Spectral Thresholds for Order and Gibbs Sampling Relaxation in a Spin System with a Given Coupling Matrix

Anonymous authors
Paper under double-blind review

Abstract

Statistical physics describes the properties of a typical spin system drawn from an ensemble, whereas in learning systems the properties of a specific system, such as a trained model, are of interest. We show that for a spin system defined by a given coupling matrix J, the spectrum of J locates the phase transition in the inverse temperature β. For $\beta < 1/\gamma(J)$, where $\gamma(J) = \lambda_{\max} - \lambda_{\min}$, the system is disordered: the mean squared overlap of two replicas is of $O(1/N)$, as for independent spins. When J has a delocalized top eigenvector and a bulk of small operator norm, the system is ordered for $\beta > 1/\lambda_{\max}$: the mean squared overlap is of order one. In a restricted Boltzmann machine (RBM) with weight matrix W and singular values $\mu_1 \geq \mu_2 \geq \cdots$, the spectrum also determines how fast Gibbs sampling relaxes. For $\beta < 1/\mu_1$ each singular value sets a time scale: the projection of the configuration on the k-th singular direction of W has a variance of $1/(1 - \beta\mu_k)$ and a corresponding decorrelation time scale, and is uncorrelated with the other projections. Experiments on synthetic couplings confirm these predictions. In an RBM trained on MNIST, the relaxation times and the onset of order follow the same formulas with μ_k multiplied by the measured average response of the units to their fields.

1 Introduction

Binary configurations with a quadratic Hamiltonian are among the most studied objects of statistical physics, from the Ising ferromagnet to mean-field spin glasses and associative memories (Edwards & Anderson, 1975; Sherrington & Kirkpatrick, 1975; Hopfield, 1982; Amit et al., 1985; Mézard et al., 1987; Talagrand, 2011; Panchenko, 2013). There the coupling matrix is drawn from an ensemble, which is a natural problem setting in physics — the properties of one particular $1\,\mu\text{m}^3$ sample of an alloy are of no practical interest; the interest is in the typical properties. In machine learning, however, a particular weight matrix (coupling) is associated with a particular neural network and can be the object of interest, for example in the study of a restricted Boltzmann machine (RBM) (Hinton, 2002; Decelle et al., 2018), a modern associative memory (Krotov & Hopfield, 2016; Ramsauer et al., 2021), or a trained layer read as a spin system (Bachtis et al., 2024). Results for a single, fixed coupling matrix are emerging (Levin et al., 2010; Eldan et al., 2022; Anari et al., 2021; Chen & Eldan, 2025; Mikulincer & Sohn, 2025).

In this work, we consider a quadratic Hamiltonian on $\{\pm 1\}^N$,

$$H(x) = -\tfrac{1}{2}\, x^\top J x, \tag{1}$$

where J is a given symmetric matrix with zero diagonal. Its Gibbs measure at inverse temperature $\beta > 0$ is $\mu_\beta(x) = e^{-\beta H(x)}/Z_\beta$. We show that the phases of μ_β, and the relaxation of the Gibbs sampling, are governed by the spectrum of J. We study how the second moment $\mathbb{E}[q_{ab}^2]$ of the Edwards–Anderson order parameter q_{ab} (Edwards & Anderson, 1975; Mézard et al., 1987) depends on β, and locate the phase transition. For Gibbs sampling from μ_β, we study its autocorrelation time, i.e., how many updates it takes to generate two effectively independent samples.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

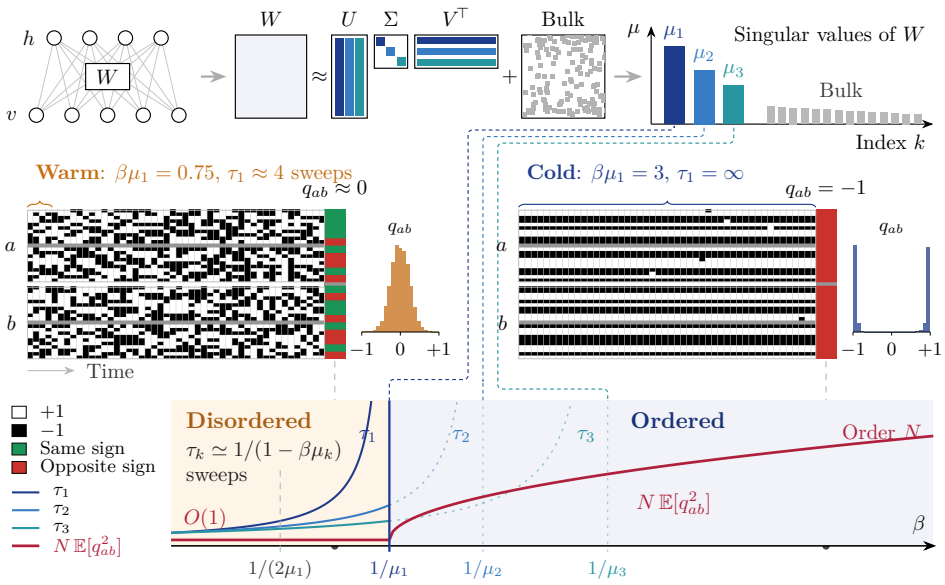

Figure 1: A schematic of the study on a bipartite spin system. **Top**: the coupling matrix W of two groups of spins and its singular values. **Middle**: two Gibbs sampling chains at a high/low temperature. **Bottom**: transition between the disordered/ordered regimes, and the relaxation times of the singular directions.

Specifically, our main results are as follows.

- For a symmetric J with eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$, let $\gamma(J) = \lambda_1 - \lambda_N$ be the spectral diameter. When $\beta < 1/\gamma(J)$, the system is disordered, $\mathbb{E}[q_{ab}^2] = O(1/N)$, and Gibbs sampling relaxes fast: every observable decorrelates within $1/(1-\beta\gamma(J))$ sweeps (Propositions 3.2 and 4.1). When $J = \lambda_1 ww^\top + J_{\text{bulk}}$ with a delocalized[1] w and $\beta\|J_{\text{bulk}}\|_{\text{op}}$ small compared with $(\beta\lambda_1 - 1)^2$, the system is ordered for $\beta > 1/\lambda_1$, $\mathbb{E}[q_{ab}^2] \geq c\,(\beta\lambda_1 - 1)^2$, with c independent of N (Theorem 3.3).
- In a bipartite system the units form two groups, $\sigma \in \{\pm 1\}^m$ and $\tau \in \{\pm 1\}^n$. Cross-group coupling is $W \in \mathbb{R}^{m\times n}$, so that $J = \left(\begin{smallmatrix} 0 & W \\ W^\top & 0 \end{smallmatrix}\right)$. Let the singular values of W be $\mu_1 \geq \mu_2 \geq \ldots$; the eigenvalues of J are $\pm\mu_k$, so $\gamma(J) = 2\mu_1$ and the general result gives disorder for $\beta < 1/(2\mu_1)$. When $W = \mu uv^\top + W_{\text{bulk}}$ has one delocalized pattern and a bulk of operator norm b, the system is disordered for $\beta < 1/(\mu + 2b)$ and, when βb is small compared with $(\beta\mu - 1)^2$, ordered for $\beta > 1/\mu$ (Proposition 3.5 and Theorem 3.4); the two regimes meet at $\beta = 1/\mu$ as $b \to 0$.
- When W consists of a few delocalized patterns and $\beta < 1/\mu_1$, Gibbs sampling has one time scale per singular value: the projection of the configuration on the k-th singular direction has variance $1/(1 - \beta\mu_k)$ to leading order (Theorem 4.4).

Numerical experiments on synthetic couplings and on an RBM trained on MNIST agree with these results. The trained weight matrix has a large bulk; for this case we give a corrected estimate of the relaxation times (Remark 4.5), which the measurements also confirm.

2 Setting

We consider N binary variables $x \in \{\pm 1\}^N$ with the Gibbs measure $\mu_\beta(x) = e^{\frac{\beta}{2}x^\top Jx}/Z_\beta$ at inverse temperature $\beta > 0$. The coupling J is a symmetric matrix with zero diagonal, and its eigenvalues are $\lambda_1 \geq \cdots \geq \lambda_N$.

[1] A unit vector $v \in \mathbb{R}^N$ is delocalized if $N\sum_i v_i^4$ is bounded by a constant independent of N, that is, v is spread over order N coordinates (Appendix A.1).

The *order parameter* compares two independent samples x^a and x^b of μ_β (Edwards & Anderson, 1975; Mézard et al., 1987). We study the second moment of the overlap,

$$q_{ab} = \frac{1}{N}\sum_{i=1}^{N} x_i^a x_i^b, \qquad \mathbb{E}[q_{ab}^2] = \frac{\|C\|_F^2}{N^2}, \tag{2}$$

where C is the covariance matrix, $C_{ij} = \mathbb{E}[x_i x_j]$. Since $C_{ii} = 1$ and $|C_{ij}| \le 1$, the order parameter satisfies $1 \le N\,\mathbb{E}[q_{ab}^2] \le N$. The lower end is attained by independent spins, for example at $J = 0$, and the upper end by a measure that sits on one configuration and its opposite. We call the range of β where $N\,\mathbb{E}[q_{ab}^2]$ is bounded by a constant independent of N the *disordered regime*, and the range where $N\,\mathbb{E}[q_{ab}^2] \ge cN$ for a constant $c > 0$ independent of N the *ordered regime*.

Gibbs sampling updates one coordinate at a time given the rest of the configuration, $x_i \sim \mu_\beta(\cdot \mid x_{\setminus i})$. An observable is a function $\varphi : \{\pm 1\}^N \to \mathbb{R}$ of the configuration. Time is counted in sweeps; in one sweep each coordinate is updated once on average. We measure how fast the chain decorrelates along φ by the *integrated autocorrelation time* $\tau_\varphi = \int_0^\infty \operatorname{Cov}(\varphi(x_0), \varphi(x_t))\,\mathrm{d}t \,/\, \operatorname{Var}_{\mu_\beta}(\varphi)$ with $x_0 \sim \mu_\beta$.

In the bipartite case $x = (\sigma, \tau)$ with $\sigma \in \{\pm 1\}^m$ and $\tau \in \{\pm 1\}^n$, and J is the block anti-diagonal matrix built from a weight matrix W. This is a binary restricted Boltzmann machine without biases (Smolensky, 1986; Hinton, 2002). The table below lists the singular values and vectors of W together with the remaining notation, and Appendix A.1 gives the full definitions.

μ_β	Gibbs measure $e^{\frac{\beta}{2}x^\top J x}/Z_\beta$	$x = (\sigma, \tau)$	bipartite configuration, $\sigma \in \{\pm 1\}^m$ and $\tau \in \{\pm 1\}^n$
$\gamma(J)$	spectral diameter $\lambda_1 - \lambda_N$	μ_k, u_k, v_k	singular values and vectors of W, or of its pattern part
$\mathbb{E}$, Var	under μ_β	$\mathcal{E}(\varphi, \varphi)$	Dirichlet form $\sum_i \mathbb{E}[\operatorname{Var}(\varphi \mid x_{\setminus i})]$
$\lambda_P(\beta)$	Poincaré constant: largest λ with $\lambda \operatorname{Var}(\varphi) \le \mathcal{E}(\varphi, \varphi)$ for every φ		

3 The order parameter

The regimes are decided by the spectrum of J, which we discuss in this section. The full proofs of the results in this and the next section are provided in Appendix A.

Lemma 3.1. *Let J be symmetric with zero diagonal and $\beta < 1/\gamma(J)$. Then $\lambda_P(\beta) \ge 1 - \beta\gamma(J)$, that is,*

$$(1 - \beta\gamma(J))\operatorname{Var}_{\mu_\beta}(\varphi) \le \mathcal{E}(\varphi, \varphi) \qquad \text{for every observable } \varphi. \tag{3}$$

This is Theorem 1 of Eldan et al. (2022) applied to $\beta J + \beta|\lambda_N| I$, which is positive semidefinite with operator norm $\beta\gamma(J)$ and defines the same measure.

Proposition 3.2 (Disorder for $\beta < 1/\gamma(J)$)**.** *Let J be symmetric with zero diagonal and $\beta < 1/\gamma(J)$. Then*

$$\|C\|_{\mathrm{op}} \le \frac{1}{1 - \beta\gamma(J)}, \qquad N\,\mathbb{E}[q_{ab}^2] \le \frac{1}{1 - \beta\gamma(J)}. \tag{4}$$

Theorem 3.3 (Order for $\beta > 1/\lambda_1$)**.** *Let $J = \lambda_1 w w^\top + J_{\mathrm{bulk}}$ with $\|w\| = 1$, and write $\theta = \beta\lambda_1 \in (1, 2]$, $\varepsilon = \theta - 1$, $\rho = \beta\|J_{\mathrm{bulk}}\|_{\mathrm{op}}$, $\kappa = \kappa_4(w) = \sum_i w_i^4$, and $X = \langle w, x\rangle$. Then for every N,*

$$N\,\mathbb{E}[q_{ab}^2] \;\ge\; N\Big(\frac{\mathbb{E}[X^2]}{N}\Big)^2, \qquad \mathbb{E}[X^2] \;\ge\; \frac{4}{3\varepsilon}\Big[\frac{\varepsilon^2}{12\kappa} - \rho N - \tfrac{1}{2}\log\frac{2}{\varepsilon}\Big], \tag{5}$$

In particular, if $\kappa \le M^2/N$ and $\rho < \varepsilon^2/(12M^2)$, then for N large enough

$$N\,\mathbb{E}[q_{ab}^2] \;\ge\; N\Big(\frac{\varepsilon}{9M^2} - \frac{4\rho}{3\varepsilon} - \frac{2\log(2/\varepsilon)}{3\varepsilon N}\Big)^2. \tag{6}$$

Proposition 3.2 gives $N\,\mathbb{E}[q_{ab}^2] = O(1)$ for $\beta < 1/\gamma(J)$, while Theorem 3.3 gives $N\,\mathbb{E}[q_{ab}^2] \geq cN$ for $\beta > 1/\lambda_1$: two replicas then agree on a macroscopic fraction of the coordinates. The restriction $\theta \leq 2$ is technical in the proof. When $\theta > 2$, the system remains in the ordered regime, and (5) holds with $\varepsilon = 1$.

The two thresholds differ by the factor $\gamma(J)/\lambda_1 = 1 + |\lambda_N|/\lambda_1$. For a coupling of the form $J = \lambda_1 w w^\top + J_{\text{bulk}}$ $|\lambda_N|$ is the width of the bulk, and the factor is close to 1; for a bipartite coupling $\lambda_N = -\lambda_1$, so the interval between the thresholds is $1/(2\mu_1) \leq \beta \leq 1/\mu_1$.

3.1 Bipartite coupling

This subsection closes that interval for a single pattern: order from $\beta = 1/\mu$ on (Theorem 3.4), and disorder up to $1/(\mu + 2b)$ (Proposition 3.5).

Theorem 3.4 (Order, one pattern plus noise)**.** *Let $W = \mu u v^\top + W_{\text{bulk}}$ with $u \in \mathbb{R}^m$, $v \in \mathbb{R}^n$ unit vectors, $\theta = \beta\mu \in (1, 2]$, $\varepsilon = \theta - 1$, $\rho = \beta\|W_{\text{bulk}}\|_{\text{op}}$, and $\kappa = \max(\kappa_4(u), \kappa_4(v))$. Let $X = u^\top\sigma$ and $Y = v^\top\tau$. Then for every m, n,*

$$\mathbb{E}[XY] \;\geq\; \frac{2}{3\varepsilon}\Big[\frac{\varepsilon^2}{6\kappa} - 2\rho\sqrt{mn} - \log\frac{2}{\varepsilon}\Big], \qquad N\,\mathbb{E}[q_{ab}^2] \;\geq\; N\Big(\frac{2\,\mathbb{E}[XY]}{N}\Big)^2 \tag{7}$$

In particular, if $\kappa \leq M^2/n$ with $m = n$ and $\rho < \varepsilon^2/(12M^2)$, then for n large enough $N\,\mathbb{E}[q_{ab}^2] \geq N\big(\varepsilon/(9M^2) - 4\rho/(3\varepsilon) - 2\log(2/\varepsilon)/(3\varepsilon n)\big)^2$.

For a single delocalized pattern the disordered regime extends from $\beta < 1/(2\mu_1)$ to $\beta < 1/(\mu + 2b)$.

Proposition 3.5 (Disorder up to the pattern strength)**.** *Let $W = \mu u v^\top + W_{\text{bulk}}$ with $\|u\|_\infty, \|v\|_\infty \leq N^{-\delta}$ for a fixed $\delta > 0$, and $b = \|W_{\text{bulk}}\|_{\text{op}}$. If $\beta < 1/(\mu + 2b)$, then*

$$\|C\|_{\text{op}} \leq \frac{2\,(1 + o(1))}{1 - \beta\mu - 2\beta b}, \qquad N\,\mathbb{E}[q_{ab}^2] \leq \frac{2\,(1 + o(1))}{1 - \beta\mu - 2\beta b}, \tag{8}$$

where $o(1)$ depends only on N and δ.

Theorem 3.4 and Proposition 3.5 together locate the transition of a single delocalized pattern at $\beta = 1/\mu$ as the bulk vanishes: disorder up to $1/(\mu + 2b)$, order from $1/\mu$ on.

4 Relaxation times

For every observable φ, the variance of $\mathbb{E}[\varphi(x_t) \mid x_0]$ over $x_0 \sim \mu_\beta$ decays at rate at least $2\lambda_P(\beta)$, so $1/\lambda_P(\beta)$ is the relaxation time of the chain.

Proposition 4.1 (Fast relaxation for $\beta < 1/\gamma(J)$)**.** *Let J be symmetric with zero diagonal and $\beta < 1/\gamma(J)$. Then the relaxation time of the chain is at most $1/(1 - \beta\gamma(J))$, and $\tau_\varphi \leq 1/(1 - \beta\gamma(J))$ for every observable φ.*

The first statement is Lemma 3.1. The bound on τ_φ follows because the chain is reversible: the autocorrelation of φ is a mixture of exponentials with rates at least $\lambda_P(\beta)$. The spectral gap also prevents macroscopic correlations and gives Proposition 3.2.

Variance also bounds the time from below.

Lemma 4.2. *For every observable φ, $\tau_\varphi \geq \mathrm{Var}_{\mu_\beta}(\varphi)/\mathcal{E}(\varphi, \varphi)$. For a linear observable $\varphi(x) = \langle a, x\rangle$, $\mathcal{E}(\varphi, \varphi) \leq \|a\|^2$, hence $\tau_\varphi \geq \mathrm{Var}_{\mu_\beta}(\varphi)/\|a\|^2$.*

The generator of the chain acts on an observable φ by $(\mathcal{L}\varphi)(x) = \sum_i(\mathbb{E}[\varphi \mid x_{\setminus i}] - \varphi(x))$. If $-\mathcal{L}\varphi = \alpha\varphi$ with $\alpha > 0$, the autocorrelation of φ is $e^{-\alpha t}$ and $\tau_\varphi = 1/\alpha$. The projections ξ_k of Section 4.1 satisfy this equation up to a small remainder, and the next proposition converts the size of the remainder into the error on τ_φ.

Proposition 4.3 (Approximate eigenfunctions)**.** *Let φ be an observable with $\mathbb{E}[\varphi] = 0$ and $\mathrm{Var}_{\mu_\beta}(\varphi) > 0$, let $\alpha > 0$, and let*

$$R = \mathcal{L}\varphi + \alpha\varphi, \qquad \eta = \frac{\|R\|_{L^2(\mu_\beta)}}{\|\varphi\|_{L^2(\mu_\beta)}}. \tag{9}$$

Then $\tau_\varphi \geq 1/(\alpha+\eta)$, and $\tau_\varphi = \alpha^{-1}(1+\delta)$ with $|\delta| \leq \eta/\lambda_P(\beta)$.

4.1 Bipartite coupling: one time scale per singular value

For the bipartite coupling, the next theorem gives the projection $\xi_k = (u_k^\top \sigma + v_k^\top \tau)/\sqrt{2}$ of the configuration on the k-th singular direction of W its own relaxation time, for every $\beta < 1/\mu_1$.

With r patterns and no noise, $W = \sum_{k\leq r} \mu_k u_k v_k^\top$, write $\theta_k = \beta\mu_k$ and $\kappa_* = \max_{k\leq r} \max\{\kappa_4(u_k), \kappa_4(v_k)\}$; a delocalized family has $\kappa_* \to 0$.

Theorem 4.4 (Directional relaxation times)**.** *Let $\theta_1 < 1$. There are constants C_1 and C_R, depending only on θ_1 and r, such that for all m, n and each $k \leq r$:*

(i) $\left|\mathrm{Var}_{\mu_\beta}(\xi_k) - \dfrac{1}{1-\theta_k}\right| \leq \dfrac{9C_1\kappa_*}{2(1-\theta_k^2)}$;

(ii) if $\kappa_* \leq (1-\theta_1^2)/(9C_1)$, then $\tau_{\xi_k} \geq (1-\theta_k)^{-1}\big(1 - 2C_R\kappa_*/(1-\theta_k)\big)$, and $\tau_{\xi_k} = (1-\theta_k)^{-1}(1+\delta_k)$ with $|\delta_k| \leq 2C_R\kappa_*/\lambda_P(\beta)$;

(iii) $\left|\mathrm{Cov}_{\mu_\beta}(\xi_j, \xi_k)\right| \leq C_1\kappa_*/(1-\theta_1)$ *for $j \neq k$.*

The Poincaré constant $\lambda_P(\beta)$ in (ii) is bounded below uniformly in N for $\beta < 1/(2\mu_1)$ and, for the flat single pattern ($m = n$, $u_1 = v_1 = n^{-1/2}\mathbf{1}$), for every $\theta_1 < 1$.

As $\kappa_* \to 0$, variance and relaxation time of ξ_k both equal $(1-\theta_k)^{-1}(1+o(1))$ in these two cases, and different directions are uncorrelated to leading order.

Proof sketch. The conditional mean of x_i is $\mathbb{E}[x_i \mid x_{\setminus i}] = \tanh(\beta(Jx)_i)$, with the local field $(Jx)_i = \sum_{j\neq i} J_{ij}x_j$, so for a linear observable $\langle a, x\rangle$ the generator gives $\sum_i a_i\big(\tanh(\beta(Jx)_i) - x_i\big)$. Taking a to be the k-th singular direction, linearizing $\tanh(\beta(Jx)_i) \approx \beta(Jx)_i$ and using the orthogonality of the patterns gives

$$\mathcal{L}\xi_k = -(1-\theta_k)\,\xi_k + \mathcal{R}_k, \qquad \|\mathcal{R}_k\|_{L^2(\mu_\beta)} \leq \sqrt{2}\,C_R\kappa_*. \tag{10}$$

So ξ_k is an approximate eigenfunction of $-\mathcal{L}$ with eigenvalue $1-\theta_k$, and part (ii) is Proposition 4.3 with $\alpha = 1-\theta_k$ and $\eta \leq 2C_R\kappa_*$. Part (i) follows from the same linearization and the stationarity identity $\mathbb{E}[\varphi\,\mathcal{L}\varphi] = -\mathcal{E}(\varphi,\varphi)$; and (iii) from the cross terms of the stationarity identity. The full proof is in Appendix A. □

In (10), when a spin is resampled it forgets its value at rate 1, while the local field pushes it towards the pattern with strength $\theta_k\xi_k$ along the k-th direction; the net restoring rate is $1-\theta_k$, and ξ_k evolves to leading order as an Ornstein–Uhlenbeck process. Its noise does not depend on θ_k, so variance and relaxation time grow together as $1/(1-\theta_k)$. As β increases the strongest pattern softens first, and its time diverges at $\beta = 1/\mu_1$; the order parameter stays $O(1/N)$ throughout, since no single direction reaches a variance of order N.

Remark 4.5. The linearization in the proof sketch replaces $\tanh(\beta(Jx)_i)$ by a line of slope 1. When the local fields $h_i = (W\tau)_i$ and $g_j = (W^\top\sigma)_j$ are not small, the slope of tanh at the field is $1 - \tanh^2$, with block averages

$$\chi_\sigma = \frac{1}{m}\sum_{i\leq m} \mathbb{E}\big[1 - \tanh^2(\beta h_i)\big], \qquad \chi_\tau = \frac{1}{n}\sum_{j\leq n} \mathbb{E}\big[1 - \tanh^2(\beta g_j)\big].$$

If the local fields and the projections $u_k^\top\sigma$, $v_k^\top\tau$ are jointly Gaussian, the same linearization with slopes χ_σ and χ_τ in place of 1 gives the relaxation time and the slow observable

$$\tau_k = \frac{1}{1 - \beta\mu_k\sqrt{\chi_\sigma\chi_\tau}}, \qquad \varphi_k = \sqrt{\chi_\tau}\,u_k^\top\sigma + \sqrt{\chi_\sigma}\,v_k^\top\tau, \tag{11}$$

which for $\chi_\sigma = \chi_\tau = 1$ are $1/(1-\theta_k)$ and $\sqrt{2}\,\xi_k$. The derivation is in Appendix A.8.

For $\beta > 1/\mu_1$ and the flat pattern, the flip $x \mapsto -x$ preserves μ_β and exchanges the sets $\{X + Y > 0\}$ and $\{X + Y < 0\}$, where $X = u^\top \sigma$ and $Y = v^\top \tau$, so a chain that samples μ_β must move between them; since a single-site update changes $X + Y$ by at most $2n^{-1/2}$, it must pass through configurations in which the two block magnetizations cancel, which have exponentially small probability, and its relaxation time is exponential in n.

Proposition 4.6 (Slow relaxation for $\beta > 1/\mu$)**.** *(a) Let W, $\theta \in (1, 2]$, ε, ρ and κ be as in Theorem 3.4. Then the relaxation time of the chain is at least*

$$\frac{1}{2e^2 N} \exp\Big(\frac{\varepsilon^2}{6\kappa} - 2\rho\sqrt{mn}\Big). \tag{12}$$

(b) Let $W = \mu u v^\top$ with $u = v = n^{-1/2}\mathbf{1}$ and $\theta = \beta\mu \in (1, 2]$. Let $X = u^\top \sigma$ and $Y = v^\top \tau$, and define $\Delta_\beta > 0$ by $\mu_\beta\big(|X + Y| \le 2n^{-1/2}\big) = e^{-n\Delta_\beta + o(n)}$, the rate at which the probability that the two block magnetizations cancel decays. Then the relaxation time of the chain is at least $e^{n\Delta_\beta - o(n)}$, and $\Delta_\beta \ge \varepsilon^2/6$.

5 Experiments

In this section we give numerical results on the second moment of the overlap between two independent samples, $\mathbb{E}[q_{ab}^2]$, and on the integrated autocorrelation time τ_{ξ_k} of the projection ξ_k of the configuration. Both quantities vary with β as Sections 3 and 4 predict, and the variation is governed by the spectrum of the coupling matrix. In some experiments we design couplings with a particular spectrum and a simple enough structure that the Gibbs average can be summed exactly; we always run Monte Carlo simulations as well, either to check those computations or where exact summation is infeasible. Error bars on the sampled values are drawn only where they would be larger than the plotting symbol, and the samplers are described in Appendix B.

Order parameter. We first verify how $\mathbb{E}[q_{ab}^2]$ varies with β on the Curie–Weiss coupling $J = (\mu/N)(\mathbf{1}\mathbf{1}^\top - I)$ (Levin et al., 2010) and a bipartite version, where $J = \frac{\mu}{n}[0, \mathbf{1}\mathbf{1}^\top; \mathbf{1}\mathbf{1}^\top, 0]$. On both couplings the energy depends on the configuration only through the number of positive spins. The order parameter $N\mathbb{E}[q_{ab}^2]$ is tractable: it depends on the total magnetization for the Curie–Weiss coupling and the number of positive spins in each block for the bipartite coupling. See Appendix B.1 for the details.

Proposition 3.2 gives disorder for $\beta < 1/\gamma(J)$ and Theorem 3.3 gives order for $\beta > 1/\lambda_1$. The Curie–Weiss coupling has $\lambda_1 = \mu(1 - 1/N)$, $\lambda_N = -\mu/N$ and $\gamma(J) = \lambda_1 - \lambda_N = \mu$. Its two thresholds are $1/\gamma(J) = 1/\mu$ and $1/\lambda_1 = 1/\mu + 1/(\mu(N-1))$, which differ by $O(1/N)$ on the β axis. The bipartite coupling has $\lambda_1 = \mu$, $\lambda_N = -\mu$, and $\gamma(J) = 2\mu$. Its two thresholds are $1/\gamma(J) = 1/(2\mu)$ and $1/\lambda_1 = 1/\mu$, and there is a gap of $1/(2\mu)$.

Figure 2 shows $N\,\mathbb{E}[q_{ab}^2]$ against $\theta = \beta\lambda_1$ for $N \in \{128, 512, 2048\}$, on the Curie–Weiss coupling in panel (a) and on the bipartite coupling in panel (b), where $\lambda_1 = \mu$ and so $\theta = \beta\mu$. With this scaling the threshold $1/\lambda_1$ of Theorem 3.3 is at $\theta = 1$ in both panels. Below it $N\,\mathbb{E}[q_{ab}^2]$ stays bounded as N grows; above it the value grows in proportion to N and reaches about $0.84N$ at $\theta \approx 2$ for all three sizes on both couplings. The circles are Monte Carlo estimates of the same quantity. In panel (a) the two thresholds almost coincide. In panel (b) the curves stay disordered up to $\theta = 1$, while $\gamma(J) = 2\mu$ puts the threshold of Proposition 3.2 at $\theta = 1/2$, where the curves show no change; the bound of Proposition 3.5 covers the interval between the two. The transition is at $1/\lambda_1$ on both couplings, and the factor of two between the thresholds in panel (b) is a property of the bipartite spectrum as in Proposition 3.5.

Directional relaxation times. We then examine Theorem 4.4 on a bipartite coupling. The theorem concerns the relaxation times of the Gibbs sampling chain along the singular directions of the weight matrix. For the projection ξ_k along the k-th singular direction with singular value μ_k (Section 4.1), the relaxation time τ_{ξ_k} is at least $(1 - \beta\mu_k)^{-1}$ at β.

We take a rank-3 $W = \sum_{k \le 3} \mu_k u_k v_k^\top$ on $m = n$ units, and $W_{\text{bulk}} = 0$. Let $\mu_1 : \mu_2 : \mu_3 = 1 : 0.8 : 0.6$, and draw u_k and v_k uniformly at random subject to orthonormality. The

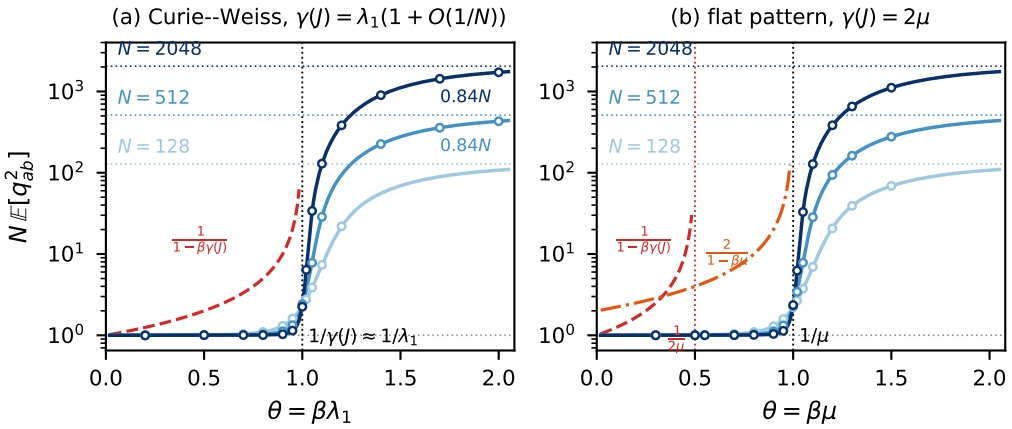

Figure 2: Phase transition of the order parameter $N\mathbb{E}[q_{ab}^2]$ demonstrated on the Curie–Weiss (a) and bipartite (b) couplings. Exact curves (lines) and Monte Carlo estimates (circles) are shown for systems of $N = 128$, 512, 2048. Dashed curves are the bounds of Propositions 3.2 and 3.5; the latter holds up to a factor $1 + o(1)$. Red is the bound of Proposition 3.2, orange the bipartite bound of Proposition 3.5; the vertical dotted lines mark where each bound diverges. The horizontal axes are scaled such that $\theta = \beta\lambda_1$. See the main text for more details.

temperature is set at $\beta\mu_1 = 0.9$. The strengths are unequal so that at one temperature the three directions have clearly different predicted times, which distinguishes one time scale per direction from a single time scale of the chain; $\beta\mu_1 = 0.9$ is below the threshold and close enough to it that the top direction is slow. We estimate the relaxation time τ_{ξ_k} by summing the autocorrelation measured along the chain up to a window set by Sokal's rule (Appendix B).

Figure 3(a) shows $1/\tau_{\xi_k}$ against $\beta\mu_k$, where the prediction is the line $1 - \beta\mu_k$, corrected for sampling the chain once per sweep (Appendix B); the points at $\beta\mu_k = 0$ are control directions orthogonal to the patterns. The predicted times $(1 - \beta\mu_k)^{-1}$ are 10, 3.57 and 2.17 sweeps for the three directions. We measure them at seven sizes, $n = 2^7$ to 2^{13}; at the largest, $n = 8192$, we obtain 10.1 ± 0.6, 3.71 ± 0.11 and 2.20 ± 0.06 sweeps. Panel (b) shows the measured time compared to the prediction. With three patterns, $\beta\mu_1 = 0.9$ lies in the interval that Proposition 3.5 covers for one pattern; that several patterns behave the same way is supported by the static scan in Appendix B.1 and discussed in Appendix A.4.

A trained weight matrix. We train a restricted Boltzmann machine on MNIST with contrastive divergence (Hinton, 2002) (784 and 64 units in the two layers, 15 epochs) and discard the biases to study the bipartite coupling of its weight matrix. When μ_1 is normalised to 1 the matrix has $\sum_k \mu_k^2 = 16.2$ and the singular values below the top are significant.

For such a W we compare the measured times with the prediction (11) of Remark 4.5. This prediction has no free parameter: χ_σ and χ_τ are averages of the slope $1 - \tanh^2$ of tanh under μ_β. Figure 4(a) compares the measured relaxation times of the eight leading directions at six values of β between 0.2 and 0.9. Figure 4(a) compares the relaxation times of the eight leading directions at six values of β between 0.2 and 0.9 with the two predictions in Theorem 4.4 and Remark 4.5, respectively. The figure shows the measured relaxation times divided by their predictions, where a dot on the line at height 1 would indicate perfect agreement. The renormalized prediction (11) is accurate on this RBM trained on real data: the green dots stay close to 1 over the whole range of β.

Figure 4(b) shows the order parameter of the same matrix against $\theta = \beta\mu_1$. The transition happens above $\theta = 1$, the order parameter reaches 10 at $\theta = 1.74$. On a finer grid of θ, the corrected coupling $\theta\sqrt{\chi_\sigma\chi_\tau}$ of Remark 4.5 rises towards 1, stays within 0.003 of 1 for

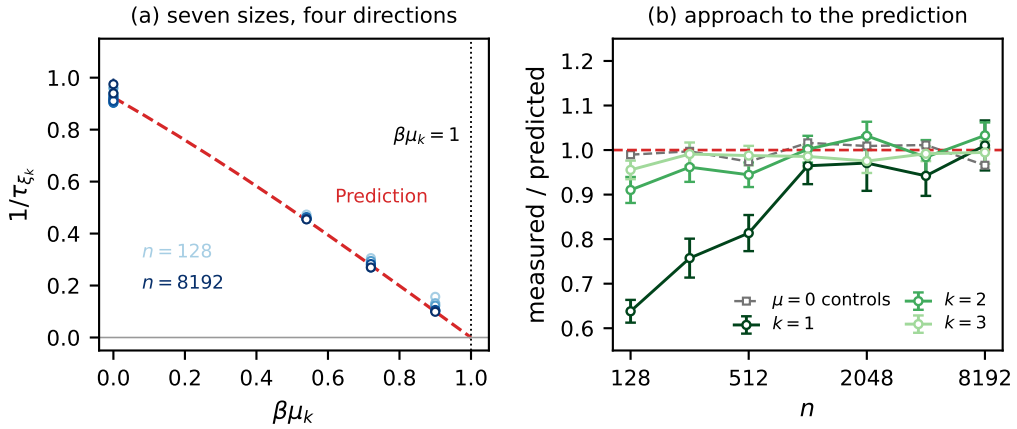

Figure 3: Relaxation times of three planted directions at $\beta\mu_1 = 0.9$. (a) $1/\tau_{\xi_k}$ against $\beta\mu_k$ for systems of size $n = 2^7$ (light) to 2^{13} (dark); the prediction of Theorem 4.4 is the curve $2\tanh\big((1-\beta\mu_k)/2\big)$, which is $1-\beta\mu_k$ corrected for sampling the chain once per sweep. (b) The measured time divided by prediction, against n; the grey squares are $\mu = 0$.

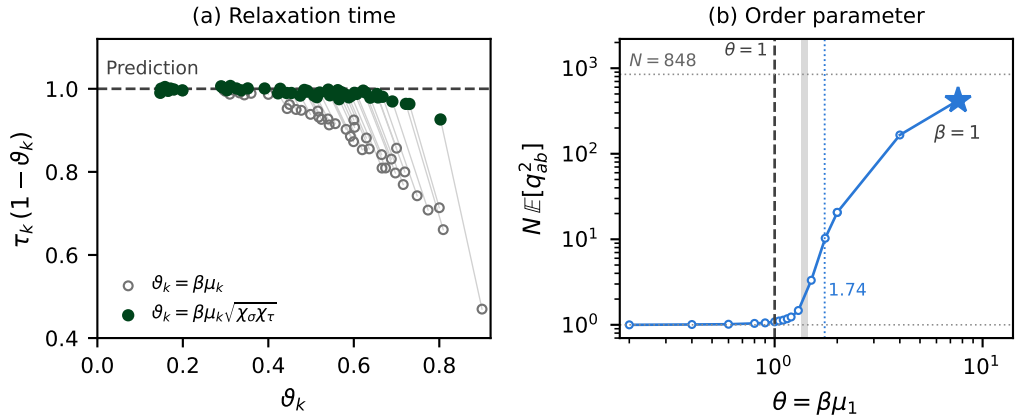

Figure 4: An RBM trained on MNIST. (a) The measured relaxation time τ_k of the leading directions, multiplied by $1-\vartheta_k$, against ϑ_k, with $\vartheta_k = \beta\mu_k$ (open grey) and with $\vartheta_k = \beta\mu_k\sqrt{\chi_\sigma\chi_\tau}$ (filled green); a thin line joins the two points of one measurement. (b) The order parameter against θ. The dotted line marks $\theta = 1.74$, the shaded band marks where $\theta\sqrt{\chi_\sigma\chi_\tau}$ is within 0.003 of 1, and the star marks $\beta = 1$.

$1.35 \le \theta \le 1.45$, and decreases beyond; the order parameter starts to grow in this interval (shaded in Figure 4(b)). The corrected coupling thus locates the onset of order as well as the relaxation times.

For reference, the model is trained at $\beta = 1$, where $\theta = \mu_1 = 7.62$ lies in the ordered regime. The training overlap $\mathbb{E}|q_{ab}| = 0.70$ and $N\,\mathbb{E}[q_{ab}^2] = 420$ of the maximum 848, and the autocorrelation of the top direction has not decayed within the 10^4 sweeps of the run.

6 Related work

Spin systems with pattern couplings. The coupling of a Hopfield network is a sum of rank-one patterns (Hopfield, 1982). A Boltzmann machine with Gaussian units in one

layer has a coupling of the same form (Barra et al., 2012). For random ± 1 patterns and for couplings that are constant on blocks of spins, the ordered phase and the fluctuations above the critical temperature are known in the limit $N \to \infty$ (Amit et al., 1985; Gentz, 1996; Fedele & Contucci, 2011; Barra et al., 2011). Theorems 3.4 and 4.4 give the ordered regime and the variance of each projection for a given W at finite size, where the patterns are real-valued and have unequal strengths. Our lower bound on $\log Z$ is the mean-field bound of Basak & Mukherjee (2017) and Jain et al. (2018), and we evaluate it at the product measure tilted along the pattern.

Restricted Boltzmann machines. Decelle et al. (2017; 2018) study the singular values of trained RBM weights and the phases that they produce. Bachtis et al. (2024) relate the singular values to the relaxation of the sampler during training. Long mixing times of the sampler in trained RBMs are also measured by Decelle et al. (2021) and Béreux et al. (2023). We state this picture as theorems for one given W at finite size. RBMs with sparse weights are studied by Tubiana & Monasson (2017).

Sampling under spectral conditions. For a given J, functional inequalities that give fast mixing of Glauber dynamics hold when the spectral diameter of βJ is below one (Bauerschmidt & Bodineau, 2019; Eldan et al., 2022; Anari et al., 2021; Chen & Eldan, 2025), and above one, sampling is hard for the class of all such matrices (Kunisky, 2024; Galanis et al., 2024). Couplings with a few outlier eigenvalues are treated in Koehler et al. (2022) and Mikulincer & Sohn (2025), and mixing times of RBMs are known for special classes of weight matrices (Kwon et al., 2024; Feng et al., 2026). We use the inequalities of Eldan et al. (2022) and Mikulincer & Sohn (2025) for the disorder results (Lemma 3.1 and Proposition 3.5), and the covariance bound that they imply (Eldan et al., 2022, Remark 14) gives, through (2), the bound on the order parameter. These works concern the mixing of the whole chain at high temperature, and our theorems also cover the ordered regime and the relaxation time of each direction.

Glauber dynamics of mean-field models. The relaxation of Glauber dynamics has been studied for the Curie–Weiss model (Levin et al., 2010; Ding et al., 2009) and for random rank-one couplings (Bovier et al., 2022), and it has been described in the limit $N \to \infty$ for the pattern overlaps of Hopfield-type couplings (Coolen, 2001). Theorem 4.4 and Proposition 4.6 describe the relaxation for a given bipartite W at finite size, below and above the transition.

Low-rank matrix estimation. In rank-one matrix estimation, the posterior distribution of the signal is a spin system of the same kind as ours, and its phase transition has been determined on average over the signal and the noise in the limit $N \to \infty$ (Lelarge & Miolane, 2019; Miolane, 2017; Baik & Lee, 2020). Our theorems concern a given coupling matrix at finite size.

7 Conclusion

For a spin system defined by a given coupling matrix J, the spectrum of J determines how fast Gibbs sampling relaxes, and it places the transition of $\mathbb{E}[q_{ab}^2]$ from order $1/N$ to order one: the system is disordered for $\beta < 1/\gamma(J)$ and, when J has a delocalized top eigenvector and a small bulk, ordered for $\beta > 1/\lambda_{\max}$. In a restricted Boltzmann machine, the projection of the configuration on the k-th singular direction of the weight matrix has a variance of $1/(1-\beta\mu_k)$ and a corresponding decorrelation time scale.

Three questions follow from these results. One is to prove that the transition stays at $1/\mu_1$ when the weight matrix has several patterns of comparable strength, as it does in all our simulations (Appendix B.1). The other two are the ordered regime of such weight matrices and the relaxation within one ordered state.

AI use statement

In this work, we used generative AI tools (Claude models by Anthropic) for the following tasks with required disclosure. Claude Fable took part in designing the strategy that uses the dynamics of the Gibbs sampler to prove the bounds on the relaxation time and the lower bound on the order parameter in the ordered regime, and in designing the experiments that test the theoretical results. Claude Opus was used to read the literature and adapt the key mathematical tools to our setting, to assist in deriving and verifying the theoretical results, and to assist in designing the experiments, analysing and interpreting their results, and writing the experiment reports. We have not used generative AI tools for translation, cleaning and reformatting data sets, and generating synthetic data sets and qualitative and thematic data analysis are not applicable to this work (some potentially relevant tasks are discussed above). Additionally, we used generative AI tools for visualization and formatting. We have reviewed all AI-assisted work. We carefully checked and verified the AI-generated proofs on paper, and we checked and double-checked the plotting code and the plotted data. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

Ethics statement

This work is theoretical. Its experiments use synthetic coupling matrices and the public MNIST dataset, and involve no human subjects or personal data. We are not aware of any ethical concern raised by this work.

Reproducibility statement

All proofs are in Appendix A, and the samplers, estimators and parameters of the experiments are described in Appendix B. The supplementary material contains the code of every experiment shown in the figures, with the random seeds used, the result files and checkpoints from which the figures are drawn, one notebook per figure that redraws it from these files, and a script that downloads MNIST.

References

Daniel J. Amit, Hanoch Gutfreund, and Haim Sompolinsky. Spin-glass models of neural networks. *Physical Review A*, 32(2):1007–1018, 1985. doi: 10.1103/PhysRevA.32.1007.

Nima Anari, Vishesh Jain, Frederic Koehler, Huy Tuan Pham, and Thuy-Duong Vuong. Entropic independence I: Modified log-Sobolev inequalities for fractionally log-concave distributions and high-temperature Ising models. *arXiv preprint arXiv:2106.04105*, 2021. First posted 2021-06-08; a related paper by the same authors under a different title, "Entropic independence: optimal mixing of down-up random walks," appeared at STOC 2022.

Dimitrios Bachtis, Giulio Biroli, Aurélien Decelle, and Beatriz Seoane. Cascade of phase transitions in the training of energy-based models. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, pp. 55591–55619. Curran Associates, Inc., 2024. doi: 10.52202/079017-1766. arXiv:2405.14689; journal version in J. Stat. Mech. 2025, 074004, doi:10.1088/1742-5468/adec64.

Jinho Baik and Ji Oon Lee. Free energy of bipartite spherical Sherrington–Kirkpatrick model. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 56(4):2897–2934, 2020. doi: 10.1214/20-AIHP1062.

Adriano Barra, Giuseppe Genovese, and Francesco Guerra. Equilibrium statistical mechanics of bipartite spin systems. *Journal of Physics A: Mathematical and Theoretical*, 44(24): 245002, 2011. doi: 10.1088/1751-8113/44/24/245002. arXiv:1012.1261.

Adriano Barra, Alberto Bernacchia, Enrica Santucci, and Pierluigi Contucci. On the equivalence of Hopfield networks and Boltzmann machines. *Neural Networks*, 34:1–9, 2012. doi: 10.1016/j.neunet.2012.06.003.

Anirban Basak and Sumit Mukherjee. Universality of the mean-field for the Potts model. *Probability Theory and Related Fields*, 168(3–4):557–600, 2017. doi: 10.1007/s00440-016-0718-0. arXiv:1508.03949.

Roland Bauerschmidt and Thierry Bodineau. A very simple proof of the LSI for high temperature spin systems. *Journal of Functional Analysis*, 276(8):2582–2588, 2019. doi: 10.1016/j.jfa.2019.01.007. arXiv:1712.03676.

Nicolas Béreux, Aurélien Decelle, Cyril Furtlehner, and Beatriz Seoane. Learning a restricted Boltzmann machine using biased Monte Carlo sampling. *SciPost Physics*, 14:032, 2023. doi: 10.21468/SciPostPhys.14.3.032. arXiv:2206.01310.

Anton Bovier, Frank den Hollander, and Saeda Marello. Metastability for Glauber dynamics on the complete graph with coupling disorder. *Communications in Mathematical Physics*, 392:307–345, 2022. doi: 10.1007/s00220-022-04351-8. arXiv:2107.04543.

Yuansi Chen and Ronen Eldan. Localization schemes: A framework for proving mixing bounds for Markov chains. *Duke Mathematical Journal*, 174(8), 2025. doi: 10.1215/00127094-2024-0063. arXiv:2203.04163; extended abstract in FOCS 2022, doi:10.1109/focs54457.2022.00018.

A. C. C. Coolen. Statistical mechanics of recurrent neural networks II. dynamics. In F. Moss and S. Gielen (eds.), *Handbook of Biological Physics, Vol. 4: Neuro-Informatics and Neural Modelling*, chapter 15, pp. 619–684. North-Holland/Elsevier, Amsterdam, 2001. doi: 10.1016/S1383-8121(01)80018-X. arXiv:cond-mat/0006011.

Aurélien Decelle, Cyril Furtlehner, and Beatriz Seoane. Equilibrium and non-Equilibrium regimes in the learning of Restricted Boltzmann Machines. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, pp. 5345–5359, 2021. arXiv:2105.13889; journal version in J. Stat. Mech. 2022, 114009, doi:10.1088/1742-5468/ac98a7.

Aurélien Decelle, Giancarlo Fissore, and Cyril Furtlehner. Spectral dynamics of learning in restricted Boltzmann machines. *EPL (Europhysics Letters)*, 119:60001, 2017. doi: 10.1209/0295-5075/119/60001.

Aurélien Decelle, Giancarlo Fissore, and Cyril Furtlehner. Thermodynamics of restricted Boltzmann machines and related learning dynamics. *Journal of Statistical Physics*, 172 (6):1576–1608, 2018. doi: 10.1007/s10955-018-2105-y.

Jian Ding, Eyal Lubetzky, and Yuval Peres. The mixing time evolution of Glauber dynamics for the mean-field Ising model. *Communications in Mathematical Physics*, 289(2):725–764, 2009. doi: 10.1007/s00220-009-0781-9. arXiv:0806.1906.

S. F. Edwards and P. W. Anderson. Theory of spin glasses. *Journal of Physics F: Metal Physics*, 5(5):965–974, 1975. doi: 10.1088/0305-4608/5/5/017.

Ronen Eldan, Frederic Koehler, and Ofer Zeitouni. A spectral condition for spectral gap: Fast mixing in high-temperature Ising models. *Probability Theory and Related Fields*, 182 (3–4):1035–1051, 2022. doi: 10.1007/s00440-021-01085-x. arXiv:2007.08200.

Micaela Fedele and Pierluigi Contucci. Scaling limits for multi-species statistical mechanics mean-field models. *Journal of Statistical Physics*, 144(6):1186–1205, 2011. doi: 10.1007/s10955-011-0334-4. arXiv:1011.3216.

Weiming Feng, Heng Guo, and Minji Yang. Rapid mixing in positively weighted restricted Boltzmann machines. *arXiv preprint arXiv:2604.00963*, 2026.

Andreas Galanis, Alkis Kalavasis, and Anthimos Vardis Kandiros. On sampling from Ising models with spectral constraints. In Amit Kumar and Noga Ron-Zewi (eds.), *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2024)*, volume 317 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 70:1–70:14, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi: 10.4230/LIPIcs.APPROX/RANDOM.2024.70. arXiv:2407.07645.

Barbara Gentz. A central limit theorem for the overlap in the Hopfield model. *The Annals of Probability*, 24(4):1809–1841, 1996. doi: 10.1214/aop/1041903207.

Geoffrey E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8):1771–1800, 2002. doi: 10.1162/089976602760128018.

John J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982. doi: 10.1073/pnas.79.8.2554.

Vishesh Jain, Frederic Koehler, and Elchanan Mossel. The mean-field approximation: Information inequalities, algorithms, and complexity. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet (eds.), *Proceedings of the 31st Conference on Learning Theory (COLT 2018)*, volume 75 of *Proceedings of Machine Learning Research*, pp. 1326–1347. PMLR, 2018. arXiv:1802.06126.

Frederic Koehler, Holden Lee, and Andrej Risteski. Sampling approximately low-rank Ising models: MCMC meets variational methods. In Po-Ling Loh and Maxim Raginsky (eds.), *Proceedings of Thirty Fifth Conference on Learning Theory (COLT 2022)*, volume 178 of *Proceedings of Machine Learning Research*, pp. 4945–4988. PMLR, 2022. arXiv:2202.08907.

Dmitry Krotov and John J. Hopfield. Dense associative memory for pattern recognition. In *Advances in Neural Information Processing Systems*, volume 29, 2016.

Dmitriy Kunisky. Optimality of Glauber dynamics for general-purpose Ising model sampling and free energy approximation. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 5013–5028, 2024. doi: 10.1137/1.9781611977912.180. arXiv:2307.12581.

Youngwoo Kwon, Qian Qin, Guanyang Wang, and Yuchen Wei. A phase transition in sampling from restricted Boltzmann machines. *arXiv preprint arXiv:2410.08423*, 2024.

Marc Lelarge and Léo Miolane. Fundamental limits of symmetric low-rank matrix estimation. *Probability Theory and Related Fields*, 173(3–4):859–929, 2019. doi: 10.1007/s00440-018-0845-x.

David A. Levin, Malwina J. Luczak, and Yuval Peres. Glauber dynamics for the mean-field Ising model: cut-off, critical power law, and metastability. *Probability Theory and Related Fields*, 146(1-2):223–265, 2010. doi: 10.1007/s00440-008-0189-z. arXiv:0712.0790.

Marc Mézard, Giorgio Parisi, and Miguel Angel Virasoro. *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications*, volume 9 of *World Scientific Lecture Notes in Physics*. World Scientific, Singapore, 1987.

Dan Mikulincer and Youngtak Sohn. Fast mixing in Ising models with a negative spectral outlier via Gaussian approximation. *arXiv preprint arXiv:2512.22803*, 2025.

Léo Miolane. Fundamental limits of low-rank matrix estimation: the non-symmetric case. *arXiv preprint arXiv:1702.00473*, 2017.

Dmitry Panchenko. *The Sherrington–Kirkpatrick Model*. Springer Monographs in Mathematics. Springer, New York, 2013.

Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Thomas Adler, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Kjetil Sandve, Victor Greiff, David Kreil, Michael Kopp, Günter Klambauer, Johannes Brandstetter, and Sepp Hochreiter. Hopfield networks is all you need. In *International Conference on Learning Representations*, 2021.

David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Physical Review Letters*, 35(26):1792–1796, 1975. doi: 10.1103/PhysRevLett.35.1792.

Paul Smolensky. Information processing in dynamical systems: Foundations of harmony theory. In David E. Rumelhart and James L. McClelland (eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1: Foundations*, pp. 194–281. MIT Press, Cambridge, MA, 1986.

Alan D. Sokal. Monte Carlo methods in statistical mechanics: foundations and new algorithms. In Cecile DeWitt-Morette, Pierre Cartier, and Antoine Folacci (eds.), *Functional Integration: Basics and Applications*, volume 361 of *NATO ASI Series B*, pp. 131–192. Springer, Boston, 1997. doi: 10.1007/978-1-4899-0319-8_6.

Michel Talagrand. *Mean Field Models for Spin Glasses: Volume I: Basic Examples*, volume 54 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge*. Springer, Berlin, 2nd edition, 2011.

Jérôme Tubiana and Rémi Monasson. Emergence of compositional representations in restricted Boltzmann machines. *Physical Review Letters*, 118:138301, 2017. doi: 10.1103/PhysRevLett.118.138301.

A Proofs

We first provide a detailed description of the setting used in the main text.

A.1 Complete setting

The system. Consider the configuration space $\{\pm 1\}^N$ of N binary variables, with the quadratic Hamiltonian and the Gibbs measure at inverse temperature $\beta > 0$,

$$H(x) = -\tfrac{1}{2}\, x^\top J x, \qquad \mu_\beta(x) = \frac{e^{-\beta H(x)}}{Z_\beta}, \qquad Z_\beta = \sum_{x \in \{\pm 1\}^N} e^{-\beta H(x)}, \tag{13}$$

where J is a symmetric matrix with zero diagonal.[2] Write $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$ for the eigenvalues of J. The zero diagonal gives $\operatorname{tr} J = \sum_k \lambda_k = 0$, so $\lambda_N < 0 < \lambda_1$ whenever $J \neq 0$. The spectral diameter is

$$\gamma(J) = \lambda_1 - \lambda_N. \tag{14}$$

It is at least λ_1, so $1/\gamma(J) \leq 1/\lambda_1$, and the two thresholds $\beta < 1/\gamma(J)$ and $\beta > 1/\lambda_1$ used below do not overlap.

The order parameter. Let x^a, x^b be two independent samples (replicas) from μ_β. Their overlap

$$q_{ab} = \frac{1}{N} \sum_{i=1}^{N} x_i^a x_i^b \tag{15}$$

is the Edwards–Anderson order parameter (Edwards & Anderson, 1975; Mézard et al., 1987). We study its second moment

$$\mathbb{E}[q_{ab}^2], \qquad x^a, x^b \sim \mu_\beta \otimes \mu_\beta \tag{16}$$

(the global flip $x \mapsto -x$ preserves μ_β, so $\mathbb{E}[q_{ab}] = 0$ for every β). Write $C_{ij} = \mathbb{E}_{\mu_\beta}[x_i x_j]$. Expanding the square of (15) and taking the expectation over the two independent samples gives

$$\mathbb{E}[q_{ab}^2] = \frac{\|C\|_F^2}{N^2}. \tag{17}$$

Since $C_{ii} = 1$, $\|C\|_F^2 \geq N$ and $\mathbb{E}[q_{ab}^2] \geq 1/N$, with equality for independent spins. We therefore consider $N\,\mathbb{E}[q_{ab}^2]$, a number between 1 and N: it is $O(1)$ when the two samples are unrelated and of order N when they agree on a macroscopic fraction of the coordinates.

Order, disorder and the transition point. For a single pair (J, N) the quantity $N\,\mathbb{E}[q_{ab}^2]$ is a number; bounded and divergent refer to $N \to \infty$. Our statements have the following shape: they hold for every J satisfying a spectral condition, with constants that depend on β and on the quantities in that condition but not on N. Disorder means an upper bound $N\,\mathbb{E}[q_{ab}^2] \leq c$ of this kind; order means a lower bound $\mathbb{E}[q_{ab}^2] \geq c > 0$ of this kind. The transition point is the position on the β axis that separates the two kinds of statement.

Gibbs sampling. The standard Markov chain for sampling from μ_β is the single-site heat bath. Each coordinate i carries an independent Poisson clock of rate 1; when the clock rings, x_i is resampled from the conditional distribution of μ_β given the other coordinates $x_{\setminus i}$, which are left unchanged:

$$P\big(x_i = +1 \mid x_{\setminus i}\big) = \frac{1}{1 + \exp\big(-2\beta\,(Jx)_i\big)}, \qquad (Jx)_i = \sum_{j \neq i} J_{ij} x_j. \tag{18}$$

[2] Since $x_i^2 = 1$, adding a diagonal matrix to J leaves μ_β unchanged; fixing the diagonal to zero loses no generality.

μ_β is the stationary distribution of this chain. The unit of time is the time in which each coordinate is updated once on average, one sweep.[3] For a real function φ of the configuration, the Dirichlet form

$$\mathcal{E}(\varphi,\varphi) = \sum_{i=1}^{N} \mathbb{E}_{\mu_\beta}\big[\mathrm{Var}(\varphi \mid x_{\setminus i})\big] \tag{19}$$

controls how fast the chain drives an initial distribution towards μ_β. The Poincaré constant is the largest λ such that $\lambda\,\mathrm{Var}_{\mu_\beta}(\varphi) \le \mathcal{E}(\varphi,\varphi)$ for all φ; for every observable φ, the variance of $\mathbb{E}[\varphi(x_t) \mid x_0]$ over $x_0 \sim \mu_\beta$ decays at rate at least 2λ, and $1/\lambda$ is the relaxation time of the chain. For the stationary chain and an observable φ, the integrated autocorrelation time

$$\tau_\varphi = \int_0^\infty \frac{\mathrm{Cov}\big(\varphi(x_0),\varphi(x_t)\big)}{\mathrm{Var}_{\mu_\beta}(\varphi)}\,\mathrm{d}t \tag{20}$$

is the time scale of the chain as seen through φ: two observations of φ separated by τ_φ sweeps are approximately uncorrelated. The relaxation time of a direction in Section 4 is τ_φ for φ the projection of the configuration on an eigenvector of J.

The bipartite case. Split the coordinates into two blocks $x = (\sigma,\tau)$, $\sigma \in \{\pm1\}^m$, $\tau \in \{\pm1\}^n$, $N = m+n$, and take a block anti-diagonal coupling,

$$J = \begin{pmatrix} 0 & W \\ W^\top & 0 \end{pmatrix}, \qquad H(\sigma,\tau) = -\sigma^\top W \tau, \qquad \mu_\beta(\sigma,\tau) = \frac{e^{\beta\sigma^\top W\tau}}{Z_\beta}, \tag{21}$$

where $W \in \mathbb{R}^{m\times n}$ is the weight matrix between the blocks. This is a binary restricted Boltzmann machine without biases (Smolensky, 1986; Hinton, 2002). Write the singular value decomposition of W as

$$W = \sum_k \mu_k u_k v_k^\top, \qquad \mu_1 \ge \mu_2 \ge \cdots \ge 0, \tag{22}$$

with (u_k) and (v_k) orthonormal. The nonzero eigenvalues of J are $\pm\mu_k$, with eigenvectors $w_k^\pm = (u_k, \pm v_k)/\sqrt{2}$; hence $\lambda_1 = \mu_1$ and $\gamma(J) = 2\mu_1$. The projection of the configuration on the k-th singular direction,

$$\xi_k = \langle x, w_k^+\rangle = \frac{u_k^\top\sigma + v_k^\top\tau}{\sqrt{2}}, \tag{23}$$

is the observable of Section 4. Besides the single-site chain, a bipartite system admits block Gibbs sampling: σ and τ are resampled alternately as whole blocks, the coordinates of one block being conditionally independent given the other. This is the chain run by contrastive divergence; the conversion between the time scales of the two chains is given in Section 5.

Structural assumption on the weights. In the bipartite case our theorems are stated for weights of the form

$$W = \sum_{k\le r} \mu_k u_k v_k^\top + W_{\mathrm{bulk}}, \qquad b = \|W_{\mathrm{bulk}}\|_{\mathrm{op}}. \tag{24}$$

We call $u_k v_k^\top$ a pattern, μ_k its strength, and W_{bulk} the bulk, or noise. A unit vector $u \in \mathbb{R}^n$ is delocalized if $n\,\kappa_4(u)$ is bounded by a constant independent of n, where $\kappa_4(u) = \sum_i u_i^4$ is the inverse participation ratio, that is, u is spread over order n coordinates; a pattern is delocalized if u_k and v_k are. Each theorem states the rate it requires (the uniform $u = n^{-1/2}\mathbf{1}$ has $\kappa_4 = 1/n$, the smallest possible value). A pattern is flat if $m = n$ and $u = v = n^{-1/2}\mathbf{1}$. Trained layers have this shape: a few leading singular values stand out, the corresponding singular vectors are spread over many units, and the remaining singular values form a bulk (Decelle et al., 2017; 2018). When $b < \mu_1/2$, λ_1 is separated from the other eigenvalues of J; such an eigenvalue is called an outlier. For a general symmetric J the corresponding form is

$$J = \lambda_1 w w^\top + J_{\mathrm{bulk}}, \qquad \|w\| = 1, \quad \|J_{\mathrm{bulk}}\|_{\mathrm{op}} \text{ small}, \tag{25}$$

with w delocalized; a single bipartite pattern is the case $w = (u,v)/\sqrt{2}$, $\lambda_1 = \mu$.

[3] This chain is known as Glauber dynamics, single-site heat bath, or the Gibbs sampler; we use this correspondence when quoting external results.

Notation. Table 1 lists the symbols that are used throughout the paper or far from their definition, with the place where each is defined. A symbol that has different meanings in different places has one row for each. For a single pattern the index $k = 1$ is dropped, so that μ, u, v, X, Y, θ and $w_\pm$ stand for μ_1, u_1, v_1, X_1, Y_1, θ_1 and $w_1^\pm$, and κ is κ_* with $r = 1$.

Table 1: Notation, with the place of definition.

Symbol	Meaning	Defined in
System and order parameter		
$\mathbb{E}$, Var	expectation and variance with respect to μ_β at the inverse temperature under consideration	Appendix A.1
q_{ab}	overlap of two independent samples	(15)
$N\,\mathbb{E}[q_{ab}^2]$	order parameter, between 1 and N	(17)
C	covariance matrix, $C_{ij} = \mathbb{E}[x_i x_j]$	(17)
$\gamma(J)$	spectral diameter $\lambda_1 - \lambda_N$	(14)
w	unit top eigenvector of a general J	(25)
J_{bulk}	bulk $J - \lambda_1 w w^\top$ of a general J (Theorem 3.3, Appendix A.5)	(25)
θ	$\beta\lambda_1$ (general J)	Theorem 3.3
ε	$\theta - 1$, excess above the threshold	Theorem 3.3
ρ	$\beta\|J_{\text{bulk}}\|_{\text{op}}$ (general J)	Theorem 3.3
X	projection $\langle w, x\rangle$ on the top eigenvector	Theorem 3.3
M	constant in the condition $\kappa \le M^2/N$	Theorem 3.3
$\kappa_4(u)$	inverse participation ratio $\sum_i u_i^4$ of a unit vector u	Appendix A.1
κ	$\kappa_4(w)$ (general J)	Theorem 3.3
κ_*	largest κ_4 over the vectors u_k, v_k of all patterns	Section 4.1
Chain and relaxation time		
τ_φ	integrated autocorrelation time of φ	(20)
$\mathcal{E}(\varphi, \varphi)$	Dirichlet form	(19)
$\lambda_P(\beta)$	Poincaré constant	Appendix A.1
P_i	conditional expectation given $x_{\setminus i}$	(28)
$\mathcal{L}$	generator of the chain	(32)
α	approximate eigenvalue of $-\mathcal{L}$	Proposition 4.3
R	residual $\mathcal{L}\varphi + \alpha\varphi$	(9)
η	relative size of the residual	(9)
$\Gamma(f, g)$	$\mathcal{L}(fg) - f\mathcal{L}g - g\mathcal{L}f$	(123)
Bipartite coupling and patterns		
σ, τ	the two blocks of $x = (\sigma, \tau)$, in $\{\pm1\}^m$ and $\{\pm1\}^n$	(21)
W	weight matrix between the two blocks	(21)
μ_k, u_k, v_k	singular values and singular vectors of W	(22)
$w_k^\pm$	eigenvectors $(u_k, \pm v_k)/\sqrt{2}$ of J	Appendix A.1
ξ_k	projection on the k-th singular direction	(23)
X_k, Y_k	$u_k^\top \sigma$ and $v_k^\top \tau$	(96)
r	number of patterns in W	(24)
W_{bulk}, b	bulk of W and its operator norm	(24)
ρ	$\beta\|W_{\text{bulk}}\|_{\text{op}}$	Theorem 3.4
J_{bulk}	block anti-diagonal matrix built from W_{bulk} (Proposition 3.5, Appendix A.4)	(56)
θ_k	$\beta\mu_k$	Section 4.1
δ_k	relative error of τ_{ξ_k}	Theorem 4.4
C_R	constant bounding the linearization remainder	(112)
C_1	error constant of the directional variance and covariance	(133)
R_k	remainder in $\mathcal{L}\xi_k = -(1 - \theta_k)\xi_k + R_k$	(113)
h_i, g_j	local fields $(W\tau)_i$ and $(W^\top\sigma)_j$	Remark 4.5
χ_σ, χ_τ	average slope $1 - \tanh^2$ over the σ and the τ block	Remark 4.5
τ_k	corrected relaxation time of the k-th direction	(11)
φ_k	corrected slow observable	(11)
Δ_β	exponential rate of the bottleneck, $\beta > 1/\mu$	Proposition 4.6
Notation of the proofs		

Symbol	Meaning	Defined in
P	positive semidefinite part of $\frac{\beta}{2}J$ after the diagonal shift	(58)
ϱ	modified log-Sobolev constant	(65)
$F(t)$, $\mathbb{E}_t$	log partition function with pattern strength t, and the expectation of its measure	(86)
θ'	comparison strength below the threshold, $1-\varepsilon/2$	Lemma A.13
P_0, $\mathbb{E}_0$	uniform measure on $\{\pm 1\}^N$ and its expectation	Appendix A.5
$\mathbb{Q}_\lambda$	product trial measure with tilt λ	Lemma A.11
a_i, b_j	$\beta(W\tau)_i$ and $\beta(W^\top\sigma)_j$	(97)
K_0, K_p	moment bounds on the projections	(100)
Notation of the experiments		
μ	strength of the Curie–Weiss and flat bipartite couplings	Section 5
ϑ_k	$\beta\mu_k$ or $\beta\mu_k\sqrt{\chi_\sigma\chi_\tau}$	Figure 4
$\hat{\tau}$	estimator of the integrated autocorrelation time	(169)

A.2 Organization of the proofs

In most parts of the main text, we roughly organize the results by the properties of the related objects: the static statements (Section 3) before the dynamical ones (Section 4), and inside each section, the general coupling comes before the bipartite one, and the disordered regime before the ordered one, except in Section 3.1, where Theorem 3.4 precedes Proposition 3.5. While in this appendix, the natural organization of the proofs is mostly determined by how the results depend on the mathematical tools and existing literature, as well as on each other. Disorder and the relaxation times rest on the chain viewed as an operator (Appendix A.3), and order and the slow relaxation above the threshold rest on bounds on the free energy; each group is therefore proved in one place, for the general and the bipartite coupling together, in Appendices A.4 to A.8.

Figure 5 shows how the statements of Sections 3 and 4 are obtained. A block is a statement of the paper or a step of a proof, with its number; an arrow from one block to another means that the proof of the second cites the first. The foundations are at the bottom. Above them the variational route, which proves order, and the dynamical route, which proves disorder and the relaxation times, do not cite each other.

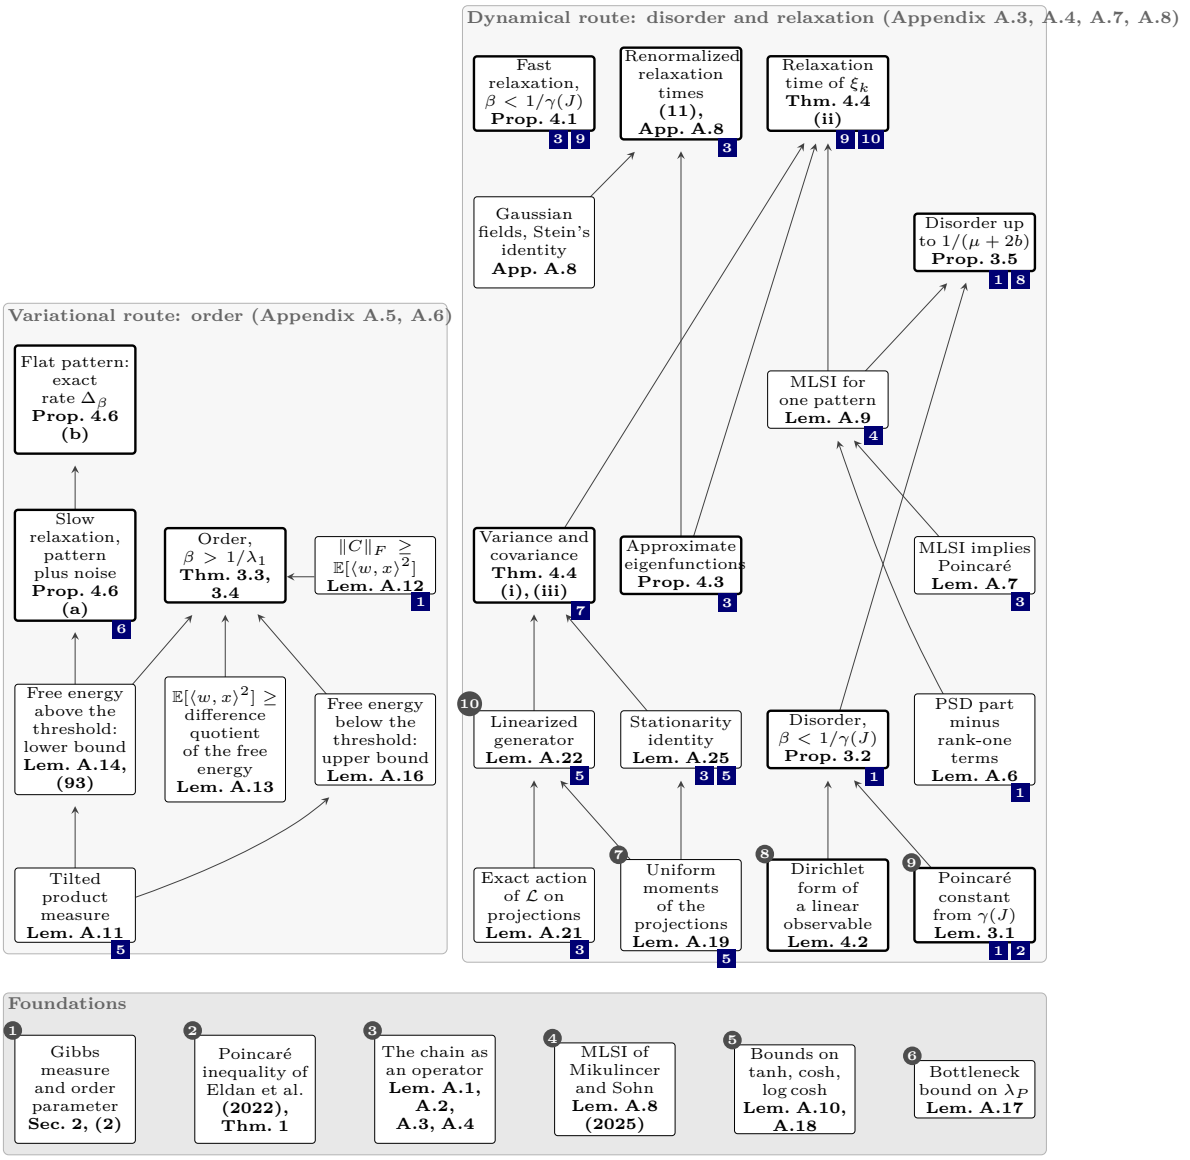

Figure 5: Organization of the proofs. An arrow from one block to another means that the proof of the second cites the first; arrows point upward, from the foundations to the results. Thick frames are statements of the main text; thin frames are lemmas of this appendix, facts of Section 2, and external theorems. The foundations are numbered in the order in which they appear in the paper and are cited by number only; above them an arrow is drawn only over a short distance, and a block cited from farther above is numbered as well. A number is a grey circle at the upper left corner of the block it names, and a blue square at the lower right corner of each block that cites it.

Foundations.

- *Gibbs measure and order parameter*, Sec. 2, (2). Since $x_i^2 = 1$, J and $J + cI$ define the same measure μ_β and the same chain, so the external inequalities can be applied to a shifted matrix that is positive semidefinite; and $\mathbb{E}[q_{ab}^2] = \|C\|_F^2/N^2$, so a bound on the correlation matrix C is a bound on the order parameter.

- *Poincaré inequality of Eldan et al.*, (2022), Thm. 1. External. For the measure proportional to $\exp(\frac{1}{2}x^\top Ax)$ with A positive semidefinite and $\|A\|_{\mathrm{op}} < 1$, the Poincaré constant of the chain is at least $1 - \|A\|_{\mathrm{op}}$.
- *The chain as an operator*, Lem. A.1, A.2, A.3, A.4. The generator $\mathcal{L}$ of the chain is self-adjoint in $L^2(\mu_\beta)$, $\mathbb{E}[\mathcal{L}\varphi] = 0$, $-\langle\varphi, \mathcal{L}\varphi\rangle = \mathcal{E}(\varphi, \varphi)$; λ_P is the smallest positive eigenvalue of $-\mathcal{L}$, and τ_φ is an average of the inverse eigenvalues weighted by the components of φ, hence $\tau_\varphi \geq \mathrm{Var}(\varphi)/\mathcal{E}(\varphi, \varphi)$.
- *MLSI of Mikulincer and Sohn*, Lem. A.8 (2025). External. The chain satisfies a modified log-Sobolev inequality (MLSI) when the matrix in the exponent is a positive semidefinite part of operator norm below $\frac{1}{2}$ minus one delocalized rank-one term.
- *Bounds on* tanh*,* cosh*,* log cosh, Lem. A.10, A.18. Second- and fourth-order bounds on $\log\cosh t$ and $t\tanh t$, $|\tanh t - t| \leq |t|^3/3$, and $\cosh t \leq e^{t^2/2}$, all valid for every real t.
- *Bottleneck bound on* λ_P, Lem. A.17. For a set A with $\mu_\beta(A) \leq \frac{1}{2}$, the indicator of A as test function gives $\lambda_P \leq N\mu_\beta(\partial A)/(2\mu_\beta(A))$, where ∂A is the set of configurations that a single flip moves into or out of A.

Variational route: order.

- *Tilted product measure*, Lem. A.11. For a unit vector w and the product measure with weights $e^{\lambda w_i x_i}$, the moment generating function of $\langle w, x\rangle$ under the uniform measure, and the mean and entropy of the tilted measure, with error terms $\lambda^3\kappa_4(w)$ and $\lambda^4\kappa_4(w)$; the entries of w may have any sign.
- *Free energy above the threshold: lower bound*, Lem. A.14, (93). By the Gibbs variational principle with the tilted measure, for $\theta > 1$, $F(\theta) \geq N\log 2 - \rho N/2 + \varepsilon^2/(12\kappa)$, where $\rho N/2$ bounds the bulk term; the equation is the bipartite form.
- *Free energy below the threshold: upper bound*, Lem. A.16. By the Gaussian identity $e^{s^2/2} = \mathbb{E}_g[e^{sg}]$, for $t < 1$, $F(t) \leq N\log 2 + \rho N/2 + \frac{1}{2}\log\frac{1}{1-t}$.
- $\mathbb{E}[\langle w, x\rangle^2] \geq$ *difference quotient of the free energy*, Lem. A.13. The free energy $F(t)$, with t the strength of the pattern, is convex and $F'(t) = \frac{1}{2}\mathbb{E}_t[\langle w, x\rangle^2]$; so this second moment is at least a difference quotient of F, and it is nondecreasing in t, which covers $\theta > 2$.
- $\|C\|_F \geq \mathbb{E}[\langle w, x\rangle^2]$, Lem. A.12. A lower bound on the order parameter needs the second moment along one unit vector w only.
- *Order,* $\beta > 1/\lambda_1$, Thm. 3.3, 3.4. The two free-energy bounds in the difference quotient give $\mathbb{E}[\langle w, x\rangle^2]$ of order N when $\kappa \leq M^2/N$ and ρ is small, and then the order parameter.
- *Slow relaxation, pattern plus noise*, Prop. 4.6 (a). The configurations in which the two block projections nearly cancel have total weight at most $2^N e^{\theta+\rho\sqrt{mn}}$, and the free-energy lower bound makes their probability exponentially small; the bottleneck bound then gives the relaxation time (12).
- *Flat pattern: exact rate* Δ_β, Prop. 4.6 (b). For the flat pattern binomial counting gives the exact rate Δ_β of that probability, and part (a) gives $\Delta_\beta \geq \varepsilon^2/6$.

Dynamical route: disorder and relaxation.

- *Dirichlet form of a linear observable*, Lem. 4.2. For $\varphi = \langle a, x\rangle$, $\mathcal{E}(\varphi, \varphi) \leq \|a\|^2$; with a Poincaré constant this bounds $a^\top Ca$. The lemma also states $\tau_\varphi \geq \mathrm{Var}(\varphi)/\mathcal{E}(\varphi, \varphi)$, which is the spectral form above.
- *Poincaré constant from* $\gamma(J)$, Lem. 3.1. This is the external Poincaré inequality applied to $\beta J + \beta|\lambda_N|I$: $\lambda_P \geq 1 - \beta\gamma(J)$.
- *Disorder,* $\beta < 1/\gamma(J)$, Prop. 3.2. A Poincaré constant bounds the variance of every linear observable, hence $\|C\|_{\mathrm{op}} \leq 1/\lambda_P$; then $\|C\|_F^2 \leq N\|C\|_{\mathrm{op}}$ and the overlap identity give the bound.
- *PSD part minus rank-one terms*, Lem. A.6. After a shift, $\frac{\beta}{2}J$ equals a positive semidefinite P with $\|P\|_{\mathrm{op}} \leq \frac{\beta}{2}(\mu_1 + 2b)$ minus one rank-one term per pattern.

- *MLSI implies Poincaré*, Lem. A.7. An MLSI with constant ϱ implies $\lambda_P \geq \varrho/2$; this is the source of the factor 2 in the bipartite bound.
- *MLSI for one pattern*, Lem. A.9. The external MLSI applies to the decomposition when $r = 1$ and gives $\lambda_P \geq (1 - o(1))(1 - \beta\mu - 2\beta b)/2$.
- *Disorder up to* $1/(\mu+2b)$, Prop. 3.5. The proof is that of the general disorder bound with this Poincaré constant.
- *Fast relaxation,* $\beta < 1/\gamma(J)$, Prop. 4.1. For $\beta < 1/\gamma(J)$ every τ_φ is at most $1/\lambda_P \leq 1/(1 - \beta\gamma(J))$.
- *Approximate eigenfunctions*, Prop. 4.3. If $\mathcal{L}\varphi + \alpha\varphi$ is small relative to φ, then τ_φ is close to $1/\alpha$; the proposition converts an approximate identity for $\mathcal{L}$ into bounds on τ_φ.
- *Uniform moments of the projections*, Lem. A.19. For $\theta_1 < 1$, by the Gaussian identity, the moments of $u_k^\top \sigma$ and $v_k^\top \tau$ under μ_β are bounded uniformly in m, n.
- *Exact action of* $\mathcal{L}$ *on projections*, Lem. A.21. $\mathcal{L}(u_k^\top \sigma) = -u_k^\top \sigma + \sum_i u_{k,i} \tanh(\beta(W\tau)_i)$, and the same with the two blocks exchanged; no approximation is made.
- *Linearized generator*, Lem. A.22. Replacing tanh by its argument gives $\mathcal{L}\xi_k = -(1 - \theta_k)\xi_k + R_k$ with $\|R_k\| \leq \sqrt{2} C_R \kappa_*$.
- *Stationarity identity*, Lem. A.25. The identity $\mathbb{E}[\mathcal{L}(fg)] = 0$, applied to a product of two projections, relates their second moment to $\mathcal{L}f$ and $\mathcal{L}g$; the lemma computes the remaining term.
- *Variance and covariance*, Thm. 4.4 (i), (iii). With the linearized generator these identities become linear equations for the second moments, whose solution is $\mathrm{Var}(\xi_k) = 1/(1 - \theta_k)$ and zero covariance, up to multiples of κ_*.
- *Relaxation time of* ξ_k, Thm. 4.4 (ii). This is the approximate-eigenfunction bound with $\alpha = 1 - \theta_k$; part (i) gives $\mathrm{Var}(\xi_k) \geq \frac{1}{2}$, and the two lower bounds on λ_P make the error small uniformly in m, n.
- *Gaussian fields, Stein's identity*, App. A.8. The local field $h_i = (W\tau)_i$ and a projection are treated as jointly Gaussian under μ_β; then $\tanh(\beta h_i)$ acts inside covariances as $\chi_i \beta h_i$ with $\chi_i = \mathbb{E}[1 - \tanh^2(\beta h_i)]$.
- *Renormalized relaxation times*, (11), App. A.8. When the bulk is not small, the exact action of $\mathcal{L}$ on the projections, which holds for every W, and the Gaussian step give the approximate eigenpair φ_k, $\alpha = 1 - \beta\mu_k\sqrt{\chi_\sigma \chi_\tau}$; the residual of the approximate-eigenfunction bound is measured on the trained matrix.

A.3 The chain as an operator

Let J be symmetric with zero diagonal, let $\beta > 0$, and write $\Omega = \{\pm 1\}^N$ for the configuration space. A function is a map $\Omega \to \mathbb{R}$, the functions form a real vector space of dimension 2^N, and $L^2(\mu_\beta)$ is this space equipped with the inner product

$$\langle \varphi, \psi \rangle = \mathbb{E}[\varphi\psi] = \sum_{x \in \Omega} \mu_\beta(x)\varphi(x)\psi(x). \tag{26}$$

Since $\mu_\beta(x) > 0$ for every x, (26) is an inner product, and $\|\varphi\|^2_{L^2(\mu_\beta)} = \langle \varphi, \varphi \rangle = \mathbb{E}[\varphi^2]$. The conditional distribution of x_i given $x_{\setminus i}$ is (18), whose mean is

$$\mathbb{E}[x_i \mid x_{\setminus i}] = \tanh\big(\beta\,(Jx)_i\big). \tag{27}$$

The conditional expectation operator. For a coordinate i and a function φ,

$$(P_i\varphi)(x) = \mathbb{E}\big[\varphi \mid x_{\setminus i}\big] = \sum_{s=\pm 1} P\big(x_i = s \mid x_{\setminus i}\big)\,\varphi\big([x_{\setminus i}, x_i = s]\big), \tag{28}$$

that is, the other $N-1$ coordinates of x are held fixed and φ is averaged over the two values of x_i with the conditional distribution (18). The function $P_i\varphi$ does not depend on x_i.

Lemma A.1. *(a)* $\mathbb{E}[P_i\varphi] = \mathbb{E}[\varphi]$. *(b)* P_i *is self-adjoint on* $L^2(\mu_\beta)$*:* $\langle \psi, P_i\varphi \rangle = \langle P_i\psi, P_i\varphi \rangle = \langle P_i\psi, \varphi \rangle$. *(c)* $\langle \varphi, \varphi \rangle - \langle \varphi, P_i\varphi \rangle = \mathbb{E}\big[\mathrm{Var}(\varphi \mid x_{\setminus i})\big]$.

Proof. (a) is the tower property: taking the conditional expectation over x_i first and then the expectation over $x_{\setminus i}$ gives the expectation of φ.

(b) $P_i\varphi$ is a function of $x_{\setminus i}$, hence a constant in the conditional expectation given $x_{\setminus i}$, so

$$\begin{aligned}\langle \psi, P_i\varphi\rangle = \mathbb{E}\big[\psi \cdot P_i\varphi\big] &= \mathbb{E}\Big[\mathbb{E}\big[\psi \cdot P_i\varphi \mid x_{\setminus i}\big]\Big] \\ &= \mathbb{E}\Big[\mathbb{E}\big[\psi \mid x_{\setminus i}\big] \cdot P_i\varphi\Big] = \mathbb{E}\big[(P_i\psi)(P_i\varphi)\big], \end{aligned} \tag{29}$$

where the second equality is the tower property and the third pulls the constant $P_i\varphi$ out of the conditional expectation. The right-hand side is symmetric in φ and ψ, so it also equals $\langle P_i\psi, \varphi\rangle$.

(c) Taking $\psi = \varphi$ in (b) gives $\langle \varphi, P_i\varphi\rangle = \mathbb{E}[(P_i\varphi)^2]$, and therefore

$$\begin{aligned}\langle \varphi, \varphi\rangle - \langle \varphi, P_i\varphi\rangle &= \mathbb{E}\big[\varphi^2\big] - \mathbb{E}\big[(P_i\varphi)^2\big] \\ &= \mathbb{E}\Big[\mathbb{E}\big[\varphi^2 \mid x_{\setminus i}\big] - \big(\mathbb{E}[\varphi \mid x_{\setminus i}]\big)^2\Big] \\ &= \mathbb{E}\big[\mathrm{Var}(\varphi \mid x_{\setminus i})\big], \end{aligned} \tag{30}$$

the middle step being the tower property applied to $\mathbb{E}[\varphi^2]$. □

The generator. From the configuration x, the clock of coordinate i rings at rate 1, and at that moment the chain jumps to $[x_{\setminus i}, x_i = s]$ with probability $P(x_i = s \mid x_{\setminus i})$. The rate of jumping from x to the configuration that differs from x exactly in coordinate i is therefore $P(x_i' = -x_i \mid x_{\setminus i})$, where x_i' denotes the resampled value, and the total rate of leaving x is the sum of these N rates. A continuous-time Markov chain on a finite state space is determined by its rate matrix $\mathcal{L}$, whose off-diagonal entries are the jump rates and whose diagonal entries are the negatives of the total rates of leaving; its transition semigroup is $e^{t\mathcal{L}}$, and for every function φ and every initial configuration x,

$$\mathbb{E}\big[\varphi(x_t) \mid x_0 = x\big] = (e^{t\mathcal{L}}\varphi)(x), \qquad \frac{\mathrm{d}}{\mathrm{d}t}\Big|_{t=0}\mathbb{E}\big[\varphi(x_t) \mid x_0 = x\big] = (\mathcal{L}\varphi)(x). \tag{31}$$

Written out on a function φ, that rate matrix is

$$\begin{aligned}(\mathcal{L}\varphi)(x) &= \sum_{i=1}^{N}\sum_{s=\pm 1} P\big(x_i = s \mid x_{\setminus i}\big)\big[\varphi([x_{\setminus i}, x_i = s]) - \varphi(x)\big] \\ &= \sum_{i=1}^{N}\big[(P_i\varphi)(x) - \varphi(x)\big], \end{aligned} \tag{32}$$

the second equality by (28) and $\sum_{s=\pm 1} P(x_i = s \mid x_{\setminus i}) = 1$. So $\mathcal{L} = \sum_{i=1}^{N}(P_i - I)$.

Lemma A.2. *(a) $\mathbb{E}[\mathcal{L}\varphi] = 0$ for every φ, and μ_β is the stationary distribution of the chain. (b) $\mathcal{L}$ is self-adjoint on $L^2(\mu_\beta)$, $-\mathcal{L}$ is positive semidefinite, and*

$$-\langle \varphi, \mathcal{L}\varphi\rangle = \sum_{i=1}^{N}\mathbb{E}\big[\mathrm{Var}(\varphi \mid x_{\setminus i})\big] = \mathcal{E}(\varphi, \varphi), \tag{33}$$

the Dirichlet form (19). *(c) $\mathcal{E}(\varphi, \varphi) = 0$ if and only if φ is constant.*

Proof. (a) By (32) and Lemma A.1 (a), $\mathbb{E}[\mathcal{L}\varphi] = \sum_i(\mathbb{E}[P_i\varphi] - \mathbb{E}[\varphi]) = 0$. By (31), $\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}_{\mu_\beta}[\varphi(x_t)] = \mathbb{E}_{\mu_\beta}[\mathcal{L}e^{t\mathcal{L}}\varphi] = 0$ for every φ, so $x_0 \sim \mu_\beta$ implies $x_t \sim \mu_\beta$.

(b) $\mathcal{L}$ is a sum of the self-adjoint operators $P_i - I$, hence self-adjoint. By (32) and Lemma A.1 (c),

$$-\langle \varphi, \mathcal{L}\varphi\rangle = \sum_{i=1}^{N}\big(\langle \varphi, \varphi\rangle - \langle \varphi, P_i\varphi\rangle\big) = \sum_{i=1}^{N}\mathbb{E}\big[\mathrm{Var}(\varphi \mid x_{\setminus i})\big] \geq 0, \tag{34}$$

and the last sum is $\mathcal{E}(\varphi,\varphi)$ by (19).

(c) If $\mathcal{E}(\varphi,\varphi)=0$, then $\mathbb{E}[\mathrm{Var}(\varphi \mid x_{\setminus i})]=0$ for each i; μ_β is positive everywhere, so $\mathrm{Var}(\varphi \mid x_{\setminus i})=0$ for every x and every i. The conditional distribution (18) gives both values $x_i=\pm 1$ a strictly positive probability, so $\varphi([x_{\setminus i},+1])=\varphi([x_{\setminus i},-1])$, that is, φ does not depend on the i-th coordinate. This holds for every i, so φ is constant. Conversely a constant has conditional variance zero. □

Spectral decomposition. By Lemma A.2 (b), $-\mathcal{L}$ has eigenpairs (ν_j, f_j), $j=0,1,\dots,2^N-1$, with the f_j orthonormal for the inner product (26),

$$-\mathcal{L}f_j=\nu_j f_j, \qquad 0=\nu_0<\nu_1\le\nu_2\le\cdots\le\nu_{2^N-1}, \qquad f_0=\mathbf{1}. \tag{35}$$

Here $\nu_0=0$ because $\mathcal{L}\mathbf{1}=\sum_i(P_i\mathbf{1}-\mathbf{1})=0$, and $\nu_1>0$ because by part (c) the eigenspace of the eigenvalue zero contains only the constants.

Lemma A.3. *Let $\varphi=\sum_{j\ge0}c_jf_j$ with $c_j=\langle\varphi,f_j\rangle$. Then $c_0=\mathbb{E}[\varphi]$ and*

$$\mathrm{Var}_{\mu_\beta}(\varphi)=\sum_{j\ge1}c_j^2, \qquad \mathcal{E}(\varphi,\varphi)=\sum_{j\ge1}\nu_jc_j^2. \tag{36}$$

Consequently $\lambda_P(\beta)=\nu_1$.

Proof. $c_0=\langle\varphi,\mathbf{1}\rangle=\mathbb{E}[\varphi]$. Orthonormality gives $\mathbb{E}[\varphi^2]=\sum_{j\ge0}c_j^2$, and subtracting $c_0^2=\mathbb{E}[\varphi]^2$ gives the first identity of (36). For the second, $-\mathcal{L}\varphi=\sum_{j\ge0}c_j\nu_jf_j$; taking the inner product with φ and using orthonormality gives $-\langle\varphi,\mathcal{L}\varphi\rangle=\sum_{j\ge1}\nu_jc_j^2$, which is $\mathcal{E}(\varphi,\varphi)$ by (33). By the two identities, $\lambda\,\mathrm{Var}_{\mu_\beta}(\varphi)\le\mathcal{E}(\varphi,\varphi)$ holds for every φ if and only if $\lambda\sum_{j\ge1}c_j^2\le\sum_{j\ge1}\nu_jc_j^2$ holds for every family of coefficients, that is, if and only if $\lambda\le\min_{j\ge1}\nu_j=\nu_1$. □

Lemma A.4. *Let $x_0\sim\mu_\beta$, let $\varphi=\sum_{j\ge0}c_jf_j$ and $\mathrm{Var}_{\mu_\beta}(\varphi)>0$. Then*

$$\mathrm{Cov}\big(\varphi(x_0),\varphi(x_t)\big)=\sum_{j\ge1}c_j^2e^{-\nu_jt}, \qquad \tau_\varphi=\frac{\sum_{j\ge1}c_j^2/\nu_j}{\sum_{j\ge1}c_j^2}. \tag{37}$$

Moreover: (a) $\tau_\varphi\ge\mathrm{Var}_{\mu_\beta}(\varphi)/\mathcal{E}(\varphi,\varphi)$; (b) $-\mathcal{L}$ is invertible on the subspace of functions with mean zero, where $(-\mathcal{L})^{-1}h=\sum_{j\ge1}(c_j/\nu_j)f_j$ for $h=\sum_{j\ge1}c_jf_j$, it satisfies $\|(-\mathcal{L})^{-1}h\|_{L^2(\mu_\beta)}\le\|h\|_{L^2(\mu_\beta)}/\nu_1$, and for φ with mean zero

$$\tau_\varphi=\frac{\langle\varphi,(-\mathcal{L})^{-1}\varphi\rangle}{\mathrm{Var}_{\mu_\beta}(\varphi)}. \tag{38}$$

Proof. By (31) and the tower property with x_0 fixed first,

$$\mathbb{E}\big[\varphi(x_0)\varphi(x_t)\big]=\mathbb{E}\Big[\varphi(x_0)\,\mathbb{E}\big[\varphi(x_t)\mid x_0\big]\Big]=\mathbb{E}_{\mu_\beta}\big[\varphi\cdot e^{t\mathcal{L}}\varphi\big]=\langle\varphi,e^{t\mathcal{L}}\varphi\rangle. \tag{39}$$

Here $e^{t\mathcal{L}}f_j=\sum_{l\ge0}\frac{t^l}{l!}\mathcal{L}^lf_j=e^{-\nu_jt}f_j$, so $\langle\varphi,e^{t\mathcal{L}}\varphi\rangle=\sum_{j\ge0}c_j^2e^{-\nu_jt}$; the term $j=0$ is $c_0^2=\mathbb{E}[\varphi]^2$, and subtracting it gives the covariance. Since $\nu_j>0$ for $j\ge1$, integrating term by term with $\int_0^\infty e^{-\nu_jt}\,\mathrm{d}t=1/\nu_j$ and dividing by $\mathrm{Var}_{\mu_\beta}(\varphi)$ gives τ_φ as stated in (37).

(a) Apply the Cauchy–Schwarz inequality to the sequences $(c_j\sqrt{\nu_j})_{j\ge1}$ and $(c_j/\sqrt{\nu_j})_{j\ge1}$:

$$\Big(\sum_{j\ge1}c_j^2\Big)^2=\Big(\sum_{j\ge1}\big(c_j\sqrt{\nu_j}\big)\big(c_j/\sqrt{\nu_j}\big)\Big)^2\le\Big(\sum_{j\ge1}\nu_jc_j^2\Big)\Big(\sum_{j\ge1}c_j^2/\nu_j\Big), \tag{40}$$

which by (36) and (37) reads $\mathrm{Var}_{\mu_\beta}(\varphi)^2\le\mathcal{E}(\varphi,\varphi)\cdot\tau_\varphi\,\mathrm{Var}_{\mu_\beta}(\varphi)$; divide by $\mathcal{E}(\varphi,\varphi)\,\mathrm{Var}_{\mu_\beta}(\varphi)>0$.

(b) A function with mean zero has $c_0=0$ by Lemma A.3, and $\nu_j>0$ for $j\ge1$, so the stated series defines the inverse. Its norm satisfies $\|(-\mathcal{L})^{-1}h\|^2_{L^2(\mu_\beta)}=\sum_{j\ge1}c_j^2/\nu_j^2\le\nu_1^{-2}\sum_{j\ge1}c_j^2$. For φ with mean zero, $\langle\varphi,(-\mathcal{L})^{-1}\varphi\rangle=\sum_{j\ge1}c_j^2/\nu_j$, which is the numerator in (37). □

Remark A.5. The chain acts on distributions by pushing them forward, $\mu_t = \mu_0 e^{t\mathcal{L}}$, and on observables by pulling them back, $\varphi_t = e^{t\mathcal{L}}\varphi$, with the same semigroup. By (37), τ_φ is the average of the times $1/\nu_j$ weighted by the variances c_j^2 of the components of φ; an observable that is an eigenfunction f_j has $\tau_\varphi = 1/\nu_j$.

Lemma 4.2. *For every observable φ, $\tau_\varphi \geq \mathrm{Var}_{\mu_\beta}(\varphi)/\mathcal{E}(\varphi,\varphi)$. For a linear observable $\varphi(x) = \langle a, x\rangle$, $\mathcal{E}(\varphi,\varphi) \leq \|a\|^2$, hence $\tau_\varphi \geq \mathrm{Var}_{\mu_\beta}(\varphi)/\|a\|^2$.*

Proof of Lemma 4.2. The first claim is Lemma A.4 (a). For the linear observable $\varphi(x) = \langle a, x\rangle$, only the term $a_i x_i$ is random conditionally on $x_{\setminus i}$, and by (27),

$$\mathrm{Var}\big(\varphi \mid x_{\setminus i}\big) = a_i^2 \,\mathrm{Var}\big(x_i \mid x_{\setminus i}\big) = a_i^2\big(1 - \tanh^2(\beta(Jx)_i)\big) \leq a_i^2. \tag{41}$$

Taking expectations and summing over i gives $\mathcal{E}(\varphi,\varphi) \leq \sum_{i=1}^N a_i^2 = \|a\|^2$ by (19), and the first claim then gives $\tau_\varphi \geq \mathrm{Var}_{\mu_\beta}(\varphi)/\|a\|^2$. □

Proposition 4.1 (Fast relaxation for $\beta < 1/\gamma(J)$)**.** *Let J be symmetric with zero diagonal and $\beta < 1/\gamma(J)$. Then the relaxation time of the chain is at most $1/(1-\beta\gamma(J))$, and $\tau_\varphi \leq 1/(1-\beta\gamma(J))$ for every observable φ.*

Proof of Proposition 4.1. By Lemma A.3 the Poincaré constant is ν_1, so the relaxation time of the chain is $1/\nu_1$, and Lemma 3.1 gives $\nu_1 \geq 1-\beta\gamma(J)$, that is, $1/\nu_1 \leq 1/(1-\beta\gamma(J))$. For an observable φ with $\mathrm{Var}_{\mu_\beta}(\varphi) > 0$, write $c_j = \langle \varphi, f_j\rangle$; each $j \geq 1$ has $1/\nu_j \leq 1/\nu_1$, so (37) gives

$$\tau_\varphi = \frac{\sum_{j\geq 1} c_j^2/\nu_j}{\sum_{j\geq 1} c_j^2} \leq \frac{1}{\nu_1} \leq \frac{1}{1-\beta\gamma(J)}. \tag{42}$$

□

Proposition 4.3 (Approximate eigenfunctions)**.** *Let φ be an observable with $\mathbb{E}[\varphi] = 0$ and $\mathrm{Var}_{\mu_\beta}(\varphi) > 0$, let $\alpha > 0$, and let*

$$R = \mathcal{L}\varphi + \alpha\varphi, \qquad \eta = \frac{\|R\|_{L^2(\mu_\beta)}}{\|\varphi\|_{L^2(\mu_\beta)}}. \tag{9}$$

Then $\tau_\varphi \geq 1/(\alpha+\eta)$, and $\tau_\varphi = \alpha^{-1}(1+\delta)$ with $|\delta| \leq \eta/\lambda_P(\beta)$.

Proof of Proposition 4.3. Let φ, α, $\mathcal{R} = \mathcal{L}\varphi + \alpha\varphi$ and $\eta = \|\mathcal{R}\|_{L^2(\mu_\beta)}/\|\varphi\|_{L^2(\mu_\beta)}$ be as in the statement. Since $\mathbb{E}[\varphi] = 0$, $\mathrm{Var}_{\mu_\beta}(\varphi) = \|\varphi\|^2_{L^2(\mu_\beta)}$, and this is positive. By Lemma A.2 (a), $\mathbb{E}[\mathcal{L}\varphi] = 0$, so $\mathbb{E}[\mathcal{R}] = \mathbb{E}[\mathcal{L}\varphi] + \alpha\mathbb{E}[\varphi] = 0$.

For (1), write $\mathcal{L}\varphi = \mathcal{R} - \alpha\varphi$ and compute

$$\begin{aligned}
\mathcal{E}(\varphi,\varphi) &= -\langle \varphi, \mathcal{L}\varphi\rangle && \text{by (33)} \\
&= -\langle \varphi, \mathcal{R} - \alpha\varphi\rangle && \text{by the definition of } \mathcal{R} \\
&= \alpha\|\varphi\|^2_{L^2(\mu_\beta)} - \langle \varphi, \mathcal{R}\rangle && \text{by bilinearity} \\
&\leq \alpha\|\varphi\|^2_{L^2(\mu_\beta)} + \|\varphi\|_{L^2(\mu_\beta)}\|\mathcal{R}\|_{L^2(\mu_\beta)} && \text{by Cauchy–Schwarz} \\
&= (\alpha+\eta)\|\varphi\|^2_{L^2(\mu_\beta)} && \text{by the definition of } \eta.
\end{aligned} \tag{43}$$

Lemma A.4 (a) and $\mathrm{Var}_{\mu_\beta}(\varphi) = \|\varphi\|^2_{L^2(\mu_\beta)}$ then give

$$\tau_\varphi \geq \frac{\mathrm{Var}_{\mu_\beta}(\varphi)}{\mathcal{E}(\varphi,\varphi)} \geq \frac{\|\varphi\|^2_{L^2(\mu_\beta)}}{(\alpha+\eta)\|\varphi\|^2_{L^2(\mu_\beta)}} = \frac{1}{\alpha+\eta}. \tag{44}$$

For (2), write the same identity as $-\mathcal{L}\varphi = \alpha\varphi - \mathcal{R}$. Both sides have mean zero, so Lemma A.4 (b) applies to them:

$$\varphi = \alpha(-\mathcal{L})^{-1}\varphi - (-\mathcal{L})^{-1}\mathcal{R}. \tag{45}$$

Taking the inner product with φ,

$$\|\varphi\|^2_{L^2(\mu_\beta)} = \alpha\langle\varphi, (-\mathcal{L})^{-1}\varphi\rangle - \langle\varphi, (-\mathcal{L})^{-1}\mathcal{R}\rangle, \tag{46}$$

and solving for the first inner product on the right,

$$\langle\varphi, (-\mathcal{L})^{-1}\varphi\rangle = \frac{\|\varphi\|^2_{L^2(\mu_\beta)} + \langle\varphi, (-\mathcal{L})^{-1}\mathcal{R}\rangle}{\alpha}. \tag{47}$$

Substituting this into (38) and using $\mathrm{Var}_{\mu_\beta}(\varphi) = \|\varphi\|^2_{L^2(\mu_\beta)}$,

$$\tau_\varphi = \frac{\langle\varphi, (-\mathcal{L})^{-1}\varphi\rangle}{\mathrm{Var}_{\mu_\beta}(\varphi)} = \frac{1}{\alpha}\big(1+\delta\big), \qquad \delta = \frac{\langle\varphi, (-\mathcal{L})^{-1}\mathcal{R}\rangle}{\|\varphi\|^2_{L^2(\mu_\beta)}}. \tag{48}$$

By the Cauchy–Schwarz inequality, the norm bound of Lemma A.4 (b) and $\lambda_P(\beta) = \nu_1$,

$$|\delta| \le \frac{\|\varphi\|_{L^2(\mu_\beta)}\|(-\mathcal{L})^{-1}\mathcal{R}\|_{L^2(\mu_\beta)}}{\|\varphi\|^2_{L^2(\mu_\beta)}} \le \frac{\|\mathcal{R}\|_{L^2(\mu_\beta)}}{\nu_1\|\varphi\|_{L^2(\mu_\beta)}} = \frac{\eta}{\nu_1} = \frac{\eta}{\lambda_P(\beta)}. \tag{49}$$

□

A.4 Disorder: proofs of Propositions 3.2 and 3.5

Proposition 3.2 (Disorder for $\beta < 1/\gamma(J)$)**.** *Let J be symmetric with zero diagonal and $\beta < 1/\gamma(J)$. Then*

$$\|C\|_{\mathrm{op}} \le \frac{1}{1-\beta\gamma(J)}, \qquad N\,\mathbb{E}[q_{ab}^2] \le \frac{1}{1-\beta\gamma(J)}. \tag{4}$$

Proof of Proposition 3.2. By (2), $\mathbb{E}[q_{ab}^2] = \|C\|_F^2/N^2$.

The matrix C is the Gram matrix of the coordinate functions $x_1, \dots, x_N$ in $L^2(\mu_\beta)$: for every $a \in \mathbb{R}^N$,

$$a^\top C a = \sum_{i,j} a_i a_j\,\mathbb{E}[x_i x_j] = \mathbb{E}\big[\langle a, x\rangle^2\big] \;\ge\; 0, \tag{50}$$

so C is positive semidefinite, and $C_{ii} = \mathbb{E}[x_i^2] = 1$ gives $\operatorname{tr} C = \sum_i C_{ii} = N$. Write $\nu_1 \ge \cdots \ge \nu_N \ge 0$ for the eigenvalues of C. The Frobenius norm is invariant under orthogonal conjugation, so $\|C\|_F^2 = \sum_k \nu_k^2$; and $0 \le \nu_k \le \nu_1$ gives $\nu_k^2 \le \nu_1\nu_k$ for every k, so

$$\|C\|_F^2 = \sum_{k=1}^N \nu_k^2 \;\le\; \nu_1 \sum_{k=1}^N \nu_k = \|C\|_{\mathrm{op}} \operatorname{tr} C = N\|C\|_{\mathrm{op}}. \tag{51}$$

Take a unit vector $a \in \mathbb{R}^N$ and the linear observable $\varphi(x) = \langle a, x\rangle$. Since $\mathbb{E}[x] = 0$,

$$\mathrm{Var}_{\mu_\beta}(\varphi) = \mathbb{E}\big[\langle a, x\rangle^2\big] = a^\top C a, \tag{52}$$

the second equality by (50). Since $\beta < 1/\gamma(J)$, Lemma 3.1 applies to φ and gives

$$\big(1-\beta\gamma(J)\big)\, a^\top C a \;\le\; \mathcal{E}(\varphi,\varphi). \tag{53}$$

The second claim of Lemma 4.2, proved in Appendix A.3, bounds the Dirichlet form of a linear observable: $\mathcal{E}(\varphi,\varphi) \le \|a\|^2 = 1$. The two displays give $a^\top C a \le 1/(1-\beta\gamma(J))$. The matrix C is symmetric and positive semidefinite, so $\|C\|_{\mathrm{op}} = \max_{\|a\|=1} a^\top C a$, and

$$\|C\|_{\mathrm{op}} \;\le\; \frac{1}{1-\beta\gamma(J)}, \tag{54}$$

which is the first bound of (4). By (2) and (51),

$$N\,\mathbb{E}[q_{ab}^2] = \frac{\|C\|_F^2}{N} \;\le\; \|C\|_{\mathrm{op}} \;\le\; \frac{1}{1-\beta\gamma(J)}, \tag{55}$$

which is the second bound of (4). □

The vectors $w_k^{\pm}$ and the matrix J_{bulk}. Let $W = \sum_{k\le r} \mu_k u_k v_k^\top + W_{\text{bulk}}$ be as in (24), with $(u_k)_{k\le r}$ orthonormal in $\mathbb{R}^m$, $(v_k)_{k\le r}$ orthonormal in $\mathbb{R}^n$, and $b = \|W_{\text{bulk}}\|_{\text{op}}$. Write $\theta_k = \beta\mu_k$, and

$$w_k^{\pm} = \frac{(u_k, \pm v_k)}{\sqrt{2}} \in \mathbb{R}^N, \qquad J_{\text{bulk}} = \begin{pmatrix} 0 & W_{\text{bulk}} \\ W_{\text{bulk}}^\top & 0 \end{pmatrix}. \tag{56}$$

Each $w_k^{\pm}$ is a unit vector, and $\langle w_j^+, w_k^+\rangle = (u_j^\top u_k + v_j^\top v_k)/2 = \delta_{jk}$, so $(w_k^+)_{k\le r}$ is orthonormal. Squaring the block matrix gives

$$J_{\text{bulk}}^2 = \begin{pmatrix} W_{\text{bulk}} W_{\text{bulk}}^\top & 0 \\ 0 & W_{\text{bulk}}^\top W_{\text{bulk}} \end{pmatrix}, \tag{57}$$

whose largest eigenvalue is the square of the largest singular value of W_{bulk}, so $\|J_{\text{bulk}}\|_{\text{op}} = b$.

Lemma A.6. *Let*

$$P = \frac{\beta}{2}\Big(\sum_{k\le r} \mu_k w_k^+ w_k^{+\top} + J_{\text{bulk}} + bI\Big). \tag{58}$$

Then

$$\frac{\beta}{2} J + \frac{\beta b}{2} I = P - \sum_{k\le r} \frac{\theta_k}{2}\, w_k^- w_k^{-\top}, \tag{59}$$

P is positive semidefinite, and $\|P\|_{\text{op}} \le \frac{\beta}{2}(\mu_1 + 2b)$. Adding $\frac{\beta b}{2} I$ to the quadratic form changes neither μ_β nor the chain (18).

Proof. Multiplying the block vectors of (56) out,

$$w_k^+ w_k^{+\top} = \frac{1}{2}\begin{pmatrix} u_k u_k^\top & u_k v_k^\top \\ v_k u_k^\top & v_k v_k^\top \end{pmatrix}, \qquad w_k^- w_k^{-\top} = \frac{1}{2}\begin{pmatrix} u_k u_k^\top & -u_k v_k^\top \\ -v_k u_k^\top & v_k v_k^\top \end{pmatrix},$$
$$w_k^+ w_k^{+\top} - w_k^- w_k^{-\top} = \begin{pmatrix} 0 & u_k v_k^\top \\ v_k u_k^\top & 0 \end{pmatrix}. \tag{60}$$

Multiplying (60) by μ_k, summing over $k \le r$ and adding J_{bulk} gives, by (21) and (24),

$$J = \sum_{k\le r} \mu_k\big(w_k^+ w_k^{+\top} - w_k^- w_k^{-\top}\big) + J_{\text{bulk}}. \tag{61}$$

Multiplying (61) by $\beta/2$, adding $\frac{\beta b}{2} I$ and using $\theta_k = \beta\mu_k$ gives (59) with P of (58).

The eigenvalues of J_{bulk} are the singular values of W_{bulk}, their negatives and 0, all in $[-b, b]$, so the eigenvalues of $J_{\text{bulk}} + bI$ lie in $[0, 2b]$: the matrix $J_{\text{bulk}} + bI$ is positive semidefinite and $\|J_{\text{bulk}} + bI\|_{\text{op}} \le 2b$. Each $\mu_k w_k^+ w_k^{+\top}$ is positive semidefinite because $\mu_k \ge 0$, so P is $\beta/2$ times a sum of positive semidefinite matrices. The family $(w_k^+)_{k\le r}$ is orthonormal, so $\sum_k \mu_k w_k^+ w_k^{+\top}$ has eigenvalues μ_k and 0 and $\|\sum_k \mu_k w_k^+ w_k^{+\top}\|_{\text{op}} = \mu_1$; the triangle inequality then gives $\|P\|_{\text{op}} \le \frac{\beta}{2}(\mu_1 + 2b)$.

For the shift, $x_i^2 = 1$ gives, for every real c and every symmetric A,

$$x^\top (A + cI) x = x^\top A x + cN, \tag{62}$$

and the constant cN cancels between the numerator and Z_β in (13), so A and $A + cI$ define the same measure. The conditional distribution (18) of x_i depends on the matrix only through its off-diagonal entries in row i, which $A \mapsto A + cI$ leaves unchanged. □

The modified log-Sobolev inequality. For $f, g : \{\pm 1\}^N \to \mathbb{R}$ the bilinear form of the chain is

$$\mathcal{E}(f, g) := -\langle f, \mathcal{L} g\rangle = \sum_{i=1}^N \mathbb{E}\big[\mathrm{Cov}(f, g \mid x_{\setminus i})\big], \tag{63}$$

where $\langle \cdot, \cdot \rangle$ is the inner product of $L^2(\mu_\beta)$, $\mathcal{L}$ is the generator (32), and the second equality is the computation of Lemma A.2(b), which uses Lemma A.1, with φ^2 replaced by fg. Taking $g = f = \varphi$ returns (19). The generator $\mathcal{L}$ is self-adjoint in $L^2(\mu_\beta)$ by Lemma A.2(b), so $\mathcal{E}(f, g) = \mathcal{E}(g, f)$. For $f : \{\pm 1\}^N \to (0, \infty)$ the entropy is

$$\mathrm{Ent}(f) := \mathbb{E}[f \log f] - \mathbb{E}[f] \log \mathbb{E}[f]. \tag{64}$$

The chain satisfies the modified log-Sobolev inequality with constant ϱ if

$$\varrho \, \mathrm{Ent}(f) \le \mathcal{E}(f, \log f) \qquad \text{for every } f > 0. \tag{65}$$

Lemma A.7. *If the chain satisfies* (65) *with constant ϱ, then $\lambda_P \ge \varrho/2$.*

Proof. We prove $(\varrho/2) \, \mathrm{Var}_{\mu_\beta}(\varphi) \le \mathcal{E}(\varphi, \varphi)$ for every observable φ. Under $\varphi \mapsto a\varphi + c$ with a, c constants both sides are multiplied by a^2, so it suffices to prove it for φ with $\mathbb{E}[\varphi] = 0$ and $\|\varphi\|_\infty \le 1$. Fix $0 < s < \frac{1}{2}$ and put $f = 1 + s\varphi$; then f takes values in $[\frac{1}{2}, \frac{3}{2}]$ and $\mathbb{E}[f] = 1$.

The first step expands the entropy. Since $\mathbb{E}[f] = 1$, (64) gives $\mathrm{Ent}(f) = \mathbb{E}[f \log f]$. For $|y| \le \frac{1}{2}$ the series $(1 + y) \log(1 + y) = y + \sum_{k \ge 1} (-1)^{k+1} y^{k+1} / (k(k + 1))$ gives

$$(1 + y) \log(1 + y) = y + \frac{y^2}{2} + r_1(y), \qquad |r_1(y)| \le \sum_{k \ge 2} \frac{|y|^{k+1}}{6} \le \frac{|y|^3}{6(1 - |y|)} \le \frac{|y|^3}{3}, \tag{66}$$

the first inequality because $k(k+1) \ge 6$ for $k \ge 2$. Taking $y = s\varphi(x)$, which satisfies $|y| \le \frac{1}{2}$, and taking the expectation over $x \sim \mu_\beta$, the linear term drops because $\mathbb{E}[s\varphi] = 0$:

$$\mathrm{Ent}(f) = \frac{s^2}{2} \mathbb{E}[\varphi^2] + \mathbb{E}[r_1(s\varphi)], \qquad \big|\mathbb{E}[r_1(s\varphi)]\big| \le \frac{s^3}{3}. \tag{67}$$

The second step expands the bilinear form. For $|y| \le \frac{1}{2}$,

$$\log(1 + y) = y - \frac{y^2}{2} + r_2(y), \qquad |r_2(y)| \le \sum_{k \ge 3} \frac{|y|^k}{3} \le \frac{|y|^3}{3(1 - |y|)} \le \frac{2|y|^3}{3}, \tag{68}$$

so $\log f = s\varphi - \frac{s^2}{2}\varphi^2 + r_2(s\varphi)$. For every g, $\mathcal{E}(1, g) = -\mathbb{E}[\mathcal{L}g] = 0$ by Lemma A.2(a). Expanding $\mathcal{E}(f, \log f)$ by bilinearity and dropping the terms with first argument 1,

$$\mathcal{E}(f, \log f) = s^2 \mathcal{E}(\varphi, \varphi) - \frac{s^3}{2} \mathcal{E}(\varphi, \varphi^2) + s \, \mathcal{E}\big(\varphi, r_2(s\varphi)\big). \tag{69}$$

The form $\mathcal{E}$ is symmetric and positive semidefinite by (63), so $|\mathcal{E}(g, h)| \le \mathcal{E}(g, g)^{1/2} \mathcal{E}(h, h)^{1/2}$; and $\mathrm{Var}(g \mid x_{\setminus i}) \le \mathbb{E}[g^2 \mid x_{\setminus i}] \le \|g\|_\infty^2$ gives $\mathcal{E}(g, g) \le N \|g\|_\infty^2$. With $\|\varphi\|_\infty \le 1$, hence $\|\varphi^2\|_\infty \le 1$, and with $\|r_2(s\varphi)\|_\infty \le 2s^3/3$ by (68), the last two terms of (69) satisfy $|\mathcal{E}(\varphi, \varphi^2)| \le N$ and $|\mathcal{E}(\varphi, r_2(s\varphi))| \le N \cdot 2s^3/3$, so

$$\big|\mathcal{E}(f, \log f) - s^2 \mathcal{E}(\varphi, \varphi)\big| \le \frac{N s^3}{2} + \frac{2N s^4}{3} \le 2N s^3. \tag{70}$$

The third step substitutes (67) and (70) into (65) and divides by s^2:

$$\frac{\varrho}{2} \mathbb{E}[\varphi^2] - \frac{\varrho s}{3} \le \mathcal{E}(\varphi, \varphi) + 2Ns. \tag{71}$$

Letting $s \downarrow 0$ gives $\frac{\varrho}{2} \mathbb{E}[\varphi^2] \le \mathcal{E}(\varphi, \varphi)$, and $\mathbb{E}[\varphi] = 0$ gives $\mathbb{E}[\varphi^2] = \mathrm{Var}_{\mu_\beta}(\varphi)$. □

Lemma A.8 (Mikulincer & Sohn (2025), Theorem 1)**.** *Let ν be the measure on $\{\pm 1\}^M$*

$$\nu(x) \propto \exp\Big(-\frac{\beta'}{M} \langle \tilde{u}, x \rangle^2 + \langle x, Ax \rangle + \langle h, x \rangle \Big), \tag{72}$$

with A positive semidefinite, $\|A\|_{\mathrm{op}} < \frac{1}{2}$, $\|\tilde{u}\|_\infty \le 1$, $h \in \mathbb{R}^M$ arbitrary, and $0 \le \beta' \le M^{1-\delta}$ for a fixed $\delta > 0$. Then the chain that resamples one coordinate at a time from the conditional distributions of ν satisfies the modified log-Sobolev inequality (65) *with a constant*

$$\varrho \ge (1 - o(1)) \big(1 - 2\|A\|_{\mathrm{op}}\big), \tag{73}$$

where $o(1)$ depends only on M and δ, and the constant is normalized by the continuous-time bilinear form (63) *in which each coordinate carries a rate-one clock.*

Lemma A.9. *Let $r = 1$ in (24), so that $W = \mu uv^\top + W_{\mathrm{bulk}}$, and write $\theta = \beta\mu$. Assume $\|u\|_\infty, \|v\|_\infty \le N^{-\delta}$ for a fixed $\delta > 0$ and $\beta(\mu + 2b) < 1$. Then the chain satisfies (65) with a constant*

$$\varrho \ge (1 - o(1))\,(1 - \beta\mu - 2\beta b), \qquad \textit{hence} \qquad \lambda_P \ge \frac{(1 - o(1))(1 - \beta\mu - 2\beta b)}{2}, \tag{74}$$

where $o(1)$ depends only on N and δ.

Proof. Write $w^- = w_1^-$ and let P be the matrix (58) with $r = 1$. By (13), (59) and the shift of Lemma A.6,

$$\mu_\beta(x) \propto \exp\Big(\frac{\beta}{2} x^\top J x\Big) = \exp\Big(\langle x, Px\rangle - \frac{\theta}{2}\langle w^-, x\rangle^2\Big). \tag{75}$$

Put $\tilde u = w^-/\|w^-\|_\infty$, so that $\|\tilde u\|_\infty = 1$ and

$$\frac{\theta}{2}\langle w^-, x\rangle^2 = \frac{\theta}{2}\|w^-\|_\infty^2 \langle \tilde u, x\rangle^2 = \frac{\beta'}{N}\langle \tilde u, x\rangle^2, \qquad \beta' = \frac{\theta}{2} N \|w^-\|_\infty^2. \tag{76}$$

Every entry of w^- is an entry of u, or the negative of an entry of v, divided by $\sqrt{2}$, so $\|w^-\|_\infty = \max(\|u\|_\infty, \|v\|_\infty)/\sqrt{2}$ is at most $N^{-\delta}/\sqrt{2}$ and

$$\beta' \le \frac{\theta}{4} N^{1-2\delta} \le N^{1-\delta} \qquad \text{once } N^\delta \ge \frac{\theta}{4}. \tag{77}$$

By (75) and (76), μ_β is the measure (72) with $M = N$, $A = P$ and $h = 0$. Its hypotheses hold: P is positive semidefinite and $\|P\|_{\mathrm{op}} \le \frac{\beta}{2}(\mu + 2b) < \frac{1}{2}$ by Lemma A.6 and $\beta(\mu + 2b) < 1$, and $\beta' \le N^{1-\delta}$ by (77). Lemma A.8 gives

$$\varrho \ge (1 - o(1))\big(1 - 2\|P\|_{\mathrm{op}}\big) \ge (1 - o(1))\,(1 - \beta\mu - 2\beta b), \tag{78}$$

and Lemma A.7 gives $\lambda_P \ge \varrho/2$. For the flat pattern, $m = n$ and $u = v = n^{-1/2}\mathbf{1}$, every entry of w^- equals $\pm(2n)^{-1/2} = \pm N^{-1/2}$, so (76) gives $\beta' = \theta/2 \le 1 \le N^{1-\delta}$ for every N and every $\delta \in (0, 1)$, and with $W_{\mathrm{bulk}} = 0$ the bound $\lambda_P \ge (1 - o(1))(1 - \theta)/2$ holds for every $\theta < 1$, which is the case used in Theorem 4.4(ii). □

Proposition 3.5 (Disorder up to the pattern strength)**.** *Let $W = \mu uv^\top + W_{\mathrm{bulk}}$ with $\|u\|_\infty, \|v\|_\infty \le N^{-\delta}$ for a fixed $\delta > 0$, and $b = \|W_{\mathrm{bulk}}\|_{\mathrm{op}}$. If $\beta < 1/(\mu + 2b)$, then*

$$\|C\|_{\mathrm{op}} \le \frac{2\,(1 + o(1))}{1 - \beta\mu - 2\beta b}, \qquad N\,\mathbb{E}[q_{ab}^2] \le \frac{2\,(1 + o(1))}{1 - \beta\mu - 2\beta b}, \tag{8}$$

where $o(1)$ depends only on N and δ.

Proof of Proposition 3.5. By Lemma A.9, $\lambda_P \ge (1 - o(1))(1 - \beta\mu - 2\beta b)/2$. Take a unit vector $a \in \mathbb{R}^N$ and the linear observable $\varphi(x) = \langle a, x\rangle$. By (52), the definition of λ_P and the second claim of Lemma 4.2,

$$a^\top C a = \mathrm{Var}_{\mu_\beta}(\varphi) \;\le\; \frac{\mathcal{E}(\varphi, \varphi)}{\lambda_P} \;\le\; \frac{\|a\|^2}{\lambda_P} = \frac{1}{\lambda_P}. \tag{79}$$

The matrix C is symmetric and positive semidefinite by (50), so $\|C\|_{\mathrm{op}} = \max_{\|a\|=1} a^\top C a$ and, with $(1 - o(1))^{-1} = 1 + o(1)$,

$$\|C\|_{\mathrm{op}} \;\le\; \frac{1}{\lambda_P} \;\le\; \frac{2\,(1 + o(1))}{1 - \beta\mu - 2\beta b}, \tag{80}$$

which is the first bound of (8). By (2) and (51),

$$N\,\mathbb{E}[q_{ab}^2] = \frac{\|C\|_F^2}{N} \;\le\; \|C\|_{\mathrm{op}} \;\le\; \frac{2\,(1 + o(1))}{1 - \beta\mu - 2\beta b}, \tag{81}$$

which is the second bound of (8). □

Several patterns. Lemma A.6 and Lemma A.7 are stated for r patterns, and the proof of Proposition 3.5 uses the Poincaré constant of the chain and nothing else about r. The one place where $r = 1$ enters is Lemma A.8, the theorem of Mikulincer & Sohn (2025), which is stated for a single negative rank-one term; its authors write that their localization and Gaussian approximation framework should extend to several outliers, at the cost of additional technical work (Mikulincer & Sohn, 2025, Section 1.4). With that extension the proof above gives disorder for $\beta < 1/(\mu_1 + 2b)$ with r patterns, which is where the transition is measured in Appendix B.1.

A.5 Order: proofs of Theorems 3.3 and 3.4

Let $J = \lambda_1 w w^\top + J_{\text{bulk}}$ with $\|w\| = 1$, and let $\theta = \beta\lambda_1$, $\varepsilon = \theta - 1$, $\rho = \beta\|J_{\text{bulk}}\|_{\text{op}}$, $\kappa = \sum_i w_i^4$ and $X = \langle w, x\rangle$ be as in Theorem 3.3. Substituting J into the exponent of (13),

$$\tfrac{\beta}{2}\, x^\top J x = \tfrac{\theta}{2}\, X^2 + \tfrac{\beta}{2}\, x^\top J_{\text{bulk}} x, \qquad \left|\tfrac{\beta}{2}\, x^\top J_{\text{bulk}} x\right| \le \tfrac{\beta}{2}\, \|J_{\text{bulk}}\|_{\text{op}} \|x\|^2 = \frac{\rho N}{2}, \tag{82}$$

where the last step uses $\|x\|^2 = N$ for $x \in \{\pm 1\}^N$. We write P_0 for the uniform measure on $\{\pm 1\}^N$ and $\mathbb{E}_0$ for its expectation; under P_0 the coordinates are independent and uniform on $\{\pm 1\}$.

Lemma A.10. *For every $t \in \mathbb{R}$: (a) $t \tanh t \le t^2$; (b) $t \tanh t \ge t^2 - t^4/3$; (c) $\log\cosh t \le t^2/2$; (d) $\log\cosh t \ge t^2/2 - t^4/12$.*

Proof. All four functions are even, so let $t \ge 0$.

(a) The function $f(t) = t - \tanh t$ has $f(0) = 0$ and $f'(t) = 1 - \operatorname{sech}^2 t = \tanh^2 t \ge 0$. Hence $\tanh t \le t$, and multiplying by $t \ge 0$ gives (a).

(b) The function $f(t) = \tanh t - t + t^3/3$ has $f(0) = 0$ and $f'(t) = \operatorname{sech}^2 t - 1 + t^2 = t^2 - \tanh^2 t$, which is nonnegative by (a). Hence $\tanh t \ge t - t^3/3$, and multiplying by $t \ge 0$ gives (b).

(c) The function $f(t) = t^2/2 - \log\cosh t$ has $f(0) = 0$ and $f'(t) = t - \tanh t \ge 0$ by (a).

(d) The function $f(t) = \log\cosh t - t^2/2 + t^4/12$ has $f(0) = 0$ and $f'(t) = \tanh t - t + t^3/3$, which is nonnegative by the proof of (b). □

For a unit vector $w \in \mathbb{R}^d$ and $\lambda \in \mathbb{R}$ let

$$\psi_w(\lambda) = \sum_{i=1}^{d} \log\cosh(\lambda w_i). \tag{83}$$

The next lemma is used with $d = N$ for Theorem 3.3, and with $d = m$ and $d = n$ for Theorem 3.4.

Lemma A.11. *Let $w \in \mathbb{R}^d$ be a unit vector, $\psi = \psi_w$, $\kappa_4(w) = \sum_i w_i^4$, and $X = \langle w, x\rangle$ for $x \in \{\pm 1\}^d$.*

(a) $\mathbb{E}_0[e^{\lambda X}] = e^{\psi(\lambda)}$, and $\psi'(\lambda) = \sum_i w_i \tanh(\lambda w_i)$.

(b) For $\lambda \ge 0$, $\lambda - \lambda^3\kappa_4(w)/3 \le \psi'(\lambda) \le \lambda$ and $\lambda^2/2 - \lambda^4\kappa_4(w)/12 \le \psi(\lambda) \le \lambda^2/2$. The upper bound $\psi(\lambda) \le \lambda^2/2$ holds for every $\lambda \in \mathbb{R}$.

(c) The measure $\mathbb{Q}_\lambda(x) = \prod_{i=1}^d \dfrac{e^{\lambda w_i x_i}}{2\cosh(\lambda w_i)}$ is a probability measure on $\{\pm 1\}^d$ with independent coordinates, $\mathbb{E}_{\mathbb{Q}_\lambda}[x_i] = \tanh(\lambda w_i)$, $\mathbb{E}_{\mathbb{Q}_\lambda}[X] = \psi'(\lambda)$, $\|\mathbb{E}_{\mathbb{Q}_\lambda}[x]\| \le \sqrt{d}$, and entropy

$$H(\mathbb{Q}_\lambda) = -\sum_x \mathbb{Q}_\lambda(x) \log \mathbb{Q}_\lambda(x) = d\log 2 + \psi(\lambda) - \lambda\psi'(\lambda). \tag{84}$$

Proof. (a) Under P_0 the coordinates are independent and $\mathbb{E}_0[e^{\lambda w_i x_i}] = \cosh(\lambda w_i)$; the product over i is $e^{\psi(\lambda)}$. Differentiating (83) term by term gives ψ'.

(b) Let $t_i = \lambda w_i$. Since $\|w\| = 1$, $\sum_i t_i^2 = \lambda^2$ and $\sum_i t_i^4 = \lambda^4 \kappa_4(w)$. Lemma A.10 (a) and (b) give $t_i^2 - t_i^4/3 \le t_i \tanh t_i \le t_i^2$. Summing over i,

$$\lambda^2 - \lambda^4 \kappa_4(w)/3 \le \lambda \psi'(\lambda) \le \lambda^2,$$

and dividing by $\lambda > 0$ gives the bounds on ψ'; for $\lambda = 0$ all three terms vanish. Summing Lemma A.10 (c) and (d) over i in the same way gives the bounds on ψ. Since ψ is even, the upper bound holds for every real λ.

(c) Each factor sums to 1 over $x_i = \pm 1$, so $\mathbb{Q}_\lambda$ is a probability measure with independent coordinates. With $t = \lambda w_i$, the mean of one coordinate is $(e^t - e^{-t})/(2\cosh t) = \tanh t$. Hence $\mathbb{E}_{\mathbb{Q}_\lambda}[X] = \sum_i w_i \tanh(\lambda w_i) = \psi'(\lambda)$ and $\|\mathbb{E}_{\mathbb{Q}_\lambda}[x]\|^2 = \sum_i \tanh^2(\lambda w_i) \le d$. The entropy of one coordinate is

$$\begin{aligned} -\sum_{s=\pm 1} \frac{e^{st}}{2\cosh t} \log \frac{e^{st}}{2\cosh t} &= \log(2\cosh t) - \sum_{s=\pm 1} \frac{e^{st}}{2\cosh t}\, st \\ &= \log 2 + \log\cosh t - t \tanh t, \end{aligned}$$

and the entropy of independent coordinates is the sum of their entropies, which is (84). □

Lemma A.12 (The order parameter dominates the energy along w)**.** *For every unit vector w and $X = \langle w, x \rangle$,*

$$\|C\|_F \;\ge\; \mathbb{E}[X^2], \qquad N\,\mathbb{E}[q_{ab}^2] \;\ge\; N\Big(\frac{\mathbb{E}[X^2]}{N}\Big)^2. \tag{85}$$

Proof. Since $\|ww^\top\|_F = \|w\|^2 = 1$, the Cauchy–Schwarz inequality for the Frobenius inner product gives

$$\|C\|_F = \|C\|_F \, \|ww^\top\|_F \;\ge\; \langle C, ww^\top \rangle = \sum_{i,j} C_{ij} w_i w_j = \mathbb{E}\big[\langle w, x\rangle^2\big] = \mathbb{E}[X^2],$$

where the third equality uses $C_{ij} = \mathbb{E}[x_i x_j]$. The second inequality of (85) follows by squaring, since $\mathbb{E}[X^2] \ge 0$, and using (2). □

This lemma holds for every J and every β; the condition $\theta > 1$ enters only through the lower bound on $\mathbb{E}[X^2]$ below.

For $t \in \mathbb{R}$ let

$$F(t) = \log \sum_{x \in \{\pm 1\}^N} \exp\Big(\tfrac{t}{2} X^2 + \tfrac{\beta}{2}\, x^\top J_{\mathrm{bulk}} x\Big), \tag{86}$$

and let $\mathbb{E}_t$ be the expectation under the measure proportional to the summand. The parameter t multiplies only the term along w; the term with J_{bulk} is fixed. By (82), $F(\theta) = \log Z_\beta$ and $\mathbb{E}_\theta = \mathbb{E}_{\mu_\beta}$.

Lemma A.13 (The energy along w dominates a difference quotient of F)**.** *F is twice differentiable with $F'(t) = \frac{1}{2}\mathbb{E}_t[X^2]$ and $F''(t) = \frac{1}{4}\operatorname{Var}_t(X^2) \ge 0$. Consequently $t \mapsto \mathbb{E}_t[X^2]$ is nondecreasing, and for $\theta' < \theta$,*

$$\mathbb{E}_\theta[X^2] \;\ge\; \frac{2\big(F(\theta) - F(\theta')\big)}{\theta - \theta'}. \tag{87}$$

Proof. Let $Z(t) = e^{F(t)}$. Differentiating the finite sum term by term, $Z'(t) = \sum_x \frac{1}{2} X^2 \exp(\cdots)$, so $F'(t) = Z'(t)/Z(t) = \frac{1}{2}\mathbb{E}_t[X^2]$. Differentiating again, $F'' = Z''/Z - (Z'/Z)^2 = \frac{1}{4}\mathbb{E}_t[X^4] - \frac{1}{4}\mathbb{E}_t[X^2]^2 = \frac{1}{4}\operatorname{Var}_t(X^2)$. Since F' is nondecreasing,

$$F(\theta) - F(\theta') = \int_{\theta'}^{\theta} F'(t)\,\mathrm{d}t \le (\theta - \theta') F'(\theta) = \tfrac{1}{2}(\theta - \theta')\,\mathbb{E}_\theta[X^2].$$

□

Lemma A.14 (Lower bound on $F(\theta)$ for $\theta > 1$)**.** *Let $\theta > 1$, $\varepsilon = \theta - 1$ and $c_\theta = 2\theta/3 + 1/6$. Then*

$$F(\theta) \;\geq\; N \log 2 - \frac{\rho N}{2} + \frac{\varepsilon^2}{8 c_\theta \kappa}. \tag{88}$$

For $\theta \in (1, 2]$, $c_\theta \leq 3/2$ and the last term is at least $\varepsilon^2/(12\kappa)$.

Proof. Step 1: for every probability measure ν on $\{\pm 1\}^N$,

$$F(\theta) \;\geq\; \mathbb{E}_\nu\big[\tfrac{\beta}{2}\, x^\top J x\big] + H(\nu). \tag{89}$$

By (13), $\log \mu_\beta(x) = \frac{\beta}{2} x^\top J x - \log Z_\beta$, so the relative entropy of ν with respect to μ_β is

$$0 \leq \sum_x \nu(x) \log \frac{\nu(x)}{\mu_\beta(x)} = -H(\nu) - \mathbb{E}_\nu\big[\tfrac{\beta}{2}\, x^\top J x\big] + \log Z_\beta,$$

and $\log Z_\beta = F(\theta)$.

Step 2: with $\nu = \mathbb{Q}_\lambda$ of Lemma A.11 (c) for the vector w and $\lambda \geq 0$,

$$F(\theta) \;\geq\; N \log 2 - \frac{\rho N}{2} + \phi(\lambda), \qquad \phi(\lambda) = \tfrac{\theta}{2}\, \psi'(\lambda)^2 - \lambda \psi'(\lambda) + \psi(\lambda), \tag{90}$$

where $\psi = \psi_w$. By (82) and $\mathbb{E}_\nu[X^2] \geq \mathbb{E}_\nu[X]^2 = \psi'(\lambda)^2$,

$$\mathbb{E}_\nu\big[\tfrac{\beta}{2}\, x^\top J x\big] = \tfrac{\theta}{2}\, \mathbb{E}_\nu[X^2] + \mathbb{E}_\nu\big[\tfrac{\beta}{2}\, x^\top J_{\mathrm{bulk}} x\big] \;\geq\; \tfrac{\theta}{2}\, \psi'(\lambda)^2 - \frac{\rho N}{2}.$$

The entropy is (84) with $d = N$. Substituting both into (89) gives (90).

Step 3: if $\lambda^2 \kappa \leq 3$, then $\phi(\lambda) \geq \frac{1}{2}\big(\varepsilon \lambda^2 - c_\theta \lambda^4 \kappa\big)$. By Lemma A.11 (b), $\psi'(\lambda) \geq \lambda(1 - \lambda^2\kappa/3) \geq 0$, so

$$\begin{aligned} &\tfrac{\theta}{2}\, \psi'(\lambda)^2 \geq \tfrac{\theta}{2}\, \lambda^2 \Big(1 - \frac{\lambda^2 \kappa}{3}\Big)^2 \geq \tfrac{\theta}{2}\, \lambda^2 \Big(1 - \frac{2\lambda^2 \kappa}{3}\Big), \\ &-\lambda \psi'(\lambda) \geq -\lambda^2, \qquad \psi(\lambda) \geq \frac{\lambda^2}{2} - \frac{\lambda^4 \kappa}{12}, \end{aligned}$$

where the second inequality of the first line is $(1 - s)^2 \geq 1 - 2s$. Adding the three bounds,

$$\phi(\lambda) \;\geq\; \frac{\theta - 1}{2}\, \lambda^2 - \Big(\frac{\theta}{3} + \frac{1}{12}\Big) \lambda^4 \kappa = \tfrac{1}{2}\big(\varepsilon \lambda^2 - c_\theta \lambda^4 \kappa\big).$$

Step 4: the choice $\lambda^2 = \varepsilon/(2 c_\theta \kappa)$ gives (88). The right-hand side of Step 3 is a concave quadratic in λ^2, maximal at this value, where it equals

$$\frac{1}{2}\Big(\frac{\varepsilon^2}{2 c_\theta \kappa} - \frac{\varepsilon^2}{4 c_\theta \kappa}\Big) = \frac{\varepsilon^2}{8 c_\theta \kappa}.$$

The condition of Step 3 holds because $\lambda^2 \kappa = \varepsilon/(2 c_\theta) = 3(\theta - 1)/(4\theta + 1) < 3/4$. For $\theta \leq 2$, $c_\theta \leq 4/3 + 1/6 = 3/2$, so $1/(8 c_\theta) \geq 1/12$. □

Remark A.15. In Step 3 the term $\frac{\theta}{2}\lambda^2$ is the energy gained along w by the tilted measure, and $\lambda^2 - \lambda^2/2$ is the entropy it loses relative to the uniform measure, both to second order in λ. Their difference $\frac{1}{2}(\theta - 1)\lambda^2$ is positive exactly when $\theta > 1$; the fourth-order terms, of size $\lambda^4 \kappa$, limit how large λ can be taken.

Lemma A.16 (Upper bound on $F(\theta')$ for $\theta' < 1$)**.** *For $0 \leq \theta' < 1$,*

$$F(\theta') \;\leq\; N \log 2 + \frac{\rho N}{2} + \frac{1}{2} \log \frac{1}{1 - \theta'}. \tag{91}$$

Proof. Step 1: $F(\theta') \le \rho N/2 + N\log 2 + \log \mathbb{E}_0[e^{\theta' X^2/2}]$. This follows from the bound in (82) and $|\{\pm 1\}^N| = 2^N$.

Step 2: $\mathbb{E}_0[e^{\theta' X^2/2}] \le (1-\theta')^{-1/2}$. Let g be a standard Gaussian variable, independent of x, with expectation $\mathbb{E}_g$. For every real s, $\mathbb{E}_g[e^{sg}] = e^{s^2/2}$; with $s = \sqrt{\theta'}X$ this reads $e^{\theta' X^2/2} = \mathbb{E}_g[e^{\sqrt{\theta'} g X}]$. The integrand is nonnegative, so the two expectations can be exchanged:

$$\begin{aligned}\mathbb{E}_0\big[e^{\theta' X^2/2}\big] &= \mathbb{E}_g\Big[\mathbb{E}_0\big[e^{\sqrt{\theta'} g\, X}\big]\Big] = \mathbb{E}_g\big[e^{\psi(\sqrt{\theta'} g)}\big] \\ &\le \mathbb{E}_g\big[e^{\theta' g^2/2}\big] = \frac{1}{\sqrt{1-\theta'}}.\end{aligned}$$

The second equality is Lemma A.11 (a), the inequality is $\psi(s) \le s^2/2$ from Lemma A.11 (b), and the last equality is the Gaussian integral $\int e^{\theta' g^2/2} e^{-g^2/2}\,\mathrm{d}g/\sqrt{2\pi}$, which converges for $\theta' < 1$. □

Theorem 3.3 (Order for $\beta > 1/\lambda_1$)**.** *Let $J = \lambda_1 w w^\top + J_{\mathrm{bulk}}$ with $\|w\| = 1$, and write $\theta = \beta\lambda_1 \in (1,2]$, $\varepsilon = \theta - 1$, $\rho = \beta\|J_{\mathrm{bulk}}\|_{\mathrm{op}}$, $\kappa = \kappa_4(w) = \sum_i w_i^4$, and $X = \langle w, x\rangle$. Then for every N,*

$$N\,\mathbb{E}[q_{ab}^2] \;\ge\; N\Big(\frac{\mathbb{E}[X^2]}{N}\Big)^2, \qquad \mathbb{E}[X^2] \;\ge\; \frac{4}{3\varepsilon}\Big[\frac{\varepsilon^2}{12\kappa} - \rho N - \tfrac{1}{2}\log\frac{2}{\varepsilon}\Big], \tag{5}$$

In particular, if $\kappa \le M^2/N$ and $\rho < \varepsilon^2/(12M^2)$, then for N large enough

$$N\,\mathbb{E}[q_{ab}^2] \;\ge\; N\Big(\frac{\varepsilon}{9M^2} - \frac{4\rho}{3\varepsilon} - \frac{2\log(2/\varepsilon)}{3\varepsilon N}\Big)^2. \tag{6}$$

Proof of Theorem 3.3. The first inequality of (5) is Lemma A.12. For the second, let $\theta \in (1,2]$ and $\theta' = 1 - \varepsilon/2$. Then $\theta' \in [1/2, 1)$, $\theta - \theta' = 3\varepsilon/2$ and $\log\frac{1}{1-\theta'} = \log\frac{2}{\varepsilon}$. By Lemmas A.13, A.14 and A.16,

$$\begin{aligned}\mathbb{E}[X^2] &\ge \frac{2\big(F(\theta) - F(\theta')\big)}{\theta - \theta'} \\ &\ge \frac{2}{3\varepsilon/2}\Big(\Big[N\log 2 - \frac{\rho N}{2} + \frac{\varepsilon^2}{12\kappa}\Big] - \Big[N\log 2 + \frac{\rho N}{2} + \frac{1}{2}\log\frac{2}{\varepsilon}\Big]\Big) \\ &= \frac{4}{3\varepsilon}\Big[\frac{\varepsilon^2}{12\kappa} - \rho N - \frac{1}{2}\log\frac{2}{\varepsilon}\Big].\end{aligned}$$

If $\kappa \le M^2/N$, then $\varepsilon^2/(12\kappa) \ge \varepsilon^2 N/(12M^2)$, and dividing by N,

$$\frac{\mathbb{E}[X^2]}{N} \;\ge\; \frac{4}{3\varepsilon}\Big[\frac{\varepsilon^2}{12M^2} - \rho - \frac{\log(2/\varepsilon)}{2N}\Big] = \frac{\varepsilon}{9M^2} - \frac{4\rho}{3\varepsilon} - \frac{2\log(2/\varepsilon)}{3\varepsilon N}.$$

If $\rho < \varepsilon^2/(12M^2)$, the difference of the first two terms is positive and the third tends to zero, so the right-hand side is positive for N large; squaring and using Lemma A.12 gives (6).

The choice of θ' balances two terms. The denominator of (87) is $\theta - \theta' = \varepsilon + (1-\theta')$, and the bound (91) contains $\log\frac{1}{1-\theta'}$. With $1-\theta'$ proportional to ε, the denominator is a multiple of ε and the logarithm does not grow with N.

The case $\theta > 2$. Fix β and J_{bulk}. By Lemma A.13, $\mathbb{E}_t[X^2]$ is nondecreasing in t, so for $\theta > 2$, $\mathbb{E}_\theta[X^2] \ge \mathbb{E}_2[X^2]$. The three lemmas above concern the function F at the parameter values $t = 2$ and $t = 1/2$, so the display in the proof above holds at $t = 2$, that is with $\varepsilon = 1$: $\mathbb{E}_\theta[X^2] \ge \frac{4}{3}\big[1/(12\kappa) - \rho N - \frac{1}{2}\log 2\big]$.

Theorem 3.4 (Order, one pattern plus noise)**.** *Let $W = \mu u v^\top + W_{\mathrm{bulk}}$ with $u \in \mathbb{R}^m$, $v \in \mathbb{R}^n$ unit vectors, $\theta = \beta\mu \in (1,2]$, $\varepsilon = \theta - 1$, $\rho = \beta\|W_{\mathrm{bulk}}\|_{\mathrm{op}}$, and $\kappa = \max(\kappa_4(u), \kappa_4(v))$. Let $X = u^\top\sigma$ and $Y = v^\top\tau$. Then for every m, n,*

$$\mathbb{E}[XY] \;\ge\; \frac{2}{3\varepsilon}\Big[\frac{\varepsilon^2}{6\kappa} - 2\rho\sqrt{mn} - \log\frac{2}{\varepsilon}\Big], \qquad N\,\mathbb{E}[q_{ab}^2] \;\ge\; N\Big(\frac{2\,\mathbb{E}[XY]}{N}\Big)^2 \tag{7}$$

In particular, if $\kappa \le M^2/n$ with $m = n$ and $\rho < \varepsilon^2/(12M^2)$, then for n large enough $N\,\mathbb{E}[q_{ab}^2] \ge N\big(\varepsilon/(9M^2) - 4\rho/(3\varepsilon) - 2\log(2/\varepsilon)/(3\varepsilon n)\big)^2$.

Proof of Theorem 3.4. The proof is the one above with the following five changes; $X = u^\top\sigma$, $Y = v^\top\tau$, $\kappa = \max(\kappa_4(u), \kappa_4(v))$ and $\rho = \beta\|W_{\mathrm{bulk}}\|_{\mathrm{op}}$ are as in the statement.

The split. Substituting $W = \mu uv^\top + W_{\mathrm{bulk}}$ into (21),

$$\beta\sigma^\top W\tau = \theta XY + \beta\sigma^\top W_{\mathrm{bulk}}\tau, \qquad \big|\beta\sigma^\top W_{\mathrm{bulk}}\tau\big| \le \beta\|W_{\mathrm{bulk}}\|_{\mathrm{op}}\|\sigma\|\|\tau\| = \rho\sqrt{mn}. \tag{92}$$

Lemma A.12. With the unit vector $w = (u, v)/\sqrt{2}$, $\langle w, x\rangle = (X + Y)/\sqrt{2}$, and $(X + Y)^2 = (X - Y)^2 + 4XY \ge 4XY$. Hence

$$\|C\|_F \;\ge\; \tfrac{1}{2}\,\mathbb{E}\big[(X + Y)^2\big] \;\ge\; 2\,\mathbb{E}[XY],$$

and if $\mathbb{E}[XY] \ge 0$, then $N\,\mathbb{E}[q_{ab}^2] \ge N(2\,\mathbb{E}[XY]/N)^2$, the second inequality of (7).

Lemma A.13. With $F(t) = \log\sum_{\sigma,\tau}\exp\big(tXY + \beta\sigma^\top W_{\mathrm{bulk}}\tau\big)$ the same computation gives $F'(t) = \mathbb{E}_t[XY]$ and $F''(t) = \mathrm{Var}_t(XY) \ge 0$, so $\mathbb{E}_\theta[XY] \ge (F(\theta) - F(\theta'))/(\theta - \theta')$.

Lemma A.14. The trial measure is the product $\nu = \mathbb{Q}_\lambda^{(u)} \otimes \mathbb{Q}_\lambda^{(v)}$ of the measures of Lemma A.11 (c) for u and for v. Under ν, σ and τ are independent, so

$$\mathbb{E}_\nu[\beta\sigma^\top W\tau] = \theta\,\psi_u'(\lambda)\psi_v'(\lambda) + \beta\,\mathbb{E}_\nu[\sigma]^\top W_{\mathrm{bulk}}\mathbb{E}_\nu[\tau] \;\ge\; \theta\,\psi_u'(\lambda)\psi_v'(\lambda) - \rho\sqrt{mn},$$

where the inequality uses $\|\mathbb{E}_\nu[\sigma]\| \le \sqrt{m}$ and $\|\mathbb{E}_\nu[\tau]\| \le \sqrt{n}$. The entropy is $N\log 2 + \psi_u - \lambda\psi_u' + \psi_v - \lambda\psi_v'$. The function ϕ becomes $\theta\psi_u'\psi_v' - \lambda(\psi_u' + \psi_v') + \psi_u + \psi_v$, every term of Step 3 doubles, and

$$F(\theta) \;\ge\; N\log 2 - \rho\sqrt{mn} + \frac{\varepsilon^2}{4c_\theta\kappa} \;\ge\; N\log 2 - \rho\sqrt{mn} + \frac{\varepsilon^2}{6\kappa}, \tag{93}$$

the second inequality for $\theta \in (1, 2]$.

Lemma A.16. Since $XY \le (X^2 + Y^2)/2$ and σ, τ are independent under P_0, $\mathbb{E}_0[e^{\theta' XY}] \le \mathbb{E}_0[e^{\theta' X^2/2}]\,\mathbb{E}_0[e^{\theta' Y^2/2}] \le (1 - \theta')^{-1}$, so $F(\theta') \le N\log 2 + \rho\sqrt{mn} + \log\frac{1}{1-\theta'}$.

With $\theta' = 1 - \varepsilon/2$ the three bounds give

$$\mathbb{E}[XY] \;\ge\; \frac{2}{3\varepsilon}\Big[\frac{\varepsilon^2}{6\kappa} - 2\rho\sqrt{mn} - \log\frac{2}{\varepsilon}\Big],$$

the first inequality of (7). For $m = n$ and $\kappa \le M^2/n$, dividing by $n = N/2$ gives $2\,\mathbb{E}[XY]/N \ge \varepsilon/(9M^2) - 4\rho/(3\varepsilon) - 2\log(2/\varepsilon)/(3\varepsilon n)$.

Where each hypothesis enters. That w (respectively u, v) is a unit vector is used in $\|ww^\top\|_F = 1$ and in $\sum_i t_i^2 = \lambda^2$. The condition $\theta > 1$ makes the leading term of Step 3 positive. The condition $\theta \le 2$ is used in $c_\theta \le 3/2$ and in $\theta' \ge 0$. The bulk enters only through the bound in (82), once in each of the two bounds on F. The condition $\kappa \le M^2/N$ is used only in the last step, to make $\varepsilon^2/(12\kappa)$ of order N. No sign condition on the entries of w, u, v is used, and in the bipartite case u and v are unrelated.

A.6 Slow relaxation: proof of Proposition 4.6

Proposition 4.6 (Slow relaxation for $\beta > 1/\mu$)**.** *(a) Let W, $\theta \in (1, 2]$, ε, ρ and κ be as in Theorem 3.4. Then the relaxation time of the chain is at least*

$$\frac{1}{2e^2 N}\exp\Big(\frac{\varepsilon^2}{6\kappa} - 2\rho\sqrt{mn}\Big). \tag{12}$$

(b) Let $W = \mu uv^\top$ with $u = v = n^{-1/2}\mathbf{1}$ and $\theta = \beta\mu \in (1, 2]$. Let $X = u^\top\sigma$ and $Y = v^\top\tau$, and define $\Delta_\beta > 0$ by $\mu_\beta\big(|X + Y| \le 2n^{-1/2}\big) = e^{-n\Delta_\beta + o(n)}$, the rate at which the probability that the two block magnetizations cancel decays. Then the relaxation time of the chain is at least $e^{n\Delta_\beta - o(n)}$, and $\Delta_\beta \ge \varepsilon^2/6$.

For a set A of configurations let ∂A be the set of configurations x for which flipping some single coordinate of x changes whether $x \in A$.

Lemma A.17. *Let A be a set of configurations with $\mu_\beta(A) \le \frac{1}{2}$. Then*

$$\lambda_P(\beta) \;\le\; \frac{N\,\mu_\beta(\partial A)}{2\,\mu_\beta(A)}. \tag{94}$$

Proof. By the definition of $\lambda_P(\beta)$ with the observable $\mathbf{1}_A$, $\lambda_P(\beta) \le \mathcal{E}(\mathbf{1}_A, \mathbf{1}_A)/\operatorname{Var}_{\mu_\beta}(\mathbf{1}_A)$. For the numerator, $\operatorname{Var}(\mathbf{1}_A \mid x_{\setminus i})$ is the variance of a variable with values in $\{0,1\}$, so it is at most $\frac{1}{4}$, and it vanishes unless the two configurations $[x_{\setminus i}, \pm 1]$ differ in their membership of A, in which case $x \in \partial A$. By (19), $\mathcal{E}(\mathbf{1}_A, \mathbf{1}_A) \le \frac{N}{4}\,\mu_\beta(\partial A)$. For the denominator, $\operatorname{Var}(\mathbf{1}_A) = \mu_\beta(A)(1-\mu_\beta(A)) \ge \frac{1}{2}\mu_\beta(A)$. □

This is the elementary direction of Cheeger's inequality: $\mu_\beta(\partial A)/\mu_\beta(A)$ bounds the bottleneck ratio of A.

Proof of (a). Let $A = \{X + Y > 0\}$ and $\delta_\infty = \max(\|u\|_\infty, \|v\|_\infty) \le 1$.

Step 1: $\partial A \subset \{|X + Y| \le 2\delta_\infty\}$. Flipping σ_i changes X by $2|u_i| \le 2\delta_\infty$ and leaves Y unchanged; flipping τ_j changes Y by at most $2\delta_\infty$. A configuration with $X + Y > 2\delta_\infty$ stays in A after any single flip, and one with $X + Y \le -2\delta_\infty$ stays outside.

Step 2: $p^* := \mu_\beta(|X + Y| \le 2\delta_\infty) \le \exp\bigl(\theta + 2\rho\sqrt{mn} - \varepsilon^2/(6\kappa)\bigr)$. On this set $XY \le (X+Y)^2/4 \le \delta_\infty^2 \le 1$, so by (92) each of its at most 2^N configurations has weight at most $e^{\theta+\rho\sqrt{mn}}$ in $Z_\beta\,\mu_\beta$. The lower bound (93) on $\log Z_\beta = F(\theta)$ gives the claim.

Step 3: $\frac{1}{2}(1-p^*) \le \mu_\beta(A) \le \frac{1}{2}$. The flip $x \mapsto -x$ preserves μ_β and maps A onto $\{X+Y < 0\}$, so $\mu_\beta(A) = \frac{1}{2}(1 - \mu_\beta(X + Y = 0))$, and $\mu_\beta(X + Y = 0) \le p^*$.

Step 4: $\lambda_P(\beta) \le 2Np^*$. If $p^* \le \frac{1}{2}$, Lemma A.17 and Steps 1 and 3 give $\lambda_P(\beta) \le Np^*/(1 - p^*) \le 2Np^*$. If $p^* > \frac{1}{2}$, then $2Np^* > 1$, while $\lambda_P(\beta) \le 1$: for the observable x_1, $\operatorname{Var}_{\mu_\beta}(x_1) = 1$ and $\mathcal{E}(x_1, x_1) = \mathbb{E}[\operatorname{Var}(x_1 \mid x_{\setminus 1})] \le 1$.

The relaxation time is $1/\lambda_P(\beta)$, and Steps 2 and 4 with $\theta \le 2$ give (12).

Proof of (b). Let $m = n$, $u = v = n^{-1/2}\mathbf{1}$ and $W_{\text{bulk}} = 0$. The block magnetizations $s = \frac{1}{n}\sum_i \sigma_i$ and $t = \frac{1}{n}\sum_j \tau_j$ take values in $S_n = \{-1 + 2k/n : k = 0, \dots, n\}$, and $X = \sqrt{n}\,s$, $Y = \sqrt{n}\,t$, so $\beta\sigma^\top W\tau = \theta n\,st$. The number of σ with magnetization s is the binomial coefficient $B_n(s) = \binom{n}{n(1+s)/2}$. With the entropy function

$$I(c) = \log 2 - \tfrac{1+c}{2}\log(1+c) - \tfrac{1-c}{2}\log(1-c), \qquad c \in [-1, 1],$$

one has $(n+1)^{-1}e^{nI(s)} \le B_n(s) \le e^{nI(s)}$. Indeed, with $p = (1+s)/2$ and $k = np$, $p^k(1-p)^{n-k} = \exp\bigl(n[p\log p + (1-p)\log(1-p)]\bigr) = e^{-nI(s)}$, and the binomial probability $B_n(s)\,p^k(1-p)^{n-k}$ is at most 1 and is the largest of the $n+1$ probabilities of the binomial distribution with parameters n and p, hence at least $1/(n+1)$.

Step 1: with $g(c) = \theta c^2 + 2I(c)$ and $\Delta_\beta = \max_{c\in[0,1]} g(c) - 2\log 2$,

$$\mu_\beta\bigl(|X + Y| \le 2n^{-1/2}\bigr) = e^{-n\Delta_\beta + o(n)}. \tag{95}$$

The event is $\{|s + t| \le 2/n\}$. Its weight in $Z_\beta\,\mu_\beta$ is $\sum_{|s+t|\le 2/n} B_n(s)B_n(t)e^{\theta n st}$. On the event $st \le (s+t)^2/4 \le n^{-2}$, so the weight is at most $2^{2n}e^{\theta/n}$. It is at least the term with $t = -s$ and $s \in \{0, 1/n\} \cap S_n$, which is at least $(n+1)^{-2}e^{2nI(s)-\theta/n} = 2^{2n}e^{-o(n)}$ because I is continuous with $I(0) = \log 2$. For Z_β: since $\theta st \le \frac{\theta}{2}(s^2 + t^2)$, $\theta st + I(s) + I(t) \le \frac{1}{2}g(s) + \frac{1}{2}g(t) \le \max g$, so $Z_\beta \le (n+1)^2 e^{n\max g}$; and for the point $s^* \in S_n$ nearest to a maximizer of g, $Z_\beta \ge B_n(s^*)^2 e^{\theta n (s^*)^2} \ge (n+1)^{-2}e^{ng(s^*)} = e^{n\max g - o(n)}$ by continuity of g. The ratio of the two estimates is (95).

Step 2: $\Delta_\beta \ge \varepsilon^2/6$. Here $\kappa = 1/n$, $\rho = 0$ and $\delta_\infty = n^{-1/2}$, so Step 2 of the proof of (a) reads $\mu_\beta(|X + Y| \le 2n^{-1/2}) \le \exp(\theta - \varepsilon^2 n/6)$; compare with (95).

Step 3. Steps 1, 3 and 4 of the proof of (a) with $p^* = e^{-n\Delta_\beta + o(n)}$ give $1/\lambda_P(\beta) \geq 1/(2Np^*) = e^{n\Delta_\beta - o(n)}$.

A.7 Directional relaxation times: proof of Theorem 4.4

Theorem 4.4 (Directional relaxation times)**.** *Let $\theta_1 < 1$. There are constants C_1 and C_R, depending only on θ_1 and r, such that for all m, n and each $k \leq r$:*

(i) $\Big|\mathrm{Var}_{\mu_\beta}(\xi_k) - \dfrac{1}{1-\theta_k}\Big| \leq \dfrac{9C_1\kappa_*}{2(1-\theta_k^2)}$;

(ii) if $\kappa_* \leq (1-\theta_1^2)/(9C_1)$, then $\tau_{\xi_k} \geq (1-\theta_k)^{-1}\big(1 - 2C_R\kappa_*/(1-\theta_k)\big)$, and $\tau_{\xi_k} = (1-\theta_k)^{-1}(1+\delta_k)$ with $|\delta_k| \leq 2C_R\kappa_*/\lambda_P(\beta)$;

(iii) $\big|\mathrm{Cov}_{\mu_\beta}(\xi_j, \xi_k)\big| \leq C_1\kappa_*/(1-\theta_1)$ *for $j \neq k$.*

The Poincaré constant $\lambda_P(\beta)$ in (ii) is bounded below uniformly in N for $\beta < 1/(2\mu_1)$ and, for the flat single pattern ($m = n$, $u_1 = v_1 = n^{-1/2}\mathbf{1}$), for every $\theta_1 < 1$.

The weight matrix is $W = \sum_{k\leq r} \mu_k u_k v_k^\top$, where $\mu_k > 0$, the vectors $u_1, \dots, u_r \in \mathbb{R}^m$ are orthonormal and the vectors $v_1, \dots, v_r \in \mathbb{R}^n$ are orthonormal, that is, (24) with $W_{\mathrm{bulk}} = 0$. We write $\theta_k = \beta\mu_k$, $\theta_1 = \max_{k\leq r}\theta_k$, $\kappa_4(w) = \sum_i w_i^4$ and $\kappa_* = \max_{l\leq r}\max\{\kappa_4(u_l), \kappa_4(v_l)\}$, and we assume $\theta_1 < 1$.

For $k \leq r$ the projections of the two blocks on the k-th pattern are

$$X_k = u_k^\top \sigma = \sum_{i=1}^m u_{k,i}\sigma_i, \qquad Y_k = v_k^\top \tau = \sum_{j=1}^n v_{k,j}\tau_j, \tag{96}$$

so that the observable (23) is $\xi_k = (X_k + Y_k)/\sqrt{2}$. We write $\vec{X} = (X_1, \dots, X_r)$ and $\vec{Y} = (Y_1, \dots, Y_r)$, with $\|\vec{Y}\|^2 = \sum_{l\leq r} Y_l^2$. The uniform measure on $\{\pm 1\}^N$ is P_0 and its expectation is $\mathbb{E}_0$; under P_0 all coordinates are independent and uniform.

The local field at a coordinate is

$$a_i = \beta(W\tau)_i = \sum_{l\leq r} \theta_l u_{l,i} Y_l \quad (i \leq m), \qquad b_j = \beta(W^\top\sigma)_j = \sum_{l\leq r} \theta_l v_{l,j} X_l \quad (j \leq n), \tag{97}$$

where the second equality in each substitutes $W = \sum_l \mu_l u_l v_l^\top$, so that $(W\tau)_i = \sum_l \mu_l u_{l,i}(v_l^\top\tau)$. The field a_i depends only on τ and the field b_j depends only on σ. For $i \leq m$ the site update (18) has $(Jx)_i = (W\tau)_i$, and for the coordinate $m + j$ of the block τ it has $(Jx)_{m+j} = (W^\top\sigma)_j$, so

$$\mathbb{E}[\sigma_i \mid x_{\setminus i}] = \tanh a_i, \qquad \mathbb{E}[\tau_j \mid x_{\setminus (m+j)}] = \tanh b_j. \tag{98}$$

Substituting W in the Hamiltonian of (21),

$$\beta\sigma^\top W\tau = \sum_{k\leq r} \theta_k X_k Y_k = \sum_{i\leq m} \sigma_i a_i = \sum_{j\leq n} \tau_j b_j. \tag{99}$$

The measure μ_β is invariant under $(\sigma, \tau) \mapsto (-\sigma, -\tau)$, and X_k, Y_k and ξ_k change sign under this map, so $\mathbb{E}[X_k] = \mathbb{E}[Y_k] = \mathbb{E}[\xi_k] = 0$, and $\mathrm{Var}(\xi_k) = \mathbb{E}[\xi_k^2] = \langle \xi_k, \xi_k\rangle$ for the inner product (26).

Lemma A.18. *For every $t \in \mathbb{R}$, (a) $|\tanh t| \leq |t|$; (b) $|\tanh t - t| \leq |t|^3/3$; (c) $\cosh t \leq e^{t^2/2}$.*

Proof. (a) The function $f(t) = t - \tanh t$ has $f(0) = 0$ and $f'(t) = 1 - \mathrm{sech}^2 t = \tanh^2 t \geq 0$, so $\tanh t \leq t$ for $t \geq 0$; tanh is odd, which gives the statement in absolute value for $t < 0$.

(b) The function $f(t) = \tanh t - t + t^3/3$ has $f(0) = 0$ and $f'(t) = \mathrm{sech}^2 t - 1 + t^2 = t^2 - \tanh^2 t \geq 0$ by (a), so $\tanh t \geq t - t^3/3$ for $t \geq 0$. Together with (a) this gives $0 \leq t - \tanh t \leq t^3/3$ for $t \geq 0$, and $t - \tanh t$ is odd.

(c) Compare the series term by term: $\cosh t = \sum_{j\ge 0} t^{2j}/(2j)!$ and $e^{t^2/2} = \sum_{j\ge 0} t^{2j}/(2^j j!)$, while $(2j)! = \prod_{i=1}^{j}(2i-1)(2i) \ge \prod_{i=1}^{j} 2i = 2^j j!$. □

Lemma A.19. *Let $\theta_1 < 1$ and, for an integer $p \ge 1$, put*

$$K_0 = \Big(\frac{1-\theta_1^2}{2}\Big)^{-r/2}, \qquad K_p = p!\Big(\frac{4}{1-\theta_1^2}\Big)^{p} K_0. \tag{100}$$

Then $\mathbb{E}[\|\vec{Y}\|^{2p}] \le K_p$ and $\mathbb{E}[\|\vec{X}\|^{2p}] \le K_p$ for all m, n; in particular $\mathbb{E}[Y_l^{2p}] \le K_p$ and $\mathbb{E}[X_l^{2p}] \le K_p$ for every $l \le r$.

Proof. Step 1: for every τ, the marginal $\mu_\beta(\tau) = \sum_\sigma \mu_\beta(\sigma,\tau)$ satisfies

$$\mu_\beta(\tau) \le 2^{-n} e^{\theta_1^2 \|\vec{Y}\|^2/2}. \tag{101}$$

By (99), $\mu_\beta(\sigma,\tau) = Z_\beta^{-1}\exp(\sum_i \sigma_i a_i)$, where each a_i depends only on τ. Summing over σ coordinate by coordinate,

$$\sum_{\sigma\in\{\pm1\}^m} e^{\sum_i \sigma_i a_i} = \prod_{i=1}^{m}\big(e^{a_i} + e^{-a_i}\big) = 2^m \prod_{i=1}^{m} \cosh a_i, \tag{102}$$

so $\mu_\beta(\tau) = Z_\beta^{-1} 2^m \prod_i \cosh a_i$. Each $\cosh a_i \ge 1$, hence $Z_\beta = \sum_\tau 2^m \prod_i \cosh a_i \ge 2^m 2^n$. By Lemma A.18 (c), $\prod_i \cosh a_i \le \exp(\frac{1}{2}\sum_i a_i^2)$, and

$$\sum_i a_i^2 = \sum_i \Big(\sum_l \theta_l u_{l,i} Y_l\Big)^2 = \sum_{l,l'} \theta_l\theta_{l'} Y_l Y_{l'} \sum_i u_{l,i}u_{l',i} = \sum_l \theta_l^2 Y_l^2 \le \theta_1^2 \|\vec{Y}\|^2, \tag{103}$$

where the third equality uses $\sum_i u_{l,i}u_{l',i} = \delta_{ll'}$. The three displays give (101). Since $P_0(\tau) = 2^{-n}$, for every nonnegative function h of τ,

$$\mathbb{E}[h(\tau)] = \sum_\tau \mu_\beta(\tau) h(\tau) \le \sum_\tau 2^{-n} e^{\theta_1^2\|\vec{Y}\|^2/2} h(\tau) = \mathbb{E}_0\big[h(\tau) e^{\theta_1^2\|\vec{Y}\|^2/2}\big]. \tag{104}$$

Step 2: for $0 < t < \frac{1}{2}$,

$$\mathbb{E}_0\big[e^{t\|\vec{Y}\|^2}\big] \le (1-2t)^{-r/2}. \tag{105}$$

Let $g \sim \mathcal{N}(0, I_r)$ with expectation $\mathbb{E}_g$. For $c \in \mathbb{R}^r$ one has $\mathbb{E}_g[e^{\langle c,g\rangle}] = e^{\|c\|^2/2}$, and $c = \sqrt{2t}\,y$ turns this into $e^{t\|y\|^2} = \mathbb{E}_g[e^{\sqrt{2t}\langle g,y\rangle}]$ for $y \in \mathbb{R}^r$. Moreover $\langle g, \vec{Y}\rangle = \sum_l g_l v_l^\top \tau = \langle w, \tau\rangle$ with $w = \sum_l g_l v_l \in \mathbb{R}^n$, and $\|w\|^2 = \|g\|^2$ because $\{v_l\}$ is orthonormal. Under P_0 the coordinates τ_j are independent and uniform, so for $c \in \mathbb{R}^n$ one has $\mathbb{E}_0[e^{\langle c,\tau\rangle}] = \prod_j \cosh c_j \le e^{\|c\|^2/2}$ by Lemma A.18 (c). The integrand is nonnegative, so the two expectations may be exchanged,

$$\mathbb{E}_0\big[e^{t\|\vec{Y}\|^2}\big] = \mathbb{E}_g\Big[\mathbb{E}_0\big[e^{\sqrt{2t}\langle w,\tau\rangle}\big]\Big] \le \mathbb{E}_g\big[e^{t\|w\|^2}\big] = \mathbb{E}_g\big[e^{t\|g\|^2}\big] = (1-2t)^{-r/2}, \tag{106}$$

the last step being r independent one-dimensional Gaussian integrals $\int e^{tg^2} e^{-g^2/2}\,\mathrm{d}g/\sqrt{2\pi} = (1-2t)^{-1/2}$, which converge for $t < \frac{1}{2}$.

Step 3: the moment bounds. Put $s = (1-\theta_1^2)/4 > 0$. Take $h = e^{s\|\vec{Y}\|^2}$ in (104) and then (105) at $t = s + \theta_1^2/2 = (1+\theta_1^2)/4 < \frac{1}{2}$, for which $1 - 2t = (1-\theta_1^2)/2$:

$$\mathbb{E}\big[e^{s\|\vec{Y}\|^2}\big] \le \mathbb{E}_0\big[e^{(s+\theta_1^2/2)\|\vec{Y}\|^2}\big] \le \Big(\frac{1-\theta_1^2}{2}\Big)^{-r/2} = K_0. \tag{107}$$

For $y \ge 0$ one has $y^p/p! \le e^y$; at $y = s\|\vec{Y}\|^2$ this reads $\|\vec{Y}\|^{2p} \le p!\, s^{-p} e^{s\|\vec{Y}\|^2}$, and taking expectations gives $\mathbb{E}[\|\vec{Y}\|^{2p}] \le p!\, s^{-p} K_0 = K_p$. For $\vec{X}$ the two blocks exchange roles: Step 1 sums over τ and uses b_j in place of a_i and the orthonormality of $\{v_l\}$ to get $\sum_j b_j^2 = \sum_l \theta_l^2 X_l^2$, and Step 2 uses the orthonormality of $\{u_l\}$, so the same bound holds. Finally $Y_l^{2p} \le \|\vec{Y}\|^{2p}$ and $X_l^{2p} \le \|\vec{X}\|^{2p}$. □

Remark A.20. The exponent $t = (1+\theta_1^2)/4$ of Step 3 is smaller than $\frac{1}{2}$ exactly when $\theta_1 < 1$, and the Gaussian integrals in (105) diverge at $t = \frac{1}{2}$. At $\theta_1 = 1$ the coefficient $\theta_1^2/2$ in the density bound (101) reaches the value at which the exponential moment of $\|\vec{Y}\|^2$ under P_0 ceases to be finite.

Lemma A.21. *For every $i \le m$ and every $j \le n$, $\mathcal{L}\sigma_i = \tanh a_i - \sigma_i$ and $\mathcal{L}\tau_j = \tanh b_j - \tau_j$, hence for every $k \le r$*

$$\mathcal{L}X_k = -X_k + \sum_{i\le m} u_{k,i} \tanh a_i, \qquad \mathcal{L}Y_k = -Y_k + \sum_{j\le n} v_{k,j} \tanh b_j. \tag{108}$$

Proof. Step 1: $\mathcal{L}\sigma_i = \tanh a_i - \sigma_i$. By (32), $\mathcal{L}\sigma_i = \sum_{i'}(P_{i'}\sigma_i - \sigma_i)$. For $i' \neq i$ the function σ_i is a function of $x_{\setminus i'}$, so $P_{i'}\sigma_i = \sigma_i$ by (28) and that term vanishes. For $i' = i$ one has $P_i\sigma_i = \tanh a_i$ by (98). The same argument with b_j in place of a_i gives $\mathcal{L}\tau_j = \tanh b_j - \tau_j$.

Step 2: (108). The operator $\mathcal{L}$ is linear, so

$$\mathcal{L}X_k = \sum_i u_{k,i}\, \mathcal{L}\sigma_i = \sum_i u_{k,i}(\tanh a_i - \sigma_i) = -X_k + \sum_i u_{k,i} \tanh a_i, \tag{109}$$

the last equality by (96). The statement for Y_k exchanges σ, u, a for τ, v, b. □

Lemma A.22. *For every $k \le r$,*

$$\mathcal{L}X_k = -X_k + \theta_k Y_k + R_k^X, \qquad \mathcal{L}Y_k = -Y_k + \theta_k X_k + R_k^Y, \tag{110}$$

where

$$R_k^X = \sum_{i\le m} u_{k,i}(\tanh a_i - a_i), \qquad R_k^Y = \sum_{j\le n} v_{k,j}(\tanh b_j - b_j), \tag{111}$$

and

$$\|R_k^X\|_{L^2(\mu_\beta)} \le C_R\kappa_*, \qquad \|R_k^Y\|_{L^2(\mu_\beta)} \le C_R\kappa_*, \qquad C_R = \frac{\theta_1^3 r^3 K_3^{1/2}}{3}. \tag{112}$$

Consequently $\xi_k = (X_k + Y_k)/\sqrt{2}$ satisfies

$$\mathcal{L}\xi_k = -(1-\theta_k)\xi_k + R_k, \quad R_k = \frac{R_k^X + R_k^Y}{\sqrt{2}}, \quad \|R_k\|_{L^2(\mu_\beta)} \le \sqrt{2}C_R\kappa_*, \quad \mathbb{E}[R_k] = 0. \tag{113}$$

Proof. Step 1: (110). Write $\tanh a_i = a_i + (\tanh a_i - a_i)$ in (108):

$$\mathcal{L}X_k = -X_k + \sum_i u_{k,i}a_i + \sum_i u_{k,i}(\tanh a_i - a_i). \tag{114}$$

The middle term is computed by substituting $a_i = \sum_l \theta_l u_{l,i} Y_l$ from (97),

$$\sum_i u_{k,i}a_i = \sum_l \theta_l Y_l \sum_i u_{k,i}u_{l,i} = \sum_l \theta_l Y_l \delta_{kl} = \theta_k Y_k, \tag{115}$$

where the second equality uses the orthonormality of $\{u_l\}$. The last term is R_k^X. The statement for Y_k exchanges σ, u, a, R^X for τ, v, b, R^Y.

Step 2: (112). By Lemma A.18 (b), $|\tanh a_i - a_i| \le |a_i|^3/3$, and by (97) with $\theta_l \le \theta_1$, $|a_i| \le \theta_1 \sum_l |u_{l,i}|\,|Y_l|$, so

$$|R_k^X| \le \frac{1}{3}\sum_i |u_{k,i}|\,|a_i|^3 \le \frac{\theta_1^3}{3}\sum_i |u_{k,i}|\Big(\sum_l |u_{l,i}|\,|Y_l|\Big)^3 = \frac{\theta_1^3}{3}\sum_{l,l',l''} |Y_l Y_{l'} Y_{l''}| \sum_i |u_{k,i}u_{l,i}u_{l',i}u_{l'',i}|. \tag{116}$$

The inner sum over i is bounded by Hölder's inequality for four sequences with all exponents equal to 4,

$$\sum_i |u_{k,i}u_{l,i}u_{l',i}u_{l'',i}| \le \kappa_4(u_k)^{1/4}\kappa_4(u_l)^{1/4}\kappa_4(u_{l'})^{1/4}\kappa_4(u_{l''})^{1/4} \le \kappa_*, \tag{117}$$

so that

$$|R_k^X| \le \frac{\theta_1^3}{3}\,\kappa_* \sum_{l,l',l''} |Y_l Y_{l'} Y_{l''}|. \tag{118}$$

Taking the $L^2(\mu_\beta)$ norm, using the triangle inequality, and applying Hölder's inequality for three functions with all exponents equal to 3 to each term,

$$\|Y_l Y_{l'} Y_{l''}\|_{L^2(\mu_\beta)} = \big(\mathbb{E}[Y_l^2 Y_{l'}^2 Y_{l''}^2]\big)^{1/2} \le \big(\mathbb{E}[Y_l^6]\,\mathbb{E}[Y_{l'}^6]\,\mathbb{E}[Y_{l''}^6]\big)^{1/6} \le K_3^{1/2}, \tag{119}$$

the last step by Lemma A.19 at $p = 3$. The triple sum has r^3 terms, so

$$\|R_k^X\|_{L^2(\mu_\beta)} \le \frac{\theta_1^3 r^3 K_3^{1/2}}{3}\,\kappa_* = C_R \kappa_*. \tag{120}$$

The bound on R_k^Y is obtained by exchanging the two blocks.

Step 3: (113). Adding the two identities of (110) and dividing by $\sqrt{2}$,

$$\mathcal{L}\xi_k = \frac{-X_k - Y_k + \theta_k Y_k + \theta_k X_k}{\sqrt{2}} + R_k = -(1-\theta_k)\xi_k + R_k. \tag{121}$$

The triangle inequality gives

$$\|R_k\|_{L^2(\mu_\beta)} \le \frac{\|R_k^X\|_{L^2(\mu_\beta)} + \|R_k^Y\|_{L^2(\mu_\beta)}}{\sqrt{2}} \le \sqrt{2} C_R \kappa_*. \tag{122}$$

Finally $\mathbb{E}[R_k] = \mathbb{E}[\mathcal{L}\xi_k] + (1-\theta_k)\mathbb{E}[\xi_k] = 0$, the first expectation by Lemma A.2 (a) and the second by the flip symmetry of μ_β. □

Remark A.23. Without the remainders, $\mathcal{L}$ acts on the span of $X_1, \dots, X_r, Y_1, \dots, Y_r$ as $-(I - \Theta)$, where Θ maps X_k to $\theta_k Y_k$ and Y_k to $\theta_k X_k$. The eigenvectors of Θ are $X_k \pm Y_k$ with eigenvalues $\pm\theta_k$, so the eigenvalue of $-(I-\Theta)$ is $-(1-\theta_k)$ on ξ_k and $-(1+\theta_k)$ on $(X_k - Y_k)/\sqrt{2}$. Of the two combinations only ξ_k has a decay rate that becomes small as θ_k approaches 1.

By Lemma A.2 (a), $\mathbb{E}[\mathcal{L}h] = 0$ for every h. Taking $h = fg$ and splitting $\mathcal{L}(fg)$ into three parts,

$$\Gamma(f,g) := \mathcal{L}(fg) - f\,\mathcal{L}g - g\,\mathcal{L}f, \qquad 0 = \mathbb{E}[f\,\mathcal{L}g] + \mathbb{E}[g\,\mathcal{L}f] + \mathbb{E}[\Gamma(f,g)], \tag{123}$$

the second identity being $\mathbb{E}[\mathcal{L}(fg)] = 0$ expanded by the definition of Γ.

Remark A.24. The second identity of (123) states that $\mathbb{E}[f(x_t)g(x_t)]$ does not change with t for the stationary chain: its derivative at $t = 0$ is $\mathbb{E}[\mathcal{L}(fg)]$, which the definition of Γ splits into the drift of f multiplied by g, the drift of g multiplied by f, and the expectation of the product of the two increments.

Lemma A.25. *(a) For all f, g,*

$$\Gamma(f,g)(x) = \sum_{i=1}^{N} \mathbb{E}\Big[\big(f(x^{(i)}) - f(x)\big)\big(g(x^{(i)}) - g(x)\big)\;\Big|\; x_{\setminus i}\Big], \tag{124}$$

where $x^{(i)} = [x_{\setminus i}, x_i']$ with x_i' drawn from the conditional distribution (18)*, the conditional expectation is over x_i', and $f(x)$, $g(x)$ are constants in it.*

(b) For $j, k \le r$,

$$\Gamma(X_j, X_k) = \sum_{i \le m} u_{j,i} u_{k,i} \big(2 - 2\sigma_i \tanh a_i\big), \qquad \Gamma(X_j, Y_k) = 0, \tag{125}$$

and $\Gamma(Y_j, Y_k)$ replaces u, σ, a by v, τ, b.

(c) Taking expectations,

$$\mathbb{E}[\Gamma(X_j, X_k)] = 2\langle u_j, u_k\rangle - 2\sum_{i \le m} u_{j,i} u_{k,i} \mathbb{E}[\tanh^2 a_i], \qquad \Big|\sum_{i \le m} u_{j,i} u_{k,i} \mathbb{E}[\tanh^2 a_i]\Big| \le \theta_1^2 r^2 K_1 \kappa_*. \tag{126}$$

Proof. (a) By (32), $\mathcal{L}h = \sum_i (P_i h - h)$, so

$$\Gamma(f,g) = \sum_i \Big(P_i(fg) - fg - f(P_i g - g) - g(P_i f - f) \Big) = \sum_i \Big(P_i(fg) - f\,P_i g - g\,P_i f + fg \Big). \tag{127}$$

By (28), $P_i(fg)(x) = \mathbb{E}[f(x^{(i)})g(x^{(i)}) \mid x_{\setminus i}]$ and $P_i g(x) = \mathbb{E}[g(x^{(i)}) \mid x_{\setminus i}]$, and likewise for $P_i f$; expanding $\mathbb{E}[(f(x^{(i)}) - f(x))(g(x^{(i)}) - g(x)) \mid x_{\setminus i}]$ into four terms produces the four terms in the bracket.

(b) Take $f = X_j$ and $g = X_k$. Resampling a coordinate $i \le m$ of the block σ gives $X_k(x^{(i)}) - X_k(x) = u_{k,i}(\sigma_i' - \sigma_i)$, and resampling a coordinate of the block τ leaves X_k unchanged, so only the terms with $i \le m$ remain in (124),

$$\Gamma(X_j, X_k) = \sum_{i \le m} u_{j,i} u_{k,i}\, \mathbb{E}\big[(\sigma_i' - \sigma_i)^2 \mid x_{\setminus i}\big]. \tag{128}$$

Here $(\sigma_i')^2 = 1$ and $\sigma_i^2 = 1$, and $\mathbb{E}[\sigma_i' \mid x_{\setminus i}] = \tanh a_i$ by (98), so the conditional expectation equals $1 - 2\sigma_i \tanh a_i + 1$. Take $f = X_j$ and $g = Y_k$: each resampling changes either σ or τ, so at least one of $X_j(x^{(i)}) - X_j(x)$ and $Y_k(x^{(i)}) - Y_k(x)$ is zero and every term vanishes.

(c) Take expectations in (125). The first term is $\sum_i u_{j,i}u_{k,i} \cdot 2 = 2\langle u_j, u_k \rangle$. In the second term a_i is a function of $x_{\setminus i}$, so the tower property gives

$$\mathbb{E}[\sigma_i \tanh a_i] = \mathbb{E}\big[\tanh a_i\, \mathbb{E}[\sigma_i \mid x_{\setminus i}]\big] = \mathbb{E}[\tanh^2 a_i], \tag{129}$$

which is the first identity of (126). For the second, by Lemma A.18 (a), $\tanh^2 a_i \le a_i^2$, and by (97) with the Cauchy–Schwarz inequality $(\sum_l z_l)^2 \le r \sum_l z_l^2$ for r terms,

$$a_i^2 = \Big(\sum_l \theta_l u_{l,i} Y_l \Big)^2 \le r \sum_l \theta_l^2 u_{l,i}^2 Y_l^2 \le \theta_1^2 r \sum_l u_{l,i}^2 Y_l^2. \tag{130}$$

Therefore

$$\Big| \sum_i u_{j,i}u_{k,i} \mathbb{E}[\tanh^2 a_i] \Big| \le \theta_1^2 r \sum_l \mathbb{E}[Y_l^2] \sum_i |u_{j,i}u_{k,i}| u_{l,i}^2 \le \theta_1^2 r \sum_l K_1 \kappa_* = \theta_1^2 r^2 K_1 \kappa_*, \tag{131}$$

where the middle step applies Hölder's inequality for three sequences with exponents $4, 4, 2$ to each l,

$$\sum_i |u_{j,i}|\, |u_{k,i}| u_{l,i}^2 \le \kappa_4(u_j)^{1/4} \kappa_4(u_k)^{1/4} \Big(\sum_i u_{l,i}^4 \Big)^{1/2} \le \kappa_*, \tag{132}$$

and uses $\mathbb{E}[Y_l^2] \le K_1$ from Lemma A.19. □

The error constant of the three statements is

$$C_1 = 2K_1^{1/2} C_R + 2\theta_1^2 r^2 K_1, \tag{133}$$

whose two parts come from the remainders of Lemma A.22 and from the second term of (126). For the first part, the Cauchy–Schwarz inequality, Lemma A.19 at $p = 1$ and (112) give, for all $l, l' \le r$,

$$\big| \mathbb{E}[X_l R_{l'}^X] \big| \le \|X_l\|_{L^2(\mu_\beta)} \|R_{l'}^X\|_{L^2(\mu_\beta)} \le K_1^{1/2} C_R \kappa_*, \tag{134}$$

and the same bound holds with X replaced by Y, with R^X replaced by R^Y, or both.

Proof of (i). For all m, n and every $k \le r$,

$$\Big| \mathrm{Var}(\xi_k) - \frac{1}{1 - \theta_k} \Big| \le \frac{9 C_1 \kappa_*}{2(1 - \theta_k^2)}. \tag{135}$$

Fix k, write $\theta = \theta_k$, and put

$$p = \mathbb{E}[X_k^2], \qquad q = \mathbb{E}[Y_k^2], \qquad s = \mathbb{E}[X_k Y_k]. \tag{136}$$

Step 1: $p = 1 + \theta s + \epsilon_1$ with $|\epsilon_1| \le C_1\kappa_*$. Take $f = g = X_k$ in (123). By (110),

$$\mathbb{E}[X_k\, \mathcal{L}X_k] = \mathbb{E}\big[X_k(-X_k + \theta Y_k + R_k^X)\big] = -p + \theta s + \mathbb{E}[X_k R_k^X]. \tag{137}$$

By (126) with $j = k$ and $\langle u_k, u_k\rangle = 1$,

$$\mathbb{E}[\Gamma(X_k, X_k)] = 2 - 2\sum_i u_{k,i}^2 \mathbb{E}[\tanh^2 a_i]. \tag{138}$$

Substituting both in (123),

$$0 = 2\big(-p + \theta s + \mathbb{E}[X_k R_k^X]\big) + 2 - 2\sum_i u_{k,i}^2 \mathbb{E}[\tanh^2 a_i], \tag{139}$$

and dividing by 2 and rearranging gives $p = 1 + \theta s + \epsilon_1$ with

$$\epsilon_1 = \mathbb{E}[X_k R_k^X] - \sum_i u_{k,i}^2 \mathbb{E}[\tanh^2 a_i]. \tag{140}$$

By (134), the bound in (126) and (133),

$$|\epsilon_1| \le K_1^{1/2} C_R \kappa_* + \theta_1^2 r^2 K_1 \kappa_* \le C_1\kappa_*. \tag{141}$$

Step 2: $q = 1 + \theta s + \epsilon_2$ with $|\epsilon_2| \le C_1\kappa_*$. Take $f = g = Y_k$ and replace X, u, a, R^X in Step 1 by Y, v, b, R^Y.

Step 3: $2s = \theta(p + q) + \epsilon_3$ with $|\epsilon_3| \le C_1\kappa_*$. Take $f = X_k$ and $g = Y_k$. By (125), $\Gamma(X_k, Y_k) = 0$. By (110),

$$\begin{aligned} \mathbb{E}[X_k\, \mathcal{L}Y_k] &= \mathbb{E}\big[X_k(-Y_k + \theta X_k + R_k^Y)\big] = -s + \theta p + \mathbb{E}[X_k R_k^Y], \\ \mathbb{E}[Y_k\, \mathcal{L}X_k] &= \mathbb{E}\big[Y_k(-X_k + \theta Y_k + R_k^X)\big] = -s + \theta q + \mathbb{E}[Y_k R_k^X]. \end{aligned} \tag{142}$$

Substituting in (123) gives $0 = -2s + \theta(p + q) + \epsilon_3$ with $\epsilon_3 = \mathbb{E}[X_k R_k^Y] + \mathbb{E}[Y_k R_k^X]$, and (134) gives $|\epsilon_3| \le 2K_1^{1/2} C_R \kappa_* \le C_1\kappa_*$.

Step 4: the three moments solved. Adding Steps 1 and 2 gives $p + q = 2 + 2\theta s + \epsilon_1 + \epsilon_2$, and substituting this in Step 3,

$$\begin{aligned} 2s &= \theta\big(2 + 2\theta s + \epsilon_1 + \epsilon_2\big) + \epsilon_3 \\ &= 2\theta + 2\theta^2 s + \theta(\epsilon_1 + \epsilon_2) + \epsilon_3, \end{aligned} \tag{143}$$

so that $2(1 - \theta^2)s = 2\theta + \theta(\epsilon_1 + \epsilon_2) + \epsilon_3$. Since $\theta < 1$ one may divide by $2(1 - \theta^2)$,

$$s = \frac{\theta}{1 - \theta^2} + \epsilon_4, \qquad \epsilon_4 = \frac{\theta(\epsilon_1 + \epsilon_2) + \epsilon_3}{2(1 - \theta^2)}, \qquad |\epsilon_4| \le \frac{3C_1\kappa_*}{2(1 - \theta^2)}. \tag{144}$$

Substituting this in Step 1,

$$p = 1 + \theta s + \epsilon_1 = 1 + \frac{\theta^2}{1 - \theta^2} + \theta\epsilon_4 + \epsilon_1 = \frac{1}{1 - \theta^2} + \epsilon_5, \qquad \epsilon_5 = \theta\epsilon_4 + \epsilon_1. \tag{145}$$

Here $|\epsilon_5| \le |\epsilon_4| + C_1\kappa_*$, and $1 - \theta^2 \le 1 < \frac{3}{2}$ gives $C_1\kappa_* \le 3C_1\kappa_*/(2(1 - \theta^2))$, so $|\epsilon_5| \le 3C_1\kappa_*/(1 - \theta^2)$. Step 2 in place of Step 1 gives $q = \frac{1}{1-\theta^2} + \epsilon_6$ with $|\epsilon_6| \le 3C_1\kappa_*/(1 - \theta^2)$.

Step 5: (135). Since $\mathbb{E}[\xi_k] = 0$,

$$\mathrm{Var}(\xi_k) = \mathbb{E}\Big[\frac{(X_k + Y_k)^2}{2}\Big] = \frac{p + q + 2s}{2} = \frac{\frac{2}{1-\theta^2} + \frac{2\theta}{1-\theta^2}}{2} + \epsilon_7 = \frac{1 + \theta}{1 - \theta^2} + \epsilon_7 = \frac{1}{1 - \theta} + \epsilon_7, \tag{146}$$

where $\epsilon_7 = (\epsilon_5 + \epsilon_6 + 2\epsilon_4)/2$, and the three bounds of Step 4 give

$$|\epsilon_7| \le \frac{1}{2}\Big(\frac{3C_1\kappa_*}{1 - \theta^2} + \frac{3C_1\kappa_*}{1 - \theta^2} + \frac{3C_1\kappa_*}{1 - \theta^2}\Big) = \frac{9C_1\kappa_*}{2(1 - \theta^2)}. \tag{147}$$

This is (135), part (i) of Theorem 4.4.

Proof of (iii). For all m, n and all $j \neq k$ with $j, k \leq r$,

$$\big|\mathrm{Cov}(\xi_j, \xi_k)\big| \leq \frac{C_1 \kappa_*}{1 - \theta_1}. \tag{148}$$

Fix $j \neq k$ and put

$$p = \mathbb{E}[X_j X_k], \qquad q = \mathbb{E}[Y_j Y_k], \qquad s = \mathbb{E}[X_j Y_k], \qquad t = \mathbb{E}[X_k Y_j]. \tag{149}$$

Step 1: the four identities

$$\begin{aligned} -2p + \theta_k s + \theta_j t &= \epsilon_1', \\ -2q + \theta_j s + \theta_k t &= \epsilon_2', \\ \theta_k p + \theta_j q - 2s &= \epsilon_3', \\ \theta_j p + \theta_k q - 2t &= \epsilon_4', \end{aligned} \tag{150}$$

with $|\epsilon_i'| \leq C_1 \kappa_*$ for $i = 1, \ldots, 4$. Take $(f, g) = (X_j, X_k)$ in (123). By (110),

$$\begin{aligned} \mathbb{E}[X_j \, \mathcal{L} X_k] &= -p + \theta_k s + \mathbb{E}[X_j R_k^X], \\ \mathbb{E}[X_k \, \mathcal{L} X_j] &= -p + \theta_j t + \mathbb{E}[X_k R_j^X], \end{aligned} \tag{151}$$

using $\mathbb{E}[X_j Y_k] = s$ in the first and $\mathbb{E}[X_k Y_j] = t$ in the second. By (126) with $\langle u_j, u_k \rangle = 0$,

$$\mathbb{E}[\Gamma(X_j, X_k)] = -2 \sum_i u_{j,i} u_{k,i} \mathbb{E}[\tanh^2 a_i], \tag{152}$$

so (123) gives the first identity with

$$\epsilon_1' = -\mathbb{E}[X_j R_k^X] - \mathbb{E}[X_k R_j^X] + 2 \sum_i u_{j,i} u_{k,i} \mathbb{E}[\tanh^2 a_i]. \tag{153}$$

Taking $(f, g) = (Y_j, Y_k)$,

$$\begin{aligned} \mathbb{E}[Y_j \, \mathcal{L} Y_k] &= -q + \theta_k t + \mathbb{E}[Y_j R_k^Y], \\ \mathbb{E}[Y_k \, \mathcal{L} Y_j] &= -q + \theta_j s + \mathbb{E}[Y_k R_j^Y], \end{aligned} \tag{154}$$

which gives the second identity. Taking $(f, g) = (X_j, Y_k)$, for which $\Gamma = 0$ by (125),

$$\begin{aligned} \mathbb{E}[X_j \, \mathcal{L} Y_k] &= -s + \theta_k p + \mathbb{E}[X_j R_k^Y], \\ \mathbb{E}[Y_k \, \mathcal{L} X_j] &= -s + \theta_j q + \mathbb{E}[Y_k R_j^X], \end{aligned} \tag{155}$$

which gives the third. Taking $(f, g) = (X_k, Y_j)$,

$$\begin{aligned} \mathbb{E}[X_k \, \mathcal{L} Y_j] &= -t + \theta_j p + \mathbb{E}[X_k R_j^Y], \\ \mathbb{E}[Y_j \, \mathcal{L} X_k] &= -t + \theta_k q + \mathbb{E}[Y_j R_k^X], \end{aligned} \tag{156}$$

which gives the fourth. Each ϵ_i' is the sum of two terms of the form $\mathbb{E}[X_j R_k^X]$, each at most $K_1^{1/2} C_R \kappa_*$ by (134), and of twice the second term of (126), which is at most $2\theta_1^2 r^2 K_1 \kappa_*$ for the first two identities and zero for the last two. By (133), $|\epsilon_i'| \leq C_1 \kappa_*$.

Step 2: $\|(p, q, s, t)\| \leq C_1 \kappa_* / (1 - \theta_1)$. Write (150) as $\mathsf{M}(p, q, s, t)^\top = (\epsilon_1', \ldots, \epsilon_4')^\top$ with

$$\mathsf{M} = \begin{pmatrix} -2I_2 & \mathsf{K} \\ \mathsf{K} & -2I_2 \end{pmatrix}, \qquad \mathsf{K} = \begin{pmatrix} \theta_k & \theta_j \\ \theta_j & \theta_k \end{pmatrix}, \tag{157}$$

which is symmetric. The eigenvectors of K are $e_\pm = (1, \pm 1)$ with eigenvalues $\theta_k \pm \theta_j$, and for each such pair $(e, \varkappa)$,

$$\mathsf{M} \begin{pmatrix} e \\ e \end{pmatrix} = \begin{pmatrix} -2e + \varkappa e \\ \varkappa e - 2e \end{pmatrix} = (\varkappa - 2) \begin{pmatrix} e \\ e \end{pmatrix}, \qquad \mathsf{M} \begin{pmatrix} e \\ -e \end{pmatrix} = \begin{pmatrix} -2e - \varkappa e \\ \varkappa e + 2e \end{pmatrix} = -(\varkappa + 2) \begin{pmatrix} e \\ -e \end{pmatrix}. \tag{158}$$

The four eigenvalues of M are therefore $-2 \pm (\theta_k + \theta_j)$ and $-2 \pm (\theta_k - \theta_j)$, whose absolute values are $2 \mp (\theta_k + \theta_j)$ and $2 \mp (\theta_k - \theta_j)$; all four are at least $2 - \theta_k - \theta_j \geq 2(1 - \theta_1) > 0$. The operator norm of the inverse of a symmetric matrix is the reciprocal of the smallest absolute value of its eigenvalues, so $\|\mathsf{M}^{-1}\|_{\mathrm{op}} \leq 1/(2(1-\theta_1))$. The right side has length at most $\sqrt{4}\, C_1\kappa_* = 2C_1\kappa_*$ by Step 1, so

$$\|(p,q,s,t)\| \leq \|\mathsf{M}^{-1}\|_{\mathrm{op}} \|(\epsilon_1', \dots, \epsilon_4')\| \leq \frac{2C_1\kappa_*}{2(1-\theta_1)} = \frac{C_1\kappa_*}{1-\theta_1}. \tag{159}$$

Step 3: (148). Since $\mathbb{E}[\xi_j] = \mathbb{E}[\xi_k] = 0$,

$$\begin{aligned} \mathrm{Cov}(\xi_j, \xi_k) = \mathbb{E}\Big[\frac{(X_j + Y_j)(X_k + Y_k)}{2}\Big] &= \frac{p+q+s+t}{2}, \\ |p+q+s+t| \leq \sqrt{4}\, \|(p,q,s,t)\| &\leq \frac{2C_1\kappa_*}{1-\theta_1}, \end{aligned} \tag{160}$$

the second line by the Cauchy–Schwarz inequality for four numbers and Step 2. Dividing by 2 gives (148), part (iii) of Theorem 4.4.

Proof of (ii). Assume $\kappa_* \leq (1-\theta_1^2)/(9C_1)$. Then

$$\mathrm{Var}(\xi_k) \geq \frac{1}{2}. \tag{161}$$

Indeed $\theta_k \leq \theta_1$ gives $1 - \theta_k^2 \geq 1 - \theta_1^2$, so the error term of (135) obeys

$$\frac{9C_1\kappa_*}{2(1-\theta_k^2)} \leq \frac{9C_1\kappa_*}{2(1-\theta_1^2)} \leq \frac{1}{2}, \tag{162}$$

and $1 - \theta_k \leq 1$ gives $\mathrm{Var}(\xi_k) \geq \frac{1}{1-\theta_k} - \frac{1}{2} \geq \frac{1}{2}$.

Apply Proposition 4.3 to $\varphi = \xi_k$, which has $\mathbb{E}[\xi_k] = 0$, with $\alpha = 1 - \theta_k > 0$ and $R = R_k$: by (113), $\mathcal{L}\xi_k + (1-\theta_k)\xi_k = R_k$. Since $\|\xi_k\| = \mathrm{Var}(\xi_k)^{1/2}$, the ratio of that proposition is

$$\eta = \frac{\|R_k\|}{\mathrm{Var}(\xi_k)^{1/2}} \leq \frac{\sqrt{2}C_R\kappa_*}{1/\sqrt{2}} = 2C_R\kappa_*, \tag{163}$$

by (113) and (161).

The lower bound of Proposition 4.3 is $\tau_{\xi_k} \geq 1/(\alpha + \eta)$, and

$$\frac{1}{\alpha+\eta} = \frac{1}{\alpha} \cdot \frac{1}{1+\eta/\alpha} \geq \frac{1}{\alpha}\Big(1 - \frac{\eta}{\alpha}\Big) \geq \frac{1}{1-\theta_k}\Big(1 - \frac{2C_R\kappa_*}{1-\theta_k}\Big), \tag{164}$$

where the first inequality is $(1+y)^{-1} \geq 1 - y$ for $y \geq 0$ and the second is (163). The two-sided statement of that proposition is

$$\tau_{\xi_k} = \frac{1}{1-\theta_k}(1+\delta_k), \qquad |\delta_k| \leq \frac{\eta}{\lambda_P(\beta)} \leq \frac{2C_R\kappa_*}{\lambda_P(\beta)}. \tag{165}$$

In two cases $\lambda_P(\beta)$ has a lower bound that does not depend on m, n. For $\beta < 1/(2\mu_1)$, Lemma 3.1 with $\gamma(J) = 2\mu_1$ gives $\lambda_P(\beta) \geq 1 - 2\beta\mu_1$. For $r = 1$, $m = n$ and $u_1 = v_1 = n^{-1/2}\mathbf{1}$, Lemma A.9 of Appendix A.4 gives $\lambda_P(\beta) \geq (1 - o(1))(1-\theta_1)/2$ for every $\theta_1 < 1$, the $o(1)$ depending only on n. The bounds (164) and (165) are part (ii) of Theorem 4.4.

Where each hypothesis enters. The orthonormality of $\{u_l\}$ and $\{v_l\}$ is used in three places: the linear term $\sum_i u_{k,i} a_i = \theta_k Y_k$ in Step 1 of Lemma A.22, the sum of squares of the local fields $\sum_i a_i^2 = \sum_l \theta_l^2 Y_l^2$ in (103), and the identity $\|w\| = \|g\|$ in Step 2 of Lemma A.19. Normalization also gives the leading term 2 of $\mathbb{E}[\Gamma(X_k, X_k)]$ and the leading term 0 of $\mathbb{E}[\Gamma(X_j, X_k)]$ in (126). The hypothesis $\theta_1 < 1$ is used in three places: the convergence of the Gaussian integrals in (105), which requires $t < \frac{1}{2}$; the division by $1 - \theta^2$ in Step 4 of the proof of (i); and the bound $2(1-\theta_1) > 0$ on the smallest absolute value of the

eigenvalues of M in Step 2 of the proof of (iii). The absence of noise is used in the fact that a_i is a linear combination of $Y_1, \dots, Y_r$: with $W_{\text{bulk}} \neq 0$ the field a_i carries the additional term $\beta(W_{\text{bulk}}\tau)_i$ and the linear term carries the additional term $\beta u_k^\top W_{\text{bulk}}\tau$, which is not a function of $Y_1, \dots, Y_r$, so the system of equations is no longer closed. Delocalization of the patterns is used only in the last step of each of the three proofs, where $\kappa_* \to 0$; the inequalities (135) and (148) hold for every m, n with no further hypothesis, and (164) and (165) hold for every m, n with $\kappa_* \le (1-\theta_1^2)/(9C_1)$, the right sides being proportional to κ_*. That r is fixed enters through the constants K_0, C_R and C_1 of (100), (112) and (133).

A.8 Relaxation times when the bulk is not small: derivation of (11)

Remark 4.5. The linearization in the proof sketch replaces $\tanh(\beta(Jx)_i)$ by a line of slope 1. When the local fields $h_i = (W\tau)_i$ and $g_j = (W^\top\sigma)_j$ are not small, the slope of tanh at the field is $1 - \tanh^2$, with block averages

$$\chi_\sigma = \frac{1}{m}\sum_{i\le m} \mathbb{E}\big[1-\tanh^2(\beta h_i)\big], \qquad \chi_\tau = \frac{1}{n}\sum_{j\le n} \mathbb{E}\big[1-\tanh^2(\beta g_j)\big].$$

If the local fields and the projections $u_k^\top\sigma$, $v_k^\top\tau$ are jointly Gaussian, the same linearization with slopes χ_σ and χ_τ in place of 1 gives the relaxation time and the slow observable

$$\tau_k = \frac{1}{1-\beta\mu_k\sqrt{\chi_\sigma\chi_\tau}}, \qquad \varphi_k = \sqrt{\chi_\tau}\, u_k^\top\sigma + \sqrt{\chi_\sigma}\, v_k^\top\tau, \tag{11}$$

which for $\chi_\sigma = \chi_\tau = 1$ are $1/(1-\theta_k)$ and $\sqrt{2}\,\xi_k$. The derivation is in Appendix A.8.

Let W be any weight matrix with singular value decomposition (22), and let $X_k = u_k^\top\sigma$, $Y_k = v_k^\top\tau$, $h_i = (W\tau)_i$ for $i \le m$ and $g_j = (W^\top\sigma)_j$ for $j \le n$. The conditional means are $\mathbb{E}[\sigma_i \mid x_{\setminus i}] = \tanh(\beta h_i)$ and $\mathbb{E}[\tau_j \mid x_{\setminus(m+j)}] = \tanh(\beta g_j)$, so by (32), exactly as in Lemma A.21,

$$-\mathcal{L}X_k = X_k - \sum_{i\le m} u_{k,i}\tanh(\beta h_i), \qquad -\mathcal{L}Y_k = Y_k - \sum_{j\le n} v_{k,j}\tanh(\beta g_j). \tag{166}$$

The derivation below replaces $\tanh(\beta h_i)$ by $\chi_i\beta h_i$, with χ_i the mean slope of tanh at the field h_i; for $\chi_i = 1$ it returns the linearization of Theorem 4.4.

The Gaussian step. Let G be a linear combination of $X_1, Y_1, X_2, Y_2, \dots$. Assume that under μ_β each pair (h_i, G) is centered and jointly Gaussian; each h_i is a sum of n terms and each G a sum of N terms. For a centered jointly Gaussian pair (h, G) and a differentiable f with bounded derivative, Stein's identity reads $\operatorname{Cov}(f(h), G) = \mathbb{E}[f'(h)]\operatorname{Cov}(h, G)$. With $f = \tanh(\beta\,\cdot)$,

$$\operatorname{Cov}\big(\tanh(\beta h_i), G\big) = \chi_i\,\beta\operatorname{Cov}(h_i, G), \qquad \chi_i = \mathbb{E}\big[1-\tanh^2(\beta h_i)\big]. \tag{167}$$

The block average. Replace each χ_i, $i \le m$, by the average $\chi_\sigma = \frac{1}{m}\sum_{i\le m}\chi_i$. Then, by (167) and $u_k^\top W = \mu_k v_k^\top$,

$$\begin{aligned}\operatorname{Cov}\Big(\sum_{i\le m} u_{k,i}\tanh(\beta h_i),\, G\Big) &= \chi_\sigma\,\beta\operatorname{Cov}\Big(\sum_{i\le m} u_{k,i}(W\tau)_i,\, G\Big)\\ &= \chi_\sigma\,\beta\mu_k\operatorname{Cov}(Y_k, G).\end{aligned}$$

The same computation in the other block, with $\chi_\tau = \frac{1}{n}\sum_{j\le n}\mathbb{E}[1-\tanh^2(\beta g_j)]$ and $Wv_k = \mu_k u_k$, gives $\chi_\tau\beta\mu_k\operatorname{Cov}(X_k, G)$. Substituting into (166): for every such G,

$$\begin{aligned}\operatorname{Cov}(-\mathcal{L}X_k, G) &= \operatorname{Cov}\big(X_k - \chi_\sigma\beta\mu_k Y_k,\, G\big),\\ \operatorname{Cov}(-\mathcal{L}Y_k, G) &= \operatorname{Cov}\big(Y_k - \chi_\tau\beta\mu_k X_k,\, G\big).\end{aligned} \tag{168}$$

The eigenpair. By (168), on an observable $\varphi = aX_k + bY_k$ the operator $-\mathcal{L}$ acts, inside covariances with G, as

$$aX_k + bY_k \;\longmapsto\; (a - b\,\chi_\tau\beta\mu_k)X_k + (b - a\,\chi_\sigma\beta\mu_k)Y_k.$$

The right-hand side equals $\alpha(aX_k + bY_k)$ if and only if $(1-\alpha)a = \chi_\tau\beta\mu_k b$ and $(1-\alpha)b = \chi_\sigma\beta\mu_k a$. Multiplying the two equations gives $(1-\alpha)^2 = (\beta\mu_k)^2\chi_\sigma\chi_\tau$, and the smaller solution is

$$\alpha = 1 - \beta\mu_k\sqrt{\chi_\sigma\chi_\tau}, \qquad a : b = \sqrt{\chi_\tau} : \sqrt{\chi_\sigma},$$

which is (11) with $\tau_k = 1/\alpha$. For $\chi_\sigma = \chi_\tau = 1$ this is the pair $(\xi_k,\, 1-\beta\mu_k)$ of Theorem 4.4.

The measured residual. Proposition 4.3 applies to $\varphi_k = \sqrt{\chi_\tau}X_k + \sqrt{\chi_\sigma}Y_k$ and $\alpha = 1 - \beta\mu_k\sqrt{\chi_\sigma\chi_\tau}$. By (166) the residual $R = \mathcal{L}\varphi_k + \alpha\varphi_k$ is the explicit function

$$R = \sqrt{\chi_\tau}\sum_{i\le m} u_{k,i}\big(\tanh(\beta h_i) - \sigma_i\big) + \sqrt{\chi_\sigma}\sum_{j\le n} v_{k,j}\big(\tanh(\beta g_j) - \tau_j\big) + \alpha\varphi_k$$

of the configuration, so $\eta = \|R\|_{L^2(\mu_\beta)}/\|\varphi_k\|_{L^2(\mu_\beta)}$ is a ratio of two averages under μ_β. On the trained weight matrix of Section 5 we estimate it from 64 block Gibbs chains of 6000 sweeps after 1500 sweeps of burn-in, for the eight leading directions and the six values of β of Figure 4(a). At $\beta\mu_1 = 0.9$ and $k = 1$, $\alpha = 0.200$ and $\eta = 0.034$ for φ_1. For ξ_1, the observable timed in Figure 4(a), the same α has $\eta = 0.046$, so the proposition gives $\tau_{\xi_1} \ge 1/(\alpha+\eta) = 4.08$; the measured time is 4.70 and (11) gives 5.01. With $\alpha = 1 - \beta\mu_1 = 0.100$ the residual of ξ_1 is $\eta = 0.12$, larger than α. Over the 48 cells $\eta/\alpha \le 0.17$ for (φ_k, α), and η increases in the order φ_k, ξ_k, $\sqrt{\chi_\sigma}X_k + \sqrt{\chi_\tau}Y_k$ in every cell.

B Experiments: protocols and further evidence

Chains and estimators. $\mathbb{E}[q_{ab}^2]$ is estimated by running pairs of independent chains and averaging q_{ab}^2 along each pair. The relaxation time of an observable is its integrated autocorrelation time (20) along the single-site chain of Section 2. It is estimated from R independent chains of T sweeps each, after a burn-in. With $\rho(t)$ the autocorrelation of the observable at lag t sweeps, pooled over the R chains and normalized by the variance, the estimate is

$$\hat\tau = \frac{1}{2} + \sum_{t=1}^{W}\rho(t), \tag{169}$$

which is the trapezoidal sum of the integral (20) at unit lag. For an autocorrelation $e^{-\alpha t}$ the sum is not $1/\alpha$ but

$$\frac{1}{2} + \sum_{t\ge 1} e^{-\alpha t} = \frac{1}{2}\coth\frac{\alpha}{2} = \frac{1}{\alpha} + \frac{\alpha}{12} + O(\alpha^3), \tag{170}$$

so the prediction of Theorem 4.4 for $\hat\tau$ is $\frac{1}{2}\coth\big((1-\beta\mu_k)/2\big)$. It exceeds $1/(1-\beta\mu_k)$ by 8% for a direction with $\mu_k = 0$ and by less than 0.1% for $\beta\mu_k = 0.9$. The runs on the trained matrix of Section 5 sample every quarter sweep, where the same sum with step $1/4$ differs from $1/\alpha$ by less than 0.6%. The window W is the smallest lag with $W \ge c\,\hat\tau(W)$, where $\hat\tau(W)$ is the sum up to W and $c = 6$ (Sokal, 1997); the sum is cut there because the terms beyond the window are dominated by noise. The standard error of $\hat\tau$ is the jackknife over chains, that is, the spread of the R estimates obtained by leaving out one chain at a time. The variance of the observable and its standard error are estimated in the same way. Some runs use the block Gibbs chain instead, which resamples all of σ given τ and then all of τ given σ; along the k-th singular direction one of its sweeps covers $\frac{1}{2}(1-\beta\mu_k)^{-1}$ sweeps of the single-site chain. We report all times in single-site sweeps.

Calibration. To test the sampling code we use systems small enough that every quantity of this paper can be computed exactly; the same code produces the estimates on large systems, where only the theory is available for comparison. On systems with $n = 6$ and $n = 8$ units per block every Gibbs expectation is a sum over all 2^N configurations, and the relaxation time of ξ_k is obtained by summing the autocorrelation series of the transition matrix of the chain. Table 2 compares these exact values with the estimates of the code that produces every sampled number in this paper. Over the twenty entries the largest relative deviation is 3.6%, for a relaxation time, and the largest deviation in units of the standard error is 3.0.

Table 2: Exact values on small systems against the sampled estimates. The deviation is estimate minus exact, relative to the exact value and in units of the standard error.

System	Observable	Exact	Estimate	Rel. dev.	Dev./SE
flat pattern, $n = 6$, $\theta = 0.49$	$\mathbb{E}[q_{ab}^2]$	0.08884	0.08869 ± 0.00019	-0.17%	-0.81
flat pattern, $n = 6$, $\theta = 0.80$	$\mathbb{E}[q_{ab}^2]$	0.11339	0.11355 ± 0.00065	$+0.14\%$	$+0.24$
	$\mathrm{Var}(\xi_1)$	2.966	2.965 ± 0.025	-0.01%	-0.01
	τ_{ξ_1}	3.468	3.405 ± 0.064	-1.81%	-0.98
flat pattern, $n = 6$, $\theta = 0.90$	$\mathrm{Var}(\xi_1)$	3.532	3.515 ± 0.027	-0.47%	-0.62
	τ_{ξ_1}	4.42	4.31 ± 0.12	-2.37%	-0.90
flat pattern, $n = 6$, $\theta = 1.20$	$\mathbb{E}[q_{ab}^2]$	0.2613	0.2623 ± 0.0019	$+0.38\%$	$+0.51$
flat pattern, $n = 8$, $\theta = 0.49$	$\mathbb{E}[q_{ab}^2]$	0.06578	0.06578 ± 0.00014	$+0.01\%$	$+0.05$
flat pattern, $n = 8$, $\theta = 0.80$	$\mathbb{E}[q_{ab}^2]$	0.08297	0.08416 ± 0.00040	$+1.44\%$	$+2.96$
	$\mathrm{Var}(\xi_1)$	3.190	3.138 ± 0.031	-1.62%	-1.69
	τ_{ξ_1}	3.63	3.65 ± 0.13	$+0.40\%$	$+0.11$
flat pattern, $n = 8$, $\theta = 0.90$	$\mathrm{Var}(\xi_1)$	3.916	3.994 ± 0.032	$+2.01\%$	$+2.44$
	$\mathrm{Var}(\xi_1)$, $2\times$ chain	3.916	3.910 ± 0.015	-0.15%	-0.38
	τ_{ξ_1}	4.74	4.76 ± 0.18	$+0.40\%$	$+0.11$
	τ_{ξ_1}, $2\times$ chain	4.744	4.572 ± 0.098	-3.62%	-1.76
flat pattern, $n = 8$, $\theta = 1.20$	$\mathbb{E}[q_{ab}^2]$	0.2247	0.2236 ± 0.0020	-0.51%	-0.58
two random patterns, $n = 8$, $\theta_1 = 0.90$	$\mathrm{Var}(\xi_1)$	2.874	2.859 ± 0.020	-0.52%	-0.76
	$\mathrm{Var}(\xi_2)$	1.943	1.954 ± 0.014	$+0.60\%$	$+0.86$
	τ_{ξ_1}	3.83	3.74 ± 0.13	-2.41%	-0.74
	τ_{ξ_2}	2.533	2.506 ± 0.059	-1.07%	-0.45

B.1 Order parameter

Details of Figure 2. For the Curie–Weiss coupling and the flat bipartite pattern of Figure 2 $\mathbb{E}[q_{ab}^2]$ can be computed exactly. On the Curie–Weiss coupling, with $S = \sum_i x_i$,

$$H(x) = -\frac{\mu}{2N}\big(S^2 - N\big), \tag{171}$$

so the energy depends on x only through the number N_+ of coordinates equal to $+1$, since $S = 2N_+ - N$. The $\binom{N}{N_+}$ configurations with the same N_+ have the same Gibbs weight, and the expectation of a function of S is a sum over $N + 1$ values,

$$\mathbb{E}[S^2] = \frac{1}{Z_\beta} \sum_{N_+=0}^{N} \binom{N}{N_+} \exp\Big(\tfrac{\beta\mu}{2N}\big((2N_+ - N)^2 - N\big)\Big)(2N_+ - N)^2. \tag{172}$$

The coordinates are exchangeable, so C_{ij} takes one value c for all $i \neq j$, and $\mathbb{E}[S^2] = \sum_{i,j} C_{ij} = N + N(N-1)\,c$ determines c. By (2),

$$N\,\mathbb{E}[q_{ab}^2] = \frac{\|C\|_F^2}{N} = \frac{N + N(N-1)\,c^2}{N} = 1 + (N-1)\,c^2. \tag{173}$$

On the flat pattern, $W = (\mu/n)\mathbf{1}\mathbf{1}^\top$ with $m = n$, so with $S_\sigma = \sum_i \sigma_i$ and $S_\tau = \sum_j \tau_j$,

$$H(\sigma, \tau) = -\frac{\mu}{n}\, S_\sigma S_\tau, \tag{174}$$

and the energy depends on (σ, τ) only through the numbers m_+ and n_+ of coordinates equal to $+1$ in σ and in τ. The expectation of a function of (S_σ, S_τ) is a sum over the $(n+1)^2$ pairs (m_+, n_+), each with weight

$$\frac{1}{Z_\beta}\binom{n}{m_+}\binom{n}{n_+}\exp\Big(\tfrac{\beta\mu}{n}(2m_+ - n)(2n_+ - n)\Big). \tag{175}$$

Coordinates are exchangeable within each block, and the two blocks are exchangeable with each other, so C_{ij} takes two values off the diagonal: $c_{\sigma\sigma}$ for two distinct coordinates of the same block, determined by $\mathbb{E}[S_\sigma^2] = n + n(n-1)\,c_{\sigma\sigma}$, and $c_{\sigma\tau} = \mathbb{E}[S_\sigma S_\tau]/n^2$ for two coordinates in different blocks. Hence, with $N = 2n$,

$$N\,\mathbb{E}[q_{ab}^2] = \frac{\|C\|_F^2}{N} = \frac{2n + 2n(n-1)\,c_{\sigma\sigma}^2 + 2n^2\,c_{\sigma\tau}^2}{2n}. \tag{176}$$

The circles in Figure 2 are Monte Carlo estimates of the quantity that the lines give exactly. The relative difference between the two is at most 1.3% for $\theta \le 0.95$ and for $\theta \ge 1.1$. It is largest for θ between 1 and 1.05, where the chain is slowest, and there it is at most 4.3%.

Proposition 3.2 gives the bound $N\,\mathbb{E}[q_{ab}^2] \le 1/(1 - \beta\gamma(J))$, drawn in red, which diverges at the threshold $1/\gamma(J)$. In panel (a) this threshold is at $\theta = 1$ up to $O(1/N)$, where the curves leave 1, so the bound covers the whole disordered regime. In panel (b) there is a gap between the two thresholds, from $1/\gamma(J)$ at $\theta = 1/2$ to $1/\lambda_1$ at $\theta = 1$. The red bound diverges at $\theta = 1/2$, but the curves stay close to 1 across the gap: the exact values are 1.008, 1.002, 1.001 at $\theta = 1/2$ and 1.30, 1.12, 1.04 at $\theta = 0.9$ for $N = 128$, 512, 2048. For the bipartite coupling Proposition 3.5 gives the bound $2/(1-\theta)$, drawn in orange, which is finite up to $\theta = 1$ and covers the gap.

General symmetric coupling with a bulk. A bulk gives J negative eigenvalues and so separates the threshold $1/\gamma(J)$ of Proposition 3.2 from $1/\lambda_1$; on the Curie–Weiss coupling of Figure 2(a) the two coincide up to $O(1/N)$. To see the system between the two thresholds on a coupling that is not bipartite, we take a coupling of the form (25) with a bulk: J is $ww^\top + A$ with its diagonal set to zero, where A is a symmetric Gaussian matrix projected on the orthogonal complement of w and scaled to $\|A\|_{\mathrm{op}} = b$. The constant in Theorem 3.3 depends on w through $N\kappa_4(w)$, so we take three values of $N\kappa_4(w)$. We use three unit vectors w: a uniform one, $w_i = N^{-1/2}$, with $N\kappa_4(w) = 1$; a medium one, $w_i \propto 1 + 0.5\,g_i$ with g_i independent standard Gaussians, with $N\kappa_4(w) = 1.74$; and a concentrated one, equal to $(8/N)^{1/2}$ on $N/8$ coordinates and zero elsewhere, with $N\kappa_4(w) = 8$. Each is combined with $b = 0$ and $b = 0.3$, at $N = 512$, 1024 and 2048.

Figure 6(a) shows the uniform w with $b = 0.3$. This coupling has $\lambda_N = -0.30$ and $\gamma(J) = 1.30$, so the threshold $1/\gamma(J)$ of Proposition 3.2 is at $\theta = \lambda_1/\gamma(J) = 0.77$. Between the two thresholds the system remains disordered: at $\theta = 0.8$ and 0.9 the values are at most 1.04 and 1.11, and they decrease with N. The curves leave 1 at $\theta = 1$. This is the same gap as in Figure 2(b), on a coupling that is not bipartite.

In the ordered regime Theorem 3.3 gives $N\,\mathbb{E}[q_{ab}^2] \ge cN$. Table 3 lists the ratio of $N\,\mathbb{E}[q_{ab}^2]$ at $N = 2048$ to its value at $N = 512$, which is 4 if the order parameter is proportional to N. For the uniform and the medium w the twelve ratios lie between 3.99 and 4.03. For the concentrated w they lie between 3.45 and 3.64, because there $N\,\mathbb{E}[q_{ab}^2]$ is below 30 and the constant 1 in it is still visible: with this constant removed, $(N\,\mathbb{E}[q_{ab}^2] - 1)/N$ changes by less than 1% over the three sizes (Figure 6(b)). The constant of proportionality depends on w. At $\theta = 2$ and $b = 0$ it is 0.840 for the uniform w and 0.0130 for the concentrated one, a factor of $64 = (N\kappa_4(w))^2$; the bound (6) depends on $M^2 = N\kappa_4(w)$ in the same way. The red dashed lines in panel (b) are the values $\big((\sum_i m_i^2)^2 - \sum_i m_i^4\big)/N^2$ obtained from the mean-field equations $m_i = \tanh(\beta w_i \bar{m})$, $\bar{m} = \sum_j w_j m_j$; at $b = 0$ the measurements agree with them to within 1.6%.

Further couplings and the location of the transition. Two further families locate the transition at $1/\mu_1$ beyond a single pattern. A Hadamard W has all its singular values equal,

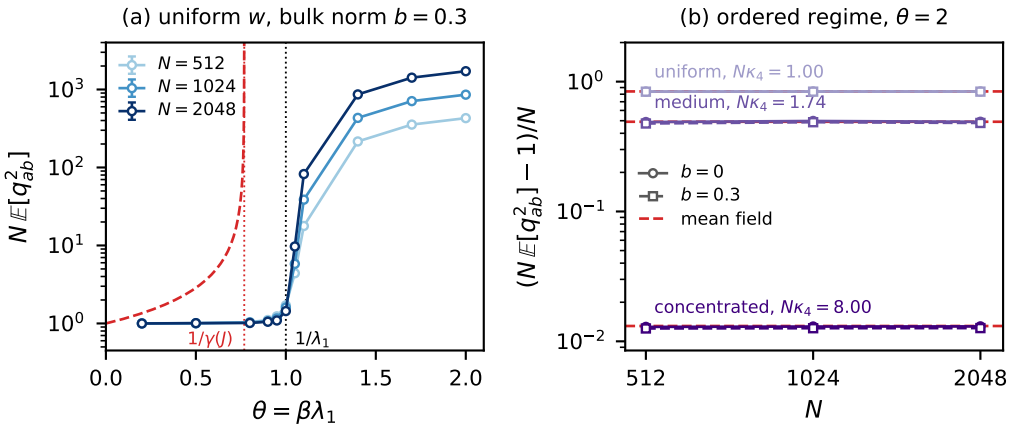

Figure 6: A general symmetric coupling $J = ww^\top + A$ with a bulk A of operator norm b. (a) Uniform w, $b = 0.3$: $N\,\mathbb{E}[q_{ab}^2]$ against $\theta = \beta\lambda_1$ for three sizes. The red dashed curve is the bound of Proposition 3.2; the vertical lines mark the two thresholds. (b) Ordered regime at $\theta = 2$: $(N\,\mathbb{E}[q_{ab}^2] - 1)/N$ against N for the three w, with $b = 0$ and $b = 0.3$; the red dashed lines are the mean-field values at $b = 0$.

Table 3: Ratio of $N\,\mathbb{E}[q_{ab}^2]$ at $N = 2048$ to its value at $N = 512$ in the ordered regime, for the six couplings. $N\kappa_4$ and $\lambda_1/\gamma(J)$ are the values at $N = 2048$.

				$N\,\mathbb{E}[q_{ab}^2]$ ratio, $N = 2048$ to 512		
w prediction	b	$N\kappa_4$	$\lambda_1/\gamma(J)$	$\theta = 1.4$ 4	1.7 4	2.0 4
uniform	0.0	1.000	1.000	4.004	4.002	4.000
uniform	0.3	1.000	0.769	4.010	3.998	3.997
medium	0.0	1.742	0.997	3.988	4.001	4.007
medium	0.3	1.742	0.769	4.021	4.018	4.026
concentrated	0.0	8.000	0.996	3.476	3.597	3.637
concentrated	0.3	8.000	0.768	3.447	3.577	3.619

the least favourable spectrum for a statement about μ_1 alone, and its order parameter is flat in n from 8 to 1024 at $\theta = 0.99$, changing by 0.7% over a factor of 128 in size (Figure 7(a)). For W built from two or three patterns of unequal strength the transition stays at $1/\mu_1$ to within 1.6%, independently of the number of patterns and of their ratio (Figure 7(b), $n = 512$); Appendix A.4 states what the proof of Proposition 3.5 gives for several patterns. The onset of order and the divergence of the directional time $1/(1-\beta\mu_1)$ are both predicted at $\theta = 1$, and we locate the transition with each observable separately. The crossing of a fixed level moves with n, so each location is extrapolated in n: the static onset, taken as the θ at which the Monte Carlo curve of Figure 2(b) reaches 3 and extrapolated in n up to 4096, has envelope [0.9968, 1.0003], and the θ at which the relaxation time of the top direction diverges, extrapolated the same way, has envelope [0.982, 1.020].

B.2 Relaxation times

Details of Figures 3 and 4. On the rank-3 coupling of Figure 3 ($\mu_1 : \mu_2 : \mu_3 = 1 : 0.8 : 0.6$, $\beta\mu_1 = 0.9$), as controls we also measure the projection $(a^\top\sigma + b^\top\tau)/\sqrt{2}$, of the same form as (23), for two pairs of unit vectors (a, b) with a orthogonal to all the u_k and b orthogonal to all the v_k. The singular value of W on these directions is 0, so the predicted time for the controls is 1, and $\frac{1}{2}\coth\frac{1}{2} = 1.08$ for the estimator (170). At $n = 8192$ we measure 10.1 ± 0.6, 3.71 ± 0.11 and 2.20 ± 0.06 for the three directions, and 1.03 ± 0.02 and 1.06 ± 0.02 for the two controls: at a single β, three directions of one matrix relax at times

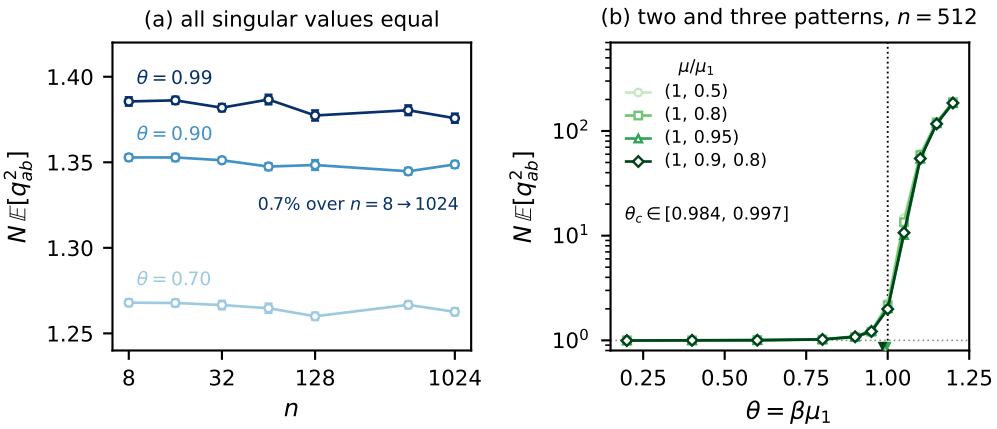

Figure 7: Bipartite couplings beyond a single pattern. (a) W with all singular values equal: $N\,\mathbb{E}[q_{ab}^2]$ against n at three values of $\theta = \beta\mu_1$ below the threshold. (b) W with two or three patterns of unequal strength, with the ratios μ_k/μ_1 given in the legend: $N\,\mathbb{E}[q_{ab}^2]$ against θ at $n = 512$; the triangles on the horizontal axis mark the transition points and the dotted line marks $\theta = 1$.

Table 4: Relaxation times and variances of the three pattern directions and of the controls (mean of the two) at $\beta\mu_1 = 0.9$. The first row gives the predictions of Theorem 4.4, for the times in the form (170) that applies to the estimator; the variance is multiplied by $1 - \beta\mu_k$, so that its prediction is 1.

	τ_{ξ_k} (sweeps)				$\mathrm{Var}(\xi_k)\,(1-\beta\mu_k)$			
n	$k=1$	$k=2$	$k=3$	controls	$k=1$	$k=2$	$k=3$	controls
prediction	10.01	3.59	2.21	1.08	1	1	1	1
128	6.4 ± 0.3	3.3 ± 0.1	2.11 ± 0.05	1.07	0.628	0.878	0.935	0.973
256	7.6 ± 0.4	3.5 ± 0.1	2.19 ± 0.06	1.08	0.728	0.928	0.956	0.978
512	8.1 ± 0.4	3.4 ± 0.1	2.18 ± 0.05	1.05	0.803	0.945	0.976	0.993
1024	9.7 ± 0.4	3.6 ± 0.1	2.18 ± 0.06	1.10	0.896	0.993	0.994	0.993
2048	9.7 ± 0.6	3.7 ± 0.1	2.16 ± 0.06	1.09	0.942	0.998	0.997	0.999
4096	9.4 ± 0.5	3.5 ± 0.1	2.19 ± 0.04	1.09	1.002	0.982	0.993	1.002
8192	10.1 ± 0.6	3.7 ± 0.1	2.20 ± 0.06	1.05	1.010	0.994	0.992	0.999

differing by almost a factor of five, and Figure 3(a) reads the same measurement as $1/\tau_{\xi_k}$ against $\beta\mu_k$, where the prediction is the curve $2\tanh\big((1-\beta\mu_k)/2\big)$, which vanishes at the threshold. In panel (b) the top direction is at 0.64 of its predicted time at $n = 128$ and 1.01 at $n = 8192$, not monotonically in between, while the two weaker directions stay within 5% from $n = 256$ on. The controls agree with their prediction 1.08 to within 3% at every size, and the variance of the same projections agrees with Theorem 4.4(i) to within 3%. Part (iii) of the theorem predicts $\mathrm{Cov}_{\mu_\beta}(\xi_1, \xi_2) = o(1)$; the measured sample correlation of ξ_1 and ξ_2 is at most 0.018 in absolute value at all seven sizes, within two standard errors of zero. Table 4 lists the measured times and variances at all seven sizes.

The matrix of Figure 4 comes from an RBM trained with biases, which are discarded after training. Its singular values below the top are large, $\sum_k \mu_k^2 = 16.2$ with μ_1 normalised to 1, and its order parameter reaches 10 only at $\theta = 1.74$. To see where the order parameter rises when these singular values are small, we also trained an RBM with the same data, architecture and schedule but with its biases held at zero throughout. Its weight matrix is nearly rank one: with μ_1 normalised to 1, $\sum_k \mu_k^2 = 1.05$ and $\mu_2 = 0.045$, against $\sum_k \mu_k^2 = 16.2$ and $\mu_2 = 0.90$ for the first matrix. Figure 8 shows the order parameters of the two matrices. The order parameter of the second reaches 10 at $\theta = 1.20$, closer to 1 than that of

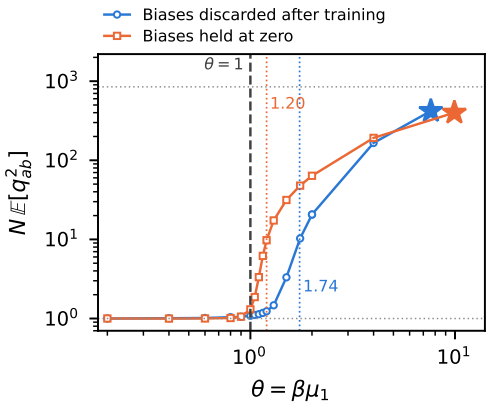

Figure 8: Order parameter of the matrix of Figure 4 (biases discarded after training) and of an RBM trained with its biases held at zero, against $\theta = \beta\mu_1$. Dotted lines mark where $N\,\mathbb{E}[q_{ab}^2] = 10$; stars mark $\beta = 1$.

the first, which reaches 10 at $\theta = 1.74$. The two matrices also differ in the delocalization of their top singular direction, $N\kappa_4(w_1^+) = 4.02$ for the second and 11.05 for the first, so this comparison does not isolate the bulk; in a scan that holds one pattern fixed and increases the bulk from zero to the pattern strength, the transition moves from $\theta = 1.03$ to $\theta = 1.45$. At $\beta = 1$ the unnormalised second matrix has $\mu_1 = 9.93$ and is also deep in the ordered regime.

Relaxation above the threshold. Proposition 4.6 states that for $\beta > 1/\mu$ the relaxation time of the flat pattern is at least $e^{n\Delta_\beta - o(n)}$. This time is too long to be measured by summing an autocorrelation series, but the chain decorrelates only by passing from one ordered state to the other, so we measure the time it needs to leave one of them and compare the growth rate of this time in n with Δ_β. At $\theta = 1.3$ we start 48 chains from the configuration with all spins $+1$ and record, for each chain, the first sweep at which the sum of the two block magnetizations per spin is below $-m_+$, where m_+ is the positive root of $m = \tanh(\theta m)$, so that $-m_+$ is the magnetization per spin of the opposite ordered state. Figure 9(a) shows the logarithm of the median of these times for $n = 24$ to 96. It grows linearly in n with slope 0.0995 ± 0.0035, and $\Delta_\beta = 0.0985$ at $\theta = 1.3$: the two differ by 1.1%, which is 0.3 standard errors. Panel (b) shows the 48 individual times behind each median.

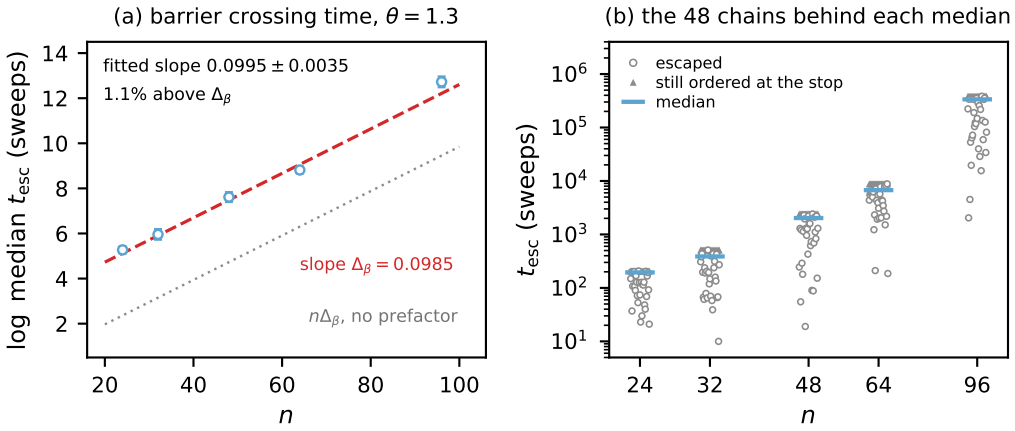

Figure 9: Time to leave an ordered state on the flat pattern at $\theta = 1.3$. (a) Logarithm of the median time against n; the red dashed line has the slope Δ_β of Proposition 4.6, and the grey dotted line is $n\Delta_\beta$ without a prefactor. (b) The 48 chains at each size; each run stops once 60% of the chains have left, and the triangles are the chains still in the initial state at that time.