}(1 + d_{2s,L}) + d_{2s,L}.$$

Since $\varepsilon < c/(2s)$, this error decays exponentially, as does $|Z-1|$. Thus $\pi = R/Z$ has exponentially small $\|\pi - 1\|_s$, $\mathbb{E}\pi^s \to 1$, and eventually $\pi \ge b_1 := (a_0 + b)/2 > b$. Normalization preserves $h(x_0) \sim M$, $(Ph)(x_0) \sim M^{-1}$ and $\pi(x_0) \to 1$.

Crucially, for $U = H/G$, use the proved dimension-dependent bound $G \ge M^{-1}$, not a fixed denominator floor. It gives

$$\begin{aligned} \|U - 1\|_s &\le M \|H - G\|_s \le 2M d_{s,L} \\ &\le 2\, 2^{1/s} e^{-(c/s - 2\varepsilon)L} + 2e^{-(c/s - \varepsilon)L}. \end{aligned} \tag{18}$$

Both rates are positive. Lemma 10.6 controls the entire actual spectrum by this ratio error.

It remains to impose exact matching. Subtract the finite Walsh sum

$$e_L = -\sum_{1 \le |S| \le k_L} (\mathbb{E}\pi\chi_S)\chi_S, \qquad \widetilde{\pi} = \pi + e_L.$$

Each coefficient is at most $\|\pi - 1\|_1$, and for all large L there are at most $k_L L^{k_L} = \exp(o(L))$ terms. Therefore $\xi_L = \|e_L\|_\infty$ is exponentially small after decreasing the positive rate. Orthogonality gives exact mean one and all requested zero coefficients. The path $\pi + \tau e_L$ has floor b eventually, and $\|\widetilde{\pi} - 1\|_s$ remains exponentially small, giving the eventual cap two.

Lemma 10.4 reconstructs the unique positive $\widetilde{h}$ with $\widetilde{h}P\widetilde{h} = \widetilde{\pi}$ and gives

$$\eta_L := \left\|\log(\widetilde{h}/h)\right\|_\infty \le \frac{\lceil (L + \log_2(1/b))/2 \rceil + 2}{2b} \xi_L.$$

This is exponentially small. Positivity of P transfers the same multiplicative bounds to $P\widetilde{h}$ and Ph, preserving the pointwise asymptotics, including the exponentially small denominator. Also $e^{-2\eta_L}U \le \widetilde{U} \le e^{2\eta_L}U$, so

$$\left\|\widetilde{U}-1\right\|_s \le e^{2\eta_L}\left\|U-1\right\|_s + e^{2\eta_L} - 1.$$

Combining this with (18) and Lemma 10.6 proves the claimed norm and full-spectrum bounds. Polynomial factors and the subexponential Walsh count are absorbed by a smaller positive exponential rate. Fill all finitely many early dimensions with $\widetilde{h} = 1$, preserving every original constraint. □

Theorem 11.1 refutes a dimension-uniform positive Ph floor and a uniform pointwise ceiling for h/Ph, even under exact matching. It does not refute an aggregate killed-return or cut-flow estimate. Small denominators and large ratios at individual vertices coexist here with a full spectrum approaching Boolean noise. The original gap question remains open; no original hypothesis was removed.

12 Regularization and exact stationary matching

The noise formula of Definition 2.46 and Proposition 2.47 in [6] makes smoothing a natural candidate operation. Smoothing the potential, however, does not commute with its nonlinear stationary-density construction. The following statements separate a preserved inequality from a failed exact equality.

Proposition 12.1. *Let $h > 0$ satisfy $\mathbb{E}hPh = 1$ and $hPh \ge b > 0$. For $0 \le \tau \le 1$ put*

$$H_\tau = P_\tau h, \qquad Z_\tau = \mathbb{E}H_\tau P H_\tau, \qquad h_\tau = H_\tau/\sqrt{Z_\tau}, \qquad \pi_\tau = h_\tau P h_\tau.$$

Then $\pi_\tau \ge b$ for every τ. No marginal-matching or moment assumption is needed for this assertion.

Proof. The noise kernels commute. Cauchy–Schwarz in each probability row of P_τ therefore gives

$$H_\tau P H_\tau = (P_\tau h)P_\tau(Ph) \ge \left(P_\tau\sqrt{hPh}\right)^2 \ge b.$$

By Walsh Parseval,

$$0 < Z_\tau = \sum_S \rho^{|S|}\tau^{2|S|}\widehat{h}(S)^2 \le \sum_S \rho^{|S|}\widehat{h}(S)^2 = 1.$$

The constant-character multiplier is one, including at $\tau = 0$. Dividing the pointwise inequality by Z_τ proves the result. □

Theorem 12.2. *Fix $0 < \rho < 1$, $0 < b < 1$, $L \ge 3$ and $1 \le k \le L-2$, and put $s = (1+\rho)/(1-\rho)$. There is a positive potential h with $\mathbb{E}hPh = 1$ such that its stationary density has exact uniform labelled marginals through order k, while*

$$\pi_\tau \ge b, \qquad \mathbb{E}\pi_\tau^s \le \sqrt{2} < 2 \quad (0 \le \tau \le 1).$$

For one specified coordinate j, the same potential satisfies $\mathbb{E}\pi_\tau\chi_{\{j\}} \ne 0$ for every $0 < \tau < 1$. Thus every nontrivial smoothing loses even a first marginal while the stationary floor and critical cap survive for this constructed family.

Proof. Choose $A \subseteq [L]$ with $m = |A| = k+1$ and $j \notin A$, and put $B = A \cup \{j\}$, $\alpha = \rho^m$ and $\beta = \rho^{m+1}$. Fix the dimension-independent number

$$\epsilon = \min\left\{\frac{1}{8}, \frac{1-b}{4}, \frac{\log 2}{64s}\right\} > 0$$

and define $w = 1+\epsilon(\chi_A+\chi_B)$. This density has mean one and $w \geq 1-2\epsilon \geq \max\{3/4, (1+b)/2\} > b$. Its only nonconstant Walsh coefficients are at A and B, whose degrees exceed k, so it has the required labelled marginals. We next realize it as the stationary density of an actual positive potential rather than treating the density and potential as independent.

Under uniform input, $y = \chi_A$ and $z = x_j$ are independent uniform signs. Since their coordinate sets are disjoint, Boolean ρ-noise induces on (y, z) the exact product kernel $Q = P_\alpha \otimes P_\rho$. In particular

$$P[f(\chi_A, x_j)] = (Qf)(\chi_A, x_j)$$

for every function f on these four points. The eigenvalues of Q are $1, \alpha, \rho, \beta$; all entries are strictly positive. Apply Lemma 10.4 on this four-point space to the path $w_t = 1 + t\epsilon(y + yz)$, $0 \leq t \leq 1$. It is bounded below by $1/2$. The unique positive solutions $g_t Q g_t = w_t$ start from $g_0 = 1$ and obey

$$\|\log g_1\|_\infty \leq \frac{\lceil \frac{1}{2}\log_2(4/(1/2))\rceil + 2}{2(1/2)}\|\epsilon(y+yz)\|_\infty = 8\epsilon =: \eta \leq \frac{\log 2}{8s}.$$

The inverse estimate uses four points, not 2^L, and is uniform as α becomes small. The lift $h = g_1(\chi_A, x_j)$ consequently satisfies $hPh = w$, $\mathbb{E}hPh = 1$, and $e^{-\eta} \leq h \leq e^{\eta}$.

Every $H_\tau = P_\tau h$ and PH_τ lies in the same interval, so $e^{-2\eta} \leq Z_\tau \leq e^{2\eta}$ and $\pi_\tau \leq e^{4\eta}$. Thus $\mathbb{E}\pi_\tau^s \leq e^{4s\eta} \leq \sqrt{2}$ for all τ. Proposition 12.1 supplies the floor for all τ. At $\tau = 1$ these are the original stationary density and its cap.

Write the unique four-character expansion

$$h = c + d\chi_{\{j\}} + e\chi_A + f\chi_B, \qquad c = \mathbb{E}h > 0.$$

Equating the three indicated coefficients in $hPh = w$ gives

$$(1+\rho)cd + (\alpha+\beta)ef = 0, \tag{19}$$
$$(1+\alpha)ce + (\rho+\beta)df = \epsilon, \tag{20}$$
$$(1+\beta)cf + (\rho+\alpha)de = \epsilon. \tag{21}$$

If $e = 0$, equation (19) forces $d = 0$, contradicting (20). If $f = 0$, the same equation again forces $d = 0$, contradicting (21). Hence $ef \neq 0$. Now

$$H_\tau = c + \tau d\chi_{\{j\}} + \tau^m e\chi_A + \tau^{m+1} f\chi_B.$$

Using (19) in the $\chi_{\{j\}}$ coefficient of $H_\tau P H_\tau$ proves the exact identity

$$\mathbb{E}\pi_\tau\chi_{\{j\}} = \frac{(\alpha+\beta)ef}{Z_\tau}(\tau^{2m+1} - \tau). \tag{22}$$

For $0 < \tau < 1$ this is nonzero, since $m \geq 2$, $ef \neq 0$ and $\alpha+\beta, Z_\tau > 0$. No perturbative remainder or numerical nonzero test is used. At $\tau = 0$ the normalized density is uniform, and at $\tau = 1$ it is the original exactly matched density. □

For fixed $C > 0$, the required $k_L = \min\{L, \lfloor C\sqrt{L} \rfloor\}$ eventually belongs to $[1, L-2]$. Apply the theorem in those dimensions and use $h = 1$ in the finitely many remaining dimensions. This gives the claimed obstruction within the original growing matching regime. No uniform lower bound on the magnitude of (22) is asserted: α tends to zero as k_L grows. Exact matching is the equality that fails. The constructed potentials have a fixed positive minimum, so they are not a collapsing-gap family.

The distinction between potential and density regularization is essential. For $w_\tau = P_\tau \pi$, stochasticity preserves mean one and the floor, Jensen's inequality and invariance of μ give $\mathbb{E}w_\tau^s \le \mathbb{E}\pi^s$, and the Walsh formula preserves every zero low-degree coefficient. Lemma 10.4 then supplies the unique actual positive potential u_τ with $u_\tau P u_\tau = w_\tau$. These elementary invariances do not compare the gaps of K_h and K_{u_τ}. A dimension-uniform spectral comparison under the original hypotheses remains unproved; it cannot be replaced by assuming a positive minimum for the reconstructed potential or a lower bound on Ph.

13 Full spectrum after stationary-density regularization

We now control the reconstructed chain on the regularized side of the comparison left open in Section 12. The result is uniform over every admissible input density; no code algebra or lower bound for its potential is imposed. It does not compare the reconstructed chain back to the original chain.

Fix $0 < \rho, b < 1$, and put $s = (1+\rho)/(1-\rho)$ and $B = 2^{1/s}$. The constants in this section are fixed before the dimension, density, matching order and smoothing parameter. Define

$$q_0 = 2\rho^{-2} + \rho^{-3} - 1, \qquad Q_* = 8s,$$
$$\sigma = \max\left\{4s, \frac{2Q_*^2}{(1-\rho^4)(q_0-1)}\right\}, \qquad F_\sigma(q) = 1 + \frac{\sigma q/(\sigma+q) - 1}{\rho^4}.$$

Starting at q_0, set $q_{n+1} = F_\sigma(q_n)$ only until its first crossing $q_{n_*} \ge Q_*$. Take $n_* = 0$ if $q_0 \ge Q_*$. The proof below establishes that this crossing is finite. Set

$$A_* = b^{-1}(B/b)^{n_*}, \qquad C_* = b^{-1}BA_*^2,$$
$$\theta_{\mathrm{in}} = \min\{1, \sqrt{s-1}\}, \qquad \theta_{\mathrm{out}} = \min\{1, (s-1)^{-1/2}\},$$
$$\tau_* = \tfrac{1}{2}\min\left\{\theta_{\mathrm{in}}\theta_{\mathrm{out}}, \sqrt{\frac{s-1}{\sigma-1}}\right\}, \qquad \gamma = \frac{1}{2s-1}.$$

In particular, $\tau_* > 0$ and $\tau_* \le 1/2$. For an integer $k \ge 0$, write

$$e_k = B\rho^{(k+1)/2}, \qquad \mathcal{E}_k = (2e_k)^\gamma (C_* + B)^{1-\gamma} + (B+1)2^{-(k+1)} + e_k^2. \tag{23}$$

Theorem 13.1 (Uniform full spectrum on the regularized side)**.** *Let $L \ge 1$, $0 \le k \le L$, and let π be a positive mean-one density on the cube with*

$$\pi \ge b, \qquad \mathbb{E}\pi^s \le 2, \qquad \mathbb{E}\pi\chi_S = 0 \quad (1 \le |S| \le k).$$

For every $0 \le \tau \le \tau_*$, put $w = P_\tau \pi$. There is a unique positive u with $uPu = w$. The actual Doob kernel $K_u(x,y) = P(x,y)u(y)/(Pu)(x)$ has stationary density w, which satisfies the

same floor, critical cap and exact matching conditions. Its eigenvalues, ordered decreasingly with multiplicities, satisfy

$$|\lambda_j(K_u) - \lambda_j(P)| \le \mathcal{E}_k \quad (1 \le j \le 2^L), \qquad |\operatorname{gap}(K_u) - (1-\rho)| \le \mathcal{E}_k.$$

The error decays exponentially in k, uniformly in L, π, τ for fixed ρ, b. At $k = k_L$ in (1) it decays exponentially in $\sqrt{L}$ for each fixed $C > 0$.

Proof. Reconstruction and a derived stronger moment. Stochasticity, Jensen's inequality and the Walsh multipliers give $\mathbb{E}w = 1$, $w \ge b$, $\mathbb{E}w^s \le 2$ and the same zero coefficients through degree k. Lemma 10.4 supplies the unique positive solution of $uPu = w$, with $\mathbb{E}uPu = 1$. Its dimension-dependent stability estimate is not used here. Ordinary hypercontractivity [2, 6] and $\tau^2 \le (s-1)/(\sigma-1)$ give

$$\|w\|_\sigma \le \|\pi\|_s \le B. \tag{24}$$

This stronger norm bound is derived, not an additional premise. It does not assert $\mathbb{E}w^\sigma \le 2$.

A finite inverse-moment crossing. Put $g = Pu$ and $Q = P_{\sqrt{\rho}}$. Then $\|Qu\|_2 = 1$ and $\|g\|_R \le 1$ for $R = 1 + \rho^{-1}$, by ordinary hypercontractivity applied to $g = Q(Qu)$. The floor $ug \ge b$ gives $u^{-1} \le g/b$, or $\|u\|_{-R} \ge b$ in generalized-mean notation. Reverse hypercontractivity, including its negative-exponent range [2], applies since $\rho^2 = (1+R)/(1+q_0)$. Thus $\|g^{-1}\|_{q_0} \le b^{-1}$. All functions are strictly positive on a finite space. The floor and rowwise Jensen also give

$$g \ge bP(1/g) \ge b/(Pg), \qquad g^{-1} \le Pg/b.$$

If $\|g^{-1}\|_q \le A_q$ and $q > \sigma/(\sigma-1)$, Hölder with $p = \sigma q/(\sigma+q) > 1$ gives $\|u\|_p = \|w/g\|_p \le BA_q$. Since $Pg = P_{\rho^2}u$, ordinary hypercontractivity then yields

$$\left\|g^{-1}\right\|_{F_\sigma(q)} \le (B/b)A_q. \tag{25}$$

This iteration uses the newly derived exponent σ, not the original critical exponent's fixed-point formula.

For $q_0 \le q < Q_*$, direct algebra and the choice of σ give

$$\begin{aligned} F_\sigma(q) - q &= \frac{(1-\rho^4)(q-1) - q^2/(\sigma+q)}{\rho^4} \\ &\ge \frac{(1-\rho^4)(q_0-1)}{2\rho^4} =: d > 0. \end{aligned}$$

Here $q_0 > 2$, while $\sigma \ge 4s > 4$, so every input before the first crossing satisfies $q > \sigma/(\sigma-1)$. If $q_0 < Q_*$, absence of a crossing would give $q_n \ge q_0 + nd$ for all n, a contradiction. In fact $n_* \le \lceil (Q_* - q_0)/d \rceil$. If $q_0 \ge Q_*$ no step is needed. At the first crossing the last valid step has already proved

$$\left\|g^{-1}\right\|_{q_{n_*}} \le A_*, \qquad q_{n_*} \ge 8s.$$

No assertion about further iterates is required. For $U = u/g = w/g^2$, $N = \mathbb{E}U$ and $r = U/N$, Hölder gives

$$\|U\|_{2s} \le \|w\|_{4s} \left\|g^{-1}\right\|_{8s}^2 \le BA_*^2, \qquad \|r\|_{2s} \le C_*, \qquad \|w\|_{2s} \le B, \tag{26}$$

where $N \ge b$ is justified independently in the next paragraph.

Actual normalization and matching rigidity. Let $a = \mathbb{E}u = \mathbb{E}g$ and $v = 1 - a^2 = \|Qu - a\|_2^2$. The floor gives $a \geq b\mathbb{E}(1/g) \geq b/a$, so $a^2 \geq b$. Cauchy–Schwarz under $w\,d\mu$ gives $N = \mathbb{E}_w(g^{-2}) \geq (\mathbb{E}_w(g^{-1}))^2 = a^2$. For the other bound, the reversible conjugate $S = M_{\sqrt{U}}PM_{\sqrt{U}}$ has Perron eigenvalue one and positive spectrum. With $X = M_{\sqrt{U}}Q$, the matrices $S = XX^*$ and $T = X^*X = QM_UQ$ are similar because X is invertible. Thus $\|T\|_{2\to 2} = 1$ and testing at the unit constant gives $N \leq 1$. This proves $b \leq a^2 \leq N \leq 1$ without any perturbation estimate.

For completeness, the matching argument of Lemma 3.2 can be used directly with exponent σ. Since $\sigma' < 2$, the degree-$> k$ part of $g = Q(Qu)$ has $L^{\sigma'}$ norm at most $\rho^{(k+1)/2}\sqrt{v}$. Exact matching gives $\mathbb{E}w\Pi_{\leq k}g = a$, hence $\mathbb{E}wg \leq a + e_k\sqrt{v}$. Cauchy–Schwarz under $u\,d\mu$ gives $1 = (\mathbb{E}ug)^2 \leq a\mathbb{E}ug^2 = a\mathbb{E}wg$. Consequently $v \leq ae_k\sqrt{v}$, including $v = 0$, and $a^2 \geq (1 + e_k^2)^{-1}$. The actual identity $\mathbb{E}\sqrt{rw} = a/\sqrt{N} \geq a$ yields

$$\|\sqrt{r} - \sqrt{w}\|_2^2 \leq 2(1 - a) \leq e_k^2, \qquad \|r - w\|_1 \leq 2e_k.$$

Interpolation with (26) gives

$$\|r - w\|_s \leq (2e_k)^\gamma (C_* + B)^{1-\gamma}, \qquad \gamma = \frac{1}{2s - 1}. \tag{27}$$

Closeness to the constant and the full spectrum. Factor $P_\tau = P_{\theta_{\text{out}}}P_\beta P_{\theta_{\text{in}}}$, where $\beta = \tau/(\theta_{\text{in}}\theta_{\text{out}}) \leq 1/2$. The input map contracts L^s to L^2: use hypercontractivity if $s < 2$, and the identity with norm monotonicity if $s \geq 2$. The output contracts L^2 to L^s: use hypercontractivity if $s > 2$, and the identity with norm monotonicity if $s \leq 2$. The Fourier support of $\pi - 1$ lies above degree k, so the middle map contracts its L^2 norm by β^{k+1}. Therefore

$$\|w - 1\|_s \leq (B + 1)2^{-(k+1)}.$$

This also holds at $\tau = 0$, where the mean-zero part vanishes. Using $U - 1 = N(r - w) + N(w - 1) + (N - 1)$, $N \leq 1$ and $1 - N \leq e_k^2/(1 + e_k^2) \leq e_k^2$ gives $\|U - 1\|_s \leq \mathcal{E}_k$. The critical sandwich estimate of Lemma 10.6 gives

$$\|T - P\|_{2\to 2} = \|QM_{U-1}Q\|_{2\to 2} \leq \|U - 1\|_s \leq \mathcal{E}_k.$$

Indeed, $2s' = 1 + \rho^{-1}$ and $Q : L^2 \to L^{2s'}$ is a contraction, so Hölder bounds $|\mathbb{E}(U - 1)(Qf)(Q\ell)|$ by $\|U - 1\|_s$ for arbitrary unit L^2 functions. This controls the full operator. The self-adjoint min–max principle compares every ordered eigenvalue of T and P; the exact similarity above identifies the spectrum of T with that of K_u. No uniform condition number for X is used. Since $\lambda_2(P) = \rho$ for $L \geq 1$, the gap bound follows. All three terms of (23) decay exponentially in k for fixed ρ, b. At $k = L$ the high-degree space is zero; at $k = 0$ the same argument uses mean one. These endpoints do not require a separate assumption. □

Proposition 13.2 (Fixed smoothing need not be a small density perturbation)**.** *Fix $C > 0$ and $0 \leq \tau < 1$. There are energy-normalized positive potentials h_L satisfying all of* (1) *such that*

$$\|\pi_L - P_\tau\pi_L\|_s \longrightarrow \varepsilon > 0.$$

Proof. Take $\varepsilon = \min\{(1 - b)/2, (B - 1)/2\} > 0$. For all sufficiently large L, one has $k_L < L$. Set $\pi_L = 1 + \varepsilon\chi_{[L]}$ in those dimensions. It has mean one, minimum at least $(1 + b)/2 > b$, maximum at most $(1 + B)/2 < B$, and no nonconstant Fourier coefficients below degree L. In particular $\mathbb{E}\pi_L^s < 2$ and the required matching holds. Lemma 10.4 realizes it as h_LPh_L for a

positive energy-normalized potential. Use $h_L = 1$ in the remaining finitely many dimensions. Since $|\chi_{[L]}| = 1$,

$$\|\pi_L - P_\tau \pi_L\|_s = \varepsilon(1 - \tau^L) \longrightarrow \varepsilon.$$

□

This refutes only the uniformly vanishing L^s premise for a transfer argument. It does not refute a spectral comparison: the parity family itself belongs to an already controlled code class. Theorem 13.1 establishes the regularized spectrum on its specified interval, which excludes $\tau = 1$. Neither a comparison with the original spectrum nor an extension to every $\tau < 1$ is inferred.

14 Failure of separate variance and form comparison

The regularized spectrum does not supply a uniform comparison of the stationary variances. The obstruction below concerns actual chains, including their nonlinear reconstructed potentials. For an energy-normalized positive potential v, write

$$\mathcal{E}_v(f) = \frac{1}{2}\mathbb{E}_{x\sim\mu} \sum_y P(x,y)v(x)v(y)(f(x) - f(y))^2.$$

This is the Dirichlet form of K_v in stationary measure $vPv\, d\mu$.

Theorem 14.1 (Separate-comparison obstruction)**.** *Fix $0 < \rho, b < 1$, $C > 0$, and $0 \le \tau < 1$. There are positive energy-normalized potentials h_L, u_L, with*

$$\pi_L = h_L P h_L, \qquad w_L = u_L P u_L = P_\tau \pi_L,$$

both satisfying (1)*, and events A_L, such that, for $f_L = 1_{A_L}$,*

$$\frac{\mathrm{Var}_{\pi_L}(f_L)}{\mathrm{Var}_{w_L}(f_L)} \longrightarrow \infty, \qquad \frac{\mathcal{E}_{h_L}(f_L)}{\mathcal{E}_{u_L}(f_L)} \longrightarrow \infty.$$

Nevertheless both indicator Rayleigh quotients tend to one, and both full-chain gaps tend to $1-\rho$. The family can be chosen with $\|\pi_L - w_L\|_s \to c > 0$ and with π_L^s not uniformly integrable, despite $\mathbb{E}\pi_L^s < 2$.

Proof. Put $B = 2^{1/s}$ and $c = \min\{(1-b)/2, (B-1)/4\} > 0$. For all sufficiently large L, set

$$d = k_L + 1, \qquad n = \lfloor L/d \rfloor, \qquad p = 2^{-n}, \qquad t = cp^{-1/s}.$$

Take n disjoint blocks of size d, let z_i be their parities, and set $A = \{z_1 = \cdots = z_n = 1\}$. Remaining coordinates are unused. Thus $\mu(A) = p$, $d \sim C\sqrt{L}$, and $n \sim \sqrt{L}/C$. Define

$$\pi = 1 + t(1_A - p).$$

It has mean one and $\pi \ge 1 - c > b$. Moreover

$$\|\pi\|_s \le 1 + t(p^{1/s} + p) = 1 + c(1 + p^{1-1/s}) \le 1 + 2c < B.$$

The expansion $1_A = p\prod_{i=1}^n (1 + z_i)$ has nonconstant Fourier support only on unions of complete blocks, whose weights are at least $d > k_L$. This proves every labelled matching condition. The density $w = P_\tau \pi$ has the same mean, stationary floor, cap and matching, by positivity, L^s

contraction and the diagonal Fourier action. Lemma 10.4 gives the unique positive solutions $hPh = \pi$ and $uPu = w$, automatically of energy one. In the finitely many excluded dimensions take $h = u = 1$ and any nontrivial event.

The exact smoothed density is

$$w(z) = 1 + tp\left[\prod_{i=1}^{n}(1 + \tau^d z_i) - 1\right].$$

For $0 \le \epsilon < 1$ the maximum absolute deviation of this product from one is $(1+\epsilon)^n - 1$: its two extremes are $(1 \pm \epsilon)^n$, whose sum is at least two. Hence

$$\|w - 1\|_\infty \le tp\big[(1+\tau^d)^n - 1\big] \le cp^{1-1/s}\big[\exp(n\tau^d) - 1\big] \longrightarrow 0.$$

Here $n\tau^d \to 0$ for each fixed $\tau < 1$, including $\tau = 0$. No assertion uniform in τ approaching one is required.

Both reconstructed potentials belong to the same block-parity algebra. Indeed translations having parity one on each block, and arbitrary translations on unused coordinates, leave π and w invariant and commute with P. Translating a positive reconstructed potential therefore gives another solution of the same equation. Uniqueness in Lemma 10.4 proves invariance. The translation orbits are exactly the parity fibers, so $h = H(z)$ and $u = J(z)$.

We record actual return probabilities for either reconstructed potential. Write H for any positive block-parity potential of energy one and $G = PH$. Put

$$\theta = \rho^d, \quad \ell = (1-\theta)^n, \quad v = (1+\theta)^n, \quad a = \mathbb{E}H.$$

The quotient noise transition from z to z' is $p\prod_i(1 + \theta z_i z_i')$. Its density relative to uniform quotient measure lies between ℓ and v. Consequently

$$a\ell \le G \le av, \qquad v^{-1} \le a^2 \le \ell^{-1}, \qquad \frac{\ell^2}{v} \le G^2 \le \frac{v^2}{\ell}.$$

The middle bounds use $1 = \mathbb{E}HG$. If q is the stationary mass of A, then $q = pH(+)G(+)$. The return probability from any point of A to the *whole fiber* A is

$$r_A = K_H(x, A) = \frac{pvH(+)}{G(+)} = \frac{vq}{G(+)^2} \le \frac{v^2}{\ell^2}q.$$

This is not a single-vertex holding probability. Since $n\rho^d \to 0$, both ℓ and v tend to one. For the two chains the stationary event masses are

$$q_\pi = p[1 + t(1-p)] \sim tp \longrightarrow 0, \qquad q_w = p\left[1 + tp\big((1+\tau^d)^n - 1\big)\right] \sim p.$$

Thus $q_\pi/q_w \sim t \to \infty$ and both actual return probabilities r_π, r_w tend to zero. Reversibility gives exactly

$$\mathrm{Var}_\pi(1_A) = q_\pi(1 - q_\pi), \quad \mathcal{E}_h(1_A) = q_\pi(1 - r_\pi),$$

and the analogous identities with w, u. Both displayed ratios in the theorem are therefore asymptotic to $t = c2^{n/s}$, whereas each Rayleigh quotient tends to one.

To control the *full* gaps, apply Proposition 5.1, which includes all nonconstant ambient sectors and unused coordinates, to H and J. In its notation the lower and upper bounds are

$$\max\{0, \rho - \varepsilon_d\} \le \lambda_2 \le \min\{1, \max\{1-\alpha, \beta\}\},$$

where, eventually $d \geq 2$,

$$\alpha = \left(\frac{1-\rho^d}{1+\rho^d}\right)^n, \quad \beta = \frac{\rho+\rho^{d-1}}{1-\rho^d}, \quad \varepsilon_d = \frac{(1-\rho^2)\rho^{d-1}}{1-\rho^d}.$$

These tend to $1, \rho, 0$, respectively. No upper bound on the heights of H or J enters this proposition. Hence both full gaps tend to $1-\rho$. This conclusion for this particular family and every fixed $\tau < 1$ does not enlarge the universal smoothing interval in Theorem 13.1.

Finally, $\|t1_A\|_s = c$ and

$$\|tP_\tau 1_A\|_s \leq tp(1+\tau^d)^n \longrightarrow 0,$$

so the reverse triangle inequality gives $\|\pi - w\|_s \to c$. Directly,

$$\mathbb{E}\pi^s = (1-p)(1-tp)^s + p[1+t(1-p)]^s \longrightarrow 1 + c^s.$$

For every fixed threshold R, the value of π^s on A eventually exceeds R, while its contribution to this expectation tends to c^s. Thus π_L^s is not uniformly integrable (interpreting the random variables through their laws on the varying cubes). □

For clarity, a usual separate-comparison route to a gap transfer would require both $\mathrm{Var}_\pi(f) \leq V \mathrm{Var}_w(f)$ and $\mathcal{E}_u(f) \leq D\mathcal{E}_h(f)$ uniformly. Theorem 14.1 disproves its variance premise. It also disproves an upper bound $\mathcal{E}_h \leq D'\mathcal{E}_u$, and hence two-sided form equivalence. It does *not* disprove $\mathcal{E}_u \leq D\mathcal{E}_h$, a direct Rayleigh-quotient comparison, or the original gap conjecture. In particular it is not a counterexample to comparison results whose hypotheses include a common invariant measure, such as the setting described in [11]. The useful remaining route must retain the joint energy and variance information, rather than impose the disproved separate variance bound.

15 Positioning, failed routes and remaining question

Ordinary and reverse hypercontractivity are classical tools. Mossel, Oleszkiewicz and Sen explicitly state the Boolean reverse inequality, including negative exponents and tensorization, in equation (1.2) [2]. The use here is the floor-derived identity $gPg \geq b$ and its particular exponent iteration; we do not claim a new hypercontractive inequality. Representation-based decomposition of reversible chains is also classical; Boyd, Diaconis, Parrilo and Xiao discuss commuting symmetries and spectral reduction [3]. Our sector estimates explicitly retain all cosets and compare their signed kernels to the quotient kernel. The binary simplex construction is not original. Entropy-volume estimates and random-linear-code existence are classical Gilbert–Varshamov tools; the simultaneous column condition used here is verified directly rather than attributed as an extra clause of the coding theorem [5].

Weak Poincare inequalities and subgeometric convergence are established tools. Andrieu, Lee, Power and Wang [11] compare kernels with a common invariant measure using tails of a pointwise comparison weight; their Theorem 36 permits oscillation as a controlling functional. Section 9 instead compares the actual and reference edge measures and then transfers variance using the stationary cap. The inverse edge moment is derived from the nonlinear floor. This specialization and its explicit constants do not constitute a claim to invent weak Poincare comparison or prove priority. The background framework and its breadth are also discussed in [12]. Our bounded reading of the former covers Proposition 35, Theorem 36 and its proof, Remark 37 and the abstract; for the latter it covers the abstract, introduction and comparison context. The proofs above

include the discrete recursion and covariance summation explicitly. Ordinary hypercontractivity uses the parameter range stated in [10].

The positive-background proof in Section 10 also uses classical hypercontractivity and Cheeger comparison. The cited lazy reversible inequality is Theorem 3.5 of Peres' notes [13]; its statement and lower-bound proof on PDF pages 6–7 were inspected. We apply it to $(I + K)/2$ and record the resulting conservative constant explicitly. The model-specific steps retain the actual nonconstant denominator, use an exact two-component compression and control every small stationary cut. The lower potential is assumed in the general theorem and derived only for its two-level consequence. No new Cheeger inequality or global priority claim is made. The bounded primary check supporting this draft reads the relevant statements just cited, not the entire related literature. Targeted searches did not establish the priority of the specific moment counterexample, the quantitative iteration or the uniform code-potential estimates. Search results about "Doob graphs" concern a different object from the Doob transform used here and do not establish overlap or novelty. The results in this draft are proved as stated; mathematical correctness, originality and relevance remain separate questions. No claim of first discovery or frontier classification is made.

The regularized-spectrum theorem combines those classical inequalities with positive symmetric reconstruction and the finite-dimensional min–max principle. Its additional stationary exponent is obtained from the original cap by noise contraction, and the subsequent inverse-moment iteration is proved to cross a fixed supercritical target in finitely many steps. The critical sandwich acts on the full operator rather than only a Fourier window. The bounded primary checks for this step cover equations (1.1)–(1.2) of [2] and the Hypercontractivity Theorem in Chapter 10.1, printed page 284, of [6]. Targeted primary-domain searches about Doob stationary-density smoothing and Boolean matrix scaling did not establish priority or absence of overlap. We make no optimal-smoothing or new hypercontractive-inequality claim. The evidence supplement preserves unsuccessful routes rather than replacing their records. Proposition 2.2 excludes divergent critical weighting moments for two-level potentials and exact-cap failure for product potentials. The finite fixed point in Section 3 obstructs only the specified iteration. Theorem 2.1 refutes the exact numerical cap, while its own moments remain bounded and its gap remains asymptotically unweighted. A generic pair of densities close in total variation is not a counterexample unless both arise from the same h and satisfy the stationary hypotheses.

The historical source 353A had a false strict-class-separation claim based on three disjoint blocks A, B, C and the code generated by $A \cup B$ and $B \cup C$. Its two quotient coordinates can indeed be written as products of the three disjoint block parities, so every such potential already belongs to the larger three-block algebra allowed in Proposition 5.1. The corrected source 353B preserves that failed claim and its adverse audit, keeps the valid core code theorem, and substitutes the simplex witness of Proposition 5.2. Lack of a disjoint basis for one chosen code is not enough to prove absence of a larger disjoint-block representation.

The single remaining research question is whether (1), at fixed ρ, b, C, forces a uniformly positive full-chain gap or allows an admissible collapsing family. One unresolved bridge is a finite dimension-independent bound on $\mathbb{E}r^s$. Under energy normalization,

$$\mathbb{E}r^s = N^{-s}\mathbb{E}\big[\pi^s g^{-2s}\big], \qquad N \geq b.$$

The portion with $g > 1$ is controlled by the stationary cap. The joint small-g contribution is not controlled at the critical exponent by the finite inverse moments above. Additional exact matching gives strong subcritical convergence and the low-to-all operator control of Section 7, but does not settle this endpoint. For arbitrary positive potentials, the actual high block beyond

degree $\lfloor ck \rfloor$ remains uncontrolled. Section 8 identifies actual high-h/Ph sets carrying the squared mass of any non-Perron mode a fixed positive distance above the noise edge; no general bound on their killed return operator is assumed or proved. Section 10 proves killed control and a full gap with an additional fixed positive minimum of the normalized potential, and derives that minimum for the two-level class. It does not derive the minimum from (1) in general. Even an endpoint moment bound would still require a justified spectral argument for the general question. The code theorem and Theorem 10.1 exclude their stated classes of collapsing mechanisms, not all stationary densities satisfying (1). Under the critical cap, collapse requires normalized potential minima tending to zero; Proposition 10.5 realizes that behavior under all original hypotheses. Theorem 10.7 rules out gap collapse for this particular family by a full operator-norm estimate. Theorem 11.1 further refutes a derived positive Ph floor and pointwise ratio ceiling. Its different family has $Ph \to 0$ at one point but exponentially stable full spectrum; a small-denominator point alone is not a killed-return or cut-flow obstruction. Theorem 12.2 excludes uncorrected potential smoothing as an invariance argument for exact matching. Theorem 13.1 controls every eigenvalue of the actual reconstructed chain after sufficiently strong stationary-density smoothing, with the extra moment derived rather than assumed. Proposition 13.2 excludes a uniformly vanishing critical-norm perturbation argument. Theorem 14.1 also excludes a uniform separate variance upper comparison and two-sided form equivalence, even though both full gaps of its actual block-spike family tend to the noise gap. It does not exclude the reverse form bound or a direct Rayleigh comparison. A remaining approach must derive joint energy and variance control on the small-Ph, high-h/Ph region from the original assumptions. No uniform comparison of the original and reconstructed gaps is inferred. The full-gap question remains open.

Related gap criteria do not by themselves certify the novelty of this specialization. Gross [14], in Notation 2.1 and Theorem 2.2, uses a Riemannian Dirichlet generator, a reference logarithmic Sobolev inequality and exponential moments of an additive Schrodinger potential to obtain transformed coercive inequalities. We have not established a reduction of the present discrete normalized kernel to that theorem. Wang [15], Theorem 1.1, characterizes spectral gaps of ergodic Markov operators using tail norms. That qualitative criterion does not replace the explicit uniform all-cut estimate proved here. Our comparison reads those statements and adjacent definitions, not either full proof or the entire related literature; no priority or non-overlap conclusion follows.

16 Model assessments and review context

The predecessor is the public Working paper ARR-2026-4T35VQ6MZB8F99F4, version 1 [9]. Its exact PDF and first model-assessment report remain preserved separately. This revision incorporates new proofs and does not inherit that PDF's score, verdict or review status. An assessment of this version, if issued, applies only to the exact artifact hash and context named in its complete accompanying report. Working-paper deposit, model assessment, human peer review and editorial acceptance are distinct.

This successor preserves the integrated positive-background and vanishing-minimum arguments of notes I–L, their actual full-chain classification M, and note N as Section 12. It retains note O as Section 13 and retains note P as Section 14, retaining the distinction between full spectral control after density regularization and the unresolved original-chain comparison. The first analytic audits of N and O passed at their respective frozen source hashes without required corrections; both complete reports are preserved. O was assessed at commit `4086adaf`, with exact raw source and report digests in the accompanying manifest. The audit checked the finite exponent

crossing, endpoint cases, full-spectrum similarity and the norm-perturbation counterexample. Its source verdict is not a grade or assessment of this integrated text. This successor integrates that previously audited result; its inclusion does not confer an assessment on the integrated manuscript. The first analytic audit of P passed at commit `011130c6` without objections; its frozen source and complete report are retained. It checked actual reconstruction, whole-fiber return probabilities, the directions of the failed comparisons, both full-chain gap limits, and failure of critical uniform integrability. This source-level audit is not an assessment of the present integration or of a PDF. M uses a classical critical sandwich estimate; its quantified application to the constructed family closes that particular collapse route. Its source and eventual PDF require new exact-artifact reports. The preceding integration in Section 10 replaced the direct two-level proof with the broader theorem and its two-level consequence; the original direct proof and constants remain preserved separately. The earlier source and its complete audit reports remain preserved separately. This successor source and any eventual PDF require their own exact-artifact reports; earlier grades are not transferred.

This revision adds note Q as Section 11. Its first analytic source audit passed at commit `c2efb6bc` without blocking or nonblocking objections; the frozen source, complete report and bounded primary-source positioning are preserved by exact digests in the accompanying manifest. It checked the two shells and return region, the floor at every point, the vanishing denominator, critical ratio estimates, exact correction, actual reconstruction and full-spectrum comparison. This is a source-level verdict, not an assessment of this new integration or a PDF. Concentration and the hypercontractive sandwich are classical; no global priority claim is made for the construction. Its integration requires its own exact-artifact report.

Earlier analytical sources 349, 350, 351, corrected 353B and 355 received separate-role source audits. The added A–G research notes and notes I and J were also audited at their frozen hashes; their full reports are retained in the evidence packet. Notes I and J received analytic passes with no blocking objections under their distinct stated hypotheses. The separate first analytic reports for K, L, M, N, O, P and Q are retained with the current source packet; each verdict must be read at its own exact hash and is not a PDF assessment. Their exact sources and full reports are identified by SHA256 digests in the accompanying manifest; those verdicts are not assessments of this integrated text or a rendered PDF. These reports concern the named source bytes and do not by themselves certify the integrated manuscript or its eventual rendered PDF. The discarded 353A scope claim, the previous column-count bound, and the terminal rejection of an overbroad follow-up README sentence are preserved. That README was corrected to restrict localization to stationary-mean-zero modes; the underlying C/D statements already had that hypothesis.

The first analytic audit of M recorded one nonblocking notation correction: its exponential error statement is an upper bound, not an equality. Theorem 10.7 incorporates that correction explicitly. The frozen note, original full report and adjacent correction are retained. The separately produced and assessed PDF from source commit `8efcb73b` does not contain notes E–Q and does not assess this successor source. Its numerical assessment is not transferred to this manuscript. This source has not yet been compiled or visually reviewed as a new PDF.

Preparation and source auditing use Codex roles in the same provider workflow, with inherited context and a shared filesystem. Exact serving backend/version and effective reasoning effort are not independently attested. No blind, human, provider-independent or organizationally independent review is claimed. The user-requested model alias is not evidence of the actual backend. Full reports retain objections and context rather than only a favorable verdict.

The new results were derived analytically. No numerical experiment, executed mathematical

search or Lean verification is claimed for them. Document production, when performed, establishes rendering and artifact provenance rather than mathematical correctness. Classical inputs are attributed with their bounded reading scope; originality and significance are not certified by source correctness. The general stationary-gap and critical-moment questions remain open.

References

[1] J. R. Lee, *Entropy, convexity, and matrix scaling*, CSE 599S, Winter 2016, updated 24 January 2016, Section 2 setup. https://homes.cs.washington.edu/~jrl/teaching/cse599swi16/notes/lecture1.pdf.

[2] E. Mossel, K. Oleszkiewicz and A. Sen, *On reverse hypercontractivity*, arXiv:1108.1210v5 (2012), introduction, equations (1.1)–(1.2). https://arxiv.org/pdf/1108.1210.

[3] S. Boyd, P. Diaconis, P. Parrilo and L. Xiao, *Symmetry Analysis of Reversible Markov Chains*, author-hosted manuscript revised 29 November 2004, introduction and symmetry setup. https://web.stanford.edu/~boyd/papers/pdf/symmetry.pdf.

[4] M. Wootters, *CS250/EE387: Error Correcting Codes*, Problem Set 1, Winter 2025, Section 2.4 (simplex-code construction). https://web.stanford.edu/class/archive/cs/cs250/cs250.1254/Homework/HW1/hw1.pdf.

[5] V. Guruswami, *Introduction to Coding Theory, Notes 2: Gilbert–Varshamov bound*, CMU, January 2010, Lemma 5 and Theorem 8, including their proofs. https://www.cs.cmu.edu/~venkatg/teaching/codingtheory/notes/notes2.pdf.

[6] R. O'Donnell, *Analysis of Boolean Functions*, corrected May 2021 arXiv edition, Definition 2.46 and Proposition 2.47, PDF page 55; Theorem 9.21 and its proof, printed p. 262; Chapter 10.1 Hypercontractivity Theorem, printed p. 284. https://arxiv.org/pdf/2105.10386.

[7] S. Heilman, E. Mossel and K. Oleszkiewicz, *Strong Contraction and Influences in Tail Spaces*, introduction and same-exponent heat-smoothing statements. https://www.stevenheilman.org/~heilman/papers/hmo.pdf.

[8] A. Volberg, *Tail spaces estimates on Hamming cube and Bernstein–Markov inequality*, arXiv:2208.00472v4, introduction and Section 2. https://arxiv.org/pdf/2208.00472.

[9] L. Eriksson, *Stationary-to-weighting moment transfer and spectral rigidity for Boolean Doob kernels*, AIRR Working paper ARR-2026-4T35VQ6MZB8F99F4, version 1 (2026). Not peer reviewed or editorially accepted by AIRR. https://airr.science/papers/ARR-2026-4T35VQ6MZB8F99F4/versions/v1/.

[10] R. O'Donnell, *Analysis of Boolean Functions, Lecture 16: The Hypercontractivity Theorem*, CMU, Theorem 1.1. https://www.cs.cmu.edu/~odonnell/boolean-analysis/lecture16.pdf.

[11] C. Andrieu, A. Lee, S. Power and A. Q. Wang, *Comparison of Markov chains via weak Poincare inequalities with application to pseudo-marginal MCMC*, Annals of Statistics 50(6), 3592–3618 (2022), Proposition 35, Theorem 36 and Remark 37. doi:10.1214/22-AOS2241. https://awllee.github.io/papers/AOS2241.pdf.

[12] C. Andrieu, A. Lee, S. Power and A. Q. Wang, *Weak Poincare Inequalities for Markov chains: theory and applications*, arXiv:2312.11689v1 (2023), abstract and introductory context. https://arxiv.org/abs/2312.11689v1.

[13] Y. Peres, *Aspects of Markov chains*, Definition 3.4 and Theorem 3.5, including the lazy reversible lower-bound proof, PDF pages 6–7. https://sites.math.duke.edu/~rtd/CPSS2009/peres6.pdf.

[14] L. Gross, *Invariance of intrinsic hypercontractivity under perturbation of Schrodinger operators*, arXiv:2412.20282v2 (2025), Notation 2.1 and Theorem 2.2. https://arxiv.org/html/2412.20282v2.

[15] F.-Y. Wang, *Criteria of Spectral Gap for Markov Operators*, arXiv:1305.4460v6 (2013), introductory tail-norm definitions and Theorem 1.1. https://arxiv.org/html/1305.4460v6.