Every phase mixture recurs
The planar Ising model with random boundary spins

Seth Douglas Nidhal Mghirbi

Version 2.1.0, November 2026

Abstract

Which thermodynamic states recur when an Ising system grows through every square and its exterior spins are independent fair signs? We prove that every convex mixture of the plus and minus phases recurs almost surely in the two-dimensional nearest-neighbour model with bulk-strength boundary bonds, at every fixed inverse temperature $K > K_c$. We also prove the $O(N^{-1/2})$ upper bound on the probability of a fixed local departure from both pure phases for every $K > \log 2$ (in particular $K \geq 1$), where N is the side length. Together with the lower bound, this gives exact order $N^{-1/2}$ for every fixed finite window X and $0 < \varepsilon < \frac{1}{2}\|\mu^+ - \mu^-\|_X$; for larger ε the probability is stretched-exponentially small. Thus the paper resolves both conjectures of van Enter, Netočný and Schaap: full-sequence recurrence throughout coexistence, and the quantitative order in its sharp threshold form for $K > \log 2$. The lower bound holds at every $K > K_c$ for odd side lengths; the sharp upper bound for $K_c < K \leq \log 2$ remains open. The proof combines dual open arcs, a central majority phase label, uniqueness at a uniform wall, and buffered boundary crossings. Joint deletion of arcs gives an exponential bound on extensive coverage, and clipping the phase odds yields the sharp upper exponent. Both parities of the side length are included. We also prove almost-sure convergence along all sizes to the favoured pure phase for each nonzero boundary bias, and recover the equal two-atom empirical metastate throughout the ordered phase. Every mixture recurs, while the sizes near mixed states have density zero.

Claude Opus 5.5, GPT-6.1 Sol, GPT-6.2 Sol and GPT-6.2 Astra assisted with this work; the authors take responsibility for every claim.

Contents

1 The question and the results **2**
2 Model and published inputs **5**
3 Dual lines and exact conditioning **7**
4 Boundary comparisons and separation near the surface **12**
5 Majority odds and phase weights **16**
6 A lower bound in every fixed window **18**
7 Recurrence, bias selection and empirical laws **20**
8 The sharp mixed-state probability **22**
9 Discussion **28**
10 How this work was done **29**
A Scales and constants **30**
References **30**

1 The question and the results

A square Ising system is given a boundary at each size determined by the restrictions of one independent sign field on $\mathbb{Z}^2$. Which infinite-volume Gibbs states can one observe along subsequences of *all* square sizes? A preference of order $\sqrt{n}$ usually selects one pure phase. The question concerns the exceptional sizes at which the two phases have nearly equal free energies. We prove that these sizes realize every mixture weight. We also determine the exact order of the fixed-window probability for $K > \log 2$.

The distinction between recurrence and frequency is essential. The empirical distribution of the finite-volume measures gives mass one half to each pure phase, yet its cluster set contains their entire segment (Figure 1). A sequence can visit every interior neighbourhood infinitely often and spend a vanishing fraction of its time there.

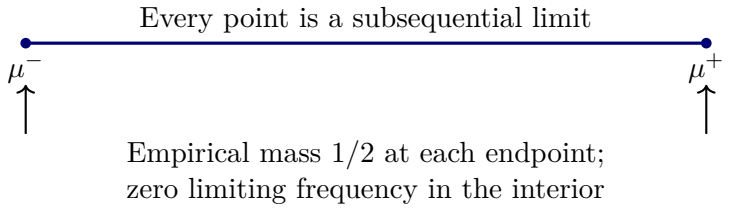

Figure 1: The cluster set and the empirical metastate describe different aspects of the same sequence. The drawing records the conclusions, rather than a finite-size probability density.

Here is the range of the results. Let $K_c = \frac{1}{2}\log(1+\sqrt{2}) = 0.44068679\ldots$ be the critical inverse temperature of the planar nearest-neighbour model [19, 20]. Throughout the main results, K is fixed and finite with $K > K_c$; the sharp upper bound in Theorem D assumes $K > \log 2$. The range $K > K_c$ is the whole region in which the plus and minus phases coexist (for $K \le K_c$ the spontaneous magnetization vanishes [17, 18] and the Gibbs state is unique [21]).

Write $B_n = [-n, n]^2 \cap \mathbb{Z}^2$, and let μ_n^η be the zero-field nearest-neighbour Gibbs measure with exterior spins η and coupling K on every bond meeting B_n. The signs η_x are independent with $\mathbb{P}_p(\eta_x = +1) = p$. Let $\mu^\pm$ be the pure phases, $a_\pm = \mu^\pm(\sigma_0 = +1)$, $m = a_+ - a_- > 0$, and

$$\nu_\lambda = \lambda\mu^+ + (1-\lambda)\mu^-, \qquad X_n = \mu_n^\eta(\sigma_0 = +1), \qquad T_n = \frac{X_n - a_-}{m}.$$

The topology on states is local weak convergence.

The main assertion concerns the complete sequence of integers.

Theorem A (Every mixture recurs)**.** *Let $K > K_c$ and $p = 1/2$. Almost surely the weak cluster set of $(\mu_n^\eta)_{n\ge1}$ is exactly $\{\nu_\lambda : 0 \le \lambda \le 1\}$.*

The quantitative ingredient is a lower bound in arbitrarily narrow fixed windows.

Theorem B (The local-window lower bound)**.** *Let $K > K_c$. For every $0 < u < v < 1$ there are $c = c(K,u,v) > 0$ and $n_0 = n_0(K,u,v)$ such that*

$$\mathbb{P}_{1/2}(u < T_n < v) \ge cn^{-1/2} \qquad (n \ge n_0).$$

For every weakly open neighbourhood $\mathcal{U}$ of any ν_λ with $0 < \lambda < 1$, the same lower bound holds for $\mathbb{P}_{1/2}(\mu_n^\eta \in \mathcal{U})$, with constants depending on K, $\mathcal{U}$ and λ.

Bias selection and the empirical law complete the picture.

Theorem C (Biased and fair laws)**.** *Let $K > K_c$. For each fixed $p > 1/2$, almost surely $\mu_n^\eta \Rightarrow \mu^+$ along all sizes; for $p < 1/2$ the limit is μ^-. For $p = 1/2$, almost surely*

$$\frac{1}{N}\sum_{n=1}^{N} \delta_{\mu_n^\eta} \Rightarrow \tfrac{1}{2}\delta_{\mu^+} + \tfrac{1}{2}\delta_{\mu^-}, \qquad \frac{1}{N}\sum_{n=1}^{N} X_n \longrightarrow \tfrac{1}{2}.$$

For the quantitative upper bound we use the local seminorm of [2, (1)],

$$\|\rho\|_X = \sup\{|\rho(f)| : \|f\|_\infty \le 1,\ f \text{ supported on } X\}, \qquad \Delta_X = \|\mu^+ - \mu^-\|_X.$$

Here X is a fixed finite set. Write $V_{2n+1} = B_n$, $V_{2n} = [-n+1, n]^2 \cap \mathbb{Z}^2$, and $\widehat{\mu}_N^\eta = \mu_{V_N}^\eta$, so N is the integer side length.

Theorem D (The sharp mixed-state probability)**.** *For the two-dimensional Ising model with fair independent bulk-strength boundary spins at any $K > \log 2$ (in particular $K \ge 1$), for every fixed finite X and $\varepsilon > 0$ there is $C = C(K, X, \varepsilon) < \infty$ such that, for all sufficiently large integer N,*

$$\mathbb{P}_{1/2}\left(\min_{s=\pm} \|\widehat{\mu}_N^\eta - \mu^s\|_X \ge \varepsilon\right) \le CN^{-1/2}. \tag{1.1}$$

For $0 < \varepsilon < \Delta_X/2$ this probability is also at least $c(K, X, \varepsilon)N^{-1/2}$, and hence has exact order $N^{-1/2}$. The lower bound itself holds at every $K > K_c$ for odd side lengths, and for both parities when $K > \log 2$. For $\varepsilon > \Delta_X/2$ the probability is at most $CN^2 e^{-cN^{2/5}}$ eventually. At equality only the upper bound is asserted. Every fixed calibrated-weight interval $0 < u < v < 1$ also has probability of exact order $N^{-1/2}$ for $K > \log 2$, with the origin calibration defined as for T_n above.

Theorem D proves the $O(N^{-1/2})$ upper bound conjectured by van Enter, Netočný and Schaap in Remark 3.6 of [2]. Together with Theorem B, it resolves their quantitative conjecture in its sharp threshold form for $K > \log 2$, while Theorem A resolves their full-sequence recurrence conjecture at every $K > K_c$. Read as a two-sided claim for every $\varepsilon > 0$, Remark 3.6 does not hold for $\varepsilon > \Delta_X/2$: a mixture has distance at most $\Delta_X/2$ from the nearer endpoint, and the local approximation error is stretched-exponentially exceptional. The sharp form is the upper bound for all $\varepsilon > 0$, with $\Theta(N^{-1/2})$ for $0 < \varepsilon < \Delta_X/2$. The upper bound for $K_c < K \le \log 2$ remains open, while the lower bound holds there for odd side lengths.

Position in the literature. Theorem A proves the two-dimensional conjecture of van Enter, Netočný and Schaap [2, Remark 3.2], stated there for sufficiently large K, at every $K > K_c$. An earlier form is due to van Enter, Medved' and Netočný [1], for weak boundary couplings. In the frozen-bulk model they found countably many null-recurrent mixtures along all cubes in dimensions two and three. They conjectured that every mixture is an almost-sure limit point at finite low temperature, and they identified a lower bound of order $n^{-(d-1)/2-\zeta}$ on fixed-interval free-energy probabilities as the missing step [1, Sections 1 and 6]. The 2005 paper's Theorem 3.1 concerns sufficiently sparse sizes, and its Theorem 3.4 gives the empirical endpoint law at sufficiently low temperature. We recover the latter conclusion at every $K > K_c$, the full range in which [2, Section 1], following [41], expected the finite-volume states to oscillate between the two phases. Their Proposition 3.5 bounds probabilities of fixed local departures from both phases by $O(n^{-1/2+\alpha})$, at a temperature threshold depending on $\alpha > 0$. Theorem B supplies the lower bound throughout coexistence, and Theorem D proves the sharp upper bound for $K > \log 2$, without an α-dependent threshold. A literature search found no earlier proof of this upper bound. The 2006 note [3, Section 3] explains the same recurrence and metastate distinction; its Remark 3.2 expects every mixture to be a

null-recurrent limit point and names the missing lower bounds on small free-energy differences as the obstacle. Higuchi's Theorem 1 [35] already gives the selection of Theorem C, on his boxes, for independent boundaries with $p < 3/8$ or $p > 5/8$ at sufficiently low temperature; his Theorem 2 shows that boundary density alone does not decide the phase above 3/8. For tuned deterministic boundary conditions every mixture is already a weak limit of square Gibbs states: Gallavotti [37] obtained the symmetric mixture from Dobrushin boundary conditions at low temperature, Messager and Miracle-Solé [38] extended this to all $K > K_c$, and Abraham and Reed [39] obtained every ν_λ by shifting the Dobrushin interface (see [40, Section 1.1.1]); Theorem A shows that boundaries which are not tuned, typical independent fair words, realize every mixture along the sequence of all squares.

Question	Prior result	Result here (ranges shown)
Cluster set	Pure endpoints along sufficiently sparse sizes [2, Theorem 3.1]; countably many mixtures for the frozen bulk [1]	Whole segment at every $K > K_c$
Interior-window probability	Upper bound $O(n^{-1/2+\alpha})$ [2, Proposition 3.5]	Lower bound at every $K > K_c$; exact order $n^{-1/2}$ for $K > \log 2$
Empirical metastate	Equal endpoint masses [2, Theorem 3.4]	Recovery throughout the ordered phase
Biased boundaries	Boundary plus-density below 3/8 selects μ^- at sufficiently low temperature [35, Theorem 1]; related settings [1, 36], [3, Section 2]	Every fixed $p \neq 1/2$: almost-sure convergence along all squares

Table 1: The model in the main results has independent exterior spins and boundary bonds of the bulk strength. Published existential thresholds in the middle column keep their original hypotheses.

Newman and Stein introduced chaotic size dependence, raising the random-boundary Ising question in [41], and developed the empirical metastate viewpoint [42]; Aizenman and Wehr provided the disorder-based metastate construction [43]. The classification of planar DLR states is due to Aizenman and Higuchi [32, 33]; we use its finite-volume form due to Coquille and Velenik [4]. Earlier work on boundary-driven phase selection includes Higuchi [35], Campanino and van Enter [36] for the Potts model at its transition temperature, and Higuchi and Yoshida [34]; see also [3, Section 2]. Endo, van Enter and Le Ny review random boundary conditions in [44]. Their 2024 work [45] obtains mixed-state metastates in one-dimensional long-range models, a related phenomenon in a different interaction geometry. Our proof rests on inputs valid throughout the ordered phase: the random-line representation of the dual high-temperature model developed by Pfister and Velenik [6, 7], surface-order large deviations of the magnetization [11, 12, 24], the spatial relaxation of [23, 24], and the wetting theory of Fröhlich and Pfister [25, 26].

Reading guide and proof idea. Section 2 fixes the model and the published inputs. Section 3 writes the model with its physical boundary in terms of open dual arcs: the boundary word determines their endpoints, reversing a boundary interval prices any family of arcs, and fixing all open arcs leaves pure-boundary Ising laws. Section 4 shows that macroscopic arcs are rare for typical boundaries, that the sign of the spin sum in a central square labels the phase, and that the two ensembles conditioned on this sign are separated near the north side. Section 5 identifies their log odds with the phase weight. The fixed-window lower bound is in Section 6: a long constant-sign boundary buffer makes each selected change

arbitrarily small, while many such changes still move the mean odds by order n. Counting their crossings of a fixed window gives $n^{-1/2}$. Section 7 proves recurrence, bias selection and frequencies. Section 8 proves Theorem D: joint deletion gives an exponential tail for the number of columns reached by short arcs, and a clipped covariance gives the sharp upper exponent. It includes the local-seminorm transfer and an even-side extension with a uniform-permutation and binomial-prefix lower bound. Readers interested in the probability argument can take Section 3, the pure-phase inputs and uniform-wall uniqueness of Section 2 as black boxes.

2 Model and published inputs

Definition 2.1 (Finite-volume states and collars)**.** Let $\Omega = \{-1,+1\}^{\mathbb{Z}^2}$ with its product topology. The effective collar is

$$C_n = \{(\pm(n+1), y) : |y| \le n\} \cup \{(x, \pm(n+1)) : |x| \le n\}.$$

For $\sigma \in \{-1,+1\}^{B_n}$ and $\eta \in \{-1,+1\}^{C_n}$ set

$$W_n^{\eta}(\sigma) = \exp\left\{ K \sum_{\substack{\{x,y\}\subset B_n \\ |x-y|_1=1}} \sigma_x\sigma_y + K \sum_{\substack{x\in B_n,\ y\in C_n \\ |x-y|_1=1}} \sigma_x\eta_y \right\}.$$

Then $\mu_n^{\eta}(\sigma) = W_n^{\eta}(\sigma)/Z_n(\eta)$. To view this as a measure on Ω, fix the spins outside B_n to the full field η. Such an extension does not affect local limits.

There are $M_n = 8n+4$ effective exterior sites. Each has exactly one incident bond into the square. At a corner of the square there are two distinct exterior neighbours. All points of C_n have sup norm $n+1$, so the collars for distinct n are disjoint. Their boundary words are therefore independent under $\mathbb{P}_p$.

Choose an enumeration $(f_j)_{j\ge1}$ of all finite-cylinder indicators and metrize the space $\mathcal{P}(\Omega)$ by

$$d(\rho,\rho') = \sum_{j\ge1} 2^{-j}|\rho(f_j) - \rho'(f_j)|.$$

This space is compact. Every weak limit of states on growing squares satisfies the infinite-volume DLR equations: for a fixed finite set V and a cylinder test f, the identity $\mu_n^{\eta}(f) = \mu_n^{\eta}(\gamma_V f)$ holds once the support and its interaction neighbours lie inside B_n. The function $\gamma_V f$ is again local and continuous. Passing to the limit gives the identity for every such f and V.

We record precisely the published estimates used below. Each applies to this zero-field, unit-bond, planar nearest-neighbour model at every $K > K_c = \frac{1}{2}\log(1+\sqrt{2})$.

Gibbs states and relaxation. Coquille and Velenik [4, Corollary 1.1] give the Aizenman–Higuchi classification

$$\mathcal{G}_K = \{\nu_\lambda : 0 \le \lambda \le 1\}. \tag{2.1}$$

The plus and minus phases are spin reversals; hence $a_+ + a_- = 1$, and $m = a_+ - a_- = \mu^+(\sigma_0)$ is the spontaneous magnetization. Yang's formula $\mu^+(\sigma_0) = (1 - \sinh^{-4} 2K)^{1/8}$ [17] (made rigorous in [18]) shows that $m > 0$, so the two phases are distinct. Equation (A.8) of [4], which records the spatial relaxation of Bricmont, Lebowitz and Pfister [23] and of Chayes,

Chayes and Schonmann [24], and its spin-flip image give, for every finite domain V and local test supported on $S \subset V$,

$$|\mu_V^{\pm}(f) - \mu^{\pm}(f)| \le \|f\|_\infty |S| e^{-c_K \operatorname{dist}(S, V^c)}. \tag{2.2}$$

Up to a factor 2, the case of a general finite V also follows from the case of boxes. In the monotone coupling of μ_V^+ with the restriction of μ^+, which it dominates, the two configurations differ at a site x with probability at most the corresponding difference for the plus law on the largest box about x inside V.

Surface-order large deviations. Let Λ_L be a square box of side L. Ioffe [11, Theorem 1.1, (1.5)] gives

$$\limsup_{L\to\infty} \frac{1}{L} \log \mu_{\Lambda_L}^+\Big(\sum_{x\in\Lambda_L} \sigma_x \le \frac{m}{2} |\Lambda_L| \Big) < 0. \tag{2.3}$$

His Hamiltonian counts each bulk bond once and gives the bonds to the plus exterior unit strength, as here, and his m^* is our m. The threshold $m/2$ lies in his range $[-m, m]$; there his rate has droplet fraction $1/4$ and is strictly negative, since the surface tension is strictly positive for $K > K_c$ [22]. Surface-order upper bounds of this kind are due to Schonmann [12] at low temperature and to Chayes, Chayes and Schonmann [24] for all $K > K_c$ (see [11, Section 1]); (2.3) is Ioffe's sharp form, and only the strict negativity of the rate is used below. The exact Wulff rate is due to Dobrushin, Kotecký and Shlosman [13] and Pfister [14] at low temperature and to Ioffe [11, 15] up to K_c, and the full phase-separation theory up to K_c is due to Ioffe and Schonmann [16].

The dual model. Put $w = e^{-2K}$ and let $K^* = \operatorname{artanh} w$ be the dual coupling of Kramers and Wannier [19]. The map $K \mapsto \operatorname{artanh} e^{-2K}$ is decreasing and fixes K_c, so $0 < K^* < K_c$ exactly when $K > K_c$. For the zero-field Ising model on the dual lattice $(\mathbb{Z} + \frac{1}{2})^2$ at the subcritical coupling K^*, Pfister and Velenik [7, Proposition 2.4 and (6.46)] bound the two-point function by $e^{-\tau(y-x)}$; by [7, Propositions 2.1 and 2.2] the decay rate τ equals the surface tension of the model at coupling K, a norm on $\mathbb{R}^2$ that is strictly positive since $K > K_c$. Hence there is $c_\tau = c_\tau(K) > 0$ with

$$\langle \sigma_x \sigma_y \rangle_{K^*} \le e^{-c_\tau \|x-y\|_\infty}. \tag{2.4}$$

Section 3 also uses the random-line representation of [7, Section 6], which applies to the zero-field free-boundary model with nonnegative couplings on any finite subgraph of the dual lattice. Its dual graphs are such subgraphs with the single coupling K^*, because every bond meeting the domain, including the boundary bonds, has the same coupling K.

The uniform wall. The last input is the half-plane model with all wall spins plus, ordinary bonds one, and wall-to-bulk coupling $h = 1$. Fröhlich and Pfister proved that its plus and minus surface states are extremal and that its surface Gibbs state is unique if and only if they coincide [25]. They also proved that for $K > K_c$ the wall free-energy difference $\tau^- - \tau^+$ equals the surface tension of the wall direction for every $h \ge 1$, and that this saturation is equivalent to uniqueness [26]; the surface tension is strictly positive for $K > K_c$ [22]. Convenient statements are [5, (3.17) and (3.25)]. In this normalization h multiplies the physical wall bonds, so our wall has exactly $h = 1$, the endpoint of that range. Thus this half-plane specification has a unique Gibbs state. Reflection or translation of the wall preserves these hypotheses.

Lemma 2.2 (Uniform identification by one marginal)**.** *Let* $\pi(t) = \min(1, \max(0, t))$. *Then*

$$\sup_\eta d(\mu_n^\eta, \nu_{\pi(T_n(\eta))}) \longrightarrow 0.$$

Proof. Otherwise take sizes tending to infinity and boundary words for which the distance is bounded away from zero. Compactness gives a subsequential limit ρ, which is DLR by the preceding argument and hence equals ν_λ by (2.1). Its origin plus-probability is $a_- + m\lambda$, so along this subsequence $T_n \to \lambda$. Continuity of $\lambda \mapsto \nu_\lambda$ contradicts the distance bound. The choice of exterior extension disappears from each fixed cylinder eventually. □

Lemma 2.3 (Order, association and propagation)**.** *Consider Ising measures with coupling K on a finite set of sites, with arbitrary real external fields; fixed and exterior spins enter as fields. Put*

$$h_K(s) = \frac{1}{1 + e^{-2Ks}}, \qquad \kappa_K = \min_{s \in \{-4,-2,0,2\}} [h_K(s+2) - h_K(s)] > 0.$$

(i) *Such a measure is stochastically increasing in the fields. In particular two such measures with ordered fixed or exterior spins admit an ordered coupling.*

(ii) *Such a measure is positively associated: increasing functions have nonnegative covariance.*

(iii) *Let $\nu \succeq \nu'$ be stochastically ordered measures on $\{-1,+1\}^{B_n}$. Suppose that under both of them the conditional plus-probability of σ_x given all other spins is h_K of the sum of the four neighbouring spins of x, exterior ones read from η. If y is a nearest neighbour of x and $d_z = \nu(\sigma_z = +1) - \nu'(\sigma_z = +1)$, with exterior spins counted as deterministic, then $d_x \geq \kappa_K d_y$.*

Proof. Parts (i) and (ii) are Holley's monotone coupling [27] and the Harris–FKG inequality [28, 29]; we include short proofs. The conditional plus-probability of a site given the others is increasing in its neighbouring spins and in its field. Run single-site heat-bath updates in a fixed cyclic order, with shared uniforms, in two systems with ordered fields, starting from ordered configurations. Monotonicity of the conditional probabilities preserves order, and full sweeps have strictly positive transition probabilities, so each marginal converges to its Gibbs measure. A subsequential limit of the joint laws is an ordered coupling. This proves (i).

For (ii) use induction on the number of sites; one site is immediate. Condition on one spin. The conditional measure of the remaining spins is again of the stated form, so the conditional covariance of two increasing functions is nonnegative by induction. Their conditional means are increasing functions of the conditioning spin by (i), and two increasing functions of one spin have nonnegative covariance. Adding the two terms of the conditional covariance identity proves association.

For (iii), the function $G = h_K(\sum_{z \sim x} \sigma_z) - \kappa_K \mathbf{1}_{\{\sigma_y = +1\}}$ is increasing: raising σ_y moves the neighbour sum from some $s \in \{-4,-2,0,2\}$ to $s+2$ and raises h_K by at least κ_K. By the hypothesis on the conditional probabilities, $d_x - \kappa_K d_y = \nu(G) - \nu'(G) \geq 0$. □

Here and below constants may depend on the fixed K, and later on a fixed buffer width. For the recurrence, lower-bound and empirical-law arguments, the temperature enters only through the published inputs above and the strict inequalities $K^* < K_c$ and $m > 0$, both valid at every $K > K_c$. The joint walk count for the sharp upper bound in Section 8 additionally requires $4e^{-2K} < 1$, or $K > \log 2$.

3 Dual lines and exact conditioning

This section writes a finite-volume model with its physical boundary in terms of the high-temperature expansion of the dual model. We use the random-line representation of Pfister and Velenik [7, Section 6] and its use with fixed boundary spins by Coquille and Velenik [4, Appendix A.2]. We need its exact form with the physical collar, including the corners,

and two consequences: a comparison that prices any family of arcs by reversing a boundary interval, and the conditional law given all open arcs.

Throughout this section $V = ([x_1, x_2] \times [y_1, y_2]) \cap \mathbb{Z}^2$ is a lattice rectangle with at least two sites in each direction. Its collar consists of the exterior sites with a neighbour in V. Each collar site has exactly one bond into V, and the two exterior neighbours of a corner site of V are distinct. Every bond meeting V has coupling K. For a collar word η the weight, the measure μ_V^η and the partition function $Z_V(\eta)$ are defined as in Section 2. Distances and diameters use the sup norm.

3.1 The physical collar and dual sources

Let G be the dual graph of V. Its vertices are the points of $(\mathbb{Z} + \frac{1}{2})^2$ in $[x_1 - \frac{1}{2}, x_2 + \frac{1}{2}] \times [y_1 - \frac{1}{2}, y_2 + \frac{1}{2}]$, and its edges join vertices at distance one. Each edge of G crosses exactly one bond meeting V, and each such bond is crossed by exactly one edge. Order the collar sites cyclically around V, the two exterior neighbours of a corner of V being consecutive. The boundary vertices of G correspond bijectively to the gaps between cyclically consecutive collar sites. At a vertex on a side of G the three crossed bonds form a path between two consecutive collar sites through two sites of V. At a corner of G the two crossed bonds join the two exterior neighbours of a corner of V through that corner. We call the boundary vertices of G positions. A position is a transition of η if its two collar sites carry different spins, and $\mathcal{S}(\eta)$ denotes the set of transitions.

For a set E of edges let ∂E be its set of vertices of odd degree. For a subgraph G' of G and a set Y of its vertices put

$$\mathcal{Z}_{G'}(Y) = \sum_{E \subset E(G') \colon \partial E = Y} w^{|E|}, \qquad w = e^{-2K}.$$

Lemma 3.1 (Collar sources)**.** *For $\sigma \in \{-1, +1\}^V$ let $E(\sigma)$ be the set of edges of G crossing bonds whose endpoints disagree, exterior spins being read from η. Then $\sigma \mapsto E(\sigma)$ is a bijection onto the edge sets E with $\partial E = \mathcal{S}(\eta)$, and*

$$Z_V(\eta) = e^{K|E(G)|}\, \mathcal{Z}_G(\mathcal{S}(\eta)). \tag{3.1}$$

If $a \neq b$ are positions and I is either of the two collar intervals between them, the reversed word η^I has $\mathcal{S}(\eta^I) = \mathcal{S}(\eta) \triangle \{a, b\}$.

Proof. Around an interior vertex of G the four crossed bonds bound a plaquette, so the number of disagreements among them is even. At a position the crossed bonds form a path between two consecutive collar sites, so the parity of its disagreements records whether the position is a transition. Hence $\partial E(\sigma) = \mathcal{S}(\eta)$. Disagreements determine σ by propagation from a collar site, so the map is injective.

For surjectivity, temporarily join cyclically consecutive collar sites by auxiliary exterior edges, using a diagonal edge at each corner of V, and label each auxiliary edge by whether its endpoints disagree in η. Given E with $\partial E = \mathcal{S}(\eta)$, label the bonds crossed by E as disagreeing and the other bonds meeting V as agreeing. The labels then sum to zero modulo 2 around every bounded face of the enlarged planar graph: the interior plaquettes, the four-edge faces along the sides and the triangles at the corners. Every cycle of this graph is a sum modulo 2 of bounded face boundaries, so the labels integrate to a spin configuration starting from one collar spin, and the auxiliary edges make it agree with η on the whole collar. The auxiliary edges carry no energy.

An agreeing bond contributes e^K to the weight and a disagreeing bond $e^{-K} = e^K w$, which gives (3.1). Reversing the collar spins of an interval changes the transitions exactly at its two end positions. □

Lemma 3.1 is the form, for the physical collar with its corners, of [4, (A.2)] and [7, Lemma 6.2]; Lemma 3.6 below makes explicit, for the same collar, the sets $\mathfrak{F}(\Gamma)$ and $\Lambda(\Gamma)$ of [4, Appendix A.2] and the factorization [7, Lemma 6.4].

3.2 Open arcs and their deletion

At each vertex the west and south directions are mates, as are the north and east directions. Every edge set E decomposes into curves by the following local rule of [7, Section 6.1] and [4, Appendix A.2]. At a vertex where E has degree four, each edge is paired with its mate. At degree three exactly one of the two pairs of mates is present; it is paired, and the remaining edge ends there. At a vertex on a side of G, where one direction is missing, this is the unique pair of mates available. At degree two the two edges are paired, and at degree one the edge ends there. Following paired edges decomposes E into edge-disjoint closed and open curves, which may revisit a vertex. Each vertex of odd degree is an endpoint of exactly one open curve. We also call open curves arcs.

For an open curve γ from a to b, let Δ_γ consist of its edges and of their mates at both ends of each edge, restricted to G. This is the edge-boundary of [7, Lemma 6.1 and (6.8)]. We write $G \setminus \Delta_\gamma$ for G with the edges of Δ_γ removed.

Lemma 3.2 (Deleting an arc)**.** *Let γ be an open curve from a to b occurring in some edge set, and let Y be a set of vertices containing a and b. An edge set E with $\partial E = Y$ has γ among its open curves if and only if E contains the edges of γ and no other edge of Δ_γ. Consequently, if $a, b \in \mathcal{S}(\eta)$,*

$$\mu_V^\eta(\gamma \text{ occurs}) = \frac{w^{|\gamma|}\, \mathcal{Z}_{G\setminus\Delta_\gamma}(\mathcal{S}(\eta) \setminus \{a,b\})}{\mathcal{Z}_G(\mathcal{S}(\eta))}. \tag{3.2}$$

Proof. If γ occurs, an occupied mate of one of its edges would be paired with that edge and would belong to γ; so the other edges of Δ_γ are absent. Conversely, suppose that the edges of γ are present and the other edges of Δ_γ absent. Where γ passes through a pair of mates, these two edges remain paired at every possible degree. Where γ passes through two edges that are not mates, the two unused directions are their mates and are absent, so the two edges are paired. At an endpoint of γ the mate of its edge is absent and the degree is odd, so any further edges there form a complementary pair of mates, paired with each other. If γ both ends at and passes through a vertex, its through pair stays paired and its terminal edge stays unpaired; two through visits use their own pairs. Hence γ is traversed as specified whatever the remaining edges are, and these form an edge set of $G \setminus \Delta_\gamma$ with odd set $Y \setminus \{a,b\}$. Summing $w^{|E|}$ over such sets and using Lemma 3.1 gives (3.2). □

For a subgraph G' of G, let $\langle\cdot\rangle_{G'}$ denote the zero-field Ising model on the vertices of G' with coupling K^* on its edges and free boundary. Writing $e^{K^*\sigma_x\sigma_y} = \cosh K^*\,(1 + w\sigma_x\sigma_y)$, since $\tanh K^* = w$, and summing over the spins gives

$$\frac{\mathcal{Z}_{G'}(Y)}{\mathcal{Z}_{G'}(\varnothing)} = \langle\sigma_Y\rangle_{G'}, \qquad \sigma_Y = \prod_{x\in Y}\sigma_x. \tag{3.3}$$

The two-source random-line weight of [7, (6.8)] is

$$\varphi_G(\gamma) = w^{|\gamma|}\frac{\mathcal{Z}_{G\setminus\Delta_\gamma}(\varnothing)}{\mathcal{Z}_G(\varnothing)}. \tag{3.4}$$

Lemma 3.2 with $Y = \{a,b\}$ and (3.3) give the random-line representation $\sum_\gamma \varphi_G(\gamma) = \langle\sigma_a\sigma_b\rangle_G$ of [7, (6.9)], the sum running over the open curves from a to b.

Lemma 3.3 (Second Griffiths inequality)**.** *If $G'' \subset G'$ are subgraphs of a finite piece of the dual lattice, then $\langle \sigma_Y \rangle_{G''} \le \langle \sigma_Y \rangle_{G'}$ for every set Y of vertices of G''.*

Proof. We include Ginibre's two-replica argument [10] for the second Griffiths inequality [8, 9]. Isolated vertices do not change correlations, so we may assume that both graphs have the same vertices. Give each edge e a coupling $J_e \ge 0$; deleting an edge sets its coupling to zero. The derivative of $\langle \sigma_Y \rangle$ in J_e is the covariance of σ_Y and σ_W, where W is the endpoint set of e. For independent replicas σ, τ, twice this covariance is the expectation of $(\sigma_Y - \tau_Y)(\sigma_W - \tau_W)$. Put $s_x = (\sigma_x + \tau_x)/2$ and $t_x = (\sigma_x - \tau_x)/2$. Each of the two differences expands with positive coefficients in monomials of odd degree in t. The replicated Boltzmann factor is $\exp\{2\sum_e J_e(s_x s_y + t_x t_y)\}$, a power series with nonnegative coefficients. Under independent uniform replicas every monomial in the s_x, t_x has nonnegative expectation, since $s_x t_x = 0$ and the powers of each variable have nonnegative expectation by symmetry. The expansion converges absolutely in finite volume, so the derivative is nonnegative. □

Lemma 3.4 (Reversal prices arcs)**.** *Let $a \ne b$ be positions, $\mathcal{L}$ a set of open curves from a to b, and I either collar interval between a and b. Then*

$$\mu_V^\eta(\textit{some } \gamma \in \mathcal{L} \textit{ occurs}) \le \frac{Z_V(\eta^I)}{Z_V(\eta)} \sum_{\gamma \in \mathcal{L}} \varphi_G(\gamma). \tag{3.5}$$

Proof. If a or b is not a transition, no open curve joins them. Otherwise the events for different $\gamma \in \mathcal{L}$ are disjoint, because a is an endpoint of exactly one open curve. By (3.2), (3.3) and Lemma 3.3, with $\mathcal{S} = \mathcal{S}(\eta)$,

$$\mu_V^\eta(\gamma \text{ occurs}) = \frac{w^{|\gamma|} \mathcal{Z}_{G \setminus \Delta_\gamma}(\varnothing) \langle \sigma_{\mathcal{S} \setminus \{a,b\}} \rangle_{G \setminus \Delta_\gamma}}{\mathcal{Z}_G(\mathcal{S})} \le \varphi_G(\gamma) \frac{\mathcal{Z}_G(\mathcal{S} \setminus \{a,b\})}{\mathcal{Z}_G(\mathcal{S})}.$$

By Lemma 3.1 the last ratio is $Z_V(\eta^I)/Z_V(\eta)$. Sum over $\mathcal{L}$. □

3.3 The cost of a large diameter

Lemma 3.5 (Diameter cost)**.** *There are $C, c_0 > 0$, depending only on K, such that for every rectangle V as above, all distinct vertices a, b of its dual graph G, and every $\ell \ge 1$,*

$$\sum_{\gamma \colon a \to b,\ \operatorname{diam} \gamma \ge \ell} \varphi_G(\gamma) \le C e^{-c_0 \ell}, \tag{3.6}$$

where the sum runs over open curves from a to b and diam *is the diameter of the vertex set.*

Proof. Let z be a vertex of G other than a and b. Split each open curve from a to b through z at its first visit to z into two edge-disjoint pieces, from a to z and from z to b. Each piece is an open curve under the local rule, and the rule reassembles the two pieces into the original curve. The product inequality of [7, Lemma 6.5], whose case of a single resulting curve is proved as [6, Lemma 5.4] (display (5.29) there is precisely the bound below), and the representation [7, (6.9)] give

$$\sum_{\gamma \colon a \to b,\ \gamma \ni z} \varphi_G(\gamma) \le \langle \sigma_a \sigma_z \rangle_G \, \langle \sigma_z \sigma_b \rangle_G. \tag{3.7}$$

This is the splitting in the proof of [7, Lemma 6.7] and the first inequality of [4, (A.6)]. For $z \in \{a, b\}$ the left side is at most $\langle \sigma_a \sigma_b \rangle_G$.

By the second Griffiths inequality, Lemma 3.3, adding vertices and edges to G can only increase its correlations. Hence each correlation in G is at most its limit along increasing dual boxes, the infinite-volume correlation of (2.4). Therefore the left side of (3.7) is at most

$e^{-c_\tau(\|a-z\|_\infty+\|z-b\|_\infty)}$ for every vertex z. A curve of diameter at least ℓ has a vertex at distance at least $\ell/2$ from a. There are $8j$ vertices of the dual lattice at distance j from a, so summing over such vertices gives (3.6) with $C=\sum_{j\geq 1} 8je^{-c_\tau j/2}$ and $c_0=c_\tau/4$, uniformly in V, a and b. □

3.4 The law given all open arcs

Lemma 3.6 (Conditioning on the open arcs)**.** *Let Γ be the family of all open curves of some configuration with collar word η, and let $\Delta_\Gamma=\bigcup_{\gamma\in\Gamma}\Delta_\gamma$. Let Θ_Γ be the set of sites of V that are endpoints of bonds crossed by edges of Δ_Γ.*

(i) *The configurations whose family of open curves is Γ are exactly those whose disagreement set contains the edges of Γ and no other edge of Δ_Γ.*

(ii) *There is a configuration $\bar\sigma$ such that these configurations are exactly those agreeing with $\bar\sigma$ on Θ_Γ.*

(iii) *Call the nearest-neighbour components of $V\setminus\Theta_\Gamma$ free components. For each of them, $\bar\sigma$ takes a single sign s on its sites and on all their neighbours outside it, collar sites included. Conditionally on Γ the spins of distinct free components are independent, and each component has the Ising law with constant boundary s.*

(iv) *Every site of Θ_Γ is within distance 2 of a vertex of a curve of Γ. Hence every connected set of sites of V at distance more than 2 from all vertices of curves of Γ lies in one free component.*

Proof. (i) By Lemma 3.2, every curve of Γ occurs exactly when its edges are present and the other edges of its Δ_γ are absent. No edge of one curve of Γ is a mate of an edge of another, since they would be paired. The curves of Γ use every vertex of odd degree, so there are then no further open curves.

(ii) The edges of Γ alone form an edge set with odd set $\mathcal{S}(\eta)$, so Lemma 3.1 gives a configuration $\bar\sigma$ whose disagreement set is exactly the edge set of Γ. Prescribing presence or absence on Δ_Γ prescribes, for each bond crossed by Δ_Γ, whether its endpoints agree. Every component of the graph of these bonds contains a bond to the collar. Indeed, follow a curve of Γ from its starting position. Its first edge crosses a collar bond, or it points inward and its tangential mate, which belongs to Δ_Γ, crosses a collar bond sharing an endpoint with the first crossed bond. At a turn consecutive crossed bonds share an endpoint. At a straight passage the two crossed bonds are opposite sides of a plaquette, and the mates of the traversed edges cross the two other sides; on a side of G one such mate, crossing the bond between the two adjacent sites of V, suffices. Each bond crossed by a mate shares an endpoint with the bond crossed by its partner. Repeated visits are covered by the same local facts. Propagation from the collar therefore fixes the spins on Θ_Γ to those of $\bar\sigma$. Conversely, spins equal to $\bar\sigma$ on Θ_Γ reproduce the prescribed disagreements on Δ_Γ, hence the family Γ by (i).

(iii) The disagreeing bonds of $\bar\sigma$ are crossed by edges of Γ, so both their endpoints lie in Θ_Γ or the collar. Every bond with an endpoint in a free component therefore agrees in $\bar\sigma$, and $\bar\sigma$ is constant on each free component together with its outside neighbours. By (ii), conditioning on Γ fixes the spins on Θ_Γ and leaves the free spins unrestricted. Restricting the Hamiltonian to the free sites gives the asserted product of constant-boundary laws.

(iv) Every edge of Δ_Γ is an edge of a curve of Γ or the mate of such an edge at one of its vertices, and the endpoints of the bond it crosses lie within distance one of that vertex. □

4 Boundary comparisons and separation near the surface

Put $k_n = \lfloor n^{4/5} \rfloor$ and, for $s \in \{+1, -1\}$,

$$\mathcal{A}^s = \Big\{ s \sum_{x \in B_{k_n}} \sigma_x > 0 \Big\}.$$

Since $|B_{k_n}|$ is odd, $\mathcal{A}^+$ and $\mathcal{A}^-$ partition the configurations. Besides the unrestricted partition function Z_n, use Z_n^s, the sum of W_n^η over $\mathcal{A}^s$. All three retain the physical boundary bonds. Superscripts + and − on these partition functions always mean the majority constraint. Write $\mu_n^{\eta,s} = \mu_n^\eta(\cdot \mid \mathcal{A}^s)$ and define

$$F_n(\eta) = \log Z_n^+(\eta) - \log Z_n^-(\eta). \tag{4.1}$$

Then, with $\mathrm{logistic}(z) = e^z/(1+e^z)$,

$$\mathrm{logistic}(F_n) = \frac{Z_n^+}{Z_n^+ + Z_n^-} = \mu_n^\eta(\mathcal{A}^+). \tag{4.2}$$

The event $\mathcal{A}^+$ is increasing and $\mathcal{A}^-$ is decreasing, so association, Lemma 2.3(ii), gives for every increasing f

$$\mu_n^{\eta,+}(f) \geq \mu_n^\eta(f) \geq \mu_n^{\eta,-}(f). \tag{4.3}$$

In particular $\mu_n^{\eta,+}$ dominates $\mu_n^{\eta,-}$ stochastically. At every site outside B_{k_n} the conditional plus-probability under either conditioned law is the unconstrained one, h_K of the neighbour sum.

We apply Section 3 to $V = B_n$, whose collar is C_n. Its dual graph G_n has M_n positions.

4.1 Concentration of interval comparisons

Order the M_n exterior bonds cyclically around the square, with the two bonds at a corner in their perimeter order. For a cyclic interval I, let η^I reverse precisely its exterior spins. The empty and full intervals are allowed. For any one of the three partition functions write

$$D_I(\eta) = \log Z(\eta^I) - \log Z(\eta), \qquad m_I = |I|.$$

Changing one exterior bit changes $\log Z$ by at most $2K$, and hence D_I by at most $4K$.

We shall also need a slightly dependent law on the boundary. Fix an odd integer $b = 2L+1$. Partition the first $b\lfloor (2\lfloor n/2 \rfloor + 1)/b \rfloor$ sites of the middle half $\{(x, n+1) : |x| \leq \lfloor n/2 \rfloor\}$ of the north collar, read from west to east, into consecutive blocks of length b. All remaining bits are singletons. In a block the $2L$ noncentral bits are independent fair signs. Call its centre active if all those bits are plus. Conditional on the noncentral bits, an active centre is plus with probability q, and an inactive centre is fair. Different blocks and singletons are independent; the latter are fair. Denote this law by $\mathbb{P}_q^{(L)}$. At $q = 1/2$ it is exactly the fair product law on every bit.

Lemma 4.1 (Interval-ratio bounds)**.** *For product boundary bias $|p - 1/2| \leq n^{-2/5}$, and also for the block law with $|q - 1/2| \leq n^{-2/5}$, the following hold. In the product case take $b = 1$.*

(i) *Let H_n be the event that $D_I \leq n^{7/10}$ for every cyclic interval I and each of Z_n, Z_n^+, Z_n^-. Then*

$$\mathbb{P}(H_n^c) \leq Cn^2 \exp\{-n^{2/5}/(288K^2b^2)\} \tag{4.4}$$

for all sufficiently large n.

(ii) *Fix $t > 0$. If $m_I \le 2\ell$ and $\ell \ge 1$, then, for all sufficiently large n uniformly in I and ℓ,*

$$\mathbb{P}(D_I > t\ell) \le \exp\{-t^2\ell/(64K^2b^2)\}. \tag{4.5}$$

(iii) *On H_n, $|F_n| \le n^{7/10}$ and $\mu_n^\eta(\mathcal{A}^s) \ge \frac{1}{2}e^{-n^{7/10}}$ for $s = \pm 1$.*

Proof. We give the calculation behind the bounded-differences inequality of Hoeffding and McDiarmid [30, 31]. If independent coordinates change a function G by at most d_j, its Doob martingale increments have conditional ranges at most d_j. For a centred variable in an interval of length d, the logarithm of its moment generating function has second derivative at most $d^2/4$ (the variance under the tilted law); integration twice gives an upper bound $s^2d^2/8$. Iterating conditional expectations and optimizing Markov's inequality yields

$$\mathbb{P}(G - \mathbb{E}G \ge v) \le \exp\left\{-\frac{2v^2}{\sum_j d_j^2}\right\}, \qquad \mathrm{Var}(G) \le \frac{1}{4}\sum_j d_j^2. \tag{4.6}$$

The same bound holds for the lower tail, and the variance bound follows as well by orthogonality of martingale increments.

For product spins, condition on the bits outside I. Reversing I replaces bias p by $1-p$. Coordinatewise coupling of the two product distributions changes an expected $|2p-1|m_I$ bits, giving

$$|\mathbb{E}(D_I \mid I^c)| \le 2K|2p-1|m_I.$$

The m_I oscillations are at most $4K$, so the upper tail beyond this conditional mean by v is at most $\exp\{-v^2/(8K^2m_I)\}$.

For blocks, condition only on complete groups disjoint from I. There are at most m_I intersected groups, including partially intersected endpoint blocks. At fairness their remaining bits are all independent fair, and reversing I preserves this conditional law, so the conditional mean of D_I is zero. Couple a general q to $1/2$, sharing all buffers and centre uniforms. Only a centre in an intersected block can change, with probability at most $|q-1/2|$. A changed centre, even outside I in an endpoint block, changes D_I by at most $4K$. Consequently

$$\begin{aligned} |\mathbb{E}(D_I \mid \text{disjoint groups})| &\le 4Km_I|q-1/2|, \\ \mathbb{P}(D_I - \mathbb{E}(D_I \mid \text{disjoint groups}) \ge v \mid \cdots) &\le e^{-v^2/(8K^2b^2m_I)}. \end{aligned} \tag{4.7}$$

A whole group has oscillation at most $4Kb$. This proves the displayed estimate by (4.6). The same mean upper bound works for the product law.

For $n \ge 4$, $M_n \le 9n$. The conditional mean is at most $36Kn^{3/5}$, eventually at most $n^{7/10}/2$. Taking $v = n^{7/10}/2$ in (4.7) and $m_I \le 9n$ gives the exponent in (4.4). There are at most $M_n(M_n+1)$ intervals and three ensembles. Empty intervals contribute zero.

For the local estimate, $m_I \le 2\ell$ and the mean is at most $8Kn^{-2/5}\ell$. Eventually it is at most $t\ell/2$. Use $v = t\ell/2$ in (4.7); this gives (4.5). The cutoffs depend on the fixed K, t, b, not on ℓ or the interval.

For (iii), spin reversal maps $\mathcal{A}^+$ to $\mathcal{A}^-$, so reversing the full collar gives $Z_n^+(\eta^{C_n}) = Z_n^-(\eta)$. Hence the full-interval values of D_{C_n} for Z_n^+ and Z_n^- are $-F_n$ and F_n, and both are at most $n^{7/10}$ on H_n. Then $\mu_n^\eta(\mathcal{A}^s) = \mathrm{logistic}(sF_n) \ge \mathrm{logistic}(-n^{7/10}) \ge \frac{1}{2}e^{-n^{7/10}}$ by (4.2). □

4.2 Macroscopic arcs and the centre

Lemma 4.2 (Macroscopic arcs)**.** *Let $\mathcal{E}_n$ be the event that some open curve of the configuration in B_n has diameter at least $n/10$. There is $c > 0$ such that for $\eta \in H_n$ and all large n*

$$\mu_n^\eta(\mathcal{E}_n) \le e^{-cn}, \qquad \mu_n^{\eta,s}(\mathcal{E}_n) \le e^{-cn} \quad (s = \pm 1). \tag{4.8}$$

Outside $\mathcal{E}_n$, for large n, the square $B_{\lfloor 2n/3 \rfloor}$ lies in one free component of Lemma 3.6. We call it the central component and write s_Γ for its boundary sign.

Proof. For distinct positions a, b of G_n choose one collar interval I between them. Lemmas 3.4 and 3.5 bound the probability that an open curve from a to b of diameter at least $n/10$ occurs by $e^{D_I} C e^{-c_0 n/10}$, with D_I for the unrestricted ensemble. On H_n this is at most $C e^{n^{7/10} - c_0 n/10}$. Summing over at most $81n^2$ pairs proves the first bound for large n, with any $c < c_0/10$. The second follows from Lemma 4.1(iii) after decreasing c, since $\mu_n^{\eta,s}(\mathcal{E}_n) \le 2e^{n^{7/10}} \mu_n^{\eta}(\mathcal{E}_n)$.

Outside $\mathcal{E}_n$ every vertex of an open curve lies within $n/10$ of its endpoint, a position at sup norm $n + \frac{1}{2}$. All sites of $B_{\lfloor 2n/3 \rfloor}$ are therefore at distance more than 2 from these vertices for large n, and Lemma 3.6(iv) applies. □

4.3 A wrong majority in a pure domain

Lemma 4.3 (Rare wrong majorities)**.** *There are $C_{\mathrm{maj}}, c_{\mathrm{maj}} > 0$ such that for every $k \ge 1$ and every finite domain $V \supset B_k$*

$$\mu_V^+\Big(\sum_{x \in B_k} \sigma_x < 0 \Big) \le C_{\mathrm{maj}} e^{-c_{\mathrm{maj}} k}, \qquad \mu_V^-\Big(\sum_{x \in B_k} \sigma_x > 0 \Big) \le C_{\mathrm{maj}} e^{-c_{\mathrm{maj}} k}. \tag{4.9}$$

Consequently, for large n and every family Γ of open curves outside $\mathcal{E}_n$,

$$\mu_n^{\eta}(\mathcal{A}^{-s_\Gamma} \mid \Gamma) \le \delta_n := C_{\mathrm{maj}} e^{-c_{\mathrm{maj}} k_n}. \tag{4.10}$$

Proof. Fix $\varepsilon > 0$ with $1 - (1+\varepsilon)^{-2} < m/2$ and let $B' = B_{\lfloor (1+\varepsilon)k \rfloor}$. If the spin sum over B_k is negative, the average of σ over B' is less than $1 - |B_k|/|B'|$, which is below $m/2$ for large k. By (2.3) the plus law on B' gives this event probability at most e^{-ck} for large k. The event is supported on B_k, at distance at least $\varepsilon k - 1$ from the complement of B', so (2.2) gives

$$\mu^+\Big(\sum_{x \in B_k} \sigma_x < 0 \Big) \le e^{-ck} + |B_k| e^{-c_K(\varepsilon k - 1)}.$$

After adjusting constants this is at most $C_{\mathrm{maj}} e^{-c_{\mathrm{maj}} k}$ for every $k \ge 1$. By the DLR equations and Lemma 2.3(i), the plus law on V dominates the restriction of μ^+ to V. The event is decreasing, which proves the first bound in (4.9); spin reversal gives the second.

For (4.10), Lemma 3.6(iii) makes the conditional law on the central component the Ising law with constant boundary s_Γ. For large n this component contains B_{k_n}, and (4.9) applies there with $k = k_n$. □

4.4 Short arcs under the majority constraint

The majority constraint is invisible near the north side. Let

$$\mathcal{N}_n = \{(x, y) \in B_n : y \ge \lceil 3n/4 \rceil\},$$

a rectangle disjoint from B_{k_n} with at least two rows, for large n.

Lemma 4.4 (Strip comparison)**.** *Let $a \ne b$ be positions on the north side of G_n, let I be the north interval of collar sites between them, let $\ell \ge 1$, and let $\mathcal{L}$ be a set of open curves from a to b of diameter at least ℓ, each of whose vertices is at distance at least 2 from the south, east and west sides of the dual graph of $\mathcal{N}_n$. Then for large n, every η and $s = \pm 1$,*

$$\mu_n^{\eta,s}(\text{some } \gamma \in \mathcal{L} \text{ occurs}) \le C e^{-c_0 \ell} \frac{Z_n^s(\eta^I)}{Z_n^s(\eta)}, \tag{4.11}$$

with C and c_0 from Lemma 3.5.

Proof. Write a configuration of B_n as (σ', ξ), with σ' on $\mathcal{N}_n$ and ξ on $B_n \setminus \mathcal{N}_n$. The collar of $\mathcal{N}_n$ consists of the north collar of B_n, the parts of its east and west collars beside $\mathcal{N}_n$, and the row of B_n just below $\mathcal{N}_n$. Given ξ, its collar word η_ξ is read from η and ξ. The weight factorizes as $W_n^\eta(\sigma', \xi) = W_{\mathcal{N}_n}^{\eta_\xi}(\sigma')R(\xi)$, where $R(\xi)$ collects the bonds not meeting $\mathcal{N}_n$. It does not depend on the north collar, whose bonds all meet $\mathcal{N}_n$. The event $\mathcal{A}^s$ depends on ξ alone.

The dual graph of $\mathcal{N}_n$ is the part of G_n at height above $\lceil 3n/4 \rceil - 1$. At vertices at distance at least 2 from its south, east and west sides, the two graphs have the same incident edges, crossing the same bonds with the same spins. Hence a curve of $\mathcal{L}$ occurs in (σ', ξ) for B_n exactly when it occurs in σ' for $\mathcal{N}_n$ with collar η_ξ. For each ξ, Lemmas 3.4 and 3.5 in $\mathcal{N}_n$ give

$$\sum_{\sigma' : \text{some } \gamma \in \mathcal{L} \text{ occurs}} W_{\mathcal{N}_n}^{\eta_\xi}(\sigma') \le Ce^{-c_0 \ell} Z_{\mathcal{N}_n}((\eta^I)_\xi).$$

Multiply by $R(\xi)$ and sum over $\xi \in \mathcal{A}^s$. The left side becomes the unnormalized weight of the event intersected with $\mathcal{A}^s$. The right side becomes $Ce^{-c_0 \ell} Z_n^s(\eta^I)$, since reversing I changes only bonds meeting $\mathcal{N}_n$. Divide by $Z_n^s(\eta)$. □

4.5 Separation at a fixed depth

For $|x| \le \lfloor n/2 \rfloor$ and a fixed even integer $D \ge 12$, let

$$i = (x, n+1), \qquad y = (x, n), \qquad v = (x, n - D).$$

The site v will measure separation between the two majority-conditioned ensembles. Define the defects

$$R_{n,v}^s(\eta) = |\mu_n^{\eta,s}(\sigma_v = +1) - a_s| \qquad (s = \pm 1).$$

Lemma 4.5 (Separation near the north side)**.** *For product or fixed-width block laws in Lemma 4.1 there are $C_b, c_b > 0$, independent of D, such that*

$$\limsup_{n \to \infty} \sup_{|x| \le \lfloor n/2 \rfloor} \sup_{bias} \mathbb{E} R_{n,v}^s \le C_b e^{-c_b D}, \tag{4.12}$$

the supremum running over the permitted bias parameters. The plus-probability gap $d_v = \mu_n^{\eta,+}(\sigma_v = +1) - \mu_n^{\eta,-}(\sigma_v = +1)$ satisfies

$$d_v \ge m - R_{n,v}^+ - R_{n,v}^-. \tag{4.13}$$

Proof. Let Q_v be the union of the vertical path from $(x, 0)$ to v and the square $v + B_{D/2}$. Call an open curve short if its diameter ℓ is less than $n/10$, and relevant if one of its vertices lies within distance 2 of Q_v. Consider a relevant short curve. Its endpoints lie within $\ell + 2$ of Q_v, whose points have abscissa within $D/2$ of x and nonnegative ordinate. For large n the south, east and west sides of G_n are farther away, so both endpoints are positions on the north side. The vertex near Q_v has ordinate at most $n - D/2 + 2$, whereas the endpoints have ordinate $n + \frac{1}{2}$; hence $\ell \ge D/2 - 2$. Every vertex of the curve is within ℓ of the vertex near Q_v, hence within horizontal distance $\ell + D/2 + 2 \le 2\ell + 4$ of x. There are therefore at most $C(\ell + 1)^2$ endpoint pairs at diameter ℓ, and the north interval I between the endpoints has $m_I \le \ell$. Every vertex lies within $\ell < n/10$ of a north position, so for large n it is far above the south side of the dual graph of $\mathcal{N}_n$, and at horizontal distance at most $n/2 + 2\ell + 4 < n - 2$ from the origin. Lemma 4.4 therefore applies to these curves, uniformly in $|x| \le \lfloor n/2 \rfloor$.

For each endpoint pair and each $\ell \ge D/2 - 2$, charge the event $\{D_I > c_0 \ell/2\}$ for Z_n^s once. Off this event, (4.11) bounds the conditional probability of the curves of diameter ℓ by

$Ce^{-c_0\ell/2}$; on it, the bound is one. Lemma 4.1(ii) with $t = c_0/2$ gives, for all large n, uniformly in x and the bias,

$$\mathbb{E}\,\mu_n^{\eta,s}(\text{some relevant short curve}) \le \sum_{\ell \ge D/2-2} C(\ell+1)^2 \left[Ce^{-c_0\ell/2} + e^{-c_0^2\ell/(256K^2b^2)}\right] \le C_b e^{-c_b D}. \tag{4.14}$$

The disorder exception is charged per endpoint pair and diameter, never per curve.

Now fix a family Γ of open curves outside $\mathcal{E}_n$ with no relevant short curve. By Lemma 3.6(iv), Q_v lies in one free component. It meets $B_{\lfloor 2n/3 \rfloor}$, so this is the central component, which therefore contains $v + B_{D/2}$. By (2.2),

$$|\mu_n^\eta(\sigma_v = +1 \mid \Gamma) - a_{s_\Gamma}| \le e^{-c_K D/2}.$$

Suppose moreover that $s_\Gamma = s$. For every $[0,1]$-valued g, the mixture identity

$$\mu_n^\eta(g \mid \Gamma, \mathcal{A}^s) - \mu_n^\eta(g \mid \Gamma) = \mu_n^\eta(\mathcal{A}^{-s} \mid \Gamma)\left[\mu_n^\eta(g \mid \Gamma, \mathcal{A}^s) - \mu_n^\eta(g \mid \Gamma, \mathcal{A}^{-s})\right]$$

and (4.10) show that conditioning further on $\mathcal{A}^s$ moves the expectation of g by at most δ_n. Hence $|\mu_n^{\eta,s}(\sigma_v = +1 \mid \Gamma) - a_s| \le e^{-c_K D/2} + \delta_n$ for such Γ.

The remaining families have small $\mu_n^{\eta,s}$-probability. For families outside $\mathcal{E}_n$ with $s_\Gamma = -s$, (4.10) bounds the unrestricted probability of their intersection with $\mathcal{A}^s$ by δ_n times their probability. Their total $\mu_n^{\eta,s}$-probability is therefore at most $\delta_n/\mu_n^\eta(\mathcal{A}^s)$, which is at most $2\delta_n e^{n^{7/10}}$ on H_n by Lemma 4.1(iii). This tends to zero because $7/10 < 4/5$. On H_n the long families have $\mu_n^{\eta,s}$-probability at most e^{-cn} by Lemma 4.2. Averaging over Γ under $\mu_n^{\eta,s}$ gives, on H_n,

$$R_{n,v}^s \le e^{-c_K D/2} + \delta_n + 2\delta_n e^{n^{7/10}} + e^{-cn} + \mu_n^{\eta,s}(\text{some relevant short curve}),$$

and $R_{n,v}^s \le 1$ on H_n^c. Take expectations, and use (4.4) and (4.14). After adjusting C_b and c_b this proves (4.12). The division by $\mu_n^\eta(\mathcal{A}^s)$ is applied only to the long and wrong-sign families, whose probabilities carry the factors e^{-cn} and δ_n.

Finally $\mu_n^{\eta,+}(\sigma_v = +1) \ge a_+ - R_{n,v}^+$ and $\mu_n^{\eta,-}(\sigma_v = +1) \le a_- + R_{n,v}^-$ give (4.13). □

5 Majority odds and phase weights

For an exterior coordinate i, write $\operatorname{osc}_i F_n$ for its value with that bit plus minus its value with that bit minus, with every other bit fixed.

Lemma 5.1 (Monotone odds and a lower response)**.** *The function F_n is odd under simultaneous spin reversal, is increasing in every exterior bit, and $0 \le \operatorname{osc}_i F_n \le 4K$. For fixed D and all n with $n - D > k_n$, at the north coordinate i and depth-D site v of Lemma 4.5,*

$$\operatorname{osc}_i F_n \ge \gamma_D d_v(\eta), \qquad \gamma_D = (1 - e^{-4K})\kappa_K^D > 0, \tag{5.1}$$

where κ_K is the heat-bath increment of Lemma 2.3. This holds for either current value of the bit in η.

Proof. Spin reversal exchanges the majority-constrained partition functions, proving oddness. Suppose first the current bit is minus. Switching it to plus multiplies each constrained weight by $e^{2K\sigma_y}$. Hence the change in F_n is

$$\log \mu_n^{\eta,+}(e^{2K\sigma_y}) - \log \mu_n^{\eta,-}(e^{2K\sigma_y}).$$

The two kernels are ordered by (4.3). As a function of the plus-probability u, the logarithm here is

$$\log[e^{-2K} + (e^{2K} - e^{-2K})u],$$

whose derivative is at least $1 - e^{-4K}$. If the current bit is plus, compute the reversed tilt with $e^{-2K\sigma_y}$; the negative of its logarithm has the same lower derivative. This gives the lower bound using the gap at y in the current kernels. For fixed D and large n the vertical path from y to v avoids B_{k_n}, so both kernels have the heat-bath conditional probabilities at its sites. Propagate the gap from v along the D edges of this path by Lemma 2.3(iii). The upper bound follows because each log partition function changes by at most $2K$. □

Lemma 5.2 (Near-fair drift)**.** *There are constants $d_* > 0$ and $c_* > 0$ such that, at product bias $p_n = 1/2 + n^{-2/5}$,*

$$\mathbb{E}_{p_n} F_n \geq d_* n^{3/5}, \qquad \mathbb{P}_{p_n}(F_n < \tfrac{1}{2} d_* n^{3/5}) \leq e^{-c_* n^{1/5}} \tag{5.2}$$

for all sufficiently large n.

Proof. Choose a fixed even $D \geq 12$ so that each expected remainder in (4.13) is eventually below $m/8$, uniformly for $p \in [1/2, p_n]$. Then $\mathbb{E}_p d_v \geq 3m/4$ at at least n north positions. Differentiating the finite product distribution gives

$$\frac{d}{dp}\mathbb{E}_p F_n = \sum_{i \in C_n} \mathbb{E}_p \operatorname{osc}_i F_n \geq \frac{3m\gamma_D}{4} n.$$

Indeed conditioning on all coordinates other than i makes their contribution the two-value difference $\operatorname{osc}_i F_n$; differentiation and summation is exact for the finite polynomial expectation. All omitted terms are nonnegative. Oddness gives $\mathbb{E}_{1/2} F_n = 0$. Integrate over an interval of length $n^{-2/5}$, with $d_* = 3m\gamma_D/4$. The lower-tail version of (4.6), using $M_n \leq 9n$ oscillations at most $4K$ and deviation $d_* n^{3/5}/2$, gives (5.2), for instance with $c_* = d_*^2/(288K^2)$. □

5.1 The majority as a phase label

The odds F_n compare complementary events, so (4.2) turns them into an exact probability. It remains to identify that probability with the phase weight at the origin.

Lemma 5.3 (Transfer of the majority odds)**.** *On H_n from Lemma 4.1, for some $C, c > 0$ and all large n,*

$$|X_n - [a_- + m \operatorname{logistic}(F_n)]| \leq C e^{-c n^{4/5}}. \tag{5.3}$$

The event and the bound apply to the product and fixed-width block laws above.

Proof. Average over the family Γ of open curves. For Γ outside $\mathcal{E}_n$ the origin lies in the central component of Lemma 4.2, at distance more than $2n/3$ from its complement. Lemma 3.6(iii) and (2.2) give

$$|\mu_n^\eta(\sigma_0 = +1 \mid \Gamma) - a_{s_\Gamma}| \leq e^{-2c_K n/3}.$$

Hence X_n differs from $a_- + m\,\mu_n^\eta(\{s_\Gamma = +1\} \setminus \mathcal{E}_n)$ by at most $e^{-2c_K n/3} + \mu_n^\eta(\mathcal{E}_n)$. By (4.10), $\mu_n^\eta(\mathcal{A}^+ \mid \Gamma)$ lies in $[1 - \delta_n, 1]$ when $s_\Gamma = +1$ and in $[0, \delta_n]$ when $s_\Gamma = -1$. Therefore $\mu_n^\eta(\mathcal{A}^+)$ differs from $\mu_n^\eta(\{s_\Gamma = +1\} \setminus \mathcal{E}_n)$ by at most $\delta_n + \mu_n^\eta(\mathcal{E}_n)$. Combine these bounds with (4.2), (4.8) and $\delta_n = C_{\mathrm{maj}} e^{-c_{\mathrm{maj}} k_n}$. □

6 A lower bound in every fixed window

The bounded one-bit response $4K$ alone cannot force visits to windows narrower than that bound. A long plus buffer reduces the response of its centre to any prescribed positive size. The derivative remains macroscopic because a positive density of centres have such buffers.

Lemma 6.1 (Arbitrarily small buffered changes)**.** *For every $\varepsilon > 0$ there is $L < \infty$ such that, for all sufficiently large n, a centre in the middle half of the north collar with L plus bits on each side satisfies*

$$0 \le \operatorname{osc}_i F_n < \varepsilon$$

uniformly over all the other exterior spins.

Proof. Compute the response with the centre itself plus. Directly inside its $2L+1$ plus wall bits take the half-box

$$V_L = \{(x+j, n-j') : -L \le j \le L,\ 0 \le j' \le L\}.$$

For fixed L and large n this box avoids B_{k_n}. Condition either majority-conditioned law on the spins outside V_L; since the majority constraint concerns only spins outside V_L, the conditional law on V_L is the ordinary Ising law with the given exterior. Ferromagnetic monotonicity brackets it by the two laws on V_L with the same plus north wall and with the other three exterior sides all plus or all minus. Let e_L be the difference of their plus-probabilities at $y = (x, n)$.

As $L \to \infty$, any subsequential limit of either extremal half-box law satisfies the half-plane specification with the uniform plus wall. For every finite set this follows by passing its finite-range DLR identity to the limit; the other three sides recede to infinity. Uniform-wall uniqueness, with $h = 1$ as checked in Section 2, makes all these limits identical. Compactness therefore implies $e_L \to 0$. The bracketing is independent of the conditioned surroundings, the majority constraint and distant exterior bits, proving the required uniformity.

With the centre plus, changing it to minus tilts by $e^{-2K\sigma_y}$. The absolute derivative of

$$\log[e^{2K} + (e^{-2K} - e^{2K})u]$$

in the one-site plus-probability u is at most $e^{4K} - 1$. The two majority-conditioned marginals at y differ by at most e_L, so $\operatorname{osc}_i F_n \le (e^{4K} - 1)e_L$. Choose L with the last quantity strictly below ε. □

Lemma 6.2 (Drift and variance for buffered centres)**.** *Fix L and use the block law $\mathbb{P}_q^{(L)}$ of Section 4. There are $c_L > 0$ and $C_L < \infty$ such that, uniformly for $|q - 1/2| \le n^{-2/5}$ and all large n,*

$$\frac{d}{dq}\mathbb{E}_q^{(L)} F_n \ge c_L n, \qquad \operatorname{Var}_q^{(L)}(F_n) \le C_L n. \tag{6.1}$$

Also $\mathbb{E}_{1/2}^{(L)} F_n = 0$.

Proof. There are $N_0 = \lfloor (2\lfloor n/2 \rfloor + 1)/b \rfloor$ complete blocks, where $b = 2L+1$, so $N_0 \ge n/(2b)$ eventually. A given centre is active with probability $a_L = 2^{-2L}$, independently of q. The concentration and surface-separation estimates were proved for these independent blocks, including partially intersected blocks, in Lemmas 4.1 and 4.5.

Choose an even depth $D = D(L) \ge 12$ so large that each of the two expected remainders at every centre is eventually below $a_L m/8$, uniformly over the specified q-window. From (4.13) and nonnegativity of the remainders,

$$\mathbb{E}_q^{(L)}[\mathbf{1}_{\{i \text{ active}\}} d_v] \ge a_L m - \mathbb{E}_q^{(L)} R_{n,v}^{+} - \mathbb{E}_q^{(L)} R_{n,v}^{-} \ge \frac{3a_L m}{4}. \tag{6.2}$$

This estimate is unconditional; activity is not imposed on a contour estimate. By (5.1), the mean active response is at least $3a_L m\gamma_D/4$. Conditional on buffers and inactive centres, differentiation in q sums the oscillations over the active centres, whose bits are independent Bernoulli(q). Averaging gives

$$\frac{d}{dq}\mathbb{E}_q^{(L)} F_n \geq N_0 \frac{3a_L m\gamma_D}{4} \geq \frac{3a_L m\gamma_D}{8b} n.$$

Thus $c_L = 3a_L m\gamma_D/(8b)$ is permissible.

Expose the independent whole blocks and singletons. Changing one group changes F_n by at most $4Kb$ and there are at most $9n$ groups. The variance bound in (4.6) gives $C_L = 36K^2b^2$. At $q = 1/2$ the whole word is fair, so oddness gives mean zero. In particular neither c_L nor C_L depends on a subsequent choice of a tilt constant. □

Proposition 6.3 (A window for the majority odds)**.** *For every $z \in \mathbb{R}$ and $\varepsilon > 0$, there are $c(z,\varepsilon) > 0$ and n_0 such that*

$$\mathbb{P}_{1/2}(|F_n - z| < \varepsilon) \geq c(z,\varepsilon) n^{-1/2} \qquad (n \geq n_0). \tag{6.3}$$

Proof. Choose L by Lemma 6.1 for jumps strictly below ε, and fix the constants in Lemma 6.2. Take

$$A = 1 + \frac{16\sqrt{C_L}}{c_L}, \qquad q_\pm = \frac{1}{2} \pm \frac{A}{\sqrt{n}}.$$

For large n these lie in $(0,1)$ and satisfy $|q_\pm - 1/2| \leq n^{-2/5}$. Integrating (6.1) from $1/2$ puts the endpoint means respectively above $c_L A\sqrt{n}$ and below $-c_L A\sqrt{n}$. Once $|z| \leq c_L A\sqrt{n}/2$, Chebyshev gives

$$\mathbb{P}_{q_+}^{(L)}(F_n \leq z) < 1/16, \qquad \mathbb{P}_{q_-}^{(L)}(F_n \geq z) < 1/16, \tag{6.4}$$

since $4C_L/(c_L^2 A^2) < 1/64$.

Couple the whole family of laws by fixing all buffers and fair inactive bits, and taking independent uniform variables U_i at the active centres: the centre is plus exactly when $U_i \leq q$. Conditional on the buffers, let N be their number and $S(q)$ their plus count. At each endpoint $S(q_\pm)$ has law $\mathrm{Bin}(N, q_\pm)$, where $N \leq 9n$. Its mean differs from $N/2$ by at most $9A\sqrt{n}$ and its variance is at most $9n/4$. With $B_* = 9A + 12$, Chebyshev and a union bound give

$$\mathbb{P}\left(\max_{q \in \{q_-, q_+\}} |S(q) - N/2| > B_*\sqrt{n} \,\middle|\, \text{buffers} \right) \leq 1/32. \tag{6.5}$$

Consequently with probability at least $1/2$ the endpoint values strictly straddle z and both counts lie in this band. All intermediate counts lie in it too, by monotonicity.

Almost surely the finitely many thresholds are distinct. On this event take the first flip in (q_-, q_+) taking F_n from a value at most z to a value above z. It is an active flip, so Lemma 6.1 makes both adjacent values lie strictly in $(z - \varepsilon, z + \varepsilon)$. This remains true if the preceding value equals z. Record such crossing flips only when their pre-flip count lies in (6.5). Let Ξ be their number. The preceding event implies $\mathbb{E}\Xi \geq 1/2$.

We bound $\mathbb{E}\Xi$ by integration of a single uniform threshold. Fix a potential centre i and its buffers. If it is active and $U_i = q$, all the other active centres at that threshold are independent Bernoulli(q). If their plus count is s, the pre-flip count condition implies

$$|s - (N-1)/2| \leq B_*\sqrt{n} + 1.$$

Their probability density relative to independent fair signs is

$$R_q = (2q)^s [2(1-q)]^{N-1-s}.$$

Using $\log(1+x) \le x$ gives

$$\log R_q \le 4(q-\tfrac{1}{2})[s-(N-1)/2] \le 4A(B_*+1) \tag{6.6}$$

throughout the recorded band. This bound is uniform in the buffer word and in N, including small N. All inactive bits and buffers have the same distribution for every q.

The crossing criterion depends only on the other bits, with bit i assigned minus before and plus after the flip. On that criterion *both* assignments have $|F_n - z| < \varepsilon$. Under fair sampling of the other bits its probability, including the activity requirement, is therefore at most the unconditional fair window probability: averaging also over the independent fair bit i preserves the window. There is no conditioning cost for any particular buffer pattern. Bound the tilted criterion by (6.6), integrate q over its interval of length $2A/\sqrt{n}$, and sum over at most $9n$ potential centres. It follows that

$$\frac{1}{2} \le \mathbb{E}\Xi \le 18Ae^{4A(B_*+1)}\sqrt{n}\,\mathbb{P}_{1/2}(|F_n - z| < \varepsilon).$$

Thus $[36Ae^{4A(B_*+1)}]^{-1}$ is a permissible positive constant in (6.3) once the large-size cutoffs above hold. Its size is immaterial for recurrence. □

Proof of Theorem B. Choose $z \in \mathbb{R}$ and $\varepsilon > 0$ so that the compact interval logistic$([z-\varepsilon, z+\varepsilon])$ lies strictly inside (u, v). On H_n, (5.3) places T_n in (u, v) whenever $|F_n - z| < \varepsilon$, for all sufficiently large n. The event H_n^c has probability $o(n^{-1/2})$ by (4.4), so Proposition 6.3 proves the lower bound after reducing its constant. The constants depend only on K, u and v.

Let $\mathcal{U}$ be a weakly open neighbourhood of ν_λ with $0 < \lambda < 1$. Continuity of $\theta \mapsto \nu_\theta$ provides $0 < u < v < 1$ and a positive distance margin such that $\lambda \in (u, v)$ and $\nu_\theta \in \mathcal{U}$ with that margin for $\theta \in [u, v]$. Lemma 2.2 then puts μ_n^η in $\mathcal{U}$ whenever $u < T_n < v$ and n is large. Hence $\mathbb{P}_{1/2}(\mu_n^\eta \in \mathcal{U}) \ge \mathbb{P}_{1/2}(u < T_n < v)$ for large n, and the already proved bound applies, with constants depending on K, $\mathcal{U}$ and λ. □

7 Recurrence, bias selection and empirical laws

7.1 Every mixture recurs

Proof of Theorem A. For every rational pair $0 < u < v < 1$, Theorem B makes $\sum_n \mathbb{P}_{1/2}(u < T_n < v) = \infty$. Write h_n for their hit probabilities. These events depend only on the disjoint effective collars C_n and are independent. For completeness, independence gives

$$\mathbb{P}(\text{no hit from } N \text{ through } M) = \prod_{n=N}^{M}(1-h_n) \le \exp\left(-\sum_{n=N}^{M} h_n\right) \longrightarrow 0.$$

Thus each rational interval is hit infinitely often, almost surely. Intersect these events over their countable basis. On this one probability-one event, every $\lambda \in (0, 1)$ is approached by T_{n_j} along an increasing subsequence, by choosing successively smaller rational intervals and a later hit in each. Lemma 2.2 implies $\mu_{n_j}^\eta \Rightarrow \nu_\lambda$. Every cluster state is DLR and hence belongs to (2.1). Finally the cluster set of a sequence in a compact metric space is closed: a diagonal choice of increasingly late points within shrinking distances of converging cluster points proves this directly. It therefore also contains ν_0 and ν_1. □

7.2 Biased boundaries

Biased part of Theorem C. At $p_n = 1/2 + n^{-2/5}$, Lemmas 5.2 and 5.3 imply, for every fixed $d > 0$ and all large n,

$$\mathbb{P}_{p_n}(X_n < a_+ - d) \le Cn^2 e^{-cn^{2/5}} + e^{-c'n^{1/5}}.$$

Indeed positive F_n of order $n^{3/5}$ has logistic value tending to one, and the transfer error vanishes. This bound is summable.

Fix $p > 1/2$. Eventually $p \geq p_n$. Couple exterior words coordinatewise by shared uniforms; ferromagnetic monotonicity makes X_n increasing. Its lower-tail probability at p is bounded by that at p_n. Summability implies that $X_n \geq a_+ - d$ eventually almost surely: the probability of any failure after size N is at most the tail sum, tending to zero. Intersect over positive rational d. For every boundary, X_n is bounded above by its all-plus boundary value, which tends to a_+ by the definition of μ^+. Hence $X_n \to a_+$ almost surely. Lemma 2.2 gives $\mu_n^\eta \Rightarrow \mu^+$. Spin reversal treats $p < 1/2$, including $p = 0$; $p = 1$ is the constant-plus boundary itself. □

7.3 Fair concentration near the pure phases

Lemma 7.1 (Fair anti-concentration)**.** *For every nonnegative sequence $r_n = o(\sqrt{n})$,*

$$\mathbb{P}_{1/2}(|F_n| \leq r_n) \longrightarrow 0.$$

Proof. Fix $\varepsilon > 0$ and select n of the north positions in Lemma 4.5. Call a position good when $d_v \geq m/2$. A bad position forces $R_{n,v}^+ + R_{n,v}^- > m/2$ by (4.13). For a sufficiently large fixed even $D = D(\varepsilon) \geq 12$, (4.12) makes each expected remainder eventually below $\varepsilon m/16$. Markov's inequality therefore shows that the event G_D of at least $n/2$ good positions satisfies

$$\mathbb{P}_{1/2}(G_D^c) \leq \frac{4}{nm} \sum_v \mathbb{E}(R_{n,v}^+ + R_{n,v}^-) \leq \varepsilon$$

for large n. No independence among these positions is needed.

Clip F_n to $[-r_n - 4K, r_n + 4K]$, obtaining an increasing function f. Whenever $|F_n(\eta)| \leq r_n$, the two values of every bit pair stay within this clipping interval; hence $\operatorname{osc}_i f = \operatorname{osc}_i F_n$ there. On that event intersected with G_D, Lemma 5.1 gives

$$\sum_i \operatorname{osc}_i f \geq (n/2)\gamma_D m/2 = n\gamma_D m/4.$$

Let S count all M_n exterior plus bits. Under the fair product law, conditioning on the other coordinates gives

$$\sum_i \mathbb{E} \operatorname{osc}_i f = 4 \operatorname{Cov}(f, S) \leq 4\sqrt{\operatorname{Var}(f) \operatorname{Var}(S)} \leq 2(r_n + 4K)\sqrt{M_n} \leq 6(r_n + 4K)\sqrt{n}.$$

All oscillations outside the specified event are nonnegative. Therefore

$$\mathbb{P}_{1/2}(|F_n| \leq r_n) \leq \varepsilon + \frac{24(r_n + 4K)}{\gamma_D m \sqrt{n}}.$$

Take $n \to \infty$ for this fixed D, and then $\varepsilon \downarrow 0$. This order of limits is what permits γ_D to be small. □

Lemma 7.2 (Law of a single random state)**.** *Under the fair boundary law,*

$$\mathcal{L}(\mu_n^\eta) \Rightarrow \tfrac{1}{2}\delta_{\mu^+} + \tfrac{1}{2}\delta_{\mu^-}.$$

Proof. Oddness of F_n and Lemma 7.1, with constant sequences r_n, put asymptotic probability $1/2$ on each of its two diverging signs. Lemma 5.3 and $\mathbb{P}(H_n^c) \to 0$ imply $\mathcal{L}(X_n) \Rightarrow \frac{1}{2}\delta_{a_+} + \frac{1}{2}\delta_{a_-}$. Uniform calibration, Lemma 2.2, transfers this to the state space: $\pi(T_n)$ has asymptotic law $\frac{1}{2}\delta_1 + \frac{1}{2}\delta_0$, and the distance of μ_n^η from $\nu_{\pi(T_n)}$ vanishes uniformly. □

Fair part of Theorem C. We give the bounded independent strong law needed for the empirical assertion. If (Y_n) are independent, centred, and bounded by a common constant, then $\mathrm{Var}(\sum_{n\le N} Y_n) \le CN$. At $N = k^2$, Chebyshev bounds the probability of $|k^{-2}\sum_{n\le k^2} Y_n| > d$ by $C/(d^2k^2)$. The summable-tail argument gives convergence to zero along squares, on intersecting over positive rational d. Between k^2 and $(k+1)^2$, boundedness changes the normalized sums by $O(1/k)$. Thus $N^{-1}\sum_{n\le N} Y_n \to 0$ almost surely.

Apply this to a rational polynomial P of finitely many fixed cylinder probabilities, evaluated at μ_n^η, and subtract its expectation. Once its cylinder supports lie in B_n, the evaluation depends only on C_n, so these bounded variables are independent after finitely many terms. By Lemma 7.2, their expected empirical mean tends to $[P(\mu^+) + P(\mu^-)]/2$. Intersect the resulting probability-one events over the countable algebra of rational cylinder polynomials. This algebra determines weak convergence of laws on $\mathcal{P}(\Omega)$: it contains the constants, separates states, and its real uniform closure is all continuous functions by the Stone–Weierstrass theorem. Equivalently, uniform continuity first reduces to finitely many cylinder coordinates and polynomial approximation on their compact cube finishes the reduction. We obtain the empirical convergence asserted in Theorem C.

For the mean of X_n, spin reversal gives $\mathbb{E}_{1/2}X_n = 1/2$ exactly at every size. The variables X_n are independent and bounded, so the same strong law directly gives their ordinary average $1/2$. This last calculation in fact needs no low-temperature condition. For the spin expectation rather than the plus-probability, the average is zero. □

8 The sharp mixed-state probability

Throughout this section the boundary spins are independent fair signs. For the upper bound assume $K > \log 2$, equivalently $4e^{-2K} < 1$; in particular the result holds for $K \ge 1$. We use the physical collar and the deterministic open-arc decomposition of Section 3. All diameters and distances are in the sup norm. Write

$$p_n = \mu_n^\eta(\mathcal{A}^+) = \mathrm{logistic}(F_n), \qquad M = 2\lfloor n/2 \rfloor + 1,$$

so $M \ge n$ is the number of selected north columns. Fix $c_{\mathrm{mac}} > 0$ for the two exponential bounds in Lemma 4.2, and retain $\delta_n = C_{\mathrm{maj}} e^{-c_{\mathrm{maj}} k_n}$ from Lemma 4.3. These are distinct decay constants.

8.1 Joint deletion and extensive coverage

Lemma 8.1 (Joint annealed price for specified arcs)**.** *Fix an ordered list of $k \ge 1$ pairs (a_j, b_j) of north positions with all $2k$ endpoints distinct and a_j to the left of b_j. Let d_j be integers with $b_j^{(1)} - a_j^{(1)} \le d_j$. Let $\mathcal{T}$ be the event that each designated pair is joined by an open arc of diameter at least d_j. There are $d_0 = d_0(K) < \infty$ and $c_{\mathrm{del}} = c_{\mathrm{del}}(K) > 0$ such that, if every $d_j \ge d_0$ and $L = \sum_j d_j$, then*

$$\mathbb{E}_{1/2}\mu_n^\eta(\mathcal{T}) \le 2e^{-c_{\mathrm{del}} L}. \tag{8.1}$$

Proof. Let I_j be the north collar interval between a_j and b_j, and put $U = I_1 \triangle \cdots \triangle I_k$. Then $|U| \le L$. Boundary transitions add modulo two, so reversing U toggles exactly the designated endpoints, including for overlapping or nested intervals.

Fix a jointly occurring tuple of actual arcs, oriented from their left endpoints. Its trails are edge-disjoint even if they meet at vertices. Delete their edges from the disagreement set E. This reduces $|E|$ by their total length and toggles parity only at their endpoints. By Lemma 3.1, the new edge set represents a configuration with collar η^U. For this fixed tuple

the map is injective: union with the deleted edges recovers E. The image may be a proper subset of the edge sets at the reversed collar, which suffices for the upper bound

$$\sum_{\sigma:\,\text{tuple occurs}} W_n^\eta(\sigma) \le w^{\sum_j |\gamma_j|} Z_n(\eta^U), \qquad w = e^{-2K}.$$

There are at most 4^h walks of length h from a specified start, and an arc of diameter at least d_j has length at least d_j. Overcounting tuples by all such walks, set

$$\beta_{\mathrm{w}} = 2K - \log 4 > 0, \qquad A_{\mathrm{w}} = (1 - e^{-\beta_{\mathrm{w}}})^{-1}, \qquad D_U = \log Z_n(\eta^U) - \log Z_n(\eta).$$

The unrestricted quenched probability is bounded by

$$\mu_n^\eta(\mathcal{T}) \le \min\{1, e^{D_U} A_{\mathrm{w}}^k e^{-\beta_{\mathrm{w}} L}\}. \tag{8.2}$$

Choose $d_0 \ge \max\{5, 4\log A_{\mathrm{w}}/\beta_{\mathrm{w}}\}$. Then $k \log A_{\mathrm{w}} \le \beta_{\mathrm{w}} L/4$, and on $D_U \le \beta_{\mathrm{w}} L/4$ the last bound is at most $e^{-\beta_{\mathrm{w}} L/2}$.

For this deterministic U, condition only on its complementary collar bits. The remaining signs are independent fair signs, and their simultaneous reversal negates D_U, so its conditional mean is zero. A bit of U changes each log partition function by at most $2K$, and hence changes D_U by at most $4K$. The martingale calculation (4.6) therefore gives

$$\mathbb{P}_{1/2}(D_U > t \mid \eta_{U^c}) \le \exp\{-t^2/(8K^2|U|)\} \qquad (t > 0).$$

When U is empty, $D_U = 0$. With $t = \beta_{\mathrm{w}} L/4$ and $|U| \le L$, use the probability cap one on the exceptional event in (8.2). Equation (8.1) follows with

$$c_{\mathrm{del}} = \min\{\beta_{\mathrm{w}}/2, \beta_{\mathrm{w}}^2/(128K^2)\}.$$

This bound is applied separately to each deterministic endpoint specification before summing over specifications. No contour occurrence has been imposed on the disorder law. □

Lemma 8.2 (Extensive coverage has exponential cost)**.** *For a fixed even $D \ge 12$, let Q_x be the lattice segment from $(x, 0)$ to $v_x = (x, n - D)$, for $|x| \le \lfloor n/2 \rfloor$. Call an open arc short if its diameter is less than $n/10$, and let B count the selected segments having a short-arc vertex within distance 2. For every fixed $\theta > 0$ there are a fixed even $D = D(K, \theta)$ and $h > 0$ such that, for all sufficiently large n,*

$$\mathbb{E}_{1/2}\mu_n^\eta(B > \theta M) \le e^{-hn}. \tag{8.3}$$

Proof. A short arc with a vertex within distance 2 of Q_x has both endpoints on the north side for large n. Indeed each endpoint is within $d + 2$ of the segment, where $d < n/10$ is the diameter, and the other sides are at distance at least $n/2$. North positions have ordinate $n + 1/2$ while the near vertex has ordinate at most $n - D + 2$, so $d \ge D - 2$. Each affected abscissa is within $d + 2$ of the left endpoint's abscissa. Thus an arc affects at most $2d + 5 \le 3d$ selected segments. These facts also hold when an arc touches the perimeter away from its endpoints.

If $B > \theta M$, record every distinct contributing arc once by its two endpoints and exact integer diameter. Endpoints of different arcs are distinct. The sum L of recorded diameters satisfies $3L \ge B > \theta M \ge \theta n$. This encoding implies the event in Lemma 8.1 with these diameters as lower bounds. Choose $D - 2 \ge d_0$.

Order the encoding by left endpoint. At each north position one can choose no arc, or an integer $d \ge D - 2$ and one of at most d right endpoints at horizontal separation at most d.

There are $2n+2 \le 4n$ north positions. Apply Lemma 8.1 to valid encodings, and only then overcount by dropping compatibility conditions. With

$$s_D = \sum_{d \ge D-2} d e^{-c_{\mathrm{del}} d/2},$$

half of the exponential price pays for $L > \theta n/3$ and half pays for this product count, giving

$$\mathbb{E}_{1/2}\mu_n^\eta(B > \theta M) \le 2e^{-c_{\mathrm{del}}\theta n/6}(1+s_D)^{4n} \le 2\exp\{-c_{\mathrm{del}}\theta n/6 + 4n s_D\}.$$

Since $s_D \downarrow 0$, choose a fixed even D with $4s_D < c_{\mathrm{del}}\theta/12$ and increase the size cutoff to absorb the factor 2. This proves (8.3). An arc affecting many segments pays for all of them through its diameter; independence of arc occurrences is not needed. □

8.2 Responses, odds and local transfer

Lemma 8.3 (A linear number of nonvanishing responses)**.** *There are a fixed even D, $a > 0$ and collar events J_n such that*

$$\mathbb{P}_{1/2}(J_n^c) \le Cn^2 e^{-cn^{2/5}} + e^{-hn/2}, \qquad \#\{i = (x, n+1) : |x| \le \lfloor n/2 \rfloor,\ \mathrm{osc}_i F_n \ge a\} \ge M/2 \tag{8.4}$$

on J_n, for all large n.

Proof. Use $\theta = m/32$ in Lemma 8.2. Set $q_\eta = \mu_n^\eta(B > \theta M)$ and $J_n = H_n \cap \{q_\eta \le e^{-hn/2}\}$. Markov's inequality and (4.4) give the probability bound. On J_n, Lemma 4.1(iii) gives, for either sign,

$$\mu_n^{\eta,s}(B > \theta M) \le 2e^{n^{7/10} - hn/2}, \qquad M^{-1}\mu_n^{\eta,s}(B) \le \theta + 2e^{n^{7/10} - hn/2}.$$

The exponential unrestricted tail is obtained before division by the majority probability.

Let $p_{s,x}$ be the $\mu_n^{\eta,s}$ probability that a short arc comes within distance 2 of Q_x. Its mean over the M columns is exactly $M^{-1}\mu_n^{\eta,s}(B)$. On H_n, the probability under $\mu_n^{\eta,s}$ of a long arc or a short-only family with $s_\Gamma = -s$ is at most

$$b_n = e^{-c_{\mathrm{mac}} n} + 2e^{n^{7/10}}\delta_n = o(1).$$

For the wrong-sign families, condition on Γ and use (4.10) to bound the unrestricted probability of their intersection with $\mathcal{A}^s$ by δ_n; division then uses Lemma 4.1(iii).

For a short-only family with no arc near Q_x and with $s_\Gamma = s$, Lemma 3.6(iv) places the entire connected segment in the central component, since $(x, 0) \in B_{\lfloor 2n/3 \rfloor}$. In a finite plus-boundary domain the one-site plus-probability is at least a_+; in a minus-boundary domain it is at most a_-. This follows by integrating the infinite pure state's finite-domain specification over its exterior, and comparing each exterior with the constant exterior by Lemma 2.3(i). Conditioning further on $\mathcal{A}^s$ changes a $[0,1]$-valued expectation by at most δ_n, by its decomposition over the two majority events and (4.10). Averaging these one-sided estimates and allowing error one for an exceptional family yields

$$d_{v_x} \ge m - 2\delta_n - 2b_n - p_{+,x} - p_{-,x}. \tag{8.5}$$

The mean of the nonnegative loss on the right is at most $2\theta + o(1) < m/4$ uniformly on J_n. Thus at least $M/2$ gaps satisfy $d_{v_x} \ge m/2$. The depth-D heat-bath paths from v_x to the north row avoid B_{k_n} for large n. Lemma 5.1 gives (8.4) with $a = \gamma_D m/2$.

We will also use the same proof with a smaller prescribed $\theta > 0$. On its corresponding event J_n, (8.5) bounds the fraction of columns with $d_{v_x} < m/2$ by $(2/m)(2\theta + o(1))$. □

Lemma 8.4 (Fixed-width odds windows)**.** *For every fixed $r \geq 0$,*

$$\mathbb{P}_{1/2}(|F_n| \leq r) \leq \mathbb{P}_{1/2}(J_n^c) + \frac{4(r+4K)\sqrt{8n+4}}{aM} = O(n^{-1/2}). \tag{8.6}$$

Every fixed bounded odds interval with nonempty interior has probability of exact order $n^{-1/2}$.

Proof. Clip F_n to $[-r-4K, r+4K]$, obtaining an increasing f. Lemma 5.1 gives $0 \leq \operatorname{osc}_i F_n \leq 4K$ everywhere. At a word with $|F_n| \leq r$, both values of every coordinate pair lie in the clipping interval, so $\operatorname{osc}_i f = \operatorname{osc}_i F_n$ there. If S counts all exterior plus bits, conditioning on the other bits under the original fair product law gives

$$\operatorname{Cov}(f, \mathbf{1}_{\{\eta_i=+1\}}) = \tfrac{1}{4}\mathbb{E}_{1/2} \operatorname{osc}_i f.$$

Since $\operatorname{Var}(S) = (8n+4)/4$ and $\operatorname{Var}(f) \leq (r+4K)^2$, nonnegativity of all oscillations and Lemma 8.3 imply

$$\frac{aM}{2}\mathbb{P}_{1/2}(J_n \cap \{|F_n| \leq r\}) \leq \sum_i \mathbb{E}_{1/2} \operatorname{osc}_i f = 4 \operatorname{Cov}(f, S) \leq 2(r+4K)\sqrt{8n+4}.$$

This is the covariance calculation of Lemma 7.1, now with an exceptional probability of smaller order than $n^{-1/2}$. The lower bound on any open subinterval is Proposition 6.3. □

Lemma 8.5 (Transfer in the local seminorm)**.** *For every fixed finite X, on H_n and for all sufficiently large n,*

$$\|\mu_n^\eta - \nu_{p_n}\|_X \leq e_n(X), \qquad e_n(X) \leq C_X e^{-c_X n^{4/5}}. \tag{8.7}$$

This transfer estimate itself holds at every $K > K_c$.

Proof. Assign a short-only arc family the sign s_Γ of its central component; assign sign plus to every long-arc family. Let u be the unrestricted thermal probability of assigned sign plus. The probability that this sign disagrees with the central majority is at most $\delta_n + \mu_n^\eta(\mathcal{E}_n)$, so the same bounds $|p_n - u|$.

For a short-only family, the fixed set X eventually lies at distance at least $n/2$ from the complement of the central component. Lemma 3.6(iii) and (2.2), uniformly over all supremum-norm-one tests supported on X, bound the conditional seminorm distance from μ^{s_Γ} by $|X|e^{-c_K n/2}$. This is precisely the arbitrary-domain estimate [4, (A.8)], with spin reversal for minus boundary. For a long-arc family the distance is at most 2. Average over families and use $\|\nu_{p_n} - \nu_u\|_X \leq 2|p_n - u|$ to obtain

$$\|\mu_n^\eta - \nu_{p_n}\|_X \leq |X|e^{-c_K n/2} + 4e^{-c_{\mathrm{mac}} n} + 2\delta_n.$$

This proves (8.7); only the inputs valid for $K > K_c$ were used. □

Proof of Theorem D for odd sides. Put $\Delta_X = \|\mu^+ - \mu^-\|_X \leq 2$. Homogeneity of the seminorm gives

$$\min_{s=\pm} \|\nu_{p_n} - \mu^s\|_X = \Delta_X \min(p_n, 1-p_n) = \frac{\Delta_X}{1+e^{|F_n|}}. \tag{8.8}$$

Choose a fixed r with $2/(1+e^r) < \varepsilon/2$ and take $e_n(X) < \varepsilon/2$. On H_n, the departure event in Theorem D then implies $|F_n| \leq r$. Equations (8.6) and (4.4) prove the upper bound; increasing its constant includes the finitely many smaller sizes.

If $0 < \varepsilon < \Delta_X/2$, choose a compact weight interval of positive length around $1/2$ on which $\Delta_X \min(t, 1-t)$ exceeds ε by a fixed positive margin. Its logit image contains a nonempty

bounded open interval. Proposition 6.3, minus $\mathbb{P}_{1/2}(H_n^c) = o(n^{-1/2})$, and Lemma 8.5 give the matching lower bound. This lower-bound argument uses only $K > K_c$.

The calibrated weight satisfies $|T_n - p_n| \le e_n(\{0\})/m$ on H_n. For $0 < u < v < 1$, a calibrated hit forces bounded odds on H_n and contains a smaller nonempty open odds interval for large n. Thus its probability has exact order $n^{-1/2}$ for $K > \log 2$. If $\varepsilon > \Delta_X/2$, (8.8) and (8.7) exclude the departure event on H_n eventually; its probability is at most $Cn^2e^{-cn^{2/5}}$. At $\varepsilon = \Delta_X/2$ there is no open interval with a strict distance margin, and only the upper bound is claimed. □

8.3 Every integer side length

Corollary 8.6 (Odd and even sides)**.** *Let $V_{2n+1} = B_n$ and $V_{2n} = [-n+1, n]^2 \cap \mathbb{Z}^2$, with fair independent spins on their physical collars and bulk-strength boundary bonds. The conclusions of Theorem D hold along every integer side length N. The lower bounds for interior odds and calibrated-weight windows, and for $0 < \varepsilon < \Delta_X/2$, hold for both parities at every $K > \log 2$, and for odd side lengths at every $K > K_c$.*

Proof. For an even side retain B_{k_n} as the odd central majority box, the north columns and Q_x above, and use the $8n$ physical collar bits. The two majority events still partition the configurations. Lemmas 3.1 and 3.6 apply to this rectangle without change. So does Lemma 8.1. In Lemma 8.2 the other sides are at distance at least $n/2 - 1$ from the segments, which still exceeds $n/10 + 2$ eventually; there are at most $2n + 2$ north positions. Thus the same extensive-coverage estimate holds.

Normalization and the central component. Define H_n by the same cutoff $n^{7/10}$ for all cyclic-interval comparisons and each of the three partition sums. The fair conditional concentration calculation in Lemma 8.1 applies also to a majority-constrained sum, since the constraint is independent of the collar. Each interval has at most $8n$ bits. A union bound over at most $3(8n+1)^2$ choices gives $\mathbb{P}_{1/2}(H_n^c) \le Cn^2e^{-cn^{2/5}}$. Full-collar reversal exchanges the constrained sums, giving $|F_n| \le n^{7/10}$ and majority probabilities at least $\frac{1}{2}e^{-n^{7/10}}$ on H_n.

For a specified pair of perimeter positions the deletion injection works on any side, using either collar interval between them. Summing walks of length at least $\lceil n/10 \rceil$ gives at most $A_{\mathrm{w}}e^{-\beta_{\mathrm{w}}n/10}$. On H_n the probability of a long arc for this pair is at most $A_{\mathrm{w}}e^{n^{7/10}-\beta_{\mathrm{w}}n/10}$. Sum over at most $(8n)^2$ pairs and divide, if needed, by the majority probability. The unrestricted and both constrained long-arc probabilities are therefore at most $e^{-c_{\mathrm{mac}}n}$ after choosing a suitable $c_{\mathrm{mac}} > 0$ and increasing the cutoff. With no long arc every arc vertex is within $n/10$ of the dual perimeter. Hence $B_{\lfloor 2n/3 \rfloor}$ is at distance greater than 2 from all open arcs and lies in one free component. Lemma 4.3 applies to every constant-boundary domain containing B_{k_n}. These facts give the central-component inputs used above for even sides.

The response and upper bound. Association, Lemma 2.3(ii), orders the two majority-conditioned laws as in (4.3). Off B_{k_n} their one-site conditionals are unchanged. Propagate the gap from v_x to the top row along D edges by Lemma 2.3(iii), giving a top-row gap at least $\kappa_K^D d_{v_x}$. Flipping a minus collar bit tilts either constrained law by $e^{2K\sigma_y}$. The logarithm of its expectation has derivative at least $1 - e^{-4K}$ in the plus-probability at y. Starting with a plus bit uses minus the logarithm of the reverse tilt and gives the same bound. The log change of either constrained sum is at most $2K$, giving upper response $4K$. Ordering and the tilt also give nonnegative responses at every other collar coordinate. Thus Lemma 5.1 holds with no parity assumption. Lemmas 8.3–8.5 and the odd-side completion now apply, with $8n$ in the covariance calculation. This proves all the even-side upper bounds and the large-ε stretched-exponential estimate.

Buffered centres and drift for the lower bound. Fix $z \in \mathbb{R}$ and $\xi > 0$. Divide the selected north bits into complete blocks of odd length $\ell = 2L + 1$, leaving any incomplete block as singletons. A centre is active if all its $2L$ buffer bits are plus. Choose $L = L(K, \xi)$ so that every active jump is less than ξ, uniformly over all other bits and large n. Here is why the choice is independent of volume parity. With the centre plus, condition either majority law outside the half-box of width $2L + 1$ and depth $L + 1$ below the block. This box avoids the central majority. Monotonicity brackets both top-site marginals between the laws with the common plus north wall and all-plus or all-minus exterior on the other three sides. As L grows, every subsequential limit is a Gibbs state for the half-plane with wall coupling $h = 1$. Uniqueness at that wall, as established in Section 2 using [5, (3.17), (3.25)], makes the difference of the two marginals tend to zero. The reverse log tilt has absolute derivative at most $e^{4K} - 1$, so the jump is less than ξ for a fixed sufficiently large L.

Let $\pi = 2^{-2L}$ and let R be the number of active centres. At fairness the activations are independent, and $\mathbb{E}R \geq \pi M/(2\ell)$ eventually. The bounded-differences estimate (4.6) on the block indicators gives

$$\mathbb{P}_{1/2}(R < \pi M/(4\ell)) \leq e^{-c_L n}.$$

Run the proof of Lemma 8.3 with θ so small that $4\theta/m < \pi/(16\ell)$. Its final observation shows that fewer than $\pi M/(8\ell)$ columns have $d_{v_x} < m/2$ on the corresponding J_n, eventually. Consequently at least $\pi M/(8\ell)$ active centres have responses at least $a = \gamma_D m/2$ except on a fair event of probability $o(1)$. Here D and a are fixed after L and before choosing a tilt constant.

Expose all buffers, singletons and inactive centres, retaining their fair law. Conditional on these data give the active centres independent plus-probability q, and denote expectation by $\mathbb{E}_q$. For any fixed A_0, this high-probability response statement remains true uniformly on $|q - 1/2| \leq A_0/\sqrt{n}$. Indeed the conditional second moment, under fairness, of the likelihood ratio is

$$(1 + 4(q - 1/2)^2)^R \leq e^{8A_0^2},$$

since $R \leq M \leq 2n$ for large n; Cauchy–Schwarz bounds the tilted failure probability by $e^{4A_0^2}$ times the square root of its fair probability. Differentiate the finite conditional product polynomial and average. The derivative sums the expected active oscillations, and hence, eventually and uniformly in that fixed window,

$$\frac{d}{dq}\mathbb{E}_q F_n \geq c'_L n, \qquad c'_L = \frac{a\pi}{16\ell} > 0. \tag{8.9}$$

The coefficient is independent of A_0. At fairness spin reversal gives $\mathbb{E}_{1/2} F_n = 0$. Complete blocks and singletons are independent groups under each q; replacing a group changes F_n by at most $4K\ell$, and there are at most $8n$ groups. The martingale variance bound in (4.6) therefore gives $\mathrm{Var}_q(F_n) \leq C_L n$ with $C_L = 32K^2\ell^2$.

A uniform permutation and an independent binomial prefix. Choose a fixed A_0 sufficiently large that $4C_L/(c'_L A_0)^2 < 1/16$. For $q_\pm = 1/2 \pm A_0/\sqrt{n}$, integration of (8.9) and Chebyshev give, for all large n,

$$\mathbb{P}_{q_-}(F_n \geq z) < 1/16, \qquad \mathbb{P}_{q_+}(F_n \leq z) < 1/16.$$

Couple these laws with one independent uniform threshold for each active centre, sharing the exposed data. Conditional on the data, the endpoint plus counts are binomial with parameters $(R, q_\pm)$. Their means are within $2A_0\sqrt{n}$ of $R/2$ and their variances are at most $n/2$. Chebyshev permits a fixed B_0 such that both counts lie within $B_0\sqrt{n}$ of $R/2$, except on

an event of probability less than 1/8. Together with the activation bound, with probability at least 1/2 the counts lie in this band, $R \geq \pi M/(4\ell)$, and the endpoint odds strictly straddle z.

In increasing threshold order the active flips add one plus at a time and increase the odds by less than ξ. The straddling event therefore yields a prefix of length j satisfying

$$|j - R/2| \leq B_0\sqrt{n}, \qquad |F_n - z| < \xi. \tag{8.10}$$

Conditional on the exposed data, the threshold order is a uniform permutation. Although straddling also uses the ordered threshold values, its implication that a qualifying prefix exists depends only on this permutation and the data. Thus the marginal probability of existence of such a prefix, with $R \geq \pi M/(4\ell)$, is at least 1/2.

Independently sample $J \sim \mathrm{Bin}(R, 1/2)$ and assign plus exactly to the first J members of the permutation. This produces the fair conditional word: conditional on $J = j$, each size-j subset is equally likely, and its unconditional probability is $\binom{R}{j}2^{-R}/\binom{R}{j} = 2^{-R}$.

For $\pi n/(4\ell) \leq R \leq 2n$, every admissible j in the band (8.10) has binomial mass at least $c_{B_0,L}/\sqrt{n}$. To verify this elementary bound, Chebyshev puts at least 3/4 of the mass within $\sqrt{R}$ of $R/2$. Its largest, central mass is therefore at least $(3/4)/(2\sqrt{R}+2)$. Adjacent binomial masses have ratio $(R-j)/(j+1)$. Moving away from a central mode, symmetry and the linear lower bound on R show that each outward ratio in the indicated band is at least $1 - C_{B_0,L}/\sqrt{n}$ for large n. There are at most $B_0\sqrt{n}+2$ steps, so multiplication loses only a fixed positive factor. This proves the mass bound.

For each permutation admitting (8.10), choose its smallest qualifying prefix length. The independent J hits that length with probability at least $c_{B_0,L}/\sqrt{n}$. Averaging over the permutations and the original exposed data yields

$$\mathbb{P}_{1/2}(|F_n - z| < \xi) \geq c' n^{-1/2}.$$

The local transfer and the strict-margin argument above give the even-side calibrated-window and separation lower bounds. □

9 Discussion

A fixed neighbourhood of an interior mixture is visited infinitely often by Theorem A, while its empirical frequency is zero by Theorem C whenever its closure avoids the pure phases. This is the null recurrence of mixed states anticipated in [1], [2, Section 12] and [3, Remark 3.2]; its zero-frequency half is [2, Theorem 3.4], which those authors already describe as the null-recurrent character of the mixed states. The cluster set is a closed segment of Gibbs measures, and the empirical metastate is the equal two-atom measure on its endpoints. Its barycentre is the symmetric Gibbs state. These three objects therefore encode different information.

Theorem D gives the $O(N^{-1/2})$ upper bound conjectured in [2, Remark 3.6] for every $K > \log 2$, including $K \geq 1$, with fair independent bulk-strength boundary spins. Together with the lower bound it gives exact order $N^{-1/2}$ for every fixed finite X and $0 < \varepsilon < \frac{1}{2}\|\mu^+ - \mu^-\|_X$, and for every fixed interior calibrated-weight interval. The full-sequence recurrence conjecture and this quantitative conjecture are therefore both resolved, the latter in that temperature range and with the necessary separation threshold. A literature search found no earlier proof of the upper bound. The result does not assert a Gaussian local limit or an asymptotic equivalent.

Read as a two-sided order statement for every $\varepsilon > 0$, Remark 3.6 of [2] does not hold above $\frac{1}{2}\|\mu^+ - \mu^-\|_X$. This qualification follows from the geometry of the mixture segment: its midpoint maximizes the distance to the nearer endpoint. Lemma 8.5 makes larger departures stretched-exponentially unlikely. The sharp statement is the upper bound for every positive separation and the matching lower bound below that midpoint value. At the threshold itself we claim only the upper bound.

The recurrence, lower-bound and empirical-law proofs work at every $K > K_c$ because their inputs are non-perturbative. Dual open arcs of large diameter have exponentially small unnormalized weight, since the dual model is subcritical. Fair boundary disorder makes every macroscopic interval reversal cost a sublinear log partition weight, so with high probability the bulk lies in one constant-sign component of the open-arc decomposition. A central majority identifies its sign with an error exponentially small in the radius k_n of the central square. For arcs near the north side one can condition on the complement of a strip, which keeps the majority event fixed and makes the same comparison available for the constrained partition functions. The heat-bath response and wall uniqueness then give the fixed-window lower probability. The same estimates show that typical fair boundaries suppress macroscopic interfaces throughout the ordered phase: by Lemma 4.2 and the summable bound (4.4), almost surely $\mu_n^\eta(\mathcal{E}_n) \le e^{-cn}$ for all large n. At sufficiently low temperature, van Enter, Netočný and Schaap proved that a typical boundary condition suppresses interfaces [2, Theorem 4.1]; see also [3, Remark 3.3].

The proof separates two scales of boundary response. Under a near-fair product bias, a linear number of separated phase responses gives a drift larger than $\sqrt{n}$. Under the buffered law, the same argument survives independent groups of fixed size, even though individual bits are dependent. Uniform-wall uniqueness makes the selected centre responses arbitrarily small. Their pivotal crossings then provide the local lower bound. The buffers and depths may be very large and depend on the target window. The constants are not asserted to be uniform as $K \to \infty$, as $K \downarrow K_c$, or as the window shrinks. For the sharp upper bound, joint deletion charges a family of distinct arcs through its total diameter before summing endpoint specifications. The resulting exponential coverage tail survives division by the majority mass and leaves a positive density of responding bits outside a summable disorder exception. Clipping the odds and taking covariance with the fair boundary plus count then gives the exponent $-1/2$. Corollary 8.6 treats both parities, including the lower bound obtained by a uniform permutation and an independent binomial prefix.

The sharp upper bound for $K_c < K \le \log 2$ remains open; the lower bound holds throughout that interval for odd side lengths. The joint deletion argument used here requires convergence of the walk sum $\sum_h 4^h e^{-2Kh}$, and so does not reach it. A density asymptotic and a limiting prefactor also remain open. Two further questions are the weak-coupling model of van Enter, Medved' and Netočný [1], whose boundary bonds are weaker than the bulk bonds, and three dimensions. In the weak-coupling model the uniform-wall uniqueness of Section 2 and the homogeneity of the dual model are no longer available. In three dimensions, [2, Remark 3.2] expects the cluster set to be the set of all translation-invariant Gibbs measures, and the planar classification and dual-line argument used here do not apply.

10 How this work was done

Claude Opus 5.5, GPT-6.1 Sol, GPT-6.2 Sol and GPT-6.2 Astra assisted with exploration, proof development and mathematical checking. The proofs were written in full. The arguments for Theorems A–C, valid at every $K > K_c$, were found with the assistance of AI models. They were checked by three non-author audits during their development and then by two independent refutation-first checks, by models from more than one developer in all, which compared the imported statements with their sources and hypotheses. The sharp upper bound of Section 8 was checked by a non-author audit and then by a separate refutation-first check from a different developer, which also compared its imports with their sources. Enumerations over all spin configurations of boxes with up to nine sites, for all boundary words or samples of them and at several temperatures, confirmed the identities of Lemmas 3.1 and 3.2, the inequality (3.5) for single curves, the inequality (3.7), and the conditional law of Lemma 3.6.

An integer enumeration of all edge sets of one dual rectangle confirmed the local contour rules and the anchoring of constrained sites to the collar. A similar enumeration on a box of twenty sites confirmed the response bounds of Lemma 5.1 and the derivative bound in the proof of Lemma 6.1. The weights in these enumerations were evaluated in floating point, and no proof depends on them. The authors take responsibility for every claim.

A Scales and constants

The temperature enters the proofs through the published inputs of Section 2 and through two strict inequalities,

$$K^{*} = \operatorname{artanh} e^{-2K} < K_c \qquad \text{and} \qquad m > 0,$$

both valid at every $K > K_c$. The first gives the decay rate $c_\tau > 0$ of (2.4) and hence $c_0 > 0$ in Lemma 3.5. The second gives the phase gap in (4.13) and the threshold $m/2$ in Lemma 4.3. At finite K the heat-bath increment κ_K and the response factor $1 - e^{-4K}$ are positive. No numerical threshold above K_c enters the recurrence, lower-bound or empirical-law proofs. For the sharp upper bound in Section 8, the walk sum requires $K > \log 2$, equivalently $\beta_{\mathrm{w}} = 2K - \log 4 > 0$. Its deletion rate c_{del}, the macroscopic-arc rate c_{mac} and the majority rate c_{maj} have different names and roles.

Four scales appear. Near-fair biases satisfy $|p - 1/2| \le n^{-2/5}$, so the mean odds in Lemma 5.2 grow like $n^{3/5}$. The interval cutoff of Lemma 4.1 is $n^{7/10}$. The majority square has radius $k_n = \lfloor n^{4/5} \rfloor$. The cutoff for macroscopic arcs is $n/10$. The proofs use the comparisons

$$\tfrac{2}{5} < \tfrac{1}{2}, \qquad \tfrac{3}{5} < \tfrac{7}{10}, \qquad \tfrac{7}{10} < \tfrac{4}{5} < 1.$$

The first places the tilted block laws $q_{\pm} = 1/2 \pm A/\sqrt{n}$ of Proposition 6.3 inside the permitted bias range. The second makes the conditional mean $36Kn^{3/5}$ of an interval comparison negligible against the cutoff $n^{7/10}$, which leaves the disorder exponent $n^{2/5}/(288K^2b^2)$ in (4.4). In the third, the majority normalization $e^{n^{7/10}}$ of Lemma 4.1(iii) is beaten by the wrong-majority error $\delta_n = C_{\mathrm{maj}} e^{-c_{\mathrm{maj}} k_n}$ and by the macroscopic-arc bound $e^{-c_0 n/10}$, and $k_n < n/2$ keeps the majority square inside the central component and away from the north strip.

Constants may depend on the fixed temperature and on the buffer width $b = 2L + 1$. For the window estimate the choices are made in the following order. First fix K and the target window (z, ε). Then choose L by Lemma 6.1, with activity probability $a_L = 2^{-2L}$. Next choose an even depth $D \ge 12$ so that each expected remainder is eventually below $a_L m/8$, giving active expected separation at least $3a_L m/4$. Then choose A and B_* from Proposition 6.3, and finally the lower size cutoff. In the fair anti-concentration argument a good position has gap at least $m/2$, and the depth is chosen after the desired exceptional probability.

For the sharp upper bound, fix $\theta = m/32$, choose the fixed depth D by Lemma 8.2, and then set $a = \gamma_D m/2$. The exceptional probability in Lemma 8.3 is smaller than $n^{-1/2}$, so no depth limit is needed after the covariance estimate. For the even-side lower bound in Corollary 8.6, choose the buffer width L for the target odds window, then a smaller θ and the corresponding D, then the drift and variance constants, and only then the fixed tilt constant A_0 and count-band constant B_0. The alternative drift proof in that corollary gives the lower bound at every $K > K_c$.

References

[1] A. C. D. van Enter, I. Medved' and K. Netočný, Chaotic size dependence in the Ising model with random boundary conditions, *Markov Process. Related Fields* **8** (2002), 479–508. arXiv:math-ph/0108011.

[2] A. C. D. van Enter, K. Netočný and H. G. Schaap, On the Ising model with random boundary condition, *J. Stat. Phys.* **118** (2005), 997–1056. doi:10.1007/s10955-004-2138-2; arXiv:math-ph/0408024.

[3] A. C. D. van Enter, K. Netočný and H. G. Schaap, Incoherent boundary conditions and metastates, *IMS Lecture Notes–Monograph Series* **48** (2006), 144–153. doi:10.1214/074921706000000176; arXiv:math-ph/0502031.

[4] L. Coquille and Y. Velenik, A finite-volume version of Aizenman–Higuchi theorem for the 2d Ising model, *Probab. Theory Related Fields* **153** (2012), 25–44. doi:10.1007/s00440-011-0339-6; arXiv:1003.6034.

[5] C.-E. Pfister and Y. Velenik, Mathematical theory of the wetting phenomenon in the 2D Ising model, *Helv. Phys. Acta* **69** (1996), 949–973. arXiv:cond-mat/9701161. Equation numbers as in the arXiv version.

[6] C.-E. Pfister and Y. Velenik, Large deviations and continuum limit in the 2D Ising model, *Probab. Theory Related Fields* **109** (1997), 435–506. doi:10.1007/s004400050139.

[7] C.-E. Pfister and Y. Velenik, Interface, surface tension and reentrant pinning transition in the 2D Ising model, *Commun. Math. Phys.* **204** (1999), 269–312; numbering as in this published version, which differs from that of the preprint arXiv:cond-mat/9704018. doi:10.1007/s002200050646.

[8] R. B. Griffiths, Correlations in Ising ferromagnets. I, *J. Math. Phys.* **8** (1967), 478–483. doi:10.1063/1.1705219.

[9] D. G. Kelly and S. Sherman, General Griffiths' inequalities on correlations in Ising ferromagnets, *J. Math. Phys.* **9** (1968), 466–484. doi:10.1063/1.1664600.

[10] J. Ginibre, General formulation of Griffiths' inequalities, *Commun. Math. Phys.* **16** (1970), 310–328. doi:10.1007/BF01646537.

[11] D. Ioffe, Exact large deviation bounds up to T_c for the Ising model in two dimensions, *Probab. Theory Related Fields* **102** (1995), 313–330. doi:10.1007/BF01192464.

[12] R. H. Schonmann, Second order large deviation estimates for ferromagnetic systems in the phase coexistence region, *Commun. Math. Phys.* **112** (1987), 409–422. doi:10.1007/BF01218484.

[13] R. Dobrushin, R. Kotecký and S. Shlosman, *Wulff construction: a global shape from local interaction*, Transl. Math. Monographs **104**, Amer. Math. Soc., Providence, RI, 1992. doi:10.1090/mmono/104.

[14] C.-E. Pfister, Large deviations and phase separation in the two-dimensional Ising model, *Helv. Phys. Acta* **64** (1991), 953–1054.

[15] D. Ioffe, Large deviations for the 2D Ising model: a lower bound without cluster expansions, *J. Stat. Phys.* **74** (1994), 411–432. doi:10.1007/BF02186818.

[16] D. Ioffe and R. H. Schonmann, Dobrushin–Kotecký–Shlosman theorem up to the critical temperature, *Commun. Math. Phys.* **199** (1998), 117–167. doi:10.1007/s002200050497.

[17] C. N. Yang, The spontaneous magnetization of a two-dimensional Ising model, *Phys. Rev.* **85** (1952), 808–816. doi:10.1103/PhysRev.85.808.

[18] G. Benettin, G. Gallavotti, G. Jona-Lasinio and A. L. Stella, On the Onsager–Yang-value of the spontaneous magnetization, *Commun. Math. Phys.* **30** (1973), 45–54. doi:10.1007/BF01646687.

[19] H. A. Kramers and G. H. Wannier, Statistics of the two-dimensional ferromagnet. Part I, *Phys. Rev.* **60** (1941), 252–262. doi:10.1103/PhysRev.60.252.

[20] L. Onsager, Crystal statistics. I. A two-dimensional model with an order-disorder transition, *Phys. Rev.* **65** (1944), 117–149. doi:10.1103/PhysRev.65.117.

[21] J. L. Lebowitz and A. Martin-Löf, On the uniqueness of the equilibrium state for Ising spin systems, *Commun. Math. Phys.* **25** (1972), 276–282. doi:10.1007/BF01877686.

[22] J. L. Lebowitz and C.-E. Pfister, Surface tension and phase coexistence, *Phys. Rev. Lett.* **46** (1981), 1031–1033. doi:10.1103/PhysRevLett.46.1031.

[23] J. Bricmont, J. L. Lebowitz and C.-E. Pfister, On the local structure of the phase separation line in the two-dimensional Ising system, *J. Stat. Phys.* **26** (1981), 313–332. doi:10.1007/BF01013174.

[24] J. T. Chayes, L. Chayes and R. H. Schonmann, Exponential decay of connectivities in the two-dimensional Ising model, *J. Stat. Phys.* **49** (1987), 433–445. doi:10.1007/BF01009344.

[25] J. Fröhlich and C.-E. Pfister, Semi-infinite Ising model. I. Thermodynamic functions and phase diagram in absence of magnetic field, *Commun. Math. Phys.* **109** (1987), 493–523. doi:10.1007/BF01206148.

[26] J. Fröhlich and C.-E. Pfister, Semi-infinite Ising model. II. The wetting and layering transitions, *Commun. Math. Phys.* **112** (1987), 51–74. doi:10.1007/BF01217679.

[27] R. Holley, Remarks on the FKG inequalities, *Commun. Math. Phys.* **36** (1974), 227–231. doi:10.1007/BF01645980.

[28] T. E. Harris, A lower bound for the critical probability in a certain percolation process, *Proc. Cambridge Philos. Soc.* **56** (1960), 13–20. doi:10.1017/S0305004100034241.

[29] C. M. Fortuin, P. W. Kasteleyn and J. Ginibre, Correlation inequalities on some partially ordered sets, *Commun. Math. Phys.* **22** (1971), 89–103. doi:10.1007/BF01651330.

[30] W. Hoeffding, Probability inequalities for sums of bounded random variables, *J. Amer. Statist. Assoc.* **58** (1963), 13–30. doi:10.1080/01621459.1963.10500830.

[31] C. McDiarmid, On the method of bounded differences, in *Surveys in Combinatorics, 1989* (J. Siemons, ed.), Cambridge Univ. Press, Cambridge, 1989, 148–188. doi:10.1017/CBO9781107359949.008.

[32] M. Aizenman, Translation invariance and instability of phase coexistence in the two dimensional Ising system, *Commun. Math. Phys.* **73** (1980), 83–94. doi:10.1007/BF01942696.

[33] Y. Higuchi, On the absence of non-translation invariant Gibbs states for the two-dimensional Ising model, in *Random Fields (Esztergom, 1979)*, Colloq. Math. Soc. János Bolyai **27**, North-Holland, Amsterdam, 1981, 517–534.

[34] Y. Higuchi and N. Yoshida, Slow relaxation of 2-D stochastic Ising models with random and non-random boundary conditions, in *New Trends in Stochastic Analysis (Charingworth, 1994)*, World Sci. Publ., River Edge, NJ, 1997, 153–167.

[35] Y. Higuchi, On limiting Gibbs states of the two-dimensional Ising models, *Publ. Res. Inst. Math. Sci.* **14** (1978), 53–69. doi:10.2977/prims/1195189280.

[36] M. Campanino and A. C. D. van Enter, Weak versus strong uniqueness of Gibbs measures: a regular short-range example, *J. Phys. A: Math. Gen.* **28** (1995), L45–L47. doi:10.1088/0305-4470/28/2/001.

[37] G. Gallavotti, The phase separation line in the two-dimensional Ising model, *Commun. Math. Phys.* **27** (1972), 103–136. doi:10.1007/BF01645615.

[38] A. Messager and S. Miracle-Solé, Correlation functions and boundary conditions in the Ising ferromagnet, *J. Stat. Phys.* **17** (1977), 245–262. doi:10.1007/BF01040105.

[39] D. B. Abraham and P. Reed, Interface profile of the Ising ferromagnet in two dimensions, *Commun. Math. Phys.* **49** (1976), 35–46. doi:10.1007/BF01608634.

[40] L. Coquille, Examples of DLR states which are not weak limits of finite volume Gibbs measures with deterministic boundary conditions, *J. Stat. Phys.* **159** (2015), 958–971. doi:10.1007/s10955-015-1211-3; arXiv:1411.3265.

[41] C. M. Newman and D. L. Stein, Multiple states and thermodynamic limits in short-ranged Ising spin-glass models, *Phys. Rev. B* **46** (1992), 973–982. doi:10.1103/PhysRevB.46.973.

[42] C. M. Newman and D. L. Stein, The metastate approach to thermodynamic chaos, *Phys. Rev. E* **55** (1997), 5194–5211. arXiv:cond-mat/9612097.

[43] M. Aizenman and J. Wehr, Rounding effects of quenched randomness on first-order phase transitions, *Commun. Math. Phys.* **130** (1990), 489–528. doi:10.1007/BF02096933.

[44] E. O. Endo, A. C. D. van Enter and A. Le Ny, The roles of random boundary conditions in spin systems, 2019 preprint; published in *In and Out of Equilibrium 3: Celebrating Vladas Sidoravicius*, *Progress in Probability* **77** (2021). arXiv:1911.10490.

[45] E. O. Endo, A. C. D. van Enter and A. Le Ny, On long range Ising models with random boundary conditions, 2024 preprint. arXiv:2405.08374.