Stationary-to-weighting moment transfer and spectral rigidity for Boolean Doob kernels

Lluis Eriksson

14 November 2026. Working paper

Abstract

For positive potentials on a Boolean cube, the stationary density of the noise Doob kernel and its normalized weighting density are different nonlinear functions of the potential. We give an admissible family for which the stationary critical moment is below two while the weighting critical moment is above two, despite a fixed positive stationary floor and exact marginals through order proportional to the square root of the dimension. This refutes transfer of that numerical cap, not the existence of a larger uniform bound. A nonlinear consequence of the stationary floor, combined with classical ordinary and reverse hypercontractivity, gives explicit inverse-potential moments and weighting-moment bounds below a computable finite threshold. Exact stationary matching then gives strong convergence of the two densities throughout this subcritical range. For potentials supported on a binary linear-code Fourier algebra, we control every invariant sector of the full Doob kernel. Under a distance–cardinality condition its gap tends uniformly to the unweighted gap, for every positive quotient potential. Direct disjoint-block estimates cover a sharper parameter regime. A standard simplex code supplies positive potentials outside the disjoint-block class. Classical random linear coding extends this mechanism to positive-rate families at every sufficiently large length, with an explicit exponential gap estimate and normalized witnesses meeting any fixed stationary floor below one and the moment cap. The general stationary-gap question and critical weighting-moment bound remain open.

1 Setting and theorem map

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

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

Both π and r have mean one. The Markov kernel

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

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

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

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

and imposes

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

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

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

The manuscript groups analytical sources 349, 350, 351, corrected 353B and 355. The total-variation estimate from source 341 is proved again in Section 3. No spectral conclusion is inferred from moment convergence alone. No assertion about a physical Yang–Mills gap or a continuum limit is made.

2 Failure of exact critical-cap transfer

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The following estimates distinguish two simpler classes: two-level potentials cannot produce divergent critical weighting moments under a fixed stationary floor and cap, while product potentials obey the stronger constant-one moment comparison.

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

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

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

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

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

3 Inverse moments and subcritical transfer

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

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

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

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

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

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

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

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

For $1 < t < t_n$,

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

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

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

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

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

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

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

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

The map has

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

4 Full-chain estimates for code-supported potentials

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

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

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

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

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

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

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

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

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

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

It remains to control the other sectors, because the quotient spectrum alone is insufficient. For each coset $t + \mathcal{C}$, let $V_t = \chi_t\{F(z) : F \text{ any quotient function}\}$. These spaces are mutually orthogonal under $\pi\, d\mu$: products of quotient functions with π remain in the $\mathcal{C}$ algebra, and a Walsh character outside $\mathcal{C}$ has zero expectation against that algebra. They span all 2^L dimensions.

Multiplication by h and by $1/g$ preserves each sector, as does noise, so each is K-invariant. Under the identification $\chi_t F \leftrightarrow F$, its operator is

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

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

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

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

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

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

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

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

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

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

5 Disjoint blocks and a strictly larger potential class

5.1 The sharper direct block regime

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

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

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

For every positive H,

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

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

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

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

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

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

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

5.2 Simplex potentials outside disjoint-block representations

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

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

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

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

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

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

Choose

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

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

The code construction is classical [3]; its elementary parameter proof is included to avoid treating a coding reference as a new theorem. The strict separation is an existence statement for the full-support positive potentials just exhibited. It does not hold for every potential in the code algebra, since constants also belong to disjoint-block algebras. Nor does this argument assert a prescribed stationary floor or moment cap for the displayed amplitude family. Its use here is to establish that the full-chain code theorem covers positive potentials beyond the disjoint-block class. Arbitrary invertible changes of the original Boolean coordinates are not permitted replacements for permutations: they generally change Hamming weights and the noise operator.

6 Positive-rate families at every sufficiently large length

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

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

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

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

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

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

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

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

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

Moreover, for every fixed $0 < b < 1$ and finite $s > 1$, there are potentials in these same algebras, at every sufficiently large length, with energy one, $\pi \geq b$ and $\mathbb{E}\pi^s \leq 2$, whose Fourier support is the full code and which admit no disjoint equal-size block-parity representation with block size at least two. In particular these witnesses satisfy (1) *for its critical exponent and every fixed $C > 0$, eventually, and have the gap estimate* (8)*.*

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

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

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

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

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

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

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

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

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

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

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

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

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

It remains to construct the admissible strict-class witnesses. Nonzero columns ensure that every coordinate is active in some codeword; distinct columns ensure that every two coordinates are separated by some codeword. For the latter, if the two columns are $v \neq w$, choose a linear functional detecting $v + w \neq 0$. For $0 < \tau < 1$, choose

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

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

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

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

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

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

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

The entropy-volume and random-linear-code ingredients are classical Gilbert–Varshamov arguments [4]; the proof is included with the precise rounding and simultaneous column requirements used here. The rate R_{code} is distinct from the inverse-moment exponent R of Section 3. The theorem supplies all-large-length positive-rate realizations of the code-restricted mechanism. It does not remove the code-algebra hypothesis, establish a general critical weighting cap, or provide a collapsing example. The strict-class conclusion is existential for the specified full-support potentials, not universal for all positive functions on the code. The prescribed-floor witness assertion excludes $b = 1$ and makes no uniform endpoint claim as ρ approaches zero or one.

7 Positioning, failed routes and remaining question

Ordinary and reverse hypercontractivity are classical tools. Mossel, Oleszkiewicz and Sen explicitly state the Boolean reverse inequality, including negative exponents and tensorization, in equation (1.2) [1]. The use here is the floor-derived identity $gPg \ge b$ and its particular exponent iteration; we do not claim a new hypercontractive inequality. Representation-based decomposition of reversible chains is also classical; Boyd, Diaconis, Parrilo and Xiao discuss commuting symmetries and spectral reduction [2]. Our sector estimates explicitly retain all cosets and compare their signed kernels to the quotient kernel. The binary simplex construction is not original. Entropy-volume estimates and random-linear-code existence are classical Gilbert–Varshamov tools; the simultaneous column condition used here is verified directly rather than attributed as an extra clause of the coding theorem [4].

The bounded primary check supporting this draft reads the relevant statements just cited, not the entire related literature. Targeted searches did not establish the priority of the specific moment counterexample, the quantitative iteration or the uniform code-potential estimates. Search results about "Doob graphs" concern a different object from the Doob transform used here and do not establish overlap or novelty. The results in this draft are proved as stated; mathematical

correctness, originality and relevance remain separate questions. No claim of first discovery or frontier classification is made.

The evidence supplement preserves unsuccessful routes rather than replacing their records. Proposition 2.2 excludes divergent critical weighting moments for two-level potentials and exact-cap failure for product potentials. The finite fixed point in Section 3 obstructs only the specified iteration. Theorem 2.1 refutes the exact numerical cap, while its own moments remain bounded and its gap remains asymptotically unweighted. A generic pair of densities close in total variation is not a counterexample unless both arise from the same h and satisfy the stationary hypotheses.

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

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

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

The portion with $g > 1$ is controlled by the stationary cap. The joint small-g contribution is not controlled at the critical exponent by the finite inverse moments above. Additional exact matching has been used to obtain strong subcritical convergence, not to settle this endpoint. Even an endpoint moment bound would still require a justified spectral argument for the general question. The code theorem excludes a concrete structured family of collapsing mechanisms, not all stationary densities satisfying (1).

8 Model assessments and review context

This source manuscript is a new consolidation, not a reviewed PDF. No PDF has been produced from 354B, and no exact-PDF model assessment, numerical grade, peer-review decision, acceptance or publication is asserted. Earlier analytical sources 349, 350, 351, corrected 353B and 355 received separate-role source audits that the programme parent accepted; those audits do not certify this newly assembled manuscript or any future rendered PDF.

The manuscript preparation is by the involved fabrication role in the same provider workflow. This role previously developed sources 349 and 350 and consolidation 352; the parent supplied the bootstrap lead in source 350, authored sources 351, 353A and 355, and supplied the corrected simplex witness in 353B. The source auditor had relevant prior exposure, including the adverse scope finding on 353A. These are disclosed dependencies, not independent-model replications. The user requested the Astra high alias; the exact serving backend is not independently attested in this preparation, and no model version is inferred from that request. Full source reports, contexts, objections, failed drafts and byte manifests are retained in the accompanying evidence packet.

Only analytical derivation and document/source handling were performed for this consolidation. No mathematical program, numerical experiment, formal proof, local compilation or runtime result is claimed. Classical external results are attributed, and their bounded reading

scope is recorded. Draft354A was frozen before source355 was incorporated and remains preserved; it was not separately manuscript-audited before this revision. The retained source hashes identify the mathematical inputs; the manuscript has its own hash and requires its own separate source audit before document production and exact-PDF review.

References

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

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

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

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