MAXIMUM INDEPENDENT SETS ON RANDOM REGULAR GRAPHS

*JIAN DING, †ALLAN SLY, AND ‡NIKE SUN

ABSTRACT. We determine the asymptotics of the independence number of the random d-regular graph for all $d \geqslant d_0$. It is highly concentrated, with constant-order fluctuations around $n\alpha_\star - c_\star \log n$ for explicit constants $\alpha_\star(d)$ and $c_\star(d)$. Our proof rigorously confirms the one-step replica symmetry breaking heuristics for this problem, and we believe the techniques will be more broadly applicable to the study of other combinatorial properties of random graphs.

1. INTRODUCTION

An *independent set* in a graph is a subset of the vertices of which no two are neighbors. Establishing asymptotics of the maximum size of an independent set (the *independence number*) on random graphs is a classical problem in probabilistic combinatorics. On the random d-regular graph $\mathcal{G}_{n,d}$, the independence number grows linearly in the number n of vertices. Upper bounds were established by Bollobás [8] and McKay [20], and lower bounds by Frieze–Suen [16], Frieze–Łuczak [17] and Wormald [24], using a combination of techniques, including first and second moment bounds, differential equations, and switchings. The bounds are quite close, with the maximal density of occupied vertices (the *independence ratio*) roughly asymptotic to $2(\log d)/d$ in the limit of large d — however, for every fixed d there remains a constant-size gap in the bounds on the independence ratio. For a more complete history and discussion of many related topics see the survey of Wormald [25].

In fact, a long-standing open problem (see [1, 3]) was to determine if there even exists a limiting independence ratio. A standard martingale bound implies that the independence number has $O(n^{1/2})$ fluctuations about its mean, so an equivalent question was to prove convergence of the expected independence ratio. This conjecture was recently resolved by Bayati–Gamarnik–Tetali [6] using interpolation methods from statistical physics. Their method is based on a sub-additivity argument which does not yield information on the limiting independence ratio or the order of fluctuations.

In this paper, we establish for all sufficiently large d the asymptotic independence ratio $\alpha_\star = \lim_n \boldsymbol{A}_n/n$, and determine also the lower-order logarithmic correction. Further, we prove tightness of the *non-normalized* independence number, proving that the random variable is much more strongly concentrated than suggested by the classical $O(n^{1/2})$ bound:

Theorem 1. The maximum size $\boldsymbol{A}_n$ of an independent set in the random d-regular graph $\mathcal{G}_{n,d}$ has constant fluctuations about

$$\bar{\boldsymbol{A}}_n \equiv n\alpha_\star - c_\star \log n \tag{1}$$

for $\alpha_\star$ and $c_\star$ explicit functions of d, provided d exceeds an absolute constant d_0.

The values $\alpha_\star$ and $c_\star$ are given in terms of a function ${}^\star\mathbf{\Phi} \equiv {}^\star\mathbf{\Phi}_d$ defined on the interval $(5/3)(\log d)/d \leqslant \alpha \leqslant 2(\log d)/d$: the function is smoothly decreasing with a unique zero $\alpha_\star$, and we set $c_\star \equiv -[2 \cdot {}^\star\mathbf{\Phi}'(\alpha_\star)]^{-1}$ (positive since the function is decreasing). Explicitly,

$$ {}^\star\mathbf{\Phi}(\alpha) = -\log[1 - q(1 - 1/\lambda)] - (d/2 - 1)\log[1 - q^2(1 - 1/\lambda)] - \alpha \log \lambda \qquad (2) $$

where q is determined from α by solving the equation

$$ \alpha = q\frac{1 + (d/2 - 1)(1 - q)^{d-1}}{1 + q - (1 - q)^{d-1}} \text{ on the interval } 1.6(\log d)/d \leqslant q \leqslant 3(\log d)/d, $$

and $\lambda = q[1 - (1 - q)^{d-1}]/(1 - q)^d$. For λ determined from α in this manner we will see that ${}^\star\mathbf{\Phi}'(\alpha) = -\log \lambda$, therefore $c_\star = [2 \log \lambda_\star]^{-1}$ for $\lambda_\star$ corresponding to $\alpha_\star$.

A natural question is whether the same behavior holds for regular graphs of low degree. Though it is certainly possible to determine an explicit d_0 from our proof, we have not done so because the calculations in the paper are already daunting, and have not been carried out with a view towards optimizing d_0. More importantly, our result is in line with the *one-step replica symmetry breaking* (1RSB) prediction, which is believed to fail on low-degree graphs where physicists expect *full replica symmetry breaking* [5]. In the latter regime no formula is predicted even at a heuristic level.

1.1. **Replica symmetry breaking.** Ideas from statistical physics have greatly advanced our understanding of random constraint satisfaction and combinatorial optimization problems [19, 22, 21]. This deep, but for the most part non-rigorous, theory has led to a detailed picture of the geometry of the space of solutions for a broad class of such problems, including exact predictions for their satisfiability threshold. While some aspects of this rich picture have been established, including celebrated results such as Aldous's solution to the random assignment problem [4] and Talagrand's proof of Parisi's formula for the Sherrington–Kirkpatrick spin-glass model [23], many of the most important ideas remain at the level of conjecture. We believe that recent developments, including our own previous work [13, 12], make it possible to establish thresholds predicted by this theory for many such models.

The natural approach to studying the independence ratio is the (first and second) moment method applied to the number of $Z_{n\alpha}$ of independent sets of fixed density α. Indeed, an analogous approach correctly determines the asymptotic independence number for the *dense* Erdős-Rényi random graph [18]. On *sparse* random graph ensembles, however, the second moment approach fails to locate the sharp transition. Due to the sparsity of the graph, almost every independent set can be locally perturbed in a linear number of places: thus the existence of a single independent set implies the existence of a *cluster* of exponentially many independent sets, all related by sequences of local perturbations. Moreover the expected cluster size remains exponentially large even beyond the first moment threshold — thus there is a regime below the first moment threshold where it overcomes the first moment, causing the second moment to be exponentially large compared with the first moment squared.

From statistical physics, the (mostly heuristic) understanding of this phenomenon is that as α exceeds roughly $(\log d)/d$, the solution space of independent sets becomes *shattered* into exponentially many well-separated clusters [9]. This geometry persists up to a further (conjectured) structural transition where the solution space *condensates* onto the largest clusters. In the non-trivial regime between the condensation and satisfiability transitions, most independent sets are concentrated within a bounded number of clusters according to the

theory from statistical physics. This within-cluster correlation then dominates the moment calculation, causing the failure of the second moment method.

In this paper, we determine the exact threshold by a novel approach which rigorizes the 1RSB heuristic from statistical physics, which suggests that we count clusters of independent sets rather than the sets themselves. Our proof has several new ideas which we now describe.

The first is that most clusters of maximal (or locally maximal) independent sets can be given the following simple combinatorial description. A cluster is encoded by a configuration of what we call the "frozen model", which takes values in $\{0, 1, \mathtt{f}\}^V$ where the $\mathtt{f}$ ("free") spins indicate the local perturbations within the cluster, and come in matched pairs corresponding to swapping occupied/unoccupied states across an edge. Each 1 indicates an occupied vertex, and can only neighbor 0 (unoccupied) vertices; and each 0 must neighbor at least two 1's. We shall prove that this effectively encodes the clusters in this model. The expression $^\star\Phi$ in equation (2) is exactly the free energy of this model, and the asymptotic density $\alpha_\star$ corresponds to the largest density for which it is positive.

Secondly, we note that this new model is itself a Gibbs measure on a random hypergraph, and its properties of local rigidity hint that applying the second moment in this model *does* locate the exact threshold. However, the actual moment calculation appears at first intractable, involving maximizations over high-dimensional simplices. By a certain "Bethe variational principle" [13, 12] we are able to characterize local maximizers via fixed points of certain tree recursions, reducing the optimization to (in the second moment) 81 real variables. With delicate *a priori* estimates we are able to establish symmetry relations among these variables which drastically reduce the dimensionality and allows us finally to pinpoint the global maximizers.

The second moment method itself only establishes the existence of clusters with asymptotically positive probability. Our final innovation is a method to improve positive probability bounds to high probability, which in this model yields the constant fluctuations. The approach is based on controlling the incremental fluctuations of the Doob martingale of a certain log-transform of the partition function.

As an illustration of the robustness of these methods, in a companion paper [15] we apply the same techniques to establish the exact satisfiability threshold for the random regular not-all-equal-SAT problem. This gives the first threshold for a sparse constraint satisfaction model with replica symmetry breaking. We expect ultimately that these methods may be extended to other combinatorial properties such as the chromatic number or maximum cut, and to the sparse Erdős-Rényi random graphs.

1.2. **Notation.** Throughout this paper, a *graph* $G = (V, M, H)$ consists of a vertex set V, a collection H of labelled half-edges (with each half-edge incident to a specified vertex), and a perfect matching M on H. Omitting the labelling of half-edges gives the undirected graph $G = (V, E)$ where $(x, y) \in E$ if and only if $(e_x, e_y) \in M$ for half-edges e_x and e_y incident to x and y respectively. A subset $S \subseteq V$ is an *independent set* of $G = (V, E)$ if $(x, y) \notin E$ for all $x, y \in S$.[1] We equivalently view an independent set as a configuration in $\{0, 1\}^V$ where 1 indicates a vertex occupied by the set.

We work with the d-regular configuration model: a *d-regular graph* with vertex set $[n] \equiv \{1, \ldots, n\}$ is a perfect matching on the set $[nd]$ of labelled half-edges, where half-edge i is incident to vertex $[i/d]$. Assume nd is even; the number of d-regular graphs on $[n]$ is the

---

[1]If E contains a self-loop (x, x) then x cannot belong to an independent set.

double factorial

$$(nd-1)!! = \prod_{j=0}^{nd/2}(nd-2j+1) = \frac{(nd)!}{(nd/2)!2^{nd/2}} = e^{O[1/(nd)]}\sqrt{2}(nd/e)^{nd/2}.$$

Under the configuration model, the *random d-regular graph* $\mathcal{G}_{n,d}$ corresponds to the uniformly random perfect matching on $[nd]$.

A graph is *simple* if it has no self-loops or multi-edges. For each simple d-regular graph with unlabelled half-edges, there are $(d!)^n$ d-regular graphs with labelled half-edges, so the graph $\mathcal{G}_{n,d}$ conditioned to be simple is the *(uniformly) random simple d-regular graph $\mathcal{G}_{n,d}^{\mathrm{s}}$.* It is well known that the probability that $\mathcal{G}_{n,d}$ is simple is bounded below as $n \to \infty$ by a constant (depending on d), so Thm. 1 also holds with $\mathcal{G}_{n,d}^{\mathrm{s}}$ in place of $\mathcal{G}_{n,d}$.

Acknowledgements.

2. Independent sets and coarsening algorithm

For non-negative functions $f(d,n)$ and $g(d,n)$ we use any of the equivalent notations $f = O_d(g)$, $g = \Omega_d(f)$, $f \lesssim_d g$, $g \gtrsim_d f$ to indicate $f \leqslant C(d)g$ for a finite constant $C(d)$ depending on d but not on n. We drop the subscript d to indicate that we can take the same constant $C(d) \equiv C$ for all $d \geqslant d_0$.

2.1. **First moment of independent sets.** Let $Z_{n\alpha}$ count the number of independent sets of cardinality $n\alpha$ on graph G.

Lemma 2.1. $\mathbb{E}Z_{n\alpha} \asymp_d n^{-1/2}\exp\{n\Phi_d(\alpha)\}$ *for Φ_d an explicit function not depending on n. The first moment threshold $\alpha_\square \equiv \alpha_\square(d) \equiv \inf\{\alpha > 0 : \Phi_d(\alpha) < 0\}$ satisfies*

$$\alpha_\square = (2/d)\cdot\left[\log d - \log\log d + \log(e/2) + O\left(\tfrac{\log\log d}{\log d}\right)\right]. \tag{3}$$

Further

$$\frac{\mathbb{E}Z_{n(\alpha_\square-\delta)}}{\exp\{n\delta[\log d - \log\log d]\}} \begin{cases} \leqslant e^{O(n\delta)} & \text{for all } 0 \leqslant \delta \leqslant \alpha_\square; \\ \geqslant e^{O(n\delta)} & \text{for } 0 \leqslant \delta \lesssim d^{-1}. \end{cases} \tag{4}$$

Proof. The expected number of independent sets of size $n\alpha$ on $\mathcal{G}_{n,d}$ is

$$\mathbb{E}Z_{n\alpha} = \binom{n}{n\alpha}\prod_{i=0}^{nd\alpha-1}\frac{nd(1-\alpha)-i}{nd-1-2i} = \binom{n}{n\alpha}\frac{(nd(1-\alpha))_{nd\alpha}}{[nd]_{nd\alpha}},$$

where $(A)_b$ and $[A]_b$ denote respectively the falling factorial and falling double factorial:

$$(A)_b \equiv \prod_{i=0}^{b-1}(A-i), \quad [A]_b \equiv \prod_{i=0}^{b-1}(A-1-2i). \tag{5}$$

We assume $\alpha \leqslant 1/2$ since otherwise $\mathbb{E}Z_{n\alpha} = 0$. For α constant in n, Stirling's formula gives $\mathbb{E}Z_{n\alpha} \asymp_d n^{-1/2}\exp\{n\Phi(\alpha)\}$ where

$$\begin{aligned}\Phi(\alpha) \equiv \Phi_d(\alpha) &\equiv H(\alpha) - d[\tfrac{1}{2}(1-2\alpha)\log(1-2\alpha) - (1-\alpha)\log(1-\alpha)] \\ &= \alpha[F(\alpha) + O(d\alpha^2)] \quad \text{with } F(\alpha) \equiv \log(e/\alpha) - (d+1)(\alpha/2).\end{aligned} \tag{6}$$

It is straightforward to check that $\Phi'(\alpha) > 0$ for small α and $\Phi''(\alpha) < 0$ for all $\alpha < \frac{1}{2}$, so Φ has a unique zero-crossing $\alpha_\square$. The function F is decreasing in α with unique zero

$$\tilde{\alpha}_\square = \frac{2}{(d+1)}W\left(\frac{(d+1)e}{2}\right) = [2 + o_d(1)]\frac{\log d}{d}$$

where W denotes the principal branch of the Lambert W function defined by $z = W(z)e^{W(z)}$ (see [10, 11] and references therein). Near $z = \infty$, the function W has absolutely convergent series expansion

$$W(z) = \log z - \log \log z + \sum_{k \geqslant 0} \frac{1}{(-\log z)^k} \overbrace{\sum_{m \geqslant 1} \frac{1}{m!} \begin{bmatrix} k+m \\ k+1 \end{bmatrix} \left(\frac{\log \log z}{\log z} \right)^m}^{O((\log \log z)/(\log z))},$$

where $\begin{bmatrix} n \\ m \end{bmatrix}$ are the Stirling cycle numbers (or unsigned Stirling numbers of the first kind), generated by $\log(1+z)^m = m! \sum_{n \geqslant 0} (-1)^{n+m} \begin{bmatrix} n \\ m \end{bmatrix} z^n/n!$. By estimating F' and F'' near $\tilde{\alpha}_\square$ we see that $\alpha_\square = \tilde{\alpha}_\square + O(d^{-2} \log d)$; (3) then follows from the above estimate on $W(z)$. The bounds (4) are easily obtained by estimating Φ' near $\alpha_\square$ and recalling that $\Phi'' < 0$. $\qquad \square$

2.2. Coarsening algorithm and frozen model.
We hereafter encode an independent set as a configuration $\underline{x} \in \{0,1\}^V$ with 1 indicating a vertex which is occupied by the independent set. We define the following algorithm to map an independent set $\underline{x} \in \{0,1\}^V$ to a *coarsened* configuration $\underline{\eta} \equiv \underline{\eta}(\underline{x}) \in \{0,1,\mathtt{f}\}^V$. In the coarsened model the spin $\mathtt{f}$ indicates vertices which are "free," as follows:

Coarsening algorithm. Set $\underline{\eta}^0 \equiv \underline{x}$.

1. *Step 1 (iterate for $0 \leqslant s < t$): formation of free pairs.*
   If there exists $v \in V$ such that $\eta_v^s = 0$ and v has a unique neighbor u with $\eta_u^s = 1$, then take the first[2] such v, set $\eta_v^{s+1} = \eta_u^{s+1} = \mathtt{f}$, and declare (uv) to be a matched edge. Set $\eta_w^{s+1} = \eta_w^s$ for all $w \neq u, v$.
   Iterate until the first time t that no such vertex v remains.
2. *Step 2: formation of single frees.*
   Set $\eta_v^{t+1} = \mathtt{f}$ whenever $\eta_v^t = 0$ and $\{u \in \partial v : \eta_u^t = 1\} = \varnothing$, otherwise set $\eta_v^{t+1} = \eta_v^t$.

Denote the terminal configuration $\underline{\eta} \equiv \underline{\eta}(\underline{x}) \equiv \underline{\eta}^{t+1}$. Write $\underline{m} \subseteq E$ for the set of all matched free pairs formed during Step 1 of the coarsening process.

The idea is that the pre-image of any $\underline{\eta}$ under the coarsening algorithm constitutes a *cluster* of independent set configurations — a set of configurations connected by (sequences of) local changes, in our setting by making neighboring $0/1$ swaps. An important property of a coarsened configuration is that every 0-vertex has at least two 1-neighbors: as this is a rigid local configuration (a $0/1$ swap across an edge cannot be made without violating the hard-core constraint), we have some indication (non-rigorously) that different clusters will be well separated in some sense. Note that Step 2 of the coarsening algorithm is needed to ensure this property even when the initial configuration is a maximum independent set: consider for example $0 - 1 - 0 - 1 - 0$ arranged in a 5-cycle such that every neighbor not on the cycle is a 0-vertex with many 1-neighbors. This can be part of a maximal configuration, but Step 1 results with $\mathtt{f} \overset{m}{-} \mathtt{f} - \mathtt{f} \overset{m}{-} \mathtt{f} - 0$ on the cycle where m indicates a matched pair and the last 0 has no 1-neighbors. The purpose of this section is show that we may discard these odd-cycle scenarios and still recover the sharp asymptotics for $\boldsymbol{A}_n$.

Let $Z_{n\alpha, \geqslant n\beta}$ denote the contribution to $Z_{n\alpha}$ from independent set configurations $\underline{x} \in \{0,1\}^V$ such that the coarsened configuration $\underline{\eta}(\underline{x}) \in \{0,1,\mathtt{f}\}^V$ has more than $n\beta$ free variables.

---

[2]First with respect to the ordering on $V = [n]$.

Proposition 2.2. *There exists an absolute constant C (not depending on d) such that if $\alpha = y(\log d)/d$ for $1 \leqslant y \leqslant 2$, then $\mathbb{E}Z_{n\alpha,\geqslant n\beta} \leqslant C^{-n\beta/2}\mathbb{E}Z_{n\alpha}$ for $\beta = d^{-y}(C\log d)$. In particular if $\alpha = \alpha_\square - O(d^{-3/2})$ then $\mathbb{E}Z_{n\alpha,\geqslant n\beta} \leqslant e^{-nc_d}$ for $\beta = C(\log d)^3/d^2$.*

Proof. Suppose $\underline{x} \in \{0,1\}^V$ such that $V_1 \equiv \{v : x_v = 1\}$ has size $n\alpha$, and write $V_0 \equiv V \backslash V_1$. By definition of the coarsening algorithm, any subset S of the f-vertices must satisfy $|\partial S \cap V_1| \leqslant |S|$. Conditioned on $\underline{x}$ being an independent set of cardinality $n\alpha$, the random graph G can be obtained by first matching all half-edges originating from V_1 to a subset of half-edges originating from V_0 uniformly at random, then placing a uniformly random matching on the remaining half-edges of V_0. For $v \in V_0$ let D_v denote the number of neighbors of v in V_1; then $(D_v)_{v \in V_0}$ is distributed as a vector of i.i.d. $\mathrm{Bin}(d, \alpha/(1-\alpha))$ random variables conditioned on $\sum_{v \in V_0} D_v = nd\alpha$. Fixing S any subset of V_0 of size $n\beta$, and recalling that $\alpha = y(\log d)/d$ with $1 \leqslant y \leqslant 2$, we have

$$(\mathbb{E}Z_{n\alpha})^{-1}\mathbb{E}Z_{n\alpha,\geqslant n\beta} \leqslant \binom{n(1-\alpha)}{n\beta}\mathbb{P}(\textstyle\sum_{v \in S} D_v \leqslant n\beta \mid \sum_{v \in V_0} D_v = nd\alpha) = n^{O(1)}\exp\{nf_d(\alpha,\beta)\},$$
$$f_d(\alpha,\beta) \equiv (1-\alpha)H(\tfrac{\beta}{1-\alpha}) - d\beta\, H(\tfrac{1}{d} \mid \tfrac{\alpha}{1-\alpha}) = \beta[\log \tfrac{y\log d}{\beta d^y} + O(1)]$$

with the $O(1)$ uniform in d. Taking $\beta = d^{-y}(C\log d)$ for C a large absolute constant then proves the first claim, $\mathbb{E}Z_{n\alpha,\geqslant n\beta} \leqslant C^{-n\beta/2}\mathbb{E}Z_{n\alpha}$. The second claim follows simply by noting that if $\alpha = \alpha_\square - O(d^{-3/2})$ then by Lem. 2.1 we have $\alpha = y(\log d)/d$ with $y\log d = 2\log d - 2\log\log d + O(1)$ with the $O(1)$ uniform in d. $\square$

Definition 2.3. We say $\underline{\eta} \in \{0,1,\mathtt{f}\}^V$ is a valid *unweighted frozen model* configuration on G if the following hold:

(a) Every $v \in V$ satisfies
    (i) If $\eta_v = 1$ then v neighbors only 0-vertices;
    (ii) If $\eta_v = 0$ then at least two edges join v to 1-vertices;
    (iii) If $\eta_v = \mathtt{f}$ then v neighbors only 0- or f-vertices, with at least one neighboring f;
(b) In the subgraph $\mathfrak{F}(\eta) \subseteq G$ induced by the f-vertices, every connected component either is not a tree, or is a tree with a (necessarily unique) perfect matching.[3]

Definition 2.4. A *(weighted) frozen model configuration* $\overline{\underline{\eta}} \equiv (\eta, \underline{m})$ on G is a unweighted frozen model configuration η together with a perfect matching $\underline{m}$ on $\mathfrak{F}(\eta)$: equivalently, every $v \in V$ satisfies

  (i) If $\eta_v = 1$ then v neighbors only 0-vertices;
  (ii) If $\eta_v = 0$ then at least two edges join v to 1-vertices;
  (iii) If $\eta_v = \mathtt{f}$ then v neighbors only 0- or f-vertices, and is matched to an f-neighbor by $\underline{m}$.

The *intensity* of $\overline{\underline{\eta}}$ is the number of 1-vertices plus the number of matched f-pairs:

$$\mathbf{i}(\overline{\underline{\eta}}) \equiv |\{v : \eta_v = 1\}| + \tfrac{1}{2}|\{v : \eta_v = \mathtt{f}\}|. \tag{7}$$

Let $\overline{\boldsymbol{Z}}_{n\alpha_1,n\beta}$ denote the *unweighted* frozen model partition function on G restricted to configurations with exactly $n\alpha_1$ 1-vertices and $n\beta$ f-vertices; define similarly $\boldsymbol{Z}_{n\alpha_1,n\beta}$ with respect to the *weighted* frozen model. In view of Lem. 2.1 and Propn. 2.2, we restrict all consideration to the *truncated* frozen model partition function

$$\boldsymbol{Z}_{n\alpha} \equiv \sum_{\substack{\alpha_1,\beta\,:\,\beta \leqslant \beta_{\max} \\ 2\alpha_1+\beta=2\alpha}} \boldsymbol{Z}_{n\alpha_1,n\beta}, \qquad \beta_{\max} \equiv d^{-3/2}. \tag{8}$$

---

[3]To be precise, since we regard G as a matching on the set $[nd]$ of labelled half-edges incident to vertices, $\mathfrak{F}$ shall be regarded as a matching on a subset of $[nd]$.

We shall always assume that the normalized intensity α lies in a restricted regime:

$$\alpha_{\mathrm{lbd}} \equiv \frac{(5/3)\log d}{d} \leqslant \alpha \leqslant \frac{2\log d}{d} \equiv \alpha_{\mathrm{ubd}}, \quad \boldsymbol{Z}_{\geqslant n\alpha} \equiv \sum_{n\alpha \leqslant n\alpha' \leqslant n\alpha_{\mathrm{ubd}}} \boldsymbol{Z}_{n\alpha'}.$$

We now describe our reduction from the independent set model to the frozen model.

Theorem 2.5. *For $d \geqslant d_0$, it holds uniformly over $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ that*

$$\lim_{n\to\infty} n^{1/2} \exp\{-n\,{}^{\star}\boldsymbol{\Phi}(\alpha)\} \, \mathbb{E}\boldsymbol{Z}_{n\alpha} = \mathscr{C}(\alpha)$$

for ${}^{\star}\boldsymbol{\Phi} \equiv {}^{\star}\boldsymbol{\Phi}_d$, $\mathscr{C} \equiv \mathscr{C}_d$ smooth functions of α. The explicit form of ${}^{\star}\boldsymbol{\Phi}$ is given in §3.4 below; in particular, on the interval $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ it is strictly decreasing with a unique zero $\alpha_{\star}$ in the interval's interior.

The proof is given in §5 (making use of §2 and §3).

Corollary 2.6. *With $\alpha_{\star}$ the unique zero of ${}^{\star}\boldsymbol{\Phi}$ on the interval $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ (Thm. 2.5), let $c_{\star} \equiv -[2 \cdot {}^{\star}\boldsymbol{\Phi}'(\alpha_{\star})]^{-1} > 0$. Then, for $\bar{\boldsymbol{A}}_n \equiv n\alpha_{\star} - c_{\star}\log n$ as in (1), $\mathbb{E}\boldsymbol{Z}_{\geqslant \bar{\boldsymbol{A}}_n + C}$ tends to zero as C tends to infinity, uniformly over $n \geqslant n_0(d)$.*

Proof. Taylor expansion gives $\mathbb{E}\boldsymbol{Z}_{\bar{\boldsymbol{A}}_n} \asymp_d 1$ while $\mathbb{E}\boldsymbol{Z}_{\bar{\boldsymbol{A}}_n + 5c_{\star}\log n} \leqslant n^{-2}$. Recall from Thm. 2.5 that ${}^{\star}\boldsymbol{\Phi}$ is decreasing on the interval $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$: since $\mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}$ is trivially a sum over at most n terms $\mathbb{E}\boldsymbol{Z}_{n\alpha'}$, this immediately implies $\mathbb{E}\boldsymbol{Z}_{\geqslant \bar{\boldsymbol{A}}_n + 5c_{\star}\log n} \leqslant n^{-1}$. For $0 \leqslant \delta \leqslant 5c_{\star}\log n$ we have $\mathbb{E}\boldsymbol{Z}_{\bar{\boldsymbol{A}}_n + \delta} \asymp_d (\lambda_{\star})^{-\delta}\exp\{O_d(n^{-1}(\log n)^2)\}$ where $\lambda_{\star} \equiv e^{1/(2c_{\star})}$. Combining these estimates gives

$$\mathbb{E}\boldsymbol{Z}_{\geqslant \bar{\boldsymbol{A}}_n + C} \asymp_d (\lambda_{\star})^{-C}\exp\{O_d(n^{-1}(\log n)^2)\}\sum_{\delta\geqslant 0}(\lambda_{\star})^{-\delta} + n^{-1} \lesssim (\lambda_{\star})^{-C},$$

implying the result. $\square$

Theorem 2. *For $d \geqslant d_0$, it holds uniformly over $n \geqslant n_0(d)$ that $\mathbb{P}(\boldsymbol{Z}_{\bar{\boldsymbol{A}}_n - C} > 0)$ tends to one as C tends to infinity.*

Proof of Thm. 1. We prove the theorem relying on Propns. 2.7 and 2.8 which will be proved in the remainder of this section.

Upper bound. Let $\mathscr{C}_{\geqslant n\alpha}$ denote the number of configurations η obtained by coarsening a *maximum*[4] independent set of size $\geqslant n\alpha$. Write $\mathcal{O}(\mathfrak{F})$ for the number of odd-size components in $\mathfrak{F}$; recall Defn. 2.3 that each odd component must contain at least one cycle. Define $\bar{\boldsymbol{Z}}_{n\alpha_1, n\beta}(O)$ to be the contribution to $\bar{\boldsymbol{Z}}_{n\alpha_1, n\beta}$ from configurations η with $\mathcal{O}[\mathfrak{F}(\eta)] = O$, and write ${}^{\mathrm{ev}}\bar{\boldsymbol{Z}}_{n\alpha_1, n\beta} \equiv \bar{\boldsymbol{Z}}_{n\alpha_1, n\beta}(O = 0)$. In any coarsening η of a *maximal* independent set $\underline{x}$, there can be no isolated f-vertices, and all tree f-components must have a perfect matching. Thus, recalling Propn. 2.2, we have that with high probability $\mathscr{C}_{\geqslant n\alpha}$ is bounded above by

$$\bar{\boldsymbol{Z}}_{\geqslant n\alpha} \equiv \left[\sum_{\substack{\alpha_1,\beta:\beta\leqslant\beta_{\max}, \\ 2n\alpha\leqslant 2n\alpha_1+n\beta}} {}^{\mathrm{ev}}\bar{\boldsymbol{Z}}_{n\alpha_1, n\beta}\right] + \left[\sum_{O>0}\sum_{\substack{\alpha_1,\beta:\beta\leqslant\beta_{\max}, \\ 2n\alpha+O\leqslant 2n\alpha_1+n\beta}} \bar{\boldsymbol{Z}}_{n\alpha_1, n\beta}(O)\right] \equiv {}^{\mathrm{ev}}\bar{\boldsymbol{Z}}_{\geqslant n\alpha} + {}^{\mathrm{odd}}\bar{\boldsymbol{Z}}_{\geqslant n\alpha} \quad (9)$$

For $k \geqslant 4$ even we shall set a threshold $\ell_k \equiv \ell_k(n\beta) \geqslant 1$ (to be determined), and further decompose

$$\overset{\equiv\,{}^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\circ}_{\geqslant n\alpha}}{{}^{\mathrm{ev}}\bar{\boldsymbol{Z}}_{\geqslant n\alpha} \equiv {}^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\mathrm{tree}}_{\geqslant n\alpha} + {}^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\infty}_{\geqslant n\alpha} + \overbrace{\sum_{A\geqslant 1} {}^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\circ}_{\geqslant n\alpha}(A)}} \quad \text{where} \quad (10)$$

---

[4]It would suffice for the independent set to be "locally maximal" in the sense that its coarsening is a configuration of the unweighted frozen model.

1. $^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\infty}_{\geqslant n\alpha}$ is the contribution from configurations $\underline{\eta}$ such that the number of f-components of size k is $\geqslant \ell_k(|\mathfrak{F}(\underline{\eta})|)$ for some $k \geqslant 4$ even;
2. $^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\circ}_{\geqslant n\alpha}(A)$ is the contribution from configurations $\underline{\eta}$ such that for all $k \geqslant 4$ even there are fewer than ℓ_k f-components of size k, where a minmum of A edges must be removed from $\mathfrak{F}$ to make it a forest; and
3. $^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\mathrm{tree}}_{\geqslant n\alpha}$ is the remainder, consisting of configurations $\underline{\eta}$ where every f-component is a tree with (unique) perfect matching, and with $< \ell_k$ size-k f-components for all $k \geqslant 4$ even.

All statements that follow are made uniformly over $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ for $d \geqslant d_0$, $n \geqslant n_0(d)$. We shall prove (Propns. 2.7 and 2.8) that $\mathbb{E}[^{\mathrm{odd}}\bar{\boldsymbol{Z}}_{\geqslant n\alpha} + {}^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\infty}_{\geqslant n\alpha} + {}^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\circ}_{\geqslant n\alpha}] \leqslant \frac{1}{2}\mathbb{E}\bar{\boldsymbol{Z}}_{\geqslant n\alpha}$, consequently

$$\tfrac{1}{2}\mathbb{E}\bar{\boldsymbol{Z}}_{\geqslant n\alpha} \leqslant \mathbb{E}[^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\mathrm{tree}}_{\geqslant n\alpha}] = \mathbb{E}\boldsymbol{Z}^{\mathrm{tree}}_{\geqslant n\alpha} \leqslant \mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}. \tag{11}$$

Recall Cor. 2.6 that $\mathbb{E}\boldsymbol{Z}_{\geqslant \bar{\boldsymbol{A}}_n + C} \to 0$ uniformly in n as $C \to \infty$, and therefore

$$\mathbb{P}(\boldsymbol{A}_n \geqslant \bar{\boldsymbol{A}}_n + C) \leqslant \mathbb{P}(\mathscr{C}_{\geqslant \bar{\boldsymbol{A}}_n + C} > 0) \leqslant o_n(1) + \mathbb{P}(\bar{\boldsymbol{Z}}_{\geqslant \bar{\boldsymbol{A}}_n + C} > 0)$$

converges to zero as $C \to \infty$, uniformly in n. (In the above, the first inequality is by the coarsening algorithm, the second inequality is by Propn. 2.2 with $o_n(1)$ indicating an error which tends to zero as $n \to \infty$, and the rightmost expression tends to zero by Markov's inequality and the chain of inequalities (11).)

Lower bound. In the weighted frozen model, by definition there are no odd-size components, so $\boldsymbol{Z}_{\geqslant n\alpha}$ has a similar decomposition as that of $^{\mathrm{ev}}\bar{\boldsymbol{Z}}$ in (10),

$$\boldsymbol{Z}_{\geqslant n\alpha} \equiv \boldsymbol{Z}^{\mathrm{tree}}_{\geqslant n\alpha} + \boldsymbol{Z}^{\infty}_{\geqslant n\alpha} + \boldsymbol{Z}^{\circ}_{\geqslant n\alpha}, \quad \boldsymbol{Z}^{\circ}_{\geqslant n\alpha} \equiv \sum_{A \geqslant 1} \boldsymbol{Z}^{\circ}_{\geqslant n\alpha}(A)$$

We shall show (Propn. 2.7(b) and Propn. 2.8(b)) that $\mathbb{E}\boldsymbol{Z}^{\infty}_{\geqslant n\alpha} \leqslant n^{-1/2}\mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}$ and $\mathbb{E}\boldsymbol{Z}^{\circ}_{\geqslant n\alpha}(A) \leqslant [d(\log n)/n]^A \mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}$. If $\bar{\underline{\eta}} = (\underline{\eta}, \underline{m})$ contributes to $\boldsymbol{Z}^{\circ}_{\geqslant n\alpha}(A)$ then one can set $\eta_u = \eta_v = 0$ on A edges $(u, v) \in \underline{m}$ (removing those edges from the matching) to obtain a configuration $\bar{\underline{\eta}}'$ contributing to $\boldsymbol{Z}^{\mathrm{tree}}_{\geqslant n\alpha - A}$, which immediately gives rise to independent sets of size $n\alpha - A$. Thus, with $\lambda_{\star} \equiv \exp\{-^{\star}\boldsymbol{\Phi}'(\alpha)\}$ as in the proof of Cor. 2.6, we have

$$\mathbb{P}(\boldsymbol{A}_n \geqslant \bar{\boldsymbol{A}}_n - 2C) \geqslant \mathbb{P}(\boldsymbol{Z}_{\geqslant \bar{\boldsymbol{A}}_n - C} > 0) - \mathbb{E}\boldsymbol{Z}^{\infty}_{\geqslant \bar{\boldsymbol{A}}_n - C} - \sum_{A > C} \mathbb{E}\boldsymbol{Z}^{\circ}_{\bar{\boldsymbol{A}}_n - C}(A)$$
$$\geqslant \mathbb{P}(\boldsymbol{Z}_{\geqslant \bar{\boldsymbol{A}}_n - C} > 0) - e^{O_d(1)}(\lambda_{\star})^C (n^{-1/2} + [d(\log n)/n]^C),$$

and by Thm. 2 this converges to one as $C \to \infty$, uniformly in n. Combining with the upper bound proves that $\boldsymbol{A}_n - \bar{\boldsymbol{A}}_n$ is a tight random variable as claimed. $\qquad\square$

The above equivalence between the maximum independent set size and the threshold of the frozen model is based on the following

Proposition 2.7. *The following hold uniformly over* $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ *for* $d \geqslant d_0$, $n \geqslant n_0(d)$, *with* c_0 *an absolute constant.*
(a) $\mathbb{E}[^{\mathrm{odd}}\bar{\boldsymbol{Z}}_{\geqslant n\alpha}] \leqslant d^{-1/2}\mathbb{E}\bar{\boldsymbol{Z}}_{\geqslant n\alpha}$.
(b) If for $k \geqslant 4$ *even we set* $\ell_k \equiv [(c_0 d\beta)^{k/8} n/(dk)] \vee 1$, *then*

$$\mathbb{E}[^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\infty}_{\geqslant n\alpha}] \leqslant n^{-1/2}\,\mathbb{E}\bar{\boldsymbol{Z}}_{\geqslant n\alpha} \quad and \quad \mathbb{E}\boldsymbol{Z}^{\infty}_{\geqslant n\alpha} \leqslant n^{-1/2}\,\mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}.$$

Proposition 2.8. *The following hold uniformly over* $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ *for* $d \geqslant d_0$, $n \geqslant n_0(d)$, *with* ℓ_k *as in the statement of Propn. 2.7.*
(a) $\mathbb{E}[^{\mathrm{ev}}\bar{\boldsymbol{Z}}^{\circ}_{\geqslant n\alpha}] \leqslant d^{-1/5}\mathbb{E}\bar{\boldsymbol{Z}}_{\geqslant n\alpha}$.
(b) $\mathbb{E}\boldsymbol{Z}^{\circ}_{\geqslant n\alpha}(A) \leqslant [dn^{-1}\log n]^A \mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}$.

2.3. **Large components and trees with matchings.** We now prove Propns. 2.7 and 2.8. Write $n_1 \equiv n\alpha_1$, $n_{\mathfrak{f}} \equiv n\beta$, and $n_0 \equiv n - n_1 - n_{\mathfrak{f}}$. We decompose

$$\mathbb{E}\bar{Z}_{n_1,n_{\mathfrak{f}}} = \sum_{|\mathfrak{F}|=n_{\mathfrak{f}}} \mathbb{E}[\bar{Z}_{n_1}(\mathfrak{F})], \quad \mathbb{E}Z_{n_1,n_{\mathfrak{f}}} = \sum_{|\mathfrak{F}|=n_{\mathfrak{f}}} \mathfrak{m}(\mathfrak{F})\,\mathbb{E}[\bar{Z}_{n_1}(\mathfrak{F})] \tag{12}$$

where $\bar{Z}_{n_1}(\mathfrak{F})$ denotes the partition function of unweighted frozen model configurations η with $|\{v : \eta_v = 1\}| = n_1$ and with f-subgraph $\mathfrak{F}(\eta) = \mathfrak{F}$, and the sum is taken over all possible f-subgraphs $\mathfrak{F}$ containing $|\mathfrak{F}| = n_{\mathfrak{f}}$ vertices. In keeping with Defn. 2.3, we regard each $\mathfrak{F}$ as a matching on a subset of $[nd]$: thus the quantity $\bar{Z}_{n_1}(\mathfrak{F})$ cannot be positive unless $\mathfrak{F}$ is present as an induced subgraph in the random graph $\mathcal{G}_{n,d}$.

Observe that $\mathbb{E}[\bar{Z}_{n_1}(\mathfrak{F})]$ depends on $\mathfrak{F}$ only through $n_{\mathfrak{f}}$ and the number $e \equiv e(\mathfrak{F})$ of internal edges in $\mathfrak{F}$: writing $E_\eta \equiv n_\eta d$ and $E \equiv nd$, we calculate

$$\mathbb{E}[\bar{Z}_{n_1}(\mathfrak{F})] = \mathbf{c}(n_1, n_{\mathfrak{f}})\,\mathbf{g}_1(E, E_1, E_{\mathfrak{f}})\,\mathbf{g}_0(n_0, E_1)\,\bar{\mathbf{g}}_{\mathfrak{f}}(E, E_1, E_{\mathfrak{f}}, e) \equiv \mathbf{g}(n_1, n_{\mathfrak{f}}, e) \tag{13}$$

where $\mathbf{c}(n_1, n_{\mathfrak{f}}) \equiv \binom{n-n_{\mathfrak{f}}}{n_1}$ enumerates the number of partitions of $V \backslash \mathfrak{F}$ into subsets V_0, V_1 of sizes n_0, n_1, and the remaining factors calculate probabilities with respect to a *fixed* partition $V_0, V_1, V_{\mathfrak{f}}$ of V (with $V_{\mathfrak{f}}$ the vertices involved in $\mathfrak{F}$), as follows:

1. $\mathbf{g}_1(E, E_1, E_{\mathfrak{f}})$ is the probability all half-edges from V_1 are matched to half-edges from V_0;
2. $\mathbf{g}_0(n_0, E_1)$ is the probability each vertex in V_0 receives at least two half-edges from V_1, with respect to a uniformly random assignment of the E_1 incoming half-edges to the $n_0 d$ available from V_0;
3. $\bar{\mathbf{g}}_{\mathfrak{f}}(E, E_1, E_{\mathfrak{f}}, e)$ is the probability, conditioned on the valid assignment of E_1 edges between V_1 and V_0, that the e internal edges of $\mathfrak{F}$ are present in the graph, while the remaining $E_{\mathfrak{f}} - 2e$ half-edges from $V_{\mathfrak{f}}$ are matched to half-edges from V_0.

After the assignment of E_1 edges between V_1 and V_0, there are $E - 2E_1 = E_0 - E_1 + E_{\mathfrak{f}}$ half-edges remaining to be assigned. Recalling the notation (5), we have explicitly

$$\mathbf{g}_1(n_0, E_1) = \frac{(E - E_1 - E_{\mathfrak{f}})_{E_1}}{[E]_{E_1}} = \frac{(E_0)_{E_1}}{[E]_{E_1}};$$

$$\bar{\mathbf{g}}_{\mathfrak{f}}(E, E_1, E_{\mathfrak{f}}, e) = \frac{1}{[E - 2E_1]_e} \times \frac{(E_0 - E_1)_{E_{\mathfrak{f}} - 2e}}{[E_0 - E_1 + E_{\mathfrak{f}} - E_1]_{E_{\mathfrak{f}} - 2e}} = \frac{(E - 2E_1 - E_{\mathfrak{f}})_{E_{\mathfrak{f}} - 2e}}{[E - 2E_1]_{E_{\mathfrak{f}} - e}}. \tag{14}$$

Since $\mathbf{c}, \mathbf{g}_1, \mathbf{g}_0$ do not depend on e we find crudely that the function $\mathbf{g}$ of (13) satisfies

$$\mathbf{g}(n_1, n_{\mathfrak{f}}, e + \Delta) \leqslant [4/(nd)]^\Delta\,\mathbf{g}(n_1, n_{\mathfrak{f}}, e) \tag{15}$$

throughout the regime $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$, indicating that excess internal edges in $\mathfrak{F}$ are costly. On the other hand let us note that $\mathbf{g}$ is much less sensitive to small shifts in mass from V_0 to $V_{\mathfrak{f}}$ or vice versa: for $\delta \leqslant n/d$ we estimate

$$\left.\begin{array}{ll}
(\circ) & \mathbf{c}(n_1, n_{\mathfrak{f}} + \delta) = \mathbf{c}(n_1, n_{\mathfrak{f}})\,e^{O(\delta)}, \\
(\circ) & \mathbf{g}_1(E, E_1, E_{\mathfrak{f}} + d\delta) = \mathbf{g}_1(E, E_1, E_{\mathfrak{f}})\,d^{O(\delta)}, \\
(\circ) & \bar{\mathbf{g}}_{\mathfrak{f}}(E, E_1, E_{\mathfrak{f}} + d\delta, e) = \bar{\mathbf{g}}_{\mathfrak{f}}(E, E_1, E_{\mathfrak{f}}, e)\,e^{O(\delta)}, \\
(\star) & \mathbf{g}_0(n_0 - \delta, E_1) = \mathbf{g}_0(n_0, E_1)\,e^{O(\delta)}
\end{array}\right\}, \text{ thus } \frac{\mathbf{g}(n_1, n_{\mathfrak{f}} + \delta, e)}{\mathbf{g}(n_1, n_{\mathfrak{f}}, e)} = d^{O(\delta)}. \tag{16}$$

The three estimates marked $(\circ)$ follow straightforwardly from the explicit expressions given above; the last estimate $(\star)$ is deferred to a later section (Propn. 2.9).

Proof of Propn. 2.7. Let $k \geqslant 3$ be fixed, and let $\mathfrak{f} \equiv \mathfrak{f}(\mathfrak{F})$ denote the components of $\mathfrak{F}$ whose size is neither 2 nor k. Write $T \equiv T(\mathfrak{F}) \equiv |\mathfrak{F}\backslash\mathfrak{f}|$, $t \equiv t(\mathfrak{F})$ for the number of vertices in $\mathfrak{F}\backslash\mathfrak{f}$ in components of size 2, and $\ell \equiv \ell(\mathfrak{F})$ for the number of components of size k, so that

$T = t + \ell k$. Let $\mathfrak{q} \equiv \mathfrak{q}(\mathfrak{F})$ denote the number of internal edges among the components of size k. Any odd-size component must contain at least one cycle, so we define the number of "excess" edges among these components to be

$$A \equiv A(\mathfrak{F}) \equiv \mathfrak{q} - \ell(k - \mathbf{1}\{k \text{ even}\}) \geqslant 0. \tag{17}$$

Write $\Omega_{\ell,A}(n_{\mathfrak{f}}; \mathfrak{f})$ for the collection of all $\mathfrak{F}$ with $n_{\mathfrak{f}}$ vertices and with the given $\mathfrak{f}, \ell, A$.

Bounds for T even. Let $\Omega(n_{\mathfrak{f}}; \mathfrak{f}) \equiv \Omega_{0,0}(n_{\mathfrak{f}}; \mathfrak{f})$ denote the collection of all $\mathfrak{F}$ with $n_{\mathfrak{f}}$ vertices and the given $\mathfrak{f}$, such that $\mathfrak{F} \backslash \mathfrak{f}$ consists of T vertices matched up into isolated edges, where we assume T is even. Then all $\mathfrak{F} \in \Omega(n_{\mathfrak{f}}; \mathfrak{f})$ have the same number of internal edges e while all $\mathfrak{F} \in \Omega_{\ell,A}(n_{\mathfrak{f}}; \mathfrak{f})$ have the same number of internal edges

$$e' = e + \mathfrak{q} - \ell k/2 = e + A + \ell(k/2 - \mathbf{1}\{k \text{ even}\}).$$

We calculate $|\Omega(n_{\mathfrak{f}}; \mathfrak{f})| = \binom{n - |\mathfrak{f}|}{T} d^T (T - 1)!!$, while

$$|\Omega_{\ell,A}(n_{\mathfrak{f}}; \mathfrak{f})| \leqslant \binom{n - |\mathfrak{f}|}{T} \binom{T}{t} d^t (t - 1)!! \frac{(\ell k)!}{\ell!(k!)^\ell} d^{2\mathfrak{q}} \binom{\ell k(k - 1)/2}{\mathfrak{q}}.$$

Combining with (15) and simplifying gives

$$\bar{\mathbf{R}}_{\ell,A}^{n_1, n_{\mathfrak{f}}}(\mathfrak{f}) \equiv \frac{|\Omega_{\ell,A}(n_{\mathfrak{f}}; \mathfrak{f})|}{|\Omega(n_{\mathfrak{f}}; \mathfrak{f})|} \frac{\mathbf{g}(n_1, n_{\mathfrak{f}}, e')}{\mathbf{g}(n_1, n_{\mathfrak{f}}, e)} \leqslant \frac{(T/2)! 2^{\ell k/2}}{(t/2)!} \cdot \frac{1}{\ell!(k!)^\ell} \cdot \binom{\ell k(k - 1)/2}{\mathfrak{q}} \cdot \left(\frac{4d}{n}\right)^{\mathfrak{q} - \ell k/2}.$$

The first factor is $\leqslant n_{\mathfrak{f}}^{\ell k/2}$ since $t + \ell k = T \leqslant n_{\mathfrak{f}}$; the second factor is $\leqslant e^{O(k\ell)}/[\ell k^k]^\ell$; and the third factor is $\leqslant e^{O(\mathfrak{q})} k^{\mathfrak{q}}$ using $\mathfrak{q} \geqslant \ell(k - 1)$. Combining with (17) gives

$$\bar{\mathbf{R}}_{\ell,A}^{n_1, n_{\mathfrak{f}}}(\mathfrak{f}) \leqslant \left(\frac{(c_0 d\beta)^{k/2}}{\ell} \left(\frac{n}{dk}\right)^{\mathbf{1}\{k \text{ even}\}}\right)^\ell \left(\frac{c_0 dk}{n}\right)^A \quad (T \text{ even, } \mathfrak{m}(\mathfrak{f}) > 0). \tag{18}$$

for some constant c_0 uniform in d. We consider also the effect of reweighting by the number of matchings on $\mathfrak{F}$, which is non-zero only if $\mathfrak{m}(\mathfrak{f}) > 0$ and k is even. If $\mathfrak{F} \in \Omega(n_{\mathfrak{f}}; \mathfrak{f})$ then clearly $\mathfrak{m}(\mathfrak{F}) = \mathfrak{m}(\mathfrak{f})$ since the collection $\mathfrak{F} \backslash \mathfrak{f}$ of isolated edges has only one matching. If $\mathfrak{F} \in \Omega_{\ell,A}(n_{\mathfrak{f}}; \mathfrak{f})$ then the trivial bound $\mathfrak{m}(\mathfrak{F}) \leqslant \mathfrak{m}(\mathfrak{f}) 2^{\mathfrak{q}}$ implies

$$\mathbf{R}_{\ell,A}^{n_1, n_{\mathfrak{f}}}(\mathfrak{f}) \equiv \bar{\mathbf{R}}_{\ell,A}^{n_1, n_{\mathfrak{f}}}(\mathfrak{f}) \frac{\max\{\mathfrak{m}(\mathfrak{F}) : \mathfrak{F} \in \Omega_{\ell,A}(n_{\mathfrak{f}}; \mathfrak{f})\}}{\min\{\mathfrak{m}(\mathfrak{F}) : \mathfrak{F} \in \Omega(n_{\mathfrak{f}}; \mathfrak{f})\}} \leqslant \left(\frac{(c_0 d\beta)^{k/2} n}{dkl}\right)^\ell \left(\frac{c_0 dk}{n}\right)^A.$$

Bounds for T odd. For T odd we instead compare $\mathfrak{F}' \in \Omega_{\ell,A}(n_{\mathfrak{f}}; \mathfrak{f})$ with $\mathfrak{F} \in \Omega(n_{\mathfrak{f}} - 1; \mathfrak{f})|$, which has cardinality $\binom{n - |\mathfrak{f}|}{T - 1} d^{T-1} (T - 2)!!$. The difference in the number of internal edges between $\mathfrak{F}'$ and $\mathfrak{F}$ is $e' - e = \frac{1}{2}(\ell k + 1) + A$, so applying (15) and (16) and calculating as above gives

$$\bar{\mathbf{R}}_{\ell,A}^{n_1, n_{\mathfrak{f}}}(\mathfrak{f}) \equiv \frac{|\Omega_{\ell,A}(n_{\mathfrak{f}}; \mathfrak{f})|}{|\Omega(n_{\mathfrak{f}} - 1; \mathfrak{f})|} \frac{\mathbf{g}(n_1, n_{\mathfrak{f}}, e')}{\mathbf{g}(n_1, n_{\mathfrak{f}} - 1, e)} \leqslant d^{O(1)} \left(\frac{(c_0 d\beta)^{k/2}}{\ell}\right)^\ell \left(\frac{c_0 dk}{n}\right)^A \quad (T \text{ odd}). \tag{19}$$

Odd-size components. Recall (9) that $\bar{\mathbf{Z}}_{n_1, n_{\mathfrak{f}}}(O)$ denotes the contribution to $\bar{\mathbf{Z}}_{n_1, n_{\mathfrak{f}}}$ from frozen model configurations η with $\mathcal{O}(\mathfrak{F}) = O$ odd-size components in $\mathfrak{F} \equiv \mathfrak{F}(\eta)$. For $k \geqslant 3$ odd and $1 \leqslant \ell \leqslant O$ let $\bar{\mathbf{Z}}_{n_1, n_{\mathfrak{f}}}(O; L_k = \ell)$ denote the contribution to $\bar{\mathbf{Z}}_{n_1, n_{\mathfrak{f}}}(O)$ from configurations with ℓ $\mathfrak{f}$-components of size k. Recalling (12) and (13),

$$\mathbb{E}[\bar{\mathbf{Z}}_{n_1, n_{\mathfrak{f}}}(O; L_k = \ell)] = \sum_{\substack{\mathfrak{f}: \mathcal{O}(\mathfrak{f}) + \ell = O, \\ A \geqslant 0}} \sum_{\mathfrak{F} \in \Omega_{\ell,A}(n_{\mathfrak{f}}; \mathfrak{f})} \mathbb{E}[\bar{\mathbf{Z}}_{n_1}(\mathfrak{F})] = \sum_{\substack{\mathfrak{f}: \mathcal{O}(\mathfrak{f}) + \ell = O, \\ A \geqslant 0}} |\Omega_{\ell,A}(n_{\mathfrak{f}}; \mathfrak{f})| \, \mathbf{g}(n_1, n_{\mathfrak{f}}, e')$$

Applying (18) (ℓ even) and (19) (ℓ odd) gives, with $\bar{n}_{\mathtt{f}} \equiv n_{\mathtt{f}} - \mathbf{1}\{\ell \text{ odd}\}$,

$$\mathbb{E}[\bar{\boldsymbol{Z}}_{n_1,n_{\mathtt{f}}}(O; L_k = \ell)] \leqslant \mathbb{E}[\bar{\boldsymbol{Z}}_{n_1,\bar{n}_{\mathtt{f}}}(L - \ell; L_k = 0)] d^{O(1)}[(c_0 d\beta)^{k/2}\ell^{-1}]^\ell \quad (\ell \geqslant 1). \qquad (20)$$

For fixed $k \geqslant 3$ odd and $\ell \geqslant 1$, summing both sides of (20) over $n_1, n_{\mathtt{f}}, O$ with $O \geqslant \ell$ and $2n\alpha + O \leqslant 2n_1 + n_{\mathtt{f}}$ gives $\mathbb{E}[\bar{\boldsymbol{Z}}_{\geqslant n\alpha}; L_k = \ell] \leqslant d^{-k\ell/5}\mathbb{E}\bar{\boldsymbol{Z}}_{\geqslant n\alpha}$ (cf. (9)). Summing over $k \geqslant 3$ odd, $\ell \geqslant 1$ then gives $\mathbb{E}[^{\text{odd}}\bar{\boldsymbol{Z}}_{\geqslant n\alpha}] \leqslant d^{-1/2}\mathbb{E}\bar{\boldsymbol{Z}}_{\geqslant n\alpha}$, concluding the proof of (a).

Even-size components. Recall $^{\text{ev}}\bar{\boldsymbol{Z}}_{n_1,n_{\mathtt{f}}}$ denotes the contribution to $\bar{\boldsymbol{Z}}_{n_1,n_{\mathtt{f}}}$ from frozen model configurations η such that $\mathfrak{F}(\eta)$ has *no* odd-size components. If L_k denotes the number of size-k components in $\mathfrak{F}$, then

$$\mathbb{E}[^{\text{ev}}\bar{\boldsymbol{Z}}_{n_1,n_{\mathtt{f}}}; L_k = \ell] = \sum_{\mathfrak{f}:\mathcal{O}(\mathfrak{f})=0} \sum_{A \geqslant 0} |\Omega_{\ell,A}(n_{\mathtt{f}}; \mathfrak{f})| \, \mathbf{g}(n_1, n_{\mathtt{f}}, \boldsymbol{e}'). \qquad (21)$$

Let $\ell_k \equiv [(c_0 d\beta)^{k/8}n/(dk)] \vee 1$ as in the statement of the lemma: then for all $\ell \geqslant \ell_k$ we have from (18) that $\bar{\mathbf{R}}_{\ell,A}^{n_1,n_{\mathtt{f}}}(\mathfrak{f}) \leqslant (c_0 d\beta)^{(3/8)k\ell}(c_0 dk/n)^A$. Further

$$(c_0 d\beta)^{(3/8)k\ell} \leqslant \begin{cases} n^{-(3/16)\log d + O(1)} & \text{if } k\ell_k > \log n, \\ [d(\log n)/n]^{3\ell} & \text{if } k\ell_k \leqslant \log n; \end{cases}$$

so in any case we have crudely that $(c_0 d\beta)^{(3/8)k\ell} < n^{-5/2}$.[5] Thus from (21) and the definition of $\bar{\mathbf{R}}$ we have

$$\mathbb{E}[^{\text{ev}}\bar{\boldsymbol{Z}}_{n_1,n_{\mathtt{f}}}; L_k = \ell] \leqslant n^{-5/2}\mathbb{E}[^{\text{ev}}\bar{\boldsymbol{Z}}_{n_1,n_{\mathtt{f}}}; L_k = 0] \quad (k \geqslant 4 \text{ even}).$$

Summing over $\ell \geqslant 1$, $k \geqslant 4$, and $n_1, n_{\mathtt{f}}$ with $2n_1 + n_{\mathtt{f}} \geqslant 2n\alpha$ proves $\mathbb{E}[^{\text{ev}}\bar{\boldsymbol{Z}}_{\geqslant n\alpha}^\infty] \leqslant n^{-1/2}\mathbb{E}\bar{\boldsymbol{Z}}_{\geqslant n\alpha}$. Essentially the same argument (with k even, $\mathbf{R}$ in place of $\bar{\mathbf{R}}$) gives $\mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}^\infty \leqslant n^{-1/2}\mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}$, concluding the proof of (b). $\qquad \square$

Proof of Propn. 2.8. Fix $k \geqslant 4$ even and let $\ell < \ell_k$ with ℓ_k be as in Propn. 2.7; it suffices to consider $k \leqslant 20\log_d n$, otherwise $\ell_k = 1$ for all $\beta \leqslant \beta_{\max}$.

Let $\mathfrak{f}^* \equiv \mathfrak{f}^*(\mathfrak{F})$ denote the $\mathtt{f}$-components in $\mathfrak{F}$ which are not of size k. Let $\Omega_{\ell,r,A'}^\circ(n_{\mathtt{f}}; \mathfrak{f}^*)$ denote the collection of $\mathfrak{F}$ with $n_{\mathtt{f}}$ vertices and the given $\mathfrak{f}^*$, such that $\mathfrak{F}\backslash\mathfrak{f}^*$ consists of $\ell \equiv \frac{1}{k}(n_{\mathtt{f}} - |\mathfrak{f}^*|)$ components with k vertices each, of which $\ell - r$ components are trees with (unique) perfect matching, and the remaining r components are not trees and have a total of $\mathfrak{q}' \equiv rk + A'$ internal edges with $A' \geqslant 0$. Let $\Omega^{\text{tree}}(n_{\mathtt{f}}; \mathfrak{f}^*) \equiv \Omega_{0,0,0}^\circ(n_{\mathtt{f}}; \mathfrak{f}^*)$ denote the collection of $\mathfrak{F}$ with $n_{\mathtt{f}}$ vertices and the given $\mathfrak{f}^*$, such that each of the ℓ components of $\mathfrak{F}\backslash\mathfrak{f}^*$ is a tree with (unique) perfect matching on k vertices. Any $\mathfrak{F} \in \Omega^{\text{tree}}(n_{\mathtt{f}}; \mathfrak{f}^*)$ has $\boldsymbol{e}$ internal edges any $\mathfrak{F} \in \Omega_{\ell,r,A'}^\circ(n_{\mathtt{f}}; \mathfrak{f}^*)$ has $\boldsymbol{e}'$ internal edges with $\boldsymbol{e}' - \boldsymbol{e} = \mathfrak{q}' - r(k-1) = r + A'$.

Regarding $[dk]$ as a set of half-edges with the i-th half-edge incident to vertex $\lceil i/k \rceil$, let $\text{TPM}_{d,k}$ denote the number of ways that $2(k-1)$ elements of $[dk]$ can be used to form a graph on $[k]$ which is a spanning tree with perfect matching (meaning a perfect matching *in the tree*, not to be confused with a perfect matching of half-edges). Then, using (15),

$$\bar{\mathbf{S}}_{\ell,r,A'}^{n_1,n_{\mathtt{f}}}(\mathfrak{f}^*) \equiv \frac{|\Omega_{\ell,r,A'}^\circ(n_{\mathtt{f}}; \mathfrak{f}^*)|}{|\Omega^{\text{tree}}(n_{\mathtt{f}}; \mathfrak{f}^*)|} \frac{\mathbf{g}(n_1, n_{\mathtt{f}}, \boldsymbol{e}')}{\mathbf{g}(n_1, n_{\mathtt{f}}, \boldsymbol{e})} \leqslant \frac{\binom{\ell}{r}(\text{TPM}_{d,k})^s d^{2\mathfrak{q}'}\binom{rk(k-1)/2}{\mathfrak{q}'}}{(\text{TPM}_{d,k})^\ell}\left(\frac{4}{nd}\right)^{r+A'}.$$

We crudely lower bound the number of spanning trees by the number of line graphs (which clearly have a perfect matching as k is even), $\text{TPM}_{d,k} \geqslant \frac{1}{2}k!\frac{[d(d-1)]^k}{(d-1)^2} \geqslant e^{O(k)}k^k d^{2(k-1)}$ with the $O(k)$ uniform in d. Consequently there exists an absolute constant C not depending on d

---

[5]As always, the statement should be understood to apply for $d \geqslant d_0$, $n \geqslant n_0(d)$.

such that $\bar{\mathbf{S}}^{n_1,n_\mathtt{f}}_{\ell,r,A'}(\mathfrak{f}^*) \leqslant (C^k \ell d/n)^r (Cdk/n)^{A'}$. By the trivial inequality $\mathfrak{m}(\mathfrak{F}) \leqslant \mathfrak{m}(\mathfrak{f}^*)\, 2^{\mathfrak{q}'}$, the same bound holds (adjusting C as needed) after reweighting by the number of matchings:

$$\mathbf{S}^{n_1,n_\mathtt{f}}_{\ell,r,A'}(\mathfrak{f}^*) \equiv \bar{\mathbf{S}}^{n_1,n_\mathtt{f}}_{\ell,r,A'}(\mathfrak{f}^*) \frac{\max\{\mathfrak{m}(\mathfrak{F}) : \mathfrak{F} \in \Omega^\circ_{\ell,r,A'}(n_\mathtt{f}; \mathfrak{f}^*)\}}{\min\{\mathfrak{m}(\mathfrak{F}) : \mathfrak{F} \in \Omega^\text{tree}(n_\mathtt{f}; \mathfrak{f}^*)\}} \leqslant \left(\frac{C^k \ell d}{n}\right)^r \left(\frac{Cdk}{n}\right)^{A'}.$$

Recalling $k \leqslant 20 \log_d n$ and $\ell < \ell_k \equiv [(c_0 d\beta)^{k/8} n/(dk)] \vee 1$ gives

$$\bar{\mathbf{S}}^{n_1,n_\mathtt{f}}_{\ell,r,A'}(\mathfrak{f}^*), \mathbf{S}^{n_1,n_\mathtt{f}}_{\ell,r,A'}(\mathfrak{f}^*) \leqslant (cd\beta)^{kr/9} [d(\log n)/n]^{A'}.$$

Recalling (12) and arguing similarly as in the proof of Propn. 2.7 we have

$$\mathbb{E}[{}^\text{ev}\bar{\boldsymbol{Z}}^\circ_{n_1,n_\mathtt{f}}] \leqslant \sum_{\substack{k \geqslant 4 \text{ even}, \, \mathfrak{f}^*:\mathcal{O}(\mathfrak{f}^*)=0, \\ 1 \leqslant r < \ell_k \qquad A' \geqslant 0}} \sum |\Omega^\circ_{\ell,r,A'}(n_\mathtt{f}; \mathfrak{f}^*)| \, \mathbf{g}(n_1, n_\mathtt{f}, \boldsymbol{e}')$$

$$\lesssim \sum_{\substack{k \geqslant 4 \text{ even}, \\ 1 \leqslant r < \ell_k}} (cd\beta)^{kr/9} \sum_{\mathfrak{f}^*:\mathcal{O}(\mathfrak{f}^*)=0} |\Omega^\text{tree}(n_\mathtt{f}; \mathfrak{f}^*)| \, \mathbf{g}(n_1, n_\mathtt{f}, \boldsymbol{e}) \leqslant d^{-1/5} \mathbb{E}[{}^\text{ev}\bar{\boldsymbol{Z}}_{n_1,n_\mathtt{f}}],$$

and summing over α_1, β proves $\mathbb{E}[{}^\text{ev}\bar{\boldsymbol{Z}}^\circ_{\geqslant n\alpha}] \leqslant d^{-1/5} \mathbb{E}[{}^\text{ev}\bar{\boldsymbol{Z}}_{\geqslant n\alpha}]$ which implies (a). Similarly

$$\mathbb{E}\boldsymbol{Z}^\circ_{\geqslant n\alpha}(A) \leqslant \sum_{\substack{k \geqslant 4 \text{ even}, \, \mathfrak{f}:\mathcal{O}(\mathfrak{f}^*)=0, \\ 1 \leqslant r < \ell_k \qquad A' \geqslant A}} \sum \mathfrak{m}(\mathfrak{F}) \, |\Omega^\circ_{\ell,r,A'}(n_\mathtt{f}; \mathfrak{f}^*)| \, \mathbf{g}(n_1, n_\mathtt{f}, \boldsymbol{e}') \leqslant [d(\log n)/n]^A \mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha},$$

proving (b). $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

2.4. **Estimates on forcing constraints.** In this subsection we estimate the probability cost of the constraint that each 0-vertex is forced by at least two neighboring 1-vertices.

Slightly more generally, for $\boldsymbol{l}$ a fixed positive integer, write $\mathbf{g}^l_0(n, E_1)$ for the probability, with respect to a uniformly random assignment of E_1 half-edges to n vertices of degree d, that each vertex receives at least $\boldsymbol{l}$ incoming half-edges (thus the function $\mathbf{g}_0$ of (13) is simply $\mathbf{g}^l_0$ with $\boldsymbol{l} = 2$). Let $X, X^1, \ldots, X^n$ be i.i.d. $\text{Bin}(d, \vartheta)$ random variables, with joint law $\mathbb{P}_\vartheta$: then

$$\mathbf{g}^l_0(n, E_1) = \mathbb{P}_\theta(X^i \geqslant \boldsymbol{l} \,\forall 1 \leqslant i \leqslant n \mid \textstyle\sum_{i=1}^n X^i = E_1).$$

where $0 < \vartheta < 1$ may be arbitrarily chosen.

We also give estimates on a two-dimensional analogue $\mathbf{g}^l_{00}$ which will be used in our second moment analysis §4: let $\mathscr{X} \equiv \{0,1\}^2 \backslash \{00\}$. If $\theta \equiv (\theta_\omega)_{\omega \in \mathscr{X}}$ is a positive vector with sum less than 1, then it defines a positive probability measure on $\{0,1\}^2$; and we take ϖ to be a $\{0,1\}^2$-valued random variable with law given by this measure. Take $(\varpi_j)_{j \geqslant 1}$ independent random variables identically distributed as ϖ, and define the multinomial random variable

$$X \equiv (X_\omega)_{\omega \in \mathscr{X}} \equiv (\textstyle\sum_{j=1}^d \mathbf{1}\{\varpi_j = \varpi\})_{\omega \in \mathscr{X}}. \tag{22}$$

Though we always use θ and X to denote three-dimensional vectors, we shall also write $\theta_{00} \equiv 1 - \sum_{\omega \in \mathscr{X}} \theta_\omega$ and $X_{00} \equiv d - \sum_{\omega \in \mathscr{X}} X_\omega$. Let X^i $(1 \leqslant i \leqslant n)$ be independent random variables identically distributed as X, and write $\mathbb{P}_\theta$ for their (joint) law. Then, with $\Omega \equiv \{x : x_{1\cdot} \wedge x_{\cdot 1} < \boldsymbol{l}\}$, we consider

$$\mathbf{g}^l_{00}(n, \underline{E}) \equiv \mathbb{P}_\theta(X^i \notin \Omega \,\forall 1 \leqslant i \leqslant n \mid \textstyle\sum_{i=1}^n X^i = \underline{E}).$$

Proposition 2.9. *Let ε_0 be a small constant uniform in d, and suppose*

$$\varepsilon_0 (\log d)/d \leqslant \{\zeta_{1\cdot}, \zeta_{\cdot 1}\} \leqslant 10(\log d)/d.$$

(a) Let a, y be defined by $\zeta_{1\cdot} \equiv a(\log d)/d$ and $\zeta_{1\cdot} \wedge \zeta_{\cdot 1} \equiv y(\log d)/d)$. Then

$$\mathbf{g}_0^l(n, nd\zeta_{1\cdot}) = \exp\{O(nd^{-a}(\log d)^l)\},$$
$$\mathbf{g}_{00}^l(n, nd\zeta) = \exp\{O(nd^{-y}(\log d)^l)\}.$$

(b) If ξ is another vector with $|\xi_\omega/\zeta_\omega - 1| \leqslant e^{-1}$ for all $\omega \in \mathscr{X}$, then

$$\mathbf{g}_0^l(n, nd\xi_{1\cdot}) = \mathbf{g}_0^l(n, nd\zeta_{1\cdot}) \exp\{O(nd|\xi[1\cdot] - \zeta_{1\cdot}|)\},$$
$$\mathbf{g}_{00}^l(n, nd\xi) = \mathbf{g}_{00}^l(n, nd\zeta) \exp\{O(nd\|\xi - \zeta\|_1)\}.$$

Lemma 2.10. *For the multinomial random variable X defined by* (22), *consider the cumulant generating function $\Lambda_\theta(\gamma) \equiv \log \mathbb{E}_\theta[e^{\langle \gamma, X \rangle} \mid X \notin \Omega]$, defined for $\gamma \in \mathbb{R}^3$. For positive vectors θ, ζ in the regime*

$$\varepsilon_0(\log d)/d \leqslant \{\theta_{1\cdot}, \theta_{\cdot 1}, \zeta_{1\cdot}, \zeta_{\cdot 1}\} \leqslant 10(\log d)/d,$$

there exists a unique $\gamma \in \mathbb{R}^3$ with $\Lambda_\theta'(\gamma) = d\zeta$; and it satisfies

$$\left| \gamma_\omega - \log \frac{\theta_{00}\zeta_\omega}{\zeta_{00}\theta_\omega} \right| \lesssim \frac{(\log d)^{l-1}}{d^y} \quad \text{for all } \omega \in \mathscr{X}.$$

Proof. If γ exists, then it is clearly unique by the strict convexity of the cumulant generating function on $\mathbb{R}^3$. To see existence, first note that

$$\Lambda_\theta'(\gamma) = \mathbb{E}_\mu[X \mid X \notin \Omega] \quad \text{for } \mu_\omega = \frac{\theta_\omega e^{\gamma_\omega}}{z_{\theta,\gamma}} \equiv \frac{\theta_\omega e^{\gamma_\omega}}{\theta_{00} + \sum_{\pi \in \mathscr{X}} \theta_\pi e^{\gamma_\pi}} \quad (\omega \in \mathscr{X}).$$

In particular, each entry of $\Lambda_\theta'(\gamma)$ lies between 0 and d. Consider the ε-perturbed ($\varepsilon > 0$ small) cumulant generating function $\Lambda_\theta(\gamma) + \frac{1}{2}\varepsilon\|\gamma\|^2$, corresponding to the random variable $X + \varepsilon^{1/2}Y$ for Y a standard Gaussian in $\mathbb{R}^3$. For any fixed $\varepsilon > 0$ this is a smooth convex function on $\mathbb{R}^3$, with gradient $\Lambda_\theta'(\gamma) + \varepsilon\gamma$ tending in norm to ∞ as $\|\gamma\| \to \infty$. It follows from Rockafellar's theorem (see e.g. [14, Lem. 2.3.12]) that there exists a unique $\gamma_\varepsilon \equiv (\gamma_{\varepsilon,\omega})_{\omega \in \mathscr{X}}$ such that $\Lambda_\theta'(\gamma_\varepsilon) + \varepsilon\gamma_\varepsilon = d\zeta$. We shall show by some rough estimates that γ_ε must remain within a compact region as ε tends to zero. We claim first that $\theta_{00} \geqslant \theta_\omega e^{\gamma_{\varepsilon,\omega}}$ for all $\omega \in \mathscr{X}$: if not, then for some $\omega \in \mathscr{X}$ we must have $z_{\theta,\gamma} \leqslant 4\theta_\omega e^{\gamma_{\varepsilon,\omega}}$, implying (in the stated regime of θ, ζ) that

$$\gamma_{\varepsilon,\omega} = \varepsilon^{-1}[d\zeta_\omega - \partial_\omega \Lambda_\theta(\gamma_\varepsilon)] \leqslant \varepsilon^{-1}[d\zeta_\omega - d/4] \ll 0,$$

contradicting the hypothesis that $\theta_{00} < \theta_\omega e^{\gamma_{\varepsilon,\omega}}$. Thus $\limsup_{\varepsilon\downarrow 0} \gamma_{\varepsilon,\omega}$ must be finite for each $\omega \in \mathscr{X}$. In the other direction, the trivial bound $z_{\theta,\gamma} \geqslant \theta_{00}$ gives

$$\gamma_{\varepsilon,\omega} = \varepsilon^{-1}[d\zeta_\omega - \partial_\omega \Lambda_\theta(\gamma_\varepsilon)] \geqslant \varepsilon^{-1}d\zeta_\omega(1 - e^{\gamma_{\varepsilon,\omega}}/\theta_{00}),$$

so clearly $\liminf_{\varepsilon\downarrow 0} \gamma_{\varepsilon,\omega}$ must also be finite. It follows by an easy compactness argument that γ_ε converges in the limit $\varepsilon \downarrow 0$ to the required solution γ of $\Lambda_\theta'(\gamma) = d\zeta$.

To control the norm of the solution γ of $\Lambda_\theta'(\gamma) = d\zeta$, we shall argue that for ζ in the stated regime, μ_ω is close to ζ_ω for each ω. The bound for $\omega = \mathtt{11}$ is easiest: consider X as the d-th step of the random walk

$$X_t \equiv (X_{t,\omega})_{\omega \in \mathscr{X}} \equiv \left(\sum_{j=1}^t \mathbf{1}\{\varpi_j = \omega\}\right)_{\omega \in \mathscr{X}}.$$

Define the stopping time $\tau \equiv \inf\{t \geqslant 0 : X_t \notin \Omega\}$, so $\{X \notin \Omega\} = \{\tau \leqslant d\}$. Since $X_{\tau,\mathtt{11}} \leqslant \boldsymbol{l}$, applying the Markov property gives

$$\begin{aligned}
d\zeta_{\mathtt{11}} &= \mathbb{E}_\mu[X_{\mathtt{11}} \mid X \notin \Omega] \leqslant \boldsymbol{l} + \mathbb{E}_\mu[X_{d,\mathtt{11}} - X_{\tau,\mathtt{11}} \mid \tau \leqslant d] \\
&\leqslant \boldsymbol{l} + \mathbb{E}_\mu[X_{d-\tau,\mathtt{11}}] \leqslant \boldsymbol{l} + d\mu_{\mathtt{11}}.
\end{aligned} \tag{23}$$

Next observe that for the multinomial random variable X defined by (22), for any $\omega \ne \pi$ the conditional expectation $\mathbb{E}_\theta[X_\omega \mid X_\pi = k] = (d-k)\theta_\omega/(1-\theta_\pi)$ is decreasing with k, so

$$\mathbb{E}_\theta[X_\omega \mid X_\pi \geqslant \ell] = \sum_{k \geqslant \ell} \mathbb{P}(X_\pi = k \mid X_\pi \geqslant \ell) \mathbb{E}_\theta[X_\omega \mid X_\pi = k] \leqslant \mathbb{E}_\theta X_\omega = d\theta_\omega. \qquad (24)$$

Define the stopping times $\tau_{1\cdot} \equiv \inf\{t \geqslant 0 : X_{1\cdot}(t) \geqslant \boldsymbol{l}\}$ and symmetrically $\tau_{\cdot 1}$; then $\tau = \tau_{1\cdot} \vee \tau_{\cdot 1}$. Since $X_{\tau_{1\cdot},10} \leqslant \boldsymbol{l}$,

$$d\zeta_{10} = \mathbb{E}_\mu[X_{10} \mid X \notin \Omega] = \frac{\mathbb{E}_\mu[X_{10}\mathbf{1}\{\tau \leqslant d\}]}{\mathbb{P}_\mu(\tau \leqslant d)} \leqslant \boldsymbol{l} + \frac{\mathbb{E}_\mu[(X_{10} - X_{\tau_{1\cdot},10})\mathbf{1}\{\tau \leqslant d\}]}{\mathbb{P}_\mu(\tau \leqslant d)}$$

$$= \boldsymbol{l} + \sum_{k<d} \sum_x \frac{\mathbb{P}_\mu(\tau_{1\cdot} = k, X_{\tau_{1\cdot}} = x, \widetilde{X}_{d-k,\cdot 1} \geqslant \boldsymbol{l} - x_{\cdot 1})}{\mathbb{P}_\mu(\tau \leqslant d)} \mathbb{E}_\mu[\widetilde{X}_{d-k,10} \mid \widetilde{X}_{d-k,\cdot 1} \geqslant \boldsymbol{l} - x_{\cdot 1}],$$

where $(\widetilde{X}_t)_{t \geqslant 0}$ is an independent realization of the random walk X. Maximizing over all possible k, x and applying 24 gives

$$d\zeta_{10} \leqslant \boldsymbol{l} + \max_{k<d, \ell \leqslant l} \mathbb{E}_\mu[\widetilde{X}_{d-k,10} \mid \widetilde{X}_{d-k,\cdot 1} \geqslant \ell] \leqslant \boldsymbol{l} + d\mu_{10},$$
$$\text{and symmetrically } d\zeta_{01} \leqslant \boldsymbol{l} + d\mu_{01}.$$

Thus for ζ in the stated regime we conclude that $d\mu_\omega \geqslant d\zeta_\omega - O(1)$ for all $\omega \in \mathscr{X}$, therefore $d\mu_{00} \leqslant d\zeta_{00} + O(1)$.

From these bounds we see clearly that, with y defined by $\zeta_{1\cdot} \wedge \zeta_{\cdot 1} \equiv y(\log d)/d$,

$$\mathbb{P}_\mu(X \in \Omega) \leqslant \mathbb{P}_\mu(X_{1\cdot} < \boldsymbol{l}) + \mathbb{P}_\mu(X_{\cdot 1} < \boldsymbol{l}) \lesssim d^{-y}(\log d)^{\boldsymbol{l}-1}.$$

Write p_μ for the law of $X \equiv X_d$ and q_μ for the law of X_{d-1}, and observe that $x_\omega p_\mu(x) = d\mu_\omega q_\mu(x - \mathbf{1}_\omega)$ (with both sides zero for $x_\omega = 0$). We then calculate

$$d\zeta_\omega = \frac{\mathbb{E}_\mu X_\omega - d\mu_\omega \sum_{x \in \Omega} q_\mu(x - \mathbf{1}_\omega)}{1 - \mathbb{P}_\mu(X \in \Omega)} = d\mu_\omega \left[1 + O(d^{-y}(\log d)^{\boldsymbol{l}-1})\right] \qquad (25)$$

for all $\omega \in \{0,1\}^2$. Summing over $\omega \in \mathscr{X}$ gives

$$z_{\theta,\gamma} = (\theta_{00}/\zeta_{00})(\zeta_{00}/\mu_{00}) = (\theta_{00}/\zeta_{00})[1 + O(d^{-1-y}(\log d)^{\boldsymbol{l}})], \qquad (26)$$

implying the stated bound on $\gamma_\omega = \log(z_{\theta,\gamma}\mu_\omega/\theta_\omega)$, $\omega \in \mathscr{X}$. $\qquad \square$

It follows (see e.g. [14, Lem. 2.3.9]) that with θ, ζ in the stated regime, the Fenchel-Legendre transform $\Lambda_\theta^*(d\zeta) \equiv \sup_\gamma[\langle \gamma, d\zeta \rangle - \Lambda_\theta(\gamma)]$ of the cumulant generating function is given by

$$\Lambda_\theta^*(d\zeta) = \langle \gamma, d\zeta \rangle - \Lambda(\gamma), \quad \gamma \text{ the solution of } \Lambda_\theta'(\gamma) = d\zeta. \qquad (27)$$

Since Λ_θ is strictly convex, we find by implicit differentiation that γ is differentiable with respect to ζ (in the stated regime). We then see from (27) that Λ_θ^* is differentiable with respect to ζ, with gradient $(\Lambda_\theta^*)'(d\zeta) = \gamma$.

Proof of Propn. 2.9. Let ν denote the normalized empirical measure of $X^1, \ldots, X^n$. The $\mathbb{P}_\theta$-probability of realizing ν is $\binom{n}{n\nu} \prod_x p_\theta(x)^{n\nu(x)}$ where p_θ denotes the law of X^1. Denote

the empirical mean $\bar{\nu} \equiv \sum_x x\nu(x)$; then $\mathbf{g}_{\mathtt{oo}}^{\mathit{l}}(n, nd\zeta)$ can be expressed as a constrained optimization over empirical measures ν with $\nu(\Omega) = 0$ and $\bar{\nu} = d\zeta$:

$$\frac{\mathbf{g}_{\mathtt{oo}}^{\mathit{l}}(n, nd\zeta)}{[1 - \mathbb{P}_\theta(\Omega)]^n} = \frac{\mathbb{P}_\theta(\sum_{i=1}^n X^i = nd\zeta, X^i \notin \Omega \,\forall i)}{\mathbb{P}_\theta(X^i \notin \Omega \,\forall i)\,\mathbb{P}_\theta(\sum_{i=1}^n X^i = nd\zeta)}$$

$$= \frac{e^{-n\langle \gamma, d\zeta \rangle} \sum_{\nu:\nu(\Omega)=0,\bar{\nu}=d\zeta} \binom{n}{n\nu} \prod_x [p_\theta(x)e^{\langle \gamma, x \rangle}]^{n\nu(x)}}{[1 - \mathbb{P}_\theta(\Omega)]^n \mathbb{P}_\theta(\sum_{i=1}^n X^i = nd\zeta)},$$

where we introduced a Lagrangian term which clearly has no effect on the constrained space. In the denominator, Stirling's approximation gives

$$\mathbb{P}_\theta(\textstyle\sum_{i=1}^n X^i = nd\zeta) \asymp_d n^{-3/2} \exp\{-ndH(\zeta \,|\, \theta)\}$$

where $H(\zeta \,|\, \theta) \equiv \sum_{\omega \in \{0,1\}^2} \zeta_\omega \log(\zeta_\omega/\theta_\omega)$ denotes the relative entropy between the probability measures induced by ζ and θ on $\{0,1\}^2$. In the numerator, if we take γ to be the unique solution γ given by Lem. 2.10 of $\Lambda'_\theta(\gamma) = d\zeta$, then we see that the sum over $\{\nu(\Omega) = 0, \bar{\nu} = d\zeta\}$ equals (up to constants depending on d) the sum over $\{\nu(\Omega) = 0\}$ divided by $n^{3/2}$: thus, recalling (27),

$$\mathbf{g}_{\mathtt{oo}}^{\mathit{l}}(n, nd\zeta) \asymp_d [1 - \mathbb{P}_\theta(\Omega)]^n \, \exp\{-n\Lambda_\theta^*(d\zeta) + ndH(\zeta \,|\, \theta)\}.$$

We can estimate this easily by taking $\theta = \zeta$. Since $e^{\langle \gamma, x \rangle} p_\theta(x) = (z_{\theta,\gamma})^d p_\mu(x)$, it follows from (25) and (26) that

$$\Lambda_\zeta(\gamma) = d \log z_{\zeta,\gamma} + \log \frac{\mathbb{P}_\mu(X \notin \Omega)}{\mathbb{P}_\zeta(X \notin \Omega)} = O(d^{-y}(\log d)^{\mathit{l}}) \quad \text{for } \Lambda'_\zeta(\gamma) = d\zeta.$$

Combining with (27) and the bound of Lem. 2.10 gives $\Lambda_\zeta^*(d\zeta) = O(d^{-y}(\log d)^{\mathit{l}})$. To compare $\mathbf{g}_{\mathtt{oo}}^{\mathit{l}}(n, nd\zeta)$ with $\mathbf{g}_{\mathtt{oo}}^{\mathit{l}}(n, nd\xi)$ where $\xi_\omega = \zeta_\omega[1 + O(\delta/n)]$, we fix $\theta = \zeta$. Recalling $(\Lambda_\theta^*)'(d\zeta) = \gamma$, the bound of Lem. 2.10 gives

$$|n[\Lambda_\zeta^*(d\xi) - \Lambda_\zeta^*(d\zeta)]| \lesssim nd^{-y}(\log d)^{\mathit{l}-1} \cdot \|d\xi - d\zeta\|_1.$$

For $|\xi_\omega/\zeta_\omega - 1| \leq e^{-1}$ it is straightforward to estimate $ndH(\xi \,|\, \zeta) \lesssim nd\|\xi - \zeta\|_1$, and combining these estimates concludes the proof of the proposition. $\qquad \square$

3. First moment of frozen model

In this section we identify the leading exponential order $^\star\mathbf{\Phi}(\alpha) = \lim_{n\to\infty} n^{-1} \log \mathbb{E}\mathbf{Z}_{n\alpha}$ of the first moment of the (truncated) frozen model partition function (8). The random d-regular graph $\mathcal{G}_{n,d}$ converges locally weakly (in the sense of [7, 2]) to the infinite d-regular tree T_d. Our calculation is based on a variational principle which relates the exponent $^\star\mathbf{\Phi}(\alpha)$ to a certain class of Gibbs measures for the frozen model on T_d which are characterized by fixed-point recursions. In fact the recursions can have multiple solutions, and much of the work goes into identifying (via *a priori* estimates) the unique fixed point which gives rise to $^\star\mathbf{\Phi}(\alpha)$. We begin by introducing the Gibbs measures which will be relevant for the variational principle.

3.1. **Frozen model tree recursions.** We shall specify a Gibbs measure ν on T_d by defining a consistent family of finite-dimensional distributions ν_t on the depth-t subtrees $T_d(t)$. A typical manner of specifying ν_t is to specify a law on some boundary conditions at depth t, and then to define ν_t as the law of the configuration on $T_d(t)$ given the (random) boundary conditions.

In our setting some difficulty is imposed by the fact that the frozen model is not a factor model (or Markov random field) in the conventional sense that $\eta|_A$ and $\eta|_B$ are conditionally independent given the configuration $\eta|_C$ on any subset C separating A from B — in particular, given the spins at level t of T_d, whether a vertex at level $t-1$ is required to take spin 1 depends on whether its neighboring 0's in level t are forced or not by 1's in level $t+1$. Also, the frozen model spins η do not encode the matching $\underline{m}$ on the f-vertices.

We shall instead specify Gibbs measures for the frozen model via a message-passing system, as follows. Each vertex v will send a message $\sigma_{v \to w}$ to each neighboring vertex $w \in \partial v$ which represents the "state of v ignoring w"; this will be a function $\dot{m}_{d-1}$ of the other $d-1$ incoming messages $(\sigma_{u \to v})_{u \in \partial v \setminus w}$. The actual state η_v of v is then a function $\dot{m}_d$ of all its incoming messages $\underline{\dot{\sigma}}_{\partial v \to v} \equiv (\sigma_{u \to v})_{u \in \partial v}$. If on the boundary of $T_d(t)$ we are given a vector $\eta^\uparrow \equiv (\sigma_{v \to w})_{v,w}$ (for v at level t, w the parent of v), then there is a *unique* completion of $\eta^\uparrow$ to a (bi-directional) message configuration on $T_d(t)$: iterating $\dot{m}$ gives all the messages upwards in the direction of the root, and once those are known we can recurse back down to determine the messages in the opposite direction. The measure ν_t can then be specified by giving the law of the boundary messages $\eta^\uparrow$ (possibly with an additional reweighting): our choice will be to take $\eta^\uparrow$ i.i.d. according to a law q; consistency of the family $(\nu_t)_t$ will then amount to fixed-point relations on q.

The message-passing rule for our frozen model is

$$\dot{m}_D : \{0,1,f\}^D \to \{0,1,f\}, \quad \dot{m}_D(\underline{\eta}) = \begin{cases} 1, & |\{i : \eta_i = 1\}| = 0, \\ f, & |\{i : \eta_i = 1\}| = 1, \\ 0, & |\{i : \eta_i = 1\}| \geq 2. \end{cases}$$

See Fig. 1 for an illustration: in each panel of the figure, the entire configuration of messages and vertex spins is uniquely determined by the messages incoming at the boundary of the subtree depicted.

We then define the ν_t-probability of a configuration on $T_d(t)$ to be (up to a normalizing factor) the product measure $q(\eta^\uparrow) \equiv \prod_i q(\eta_i^\uparrow)$ tilted by the factor $\lambda > 1$ raised to the number of *upward* (in the direction away from the boundary) 1-messages from vertices in $T_d(t-1)$, *including* $\eta_o^\uparrow \equiv \dot{m}_d((\sigma_{v \to o})_{v \in \partial o}) \equiv \eta_o$. Thus

$$\nu_1(\eta_o = 1) = \frac{\lambda(1-q_1)^d}{(\lambda-1)(1-q_1)^d + 1}, \quad \nu_1(\eta_o = f) = \frac{dq_1(1-q_1)^{d-2}}{(\lambda-1)(1-q_1)^d + 1},$$

with the remaining probability going to $\eta_o = 0$. The ν_t are consistent if and only if q satisfies the *frozen model recursions*

$$q_1 = \frac{\lambda(1-q_1)^{d-1}}{(\lambda-1)(1-q_1)^{d-1} + 1}, \quad q_f = \frac{(d-1)q_1(1-q_1)^{d-2}}{(\lambda-1)(1-q_1)^{d-1} + 1} = \frac{(d-1)q_1^2}{\lambda(1-q_1)}, \tag{28}$$

and $q_0 = 1 - q_1 - q_f$. (In contrast, the hard-core model with fugacity parameter λ, which allows for 0-vertices with any number of neighboring 1-vertices, has recursion $q = \lambda(1 - q)^{d-1}/[\lambda(1-q)^{d-1} + 1]$ on T_d.)

(A) Root spin $\eta_o = 1$ (B) Root spin $\eta_o = \mathtt{f}$

(C) Root spin $\eta_o = 0$ (D) Root spin $\eta_o = 0$
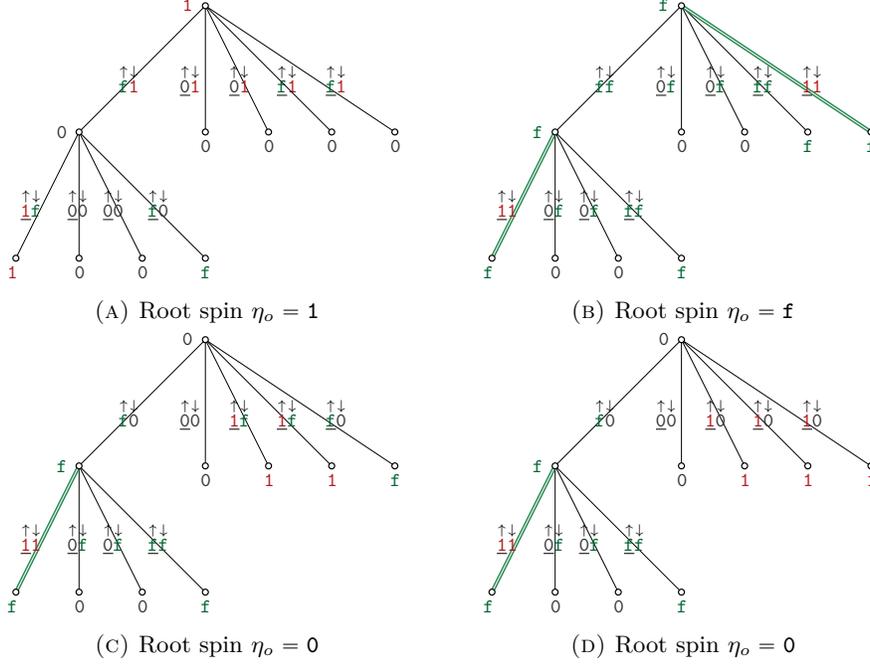
FIGURE 1. Completion of incoming boundary messages (underlined)

A solution $q \equiv (q_1, q_{\mathtt{f}}, q_0)$ of the frozen model recursions (28) corresponds to a root q_1 of the function $f(q) = (1-q)^{d-1}(\lambda + q - \lambda q) - q$ which, for $0 \leqslant q \leqslant 1$ and $\lambda > 1$, is decreasing in q and increasing in λ. Therefore (28) is solved by a unique measure q with q_1 increasing in λ. Note that if $q_1 = y(\log d)/d$ with $y \asymp 1$ then

$$\lambda = \frac{q_1[1 - (1-q_1)^{d-1}]}{(1-q_1)^d} = e^{o_d(1)}d^{y-1}y \log d. \tag{29}$$

In the following, to emphasize the dependence on λ we sometimes write $\nu \equiv \nu^\lambda$, $q \equiv q^\lambda$.

3.2. **Auxiliary model and Bethe variational principle.** On the tree T_d, the frozen model configuration $\overline{\eta} = (\eta, \underline{m})$ can be uniquely recovered from the configuration $\underline{\sigma}$ of messages on all the directed edges: each vertex spin η_v is determined by applying $\dot{\mathtt{m}}_d$ to the incoming messages, and $\underline{m}$ consists of the edges (uv) with $\sigma_{u \to v} = \sigma_{v \to u} = 1$. We refer to $\underline{\sigma}$ as the *auxiliary* configuration, and we now observe that we can define a model on auxiliary configurations on d-regular graphs which is in bijection with the frozen model but has the advantage of being a *factor model* in a relatively simple sense.

Recall our definition of the random d-regular graph $\mathcal{G}_{n,d}$ as given by a uniformly random matching M on half-edges $H = [nd]$ incident to vertices $V = [n]$. For convenience, we now bisect each edge in $\mathcal{G}_{n,d}$ by a new *clause* vertex a, and refer to the resulting graph as the $(d, 2)$-*regular bipartite factor graph*: this graph has vertex set $V \cup F$ with bipartition into the set V of *variables* (vertices in the original graph) and the set F of *clauses* (edges in the original graph). The new graph has edge set E where $(ia) \in E$ ($i \in V$, $a \in F$) indicates that in the original graph vertex i is incident to edge a. These edges are *labelled*, thus the new bipartite graph is simply equivalent to the original graph $G = (V, M, H)$ together with a labelling on M as well as an ordering within each pair formed by M: this contributes a factor $(nd/2)!2^{nd/2}$ to the enumeration but clearly the problem remains unchanged, so we shall continue to use the notation $\mathcal{G}_{n,d}$ for the $(d, 2)$-regular bipartite factor graph.

We shall define a model on auxiliary configurations $\underline{\sigma}$ on the bipartite factor graph where $\sigma_{u\to a} = \sigma_{a\to v} \equiv \sigma_{u\to v}$ for variables $u, v \in V$ joined by clause $a \in F$ (that is, the clause acts trivially by passing on the incoming message). The spins of the auxiliary model are the bidirectional messages $\sigma_{(va)} \equiv \sigma_{(av)} \equiv (\sigma_{v\to a}, \sigma_{a\to v})$, taking values in the alphabet $\mathscr{M} \equiv \{0, 1, \mathtt{f}\}^2$. Write $\dot{\underline{\sigma}}_v$ for the d-tuple of spins on the edges incident to variable $v \in V$, and write $\hat{\underline{\sigma}}_a$ for the pair of spins on the edges incident to clause $a \in F$. We hereafter refer to the *intensity* $\mathbf{i}(\underline{\sigma})$ of $\underline{\sigma}$ to mean the intensity of its corresponding frozen model configuration, $\mathbf{i}(\underline{\sigma}) \equiv |\{v : \dot{\mathtt{m}}_d(\underline{\sigma}_v) = \mathtt{1}\}| + \frac{1}{2}|\{v : \dot{\mathtt{m}}_d(\underline{\sigma}_v) = \mathtt{f}\}|$.

The weight of configuration $\underline{\sigma} \in \mathscr{M}^E$ under the *auxiliary model* is given by

$$\Psi^\lambda(\underline{\sigma}) \equiv \prod_{v\in V} \dot{\varphi}^\lambda(\dot{\underline{\sigma}}_v) \prod_{a\in F} \hat{\varphi}^\lambda(\hat{\underline{\sigma}}_a) \equiv \lambda^{\mathbf{i}(\underline{\sigma})} \prod_{v\in V} \dot{\varphi}(\dot{\underline{\sigma}}_v) \prod_{a\in F} \hat{\varphi}(\hat{\underline{\sigma}}_a) \equiv \lambda^{\mathbf{i}(\underline{\sigma})} \Psi(\underline{\sigma})$$

for factors defined as follows: in the unweighted version, the variable factor weight $\dot{\varphi}(\dot{\underline{\sigma}}_v)$ is simply the indicator that each outgoing message $\sigma_{v\to a}$ is determined by the message-passing rule $\dot{\mathtt{m}}_{d-1}$ from the incoming messages $\sigma_{b\to v}$, $b \in \partial v \backslash a$. We define the clause factor $\hat{\varphi}$ analogously, so (recalling that σ denotes pair consisting of the variable-to-clause message followed by the clause-to-variable message) $\hat{\varphi}(\sigma, \acute{\sigma}) = \mathbf{1}\{\sigma = \mathsf{R}\acute{\sigma}\}$ where R is the reflection map on $\mathscr{M}$, $\mathsf{R} : \eta\eta' \mapsto \eta'\eta$. The weighted version gives a factor λ to vertex configurations $\dot{\mathtt{m}}_d(\dot{\underline{\sigma}}_{\partial v\to v}) = \mathtt{1}$ and to edge configurations $(\sigma, \sigma') = (\mathtt{11}, \mathtt{11})$:

$$\hat{\varphi}^\lambda(\sigma, \sigma') = \hat{\varphi}(\sigma, \sigma')\lambda^{\mathbf{1}\{\sigma=\mathtt{11}\}} = \mathbf{1}\{\sigma' = \mathsf{R}\sigma\}\lambda^{\mathbf{1}\{\sigma=\mathtt{11}\}},$$

$$\dot{\varphi}^\lambda(\dot{\underline{\sigma}}) = \dot{\varphi}(\dot{\underline{\sigma}})\lambda^{\mathbf{1}\{\eta(\dot{\underline{\sigma}})=\mathtt{1}\}} = \begin{cases} \lambda, & \dot{\underline{\sigma}} \in \mathrm{Per}[(\mathtt{10}^j, \mathtt{1f}^{d-j})_{0\leqslant j\leqslant d}] \quad \text{(1-variable)}, \\ 1, & \dot{\underline{\sigma}} \in \mathrm{Per}[(\mathtt{11}, \mathtt{f0}^j, \mathtt{ff}^{d-1-j})_{0\leqslant j\leqslant d-1}] \quad \text{(f-variable)}, \\ 1, & \dot{\underline{\sigma}} \in \mathrm{Per}[(\mathtt{f1}^2, \mathtt{00}^j, \mathtt{0f}^{d-2-j})_{0\leqslant j\leqslant d-2}] \quad \text{(0-variable, \textit{susceptible})}, \\ 1, & \dot{\underline{\sigma}} \in \mathrm{Per}[(\mathtt{01}^k, \mathtt{00}^j, \mathtt{0f}^{d-k-j})_{0\leqslant j\leqslant d-k, 3\leqslant k\leqslant d}] \quad \text{(0-variable, \textit{robust})}, \end{cases} \tag{30}$$

where we have classified a variable with spin $\mathtt{0}$ in the frozen model as *susceptible* (hereafter $\mathtt{0}^\mathtt{s}$) if it has exactly two neighboring $\mathtt{1}$'s, *robust* otherwise (hereafter $\mathtt{0}^\mathtt{r}$).

Remark 3.1. The frozen model is in exact bijection with the auxiliary model.[6] Indeed, let $\mathscr{YM} \equiv \{\mathtt{00}, \mathtt{01}, \mathtt{10}, \mathtt{0f}, \mathtt{f0}, \mathtt{ff}, \overline{\mathtt{ff}}\}$ and consider the mapping $\boldsymbol{v} : \mathscr{M} \to \mathscr{YM}$ which takes $\mathtt{11} \mapsto \overline{\mathtt{ff}}$, $\{\mathtt{10}, \mathtt{1f}\} \mapsto \mathtt{10}$, $\{\mathtt{01}, \mathtt{f1}\} \mapsto \mathtt{01}$, and acts as the identity on the remaining spins: $\boldsymbol{v}(\sigma)$ gives the two frozen model spins which must be incident to an edge with auxiliary model spin σ, with $\overline{\mathtt{ff}}$ indicating a matched pair of $\mathtt{f}$'s. Coordinate-wise application of $\boldsymbol{v}$ to an auxiliary configuration $\underline{\sigma}$ produces the configuration $\boldsymbol{v}(\underline{\sigma})$ from which one can directly read the corresponding frozen configuration $\overline{\underline{\eta}}$. This mapping is clearly invertible, by changing the $\mathtt{01}$-messages incident to $\mathtt{0}^\mathtt{s}$-vertices back to $\mathtt{f1}$. Let us note that the distribution of $\boldsymbol{v}(\underline{\sigma})$ is also a factor model, with factors $^\mathtt{v}\varphi$ obtained by applying φ to the pre-image under $\boldsymbol{v}$; we refer to this as the *vertex-auxiliary* model.

The primary purpose of defining the auxiliary model is that it gives us the following approach for calculating $\mathbb{E}\boldsymbol{Z}_{n\alpha}$. Given an auxiliary configuration $\underline{\sigma}$, write $m \equiv nd/2$ and consider the normalized empirical measures

$$\dot{\boldsymbol{h}}(\dot{\underline{\sigma}}) \equiv n^{-1} \sum_{v\in V} \mathbf{1}\{\dot{\underline{\sigma}}_v = \dot{\underline{\sigma}}\} \quad (\dot{\underline{\sigma}} \in \mathscr{M}^d) \qquad \text{variable empirical measure;}$$
$$\hat{\boldsymbol{h}}(\hat{\underline{\sigma}}) \equiv m^{-1} \sum_{a\in F} \mathbf{1}\{\hat{\underline{\sigma}}_a = \hat{\underline{\sigma}}\} \quad (\hat{\underline{\sigma}} \in \mathscr{M}^2) \qquad \text{clause empirical measure.}$$

---

[6]The correspondence remains valid even when the graph has multi-edges, provided we count neighbors with edge multiplicity — e.g. if a $\mathtt{0}$ neighbors a single $\mathtt{1}$-variable via a doubled edge, we consider it as neighboring two distinct $\mathtt{1}$-variables.

We regard $\boldsymbol{h} \equiv (\dot{\boldsymbol{h}}, \hat{\boldsymbol{h}})$ as a vector indexed by $(\operatorname{supp} \dot{\varphi}, \operatorname{supp} \hat{\varphi})$. For $\sigma \in \mathcal{M}$ and $\underline{\dot{\sigma}} \in \operatorname{supp} \dot{\varphi}$ let $\dot{H}_{\sigma,\underline{\dot{\sigma}}}$ denote the number of appearances of σ in $\underline{\dot{\sigma}}$, and similarly write $\hat{H}_{\sigma,\underline{\hat{\sigma}}}$ for the number of appearances of σ in $\underline{\hat{\sigma}}$. Then $\dot{H} \in \mathbb{Z}^{\bar{s} \times \dot{s}}$, $\hat{H} \in \mathbb{Z}^{\bar{s} \times \bar{s}}$ where $\dot{\boldsymbol{s}} \equiv |\operatorname{supp} \dot{\varphi}|$, $\bar{s} \equiv |\operatorname{supp} \hat{\varphi}| = |\mathcal{M}|$; and for $\boldsymbol{h}$ to correspond to a valid configuration $\underline{\sigma}$, the variable and clause empirical measures must give rise to the same edge marginals

$$\bar{h} = d^{-1} \dot{H} \dot{\boldsymbol{h}} = 2^{-1} \hat{H} \hat{\boldsymbol{h}}, \quad \bar{h}(\sigma) \equiv (nd)^{-1} \sum_{(va) \in H} \mathbf{1}\{\sigma_{(va)} = \sigma\}.$$

Definition 3.2. Let $\boldsymbol{\Delta}$ denote the space of probability measures $\boldsymbol{h} \equiv (\dot{\boldsymbol{h}}, \hat{\boldsymbol{h}})$ on $\operatorname{supp} \varphi$ (i.e., $\dot{\boldsymbol{h}}$ is a probability measure on $\operatorname{supp} \dot{\varphi}$ while $\hat{\boldsymbol{h}}$ is a probability measure on $\operatorname{supp} \hat{\varphi}$) such that

(i) $(\dot{\boldsymbol{h}}, \frac{d}{2}\hat{\boldsymbol{h}})$ lies in the kernel of matrix $H_{\boldsymbol{\Delta}} \equiv (\dot{H} \ \ -\hat{H}) \in \mathbb{Z}^{\bar{s} \times (\dot{s} + \bar{s})}$, and
(ii) $\boldsymbol{h}(\dot{\mathfrak{m}}_d(\underline{\dot{\sigma}}) = \mathtt{f}) \leqslant \beta_{\max}$ (cf. (8)).

Let $\boldsymbol{\Delta}_\alpha$ denote the subspace of measures $\boldsymbol{h} \in \boldsymbol{\Delta}$ with normalized intensity

$$\mathbf{i}(\boldsymbol{h}) \equiv \dot{\boldsymbol{h}}(\eta(\underline{\dot{\sigma}}) = 1) + \tfrac{1}{2}\dot{\boldsymbol{h}}(\eta(\underline{\dot{\sigma}}) = \mathtt{f}) = \bar{h}(\mathtt{1Z}) + (d/2)\,\bar{h}(\mathtt{11}) = \alpha.$$

We shall show (Lem. 5.4) that $H_{\boldsymbol{\Delta}}$ is surjective, therefore $\boldsymbol{\Delta}$ is an $(\dot{\boldsymbol{s}} - 1)$-dimensional space with $\boldsymbol{\Delta}_\alpha$ an $(\dot{\boldsymbol{s}} - 2)$-dimensional subspace.

The expected number of auxiliary configurations on $\mathcal{G}_{n,d}$ with empirical measure $\boldsymbol{h}$ is

$$\mathbb{E}\boldsymbol{Z}(\boldsymbol{h}) = \frac{\binom{n}{n\dot{\boldsymbol{h}}} \binom{m}{m\hat{\boldsymbol{h}}}}{\binom{nd}{nd\bar{h}}} \dot{\varphi}^{n\dot{\boldsymbol{h}}} \hat{\varphi}^{m\hat{\boldsymbol{h}}} \equiv \frac{n!\,m!}{(nd)!/\{\prod_\sigma (nd\bar{h}(\sigma))!\}} \prod_{\underline{\dot{\sigma}}} \frac{\dot{\varphi}(\underline{\dot{\sigma}})^{n\dot{\boldsymbol{h}}(\underline{\dot{\sigma}})}}{(n\dot{\boldsymbol{h}}(\underline{\dot{\sigma}}))!} \prod_{\underline{\hat{\sigma}}} \frac{\hat{\varphi}(\underline{\hat{\sigma}})^{m\hat{\boldsymbol{h}}(\underline{\hat{\sigma}})}}{(m\hat{\boldsymbol{h}}(\underline{\hat{\sigma}}))!}.$$

Stirling's formula gives $\mathbb{E}\boldsymbol{Z}(\boldsymbol{h}) = n^{O(1)} \exp\{n\boldsymbol{\Phi}(\boldsymbol{h})\}$ where

$$\boldsymbol{\Phi}(\boldsymbol{h}) \equiv \sum_{\underline{\dot{\sigma}}} \dot{\boldsymbol{h}}(\underline{\dot{\sigma}}) \log \frac{\dot{\varphi}(\underline{\dot{\sigma}})}{\dot{\boldsymbol{h}}(\underline{\dot{\sigma}})} + (d/2) \sum_{\underline{\hat{\sigma}}} \hat{\boldsymbol{h}}(\underline{\hat{\sigma}}) \log \frac{\hat{\varphi}(\underline{\hat{\sigma}})}{\hat{\boldsymbol{h}}(\underline{\hat{\sigma}})} - d \sum_\sigma \bar{h}(\sigma) \log \frac{1}{\bar{h}(\sigma)}. \tag{31}$$

If further $\min \boldsymbol{h} \gtrsim_d 1$ as $n \to \infty$, then

$$\mathbb{E}\boldsymbol{Z}(\boldsymbol{h}) = \frac{e^{O_d(n^{-1})} \mathscr{P}(\boldsymbol{h})}{(2\pi n)^{(\dot{s}-1)/2}} \exp\{n\,\boldsymbol{\Phi}(\boldsymbol{h})\}, \quad \mathscr{P}(\boldsymbol{h}) \equiv \left[\frac{\prod_\sigma d\bar{h}(\sigma)}{2 \prod_{\underline{\dot{\sigma}}} \dot{\boldsymbol{h}}(\underline{\dot{\sigma}}) \prod_{\underline{\hat{\sigma}}} (\frac{1}{2} d\hat{\boldsymbol{h}}(\underline{\hat{\sigma}}))}\right]^{1/2}. \tag{32}$$

Clearly an analogous expansion holds for the expectation of the λ-weighted partition function $\boldsymbol{Z}^\lambda(\boldsymbol{h}) = \lambda^{n\,\mathbf{i}(\boldsymbol{h})} \boldsymbol{Z}(\boldsymbol{h})$; we write $\boldsymbol{\Phi}^\lambda(\boldsymbol{h}) = \boldsymbol{\Phi}(\boldsymbol{h}) + \mathbf{i}(\boldsymbol{h}) \log \lambda$ for the associated rate function (with φ^λ in place of φ).

The first moment of frozen model configurations at intensity α is $\mathbb{E}\boldsymbol{Z}_{n\alpha} = \sum_{\boldsymbol{h} \in \boldsymbol{\Delta}_\alpha} \mathbb{E}\boldsymbol{Z}(\boldsymbol{h})$.[7] The aim of this section is to compute the exponent $^\star\boldsymbol{\Phi}(\alpha) = \lim_n n^{-1} \log \mathbb{E}\boldsymbol{Z}_{n\alpha}$ by determining the maximizer $^\star\boldsymbol{h}_\alpha \equiv (^\star\dot{\boldsymbol{h}}_\alpha, {}^\star\hat{\boldsymbol{h}}_\alpha)$ of $\boldsymbol{\Phi}$ on $\boldsymbol{\Delta}_\alpha$. Observe it is clear from the functional form of $\boldsymbol{\Phi}$ that $^\star\dot{\boldsymbol{h}}_\alpha$ and $^\star\hat{\boldsymbol{h}}_\alpha$ must be symmetric functions on $\mathcal{M}^d$ and $\mathcal{M}^2$ respectively.

The fugacity parameter λ serves the purpose of a Lagrange multiplier: if $\boldsymbol{h} \in \boldsymbol{\Delta}_\alpha$ is a stationary point of $\boldsymbol{\Phi}$ restricted to $\boldsymbol{\Delta}_\alpha$, then for some λ it must be a stationary point of $\boldsymbol{\Phi}^\lambda$ on the *unrestricted* space $\boldsymbol{\Delta}$. Strictly positive measures $\boldsymbol{h}$ which are stationary for $\boldsymbol{\Phi}^\lambda$ as a function on $\boldsymbol{\Delta}$ (unrestricted) correspond to a generalization of the tree Gibbs measures considered in §3.1, where the boundary conditions are specified by a law on incoming *and* *outgoing* messages, as we now describe. Let $T_{d,2}$ denote the infinite tree given by bisecting

---

[7] Though we omit it from the notation, the sum should be taken over measures $\boldsymbol{h}$ such that $n\dot{\boldsymbol{h}}$ and $\frac{nd}{2}\hat{\boldsymbol{h}}$ are integer-valued.

each edge of T_d by a new (clause) vertex; this is the local weak limit of the random $(d, 2)$-regular bipartite factor graph. The vertices at level t of $T_{d,2}$ are variables for t odd, clauses for t even. If $\underline{\sigma}$ is a message configuration on the edges of $T_{d,2}(t)$ — including the edges $E(t-1, t)$ joining levels $t-1$ and t — then let $\Psi_t^\lambda(\underline{\sigma})$ denote the product of the factor weights $\dot\varphi^\lambda(\underline{\dot\sigma}_v)$, $\hat\varphi^\lambda(\underline{\hat\sigma}_a)$ over all $v, a \in T_{d,2}(t-1)$. For probability measures $\dot h, \hat h$ on $\mathscr{M}$ we define the measures

$$
\begin{aligned}
Z_t \boldsymbol{\nu}_t(\underline{\sigma}) &= \Psi_t^\lambda(\underline{\sigma}) \textstyle\prod_{e \in E(t-1,t)} \dot h_{\sigma_e} \quad (t \text{ even}), \\
Z_t \boldsymbol{\nu}_t(\underline{\sigma}) &= \Psi_t^\lambda(\underline{\sigma}) \textstyle\prod_{e \in E(t-1,t)} \hat h_{\sigma_e} \quad (t \text{ odd}),
\end{aligned}
\tag{33}
$$

with Z_t the normalizing constant which makes $\boldsymbol{\nu}_t$ a probability measure.[8] The family $(\boldsymbol{\nu}_t)_t$ is consistent if and only if $h \equiv (\dot h, \hat h)$ satisfies the *Bethe recursions*

$$
\begin{aligned}
\dot z \dot h_\sigma &= \textstyle\sum_{\underline{\dot\sigma}\,:\,\sigma_1 = \sigma} \dot\varphi^\lambda(\underline{\dot\sigma}) \prod_{i=2}^d \hat h_{\sigma_i}, &\quad &\text{variable Bethe recursions;} \\
\hat z \hat h_\sigma &= \textstyle\sum_{\underline{\hat\sigma}\,:\,\sigma_1 = \sigma} \hat\varphi^\lambda(\underline{\hat\sigma}) \dot h_{\sigma_2} &\quad &\text{clause Bethe recursions.}
\end{aligned}
\tag{34}
$$

(with $\dot z, \hat z$ the normalizing constants). We show below that (as expected from our definition) these recursions are a generalization of the frozen model recursions (28). Thus a solution h of (34) specifies a Gibbs measure $\boldsymbol{\nu} \equiv \boldsymbol{\nu}^\lambda$ for the auxiliary model on $T_{d,2}$ which generalizes the measures $\nu \equiv \nu^\lambda$ described in §3.1.

The Bethe recursions read explicitly as follows. Write $\mathtt{Z} \equiv \{\mathtt{0}, \mathtt{f}\}$, e.g. $\hat h_{\mathtt{1Z}} \equiv \hat h_{\mathtt{10}} + \hat h_{\mathtt{1f}}$. The clause Bethe recursions are simply

$$
\hat z \hat h_\sigma = \dot h_{\mathtt{R}\sigma} \lambda^{\mathbf{1}\{\sigma = \mathtt{11}\}} \quad \text{with } \hat z = 1 + (\lambda - 1)\dot h_{\mathtt{11}}.
\tag{35}
$$

The variable Bethe recursions are

(A) $\dot z \dot h_{\mathtt{10}} = \dot z \dot h_{\mathtt{1f}} = \lambda(\hat h_{\mathtt{1Z}})^{d-1}, \qquad \lambda \dot z \dot h_{\mathtt{11}} = \lambda(\hat h_{\mathtt{fZ}})^{d-1},$
$\qquad \dot z \dot h_{\mathtt{f0}} = \dot z \dot h_{\mathtt{ff}} = (d-1)\hat h_{\mathtt{11}}(\hat h_{\mathtt{fZ}})^{d-2}, \qquad \dot z \dot h_{\mathtt{f1}} = (d-1)\hat h_{\mathtt{f1}}(\hat h_{\mathtt{0Z}})^{d-2},$

(B) $\dot z \dot h_{\mathtt{01}} = (\hat h_{\mathtt{0.}})^{d-1} - (\hat h_{\mathtt{0Z}})^{d-1} - (d-1)\hat h_{\mathtt{01}}(\hat h_{\mathtt{0Z}})^{d-2},$

(C) $\dot z \dot h_{\mathtt{00}} = \dot z \dot h_{\mathtt{0f}} = (\hat h_{\mathtt{0.}})^{d-1} - (\hat h_{\mathtt{0Z}})^{d-1} - (d-1)\hat h_{\mathtt{01}}(\hat h_{\mathtt{0Z}})^{d-2}$
$\qquad\qquad + \frac{1}{2}(d-1)(d-2)[(\hat h_{\mathtt{f1}})^2 - (\hat h_{\mathtt{01}})^2](\hat h_{\mathtt{0Z}})^{d-3}$

We immediately have $\hat h_{\mathtt{01}} = \hat h_{\mathtt{f1}}$, $\hat h_{\mathtt{0f}} = \hat h_{\mathtt{ff}}$, and $\hat h_{\mathtt{00}} = \hat h_{\mathtt{f0}}$, and then comparing (B) and (C) gives $\hat h_{\mathtt{10}} = \hat h_{\mathtt{00}} = \hat h_{\mathtt{f0}}$. It then follows from (A) that the following are equivalent (with the symbol ✓ indicating the identities we already know):

$$
\dot h_{\mathtt{f1}} = \dot h_{\mathtt{f0}} \overset{\checkmark}{=} \dot h_{\mathtt{ff}}, \quad \dot h_{\mathtt{1f}} = \hat h_{\mathtt{0f}} \overset{\checkmark}{=} \hat h_{\mathtt{ff}}, \quad \lambda \dot h_{\mathtt{11}} = \dot h_{\mathtt{10}} \overset{\checkmark}{=} \dot h_{\mathtt{1f}}, \quad \hat h_{\mathtt{11}} = \hat h_{\mathtt{01}} \overset{\checkmark}{=} \hat h_{\mathtt{f1}}.
\tag{36}
$$

If (36) *holds, then* (34) *reduces to the frozen model recursions* (28) *with*

$$
\begin{aligned}
3\hat h_{\eta'\eta} &= q_\eta, \quad \text{therefore (substituting into (35))} \\
3\dot h_{\eta\eta'} &= \lambda^{-\mathbf{1}\{\eta\eta' = \mathtt{11}\}} q_\eta \hat z \quad \text{with } \hat z = [1 - \tfrac{1}{3}q_1(1 - 1/\lambda)]^{-1}
\end{aligned}
\tag{37}
$$

proving our claim that the measures $\boldsymbol{\nu}$ generalize the measures ν of §3.1. The connection between these Gibbs measures and the rate function $\boldsymbol{\Phi}^\lambda$ is given by the following variational principle:

Lemma 3.3. *If a measure $\boldsymbol{h}$ in the interior $\boldsymbol{\Delta}^\circ$ of $\boldsymbol{\Delta}$ is stationary for $\boldsymbol{\Phi}^\lambda$, then $\boldsymbol{h}$ corresponds to a solution $h \equiv h^\lambda$ of the Bethe recursions* (34) *via*

$$
\dot z \dot{\boldsymbol{h}}(\underline{\dot\sigma}) = \dot\varphi^\lambda(\underline{\dot\sigma}) \textstyle\prod_{i=1}^d \hat h_{\sigma_i}, \quad \hat z \hat{\boldsymbol{h}}(\underline{\hat\sigma}) = \hat\varphi^\lambda(\underline{\hat\sigma}) \textstyle\prod_{i=1}^2 \dot h_{\sigma_i}, \quad \bar z \bar h(\sigma) = \hat h_\sigma \dot h_\sigma
\tag{38}
$$

---

[8]We suppress the λ-dependence from the notation except to differentiate φ^λ from $\varphi \equiv \varphi^\lambda|_{\lambda=1}$.

with $\dot{\boldsymbol{z}}, \hat{\boldsymbol{z}}, \bar{z}$ normalizing constants satisfying $\bar{z} = \dot{\boldsymbol{z}}/\dot{z} = \hat{\boldsymbol{z}}/\hat{z}$ for $\dot{z}, \hat{z}$ as in (34).

Proof. Follows from [15] upon verifying that $\dot{H}, \hat{H}$ are surjective; we will prove a stronger condition than surjectivity in Lem. 5.4 and (70). □

Recall now that our aim is to locate the global maximizer of $\boldsymbol{\Phi}$ on $\boldsymbol{\Delta}_\alpha$ as a stationary point of $\boldsymbol{\Phi}^\lambda$ for some value of λ.

Theorem 3.4. *Let $^\star\boldsymbol{h}^\lambda$ denote the unique stationary point of $\boldsymbol{\Phi}^\lambda$ which corresponds via (38) to the only solution $h \equiv h^\lambda$ of the Bethe recursions (34) which satisfies (36), that is, $\hat{h}^\lambda_{\eta'\eta} = \frac{1}{3}q^\lambda_\eta$. The unique maximizer $^\star\boldsymbol{h}_\alpha$ of $\boldsymbol{\Phi}$ on $\boldsymbol{\Delta}_\alpha$ is given by $^\star\boldsymbol{h}^\lambda$ for the unique value of λ at which $^\star\boldsymbol{h}^\lambda(\dot{\mathrm{m}}_d(\underline{\dot{\sigma}}) = 1) = \alpha$.*

In view of Lem. 3.3 and our preceding discussion of Gibbs measures and Lagrange multipliers, Thm. 3.4 will follow by showing

1. Any global maximizer $^\star\boldsymbol{h}_\alpha$ of $\boldsymbol{\Phi}$ on $\boldsymbol{\Delta}_\alpha$ lies in the interior $\boldsymbol{\Delta}_\alpha^\circ$. Thus for some[9] λ it is a stationary point of $\boldsymbol{\Phi}^\lambda$, and hence corresponds via (38) to a solution $h \equiv h_\alpha$ of the Bethe recursions (34).
2. Any such Bethe solution $h \equiv h_\alpha$ satisfies the symmetries (36), and hence is identified with the unique solution of the frozen model recursions (28). Consequently λ is given by the unique solution of $^\star\boldsymbol{h}^\lambda(\dot{\mathrm{m}}_d(\underline{\dot{\sigma}}) = 1) = \alpha$.

Ruling out boundary maximizers for $\boldsymbol{\Phi}$ is relatively easy, so we defer the proof to §4 where we will use the same argument to rule out boundary maximizers for the second-moment exponent $_2\boldsymbol{\Phi}$. We turn now to the more delicate task of proving the symmetries (36).

3.3. **Bethe recursion symmetries.** Suppose $^\star\boldsymbol{h}_\alpha$ is an interior maximizer for $\boldsymbol{\Phi}$ on $\boldsymbol{\Delta}_\alpha$, and so corresponds to a Bethe solution $h \equiv h^\lambda$. Let $\acute{T}_{d,2}$ denote $T_{d,2}$ with a subtree incident to the root removed, such that one clause $a \in \partial o$ is incident to an unmatched half-edge $\acute{e}$ (Fig. 2). Consider defining a Gibbs measure for φ^λ on $\acute{T}_{d,2}$ in the manner of (33), with boundary law given by h. Then the marginal law of $\sigma_{\acute{e}}$ will be $\hat{h}$, and the marginal law of the d-tuple of spins incident to any given vertex will be $\dot{\boldsymbol{h}}$. Further, the Gibbs measure on $\acute{T}_{d,2}$ can be generated in Markovian fashion, starting with spin $\sigma_{\acute{e}}$ distributed according to $\hat{h}$, generating the messages on the other $d-1$ edges incident to o according to the conditional measure $\dot{\boldsymbol{h}}(\underline{\dot{\sigma}} \,|\, \sigma_1 = \sigma_{\acute{e}})$, and continuing iteratively down the tree.

Write $\sigma_{\acute{e}} \equiv \boldsymbol{io}$ where $\boldsymbol{i}$ is the incoming message and $\boldsymbol{o}$ is the outgoing (in Fig. 2, $\boldsymbol{o}$ is directed upwards, $\boldsymbol{i}$ downwards). Given any valid auxiliary configuration $\underline{\sigma}$ on the edges of $\acute{T}_{d,2}$, changing $\boldsymbol{i}$ and passing the changed message through the tree (via $\dot{\mathrm{m}}_{d-1}$) produces a new auxiliary configuration $\underline{\sigma}'$ (Fig. 2). The symmetries (36) will follow by showing that for any fixed $\boldsymbol{o}$, the effect of changing $\boldsymbol{i}$ is *measure-preserving* under the Gibbs measure $\boldsymbol{\nu}_\alpha$ corresponding to $^\star\boldsymbol{h}_\alpha$. From our definition of the Gibbs measure via the boundary law, the measure-preserving property will follow by showing that the effect of changing $\boldsymbol{i}$ almost surely does not percolate down the tree.

Indeed, recall that we already saw directly from the Bethe recursions that $\hat{h}_{0\eta} = \hat{h}_{\mathtt{f}\eta}$ for any η: this corresponds to the fact that $\dot{\mathrm{m}}_{d-1}$ does not differentiate between incoming messages 0 and f, so changing $\boldsymbol{i}$ from 0 to f or vice versa has no effect at all below $\acute{e}$. We also found that $\hat{h}_{\eta 0}$ does not depend on η: if $\boldsymbol{o} = 0$ then all messages outgoing from the root must be 0 or f, so changing $\boldsymbol{i}$ can have an effect at most one level down.

---

[9]Eventually we will find $\lambda = \lambda_\star = \exp\{-(^\star\boldsymbol{\Phi})'(\alpha)\}$.
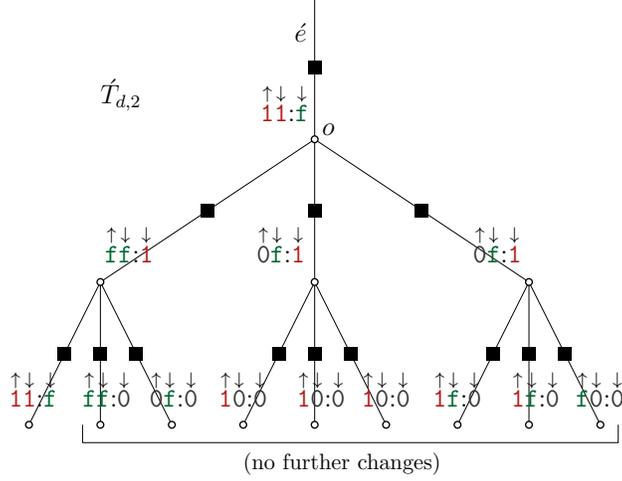
FIGURE 2. Change of message incoming down to $\acute{e}$ is passed down $\acute{T}_{d,2}$
($\eta\acute{\eta}{:}\grave{\eta}$ means message η up, message $\acute{\eta}$ down in $\underline{\sigma}$, message $\grave{\eta}$ down in $\underline{\sigma}'$)

Proof of Thm. 3.4. We shall assume that the maximizer $^\star\boldsymbol{h}_\alpha$ of $\boldsymbol{\Phi}$ on $\boldsymbol{\Delta}_\alpha$ lies in the interior $\boldsymbol{\Delta}_\alpha^\circ$, deferring the proof to §4 (see Propn. 4.5 and Cor. 4.8).

From the above discussion it remains to show the corresponding Bethe solution $h \equiv h_\alpha$ satisfies $\hat{h}_{11} = \hat{h}_{\mathtt{f1}}$: meaning that in the Gibbs measure $\boldsymbol{\nu}_\alpha$ corresponding to $^\star\boldsymbol{h}_\alpha$, changing $\boldsymbol{io}$ from $\mathtt{11}$ to $\mathtt{f1}$ or vice versa has a finite-range effect. Fig. 2 shows that the effect of changing $\mathtt{11}$ to $\mathtt{f1}$ can only propagate through components of $\mathtt{f}$-variables, while the effect of changing $\mathtt{f1}$ to $\mathtt{11}$ can only propagate through chains of alternating $\mathtt{1}$- and $\mathtt{0^s}$-variables (cf. (30)).[10]

We claim that both propagations are subcritical under $\boldsymbol{\nu}_\alpha$. By (38), for any $2 \leqslant j \leqslant d$ the density of $\mathtt{0}$-variables with exactly j neighboring $\mathtt{1}$-variables is, up to a normalizing constant, $\binom{d}{j}\dot{h}_{\mathtt{10}}^j \dot{h}_{\mathtt{Z0}}^{d-j}$ (where we have used that $\dot{h}_{\mathtt{10}} = \dot{h}_{\mathtt{1f}}$). Thus, among all the edges incident to $\mathtt{0}$-variables, the proportion coming from $\mathtt{1}$-variables is

$$\frac{\alpha_{\mathtt{1}}}{1 - \alpha_{\mathtt{1}} - \beta} = \frac{\frac{1}{d}\sum_{j \geqslant 2} j \binom{d}{j} u^j}{\sum_{j \geqslant 2} \binom{d}{j} u^j} \quad \text{where } u \equiv \dot{h}_{\mathtt{10}}/\dot{h}_{\mathtt{Z0}}. \tag{39}$$

This is an increasing function of u (the derivative in u is the variance of a conditioned binomial random variable), so in the regime $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ we must have

$$\dot{h}_{\mathtt{10}}/\dot{h}_{\cdot\mathtt{0}} = \alpha\left[1 + O(\tfrac{\log d}{d})\right]. \tag{40}$$

Comparing the number of $\mathtt{1}$-to-$\mathtt{0^s}$ edges with the number of other $\mathtt{1}$-to-$\mathtt{0}$ edges gives (applying (38) and (40))

$$\frac{\dot{h}_{\mathtt{f1}}}{\dot{h}_{\mathtt{01}}} = \frac{nd\hat{\boldsymbol{h}}(\mathtt{1f},\mathtt{f1})}{nd\hat{\boldsymbol{h}}(\mathtt{10},\mathtt{01})} = \frac{2n\dot{\boldsymbol{h}}(\eta = \mathtt{0^s})}{nd\alpha - 2n\dot{\boldsymbol{h}}(\eta = \mathtt{0^s})} = \frac{2\binom{d}{2}u^2}{\displaystyle\sum_{j \geqslant 3} j\binom{d}{j}u^j} \leqslant d^{-3/2}. \tag{41}$$

---

[10]Here we write η-variables to mean variables with spin η in the corresponding frozen model configuration.

Thus, conditioned on $\boldsymbol{io} = \mathtt{f1}$, the expected number of $a \in \partial o$ with $\sigma_{a \to o} = \mathtt{f}$ is less than one, from which it follows that the effect of changing $\mathtt{f1}$ to $\mathtt{11}$ does not percolate.[11] Likewise

$$\frac{\beta}{1 - \alpha - \beta} \geqslant \frac{\frac{1}{d}\sum_{j \geqslant 2}(d - j)(\dot{h}_{\mathtt{f0}}/\dot{h}_{\mathtt{Z0}})\binom{d}{j}u^j}{\sum_{j \geqslant 2}\binom{d}{j}u^j} = \dot{h}_{\mathtt{f0}}/\dot{h}_{\mathtt{Z0}}\left[1 - O((\log d)/d)\right]$$

so $\dot{h}_{\mathtt{f0}}/\dot{h}_{\mathtt{00}} \lesssim d^{-3/2}$ which implies that the effect of changing $\mathtt{11}$ to $\mathtt{f1}$ does not percolate, concluding the proof. $\qquad\square$

3.4. **Explicit form of first moment exponent.** We conclude this section by giving the explicit form of $^\star\boldsymbol{\Phi}(\alpha) \equiv {}^\star\boldsymbol{\Phi}_d(\alpha)$ when $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$:

Proposition 3.5. *For $d \geqslant d_0$, $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$, $^\star\boldsymbol{\Phi}(\alpha) \equiv {}^\star\boldsymbol{\Phi}_d(\alpha)$ is given by*

$$^\star\boldsymbol{\Phi}(\alpha) = -\log[1 - q(1 - 1/\lambda)] - (d/2 - 1)\log[1 - q^2(1 - 1/\lambda)] - \alpha\log\lambda \qquad (42)$$

where λ and q are determined from α via

$$\begin{cases} \lambda = q\dfrac{1 - (1 - q)^{d-1}}{(1 - q)^d}, \ \textit{and } q \equiv x(\log d)/d \textit{ is the unique solution of} \\ \alpha = q\dfrac{1 + (d/2 - 1)(1 - q)^{d-1}}{1 + q - (1 - q)^{d-1}} \ \textit{ on the interval } 1.6 \leqslant x \leqslant 3. \end{cases} \qquad (43)$$

On the interval $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ the function $^\star\boldsymbol{\Phi}$ is strictly decreasing, with a unique zero $\alpha_\star$ in the interval's interior. The gap between $\alpha_\star$ and the first moment threshold $\alpha_\square$ of the original independent set partition function (3) is given by

$$\alpha_\square - \alpha_\star = \left(\frac{e\log d}{2d}\right)^2\left[1 + O\left(\frac{1}{\log d}\right)\right]. \qquad (44)$$

Proof. In the following, let $q \equiv q^\lambda$ denote the solution $(q_1)^\lambda$ of the recursions (28).
Explicit Bethe prediction. Substituting (38) into (31) and rearranging gives

$$\begin{aligned} ^\star\boldsymbol{\Phi}^\lambda = \boldsymbol{\Phi}^\lambda(^\star\boldsymbol{h}^\lambda) &= \log\dot{z} + (d/2)\log\hat{z} - d\sum_\sigma \bar{h}(\sigma)\log[\dot{h}_\sigma\hat{h}_\sigma/\bar{h}(\sigma)] \\ &= \log\dot{z} + (d/2)\log\hat{z} - d\log\bar{z}. \end{aligned}$$

We use (37) to calculate $\bar{z}, \dot{z}, \hat{z}$ in terms of q, λ:

$$\bar{z} = \hat{z}\sum_\sigma \frac{\hat{h}_\sigma\hat{h}_{\mathrm{R}\sigma}}{\lambda^{\mathbf{1}\{\sigma=\mathtt{11}\}}} = \hat{z}\sum_{\eta,\eta'}\frac{(q_\eta/3)(q_{\eta'}/3)}{\lambda^{\mathbf{1}\{\sigma=\mathtt{11}\}}} = \frac{\hat{z}}{9}\left[1 - q^2(1 - 1/\lambda)\right];$$

$$\hat{\hat{z}} = \sum_\sigma \lambda^{\mathbf{1}\{\sigma=\mathtt{11}\}}\hat{h}_\sigma\hat{h}_{\mathrm{R}\sigma} = \hat{z}^2\sum_\sigma \frac{\hat{h}_\sigma\hat{h}_{\mathrm{R}\sigma}}{\lambda^{\mathbf{1}\{\sigma=\mathtt{11}\}}} = \bar{z}\hat{z};$$

$$\dot{z} = \lambda(\hat{h}_{\mathtt{1Z}})^d + d\hat{h}_{\mathtt{11}}(\hat{h}_{\mathtt{fZ}})^{d-1} + \binom{d}{2}(\hat{h}_{\mathtt{f1}})^2(\hat{h}_{\mathtt{0Z}})^{d-2} + \sum_{k \geqslant 3}(\hat{h}_{\mathtt{01}})^k(\hat{h}_{\mathtt{0Z}})^{d-k}$$

$$= \frac{1}{3^d}\left[1 + (\lambda - 1)(1 - q)^d\right] = \frac{1 - q^2(1 - 1/\lambda)}{3^d[1 - q(1 - 1/\lambda)]} = \bar{z}\dot{z}.$$

Substituting back into (38) gives $\dot{\boldsymbol{h}}(\dot{\mathtt{m}}_d(\underline{\dot{\sigma}}) = \mathtt{1}) = \dot{z}^{-1}\lambda(\hat{h}_{\mathtt{10}})^d = q(1 - q)/[1 - q^2(1 - 1/\lambda)]$ and $\dot{\boldsymbol{h}}(\dot{\mathtt{m}}_d(\underline{\dot{\sigma}}) = \mathtt{f}) = \dot{z}^{-1}d\hat{h}_{\mathtt{11}}(\hat{h}_{\mathtt{fZ}})^{d-1} = dq/[\lambda(1 - q)]\dot{\boldsymbol{h}}(\dot{\mathtt{m}}_d(\underline{\dot{\sigma}}) = \mathtt{1})$, therefore

$$\alpha = \dot{\boldsymbol{h}}(\dot{\mathtt{m}}_d(\underline{\dot{\sigma}}) = \mathtt{1}) + \tfrac{1}{2}\dot{\boldsymbol{h}}(\dot{\mathtt{m}}_d(\underline{\dot{\sigma}}) = \mathtt{f}) = q\frac{1 - q + dq/(2\lambda)}{1 - q^2(1 - 1/\lambda)}.$$

---

[11]Note the main estimate we required for this argument as that $\dot{\boldsymbol{h}}(\eta = \mathtt{0^s}) \ll d^{-1}$, which is intuitively clear and can be derived by other means. We have taken an approach which circumvents further *a priori* estimates.

The recursion (28) also gives the expressions in (43) for λ and α solely in terms of q. The mapping $q \mapsto \alpha$ is not one-to-one on the entire interval $0 \leqslant q \leqslant 1$, but recalling (40) we must have $q = x(\log d)/d$ for $1.6 \leqslant x \leqslant 3$, and on this interval it is easily verified that the mapping is indeed one-to-one, with $\lambda = d^x q[1 + O(d^{-1}(\log d)^2)]$ and $D_q\,\alpha = 1 + O(d^{-1/2})$ where D_q denotes the total derivative with respect to q. This completes the verification of (43); it then follows from Thm. 3.4 that $^\star\Phi(\alpha) = {}^\star\Phi^\lambda(^\star h^\lambda) - \alpha \log \lambda$ is given by (42). We note here that $^\star\Phi'(\alpha) = -\log \lambda$, therefore $^\star\Phi''(\alpha) = -(D_q\alpha)^{-1}D_q(\log \lambda) = -d[1 + O((\log d)^{-1})]$.

Comparison of first-moment exponents. By contrast, the original independent set partition function has first-moment exponent $\Phi(\alpha)$ calculated in (6). This exponent also has a Bethe variational characterization, which can be expressed in terms of the fixed point $\acute{q}$ of the hard-core tree recursions:

$$\Phi(\alpha) = \log[\acute{\lambda}(1 - \acute{q})^d + 1] - (d/2)\log(1 - \acute{q}^2) - \alpha \log \acute{\lambda}$$

$$\text{where } \acute{q} = \frac{\acute{\lambda}(1 - \acute{q})^{d-1}}{1 + \acute{\lambda}(1 - \acute{q})^{d-1}} \text{ and } \alpha = \frac{\acute{\lambda}(1 - \acute{q})^d}{1 + \acute{\lambda}(1 - \acute{q})^d}.$$

This formula can be derived loosely in the same manner as (42); its validity can be checked simply by verifying that it agrees with (6). The relation between $\alpha, \acute{q}$ is given by $\alpha = \acute{\alpha}(\acute{q}) = \acute{q}/(1 + \acute{q})$, $\acute{q} = \acute{\alpha}^{-1}(\alpha) = \alpha/(1 - \alpha)$, so we compare $^\star\Phi(\alpha)$ and $\Phi(\alpha)$ by expressing both in terms of q: $^\star\Phi(\alpha)$ is given by (42) with α, λ defined in terms of q by (43), while

$$\Phi(\alpha) = -\log(1 - \acute{q}) - (d/2 - 1)\log(1 - \acute{q}^2) - \alpha \log \acute{\lambda},$$
$$\text{where } \acute{\lambda} = \acute{q}/(1 - \acute{q})^d \text{ and } \acute{q} = \acute{\alpha}^{-1}(\alpha(q)).$$

Let us emphasize that q and $\acute{q}$ are not equal but rather are related through the same α; explicitly $\acute{q} = q(1 - q + \frac{dq}{2\lambda})/(1 - q - \frac{d-2}{2\lambda}q^2)$. A little algebra then gives

$$(\Phi - \Phi)(\alpha) = \log \frac{1 - q + q/\lambda}{1 - q} + \log \frac{1 - q}{1 - \acute{q}} + (d/2 - 1)\log \frac{1 - q^2 + q^2/\lambda}{1 - \acute{q}^2} - \alpha \log \frac{\acute{\lambda}}{\lambda}$$

$$= \log \frac{1 - q + q/\lambda}{1 - q} + \log \frac{(1 - q)(1 - q - \frac{d-2}{2\lambda}q^2)}{(1 - q)^2 - \frac{d-1}{\lambda}q^2} + (d/2 - 1)\log \frac{(1 - q)^2(1 - \frac{d-2}{2\lambda}\frac{q^2}{1-q})^2}{(1 - q)^2 - \frac{d-1}{\lambda}q^2} - \alpha \log \frac{\acute{\lambda}}{\lambda}$$

$$= \log\left(1 + \frac{q}{\lambda(1 - q)}\right) - \log\left(1 - \frac{d - 2}{2\lambda}\frac{q^2}{1 - q}\right) + d\log \frac{1 - \frac{d-2}{2\lambda}\frac{q^2}{1-q}}{[1 - \frac{d-1}{\lambda}\frac{q^2}{(1-q)^2}]^{1/2}} - \alpha \log \frac{\acute{\lambda}}{\lambda}.$$

Taylor expanding (recalling $q = x(\log d)/d$, $\lambda = d^x q[1 + O(d^{-1}(\log d)^2)]$) gives

$$\alpha \log \frac{\acute{\lambda}}{\lambda} = \alpha \log \frac{\acute{q}}{q} + d\alpha \log \frac{1 - q}{1 - \acute{q}} - \overbrace{\alpha \log[1 - (1 - q)^d]}^{O(q^2/\lambda)}$$

$$= \alpha \log\left(1 + \frac{dq}{2\lambda}\frac{1 + (1 - \frac{2}{d})q}{1 - q - \frac{d-2}{2\lambda}q^2}\right) + d\alpha \log\left(1 + \frac{dq^2}{2\lambda}\frac{1 + (1 - \frac{2}{d})q}{(1 - q)^2(1 - \frac{(d-1)q^2}{\lambda(1-q)^2})}\right) - O\left(\frac{q^2}{\lambda}\right)$$

$$= \frac{dq^2}{2\lambda}(dq + 1) + O\left(\frac{d^2q^4}{\lambda} + \frac{d^3q^4}{\lambda^2}\right) = \frac{dq^2}{2\lambda}(dq + 1) + O\left(\frac{(\log d)^3}{d^{x+1}} + \frac{(\log d)^2}{d^{2x-1}}\right).$$

Substituting into the above expression for $(\Phi - \Phi)(\alpha)$ and expanding the other terms gives

$$(\Phi - \Phi)(\alpha) = \frac{q}{\lambda} + \frac{dq^2}{2\lambda}(dq + 2) + O\left(\frac{dq^3}{\lambda} + \frac{d^2q^3}{\lambda^2}\right) - \alpha \log \frac{\acute{\lambda}}{\lambda}$$

$$= \frac{q}{\lambda} + \frac{dq^2}{2\lambda} + O\left(\frac{d^2q^4}{\lambda} + \frac{d^3q^4}{\lambda^2}\right) = \frac{q}{2\lambda}(dq + 2) + O\left(\frac{(\log d)^3}{d^{x+1}} + \frac{(\log d)^2}{d^{2x-1}}\right) \asymp d^{-x}\log d. \quad (45)$$

Comparison of first-moment thresholds. It is clear from the above that $^\star\mathbf{\Phi}$ has a unique zero $\alpha_{\mathrm{lbd}} < \alpha_\star < \alpha_\square$ with $\alpha_\square$ the first moment threshold of the original independent partition function, determined in Lem. 2.1. Let $q_\square, \lambda_\square$ denote the solution of (43) for $\alpha = \alpha_\square$, and note $q_\square = \alpha_\square[1 + O(d^{-1}(\log d)^2)]$. Consider $0 \leqslant \delta \leqslant d^{-2}(\log d)^3$: applying (45) with the above estimate of $^\star\mathbf{\Phi}''(\alpha)$ gives

$$\begin{aligned} ^\star\mathbf{\Phi}(\alpha_\square - \delta) &= {}^\star\mathbf{\Phi}(\alpha_\square) + \delta \log \lambda_\square + O(d\delta^2) \\ &= -(2\lambda_\square)^{-1}\alpha_\square(d\alpha_\square + 2)[1 + O(d^{-1}(\log d)^3)] + \delta \log \lambda_\square, \end{aligned}$$

so the gap between the threshold is given by

$$\alpha_\square - \alpha_\star = [1 + O(d^{-1}(\log d)^3)]\frac{\alpha_\square(d\alpha_\square + 2)}{2\lambda_\square \log \lambda_\square} \asymp d^{-2}(\log d)^2.$$

For a more precise estimate, recall from Lem. 2.1 that $\alpha_\square = \widetilde{\alpha}_\square[1 + O(d^{-1}\log d)]$ where $\widetilde{\alpha}_\square$ is the zero of $F(\alpha) \equiv \log(e/\alpha) - (d+1)\alpha/2$. Then

$$\begin{aligned} \log \lambda_\square &= F(q_\square) + (dq_\square/2) + O(d^{-1}(\log d)^2) \\ &= (dq_\square/2) + O(d^{-1}(\log d)^3) = (d\widetilde{\alpha}_\square/2)\left[1 + O(d^{-1}(\log)^2)\right]. \end{aligned}$$

Substituting into the above gives (44), concluding the proof. □

4. Second moment of frozen model

In this section we compute the exponential growth rate $^\star_2\mathbf{\Phi}(\alpha) = \lim_{n\to\infty} n^{-1} \log \mathbb{E}[\mathbf{Z}_{n\alpha}^2]$ of the second moment of the (truncated) frozen model partition function (8). This is done in the same framework as in introduced in §3, regarding the second moment as the first moment of the partition function of *pair* frozen model configurations $(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2)$ on the same underlying graph. The corresponding model of pair auxiliary configurations $\underline{\tau} \equiv (\underline{\sigma}^1, \underline{\sigma}^2)$ has factors $\dot{\varphi}_2 \equiv \dot{\varphi} \otimes \dot{\varphi}$ and $\hat{\varphi}_2 \equiv \hat{\varphi} \otimes \hat{\varphi}$, and rate function $_2\mathbf{\Phi}$ on the space $_2\mathbf{\Delta}_\alpha$ of empirical measures $_2\boldsymbol{h}$ with both marginals in $\mathbf{\Delta}_\alpha$.

Theorem 4.1. *The rate function $_2\mathbf{\Phi}$ on $_2\mathbf{\Delta}_\alpha$ attains its maximum only at the product measure $^\star_2\boldsymbol{h}_\alpha \equiv {}^\star\boldsymbol{h}_\alpha \otimes {}^\star\boldsymbol{h}_\alpha$ or at the measure $^{\cdot\cdot}_2\boldsymbol{h}_\alpha$ with marginals $^\star\boldsymbol{h}_\alpha$ which is supported on pair configurations $\underline{\tau} \equiv (\underline{\sigma}, \underline{\sigma})$.*

4.1. **Intermediate overlap regime.** Write $_2\boldsymbol{Z}_{n\alpha}(n\rho)$ for the contribution to $(\boldsymbol{Z}_{n\alpha})^2$ from *pair frozen model* configurations $(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2)$ on G in which exactly $n\rho$ vertices take spin $\mathtt{1}$ in both configurations. In this subsection we show that the value of ρ maximizing $\mathbb{E}[_2\boldsymbol{Z}_{n\alpha}(n\rho)]$ must either be very small or very close to α. This will be done by a rather crude comparison with the contribution $_2Z_{n\alpha}(n\rho)$ to $(Z_{n\alpha})^2$ from *pair independent set* configurations (x^1, x^2) on G with overlap $n\rho$. At the same time we shall prove that the global maximizer of $\mathbf{\Phi}$ on $\mathbf{\Delta}_\alpha$ lies in the interior $\mathbf{\Delta}_\alpha^\circ$ (as required in the proof of Thm. 3.4), and that any global maximizer of $_2\mathbf{\Phi}$ on $_2\mathbf{\Delta}_\alpha$ *away from* $^{\cdot\cdot}_2\boldsymbol{h}_\alpha$ lies in the interior $_2\mathbf{\Delta}_\alpha^\circ$. More refined estimates for the regimes near $^\star_2\boldsymbol{h}_\alpha$ and $^{\cdot\cdot}_2\boldsymbol{h}_\alpha$ are reserved for §4.3 and 4.4.

Given a pair frozen configuration $\overline{\boldsymbol{\eta}}^i \equiv (\underline{\eta}^i, \underline{m}^i)$ $(i = 1, 2)$ on G, define

$$\underline{\omega} \equiv \underline{\omega}(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2) \in \mathscr{P}^V, \quad \mathscr{P} \equiv \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}^2 \cup \{\mathtt{f!f}\} \tag{46}$$

by setting $\omega_v = (\eta_v^1, \eta_v^2)$ *unless* $(\eta_v^1, \eta_v^2) = \mathtt{ff}$ and v is matched to different vertices under $\underline{m}^1$ and $\underline{m}^2$, in which case we set $\omega_v = \mathtt{f!f}$. For $\eta \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}$ write $\mathscr{P}_\eta$ for the spins $\omega \in \mathscr{P}$ with either coordinate equal to η. For $S \subseteq V$ write

$$\begin{aligned} &\text{(for } \omega \in \mathscr{P}) \quad S_\omega \equiv \{v \in S : \omega_v = \omega\}, \quad n_\omega \equiv |V_\omega|; \\ &\text{(for } \eta \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}) \quad S_\eta \equiv \{v \in S : \omega_v \in \mathscr{P}_\eta\}, \quad n_\eta \equiv |V_\eta|. \end{aligned} \tag{47}$$

Let $V_{0f}^\bullet$ denote the set of $0f$-vertices whose matched partner does not have spin $1f$, and $n_{0f}^\bullet \equiv |V_{0f}^\bullet| = n_{0f} - n_{1f}$. Define symmetrically $V_{f0}^\bullet$ and $n_{f0}^\bullet \equiv |V_{f0}^\bullet|$. We decompose

$$\mathbb{E}[_2 Z_{n\alpha}(n\rho)] = \sum_{\pi \in M(n\alpha, n\rho)} \mathbb{E}[_2 Z(\pi)] \quad \text{and} \quad \mathbb{E}[_2 \boldsymbol{Z}_{n\alpha}(n\rho)] = \sum_{\overline{\pi} \in \mathcal{M}(n\alpha, n\rho)} \mathbb{E}[_2 \boldsymbol{Z}(\overline{\pi})]$$

where $M(n\alpha, n\rho)$ is the appropriate space of empirical measures on $\{0, 1\}^2$ and $\mathcal{M}(n\alpha, n\rho)$ is the appropriate space of empirical measures on $\mathscr{P}$.

Proposition 4.2. *Suppose $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ and $0 \leqslant \rho \leqslant \alpha - \varepsilon_0(\log d)/d$ for ε_0 a small positive constant uniform in d. For $d \geqslant d_0(\varepsilon_0)$, $n \geqslant n_0(d)$,*

$$\mathbb{E}[_2 \boldsymbol{Z}_{n\alpha}(n\rho)] = \mathbb{E}[_2 Z_{n\alpha}(n\rho)] \exp\{O[nd^{-1}(d^{-0.49} + d^{-3\varepsilon_0/4})]\},$$

(recalling that $\boldsymbol{Z}$ refers to the frozen model while Z refers to the independent set model).

Proof. Let $\pi \in \mathcal{M}(n\alpha, n\rho)$; we decompose the expected contribution from π to the pair frozen model partition function as

$$\mathbb{E}[_2 \boldsymbol{Z}(\pi)] = \mathbf{c}(\pi) \, \mathbb{E}[\mathbf{J}_f^\pi \mathbf{J}_1^\pi \mathbf{J}_0^\pi] \tag{48}$$

where $\mathbf{c}(\pi)$ denotes the multinomial coefficient $\binom{n}{n\pi}$, and the $\mathbf{J}_\eta^\pi$ are defined with respect to a *fixed* spin configuration $\underline{\omega} \in \mathscr{P}^V$ with empirical measure π, as follows: $\mathbf{J}_f^\pi$ is the number of matchings on the f-vertices; $\mathbf{J}_1^\pi$ is the indicator that 1-vertices neighbor only 0-vertices and further that every vertex in $V_{0f} \cup V_{f0}$ is forced; and $\mathbf{J}_0^\pi$ is the indicator that every vertex in $V_{00} \cup V_{01} \cup V_{10}$ is "forced," that is, has at least two 1-neighbors in each coordinate in which it takes spin 0.

Matchings on f-vertices.

Let $\mathbf{j}_f(\pi) \equiv \mathbb{E}[\mathbf{J}_f^\pi]$ denote the expected number of matchings on the n_f free vertices in the spin configuration $\underline{\omega}$. We claim that for any valid π,

$$(2n)^{-1} \mathbf{k}_f(\pi) \leqslant \mathbf{j}_f(\pi) \leqslant \mathbf{k}_f(\pi) \quad \text{where}$$
$$\mathbf{k}_f(\pi) \equiv \frac{d^{n_f + n_{f!f}}}{[d/(d-1)]^{n_{f!f}}} \frac{(n_{ff} - 1)!!(n_{0f})_{n_{1f}}(n_{f0})_{n_{f1}}(n_{0f}^\bullet + n_{f!f} - 1)!!(n_{f0}^\bullet + n_{f!f} - 1)!!}{[nd]_{(n_f + n_{f!f})/2}}. \tag{49}$$

(For each vertex in $V_f \backslash V_{f!f}$ we distinguish one half-edge to participate in the matching. For each $f!f$-vertex we distinguish an ordered pair of half-edges, the first to participate in $\underline{m}^1$ and the second in $\underline{m}^2$. There are $n_f + n_{f!f}$ distinguished half-edges in total.) If $n_f < 2$ then $\mathbf{j}_f(\pi) = \mathbf{k}_f(\pi)$; in general, $\mathbf{k}_f$ upper bounds $\mathbf{j}_f$ because it counts matchings without enforcing the constraint that two $f!f$-vertices cannot be matched in both $\underline{m}^1$ and $\underline{m}^2$.

For the lower bound when $n_{f!f} \geqslant 2$, suppose without loss that $n_{f0}^\bullet \geqslant n_{0f}^\bullet$, and note this implies $a^\bullet \equiv n_{f0}^\bullet + n_{f!f}$ must be at least 4 — otherwise $n_{f!f} = 2$ while $n_{0f}^\bullet = n_{f0}^\bullet = 0$ so there is no valid matching on the $f!f$-vertices. Say that all the distinguished half-edges have been assigned except for the $a^\bullet$ half-edges incident to $V_{f0}^\bullet \cup V_{f!f}$ which were chosen to participate in $\underline{m}^1$. Now match these remaining half-edges one pair at a time, but avoid forming pairs already present in $\underline{m}^2$. The number of choices for the first $(a^\bullet/2) - 1$ pairs is $\geqslant (a^\bullet - 2)(a^\bullet - 4) \cdots 2 \geqslant (a^\bullet - 3)!! \geqslant (a^\bullet - 1)!!/n$.

The procedure succeeds if and only if the final pair remaining is not already present in $\underline{m}^2$. To bound the probability that it fails, note that if given a failed matching in which the final pair is already present in $\underline{m}^2$, we can choose any of the first $(a^\bullet/2) - 1$ pairs, and switch the half-edges in one of two ways to produce a valid matching (Fig. 3). Thus each failed matching maps to $a^\bullet - 2$ valid matchings.

(A) Matching of final $a^\bullet$ half-edges fails because final (leftmost) pair already present in $\underline{m}^2$

(B) One valid switching

(C) A distinct valid switching
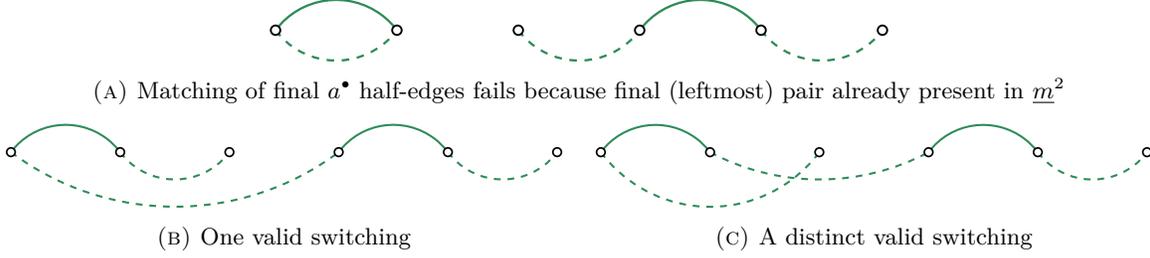
FIGURE 3. Two valid switchings of a failed matching (dashed lines $\underline{m}^1$, solid lines $\underline{m}^2$).

In the reverse direction, if the final pair is given then the failed matching can be uniquely recovered from the valid matching, so each valid matching has at most $n_{\mathtt{f}!\mathtt{f}}/2$ preimages. Thus the ratio of failed to valid matchings is (recalling $a^\bullet \geqslant 4$) at most $(n_{\mathtt{f}!\mathtt{f}}/2)/(a^\bullet - 2) \leqslant (2 - 4/a^\bullet)^{-1} \leqslant 1$. This proves that the matching procedure succeeds with probability at least $^1\!/_2$, and the lower bound in (49) follows. Combining with $\mathbf{c}(\pi)$ gives

$$n^{O(1)}\,\mathbf{c}(\pi)\,\mathbf{j_f}(\pi) = \frac{d^{(n_{\mathtt{f}}+n_{\mathtt{f}!\mathtt{f}})/2}\,(1-1/d)^{n_{\mathtt{f}!\mathtt{f}}}\,2^{-n_{\mathtt{f}\mathtt{f}}/2}\,n!\,(n_{0\mathtt{f}}^\bullet + n_{\mathtt{f}!\mathtt{f}} - 1)!!\,(n_{\mathtt{f}0}^\bullet + n_{\mathtt{f}!\mathtt{f}} - 1)!!}{\left(\prod_{\omega \in \{0,1\}^2} n_\omega!\right)(n_{\mathtt{f}\mathtt{f}}/2)!\,n_{1\mathtt{f}}!\,n_{\mathtt{f}1}!\,n_{0\mathtt{f}}^\bullet!\,n_{\mathtt{f}0}^\bullet!\,\prod_{i=0}^{(n_{\mathtt{f}}+n_{\mathtt{f}!\mathtt{f}})/2-1}(n - 1/d - 2i/d)}.$$

Edges from 1-vertices.
Fix any matching $\underline{m}$ on the $\mathtt{f}$-vertices, and let $\boldsymbol{E}_\omega$ denote the number of unmatched half-edges incident to V_ω given the matching on the $\mathtt{f}$-vertices; the total is $\boldsymbol{E} = \sum_\omega \boldsymbol{E}_\omega = nd - n_{\mathtt{f}} - n_{\mathtt{f}!\mathtt{f}}$. Consider the assignment of the remaining $\boldsymbol{E}_1 = n_1 d - n_{1\mathtt{f}} - n_{\mathtt{f}1}$ unmatched half-edges from vertices with spin 1 in either coordinate. Recalling $\mathtt{Z} \equiv \{0,\mathtt{f}\}$, define $\chi_{K,F}$ to be the indicator that there are K edges between V_{10} and V_{01}, F_1 edges between $V_{1\mathtt{Z}}$ and $V_{0\mathtt{f}}$, and F_2 edges between $V_{\mathtt{Z}1}$ and $V_{\mathtt{f}0}$; and let $\mathbf{j}_{1,K,F}(\pi,\underline{m}) \equiv \mathbb{E}[\mathbf{J}_1^\pi \chi_{K,F} \,|\, \underline{m}]$. Then, writing $F_\bullet \equiv F_1 + F_2$,

$$\mathbf{j}_{1,K,F}(\pi,\underline{m}) = \frac{(\boldsymbol{E}_{1\mathtt{Z}})_{F_1}(\boldsymbol{E}_{0\mathtt{f}})_{F_1}}{(F_1)!}\,\frac{(\boldsymbol{E}_{\mathtt{Z}1})_{F_2}(\boldsymbol{E}_{\mathtt{f}0})_{F_2}}{(F_2)!}\,\frac{(\boldsymbol{E}_{10})_K(\boldsymbol{E}_{01})_K}{K!}\,\frac{(\boldsymbol{E}_{00})_{\boldsymbol{E}_1 - 2K - F_\bullet}}{[\boldsymbol{E}]_{\boldsymbol{E}_1 - K}}$$

$$= \exp\{O(F_\bullet \tfrac{\log d}{d})\}\,\underbrace{\frac{(\boldsymbol{E}_{1\mathtt{Z}})_{F_1}(\boldsymbol{E}_{0\mathtt{f}})_{F_1}(\boldsymbol{E}_{\mathtt{Z}1})_{F_2}(\boldsymbol{E}_{\mathtt{f}0})_{F_2}}{(F_1)!(F_2)!(\boldsymbol{E}_{00} - \boldsymbol{E}_1)^{F_\bullet}}}_{\mathbf{b}_{1,F}(\pi,\underline{m})} \times \underbrace{\frac{(\boldsymbol{E}_{10})_K(\boldsymbol{E}_{01})_K(\boldsymbol{E}_{00})_{\boldsymbol{E}_1 - 2K}}{K![\boldsymbol{E}]_{\boldsymbol{E}_1 - K}}}_{\tilde{\mathbf{a}}_{1,K}(\pi,\underline{m})}.$$

Let $\mathbf{a}_{1,K}(\pi,\underline{m}) \equiv \tilde{\mathbf{a}}_{1,K}(\pi,\underline{m})/\mathbf{g}_1(\boldsymbol{E},\boldsymbol{E}_1,\boldsymbol{E}_{00})$ where $\mathbf{g}_1$ is as in (14). We then estimate

$$\mathbf{a}_{1,K}(\pi,\underline{m})\left[\frac{(\boldsymbol{E}_{10}\boldsymbol{E}_{01})^K}{K!(\boldsymbol{E}_{00} - \boldsymbol{E}_1)^K}\right]^{-1} \begin{cases} \leqslant \exp\{O(K\beta_{\max})\} & \text{for all } K; \\ \geqslant \exp\{O(K\tfrac{\log d}{d})\} & \text{for } K \lesssim K^\star \equiv \boldsymbol{E}_{10}\boldsymbol{E}_{01}/\boldsymbol{E}. \end{cases} \tag{50}$$

Similarly, writing $F^\star \equiv (F_1, F_2)^\star \equiv (\boldsymbol{E}_{1\mathtt{Z}}\boldsymbol{E}_{0\mathtt{f}}/\boldsymbol{E}, \boldsymbol{E}_{\mathtt{Z}1}\boldsymbol{E}_{\mathtt{f}0}/\boldsymbol{E})$, we have

$$\mathbf{b}_{1,F}(\pi,\underline{m})\left[\frac{(\boldsymbol{E}_{1\mathtt{Z}}\boldsymbol{E}_{0\mathtt{f}})^{F_1}(\boldsymbol{E}_{\mathtt{Z}1}\boldsymbol{E}_{\mathtt{f}0})^{F_2}}{(F_1)!(F_2)!(\boldsymbol{E}_{00} - \boldsymbol{E}_1)^{F_\bullet}}\right]^{-1} \begin{cases} \leqslant 1 & \text{for all } F, \\ \geqslant \exp\{O(F_\bullet \tfrac{\log d}{d})\} & \text{for } F \lesssim F^\star. \end{cases}$$

Forcing of 0-vertices.
Recall (13) that $\mathbf{g}_0(n, nd\zeta_{1\cdot})$ denotes the probability, with respect to a uniformly random assignment of $nd\zeta_{1\cdot}$ half-edges to n degree-d vertices, that each vertex is forced (that is, receives at least two of the incoming half-edges). Similarly, define the following:

1. For $n^\bullet \leqslant n$, suppose each variable $1 \leqslant i \leqslant n^\bullet$ has degree d, while each variable $n^\bullet < i \leqslant n$ has already received one forcing half-edge, and now has degree $d - 1$. Let $\mathbf{g}_0^\bullet(n, n^\bullet, nd\zeta_{1\cdot})$ denote the probability, with respect to a uniformly random assignment of $nd\zeta_{1\cdot}$ half-edges to the n vertices, that each vertex is forced (thus, each vertex $n^\bullet < i \leqslant n$ is required only to receive at least one incoming half-edge).

2. Writing $\zeta \equiv (\zeta_{11}, \zeta_{10}, \zeta_{01})$, let $\mathbf{g}_{00}(n, nd\zeta)$ denote the probability, with respect to a uniformly random assignment of $nd\zeta$ half-edges to n degree-d vertices, that each vertex is forced in both coordinates (that is, receives at least two of the incoming half-edges both from $\{11, 10\}$ and from $\{11, 01\}$).

Then $\mathbf{j}_{0,K,F}(\pi, \underline{m}) \equiv \mathbb{E}[\mathbf{J}_0^\pi \,|\, \underline{m}, \mathbf{J}_1^\pi \chi_{K,F} > 0]$ can be written as

$$\mathbf{j}_{0,K,F}(\pi, \underline{m}) = \overbrace{\mathbf{g}_0(n_{10}, K)\,\mathbf{g}_0(n_{01}, K)}^{(\text{A})} \times \overbrace{\mathbf{g}_0^\bullet(n_{0\mathtt{f}}, n_{0\mathtt{f}}^\bullet, F_1)\,\mathbf{g}_0^\bullet(n_{\mathtt{f}0}, n_{\mathtt{f}0}^\bullet, F_2)}^{(\text{B})}$$
$$\times \underbrace{\mathbf{g}_{00}(n_{00}, (\boldsymbol{E}_{11}, \boldsymbol{E}_{1\mathtt{Z}} - K - F_1, \boldsymbol{E}_{\mathtt{Z}1} - K - F_2))}_{(\text{C})} .$$

The functions $\mathbf{g}_0$ and $\mathbf{g}_{00}$ were estimated in Propn. 2.9. It is clear that $\mathbf{g}_0^\bullet$ satisfies the same estimates as $\mathbf{g}_0$, so we conclude

(A) is $\exp\{O((n_{10} + n_{01})d^{-4\varepsilon_0/5})\}$ for $K \geqslant {}^9\!/_{10} \cdot \boldsymbol{E}_{10}\boldsymbol{E}_{01}/\boldsymbol{E}$;
(B) is $\exp\{O((n_{0\mathtt{f}} + n_{\mathtt{f}0})d^{-4\varepsilon_0/5})\}$ for $F_1 \geqslant {}^9\!/_{10} \cdot \boldsymbol{E}_{1\mathtt{Z}}\boldsymbol{E}_{0\mathtt{f}}/\boldsymbol{E}$, $F_2 \geqslant {}^9\!/_{10} \cdot \boldsymbol{E}_{\mathtt{Z}1}\boldsymbol{E}_{\mathtt{f}0}/\boldsymbol{E}$;
(C) is $\exp\{O(nd^{-1.65})\}$ for $K \lesssim n_{10} \wedge n_{01}$.

Comparison of frozen model with independent set.
From the preceding estimates, $\arg\max_K \mathbf{a}_{1,K}(\pi, \underline{m})$ must be close to $K^\star$. If K deviates from $K^\star$ by more than a $(1 \pm \varepsilon)$-factor for ε a small constant uniform in d, then the resulting decrease in $\mathbf{a}_{1,K}(\pi, \underline{m})$ will be much larger than any possible gain in $\mathbf{j}_{0,K,F}(\pi, \underline{m})$:

$$\mathbf{a}_{1,K}(\pi, \underline{m}) / \mathbf{a}_{1,K^\star}(\pi, \underline{m}) \leqslant \exp\{-\tfrac{2}{3}\varepsilon^2 \cdot K^\star\} \leqslant \exp\{-\tfrac{1}{2}\varepsilon^2 \cdot n(\varepsilon_0 \log d)^2/d\},$$
$$\mathbf{j}_{0,K,F}(\pi, \underline{m}) / \mathbf{j}_{0,K^\star,F}(\pi, \underline{m}) \leqslant 1 / \mathbf{j}_{0,K^\star,L}(\pi, \underline{m}) \leqslant \exp\{O(n(\varepsilon_0 \log d)/d)\}.$$

Therefore we may restrict consideration to ${}^9\!/_{10} \cdot K^\star \equiv \underline{K} \leqslant K \leqslant \overline{K} \equiv {}^{11}\!/_{10} \cdot K^\star$, since the contribution to $\mathbb{E}[{}_2\boldsymbol{Z}(\pi)]$ from K outside this regime must be an exponentially small fraction of the total. By the same reasoning we may similarly restrict $F \equiv (F_1, F_2)$ to lie in the regime ${}^9\!/_{10} \cdot F^\star \equiv \underline{F} \leqslant F \leqslant \overline{F} \equiv {}^{11}\!/_{10} \cdot F^\star$, therefore

$$\mathbb{E}[{}_2\boldsymbol{Z}(\pi)] \asymp \mathbf{c}(\pi)\,\mathbf{j}_{\mathtt{f}}(\pi)\,\mathbb{E}_{\underline{m}}\left[\sum_{\underline{K} \leqslant K \leqslant \overline{K}} \sum_{\underline{F} \leqslant F \leqslant \overline{F}} \mathbf{j}_{1,K,F}(\pi, \underline{m})\,\mathbf{j}_{0,K,F}(\pi, \underline{m}) \right] \tag{51}$$

where $\mathbb{E}_{\underline{m}}$ denotes the expectation over the random matching $\underline{m}$, conditioned on $\mathbf{J}_{\mathtt{f}}^\pi$ being positive. For any $\pi \in \mathcal{M}(n\alpha, n\rho)$, consider the measure $\overline{\pi} \in M(n\alpha, n\rho)$ defined by

$$\overline{\pi}_{11} = \pi_{11}, \quad \overline{\pi}_{10} = \pi_{10} + \pi_{1\mathtt{f}} + \tfrac{1}{2}\pi_{\mathtt{f}\cdot}, \quad \overline{\pi}_{01} = \pi_{01} + \pi_{\mathtt{f}1} + \tfrac{1}{2}\pi_{\cdot\mathtt{f}},$$

with the remaining probability going to $\overline{\pi}_{00}$. Then

$$\frac{\mathbb{E}[{}_2\boldsymbol{Z}(\pi)]}{\mathbb{E}[{}_2Z(\overline{\pi})]} = \frac{\mathbf{c}(\pi)\,\mathbf{j}_{\mathtt{f}}(\pi)}{\mathbf{c}(\overline{\pi})} \frac{\mathbb{E}_{\underline{m}}[\sum_{K,F} \mathbf{j}_{1,K,F}(\pi, \underline{m})\mathbf{j}_{0,K,F}(\pi, \underline{m})]}{\sum_K \mathbf{j}_{1,K,0}(\overline{\pi})},$$

and it follows from the above estimates that this is $d^{O(n\beta_{\max})} \exp\{O(nd^{-1.65} + n_1 d^{-3\varepsilon_0/4})\}$, implying the result. $\qquad\square$

Corollary 4.3. *Suppose $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ and $\varepsilon_0 > 0$ a small positive constant uniform in d. For $d \geqslant d_0(\varepsilon_0)$, $_2\Phi$ can only attain its global maximum on $_2\Delta_\alpha$ either in the near-independence regime $^\star_2\Delta_\alpha \subseteq {}_2\Delta_\alpha$ of measures with overlap $\rho \leqslant d^{-1}$, or in the near-identical regime $^{\cdot\cdot}_2\Delta_\alpha \subseteq {}_2\Delta_\alpha$ of measures with overlap $\alpha - \varepsilon_0(\log d)/d \leqslant \rho \leqslant \alpha$.*

Proof. We shall prove that for $d \geqslant d_0(\varepsilon_0)$ and $n \geqslant n_0(d)$, the ratio of $\mathbb{E}[_2Z_{n\alpha}(n\rho)]$ to $\max_r \mathbb{E}[_2Z_{n\alpha}(nr)]$ is exponentially small in n for all $d^{-1.45} \leqslant \rho \leqslant \alpha - \varepsilon_0(\log d)/d$. We begin with an analogous calculation in the independent set model: there is no matching on the f-vertices and no forcing constraint on the 0-vertices, so

$$\frac{\mathbb{E}[_2Z_{n\alpha}(n\rho)]}{\mathbb{E}[Z_{n\alpha}]} = \sum_{\pi \in M(n\alpha, n\rho)} \frac{\mathbf{c}(\pi)\mathbb{E}[\mathbf{J}_1^\pi]}{\binom{n}{n\alpha}\mathbf{g}_1(nd, nd\alpha, 0)} = \sum_{\pi \in M(n\alpha, n\rho)} \frac{\mathbf{c}(\pi)\mathbf{g}_1(nd, n_1 d, 0)\sum_K \mathbf{a}_{1,K}(\pi)}{\binom{n}{n\alpha}\mathbf{g}_1(nd, nd\alpha, 0)},$$

where $\mathbf{a}_{1,K}(\pi) \equiv \mathbf{a}_{1,K}(\pi, \varnothing)$ for $\mathbf{a}_{1,K}(\pi, \underline{m})$ as in (50). For $n_1 \lesssim n(\log d)/d$ we have

$$\mathbf{g}_1(nd, n_1 d, 0) = \exp\{-\tfrac{1}{2}n_1^2 d/n + O(nd^{-2}(\log d)^3)\}.$$

The estimates of Propn. 4.2 imply

$$\sum_K \mathbf{a}_{1,K}(\pi) = \exp\{O(K\tfrac{\log d}{d})\} \sum_K \frac{(n_{10}n_{01}d)^K}{K!(n_{00} - n_1)^K} = \exp\{n_{10}n_{01}d/n + O(nd^{-2}(\log d)^3)\}.$$

Combining gives

$$\mathbb{E}[_2Z_{n\alpha}(n\rho)] = \mathbb{E}[Z_{n\alpha}]\exp\{ng_\alpha(\rho) + O(nd^{-2}(\log d)^3)\} \quad \text{where}$$
$$g_\alpha(\rho) \equiv \alpha H(\tfrac{\rho}{\alpha}) + (1-\alpha)H(\tfrac{\alpha-\rho}{1-\alpha}) - \tfrac{d}{2}n\alpha^2 + \tfrac{d}{2}n\rho^2.$$

The function g_α has first derivative $(g_\alpha)'(\rho) = 2\log(\alpha - \rho) - \log\rho - \log(1 - 2\alpha + \rho) + d\rho$; differentiating again gives $(g_\alpha)''(\rho) = d - 2(\alpha - \rho)^{-1} - \rho^{-1} - (1 - 2\alpha + \rho)^{-1}$, so we see that g_α is strictly convex on the interval $2/d \leqslant \rho \leqslant \alpha - 3/d$. From the expression for $(g_\alpha)'$ we see that the (unique) minimizer $\rho_\circ$ of g_α on this interval must lie near $(\log d)/d$, and so applying Propn. 4.2 gives

$$\frac{\sup_{2d^{-1.48} \leqslant \rho \leqslant \rho_\circ} \mathbb{E}[_2\mathbf{Z}_{n\alpha}(n\rho)]}{\mathbb{E}[_2\mathbf{Z}_{n\alpha}(nd^{-1.48})]} = \exp\{O(nd^{-1.49})\}\frac{\sup_{2d^{-1.48} \leqslant \rho \leqslant \rho_\circ} \mathbb{E}[_2\mathbf{Z}_{n\alpha}(n\rho)]}{\mathbb{E}[_2\mathbf{Z}_{n\alpha}(nd^{-1.48})]}; \quad (52)$$

$$\frac{\sup_{\rho_\circ \leqslant \rho \leqslant \alpha - \varepsilon_0(\log d)/d} \mathbb{E}[_2\mathbf{Z}_{n\alpha}(n\rho)]}{\mathbb{E}[_2\mathbf{Z}_{n\alpha}(n\alpha - n\varepsilon_0(\log d)/(2d))]} = \exp\{O(n/d)\}\frac{\sup_{\rho_\circ \leqslant \rho \leqslant \alpha - \varepsilon_0(\log d)/d} \mathbb{E}[_2\mathbf{Z}_{n\alpha}(n\rho)]}{\mathbb{E}[_2\mathbf{Z}_{n\alpha}(n\alpha - n\varepsilon_0(\log d)/(2d))]}. \quad (53)$$

We estimate $(g_\alpha)'(\rho) \leqslant -(\log d)/10$ for $d^{-1.85}\rho \leqslant 100/d$, and similarly $(g_\alpha)'(\rho) \geqslant (\log d)/10$ for $1.6(\log d)/d \leqslant \rho \leqslant \alpha - d^{-1.25}$. Thus

(52) is $\leqslant \exp\{n\left[g_\alpha(2d^{-1.48}) - g_\alpha(d^{-1.48})\right] + O(nd^{-1.49})\} \leqslant \exp\{-nd^{-1.48}\}$,

(53) is $\leqslant \exp\{n[g_\alpha(n\alpha - n\varepsilon_0(\log d)/d) - g_\alpha(n\alpha - n\varepsilon_0(\log d)/(2d))] + O(n/d)\}$

$\leqslant \exp\{-n\varepsilon_0(\log d)/d\}.$

These estimates cover the entire interval $d^{-1.45} \leqslant \rho \leqslant \alpha - \varepsilon_0(\log d)/d$, implying the result. $\quad\square$

4.2. **Boundary estimates.**

Proposition 4.4. *For $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ the following hold:*
(a) For $d \geqslant d_0$, $n \geqslant n_0(d)$, $\arg\max_{0 \leqslant \beta \leqslant \beta_{\max}} \mathbb{E}[\mathbf{Z}_{n(\alpha - \beta/2), n\beta}]$ is positive.

(b) Suppose $0 \leqslant \rho \leqslant \alpha - \varepsilon_0(\log d)/d$ for ε_0 a small positive constant uniform in d. Then for $d \geqslant d_0(\varepsilon_0)$, $n \geqslant n_0(d)$, any measure $\pi \in \mathcal{M}(n\alpha, n\rho)$ maximizing $\mathbb{E}[_2\boldsymbol{Z}(\pi)]$ has full support on $\mathscr{P}$; further $\pi_{\mathtt{0f}}^{\bullet} \equiv \pi_{\mathtt{0f}} - \pi_{\mathtt{1f}}$ and $\pi_{\mathtt{f0}}^{\bullet} \equiv \pi_{\mathtt{f0}} - \pi_{\mathtt{f1}}$ must be positive.

Proof. Suppose that $\pi_\omega \equiv n_\omega/n$ is the maximizer of $\mathbb{E}[_2\boldsymbol{Z}(\pi)]$ on $\mathcal{M}(n\alpha, n\rho)$. We prove (b) by a series of comparison estimates. Write $\delta \equiv n/e^{d^2}$.

1. Suppose $\pi_{\mathtt{0f}} = 0$, and consider the measure

$$\overline{\pi} = \pi + (\delta/n)\,(2 \cdot \mathbf{1}\{\mathtt{0f}\} - \mathbf{1}\{\mathtt{00}, \mathtt{01}\}) \in \mathcal{M}(n\alpha, n\rho). \tag{54}$$

Let $\underline{m}$ denote any matching on the $\mathtt{f}$-vertices for the π-configuration; clearly $\boldsymbol{E}_{\mathtt{0f}} = F_1 = 0$. Likewise let $\overline{m}$ denote any matching on the $\mathtt{f}$-vertices for the $\overline{\pi}$-configuration; note that in this case $0 \leqslant F_1 \leqslant \overline{\boldsymbol{E}}_{\mathtt{0f}} = 2\delta d - O(\delta)$. We estimate

$$\frac{\mathbf{j}_{1,K,(F_1,F_2)}(\overline{\pi}, \overline{m})}{\mathbf{j}_{1,K,(0,F_2)}(\pi, \underline{m})} = \frac{(\boldsymbol{E}_{\mathtt{1Z}})_{F_1}(\overline{\boldsymbol{E}}_{\mathtt{0f}})_{F_1}}{F_1!\,\boldsymbol{E}^{F_1}\,d^{O(\delta)}} \frac{(\boldsymbol{E}_{\mathtt{Z1}} - \delta d)_{F_2}}{(\boldsymbol{E}_{\mathtt{Z1}})_{F_2}} \frac{(\boldsymbol{E}_{\mathtt{01}} - \delta d)_K}{(\boldsymbol{E}_{\mathtt{01}})_K} \overset{\star}{=} d^{O(\delta)},$$

$$\frac{\mathbf{j}_{0,K,(F_1,F_2)}(\overline{\pi}, \overline{m})}{\mathbf{j}_{0,K,(0,F_2)}(\pi, \underline{m})} = e^{O(\delta)}\,\mathbf{g}_0^{\bullet}(\overline{n}_{\mathtt{0f}}, \overline{n}_{\mathtt{0f}}, F_1) = e^{O(\delta)}\,\mathbf{g}_0(\delta, F_1) \overset{\star}{=} e^{O(\delta)}$$

where $\star$ holds for K, F appearing in the sum (51) with respect to the $\overline{\pi}$-configuration. Thus

$$\frac{\mathbb{E}[_2\boldsymbol{Z}(\overline{\pi})]}{\mathbb{E}[_2\boldsymbol{Z}(\pi)]} = \frac{\mathbf{c}(\overline{\pi})\,\mathbf{j}_{\mathtt{f}}(\overline{\pi})}{\mathbf{c}(\pi)\,\mathbf{j}_{\mathtt{f}}(\pi)} \underbrace{\frac{\mathbb{E}_{\overline{m}}[\sum_{K,F}\mathbf{j}_{1,K,F}(\overline{\pi}, \overline{m})\,\mathbf{j}_{0,K,F}(\overline{\pi}, \overline{m})]}{\mathbb{E}_{\underline{m}}[\sum_{K,F_2}\mathbf{j}_{1,K,(0,F_2)}(\pi, \underline{m})\,\mathbf{j}_{0,K,(0,F_2)}(\pi, \underline{m})]}}_{\geqslant d^{O(\delta)}} \geqslant \frac{(n/\delta)^{\delta}}{d^{O(\delta)}} \gg 1,$$

contradicting the assumption that π was a maximizer $\mathbb{E}[_2\boldsymbol{Z}(\pi)]$ on $\mathcal{M}(n\alpha, n\rho)$.

2. Since the preceding estimate shows that $\pi_{\mathtt{0f}}$ and $\pi_{\mathtt{f0}}$ must be positive, we must have $\delta' \equiv (n_{\mathtt{0f}} \wedge n_{\mathtt{f0}})/e^{d^2} \gg 0$ scaling linearly with n. Suppose $\pi_{\mathtt{0f}}^{\bullet} = 0$: if we let $\overline{\pi}$ be defined by (54) with δ' in place of δ, then a simple term-by-term comparison gives

$$\frac{\mathbb{E}[_2\boldsymbol{Z}(\overline{\pi})]}{\mathbb{E}[_2\boldsymbol{Z}(\pi)]} = e^{O(\delta')}\frac{\mathbf{c}(\overline{\pi})\,\mathbf{j}_{\mathtt{f}}(\overline{\pi})}{\mathbf{c}(\pi)\,\mathbf{j}_{\mathtt{f}}(\pi)} \geqslant \frac{(n/\delta')^{\delta'}}{d^{O(\delta')}} \gg 1,$$

again a contradiction. Similarly, if $\pi_{\mathtt{ff}} = 0$, then $\mathbb{E}[_2\boldsymbol{Z}(\pi)] \ll \mathbb{E}[_2\boldsymbol{Z}(\overline{\pi})]$

$$\overline{\pi} = \pi + (\delta'/n)\,(2 \cdot \mathbf{1}\{\mathtt{ff}\} - \mathbf{1}\{\mathtt{10}, \mathtt{01}\}) \in \mathcal{M}(n\alpha, n\rho).$$

This proves that $\pi_{\mathtt{0f}}^{\bullet}$, $\pi_{\mathtt{f0}}^{\bullet}$, and $\pi_{\mathtt{ff}}$ must be positive.

3. Let $\delta'' \equiv (n_{\mathtt{0f}}^{\bullet} \wedge n_{\mathtt{f0}}^{\bullet} \wedge n_{\mathtt{ff}})/e^{d^2}$; by the preceding estimates we have $\delta'' \gg 0$ scaling linearly with n. By similar calculations as before, if $\pi_{\mathtt{f!f}} = 0$ then $\mathbb{E}[_2\boldsymbol{Z}(\pi)] \ll \mathbb{E}[_2\boldsymbol{Z}(\overline{\pi})]$ for

$$\overline{\pi} = \pi + (\delta''/n)\,(2 \cdot \mathbf{1}\{\mathtt{f!f}\} - \mathbf{1}\{\mathtt{0f}, \mathtt{f0}\}) \in \mathcal{M}(n\alpha, n\rho).$$

If $\pi_{\mathtt{1f}} = 0$, then $\mathbb{E}[_2\boldsymbol{Z}(\pi)] \ll \mathbb{E}[_2\boldsymbol{Z}(\overline{\pi})]$ for

$$\overline{\pi} = \pi + (\delta''/n)\,(\mathbf{1}\{\mathtt{1f}, \mathtt{0f}\} - 2 \cdot \mathbf{1}\{\mathtt{ff}\}) \in \mathcal{M}(n\alpha, n\rho).$$

This concludes the proof of (b). The first-moment estimate (a) follows simply by restricting consideration to π supported on $\{\mathtt{00}, \mathtt{11}, \mathtt{ff}\}$. $\qquad \square$

Proposition 4.5. *For $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ the following hold:*

(a) For $d \geqslant d_0$, any global maximizer of $\boldsymbol{\Phi}$ on $\boldsymbol{\Delta}_\alpha$ must be strictly positive on $\operatorname{supp}\varphi$.

(b) Fix any small constant $\varepsilon_0 > 0$ uniform in d. For $d \geqslant d_0(\varepsilon_0)$, any global maximizer of $_2\boldsymbol{\Phi}$ on $_2\boldsymbol{\Delta}_\alpha$ which lies outside of $\ddot{_2}\boldsymbol{\Delta}_\alpha$ must be strictly positive on $\operatorname{supp}\varphi_2$.

Proof. We shall prove (b); the proof of (a) is similar but simpler.

Boundary derivative of rate function.
As we have noted before, the functional form of $_2\Phi$ implies that the optimal $\boldsymbol{h}$ in $_2\Delta_\alpha$ must be symmetric, with

$$_2\Phi(\boldsymbol{h}) = \sum_{\dot{\underline{\tau}}} \dot{h}(\dot{\underline{\tau}}) \log \dot{\varphi}_2(\dot{\underline{\tau}}) + H(\dot{\boldsymbol{h}}) - \tfrac{d}{2}H(\bar{h})$$

For $\boldsymbol{\delta} \equiv (\dot{\boldsymbol{\delta}}, \hat{\boldsymbol{\delta}})$ such that $\boldsymbol{h} + t\boldsymbol{\delta}$ is also symmetric and lies in $_2\Delta_\alpha$ for $t \geqslant 0$ small, consider

$$\mathbf{T}(_2\Phi)(\boldsymbol{h}; \boldsymbol{\delta}) \equiv \lim_{t\downarrow 0} \frac{_2\Phi(\boldsymbol{h} + t\boldsymbol{\delta}) - {}_2\Phi(\boldsymbol{h})}{t \log(1/t)} = \dot{\boldsymbol{\delta}}[(\operatorname{supp} \dot{\boldsymbol{h}})^c] - (d/2)\, \bar{\delta}[(\operatorname{supp} \bar{h})^c].$$

To show that $\boldsymbol{h}$ is not a maximizer it suffices to exhibit $\mathbf{T}(_2\Phi)(\boldsymbol{h}; \boldsymbol{\delta}) > 0$ for some $\boldsymbol{\delta}$. In particular, it follows by convexity that for any $\boldsymbol{h} \in {}_2\Delta_\alpha$, $\boldsymbol{h} + t(^\star_2\boldsymbol{h}_\alpha - \boldsymbol{h}) \in {}_2\Delta^\circ_\alpha$ for $t > 0$ small and $^\star_2\boldsymbol{h}_\alpha$ as in the statement of Thm. 4.1. Therefore, if $\boldsymbol{h}$ is a maximizer such that the edge marginal has full support $\operatorname{supp} \bar{h} = \mathscr{M}^2$, then necessarily $\operatorname{supp} \boldsymbol{h} = \operatorname{supp} \varphi_2$, since otherwise $\mathbf{T}(_2\Phi)(\boldsymbol{h}; ^\star_2\boldsymbol{h}_\alpha - \boldsymbol{h}) > 0$.

Positive maximizer for vertex-auxiliary model.
Clearly the preceding calculation applies equally well to the rate function $^\mathsf{v}_2\Phi$ of the pair version of the *vertex-auxiliary* model introduced in Rmk. 3.1, with factors $^\mathsf{v}\varphi_2 \equiv {}^\mathsf{v}\varphi \otimes {}^\mathsf{v}\varphi$. We now show that the maximizer $^\mathsf{v}\boldsymbol{h}$ of $^\mathsf{v}_2\Phi$ in the regime $\rho \leqslant \alpha - \varepsilon_0(\log d)/d$ satisfies $\operatorname{supp} {}^\mathsf{v}\bar{h} = \mathscr{M}^2$, hence (from the above) $\operatorname{supp} {}^\mathsf{v}\boldsymbol{h} = \operatorname{supp} {}^\mathsf{v}\varphi_2$.

By Propn. 4.4b, $^\mathsf{v}\boldsymbol{h}$ must satisfy $^\mathsf{v}\bar{h}\binom{r}{s} > 0$ for r, s any of the sets in the partition $\{\mathtt{10}\}$, $\{\mathtt{01}\}$, $\{\overline{\mathtt{ff}}\}$, $\mathtt{0Z} = \{\mathtt{00}, \mathtt{0f}\}$, $\mathtt{fZ} = \{\mathtt{f0}, \mathtt{ff}\}$ (e.g. $^\mathsf{v}\bar{h}\binom{\overline{\mathtt{ff}}}{\mathtt{fZ}} > 0$ due to the presence of $\mathtt{f!f}$-variables. If $\tau \equiv (\sigma^1, \sigma^2) \notin \operatorname{supp} {}^\mathsf{v}\bar{h}$, then $\acute{\tau} \equiv (\acute{\sigma}^1, \acute{\sigma}^2) \in \operatorname{supp} {}^\mathsf{v}\bar{h}$ for some choice of $\acute{\sigma}^i$ belonging to the same subset as σ^i in this partition. Therefore $(\acute{\tau}, \underline{\tau}) \in \operatorname{supp} {}^\mathsf{v}\boldsymbol{h}$ for some $\underline{\tau} \in \mathscr{M}^{2(d-1)}$, so $(\tau, \underline{\tau}) \in \operatorname{supp} {}^\mathsf{v}\dot{\varphi}_2 \backslash \operatorname{supp} {}^\mathsf{v}\dot{\boldsymbol{h}}$. By the symmetry $^\mathsf{v}\bar{h}(\sigma) = {}^\mathsf{v}\bar{h}(\mathsf{R}\sigma)$, there must exist $\underline{\tau}' \in \mathscr{M}^{2(d-1)}$ such that $(\mathsf{R}\acute{\tau}, \underline{\tau}') \in \operatorname{supp} {}^\mathsf{v}\boldsymbol{h}$. If τ belongs to $\mathtt{ZZ} \times (\mathtt{ZZ} \cup \{\overline{\mathtt{ff}}, \mathtt{10}, \mathtt{01}\})$, then $(\mathsf{R}\tau, \underline{\tau}')$ must be in $\operatorname{supp} {}^\mathsf{v}\dot{\varphi}_2 \backslash \operatorname{supp} {}^\mathsf{v}\dot{\boldsymbol{h}}$. Taking

$$^\mathsf{v}\dot{\boldsymbol{\delta}} = \mathbf{1}_{(\tau, \underline{\tau})} - \mathbf{1}_{(\acute{\tau}, \underline{\tau})} + \mathbf{1}_{(\mathsf{R}\tau, \underline{\tau}')} - \mathbf{1}_{(\mathsf{R}\acute{\tau}, \underline{\tau}')}, \quad ^\mathsf{v}\bar{\delta} = \tfrac{1}{d}(\mathbf{1}_\tau - \mathbf{1}_{\acute{\tau}} + \mathbf{1}_{\mathsf{R}\tau} - \mathbf{1}_{\mathsf{R}\acute{\tau}})$$

gives $\mathbf{T}(^\mathsf{v}_2\Phi)(^\mathsf{v}\boldsymbol{h}; {}^\mathsf{v}\boldsymbol{\delta}) = 2 - 1 > 0$, proving our claim that the vertex-auxiliary maximizer must have $\operatorname{supp} {}^\mathsf{v}\bar{h} = \mathscr{M}^2$, hence $\operatorname{supp} {}^\mathsf{v}\boldsymbol{h} = \operatorname{supp} {}^\mathsf{v}\varphi_2$.

Positive maximizer for auxiliary model.
It directly follows from $\operatorname{supp} {}^\mathsf{v}\bar{h} = \mathscr{M}^2$ that the maximizer $\boldsymbol{h}$ for the original auxiliary model on $_2\Delta_\alpha \backslash \ddot{2}\Delta_\alpha$ satisfies $\bar{h}\binom{r}{s} > 0$ for r, s any of the sets in the partition $\{\mathtt{00}\}$, $\{\mathtt{0f}\}$, $\{\mathtt{f0}\}$, $\{\mathtt{ff}\}$, $\{\mathtt{11}\}$, $\mathtt{1Z}$, $\mathtt{Z1}$. Further, since $\operatorname{supp} {}^\mathsf{v}\dot{\boldsymbol{h}} = \operatorname{supp} {}^\mathsf{v}\dot{\varphi}_2$ we may replace $\mathtt{Z1}$ in the above list by $\{\mathtt{01}\}$, $\{\mathtt{f1}\}$. By symmetry, the only remaining possibility for $\tau \notin \operatorname{supp} \bar{h}$ is that either τ or $\mathsf{R}\tau$ belongs to $\mathtt{1Z} \times \mathtt{Z1}$. Consider the Gibbs measure $^\mathsf{v}\boldsymbol{\nu}$ associated to $^\mathsf{v}\boldsymbol{h}$ on the subtree shown in Fig. 4: clearly, all combinations of j, j' appear with positive $^\mathsf{v}\boldsymbol{\nu}$-probability. Thus $\mathtt{1Z} \times \mathtt{Z1}$, and symmetrically $\mathtt{Z1} \times \mathtt{1Z}$, must belong in the support of $\bar{h}$, concluding our proof that $\operatorname{supp} \boldsymbol{h} = \operatorname{supp} \varphi_2$. $\qquad\square$

Recall (8) and Defn. 3.2 that we have truncated the frozen model and the spaces $\Delta_\alpha, {}_2\Delta_\alpha$ by restricting the density of $\mathtt{f}$-variables, so in order to establish Φ and $_2\Phi$ have *interior* global maximizers in Δ_α and $_2\Delta_\alpha$ respectively, we must verify in addition to Propn. 4.5 that the maximizer does not occur near the boundary with density $\beta_{\max}$ of $\mathtt{f}$-variables; this will be done in §4.3.
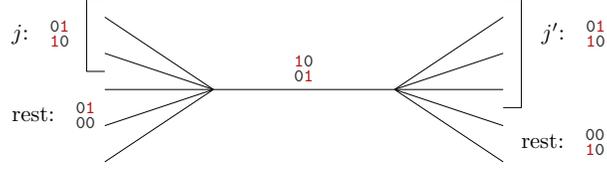
FIGURE 4. Any combination of j, j' must appear in vertex-auxiliary model

4.3. Near-independence regime.

In this subsection we complete our analysis of the near-independence regime $_2^\star\mathbf{\Delta}_\alpha$ to prove

Proposition 4.6. *The unique global maximizer of the restriction of $_2\mathbf{\Phi}$ to $_2^\star\mathbf{\Delta}_\alpha$ is $_2^\star\mathbf{h}_\alpha$.*

Recall that in the pair frozen model, $\mathcal{M}(n\alpha, n\rho)$ denotes the space of empirical measures π on $\mathscr{P}$ which contribute to $\mathbb{E}[\mathbf{Z}_{n\alpha}^2]$ and have $\pi_{11} = \rho$.

Lemma 4.7. *For $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ let y be defined by $\alpha \equiv y(\log d)/d$. If ε is any small positive constant uniform in d, then the following hold for $d \geqslant d_0(\varepsilon)$, $n \geqslant n_0(d)$:*

(a) The maximizer of $\mathbb{E}[\mathbf{Z}_{\alpha-\beta/2,\beta}]$ over $\beta \leqslant \beta_{\max}$ must satisfy $d^{y-1}\beta/\alpha \leqslant d^\varepsilon$.
(b) For $\rho \leqslant d^{-1}$, the maximizer of $\mathbb{E}[_2\mathbf{Z}(\pi)]$ over $\pi \in \mathcal{M}(n\alpha, n\rho)$ must satisfy

$$\frac{\pi_{\mathtt{1f}}}{\pi_{\mathtt{0f}}\pi_{\mathtt{10}}}, \frac{\pi_{\mathtt{f1}}}{\pi_{\mathtt{f0}}\pi_{\mathtt{01}}}, \frac{d\pi_{\mathtt{ff}}}{\pi_{\mathtt{0f}}\pi_{\mathtt{f0}}}, \frac{\pi_{\mathtt{f!f}}}{\pi_{\mathtt{0f}}\pi_{\mathtt{f0}}}, \frac{d^{y-1}\pi_{\mathtt{0f}}}{\pi_{\mathtt{10}}} \leqslant d^\varepsilon.$$

Proof. Suppose the measure $\pi_\omega \equiv n_\omega/n$ maximizes $\mathbb{E}[_2\mathbf{Z}(\pi)]$ over $\mathcal{M}(n\alpha, n\rho)$. As in the proof of Propn. 4.4, we shall estimate π by comparing it with various nearby measures $\overline{\pi}_\omega \equiv \overline{n}_\omega/n$. Note that if we fix any small positive constant ε which is uniform in d, then (51) continues to hold if the sum is restricted to

$$1 - \varepsilon \leqslant \{K/K^\star, F/F^\star\} \leqslant 1 + \varepsilon. \tag{55}$$

Let $\delta \equiv e^{-d^2}\min\{n_\omega : \omega \in \mathscr{P}\}$; then Propn. 4.4 implies that δ must scale linearly with n.

1. *Comparison of $\mathtt{1f}$ to $\mathtt{0f}$.*
   Let $\overline{\pi} = \pi + (\delta/n)(\mathbf{1}\{\mathtt{10}, \mathtt{0f}\} - \mathbf{1}\{\mathtt{1f}, \mathtt{00}\})$. Let $\underline{m}$ (resp. $\overline{m}$) be any matching of the $\mathtt{f}$-vertices for the π-configuration (resp. $\overline{\pi}$-configuration). For $\rho \lesssim d^{-1}$ and K, F satisfying (55),

$$\frac{\mathbf{j}_{1,K,F}(\overline{\pi}, \overline{m})}{\mathbf{j}_{1,K,F}(\pi, \underline{m})} = e^{O(\delta)}\frac{(\mathbf{E}_{\mathtt{0f}} + \delta d)_{F_1}}{(\mathbf{E}_{\mathtt{0f}})_{F_1}}\frac{(\mathbf{E}_{\mathtt{10}} + \delta d)_K}{(\mathbf{E}_{\mathtt{10}})_K}\frac{(\mathbf{E}_{\mathtt{00}} - \delta d)_{E_1-2K-F_\bullet}}{(\mathbf{E}_{\mathtt{00}})_{E_1-2K-F_\bullet}} \geqslant \frac{e^{O(\delta)}}{d^{O(\delta\varepsilon)}},$$

   where we used that $(A + c)_b/(A)_b = \exp\{(bc/A)[1 + O(x)]\}$ for $|b|, |c| \leqslant xA$. Thus

$$\frac{\mathbb{E}[_2\mathbf{Z}(\overline{\pi})]}{\mathbb{E}[_2\mathbf{Z}(\pi)]} \geqslant \frac{\mathbf{c}(\overline{\pi})\,\mathbf{j}_{\mathtt{f}}(\overline{\pi})}{\mathbf{c}(\pi)\,\mathbf{j}_{\mathtt{f}}(\pi)}\frac{e^{O(\delta)}}{d^{O(\delta\varepsilon)}} \geqslant \left(\frac{\pi_{\mathtt{1f}}}{\pi_{\mathtt{10}}\pi_{\mathtt{0f}}}\frac{e^{O(1)}}{d^{O(\varepsilon)}}\right)^\delta,$$

   so if π is the maximizer then it must satisfy $\pi_{\mathtt{1f}} \leqslant d^{O(\varepsilon)}\pi_{\mathtt{0f}}\pi_{\mathtt{10}}$.
2. *Comparison of $\mathtt{ff}, \mathtt{f!f}$ to $\mathtt{0f}, \mathtt{f0}$.*
   Let $\overline{\pi} = \pi + (\delta/n)(\mathbf{1}\{\mathtt{0f}, \mathtt{f0}\} - \mathbf{1}\{\mathtt{00}, \mathtt{ff}\})$. By a similar calculation as above we find

$$\frac{\mathbb{E}[_2\mathbf{Z}(\overline{\pi})]}{\mathbb{E}[_2\mathbf{Z}(\pi)]} \geqslant \frac{\mathbf{c}(\overline{\pi})\,\mathbf{j}_{\mathtt{f}}(\overline{\pi})}{\mathbf{c}(\pi)\,\mathbf{j}_{\mathtt{f}}(\pi)}\frac{e^{O(\delta)}}{d^{O(\delta\varepsilon)}} \geqslant \left(\frac{d\pi_{\mathtt{ff}}}{\pi_{\mathtt{0f}}\pi_{\mathtt{f0}}}\frac{e^{O(1)}}{d^{O(\varepsilon)}}\right)^{\delta/2},$$

proving that the maximizer π must satisfy $d\pi_{\mathtt{ff}} \leqslant d^{O(\varepsilon)}\pi_{\mathtt{0f}}\pi_{\mathtt{f0}}$. The calculation with $\mathtt{f!f}$ in place of $\mathtt{ff}$ is similar and shows that the maximizer π must satisfy $\pi_{\mathtt{f!f}} \leqslant d^{O(\varepsilon)}\pi_{\mathtt{0f}}\pi_{\mathtt{f0}}$.

3. *Comparison of* $\mathtt{0f}$ *to* $\mathtt{01}$.

Let $\overline{\pi} \equiv \pi + (\delta/n)\,(\mathbf{1}\{\mathtt{00},\mathtt{01}\} - 2\cdot\mathbf{1}\{\mathtt{0f}\})$. Then

$$\frac{\mathbf{j}_{1,K,F}(\overline{\pi},\overline{m})}{\mathbf{j}_{1,K,F}(\pi,m)} = \frac{(\boldsymbol{E}_{\mathtt{0f}}-2\delta d)_{F_1}}{(\boldsymbol{E}_{\mathtt{0f}})_{F_1}}\frac{(\boldsymbol{E}_{\mathtt{Z1}}+\delta d)_{F_2}}{(\boldsymbol{E}_{\mathtt{Z1}})_{F_2}}\frac{(\boldsymbol{E}_{\mathtt{01}}+\delta d)_K}{(\boldsymbol{E}_{\mathtt{01}})_K}\frac{(\boldsymbol{E}_{\mathtt{00}})^{\delta d}\,e^{O(\delta)}}{(\boldsymbol{E}_{\mathtt{00}}-\boldsymbol{E}_1)^{\delta d}} \geqslant \frac{d^{\delta y}e^{O(\delta)}}{d^{O(\delta\varepsilon)}}.$$

From the preceding estimates, $\pi_{\mathtt{1f}} + \pi_{\mathtt{f!f}} \lesssim \pi_{\mathtt{0f}}(\log d)/d$, and so

$$\frac{\mathbb{E}[_2\boldsymbol{Z}(\overline{\pi})]}{\mathbb{E}[_2\boldsymbol{Z}(\pi)]} \geqslant \frac{\mathbf{c}(\overline{\pi})\,\mathbf{j}_{\mathtt{f}}(\overline{\pi})}{\mathbf{c}(\pi)\,\mathbf{j}_{\mathtt{f}}(\pi)}\frac{d^{\delta y}e^{O(\delta)}}{d^{O(\delta\varepsilon)}} \geqslant \left(\frac{n\,n_{\mathtt{0f}}^{\bullet}}{n_{\mathtt{00}}\,n_{\mathtt{01}}}\frac{n_{\mathtt{0f}}^{\bullet}}{n_{\mathtt{0f}}^{\bullet}+n_{\mathtt{f!f}}}\frac{d^y\,e^{O(1)}}{d^{1+O(\varepsilon)}}\right)^{\delta} \geqslant \left(\frac{\pi_{\mathtt{0f}}}{\pi_{\mathtt{01}}}\frac{d^y e^{O(1)}}{d^{1+O(\varepsilon)}}\right)^{\delta},$$

implying that the maximizer π must satisfy $d^{y-1}\pi_{\mathtt{0f}} \leqslant d^{O(\varepsilon)}\pi_{\mathtt{01}}$.

Adjusting ε as needed, this concludes the proof of (b). To prove (a), consider π supported on $\{\mathtt{00},\mathtt{11},\mathtt{ff}\}$ with $\pi_1 = \alpha = y(\log d)/d$ and $\pi_{\mathtt{f}} = \beta$, and let $\overline{\pi} = \pi + (\delta/n)\,(\mathbf{1}\{\mathtt{00},\mathtt{11}\} - 2\cdot\mathbf{1}\{\mathtt{ff}\})$. Clearly we must always have $K = F_1 = F_2 = 0$, and we estimate

$$\frac{\mathbb{E}[_2\boldsymbol{Z}(\overline{\pi})]}{\mathbb{E}[_2\boldsymbol{Z}(\pi)]} \geqslant \frac{\mathbf{c}(\overline{\pi})\,\mathbf{j}_{\mathtt{f}}(\overline{\pi})}{\mathbf{c}(\pi)\,\mathbf{j}_{\mathtt{f}}(\pi)}e^{O(\delta)}d^{\delta y} \geqslant \left(\frac{\pi_{\mathtt{f}}}{\pi_1}\frac{d^{y-1}}{e^{O(1)}}\right)^{\delta},$$

so that the maximizer must satisfy $d^{y-1}\pi_{\mathtt{f}} \leqslant d^{O(\varepsilon)}\pi_1$, implying the result. $\qquad\square$

Corollary 4.8. *Let $\varepsilon_0 > 0$ be a small constant uniform in d. For $d \geqslant d_0(\varepsilon_0)$, the following hold for all $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$:*

(a) Any global maximizer of $\boldsymbol{\Phi}$ on $\boldsymbol{\Delta}_\alpha$ lies in the interior $\boldsymbol{\Delta}_\alpha^\circ$.

(b) Any global maximizer of $_2\boldsymbol{\Phi}$ on $_2^\star\boldsymbol{\Delta}_\alpha$ must be an interior stationary point.

Proof. Propn. 4.5a and Lem. 4.7a combine to give (a), while (b) follows by combining Cor. 4.3, Propn. 4.5b, and Lem. 4.7b. $\qquad\square$

Cor. 4.8 was required in the proof of Thm. 3.4; it also implies (with Lem. 3.3) that any maximizer $\boldsymbol{h}$ on $_2\boldsymbol{\Phi}$ on $_2^\star\boldsymbol{\Delta}_\alpha$ corresponds to a solution h of the pair Bethe recursions for some $\lambda \equiv (\lambda_1,\lambda_2)$ ((38) and (34) with $\varphi_2^\lambda \equiv \varphi^{\lambda_1} \otimes \varphi^{\lambda_2}$ in place of φ^λ). It remains to identify this Bethe solution with the one corresponding to $_2^\star\boldsymbol{h}_\alpha = {}^\star\boldsymbol{h}_\alpha \otimes {}^\star\boldsymbol{h}_\alpha$.

We first show an analogous result in the simpler pair frozen model recursions, as follows. The unique fixed point $q = q^\lambda$ of (28) defines a Gibbs measure ν on frozen model configurations $\overline{\eta}$ on the infinite d-regular tree $T_d(t)$. We shall define now a Gibbs measure $_2\nu$ on *pair* frozen configurations $(\overline{\eta}^1, \overline{\eta}^2)$: the message-passing rule is given simply by applying $\dot{\mathsf{m}}_{d-1}$ separately in each coordinate. We weight according to λ_1 raised to the number of upward messages in $T_d(t-1)$ which are $\mathtt{1}$ in the first coordinate, times λ_2 raised to the number of upward messages which are $\mathtt{1}$ in the second coordinate. However, we will allow for more general boundary laws in which the incoming messages are not necessarily independent across the two copies: we will take the boundary messages $\tau_{v\to w} \equiv (\sigma_{v\to w}^1, \sigma_{v\to w}^2)$ to be distributed according to a measure $\boldsymbol{q}$ on $\{\mathtt{0},\mathtt{1},\mathtt{f}\}^2$ (still i.i.d. across the different boundary edges $v \to w$). Recall $\mathtt{Z} \equiv \{\mathtt{0},\mathtt{f}\}$, e.g. $\boldsymbol{q}_{\mathtt{ZZ}} \equiv \boldsymbol{q}(\{\mathtt{0},\mathtt{f}\}^2) = 1 - 2q_1 + \boldsymbol{q}_{11}$, and $\boldsymbol{q}_{\mathtt{1Z}} = q_1 - \boldsymbol{q}_{11} = \boldsymbol{q}_{\mathtt{Z1}}$. The reader can check that for q to define a consistent family $(_2\nu_t)_t$ of finite-dimensional distributions, its marginals $\boldsymbol{q}_{\eta\cdot} \equiv \sum_{\eta'}\boldsymbol{q}_{\eta\eta'}$ and $\boldsymbol{q}_{\cdot\eta} \equiv \sum_{\eta'}\boldsymbol{q}_{\eta'\eta}$ must satisfy the single-copy recursions (28) (with respect to λ_1 and λ_2 respectively), and further we must have

$$\begin{aligned} &z\boldsymbol{q}_{11} = \lambda_1\lambda_2(\boldsymbol{q}_{\mathtt{ZZ}})^{d-1}, \quad z\boldsymbol{q}_{\mathtt{1f}} = \lambda_1(d-1)\boldsymbol{q}_{\mathtt{Z1}}(\boldsymbol{q}_{\mathtt{ZZ}})^{d-2}, \quad z\boldsymbol{q}_{\mathtt{f1}} = \lambda_2(d-1)\boldsymbol{q}_{\mathtt{1Z}}(\boldsymbol{q}_{\mathtt{ZZ}})^{d-2} \\ &z\boldsymbol{q}_{\mathtt{ff}} = (d-1)\boldsymbol{q}_{11}(\boldsymbol{q}_{\mathtt{ZZ}})^{d-2} + (d-1)(d-2)\boldsymbol{q}_{\mathtt{1Z}}\boldsymbol{q}_{\mathtt{Z1}}(\boldsymbol{q}_{\mathtt{ZZ}})^{d-3}, \end{aligned} \tag{56}$$

with the rest being determined by margin constraints. Clearly $q^{\lambda_1} \otimes q^{\lambda_2}$ solves (56), and the following lemma identifies a regime in which it is the unique solution:

Lemma 4.9. *For $d^{1/10} \leqslant \{\lambda_1, \lambda_2\} \leqslant d^2$, the recursions* (56) *have the unique solution $q^{\lambda_1} \otimes q^{\lambda_2}$ in the regime $\boldsymbol{q}_{11} \leqslant {}^{9}\!/\!{}_{10} \cdot (\log d)/d$.*

Proof. Write $_iq$ for the solutions of (28) with respect to λ_i $(i = 1, 2)$. A measure $\boldsymbol{q}$ solving (56) corresponds to a root $\boldsymbol{q}_{11}$ of the function

$$f(x) \equiv x\left[1 + (\lambda_1 - 1)(1 - {}_1q_1)^{d-1} + (\lambda_2 - 1)(1 - {}_2q_1)^{d-1}\right] \\ -w(x)^{d-1}\left[\lambda_1\lambda_2 - x(\lambda_1 - 1)(\lambda_2 - 1)\right]$$

where $w(x) \equiv 1 - {}_1q_1 - {}_2q_1 + x$, so that $w(\boldsymbol{q}_{11}) = \boldsymbol{q}_{\mathtt{ZZ}}$. We calculate

$$f'(x) \equiv 1 + (\lambda_1 - 1)(1 - {}_1q_1)^{d-1} + (\lambda_2 - 1)(1 - {}_2q_1)^{d-1} \\ -w(x)^{d-2}\left[(d-1)(\lambda_1\lambda_2 - x(\lambda_1 - 1)(\lambda_2 - 1)) - w(x)(\lambda_1 - 1)(\lambda_2 - 1)\right].$$

Recall (29) that λ is monotone in $q_1 \equiv q_1^\lambda$, so in the regime $d^{1/10} \leqslant \lambda_i \leqslant d^2$ we must have $_iq_1 = y_i(\log d)/d$ with ${}^{21}\!/\!{}_{20} \leqslant y_i \leqslant 3$, and consequently

$$(\lambda_i - 1)(1 - {}_iq_1)^{d-1} = {}_iq_1/(1 - {}_iq_1) - (1 - {}_iq_1)^{d-2} = O((\log d)/d) \quad \text{for } i = 1, 2.$$

If $x = \rho(\log d)/d$ then $w(x)^{d-2} \asymp d^{\rho-y_1-y_2}$, so for $\rho \leqslant {}^{9}\!/\!{}_{10}$ we find $f'(x) = 1 - O(d^{-1/20})$, which shows clearly that f has the unique root $\boldsymbol{q}_{11} = ({}_1q_1)({}_2q_1)$ in the stated regime. $\quad\square$

Proof of Propn. 4.6. As noted above, Cor. 4.8b implies that any maximizer $\boldsymbol{h}$ of $_2\boldsymbol{\Phi}$ on $_2^{\star}\boldsymbol{\Delta}_\alpha$ corresponds to a solution h of the pair Bethe recursions (34) with respect to some $\varphi^{\lambda_1} \otimes \varphi^{\lambda_2}$. We now show that h must satisfy the Bethe symmetries $\hat{h}(\boldsymbol{io}) = \hat{h}(\boldsymbol{i'o})$, where $\boldsymbol{i}$ or $\boldsymbol{i'}$ now indicates the incoming *pair* of variable-to-clause messages, and $\boldsymbol{o}$ the outgoing pair of clause-to-variable messages.

We apply the same argument used to show the first-moment symmetries $\hat{h}_{\eta'\eta} = \hat{h}_{\eta''\eta}$ in the proof of Thm. 3.4: that is, we shall argue that the effect of changing the (pair) message $\boldsymbol{i}$ incoming to $\hat{T}_{d,2}$ does not percolate down the tree. Again, if $\boldsymbol{i}$ and $\boldsymbol{i'}$ differ only by changing $\mathtt{0}$ to $\mathtt{f}$ or vice versa in either coordinate, the effect does not propagate at all so we immediately find $\hat{h}_{\boldsymbol{io}} = \hat{h}_{\boldsymbol{i'o}}$. Further, whenever the outgoing message is $\mathtt{0}$ the value of $\hat{h}$ does not depend on the incoming message. Analogously to the equivalences (36), all remaining symmetries will follow by showing $\hat{h}(\begin{smallmatrix}\sigma\\\mathtt{11}\end{smallmatrix}) = \hat{h}(\begin{smallmatrix}\sigma\\\mathtt{f1}\end{smallmatrix})$ for any σ: as in the proof of Thm. 3.4, this follows by showing

$$d\,\dot{h}(\begin{smallmatrix}\sigma\\\mathtt{f1}\end{smallmatrix}) \ll \dot{h}(\begin{smallmatrix}\sigma\\\mathtt{01}\end{smallmatrix}) \quad \text{and} \quad d\,\dot{h}(\begin{smallmatrix}\sigma\\\mathtt{ff}\end{smallmatrix}) \ll \dot{h}(\begin{smallmatrix}\sigma\\\mathtt{0f}\end{smallmatrix}) \quad \text{for all } \sigma \in \mathscr{M}.$$

1. The estimates $d\,\dot{h}(\begin{smallmatrix}\sigma\\\mathtt{f1}\end{smallmatrix}) \ll \dot{h}(\begin{smallmatrix}\sigma\\\mathtt{01}\end{smallmatrix})$ are all proved in essentially the same manner which is similar to the proof of $\dot{h}_{\mathtt{f1}}/\dot{h}_{\mathtt{01}} \ll d^{-1}$ (see (41)), so we sketch the case $\sigma = \mathtt{f0}$ (which is slightly delicate) and leave the rest to the reader.

   By (38) and the $\mathtt{0}$-$\mathtt{f}$ symmetry noted above, $\dot{h}(\begin{smallmatrix}\sigma\\\mathtt{f1}\end{smallmatrix})/\dot{h}(\begin{smallmatrix}\sigma\\\mathtt{01}\end{smallmatrix}) = \hat{\boldsymbol{h}}(\begin{smallmatrix}\sigma & \mathtt{R}\sigma\\\mathtt{f1} & \mathtt{1f}\end{smallmatrix})/\hat{\boldsymbol{h}}(\begin{smallmatrix}\sigma & \mathtt{R}\sigma\\\mathtt{01} & \mathtt{10}\end{smallmatrix})$, the ratio between the number of $\mathtt{f0^s}$–$\mathtt{01}$ edges to the number of $\mathtt{f0^r}$–$\mathtt{01}$ edges. Let $\mu = n_{\mathtt{f1}}/n_{\mathtt{f0}}$ denote the proportion of $\mathtt{f0}$-variables which are matched to $\mathtt{f1}$-variables, and as before let F_2 denote the number of non-matching edges between $V_{\mathtt{1}\cdot}$ and $V_{\mathtt{f0}}$. Again using the $\mathtt{0}$-$\mathtt{f}$ symmetry we calculate

$$\begin{aligned} F_2/n_{\mathtt{f0}} &= \mu\,\mathbb{E}[Y\,|\,Y \geqslant 1] + (1 - \mu)\,\mathbb{E}[Y\,|\,Y \geqslant 2], \\ Y &\sim \mathrm{Bin}(d-1, p) \text{ where } p \equiv \dot{h}(\begin{smallmatrix}\mathtt{Zf}\\\mathtt{10}\end{smallmatrix})/\dot{h}(\begin{smallmatrix}\mathtt{Zf}\\\mathtt{\cdot0}\end{smallmatrix}). \end{aligned} \tag{57}$$

Recall (51) that $F_2/n_{\mathtt{f0}}$ must be close to $d\pi_{\mathtt{01}}$; a slight variation of the argument following (39) then gives $p \approx \pi_{\mathtt{01}}$. We therefore obtain the desired bound

$$\frac{\dot{h}\left(\begin{smallmatrix}\mathtt{f0}\\\mathtt{f1}\end{smallmatrix}\right)}{\dot{h}\left(\begin{smallmatrix}\mathtt{f0}\\\mathtt{01}\end{smallmatrix}\right)} = \frac{\mu\,\mathbb{E}[Y\mathbf{1}_{\{Y=1\}}\,|\,Y \geqslant 1] + (1-\mu)\,\mathbb{E}[Y\mathbf{1}_{\{Y=2\}}\,|\,Y \geqslant 2]}{\mu\,\mathbb{E}[Y\mathbf{1}_{\{Y>1\}}\,|\,Y \geqslant 1] + (1-\mu)\,\mathbb{E}[Y\mathbf{1}_{\{Y>2\}}\,|\,Y \geqslant 2]} \leqslant d^{-3/2}.$$

2. The estimates $d\,\dot{h}\left(\begin{smallmatrix}\sigma\\\mathtt{ff}\end{smallmatrix}\right) \ll \dot{h}\left(\begin{smallmatrix}\sigma\\\mathtt{0f}\end{smallmatrix}\right)$ can be proved by estimating the ratios of edge densities $\hat{h}\left(\begin{smallmatrix}\sigma & \mathtt{R}\sigma\\\mathtt{ff} & \mathtt{ff}\end{smallmatrix}\right)/\hat{h}\left(\begin{smallmatrix}\sigma & \mathtt{R}\sigma\\\mathtt{0f} & \mathtt{f0}\end{smallmatrix}\right)$. These are relatively straightforward from the estimates of Lem. 4.7 and are left to the reader.

This proves the symmetries $\hat{h}(\boldsymbol{io}) = \hat{h}(\boldsymbol{i'o})$, so we conclude that h must in fact correspond to a solution $\boldsymbol{q}$ of the pair frozen model recursions (56) via $9\,\hat{h}(\boldsymbol{io}) = \boldsymbol{q_o}$ (cf. (37)). It remains to verify that $\boldsymbol{q}$ falls within the regime of Lem. 4.9. As before, let K denote the number of $\mathtt{01}$–$\mathtt{10}$ edges: from (38) and the h-to-$\boldsymbol{q}$ correspondence,

$$K/n_{\mathtt{10}} = \frac{\sum_{j\geqslant 2} j\binom{d}{j}(\boldsymbol{q}_{\mathtt{Z1}}/\boldsymbol{q}_{\mathtt{ZZ}})^j}{\sum_{j\geqslant 2}\binom{d}{j}(\boldsymbol{q}_{\mathtt{Z1}}/\boldsymbol{q}_{\mathtt{ZZ}})^j}, \quad \text{so (51) implies } \boldsymbol{q}_{\mathtt{Z1}} \approx \pi_{\mathtt{01}}\boldsymbol{q}_{\mathtt{ZZ}} \text{ and } \boldsymbol{q}_{\mathtt{1Z}} \approx \pi_{\mathtt{10}}\boldsymbol{q}_{\mathtt{ZZ}}. \qquad (58)$$

The parameter p in (57) can be expressed as $\boldsymbol{q}_{\mathtt{Z1}}/\boldsymbol{q}_{\mathtt{Z}\cdot}$, so (58) implies $p \approx \pi_{\mathtt{10}}$, therefore $\mathbb{P}(Y \geqslant 2) \asymp 1$ for $Y \sim \mathrm{Bin}(d-1, p)$. Another application of (38) gives

$$\frac{n_{\mathtt{f1}}}{n_{\mathtt{f0}}^{\bullet}} = \frac{\mu}{1-\mu} = \frac{\boldsymbol{q}_{\mathtt{11}}\,\mathbb{P}(Y \geqslant 1)}{\boldsymbol{q}_{\mathtt{1Z}}\,\mathbb{P}(Y \geqslant 2)} \asymp \frac{\boldsymbol{q}_{\mathtt{11}}}{\boldsymbol{q}_{\mathtt{1Z}}},$$

so Lem. 4.7b implies (crudely) that $\boldsymbol{q}_{\mathtt{11}}/\boldsymbol{q}_{\mathtt{1Z}} \leqslant d^{-1/2}$. From these estimates we may apply Lem. 4.9 to conclude that $\boldsymbol{q} = q^{\lambda_1} \otimes q^{\lambda_2}$. By the one-to-one correspondence (Propn. 3.5) between α and q in this regime we conclude $\boldsymbol{h} = {}^{\star}\boldsymbol{h}_\alpha$ as claimed. $\qquad\square$

4.4. **A priori rigidity estimate.** In this section we analyze near-identical frozen model configurations to prove

Proposition 4.10. *The contribution to* $\mathbb{E}[\boldsymbol{Z}_{n\alpha}^2]$ *from configurations* $\boldsymbol{h} \in \ddot{}_2\boldsymbol{\Delta}_\alpha$ *is* $\asymp_d \mathbb{E}\boldsymbol{Z}_{n\alpha}$.

The proof of Propn. 4.10 is based on an *a priori* estimate showing that frozen model configurations are sufficiently rigid that one typically does not find a large cluster of configurations near a given one. For application in our proof of the tightness of $\boldsymbol{A}_n$ we shall prove this estimate for graphs drawn from the following slight generalization of the configuration model which allows for some unmatched edges (Fig. 5). Let V be a set of n vertices, each incident to d half-edges. Let $\mathcal{Y}^\star$ be a disjoint set of $\boldsymbol{y} \lesssim_d \log n$ vertices, each incident to a single half-edge. Finally let F be a set of $(nd + \boldsymbol{y})/2$ clauses, each incident to 2 half-edges. Let G^∂ be the graph formed by taking a random matching between the half-edges incident to $V^\partial \equiv V \cup \mathcal{Y}^\star$ with the half-edges incident to F. We shall write $F^\partial \equiv \partial\mathcal{Y}^\star \subseteq F$, $\mathcal{Y} \equiv V \cap \partial F^\partial$, and $\mathcal{R} \equiv V\backslash\mathcal{Y}$.

Definition 4.11. An auxiliary model configuration on G^∂ is a message configuration $\underline{\sigma}^\partial$ on the edges of G^∂ (including variable-to-clause messages from $\mathcal{Y}^\star$ to F^∂) such that the configuration of messages incident to any variable or clause within G^∂ is valid. A frozen model configuration $\overline{\boldsymbol{\eta}}$ on G^∂ is a spin configuration $\eta \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}^{V^\partial}$ together with a subset $\underline{m} \subseteq F$ such that every $v \in V$ satisfies properties (i)-(iii) of Defn. 2.4[12], and further the density of $\mathtt{f}$-variables in each coordinate is $\leqslant \beta_{\max}$.

---

[12]The properties need not be satisfied by vertices on $\mathcal{Y}^\star$.

Arguing as in Rmk. 3.1, there is a one-to-one correspondence between the frozen and auxiliary models on G^∂. As in (46), we associate to a pair frozen configuration $\overline{\boldsymbol{\eta}}^i \equiv (\eta^i, \underline{m}^i)$ $(i = 1, 2)$ on G^∂ a spin configuration $\underline{\omega} \equiv \omega(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2) \in \mathscr{P}^{V^\partial}$. As in (47), for $\omega \in \mathscr{P}$ and $S \subseteq V^\partial$ we write S_ω for the subset of vertices in S with spin ω; in particular, write $S_{\neq}$ for the subset of vertices in S with spin in $\mathscr{P}_{\neq} \equiv \mathscr{P} \backslash \{\mathtt{00}, \mathtt{11}, \mathtt{ff}\}$, and $S_= \equiv S \backslash S_{\neq}$. Let

$$\boldsymbol{r}_\omega \equiv |\mathcal{R}_\omega|, \qquad \boldsymbol{y}_\omega \equiv |\mathcal{Y}_\omega|, \qquad n_\omega \equiv |V_\omega| = \boldsymbol{r}_\omega + \boldsymbol{y}_\omega,$$
$$\boldsymbol{\rho} \equiv \boldsymbol{r}_{\neq}, \qquad \boldsymbol{\varepsilon} \equiv \boldsymbol{y}_{\neq}, \qquad \Delta \equiv n_{\neq} = \boldsymbol{\rho} + \boldsymbol{\varepsilon}.$$

For $S, S' \subseteq V^\partial$ write $F(S, S')$ for the set of clauses $a \in F$ joining a vertex in S to a vertex in S'; write $F(S) \equiv F(S, S)$ for the clauses internal to S.
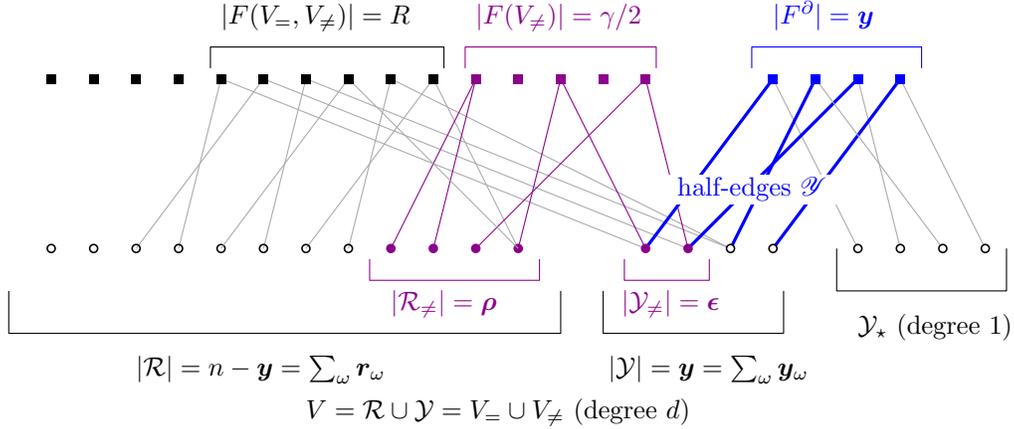
FIGURE 5. Graph G^∂ with dangling edges

Write $\mathscr{Y}$ for the half-edges joining $\mathcal{Y}$ to F^∂. In the following, we fix boundary conditions given by a auxiliary pair configuration $\underline{\tau}_{\mathscr{Y}} \equiv (\underline{\sigma}^1_{\mathscr{Y}}, \underline{\sigma}^2_{\mathscr{Y}})$ on $\mathscr{Y}$. Let $_2\boldsymbol{Z}^\partial[\pi|\underline{\tau}_{\mathscr{Y}}]$ denote the cardinality of the set $_2\boldsymbol{O}^\partial[\pi|\underline{\tau}_{\mathscr{Y}}]$ of pair frozen model configurations on G^∂ which are consistent with $\underline{\tau}_{\mathscr{Y}}$ and have empirical measure $\pi_\omega = \boldsymbol{r}_\omega/\boldsymbol{r}$ when restricted to $\mathcal{R}$. We decompose

$$_2\boldsymbol{Z}^\partial[\pi|\underline{\tau}_{\mathscr{Y}}] = \sum_{A \geqslant 0} {}_2\boldsymbol{Z}^\partial_A[\pi|\underline{\tau}_{\mathscr{Y}}] = \sum_{A \geqslant 0} |_2\boldsymbol{O}^\partial_A[\pi|\underline{\tau}_{\mathscr{Y}}]| \tag{59}$$

where $_2\boldsymbol{O}^\partial_A[\pi|\underline{\tau}_{\mathscr{Y}}] \subseteq {}_2\boldsymbol{O}^\partial[\pi|\underline{\tau}_{\mathscr{Y}}]$ denotes the subset of configurations which have $\gamma/2 \equiv |F(V_{\neq})| = \boldsymbol{\rho} + A$ internal edges among the unequal spins.[13] We shall compare the expectation of $_2\boldsymbol{Z}^\partial_A[\pi|\underline{\tau}_{\mathscr{Y}}]$ with that of $\boldsymbol{Z}^\partial[\pi^1|\underline{\sigma}^1_{\mathscr{Y}}]$ — the number of frozen model configurations on G^∂ which are consistent with $\underline{\sigma}^1_{\mathscr{Y}}$ and have empirical measure given by the projection π^1 of π onto the first coordinate:

$$|\mathcal{R}|\pi^1_\eta \equiv |\mathcal{R}|\pi_{\eta \cdot} = \boldsymbol{r}_{\eta \cdot} \equiv \boldsymbol{r}_{\eta\eta} + \boldsymbol{\rho}_\eta \quad (\eta \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}). \tag{60}$$

Note $\boldsymbol{Z}^\partial[\pi^1|\underline{\sigma}^1_{\mathscr{Y}}] = \mathsf{proj}^1(_2\boldsymbol{O}^\partial[\pi|\underline{\tau}_{\mathscr{Y}}])$ where proj^1 is the projection mapping $(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2) \mapsto \overline{\boldsymbol{\eta}}^1$.

Proposition 4.12. *There exists a small constant $\varepsilon_0 > 0$ (uniform in $d \geqslant d_0$) such that for any empirical measure π on $\mathscr{P}$ whose projections π^1, π^2 have intensities $\alpha_{\mathrm{lbd}} \leqslant \alpha^1, \alpha^2 \leqslant \alpha_{\mathrm{ubd}}$,*

$$\mathbb{E}[_2\boldsymbol{Z}^\partial[\pi|\underline{\tau}_{\mathscr{Y}}]] \leqslant \frac{\mathbb{E}[\boldsymbol{Z}^\partial[\pi^1|\underline{\sigma}^1_{\mathscr{Y}}]] + \mathbb{E}[\boldsymbol{Z}^\partial[\pi^2|\underline{\sigma}^2_{\mathscr{Y}}]]}{d^{(\boldsymbol{\rho}-\boldsymbol{\varepsilon})/10}}$$

provided $\Delta \equiv \boldsymbol{\rho} + \boldsymbol{\varepsilon} \leqslant n\varepsilon_0(\log d)/d$.

---

[13]It is straightforward to see that $A \geqslant 0$; we show a stronger inequality in Lem. 6.5.

Proof. Fix $\underline{\omega} \in \mathscr{P}^{V^{\partial}}$ which is consistent with $\underline{\tau}_{\mathscr{Y}}$ and whose restriction to $\mathcal{R}$ has empirical measure π, and write $\underline{\eta}^1 \equiv \eta^1(\underline{\omega}) \in \{0, 1, \mathbf{f}\}^{V^{\partial}}$ for its projection onto the first coordinate. Write $\alpha_1 \equiv \pi_{1\cdot}$, $\beta \equiv \pi_{\mathbf{f}\cdot}$. Write $\boldsymbol{E}_{\omega}$ for the number of half-edges incident to V_{ω} and not matched to a half-edge from $\mathcal{Y}^{\star}$ (for example $\boldsymbol{E}_{\neq} = \boldsymbol{\rho}d + \boldsymbol{\varepsilon}(d-1)$); the total number of such half-edges is $\boldsymbol{E} \equiv \sum_{\omega} \boldsymbol{E}_{\omega} = nd - \boldsymbol{y}$. Let $R \equiv |F(V_{=}, V_{\neq})| = \boldsymbol{E}_{\neq} - \gamma$.

Assignment of half-edges from $V_{\neq}$.
We first estimate the probability of the event $F(V_{\neq}, V_{11}) = \varnothing$: to this end, define

$$\mathscr{D}_{A,B}(\underline{\omega}) \equiv \{|F(V_{\neq})| = \gamma/2 \equiv \boldsymbol{\rho} + A, |F(V_{\neq}, V_{\mathbf{ff}})| = B, F(V_{\neq}, V_{11}) = \varnothing\}.$$

Then, writing $\text{bin}_{d,p}(j) \equiv \binom{d}{j}p^j(1-p)^{d-j}$ and recalling the notation (5), we have

$$\mathbb{P}[\mathscr{D}_{A,B}(\underline{\omega})] \leqslant \underbrace{\frac{(\boldsymbol{E}_{\neq})_{\gamma}\,(\gamma-1)!!}{\gamma!\,[\boldsymbol{E}]_{\gamma/2}}}_{\leqslant e^{O(\gamma)}(\Delta^2 d/(n\gamma))^{\gamma/2}} \times \underbrace{\frac{(\boldsymbol{E}_{\mathbf{ff}} + \boldsymbol{E}_{00})_R}{[\boldsymbol{E} - \gamma]_R}}_{\leqslant d^{O(\Delta\varepsilon_0)}e^{-\boldsymbol{E}_{11}R/\boldsymbol{E}}} \times \underbrace{\frac{\binom{R}{B}(\boldsymbol{E}_{\mathbf{ff}})_B(\boldsymbol{E}_{00})_{R-B}}{(\boldsymbol{E}_{\mathbf{ff}} + \boldsymbol{E}_{00})_R}}_{\leqslant \text{bin}_{R,2\beta}(B)},$$

so altogether $\mathbb{P}[\mathscr{D}_{A,B}(\underline{\omega})] \lesssim e^{O(A)}d^{O(\Delta\varepsilon_0)}(\Delta^2 d/(n\gamma))^{\gamma/2}e^{-\boldsymbol{E}_{11}R/\boldsymbol{E}}\text{bin}_{R,2\beta}(B)$.

Assignment of remaining half-edges.
Let $R' \equiv \boldsymbol{E}_{\neq} + R = 2\boldsymbol{E}_{\neq} - \gamma \leqslant 2\Delta d$ denote the number of half-edges which are either incident to $V_{\neq}$ or matched with a half-edge from $V_{\neq}$. Let $G_{=}$ denote the induced subgraph on $V_{=}$, and let $\mathcal{J}_{=}(\underline{\omega})$ be the number of frozen model configurations on $G_{=}$ which are consistent with $\underline{\omega}$. Meanwhile let $\mathcal{J}(\underline{\eta}^1)$ be the number of frozen model configurations on G^{∂} which are consistent with $\underline{\eta}^1$. We compare $\mathbb{E}[\mathcal{J}_{=}(\underline{\omega}) \,|\, \mathscr{D}_{A,B}(\underline{\omega})]$ to $\mathbb{E}[\mathcal{J}(\underline{\eta}^1)]$ by modifying the proof of Propn. 4.2: the boundary $\mathcal{Y}^{\star}$ causes some minor complication, but otherwise the calculation is much simplified by the restriction to vertices of equal spin.

Let $\widetilde{V}_{\mathbf{ff}}$ denote the subset of vertices in $V_{\mathbf{ff}}$ not matched to $\mathcal{Y}^{\star}$ under $\underline{\tau}_{\mathscr{Y}}$. Since there are by definition no edges between $\mathcal{R}$ and $\mathcal{Y}^{\star}$, clearly $\widetilde{V}_{\mathbf{ff}} \cap \mathcal{R} = \mathcal{R}_{\mathbf{ff}}$, so we write

$$\widetilde{\mathcal{Y}}_{\mathbf{ff}} \equiv \widetilde{V}_{\mathbf{ff}} \cap \mathcal{Y}, \quad \widetilde{\boldsymbol{y}}_{\mathbf{ff}} \equiv |\widetilde{\mathcal{Y}}_{\mathbf{ff}}|, \quad \widetilde{n}_{\mathbf{ff}} \equiv |\widetilde{V}_{\mathbf{ff}}| \equiv r_{\mathbf{ff}} + \widetilde{\boldsymbol{y}}_{\mathbf{ff}}.$$

Similarly, let $\widetilde{V}_{\mathbf{f}\cdot}$ denote the subset of vertices in $V_{\mathbf{f}\cdot}$ not matched to $\mathcal{Y}^{\star}$ under $\underline{\sigma}^1_{\mathscr{Y}}$: then $\widetilde{V}_{\mathbf{f}\cdot} \cap \mathcal{R} = \mathcal{R}_{\mathbf{f}\cdot}$, $\widetilde{V}_{\mathbf{f}\cdot} \cap \mathcal{R}_{=} = \mathcal{R}_{\mathbf{ff}}$, and $\widetilde{V}_{\mathbf{f}\cdot} \cap \mathcal{Y}_{=} = \widetilde{\mathcal{Y}}_{\mathbf{ff}}$, so we write

$$\widetilde{\mathcal{Y}}_{\mathbf{f}\cdot} \equiv \widetilde{V}_{\mathbf{f}\cdot} \cap \mathcal{Y}, \quad \widetilde{\boldsymbol{y}}_{\mathbf{f}\cdot} \equiv |\widetilde{\mathcal{Y}}_{\mathbf{f}\cdot}|, \quad \widetilde{n}_{\mathbf{f}\cdot} \equiv |\widetilde{V}_{\mathbf{f}\cdot}| \equiv r_{\mathbf{f}\cdot} + \widetilde{\boldsymbol{y}}_{\mathbf{f}\cdot},$$
$$\widetilde{\varepsilon}_{\mathbf{f}\cdot} \equiv |\widetilde{V}_{\mathbf{f}\cdot} \cap \mathcal{Y}_{\neq}| = \widetilde{\boldsymbol{y}}_{\mathbf{f}\cdot} - \widetilde{\boldsymbol{y}}_{\mathbf{ff}}, \quad \widetilde{\Delta}_{\mathbf{f}\cdot} \equiv |\widetilde{V}_{\mathbf{f}\cdot} \cap V_{\neq}| = \rho_{\mathbf{f}} + \widetilde{\varepsilon}_{\mathbf{f}\cdot} = \widetilde{n}_{\mathbf{f}\cdot} - \widetilde{n}_{\mathbf{ff}}.$$

Consider the expected number of matchings on $\widetilde{V}_{\mathbf{ff}}$ conditioned on $\mathscr{D}_{A,B}(\underline{\omega})$, divided by the expected number of matchings on $\widetilde{V}_{\mathbf{f}\cdot}$: this ratio is given by

$$\frac{\left(\prod_{v \in \widetilde{V}_{\mathbf{ff}}} |\partial v \backslash (\mathcal{Y}^{\star} \cup V_{\neq})|\right)(\widetilde{n}_{\mathbf{ff}} - 1)!!\,/\,[\boldsymbol{E} - R']_{\widetilde{n}_{\mathbf{ff}}/2}}{\left(\prod_{v \in \widetilde{V}_{\mathbf{f}\cdot}} |\partial v \backslash \mathcal{Y}^{\star}|\right)(\widetilde{n}_{\mathbf{f}\cdot} - 1)!!\,/\,[\boldsymbol{E}]_{\widetilde{n}_{\mathbf{f}\cdot}/2}} \leqslant \mathscr{J}_{\mathbf{f}} \equiv \frac{\mathbf{g}_{\mathbf{f}}(\widetilde{n}_{\mathbf{ff}}, \boldsymbol{E} - R')}{\mathbf{g}_{\mathbf{f}}(\widetilde{n}_{\mathbf{f}\cdot}, \boldsymbol{E})}\,[d/(d-1)]^{\widetilde{\varepsilon}_{\mathbf{f}\cdot}},$$

where $\mathbf{g}_{\mathbf{f}}(n_{\mathbf{f}}, E) \equiv d^{n_{\mathbf{f}}}(n_{\mathbf{f}} - 1)!!/[E]_{n_{\mathbf{f}}/2}$ (cf. (49)), and the factor $[d/(d-1)]^{\widetilde{\varepsilon}_{\mathbf{f}\cdot}}$ accounts for the vertices in $v \in \widetilde{\mathcal{Y}}_{\mathbf{f}\cdot}\backslash\widetilde{V}_{\mathbf{ff}}$ which contribute factors $d-1$ in the denominator on the left-hand side. We estimate $\mathscr{J}_{\mathbf{f}} \leqslant e^{O(\Delta)}(d\beta)^{-\widetilde{\Delta}_{\mathbf{f}\cdot}/2}$.

Because of the restriction to $V_=$, in the current setting the method of Propn. 4.2 reduces to a first-moment calculation, yielding

$$\frac{\mathbb{E}[\mathcal{J}_=(\underline{\omega}) \,|\, \mathscr{D}_{A,B}(\underline{\omega})]}{\mathbb{E}[\mathcal{J}(\underline{\eta}^1)]} \leqslant \mathscr{J}_\mathtt{f} \times \underbrace{\frac{\mathbf{g}_1(\boldsymbol{E} - R' - \widetilde{n}_\mathtt{ff}, \boldsymbol{E}_{11}, \boldsymbol{E}_\mathtt{ff} - B - \widetilde{n}_\mathtt{ff})}{\mathbf{g}_1(\boldsymbol{E} - \widetilde{n}_{\mathtt{f}\cdot}, \boldsymbol{E}_{1\cdot}, \boldsymbol{E}_{\mathtt{f}\cdot} - \widetilde{n}_{\mathtt{f}\cdot})}}_{\mathscr{J}_1(B)} \times \mathscr{J}_0$$

where $\mathscr{J}_0$ is the (conditional) probability of all vertices in V_{00} being forced, divided by the (conditional) probability of all vertices in $V_{0\cdot}$ being forced: explicitly, for $v \in V$ set $\ell^v = 1$ if v has a neighbor $w \in \mathcal{Y}^\star$ with $\sigma_w^1 = \mathtt{1}$, otherwise set $\ell^v = 0$. Fix $0 < \vartheta < 1$ and define independent random variables $X^v \sim \mathrm{Bin}(|\partial v \backslash \mathcal{Y}^\star|, \vartheta)$, with joint law $\mathbb{P}_\vartheta$: then

$$\mathscr{J}_0 \leqslant \frac{\mathbb{P}_\vartheta(X^i + \ell^i \geqslant 2\,\forall i \in V_{00}\backslash\partial V_{\neq} \,|\, \sum_{i\in V_{00}\backslash\partial V_{\neq}} X^i = \boldsymbol{E}_{11})}{\mathbb{P}_\vartheta(X^i + \ell^i \geqslant 2\,\forall i \in V_{0\cdot} \,|\, \sum_{i\in V_{0\cdot}} X^i = \boldsymbol{E}_{1\cdot})},$$

where we have ignored the forcing constraint on the vertices in V_{00} neighboring $V_{\neq}$. A trivial variation on the proof of Propn. 2.9 gives $\mathscr{J}_0 = e^{O(\Delta)}$.

Recall $R' \equiv \boldsymbol{E}_{\neq} + R$; there remain $\boldsymbol{E}_{00} - R + B$ available half-edges from V_{00}. Therefore

$$\mathscr{J}_1(B) = \frac{(\boldsymbol{E}_{00} - R + B)_{\boldsymbol{E}_{11}} \times [\boldsymbol{E} - \widetilde{n}_{\mathtt{f}\cdot}]_{\boldsymbol{E}_{1\cdot}}}{[\boldsymbol{E}_= - R - \widetilde{n}_\mathtt{ff}]_{\boldsymbol{E}_{11}} \times (\boldsymbol{E}_{0\cdot})_{\boldsymbol{E}_{1\cdot}}} = d^{O(\Delta\varepsilon_0)} \frac{\exp\{\boldsymbol{E}_{11}(B + \boldsymbol{E}_{\neq})/\boldsymbol{E}\}}{\exp\{\boldsymbol{E}_{11}(\boldsymbol{E}_{0\cdot} - \boldsymbol{E}_{00})/\boldsymbol{E}\}}.$$

Writing $\Delta_0 \equiv n_0 - n_{00}$ and recalling $\boldsymbol{E}_{\neq} - R = \gamma = 2(\boldsymbol{\rho} + A)$ then gives

$$\sum_B \frac{\mathrm{bin}_{R,2\beta}(B)\,\mathscr{J}_1(B)}{\exp\{\boldsymbol{E}_{11}R/\boldsymbol{E}\}} = d^{O(\Delta\varepsilon_0)} e^{O(A)} \sum_B \frac{\mathrm{bin}_{R,2\beta}(B)\exp\{\boldsymbol{E}_{11}B/\boldsymbol{E}\}}{\exp\{\boldsymbol{E}_{11}(\boldsymbol{E}_{0\cdot} - \boldsymbol{E}_{00})/\boldsymbol{E}\}} = \frac{d^{O(\Delta\varepsilon_0)} e^{O(A)}}{\exp\{\alpha_1 d\Delta_0\}}.$$

Combinatorial factors.

The preceding estimates were for a fixed configuration $\underline{\omega} \in \mathscr{P}^{V^\partial}$ consistent with $\underline{\tau}_\mathscr{Y}, \pi$. Accounting for the permutations of $\underline{\omega}$ and $\underline{\eta}^1$ gives (recalling $\boldsymbol{r} = |\mathcal{R}|$)

$$R_A^1[\pi|\underline{\tau}_\mathscr{Y}] \equiv \frac{\mathbb{E}[_2\boldsymbol{Z}_A^\partial[\pi|\underline{\tau}_\mathscr{Y}]]}{\mathbb{E}[\boldsymbol{Z}^\partial[\pi^1|\underline{\sigma}_\mathscr{Y}^1]]} \leqslant e^{O(\Delta+A)} \frac{\binom{\boldsymbol{r}}{\boldsymbol{\rho},\boldsymbol{r}_{00},\boldsymbol{r}_{11},\boldsymbol{r}_\mathtt{ff}}}{\binom{\boldsymbol{r}}{\boldsymbol{r}_{0\cdot},\boldsymbol{r}_{1\cdot},\boldsymbol{r}_{\mathtt{f}\cdot}}} \frac{\sum_B \mathbb{E}[\mathcal{J}_=(\underline{\omega}); \mathscr{D}_{A,B}(\underline{\omega})]}{\mathbb{E}[\mathcal{J}(\underline{\eta}^1)]}.$$

The ratio of multinomial coefficients is $\lesssim e^{O(\boldsymbol{r})}(n/\boldsymbol{\rho})^{\boldsymbol{\rho}}(\alpha_1)^{\boldsymbol{\rho}_1}\beta^{\boldsymbol{\rho}_\mathtt{f}}$. Combining with the preceding estimates gives

$$R_A^1[\pi|\underline{\tau}_\mathscr{Y}] \leqslant d^{O(\Delta\varepsilon_0)} e^{O(A)} (\boldsymbol{\rho}/\gamma)^{\boldsymbol{\rho}} \Big(\frac{\Delta^2 d}{n\gamma}\Big)^A \frac{d^{\boldsymbol{\rho}}(\alpha_1)^{\boldsymbol{\rho}_1}\beta^{\boldsymbol{\rho}_\mathtt{f}}}{\exp\{\alpha_1 d\Delta_0\}(d\beta)^{\widetilde{\Delta}_{\mathtt{f}\cdot}/2}}, \tag{61}$$

where $\gamma/2 = \boldsymbol{\rho} + A$ and $\Delta \leqslant \boldsymbol{\rho} e^{\varepsilon/\boldsymbol{\rho}}$ were used to bound $(\Delta^2/(\boldsymbol{\rho}\gamma))^{\boldsymbol{\rho}} \leqslant e^{O(\boldsymbol{\rho})}(\boldsymbol{\rho}/\gamma)^{\boldsymbol{\rho}}$. Using $\alpha_1 \geqslant (3/2)(\log d)/d$ and $\beta \leqslant \beta_{\max}$ then gives

$$R_A^1[\pi|\underline{\tau}_\mathscr{Y}] \leqslant \frac{d^{O(\Delta\varepsilon_0)}(\log d)^{O(\Delta)}(\boldsymbol{\rho}/\gamma)^{\boldsymbol{\rho}}(\frac{\Delta^2 d}{n\gamma})^A}{\exp\{[\frac{3}{2}\Delta_0 - \boldsymbol{\rho}_0 + \frac{1}{4}(\boldsymbol{\rho}_\mathtt{f} - \widetilde{\varepsilon}_{\mathtt{f}\cdot})]\log d\}}.$$

Recall $\boldsymbol{r}_{1\mathtt{f}} \leqslant n_{0\mathtt{f}}$ and $\boldsymbol{r}_{\mathtt{f}1} \leqslant n_{\mathtt{f}0}$; by symmetry we shall suppose $n_{01} + n_{0\mathtt{f}} \geqslant n_{10} + n_{\mathtt{f}0}$. Then

$$\tfrac{3}{2}\Delta_0 - \boldsymbol{\rho}_0 \geqslant \tfrac{1}{2}\Delta_0 \geqslant \tfrac{1}{4}(n_{01} + n_{0\mathtt{f}} + n_{10} + n_{\mathtt{f}0}) \geqslant \tfrac{1}{8}(\boldsymbol{\rho} - \boldsymbol{r}_{\mathtt{f}!\mathtt{f}}),$$

so combining with $\boldsymbol{\rho}_\mathtt{f} - \widetilde{\varepsilon}_{\mathtt{f}\cdot} \geqslant \boldsymbol{\rho}_\mathtt{f} - \varepsilon$ gives, for some constant C uniform in d,

$$\sum_{A\geqslant 0} R_A^1[\pi|\underline{\tau}_\mathscr{Y}] \leqslant \frac{d^{O(\Delta\varepsilon_0)}(\log d)^{O(\Delta)}}{d^{(\boldsymbol{\rho}-\varepsilon)/8}} \sum_{A\geqslant 0} (\boldsymbol{\rho}/\gamma)^{\boldsymbol{\rho}} \Big(\frac{C\Delta^2 d}{n\gamma}\Big)^A \leqslant \frac{d^{O(\Delta\varepsilon_0)}(\log d)^{O(\Delta)}}{d^{(\boldsymbol{\rho}-\varepsilon)/8}}.$$

To see the last inequality, note the sum over $A \geqslant 0$ is clearly $\lesssim 1$ if $C\Delta^2 d \leqslant 2n\gamma$. If $C\Delta^2 d > 2n\gamma$, then recalling $\boldsymbol{\rho} \leqslant \gamma/2$ and optimizing over γ gives $(\boldsymbol{\rho}/\gamma)^{\boldsymbol{\rho}}(C\Delta^2 d/(n\gamma))^A \leqslant d^{O(\Delta\varepsilon_0)}$. It follows that there exists a small constant ε_0 (uniform in d) such that

$$\sum_{A \geqslant 0} R_A^1[\pi|_{\mathcal{T}_{\mathscr{Y}}}] \leqslant d^{-(\rho-\varepsilon)/10} \quad \text{for } \Delta = \boldsymbol{\rho} + \boldsymbol{\varepsilon} \leqslant n\varepsilon_0(\log d)/d,$$

implying the result. $\qquad\square$

Proof of Propn. 4.10. Follows from Propn. 4.12 applied to our original random graph $\mathcal{G}_{n,d}$ with no unmatched edges. $\qquad\square$

Proof of Propn. 4.1. Follows by combining Propns. 4.6 and 4.10. $\qquad\square$

5. Negative-definiteness of free energy Hessians

In this section we prove Thm. 2.5 as well as

Theorem 5.1. *It holds uniformly over $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$ that $\mathbb{E}[\boldsymbol{Z}_{n\alpha}^2] \lesssim_d (\mathbb{E}\boldsymbol{Z}_{n\alpha})^2 + \mathbb{E}\boldsymbol{Z}_{n\alpha}$.*

Proposition 5.2. *For $d \geqslant d_0$, the Hessians $H\boldsymbol{\Phi}({}^{\star}\boldsymbol{h}_\alpha)$ and $H({}_2\boldsymbol{\Phi})({}_2^{\star}\boldsymbol{h}_\alpha)$ as functions on $\boldsymbol{\Delta}_\alpha$ and ${}_2\boldsymbol{\Delta}_\alpha$ respectively are negative-definite.*

The calculation of this section is similar to that of [15, §7]. Let $\boldsymbol{h} \in \boldsymbol{\Delta}_\alpha^\circ$ with $\dot{\boldsymbol{h}}$ and $\hat{\boldsymbol{h}}$ both symmetric, and let $\boldsymbol{\delta}$ be any signed measure on $\operatorname{supp}\varphi$ (not necessarily symmetric) with $\boldsymbol{h} + s\boldsymbol{\delta} \in \boldsymbol{\Delta}_\alpha^\circ$ for sufficiently small $|s|$. Then

$$\partial_s^2 \boldsymbol{\Phi}(\boldsymbol{h} + s\boldsymbol{\delta})|_{s=0} = -\langle(\dot{\boldsymbol{\delta}}/\dot{\boldsymbol{h}})^2\rangle_{\dot{\boldsymbol{h}}} - (d/2)\langle(\hat{\boldsymbol{\delta}}/\hat{\boldsymbol{h}})^2\rangle_{\hat{\boldsymbol{h}}} + d\langle(\bar{\delta}/\bar{h})^2\rangle_{\bar{h}}, \tag{62}$$

where a/b denotes the vector given by coordinate-wise division of a by b, and $\langle\cdot\rangle_h$ denotes integration with respect to measure h, e.g. $\langle(\bar{\delta}/\bar{h})^2\rangle_{\bar{h}} = \sum_\sigma \bar{\delta}(\sigma)^2/\bar{h}(\sigma)$. Consider maximizing (62) over $\boldsymbol{\delta}$ subject to fixed marginals $\bar{\delta}$: we find that the optimal $\hat{\boldsymbol{\delta}}$ will be symmetric, with $\hat{\boldsymbol{\delta}}(\sigma, \mathsf{R}\sigma) = \bar{\delta}(\sigma) = \bar{\delta}(\mathsf{R}\sigma)$. The optimal $\dot{\boldsymbol{\delta}}$ will be of form $d\,\dot{\boldsymbol{\delta}}(\underline{\dot{\sigma}}) = \dot{\boldsymbol{h}}(\underline{\dot{\sigma}})\sum_{i=1}^d \dot{\chi}_{\sigma_i}$ with $\dot{\chi}$ chosen to satisfy the margin constraint — which, after a little algebra, becomes the system of equations

$$\bar{H}^{-1}\bar{\delta} = d^{-1}\left[I + (d-1)\dot{M}\right]\dot{\chi} \tag{63}$$

where $\bar{H} \equiv \operatorname{diag}(\bar{h})$ and $\dot{M}$ denotes the stochastic matrix with entries

$$\dot{M}_{\sigma,\sigma'} \equiv \frac{1}{\bar{h}(\sigma)}\sum_{\underline{\dot{\sigma}}} \dot{\boldsymbol{h}}(\underline{\dot{\sigma}})\mathbf{1}\{(\sigma_1, \sigma_2) = (\sigma, \sigma')\}. \tag{64}$$

If such $\dot{\chi}$ exists, then the minimal value of $\langle(\dot{\boldsymbol{\delta}}/\dot{\boldsymbol{h}})^2\rangle_{\dot{\boldsymbol{h}}}$ subject to marginals $\bar{\delta}$ is $\langle\bar{\delta}, \dot{\chi}\rangle$ (which clearly remains invariant under translations of $\dot{\chi}$ by vectors in the kernel of $I + (d-1)\dot{M}$).

Throughout the following we take $\boldsymbol{h}$ to be ${}^{\star}\boldsymbol{h}_\alpha$ (first moment) or ${}_2^{\star}\boldsymbol{h}_\alpha$ (second moment).

Lemma 5.3. *The eigenvalues of $\dot{M}$ counted with geometric multiplicity are*

$$\operatorname{eigen}(\dot{M}) = (1, 1, 1, 0, 0, 0, -(d-1)^{-1}, \lambda_1, \lambda_2)$$

where $|\lambda_1| \leqslant d^{-1.9}$ and $0 < |\lambda_2 - (d-1)^{-1}| \leqslant d^{-1.2}$. The right eigenvector $\bar{x}$ corresponding to eigenvalue $-(d-1)^{-1}$ is given by $\bar{x}_\sigma = (d-1)\mathbf{1}\{\sigma = \mathtt{11}\} - \mathbf{1}\{\sigma = \mathtt{f0} \text{ or } \mathtt{ff}\}$.

Proof. The matrix $\dot{M}$ is $\bar{h}$-reversible and block diagonal with blocks $\dot{M}_1$ acting on $\{10, 1f\}$, $\dot{M}_f$ acting on $\{11, f0, ff\}$, and $\dot{M}_0$ acting on $\{f1, 01, 00, 0f\}$. We compute

$$\dot{M}_1 = \begin{array}{c} \\ 10 \\ 1f \end{array}\!\!\begin{array}{cc} 10 & 1f \\ \left(\begin{array}{cc} \frac{q_0}{1-q_1} & \frac{q_f}{1-q_1} \\ \frac{q_0}{1-q_1} & \frac{q_f}{1-q_1} \end{array}\right) \end{array}, \quad \dot{M}_f = \begin{array}{c} \\ 11 \\ f0 \\ ff \end{array}\!\!\begin{array}{ccc} 11 & f0 & ff \\ \left(\begin{array}{ccc} 0 & \frac{q_0}{1-q_1} & \frac{q_f}{1-q_1} \\ \frac{1}{d-1} & \frac{d-2}{d-1}\frac{q_0}{1-q_1} & \frac{d-2}{d-1}\frac{q_f}{1-q_1} \\ \frac{1}{d-1} & \frac{d-2}{d-1}\frac{q_0}{1-q_1} & \frac{d-2}{d-1}\frac{q_f}{1-q_1} \end{array}\right) \end{array}$$

so $\mathsf{eigen}(\dot{M}_1) = \{1, 0\}$ and $\mathsf{eigen}(\dot{M}_f) = \{1, 0, -(d-1)^{-1}\}$; the right eigenvector of $\dot{M}_f$ corresponding to eigenvalue $-(d-1)^{-1}$ is $(d-1, -1, -1)$. We also compute

$$\dot{M}_0 = \begin{array}{c} \\ f1 \\ 01 \\ 00 \\ 0f \end{array}\!\!\begin{array}{cccc} f1 & 01 & 00 & 0f \\ \left(\begin{array}{cccc} \frac{1}{d-1} & 0 & \frac{d-2}{d-1}\frac{q_0}{1-q_1} & \frac{d-2}{d-1}\frac{q_f}{1-q_1} \\ 0 & \varepsilon + q_1(1-\varepsilon) & q_0(1-\varepsilon) & q_f(1-\varepsilon) \\ \varepsilon & q_1(1-\varepsilon) & q_0(1-\varepsilon) & q_f(1-\varepsilon) \\ \varepsilon & q_1(1-\varepsilon) & q_0(1-\varepsilon) & q_f(1-\varepsilon) \end{array}\right) \end{array}$$

where (using $\bar{h}$-reversibility of $\dot{M}$, or alternatively the frozen model recursion)

$$\varepsilon = \frac{(d-2)q_1^2(1-q_1)^{d-3}}{1 - [1 + (d-2)q_1](1-q_1)^{d-2}} = \frac{d-2}{d-1}\frac{q_1 q_f}{(1-q_1)q_0} \asymp q_1 q_f.$$

Write $\bar{H}_0$ for the restriction of $\bar{H}$ to $\{f1, 01, 00, 0f\}$; then the symmetric matrix $\bar{H}_0^{1/2}\dot{M}_0\bar{H}_0^{-1/2}$ has spectral decomposition $\sum_{i=1}^4 \lambda_i e_i e_i^t$ with $\{\lambda_i\} = \mathsf{eigen}(\dot{M}_0)$. We can take

$$\lambda_3 = 1, \; e_3 = \bar{h}^{1/2} \quad \text{and} \quad \lambda_4 = 0, \; e_4 = (0, 0, q_f(1-q_1)^{-1}, -q_0(1-q_1)^{-1})^{1/2}.$$

For the other two eigenvalues, consider the following "almost" eigenvalue equations:

$$v_1 = ((d-2)q_1, 1-q_1, -q_1, -q_1), \quad \dot{M}_0 v_1 = \varepsilon(0, 1-q_1, (d-2)q_1, (d-2)q_1),$$

so v_1 is almost in the kernel of $\dot{M}_0$; and

$$v_2 = (1, 0, 0, 0), \quad (\dot{M}_0 - (d-1)^{-1}I)v_2 = \varepsilon(0, 0, 1, 1),$$

so v_2 is almost an eigenvector of $\dot{M}_0$ with eigenvalue $(d-1)^{-1}$. We have $\langle \bar{H}_0^{1/2}v_1, e_4\rangle = 0$ so $\bar{H}_0^{1/2}v_1 = \sum_i a_i e_i$ with $a_4 = 0$. We then calculate

$$\frac{\sum_i \lambda_i^2 a_i^2}{\sum_i a_i^2} = \frac{\|\bar{H}_0^{1/2}\dot{M}_0 v_1\|^2}{\|\bar{H}_0^{1/2}v_1\|^2} = \frac{\varepsilon^2 q_0 q_1(1-q_1)[q_1 + \frac{1-q_1}{(d-2)^2}]}{q_1[q_f q_1^2 + \frac{q_0(1-q_1)}{(d-2)^2}]} \asymp q_1(d\varepsilon)^2 \asymp q_1(q_f \log d)^2.$$

From Propn. 3.5 we have $\lambda \geqslant d^{0.59}$, and substituting into (28) gives $q_f \asymp dq_1^2/\lambda \leqslant d^{-1.58}$. It then follows from the above that at least one of the other two eigenvalues, say λ_1, must have absolute value $\leqslant d^{-2}$. Similarly, representing $\bar{H}_0^{1/2}v_2 = \sum_i b_i e_i$ gives

$$\frac{\sum_i (\lambda_i - (d-1)^{-1})^2 b_i^2}{\sum_i b_i^2} = \frac{\|\bar{H}_0^{1/2}(\dot{M}_0 - (d-1)^{-1}\cdot I)v_2\|^2}{\|\bar{H}_0^{1/2}v_2\|^2} = \frac{\varepsilon^2 q_0(1-q_1)}{q_1 q_f} \asymp q_1 q_f \leqslant d^{-2.5},$$

so the last eigenvalue λ_2 must satisfy $|\lambda_2 - (d-1)^{-1}| \leqslant d^{-1.2}$. Note however that

$$\det[\dot{M}_0 - (d-1)^{-1}\cdot I] = \frac{(d-2)\varepsilon}{(d-1)^3}[d(q_1 + \varepsilon - q_1\varepsilon) - (1 + q_1 + \varepsilon - q_1\varepsilon)] \asymp \varepsilon dq_1$$

so λ_2 does not exactly equal $(d-1)^{-1}$. $\qquad\square$

Proof of Propn. 5.2. Consider the first moment Hessian $H\mathbf{\Phi}(^\star\boldsymbol{h}_\alpha)$ on $\boldsymbol{\Delta}_\alpha$. For convenience, rewrite (63) as $\bar{H}^{-1/2}\bar{\delta} = \dot{S}\bar{H}^{1/2}\dot{\chi}$ with $\dot{S}$ the symmetric matrix $\dot{S} \equiv d^{-1}\bar{H}^{1/2}[I + (d-1)\dot{M}]\bar{H}^{-1/2}$. By the requirement that $\boldsymbol{h}, \boldsymbol{h} + \eta\boldsymbol{\delta}$ both lie in $\boldsymbol{\Delta}_\alpha^\circ$ for $|\eta|$ small, the permissible vectors $\bar{\delta}$ span a linear subspace $\overline{W}$. In particular, necessarily $\langle\bar{\delta}, \bar{x}\rangle = 0$ for the vector $\bar{x}$ which was determined in Lem. 5.3 to span the kernel of $\dot{L} \equiv d^{-1}[I + (d-1)\dot{M}]$. For such $\bar{\delta}$, (63) is easily solved by considering the invariant action of $\dot{S}$ on the subspace $(\bar{H}^{1/2}\bar{x})^\perp$: if $U \in \mathbb{R}^{9\times 8}$ with columns giving an orthonormal basis for this subspace, then for all $\bar{\delta}$ orthogonal to $\bar{x}$ (in particular, for all $\bar{\delta} \in \overline{W}$) we may define $\bar{H}^{1/2}\dot{\chi} = U(U^t\dot{S}U)^{-1}U^t\bar{H}^{-1/2}\bar{\delta}$, consequently the maximal value of $\partial_\eta^2\mathbf{\Phi}(\boldsymbol{h} + \eta\boldsymbol{\delta})|_{\eta=0}$ subject to marginals $\bar{\delta}$ is

$$-\langle\bar{\delta}, \dot{\chi}\rangle + (d/2)\langle(\bar{\delta}/\bar{h})^2\rangle_{\bar{h}} = -(U^t\bar{H}^{-1/2}\bar{\delta})^t\dot{Q}(U^t\bar{H}^{-1/2}\bar{\delta}), \quad \dot{Q} \equiv (U^t\dot{S}U)^{-1} - (d/2)\,I$$

The eigenvalues of $\dot{Q}$ are given by

$$\frac{d}{1 + (d-1)\lambda} - \frac{d}{2} \quad \text{for } \lambda \in \mathsf{eigen}(\dot{M})\backslash\{-(d-1)^{-1}\},$$

so $\dot{Q}$ is non-singular since we saw in Lem. 5.3 that $(d-1)^{-1} \notin \mathsf{eigen}(\dot{M})$. Since we proved in Thm. 3.4 that $^\star\boldsymbol{h}_\alpha \in \boldsymbol{\Delta}_\alpha^\circ$ is the global maximizer of $\mathbf{\Phi}$ on $\boldsymbol{\Delta}_\alpha$, the restriction of the above quadratic form to the space of permissible $\bar{\delta}$ (formally, to $U^t\bar{H}^{-1/2}\overline{W}$) must be negative semi-definite, so by non-singularity we see that it is in fact negative-definite.

The proof for the second moment Hessian $H(_2\mathbf{\Phi})(^\star_2\boldsymbol{h}_\alpha)$ on $_2\boldsymbol{\Delta}_\alpha$ is similar; note $^\star_2\boldsymbol{h}_\alpha = {}^\star\boldsymbol{h}_\alpha \otimes {}^\star\boldsymbol{h}_\alpha$ implies $\dot{M}_2 = \dot{M} \otimes \dot{M}$, with eigenvalues $\{\lambda\lambda' : \lambda, \lambda' \in \mathsf{eigen}(\dot{M})\}$. The kernel of $d^{-1}[I + (d-1)\dot{M}_2]$ is spanned by vectors $\bar{x} \otimes \bar{y}$ or $\bar{y} \otimes \bar{x}$ with $\bar{x}$ as before and $\bar{y}$ any right eigenvector of M_2 with eigenvalue 1, and again the permissible measures $\bar{\delta}$ must be orthogonal to the kernel. Negative-definiteness then follows as above from the observation that $(d-1)^{-1} \notin \mathsf{eigen}(\dot{M}_2)$. $\qquad\square$

Lemma 5.4. *For any* $\sigma, \sigma' \in \mathscr{M}$ *there exists a signed integer measure* $\boldsymbol{\delta} \equiv \boldsymbol{\delta}_{\sigma-\sigma'} = (\dot{\boldsymbol{\delta}}, \hat{\boldsymbol{\delta}})$ *with* $\mathrm{supp}\,\boldsymbol{\delta} \subseteq \mathrm{supp}\,\varphi$ *such that*

$$0 = \langle\dot{\boldsymbol{\delta}}, 1\rangle = \langle\hat{\boldsymbol{\delta}}, 1\rangle = \mathbf{i}(\boldsymbol{\delta}) \quad and \quad \dot{H}\dot{\boldsymbol{\delta}} - \hat{H}\hat{\boldsymbol{\delta}} = \mathbf{1}_\sigma - \mathbf{1}_{\sigma'}.$$

The analogous condition holds for the support of the second-moment factors φ_2.

Proof. Define a graph on $\mathscr{M}^2$ by placing an edge between τ and τ' if and only if there exists $\boldsymbol{\delta}_{\tau-\tau'}$ satisfying the requirements in the statement of the lemma. We shall show this graph is connected (hence complete). By considering $\mathrm{supp}\,\dot{\varphi}_2 = (\mathrm{supp}\,\dot{\varphi})^2$ (see (30)) it is clear that $\{\substack{10\\\sigma}, \substack{1f\\\sigma}\}$, $\{\substack{f0\\\sigma}, \substack{ff\\\sigma}\}$, and $\{\substack{00\\\sigma}, \substack{01\\\sigma}, \substack{0f\\\sigma}\}$ are connected subsets of $\mathscr{M}^2$ for any $\sigma \in \mathscr{M}$ — for example, $\substack{10\\\sigma}$ and $\substack{1f\\\sigma}$ can be connected via

$$\boldsymbol{\delta}_{\substack{10\\\sigma} - \substack{1f\\\sigma}} = (\dot{\boldsymbol{\delta}}, \hat{\boldsymbol{\delta}}) = \left(\mathbf{1}\{(\substack{10\ 10^{d-1}\\\sigma\quad\underline{\sigma}})\} - \mathbf{1}\{(\substack{1f\ 10^{d-1}\\\sigma\quad\underline{\sigma}})\}, 0\right)$$

for any choice of $\underline{\sigma} \in \mathscr{M}^{d-1}$ with $(\sigma, \underline{\sigma}) \in \mathrm{supp}\,\dot{\varphi}$. Then $\substack{00\\\sigma}$ and $\substack{10\\\sigma}$ can be connected via

$$\boldsymbol{\delta}_{\substack{00\\\sigma} - \substack{10\\\sigma}} = \boldsymbol{\delta}_{\substack{01\\R\sigma} - \substack{00\\R\sigma}} + \left(0, \mathbf{1}\{(\substack{10\ 01\\\sigma\ R\sigma})\} - \mathbf{1}\{(\substack{00\ 00\\\sigma\ R\sigma})\}\right);$$

the spins $\substack{00\\\sigma}$ and $\substack{f0\\\sigma}$ can be similarly connected. We can then connect $\substack{f1\\\sigma}$ and $\substack{f0\\\sigma}$ via

$$\boldsymbol{\delta}_{\substack{f1\\\sigma} - \substack{01\\\sigma}} = \left(\mathbf{1}\{(\substack{f1\ f1\ 00\ 00^{d-3}\\\sigma\ \sigma_2\ \sigma_3\quad\underline{\sigma}})\} - \mathbf{1}\{(\substack{01\ 01\ 01\ 00^{d-3}\\\sigma\ \sigma_2\ \sigma_3\quad\underline{\sigma}})\}, \mathbf{1}\{(\substack{f1\ 1f\\\sigma_2\ R\sigma_2})\} - \mathbf{1}\{(\substack{01\ 10\\\sigma_2\ R\sigma_2})\}\right)$$
$$+ \boldsymbol{\delta}_{\substack{1f\\R\sigma_2} - \substack{10\\R\sigma_2}} + \boldsymbol{\delta}_{\substack{01\\\sigma_3} - \substack{00\\\sigma_3}}$$

for any choice of σ_2, σ_3, and $\underline{\sigma}$ such that $(\sigma, \sigma_2, \sigma_3, \underline{\sigma}) \in \operatorname{supp} \dot\varphi$. Lastly, to connect $\overset{11}{\sigma}$ and $\overset{\texttt{f1}}{\sigma}$, find σ_2 and $\underline{\sigma} \equiv (\sigma_3, \ldots, \sigma_d)$ with $(\sigma, \sigma_2, \underline{\sigma}) \in \operatorname{supp} \dot\varphi$, and take

$$\boldsymbol{\delta}_{\overset{11}{\sigma} - \overset{\texttt{f1}}{\sigma}} = \left(\mathbf{1}\{(\overset{11}{\sigma}\,\overset{\texttt{f0}}{\sigma_2}\,\overset{\texttt{f0}^{d-2}}{\underline{\sigma}})\} - \mathbf{1}\{(\overset{\texttt{f1}}{\sigma}\,\overset{\texttt{f1}}{\sigma_2}\,\overset{00^{d-2}}{\underline{\sigma}})\}, 0 \right) + \boldsymbol{\delta}_{\overset{\texttt{f1}}{\sigma_2} - \overset{\texttt{f0}}{\sigma_2}} + \sum_{j=3}^d \boldsymbol{\delta}_{\overset{00}{\sigma_j} - \overset{\texttt{f0}}{\sigma_j}}$$

This shows that for any $s, \sigma \in \mathscr{M}$, $\overset{s}{\sigma}$ is connected to $\overset{00}{\sigma}$, which in turn by symmetry is connected to $\overset{00}{00}$. This proves connectivity of $\mathscr{M}^2$; connectivity of $\mathscr{M}$ trivially follows. $\square$

Proof of Thms. 2.5 and 5.1. Recall (Defn. 3.2) that $\mathbb{E}\boldsymbol{Z}_{n\alpha}$ is the sum of $\mathbb{E}\boldsymbol{Z}(\boldsymbol{h})$ over probability measures $\boldsymbol{h} \equiv (\dot{\boldsymbol{h}}, \hat{\boldsymbol{h}})$ on $\operatorname{supp}\psi$ such that $\boldsymbol{g} \equiv (\dot{\boldsymbol{g}}, \hat{\boldsymbol{g}}) \equiv (n\dot{\boldsymbol{h}}, \frac{nd}{2}\hat{\boldsymbol{h}})$ is integer-valued, and lies in the kernel of $H_{\boldsymbol{\Delta}} \equiv (\dot{H} \quad -\hat{H})$ with intensity $\mathbf{i}(\boldsymbol{g}) = n\alpha$. Let $\mathbb{E}[^\star\boldsymbol{Z}_{n\alpha}]$ denote the restriction of this sum to measures $\boldsymbol{g}$ within a euclidean ball of radius $n^{1/2} \log n$ centered at $^\star\boldsymbol{g}_\alpha$; then Propn. 5.2 implies $\mathbb{E}\boldsymbol{Z}_{n\alpha} = e^{o(1)} \mathbb{E}[^\star\boldsymbol{Z}_{n\alpha}]$.

Lem. 5.4 shows that the integer matrix $H_{\boldsymbol{\Delta}}$ defines a surjection from $L' \equiv \{\boldsymbol{\delta} \in \mathbb{R}^{\operatorname{supp}\psi} : \langle \dot{\boldsymbol{\delta}}, 1 \rangle = \langle \hat{\boldsymbol{\delta}}, 1 \rangle = \mathbf{i}(\boldsymbol{\delta}) = 0\}$ to $\{\bar{\delta} \in \mathbb{R}^{\mathscr{M}} : \langle \bar{\delta}, 1 \rangle = 0\}$, so we see that $L \equiv L' \cap (\ker H_{\boldsymbol{\Delta}}) \cap \mathbb{Z}^{\operatorname{supp}\psi}$ is an $(\dot{\boldsymbol{s}} - 2)$-dimensional lattice with spacings $\asymp_d 1$. The measures $\boldsymbol{g}$ contributing to $\mathbb{E}[^\star\boldsymbol{Z}_{n\alpha}]$ are given by the intersection of the euclidean ball $\{\|\boldsymbol{g} - {}^\star\boldsymbol{g}_\alpha\| \leqslant n^{1/2} \log n\}$ with an affine translation of L. The expansion (32) then shows that $n^{1/2}\mathbb{E}[^\star\boldsymbol{Z}_{n\alpha}]$ defines a convergent Riemann sum, implying Thm. 2.5.

The same argument implies that the contribution $\mathbb{E}[^\star_2\boldsymbol{Z}_{n\alpha}]$ to the second moment from measures in $^\star_2\boldsymbol{\Delta}_\alpha$ is $\asymp_d n^{-1} \exp\{n\,[_2\Phi(^\star_2\boldsymbol{h}_\alpha)]\} \asymp_d (\mathbb{E}\boldsymbol{Z}_{n\alpha})^2$, so combining with Thm. 4.1 and Propn. 4.10 gives $\mathbb{E}[\boldsymbol{Z}_{n\alpha}^2] \asymp_d (\mathbb{E}\boldsymbol{Z}_{n\alpha})^2 + \mathbb{E}\boldsymbol{Z}_{n\alpha}$. $\square$

6. Constant fluctuations

In this section we prove Thm. 2 by a variance decomposition of $\boldsymbol{Z}_{n\alpha}$. An application of our method in a simpler setting appears in [15, §8], where the decomposition is introduced in greater detail. In what follows we write $\boldsymbol{Z}_{n\alpha}$ for the partition function of auxiliary model configurations $\underline{\sigma}$ with intensity $\mathbf{i}(\underline{\sigma}) = n\alpha$ and empirical measure $\boldsymbol{h}$ within distance $n^{-1/2}\log n$ of the first moment maximizer $^\star\boldsymbol{h}_\alpha$. In this section we restrict attention to the frozen model at intensities $n\alpha_{\mathrm{lbd}} \leqslant n\alpha \leqslant \bar{\boldsymbol{A}}_n$, and we write $f = O_d(g)$ or $f \lesssim_d g$ to indicate $f \leqslant C(d)\,g$ where $C(d)$ can be chosen uniformly over α in this regime.

Proposition 6.1. *It holds uniformly over $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \bar{\boldsymbol{A}}_n/n$ that*

$$\operatorname{Var} \log(\boldsymbol{Z}_{n\alpha} + \varepsilon\,\mathbb{E}\boldsymbol{Z}_{n\alpha}) \lesssim_d 1 + c_{d,\varepsilon}(\mathbb{E}\boldsymbol{Z}_{n\alpha})^{-1}$$

Proof of Thm. 2. Write $\mathbf{L}_{n\alpha}^\varepsilon \equiv \log(\boldsymbol{Z}_{n\alpha} + \varepsilon\,\mathbb{E}\boldsymbol{Z}_{n\alpha})$. Since Thm. 2.5 and Thm. 5.1 together imply $\mathbb{E}[(\boldsymbol{Z}_{n\alpha})^2] \asymp_d (\mathbb{E}\boldsymbol{Z}_{n\alpha})^2$ uniformly over $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \bar{\boldsymbol{A}}_n/n$, there must exist a constant $\delta \equiv \delta(d) > 0$ for which $\mathbb{P}(\boldsymbol{Z}_{n\alpha} \geqslant \delta\,\mathbb{E}\boldsymbol{Z}_{n\alpha}) \geqslant \delta$, hence $\mathbb{E}\mathbf{L}_{n\alpha}^\varepsilon \geqslant \delta\log\delta + (1-\delta)\log\varepsilon + \log\mathbb{E}\boldsymbol{Z}_{n\alpha}$. On the other hand, $\boldsymbol{Z}_{n\alpha} \leqslant \varepsilon\,\mathbb{E}\boldsymbol{Z}_{n\alpha}$ if and only if $\mathbf{L}_{n\alpha}^\varepsilon \leqslant \log(2\varepsilon\,\mathbb{E}\boldsymbol{Z}_{n\alpha})$, which for small $\varepsilon > 0$ is much less than the lower bound on $\mathbb{E}\mathbf{L}_{n\alpha}^\varepsilon$. Applying Chebychev's inequality and Propn. 6.1 therefore gives

$$\mathbb{P}(\boldsymbol{Z}_{n\alpha} \leqslant \varepsilon\,\mathbb{E}\boldsymbol{Z}_{n\alpha}) \leqslant \frac{\operatorname{Var}\mathbf{L}_{n\alpha}^\varepsilon}{(\mathbb{E}\mathbf{L}_{n\alpha}^\varepsilon - \log(2\varepsilon\,\mathbb{E}\boldsymbol{Z}_{n\alpha}))^2} \lesssim_d \frac{1 + c_{d,\varepsilon}(\mathbb{E}\boldsymbol{Z}_{n\alpha})^{-1}}{(\delta\log(\delta/\varepsilon) - \log 2)^2}.$$

With $c_\star$ as above (see Cor. 2.6), define $C_{d,\varepsilon} \equiv 4c_\star[(\log c_{d,\varepsilon}) \vee 0]$: it then follows from Thm. 2.5 that for all $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \bar{\boldsymbol{A}}_n - C_{d,\varepsilon}$ that

$$c_{d,\varepsilon}(\mathbb{E}\boldsymbol{Z}_{n\alpha})^{-1} \lesssim_d 1, \quad \text{therefore } \mathbb{P}(\boldsymbol{Z}_{n\alpha} = 0) \lesssim (\delta\log(\delta/\varepsilon) - \log 2)^{-2}.$$

The right-hand side tends to zero as ε decreases to zero, proving the theorem. $\square$

We prove Propn. 6.1 by controlling the increments of the Doob martingale of the random variable $\mathbf{L}_{n\alpha}^{\varepsilon} \equiv \log(\mathbf{Z}_{n\alpha} + \varepsilon \mathbb{E}\mathbf{Z}_{n\alpha})$ with respect to the edge-revealing filtration $(\mathscr{F}_i)_{1 \leqslant i \leqslant m}$, $m \equiv (nd/2)$, for the graph $G_n \sim \mathcal{G}_{n,d}$. For $i \leqslant m - 2$ let G° denote the graph with edges $a_1, \ldots, a_{i-1}$. Randomly partition the remaining unmatched half-edges into disjoint sets $\mathscr{K}$ and $\mathscr{W}$ with $|\mathscr{K}| = 4$. We showed in [15, §8] by a coupling argument that for any $m' \leqslant m - 2$,

$$\begin{aligned} &\operatorname{Var}\mathbf{L}_{n\alpha}^{\varepsilon} \leqslant m' \max_{A, \acute{A}} \mathbb{E}[|\mathbf{V}_{n\alpha,m'}^{\varepsilon}|^2] + \sum_{i=m'+1}^{m-1} \max_{A, \acute{A}} \mathbb{E}[|\mathbf{V}_{n\alpha,i}^{\varepsilon}|^2] \\ &\text{for } \mathbf{V}_{n\alpha,i}^{\varepsilon} \equiv \mathbb{E}_{i-1}[\log(\mathbf{Z}_{n\alpha} + \varepsilon \mathbb{E}\mathbf{Z}_{n\alpha}) - \log(\acute{\mathbf{Z}}_{n\alpha} + \varepsilon \mathbb{E}\mathbf{Z}_{n\alpha})] \end{aligned} \tag{65}$$

where $\mathbb{E}_{i-1}$ is expectation conditioned on the graph G°; and $\mathbf{Z}_{n\alpha} \equiv \mathbf{Z}_{n\alpha}(A), \acute{\mathbf{Z}}_{n\alpha} \equiv \mathbf{Z}_{n\alpha}(\acute{A})$ are the partition functions for coupled completions $G, \acute{G}$ of G° to d-regular graphs: G is formed by placing random edges A_W on $\mathscr{W}$ and the given edges A on $\mathscr{K}$, while $\acute{G}$ is formed by placing the *same* edges A_W on $\mathscr{W}$ and the given edges $\acute{A}$ on $\mathscr{K}$. (Note that $\mathbf{V}_{n\alpha,i}^{\varepsilon}$ depends on $A, \acute{A}$ though we suppress this from the notation.)

Assume throughout that $m' \leqslant m - 2$ with $m - m' \lesssim_d \log n$. For such m', we showed in [15, §8] that the sum appearing in (65) has two dominant components: the first is an "independent-copies contribution" coming from pair configurations with empirical measure near $\overset{\star}{_2}\boldsymbol{h}_\alpha \equiv {}^\star\boldsymbol{h}_\alpha \otimes {}^\star\boldsymbol{h}_\alpha$; we showed that this contribution can be controlled under some abstract conditions (Lem. 6.2 and Cor. 6.3 below). The other component is an "identical-copies contribution" coming from closely correlated pair configurations: this was controlled in [15, §8] under the assumption that the first moment is exponentially large, and the main work of this section is to control the identical-copies contribution assuming only that the first moment is bounded below by a large constant. Before turning to this we first show in §6.1 that the independent-copies contribution is controlled by a straightforward application of the method of [15, §8].

6.1. Independent-copies contribution.

Let $B_\ell^\circ(\mathscr{K})$ denote the ball of graph distance ℓ about $\mathscr{K}$ in the graph G°; the leaves of $B_\ell^\circ(\mathscr{K})$ are half-edges incident to variables (ℓ odd) or clauses (ℓ even). Set

$$T = B_{2t'}^\circ(\mathscr{K}), \quad t' = t \wedge \min\{\ell : B_{2\ell}^\circ \text{ has fewer than } |\mathscr{K}| \text{ connected components}\}. \tag{66}$$

$\mathscr{W}$ can only intersect T in its leaves; and we shall let $\mathscr{U}$ denote the leaves of T without $\mathscr{W}$.

Regarding G as a $(d, 2)$-regular graph of variables V (degree d) and clauses F (degree 2), for $S \subseteq V \cup F$ and any configuration $\underline{\sigma}_S$ on the edges incident to the vertices $s \in S$, let

$$\Psi_{S,l}^\lambda(\underline{\sigma}_S) \equiv \mathbf{1}\{\mathbf{i}(\underline{\sigma}_S) = l\} \prod_{v \in S \cap V} \dot{\varphi}^\lambda(\underline{\dot{\sigma}}_v) \prod_{a \in S \cap F} \hat{\varphi}^\lambda(\underline{\hat{\sigma}}_a) \in \{0, \lambda^l\};$$

we write $\Psi_{S,l}$ for the unweighted version given by taking $\lambda = 1$. Then, with $\mathscr{Y} \equiv \mathscr{U} \cup \mathscr{W}$,

$$\mathbf{Z}_{n\alpha} = \sum_{r,j} \sum_{\underline{\sigma}_{\mathscr{Y}}} \Psi_{A_W,j}(\underline{\sigma}_{\mathscr{W}}) \kappa_{r-j}(\underline{\sigma}_{\mathscr{U}}) \mathbf{Z}_{n\alpha-r}^\partial[\underline{\sigma}_{\mathscr{Y}}] = \lambda^{-n\alpha} \sum_{r,j} \sum_{\underline{\sigma}_{\mathscr{U}}} \Psi_{A_W,j}^\lambda(\underline{\sigma}_{\mathscr{W}}) \kappa_{r-j}^\lambda(\underline{\sigma}_{\mathscr{U}}) \mathbf{Z}_{n\alpha-r}^{\partial,\lambda}[\underline{\sigma}_{\mathscr{Y}}]$$

where κ_l (resp. $\acute{\kappa}_l$) is the number of configurations on $T \cup A$ (resp. $T \cup \acute{A}$) with intensity l, with λ-weighted versions denoted by $\kappa_l^\lambda, \acute{\kappa}_l^\lambda$. Though we suppress it from the notation, let us emphasize that κ_l depends on T and also on $\underline{\sigma}_{\mathscr{W}}$, since T may intersect $\mathscr{W}$. Writing $\underline{\sigma}_T \sim \underline{\sigma}_{\mathscr{Y}}$ to indicate that $\underline{\sigma}_T$ and $\underline{\sigma}_{\mathscr{Y}}$ agree on the leaves of T,

$$\kappa_l(\underline{\sigma}_{\mathscr{U}}) \equiv \kappa_{T,l}(\underline{\sigma}_{\mathscr{U}}|\underline{\sigma}_{\mathscr{W}}) \equiv \sum_{\underline{\sigma}_T \sim \underline{\sigma}_{\mathscr{Y}}} \Psi_{T \cup A,l}(\underline{\sigma}_T).$$

Let $\boldsymbol{Z}_{n\alpha}^{\partial,\lambda}$ denote the sum of $\boldsymbol{Z}_{n\alpha}^{\partial,\lambda}[\underline{\sigma}_{\mathscr{Y}}]$ over all $\underline{\sigma}_{\mathscr{Y}}$, and $\boldsymbol{Z}_{n\alpha}^{\partial,\lambda}[\cdot,\underline{\sigma}_{\mathscr{W}}]$ the sum over the $\mathscr{U}$-coordinates only. Write $\mathbb{E}_T$ for expectation over the law of the graph G^∂ conditioned on T. As shown in [15, §8], Lem. 5.4 can be used to prove

$$\mathbb{E}_T[\boldsymbol{Z}_{n\alpha-r}^{\partial,\lambda}[\underline{\sigma}_{\mathscr{U}},\underline{\sigma}_{\mathscr{W}}]] \sim \boldsymbol{p}^\lambda(\underline{\sigma}_{\mathscr{U}})\,\mathbb{E}_T[\boldsymbol{Z}_{n\alpha-r}^{\partial,\lambda}[\cdot,\underline{\sigma}_{\mathscr{W}}]] \sim \boldsymbol{p}^\lambda(\underline{\sigma}_{\mathscr{U}})\,\mathbb{E}_T[\boldsymbol{Z}_{n\alpha-r}^{\partial,\lambda}] \tag{67}$$

where $\boldsymbol{p}^\lambda$ indicates the product measure with marginals $\dot{h}^\lambda$ coming from the Bethe solution corresponding to $^\star\boldsymbol{h}^\lambda = {}^\star\boldsymbol{h}_\alpha$. We then decompose $\boldsymbol{Z}_{n\alpha}$ according to a Fourier basis for $L^2(\mathscr{M}^{\mathscr{U}},\boldsymbol{p}^\lambda)$: take $(\boldsymbol{b}_1^\lambda,\ldots,\boldsymbol{b}_{|\mathscr{M}|}^\lambda)$ to be an orthonormal basis for $L^2(\mathscr{M},\dot{h}^\lambda)$ with $\boldsymbol{b}_1^\lambda \equiv 1$. Then the functions $\boldsymbol{b}_{\underline{s}}^\lambda(\underline{\sigma}_{\mathscr{U}}) \equiv \prod_{u\in\mathscr{U}} \boldsymbol{b}_{s(u)}^\lambda(\sigma_u)$ ($\underline{s} \in [|\mathscr{M}|]^{\mathscr{U}}$) form an orthonormal basis for $L^2(\mathscr{M}^{\mathscr{U}},\boldsymbol{p}^\lambda)$. By Plancherel's identity

$$\boldsymbol{Z}_{n\alpha} = \lambda^{-n\alpha}\sum_{j,l}\sum_{\underline{\sigma}_{\mathscr{W}}}\Psi_{A_W,j}^\lambda(\underline{\sigma}_{\mathscr{W}})\sum_{\underline{s}}(\kappa_l^\lambda)_{\underline{s}}^\wedge(\mathbf{F}_{n\alpha-j-l}^\lambda)_{\underline{s}}^\wedge \tag{68}$$

where $^\wedge$ indicates the Fourier transform with respect to the basis $\boldsymbol{b}^\lambda$, and

$$\mathbf{F}_{n\alpha-r}^\lambda(\underline{\sigma}_{\mathscr{U}}) \equiv \mathbf{F}_{n\alpha-r}^\lambda(\underline{\sigma}_{\mathscr{U}}|\underline{\sigma}_{\mathscr{W}}) \equiv \frac{\boldsymbol{Z}_{n\alpha-r}^{\partial,\lambda}[\underline{\sigma}_{\mathscr{U}},\underline{\sigma}_{\mathscr{W}}]}{\boldsymbol{p}^\lambda(\underline{\sigma}_{\mathscr{U}})}.$$

By the method of [15, §8] applied to the decomposition (68), the independent-copies contribution to $\operatorname{Var}\mathbf{L}_{n\alpha}^\varepsilon$ is controlled subject to a few conditions which are proved in the following lemma. For $\underline{s} \in [|\mathscr{M}|]^{\mathscr{U}}$ write $|\underline{s}| \equiv |\{u \in \mathscr{U} : s_u \neq 1\}|$, and write $\varnothing$ for the identically-1 vector in $[|\mathscr{M}|]^{\mathscr{U}}$.

Lemma 6.2. *The auxiliary model satisfies the following:*

(a) (Symmetries) On the event $\mathbf{T}$ that T consists of 4 disjoint tree components with $T \cap \mathscr{W} = \varnothing$, we have $(\kappa_l^\lambda)_{\underline{s}}^\wedge = (\acute{\kappa}_l^\lambda)_{\underline{s}}^\wedge$ for all $|\underline{s}| \leq 1$; and the zeroth order Fourier coefficient $(\kappa_l^\lambda)_\varnothing^\wedge = \sum_{\underline{\sigma}_{\mathscr{U}}}\kappa_l^\lambda(\underline{\sigma}_{\mathscr{U}})\boldsymbol{p}^\lambda(\underline{\sigma}_{\mathscr{U}})$ restricted to the event $\mathbf{T}$ takes a constant value $(\overline{\kappa}_l^\lambda)_\varnothing^\wedge$ which does not depend on the edges A. On the event $\mathbf{C}^\circ$ that T either contains a single cycle or has a single intersection with $\mathscr{W}$ (but not both), we have $(\kappa_l^\lambda)_\varnothing^\wedge = (\acute{\kappa}_l^\lambda)_\varnothing^\wedge$.

(b) (Correlation decay) For $\alpha \geq \alpha_{\mathrm{lbd}}$, the Gibbs measure ν corresponding to $^\star\boldsymbol{h}_\alpha$ on the infinite d-regular tree has correlation decay at rate faster than square root of the branching rate: if edges e, e' are separated by t intermediate vertices then $|\operatorname{Cov}_\nu(f(\sigma_e),g(\sigma_{e'}))| \lesssim_d \|fg\|_\infty d^{-(3/2)t} \ll \|fg\|_\infty(d-1)^{-t/2}$.

Proof. (a) Follows by the argument of [15, §8] (simpler in the current setting since random literals are not involved).

(b) Let N denote the event that for some intermediate vertex v along the path γ joining e, e' in the tree, there exist at least two vertices $w \in \partial v\backslash\gamma$ with $\sigma_{w\to v} = 1$. The spins $\sigma_e, \sigma_{e'}$ are exactly independent on the event N, which has probability $\nu(N) \geq 1 - d^{-(3/2)t}$. $\square$

As shown in [15, §8], the conditions of Lems. 5.4 and 6.2 taken together give control over the independent-copies component of $\operatorname{Var}\mathbf{L}_{n\alpha}^\varepsilon$. Explicitly, let

$$\widetilde{\kappa}_l \equiv \kappa_l + \acute{\kappa}_l + (\kappa_l + \acute{\kappa}_l)_\varnothing^\wedge = \kappa_l + \acute{\kappa}_l + \lambda^{-l}(\kappa_l^\lambda + \acute{\kappa}_l^\lambda)_\varnothing^\wedge,$$

and let $\overline{\psi}_j(\underline{\sigma}_{\mathscr{W}})$ denote the expectation of $\Psi_{A_W,j}(\underline{\sigma}_{\mathscr{W}})$ over the random edges A_W. Then define

$$\boldsymbol{v}_{n\alpha,i}^T \equiv \sum_{r^1,r^2}\sum_{j^1,j^2}\sum_{\underline{\tau}_{\mathscr{Y}}}\overline{\psi}_{j^1}(\underline{\sigma}_{\mathscr{W}}^1)\overline{\psi}_{j^2}(\underline{\sigma}_{\mathscr{W}}^2)\widetilde{\kappa}_{r^1-j^1}(\underline{\sigma}_{\mathscr{U}}^1)\widetilde{\kappa}_{r^2-j^2}(\underline{\sigma}_{\mathscr{U}}^2)\ddot{\,}_2\boldsymbol{Z}_{n\alpha-(r^1,r^2)}^\partial[\underline{\tau}_{\mathscr{Y}}] \tag{69}$$

where $\ddot{{}_2}\boldsymbol{Z}^{\partial}_{n\alpha-(r^1,r^2)}$ refers to the contribution to the pair partition function on $G^{\partial} \equiv G^{\circ}\backslash T$ from empirical measures within distance $n^{-1/5}$ of $\ddot{{}_2}\boldsymbol{h}_{\alpha}$, with intensity $n\alpha - r^i$ in the i-th coordinate for $i = 1, 2$. Then [15, §8] together with Lems. 5.4 and 6.2 implies the following

Corollary 6.3. *For $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \bar{\boldsymbol{A}}_n/n$ and T as in* (66) *with $t \equiv t(d, \varepsilon) = 4\log_d(1/\varepsilon)$,*

$$\operatorname{Var}\mathbf{L}^{\varepsilon}_{n\alpha} \lesssim_d 1 + o_{\varepsilon}(1) + \underbrace{\frac{m'\mathbb{E}\boldsymbol{v}^T_{n\alpha,m'} + \sum_{i=m'+1}^{m-2}\mathbb{E}\boldsymbol{v}^T_{n\alpha,i}}{(\varepsilon\,\mathbb{E}\boldsymbol{Z}_{n\alpha})^2}}_{\textit{identical-copies component}}$$

for any m' with $m - m' \lesssim_d \log n$.

The remainder of this section is devoted to the handling of the identical-copies component. Let us first note that Lem. 5.4 can also be used to see how (67) varies with $n\alpha - r$: it is clear from the lemma that we can find signed integer measures $\boldsymbol{\delta}^{\mathrm{v}}$, $\boldsymbol{\delta}^{\mathrm{c}}$, $\boldsymbol{\delta}^{\mathrm{i}}$, all on the support $\operatorname{supp}\varphi$ of the first-moment auxiliary model factors (30), such that

$$\begin{array}{llll} \langle\dot{\boldsymbol{\delta}}^{\mathrm{v}}, 1\rangle = 1, & \langle\hat{\boldsymbol{\delta}}^{\mathrm{v}}, 1\rangle = 0, & \mathbf{i}(\hat{\boldsymbol{\delta}}^{\mathrm{v}}) = 0, & \dot{H}\dot{\boldsymbol{\delta}}^{\mathrm{v}} - \hat{H}\hat{\boldsymbol{\delta}}^{\mathrm{v}} = d\,\mathbf{1}_{00} \\ \langle\dot{\boldsymbol{\delta}}^{\mathrm{c}}, 1\rangle = 0, & \langle\hat{\boldsymbol{\delta}}^{\mathrm{c}}, 1\rangle = 1, & \mathbf{i}(\hat{\boldsymbol{\delta}}^{\mathrm{c}}) = 0, & \dot{H}\dot{\boldsymbol{\delta}}^{\mathrm{c}} - \hat{H}\hat{\boldsymbol{\delta}}^{\mathrm{c}} = 2\,\mathbf{1}_{00} \\ \langle\dot{\boldsymbol{\delta}}^{\mathrm{i}}, 1\rangle = 0, & \langle\hat{\boldsymbol{\delta}}^{\mathrm{i}}, 1\rangle = 0, & \mathbf{i}(\hat{\boldsymbol{\delta}}^{\mathrm{i}}) = 1, & \dot{H}\dot{\boldsymbol{\delta}}^{\mathrm{c}} - \hat{H}\hat{\boldsymbol{\delta}}^{\mathrm{i}} = 0. \end{array} \tag{70}$$

For example, taking $\boldsymbol{\delta}_{\mathtt{01-10}}$ as given by Lem. 5.4, we can take $\boldsymbol{\delta}^{\mathrm{i}}$ to be

$$\boldsymbol{\delta}^{\mathrm{i}} = (\mathbf{1}\{(\mathtt{10}^d)\} - \mathbf{1}\{(\mathtt{01}^d)\}, 0) + d\,\boldsymbol{\delta}_{\mathtt{01-10}}.$$

If G^{∂} has n^{∂} variables and m^{∂} clauses, then set

$$\boldsymbol{\delta}^{\sigma_{\mathscr{Y}}}_r = (n - n^{\partial})\boldsymbol{\delta}^{\mathrm{v}} + (m - m^{\partial})\boldsymbol{\delta}^{\mathrm{c}} + \sum_{y\in\mathscr{Y}}\boldsymbol{\delta}_{\sigma_y-\mathtt{00}} + r\boldsymbol{\delta}^{\mathrm{i}}.$$

Away from the simplex boundary, empirical measures contributing to $\boldsymbol{Z}^{\partial}_{n\alpha-r}[\underline{\sigma}_{\mathscr{Y}}]$ can be parametrized as $\boldsymbol{g} - \boldsymbol{\delta}^{\sigma_{\mathscr{Y}}}_r$ where $\boldsymbol{g}$ runs over empirical measures contributing to $\boldsymbol{Z}_{n\alpha}$. Write $\underline{\sigma}_{\mathscr{Y}} \equiv (\underline{\boldsymbol{o}}, \underline{\boldsymbol{i}})$ where $\underline{\boldsymbol{o}} \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}^{\mathscr{Y}}$ are the outgoing messages of $\underline{\sigma}^i_{\mathscr{Y}}$, and $\underline{\boldsymbol{i}} \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}^{\mathscr{Y}}$ the incoming. Then

$$\frac{\mathbb{E}\boldsymbol{Z}^{\partial}_{n\alpha-r}[\underline{\sigma}_{\mathscr{Y}}]}{\lambda^r\mathbb{E}\boldsymbol{Z}_{n\alpha}} = \frac{\mathbb{E}\boldsymbol{Z}^{\partial,\lambda}_{n\alpha-r}[\underline{\sigma}_{\mathscr{Y}}]}{\mathbb{E}\boldsymbol{Z}^{\lambda}_{n\alpha}} \sim \boldsymbol{c}^{-1}\,\boldsymbol{p}^{\lambda}(\underline{\sigma}_{\mathscr{Y}}) = \frac{\boldsymbol{c}^{-1}\,\boldsymbol{q}(\underline{\boldsymbol{o}})\lambda^{-|\{y:\sigma_y=\mathtt{11}\}|}}{[3 - (1 - \lambda^{-1})q_1]^{|\mathscr{Y}|}}. \tag{71}$$

for $\boldsymbol{c}$ a proportionality constant depending only on $n - n^{\partial}$ and $m - m^{\partial}$, and $\boldsymbol{q}$ the product measure with marginals $q \equiv (q_0, q_1, q_{\mathtt{f}})$.

6.2. **Refined rigidity estimate.** In this subsection we prove a refined version of Propn. 4.12 for small Δ.

Definition 6.4. Given a pair frozen configuration $(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2)$ on G^{∂}, let $\mathfrak{T}_{\mathtt{1f}}$ denote the collection of tree components of $V^{\partial}_{\neq}$ which intersect $\mathcal{Y}$ in a single $\mathtt{1f}$-vertex with matched partner in $\mathcal{Y}^{\star}$, and consist entirely of $\mathtt{1f}$-$\mathtt{0f}$ matched pairs. Let $\mathfrak{T}_{\mathtt{0f}}$ denote the collection of tree components of $V^{\partial}_{\neq}$ which intersect $\mathcal{Y}$ in a single $\mathtt{0f}$-vertex with matched partner in $\mathcal{R}$, and consist entirely of $\mathtt{1f}$-$\mathtt{0f}$ matched pairs except for one additional vertex in $\mathcal{Y}^{\star}$. Define symmetrically $\mathfrak{T}_{\mathtt{f1}}$ and $\mathfrak{T}_{\mathtt{f0}}$, and let $\mathfrak{T} \equiv \mathfrak{T}(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2)$ denote the union of $\mathfrak{T}_{\mathtt{1f}}$, $\mathfrak{T}_{\mathtt{0f}}$, $\mathfrak{T}_{\mathtt{f1}}$, $\mathfrak{T}_{\mathtt{f0}}$.

Lemma 6.5. *Let $(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2)$ be a frozen model pair configuration on G^{∂}.*
(a) Any tree component T of $V^{\partial}_{\neq}$ with $|T \cap \mathcal{Y}^{\star}| \leqslant 1$ must be in $\mathfrak{T}$.

(b) With $|F(V_{\neq})| = \gamma/2 \equiv \rho + A$ as before, we have

$$A'' \equiv A - \tfrac{1}{2}|\mathcal{Y}_{\neq} \backslash \mathfrak{T}| \geqslant 0. \tag{72}$$

Further, let $n^\bullet_{\texttt{0f}}$ denote the cardinality of the set $V^\bullet_{\texttt{0f}}$ of vertices in $V_{\texttt{0f}}$ whose matched partner does not have spin $\texttt{1f}$; symmetrically define $n^\bullet_{\texttt{f0}} = |V^\bullet_{\texttt{f0}}|$. Then

$$2A'' \geqslant [|F(V_{\texttt{01}} \cup V_{\texttt{10}})| - 2r_{\texttt{01}} - y_{\texttt{01}}] + [|F(V_{\texttt{01}} \cup V_{\texttt{10}})| - 2r_{\texttt{10}} - y_{\texttt{10}}] + n^\bullet_{\texttt{0f}} + n^\bullet_{\texttt{f0}} \tag{73}$$

with both $|F(V_{\texttt{01}} \cup V_{\texttt{10}})| - 2r_{\texttt{01}} - y_{\texttt{01}}$ and $|F(V_{\texttt{01}} \cup V_{\texttt{10}})| - 2r_{\texttt{10}} - y_{\texttt{10}}$ non-negative.

Proof. Recall Defn. 4.11 that $\overline{\eta} \equiv (m, \eta)$ is a frozen model configuration on G^∂ if and only if every vertex in $V = V^\partial \backslash \mathcal{Y}^\star$ satisfies properties (i)-(iii) of Defn. 2.4; the properties need not be satisfied on $\mathcal{Y}^\star$.

(a) Let $\mathfrak{C}$ be a component of $V^\partial_{\neq}$, and write $\chi \equiv \chi_{\mathfrak{C}} \equiv |F(\mathfrak{C})| - |\mathfrak{C}|$. Assume $\mathfrak{C} \backslash \mathcal{Y}^\star \neq \varnothing$ since the result is otherwise trivial. Since $V^\partial_{\neq}$ can have no neighbors in $V^\partial_{\texttt{11}}$, every $v \in \mathfrak{C} \cap V$ must satisfy properties (i)-(iii) *within $\mathfrak{C}$ alone* (in both coordinates), i.e. in the absence of the rest of the graph. Write $d_v(\mathfrak{C}) \equiv |F(v, \mathfrak{C})|$, and note that if v_0 is a leaf vertex of $\mathfrak{C}$ (meaning $d_v(\mathfrak{C}) = 1$) then $v_0 \in V_{\texttt{1f}} \cup V_{\texttt{f1}} \cup \mathcal{Y}^\star$.

Suppose then that $v_0 \in V_{\texttt{1f}}$ is a leaf vertex of $\mathfrak{C}$; it must be paired under $\underline{m}^2$ to a partner $v_1 \in \mathfrak{C}$ with spin $\texttt{0f}$. If $v_1 \notin \mathcal{Y}^\star$, then (by property (ii)) it must have at least one other neighbor $v_2 \in \mathfrak{C}$ with spin $\texttt{1f}$ or $\texttt{10}$. Since v_0 is a leaf and v_1 is a $\texttt{0f}$-vertex matched to v_0, the presence of these two vertices in $\mathfrak{C}$ is not needed to ensure that the other vertices satisfy the properties (i)-(iii). Therefore, in the case that v_1 has *exactly two* neighbors in $\mathfrak{C}$ with each joined to v_1 by a single edge, we may choose to remove both v_0, v_1 from $\mathfrak{C}$ causing no change in χ. We iterate this procedure on leaf vertices in $V_{\texttt{1f}}$ or $V_{\texttt{f1}}$ to produce a reduced graph $\mathfrak{R} \equiv \mathfrak{R}(\mathfrak{C}) \subseteq \mathfrak{C}$ which is one of the following:

A. $\mathfrak{R} = \{v_0, v_1\}$ where $v_0 \in \mathcal{Y}$ has spin $\texttt{1f}$ or $\texttt{f1}$, and $v_1 \in \mathcal{Y}^\star$ is its matched partner.
B. $\mathfrak{R} = \{v_0\}$ where $v_0 \in \mathcal{Y}^\star$ has spin $\texttt{10}$, $\texttt{1f}$, $\texttt{01}$, or $\texttt{f1}$.
C. $\mathfrak{R} \backslash \mathcal{Y}^\star \neq \varnothing$; and any $v \in V$ which is a leaf vertex of $\mathfrak{R}$ is either a $\texttt{1f}$-vertex matched to a $\texttt{0f}$-vertex u with $d_u(\mathfrak{R}) \geqslant 3$, or symmetrically an $\texttt{f1}$-vertex matched to an $\texttt{f0}$-vertex w with $d_w(\mathfrak{R}) \geqslant 3$.

In cases A and B, $\chi_{\mathfrak{C}} = \chi_{\mathfrak{R}} = -1$ so $\mathfrak{C}$ is a tree consisting entirely of $\texttt{1f}$-$\texttt{0f}$ and $\texttt{f1}$-$\texttt{f0}$ matched pairs except in case B for the vertex on $\mathcal{Y}^\star$. We claim further that all the matched pairs must be of the same type: indeed, since the only edges between $V^\partial_{\texttt{1f}} \cup V^\partial_{\texttt{0f}}$ and $V^\partial_{\texttt{f1}} \cup V^\partial_{\texttt{f0}}$ must occur between $V^\partial_{\texttt{0f}}$ and $V^\partial_{\texttt{f0}}$, every $v \in \mathfrak{C} \backslash V_{\texttt{f}\cdot}$ satisfies properties (i)-(iii) in the absence of $V_{\texttt{f}\cdot\cdot}$. Therefore apply the leaf-removal procedure separately to each component $\mathfrak{c}$ of $\mathfrak{C} \backslash V_{\texttt{f}\cdot}$ or of $\mathfrak{C} \backslash V_{\cdot\texttt{f}}$: if $\mathfrak{R}(\mathfrak{c}) \neq \varnothing$ then it contains some $v \in \mathcal{Y}^\star$ which must also appear in $\mathfrak{R} = \mathfrak{R}(\mathfrak{C})$. All the $\mathfrak{c}$ are disjoint, and $|\mathfrak{R} \cap \mathcal{Y}^\star| = 1$ by assumption, so only a single component $\mathfrak{c}$ can be non-empty. This proves the claim, and we conclude $\mathfrak{C} \in \mathfrak{T}_{\texttt{1f}} \cup \mathfrak{T}_{\texttt{f1}}$ in case A, $\mathfrak{C} \in \mathfrak{T}_{\texttt{0f}} \cup \mathfrak{T}_{\texttt{f0}}$ in case B (since we assumed $\mathfrak{C} \backslash \mathcal{Y}^\star \neq \varnothing$).

In case C, write $m_v \equiv 2 - \mathbf{1}\{v \in \mathcal{Y}^\star\}$, and consider

$$\sum_{v \in \mathfrak{R}} (d_v(\mathfrak{R}) - m_v) = 2\chi_{\mathfrak{C}} + |\mathfrak{R} \cap \mathcal{Y}^\star| \tag{74}$$

Any $v \in \mathfrak{R}$ with $d_v(\mathfrak{R}) - m_v < 0$ must be in $V_{\texttt{1f}} \cup V_{\texttt{f1}}$, and is compensated in the above sum by a matched partner $u \in \mathfrak{R}$ with $d_u(\mathfrak{R}) \geqslant 3$. Therefore the sum restricted only to vertices belonging in $\texttt{f1}$-$\texttt{f0}$ or $\texttt{1f}$-$\texttt{0f}$ matched pairs inside $\mathfrak{R}$ is non-negative. Any other vertex gives a

non-negative contribution to the sum. Thus $\chi_{\mathfrak{C}} \geqslant -\frac{1}{2}|\mathfrak{R} \cap \mathcal{Y}^\star| \geqslant -\frac{1}{2}|\mathfrak{C} \cap \mathcal{Y}^\star|$, so in particular if $|\mathfrak{C} \cap \mathcal{Y}^\star| = 1$ then (recalling we have assumed $\mathfrak{C}\backslash\mathcal{Y}^\star \neq \varnothing$) $\mathfrak{C}$ cannot be a tree.

(b) Let $\mathbf{R}$ denote the union of the $\mathfrak{R}(\mathfrak{C})$ over the components $\mathfrak{C}$ of $V_{\neq}^\partial\backslash\mathfrak{T}$ with $\mathfrak{C}\backslash\mathcal{Y}^\star \neq \varnothing$: it follows from the proof of (a) that

$$0 \leqslant \sum_{v \in \mathbf{R}}(d_v(\mathbf{R}) - m_v) = 2\chi_{\mathbf{R}\backslash\mathcal{Y}^\star} + |\mathbf{R} \cap \mathcal{Y}^\star| \leqslant 2A - |\mathcal{Y}_{\neq}\backslash\mathfrak{T}| = 2A'' \leqslant 2A.$$

We furthermore decompose

$$2A'' \geqslant \underbrace{\sum_{v \in V_{01}}(d_v(V_{10}) - m_v)}_{=F(V_{01},V_{10})-2r_{01}-y_{01}\geqslant 0} + \underbrace{\sum_{v \in V_{10}}(d_v(V_{01}) - m_v)}_{=F(V_{01},V_{10})-2r_{10}-y_{10}\geqslant 0} + \underbrace{\sum_{v \in V_{0f}^\bullet}(d_v(\mathbf{R}) - m_v)}_{\geqslant |V_{0f}^\bullet|=n_{0f}^\bullet} + \underbrace{\sum_{v \in V_{f0}^\bullet}(d_v(\mathbf{R}) - m_v)}_{\geqslant |V_{f0}^\bullet|=n_{f0}^\bullet},$$

concluding the proof. □

Suppose $(\overline{\eta}^1, \overline{\eta}^2)$ is a pair frozen model configuration on G^∂ which is consistent with auxiliary configuration $\underline{\tau}_{\mathcal{Y}} \equiv (\sigma_{\mathcal{Y}}^1, \sigma_{\mathcal{Y}}^2)$ on the half-edges $\mathcal{Y}$ joining $\mathcal{Y}$ to the clauses $F^\partial = F(\mathcal{Y}, \mathcal{Y}^\star)$. Write $\underline{\tau}_{\mathcal{Y}} \equiv (\underline{o}, \underline{i}) \equiv (\underline{o}^i, \underline{i}^i)_{i=1}^2$ where $\underline{o}^i \in \{0, 1, \mathtt{f}\}^{\mathcal{Y}}$ are the outgoing messages of $\sigma_{\mathcal{Y}}^i$, and $\underline{i}^i \in \{0, 1, \mathtt{f}\}^{\mathcal{Y}}$ the incoming. If $y \in \mathcal{Y}$ is at the root of a tree $T \in \mathfrak{T}$ then $\underline{o}_y \equiv o_y^1 o_y^2$ is either $\mathtt{11}$ ($T \in \mathfrak{T}_{\mathtt{1f}} \cup \mathfrak{T}_{\mathtt{f1}}$) or $\mathtt{ff}$ ($T \in \mathfrak{T}_{\mathtt{0f}} \cup \mathfrak{T}_{\mathtt{f0}}$). Further, $v \in \mathcal{Y}$ is incident to $y \in \mathcal{Y}$ with outgoing messages $\underline{o}_y^1 \neq \underline{o}_y^2$ only if either $v \in \mathcal{Y}_{\neq}\backslash\mathfrak{T}$, or $v \in \mathcal{Y}_{00}$ but is neighboring to $\mathcal{R}_{\neq}$. Thus the Hamming distance $\underline{d}(\underline{o})$ between $\underline{o}^1$ and $\underline{o}^2$ satisfies

$$\underline{d}(\underline{o}) \leqslant |\mathcal{Y}_{\neq}\backslash\mathfrak{T}| + k' \quad \text{where } k' \equiv |\{v \in \mathcal{Y}_{00} : F(V_{\neq}, v) \neq \varnothing\}|. \tag{75}$$

Lemma 6.6. *Suppose $(\overline{\eta}^1, \overline{\eta}^2)$ is a pair frozen model configuration on G^∂ contributing to $\ddot{_2}\mathbf{Z}_{n\alpha-(r^1,r^2)}^\partial[\underline{\tau}_{\mathcal{Y}}]$. For $i = 1, 2$ let $2j^i \equiv |\{w \in \mathcal{W} : \sigma_w^i = \mathtt{11}\}|$, $l^i \equiv |\{u \in \mathcal{U} : \sigma_u^i = \mathtt{11}\}|$, and $s^i \equiv r^i - j^i$. Then*

$$|\mathbf{J} - \mathbf{I}| \leqslant \underline{d}(\underline{o}) + 2A. \quad \text{for } \mathbf{J} \equiv j^1 - j^2, \; \mathbf{I} \equiv s^1 - l^1 - (s^2 - l^2).$$

Proof. For $\omega \in \{\mathtt{f1}, \mathtt{f0}, \mathtt{1f}, \mathtt{0f}, \mathtt{ff}\}$ let $\underline{e}_\omega$ count the number of vertices in $\mathcal{Y}$ with spin ω and matched partner in $\mathcal{Y}^\star$. For $i = 1, 2$ let $\underline{e}_{\neq}^i$ count the number of $\mathtt{f!f}$-vertices in $\mathcal{Y}$ whose matched partner in the i-th coordinate is in $\mathcal{Y}^\star$. Note $\underline{e}_{\mathtt{f1}}$ also counts the set of $y \in \mathcal{Y}$ with $o_y^1 = 1 = o_y^2$ but $i_y^1 = 1 \neq i_y^2$, so

$$|(2j^1 + l^1) - (2j^2 + l^2) - (\underline{e}_{\mathtt{f1}} - \underline{e}_{\mathtt{1f}})| \leqslant \underline{d}(\underline{o}). \tag{76}$$

Meanwhile $n\alpha - r^1 = n_{1\cdot} + \frac{1}{2}(n_{\mathtt{f}\cdot} - \underline{e}_{\mathtt{f0}} - \underline{e}_{\mathtt{f1}} - \underline{e}_{\mathtt{ff}} - \underline{e}_{\neq}^1)$ and similarly for $n\alpha - r^2$. Recall $n_{0\mathtt{f}}^\bullet$ denotes the number of vertices in $V_{0\mathtt{f}}$ with matched partner in $V_{\mathtt{Zf}} \cup \mathcal{Y}_{\mathtt{Zf}}^\star$, and let $n_{0\mathtt{f}}^\circ \leqslant n_{0\mathtt{f}}^\bullet$ denote the number of vertices in $V_{0\mathtt{f}}$ with matched partner in $V_{\mathtt{Zf}}$: then $n_{0\mathtt{f}} - n_{1\mathtt{f}} = \underline{e}_{0\mathtt{f}} - \underline{e}_{1\mathtt{f}} + n_{0\mathtt{f}}^\circ$, so taking the difference between the $n\alpha - r^i$ gives

$$r^1 - r^2 = n_{01} - n_{10} + \tfrac{1}{2}(\underline{e}_{\neq}^1 - \underline{e}_{\neq}^2 + n_{0\mathtt{f}}^\circ - n_{\mathtt{f}0}^\circ) + \underline{e}_{\mathtt{f1}} - \underline{e}_{\mathtt{1f}}.$$

Applying (73) together with the trivial fact that $a \vee b \geqslant |a - b|$ for $a, b \geqslant 0$ gives

$$|n_{10} - n_{01}| \leqslant |r_{10} + \tfrac{1}{2}y_{10} - r_{01} - \tfrac{1}{2}y_{01}| + \tfrac{1}{2}|y_{10} - y_{01}| \leqslant A'' + \tfrac{1}{2}|\mathcal{Y}_{\neq}\backslash\mathfrak{T}| = A, \tag{77}$$

and similarly $\frac{1}{2}|\underline{e}_{\neq}^1 - \underline{e}_{\neq}^2 + n_{0\mathtt{f}}^\circ - n_{\mathtt{f}0}^\circ| \leqslant \frac{1}{2}[|\mathcal{Y}_{\neq}\backslash\mathfrak{T}| + n_{0\mathtt{f}}^\bullet \vee n_{\mathtt{f}0}^\bullet] \leqslant A$, so

$$|r^1 - r^2 - (\underline{e}_{\mathtt{f1}} - \underline{e}_{\mathtt{1f}})| \leqslant 2A. \tag{78}$$

Comparing (78) with (76) gives $|\underline{e}_{\mathtt{f1}} - \underline{e}_{\mathtt{1f}} - [2s^1 - l^1 - (2s^2 - l^2)]| \leqslant \underline{d}(\underline{o}) + 4A$, and combining with (76) proves the lemma. □

The following is our refinement of Propn. 4.12 in the regime of small Δ:

Proposition 6.7. *Suppose* $\Delta \leqslant n^{1/5}$, *and let* $\underline{\tau}_{\mathscr{Y}} = (\underline{o}, \underline{i})$ *with* $\boldsymbol{J}$, $\boldsymbol{I}$ *as in Lem.* 6.6. *Then*

$$\mathbb{E}[{}_2^{\cdot}\boldsymbol{Z}^{\partial}_{n\alpha - (r^1, r^2)}[\underline{\tau}_{\mathscr{Y}}]] \leqslant n^{-\Xi} \, \mathbb{E}[\boldsymbol{Z}^{\partial}_{n\alpha - r^1}[\underline{\sigma}^1_{\mathscr{Y}}]]$$

for $\Xi \equiv \Xi(\underline{o}, \boldsymbol{J}, \boldsymbol{I}) \equiv \frac{1}{20}\boldsymbol{d}(\underline{o}) + 0 \vee \frac{1}{20}(|\boldsymbol{J}| - |\boldsymbol{I}|)$.

Proof. Recall (59) that ${}_2\boldsymbol{O}^{\partial}_A[\pi|\underline{\tau}_{\mathscr{Y}}]$ denotes the set of pair frozen model configurations $(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2)$ on G^{∂} which are consistent with boundary conditions $\underline{\tau}_{\mathscr{Y}}$, have empirical measure $\pi_{\omega} = \boldsymbol{r}_{\omega}/\boldsymbol{r}$ when restricted to $\mathcal{R}$, and have $\boldsymbol{\rho} + A$ internal edges among the unequal spins. Analogously we now let ${}_2\boldsymbol{O}^{\partial}_{A,k'}[\pi|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}]$ denote the set of pair frozen model configurations $(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2)$ on G^{∂} which are consistent with boundary conditions $\underline{\tau}_{\mathscr{Y}}$ *and have the given* $\mathfrak{T}$ (see Defn. 6.4), have empirical measure $\pi_{\omega} = \boldsymbol{r}_{\omega}/\boldsymbol{r}$ when restricted to $\mathcal{R}$, have $\boldsymbol{\rho} + A$ internal edges among the unequal spins, and lastly have k' as in (75). Let $\mathsf{proj}^1_{\mathfrak{T}}$ denote the projection mapping $(\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\eta}}^2) \mapsto (\overline{\boldsymbol{\eta}}^1, \overline{\boldsymbol{\mu}}^2)$ where $\overline{\boldsymbol{\mu}}^2$ equals $\overline{\boldsymbol{\eta}}^2$ on $\mathfrak{T}$ but equals $\overline{\boldsymbol{\eta}}^1$ on $V \backslash \mathfrak{T}$. We shall compare

$${}_2\boldsymbol{Z}^{\partial}_{A,k'}[\pi|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}] \equiv \left|{}_2\boldsymbol{O}^{\partial}_{A,k'}[\pi|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}]\right|, \quad {}_2\boldsymbol{Y}^{\partial}[\pi|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}] \equiv \left|\bigcup_{A,k'} \mathsf{proj}^1_{\mathfrak{T}}\big({}_2\boldsymbol{O}^{\partial}_{A,k'}[\pi|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}]\big)\right|.$$

The ratio of their expected values can be upper bounded by repeating the derivation of (61), but now assuming $\Delta \leqslant n^{1/5}$ and fixing the spins on $\mathfrak{T}$ as well as on $\mathcal{Y}^{\star} \cup \mathcal{Y}$. Further accounting for the assignment of $A' \geqslant k'$ edges from $V_{\neq}$ to $\mathcal{Y}_{00}$ yields

$$\mathscr{R}^1_{A,k'}[\pi|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}] \equiv \frac{\mathbb{E}[{}_2\boldsymbol{Z}^{\partial}_{A,k'}[\pi|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}]]}{\mathbb{E}[{}_2\boldsymbol{Y}^{\partial}[\pi|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}]]} \leqslant \frac{1}{n^{k'/2}} \frac{(\log d)^{O(\Delta'')} d^{(\widetilde{\varepsilon}_{\mathsf{f}\cdot})''/4}}{d^{(\Delta_0/2 + \rho_{\mathsf{f}}/4)''}} \left(\frac{C(\Delta'')^2 d}{n\gamma''}\right)^A,$$

where we use $''$ to denote counts off of $\mathfrak{T}$. Recall (77) that $|n_{01} - n_{10}| \leqslant A$, and note also $n_{01} = n''_{01}$ and $n''_{0\mathsf{f}} \geqslant \boldsymbol{r}''_{1\mathsf{f}}$, so

$$\tfrac{1}{2}\Delta''_0 + \tfrac{1}{4}\boldsymbol{\rho}''_{\mathsf{f}} = \tfrac{1}{2}(n_{01} + n''_{0\mathsf{f}}) + \tfrac{1}{4}\boldsymbol{\rho}''_{\mathsf{f}} \geqslant \tfrac{1}{4}[n_{01} + n_{10} + n''_{0\mathsf{f}} + \boldsymbol{r}''_{1\mathsf{f}} + \boldsymbol{\rho}_{\mathsf{f}}] - \tfrac{1}{4}A \geqslant \tfrac{1}{4}(\boldsymbol{\rho}'' - A).$$

Meanwhile $(\widetilde{\varepsilon}_{\mathsf{f}\cdot})'' \leqslant \varepsilon'' = |\mathcal{Y}_{\neq} \backslash \mathfrak{T}| \leqslant 2A$ by (72), so

$$\mathscr{R}^1_{A,k'}[\pi''|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}] \leqslant \frac{(\log d)^{O(\Delta'')} d^{(\widetilde{\varepsilon}_{\mathsf{f}\cdot})''/4}}{d^{\rho''/4} n^{k'/2}} (d^{1/4}/n^{4/7})^A \leqslant \frac{(d/n^{4/7})^A}{d^{\rho''/5} n^{k'/2}} \leqslant \frac{1}{d^{\rho''/5} n^{(A+k')/2}}.$$

After applying $\mathsf{proj}^1_{\mathfrak{T}}$ on the spins of $V \backslash \mathfrak{T}$, the spins on $\mathfrak{T}$ are uniquely determined by the incoming messages $\underline{i}$, therefore $\sum_{\mathfrak{T}} {}_2\boldsymbol{Y}^{\partial}[\pi|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}] \leqslant \boldsymbol{Z}^{\partial}[\pi^1|\underline{\sigma}^1_{\mathscr{Y}}]$. From (75) and (72), $\boldsymbol{d}(\underline{o}) \leqslant 2A + k'$. Combining Lem. 6.6 gives $|\boldsymbol{J}| + \boldsymbol{d}(\underline{o}) \leqslant |\boldsymbol{I}| + 6(A + k')$, therefore

$$\tfrac{1}{2}(A + k') \geqslant [\tfrac{1}{2}\boldsymbol{d}(\underline{o})] \vee [\tfrac{1}{12}(\boldsymbol{d}(\underline{o}) + |\boldsymbol{J}| - |\boldsymbol{I}|)] \geqslant \tfrac{1}{12}\boldsymbol{d}(\underline{o}) + \tfrac{1}{12}(|\boldsymbol{J}| - |\boldsymbol{I}|) \vee 0.$$

The result follows by decomposing $\mathbb{E}[{}_2^{\cdot}\boldsymbol{Z}^{\partial}_{n\alpha - (r^1, r^2)}[\underline{\tau}_{\mathscr{Y}}]]$ as a sum over $\mathbb{E}[{}_2\boldsymbol{Z}^{\partial}_{A,k'}[\pi|\underline{\tau}_{\mathscr{Y}}, \mathfrak{T}]]$. $\square$

6.3. **Exponential cost of second matching.** For $i = 1, 2$ let A^i_W be a random matching on $\mathscr{W}$, and write $A^i_W \in \mathfrak{M}(\underline{o}^i_{\mathscr{W}}, j^i)$ if A^i_W forms exactly j^i pairs (w, w') with $\boldsymbol{o}^i_w = \boldsymbol{o}^i_{w'} = 1$. Note that $\underline{o}_{\mathscr{W}}$ and A_W together specify a unique configuration $\underline{i}(\underline{o}_{\mathscr{W}}, A_W)$ of incoming messages. We therefore rewrite (69) as

$$\boldsymbol{v}^T_{n\alpha, i} = \sum_{\substack{r^1, j^1, \underline{o}^1, \underline{i}^1_{\mathscr{W}}, \\ A^1_W \in \mathfrak{M}(\underline{o}^1_{\mathscr{W}}, j^1)}} \mathbb{P}(A^1_W)\widetilde{\kappa}_{r^1 - j^1}(\underline{\sigma}^1_{\mathscr{U}}) \sum_{\substack{r^2, j^2, \underline{o}^2, \underline{i}^2_{\mathscr{W}}, \\ A^2_W \in \mathfrak{M}(\underline{o}^2_{\mathscr{W}}, j^2)}} \mathbb{P}(A^2_W)\widetilde{\kappa}_{r^2 - j^2}(\underline{\sigma}^2_{\mathscr{W}}){}_2^{\cdot}\boldsymbol{Z}^{\partial}_{n\alpha - (r^1, r^2)}[\underline{\tau}_{\mathscr{Y}}].$$

where $\underline{\tau}_{\mathscr{Y}} \equiv (\underline{o}, \underline{i})$ with $\underline{i}$ uniquely determined by $\underline{i}_{\mathscr{U}}$ and the matchings A^i_W.

Proposition 6.8. *It holds uniformly over all realizations T of* (66) *that*

$$\mathbb{E}_T[\boldsymbol{v}^T_{n\alpha,i}] \lesssim_{d,t} e^{-c_d\boldsymbol{w}}\mathbb{E}\boldsymbol{Z}_{n\alpha}.$$

Proof. By applying Propn. 6.7 together with the fact that the total number of possibilities of T, $r^i - j^i$, $\underline{\sigma}^1_{\mathscr{U}}$, $\widetilde{\kappa}$ is $\lesssim_{d,t} 1$, we find

$$\mathbb{E}_T[\boldsymbol{v}^T_{n\alpha,i}] \lesssim_{d,t} \sum_{\substack{j^1, \underline{\boldsymbol{o}}^1_{\mathscr{W}} \\ A^1_W \in \mathfrak{M}(\underline{\boldsymbol{o}}^1_{\mathscr{W}}, j^1)}} \mathbb{P}(A^1_W)\mathbb{E}\boldsymbol{Z}^{\mathscr{W}}_{n\alpha - j^1}[\underline{\sigma}^1_{\mathscr{W}}] \overbrace{\sum_{j^2, \underline{\boldsymbol{o}}^2_{\mathscr{W}}} \frac{\overline{\psi}_{j^2}(\underline{\boldsymbol{o}}^2_{\mathscr{W}})}{n^{d(\underline{\boldsymbol{o}}_{\mathscr{W}})/20 + 0 \vee [|\boldsymbol{J}| - c_{d,t}]/20}}}^{\boldsymbol{S} \equiv \boldsymbol{S}(j^1, \underline{\boldsymbol{o}}^1_{\mathscr{W}})} \tag{79}$$

where $\boldsymbol{Z}^{\mathscr{W}}$ refers to the partition function of the graph $G\backslash A_W$ with $\boldsymbol{w} \equiv |\mathscr{W}|$ unmatched edges, $c_{d,t}$ denotes the maximum possible value of $|\boldsymbol{I}|$, and $\overline{\psi}_{j^2}(\underline{\boldsymbol{o}}^2_{\mathscr{W}}) \equiv \mathbb{P}(A^2_W \in \mathfrak{M}(\underline{\boldsymbol{o}}^2_{\mathscr{W}}, j^2))$. If $\underline{\boldsymbol{o}}^2_{\mathscr{W}}$ equals 1 in ℓ coordinates, then

$$\overline{\psi}_j(\underline{\boldsymbol{o}}^2_{\mathscr{W}}) = \binom{\ell}{2j}\frac{(2j-1)!!(\boldsymbol{w}-\ell)_{\ell-2j}(\boldsymbol{w}-2\ell+2j-1)!!}{(\boldsymbol{w}-1)!!} = 2^{\ell-2j}\binom{\boldsymbol{w}/2}{j, \ell-2j, \boldsymbol{w}/2-\ell+j}\bigg/\binom{\boldsymbol{w}}{\ell}.$$

This is positive for $0 \le \ell \le \boldsymbol{w}$ and $0 \vee (2\ell - \boldsymbol{w}) \le 2j \le \ell$; and if $\psi^{\boldsymbol{w}}_{\ell,j} > 0$ then crudely $\psi^{\boldsymbol{w}}_{\ell',j'}/\psi^{\boldsymbol{w}}_{\ell,j} \le \boldsymbol{w}^{10[|j'-j|+|\ell'-\ell|]}$ for any other j', ℓ'. Since $\boldsymbol{w} \lesssim_d \log n$, we see that the main effect in sum defining $\boldsymbol{S} \equiv \boldsymbol{S}(j^1, \underline{\boldsymbol{o}}^1_{\mathscr{W}})$ is the polynomial decay in $d(\underline{\boldsymbol{o}}_{\mathscr{W}})$ and in $|\boldsymbol{J}| \ge c_{d,t}$:

$$\boldsymbol{S}(j^1, \underline{\boldsymbol{o}}^1_{\mathscr{W}}) \lesssim \sum_{|j^2 - j^1| \le c_{d,t}} \overline{\psi}_{j^2}(\underline{\boldsymbol{o}}^1_{\mathscr{W}}).$$

Since $A^1_W \in \mathfrak{M}(\underline{\boldsymbol{o}}^1_{\mathscr{W}}, j^1)$, $\underline{\sigma}^1_{\mathscr{W}}$ must equal 11 in exactly $2j^1$ coordinates, and then (71) gives

$$\mathbb{E}\boldsymbol{Z}^{\mathscr{W}}_{n\alpha - j^1}[\underline{\sigma}^1_{\mathscr{W}}] \sim \widetilde{\boldsymbol{c}}^{-1}\boldsymbol{q}(\underline{\boldsymbol{o}}^1_{\mathscr{W}})\lambda^{-j^1}\mathbb{E}\boldsymbol{Z}_{n\alpha} \tag{80}$$

for $\widetilde{\boldsymbol{c}}$ a proportionality constant depending on n, $|\mathscr{W}|$, and λ. Therefore

$$\frac{\mathbb{E}_T[\boldsymbol{v}^T_{n\alpha,i}]}{\mathbb{E}\boldsymbol{Z}_{n\alpha}} \lesssim_{d,t} \widetilde{\boldsymbol{c}}^{-1}\sum_{\underline{\boldsymbol{o}}^1_{\mathscr{W}}}\boldsymbol{q}(\underline{\boldsymbol{o}}^1_{\mathscr{W}})\sum_{j^1}\frac{\overline{\psi}_{j^1}(\underline{\boldsymbol{o}}^1_{\mathscr{W}})}{\lambda^{j^1}}\sum_{|j^1-j^2|\le c_{d,t}}\overline{\psi}_{j^2}(\underline{\boldsymbol{o}}^1_{\mathscr{W}}) \le \widetilde{\boldsymbol{c}}^{-1}\sum_{\underline{\boldsymbol{o}}^1_{\mathscr{W}}}\boldsymbol{q}(\underline{\boldsymbol{o}}^1_{\mathscr{W}})\sum_{j^1}\frac{\overline{\psi}_{j^1}(\underline{\boldsymbol{o}}^1_{\mathscr{W}})}{\lambda^{j^1}} \sim 1,$$

where the last step is by a second application of (80). Stirling's formula gives

$$\psi^{\boldsymbol{w}}_{\ell,j} = \boldsymbol{w}^{O(1)}\exp\{\boldsymbol{w}\,f(\ell/\boldsymbol{w}, j/\boldsymbol{w})\}, \quad 1 \sim \widetilde{\boldsymbol{c}}^{-1}\boldsymbol{w}^{O(1)}\sum_{(\ell/\boldsymbol{w}, j/\boldsymbol{w})\in\mathscr{C}}\exp\{\boldsymbol{w}\,g(\ell/\boldsymbol{w}, j/\boldsymbol{w})\}$$

where $\mathscr{C}$ denotes the convex set $\{0 \le x_\ell \le 1, 0 \vee (2x_\ell - 1) \le 2x_j \le x_\ell\}$, and

$$f(x_\ell, x_j) \equiv x_\ell \log 2 - 2x_j \log 2 + \tfrac{1}{2}H(2x_j, 2x_\ell - 4x_j, 1 - 2x_\ell + 2x_j) - H(x_\ell),$$
$$g(x_\ell, x_j) \equiv -H(x_\ell \,|\, q_1) - x_j \log \lambda + f(x_\ell, x_j).$$

Because of the entropy terms, the compact set $\arg\max_x g(x)$ must lie some positive distance ε'_d away from the boundary of $\mathscr{C}$. By computing the Hessian, we find g is strictly concave in the interior of $\mathscr{C}$, so it has a unique maximizer $(x^\lambda_\ell, x^\lambda_j)$. Now compare

$$x^\lambda_j(x_\ell) = \arg\max_{x_j}[-x_j \log \lambda + f(x_\ell, x_j)] \quad \text{against} \quad x^1_j(x_\ell) = \arg\max_{x_j} f(x_\ell, x_j):$$

clearly they must be separated by some constant distance $5\varepsilon_d$ uniformly over all x_ℓ near x^λ_ℓ. Then, by concavity, there exists a constant $c_d > 0$ such that $e^{\boldsymbol{w}\,f(x_\ell, x_j)} \le e^{-2c_d\boldsymbol{w}}$ for all

$\|(x_\ell, x_j) - (x_\ell^\lambda, x_j^\lambda)\|_1 \leqslant \varepsilon_d$, while $\widetilde{\boldsymbol{c}}^{-1} e^{\boldsymbol{w} \, g(x_\ell, x_j)} \leqslant e^{-2c_d \boldsymbol{w}}$ for all $\|(x_\ell, x_j) - (x_\ell^\lambda, x_j^\lambda)\|_1 \geqslant \varepsilon_d$, so

$$\frac{\mathbb{E}_T[\boldsymbol{v}_{n\alpha,i}^T]}{\mathbb{E}\boldsymbol{Z}_{n\alpha}} \lesssim_{d,t} \boldsymbol{w}^{O(1)} \max_{(x_\ell, x_j) \in \mathscr{C}} \widetilde{\boldsymbol{c}}^{-1} \exp\{\boldsymbol{w} \, g(x_\ell, x_j) + \boldsymbol{w} \, f(x_\ell, x_j)\} \leqslant e^{-c_d \boldsymbol{w}}$$

as claimed. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $\square$

Proof of Propn. 6.1. Follows from Cor. 6.3 and Propn. 6.8 with $m' = m - (2/c_d) \log n$. $\quad$ $\square$

REFERENCES

[1] D. Aldous. Open problems. Posted online, 2003.
[2] D. Aldous and R. Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12(54):1454–1508 (electronic), 2007.
[3] D. Aldous and J. M. Steele. The objective method: probabilistic combinatorial optimization and local weak convergence. In *Probability on discrete structures*, volume 110 of *Encyclopaedia Math. Sci.*, pages 1–72. Springer, Berlin, 2004.
[4] D. J. Aldous. The $\zeta(2)$ limit in the random assignment problem. *Random Struct. Algor.*, 18(4):381–418, 2001.
[5] J. Barbier, F. Krzakała, L. Zdeborová, and P. Zhang. The hard-core model on random graphs revisited. Preprint at http://arxiv.org/abs/1306.4121, 2013.
[6] M. Bayati, D. Gamarnik, and P. Tetali. Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. In *Proc. ACM Symp. (STOC)*, pages 105–114. ACM, New York, 2010.
[7] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6(23):1–13 (electronic), 2001.
[8] B. Bollobás. The independence ratio of regular graphs. *Proc. Amer. Math. Soc.*, 83(2):433–436, 1981.
[9] A. Coja-Oghlan and C. Efthymiou. On independent sets in random graphs. In *Proc. ACM-SIAM Symp. (SODA)*, pages 136–144. SIAM, 2011.
[10] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the Lambert W function. *Adv. Comput. Math.*, 5(4):329–359, 1996.
[11] R. M. Corless, D. J. Jeffrey, and D. E. Knuth. A sequence of series for the Lambert W function. In *Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation (Kihei, HI)*, pages 197–204 (electronic), New York, 1997. ACM.
[12] A. Dembo, A. Montanari, A. Sly, and N. Sun. The replica symmetric solution for Potts models on d-regular graphs. Preprint at http://arxiv.org/abs/1207.5500, 2012.
[13] A. Dembo, A. Montanari, and N. Sun. Factor models on locally tree-like graphs. *Ann. Probab (to appear)*, 2011.
[14] A. Dembo and O. Zeitouni. *Large deviations techniques and applications*, volume 38 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin, 2010. Corrected reprint of the second (1998) edition.
[15] J. Ding, A. Sly, and N. Sun. Satisfiability threshold for random regular NAE-SAT. Preprint, 2013.
[16] A. Frieze and S. Suen. Analysis of two simple heuristics on a random instance of k-SAT. *J. Algorithms*, 20(2):312–355, 1996.
[17] A. M. Frieze and T. Łuczak. On the independence and chromatic numbers of random regular graphs. *J. Combin. Theory Ser. B*, 54(1):123–132, 1992.
[18] G. R. Grimmett and C. J. H. McDiarmid. On colouring random graphs. *Math. Proc. Cambridge Philos. Soc.*, 77:313–324, 1975.
[19] F. Krzakała, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, 104(25):10318–10323, 2007.
[20] B. D. McKay. Independent sets in regular graphs of high girth. In *Proceedings of the Singapore conference on combinatorial mathematics and computing (Singapore, 1986)*, volume 23, pages 179–185, 1987.
[21] M. Mézard and A. Montanari. *Information, physics, and computation.* Oxford Graduate Texts. Oxford University Press, Oxford, 2009.

[22] M. Mézard, G. Parisi, and M. A. Virasoro. *Spin glass theory and beyond*, volume 9 of *World Scientific Lecture Notes in Physics*. World Scientific Publishing Co. Inc., Teaneck, NJ, 1987.

[23] M. Talagrand. The Parisi formula. *Ann. Math. (2)*, 163(1):221–263, 2006.

[24] N. C. Wormald. Differential equations for random processes and random graphs. *Ann. Appl. Probab.*, 5(4):1217–1235, 1995.

[25] N. C. Wormald. Models of random regular graphs. In *Surveys in combinatorics, 1999 (Canterbury)*, volume 267 of *London Math. Soc. Lecture Note Ser.*, pages 239–298. Cambridge Univ. Press, Cambridge, 1999.