Product Marginals and Density-Moment Thresholds for Ising Transfer Kernels

Lluis Eriksson
Independent researcher

13 November 2026

Abstract

Fix a positive Ising coupling. We construct strictly positive Boolean weights whose labelled marginals on every set of at most a fixed positive fraction of coordinates equal uniform product marginals, while exponentially many normalized transfer eigenvalues approach one exponentially fast. The Perron stationary law has the same marginals. Binary codes with positive primal and dual distances, a positive product admixture, and finite block estimates supply the construction. An exact parity spectrum distinguishes transfer relaxation from single-site heat-bath relaxation. Conversely, for Boolean noise parameter $0 < \rho < 1$, density moments of order $s > (1+\rho)/(1-\rho)$ imply a dimension-independent normalized transfer gap, even allowing zero weights. Below this exponent, every norm budget $B > 1$ admits positive densities with exponentially collapsing gaps. At the critical exponent $s = (1+\rho)/(1-\rho)$ the exact worst-case norm boundary is $B_* = 2^{1/s}$: smaller budgets give a positive gap, whereas at B_* and above collapse is possible with any prescribed density floor below one. We also identify the sharp worst-case boundary for the asymptotic escaped critical moment $\Theta = \lim_{K\to\infty} \sup_L \mathbb{E}[r_L^s \mathbf{1}_{r_L > K}]$. The condition $\Theta < 1$ gives an explicit uniform gap, while a collapsing family has $\Theta = 1$. More strongly, a single collapsing family has tail moment above one at most one and total moment at most two in every dimension; it meets the same tail cap simultaneously at every larger cutoff. For fixed $0 \le K < 1$, this tail cap instead guarantees a positive gap. Nevertheless stable one-vertex spikes realize every finite $\Theta > 0$, so critical-power uniform integrability is sufficient but not necessary. These conclusions use classical hypercontractivity, sign compression, balanced product tilts and rank-one perturbations. For $k_s = (s-1)/s < K < 1$, finite parity backgrounds also give collapse at the attained tail budget $1 + C_s(1-K)K^{s-1} < 2$, with density floor exactly at least K, where $C_s = s^s/(s-1)^{s-1}$. We prove that this is the exact universal strict-tail threshold: every smaller budget yields an explicit uniform gap for arbitrary nonnegative densities, without a density floor. The proof combines reverse hypercontractivity, normalized sign compression and truncation; a sharp scalar total-moment envelope supplies a simpler intermediate bound. External priority is unestablished. Global many-body weights are allowed; no fixed-range Gibbs, local-dynamics or physical-gap conclusion follows.

The constructions serve different conclusions. The code family in Theorem 1 preserves labelled marginals and produces many slow transfer modes. The product tilts in Theorems 10 and 13 establish the moment-budget boundary, without the code family's marginal matching. The critical-tail section gives a sufficient tail condition, stable counterexamples to its necessity, and collapse at the finite-cutoff endpoint in Theorem 17. Section 12 then optimizes finite parity backgrounds to extend the sub-two tail-budget collapse region below cutoff one. Section 13 attains the constructive budget and allows the floor to equal the cutoff. Theorem 22 gives the exact scalar total-moment envelope. Theorem 23 then proves necessity of the constructed tail budget for every collapsing family. Readers interested in this sharp tail classification can start at Sections 14–15, using Theorem 21 for the matching endpoint. These are separate families; their properties are not asserted jointly.

1 The question and the transfer convention

Local marginal agreement and rapid relaxation are different kinds of information. The distinction is familiar in limited independence and Markov-chain theory, but a comparison must specify the dynamics: a slow single-site heat-bath chain does not automatically give a slow transfer operator. We make this distinction within a single fixed-coupling transfer family and require exact, rather than approximate, agreement of marginal input.

Let $\Omega_L = \mathbb{F}_2^L$, with Hamming distance $d(x,y)$. Fix $\beta > 0$ and put $q = e^{-2\beta}$. Our unweighted kernel is

$$M_\beta(x,y) = q^{d(x,y)} = e^{-\beta L} \exp\left(\beta \sum_{i=1}^{L} (-1)^{x_i + y_i}\right). \tag{1}$$

For a strictly positive probability law μ on Ω_L, define

$$A_\mu = \operatorname{diag}(\sqrt{\mu}) M_\beta \operatorname{diag}(\sqrt{\mu}). \tag{2}$$

The kernel M_β is a tensor power of $\left(\begin{smallmatrix} 1 & q \\ q & 1 \end{smallmatrix}\right)$ and is positive definite. The same is true of A_μ, which also has strictly positive entries. Write its eigenvalues in decreasing order as $\lambda_1 > \lambda_2 \geq \cdots > 0$, counting multiplicities, and set

$$\operatorname{gap}(A_\mu) = 1 - \lambda_2/\lambda_1.$$

Perron–Frobenius strictness gives a positive gap at each finite L; a bound uniform in L is a separate matter. Multiplying either kernel by a positive scalar does not change any eigenvalue ratio.

Let $u_L = \operatorname{Unif}(\Omega_L)$. Tensorization gives

$$\lambda_2(A_{u_L})/\lambda_1(A_{u_L}) = \frac{1-q}{1+q} = \tanh\beta. \tag{3}$$

For a law ν, the literal marginal input of order r is

$$\mathcal{M}_r(\nu) = (\nu_S : S \subseteq [L],\ |S| \leq r), \tag{4}$$

where coordinate labels are retained. Full conditional distributions, entrywise weight bounds, and the transfer matrix itself are not included in this tuple.

Theorem 1. *For each fixed $\beta > 0$ there are constants $\alpha, R, \gamma > 0$ such that, for all sufficiently large integers L, a strictly positive law μ_L on Ω_L has the following properties. Put $r_L = \lfloor \alpha L \rfloor$ and $m_L = 2^{\lceil RL \rceil}$.*

1. *$\mathcal{M}_{r_L}(\mu_L) = \mathcal{M}_{r_L}(u_L)$.*
2. *For every $2 \leq j \leq m_L$,*
$$0 < 1 - \lambda_j(A_{\mu_L})/\lambda_1(A_{\mu_L}) \leq 7e^{-\gamma L}. \tag{5}$$
3. *Let $v > 0$ be the unit Euclidean Perron vector of A_{μ_L}. The reversible Perron Doob chain*
$$P_L(x,y) = \frac{A_{\mu_L}(x,y) v_y}{\lambda_1 v_x}, \qquad \pi_L(x) = v_x^2, \tag{6}$$
has $\mathcal{M}_{r_L}(\pi_L) = \mathcal{M}_{r_L}(u_L)$ and the same normalized spectrum as A_{μ_L}.

One admissible choice is

$$R = \min\left\{\frac{1-h(1/4)}{2}, \frac{\beta}{4\log 2}\right\}, \quad h(\alpha) < R/2, \quad \gamma = \beta/2 - R\log 2 \geq \beta/4, \tag{7}$$

with $0 < \alpha < 1/2$ *and binary entropy* $h(p) = -p\log_2 p - (1-p)\log_2(1-p)$.

There are at least $m_L - 1$ slow nonleading modes in this statement. We do not assert that their eigenvalues coincide. The positivity and strict gap at each finite size are compatible with their exponentially small separation from the Perron root. The theorem does not impose a uniform lower bound on the weights or a bounded interaction range.

2 Codes with two distance conditions

We include the coding argument to make all hypotheses of the spectral construction explicit. Simultaneously good primal and dual codes are classical; stronger algorithmic constructions are available [5]. The relation between code projections and limited independence is standard background for the marginal comparison [4].

Lemma 2. *Suppose* $0 < \alpha, \delta < 1/2$ *and* $h(\alpha) < R < 1 - h(\delta)$. *Put* $k = \lceil RL \rceil$. *If* $k < L$ *and*

$$2^{k-L(1-h(\delta))} + 2^{-k+Lh(\alpha)} < 1, \tag{8}$$

there is a binary linear code $C \subseteq \mathbb{F}_2^L$ *of dimension* k *such that every nonzero codeword has weight at least* δL *and* $C^\perp$ *has no nonzero word of weight at most* $\lfloor \alpha L \rfloor$. *Every coordinate projection of* C *onto at most* $\lfloor \alpha L \rfloor$ *coordinates is surjective. The displayed conditions hold for all sufficiently large* L.

Proof. Choose a uniformly random binary $k \times L$ matrix G. For each fixed nonzero row vector $z \in \mathbb{F}_2^k$, the word zG is uniform in $\mathbb{F}_2^L$. The probability that some such zG has weight less than δL is at most

$$2^k 2^{-L} \sum_{j \leq \delta L} \binom{L}{j} \leq 2^{k-L(1-h(\delta))}.$$

Failure of full row rank is included in this event, since it produces a zero word from a nonzero message. For a fixed nonzero column vector $w \in \mathbb{F}_2^L$, the vector Gw is uniform in $\mathbb{F}_2^k$. Thus the probability of a nonzero dual word with weight at most $\lfloor \alpha L \rfloor$ is at most

$$2^{-k} \sum_{1 \leq j \leq \alpha L} \binom{L}{j} \leq 2^{-k+Lh(\alpha)}.$$

Their sum is less than one, so both failures can be avoided simultaneously.

For the entropy estimate, if $0 < p < 1/2$, set $z = p/(1-p) < 1$. Since $z^j \geq z^{pL}$ for $j \leq pL$, the binomial theorem gives

$$\sum_{j \leq pL} \binom{L}{j} \leq z^{-pL}(1+z)^L = 2^{Lh(p)}.$$

This bound applies to the inclusive sums and hence to the strict-weight failure event as well. No integrality assumption on δL is needed.

If the columns indexed by S were dependent, their relation would give a nonzero word of $C^\perp$ supported in S. Consequently G_S has rank $|S|$ whenever $|S| \leq \lfloor \alpha L \rfloor$. Uniform messages therefore give uniform projections onto $\mathbb{F}_2^S$. Finally, the exponents in (8) tend to minus infinity, since $k = RL + O(1)$ and both entropy inequalities are strict. □

3 A finite positive concentration estimate

The spectral step needs only the size and minimum distance of the code. The dual distance will enter through the marginal statement.

Proposition 3. *Let $C \subseteq \mathbb{F}_2^L$ be a binary linear code with size $m \geq 2$ and minimum nonzero weight d_C. Suppose its coordinate projections onto sets of size at most r are surjective. For $0 < \eta \leq 1/2$, let*

$$\mu = (1-\eta)\,\mathrm{Unif}(C) + \eta u_L, \qquad s = (m-1)e^{-2\beta d_C}. \tag{9}$$

Then μ is strictly positive, $\mathcal{M}_r(\mu) = \mathcal{M}_r(u_L)$, and every $2 \leq j \leq m$ satisfies

$$0 < 1 - \lambda_j(A_\mu)/\lambda_1(A_\mu) \leq 2s + 2m(\sqrt{\eta} + \eta). \tag{10}$$

Proof. Uniform codewords have uniform projections by surjectivity, and a mixture of two laws with the same marginals has those same marginals. The positive product part gives strict positivity at every configuration.

Every codeword has probability

$$a = (1-\eta)/m + \eta/2^L \geq 1/(2m).$$

The principal block of A_μ indexed by C is aM_{CC}. Its diagonal is a, and the off-diagonal row sums of M_{CC} are at most s: distinct codewords differ by a nonzero codeword of weight at least d_C. Thus all eigenvalues of M_{CC} lie in $[1-s, 1+s]$. For example, the norm of its symmetric off-diagonal part is bounded by the maximum absolute row sum. Principal interlacing gives

$$\lambda_j(A_\mu) \geq \lambda_m(A_\mu) \geq a(1-s) \quad (1 \leq j \leq m). \tag{11}$$

Independently, a coordinate Rayleigh vector gives $\lambda_1(A_\mu) \geq a$.

Write A_μ in C, C^c blocks. The outside principal block is positive semidefinite, and its trace is $\sum_{C^c} \mu = \eta(1 - m/2^L) \leq \eta$, so its norm is at most η. For the rectangular cross block B, all entries of M_β are at most one, and therefore

$$\|B\|^2 \leq \|B\|_{\mathrm{F}}^2 \leq \left(\sum_{C} \mu\right)\left(\sum_{C^c} \mu\right) \leq \eta.$$

The symmetric block matrix with off-diagonal blocks B, B^* and zero diagonal has norm $\|B\|$. Consequently

$$\left\|A_\mu - \mathrm{diag}(aM_{CC}, 0)\right\| \leq \sqrt{\eta} + \eta, \qquad \lambda_1(A_\mu) \leq a(1+s) + \sqrt{\eta} + \eta.$$

Subtract (11), divide by $\lambda_1 \geq a$, and use $a \geq 1/(2m)$ to obtain (10). Perron strictness proves the strict lower inequality. The upper bound is valid even when it is larger than one; no small-size success is claimed. □

Proof of Theorem 1, parts one and two. Use (7) and Lemma 2 with $\delta = 1/4$. The choice of R satisfies $0 < R < 1 - h(1/4)$ and $R \log 2 \leq \beta/4$. Choose codes for sufficiently large L, and put

$$m = 2^{\lceil RL \rceil}, \qquad \eta = e^{-2\beta L}/m^2.$$

This number lies in $(0, 1/4]$. Since $d_C \geq L/4$ and $m \leq 2e^{R(\log 2)L}$,

$$2(m-1)e^{-2\beta d_C} \leq 4e^{-(\beta/2 - R\log 2)L} = 4e^{-\gamma L}.$$

Also, using $m \geq 2$ and $\gamma \leq \beta$,

$$2m(\sqrt{\eta} + \eta) = 2e^{-\beta L} + 2e^{-2\beta L}/m \leq 3e^{-\beta L} \leq 3e^{-\gamma L}.$$

Proposition 3 now proves both claims. The condition $k < L$ ensures $m < 2^L$, so every asserted eigenvalue index exists. □

4 Stationary marginals survive the Perron transform

The weighting law in a transfer congruence need not equal the stationary law of its associated Markov chain. Matching the former alone would leave a significant interpretive gap. Code symmetry pays the latter exactly.

Lemma 4. *For the mixture (9), the Perron stationary law $\pi(x) = v_x^2$ satisfies $\mathcal{M}_r(\pi) = \mathcal{M}_r(u_L)$.*

Proof. For every $c \in C$, translation $x \mapsto x + c$ preserves μ, Hamming distance, and hence A_μ. The corresponding permutation commutes with A_μ and takes a positive unit Perron vector to another such vector. Uniqueness forces $v_{x+c} = v_x$, so π is invariant under translations by C. If $|S| \leq r$, every translation on $\mathbb{F}_2^S$ lifts to some $c \in C$ by the assumed surjectivity. The S-marginal of π is therefore invariant under all translations and must be uniform. □

Completion of Theorem 1. The Perron eigenvector equation gives row sums one for (6). Symmetry gives detailed balance:

$$\pi(x)P_L(x, y) = v_x A_\mu(x, y)v_y/\lambda_1 = \pi(y)P_L(y, x).$$

Conjugation by $\mathrm{diag}(v)$ shows that P_L is similar to A_μ/λ_1. It has strictly positive entries and the claimed spectrum. Apply Lemma 4 to obtain its stationary marginal assertion. For the uniform weight, the Perron vector and stationary law are uniform, and (3) supplies the comparison chain. □

5 An exact parity control with different dynamics

The positive parity family

$$\mu_t(x) = 2^{-L}(1 + t\chi_{[L]}(x)), \qquad |t| < 1, \qquad \chi_S(x) = (-1)^{\sum_{i \in S} x_i}, \tag{12}$$

has all proper marginals equal to product marginals. Its single-site heat-bath gap tends to zero as $|t| \to 1$ at fixed size $L \geq 2$ and fixed finite positive site rates [7]. Nevertheless, its transfer ratio at fixed β is bounded uniformly away from one, as follows.

Proposition 5. *Let $L \geq 2$, $\rho = \tanh\beta$, and $|t| < 1$. Up to a common positive scalar, the spectrum of A_{μ_t} is the union over unordered complement pairs $\{S, S^c\}$ of the two roots*

$$\frac{z_s \pm \sqrt{z_s^2 - 4\Delta}}{2}, \quad z_s = \rho^s + \rho^{L-s}, \quad s = |S|, \quad \Delta = (1 - t^2)\rho^L. \tag{13}$$

The Perron root is the upper root at $s = 0$, and the second eigenvalue is the upper root at $s = 1$. In particular,

$$\frac{\lambda_2(A_{\mu_t})}{\lambda_1(A_{\mu_t})} \leq \frac{\rho + \rho^{L-1}}{1 + \rho^L} \leq \tanh(2\beta) < 1. \tag{14}$$

The last bound is sharp as a supremum over $L \geq 2$ and $|t| < 1$.

Proof. The row-normalized product transfer $P_\beta = M_\beta/(1+q)^L$ acts on χ_S by $\rho^{|S|}$. Write $D_t = \operatorname{diag}\sqrt{1+t\chi_{[L]}}$. The weighted congruence, apart from a common scalar, is $D_t P_\beta D_t$, which is similar to $P_\beta D_t^2$ since D_t is invertible. Multiplication by $\chi_{[L]}$ exchanges χ_S and χ_{S^c}. The coefficient block on that ordered pair is

$$\begin{pmatrix} \rho^s & t\rho^s \\ t\rho^{L-s} & \rho^{L-s} \end{pmatrix}.$$

Its trace is z_s and its determinant is Δ, yielding (13). The 2^{L-1} complement pairs exhaust the character basis, with repetitions counted as usual.

The largest z_s is $z_0 = 1+\rho^L$, and the largest among proper pairs is $z_1 = \rho + \rho^{L-1}$. The upper root increases with z_s. The lower root in the empty/full pair is at most ρ^L, whereas the upper root for $s=1$ is at least ρ. Every other lower root is no larger than its own upper root. This identifies the two largest eigenvalues.

For fixed $\Delta > 0$, the upper root divided by z is

$$\frac{1+\sqrt{1-4\Delta/z^2}}{2},$$

which increases with $z \geq 2\sqrt{\Delta}$. Comparing $z_1 \leq z_0$ proves the first inequality in (14). For the second, write $b = \rho^{L-2} \leq 1$; the ratio is $\rho(1+b)/(1+\rho^2 b)$, increasing in b. Its maximum is $2\rho/(1+\rho^2) = \tanh(2\beta)$. For $L=2$ and $t \to 1$, $\Delta \to 0$, so equality is approached. □

This control rules out a tempting shortcut: the known parity heat-bath obstruction cannot be transferred by changing the name of the operator. The code construction uses a different geometric ingredient, a set of exponentially many mutually separated configurations while retaining surjective coordinate projections.

6 What marginal certificates cannot establish

Corollary 6. *Fix the constants of Theorem 1. Suppose F is a valid lower bound on the transfer gap for every strictly positive weight, and factors through the tuple*

$$(L, \beta, \mathcal{M}_{r_L}(\mu), \mathcal{M}_{r_L}(\pi)).$$

At the product input and for all sufficiently large L, its value is at most $7e^{-\gamma L}$. The same conclusion holds if only one of the two marginal tuples is supplied.

Proof. At the same L and β, the constructed weight and its stationary law supply exactly the same input as the uniform weight and uniform stationary law. A function of that input takes the same value on both. Validity on the constructed example and (5) for $j=2$ force the asserted bound. □

The corollary does not assert that every such functional is zero or that the product chain is slow. Its actual gap is $1-\tanh\beta$. Nor does it concern a certificate using full conditionals, uniform density bounds, the operator entries, or code membership: those are additional information. In particular, the weights here can have exponentially small probabilities and global many-body interactions. The theorem says nothing against a result restricted to uniformly controlled, fixed-range Gibbs weights. The Perron Doob chain updates a full configuration and is not a local single-site dynamics. We make no claim about continuum Yang–Mills, a physical mass gap, entropy tensorization, or a specified spatial mixing rate. The many-mode conclusion is spectral; no separate theorem about metastable sets is implicit.

7 Upper density control without a lower bound

The obstruction above allows the maximum density relative to product to increase with dimension. A uniform upper bound changes the conclusion; a positive lower bound is not required. Write $\rho = \tanh\beta$ and let P_ρ be the Boolean noise operator in $L^2(u_L)$, whose Walsh multipliers are $\rho^{|S|}$. For a nonnegative density r with $\mathbb{E}_{u_L} r = 1$, put

$$T_r = \operatorname{diag}(\sqrt{r}) P_\rho \operatorname{diag}(\sqrt{r}).$$

Since $P_\rho = M_\beta/(1+e^{-2\beta})^L$, for $\mu = 2^{-L}r$ we have $A_\mu = 2^{-L}(1+e^{-2\beta})^L T_r$. Thus the eigenvalue ratios agree. Here zeros of r are allowed; we keep the zero eigenvalues outside its support in the full 2^L-dimensional list.

Theorem 7. *Let $L \geq 1$, $0 < \rho < 1$, $B \geq 1$, and $0 \leq r \leq B$ with $\mathbb{E}_{u_L} r = 1$. Set*

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

The two largest eigenvalues of T_r satisfy $1 - \lambda_2/\lambda_1 \geq c(B,\rho) > 0$, independently of L. In particular, no positive lower bound on r is needed.

Proof. All norms and scalar products use u_L. The noise operator is positive semidefinite, fixes constants and is a contraction. Hence $\lambda_1 \leq B$. The unit vector $\sqrt{r}$ has Rayleigh quotient $\langle r, P_\rho r\rangle \geq (\mathbb{E} r)^2 = 1$, by Walsh expansion, so $1 \leq \lambda_1 \leq B$. On $S = \{r > 0\}$ the matrix is strictly positive; its Perron vector is positive on S and its leading eigenvalue is simple.

We use two classical Boolean hypercontractive inequalities. Put $q_* = 1 + \rho^{-2}$ and $p_* = 1 - \rho$. Then

$$\|P_\rho f\|_{q_*} \leq \|f\|_2, \qquad \langle f, P_\rho g\rangle \geq \|f\|_{p_*}\|g\|_{p_*} \quad (f, g \geq 0). \tag{15}$$

For clarity about conventions, equations (1.1)–(1.2) and tensorization in [10] use time $t = -\log\rho$ for the uniform Boolean semigroup. For the second inequality set $p_*' = p_*/(p_*-1) = -(1-\rho)/\rho$. Then $\frac{1}{2}\log((1-p_*')/(1-p_*)) = -\log\rho$. The reverse inequality gives $\|P_\rho g\|_{p_*'} \geq \|g\|_{p_*}$; reverse Hölder gives the displayed bilinear bound. This is first applied to positive functions, then to $f+\varepsilon, g+\varepsilon$ and the limit $\varepsilon \downarrow 0$. Finite sums and $p_* > 0$ justify that limit; we do not take a negative moment of a function with zeros.

If $\lambda_2 < \lambda_1/2$, the gap exceeds $1/2$. Otherwise $\lambda_2 \geq 1/2$. Choose a real unit second eigenvector w orthogonal to the Perron vector. It vanishes off S and has both signs. Write $w = sh_+ - th_-$, where $s, t > 0$, $s^2 + t^2 = 1$, and h_+, h_- are nonnegative orthonormal functions with disjoint supports. Positivity and the eigenrelation give pointwise $\lambda_2 h_\pm \leq T_r h_\pm$: on the corresponding sign support this follows from the eigenrelation, and elsewhere from positivity. Integration against the unit $h_\pm$ gives $\lambda_2 \leq \langle \sqrt{r}h_\pm, P_\rho\sqrt{r}h_\pm\rangle$. For $f_\pm = \sqrt{r}\,h_\pm$ this yields

$$\lambda_2 f_\pm \leq r P_\rho f_\pm, \qquad \|f_\pm\|_{q_*} \leq \frac{B}{\lambda_2}\|f_\pm\|_2, \qquad \|f_\pm\|_2^2 \geq \langle f_\pm, P_\rho f_\pm\rangle \geq \lambda_2.$$

The middle inequality uses the first part of (15). No division by r occurs.

Let $\theta = p_*(q_*-2)/(2(q_*-p_*))$. Ordinary Hölder applied to powers of f gives $\|f\|_2 \leq \|f\|_{p_*}^{\theta}\|f\|_{q_*}^{1-\theta}$, since $1/2 = \theta/p_* + (1-\theta)/q_*$. This does not require a normed space at exponent $p_* < 1$: the Hölder weights $2\theta/p_*$ and $2(1-\theta)/q_*$ are positive and sum to one. Thus each sign part satisfies

$$\|f_\pm\|_{p_*} \geq \sqrt{\lambda_2}(B/\lambda_2)^{-e_\rho}, \qquad \frac{1-\theta}{\theta} = \frac{q_*(2-p_*)}{p_*(q_*-2)} = e_\rho.$$

By the second part of (15), the cross entry obeys

$$d = \langle h_+, T_r h_- \rangle = \langle f_+, P_\rho f_- \rangle \geq \lambda_2 (B/\lambda_2)^{-2e_\rho}.$$

The compression of T_r to the basis (h_+, h_-) is $\left(\begin{smallmatrix} a & d \\ d & b \end{smallmatrix}\right)$. Its eigenvector $(s, -t)$ has eigenvalue λ_2, giving $a - \lambda_2 = dt/s$ and $b - \lambda_2 = ds/t$. Its other eigenvalue is $\lambda_2 + d/(st)$, because $s^2 + t^2 = 1$. The variational principle and $st \leq 1/2$ now imply $\lambda_1 - \lambda_2 \geq d/(st) \geq 2d$. Consequently

$$1 - \frac{\lambda_2}{\lambda_1} \geq \frac{2\lambda_2}{\lambda_1} (B/\lambda_2)^{-2e_\rho} \geq B^{-1} (2B)^{-2e_\rho}.$$

This does not deteriorate when one sign part has small norm. In the previously omitted case, $c(B, \rho) \leq 1/4 < 1/2$, since $e_\rho \geq 1$. If the support has one point, $\lambda_2 = 0$ and the gap is one, so that case is included as well. □

Corollary 8. *At fixed ρ, suppose a sequence has $\operatorname{gap}(T_{r_L}) \leq K e^{-\gamma L}$, with $K, \gamma > 0$. Then*

$$\max r_L \geq \left(\frac{2^{-2e_\rho}}{K} \right)^{1/(2e_\rho+1)} e^{\gamma L/(2e_\rho+1)}.$$

Proof. Apply Theorem 7 with $B = \max r_L \geq 1$ and rearrange. □

For Theorem 1, take $K = 7$ and its positive γ. The code mixture has maximum density $2^L(1-\eta)/|C| + \eta$, consistent with this necessary growth. Upper density growth alone is not asserted to be sufficient for collapse. Tiny minimum density alone cannot cause collapse under a uniform maximum-density bound. These are statements about this transfer family, not single-site heat-bath dynamics. The constants are not claimed optimal.

The moment analysis separates universal guarantees from adverse examples. The two strict exponent regimes are followed by auxiliary critical estimates and an explicit rate, then the exact critical-budget theorem. The final critical-tail section gives the sharp escaped-moment and fixed-cutoff boundaries, and a stable counterexample to necessity of uniform integrability. Earlier partial estimates are retained for their explicit rate and proof dependencies; they are not additional independent discoveries.

8 Density-moment threshold: the two strict regimes

A pointwise density cap is stronger than necessary. We next replace it by a dimension-independent finite moment. A second construction shows why one cannot simply assume that any finite moment suffices. All norms in this section use the uniform probability law u_L; the operator is the same T_r as in Section 7.

Theorem 9. *Fix $0 < \rho < 1$, $s > (1+\rho)/(1-\rho)$ and $B \geq 1$. For every $L \geq 1$ and density $r \geq 0$ with $\mathbb{E} r = 1$ and $\|r\|_s \leq B$,*

$$\operatorname{gap}(T_r) \geq \min\{1/2, (2B)^{-2e}\} > 0, \qquad e = \frac{2\rho s}{(1-\rho)(s(1-\rho) - 1 - \rho)}.$$

Neither a pointwise upper bound nor a positive lower bound is assumed.

Proof. Set $A = 1+\rho$, $p = 1-\rho$, $t = 1+1/\rho$ and

$$Q = \frac{st}{s+t} = \frac{s(1+\rho)}{s\rho+1+\rho}.$$

The condition on s is exactly $Q > A$, so $Q > A > 1 > p > 0$. Ordinary Boolean hypercontractivity gives

$$\|P_{\sqrt{\rho}} f\|_2 \le \|f\|_A, \qquad \|P_\rho f\|_t \le \|f\|_A.$$

Indeed the noise thresholds are respectively $\sqrt{(A-1)/(2-1)} = \sqrt{\rho}$ and $\sqrt{(A-1)/(t-1)} = \rho$. The reverse bilinear inequality in (15) applies at exponent p.

As before, testing the unit vector $\sqrt{r}$ gives $\lambda_1 \ge 1$. For any real unit vector h, put $f = \sqrt{r}\,h$ and $A_0 = 2s/(s+1) > A$. Hölder and probability-space monotonicity give

$$\langle h, T_r h\rangle = \|P_{\sqrt{\rho}} f\|_2^2 \le \|f\|_A^2 \le \|f\|_{A_0}^2 \le \|\sqrt{r}\|_{2s}^2 \|h\|_2^2 \le B,$$

where $1/A_0 = 1/(2s) + 1/2$. Thus $1 \le \lambda_1 \le B$ without using a maximum-density estimate.

If $\lambda_2 < \lambda_1/2$, the gap is greater than $1/2$. Otherwise $\lambda_2 \ge 1/2$. On the nonempty support of r the matrix has strictly positive entries. A real unit second eigenvector vanishes off this support and has both signs, by orthogonality to the positive Perron vector. Write it as $w = ch_+ - dh_-$ with $c, d > 0$, $c^2 + d^2 = 1$ and nonnegative orthonormal disjoint sign parts. Positivity and the eigenrelation imply $\lambda_2 h_\pm \le T_r h_\pm$ pointwise, including their zero sets. For either $f = \sqrt{r}\,h_\pm$, multiplication by $\sqrt{r}$ gives $\lambda_2 f \le r P_\rho f$. Hölder with $1/Q = 1/s + 1/t$ yields

$$\|f\|_Q \le (B/\lambda_2)\|f\|_A.$$

Integration against the unit sign part also gives $\lambda_2 \le \langle f, P_\rho f\rangle \le \|f\|_A^2$. Let $0 < \theta < 1$ solve $1/A = \theta/p + (1-\theta)/Q$. Hölder applied to powers, with positive weights $A\theta/p$ and $A(1-\theta)/Q$ summing to one, gives $\|f\|_A \le \|f\|_p^\theta \|f\|_Q^{1-\theta}$. No triangle inequality at exponent $p < 1$ is used. Combining these bounds,

$$\|f\|_p \ge \sqrt{\lambda_2}(B/\lambda_2)^{-e}, \qquad e = \frac{1-\theta}{\theta} = \frac{Q(A-p)}{p(Q-A)} = \frac{2\rho s}{(1-\rho)(s(1-\rho)-1-\rho)}.$$

Reverse hypercontractivity now bounds the cross entry by $D = \langle h_+, T_r h_-\rangle \ge \lambda_2 (B/\lambda_2)^{-2e}$. The two-dimensional compression, exactly as in the proof of Theorem 7, has eigenvalues λ_2 and $\lambda_2 + D/(cd)$. Since $cd \le 1/2$, the variational principle gives

$$\mathrm{gap}(T_r) \ge \frac{2D}{\lambda_1} \ge \frac{2\lambda_2}{\lambda_1}(B/\lambda_2)^{-2e} \ge (2B)^{-2e}.$$

Here $2\lambda_2/\lambda_1 \ge 1$ and $\lambda_2 \ge 1/2$. Combining the two cases proves the claim. One-point support has $\lambda_2 = 0$ and is already covered. No step divides by r. □

Theorem 10. *Fix $0 < \rho < 1$, $1 < s < s_c := (1+\rho)/(1-\rho)$ and $B > 1$. There are strictly positive mean-one densities r_L, for every $L \ge 1$, with $\|r_L\|_s \le B$ and $r_L \ge 1/B$, and constants $\kappa > 1$ and $0 < \omega < \kappa$, independent of L, such that*

$$0 < \mathrm{gap}(T_{r_L}) \le 2(\omega/\kappa)^L + \varepsilon^{-1}\kappa^{-L}, \qquad \varepsilon = (B-1)/2.$$

Thus a uniform positive lower density does not prevent exponential gap collapse throughout the strict subcritical moment regime.

Proof. Relabel the cube as $\{-1,1\}^L$. For a fixed tilt $0<\delta<1$, put

$$n=\left(\frac{(1+\delta)^s+(1-\delta)^s}{2}\right)^{1/s},\qquad w_\pm(x)=\frac{1\pm\delta x}{n},\qquad m=n^{-1}.$$

Strict convexity gives $n>1$, while $\|w_\pm\|_s=1$ and $\mathbb{E}w_\pm=m$. Define $R_{\pm,L}(x)=\prod_i w_\pm(x_i)$ and

$$Z_L=1+2\varepsilon m^L,\qquad r_L=\frac{1+\varepsilon(R_{+,L}+R_{-,L})}{Z_L}.$$

Product integration gives $\|R_{\pm,L}\|_s=1$ and $\mathbb{E}R_{\pm,L}=m^L$. Hence $r_L>0$, $\mathbb{E}r_L=1$, $\|r_L\|_s\le(1+2\varepsilon)/Z_L\le B$, and $r_L\ge 1/Z_L\ge 1/B$.

Write $Q=P_{\sqrt{\rho}}$. Since $Q^2=P_\rho$, the matrices $T_{r_L}=AA^*$ and $S_L=A^*A$, with $A=M_{\sqrt{r_L}}Q$, have identical spectra, including multiplicities. Here M_f denotes multiplication by f. Direct expansion gives

$$Z_LS_L=P_\rho+\varepsilon(X_L+Y_L),\qquad X_L=H_+^{\otimes L},\quad Y_L=H_-^{\otimes L},$$

where the one-site matrices, in the orthonormal Walsh basis $(1,x)$, are

$$H_\pm=\frac{1}{n}\begin{pmatrix}1 & \pm\delta\sqrt{\rho}\\ \pm\delta\sqrt{\rho} & \rho\end{pmatrix}.$$

They are positive definite and conjugate by $\mathrm{diag}(1,-1)$. Their common simple top eigenvalue is

$$\kappa(\delta)=\frac{1+\rho+\sqrt{(1-\rho)^2+4\rho\delta^2}}{2n(\delta)}.$$

At zero, $\kappa(0)=1$, $(\log\kappa)'(0)=0$, and differentiation yields

$$(\log\kappa)''(0)=\frac{2\rho}{1-\rho}-(s-1)>0.$$

All functions are smooth near zero. Taylor's theorem therefore gives $\kappa>1$ for every sufficiently small positive δ. Fix one such $\delta<1$. This choice depends on ρ,s and is independent of L; no uniform exponential rate as $s\uparrow s_c$ is asserted.

Set $\omega=\|H_+^{1/2}H_-^{1/2}\|$. Positive definiteness and submultiplicativity give $0<\omega\le\kappa$. If equality held, norm attainment at a unit vector v would force equality in both steps of

$$\|H_+^{1/2}H_-^{1/2}v\|\le\sqrt{\kappa}\,\|H_-^{1/2}v\|\le\kappa.$$

The second equality makes v a top eigenvector of H_-; the first then makes it a top eigenvector of H_+. But H_+-H_- has zero diagonal and nonzero equal off-diagonal entries, so is invertible. It cannot annihilate their common top eigenvector. Thus $\omega<\kappa$.

For $N=2^L\ge 2$, use the N-by-$2N$ matrix $C=[\sqrt{X_L},\sqrt{Y_L}]$. Its two Gram matrices satisfy

$$CC^*=X_L+Y_L,\qquad C^*C=\begin{pmatrix}X_L & \sqrt{X_L}\sqrt{Y_L}\\ \sqrt{Y_L}\sqrt{X_L} & Y_L\end{pmatrix}.$$

The block diagonal part has its two largest eigenvalues equal to κ^L. Tensorization of positive square roots and of operator norms gives $\|\sqrt{X_L}\sqrt{Y_L}\|=\omega^L$, which is also the norm of the self-adjoint off-diagonal perturbation. The variational eigenvalue bound consequently gives

$$\lambda_1(X_L+Y_L)\le\kappa^L+\omega^L,\qquad \lambda_2(X_L+Y_L)\ge\kappa^L-\omega^L>0.$$

Indeed C^*C and CC^* have the same N positive eigenvalues, with N additional zeros only in C^*C, so both leading bounds transfer even when $L = 1$. Also $X_L + Y_L \geq X_L$ gives $\lambda_1(X_L + Y_L) \geq \kappa^L$.

Finally P_ρ is positive semidefinite of norm one. For $F_L = Z_L S_L$, the variational principle and norm bound imply

$$\lambda_1(F_L) \leq \varepsilon(\kappa^L + \omega^L) + 1, \qquad \lambda_2(F_L) \geq \varepsilon(\kappa^L - \omega^L), \qquad \lambda_1(F_L) \geq \varepsilon\kappa^L.$$

Subtracting and dividing gives the claimed upper bound. The scalar Z_L cancels from the ratio, and S_L and the actual T_{r_L} have the same spectrum. Its entrywise positivity gives strict Perron separation at every finite L. Both terms in the upper bound decay exponentially, although the bound can exceed one at small L. □

Theorems 9 and 10 separate the two strict regimes at the same value s_c. The next two sections treat the critical exponent $s = s_c$: first quantitative bounds in two budget ranges, then the exact budget boundary including equality. At $B = 1$, Jensen forces $r = 1$ and the actual gap is $1 - \rho$; the sufficient lower bound is not claimed optimal. The tilted family does not claim product marginals for its weighting or stationary law. It isolates the moment obstruction for the same transfer operator, whereas Theorem 1 isolates the marginal-information obstruction. The earlier two-point-spike construction in the retained analytical sources covers the narrower interval $1 < s < \log 2/\log(2/(1+\rho))$; it remains valid but is superseded here by the product-tilt construction.

9 Auxiliary critical estimates and an explicit rate

At $s = s_c = (1+\rho)/(1-\rho)$ the size of the moment budget matters. We give a sufficient stability condition near $B = 1$ and a counterexample for every $B > 2$. These preliminary bounds are retained for their quantitative content. Section 10 then resolves the intervening budget range and the exact boundary. All norms in this section use the uniform probability measure u_L.

Theorem 11 (Small critical budgets)**.** *Fix $0 < \rho < 1$ and $s = s_c$. For any $L \geq 1$ and any density $r \geq 0$ with $\mathbb{E}r = 1$ and $\|r\|_s \leq B$, $B \geq 1$,*

$$\operatorname{gap}(T_r) \geq \max\{0, 1 - \rho - d_s(B)\}, \qquad d_s(B) = \begin{cases} (B^s - 1)^{1/s}, & s \geq 2, \\ \sqrt{\dfrac{2(B^s-1)}{s(s-1)}}\,(B+1)^{(2-s)/2}, & 1 < s < 2. \end{cases}$$

In particular, each fixed ρ admits a nontrivial interval of budgets $B > 1$ with a positive dimension-independent gap. No pointwise density bounds are required.

Proof. Put $Q = P_{\sqrt{\rho}}$ and $S_r = QM_rQ$, which has the spectrum of T_r. Ordinary Boolean hypercontractivity, recalled in [10, Eq. (1.1)], gives $\|Qf\|_q \leq \|f\|_2$ for $q = 2s/(s-1)$, since $\rho = (s-1)/(s+1) = 1/(q-1)$. Hölder's inequality with exponents s, q, q yields

$$|\langle f, QM_{r-1}Qg\rangle| \leq \|r - 1\|_s \|Qf\|_q \|Qg\|_q \leq \|r - 1\|_s \|f\|_2 \|g\|_2.$$

Thus $\|S_r - P_\rho\| \leq \|r - 1\|_s$ and $\lambda_2(S_r) \leq \rho + \|r - 1\|_s$. Also $\lambda_1(S_r) \geq \langle 1, S_r 1\rangle = 1$. It remains to show $\|r - 1\|_s \leq d_s(B)$.

Define $\phi_s(t) = t^s - 1 - s(t-1)$ for $t \geq 0$, so that $\mathbb{E}\phi_s(r) \leq B^s - 1$. For $s \geq 2$ and $u \geq 0$,

$$\phi_s(1+u) = s\int_0^u ((1+v)^{s-1} - 1)\,dv \geq u^s,$$

using $(1+v)^{s-1} \geq 1+v^{s-1}$. For $0 \leq u \leq 1$, the derivative of $\phi_s(1-u)-u^s$ is $s[1-(1-u)^{s-1}-u^{s-1}] \geq 0$ and its value at zero is zero. Hence $\phi_s(t) \geq |t-1|^s$, proving the first estimate by integration. For $1 < s < 2$, the integral Taylor remainder gives

$$\begin{aligned}\phi_s(t) &= s(s-1)(t-1)^2 \int_0^1 (1-v)[1+v(t-1)]^{s-2}\, dv \\ &\geq \frac{s(s-1)}{2} \frac{(t-1)^2}{(t+1)^{2-s}}.\end{aligned}$$

At $t = 0$ the identity follows by continuity; the integral is finite for $s > 1$. For $W = (r-1)^2/(r+1)^{2-s}$, Hölder gives

$$\mathbb{E}|r-1|^s \leq (\mathbb{E}W)^{s/2} (\mathbb{E}(r+1)^s)^{(2-s)/2}.$$

Use $\mathbb{E}W \leq 2(B^s-1)/(s(s-1))$ and $\|r+1\|_s \leq B+1$. Both formulas for $d_s(B)$ tend to zero as $B \downarrow 1$. Positive semidefiniteness supplies the trivial zero lower bound. □

Theorem 12 (Balanced critical tilt)**.** *Fix $0 < \rho < 1$, $s = s_c$, and $B > 2$. There are constants $c, d > 0$, independent of L, and strictly positive densities r_L, for every $L \geq 1$, such that*

$$\mathbb{E}r_L = 1, \quad \|r_L\|_s \leq B, \quad r_L \geq 1/B, \qquad 0 < \operatorname{gap}(T_{r_L}) \leq 2\sqrt{B-1}\, e^{-cd^2 L^{1/3}}.$$

Proof. For a uniform sign $x \in \{-1, 1\}$ and $|\delta| < 1$ define

$$n(\delta) = \left(\frac{(1+\delta)^s + (1-\delta)^s}{2} \right)^{1/s}, \quad w_\delta(x) = \frac{1+\delta x}{n(\delta)}, \quad m(\delta) = n(\delta)^{-1}.$$

The one-site matrix is

$$H(\delta) = \frac{1}{n(\delta)} \begin{pmatrix} 1 & \delta\sqrt{\rho} \\ \delta\sqrt{\rho} & \rho \end{pmatrix}, \qquad D = \operatorname{diag}(1, \rho).$$

It is positive definite with top eigenvalue

$$\kappa(\delta) = \frac{1+\rho+\sqrt{(1-\rho)^2+4\rho\delta^2}}{2n(\delta)}.$$

For $\tau(\delta) = \|D^{1/2} H(\delta)^{1/2}\|$, the square is the top eigenvalue of $D^{1/2} H(\delta) D^{1/2}$, so

$$\tau(\delta)^2 = \frac{1+\rho^2+\sqrt{(1-\rho^2)^2+4\rho^2\delta^2}}{2n(\delta)}.$$

These positive functions have even analytic logarithms near zero. Taylor expansion with bounded fourth derivative gives, at $s = s_c$,

$$\begin{aligned}\log n(\delta) &= \frac{\rho}{1-\rho}\delta^2 + O(\delta^4), \\ \log \kappa(\delta) &= O(\delta^4), \\ \log \tau(\delta) &= -\frac{\rho}{2(1-\rho^2)}\delta^2 + O(\delta^4).\end{aligned}$$

Indeed, the respective square-root numerators divided by two are $1 + \rho\delta^2/(1-\rho) + O(\delta^4)$ and $1 + \rho^2\delta^2/(1-\rho^2) + O(\delta^4)$. Consequently there are $0 < \delta_* < 1$, $C \geq 1$ and $c > 0$ such that, for $|\delta| \leq \delta_*$,

$$|\log \kappa(\delta)| \leq C\delta^4, \qquad \log \tau(\delta) \leq -c\delta^2.$$

For example, $c = \rho/[4(1-\rho^2)]$ works after reducing δ_*. Choose a fixed $0 < d \le \delta_*$ with $Cd^4 \le \log(B-1)$, possible because $B > 2$. Put

$$\delta_L = dL^{-1/3}, \quad a_L = \kappa(\delta_L)^{-L}, \quad R_L = \prod_{i=1}^{L} w_{\delta_L}(x_i), \quad Z_L = 1 + a_L m(\delta_L)^L, \quad r_L = \frac{1 + a_L R_L}{Z_L}.$$

The coefficient balances the two top block eigenvalues exactly. Since $|\log a_L| \le Cd^4 L^{-1/3} \le \log(B-1)$, we have $a_L \le B-1$ and $a_L \to 1$. Product measure gives $\|R_L\|_s = 1$ and $\mathbb{E}R_L = m(\delta_L)^L$, with $0 < m(\delta_L) < 1$. Therefore $\mathbb{E}r_L = 1$, $\|r_L\|_s \le (1+a_L)/Z_L \le B$, and $r_L \ge 1/Z_L \ge 1/B$ for every $L \ge 1$.

Let $X_L = H(\delta_L)^{\otimes L}$. Spectral equivalence gives $F_L = Z_L S_{r_L} = P_\rho + a_L X_L$. Both summands have top eigenvalue one. For the N by $2N$ matrix $\mathcal{C}_L = [P_\rho^{1/2}, \sqrt{a_L} X_L^{1/2}]$, $N = 2^L$, the Gram matrix $\mathcal{C}_L^* \mathcal{C}_L$ has those summands on its diagonal and off-diagonal block of norm

$$\eta_L = \sqrt{a_L}\, \tau(\delta_L)^L \le \sqrt{B-1}\, e^{-cd^2 L^{1/3}}.$$

Its block diagonal has two leading eigenvalues exactly one. The off-diagonal self-adjoint perturbation has norm η_L, whence $\lambda_1 \le 1 + \eta_L$ and $\lambda_2 \ge 1 - \eta_L$. The N positive eigenvalues of $\mathcal{C}_L \mathcal{C}_L^* = F_L$ are precisely the positive eigenvalues of $\mathcal{C}_L^* \mathcal{C}_L$; the latter has N additional zeros. Thus the bounds apply to the first two eigenvalues of F_L, including at $L = 1$. Since $F_L \ge P_\rho$, $\lambda_1(F_L) \ge 1$ and $\mathrm{gap}(F_L) \le 2\eta_L$. Scalar normalization and AA^*/A^*A preserve the ratio. Strict entrywise positivity of T_{r_L} gives its strictly positive finite-dimensional gap. □

The uniform component is a spectral channel here, rather than an additive error to discard. At the critical exponent the tilted top eigenvalue has zero quadratic variation, while its cross norm with the uniform channel has negative quadratic variation. The dimension-dependent tilt separates these orders, and balancing removes the residual diagonal mismatch. The small-budget estimate uses classical hypercontractivity and elementary perturbation theory; the large-budget family asserts neither matching marginals nor fixed-range interactions. The sufficient budget $B > 2$ is not optimal; the next section identifies the exact threshold without claiming a rate there.

10 The exact critical moment budget

The preceding balanced construction has a sharper limiting moment than the triangle inequality used to obtain its all-dimension rate bound. Combining that limit with a sign decomposition of a second eigenvector identifies the entire critical budget boundary. Throughout this section, all norms and expectations use uniform probability on $\{-1, 1\}^L$.

Theorem 13 (Exact critical budget)**.** *Fix $0 < \rho < 1$, $s = (1+\rho)/(1-\rho)$, and $B_* = 2^{1/s}$. For every $L \ge 1$ and $r \ge 0$ with $\mathbb{E}r = 1$,*

$$2\lambda_2(T_r)^s \le \mathbb{E}r^s. \tag{16}$$

Consequently $\|r\|_s \le B < B_*$ implies

$$\mathrm{gap}(T_r) \ge 1 - B/B_* > 0.$$

For each fixed $b \in (0,1)$ there are strictly positive densities $\widetilde{r}_L$, for every $L \ge 1$, such that

$$\mathbb{E}\widetilde{r}_L = 1, \qquad \|\widetilde{r}_L\|_s \le B_*, \qquad \widetilde{r}_L \ge b, \qquad \mathrm{gap}(T_{\widetilde{r}_L}) \longrightarrow 0.$$

Thus the infimum of the gap over all dimensions and densities with $\|r\|_s \le B$ is positive exactly for $1 \le B < B_*$. The same classification holds with any prescribed fixed floor $r \ge b$, $0 < b < 1$. The positive lower bound displayed above is not asserted optimal; no convergence rate at $B = B_*$ is asserted.

Proof. Two sign regions. Put $Q = P_{\sqrt{\rho}}$ and $p = 1 + \rho = 2s/(s+1)$. Endpoint Boolean hypercontractivity gives $\|Qg\|_2 \le \|g\|_p$. For h supported on C, Hölder's inequality gives

$$\begin{aligned}\langle h, T_r h\rangle &= \|Q(\sqrt{r}\,h)\|_2^2 \\ &\le \|\sqrt{r}\,h\|_p^2 \le \left(\mathbb{E}[r^s \mathbf{1}_C]\right)^{1/s} \|h\|_2^2. \end{aligned} \tag{17}$$

Indeed, before raising to $2/p$, the conjugate exponents are $2/(2-p)$ and $2/p$, and $p/(2-p) = s$.

On its nonempty active support $A = \{r > 0\}$, the matrix T_r is strictly entrywise positive. It has a simple positive Perron root and a Perron vector strictly positive on A. Its other coordinates form a zero block. If $\lambda_2 = 0$, (16) is immediate. Otherwise a real second eigenvector f vanishes off A and is orthogonal to the Perron vector. Both its positive and negative parts f_+, f_- are therefore nonzero, including when λ_2 has multiplicity greater than one. Entrywise positivity gives

$$\lambda_2 \|f_+\|_2^2 = \langle f_+, T_r f_+\rangle - \langle f_+, T_r f_-\rangle \le \langle f_+, T_r f_+\rangle.$$

Applying (17) on $\{f > 0\}$ yields $\lambda_2^s \le \mathbb{E}[r^s \mathbf{1}_{\{f>0\}}]$. The negative part gives the analogous bound on $\{f < 0\}$. Adding over these disjoint regions proves (16). No mean-zero condition under uniform measure is imposed on f. Finally $\sqrt{r}$ is a unit vector and

$$\lambda_1(T_r) \ge \langle r, P_\rho r\rangle \ge (\mathbb{E} r)^2 = 1$$

by the nonnegative Walsh eigenvalues. This proves the positive gap bound below B_*. At $B = 1$, strict Jensen forces $r = 1$ and the actual gap is $1 - \rho$.

The balanced family's exact moment. Use the single-site quantities n, w, m, H, κ, τ from Theorem 12, with any sufficiently small fixed $d > 0$ and $\delta_L = dL^{-1/3}$. This step does not impose the earlier condition involving $B > 2$. Define

$$a_L = \kappa(\delta_L)^{-L}, \quad R_L = \prod_{i=1}^{L} w_{\delta_L}(x_i), \quad Z_L = 1 + a_L m(\delta_L)^L, \quad r_L = (1 + a_L R_L)/Z_L.$$

The critical expansions already proved give $a_L \to 1$ and $m(\delta_L)^L \to 0$, hence $Z_L \to 1$. For the latter limit, $\log m(\delta) = -(s-1)\delta^2/2 + O(\delta^4)$ and $L\delta_L^2 \to \infty$, whereas $L\delta_L^4 \to 0$. The overlap $\eta_L = \sqrt{a_L}\tau(\delta_L)^L$ tends to zero. More explicitly, the quadratic coefficient of $\log \tau$ is

$$\frac{1}{2}\left[\frac{\rho^2}{1-\rho^2} - \frac{s-1}{2}\right] = -\frac{\rho}{2(1-\rho^2)}.$$

The same Gram comparison, now without a prescribed all-L budget, gives $1 \le \lambda_1(S_L) \le 1 + \eta_L$ and $\lambda_2(S_L) \ge 1 - \eta_L$ for $S_L = P_\rho + a_L H(\delta_L)^{\otimes L}$. For large L, the second bound is positive, so it corresponds to a nonzero eigenvalue rather than an added Gram zero. Since $QM_{r_L}Q = S_L/Z_L$, both of its top eigenvalues tend to one. They are the actual eigenvalues of T_{r_L}: for $X = M_{\sqrt{r_L}}Q$, the square matrices X^*X and XX^* have identical spectra, including zeros and multiplicities.

For each fixed $0 < t < s$, direct single-site expansion yields

$$\log \mathbb{E} w_\delta^t = \frac{t(t-s)}{2}\delta^2 + O(\delta^4).$$

Thus $\mathbb{E}R_L^t \to 0$, while $\mathbb{E}R_L^s = 1$ exactly. The remainders may depend on fixed t; no uniformity near the endpoints is needed. Set $\theta = \min\{1, s/2\}$. Both θ and $s - \theta$ belong to $(0, s)$, and

$$0 \le (1+y)^s - 1 - y^s \le s2^{s-1}(y^\theta + y^{s-\theta}), \qquad y \ge 0. \tag{18}$$

For $0 \le y \le 1$, the mean-value theorem bounds the left side by $s2^{s-1}y \le s2^{s-1}y^\theta$. For $y \ge 1$, apply that bound to $1/y$ and multiply by y^s, obtaining $s2^{s-1}y^{s-1} \le s2^{s-1}y^{s-\theta}$. The lower bound follows from $s > 1$. Substitute $y = a_L R_L$ in (18) and take expectations. The two interior moments vanish and $a_L \to 1$, so

$$\mathbb{E}(1 + a_L R_L)^s = 1 + a_L^s + o(1) \longrightarrow 2.$$

Dividing by $Z_L^s \to 1$ proves $B_L := \|r_L\|_s \to 2^{1/s} = B_*$.

The exact boundary and a prescribed floor. Each r_L is nonconstant, so $B_L > 1$. Put

$$\varepsilon_L = \begin{cases} 0, & B_L \le B_*, \\ (B_L - B_*)/(B_L - 1), & B_L > B_*, \end{cases} \qquad \widetilde{r}_L = (1 - \varepsilon_L)r_L + \varepsilon_L.$$

The only denominator used is positive. Since $B_L \to B_* > 1$, we have $0 \le \varepsilon_L < 1$ and $\varepsilon_L \to 0$. Normalization and positivity persist, and Minkowski gives

$$\|\widetilde{r}_L\|_s \le (1 - \varepsilon_L)B_L + \varepsilon_L \le B_*.$$

Because $Z_L > 1$, both r_L and 1 are at least $1/Z_L$, so the mixture has that same floor.

For signed h, the dual endpoint estimate $\|Qf\|_q \le \|f\|_2$, $q = 2s/(s-1)$, and Hölder with $1/s + 2/q = 1$ give $\|QM_hQ\| \le \|h\|_s$. Hence

$$\|QM_{\widetilde{r}_L - r_L}Q\| \le \varepsilon_L \|1 - r_L\|_s \le \varepsilon_L(1 + B_L) \longrightarrow 0.$$

Weyl's bound and the exact X^*X/XX^* spectral identity imply that both top eigenvalues of $T_{\widetilde{r}_L}$ still tend to one. Its normalized gap therefore tends to zero. This argument requires neither pointwise upper control nor a square-root multiplier estimate.

Given fixed $0 < b < 1$, eventually $1/Z_L \ge b$. Replace the finite initial segment by the constant density 1. Every stated constraint then holds for every $L \ge 1$, with the same limiting gap. The same sequence is admissible for all $B \ge B_*$. Each finite transfer matrix remains strictly positive, so its individual gap is positive; only the dimension-free infimum vanishes. A floor $b = 1$ would force the uniform density and is explicitly excluded. □

The sign-region argument is classical. In the additive lattice setting, Biskup and König [3, Lemma 4.5 and Eq. (4.29)] bound the second eigenvalue using the two sign regions and split a potential budget between them. Here endpoint hypercontractivity gives the restricted s-moment estimate, while the balanced construction and vanishing dilution attain the boundary asymptotically. This comparison credits the mechanism; it does not establish historical priority of the exact multiplicative-cube classification. The new family makes no matching-marginal or fixed-range interaction claim.

11 The escaped critical-moment boundary

Fix $0 < \rho < 1$ and $s = (1+\rho)/(1-\rho)$ throughout this section. Expectations and norms refer to the uniform Boolean measure. Write $Q = P_{\sqrt{\rho}}$ and $T_r = M_{\sqrt{r}}P_\rho M_{\sqrt{r}}$. We also use $\mathrm{gap}(r) := \mathrm{gap}(T_r)$

when the density shorthand is convenient. Here $r \geq 0$ and $\mathbb{E}r = 1$. For a family indexed by $L \geq 1$, define

$$H(K) = \sup_L \mathbb{E}[r_L^s \mathbf{1}_{r_L > K}], \qquad \Theta = \lim_{K\to\infty} H(K) = \inf_{K\geq 1} H(K).$$

The nonnegative monotone limit is allowed to be infinite. Uniform integrability of the critical powers is the special case $\Theta = 0$.

Theorem 14 (Finite cutoff with tail budget below one)**.** *Suppose $K \geq 1$, $0 \leq \tau < 1$, and $\mathbb{E}[r^s \mathbf{1}_{r>K}] \leq \tau$. Put*

$$t = \tau^{1/(2s)}, \quad m = \frac{1+t}{2}, \quad d = \frac{1-t}{2}, \quad B_K = (K^{s-1} + \tau)^{1/s}.$$

Then

$$\operatorname{gap}(T_r) \geq \min\left\{1 - m^2, \frac{2d^{2s^2} K^{1-s^2}}{B_K}\right\} > 0.$$

In particular, $\Theta < 1$ implies a dimension-independent positive gap. No optimality is claimed for the displayed positive constant.

Proof. For $w \geq 0$ and $\|h\|_2 = 1$, Hölder and forward hypercontractivity, with $p = 1 + \rho$ and $p/(2-p) = s$, give

$$\|Q\sqrt{w}h\|_2^2 \leq \|\sqrt{w}h\|_p^2 \leq \|w\|_s.$$

Since $r^s \mathbf{1}_{r\leq K} \leq K^{s-1} r$, the total moment is at most $K^{s-1} + \tau$. Thus $\lambda_1(T_r) \leq B_K$. The unit vector $\sqrt{r}$ also gives $\lambda_1 \geq \langle r, P_\rho r\rangle \geq 1$ by the nonnegative Walsh spectrum. If the active support has one point, $\lambda_2 = 0$ and the gap is one. If $\lambda_2 < m^2$, the gap is greater than $1 - m^2$.

Otherwise choose a real second eigenvector, orthogonal to the positive active Perron vector. Write it as $uh_+ - vh_-$, with $u, v > 0$, $u^2 + v^2 = 1$, and disjointly supported nonnegative unit vectors h_+, h_-. Their uniform means need not be equal. Set $D = \langle h_+, T_r h_-\rangle > 0$. The eigenvector equations give

$$A_+ = \lambda_2 + (v/u)D, \qquad A_- = \lambda_2 + (u/v)D, \quad A_\pm = \langle h_\pm, T_r h_\pm\rangle.$$

For each sign split $f = \sqrt{r}h_\pm$ at $r = K$. Then $\|Qf\|_2 \geq m$ and $\|Qf_{\text{high}}\|_2 \leq t$, so

$$\|Qf_{\text{low}}\|_2 \geq d, \qquad \|f_{\text{low}}\|_p \geq d, \qquad \|f_{\text{low}}\|_2 \leq \sqrt{K}.$$

For $a = 1-\rho$ and $\theta = s^{-2}$ the identity $1/p = \theta/a + (1-\theta)/2$ and power-Hölder give $\|g\|_p \leq \|g\|_a^\theta \|g\|_2^{1-\theta}$. This applies ordinary Hölder to powers and requires no quasi-norm triangle inequality. The positive lower p-norm permits division; hence $\|f_{\text{low}}\|_a \geq d^{1/\theta} K^{-(1-\theta)/(2\theta)}$. Positivity and reverse hypercontractivity yield

$$D \geq \langle f_{+,\text{low}}, P_\rho f_{-,\text{low}}\rangle \geq d^{2s^2} K^{1-s^2}.$$

The sign compression has eigenvalues λ_2 and $\lambda_2 + D/(uv)$. Its largest eigenvalue is at most λ_1, so $\lambda_1 - \lambda_2 \geq 2D$. Dividing by $\lambda_1 \leq B_K$ proves the bound, including repeated second eigenvalues. Finally, if $\Theta < 1$, choose $\Theta < \tau < 1$ and a single finite K with $H(K) \leq \tau$. The same constants then apply to every L. □

Uniform integrability is sufficient but stronger than needed. For example, a bound

$$\sup_L \mathbb{E}[r_L^s \psi(r_L)] \leq C < \infty$$

with a nonnegative nondecreasing function ψ diverging at infinity implies $H(K) \le C/\psi(K)$ when $\psi(K) > 0$. This is a weighted-moment condition; an arbitrary such ψ need not be a Young function or define a Luxemburg norm. A fixed $s+\varepsilon$ moment bound also implies $H(K) \le K^{-\varepsilon} \sup_L \mathbb{E} r_L^{s+\varepsilon}$. Neither condition follows from a mere critical-moment bound.

A total bound $\sup_L \|r_L\|_s \le B$ does imply

$$\Theta \le B^s - 1. \tag{19}$$

Indeed $\mathbb{E}[r_L \mathbf{1}_{r_L > K}] \le B^s K^{1-s}$, while Jensen for $X_L = r_L \mathbf{1}_{r_L \le K}$ gives $\mathbb{E} X_L^s \ge (1 - B^s K^{1-s})_+^s$. Subtraction and the limit $K \to \infty$ prove (19). For $B < 2^{1/s}$ this recovers a positive gap. The direct exact-budget theorem already gives the stronger explicit lower bound $1 - B/2^{1/s}$.

Proposition 15 (Escaped moment with a stable gap)**.** *For any $M > 1$, $B > M^{1/s}$, and $0 < b < 1$, there are densities r_L with $\mathbb{E} r_L = 1$, $r_L \ge b$, $\|r_L\|_s \le B$, and*

$$\mathbb{E} r_L^s \to M, \qquad \mathbb{E}[r_L^s \mathbf{1}_{r_L > K}] \to M - 1 \quad (K > 1 \text{ fixed}),$$

while $\mathrm{gap}(T_{r_L}) \to 1 - \rho$ and $\mathrm{gap}(T_{r_L}) \ge (1-\rho)/2$ for every $L \ge 1$.

Proof. Set $N = 2^L$, $c = (M-1)^{1/s}$, choose a vertex z_L, and first put

$$q_L = \frac{1 + cN^{1/s} \mathbf{1}_{z_L}}{Z_L}, \qquad Z_L = 1 + cN^{1/s-1}.$$

These densities have mean one, lower bound $1/Z_L \to 1$, and moments

$$\mathbb{E} q_L^s = \frac{1 - N^{-1} + (N^{-1/s} + c)^s}{Z_L^s} \longrightarrow 1 + c^s = M.$$

For each fixed $K > 1$ only the exceptional vertex eventually exceeds K, and its contribution tends to $c^s = M - 1$.

The operators T_{q_L} and $S_L = Q M_{q_L} Q$ are isospectral, including multiplicities, since they are AA^* and A^*A for $A = M_{\sqrt{q_L}} Q$. In probability-normalized L^2, the vector $e_z = \sqrt{N} \mathbf{1}_{z_L}$ is a unit vector and $M_{\mathbf{1}_{z_L}} = |e_z\rangle\langle e_z|$. Therefore

$$S_L = \frac{P_\rho + cN^{1/s} Q|e_z\rangle\langle e_z|Q}{Z_L}, \qquad \|Q|e_z\rangle\langle e_z|Q\| = \langle e_z, P_\rho e_z\rangle = \left(\frac{1+\rho}{2}\right)^L.$$

It follows that $\|S_L - P_\rho\| \le \delta_L$, where

$$\delta_L = \frac{c d_\rho^L + c a_\rho^L}{Z_L}, \quad d_\rho = 2^{-(1-1/s)}, \quad a_\rho = \frac{1+\rho}{2^{2\rho/(1+\rho)}}.$$

Both bases lie strictly between zero and one. For a_ρ, strict convexity of $t \log t$ on $[1,2]$ gives $(1+\rho)\log(1+\rho) < 2\rho \log 2$.

The largest two eigenvalues of P_ρ are 1 and ρ, counting multiplicity. The min-max principle bounds the changes in both by δ_L. Also $\langle 1, S_L 1\rangle = 1$, so $\lambda_1(S_L) \ge 1$. Consequently

$$\left| \frac{\lambda_2(S_L)}{\lambda_1(S_L)} - \rho \right| \le (1+\rho)\delta_L.$$

Choose L_0 so that for every $L \ge L_0$ the floor and norm bounds hold and $c(d_\rho^L + a_\rho^L) \le (1-\rho)/(2(1+\rho))$. Existence uses $B^s > M$ and the strict exponential contractions. Finally set $r_L = 1$ for $L < L_0$ and $r_L = q_L$ otherwise. The prefix has gap $1 - \rho$; the tail has the asserted common bound. All limits persist. □

Proposition 16 (Sharpness and nonnecessity of the tail criterion)**.** *The universal guarantee* $\Theta < 1$ *cannot be extended to* $\Theta \le 1$*: the exact-budget family of Theorem 13 has* $\Theta = 1$ *and gaps tending to zero. For every finite* $\Theta_0 > 0$*, there are nevertheless families with* $\Theta = \Theta_0$ *and a uniform positive gap, even with an arbitrarily prescribed density floor below one. Thus* $\Theta \ge 1$ *is necessary for collapse, but not sufficient.*

Proof. First suppose $r_L \to 1$ in probability and $\mathbb{E}r_L^s \to M < \infty$. For each fixed $K > 1$, the function $F_K(x) = x^s \mathbf{1}_{x \le K}$ is bounded and continuous at one, so $\mathbb{E}F_K(r_L) \to 1$. This follows by splitting a neighborhood of one from its probability-small complement and works on the varying cubes. Therefore the tail moments tend to $M - 1$ for every such K, giving $\Theta \ge M - 1$. For the upper bound fix $\varepsilon > 0$ and $K_0 > 1$. For all sufficiently large L, the total moment is at most $M + \varepsilon$ and $\mathbb{E}F_{K_0}(r_L) \ge 1 - \varepsilon$. For every $K \ge K_0$ their tail moments are then at most $M - 1 + 2\varepsilon$. A larger cutoff kills the tails of the finitely many remaining finite functions. Taking the supremum over all L, then $K \to \infty$ and $\varepsilon \downarrow 0$, proves $\Theta = M - 1$. This finite-prefix step is required; the fixed-K limit alone would not justify exchanging the supremum and cutoff limit.

The diluted balanced family of Theorem 13 satisfies $r_L \to 1$ in probability: its tilt has every smaller positive moment tending to zero, while the balancing and dilution parameters tend to one and zero, respectively. Also $\lambda_2(T_{r_L}) \to 1$ and $2\lambda_2^s \le \mathbb{E}r_L^s \le 2$, so the critical moment tends to two. Hence $\Theta = 1$, while that family's gap tends to zero. Finally apply Proposition 15 with $M = 1 + \Theta_0$. Its densities converge to one in probability, have critical moment tending to M, and keep the positive gap there proved. Thus $\Theta = \Theta_0$. Its strict total-budget slack and finite-prefix replacement remain part of the construction. □

Theorem 17 (The fixed-cutoff endpoint)**.** *For every* $0 < b < 1$*, there is a single family of densities* r_L *with*

$$\mathbb{E}r_L = 1, \qquad r_L \ge b, \qquad \mathbb{E}r_L^s \le 2, \qquad \mathbb{E}[r_L^s \mathbf{1}_{r_L > 1}] \le 1$$

for every $L \ge 1$*, but* $\lambda_1(T_{r_L}), \lambda_2(T_{r_L}) \to 1$*. Moreover* $r_L \to 1$ *in probability,* $\mathbb{E}r_L^s \to 2$*, and* $\Theta = 1$*. Thus the finite-cutoff guarantee in Theorem 14 cannot include* $\tau = 1$ *for* $K \ge 1$*, even under the exact total moment cap and a prescribed positive floor. The same family satisfies the tail cap at every* $K \ge 1$ *simultaneously; it depends on* ρ *and* b*, but not on* K*.*

Proof. Use the undiluted product tilt R_L from the proof of Theorem 13. It satisfies

$$R_L > 0, \qquad \mathbb{E}R_L^s = 1, \qquad \mathbb{E}R_L^t \longrightarrow 0 \quad (0 < t < s).$$

Write $m_L = \mathbb{E}R_L \le 1$. The same proof supplies coefficients $\beta_L \to 1$ such that $w_L = (1 + \beta_L R_L)/Z_L$, with $Z_L = 1 + \beta_L m_L \to 1$, has both leading eigenvalues tending to one. Here β_L denotes the balancing coefficient previously called a_L; the new coefficient below is different. The product definition gives additional information at its own mean. For $\delta_L = dL^{-1/3}$ and the single-site normalizer $n(\delta_L)$, $m_L = n(\delta_L)^{-L} > 0$, so exactly

$$\frac{R_L}{m_L} = \prod_{i=1}^{L}(1 + \delta_L x_i), \qquad \mathbb{E}(R_L/m_L)^t = \left(\frac{(1+\delta_L)^t + (1-\delta_L)^t}{2}\right)^L.$$

Fix $0 < t < 1$. The logarithm of the bracket is an even smooth function on any fixed compact interval $|\delta| \le d < 1$. Taylor's theorem, with bounded fourth derivative there, yields

$$\log \frac{(1+\delta)^t + (1-\delta)^t}{2} = \frac{t(t-1)}{2}\delta^2 + O_{t,d}(\delta^4).$$

After multiplication by L, the leading term tends to minus infinity as $L^{1/3}$ and the remainder tends to zero. Markov's inequality gives $\mathbb{P}(R_L > m_L) \le \mathbb{E}(R_L/m_L)^t \to 0$. This relative-mean estimate does not follow from $R_L \to 0$ alone. Set

$$p_L = \mathbb{P}(R_L > m_L)^{1/s}, \quad a_L = 1 - p_L, \quad A_L = 1 - a_L m_L, \quad q_L = A_L + a_L R_L.$$

All three coefficients belong to $[0,1]$, $\mathbb{E}q_L = 1$, and $a_L, A_L \to 1$, whence $q_L \to 1$ in probability. The identity $q_L - 1 = a_L(R_L - m_L)$ shows that $E_L = \{q_L > 1\}$ equals $\{R_L > m_L\}$ when $a_L > 0$ and is empty when $a_L = 0$. In both cases the event inclusion holds. Minkowski gives

$$\|q_L \mathbf{1}_{E_L}\|_s \le A_L \mathbb{P}(E_L)^{1/s} + a_L \|R_L \mathbf{1}_{E_L}\|_s \le p_L + a_L = 1.$$

No division by a_L occurs; its zero endpoint gives $q_L = 1$.

Put $S(v) = QM_vQ$. Endpoint forward hypercontractivity maps L^2 to $L^{2s/(s-1)}$ with norm at most one. Hölder therefore gives $\|QM_hQ\| \le \|h\|_s$ for every real h. In particular, $\|QM_{R_L}Q\| \le 1$, and

$$\|S(q_L) - S(w_L)\| \le |A_L - Z_L^{-1}| + |a_L - \beta_L/Z_L| \longrightarrow 0.$$

The square matrices $S(v)$ and T_v are X^*X and XX^* for $X = M_{\sqrt{v}}Q$, so are isospectral with multiplicities. The min-max perturbation bound thus transfers both actual leading eigenvalue limits to T_{q_L}.

We next prove $M_L := \mathbb{E}q_L^s \to 2$. The lower bound is $M_L \ge A_L^s + a_L^s \mathbb{E}R_L^s \to 2$. For a fixed $T > 0$, $q_L^s \mathbf{1}_{R_L \le T}$ converges in probability to one and is bounded by $(1+T)^s$, so its expectation tends to one. On $R_L > T$, $q_L^s \le (1 + 1/T)^s R_L^s$. Consequently $\limsup_L M_L \le 1 + (1+1/T)^s$. Letting $T \to \infty$ proves the claim. This auxiliary truncation is used only for the moment limit.

Set $\varepsilon_L = 0$ when $M_L \le 2$, and otherwise set $\varepsilon_L = (M_L - 2)/(M_L - 1)$. Then $0 \le \varepsilon_L < 1$ and $\varepsilon_L \to 0$. The mean-one density $v_L = (1 - \varepsilon_L)q_L + \varepsilon_L$ satisfies, by convexity,

$$\mathbb{E}v_L^s \le (1 - \varepsilon_L)M_L + \varepsilon_L \le 2.$$

Since $\varepsilon_L < 1$, the events $v_L > 1$ and $q_L > 1$ coincide. On this event $v_L \le q_L$, so pointwise $v_L^s \mathbf{1}_{v_L > 1} \le q_L^s \mathbf{1}_{q_L > 1}$. The strict tail cap at one is preserved. Also $S(v_L) - S(q_L) = \varepsilon_L(P_\rho - S(q_L)) \to 0$ in operator norm, since $\|S(q_L)\| \to 1$. Both eigenvalue limits persist, as does convergence in probability to one. The two-sign estimate (16) gives $2\lambda_2(T_{v_L})^s \le \mathbb{E}v_L^s \le 2$, and hence the critical moment tends to two.

Since $v_L \ge A_L \to 1$, replace a finite initial segment by $r_L = 1$ and use $r_L = v_L$ thereafter to impose $r_L \ge b$ in every dimension. The prefix has empty strict tail above one, and all limits are unchanged. For every $K \ge 1$, the tail integrand is pointwise at most the one at cutoff one. Thus this same family works simultaneously for all such K. For each fixed $J > 1$, the bounded truncation $r_L^s \mathbf{1}_{r_L \le J}$ has expectation tending to one. Thus the J-tail tends to one, giving $H(J) \ge 1$. The all-dimension tail bound at one also gives $H(J) \le 1$ for every $J > 1$. Therefore $H(J) = 1$ for every $J > 1$ and $\Theta = 1$. □

Each finite matrix in this construction has a positive Perron gap; the gaps tend to zero with dimension. The two moment caps are inequalities in each dimension, with total moment equality asserted only in the limit. The mean-relative probability coefficient resolves the cutoff-one case without a dimension-dependent cutoff.

Corollary 18 (Cutoffs below one)**.** *For $0 \le K < 1$, a nonnegative mean-one density satisfying $\mathbb{E}[r^s \mathbf{1}_{r > K}] \le 1$ has*

$$\operatorname{gap}(T_r) \ge 1 - \left(\frac{1 + K^{s-1}}{2}\right)^{1/s} > 0.$$

Proof. Pointwise $r^s \mathbf{1}_{r \le K} \le K^{s-1} r$, including $K = 0$. Thus $\mathbb{E} r^s \le 1 + K^{s-1} < 2$. The two-sign bound and $\lambda_1 \ge 1$ give $\lambda_2/\lambda_1 \le \lambda_2 \le (\mathbb{E} r^s/2)^{1/s}$. □

No optimality is claimed for this lower constant or for a general tail-budget threshold below cutoff one. The displayed bound is uniform in dimension for fixed $K < 1$, not uniform as K approaches one.

These are worst-case statements about a specialized fixed Boolean congruence. Scalar tails under the base measure are not the operator-tail functionals under a Perron invariant measure in [14]. The cutoff bound refines the same sign-truncation proof; it is not an independent method or a claim of historical priority. The qualitative boundary is sharp, while the positive constant is not optimized.

12 Optimized finite backgrounds below cutoff one

Fix $0 < \rho < 1$ and, throughout this section, set

$$s = \frac{1+\rho}{1-\rho}, \qquad k_s = \frac{s-1}{s}, \qquad C_s = \frac{s^s}{(s-1)^{s-1}}.$$

The cutoff-one result does not classify a fixed cutoff $K < 1$ with a tail budget greater than one. We now combine long parities with the balanced critical tilt to give a sub-two collapse region. The sharp optimization below concerns a scalar background law, not the optimal spectral boundary among all densities.

Proposition 19 (Scalar background tail optimum)**.** *For $s > 1$ and $0 < K < 1$, among nonnegative random variables A with $\mathbb{E}A = 1$,*

$$\inf \mathbb{E}[A^s \mathbf{1}_{A>K}] = m_s(K), \qquad m_s(K) = \begin{cases} 1, & 0 < K \le k_s, \\ C_s(1-K)K^{s-1}, & k_s < K < 1. \end{cases}$$

For $K \le k_s$ the constant law $A = 1$ attains the infimum. For $K > k_s$ an attaining law has values K and $v_* = sK/(s-1)$, with high probability $p_* = (s-1)(1-K)/K$. Replacing the strict event by $A \ge K$ leaves the infimum unchanged, but that infimum is not attained when $K > k_s$.

Proof. For $K \le k_s$, convexity gives $y^s \ge 1 + s(y-1)$. If $y \le K$, the right side is nonpositive. Hence $y^s \mathbf{1}_{y>K} \ge 1 + s(y-1)$ everywhere. Taking expectations gives the lower bound one, attained by $A = 1$. The affine minorant is integrable even if the left side has infinite expectation.

For $K > k_s$, the derivative of $y^s/(y-K)$ on $y > K$ has the sign of $(s-1)y - sK$. Its unique minimum is at v_* and its value is $c = C_s K^{s-1}$. Thus

$$y^s \mathbf{1}_{y>K} \ge c(y-K) \qquad (y \ge 0).$$

Integration gives $c(1-K)$. The stated probabilities lie in $(0,1)$, have mean $K + p_*(v_* - K) = 1$, and put mass only at the two equality points of the minorant. At $K = k_s$ this law degenerates to $A = 1$.

For the non-strict event the same lower bound holds. When $K > k_s$, choose $u < K$ tending to K with $u > k_s$, high value $su/(s-1)$, and high probability $(s-1)(1-u)/u$. The high value exceeds K eventually, and the cost tends to $C_s(1-K)K^{s-1}$. Thus the infimum is unchanged. Equality in the affine minorant for the non-strict integrand is possible only at v_*: below K its right side is strictly negative, and at K its left side is $K^s > 0$. Concentration at $v_* > 1$ cannot have mean one, proving nonattainment. For $K \le k_s$, $A = 1$ still attains one. □

This is an elementary instance of classical moment optimization by pointwise affine bounds; see Bertsimas and Popescu [1] for the general moment-duality setting. Neither that method nor the formula's external priority is claimed as new here.

Theorem 20 (Sub-two tail budgets with finite parity backgrounds)**.** *Fix $k_s < K < 1$, $0 < b < K$, and*

$$1 + C_s(1-K)K^{s-1} < \tau < 2.$$

There are strictly positive densities r_L on every Boolean cube of integer dimension $L \geq 1$ satisfying

$$\mathbb{E} r_L = 1, \qquad r_L \geq b, \qquad \mathbb{E}[r_L^s \mathbf{1}_{r_L > K}] \leq \tau,$$

such that both leading eigenvalues of T_{r_L} tend to one and $\mathrm{gap}(r_L) \to 0$. The densities use global interactions. No marginal matching, locality, equality-budget attainability or optimal universal spectral threshold is asserted.

Proof. A fixed dyadic background. Write $g(u) = C_s(1-u)u^{s-1}$. On $(k_s, 1)$,

$$g'(u) = C_s u^{s-2}((s-1) - su) < 0,$$

and its endpoint values are one and zero. Strict budget slack permits a fixed $\max(b, k_s) < u < K$ with $1 + g(u) < \tau$. The optimal high probability at lower value u is $p_u = (s-1)(1-u)/u$. For $p \in (0,1)$ define $v(p) = u + (1-u)/p$. The mean $(1-p)u + pv(p)$ is exactly one and $pv(p)^s$ is continuous in p. Choose a fixed dyadic $p = j/2^k$ sufficiently close to p_u that

$$1 + pv^s < \tau, \qquad v = v(p) > 1 > K, \qquad k \geq 1, \quad 1 \leq j < 2^k.$$

On k unbiased signs choose a function F equal to v at j patterns and equal to u elsewhere. Then $\mathbb{E}F = 1$. Replace each sign by the parity of a disjoint block of M unbiased bits, obtaining a_M on kM bits. Those parities are independent unbiased signs, so a_M has exactly the same two-point law. All u, v, p, k are fixed before the dimension grows.

An operator-norm estimate for every character. Let $Q_d = P_{\sqrt{\rho},d}$ and let χ_U be a Walsh character. For any Walsh basis vector χ_S,

$$Q_d M_{\chi_U} Q_d \chi_S = \rho^{(|S| + |S \triangle U|)/2} \chi_{S \triangle U} = \rho^{|U|/2 + |S \setminus U|} \chi_{S \triangle U}.$$

This is a weighted permutation of an orthonormal basis. Its norm is exactly $\rho^{|U|/2}$, attained for $S \subset U$. Writing the finite Walsh expansion $F = 1 + \sum_{I \neq \varnothing} \widehat{F}(I)\chi_I$ and putting $D_F = \sum_{I \neq \varnothing} |\widehat{F}(I)|$, block substitution gives

$$\|Q_{kM} M_{a_M} Q_{kM} - P_{\rho, kM}\| \leq \sum_{I \neq \varnothing} |\widehat{F}(I)| \rho^{M|I|/2} \leq D_F \rho^{M/2} \longrightarrow 0. \tag{20}$$

Large fixed amplitudes or coefficients do not obstruct this limit.

The inherited critical family and actual eigenvalues. Use the balanced critical family in Theorem 13: with fixed small $d > 0$, $\delta_n = dn^{-1/3}$ and

$$N(\delta) = \left(\frac{(1+\delta)^s + (1-\delta)^s}{2} \right)^{1/s}, \qquad R_n(y) = \prod_{i=1}^{n} \frac{1 + \delta_n y_i}{N(\delta_n)},$$

put

$$\kappa(\delta) = \frac{1 + \rho + \sqrt{(1-\rho)^2 + 4\rho\delta^2}}{2N(\delta)}, \quad \beta_n = \kappa(\delta_n)^{-n}, \quad Z_n = 1 + \beta_n \mathbb{E} R_n, \quad w_n = \frac{1 + \beta_n R_n}{Z_n}.$$

The established balanced Gram estimates give the actual limits $\lambda_1(T_{w_n}), \lambda_2(T_{w_n}) \to 1$, not merely two trial-vector bounds. Also $\beta_n, Z_n \to 1$, $\mathbb{E}R_n^s = 1$, and $\mathbb{E}R_n^t \to 0$ for every fixed $0 < t < s$. The latter follows from the one-site expansion

$$\log \mathbb{E}\left[\frac{1+\delta Y}{N(\delta)}\right]^t = \frac{t(t-s)}{2}\delta^2 + O(\delta^4),$$

since $n\delta_n^2 \to \infty$ and $n\delta_n^4 \to 0$.

For sufficiently large L take $M = \lfloor L/(2k) \rfloor$ and $n = L - kM$. Both tend to infinity. On the two independent blocks set

$$r_L(x,y) = \frac{a_M(x) + \beta_n R_n(y)}{Z_n}.$$

Its mean is one and its minimum is at least u/Z_n, eventually greater than b. The matrices $S_r = Q_L M_r Q_L$ and T_r have the same full spectrum by the AA^* and A^*A identity with $A = M_{\sqrt{r}} Q_L$. Linearity and tensorization show

$$\|S_{r_L} - P_{\rho,kM} \otimes S_{w_n}\| \le D_F \rho^{M/2}/Z_n \longrightarrow 0.$$

The tensor comparison has largest eigenvalue $\lambda_1(S_{w_n})$ and second eigenvalue between $\lambda_2(S_{w_n})$ and $\lambda_1(S_{w_n})$: use the constant background eigenvector and the first two tilt eigenvectors for the lower bound. Positivity and the fact that the top background eigenvalue is one give the upper bound. Both leading eigenvalues therefore tend to one. The self-adjoint eigenvalue perturbation inequality transfers both limits to S_{r_L} and hence T_{r_L}. Each finite matrix remains strictly positive, with a simple Perron root.

Noninteger moments and the moving strict cutoff. Let $\theta = \min(1, s/2)$. For $c > 0, z \ge 0$,

$$0 \le (c+z)^s - c^s - z^s \le s2^{s-1}\big(c^{s-\theta}z^\theta + c^\theta z^{s-\theta}\big). \tag{21}$$

To prove this, divide by c^s and put $y = z/c$. For $y \le 1$, the mean value theorem bounds $(1+y)^s - 1 - y^s$ by $s2^{s-1}y \le s2^{s-1}y^\theta$. For $y \ge 1$ interchange c and z. Nonnegativity follows from convexity. Both powers $\theta, s-\theta$ lie strictly between zero and s. For $z = \beta_n R_n$, the mixed expectations vanish and $\mathbb{E}z^s \to 1$. Thus $\mathbb{E}(c + \beta_n R_n)^s \to c^s + 1$.

In the high sector, $v/Z_n > K$ eventually, so the entire sector contributes and its conditional tail moment tends to $v^s + 1$. In the low sector the event is exactly $R_n > c_n$, where $c_n = (KZ_n - u)/\beta_n \to K - u > 0$. Fix positive finite c_-, c_+ containing these thresholds eventually. For $0 < t < s$,

$$\mathbb{P}(R_n > c_n) \le c_-^{-t}\mathbb{E}R_n^t \to 0, \qquad \mathbb{E}[R_n^s \mathbf{1}_{R_n \le c_n}] \le c_+^{s-t}\mathbb{E}R_n^t \to 0.$$

Consequently the strict event retains asymptotic s-moment one. In (32) restricted to this event, the constant term vanishes with its probability, and the mixed terms are bounded by their vanishing unrestricted expectations. After normalization, the low-sector tail moment tends to one. Averaging the fixed sector probabilities,

$$\mathbb{E}[r_L^s \mathbf{1}_{r_L > K}] \longrightarrow (1-p) + p(v^s + 1) = 1 + pv^s < \tau.$$

This calculation does not infer moment convergence from convergence in probability. Indeed the full moment tends to $1 + (1-p)u^s + pv^s > 2$, by strict Jensen for the nonconstant background.

Choose a common L_0 beyond which both blocks are positive, the floor holds and the tail cost is at most τ. For $1 \le L < L_0$ set $r_L = 1$. Its tail cost is one and its floor is valid since $b < K < 1$. This finite replacement preserves all limits and proves the assertion in every dimension. No effective value of L_0 is claimed. □

Comparison of sufficient and exact spectral thresholds

The half-parity background with low value approaching K and high value approaching $2-K$ has limiting scalar cost $(2-K)^s/2$. It gives collapse whenever

$$2-2^{1/s}<K<1, \qquad 0<b<K, \qquad 1+(2-K)^s/2<\tau<2.$$

This is a specialization of the same proof with $p=1/2$ and $u<K$ chosen with strict slack. Proposition 19 shows that the optimized cost is no larger. Equality occurs only when $2-K=sK/(s-1)$, that is $K=2(s-1)/(2s-1)$; outside this isolated value the cost improvement is strict. Strict convexity of 2^t below its chord on $0<t<1$ gives $2^{1/s}<1+1/s$. Hence the new interval $k_s<K<1$ strictly extends the half-parity sub-two interval.

On the other hand, any mean-one density with tail budget at most τ obeys $\mathbb{E}r^s\le K^{s-1}+\tau$, since $r^s\mathbf{1}_{r\le K}\le K^{s-1}r\mathbf{1}_{r\le K}$. The established total-moment estimate therefore gives

$$\operatorname{gap}(r)\ge 1-\left(\frac{K^{s-1}+\tau}{2}\right)^{1/s}>0 \quad \text{if } \tau<2-K^{s-1}.$$

This sufficient estimate alone does not identify the universal spectral threshold. The construction justifies adding an escaped unit to the scalar optimum only for that particular family. Theorem 21 attains this budget with floor $b=K$, while the separate necessary estimate in Theorem 23 proves its universal optimality. For $K<b<1$ the event contains every point and reduces to the total-moment condition. At $b=1$ mean one forces $r=1$ and gap $1-\rho$; at $b>1$ the class is empty. The earlier source309A omitted $b<1$ in its additional comparison; its failed gate and the corrected309B report remain preserved, not relabelled.

13 Attaining the constructive cutoff budget

The preceding fixed-background argument uses strict slack in both the floor and the tail budget. We now remove both restrictions. The cutoff event remains strictly $r>K$; mass exactly at K is essential.

Theorem 21. *Fix $0<\rho<1$, $s=(1+\rho)/(1-\rho)$, and*

$$k_s=\frac{s-1}{s}<K<1, \qquad g(K)=\frac{s^s}{(s-1)^{s-1}}(1-K)K^{s-1}.$$

For every integer $L\ge 1$ there is a strictly positive density r_L on the uniform Boolean cube such that

$$\mathbb{E}r_L=1, \qquad r_L\ge K, \qquad \mathbb{E}[r_L^s\mathbf{1}_{r_L>K}]\le 1+g(K),$$

and $\lambda_1(T_{r_L})\to 1$, $\lambda_2(T_{r_L})\to 1$. Consequently $\operatorname{gap}(T_{r_L})\to 0$. The same family meets every prescribed floor $0<b\le K$ and every budget $\tau\ge 1+g(K)$. This is an attained constructive upper budget, not a claim that $1+g(K)$ is the optimal threshold over all collapsing families.

Proof. All expectations below use normalized uniform probability. Write $Q_d=P_{\sqrt{\rho},d}$. For positive r, the full spectra of $T_r=M_{\sqrt{r}}P_{\rho,d}M_{\sqrt{r}}$ and $S_r=Q_dM_rQ_d$ agree, by the AA^*/A^*A identity with $A=M_{\sqrt{r}}Q_d$.

The inherited balanced tilt. Use the explicit family in the proof of Theorem 13. Fix a sufficiently small $d>0$ and set

$$\delta_n=dn^{-1/3}, \qquad N(\delta)=\left(\frac{(1+\delta)^s+(1-\delta)^s}{2}\right)^{1/s}, \qquad R_n(y)=\prod_{i=1}^n\left(\frac{1+\delta_n y_i}{N(\delta_n)}\right).$$

In particular, the normalization is applied to every product factor. Define

$$H_n = Q_n M_{R_n} Q_n, \qquad \kappa(\delta) = \frac{1+\rho+\sqrt{(1-\rho)^2+4\rho\delta^2}}{2N(\delta)}, \qquad b_n = \kappa(\delta_n)^{-n}, \qquad F_n = P_{\rho,n} + b_n H_n.$$

The earlier proof gives

$$b_n \to 1, \quad \|H_n\| = \kappa(\delta_n)^n \to 1, \quad \lambda_1(F_n) \to 1, \quad \lambda_2(F_n) \to 1, \quad \mathbb{E}R_n^s = 1. \tag{22}$$

These are conclusions for that explicit family, not additional hypotheses. For completeness, its interior moments have the exact expression

$$a_n(t) := \mathbb{E}R_n^t = \left(\frac{(1+\delta_n)^t + (1-\delta_n)^t}{2N(\delta_n)^t}\right)^n \longrightarrow 0 \qquad (0 < t < s).$$

Indeed their logarithm is $n[t(t-s)\delta_n^2/2 + O(\delta_n^4)]$. At the critical exponent, $\log \kappa(\delta) = O(\delta^4)$. The block Gram comparison proved earlier has two diagonal top eigenvalues one and an off-diagonal norm tending to zero; it controls the actual two leading eigenvalues of F_n. Its second lower bound is eventually positive, so it is not an added zero of the Gram matrix.

Growing finite background laws. The scalar optimizer from Proposition 19 has

$$v_* = \frac{sK}{s-1} > 1, \qquad p_* = \frac{(s-1)(1-K)}{K} \in (0,1), \qquad p_* v_*^s = g(K).$$

For sufficiently large L choose

$$k = \lfloor L^{1/3} \rfloor, \quad M = \left\lfloor \frac{L}{2k} \right\rfloor, \quad n = L - kM, \qquad p = p_L = \frac{\lfloor 2^k p_* \rfloor}{2^k}, \qquad v = v_L = K + \frac{1-K}{p}.$$

Then $p \in [p_*/2, p_*]$, $p \to p_*$, $v \to v_*$ and $v \le V := 2v_*$. Also $k, M, n \to \infty$, $n \ge L/2$ and $M/k \to \infty$. Choose exactly $2^k p$ patterns in $\{-1,1\}^k$, and let f_L take value v there and K elsewhere. Replace its k signs by the parities of k disjoint blocks of M bits. Call the resulting background A_L, with high-sector indicator I_L. Independent unbiased parities give

$$\mathbb{E}I_L = p, \qquad A_L = K + (v-K)I_L, \qquad \mathbb{E}A_L = 1.$$

For a nonempty Walsh support U,

$$Q_d M_{\chi_U} Q_d \chi_S = \rho^{|U|/2 + |S\setminus U|} \chi_{S \triangle U}.$$

Thus this weighted permutation has norm exactly $\rho^{|U|/2}$. Each Walsh coefficient of f_L has absolute value at most $\mathbb{E}f_L = 1$, and there are at most $2^k - 1$ nonconstant ones. Every substituted support has size at least M, so

$$\|Q_{kM} M_{A_L} Q_{kM} - P_{\rho,kM}\| \le (2^k - 1)\rho^{M/2} \longrightarrow 0. \tag{23}$$

The limit follows from $k \log 2 + (M/2)\log \rho \to -\infty$. In particular, the growing number of coefficients is explicitly controlled. The scalar affine minorant gives

$$d_L := p_L v_L^s - g(K) \ge 0, \qquad d_L \to 0. \tag{24}$$

No uniform finite cube is assumed to realize a non-dyadic p_* exactly.

Clipping, exact mean, and the amplitude reserve. Put $\theta = \min(1, s/2)$, so both θ and $s - \theta$ lie in $(0, s)$. For $n = n_L$ define

$$c_n = a_n(\theta)^{1/(2\theta)}, \qquad U_n = (R_n - c_n)_+, \qquad m'_n = \mathbb{E}U_n.$$

Then $c_n \to 0$, $m'_n \le a_n(1) \to 0$, and

$$\mathbb{P}(U_n > 0) \le \sqrt{a_n(\theta)}, \qquad \|Q_n M_{U_n - R_n} Q_n\| \le c_n. \tag{25}$$

The first inequality is Markov's inequality at c_n; the second uses $0 \le R_n - U_n \le c_n$ and $\|Q_n\| = 1$. Equality $R_n = c_n$ is in the complement of the strict event.

For real $s > 1$, $c > 0$ and $z \ge 0$, the mixed-power bound is

$$0 \le (c + z)^s - c^s - z^s \le B_s(c^{s-\theta} z^\theta + c^\theta z^{s-\theta}), \qquad B_s = s2^{s-1}. \tag{26}$$

To see the upper bound, divide by c^s. For $z/c \le 1$ the mean value theorem bounds the difference by $B_s(z/c)$, which is at most $B_s(z/c)^\theta$. For $z/c \ge 1$ exchange c, z. The lower bound follows by integrating a positive derivative. Thus this estimate does not require an integer exponent. Set

$$\begin{aligned} E_L &= K^s \sqrt{a_n(\theta)} \\ &\quad + B_s\big[(K^{s-\theta} + V^{s-\theta}) a_n(\theta) + (K^\theta + V^\theta) a_n(s - \theta)\big], \\ e_L &= d_L + E_L. \end{aligned}$$

Both are nonnegative and tend to zero. Eventually $e_L < 1/2$ and $m'_n < 1 - K$. On that range let

$$\beta_L = (1 - e_L)^{1/s}, \qquad \gamma_L = \frac{\beta_L m'_n}{p_L}, \qquad r_L(x, y) = A_L(x) - \gamma_L I_L(x) + \beta_L U_n(y). \tag{27}$$

Here $0 < \beta_L \le 1$, $\beta_L \to 1$ and $\gamma_L \to 0$. The mean is exactly $1 - p_L\gamma_L + \beta_L m'_n = 1$. The low sector has density $K + \beta_L U_n \ge K$; on the high sector,

$$v_L - \gamma_L = K + \frac{1 - K - \beta_L m'_n}{p_L} > K.$$

Thus mean compensation has not lowered the floor. It acts only on the high sector, and no global normalization is applied.

On the low sector the strict event is precisely $U_n > 0$. Apply (26) with $c = K$ and $z = \beta_L U_n$, restricted to that event. Its constant contribution is at most $K^s \mathbb{P}(U_n > 0)$. The high sector is entirely above K; discarding the nonnegative subtraction γ_L gives an upper bound from (26) with $c = v_L$. Averaging these bounds with probabilities $1 - p_L, p_L$, the coefficient of $\beta_L^s \mathbb{E}U_n^s$ is exactly one. Use $U_n \le R_n$, $v_L \le V$, $\beta_L \le 1$, $\mathbb{E}R_n^s = 1$, and (25). The two mixed terms are bounded by those defining E_L. Hence

$$\mathbb{E}[r_L^s \mathbf{1}_{r_L > K}] \le p_L v_L^s + \beta_L^s + E_L = g(K) + d_L + 1 - e_L + E_L = 1 + g(K). \tag{28}$$

The reserve pays for the dyadic excess and every retained tail error. No convergence of s-moments was inferred from convergence in probability.

The actual leading eigenvalues. Write $P_b = P_{\rho,kM}$, $P_t = P_{\rho,n}$ and similarly Q_b, Q_t. By linearity and tensorization,

$$S_{r_L} = (Q_b M_{A_L} Q_b - \gamma_L Q_b M_{I_L} Q_b) \otimes P_t + \beta_L P_b \otimes (Q_t M_{U_n} Q_t).$$

Compare with $B_L = P_b \otimes F_n$. Since $0 \le I_L \le 1$, (23) and (25) imply

$$\|S_{r_L} - B_L\| \le (2^k - 1)\rho^{M/2} + \gamma_L + \beta_L c_n + |\beta_L - b_n|\|H_n\| \longrightarrow 0.$$

All four errors vanish by the established limits. Both tensor factors of B_L are positive definite and $\lambda_1(P_b) = 1$. Therefore

$$\lambda_1(B_L) = \lambda_1(F_n), \qquad \lambda_2(F_n) \le \lambda_2(B_L) \le \lambda_1(F_n).$$

The lower bound uses two actual product eigenvectors with constant background factor. In view of (22), both leading eigenvalues of B_L tend to one. Weyl's self-adjoint perturbation bound transfers both limits to S_{r_L}, and the full-spectrum identity transfers them to T_{r_L}. This controls the top eigenvalue as well as the second; two trial vectors alone would not suffice.

Choose one L_0 so that every eventual condition above holds for all $L \ge L_0$. For $1 \le L < L_0$ set $r_L = 1$. Its floor is at least K, its mean is one, and its strict K-tail cost is $1 \le 1 + g(K)$. This supplies every dimension without changing the limits. No effective estimate of L_0 is asserted. □

For the non-strict event $r_L \ge K$, every point contributes and the tail becomes the total moment. The theorem therefore cannot be relabeled with a non-strict event. The scalar optimization, clipping and perturbation tools are classical; this construction and its source audit do not establish external priority. The next two sections improve the sufficient estimate and prove that the constructed upper budget is the exact universal strict-tail threshold. No marginal matching, locality, heat-bath gap or physical conclusion is added to the separate families proved earlier.

14 The exact scalar envelope at a strict cutoff

This section sharpens the elementary sufficient estimate above. It is useful for explicit bounds, but scalar optimality alone does not settle the spectral question. Section 15 supplies that step. Throughout these two sections fix $0 < \rho < 1$, $s = (1+\rho)/(1-\rho)$, $k_s = (s-1)/s < K < 1$, and put

$$a = 1 - K, \quad p_* = (s-1)a/K, \quad C_s = \frac{s^s}{(s-1)^{s-1}}, \quad g = C_s a K^{s-1}, \quad h(p) = \frac{(a+Kp)^s}{p^{s-1}}.$$

Then $0 < p_* < 1$ and $0 < g < 1$. For $\tau \ge g$ let p_τ be the unique solution of $h(p_\tau) = \tau$ in $(0, p_*]$, and set $U(\tau) = \tau + K^s(1 - p_\tau)$.

Theorem 22. *Among nonnegative random variables on arbitrary probability spaces with $\mathbb{E}r = 1$ and $\mathbb{E}[r^s \mathbf{1}_{r>K}] \le \tau$, the class is empty for $\tau < g$. For $\tau \ge g$ the maximum of $\mathbb{E}r^s$ is $U(\tau)$, attained by the law with value K of probability $1 - p_\tau$ and value $K + a/p_\tau$ of probability p_τ. At $\tau = g$ this is the unique admissible law; a uniform L-cube realizes it precisely when $2^L p_*$ is integral. For $\tau > g$, the supremum over all finite uniform Boolean cubes is $U(\tau)$, approached by positive mean-one densities with floor K and strict tail at most τ.*

Proof. Logarithmic differentiation gives

$$\frac{h'(p)}{h(p)} = \frac{Kp - (s-1)a}{p(a+Kp)}.$$

Thus h decreases from infinity to $h(p_*) = g$ on $(0, p_*]$, then increases to $h(1) = 1$. For an admissible r, write $p = \mathbb{P}(r > K) > 0$ and $m = \mathbb{E}[r\mathbf{1}_{r>K}]$. The complementary mean is at most $K(1-p)$, so $m \ge a + Kp$. Conditional Jensen yields

$$\tau \ge \mathbb{E}[r^s \mathbf{1}_{r>K}] \ge m^s / p^{s-1} \ge h(p) \ge g.$$

It also forces $p \ge p_\tau$. Consequently

$$\mathbb{E}r^s \le \tau + K^s(1-p) \le U(\tau).$$

The stated two-point law has mean one, strict tail $h(p_\tau) = \tau$ and total moment $U(\tau)$, proving attainment. If $\tau = g$, every inequality above is equality: $p = p_*$, the lower value is almost surely K, and strict convexity forces a constant upper value. This proves uniqueness and the finite-cube integrality condition. For $\tau > g$, choose $p_L = (\lfloor 2^L p_\tau \rfloor + 1)/2^L$. Eventually $p_\tau < p_L < p_*$ and $p_L \to p_\tau$. Assign value $K + a/p_L$ to exactly $2^L p_L$ vertices and K to the others. Its mean is one and its strict tail $h(p_L) < \tau$, whereas its total moment tends to $U(\tau)$. Exact attainment on a specified cube requires the high probability p_τ to be realizable there; no irrational probability is assigned to a finite uniform event. □

The earlier sign compression, $2\lambda_2^s \le \mathbb{E}r^s$ and $\lambda_1 \ge 1$, now gives

$$\operatorname{gap}(T_\tau) \ge 1 - [U(\tau)/2]^{1/s} > 0 \qquad (U(\tau) < 2).$$

There is a unique $q_0 \in (0, p_*)$ with $h(q_0) + K^s(1-q_0) = 2$, since this expression decreases strictly from infinity to a value below two. Its tail budget $\tau_c = 2 - K^s(1-q_0)$ satisfies $2 - K^{s-1} < 2 - K^s < \tau_c < 2$; $U(\tau) < 2$ exactly when $g \le \tau < \tau_c$. For $s = 2$ direct expansion gives

$$q_0 = \frac{(1-K)^2}{2 - 2K + K^2}, \qquad \tau_c = 2 - \frac{K^2}{2 - 2K + K^2}.$$

In particular $K = 3/4$ gives $\tau_c = 25/17$, improving the elementary $5/4$. This is only a sufficient spectral bound: a scalar extremizer with total moment two need not saturate the two eigenvalue inequalities. The moment optimization is an application of classical conditional Jensen and extremal-law reasoning; see [1] for the general method, without an assertion of priority for this formula.

15 The universal strict-tail threshold

The missing link is a necessary escaping-moment estimate for an arbitrary collapsing family. Reverse hypercontractivity forces one normalized sign part of a second eigenvector to have a small low quasi-norm. Truncation then forces its energy into the high-density region. This proves the additional unit of tail cost, without assuming the constructed density floor or an a priori decomposition into background and tilt.

Theorem 23. *Fix the parameters above. For every $g \le \tau < 1 + g$ there is an explicit constant $\gamma_0 > 0$, depending only on ρ, K, τ, such that every nonnegative mean-one density on every finite uniform Boolean cube with $\mathbb{E}[r^s \mathbf{1}_{r>K}] \le \tau$ satisfies $\operatorname{gap}(T_r) > \gamma_0$. No positive lower bound on r is required. For $\tau < g$ the class is empty. At $\tau = 1 + g$, and hence at every larger budget, Theorem 21 gives a positive family with floor K and gap tending to zero. Thus $1 + g$ is the exact universal strict-tail collapse threshold for $k_s < K < 1$.*

Proof. Write $B = \tau + K^s$, $p = 1 - \rho$, $q = 1 + \rho$ and $\theta = 1/s^2$. Classical forward and reverse Boolean hypercontractivity [8, Theorem 3.1 and Corollary 3.3] give

$$\langle F, P_\rho F \rangle = \|P_{\sqrt{\rho}} F\|_2^2 \le \|F\|_q^2, \tag{29}$$

$$\langle F, P_\rho G \rangle \ge \|F\|_p \|G\|_p \quad (F, G \ge 0). \tag{30}$$

The reverse inequality has parameter condition $\rho \le \sqrt{(1-p)(1-p)}$, with equality here, and permits zeros. The quantities at exponent $p < 1$ are quasi-norms; no triangle inequality for them will be used. Hölder with $q/(2-q) = s$ gives

$$\|\sqrt{z}\, h\|_q^2 \le (\mathbb{E} z^s)^{1/s} \|h\|_2^2 \quad (z \ge 0).$$

Since $\mathbb{E} r^s \le B$, this implies $\lambda_1 \le B^{1/s}$. The unit vector $\sqrt{r}$ has Rayleigh quotient $\langle r, P_\rho r\rangle \ge (\mathbb{E} r)^2 = 1$, so $\lambda_1 \ge 1$.

Put $\gamma = 1 - \lambda_2/\lambda_1$. If $\lambda_2 = 0$, the gap is one and the assertion is immediate. Otherwise work on the active set $\{r > 0\}$. The restricted matrix has strictly positive entries and a strictly positive Perron vector. A normalized real second eigenvector vanishes off this set, is orthogonal to the Perron vector, and therefore has both signs. Write it as $v = f_+ - f_-$, with disjoint nonnegative parts, $a = \|f_+\|_2 > 0$, $b = \|f_-\|_2 > 0$, $a^2 + b^2 = 1$, and $c = \langle f_+, T_r f_-\rangle > 0$. The eigenvector equation implies

$$\langle f_+, T_r f_+\rangle = \lambda_2 a^2 + c, \qquad \langle f_-, T_r f_-\rangle = \lambda_2 b^2 + c.$$

The compression to the orthonormal pair $u = f_+/a, w = f_-/b$ is

$$\lambda_2 I + c \begin{pmatrix} a^{-2} & (ab)^{-1} \\ (ab)^{-1} & b^{-2} \end{pmatrix}.$$

Its upper eigenvalue is $\lambda_2 + c/(a^2 b^2) \le \lambda_1$. Thus the *normalized* cross term obeys

$$d = \langle u, T_r w\rangle = c/(ab) \le (\lambda_1 - \lambda_2) ab \le B^{1/s}\gamma/2.$$

Both sign Rayleigh quotients are at least λ_2. Apply (30) to $\sqrt{r}\, u$ and $\sqrt{r}\, w$. Choose the sign vector h with the smaller p quasi-norm after multiplication by $\sqrt{r}$. For $F = \sqrt{r}\, h$ we obtain

$$\|F\|_p \le (B^{1/s}\gamma/2)^{1/2}, \qquad \langle F, P_\rho F\rangle \ge \lambda_2 \ge 1 - \gamma. \tag{31}$$

The argument tolerates arbitrarily unbalanced sign norms a, b.

Fix $M > K$, let $F_0 = F\mathbf{1}_{r \le M}$, $F_1 = F\mathbf{1}_{r > M}$, and $H_M = \mathbb{E}[r^s \mathbf{1}_{r > M}]$. One checks $1/q = \theta/p + (1-\theta)/2$. Ordinary Hölder applied to $Z^q = (Z^p)^{\theta q/p}(Z^2)^{(1-\theta)q/2}$ proves

$$\|Z\|_q \le \|Z\|_p^\theta \|Z\|_2^{1-\theta}.$$

The two powers are positive and sum to one, so this use of Hölder has exponents greater than one even though $p < 1$. Since $\|F_0\|_2 \le \sqrt{M}$ and $\|F_0\|_p \le \|F\|_p$, (29) and (31) yield

$$E_0 := \langle F_0, P_\rho F_0\rangle \le M^{1-\theta}(B^{1/s}\gamma/2)^\theta.$$

The same forward estimate and Hölder yield $E_1 := \langle F_1, P_\rho F_1\rangle \le H_M^{1/s}$. The Hilbert-space triangle after applying $P_{\sqrt{\rho}}$ gives $\sqrt{\langle F, P_\rho F\rangle} \le \sqrt{E_0} + \sqrt{E_1}$. Consequently

$$H_M \ge \left[\sqrt{1-\gamma} - M^{(1-\theta)/2}(B^{1/s}\gamma/2)^{\theta/2}\right]_+^{2s}. \tag{32}$$

For a collapsing family with fixed B, this implies $\liminf H_M \ge 1$ for each fixed M.

Put $A = C_s K^{s-1}$, so $g = A(1-K)$. The pointwise scalar inequality

$$z^s \mathbf{1}_{z > K} \ge A(z - K) \qquad (z \ge 0)$$

holds because its right side is nonpositive below K, whereas $z^s/(z-K)$ has its global minimum A at $z = sK/(s-1)$ above K. Apply it to $z = r\mathbf{1}_{r\le M}$ on the original probability space. Since $\mathbb{E}[r\mathbf{1}_{r>M}] \le B/M^{s-1}$, the middle strict tail costs at least $g - AB/M^{s-1}$. It is disjoint from the high tail, so

$$\tau \ge H_M + g - AB/M^{s-1}. \tag{33}$$

For any collapsing family, first fix M, use (32), and then send $M \to \infty$ in (33). This gives $\tau \ge 1 + g$. No density floor or exchange of moment limits is involved.

For an explicit positive bound, let $\varepsilon = 1 + g - \tau \in (0, 1]$ and

$$\begin{aligned} M &= \max\{K+1, (4AB/\varepsilon)^{1/(s-1)}\}, \\ D &= 1 - (1 - 3\varepsilon/4)^{1/(2s)}, \\ \gamma_0 &= \min\{D/4,\ 2B^{-1/s}(D/4)^{2/\theta}M^{-(1-\theta)/\theta}\}. \end{aligned}$$

These constants are finite and positive, with $\gamma_0 < 1$. If $\gamma \le \gamma_0$, the subtracted core term in (32) is at most $D/4$ and $\sqrt{1-\gamma} \ge 1 - \gamma \ge 1 - D/4$. Hence

$$H_M \ge (1 - D/2)^{2s} > (1 - D)^{2s} = 1 - 3\varepsilon/4.$$

But (33) gives $H_M \le \tau - g + \varepsilon/4 = 1 - 3\varepsilon/4$, a contradiction. Thus $\gamma > \gamma_0$. The empty regime $\tau < g$ was proved in Theorem 22; equality $\tau = g$ is included whenever the finite-cube class exists. Finally Theorem 21 uses exactly this ρ, s, K and strict event, and supplies collapse at $1 + g$ with floor K. This completes the sharpness assertion. □

For $s = 2$, $\rho = 1/3$ and $K = 3/4$, the exact threshold is $1 + 4(1/4)(3/4) = 7/4$. The intermediate scalar bound $25/17$ remains valid and can be numerically much simpler than the conservative explicit constant above. Neither scalar optimality nor the construction alone proved the universal threshold: the necessary escape estimate (32) closes that gap. External priority of this specific synthesis is unestablished. Its inputs are classical; related transmission of log-Sobolev and spectral-gap inequalities for perturbed diffusion measures is studied, for example, in [9]. Those generator/measure results are not silently identified with the finite transfer-matrix statement here. This theorem does not add marginal matching, finite-range interactions or a physical interpretation.

16 Comparison with hyperbounded and perturbation results

The value s_c is the classical Boolean hypercontractive boundary, not a new integrability scale. Indeed, for $a_s = 2s/(s+1)$, the norm-one estimate $\|P_{\sqrt{\rho}}f\|_2 \le \|f\|_{a_s}$ holds when $s \ge s_c$ by the standard inequality recalled in [10]. Its local sharpness is seen from $f = 1 + ux$ on one bit: $\|P_{\sqrt{\rho}}f\|_2^2 = 1 + \rho u^2$, whereas $\|f\|_{a_s}^2 = 1 + (a_s - 1)u^2 + o(u^2)$. Tensorization amplifies a violated one-bit norm bound. Such top-norm growth alone does not imply a closing normalized spectral gap; Theorem 10 adds two antipodal channels and their Gram comparison to obtain a near-degenerate second eigenvalue.

The inequalities and matrix arguments used above are classical. Related literature already supplies both qualitative gap criteria and, under different hypotheses, dimension-independent quantitative perturbation bounds. Miclo's Theorem 1 in [13] gives a gap for ergodic self-adjoint hyperbounded Markov operators. Section 6 and Proposition 11 also discuss a norm-budget transition at two for the fourth power of an L^2-to-L^4 Markov norm; this is an important precedent for budget thresholds, not the raw density norm used here. Hyperboundedness alone does not give a general quantitative gap. Wang's Theorem 1.1 in [14] gives a tail-norm criterion for the gap of an invariant

Markov operator's symmetrization. These fixed-space conclusions must not be confused with a bound uniform over a family of increasing dimensions.

Here T_r is generally not Markov in u_L. Its Perron Doob transform is Markov in its own stationary law, which generally differs from ru_L. The ground-state multiplication is an L^2 isometry between these spaces, but does not automatically identify their higher norms uniformly in L. Thus an implication from a qualitative hyperboundedness theorem would still need a quantitative change-of-measure argument.

Gross's Theorem 2.2 in [11] does give dimension-independent gaps for additive Schrödinger perturbations on Riemannian probability manifolds with a base logarithmic Sobolev inequality and suitable moments of both e^V and e^{-V}. We have not identified that additive operator with the logarithm of this discrete transfer congruence. Our moment hypothesis in Theorem 9 supplies no inverse-density moment and permits zeros. In contrast, the counterexamples in Theorems 10 and 12 have $r \geq 1/B$ and therefore bounded inverse moments. For these families, missing inverse control is not a distinction from Gross; the operator/form and exponent correspondences are the unestablished steps. Glück [12], especially Theorem 5.1, gives hyperboundedness criteria and uniform essential spectral-radius estimates; those estimates alone do not control the second eigenvalue of a finite matrix. These differences do not prove originality or exclude a stronger implication. Aida's directly accessible author preprint [2] uses an additive Schrödinger form with a derivation and chain-rule hypotheses, together with ground-state and weak-Poincaré inputs. Its setup and Theorem 10 were read in a targeted comparison; no identification of those hypotheses with this multiplicative Boolean transfer congruence has been established. No external priority, improvement over this literature, or exhaustive novelty clearance is asserted.

17 Relation to predecessors and reproducibility

The congruence analysis in [6] permits positive weights concentrating on an antipodal pair, with a transfer ratio tending to one as volume grows. Such a pair fails to preserve pair marginals. The parity analysis in [7] preserves every proper marginal but concerns a heat-bath generator; Proposition 5 shows precisely why it does not supply the transfer conclusion. The present synthesis obtains linear-order marginal matching for both the weight and the Perron stationary law, together with exponentially many slow transfer modes.

The code ingredients and the matrix methods are classical. In particular, Lemma 2 is an elementary existence argument, not a new code construction and not an efficient encoding claim. The citation to [4] supplies general limited-independence background, not a direct historical attribution for the dual-distance projection equivalence proved here. Work on simultaneously good primal and dual codes, including [5], is stronger in algorithmic aspects. Our targeted primary-source check did not identify the exact joint conclusion above; this is not proof of historical priority. External originality and the importance of this specialization remain unestablished.

All proofs are finite-dimensional and analytical. No numerical experiment or formal verification is needed or claimed. A source package retains the preregistered proof checks, complete analytical audit, source hashes and this manuscript. The permanent source tree is `papers/marginal-transfer` at the source commit below. The earlier spike proof is retained there in `evidence296/finite-moment296A.txt`; the balanced critical proof and its complete source audit are retained in `evidence298`. The exact-budget proof, registered checks and full analytical source audit are retained in `evidence299`. The full analytical reports, contexts, assignments and source bindings for the grouped tail results are in `evidence305` and `evidence307`. Their public-subset manifests explicitly identify private reference PDFs and an unrelated local directory index omitted from the original audit manifests;

no redistribution of these files is implied. The exact v4 PDF assessment is preserved there solely as historical evidence, with all seven objections and a separate author response. Historical partial results and exact-PDF reviews keep their original artifact identities; this revision neither changes those reports nor inherits their grades. The new background construction is grouped with these results because it uses the same transfer operator, critical exponent and balanced tilt. The complete source309A failure, corrected309B audit and source310 audit are retained in `evidence311`, with their original contexts and manifests. Three full third-party source captures from audit310 remain private and are explicitly excluded from the public subset; their original hashes are retained without implying redistribution rights. The original v5 PDF and its complete first exact-PDF assessment, including all six reservations, are historical evidence only. The expanded scalar attainment discussion follows the recorded source310 audit clarification. No v5 score is transferred to this manuscript. The attained-cutoff construction is retained as source313, with its complete analytical report, assignment, context and byte manifests in `evidence314`. It removes two restrictions of the preceding construction by clipping the tilt, compensating its mean only on the high sector, and reserving amplitude for finite tail errors. The source311 document and production-runner audits are also retained; that intermediate manuscript was not compiled or given a PDF score. A redundant complete frozen311 source archive remains private on D; the public subset declares its exact hash and original Git source, while retaining the full reports and original manifest. No report is shortened. The sharp scalar envelope (source315) and reverse-hypercontractive necessity argument (source317) are grouped here with the attained construction. Their complete source reports, contexts and fixed judges are retained in `evidence318`. The earlier exact v6 PDF, its full first-round assessment and all six reservations are historical evidence, not a rating of this revision. Its two explicitly forced equation numbers are replaced by ordinary sequential numbering; the corresponding mathematical displays are unchanged. The previous claim that the universal strict-tail interval was unresolved is superseded by Theorem 23. Document reproduction is a separate check of exact bytes and layout, not an independent proof of the theorems.

Model assessments and review context

OpenAI/Codex assisted the mathematical synthesis and manuscript preparation. The fabrication session is involved in the work. A distinct involved Codex/OpenAI session audits the frozen mathematical source; exact served backend versions and reasoning settings are not exposed for these source audits, so no particular model version or provider independence is asserted. Analytical source review is not human peer review and supplies no inherited PDF rating. The tail audits300–303 used one distinct auditor context with its earlier readings retained, not four blind or provider-independent reviews. The cutoff-one source306 received a separate fresh-assignment AI audit; shared provider and filesystem do not establish technical isolation or independence. The historical v4 PDF review records a requested configuration separately from what was actually exposed; its rating applies only to its original PDF. Source309's first audit found a missing upper restriction on a floor parameter in an additional comparison; that failed version remains unchanged. Its corrected version was reread in the same involved auditor context. Source310 received a separate fresh-assignment source audit; shared provider and filesystem again preclude an independence claim. The v5 exact-PDF review requested gpt-6-astra with high reasoning effort, but the served backend and effort were not independently attested. Its full report and six reservations remain bound to its original PDF, not this revision. No numerical assessment of the present revision has yet been incorporated or inherited. Source313 received a separate fresh-assignment analytical audit. Its reader checked the inherited explicit balanced family rather than relying on previous acceptance labels. The

shared provider and filesystem again do not establish independence. Its minor request to bracket each normalized product factor is incorporated here. The prior311 document audit's notation and navigation observations are also addressed, with all twenty-one prior proofs and twenty prior statement environments preserved. None of these source checks is an exact-PDF assessment or human review. The scalar-envelope and universal-tail source checks use a distinct involved auditor context with prior source313 readings retained. They are not blind or provider-independent reviews. The historical v6 PDF review requested gpt-6-astra with high reasoning effort, with served backend and effort not independently attested. Its full report remains bound to its original PDF. The new mathematical claims require assessment on this revision's exact final PDF and do not inherit its numerical rating. Any assessment attached to this manuscript must identify the exact final PDF hash, its actual review context, complete report, scores and objections. The manuscript does not incorporate a grade from another version or paper. The author owns and administers AIRR. A working-paper deposit is not an editorial acceptance or evidence of human peer review.

Reproducible source commit:
`d7e35107e2826d7a9342bdce30b50a3a725ec3dc`

References

[1] D. Bertsimas and I. Popescu, Optimal Inequalities in Probability Theory: A Convex Optimization Approach, *SIAM Journal on Optimization* **15**(3) (2005), 780–804. Publisher version.

[2] S. Aida, An estimate of the gap of spectrum of Schrödinger operators which generate hyperbounded semigroup, author preprint, 9 February 2001. Author preprint.

[3] M. Biskup and W. König, Eigenvalue order statistics for random Schrödinger operators with doubly-exponential tails, author manuscript, 21 April 2015, Lemma 4.5 and Eq. (4.29). Author manuscript.

[4] N. Alon, O. Goldreich and Y. Mansour, Almost k-wise independence versus k-wise independence, *Information Processing Letters* **88** (2003), 107–110. Author manuscript.

[5] M. Brehm and N. Resch, Linear Time Encodable Binary Code Achieving GV Bound with Linear Time Encodable Dual Achieving GV Bound, *ITCS 2026*, LIPIcs **362**, article 28. Publisher version.

[6] L. Eriksson, Congruence Rigidity and the Fusion Bound: What a Positive Weight Can and Cannot do to the Spectrum of a Transfer Kernel, AIRR, ARR-2026-7S1XA3PZEB9QM8FQ, version 1. Public record.

[7] L. Eriksson, Strictly Positive Parity Laws and the Limits of Marginal Certificates: Exact Heat-Bath Spectrum and Optimal Variance Tensorization, AIRR, ARR-2026-5HRC3PS9HZ9WDAA5, version 2 (2026). Public working paper.

[8] E. Mossel, R. O'Donnell, O. Regev, J. E. Steif and B. Sudakov, Non-interactive correlation distillation, inhomogeneous Markov chains, and the reverse Bonami-Beckner inequality, *Israel Journal of Mathematics* **154** (2006), 299–336. Author manuscript.

[9] P. Cattiaux, Hypercontractivity for perturbed diffusion semigroups, *Annales de la Faculté des Sciences de Toulouse* **14**(4) (2005), 609–628. Publisher archive.

[10] E. Mossel, K. Oleszkiewicz and A. Sen, On Reverse Hypercontractivity, arXiv:1108.1210v5 (2012). Primary manuscript.

[11] L. Gross, Invariance of intrinsic hypercontractivity under perturbation of Schrödinger operators, arXiv:2412.20282v2 (2025). Primary manuscript.

[12] J. Glück, Spectral gaps for hyperbounded operators, arXiv:1802.09422v6 (2019). Primary manuscript.

[13] L. Miclo, On hyperboundedness and spectrum of Markov operators. Author manuscript.

[14] F.-Y. Wang, Criteria of Spectral Gap for Markov Operators, arXiv:1305.4460v6 (2013). Primary manuscript.