A sharp stationary moment threshold away from equality for Boolean Doob chains with two bases

Lluis Eriksson

19 November 2026

Abstract

For Boolean noise of correlation one half, consider a reversible Doob transform whose stationary law factors through a binary linear quotient with two disjoint column bases. Above stationary moment order 15/7, exact matching of increasing-order coordinate marginals gives full ordered spectral recovery by the previously established packing criterion. For every $2 \le p < 15/7$ we construct chains in the same class with exact matching through a linear fraction of the coordinates and exponential convergence of the stationary density to one in L^p, yet with exponentially small spectral gap. The construction uses ordinary noise of correlation $1/4$ on a duplicated-coordinate quotient. A rare sphere traps the chain while most of the potential is supported near typical magnetizations. A convex correction on the two endpoints of each edge enforces exact moment matching even when the minimum density tends to zero. Conversely, at the second moment a uniformly bounded inverse stationary moment of order above 35/9 suffices for recovery at arbitrary rank. Without inverse control the equality case $p = 15/7$ remains open.

1 Setting and result

On $\{-1,1\}^L$ let μ be uniform and let P_ρ flip coordinates independently with probability $(1-\rho)/2$. For a strictly positive potential h write

$$K_h(x,y) = P_\rho(x,y)h(y)/(P_\rho h)(x), \qquad \pi_h = hP_\rho h/\mathbb{E}_\mu(hP_\rho h).$$

The stationary measure is $\nu = \pi_h\mu$. Its reversible edge law is proportional to $\mu(x)P_\rho(x,y)h(x)h(y)$. Matching through degree k means that every nonconstant Walsh character of degree at most k has zero ν-mean. Equivalently, the marginals on any k or fewer coordinates are uniform. Code and matrix terminology uses the identification of signs with bits. The spectral gap is $1-\lambda_2$, with eigenvalues in nonincreasing order and counted with multiplicity.

Theorem 1. *At ambient correlation $1/2$, for each fixed $2 \le p < 15/7$ there exist $\beta, c > 0$ and positive potentials in every sufficiently large even dimension $L = 2r$ such that*

$$\|\pi_h - 1\|_p \le e^{-cL}, \qquad \mathbb{E}_\nu \chi_T = 0 \quad (1 \le |T| \le \lfloor \beta L \rfloor), \qquad 1 - \lambda_2(K_h) \le e^{-cL}.$$

The stationary law factors through the matrix $[I_r\ I_r]$. Its displayed columns form two disjoint bases. In particular the second eigenvalue differs from the ordinary ambient value $1/2$ by at least $1/2 - e^{-cL}$.

The positive result of [1, section 7] applies to any binary quotient admitting two disjoint column bases. For $p > 15/7$, a uniformly bounded stationary L^p norm and matching through order $k \to \infty$ imply ordered spectral error $O(e^{-c_p k})$, uniformly in rank. It assumes positivity in each finite dimension. Theorem 1 proves that its strict moment threshold cannot be lowered throughout this class. The construction allows a vanishing stationary minimum. Theorem 4 below proves recovery at the second moment under a finite inverse moment, in particular under a fixed positive floor.

Our construction works more generally at quotient noise ρ. Put

$$j = \frac{1-\rho}{1+\rho}, \qquad b = \frac{\rho^2}{1-\rho^2}.$$

Theorem 2. *Fix $0 < \rho < 1$ and $1 < p < \infty$ with $1/p > j - 2b$. For all sufficiently large r, on the r-cube at correlation ρ, there are strictly positive potentials for which the three estimates of Theorem 1 hold with r in place of L.*

At $\rho = 1/4$ we have $j = 3/5$, $b = 1/15$ and $j - 2b = 7/15$. For the lift to $2r$ coordinates, let $y_i = x_i x_{i+r}$ and pull the quotient potential back along this map. Each y_i experiences independent correlation $(1/2)^2$. The stationary density pulls back exactly. A Walsh coefficient of the pullback vanishes unless its index is a union of complete pairs; complete pairs give exactly the corresponding quotient coefficient. Matching therefore holds to the stated ambient linear order, decreasing β if needed. The centered indicator of the inverse image of a quotient trap has the same variance and Dirichlet energy as its quotient indicator. This proves Theorem 1 from Theorem 2.

2 Ordinary returns and transitions to typical magnetizations

Write $m(x) = r^{-1} \sum x_i$ and $I(d) = ((1+d)\log(1+d) + (1-d)\log(1-d))/2$. For a small positive constant d, let S be the sphere at magnetization $d_r = d + O(1/r)$ obtained by rounding the plus count to the nearest integer. Let $B = \{|m| \le \alpha d\}$ with $0 < \alpha < \rho$, and write $q = P_\rho(x, S)$ for $x \in S$, independent of x by symmetry. We use three estimates:

$$\mu(S) \le e^{-rI(d)+O(1)}, \quad q \ge \exp[-rJ(d) - O(\log r)], \tag{1}$$

$$\sup_S P_\rho(x, B) \le \exp\left[-r\left(\frac{(\rho-\alpha)^2}{1-\rho^2}\frac{d^2}{2} - O_\rho(d^3)\right) + O(1)\right], \tag{2}$$

$$\sup_B P_\rho(y, S) \le \exp[-r(1-\rho\alpha)^2 d^2/2 + O(1)]. \tag{3}$$

Here $J(d) = (1+\rho)I(2d/(1+\rho))/2 - I(d) = jI(d) + O_\rho(d^4)$. Bounds are applied with d, α fixed before the dimension grows.

For (1), the first estimate is the usual binomial bound. In a stationary noise edge the probabilities of $++, +-, -+, --$ are (a, z, z, a) with $a = (1+\rho)/4$ and $z = (1-\rho)/4$. Choose the same nearest integer to rz for each off-diagonal count and fix the diagonals by the two prescribed sphere margins. Stirling's formula bounds the probability of this type below by $\exp[-r(1+\rho)I(2d/(1+\rho))/2 - O(\log r)]$. All four limiting cell proportions are positive. Divide by $\mu(S)$ and use transitivity. This return estimate also occurs in the earlier shell construction [2]. For (2), condition on the input x. The output signs have means ρx_i and each variance $1-\rho^2$. Their centered log moment generating functions at $-u$ are $(1-\rho^2)u^2/2 + O_\rho(u^3)$ uniformly for small u. Apply the exponential Markov bound with $u = (\rho-\alpha)d/(1-\rho^2)$. Finally, from B the mean output magnetization is at most $\rho\alpha d$. Hoeffding's bound for independent signs at threshold d_r gives (3).

3 A trap with negligible stationary moment cost

Choose a constant ℓ satisfying

$$\max(0, (j-1/p)/2) < \ell < \min(j, b).$$

Such a choice follows from the strict hypothesis and $j > 0$. Choose α small enough that $(\rho-\alpha)^2/(1-\rho^2) > \ell$ and $(1-\rho\alpha)^2 > j-\ell$. Since $I(d) = d^2/2 + O(d^4)$, choose $d > 0$ small enough that the following inequalities all have strict exponential slack:

$$J(d) - 2\ell I(d) < I(d)/p, \quad \sup_S P_\rho(x, B) \le e^{-r(\ell I(d)+\eta)}, \quad e^{r(J(d)-\ell I(d))} \sup_B P_\rho(y, S) \le e^{-\eta r}$$

for some $\eta > 0$ and all large r. The set B has complement of exponentially small uniform measure.

Put $P = P_\rho$, $H_0 = 1$ on B and $H_0 = e^{-Dr}$ elsewhere, where D is a fixed constant larger than $2 + |J(d)| + \ell I(d)$. Set

$$M = e^{-\ell I(d) r}/q, \qquad H = H_0 + M\mathbf{1}_S, \qquad W = HPH.$$

For $x \in S$, $PH_0(x)/(Mq) \le e^{-cr}$ for some $c > 0$. The numerator bounds all contributions from outside S. Consequently each row of the transformed kernel starting on S has exit probability at most e^{-cr}.

We verify $\|W - 1\|_p \le Ce^{-cr}$, allowing c to decrease from one estimate to the next. The background obeys $\|H_0PH_0 - 1\|_p \le 2\mu(B^c)^{1/p}$, by contraction of P and $0 < H_0 \le 1$. Off S the extra term is $H_0MP\mathbf{1}_S$. On B it is exponentially small by (1) and (3), absorbing the polynomial factor in the slack. Elsewhere off S it is at most $e^{-Dr}M$, also exponentially small. On S we have $PH \le 2Mq$ and $H \le 2M$ eventually (use $q \le 1$). Hence

$$\|\mathbf{1}_S W\|_p \le 4\,\mu(S)^{1/p} \frac{e^{-2\ell I(d) r}}{q} \le \exp\{r[J(d) - 2\ell I(d) - I(d)/p] + O(\log r)\}.$$

The stationary density $W/\mathbb{E}W$ therefore converges to one exponentially in L^p. The very small background outside B makes an additive correction to the density potentially incompatible with positivity. We correct the potential itself.

4 Exact marginals by tilting both ends of an edge

Lemma 3. *Let P be a reversible positive semidefinite Markov operator for a finite uniform measure μ, let $H > 0$, and let π be its normalized stationary density as above. Let $\phi_1, \dots, \phi_N$ be distinct nonconstant orthonormal Walsh characters on a cube. If $N\varepsilon$ is sufficiently small, where $\varepsilon = \|\pi - 1\|_1$, there exists f in their span with $\|f\|_\infty \le CN\varepsilon$ such that He^f has zero stationary mean for every listed character. Constants are absolute.*

Proof. Use the normalized symmetric edge law J of H. For $\Phi = (\phi_i)_i$ define $\psi(s) = \log \mathbb{E}_J e^{s\cdot(\Phi(X)+\Phi(Y))}$. Its gradient at zero has Euclidean norm at most $2\sqrt{N}\,\varepsilon$. The tilt by s is exactly the edge law of $He^{s\cdot\Phi}$, with marginal ν_s. For any Euclidean unit v, put $F = v \cdot \Phi$. Then $|F| \le \sqrt{N}$, $\mathbb{E}_\mu F = 0$, $\mathbb{E}_\mu F^2 = 1$, and

$$v^\top \nabla^2 \psi(s) v = 2\,\mathrm{Var}_{\nu_s} F + 2\,\mathrm{Cov}_{J_s}(F(X), F(Y)) \;\ge\; 2\,\mathrm{Var}_{\nu_s} F.$$

The covariance is nonnegative: the centered quadratic form of the edge law is that of P applied to the potential times the centered function, divided by its normalization. At zero, the variance is at least $1 - N\varepsilon - N\varepsilon^2$. In the Euclidean ball of radius $1/(4\sqrt{N})$ the marginal density relative to ν_0 is at least e^{-1}, by the two endpoint factors and their normalization. The variance comparison follows by minimizing the squared distance to a constant. Thus the Hessian is bounded below by an absolute positive multiple of the identity throughout this ball. For $N\varepsilon$ sufficiently small, Taylor's inequality gives $\psi(s) > \psi(0)$ on its boundary. The minimum is interior, and strong convexity bounds its distance from zero by $C\sqrt{N}\,\varepsilon$. At that minimum the gradient is twice the vector of stationary character means. Taking $f = s \cdot \Phi$ proves the bound. □

Use all characters of degrees 1 through $\lfloor \beta r \rfloor$. Their number is at most $\exp[rH_{\mathrm{bin}}(\beta)]$ for $0 < \beta < 1/2$ (natural entropy). Choose β sufficiently small using the exponential L^1 estimate. Then $\|f\|_\infty$ is exponentially small. The new stationary density differs from the old by a pointwise factor between $e^{-4\|f\|_\infty}$ and $e^{4\|f\|_\infty}$. The L^p conclusion follows. Each transition changes by at most $e^{2\|f\|_\infty}$ as a ratio; the small exit probability survives. The corrected stationary mass of S is at most $\|\pi_{\mathrm{new}}\|_p \mu(S)^{1-1/p} = o(1)$. Testing the gap variational formula with the centered indicator of S gives a gap at most $\sup_{x \in S} K(x, S^c)/(1 - \nu_{\mathrm{new}}(S))$. This proves Theorem 2.

5 Recovery with a finite inverse stationary moment

The counterexample uses a deep depression of the stationary density. A finite inverse moment suffices to exclude it even when the pointwise minimum tends to zero. We record the positive result in the same two-base class.

Theorem 4. *Let the ambient correlation be* $1/2$. *Suppose that* π *factors through a linear quotient with two disjoint full column bases, matches through degree* k, *and satisfies* $\|\pi\|_2 \le B$ *and* $\|\pi^{-1}\|_s \le D$ *for fixed* $s > 35/9$. *Then the maximum difference between corresponding ordered eigenvalues of* K_h *and* $P_{1/2}$ *is at most* Ce^{-ck}, *with constants depending only on* B, D, s. *In particular a fixed positive floor suffices at the second moment for any rank.*

Proof. Normalize h by $\mathbb{E}hPh = 1$, and set $P = P_{1/2}$, $Q = P_{1/\sqrt{2}}$, $g = Ph$, $F = Qh$, $U = h/g = \pi/g^2$. Then $\|F\|_2 = 1$. The positive scaling h inherits the quotient symmetry from π by uniqueness: the positive diagonal of P makes the stationary-scaling functional strictly convex in $\log h$ [2]. Two disjoint bases give the quotient contraction $P_\lambda : L^a \to L^{1+(a-1)\lambda^{-4}}$ for $a > 1$. This is ordinary forward hypercontractivity in the first basis followed by the same estimate in the second, with any remaining column noise a contraction [1]. Thus $\|g\|_5 = \|QF\|_5 \le 1$. Jensen gives $g^{-1} = 1/P(\pi/g) \le P(g/\pi)$. Set $1/a = 1/5 + 1/s$. Then

$$\|g^{-1}\|_q \le D, \qquad q = 1 + 16(a-1) > 20.$$

This gain is the reason for the inverse moment hypothesis. By Hölder, U is bounded in L^v for $1/v = 1/2 + 2/q < 3/5$.

For clarity, the comparison from moments to the spectrum can be carried out on the full cube. Put $a_0 = \mathbb{E}h$ and $V = \|F - a_0\|_2^2 = 1 - a_0^2$. Cauchy–Schwarz gives $1 = (\mathbb{E}hg)^2 \le a_0 \mathbb{E}(\pi g)$. Since $g = QF$ and $\pi - 1$ has only degrees $> k$, we have $V \le a_0 B\, 2^{-(k+1)/2}\sqrt{V}$. Thus $\|g - 1\|_2 \le C_B 2^{-k/2}$. On $E = \{|g - 1| > 2^{-k/8}\}$, uniform measure is at most $C_B 2^{-3k/4}$. Consequently $\|(U - \pi)\mathbf{1}_E\|_{5/3} \le C\mu(E)^{3/5 - 1/v}$ (here $v < 2$). Off E, $|U - \pi| \le C2^{-k/8}\pi$ for large k.

For each nonnegative quotient weight W, $\|QM_WQ\| \le \|W\|_{5/3}$ on the full cube. Indeed, its positive leading eigenvector for $W \neq 0$ is invariant under kernel translations by uniqueness, so

the quotient $L^2 \to L^5$ estimate bounds its quadratic form. It follows that $\|QM_{U-\pi}Q\|$ decays exponentially. For the signed comparison choose $\lambda = 3/4$, $\theta = 1/(\sqrt{2}\lambda) < 1$. Since $1 + \lambda^{-4} \geq 4$, the same positive multiplier argument bounds $\|P_\lambda M_{\pi-1} P_\lambda\|$ by $\|\pi - 1\|_2$ using $|\pi - 1|$. Split each P_θ in $Q = P_\theta P_\lambda$ at degree $\lfloor k/2 \rfloor$. The low–low form vanishes by matching, and each other block carries a decaying factor $\theta^{\lfloor k/2 \rfloor + 1}$. Thus $\|QM_U Q - P\|$ is exponentially small. The reversible conjugate of K_h is $M_{\sqrt{U}} P M_{\sqrt{U}}$. It has the same eigenvalues as $QM_U Q$ by the two Gram products of $M_{\sqrt{U}} Q$. Min–max completes the comparison. □

The number 35/9 is sufficient; we do not assert necessity. In particular the traps of Theorem 1 at $p = 2$ cannot have a uniformly bounded inverse stationary moment of any fixed order above 35/9.

6 Scope and antecedents

The rare-sphere joint type and calibration by its return probability are shared with [2]. There, compensation bands preserve a uniform floor and the all-correlation conclusion reaches some moment above $(1 + \rho)/(1 - \rho)$. Here the potential on the bulk complement can be arbitrarily small. This permits the full one-step cost b and the range $1/p > j - 2b$. The correction in Lemma 3 works with the edge law even when the density has a very small minimum. On a quotient obtained by repeating each coordinate t times, ambient correlation $1/2$ gives $\rho = 2^{-t}$ and $j - 2b = (2^t - 1)/(2^t + 1) - 2/(4^t - 1)$. Our sharp-threshold claim concerns two bases and fixed $p \geq 2$ away from equality. At equality the proof has no strict interval for ℓ. The available rank-dependent recovery criterion [3] is consistent with these traps; it imposes an additional condition on the rank relative to the matching order.

Classical results of Kirshner and Samorodnitsky [4] describe near extremality of spheres under ordinary Boolean noise. Reverse hypercontractivity [5] underlies the earlier sufficient packing estimate. These are inputs and comparisons for finite chains. Our result concerns the transformed stationary law, exact marginals and the gap of its actual kernel. We make no literature-wide priority claim.

Model assessments and review context

OpenAI Codex assisted with the construction, proof and manuscript in the programme research chat. Any assessment accompanying a deposit refers to the bytes and version identified in that record. The arguments here are analytic. The author administers AIRR.

References

[1] L. Eriksson, *Stationary concentration and local return in Boolean Doob chains*, AIRR working paper, v1, section 7.

[2] L. Eriksson, *A metastable Hamming shell with exactly matched stationary marginals*, public v1; subsequent source, section 6.

[3] L. Eriksson, *Controlling the denominator by the weight enumerator*, source revision, 2026.

[4] N. Kirshner and A. Samorodnitsky, *A moment ratio bound for polynomials and some extremal properties of Krawchouk polynomials and Hamming spheres*, arXiv:1909.11929.

[5] E. Mossel, K. Oleszkiewicz and A. Sen, *On reverse hypercontractivity*, arXiv:1108.1210.