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 weights on Boolean cubes whose labelled marginals on every set of at most a fixed positive fraction of the coordinates equal the uniform product marginals, while exponentially many normalized eigenvalues of the weighted transfer kernel approach one exponentially fast in the number of coordinates. The stationary law of the Perron Doob transform has exactly the same marginals as well. The unweighted product kernel has a dimension-independent gap. The construction combines binary codes with positive primal and dual distances, a positive product admixture, and a finite block-norm estimate. We prove the coding input by a direct entropy union bound and give explicit finite spectral constants. As a negative control, we calculate the full transfer spectrum of a single positive parity tilt: unlike its single-site heat-bath gap, its transfer gap stays uniformly positive at fixed coupling. Conversely, for the Boolean noise parameter $0 < \rho < 1$, density moments of order $s > (1+\rho)/(1-\rho)$ give an explicit dimension-independent transfer gap, without pointwise upper or positive lower bounds and even allowing zero weights. For every $1 < s < (1+\rho)/(1-\rho)$ and every $B > 1$, we construct normalized densities with $\|r_L\|_s \le B$ and $r_L \ge 1/B$ whose gaps collapse exponentially. Two antipodal product tilts and a Gram-matrix comparison supply this converse. The critical exponent has budget-dependent behavior: budgets sufficiently close to one give a uniform positive gap, whereas every $B > 2$ admits a strictly positive normalized family with $r_L \ge 1/B$ and stretched-exponentially closing gaps. The latter construction balances one product tilt against the uniform channel. The optimal critical budget remains unresolved. The ingredients are classical and external priority of these joint conclusions is not established. Global many-body weights are allowed; no fixed-range Gibbs or local-dynamics claim follows.

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

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

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 $u_L = \mathrm{Unif}(\Omega_L)$. 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 [2]. The relation between code projections and limited independence is standard background for the marginal comparison [1].

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

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

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

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

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 $z^j \ge z^{pL}$ for $j \le pL$, the binomial theorem gives

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

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 $|S| \le \lfloor \alpha L \rfloor$. 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 $m \ge 2$ and minimum nonzero weight d_C. Suppose its coordinate projections onto sets of size at most r are surjective. For $0 < \eta \le 1/2$, 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 $2 \le j \le m$ satisfies

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

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 $|S| \le r$, 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 $L \ge 2$ and fixed finite positive site rates [4]. Nevertheless, its transfer ratio at fixed β is bounded uniformly away from one, as follows.

Proposition 5. *Let $L \ge 2$, $\rho = \tanh\beta$, 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,

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

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

Proof. The row-normalized product transfer $P_\beta = M_\beta/(1+q)^L$ acts on χ_S by $\rho^{|S|}$. Write $D_t = \mathrm{diag}\sqrt{1 + t\chi_{[L]}}$. 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 [5] 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 $st \le 1/2$ now imply $\lambda_1 - \lambda_2 \ge d/(st) \ge 2d$. Consequently

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

This does not deteriorate when one sign part has small norm. In the previously omitted case, $c(B, \rho) \le 1/4 < 1/2$, since $e_\rho \ge 1$. 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 $\operatorname{gap}(T_{r_L}) \le Ke^{-\gamma L}$, with $K, \gamma > 0$. Then*

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

Proof. Apply Theorem 7 with $B = \max r_L \ge 1$ 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.

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 $B \ge 1$. For every $L \ge 1$ and density $r \ge 0$ with $\mathbb{E}r = 1$ and $\|r\|_s \le B$,*

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

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

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

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

If $\lambda_2 < \lambda_1/2$, the gap is greater than $1/2$. Otherwise $\lambda_2 \geq 1/2$. 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 $\lambda_2 h_\pm \leq T_r h_\pm$ pointwise, including their zero sets. For either $f = \sqrt{r}\, h_\pm$, multiplication by $\sqrt{r}$ gives $\lambda_2 f \leq r P_\rho f$. Hölder with $1/Q = 1/s + 1/t$ yields

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

Integration against the unit sign part also gives $\lambda_2 \leq \langle f, P_\rho f\rangle \leq \|f\|_A^2$. 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 $\|f\|_A \leq \|f\|_p^\theta \|f\|_Q^{1-\theta}$. No triangle inequality at exponent $p < 1$ is used. Combining these bounds,

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

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

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

Here $2\lambda_2/\lambda_1 \geq 1$ and $\lambda_2 \geq 1/2$. 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 $L \geq 1$, with $\|r_L\|_s \leq B$ and $r_L \geq 1/B$, and constants $\kappa > 1$ and $0 < \omega < \kappa$, independent of L, such that*

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

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

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

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 $X_L + Y_L \ge X_L$ gives $\lambda_1(X_L + Y_L) \ge \kappa^L$.

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

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

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 section distinguishes small and large moment budgets at the critical exponent $s = s_c$. 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 The critical exponent: small and large moment budgets

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 leave intermediate budgets unresolved; no optimal critical budget is asserted. 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 [5, 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 $\mathbb{E}W \le 2(B^s-1)/(s(s-1))$ and $\|r+1\|_s \le B+1$. 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 $L\ge 1$, such that*

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

Proof. For $|\delta|<1$ 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 $0<\delta_*<1$, $C\ge 1$ and $c>0$ such that, for $|\delta|\le\delta_*$,

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

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 $|\log a_L| \le Cd^4L^{-1/3} \le \log(B-1)$, 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. No optimality of the sufficient budget $B > 2$ is claimed.

10 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 [5]. 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 [8] gives a gap for ergodic self-adjoint hyperbounded Markov operators. Its Proposition 11 explains why a general quantitative gap does not follow from hyperboundedness alone. Wang's Theorem 1.1 in [9] 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 [6] 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 counterexample in Theorems 10 and 12 have $r \ge 1/B$ 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 [7] 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. The retained comparison also records incomplete access to Aida's original gap theorem. No external priority, improvement over this literature, or exhaustive novelty clearance is asserted.

11 Relation to predecessors and reproducibility

The congruence analysis in [3] 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 [4] 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 [1] 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 [2], 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`. 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, so no particular model version or provider independence is asserted. Analytical source review is not human peer review and supplies no inherited PDF 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:
`45d91a667750b295ce096fd495863583ba56eced`

References

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

[2] 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.

[3] 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.

[4] 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.

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

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

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

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

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