Acta Math., 217 (2016), 263–340
DOI: 10.1007/s11511-017-0145-9

Maximum independent sets on random regular graphs

by

Jian Ding

University of Chicago
Chicago, IL, U.S.A.

Allan Sly

University of California, Berkeley
Berkeley, CA, U.S.A.
and
Australian National University
Canberra, Australia

Nike Sun

Stanford University
Stanford, CA, U.S.A.

1. Introduction

An *independent set* in a graph is a subset of vertices not containing any two neighbors. Establishing asymptotics of the maximum independent set size (the *independence number*) on random graphs is a classical problem in probabilistic combinatorics. On the random d-regular graph $\mathscr{G}_{d,n}$, the independence number grows linearly in the number n of vertices. Upper bounds were established by Bollobás [13] and McKay [35], and lower bounds by Frieze–Suen [26], Frieze–Łuczak [25] and Wormald [43]. These were obtained by using a combination of techniques, including first and second moment bounds, differential equations, and the switching method. (For a more complete history and discussion of many related topics see the survey of Wormald [44].) The bounds are close, with the maximal density of occupied vertices (the *independence ratio*) asymptotic to $2(\log d)/d$ in the limit of large d(1)—however, for every fixed d, a gap in the bounds remains.

By a classical martingale argument, the independence ratio is well-concentrated about its mean, with only $O(n^{-1/2})$ fluctuations. Nevertheless, it was a long-standing open problem (see [7], [9]) to determine if there even exists a limiting independence ratio. This was recently established by Bayati–Gamarnik–Tetali [11] using a super-additivity

Research supported by NSF grant DMS-1313596 (J. D.), Sloan Research Fellowship (A. S.), NDSEG and NSF GRF (N. S.).

(1) Throughout this paper, log denotes the natural logarithm.

argument. Though their method proves the existence of thresholds in a quite general class of problems, it yields no information either on the threshold value or the order of fluctuations.

In this paper, we determine for all large d the value of the limiting independence ratio.

THEOREM 1. *The independence ratio* IR_n *of the random* d*-regular graph* $\mathscr{G}_{d,n}$ *converges in probability to an explicit constant* $\alpha_\star \equiv \alpha_\star(d)$, *defined below.*

(We write $a \equiv b$ to indicate that a and b are defined to be equal.) The constant $\alpha_\star$ can be found as follows: solve for the largest root $q = q_\star \leqslant 2(\log d)/d$ of the function

$$\mathbf{f}(q) \equiv -\log\left[1-q\left(1-\frac{1}{\lambda}\right)\right]-\left(\frac{d}{2}-1\right)\log\left[1-q^2\left(1-\frac{1}{\lambda}\right)\right]-\alpha\log\lambda, \tag{1}$$

where λ and α are defined in terms of q by

$$\lambda(q) \equiv q\frac{1-(1-q)^{d-1}}{(1-q)^d} \quad \text{and} \quad \alpha(q) \equiv q\frac{1-q+dq/2\lambda}{1-q^2(1-1/\lambda)}. \tag{2}$$

Then $\alpha_\star = \alpha(q_\star)$. This value was predicted in the statistical physics literature using the heuristic methods of the one-step replica symmetry breaking framework [29], [41]. By confirming this prediction, our theorem solves a long-standing and fundamental problem in probabilistic combinatorics.

Remark 1.1. Let $\mathsf{MIS}_n \equiv n \cdot \mathsf{IR}_n$ denote the non-normalized maximum independent set size in $\mathscr{G}_{d,n}$. Theorem 1 asserts that $\mathsf{MIS}_n = [1+o_n(1)]n\alpha_\star$ with high probability. The analysis of this paper, together with the well-known bound([2])

$$\mathbb{P}(|\mathsf{MIS}_n - \mathbb{E}\mathsf{MIS}_n| \geqslant x) \leqslant \exp\left(-\frac{x^2}{nd}\right),$$

will show that

$$\lim_{C\to\infty}\limsup_{n\to\infty}\mathbb{P}(|\mathsf{MIS}_n - n\alpha_\star| \geqslant Cn^{1/2}) = 0.$$

In fact, by adapting methods from [24], we can obtain the stronger result

$$\lim_{C\to\infty}\limsup_{n\to\infty}\mathbb{P}(|\mathsf{MIS}_n - n\alpha_\star + c_\star\log n| \geqslant C) = 0, \tag{3}$$

for $\alpha_\star$ as before and $c_\star = [2\log\lambda(q_\star)]^{-1}$. That is to say, MIS_n has only constant-order fluctuations around $\boldsymbol{a}_n \equiv n\alpha_\star - c_\star\log n$. The proof of (3) is omitted here, but will be given in full in the online manuscript.

([2]) This can be seen by taking the Doob martingale of MIS_n with respect to the edge-revealing filtration of $\mathscr{G}_{d,n}$ and applying the Azuma–Hoeffding inequality.

1.1. Constraint satisfaction problems

The problem of finding an independent set of given size in a graph is a classic example of a *constraint satisfaction problem* (CSP), and is well known to be NP-hard [31]. The limiting independence ratio $\alpha_\star$ in a sparse random graph ensemble is one example of an extensively studied phenomenon of *satisfiability thresholds* in random CSPs. (The analogous "constraint parameter" α in the *random k-satisfiability* (k-SAT) *problem*, for example, is the ratio of clauses to variables.) For a recent survey see [3].

A major early advance in this area was the realization (see [37], [27]) that many hard CSPs can be recast as *spin glasses*, and analyzed by methods of statistical mechanics. Since these pioneering works, statistical physicists have developed an extensive, but largely non-rigorous, theory around CSPs as models of disordered systems (see [33], [36] and references therein). For a broad class of CSPs, the theory yields a detailed understanding of the phase diagram, including exact predictions of the satisfiability threshold.

Some aspects of this rich picture have been established, including such celebrated results as Aldous's solution to the random assignment problem [6] and Talagrand's proof of the Parisi formula for the Sherrington–Kirkpatrick model [42], both on the complete graph. However, many of the most important ideas remain at the level of conjecture, particularly in sparse ensembles (dilute spin glasses). A principal aim of this work is to advance the mathematical understanding by establishing for the first time an exact satisfiability threshold for a sparse random CSP exhibiting *replica symmetry breaking* (RSB)—describing a broad class of problems which includes independent set, k-SAT (for $k\geqslant 3$), coloring, and MAX-CUT. We believe that the approach developed here gives some rigorous validation to the physics formalism and supplies a framework for establishing other thresholds of interest.

In the remainder of this introductory section, we review the concept of RSB and explain why it obstructs the standard techniques for locating thresholds, then present a brief overview of our solution. As an illustration of the robustness of our methods, in a companion paper [24] we apply the same techniques to determine the exact satisfiability threshold for another CSP in this class, the random regular NAE-SAT problem. Our intention is that these methods may eventually be extended to other combinatorial properties such as the chromatic number or maximum cut, and to the sparse Erdős–Rényi random graphs.

1.2. Moments and non-concentration

The natural approach to studying the independence ratio is the (first and second) moment method applied to the number $Z_{n\alpha}$ of independent sets of fixed density α. Indeed, this approach successfully determines the asymptotics of the independence number for the *dense* Erdős–Rényi random graph [28]. On *sparse* random graphs, the moment method *fails* to locate the sharp transition. The first moment of $Z_{n\alpha}$ over $\mathscr{G}_{d,n}$ is straightforward to calculate [13], and scales exponentially in n:

$$\mathbb{E}Z_{n\alpha} = n^{O(1)} \exp(n\phi(\alpha)),$$

where ϕ is a smooth function of α (we supply the explicit calculation in §1.5 below). The *first moment threshold* $\alpha_\square \equiv \sup\{\alpha \geqslant 0 : \phi(\alpha) \geqslant 0\}$ occurs at [35]

$$\alpha_\square = \frac{2}{d}\left[\log d - \log\log d + \log\left(\frac{e}{2}\right) + O\left(\frac{\log\log d}{\log d}\right)\right]. \tag{4}$$

Since $\mathbb{P}(Z_{n\alpha} > 0) \leqslant \mathbb{E}Z_{n\alpha}$ is exponentially small above $\alpha_\square$, clearly

$$\limsup_{n\to\infty} \mathsf{IR}_n \leqslant \alpha_\square$$

with high probability.

A standard approach for lower bounds is the second moment method: for any non-negative random variable Z, the Cauchy–Schwarz inequality gives

$$\mathbb{P}(Z > 0) \geqslant \frac{(\mathbb{E}Z)^2}{\mathbb{E}[Z^2]}; \tag{5}$$

apply this with $Z = Z_{n\alpha}$ to bound IR_n from below. On $\mathscr{G}_{d,n}$, however, this approach fails to match the first moment bound: there is a regime $\alpha_2 < \alpha < \alpha_\square$, where $Z_{n\alpha}$ is highly non-concentrated, with $\mathbb{E}[(Z_{n\alpha})^2] \gg (\mathbb{E}Z_{n\alpha})^2 \gg 1$; and in fact $\limsup_{n\to\infty} \mathsf{IR}_n < \alpha_\square$ (see e.g. [35]).

This non-concentration of $Z_{n\alpha}$ is caused by a particular geometry within the space of independent sets: we will see that, due to the sparsity of the graph, most independent sets can be locally perturbed in a linear number of places. Specifically, given an independent set S, wherever an unoccupied vertex u has a single occupied neighbor v, the states of the neighbors u and v can be exchanged to obtain a new independent set S'. We say in this case that u and v are "free" (with respect to S), and that S' is connected to S by an edge-swap.

It has been previously observed that edge-swaps are a source of correlation in the space of independent sets (see e.g. [35, Lemma 2.1]). If, given S, we have a subset $U \subseteq V$ of free unoccupied vertices such that every pair $u \neq u'$ in U lies at graph distance $d(u, u') \geqslant 2$, then there are at least $2^{|U|}$ distinct independent sets S' which are connected to S by edge-swaps. If U is of linear size, this means we have found a *cluster* of exponentially many closely correlated independent sets, all of the same size.

We can decompose $Z_{n\alpha}=Z'_{n\alpha}+Z''_{n\alpha}$, where $Z'_{n\alpha}$ counts all independent sets where it is possible to find U as above with $|U|\geqslant nc$. We claim there is a constant c (depending only on d) so that $Z''_{n\alpha}$ (which counts independent sets without $|U|\geqslant nc$) has expectation

$$\mathbb{E}Z''_{n\alpha} \ll \mathbb{E}Z_{\alpha}, \tag{6}$$

up to and even beyond the first moment threshold. This implies

$$\mathbb{E}(Z_{n\alpha})^2 \geqslant \mathbb{E}(Z'_{n\alpha})^2 \geqslant 2^{nc}\mathbb{E}Z'_{n\alpha} \geqslant 2^{nc/2}\mathbb{E}Z_{n\alpha}.$$

Meanwhile, by definition of the first moment threshold $\alpha_\square$, there is a regime of α slightly below $\alpha_\square$ where $2^{nc/2}\gg\mathbb{E}Z_{n\alpha}\gg 1$, implying that $\mathbb{E}(Z_{n\alpha})^2\gg(\mathbb{E}Z_{n\alpha})^2$.

Let us now argue (6). Consider any fixed subset $S\subseteq V$ of size $n\alpha$. In the random regular graph $\mathscr{G}_{d,n}$, conditioned on the event that S is an independent set, the $nd\alpha$ half-edges leaving S are matched to a uniformly random subset of the $nd(1-\alpha)$ half-edges leaving $V_0\equiv V\setminus S$. For $v\in V_0$ let D_v count the number of edges from S to v, and note that S can contribute to $Z''_{n\alpha}$ only if

$$|\{v\in V_0 : D_v=1\}| \leqslant 2ncd.$$

For any $p\in(0,1)$, $(D_v)_{v\in V_0}$ has the same distribution as a vector $(B_v)_{v\in V_0}$ of indipendent and identically distributed (i.i.d.) $\mathrm{Bin}(d,p)$ random variables conditioned to have sum $nd\alpha$. Let us choose $p=\alpha/(1-\alpha)$, so that the B_v sum to $nd\alpha$ with probability $n^{O(1)}$. Then

$$\begin{aligned}\mathbb{E}Z''_{n\alpha} &= \mathbb{E}Z_{n\alpha}\frac{\mathbb{P}\big(|\{v\in V_0 : D_v=1\}|\leqslant 2ncd; \sum_v B_v = nd\alpha\big)}{\mathbb{P}\big(\sum_v B_v=nd\alpha\big)}\\ &\leqslant n^{O(1)}\mathbb{E}Z_{n\alpha}\mathbb{P}(|\{v\in V_0 : X_v=1\}|\leqslant 2ncd).\end{aligned}$$

It is clear that the right-hand side can be made exponentially smaller than $\mathbb{E}Z_{n\alpha}$ by choosing c appropriately (depending on d). Indeed, our argument implies that the ratio between $\mathbb{E}[Z^2_{n\alpha}]$ and $(\mathbb{E}Z_{n\alpha})^2$ is exponentially large in n.

1.3. Replica symmetry breaking

Statistical physicists have developed a deep heuristic for these problems, the so-called RSB formalism, which posits a few structural assumptions on the CSP solution space, and from them deduces precise quantitative predictions (for the sparse setting see [38]). These methods have yielded an understanding of the non-concentration phenomenon as fitting into a much fuller picture. The independent set problem is one of a broad class of CSPs conjectured to have the following "phase diagram" [33]: as the constraint parameter α exceeds a certain threshold α_{s}, the solution space becomes *shattered* into exponentially many well-separated clusters, with each individual cluster comprising an exponentially small fraction of the total mass. This geometry persists up to a further structural transition α_{c} where the solution space *condensates* onto the largest clusters. The regime $\alpha<\alpha_{\mathrm{c}}$ exhibits correlation decay properties [33, equation (5)] known loosely as *replica symmetry*. The condensation transition α_{c} marks the onset of *long-range correlations* (RSB). In the regime $\alpha_{\mathrm{c}}<\alpha<\alpha_\star$, most solutions are concentrated within a bounded number of clusters: the within-cluster correlation then dominates the moment calculation, causing the failure of the second moment method.

The satisfiability transition has been exactly located in some sparse random CSPs *without* an RSB regime, e.g. 2-SAT [16], 1-in-k-SAT [1], and k-XOR-SAT [23], [40]; for 2-SAT, even the finite-size scaling has been successfully characterized [14]. In contrast, no exact satisfiability threshold has been previously located in a sparse random CSP exhibiting RSB, despite a long series of works giving improving bounds in various models within this class, including k-SAT (for $k\geqslant 3$) [32], [4], [18], coloring [2], [17], and independent set.

Many of these models are believed to exhibit a *one-step replica symmetry breaking* (1RSB), in which *clusters are replica symmetric* though the individual solutions are not. In 1RSB models, physicists can predict exact satisfiability thresholds by applying replica symmetric heuristics at the cluster level [36, Chapter 19]. Our main result Theorem 1 confirms the 1RSB prediction for independent sets, which has been derived by [29], [41]. Further, our proof gives some validation to the 1RSB formalism, by locating $\alpha_\star$ as the first moment threshold in a model of independent set clusters.

A natural question is whether the 1RSB prediction holds for independent sets on low-degree regular graphs. Though it is in principle possible to determine from our proof an explicit lower bound d_0 in Theorem 1, 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, calculations suggest that 1RSB at some point breaks down and the low-degree graphs instead exhibit *full* RSB [10], meaning that one expects to find an infinite-depth hierarchy of clusters as is observed in the low-temperature Sherrington–Kirkpatrick model (see [39]). In this regime no formula is predicted even at a heuristic level.

1.4. Overview of our approach

We determine the sharp threshold for independent sets on random regular graphs by a novel approach which rigorizes the 1RSB heuristic from statistical physics, by applying the moment method to count clusters of independent sets rather than the sets themselves. We briefly describe here the key new ideas in our proof; a more detailed outline is presented in §2.

Firstly, we establish a simple combinatorial description for clusters of maximal (or locally maximal) independent sets. On a graph $G=(V,E)$, a (locally) maximal independent set cluster is encoded by what we call a "frozen configuration": an ordered pair $\underline{\bar{\boldsymbol{\eta}}}\equiv(\underline{\eta},\underline{m})$ with $\underline{\eta}\in\{\mathtt{0},\mathtt{1},\mathtt{f}\}^V$ and $\underline{m}\subseteq E$, where

- a $\mathtt{1}$-vertex (interpreted as "occupied throughout the cluster") can only have $\mathtt{0}$'s as neighbors;
- a $\mathtt{0}$-vertex (interpreted as "unoccupied throughout") must have at least two $\mathtt{1}$'s as neighbors;
- $\underline{m}$ is a perfect matching on the $\mathtt{f}$-vertices (interpreted as "taking both states" or "free"), and the total density of $\mathtt{f}$-vertices is $\leqslant d^{-3/2}$.

We shall prove that this model effectively encodes the independent set clusters.

Having established this combinatorial description, the main technical component of this work is to locate the sharp threshold in the cluster model. Importantly, the requirement that each $\mathtt{0}$ neighbor at least *two* $\mathtt{1}$'s makes frozen configurations much more difficult to perturb locally compared with the original independent sets. This local rigidity hints that applying the moment method in this model *does* locate the exact threshold—that is to say, the asymptotic density $\alpha_\star$ appearing in Theorem 1 is the first moment threshold of the cluster model, as is precisely in line with the 1RSB heuristic.

To prove this, we compute the first and second moments of our cluster model up to constants. This calculation is technically challenging, and our solution builds on some of our previous work [21], [20]. Like the original independent set model, the cluster model defines a Gibbs measure (Markov random field) on the random graph—albeit a slightly non-standard one, the most natural form having spins on directed edges rather than on vertices. The (first or second) moment can be understood as an optimization of a rate function $\boldsymbol{\Phi}$ over a simplex of empirical measures, which turns out to have high dimension, growing with d, due to the unusual form of the Gibbs measure. However, by a certain "Bethe variational principle" we are able to characterize local maximizers of $\boldsymbol{\Phi}$ via fixed points of certain tree recursions, reducing the optimization to a dimension constant in d. Even so we are tasked with eliminating a possible multitude of local maximizers, particularly in the second moment which reduces to a fixed-point problem in 81 real variables. This is resolved in the physics folklore by invoking a "causality principle" (cf. [36, equation (19.26)]) which imposes symmetries among the variables, drastically reducing the dimension. Through delicate a-priori estimates we rigorously establish these symmetries, and thereby pinpoint the global maximizer of $\mathbf{\Phi}$. This computes the moments up to polynomial corrections, and we improve the calculation to within constant factors by establishing that $\mathbf{\Phi}$ has negative-definite Hessian at its global maximizer.

The moment calculation itself only establishes the existence of clusters with asymptotically positive probability, but combined with the classical martingale bound this implies that $\mathsf{MIS}_n = n\alpha_\star + O(n^{1/2})$ with high probability, which in turn implies Theorem 1. Our final innovation is a method to improve positive probability bounds to high probability, which proves the stronger result stated in Remark 1.1. The approach is based on controlling the incremental fluctuations of the Doob martingale of a log-transform of the partition function.

The same broad outline of proof applies to our computation [24] of the satisfiability threshold in random regular NAE-SAT. The proof of this paper is more difficult because the cluster representation has additional complications, and also because we require more precise estimates to achieve a sharper bound on the fluctuations (Remark 1.1).

1.5. Configuration model for random regular graphs

Unless indicated otherwise, graphs are permitted to have self-loops and multi-edges. On any graph $G=(V,E)$, say that $\underline{x} \in \{\mathtt{0},\mathtt{1}\}^V$ is an *independent set* if $x_u x_v = \mathtt{0}$ on every edge $(uv) \in E$—in particular, a vertex with a self-loop cannot belong to any independent set.

Let $\mathscr{G}_{d,n}$ denote the uniformly random d-regular graph on n vertices, sampled according to the standard *configuration model*—that is, start with n isolated vertices each equipped with d labeled half-edges, and form the graph by taking a uniformly random matching on the nd half-edges (where nd is assumed to be even). We denote the falling factorial by $(A)_b$ and the falling double factorial by $[A]_b$, namely

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

The first moment of $Z_{n\alpha}$ over $\mathscr{G}_{d,n}$ is given by

$$\mathbb{E} Z_{n\alpha} = \binom{n}{n\alpha} \frac{(nd(1-\alpha)_{nd\alpha})_{nd\alpha}}{[nd]_{nd\alpha}} = n^{O(1)} \underbrace{\left[\frac{(1-\alpha)^{(d-1)(1-\alpha)}}{\alpha^\alpha (1-2\alpha)^{(d/2)(1-2\alpha)}} \right]^n}_{\equiv \exp(\phi(\alpha))}, \tag{7}$$

where ϕ is calculated from Stirling's formula.

Conditioned on the event that $\mathscr{G}_{d,n}$ is free of self-loops and multi-edges, it has the law of the uniformly random *simple* d-regular graph, $\mathcal{G}_{d,n}$. This event occurs (for fixed d)

with uniformly positive probability in the limit $n \to \infty$ (see e.g. [30]), so once Theorem 1 is established for $\mathscr{G}_{d,n}$ it immediately follows for $\mathcal{G}_{d,n}$ as well. We therefore work throughout with the configuration model $\mathscr{G}_{d,n}$.

Notational conventions

For non-negative quantities $f = f_{d,n}$ and $g = g_{d,n}$ we use any of the equivalent notations $f = O_d(g)$, $g = \Omega_d(f)$, $f \lesssim_d g$, and $g \gtrsim_d f$ to indicate that for each fixed $d \geqslant d_0$,

$$\limsup_{n \to \infty} \frac{f}{g} < \infty$$

(with the convention $0/0 \equiv 1$). We drop the subscript d to indicate

$$\limsup_{d \to \infty} \left(\limsup_{n \to \infty} \frac{f}{g} \right) < \infty.$$

We write $f \asymp_d g$ to indicate that both $f \lesssim_d g$ and $g \lesssim_d f$ hold, and drop the subscript d if both $f \lesssim g$ and $g \lesssim f$ hold. Lastly, we use any of the equivalent notations $f = o(g)$, $g = \omega(f)$, $f \ll g$, and $g \gg f$ to indicate that for each fixed $d \geqslant d_0$,

$$\limsup_{n \to \infty} \frac{f}{g} = 0.$$

Acknowledgements

We are grateful to Sourav Chatterjee, Amir Dembo, Persi Diaconis, Elchanan Mossel, and Andrea Montanari for helpful conversations. We thank the anonymous referee for a careful reading of the manuscript and many helpful comments.

2. Model of independent set clusters

In this section we present the combinatorial characterization of independent set clusters.

2.1. Combinatorial representation of clusters

Our first task is to obtain a simple encoding of these clusters. In the geometry of each individual cluster, the main source of complication is that making one edge-swap can free up room for another edge-swap which was not previously permitted, so the minimal chain of edge-swaps joining two independent sets in the same cluster may be extremely long. Indeed, since the local structure of $\mathscr{G}_{d,n}$ is that of the d-regular tree, the propagation of edge-swaps (at least at short distances) behaves as a branching process. A useful heuristic is that, for independent sets at density $y(\log d)/d$ with $y\leqslant 2$, the branching rate will approximately be

$$(d-1)\mathbb{P}\left(\mathrm{Bin}\left(d-1,\frac{y\log d}{d}\right)=1\right)=\frac{y\log d}{d^{y-1}}\left[1+O\left(\frac{(\log d)^2}{d}\right)\right]. \tag{8}$$

The transition between supercritical and subcritical branching occurs at $y\approx 1$. In the near-maximal regime $y\approx 2$ the branching will be quite subcritical: this key fact makes it possible to understand the cluster geometry in a relatively straightforward way.

Our first step in modeling the clusters is to consider the following procedure (cf. [5]): it takes as input a graph $G=(V,E)$ together with an independent set $\underline{x}\in\{\mathtt{0},\mathtt{1}\}^V$. The output is a configuration $\underline{\eta}\in\{\mathtt{0},\mathtt{1},\mathtt{f}\}^V$ ($\mathtt{f}$ denoting "free") together with a subset $\underline{m}\subseteq E$ forming a perfect matching on a subset of the $\mathtt{f}$-vertices (indicating edge-swaps).

Definition 2.1. Given an input graph $G=(V,E)$ and an independent set $\underline{x}\in\{\mathtt{0},\mathtt{1}\}^V$ of G, the *coarsening algorithm* proceeds as follows:

(1) *Form free pairs* (iterate for $0\leqslant s<t$): take the first vertex u with $\eta_u=\mathtt{0}$ and a unique neighbor v with $\eta_v=\mathtt{1}$. Set $\eta_u=\eta_v=\mathtt{f}$ and add (uv) to $\underline{m}$. Iterate until the first time t that no such vertex u remains.

(2) *Identify single frees*: for all vertices v which have state $\mathtt{0}$ and no neighboring $\mathtt{1}$'s under $\underline{\eta}$, update η_v to state $\mathtt{f}$.

At the end of this procedure output the pair $\underline{\bar{\boldsymbol{\eta}}}\equiv(\underline{\eta},\underline{m})$.

(In the first step, vertices are processed in order with respect to the given ordering on the vertices $V=[n]\equiv\{1,...,n\}$.) We shall see that the correspondence between clusters and coarsenings is sufficiently close to a bijection for our purposes. The following proposition, whose proof is given in §2.3, is a rigorous version of the heuristic estimate (8) on the subcritical propagation of the edge-swaps.

PROPOSITION 2.2. *Recall from* (4) *that* $\alpha_\square$ *denotes the first moment threshold for independent sets on* $\mathscr{G}_{d,n}$, *and let* $\alpha_-\equiv\alpha_\square-d^{-5/3}$. *Let* $\boldsymbol{F}$ *denote the event that any independent set on* $\mathscr{G}_{d,n}$ *with density* $\geqslant\alpha_-$ *has, after coarsening, a density of frees* $\geqslant d^{-3/2}$. *Then, for* $d\geqslant d_0$, $\lim_{n\to\infty}\mathbb{P}(\boldsymbol{F})=0$.

We can model the clusters of large independent sets in a simple manner in large part because the frees are so subcritical. The first coarsening step (forming pairs) is the main one conceptually, and the intuition is that the resulting frees—since their density is so low—will occur predominantly in isolated pairs or small trees of linked pairs, reflecting a relatively tractable geometry within the cluster. The second coarsening step (forming single frees) is a clean-up procedure: if we only consider the coarsening applied to *maximum* independent sets, this step becomes relevant only when a chain of edge-swaps can be made along an *odd-length cycle*. (This is also a case in which configurations in the same cluster may have different coarsenings.) We shall show that we can discard these odd-cycle scenarios and still recover the sharp asymptotics for MIS_n.

Definition 2.3. A *frozen configuration* on a graph $G=(V,E)$ is a pair $\underline{\bar{\boldsymbol{\eta}}}\equiv(\underline{\eta},\underline{m})$, with $\underline{\eta}\in\{\mathtt{0},\mathtt{1},\mathtt{f}\}^V$ and $\underline{m}\subseteq E$, satisfying the following:

(1) The $\mathtt{1}$-vertices neighbor only $\mathtt{0}$-vertices;

(2) Each $\mathtt{0}$-vertex has at least two $\mathtt{1}$-neighbors;

(3) The total density of $\mathtt{f}$-vertices is $\leqslant\beta_{\max}\equiv d^{-3/2}$;

(4) The edges $\underline{m}$ form a perfect matching on the $\mathtt{f}$-vertices.

The *frozen model* is the counting measure on frozen configurations.

The frozen model is our independent set cluster representation: clearly, a frozen configuration is an idealized coarsening of an independent set, in which the second step of forming single frees is not needed. The truncation at $\beta_{\max}=d^{-3/2}$ is justified by Proposition 2.2. Write $V_\eta\equiv\{v:\eta_v=\eta\}$, $\eta\in\{\mathtt{0},\mathtt{1},\mathtt{f}\}$, and define the *intensity* of $\underline{\bar{\boldsymbol{\eta}}}$ by

$$\mathbf{i}(\underline{\bar{\boldsymbol{\eta}}})=|V_{\mathtt{1}}|+\tfrac{1}{2}|V_{\mathtt{f}}| \tag{9}$$

(corresponding to the size of the original independent set). We shall always assume that the intensity lies in a restricted regime:

$$\alpha_{\mathrm{lbd}}\equiv\frac{5\log d}{3d}\leqslant\alpha\leqslant\frac{2\log d}{d}\equiv\alpha_{\mathrm{ubd}};$$

and occasionally we will restrict further $\alpha\geqslant\alpha_-\equiv\alpha_\square-d^{-5/3}$ (as defined in Proposition 2.2). Let $\boldsymbol{Z}_{n\alpha}$ count the frozen configurations on $\mathscr{G}_{d,n}$ of intensity exactly $n\alpha$, while $\boldsymbol{Z}_{\geqslant n\alpha}$ counts the configurations of intensity at least $n\alpha$. Explicitly, writing $\boldsymbol{Z}_{n_{\mathtt{1}},n_{\mathtt{f}}}$ for the number of frozen configurations $\underline{\bar{\boldsymbol{\eta}}}\equiv(\underline{\eta},\underline{m})$ on $\mathscr{G}_{d,n}$ with $|V_{\mathtt{1}}|=n_{\mathtt{1}}$ and $|V_{\mathtt{f}}|=n_{\mathtt{f}}$, we have

$$\boldsymbol{Z}_{n\alpha}\equiv\sum_{n_{\mathtt{1}},n_{\mathtt{f}}\geqslant0}\mathbf{1}\left\{\begin{array}{c}n_{\mathtt{f}}\leqslant n\beta_{\max},\\ 2n\alpha=2n_{\mathtt{1}}+n_{\mathtt{f}}\end{array}\right\}\boldsymbol{Z}_{n_{\mathtt{1}},n_{\mathtt{f}}}\quad\text{and}\quad\boldsymbol{Z}_{\geqslant n\alpha}\equiv\sum_{\alpha\leqslant\alpha'}\boldsymbol{Z}_{n\alpha'} \tag{10}$$

The following theorem validates the 1RSB hypothesis by establishing that the threshold for the existence of clusters occurs precisely at the first moment threshold.

THEOREM 2. *For $d\geqslant d_0$ and $\boldsymbol{a}_n=n\alpha_\star+c_\star\log n$ as in Remark 1.1, the following statements hold*:

(a) *The gap $\alpha_\square-\alpha_\star$ is of order $[(\log d)/d]^2$, so $\alpha_-\equiv\alpha_\square-d^{-5/3}$ lies below $\alpha_\star$.*

(b) *For any constant* C, $\limsup_{n\to\infty} \mathbb{E}\boldsymbol{Z}_{\geqslant \boldsymbol{a}_n - C}$ *is finite. We have the upper bound*

$$\lim_{C\to\infty} \limsup_{n\to\infty} \mathbb{E}\boldsymbol{Z}_{\geqslant \boldsymbol{a}_n + C} = 0,$$

as well as the lower bound

$$\liminf_{n\to\infty} \mathbb{P}(\boldsymbol{Z}_{\geqslant \boldsymbol{a}_n} \geqslant 1) > 0.$$

In the rest of this section we explain how this implies our main result Theorem 1. The proof of Theorem 2 occupies the remainder of the paper.

2.2. Independent set threshold from cluster threshold

The coarsening algorithm outputs a pair $(\underline{\eta}, \underline{m})$, where $\underline{\eta} \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}^V$ and $\underline{m}$ is a matching on (a subset of) the $\mathtt{f}$-vertices. Discarding the matching $\underline{m}$ leads to the following definition.

Definition 2.4. On a graph $G=(V,E)$, given a spin configuration $\underline{\eta} \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}^V$, write $\mathfrak{F}(\underline{\eta})$ for the subgraph induced by the $\mathtt{f}$-vertices. An *unweighted coarsening* on a graph $G=(V,E)$ is a spin configuration $\underline{\eta} \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}^V$ satisfying conditions (1)–(3) in Definition 2.3, such that every acyclic component of $\mathfrak{F}(\underline{\eta})$ has a (necessarily unique) perfect matching.

Let $\mathcal{Z}_{n_1, n_\mathtt{f}}$ count unweighted coarsenings $\underline{\eta}$ of $\mathscr{G}_{d,n}$ with $|V_\mathtt{1}| = n_\mathtt{1}$ and $|V_\mathtt{f}| = n_\mathtt{f}$. Let also $\mathcal{Z}_{n_1, n_\mathtt{f}}(\mathcal{O})$ denote the contribution from those $\underline{\eta}$ with exactly $\mathcal{O}$ odd-sized components; from the preceding definition, each such component must contain a cycle. For $\alpha \geqslant \alpha_\square - d^{-5/3}$, if $\mathsf{MIS}_n \geqslant n\alpha$ then either the event $\boldsymbol{F}$ of Proposition 2.2 occurs, or the random variable

$$\mathcal{Z}_{\geqslant n\alpha} \equiv \sum_{\mathcal{O}\geqslant 0} \sum_{n_1, n_\mathtt{f} \geqslant 0} \mathbf{1}\left\{ \begin{array}{c} n_\mathtt{f} \leqslant n\beta_{\max}, \\ 2n\alpha + \mathcal{O} \leqslant 2n_\mathtt{1} + n_\mathtt{f} \end{array} \right\} \mathcal{Z}_{n_1, n_\mathtt{f}}(\mathcal{O}) \tag{11}$$

is strictly positive. Recalling the remarks below the statement of Proposition 2.2, let

$$\begin{aligned} \mathcal{Z}^{\text{tree}} &\equiv \text{contribution to } \mathcal{Z} \text{ from configurations whose } \mathtt{f}\text{-subgraph is acyclic;} \\ \boldsymbol{Z}^{\text{unic}} &\equiv \text{contribution to } \boldsymbol{Z} \text{ from configurations whose } \mathtt{f}\text{-subgraph contains} \\ &\qquad \text{only trees and unicycles.} \end{aligned} \tag{12}$$

The following proposition will be proved in §2.3.

PROPOSITION 2.5. *For* $d \geqslant d_0$, *the following hold uniformly over* $\alpha_{\text{lbd}} \leqslant \alpha \leqslant \alpha_{\text{ubd}}$:
(a) $\mathbb{E}\mathcal{Z}^{\text{tree}}_{\geqslant n\alpha} \geqslant \frac{1}{2}\mathbb{E}\mathcal{Z}_{\geqslant n\alpha}$ *for large* n (*depending only on* d); *and*
(b) $\mathbb{E}\boldsymbol{Z}^{\text{unic}}_{\geqslant n\alpha} = [1 - o(1)]\,\mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}$, *where* $o(1)$ *denotes an error tending to zero as* $n \to \infty$.

Our main result follows from Theorem 2 combined with Propositions 2.2 and 2.5.

Proof of Theorem 1. Recall the notation $\alpha_- \equiv \alpha_\square - d^{-5/3}$. For the upper bound, note that for $\alpha \geqslant \alpha_-$ we have $\{\mathsf{MIS}_n \geqslant n\alpha\} \subseteq \boldsymbol{F} \cup \{\mathcal{Z}_{\geqslant n\alpha} > 0\}$, where $\mathbb{P}(\boldsymbol{F}) = o(1)$ by Proposition 2.2. Markov's inequality and Proposition 2.5 (a) give, for n large, that

$$\mathbb{P}(\mathcal{Z}_{\geqslant n\alpha} > 0) \leqslant \mathbb{E}\mathcal{Z}_{\geqslant n\alpha} \leqslant 2\mathbb{E}\mathcal{Z}^{\text{tree}}_{\geqslant n\alpha} \leqslant 2\mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha},$$

where $\mathcal{Z}^{\text{tree}}_{\geqslant n\alpha} \leqslant \boldsymbol{Z}_{\geqslant n\alpha}$ by comparing Definitions 2.3 and 2.4. Since Theorem 2 (a) gives $\alpha_\star \geqslant \alpha_-$, the above applies with $n\alpha = \boldsymbol{a}_n$, and we conclude that

$$\mathbb{P}(\mathsf{MIS}_n \geqslant \boldsymbol{a}_n + C) \leqslant o(1) + 2\mathbb{E}\boldsymbol{Z}_{\geqslant \boldsymbol{a}_n + C}.$$

It then follows from the upper bound in Theorem 2 (b) that

$$\lim_{C\to\infty} \limsup_{n\to\infty} \mathbb{P}(\mathsf{MIS}_n \geqslant \boldsymbol{a}_n + C) = 0.$$

For the lower bound, note that $\{\mathsf{MIS}_n \geqslant n\alpha\} \supseteq \{\boldsymbol{Z}^{\text{unic}}_{\geqslant n\alpha} > 0\}$ and by Proposition 2.5 (b) we have that $\mathbb{P}(\boldsymbol{Z}^{\text{unic}}_{\geqslant n\alpha} < \boldsymbol{Z}_{\geqslant n\alpha}) = o(1)\,\mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}$. This implies

$$\mathbb{P}(\mathsf{MIS}_n \geqslant n\alpha) \geqslant \mathbb{P}(\boldsymbol{Z}^{\text{unic}}_{\geqslant n\alpha} > 0) \geqslant \mathbb{P}(\boldsymbol{Z}_{\geqslant n\alpha} > 0) - o(1)\mathbb{E}\boldsymbol{Z}_{\geqslant n\alpha}.$$

Taking $n\alpha = \boldsymbol{a}_n$ and applying Theorem 2 (b) on the right-hand side above gives

$$\liminf_{n\to\infty} \mathbb{P}(\mathsf{MIS}_n \geqslant \boldsymbol{a}_n) > 0.$$

Recalling the classical bound $\mathbb{P}(|\mathsf{MIS}_n - \mathbb{E}\mathsf{MIS}_n| \geqslant x) \leqslant \exp(-x^2/nd)$ (see Remark 1.1), we conclude that $\mathbb{E}\mathsf{MIS}_n$ must lie within $O_d(n^{1/2})$ of $\boldsymbol{a}_n$, and the result follows. □

2.3. Large components and trees with matchings

We now turn to the proof of Proposition 2.5. First let us fix some notation. For any measures p and q defined on a discrete space $\mathscr{S}$ we denote the entropy by $H(p)$, and the relative entropy by $H(q\,|\,p)$:

$$H(p) = -\sum_{x\in\mathscr{S}} p(x)\log p(x) \quad \text{and} \quad H(q\,|\,p) = \sum_{x\in\mathscr{S}} q(x)\log\frac{q(x)}{p(x)}.$$

If p and q are probability measures on the binary set $\{0,1\}$, then we may abuse notation and represent the measures p and q by the scalars $x = p(1)$ and $y = q(1)$:

$$H(x) = -x\log x - (1-x)\log(1-x) \quad \text{and} \quad H(y\,|\,x) = y\log\frac{y}{x} + (1-y)\log\frac{1-y}{1-x}.$$

In a graph G, if U is any subgraph or subset of vertices, we write $|U|$ for the number of vertices in U, and ∂U for the set of vertices in $G\setminus U$ that are neighboring to U. We write $\mathtt{E}(U)$ for the number of internal edges of U (counting self-loops and multi-edges).

LEMMA 2.6. *For any $s>1$, there is a bounded constant $C=C(s)$ such that, with high probability, the random d-regular graph contains no subset of n/d^s vertices having more than Cn/d^s internal edges.* (*In particular, $C(s)=(s+1)/(s-1)$ suffices.*)

Proof. Fix a subset S of $n\beta$ vertices, where $\beta=1/d^s$. Generate the random graph according to the configuration model. The number of internal edges of S is stochastically dominated by a binominal random variable,

$$X\sim \mathrm{Bin}\biggl(nd\beta, \frac{nd\beta}{nd-2nd\beta}\biggr),$$

which has mean $\mathbb{E}X\leqslant 2nd\beta^2$. By the Chernoff bound,

$$\mathbb{P}(X\geqslant Cn\beta)\leqslant \exp\biggl(-n\beta\biggl(C\log\frac{1}{d\beta}-O(1)\biggr)\biggr).$$

The number of subsets S of size $n\beta$ is given by

$$\binom{n}{n\beta}=\exp\biggl(n\beta\biggl(\log\frac{1}{\beta}+O(1)\biggr)\biggr).$$

If we set $C=(s+1)/(s-1)$ and take a union bound over the subsets S, we find that the chance for any S of size n/d^s to have more than Cn/d^s internal edges is bounded from above by

$$\exp\biggl(\frac{n}{d^s}\biggl(C\log d-(C-1)s\log d+O(1)\biggr)\biggr)\leqslant \exp\biggl(-\frac{n[\log d-O(1)]}{d^s}\biggr).$$

For d exceeding an absolute constant, the above is clearly $o_n(1)$, proving the lemma. □

Proof of Proposition 2.2. Fix a subset V_1 of $n\alpha$ vertices, and let $V_0=V\setminus V_1$. Supposing that V_1 forms an independent set in the graph, perform the coarsening algorithm and let S' denote the resulting subset of vertices set to $\mathtt{f}$. For any subset $S_0\subseteq S'\cap V_0$, let S_1 denote the subset of vertices in V_1 neighboring to S_0; the definition of coarsening implies that $S_1\subseteq S'$ and $|S_1|\leqslant|S_0|$.

For each $v\in V_0$ let d_v count the number of edges between v and V_1. As before, for any $S_0\subseteq S'\cap V_0$, let S_1 denote the neighbors of S_0 in V_1. Then either the vector $(d_v)_{v\in S_0}$ must have sum at most $10|S_0|$, or the number of internal edges in $S=S_0\cup S_1$ must be more than $10|S_0|$, which in turn is at least $5|S|$ since $|S_1|\leqslant|S_0|$. We saw in Lemma 2.6 that, with high probability, there is no subset S of size $|S|=n/d^{3/2}$ with more than $5|S|$ internal edges. It therefore remains to rule out the case where $(d_v)_{v\in S_0}$ has sum at most $10|S_0|$.

By the same considerations as in the proof of (6), the vector $(d_v)_{v\in V_0}$ has the same distribution as a vector of i.i.d. $\mathrm{Bin}(d,\alpha/(1-\alpha))$ random variables D_i, $1\leqslant i\leqslant n(1-\alpha)$, conditioned to have sum $nd\alpha$. With $\beta=1/d^{3/2}$, we have

$$\mathbb{P}\biggl(\sum_{i=1}^{n\beta}D_i\leqslant 10n\beta\,\biggl|\,\sum_{i=1}^{n(1-\alpha)}D_i=nd\alpha\biggr)\leqslant n^{O(1)}\mathbb{P}\biggl(\sum_{i=1}^{n\beta}D_i\leqslant 10n\beta\biggr).$$

For $\alpha = y(\log d)/d$, it follows by a Chernoff bound that the right-hand side above is less than or equal to

$$\exp(-n\beta(y \log d - O(\log\log d))).$$

For $\beta = 1/d^{3/2}$, the number of subsets $S_0 \subseteq V_0$ of size $n\beta$ is bounded from above by

$$\binom{n(1-\alpha)}{n\beta} \leqslant \binom{n}{n\beta} \leqslant \exp\left(n\beta\left(\frac{3\log d}{2} + O(1)\right)\right).$$

It follows by a union bound that, with high probability, no such S_0 occurs for the entire range

$$\alpha \in \left[\frac{1.55\log d}{d}, \frac{2\log d}{d}\right],$$

which is clearly a superset of the range $[\alpha_\square - d^{-5/3}, \alpha_\square]$. Finally, we saw previously that $\mathsf{MIS}_n \leqslant \alpha_\square < 2(\log d)/d$ with high probability, so the result of the proposition follows. $\square$

We turn now to the proof of Proposition 2.5. Recall from Definition 2.4 that for $\underline{\eta} \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}^V$, we use $\mathfrak{F}(\underline{\eta})$ to denote the subgraph induced by the $\mathtt{f}$-vertices. That is to say, $\mathfrak{F}(\underline{\eta}) = (V_\mathfrak{F}, E_\mathfrak{F})$, where $V_\mathfrak{F} \subseteq V$ is the subset of $\mathtt{f}$-vertices, and $E_\mathfrak{F}$ is the subset of edges with both endpoints in $V_\mathfrak{F}$. Recall that in the d-regular configuration model, each edge is a matching of two half-edges from $[nd]$. Let $H_\mathfrak{F} \subseteq [nd]$ denote the half-edges involved in $E_\mathfrak{F}$. We can then encode the subgraph $\mathfrak{F}$ alone (removing all other edges of the graph) as the subset $H_\mathfrak{F}$ together with a perfect matching on $H_\mathfrak{F}$. Note that given $H_\mathfrak{F}$ we can determine $V_\mathtt{f} = V_\mathfrak{F}$ to be the subset of vertices having an incident half-edge in $H_\mathfrak{F}$.

Let $\mathcal{Z}_{n_1}(\mathfrak{F})$ count unweighted coarsenings $\underline{\eta}$ of $\mathscr{G}_{d,n}$ that have $|V_1| = n_1$ and $\mathfrak{F}(\underline{\eta}) = \mathfrak{F}$. In particular, we must have $\mathcal{Z}_{n_1}(\mathfrak{F}) = 0$ except on the event $B_\mathfrak{F}$ that $\mathfrak{F}$ is precisely the subgraph induced by its incident vertices $V_\mathfrak{F}$. Then

$$\mathcal{Z}_{n_1,n_\mathtt{f}} = \sum_{|\mathfrak{F}|=n_\mathtt{f}} \mathcal{Z}_{n_1}(\mathfrak{F}) \quad \text{and} \quad \boldsymbol{Z}_{n_1,n_\mathtt{f}} = \sum_{|\mathfrak{F}|=n_\mathtt{f}} \mathfrak{m}(\mathfrak{F})\, \mathcal{Z}_{n_1}(\mathfrak{F}), \tag{13}$$

where $\mathfrak{m}(\mathfrak{F})$ counts the number of perfect matchings of $V_\mathtt{f}$ contained in $\mathfrak{F}$.([3]) For a given $\mathfrak{F}$, write $\mathbf{c}(n, n_1, n_\mathtt{f}) \equiv \binom{n-n_\mathtt{f}}{n_1}$ for the number of choices for V_1. Fixing any such V_1,

$$\mathbb{E}[\mathcal{Z}_{n_1}(\mathfrak{F})] = \mathbf{c}(n, n_1, n_\mathtt{f})\mathbb{P}(B_1 \cap B_0 \cap B_\mathtt{f}),$$

where

$$B_1 \equiv \{\text{all edges leaving } V_1 \text{ go to } V_0\},$$

$$B_0 \equiv \{\text{each vertex in } V_0 \text{ has at least two edges coming from } V_1\},$$

$$B_\mathtt{f} \equiv \{\mathfrak{F} \text{ is the subgraph induced by } V_\mathtt{f}\}.$$

([3]) We emphasize that $\mathfrak{m}(\mathfrak{F})$ refers to perfect matchings on the vertices of $V_\mathtt{f}$, in contrast with $\mathfrak{F}$ which is a perfect matching on a subset of labeled half-edges incident to $V_\mathtt{f}$.

We now observe that the quantity $\mathbb{E}[\mathcal{Z}_{n_1}(\mathfrak{F})]$ depends on $\mathfrak{F}$ only through the number of vertices $n_{\mathtt{f}}=|\mathfrak{F}|$, and the number of internal edges $\mathtt{E}\equiv\mathtt{E}(\mathfrak{F})$. To see this, abbreviate $E_\eta\equiv n_\eta d$ and $\boldsymbol{E}\equiv nd$, and first note that

$$\mathbb{P}(B_1)=\mathbf{g}_1(\boldsymbol{E},E_1,E_{\mathtt{f}})\equiv\prod_{i=0}^{E_1-1}\frac{E_0-i}{\boldsymbol{E}-1-2i}=\frac{(E_0)_{E_1}}{[\boldsymbol{E}]_{E_1}}. \tag{14}$$

Next, it is clear that $\mathbb{P}(B_0\,|\,B_1)$ can be expressed as a function $\mathbf{g}_0$ of (n_0,E_1) only. Lastly, if we condition on $B_1\cap B_0$, then all E_1 half-edges leaving V_1 have been assigned, matching with a subset of E_1 half-edges leaving V_0. Out of the $E_{\mathtt{f}}$ half-edges leaving $V_{\mathtt{f}}$, the specification of $\mathfrak{F}$ fixes the assignment of $2\mathtt{E}$ half-edges. Then $E_{\mathtt{f}}-2\mathtt{E}$ half-edges remain to be assigned, and these must match to half-edges from V_0 (by the assumption that $\mathfrak{F}$ is the induced subgraph on $V_{\mathtt{f}}$). Therefore

$$\mathbb{P}(B_{\mathfrak{F}}\,|\,B_1\cap B_0)=\bar{\mathbf{g}}_{\mathtt{f}}(\boldsymbol{E},E_1,E_{\mathtt{f}},\mathtt{E})\equiv\frac{(\boldsymbol{E}-2E_1-E_{\mathtt{f}})_{E_{\mathtt{f}}-2\mathtt{E}}}{[\boldsymbol{E}-2E_1]_{E_{\mathtt{f}}-\mathtt{E}}}.$$

This verifies the claim that $\mathbb{E}[\mathcal{Z}_{n_1}(\mathfrak{F})]$ depends only on $n_{\mathtt{f}}=|\mathfrak{F}|$ and $\mathtt{E}=\mathtt{E}(\mathfrak{F})$:

$$\mathbb{E}[\mathcal{Z}_{n_1}(\mathfrak{F})]=\mathbf{c}(n,n_1,n_{\mathtt{f}})\mathbf{g}_1(\boldsymbol{E},E_1,E_{\mathtt{f}})\mathbf{g}_0(n_0,E_1)\bar{\mathbf{g}}_{\mathtt{f}}(\boldsymbol{E},E_1,E_{\mathtt{f}},\mathtt{E})\equiv\mathbf{g}(n,n_1,n_{\mathtt{f}},\mathtt{E}). \tag{15}$$

As $\mathbf{c}$, $\mathbf{g}_1$, and $\mathbf{g}_0$ do not depend on $\mathtt{E}$, we find crudely that the function $\mathbf{g}$ of (15) satisfies

$$\mathbf{g}(n,n_1,n_{\mathtt{f}},\mathtt{E}+\delta)\leqslant\left(\frac{4}{nd}\right)^{\delta}\mathbf{g}(n,n_1,n_{\mathtt{f}},\mathtt{E}) \tag{16}$$

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 $V_{\mathtt{f}}$ or vice versa: for $\delta\leqslant n/d$ we estimate

(1) $\mathbf{c}(n,n_1,n_{\mathtt{f}}+\delta)=\mathbf{c}(n,n_1,n_{\mathtt{f}})e^{O(\delta)}$,
(2) $\mathbf{g}_1(\boldsymbol{E},E_1,E_{\mathtt{f}}+d\delta)\leqslant\mathbf{g}_1(\boldsymbol{E},E_1,E_{\mathtt{f}})$,
(3) $\bar{\mathbf{g}}_{\mathtt{f}}(\boldsymbol{E},E_1,E_{\mathtt{f}}+d\delta,\mathtt{E})=\bar{\mathbf{g}}_{\mathtt{f}}(\boldsymbol{E},E_1,E_{\mathtt{f}},\mathtt{E})e^{O(\delta)}$,
(4) $\mathbf{g}_0(n_0-\delta,E_1)=\mathbf{g}_0(n_0,E_1)e^{O(\delta)}$,

so

$$\frac{\mathbf{g}(n,n_1,n_{\mathtt{f}}+\delta,\mathtt{E})}{\mathbf{g}(n,n_1,n_{\mathtt{f}},\mathtt{E})}\leqslant e^{O(\delta)}. \tag{17}$$

Estimates (1)–(3) follow straightforwardly from the explicit expressions given above; the proof of (4) is deferred to §2.4.

Fix any $k\geqslant2$ (odd or even), and let $\mathfrak{f}$ be a subgraph having no size-k components and no components consisting of a single isolated edge. Let $\Omega^k_{l,A}(n_{\mathtt{f}};\mathfrak{f})$ denote the collection of subgraphs $\mathfrak{F}$ such that

- $\mathfrak{F}$ contains $\mathfrak{f}$ and has $|\mathfrak{F}|=n_{\mathtt{f}}$ vertices;

• $\mathfrak{F}\setminus\mathfrak{f}$ has exactly l components of size k (not isolated edges), with all remaining components isolated edges; and

• the size-k components have $\mathtt{q}\equiv l(k-\mathbf{1}\{k \text{ even}\})+A$ internal edges.

Recall from Definition 2.4 that any odd-sized component must contain a cycle, so we see that A is non-negative. In the special case $k=2$, $\mathfrak{F}\setminus\mathfrak{f}$ consists of a collection of isolated edges along with l size-2 components each having at least two internal edges, so $A\geqslant l$.

We will argue below that in the first moment calculation, it is more efficient to break up the components of $\mathfrak{F}\setminus\mathfrak{f}$ into isolated edges. To this end, let $\bar{n}_{\mathtt{f}}\equiv n_{\mathtt{f}}+\mathbf{1}\{T \text{ odd}\}$, and note that each $\mathfrak{F}\in\Omega^k_{0,0}(\bar{n}_{\mathtt{f}};\mathfrak{f})$ has the same number $\mathtt{E}$ of internal edges, while each $\mathfrak{F}'\in\Omega^k_{l,A}(n_{\mathtt{f}};\mathfrak{f})$ has the same number $\mathtt{E}''$ of internal edges:

$$\mathtt{E}=\mathtt{E}(\mathfrak{f})+\tfrac{1}{2}(\bar{n}_{\mathtt{f}}-|\mathfrak{f}|)=\left\lceil\tfrac{1}{2}T\right\rceil \quad\text{and}\quad \mathtt{E}''=\mathtt{E}(\mathfrak{f})+\tfrac{1}{2}t+\mathtt{q}=\mathtt{E}+\mathtt{q}-\tfrac{1}{2}lk-\tfrac{1}{2}\mathbf{1}\{T \text{ odd}\}.$$

To compare the two scenarios, define

$$\mathbf{R}^{n_1,n_{\mathtt{f}}}_{k,l,A}(\mathfrak{f})\equiv\frac{|\Omega^k_{l,A}(n_{\mathtt{f}};\mathfrak{f})|}{|\Omega^k_{0,0}(\bar{n}_{\mathtt{f}};\mathfrak{f})|}\frac{\mathbf{g}(n,n_1,n_{\mathtt{f}},\mathtt{E}'')}{\mathbf{g}(n,n_1,\bar{n}_{\mathtt{f}},\mathtt{E})}\left[\frac{\max\{\mathfrak{m}(\mathfrak{F}'):\mathfrak{F}'\in\Omega^k_{l,A}(n_{\mathtt{f}};\mathfrak{f})\}}{\min\{\mathfrak{m}(\mathfrak{F}):\mathfrak{F}\in\Omega^k_{0,0}(\bar{n}_{\mathtt{f}};\mathfrak{f})\}}\right],$$

where we include the ratio in square brackets in order to simultaneously address unweighted coarsenings and the frozen model. This last ratio will not have a substantial effect: first note that if $\mathfrak{m}(\mathfrak{f})=0$, then we must also have $\mathfrak{m}(\mathfrak{F}')=0$ for any $\mathfrak{F}'\in\Omega^k_{l,A}(n_{\mathtt{f}};\mathfrak{f})$; in this case we shall define the ratio to be 1. In general, if $\mathfrak{F}\in\Omega^k_{0,0}(\bar{n}_{\mathtt{f}};\mathfrak{f})$, then $\mathfrak{m}(\mathfrak{F})=\mathfrak{m}(\mathfrak{f})$, while if $\mathfrak{F}'\in\Omega^k_{l,A}(n_{\mathtt{f}};\mathfrak{f})$, then $\mathfrak{m}(\mathfrak{F})\leqslant 2^{\mathtt{q}}\mathfrak{m}(\mathfrak{f})$. Therefore in any case the ratio in square brackets is $\leqslant 2^{\mathtt{q}}$, and we will see in the calculation below that $e^{O(\mathtt{q})}$ factors can be ignored.

LEMMA 2.7. *For $d\geqslant d_0$, $n_{\mathtt{f}}\leqslant n\beta_{\max}$, $(2n_1+n_{\mathtt{f}})/2n\in[\alpha_{\mathrm{lbd}},\alpha_{\mathrm{ubd}}]$, and $n\geqslant n_0(d)$, there exists an absolute constant c such that*

$$\mathbf{R}^{n_1,n_{\mathtt{f}}}_{k,l,A}(\mathfrak{f})\leqslant\left(\frac{(cd\beta_{\max})^{k/2}}{l}\left(\frac{n}{dk}\right)^{\mathbf{1}\{k\ \mathit{even}\}}\right)^{l}\left(\frac{cdk}{n}\right)^{A}. \tag{18}$$

Proof. Given $\mathfrak{f}$, an element of $\Omega^k_{l,A}(n_{\mathtt{f}};\mathfrak{f})$ is obtained as follows: first, from the $n-|\mathfrak{f}|$ vertices that lie outside $\mathfrak{f}$, choose $T\equiv n_{\mathtt{f}}-|\mathfrak{f}|$ vertices to belong to $\mathfrak{F}\setminus\mathfrak{f}$. From these we choose a further subset of $t\equiv T-lk$ vertices to belong to the isolated edge components. For each of these t vertices we choose one of the d incident half-edges, then take a perfect matching on the t chosen half-edges to form the isolated edge components. Next we turn to the remaining lk chosen vertices, and divide these into l groups of size k; the number of ways to do this is $(lk)!/l!(k!)^l$. To determine the internal edges among these components, we can first choose an ordered list of half-edges $a_1, \ldots, a_{\mathtt{q}}$ from the lkd half-edges available. Each half-edge a_i must match to a partner b_i which is one of the $\leqslant kd$ half-edges available within the same component of k vertices. There are $\mathtt{q}!2^{\mathtt{q}}$ lists of ordered pairs $(a_i, b_i)_{i\leqslant \mathtt{q}}$ yielding the same set of internal edges, so altogether we have

$$\begin{aligned}|\Omega^k_{l,A}(n_{\mathtt{f}};\mathfrak{f})| &\leqslant \binom{n-|\mathfrak{f}|}{T} \underbrace{\left[\binom{T}{t} d^t (t-1)!!\right]}_{\text{size-2 components}} \underbrace{\left[\frac{(lk)!}{l!(k!)^l}\frac{(lkd)^{\mathtt{q}}(kd)^{\mathtt{q}}}{\mathtt{q}!2^{\mathtt{q}}}\right]}_{\text{size-}k\text{ components}} \\ &= \frac{(n-|\mathfrak{f}|)_T}{(\frac{1}{2}t)!}\frac{d^{t+2\mathtt{q}}}{2^{t/2+\mathtt{q}}}\frac{(lk)^{\mathtt{q}}}{\mathtt{q}!}\frac{k^{\mathtt{q}}}{l!(k!)^l} \leqslant e^{O(\mathtt{q})}\frac{(n-|\mathfrak{f}|)_T}{(\frac{1}{2}t)!}\frac{d^{t+2\mathtt{q}}}{2^{t/2}}\frac{k^{\mathtt{q}}}{(lk^k)^l},\end{aligned} \tag{19}$$

where the last bound uses $\mathtt{q}\geqslant l(k-1)\geqslant \frac{1}{2}lk$. For T even, combining (16) and (19) gives

$$\mathbf{R}^{n_1,n_{\mathtt{f}}}_{k,l,A}(\mathfrak{f}) \leqslant \frac{|\Omega^k_{l,A}(n_{\mathtt{f}};\mathfrak{f})| 2^{\mathtt{q}}(4/nd)^{\mathtt{q}-lk/2}}{\binom{n-|\mathfrak{f}|}{T} d^T (T-1)!!} \leqslant e^{O(\mathtt{q})}\left[\frac{(\frac{1}{2}T)!2^{lk/2}}{(\frac{1}{2}t)!}\right]\frac{k^{\mathtt{q}}}{(lk^k)^l}\left(\frac{4d}{n}\right)^{\mathtt{q}-lk/2}.$$

Since $T=t+lk\leqslant n\beta_{\max}$, the first factor is $\leqslant T^{lk/2}\leqslant (n\beta_{\max})^{lk/2}$. It follows that for some absolute constant c (not depending on d),

$$\mathbf{R}^{n_1,n_{\mathtt{f}}}_{k,l,A}(\mathfrak{f}) \overset{(T\text{ even})}{\leqslant} \left(\frac{(cd\beta_{\max})^{k/2}}{l}\left(\frac{n}{dk}\right)^{\mathbf{1}\{k\text{ even}\}}\right)^l \left(\frac{cdk}{n}\right)^A.$$

For T odd, combining (16), (17), and (19) (and adjusting the constant c as needed) gives

$$\mathbf{R}^{n_1,n_{\mathtt{f}}}_{k,l,A}(\mathfrak{f}) \overset{(T\text{ odd})}{\leqslant} \frac{|\Omega^k_{l,A}(n_{\mathtt{f}};\mathfrak{f})| O[(4/nd)^{\mathtt{q}-lk/2-1/2}]}{\binom{n-|\mathfrak{f}|}{T+1} d^{T+1}(T!!)} \leqslant \left(\frac{(cd\beta_{\max})^{k/2}}{l}\right)^l \left(\frac{cdk}{n}\right)^A,$$

concluding the proof. $\square$

LEMMA 2.8. *Write $\mathcal{Z}''_{\geqslant n\alpha}$ for the contribution to $\mathcal{Z}_{\geqslant n\alpha}$ from configurations having at least one $\mathtt{f}$-component which either has an odd number of vertices, or has two vertices with more than one internal edge (meaning it must contain a self-loop or doubled edge). For $d\geqslant d_0$, $\alpha_{\mathrm{lbd}}\leqslant\alpha\leqslant\alpha_{\mathrm{ubd}}$, and $n\geqslant n_0(d)$, we have*

$$\mathbb{E}[\mathcal{Z}''_{\geqslant n\alpha}] \leqslant d^{-1/3}\mathbb{E}\mathcal{Z}_{\geqslant n\alpha}.$$

Proof. Recall the notation of (13) and (15). Let $\mathcal{Z}_{n_1,n_{\mathtt{f}}}(L_k=l)$ denote the contribution to $\mathcal{Z}_{n_1,n_{\mathtt{f}}}$ from configurations whose $\mathtt{f}$-subgraph has exactly l components of size k that are not isolated edges. For the case $k=2$, applying Lemma 2.7 gives

$$\frac{\sum_{l\geqslant 1}\mathbb{E}[\mathcal{Z}_{n_1,n_{\mathtt{f}}}(L_k=l)]}{\mathbb{E}[\mathcal{Z}_{n_1,n_{\mathtt{f}}}(L_k=0)]} \leqslant \frac{\sum_{l\geqslant 1}\sum_{\mathfrak{f}}\sum_{A\geqslant l}|\Omega^k_{0,0}(n_{\mathtt{f}};\mathfrak{f})|\mathbf{g}(n,n_1,n_{\mathtt{f}},\mathtt{E}'')}{\sum_{\mathfrak{f}}|\Omega^k_{0,0}(n_{\mathtt{f}};\mathfrak{f})|\mathbf{g}(n,n_1,n_{\mathtt{f}},\mathtt{E})} \lesssim d\beta_{\max}. \tag{20}$$

Next assume that k is odd, and let $\mathcal{Z}_{n_1,n_\mathtt{f}}(\mathcal{O};L_k=l)$ denote the contribution to $\mathcal{Z}_{n_1,n_\mathtt{f}}(\mathcal{O})$ from configurations having exactly l $\mathtt{f}$-components of size k (where $l\leqslant\mathcal{O}$). Then

$$\frac{\mathbb{E}[\mathcal{Z}_{n_1,n_\mathtt{f}}(\mathcal{O};L_k=l)]}{\mathbb{E}[\mathcal{Z}_{n_1,\bar{n}_\mathtt{f}}(\mathcal{O}-l;L_k=0)]}=\frac{\sum_{\mathfrak{f}:\mathcal{O}(\mathfrak{f})+l=\mathcal{O}}\sum_{A\geqslant 0}|\Omega^k_{l,A}(n_\mathtt{f};\mathfrak{f})|\mathbf{g}(n,n_1,n_\mathtt{f},\mathtt{E}'')}{\sum_{\mathfrak{f}:\mathcal{O}(\mathfrak{f})+l=\mathcal{O}}|\Omega^k_{0,0}(\bar{n}_\mathtt{f};\mathfrak{f})|\mathbf{g}(n,n_1,\bar{n}_\mathtt{f},\mathtt{E})},$$

and applying Lemma 2.7 gives (for $l\geqslant 1$)

$$\mathbb{E}[\mathcal{Z}_{n_1,n_\mathtt{f}}(\mathcal{O};L_k=l)]\leqslant(cd\beta_{\max})^{kl/2}\mathbb{E}[\mathcal{Z}_{n_1,\bar{n}_\mathtt{f}}(\mathcal{O}-l;L_k=0)].$$

For $k\geqslant 3$ odd and $l\geqslant 1$, summing both sides of the above over n_1, $n_\mathtt{f}$, and $\mathcal{O}$ with $\mathcal{O}\geqslant l$ and $2n\alpha+\mathcal{O}\leqslant 2n_1+n_\mathtt{f}$ gives $\mathbb{E}[\mathcal{Z}_{\geqslant n\alpha};L_k=l]\leqslant d^{-lk/5}\mathbb{E}\mathcal{Z}_{\geqslant n\alpha}$ (cf. (11)). The claimed bound now follows readily by summing over odd $k\geqslant 3$ and $l\geqslant 1$, and combining with (20). □

We turn now to the components of even size $k\geqslant 4$. Recall that L_k counts the number of size-k components in the $\mathtt{f}$-subgraph, and define

$$\begin{aligned} l_k &\equiv[(cd\beta_{\max})^{k/8}n/dk]\vee 1,\\ \boldsymbol{z} &\equiv \boldsymbol{Z}(L_k<l_k \text{ for all } k\geqslant 4 \text{ even});\\ z^{\mathrm{ev}} &\equiv \mathcal{Z}^{\mathrm{ev}}(L_k<l_k \text{ for all } k\geqslant 4 \text{ even}),\ \text{where } \mathcal{Z}^{\mathrm{ev}}\equiv\mathcal{Z}(\mathcal{O}=0).\end{aligned}\tag{21}$$

LEMMA 2.9. *For $d\geqslant d_0$, $\alpha_{\mathrm{lbd}}\leqslant\alpha\leqslant\alpha_{\mathrm{ubd}}$, and $n\geqslant n_0(d)$,*

$$\mathbb{E}[z^{\mathrm{ev}}_{\geqslant n\alpha}]\geqslant(1-n^{-1/2})\mathbb{E}[\mathcal{Z}^{\mathrm{ev}}_{\geqslant n\alpha}]\quad\textit{and}\quad\mathbb{E}[\boldsymbol{z}_{\geqslant n\alpha}]\geqslant(1-n^{-1/2})\mathbb{E}[\boldsymbol{Z}_{\geqslant n\alpha}].$$

Proof. For $k\leqslant n\beta_{\max}$ and $l\geqslant l_k$, (18) gives $\mathbf{R}^{n_1,n_\mathtt{f}}_{k,l,A}(\mathfrak{f})\leqslant(cd\beta_{\max})^{(3/8)lk+A}$, so

$$\begin{aligned}\frac{\mathbb{E}[\mathcal{Z}^{\mathrm{ev}}_{n_1,n_\mathtt{f}}(L_k\geqslant l_k)]}{\mathbb{E}[\mathcal{Z}^{\mathrm{ev}}_{n_1,n_\mathtt{f}}(L_k=0)]}&\leqslant\frac{\sum_{\mathcal{O}(\mathfrak{f})=0}\sum_{l\geqslant l_k}\sum_{A\geqslant 0}|\Omega^k_{l,A}(n_\mathtt{f};\mathfrak{f})|\mathbf{g}(n,n_1,n_\mathtt{f},\mathtt{E}'')}{\sum_{\mathcal{O}(\mathfrak{f})=0}|\Omega^k_{0,0}(n_\mathtt{f};\mathfrak{f})|\mathbf{g}(n,n_1,n_\mathtt{f},\mathtt{E})}\\ &\leqslant 2(cd\beta_{\max})^{(3/8)l_k k}.\end{aligned}$$

We claim that the above is $\leqslant n^{-5/2}$. If $l_k k>\log n$ this clearly holds. If $l_k k\leqslant\log n$, then by definition of l_k we have $(cd\beta_{\max})^{k/8}\leqslant d(\log n)/n$, so the above is $\leqslant[d(\log n)/n]^{3l_k}\ll n^{-5/2}$, verifying the claim. Summing over even $k\geqslant 4$, and n_1 and $n_\mathtt{f}$ with $2n_1+n_\mathtt{f}\geqslant 2n\alpha$ proves the bound for $\mathcal{Z}$. The bound for $\boldsymbol{Z}$ follows in exactly the same manner since the effect of reweighting by the number of matchings was already accounted for in (18). □

Proof of Proposition 2.5. Fix $k\geqslant 4$ even and assume the quantity l_k of Lemma 2.9 to be strictly greater than 1. For a subgraph $\mathfrak{f}^*$ having no components of size k, let $\Xi^k_{l,r,A'}(n_\mathtt{f};\mathfrak{f}^*)$ denote the collection of subgraphs $\mathfrak{F}$ such that

(1) $\mathfrak{F}$ contains $\mathfrak{f}$ and has $|\mathfrak{F}|=n_\mathtt{f}$ vertices; and

(2) $\mathfrak{F}\backslash\mathfrak{f}^*$ consists of l size-k components, of which $l-r$ components are trees with (unique) perfect matching, while the remaining r components are not trees and have $\mathtt{q}'\equiv rk+A'$ internal edges.

Note that each $\mathfrak{F}\in\Xi^k_{l,0,0}(n_{\mathtt{f}};\mathfrak{f}^*)$ has the same number $\mathtt{E}$ of internal edges, while each $\mathfrak{F}\in\Xi^k_{l,r,A'}(n_{\mathtt{f}};\mathfrak{f}^*)$ has the same number $\mathtt{E}''$ of internal edges:

$$\mathtt{E}=\mathtt{E}(\mathfrak{f})+l(k-1) \quad \text{and} \quad \mathtt{E}''=\mathtt{E}(\mathfrak{f})+(l-r)(k-1)+\mathtt{q}'=\mathtt{E}+r+A'.$$

Similarly as before, define the ratio

$$\mathbf{S}^{n_1,n_{\mathtt{f}}}_{k,l,r,A'}(\mathfrak{f}^*)\equiv\frac{|\Xi^k_{l,r,A'}(n_{\mathtt{f}};\mathfrak{f}^*)|}{|\Xi^k_{l,0,0}(n_{\mathtt{f}};\mathfrak{f}^*)|}\,\frac{\mathbf{g}(n_1,n_{\mathtt{f}},\mathtt{E}'')}{\mathbf{g}(n_1,n_{\mathtt{f}},\mathtt{E})}\,\frac{1\vee\max\{\mathfrak{m}(\mathfrak{F}):\mathfrak{F}\in\Xi^k_{l,r,A'}(n_{\mathtt{f}};\mathfrak{f}^*)\}}{1\vee\min\{\mathfrak{m}(\mathfrak{F}):\mathfrak{F}\in\Xi^k_{l,0,0}(n_{\mathtt{f}};\mathfrak{f}^*)\}}.$$

Given k isolated vertices each equipped with d half-edges, let $\text{TPM}_{d,k}$ denote the number of ways to match up $k-1$ pairs of half-edges such that the resulting graph is a tree on k vertices with a (unique) perfect matching.(4) We crudely bound the number of spanning trees from below by the number of line graphs, which clearly have a perfect matching as k is even: $\text{TPM}_{d,k}\geqslant\frac{1}{2}k!d^k(d-1)^{k-2}\geqslant e^{O(k)}k^kd^{2(k-1)}$. Combining with (16) gives

$$\mathbf{S}^{n_1,n_{\mathtt{f}}}_{k,l,r,A'}(\mathfrak{f}^*)\leqslant 2^{\mathtt{q}'}\,\frac{\binom{l}{r}(\text{TPM}_{d,k})^{l-r}}{(\text{TPM}_{d,k})^l}\,\frac{(rkd)^{\mathtt{q}'}(kd)^{\mathtt{q}'}}{2^{\mathtt{q}'}(\mathtt{q}')!}\left(\frac{4}{nd}\right)^{r+A'}$$

(similarly to the derivation of (19)). Simplifying, we see that

$$\mathbf{S}^{n_1,n_{\mathtt{f}}}_{k,l,r,A'}(\mathfrak{f}^*)\leqslant\left(\frac{C^kld}{n}\right)^r\left(\frac{Cdk}{n}\right)^{A'}$$

for C an absolute constant. Recall that we assumed $l_k=(cd\beta_{\max})^{k/8}n/dk>1$; as a result it suffices to consider $k\leqslant C^{-1}\log n$, so

$$\mathbf{S}^{n_1,n_{\mathtt{f}}}_{k,l,r,A'}(\mathfrak{f}^*)\leqslant(cd\beta_{\max})^{kr/9}\left[\frac{d(\log n)}{n}\right]^{A'}.$$

Decompose the quantities appearing in (21) as

$$z^{\text{ev}}\equiv z^{\text{tree}}+z^{\text{cyc}} \quad \text{and} \quad \boldsymbol{z}=\boldsymbol{z}^{\text{unic}}+\boldsymbol{z}^{\text{bic}},$$

where z^{tree} and $\boldsymbol{z}^{\text{unic}}$ are defined analogously to (12). Then

$$\frac{\mathbb{E}[z^{\text{cyc}}_{n_1,n_{\mathtt{f}}}]}{\mathbb{E}[z^{\text{ev}}_{n_1,n_{\mathtt{f}}}]}\leqslant\sum_{k\geqslant 4\text{ even}}\frac{\sum_{\mathfrak{f}^*:\mathcal{O}(\mathfrak{f}^*)=0}\sum_{1\leqslant r\leqslant l\leqslant l_k}\sum_{A'\geqslant 0}|\Xi^k_{l,r,A'}(n_{\mathtt{f}};\mathfrak{f}^*)|\mathbf{g}(n,n_1,n_{\mathtt{f}},\mathtt{E}'')}{\sum_{\mathfrak{f}^*:\mathcal{O}(\mathfrak{f}^*)=0}\sum_{1\leqslant l\leqslant l_k}|\Xi^k_{l,0,0}(n_{\mathtt{f}};\mathfrak{f}^*)|\mathbf{g}(n,n_1,n_{\mathtt{f}},\mathtt{E})}\leqslant d^{-1/5},$$

thus $\mathbb{E}[z^{\text{cyc}}_{\geqslant n\alpha}]\leqslant d^{-1/5}\mathbb{E}[z^{\text{ev}}_{\geqslant n\alpha}]$. A very similar calculation shows $\mathbb{E}[\boldsymbol{z}^{\text{bic}}_{\geqslant n\alpha}]=o(1)\,\mathbb{E}[\boldsymbol{z}_{\geqslant n\alpha}]$ (in this case we sum over $A'\geqslant 1$ instead of $A'\geqslant 0$). Combining with Lemmas 2.8 and 2.9 concludes the proof. $\square$

(4) This means a perfect matching of the tree's vertices, not to be confused with the matching of half-edges.

2.4. Estimates on forcing constraints

In this subsection we calculate the probability cost of the constraint that each $\mathtt{0}$-vertex is forced by at least two neighboring $\mathtt{1}$-vertices—that is, the function $\mathbf{g}_{\mathtt{0}}(n_{\mathtt{0}}, E_{\mathtt{1}})$ appearing in (15) is equal to $\mathbb{P}(B_{\mathtt{0}}\,|\,B_{\mathtt{1}})$. To be slightly more general, for a fixed positive integer $\boldsymbol{\ell}$, write $\mathbf{g}_{\mathtt{0},\boldsymbol{\ell}}(n, E_{\mathtt{1}})$ for the probability, with respect to a uniformly random assignment of $E_{\mathtt{1}}$ half-edges to n vertices of degree d, that each vertex receives at least $\boldsymbol{\ell}$ incoming half-edges (so $\mathbf{g}_{\mathtt{0}} \equiv \mathbf{g}_{\mathtt{0},\boldsymbol{\ell}=2}$).

We also analyze a bivariate analogue $\mathbf{g}_{\mathtt{00},\boldsymbol{\ell}}$ which will be used in our second moment analysis §4. Let $\mathscr{X} \equiv \{\mathtt{0},\mathtt{1}\}^2 \setminus \{\mathtt{00}\}$ and write $\theta \equiv (\theta_\omega)_{\omega\in\mathscr{X}}$. If the three entries of θ are positive and have sum strictly less than 1, then

$$\left(\theta_{\mathtt{11}}, \theta_{\mathtt{10}}, \theta_{\mathtt{01}}, 1 - \sum_{\omega\in\mathscr{X}} \theta_\omega\right)$$

defines a probability measure on $\{\mathtt{0},\mathtt{1}\}^2$ with full support. We take ϖ to be a $\{\mathtt{0},\mathtt{1}\}^2$-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

$$X \equiv (X_\omega)_{\omega\in\mathscr{X}} \equiv (|\{1 \leqslant j \leqslant d : \varpi_j = \omega\}|)_{\omega\in\mathscr{X}}. \tag{22}$$

Though we always use θ and X to denote three-dimensional vectors indexed by $\omega\in\mathscr{X}$, we also use the abbreviations

$$\theta_{\mathtt{00}} \equiv 1 - \sum_{\omega\in\mathscr{X}} \theta_\omega \quad \text{and} \quad X_{\mathtt{00}} \equiv d - \sum_{\omega\in\mathscr{X}} X_\omega.$$

Let p_θ denote the law of X. Let X^i, $1\leqslant i\leqslant n$, be independent random variables identically distributed according to p_θ, and write $\mathbb{P}_\theta$ for their (joint) law:

$$\mathbb{P}_\theta(\underline{x}) \equiv \prod_{i=1}^n p_\theta(x^i), \quad p_\theta(x^i) \equiv \frac{d!(\theta_{\mathtt{11}})^{x_{\mathtt{11}}}(\theta_{\mathtt{10}})^{x_{\mathtt{10}}}(\theta_{\mathtt{01}})^{x_{\mathtt{01}}}(\theta_{\mathtt{00}})^{x_{\mathtt{00}}}}{x_{\mathtt{11}}!x_{\mathtt{10}}!x_{\mathtt{01}}!x_{\mathtt{00}}!}.$$

Write $\llcorner\!\lrcorner$ for $\{\mathtt{0},\mathtt{1},\mathtt{f}\}$, e.g. $x_{\mathtt{1}\llcorner\!\lrcorner} \equiv x_{\mathtt{10}} + x_{\mathtt{11}} + x_{\mathtt{1f}}$. Then, with $\Omega \equiv \{x : x_{\mathtt{1}\llcorner\!\lrcorner} \wedge x_{\llcorner\!\lrcorner\mathtt{1}} < \boldsymbol{\ell}\}$,

$$\mathbf{g}_{\mathtt{0},\boldsymbol{\ell}}(n, E_{\mathtt{1}}) \equiv \mathbb{P}_\theta\left(X^i_{\mathtt{1}\llcorner\!\lrcorner} \geqslant \boldsymbol{\ell} \ \text{ for } i = 1, ..., n \,\middle|\, \sum_{i=1}^n X^i_{\mathtt{1}\llcorner\!\lrcorner} = E_{\mathtt{1}}\right),$$
$$\mathbf{g}_{\mathtt{00},\boldsymbol{\ell}}(n, \underline{E}) \equiv \mathbb{P}_\theta\left(X^i \notin \Omega \ \text{ for } i = 1, ..., n \,\middle|\, \sum_{i=1}^n X^i = \underline{E}\right).$$

Further, write $X\sim\nu$ to indicate that $(X^i)_{i\leqslant n}$ has empirical measure ν, and likewise write $X_{1\_}\sim\nu'$ to indicate that $X_{1\_}\equiv(X^i_{1\_})_{i\leqslant n}$ has empirical measure ν'. Define

$$\mathbf{p}_{0,\boldsymbol{\ell}}(n,E_1,\nu')\equiv\mathbb{P}_\theta\biggl(X_{1\_}\sim\nu'\biggm|\sum_{i=1}^n X^i_{1\_}=E_1, X^i_{1\_}\geqslant\boldsymbol{\ell}\text{ for }i=1,\dots,n\biggr),$$
$$\mathbf{p}_{00,\boldsymbol{\ell}}(n,\underline{E},\nu)\equiv\mathbb{P}_\theta\biggl(X\sim\nu\biggm|\sum_{i=1}^n X^i=\underline{E}, X^i\notin\Omega\text{ for }i=1,\dots,n\biggr).$$

PROPOSITION 2.10. *Let ε be a small constant uniform in d, and suppose that*

$$\frac{\varepsilon\log d}{d}\leqslant\zeta_{1\_},\zeta_{\_1}\leqslant\frac{10\log d}{d}.$$

(a) *Let a and y be defined by $\zeta_{1\_}\equiv a(\log d)/d$ and $\zeta_{1\_}\wedge\zeta_{\_1}\equiv y(\log d)/d$. Then*

$$\mathbf{g}_{0,\boldsymbol{\ell}}(n,nd\zeta_{1\_})=\exp(O(nd^{-a}(\log d)^{\boldsymbol{\ell}})),$$
$$\mathbf{g}_{00,\boldsymbol{\ell}}(n,nd\zeta)=\exp(O(nd^{-y}(\log d)^{\boldsymbol{\ell}})).$$

Further, $\mathbf{p}_{0,\boldsymbol{\ell}}(n,nd\zeta_{1\_},\nu')$ and $\mathbf{p}_{00,\boldsymbol{\ell}}(n,nd\zeta,\nu)$ are exponentially small in n, unless

$$\begin{aligned}\nu'(x)&=\exp\biggl(\frac{(\log d)^{O(1)}}{d^y}\biggr)\mathbf{1}\{x\geqslant\boldsymbol{\ell}\}\mathbb{P}(\mathrm{Bin}(d,\zeta_{1\_})=x),\\ \nu(x)&=\exp\biggl(\frac{(\log d)^{O(1)}}{d^y}\biggr)\mathbf{1}\{x\in\Omega\}p_\zeta(x).\end{aligned}\tag{23}$$

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

$$\mathbf{g}_{0,\boldsymbol{\ell}}(n,nd\xi_{1\_})=\mathbf{g}_{0,\boldsymbol{\ell}}(n,nd\zeta_{1\_})\exp(O(nd|\xi_{1\_}-\zeta_{1\_}|)),$$
$$\mathbf{g}_{00,\boldsymbol{\ell}}(n,nd\xi)=\mathbf{g}_{00,\boldsymbol{\ell}}(n,nd\zeta)\exp(O(nd\|\xi-\zeta\|_1)).$$

LEMMA 2.11. *For the multinomial random variable X defined by (22), consider the cumulant generating function $\Lambda_\theta(\gamma)\equiv\log\mathbb{E}_\theta[e^{\langle\gamma,X\rangle}\,|\,X\notin\Omega]$, defined for $\gamma\in\mathbb{R}^3$. For positive vectors θ and ζ in the regime*

$$\frac{\varepsilon\log d}{d}\leqslant\theta_{1\_},\theta_{\_1},\zeta_{1\_},\zeta_{\_1}\leqslant\frac{10\log d}{d},$$

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

$$\biggl|\gamma_\omega-\log\frac{\theta_{00}\zeta_\omega}{\zeta_{00}\theta_\omega}\biggr|\lesssim\frac{(\log d)^{\boldsymbol{\ell}-1}}{d^y}\quad\textit{for all }\omega\in\mathscr{X}.$$

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

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

In particular, each entry of $\Lambda'_\theta(\gamma)$ lies between 0 and d. Now, for $\varepsilon > 0$ small, consider the ε-perturbed 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. [22, Lemma 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 $\theta_{\mathtt{00}} \geqslant \theta_\omega e^{\gamma_{\varepsilon,\omega}}$ 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

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

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

$$\gamma_{\varepsilon,\omega} = \varepsilon^{-1}(d\zeta_\omega - \partial_\omega \Lambda_\theta(\gamma_\varepsilon)) \geqslant \varepsilon^{-1} d\zeta_\omega \left(1 - \frac{e^{\gamma_{\varepsilon,\omega}}}{\theta_{\mathtt{00}}}\right),$$

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 dth step of the random walk

$$X_t \equiv (X_{t,\omega})_{\omega\in\mathscr{X}} \equiv (|1 \leqslant j \leqslant t : \varpi_j = \omega|)_{\omega\in\mathscr{X}}.$$

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

$$d\zeta_{\mathtt{11}} = \mathbb{E}_\mu[X_{\mathtt{11}} \,|\, X \notin \Omega] \leqslant \boldsymbol{\ell} + \mathbb{E}_\mu[X_{d,\mathtt{11}} - X_{\tau,\mathtt{11}} \,|\, \tau \leqslant d] \leqslant \boldsymbol{\ell} + \mathbb{E}_\mu[X_{d-\tau,\mathtt{11}}] \leqslant \boldsymbol{\ell} + d\mu_{\mathtt{11}}. \qquad (24)$$

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

$$\mathbb{E}_\theta[X_\omega \,|\, X_\pi \geqslant l] = \sum_{k\geqslant l} \mathbb{P}(X_\pi = k \,|\, X_\pi \geqslant l)\,\mathbb{E}_\theta[X_\omega \,|\, X_\pi = k] \leqslant \mathbb{E}_\theta X_\omega = d\theta_\omega. \qquad (25)$$

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

$$\begin{aligned} d\zeta_{10} &= \mathbb{E}_\mu[X_{10}\,|\,X\notin\Omega] = \frac{\mathbb{E}_\mu[X_{10}\mathbf{1}\{\tau\leqslant d\}]}{\mathbb{P}_\mu(\tau\leqslant d)} \leqslant \boldsymbol{\ell}+\frac{\mathbb{E}_\mu[(X_{10}-X_{\tau_{1\text{-}},10})\mathbf{1}\{\tau\leqslant d\}]}{\mathbb{P}_\mu(\tau\leqslant d)} \\ &= \boldsymbol{\ell}+\sum_{k<d}\sum_x \frac{\mathbb{P}_\mu(\tau_{1\text{-}}=k, X_{\tau_{1\text{-}}}=x, \widetilde{X}_{d-k,\text{-}1}\geqslant\boldsymbol{\ell}-x_{\text{-}1})}{\mathbb{P}_\mu(\tau\leqslant d)}\mathbb{E}_\mu[\widetilde{X}_{d-k,10}\,|\,\widetilde{X}_{d-k,\text{-}1}\geqslant\boldsymbol{\ell}-x_{\text{-}1}], \end{aligned}$$

where $(\widetilde{X}_t)_{t\geqslant 0}$ is an independent realization of the random walk X. Maximizing over all possible k and x and applying (25) gives

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

and symmetrically

$$d\zeta_{01}\leqslant\boldsymbol{\ell}+d\mu_{01}.$$

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

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

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

Write p_μ for the law of $X\equiv X_d$ and q_μ for that 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

$$d\zeta_\omega=\frac{\mathbb{E}_\mu X_\omega-d\mu_\omega\sum_{x\in\Omega}q_\mu(x-\mathbf{1}_\omega)}{1-\mathbb{P}_\mu(X\in\Omega)}=d\mu_\omega\,[1+O(d^{-y}(\log d)^{\boldsymbol{\ell}-1})] \tag{26}$$

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

$$z_{\theta,\gamma}=\frac{\theta_{00}}{\zeta_{00}}\frac{\zeta_{00}}{\mu_{00}}=\frac{\theta_{00}}{\zeta_{00}}[1+O(d^{-1-y}(\log d)^{\boldsymbol{\ell}})], \tag{27}$$

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

It follows (see e.g. [22, Lemma 2.3.9]) that, with θ and 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

$$\Lambda^*_\theta(d\zeta)=\langle\gamma,d\zeta\rangle-\Lambda(\gamma), \tag{28}$$

where γ is the solution of $\Lambda'_\theta(\gamma)=d\zeta$. 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 (28) that Λ^*_θ is differentiable with respect to ζ, with gradient $(\Lambda^*_\theta)'(d\zeta)=\gamma$.

Proof of Proposition 2.10. Let $\mathbf{g}_{00,\boldsymbol{\ell}}(n,nd\zeta,\nu)$ be the contribution to $\mathbf{g}_{00,\boldsymbol{\ell}}(n,nd\zeta)$ from the event that $X\sim\nu$. This can be positive only if $\nu(\Omega)=0$ and $\bar{\nu}=d\zeta$, in which case

$$\mathbf{g}_{00,\boldsymbol{\ell}}(n,nd\zeta,\nu)=\mathbf{g}_{00,\boldsymbol{\ell}}(n,nd\zeta)\mathbf{p}_{00,\boldsymbol{\ell}}(n,nd\zeta,\nu)=\frac{\mathbb{P}_\theta(X\sim\nu)}{\mathbb{P}_\theta\left(\sum_{i=1}^n X^i=nd\zeta\right)}.$$

By Stirling's formula, the denominator is $\asymp_d n^{-3/2}\exp(-ndH(\zeta\,|\,\theta))$. Fixing θ, we optimize the numerator over ν, introducing a Lagrangian term for the constraint on $\bar{\nu}$:

$$\begin{aligned}\mathbb{P}_\theta(X\sim\nu)&=\binom{n}{n\nu}\prod_x p_\theta(x)^{n\nu(x)}\exp\left(n\left\langle\gamma,\sum_x x\nu(x)-d\zeta\right\rangle\right)\\&=n^{O(1)}\exp\left(-n\left[\langle\gamma,d\zeta\rangle+\sum_x\nu(x)(\log\nu(x)-\log(p_\theta(x)\exp(\langle\gamma,x\rangle)))\right]\right).\end{aligned}$$

For any fixed θ and γ, over the space $\{\nu:\nu(\Omega)=0\}$ (without the constraint on $\bar{\nu}$), the exponent of $\mathbb{P}_\theta(X\sim\nu)$ is maximized at

$$\nu(x)=\frac{\mathbf{1}\{x\notin\Omega\}p_\theta(x)\exp(\langle\gamma,x\rangle)}{\sum_{x'\notin\Omega}p_\theta(x')\exp(\langle\gamma,x'\rangle)}.\tag{29}$$

If we take γ to be the unique solution γ given by Lemma 2.11 of $\Lambda_\theta'(\gamma)=d\zeta$, then (29) gives a measure ν that satisfies $\bar{\nu}=d\zeta$. Thus, recalling (28),

$$\mathbf{g}_{00,\boldsymbol{\ell}}(n,nd\zeta)\asymp_d[1-\mathbb{P}_\theta(\Omega)]^n\ \exp(-n\Lambda_\theta^*(d\zeta)+ndH(\zeta\,|\,\theta)).$$

We can estimate this by taking $\theta=\zeta$. Since $p_\theta(x)\exp(\langle\gamma,x\rangle)=p_\mu(x)(z_{\theta,\gamma})^d$, it follows from (26) and (27) that

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

Combining with (28) and the bound of Lemma 2.11 gives $\Lambda_\zeta^*(d\zeta)=O(d^{-y}(\log d)^{\boldsymbol{\ell}})$. Moreover, the estimate for $\mathbf{p}_{00,\boldsymbol{\ell}}(n,nd\zeta,\nu)$ in Proposition 2.10 (a) follows from (29) together with the estimate on γ given by Lemma 2.11. To compare $\mathbf{g}_{00,\boldsymbol{\ell}}(n,nd\zeta)$ with $\mathbf{g}_{00,\boldsymbol{\ell}}(n,nd\xi)$, where $\xi_\omega=\zeta_\omega[1+O(\delta/n)]$, let $\theta=\zeta$. Recalling $(\Lambda_\theta^*)'(d\zeta)=\gamma$, the bound of Lemma 2.11 gives

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

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

The remainder of this paper is occupied with the proof of Theorem 2. From now on it is assumed, even where not explicitly stated, that $\alpha_{\mathrm{lbd}}\leqslant\alpha\leqslant\alpha_{\mathrm{ubd}}$, $d\geqslant d_0$, and $n\geqslant n_0$, where n_0 may depend on d.

3. First moment of frozen model

In this section we identify the leading exponential order of the first moment of the frozen model partition function (10),

$$\boldsymbol{\phi}(\alpha)=\lim_{n\to\infty} n^{-1}\log \mathbb{E}\boldsymbol{Z}_{n\alpha}.$$

Recall the definitions in equations (1) and (2).

THEOREM 3.1. *For $d\geqslant d_0$ and $\alpha_{\mathrm{lbd}}\leqslant\alpha\leqslant\alpha_{\mathrm{ubd}}$, $\boldsymbol{\phi}(\alpha)$ is given by $\mathbf{f}(q)$ evaluated at the unique solution $q\in[1.6(\log d)/d, 3(\log d)/d]$ of the equation $\alpha(q)=\alpha$. The function $\boldsymbol{\phi}(\alpha)$ is strictly decreasing on this interval of α, with a unique zero $\alpha_\star$ in the interval's interior. The gap between $\alpha_\star$ and the first moment threshold $\alpha_\square$ from (4) is given by*

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

Note that Theorem 3.1 implies Theorem 2 (a).

3.1. Gibbs measures on message configurations

In [20] and [21] we established a variational principle for calculating the expectation over $\mathscr{G}_{d,n}$ of the partition function Z for a *Gibbs measure* (also termed Markov random field or factor model). By decomposing Z as a sum of the contributions $Z(\boldsymbol{h})$ from all possible edge empirical measures $\boldsymbol{h}$, each $\mathbb{E}Z(\boldsymbol{h})$ can be calculated simply as a product of multinomial coefficients, and is found to scale exponentially in n (up to $n^{O(1)}$ corrections) with a rate $\boldsymbol{\Phi}(\boldsymbol{h})$ which varies smoothly with $\boldsymbol{h}$. Thus, disregarding polynomial corrections, the moment calculation reduces to optimization of $\boldsymbol{\Phi}$ over the simplex of edge empirical measures. We then showed that interior stationary points of $\boldsymbol{\Phi}$ are in bijective correspondence with fixed points of the *belief propagation* or *Bethe recursion* for the corresponding model on the infinite d-regular tree—thereby reducing an optimization to a fixed point problem, generally in a lower-dimensional space.

To count frozen configurations of a particular intensity, we shall introduce in the frozen model a *fugacity* parameter $\lambda>0$, where $\log\lambda$ will serve as a Lagrange multiplier: under the λ-*weighted frozen model*, each valid frozen configuration $\underline{\bar{\boldsymbol{\eta}}}\equiv(\underline{\eta},\underline{m})$ receives weight $\lambda^{\mathbf{i}(\underline{\bar{\boldsymbol{\eta}}})}$ where $\mathbf{i}(\underline{\bar{\boldsymbol{\eta}}})=|V_{\mathtt{1}}|+\frac{1}{2}|V_{\mathtt{f}}|$ as defined in (9).(5) The partition function of this

(5) We call λ the fugacity in accordance with standard terminology—the *hard-core model* at fugacity λ is the measure supported on independent sets of a given graph, in which each set S appears with weight $\lambda^{|S|}$.

model restricted to configurations of intensity $n\alpha$ is simply $\boldsymbol{Z}^{\lambda}_{n\alpha} \equiv \lambda^{n\alpha}\boldsymbol{Z}_{n\alpha}$; and the unrestricted partition function is

$$\boldsymbol{Z}^{\lambda} = \sum_{\alpha} \lambda^{n\alpha}\boldsymbol{Z}_{n\alpha}.$$

Given α we will adjust λ appropriately so that the dominant contribution to $\mathbb{E}\boldsymbol{Z}^{\lambda}$ comes from the configurations of intensity $n\alpha$.

Some difficulty is also posed by the fact that the frozen model is not a Gibbs measure in the most conventional sense of variable spins coupled only by edge interactions. However, we describe two (closely related) ways to recast the frozen model as a Gibbs measure, although with spins located on the edges of the graph. One way to do this, which has been introduced in the physics literature [38] (see also [15], [34] and [36, Chapter 19]), is by way of *message configurations* or *auxiliary configurations* $\underline{\sigma}$,([6]) where each directed edge $v\to w$ carries the message

$$\sigma_{v\to w}\colon \text{“state of } v \text{ in absence of } w\text{.”}$$

Note that, in a valid message configuration, $\sigma_{v\to w}$ will be a certain function of the messages $(\sigma_{u\to v})_{u\in\partial v\setminus w}$ incoming to v from its other neighbors. The actual state η_v of v will then be a function of all its incoming messages $\underline{\dot{\sigma}}_{\partial v\to v} \equiv (\sigma_{u\to v})_{u\in\partial v}$. The model is formally defined as follows. Let $\{\mathtt{0},\mathtt{1},\mathtt{f}\}^*$ be the union of $\{\mathtt{0},\mathtt{1},\mathtt{f}\}^l$ over $l \geqslant 0$. We then define the message-passing rule

$$\dot{\mathtt{m}}\colon \{\mathtt{0},\mathtt{1},\mathtt{f}\}^* \longrightarrow \{\mathtt{0},\mathtt{1},\mathtt{f}\}, \quad \dot{\mathtt{m}}(\underline{\eta}) = \begin{cases} \mathtt{1}, & \text{if } |\{i : \eta_i = \mathtt{1}\}| = 0, \\ \mathtt{f}, & \text{if } |\{i : \eta_i = \mathtt{1}\}| = 1, \\ \mathtt{0}, & \text{if } |\{i : \eta_i = \mathtt{1}\}| \geqslant 2. \end{cases}$$

Definition 3.2. A valid *message configuration* on $\mathscr{G}_{d,n} = (V, E)$ is a vector $(\sigma_{v\to w})$, indexed by all pairs of neighbors $v, w \in V$, that satisfies the message-passing rules

$$\sigma_{v\to w} = \dot{\mathtt{m}}[(\sigma_{u\to v})_{u\in\partial v\setminus w}] \quad \text{for all } (vw) \in E.$$

Its corresponding frozen configuration $\underline{\bar{\boldsymbol{\eta}}} \equiv (\underline{\eta}, \underline{m})$ is given by setting

$$\begin{aligned} \eta_v &= \dot{\mathtt{m}}[(\sigma_{u\to v})_{u\in\partial v}] \text{ for all } v \in V, \\ \underline{m} &= \{(uv) : \sigma_{u\to v} = \sigma_{v\to u} = \mathtt{1}\}. \end{aligned}$$

([6]) In the literature, this formalism is sometimes also called “warning propagation,” and the configurations $\underline{\sigma}$ are called “warning configurations.”

Let us check that if a message configuration maps to $\underline{\bar{\boldsymbol{\eta}}}$ under Definition 3.2, then $\underline{\bar{\boldsymbol{\eta}}}$ is in fact a valid frozen configuration according to Definition 2.3, assuming

$$\{v : \eta_v = \mathtt{f}\} \leqslant n\beta_{\max}.$$

– If $\eta_v = \mathtt{1}$, then all messages incoming to v are $\mathtt{0}$ or $\mathtt{f}$, so all outgoing ones are $\mathtt{1}$.

– If $\eta_v = \mathtt{f}$, then all incoming messages are $\mathtt{0}$ or $\mathtt{f}$ except for a single incoming $\mathtt{1}$ from some $u \in \partial v$, so all outgoing messages are $\mathtt{f}$ except for the return message $\sigma_{v\to u} = \mathtt{1}$.

– If $\eta_v = \mathtt{0}$, then there are at least two incoming $\mathtt{1}$-messages. If there are more than two then all outgoing messages are $\mathtt{0}$, and we term this a *robust* zero, labeled $\mathtt{0}^{\mathtt{r}}$. If v receives exactly two incoming $\mathtt{1}$-messages from $u, w \in \partial v$, then all messages outgoing from v are $\mathtt{0}$ except $\sigma_{v\to u} = \sigma_{v\to w} = \mathtt{f}$. We term this a *susceptible* zero, labeled $\mathtt{0}^{\mathtt{s}}$.

We see from the above that if $\sigma_{u\to v} = \mathtt{1}$ and $\sigma_{v\to u} \neq \mathtt{1}$, then $\eta_u = \mathtt{1}$ and $\eta_v = \mathtt{0}$. If $\eta_v = \mathtt{0}$ then it has no outgoing $\mathtt{1}$-messages, so we see that for $u \in \partial v$, $\eta_u = \mathtt{1}$ if and only if $\sigma_{u\to v} = \mathtt{1}$. This verifies the condition that $|\{u \in \partial v : \eta_u = \mathtt{1}\}| \geqslant 2$ whenever $\eta_v = \mathtt{0}$, as well as the condition that $\{|\{v \in \partial u : \eta_v \neq \mathtt{0}\}| = 0\}$ whenever $\eta_u = \mathtt{1}$. Lastly it is clear that whenever $\eta_u = \mathtt{f}$, u is matched under $\underline{m}$ to the unique neighboring $\mathtt{f}$-vertex v for which $\sigma_{u\to v} = \sigma_{v\to u} = \mathtt{1}$.

We will see below that this is one way of expressing the frozen model as a Gibbs measure. A second way, which is slightly more direct but less standard, is to start with a frozen configuration $\underline{\bar{\boldsymbol{\eta}}}$, and simply set $\sigma_{v\to w} = \eta_v$ except for when v is matched to w, in which case set $\sigma_{v\to w} = \bar{\mathtt{f}}$. This defines a different vector $(\sigma_{v\to w})$, which we term a *vertex-message configuration*. It is clearly simply a rewriting of the original configuration $\underline{\bar{\boldsymbol{\eta}}}$.

PROPOSITION 3.3. *There is a one-to-one correspondence between* (i) *frozen configurations*, (ii) *vertex-message configurations*, *and* (iii) *message configurations.*

Proof. It is clear from the definition that frozen configurations are in bijection with vertex-message configurations: for any edge (vw), the vertex spin η_v can be read off from $\sigma_{v\to w}$, and (vw) participates in a matching if and only if $\sigma_{v\to w} = \bar{\mathtt{f}}$. In the discussion following Definition 3.2 we saw that message configurations map to frozen configurations (equivalently, to vertex-message configurations). To invert the mapping, (i) change every $\bar{\mathtt{f}}$-message in the vertex-message configuration to $\mathtt{1}$, and (ii) change the $u \to v$ message from $\mathtt{0}$ (in the vertex-message configuration) to $\mathtt{f}$ whenever $\eta_u = \mathtt{0}^{\mathtt{s}}$ and $\eta_v = \mathtt{1}$.([7]) $\square$

Let us introduce some convenient notation. We bisect each edge in $\mathscr{G}_{d,n}$ 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

([7]) 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.

the original graph) and the set F of *clauses* (edges in the original graph). We will denote variables generally by u,v,w, and clauses by a,b,c. The new graph has edge set $\boldsymbol{E}$, where $(av){=}(va){\in}\boldsymbol{E}$ indicates that in the original graph, vertex v is incident to edge a. These edges are labeled, so the enumeration of graphs is now $(nd)!$ rather than $(nd{-}1)!!$, but clearly the problem remains unchanged. We denote the bipartite graph $(V,F,\boldsymbol{E})$.

Given a valid auxiliary configuration on $\mathscr{G}_{d,n}$ we define a auxiliary configuration on the bipartite graph simply by setting

$$\sigma_{v\to w}\equiv\sigma_{v\to a}\equiv\sigma_{a\to v},$$

whenever the variables $v,w{\in}V$ are joined by the clause $a{\in}F$. That is to say, in the bipartite graph, variables pass messages using the same rule $\dot{\mathtt{m}}$ as above, while clauses act trivially by passing on the same message. We write $\sigma_{av}{\equiv}\sigma_{va}$ for the pair of messages $(\sigma_{v\to a},\sigma_{a\to v})$ on the edge (av), where in the pair we always write the variable-to-clause message first, and the clause-to-variable messages second. Let $\mathscr{M}$ denote the space of all possible values for σ_{av} in an auxiliary configuration, and let $\mathscr{M}_{\mathrm{v}}$ denote the space of values in a vertex-auxiliary configuration. The mapping from message configurations to vertex-message configurations is given simply by coordinatewise application of the map $\mathsf{proj}_{\mathrm{v}}\colon\mathscr{M}\to\mathscr{M}_{\mathrm{v}}$ which takes

$$\mathtt{11}\longmapsto\overline{\mathtt{ff}},\quad\{\mathtt{10},\mathtt{1f}\}\longmapsto\mathtt{10},\quad\{\mathtt{01},\mathtt{f1}\}\longmapsto\mathtt{01},\tag{30}$$

and acts as the identity on the remaining spins.

Remark 3.4. In the remainder of the paper we will work with both the original graph $\mathscr{G}_{d,n}$ and its (equivalent) bipartite version $(V,F,\boldsymbol{E})$. To avoid confusion, we refer to elements of $\boldsymbol{E}$ as *variable-clause edges*. When we say simply “edges,” we mean the edges of the original graph $\mathscr{G}_{d,n}$.

Given a message configuration $\underline{\sigma}$ on $(V,F,\boldsymbol{E})$, write $\underline{\dot{\sigma}}_v$ for the d-tuple of spins incident to variable $v{\in}V$, and write $\underline{\hat{\sigma}}_a$ for the pair of spins incident to clause $a{\in}F$. The counting measure on valid message configurations $\underline{\sigma}$ of $(V,F,\boldsymbol{E})$ is given by

$$\Psi(\underline{\sigma})\equiv\prod_{v\in V}\dot{\varphi}(\underline{\dot{\sigma}}_v)\prod_{a\in F}\hat{\varphi}(\underline{\hat{\sigma}}_a),\tag{31}$$

where the variable factor $\dot{\varphi}(\underline{\dot{\sigma}}_v)$ checks the message-passing rule at v, and the clause factor $\hat{\varphi}(\underline{\hat{\sigma}}_a)$ checks the message-passing rule at a. Explicitly,

$$\begin{aligned}
\hat{\varphi}(\sigma,\sigma') &= \mathbf{1}\{\sigma'=\mathsf{R}\sigma\}, \quad \text{where } \mathsf{R}\colon \mathscr{M}\to\mathscr{M},\ \eta\eta'\mapsto\eta'\eta \text{ (reflection map)},\\
\dot{\varphi}(\underline{\dot{\sigma}}_v) &= \prod_{w\in\partial v}\mathbf{1}\{\sigma_{v\to w}=\dot{\mathtt{m}}[(\sigma_{u\to v})_{u\in\partial v\setminus w}]\}\\
&= \begin{cases}
1, & \underline{\dot{\sigma}}\in\mathrm{Per}[(\mathtt{10}^j,\mathtt{1f}^{d-j})_{0\leqslant j\leqslant d}] \quad \eta_v=\mathtt{1},\\
1, & \underline{\dot{\sigma}}\in\mathrm{Per}[(\mathtt{11},\mathtt{f0}^j,\mathtt{ff}^{d-1-j})_{0\leqslant j\leqslant d-1}] \quad \eta_v=\mathtt{f},\\
1, & \underline{\dot{\sigma}}\in\mathrm{Per}[(\mathtt{f1}^2,\mathtt{00}^j,\mathtt{0f}^{d-2-j})_{0\leqslant j\leqslant d-2}] \quad \eta_v=\mathtt{0}^{\mathtt{s}} \text{ (susceptible)},\\
1, & \underline{\dot{\sigma}}\in\mathrm{Per}[(\mathtt{01}^k,\mathtt{00}^j,\mathtt{0f}^{d-k-j})_{0\leqslant j\leqslant d-k,3\leqslant k\leqslant d}] \quad \eta_v=\mathtt{0}^{\mathtt{r}} \text{ (robust)},\\
0, & \text{otherwise},
\end{cases}
\end{aligned} \tag{32}$$

where Per denotes the set of all permutations of the tuples listed. The measure (31) is a *Gibbs measure* or *factor model*, meaning that it is specified by a product of local factors. Note that the vertex-auxiliary configurations also define a Gibbs measure, with factors $\varphi_{\mathrm{v}}\equiv(\dot{\varphi}_{\mathrm{v}},\hat{\varphi})_{\mathrm{v}}$ obtained by applying the weights $\varphi\equiv(\dot{\varphi},\hat{\varphi})$ of (32) to the pre-image under $\mathsf{proj}_{\mathrm{v}}$. Extending the above discussion to the λ-weighted model we have the following result.

PROPOSITION 3.5. *The λ-weighted frozen model is in measure-preserving one-to-one correspondence with the Gibbs measure on message configurations with weights*

$$\Psi^{\lambda}(\underline{\sigma})\equiv\prod_{v\in V}\dot{\varphi}^{\lambda}(\underline{\dot{\sigma}}_v)\prod_{a\in F}\hat{\varphi}^{\lambda}(\underline{\hat{\sigma}}_a)\equiv\lambda^{\mathbf{i}(\underline{\sigma})}\prod_{v\in V}\dot{\varphi}(\underline{\dot{\sigma}}_v)\prod_{a\in F}\hat{\varphi}(\underline{\hat{\sigma}}_a)\equiv\lambda^{\mathbf{i}(\underline{\sigma})}\Psi(\underline{\sigma}), \tag{33}$$

where $\mathbf{i}(\underline{\sigma})\equiv\mathbf{i}[\underline{\bar{\boldsymbol{\eta}}}(\underline{\sigma})]$ and $\varphi^{\lambda}\equiv(\dot{\varphi}^{\lambda},\hat{\varphi}^{\lambda})$ are λ-weighted versions of the indicator functions $\varphi\equiv(\dot{\varphi},\hat{\varphi})$ defined in (32):

$$\hat{\varphi}^{\lambda}(\sigma,\sigma')\equiv\lambda^{\mathbf{1}\{\sigma=\mathtt{11}\}}\hat{\varphi}(\sigma,\sigma') \quad \textit{and} \quad \dot{\varphi}^{\lambda}(\underline{\dot{\sigma}})\equiv\lambda^{\mathbf{1}\{\eta(\underline{\dot{\sigma}})=\mathtt{1}\}}\dot{\varphi}(\underline{\dot{\sigma}}).$$

Definition 3.6. The λ-*weighted auxiliary model* is the Gibbs measure (33) restricted to configurations with $\leqslant n\beta_{\max}$ $\mathtt{f}$-variables. Its image under $\mathsf{proj}_{\mathrm{v}}$ is the λ-*weighted vertex-auxiliary model*. The model partition function is the normalizing constant

$$\boldsymbol{Z}^{\lambda}=\sum_{\underline{\sigma}\in\mathscr{M}^{\boldsymbol{E}}}\mathbf{1}\{|\{v:\eta_v=\mathtt{f}\}|\leqslant n\beta_{\max}\}\,\Psi^{\lambda}(\underline{\sigma}).$$

3.2. Bethe variational principle

Write $m \equiv \frac{1}{2}nd$ for the number of clauses. Given a message configuration $\underline{\sigma} \in \mathscr{M}^{\boldsymbol{E}}$, consider the normalized empirical measures

$$\begin{aligned}
\dot{\boldsymbol{h}}(\underline{\dot\sigma}) &\equiv n^{-1} \sum_{v \in V} \mathbf{1}\{\underline{\dot\sigma}_v = \underline{\dot\sigma}\}, && \underline{\dot\sigma} \in \mathscr{M}^d && \text{(variable empirical measure);}\\
\hat{\boldsymbol{h}}(\underline{\hat\sigma}) &\equiv m^{-1} \sum_{a \in F} \mathbf{1}\{\underline{\hat\sigma}_a = \underline{\hat\sigma}\}, && \underline{\hat\sigma} \in \mathscr{M}^2 && \text{(clause empirical measure);}\\
\bar{h}(\sigma) &\equiv (nd)^{-1} \sum_{e \in \boldsymbol{E}} \mathbf{1}\{\sigma_e = \sigma\}, && \sigma \in \mathscr{M} && \text{(half-edge empirical measure).}
\end{aligned}$$

We regard $\boldsymbol{h} \equiv (\dot{\boldsymbol{h}}, \hat{\boldsymbol{h}})$ as a probability measure on $\operatorname{supp}\varphi$ (meaning that $\dot{\boldsymbol{h}}$ is a probability measure on $\operatorname{supp}\dot\varphi$, while $\hat{\boldsymbol{h}}$ is a probability measure on $\operatorname{supp}\hat\varphi$). Let

$$\begin{aligned}
\dot{H}_{\sigma,\underline{\dot\sigma}} &\equiv \text{number of appearances of } \sigma \text{ in } \underline{\dot\sigma},\\
\dot{H}_{\sigma,\underline{\dot\sigma}} &\equiv \text{number of appearances of } \sigma \text{ in } \underline{\hat\sigma}.
\end{aligned}$$

Write $\dot{\boldsymbol{s}} \equiv |\operatorname{supp}\dot\varphi|$ and $\bar{s} \equiv |\operatorname{supp}\hat\varphi| = |\mathscr{M}|$, and write

$$\begin{aligned}
\dot{H} &\equiv (\dot{H}_{\sigma,\underline{\dot\sigma}} : \sigma \in \mathscr{M}, \underline{\dot\sigma} \in \operatorname{supp}\dot\varphi) \in \mathbb{Z}^{\bar{s} \times \dot{\boldsymbol{s}}},\\
\widehat{H} &\equiv (\widehat{H}_{\sigma,\underline{\hat\sigma}} : \sigma \in \mathscr{M}, \underline{\hat\sigma} \in \operatorname{supp}\hat\varphi) \in \mathbb{Z}^{\bar{s} \times \bar{s}}.
\end{aligned}$$

For $\boldsymbol{h}$ to correspond to a valid configuration $\underline{\sigma}$, the variable and clause empirical measures must give rise to the same non-normalized marginals, $nd\bar{h} = n\dot{H}\dot{\boldsymbol{h}} = m\widehat{H}\hat{\boldsymbol{h}}$.

Definition 3.7. Let $\boldsymbol{\Delta}$ denote the simplex of probability measures $\boldsymbol{h}$ on $\operatorname{supp}\varphi$ such that $\dot{\boldsymbol{h}}(\eta = \mathtt{f}) \leqslant \beta_{\max}$, and $\big(\dot{\boldsymbol{h}}, \frac{1}{2}d\hat{\boldsymbol{h}}\big)$ lies in the kernel of $H_{\boldsymbol{\Delta}} \equiv (\,\dot{H}\ -\widehat{H}\,) \in \mathbb{Z}^{\bar{s} \times (\dot{\boldsymbol{s}} + \bar{s})}$. For $\boldsymbol{h} \in \boldsymbol{\Delta}$, we say that the *normalized intensity* of $\boldsymbol{h}$ is

$$\mathbf{i}(\boldsymbol{h}) \equiv \dot{\boldsymbol{h}}(\eta = \mathtt{1}) + \tfrac{1}{2}\dot{\boldsymbol{h}}(\eta = \mathtt{f}) = \bar{h}(\mathtt{1Z}) + \tfrac{1}{2}d\bar{h}(\mathtt{11}). \tag{34}$$

Let $\boldsymbol{\Delta}[\alpha]$ denote the subsimplex of measures $\boldsymbol{h} \in \boldsymbol{\Delta}$ with normalized intensity $\mathbf{i}(\boldsymbol{h}) = \alpha$. We shall show (Lemma 5.4) that $H_{\boldsymbol{\Delta}}$ is surjective, implying that $\boldsymbol{\Delta}$ is an $(\dot{\boldsymbol{s}} - 1)$-dimensional simplex with $\boldsymbol{\Delta}[\alpha]$ an $(\dot{\boldsymbol{s}} - 2)$-dimensional subsimplex.

Let $\boldsymbol{Z}(\boldsymbol{h})$ denote the contribution to the partition function from message configurations $\underline{\sigma} \in \mathscr{M}^{\boldsymbol{E}}$ with empirical measure $\boldsymbol{h}$, so that $\boldsymbol{Z}_{n\alpha} = \sum_{\boldsymbol{h} \in \boldsymbol{\Delta}[\alpha]} \boldsymbol{Z}(\boldsymbol{h})$. Calculating configuration model probabilities (cf. [20, §2.1]) gives, with the usual multi-index notation,

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

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

$$\boldsymbol{\Phi}(\boldsymbol{h})\equiv\sum_{\underline{\dot\sigma}}\dot{\boldsymbol{h}}(\underline{\dot\sigma})\log\frac{\dot\varphi(\underline{\dot\sigma})}{\dot{\boldsymbol{h}}(\underline{\dot\sigma})}+\frac{d}{2}\sum_{\underline{\hat\sigma}}\hat{\boldsymbol{h}}(\underline{\hat\sigma})\log\frac{\hat\varphi(\underline{\hat\sigma})}{\hat{\boldsymbol{h}}(\underline{\hat\sigma})}-d\sum_{\sigma}\bar h(\sigma)\log\frac{1}{\bar h(\sigma)}. \tag{35}$$

In the above, the first sum goes over all $\underline{\dot\sigma}$ in $\operatorname{supp}\dot\varphi$, the second sum goes over all $\underline{\hat\sigma}$ in $\operatorname{supp}\hat\varphi$, and the final sum goes over all σ in $\mathscr{M}$. If further $\min\boldsymbol{h}\gtrsim_d 1$, then

$$\mathbb{E}\boldsymbol{Z}(\boldsymbol{h})=\frac{e^{O_d(n^{-1})}\mathscr{P}(\boldsymbol{h})}{(2\pi n)^{(\boldsymbol{s}-1)/2}}\exp(n\boldsymbol{\Phi}(\boldsymbol{h})),\quad \mathscr{P}(\boldsymbol{h})\equiv\left(\frac{\prod_\sigma d\bar h(\sigma)}{2\prod_{\underline{\dot\sigma}}\dot{\boldsymbol{h}}(\underline{\dot\sigma})\prod_{\underline{\hat\sigma}}\frac{1}{2}d\hat{\boldsymbol{h}}(\underline{\hat\sigma})}\right)^{1/2}. \tag{36}$$

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). We shall compute the first moment exponent $\boldsymbol{\phi}(\alpha)$ by using the Lagrangian method to locate

$$^\star\boldsymbol{h}[\alpha]\equiv\arg\max\nolimits_{\boldsymbol{h}\in\boldsymbol{\Delta}[\alpha]}\boldsymbol{\Phi}(\boldsymbol{h}). \tag{37}$$

If $\boldsymbol{h}$ is a stationary point of $\boldsymbol{\Phi}$ restricted to $\boldsymbol{\Delta}[\alpha]$, then for some $\lambda=\lambda(\alpha)$ it must be a stationary point of $\boldsymbol{\Phi}^\lambda$ on the unrestricted space $\boldsymbol{\Delta}$.

Remark 3.8. Observe from the functional form of $\boldsymbol{\Phi}^\lambda$ that $^\star\dot{\boldsymbol{h}}={}^\star\dot{\boldsymbol{h}}[\alpha]$ and $^\star\hat{\boldsymbol{h}}\equiv{}^\star\hat{\boldsymbol{h}}[\alpha]$ must be symmetric functions on $\mathscr{M}^d$ and $\mathscr{M}^2$, respectively, that is,

$$^\star\dot{\boldsymbol{h}}(\underline{\dot\sigma})=\dot{\boldsymbol{h}}(\sigma_{\pi(1)},\dots,\sigma_{\pi(d)})$$

for any permutation π of $[d]$, and similarly $^\star\hat{\boldsymbol{h}}(\sigma,\mathsf{R}\sigma)={}^\star\hat{\boldsymbol{h}}(\mathsf{R}\sigma,\sigma)$, where R is the reflection map from (32). Thus we have $\bar h=\hat{\boldsymbol{h}}(\sigma,\mathsf{R}\sigma)$, so we can simplify $\boldsymbol{\Phi}^\lambda$ as

$$\boldsymbol{\Phi}^\lambda(\boldsymbol{h})=\sum_{\underline{\dot\sigma}}\dot{\boldsymbol{h}}(\underline{\dot\sigma})\log\frac{\dot\varphi^\lambda(\underline{\dot\sigma})}{\dot{\boldsymbol{h}}(\underline{\dot\sigma})}-\frac{d}{2}\sum_\sigma\bar h(\sigma)\log\frac{1}{\bar h(\sigma)\hat\varphi^\lambda(\sigma,\mathsf{R}\sigma)}. \tag{38}$$

In the unweighted setting $\lambda=1$, we simply have

$$\boldsymbol{\Phi}(\boldsymbol{h})=\sum_{\underline{\dot\sigma}}\dot{\boldsymbol{h}}(\underline{\dot\sigma})\log\frac{1}{\dot{\boldsymbol{h}}(\underline{\dot\sigma})}-\frac{d}{2}\sum_\sigma\bar h(\sigma)\log\frac{1}{\bar h(\sigma)}=H(\dot{\boldsymbol{h}})-\frac{d}{2}H(\bar h),$$

under the assumption that $\operatorname{supp}\dot{\boldsymbol{h}}\subseteq\operatorname{supp}\dot\varphi$.

The *belief propagation* or *Bethe recursion* for the λ-weighted auxiliary model acts on probability measures $\dot h$ and $\hat h$ on $\mathscr{M}$, with $\dot h$ mapping to $\hat h$ and vice versa:

$$\underbrace{\dot z\dot h_\sigma=\sum_{\underline{\dot\sigma}:\sigma_1=\sigma}\dot\varphi^\lambda(\underline{\dot\sigma})\prod_{i=2}^d\hat h_{\sigma_i},}_{\substack{\text{variable Bethe recursions,}\\ \text{with normalizing constant } \dot z}}\quad\underbrace{\hat z\hat h_\sigma=\sum_{\underline{\hat\sigma}:\sigma_1=\sigma}\hat\varphi^\lambda(\underline{\hat\sigma})\dot h_{\sigma_2}.}_{\substack{\text{clause Bethe recursions,}\\ \text{with normalizing constant } \hat z}} \tag{39}$$

LEMMA 3.9. *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 (39) via*

$$\dot{\boldsymbol{z}}\dot{\boldsymbol{h}}(\underline{\dot{\sigma}}) = \dot{\varphi}^\lambda(\underline{\dot{\sigma}}) \prod_{i=1}^{d} \hat{h}_{\sigma_i}, \quad \hat{\boldsymbol{z}}\hat{\boldsymbol{h}}(\underline{\hat{\sigma}}) = \hat{\varphi}^\lambda(\underline{\hat{\sigma}}) \prod_{i=1}^{2} \dot{h}_{\sigma_i} \quad \text{and} \quad \bar{z}\bar{h}(\sigma) = \hat{h}_\sigma \dot{h}_\sigma, \tag{40}$$

with $\dot{\boldsymbol{z}}$, $\hat{\boldsymbol{z}}$ and $\bar{z}$ normalizing constants satisfying $\bar{z} = \dot{\boldsymbol{z}}/\dot{z} = \hat{\boldsymbol{z}}/\hat{z}$ for $\dot{z}$ and $\hat{z}$ as in (39).

Proof. First consider optimizing (40) over $\dot{\boldsymbol{h}}$, subject to fixed marginals $\bar{h}$, but we require $\bar{h}$ to be a feasible edge marginal, meaning we have $\dot{H}\dot{\boldsymbol{h}} = \bar{h} = \widehat{H}\hat{\boldsymbol{h}}$ for some $\boldsymbol{h}$. This imposes linear constraints on $\bar{h}$: in this particular setting we must have

$$\bar{h}(\sigma) = \bar{h}(\mathsf{R}\sigma) \quad \text{and} \quad (d-1)\bar{h}(\mathtt{11}) = \bar{h}(\mathtt{f0}) + \bar{h}(\mathtt{ff}). \tag{41}$$

Let M be any minimal subset of $\mathscr{M}$ satisfying the property that if $\bar{h}(\sigma)$ is given for all $\sigma \in M$, the remaining values $\bar{h}(\sigma')$, $\sigma' \in \mathscr{M} \setminus M$, are determined using (41). Let $\dot{H}_\circ$ denote the submatrix of $\dot{H}$ formed by the rows with indices in M, and note that, by definition, the matrix $\dot{H}_\circ$ is surjective.

Now consider differentiating (38) in the direction of some signed measure $\boldsymbol{\delta} \equiv (\dot{\boldsymbol{h}}, \hat{\boldsymbol{h}})$ such that $\boldsymbol{h} + t\boldsymbol{\delta} \in \boldsymbol{\Delta}$ for small enough t: this means that we must have

$$\langle \bar{\delta}, 1 \rangle = 0 \quad \text{and} \quad \bar{\delta}(\sigma) = \bar{\delta}(\mathsf{R}\sigma). \tag{42}$$

Since $\boldsymbol{h}$ is stationary, we must have

$$0 = \partial_t \boldsymbol{\Phi}^\lambda(\boldsymbol{h} + t\boldsymbol{\delta}) = \sum_{\underline{\dot{\sigma}}} \dot{\boldsymbol{\delta}}(\underline{\dot{\sigma}}) \log \frac{\dot{\varphi}^\lambda(\underline{\dot{\sigma}})}{\dot{\boldsymbol{h}}(\underline{\dot{\sigma}})} + \frac{d}{2} \sum_{\sigma} \bar{\delta}(\sigma) \log[\bar{h}(\sigma)\hat{\varphi}^\lambda(\sigma, \mathsf{R}\sigma)]. \tag{43}$$

First choose $\boldsymbol{\delta}$ such that $\bar{\delta} = 0$—equivalently, such that $\dot{H}_\circ \bar{\delta} = 0$. Then (43) simplifies to

$$0 = \sum_{\underline{\dot{\sigma}}} \dot{\boldsymbol{\delta}}(\underline{\dot{\sigma}}) \dot{\boldsymbol{a}}(\underline{\dot{\sigma}}), \quad \text{where } \dot{\boldsymbol{a}}(\underline{\dot{\sigma}}) \equiv \log \frac{\dot{\varphi}^\lambda(\underline{\dot{\sigma}})}{\dot{\boldsymbol{h}}(\underline{\dot{\sigma}})}.$$

Since $\bar{\delta} = 0$, we furthermore have for any vector $\mu \in \mathbb{R}^M$ that

$$0 = \sum_{\underline{\dot{\sigma}}} \dot{\boldsymbol{\delta}}(\underline{\dot{\sigma}}) \dot{\boldsymbol{\varepsilon}}(\underline{\dot{\sigma}}), \quad \text{where } \dot{\boldsymbol{\varepsilon}}(\underline{\dot{\sigma}}) \equiv \log \frac{\dot{\varphi}^\lambda(\underline{\dot{\sigma}})}{\dot{\boldsymbol{h}}(\underline{\dot{\sigma}})} + \sum_{i=1}^{d} \mu_{\sigma_i} \mathbf{1}\{\sigma_i \in M\}.$$

Noting that $\dot{\boldsymbol{\varepsilon}} = a + (\dot{H}_\circ)^t \mu$, we now solve for $\dot{\boldsymbol{\varepsilon}}$ to have zero marginals:

$$0 = \dot{H}_\circ \dot{\boldsymbol{\varepsilon}} = \dot{H}_\circ a + \dot{H}_\circ (\dot{H}_\circ)^t \mu.$$

Since $\dot{H}_\circ$ is surjective, this has a unique solution μ. If we set $\dot{\boldsymbol{\delta}}=\dot{\boldsymbol{\varepsilon}}$ for this value of μ and substitute back into (43), we find that zero equals a sum of squares, so the only possibility is to have $\dot{\boldsymbol{\varepsilon}}$ identically zero. This proves that we can write

$$\dot{\boldsymbol{z}}\dot{\boldsymbol{h}}(\underline{\dot{\sigma}})=\dot{\varphi}^\lambda(\underline{\dot{\sigma}})\prod_{i=1}^d \hat{h}_{\sigma_i} \tag{44}$$

for some probability measure $\hat{h}$ on $\mathscr{M}$. Taking the marginal of (44) gives

$$\dot{\boldsymbol{z}}\bar{h}(\sigma)=\hat{h}_\sigma\sum_{\underline{\dot{\sigma}}}\mathbf{1}\{\sigma_1=\sigma\}\dot{\varphi}^\lambda(\underline{\dot{\sigma}})\prod_{i=2}^d \hat{h}_{\sigma_i}=\dot{z}\dot{h}_\sigma\hat{h}_\sigma. \tag{45}$$

We now return to the derivative (43) for general $\boldsymbol{\delta}$. Substituting in (44) and (45) gives

$$0=\sum_{\underline{\dot{\sigma}}}\dot{\boldsymbol{\delta}}(\underline{\dot{\sigma}})\log\frac{\dot{\boldsymbol{z}}}{\prod_{i=1}^d \hat{h}_{\sigma_i}}+\frac{d}{2}\sum_\sigma\bar{\delta}(\sigma)\log\frac{\hat{\varphi}(\sigma,\mathsf{R}\sigma)\dot{z}\dot{h}_\sigma\hat{h}_\sigma}{\dot{\boldsymbol{z}}}=\sum_\sigma\bar{\delta}(\sigma)\log\frac{\hat{\varphi}^\lambda(\sigma,\mathsf{R}\sigma)\dot{h}_\sigma}{\hat{h}_\sigma},$$

using $\langle\bar{\delta},1\rangle=0$ to eliminate $\dot{\boldsymbol{z}}$ and $\dot{z}$. As $\bar{h}$ is invariant under R, we see using (45) that

$$\frac{\dot{h}_\sigma}{\hat{h}_\sigma}=\frac{\dot{h}_\sigma\hat{h}_\sigma}{\hat{h}_\sigma\hat{h}_\sigma}=\frac{\dot{h}_{\mathsf{R}\sigma}\hat{h}_{\mathsf{R}\sigma}}{\hat{h}_\sigma\hat{h}_\sigma}.$$

Then, since $\bar{\delta}$ is also invariant under R, we can re-express the stationary equation as

$$0=\sum_\sigma\bar{\delta}(\sigma)\log\frac{\hat{\varphi}^\lambda(\sigma,\mathsf{R}\sigma)\dot{h}_{\mathsf{R}\sigma}\hat{h}_{\mathsf{R}\sigma}}{\hat{h}_\sigma\hat{h}_{\mathsf{R}\sigma}}=\sum_\sigma\bar{\delta}(\sigma)\underbrace{\log\frac{\hat{\varphi}^\lambda(\sigma,\mathsf{R}\sigma)\dot{h}_{\mathsf{R}\sigma}}{\hat{h}_\sigma}}_{\boldsymbol{\varkappa}(\sigma)},$$

where $\boldsymbol{\varkappa}$ is also invariant under R. Again, using $\langle\bar{\delta},1\rangle=0$, we have

$$0=\sum_\sigma\bar{\delta}(\sigma)\widetilde{\boldsymbol{\varkappa}}(\sigma)\quad\text{where }\ \widetilde{\boldsymbol{\varkappa}}(\sigma)\equiv\boldsymbol{\varkappa}(\sigma)-\frac{1}{|\mathscr{M}|}\sum_{\sigma'}\boldsymbol{\varkappa}(\sigma').$$

Note that $\bar{\delta}=\widetilde{\boldsymbol{\varkappa}}$ satisfies the requirements of (42), and substituting into the above we see that the only possibility is to have $\widetilde{\boldsymbol{\varkappa}}(\sigma)=0$ for all $\sigma\in\mathscr{M}$. This proves that

$$\hat{z}\hat{h}_\sigma=\dot{h}_{\mathsf{R}\sigma}\hat{\varphi}^\lambda(\sigma,\mathsf{R}\sigma). \tag{46}$$

Now (45) and (46) together constitute the belief propagation equations. Note also that combining (45) and (46) gives

$$\hat{\boldsymbol{h}}(\sigma,\mathsf{R}\sigma)=\bar{h}(\sigma)=\frac{\dot{z}\dot{h}_\sigma\hat{h}_\sigma}{\dot{\boldsymbol{z}}}=\frac{\dot{z}\dot{h}_\sigma\dot{h}_{\mathsf{R}\sigma}\hat{\varphi}^\lambda(\sigma,\mathsf{R}\sigma)}{\dot{\boldsymbol{z}}\hat{z}}.$$

Together with (44) and (45), this proves (40) with $\bar{z}=\dot{\boldsymbol{z}}/\dot{z}=\hat{\boldsymbol{z}}/\hat{z}$. ☐

While the optimization (37) is over the space $\mathbf{\Delta}[\alpha]$ whose dimension grows with d, the Bethe recursions act on probability measures $\dot{h}$ and $\hat{h}$ over $\mathscr{M}$ whose dimension clearly does not depend on d. (Boundary maximizers will be ruled out by a-priori estimates in §4.2.) Even with this reduction, however, we are tasked with eliminating a potential multiplicity of fixed points (local maximizers). In the physics literature this has been resolved by prescribing that $\hat{h}_{\eta'\eta}$ depends only on the clause-to-variable message η, justified by appealing to some notion of "causality" (cf. [36, equation (19.26)]). We prove a rigorous version here via the interpretation for Bethe recursion fixed points in terms of (infinite-volume) Gibbs measures on the d-regular tree T_d, which we now describe.

We have commented already that the local structure of the random d-regular graph $\mathscr{G}_{d,n}$ is that of T_d. More precisely, it is well known that $\mathscr{G}_{d,n}$ converges locally to T_d in the sense of [12] and [8] (see e.g. [19]). Bisect each edge of T_d by a new clause vertex (as done above for $\mathscr{G}_{d,n}$), and refer to the resulting graph as the *$(d,2)$-regular tree* $T\equiv T_{d,2}$. Thus the leaf vertices of the depth-t subtree $T(t)$ are variables for t even, clauses for t odd; we write $\boldsymbol{E}(t-1,t)$ for the half-edges joining levels $t-1$ and t. We let $\underline{\sigma}$ denote a message configuration on T, with $\mathscr{M}$-valued spins located on the edges of the tree.

Given a finite subtree U of T, let U° denote the interior vertices of U (that is, the variables and clauses having no neighbors outside U). Let δU denote the internal edge boundary of U, that is, the edges (xy) where $x\in U^\circ$ and $y\in U\setminus U^\circ$. Let $\delta_V U\subseteq\delta U$ denote the subset of such (xy) where $x\in V$, and let $\delta_F U$ denote the rest, with $x\in F$. Given probability measures $\dot{h}$ and $\hat{h}$ on $\mathscr{M}$ we define the measures

$$Z_U^\lambda \boldsymbol{\nu}_U^\lambda(\underline{\sigma}_U)=\prod_{v\in V\cap U^\circ}\dot{\varphi}(\underline{\dot{\sigma}}_v)\prod_{a\in F\cap U^\circ}\hat{\varphi}(\underline{\hat{\sigma}}_a)\prod_{e\in\delta_V U}\dot{h}_{\sigma_e}\prod_{e\in\delta_F U}\hat{h}_{\sigma_e}, \tag{47}$$

with Z_U^λ the normalizing constant that makes $\boldsymbol{\nu}_U^\lambda$ a probability measure over valid message configurations $\underline{\sigma}$ on U. Then $h\equiv(\dot{h},\hat{h})$ satisfies the Bethe recursion (39) if and only if $(\boldsymbol{\nu}_U^\lambda)_U$ is a consistent family of finite-dimensional distributions, meaning that $\boldsymbol{\nu}_S^\lambda$ is a projection of $\boldsymbol{\nu}_U^\lambda$ for any $S\subseteq U$. It follows by the Kolmogorov extension theorem that the collection $(\boldsymbol{\nu}_U^\lambda)_U$ defines an infinite-volume Gibbs measure $\boldsymbol{\nu}^\lambda$ on T with finite-dimensional marginals $\boldsymbol{\nu}_U^\lambda$.

3.3. Bethe recursion symmetries

Observe that on any subtree U of T, the entire message configuration $\underline{\sigma}_U$ on U is completely determined by the *incoming* boundary messages $\sigma_{y\to x}$ from $y\in U\setminus U^\circ$ for $x\in U^\circ$—given the incoming messages, $\underline{\sigma}_U$ can be recovered by iterating $\dot{\mathtt{m}}$ from the boundary inwards. Thus a natural special case of (47) is to take $\hat{h}_{\eta'\eta}$ to depend only on the clause-to-variable message η, that is, $3\,\hat{h}_{\eta'\eta}=q_\eta$ for a probability measure q on $\{\mathtt{0},\mathtt{1},\mathtt{f}\}$. Then

$$\boldsymbol{\nu}^\lambda_1(\eta_o=\mathtt{1})=\frac{\lambda(1-q_{\mathtt{1}})^d}{(\lambda-1)(1-q_{\mathtt{1}})^d+1}\quad\text{and}\quad \boldsymbol{\nu}^\lambda_1(\eta_o=\mathtt{f})=\frac{dq_{\mathtt{1}}(1-q_{\mathtt{1}})^{d-2}}{(\lambda-1)(1-q_{\mathtt{1}})^d+1},$$

with the remaining probability going to $\eta_o=\mathtt{0}$. The $(\boldsymbol{\nu}^\lambda_U)_U$ are consistent if and only if q satisfies the *frozen model recursions*

$$q_{\mathtt{1}}=\frac{\lambda(1-q_{\mathtt{1}})^{d-1}}{(\lambda-1)(1-q_{\mathtt{1}})^{d-1}+1},\quad q_{\mathtt{f}}=\frac{(d-1)q_{\mathtt{1}}(1-q_{\mathtt{1}})^{d-2}}{(\lambda-1)(1-q_{\mathtt{1}})^{d-1}+1}=\frac{(d-1)(q_{\mathtt{1}})^2}{\lambda(1-q_{\mathtt{1}})},\tag{48}$$

and $q_{\mathtt{0}}=1-q_{\mathtt{1}}-q_{\mathtt{f}}$.[8]

Since a solution of (48) is fully determined by $q_{\mathtt{1}}$, we hereafter abuse notation and use q to denote both the measure $(q_{\mathtt{1}},q_{\mathtt{f}},q_{\mathtt{0}})$ and the value $q_{\mathtt{1}}$. The value $q=q_{\mathtt{1}}$ must be a root 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 f has a unique root $0<q<1$ which is increasing in λ, and (2) expresses λ in terms of q. In the following, to emphasize the dependence on λ we sometimes write $q\equiv q^\lambda$.

The full Bethe recursions (39)—a generalization of (48)—read explicitly as follows. The clause Bethe recursions are simply

$$\hat{z}\hat{h}_\sigma=\dot{h}_{\mathsf{R}\sigma}\lambda^{\mathbf{1}\{\sigma=\mathtt{11}\}},\quad\text{with }\ \hat{z}=1+(\lambda-1)\dot{h}_{\mathtt{11}}.\tag{49}$$

Write $\mathtt{Z}\equiv\{\mathtt{0},\mathtt{f}\}$, for example $\hat{h}_{\mathtt{1Z}}\equiv\hat{h}_{\mathtt{10}}+\hat{h}_{\mathtt{1f}}$. Recall also from §2.4 that we use $\_$ to denote $\{\mathtt{0},\mathtt{1},\mathtt{f}\}$. The variable Bethe recursions are

$$\begin{aligned}
&(\text{A})\quad \dot{z}\dot{h}_{\mathtt{10}}=\dot{z}\dot{h}_{\mathtt{1f}}=\lambda(\hat{h}_{\mathtt{1Z}})^{d-1},\qquad \lambda\dot{z}\dot{h}_{\mathtt{11}}=\lambda(\hat{h}_{\mathtt{fZ}})^{d-1},\\
&\phantom{(\text{A})}\quad \dot{z}\dot{h}_{\mathtt{f0}}=\dot{z}\dot{h}_{\mathtt{ff}}=(d-1)\hat{h}_{\mathtt{11}}(\hat{h}_{\mathtt{fZ}})^{d-2},\qquad \dot{z}\dot{h}_{\mathtt{f1}}=(d-1)\hat{h}_{\mathtt{f1}}(\hat{h}_{\mathtt{0Z}})^{d-2},\\
&(\text{B})\quad \dot{z}\dot{h}_{\mathtt{01}}=(\hat{h}_{\mathtt{0}_})^{d-1}-(\hat{h}_{\mathtt{0Z}})^{d-1}-(d-1)\hat{h}_{\mathtt{01}}(\hat{h}_{\mathtt{0Z}})^{d-2},\\
&(\text{C})\quad \dot{z}\dot{h}_{\mathtt{00}}=\dot{z}\dot{h}_{\mathtt{0f}}=(\hat{h}_{\mathtt{0}_})^{d-1}-(\hat{h}_{\mathtt{0Z}})^{d-1}-(d-1)\hat{h}_{\mathtt{01}}(\hat{h}_{\mathtt{0Z}})^{d-2}\\
&\qquad\qquad\qquad\qquad +\tfrac{1}{2}(d-1)(d-2)[(\hat{h}_{\mathtt{f1}})^2-(\hat{h}_{\mathtt{01}})^2](\hat{h}_{\mathtt{0Z}})^{d-3}.
\end{aligned}$$

We immediately have $\hat{h}_{\mathtt{01}}=\hat{h}_{\mathtt{f1}}$, $\hat{h}_{\mathtt{0f}}=\hat{h}_{\mathtt{ff}}$, and $\hat{h}_{\mathtt{00}}=\hat{h}_{\mathtt{f0}}$. Comparing (B) and (C) then 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):

$$\dot{h}_{\mathtt{f1}}=\dot{h}_{\mathtt{f0}}\overset{\checkmark}{=}\dot{h}_{\mathtt{ff}},\quad \hat{h}_{\mathtt{1f}}=\hat{h}_{\mathtt{0f}}\overset{\checkmark}{=}\hat{h}_{\mathtt{ff}},\quad \lambda\dot{h}_{\mathtt{11}}=\dot{h}_{\mathtt{10}}\overset{\checkmark}{=}\dot{h}_{\mathtt{1f}},\quad \hat{h}_{\mathtt{11}}=\hat{h}_{\mathtt{01}}\overset{\checkmark}{=}\hat{h}_{\mathtt{f1}}.\tag{50}$$

---

[8] For comparison, the corresponding recursion for the hard-core model at fugacity λ is

$$q=\frac{\lambda(1-q)^{d-1}}{\lambda(1-q)^{d-1}+1}.$$

If (50) holds, then (39) reduces to the frozen model recursions (48) with

$$3\hat{h}_{\eta'\eta}=q_\eta,$$

and therefore, substituting into (49),

$$3\dot{h}_{\eta\eta'}=\lambda^{-\mathbf{1}\{\eta\eta'=11\}}q_\eta\hat{z} \quad \text{with } \hat{z}=\left[1-\frac{q_1}{3}\left(1-\frac{1}{\lambda}\right)\right]^{-1}. \tag{51}$$

Write $\mathbf{\Delta}^\circ[\alpha]\equiv\mathbf{\Delta}^\circ\cap\mathbf{\Delta}[\alpha]$, and suppose $\boldsymbol{h}\in\mathbf{\Delta}^\circ[\alpha]$ is a stationary point for the restriction of $\mathbf{\Phi}$ to $\mathbf{\Delta}[\alpha]$. By the method of Lagrange multipliers, there exists λ (unique given $\boldsymbol{h}$) such that $\boldsymbol{h}$ is an interior stationary point for $\mathbf{\Phi}^\lambda$ on the larger space $\mathbf{\Delta}$: therefore $\boldsymbol{h}$ corresponds, via (40), to a solution h of the Bethe recursions (39) for the λ-weighted model.

Proposition 3.10. *For $d\geqslant d_0$ and $\alpha_{\mathrm{lbd}}\leqslant\alpha\leqslant\alpha_{\mathrm{ubd}}$, let h be any solution of the Bethe recursions (39) which corresponds in the manner described above to an interior stationary point $\boldsymbol{h}$ for the restriction of $\mathbf{\Phi}$ to $\mathbf{\Delta}[\alpha]$. Then h satisfies the symmetries (50), and corresponds via (51) to a solution $q=q_1$ of (48) with $q=\alpha+O(d^{-1})$.*

Ruling out boundary maximizers for $\mathbf{\Phi}$ on $\mathbf{\Delta}[\alpha]$ 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 $\mathbf{\Phi}_2$. We therefore conclude the following.

Theorem 3.11. *For $d\geqslant d_0$ and $\alpha_{\mathrm{lbd}}\leqslant\alpha\leqslant\alpha_{\mathrm{ubd}}$, the restriction of $\mathbf{\Phi}$ to $\mathbf{\Delta}[\alpha]$ has a unique global maximizer ${}^\star\boldsymbol{h}[\alpha]$, which is also the unique interior stationary point. It correponds, via (40) and (51), to the unique $q\in[1.6(\log d)/d, 3(\log d)/d]$ for which $\alpha(q)$, as defined in (2), evaluates to α.*

Proof. We assume that any global maximizer $\boldsymbol{h}$ of $\mathbf{\Phi}$ on $\mathbf{\Delta}[\alpha]$ lies in the interior $\mathbf{\Delta}^\circ[\alpha]$, deferring the proof to §4 (see Proposition 4.13 and Corollary 4.17). Therefore, as described above, $\boldsymbol{h}$ corresponds via (40) to a solution h of the Bethe recursions (39) for a particular λ. Proposition 3.10 implies that h corresponds via (51) to a solution $q=q_1$ of the frozen model recursions (48) for this value of λ. Rearranging (48) shows that λ must equal the expression $\lambda(q)$ defined by (2). In §3.4 below we shall calculate (from (40) and (51)) that $\mathbf{i}(\boldsymbol{h})$ equals the expression $\alpha(q)$ of (2). The mapping $q\mapsto\alpha(q)$ is not one-to-one on the entire interval $0\leqslant q\leqslant 1$, but if $q=x(\log d)/d$ with $1.6\leqslant x\leqslant 3$ then we can easily compute the derivative $\alpha'(q)=1+O(d^{-1/2})$. Since Proposition 3.10 also gave the estimate $q=\alpha+O(d^{-1})$, this clearly uniquely identifies q. Thus we have located the unique global maximizer of $\mathbf{\Phi}$ on $\mathbf{\Delta}[\alpha]$ as the unique interior stationary point. □

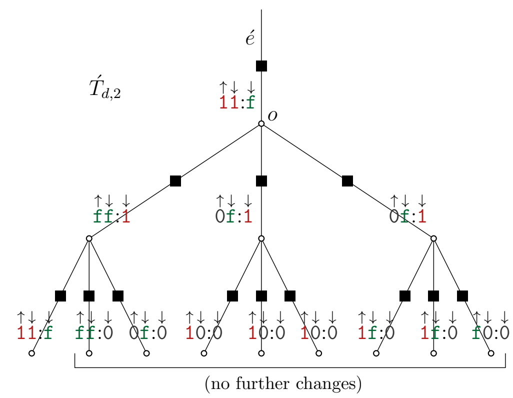

Figure 1. Change of message incoming down to $\acute{e}$ is passed down $\acute{T}$ ($\eta\acute{\eta}$:$\grave{\eta}$ means message η up, message $\acute{\eta}$ down in $\underline{\sigma}$, message $\grave{\eta}$ down in $\underline{\sigma}'$).

It remains to prove Proposition 3.10. Recall that $T \equiv T_{d,2}$ denotes the $(d,2)$ -regular tree. Let $\acute{T}$ denote T with the subtree descendant from one of the root neighbors removed, so that $\acute{T}$ has one clause a which is joined to the root variable o and to an unmatched half-edge $\acute{e}$ (Figure 1). Define a Gibbs measure $\acute{\boldsymbol{\nu}}^{\lambda}$ on $\acute{T}$ in the manner of (47), via its finite-dimensional marginals $\acute{\boldsymbol{\nu}}^{\lambda}_{U}$, with factor weights φ^{λ} and boundary law given by h. If h solves the Bethe recursions (39), then the resulting marginal law on $\sigma_{\acute{e}}$ will be simply $\hat{h}$, and the marginal law of the d-tuple of spins incident to any given vertex will be $\dot{\boldsymbol{h}}$.

Write $\sigma_{\acute{e}} \equiv \boldsymbol{io}$ where $\boldsymbol{i}$ is the incoming (variable-to-clause) message and $\boldsymbol{o}$ is the outgoing (in Figure 1, $\boldsymbol{o}$ is directed upwards, $\boldsymbol{i}$ downwards). Given any message configuration $\underline{\sigma}$ on $\acute{T}$, changing $\boldsymbol{i}$ and passing the changed message through the tree (via $\dot{\mathtt{m}}$) produces a new configuration $\underline{\sigma}'$ (Figure 1). Proposition 3.10 will follow by showing that for any fixed $\boldsymbol{o}$, the effect of changing $\boldsymbol{i}$ is *measure-preserving* under the Gibbs measure $\acute{\boldsymbol{\nu}}^{\lambda}$. Since the finite-dimensional marginals of $\acute{\boldsymbol{\nu}}^{\lambda}$ are defined (in the manner of (47)) from 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 $\hat{h}_{\mathtt{0}\eta} = \hat{h}_{\mathtt{f}\eta}$ for any η: this corresponds to the fact that $\dot{\mathtt{m}}$ does not differentiate between incoming messages $\mathtt{0}$ and $\mathtt{f}$, so changing $\boldsymbol{i}$ from $\mathtt{0}$ to $\mathtt{f}$ or vice versa has no effect at all below $\acute{e}$. We also found that $\hat{h}_{\eta\mathtt{0}}$ does not depend on η: if $\boldsymbol{o} = \mathtt{0}$ then all messages outgoing from the root must be $\mathtt{0}$ or $\mathtt{f}$, so changing $\boldsymbol{i}$ can have an effect at most one level down.

Proof of Proposition 3.10. To prove (50) it remains to verify that $\hat{h}_{\mathtt{11}} = \hat{h}_{\mathtt{f1}}$. Changing $\boldsymbol{io}$ from $\mathtt{11}$ to $\mathtt{f1}$ induces a mapping ι on messages configurations $\underline{\sigma}$ of $\acute{T}$; likewise changing $\boldsymbol{io}$ from $\mathtt{f1}$ to $\mathtt{11}$ induces a mapping $\bar{\iota}$. From the above discussion, we must prove that the maps ι and $\bar{\iota}$ have a finite-range effect.

Figure 1 shows that the effect of changing 11 to f1 can only propagate through components of f-variables, while the effect of changing f1 to 11 can only propagate through chains of alternating 1-variables and $0^{\mathtt{s}}$-variables (with $0^{\mathtt{s}}$ as defined in (32)).

We claim that both propagations are subcritical under $\acute{\boldsymbol{\nu}}^\lambda$. Observe that a sample from $\acute{\boldsymbol{\nu}}^\lambda$ can be generated in Markovian fashion: start with spin $\sigma_{\acute{e}}$ distributed according to $\hat{h}$, generate the messages on the other $d-1$ half-edges incident to o according to the conditional measure $\dot{\boldsymbol{h}}(\underline{\dot{\sigma}}\,|\,\sigma_1=\sigma_{\acute{e}})$, and continue iteratively down the tree. It follows from (40) that if we condition on $\sigma_{\acute{e}}=\mathtt{f1}$, the expected number of children with spin f1 will be $(d-1)\hat{h}_{\mathtt{1f}}/\hat{h}_{\mathtt{1Z}}$; likewise, if we condition on $\sigma_{\acute{e}}=\mathtt{11}$, the expected number of children with spin ff will be $(d-1)\hat{h}_{\mathtt{ff}}/\hat{h}_{\mathtt{fZ}}=(d-1)\hat{h}_{\mathtt{0f}}/\hat{h}_{\mathtt{0Z}}$. We now show that both these quantities are less than one, meaning that with probability 1 (under $\acute{\boldsymbol{\nu}}^\lambda$) the changes do not percolate through the tree. This implies the claimed result: writing $\acute{\boldsymbol{\nu}}^\lambda_t$ for the marginal of $\acute{\boldsymbol{\nu}}^\lambda$ on the depth-t subtree $\acute{T}(t)$ of $\acute{T}$, we have

$$\hat{h}_{\mathtt{11}}=\acute{\boldsymbol{\nu}}^\lambda(\sigma_{\acute{e}}=\mathtt{11})=\sum_{\underline{\sigma}}\mathbf{1}\{\sigma_{\acute{e}}=\mathtt{11}\}\acute{\boldsymbol{\nu}}^\lambda_t(\underline{\sigma})$$

where the sum is taken over message configurations $\underline{\sigma}$ on $\acute{T}(t)$. From the above description of how ι propagates, and from the definition of $\acute{\boldsymbol{\nu}}^\lambda_t$, we have $\acute{\boldsymbol{\nu}}^\lambda_t(\underline{\sigma})=\acute{\boldsymbol{\nu}}^\lambda_t(\iota\underline{\sigma})$ as long as $\underline{\sigma}$ and $\iota\underline{\sigma}$ agree at depth t, which holds with probability $1-o_t(1)$ under $\acute{\boldsymbol{\nu}}^\lambda_t$. Therefore

$$\hat{h}_{\mathtt{11}}=o_t(1)+\sum_{\underline{\sigma}}\mathbf{1}\{\sigma_{\acute{e}}=\mathtt{11}\}\acute{\boldsymbol{\nu}}^\lambda_t(\iota\underline{\sigma})\leqslant o_t(1)+\hat{h}_{\mathtt{1f}},$$

and taking $t\to\infty$ proves $\hat{h}_{\mathtt{11}}\leqslant\hat{h}_{\mathtt{1f}}$. Repeating the same argument with $\bar{\iota}$ in place of ι proves the reverse inequality $\hat{h}_{\mathtt{1f}}\leqslant\hat{h}_{\mathtt{11}}$, concluding the proof.

It remains to verify the subcriticality of the propagations. Recall $V_\eta\equiv\{v\in V:\eta_v=\eta\}$. For $S,S'\subseteq V$ write $\mathtt{E}(S,S')$ for the number of clauses joining S to S'. The number $\mathtt{E}(V_{\mathtt{0}},V_{\mathtt{0}})$ of edges internal to $V_{\mathtt{0}}$ is $n\left[\frac{1}{2}d-O(\log d)\right]$. By (39), (40), and the trivial symmetry $\hat{h}_{\mathtt{0}\eta}=\hat{h}_{\mathtt{f}\eta}$ already noted,

$$\frac{\hat{h}_{\mathtt{0f}}}{\hat{h}_{\mathtt{00}}}=\frac{\dot{h}_{\mathtt{f0}}}{\dot{h}_{\mathtt{00}}}=\frac{\dot{h}_{\mathtt{f0}}\dot{h}_{\mathtt{0f}}}{\dot{h}_{\mathtt{00}}\dot{h}_{\mathtt{00}}}=\frac{\hat{\boldsymbol{h}}(\mathtt{f0},\mathtt{0f})}{\hat{\boldsymbol{h}}(\mathtt{00},\mathtt{00})}=\frac{\mathtt{E}(V_{\mathtt{f}},V_{\mathtt{0}})}{2\,\mathtt{E}(V_{\mathtt{0}},V_{\mathtt{0}})}\leqslant\frac{nd\beta_{\max}}{n[d-O(\log d)]}\leqslant 2\beta_{\max}.$$

Next, since $\partial V_{\mathtt{1}}\subseteq V_{\mathtt{0}}$, $\mathtt{E}(V_{\mathtt{1}},V_{\mathtt{0}})=\mathtt{E}(V_{\mathtt{1}},V)=d|V_{\mathtt{1}}|=nd(\alpha-O(\beta_{\max}))$. Also, since each vertex in $V_{\mathtt{0}^{\mathtt{s}}}$ can have at most two neighbors in $V_{\mathtt{1}}$, crudely $\mathtt{E}(V_{\mathtt{1}},V_{\mathtt{0}^{\mathtt{s}}})\leqslant 2|V_{\mathtt{0}^{\mathtt{s}}}|\leqslant 2n$. Therefore

$$\frac{\hat{h}_{\mathtt{01}}}{\hat{h}_{\mathtt{00}}}=\frac{\hat{\boldsymbol{h}}(\mathtt{10},\mathtt{01})}{\hat{\boldsymbol{h}}(\mathtt{00},\mathtt{00})}=\frac{\mathtt{E}(V_{\mathtt{1}},V_{\mathtt{0}^{\mathtt{s}}})}{2\,\mathtt{E}(V_{\mathtt{0}},V_{\mathtt{0}})}=\frac{nd(\alpha-O(\beta_{\max}))-O(n)}{n[d-O(\log d)]}=\alpha-O(d^{-1}). \tag{52}$$

Applying (39) again gives

$$\frac{\hat{h}_{\mathtt{1f}}}{\hat{h}_{\mathtt{10}}}=\frac{\dot{h}_{\mathtt{1f}}}{\dot{h}_{\mathtt{10}}}=\frac{(d-1)\hat{h}_{\mathtt{01}}(\hat{h}_{\mathtt{0Z}})^{d-1}}{\sum_{j\geqslant 2}(\hat{h}_{\mathtt{01}})^j(\hat{h}_{\mathtt{0Z}})^{d-1-j}}=\mathbb{P}\left(\mathrm{Bin}\left(d-1,\frac{\hat{h}_{\mathtt{01}}}{\hat{h}_{\mathtt{0}\_}}\right)=1\right)\leqslant d^{-1.6}.$$

Therefore, writing $f \ll_d g$ to indicate $\lim_{d\to\infty} f/g=0$, we have shown $\hat{h}_{\mathtt{0f}}/\hat{h}_{\mathtt{00}} \ll_d d^{-1}$ and $\hat{h}_{\mathtt{1f}}/\hat{h}_{\mathtt{10}} \ll_d d^{-1}$, implying (50). Thus h corresponds via (51) to a solution q of (48), and we conclude a posteriori from the two preceding estimates that $q_1 = \hat{h}_{\mathtt{01}}/\hat{h}_{\mathtt{0}\_} = \alpha + O(d^{-1})$, concluding the proof. □

3.4. Explicit form of first moment exponent

We conclude this section by giving the explicit form of $\boldsymbol{\phi}(\alpha)$.

Proof of Theorem 3.1. The result follows from Theorem 3.11 together with a few straightforward calculations. Recall from the proof of Theorem 3.11 that if $\boldsymbol{h}$ is an interior global maximizer for $\boldsymbol{\Phi}$ on $\boldsymbol{\Delta}[\alpha]$, then, by Proposition 3.10, $\boldsymbol{h}$ corresponds to a solution q of the frozen model recursions for this value of λ, from which we conclude that $\lambda = \lambda(q)$ as in (2). Writing $q = x(\log d)/d$ we estimate

$$\lambda = \lambda(q) = d^x q[1 + O(d^{-1}(\log d)^2)].$$

Since $\boldsymbol{h} \in \boldsymbol{\Delta}[\alpha]$, we have

$$\alpha = \dot{\boldsymbol{h}}(\dot{\mathtt{m}}(\underline{\dot{\sigma}}) = \mathtt{1}) + \tfrac{1}{2}\dot{\boldsymbol{h}}(\dot{\mathtt{m}}(\underline{\dot{\sigma}}) = \mathtt{f}), \tag{53}$$

where the right-hand side can be explicitly calculated from the relations (40) and (51) for $\boldsymbol{h}$.

Explicit Bethe prediction. Substituting (40) into (35) and rearranging gives

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

We use (40) and (51) to calculate $\bar{z}$, $\dot{\boldsymbol{z}}$, and $\hat{\boldsymbol{z}}$ in terms of q and λ:

$$\bar{z} = \frac{\hat{z}}{9}\left[1 - q^2\left(1 - \frac{1}{\lambda}\right)\right], \quad \hat{\boldsymbol{z}} = \bar{z}\hat{z}, \quad \dot{\boldsymbol{z}} = \frac{1 - q^2(1 - 1/\lambda)}{3^d[1 - q(1 - 1/\lambda)]} = \bar{z}\dot{z}.$$

Recalling (53) we can also calculate α in terms of q and λ:

$$\alpha = q\frac{1 - q + dq/2\lambda}{1 - q^2(1 - 1/\lambda)},$$

completing the verification of (2). The recursion (48) also gives the expressions in (2) for λ and α solely in terms of q. It then follows from Theorem 3.11 that

$$\boldsymbol{\phi}(\alpha) = \boldsymbol{\Phi}^\lambda({}^\star\boldsymbol{h}[\alpha]) - \alpha\log\lambda$$

is given by (1). We note also that $\boldsymbol{\phi}'(\alpha)=-\log\lambda$, and therefore

$$\boldsymbol{\phi}''(\alpha)=-[\alpha'(q)]^{-1}[(\log\lambda)'(q)]=-d[1+O((\log d)^{-1})].$$

Comparison of first-moment exponents. By contrast, the original independent set partition function has first-moment exponent $\phi(\alpha)$ calculated in (7).([9]) This exponent also has a Bethe variational characterization, which can be expressed in terms of the fixed point q' of the hard-core tree recursions:

$$\phi(\alpha)=\log(1+q')-\tfrac{1}{2}d\log(1-(q')^2)-\alpha\log\lambda',$$

where

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

This formula can be derived heuristically in the same manner as (1); its correctness can be checked simply by verifying that it agrees with (7). We then compare $\boldsymbol{\phi}(\alpha)$ and $\phi(\alpha)$ by expressing both in terms of q: then q and q' are related via $\alpha(q)=\alpha=\alpha'(q')$, explicitly

$$q'=\frac{\alpha}{1-\alpha}=q\frac{1-q+dq/2\lambda}{1-q-(d-2)q^2/(2\lambda)}.$$

A little algebra then gives

$$\begin{aligned}(\phi-\boldsymbol{\phi})(\alpha)&=\log\frac{1-q+q/\lambda}{1-q'}+\left(\frac{d}{2}-1\right)\log\frac{1-q^2+q^2/\lambda}{1-(q')^2}-\alpha\log\left(\frac{\lambda'}{\lambda}\right)\\&=\log\left[\left(1+\frac{q/\lambda}{1-q}\right)\left(1-\frac{(d-2)q^2}{2\lambda(1-q)}\right)^{-1}\right]\\&\qquad+d\log\left[\left(1-\frac{(d-2)q^2}{2\lambda(1-q)}\right)\left(1-\frac{(d-1)q^2}{\lambda(1-q)^2}\right)^{-1/2}\right]-\alpha\log\left(\frac{\lambda'}{\lambda}\right).\end{aligned}$$

Recalling that $q=x(\log d)/d$ and $\lambda=d^xq[1+O(d^{-1}(\log d)^2)]$, we expand

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

Substituting into the above expression for $(\phi-\boldsymbol{\phi})(\alpha)$ and expanding the other terms gives

$$(\phi-\boldsymbol{\phi})(\alpha)=\frac{q(dq+2)}{2\lambda}+O\left(\frac{(d+\lambda)(dq)^2q^2}{\lambda^2}\right)\asymp d^{-x}\log d. \tag{54}$$

([9]) Recall that in the original independent set model α refers to the set density, while in the frozen model it refers to the intensity (9).

Comparison of first-moment thresholds. It is clear from the above that $\boldsymbol{\phi}$ has a unique zero $\alpha_{\mathrm{lbd}} < \alpha_\star < \alpha_\square$, with $\alpha_\square$ the first moment threshold of the original independent partition function. Let $q_\square$ and $\lambda_\square$ denote the solution of (2) for $\alpha = \alpha_\square$, and note that

$$\lambda_\square = \frac{e^2}{2}\frac{d}{\log d} \quad \text{and} \quad q_\square = \alpha_\square[1 + O(d^{-1}(\log d)^2)].$$

Consider $0 \leqslant \delta \lesssim (\log d)^2/d^2$: applying (54) with the above estimate of $\boldsymbol{\phi}''(\alpha)$ gives

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

so the gap between the thresholds is given by

$$\alpha_\square - \alpha_\star = [1 + O(d^{-1}(\log d)^3)]\frac{\alpha_\square(d\alpha_\square + 2)}{2\lambda_\square \log \lambda_\square} = \left(\frac{2}{e}\frac{\log d}{d}\right)^2\left[1 + O\left(\frac{\log\log d}{\log d}\right)\right],$$

concluding the proof. $\square$

4. Second moment of frozen model

In this section we compute the exponential growth rate $\boldsymbol{\phi}_2(\alpha) \equiv \lim_{n\to\infty} n^{-1} \log \mathbb{E}[\boldsymbol{Z}_{n\alpha}^2]$ of the second moment of the frozen model partition function (10). This will be done within the framework introduced in §3, regarding the second moment as the first moment of the *pair* model: given a graph $G = (V, E)$, a *pair frozen configuration* is a pair $(\underline{\bar{\boldsymbol{\eta}}}^1, \underline{\bar{\boldsymbol{\eta}}}^2)$ where each $\underline{\bar{\boldsymbol{\eta}}}^i \equiv (\underline{\eta}^i, \underline{m}^i)$ is a valid frozen configuration (Definition 2.3) on the same underlying graph. Thus $\mathbb{E}[\boldsymbol{Z}_{n\alpha}^2]$ is the first moment of pair frozen configurations at intensity $n\alpha$ (see (9)).

Recall from Proposition 3.3 that each frozen configuration $\underline{\bar{\boldsymbol{\eta}}}^i$ corresponds bijectively to a message configuration $\underline{\sigma}^i$. We say that $\underline{\tau} \equiv (\underline{\sigma}^1, \underline{\sigma}^2)$ is a *pair message configuration.* Following the discussion of §3.1, the pair frozen model can be recast as a Gibbs measure on pair message configurations, expressed as in (31), but with $\dot{\varphi}$ and $\hat{\varphi}$ replaced by the factors

$$\dot{\varphi}_2 \equiv \dot{\varphi} \otimes \dot{\varphi} \quad \text{and} \quad \hat{\varphi}_2 \equiv \hat{\varphi} \otimes \hat{\varphi},$$

respectively. We refer to this as the *pair auxiliary model.* Following §3.2, we decompose

$$\mathbb{E}[\boldsymbol{Z}_{n\alpha}^2] = \sum_{\boldsymbol{h}} \mathbb{E}[\boldsymbol{Z}^2(\boldsymbol{h})] = \sum_{\boldsymbol{h}} \exp(n\boldsymbol{\Phi}_2(\boldsymbol{h})), \tag{55}$$

where $\boldsymbol{h} \equiv (\dot{\boldsymbol{h}}, \hat{\boldsymbol{h}})$ now denotes a pair empirical measure:

$$\begin{aligned}&\dot{\boldsymbol{h}} \text{ is a measure on pairs } (\underline{\dot{\sigma}}^1, \underline{\dot{\sigma}}^2) \text{ with } \underline{\dot{\sigma}}^i \in \mathscr{M}^d, \\ &\hat{\boldsymbol{h}} \text{ is a measure on pairs } (\underline{\hat{\sigma}}^1, \underline{\hat{\sigma}}^2) \text{ with } \underline{\hat{\sigma}}^i \in \mathscr{M}^2.\end{aligned} \tag{56}$$

Write $\dot{\boldsymbol{h}}^i$ for the marginal of $\dot{\boldsymbol{h}}$ on $\underline{\dot{\sigma}}^i$, write $\hat{\boldsymbol{h}}^i$ for the marginal of $\hat{\boldsymbol{h}}$ on $\underline{\hat{\sigma}}^i$, and write $\boldsymbol{h}^i$ for the pair $(\dot{\boldsymbol{h}}^i, \hat{\boldsymbol{h}}^i)$. The rate function $\boldsymbol{\Phi}_2$ is the obvious analogue of (35) for the pair auxiliary model. The sum in (55) is taken over the space $\boldsymbol{\Delta}_2[\alpha]$ of pair empirical measures $\boldsymbol{h}$ such that $\boldsymbol{h}^i \in \boldsymbol{\Delta}[\alpha]$ for both $i=1,2$. Let ${}^{\otimes}\boldsymbol{h}[\alpha]$ denote the product measure ${}^{\star}\boldsymbol{h}[\alpha] \otimes {}^{\star}\boldsymbol{h}[\alpha]$, and let ${}^{\mathrm{id}}\boldsymbol{h}[\alpha]$ denote the measure with marginals ${}^{\star}\boldsymbol{h}[\alpha]$ which is supported on pair configurations $\underline{\tau} \equiv (\underline{\sigma}, \underline{\sigma})$.

THEOREM 4.1. *For $d \geqslant d_0$ and $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$, the restriction of $\boldsymbol{\Phi}_2$ to $\boldsymbol{\Delta}_2[\alpha]$ satisfies*

$$\boldsymbol{\Phi}_2(\boldsymbol{h}) < \max\{\boldsymbol{\Phi}(\boldsymbol{h}') : \boldsymbol{h}' \in \boldsymbol{\Delta}_2[\alpha]\} \quad \textit{for all } \boldsymbol{h} \in \boldsymbol{\Delta}_2[\alpha] \setminus \{{}^{\otimes}\boldsymbol{h}[\alpha], {}^{\mathrm{id}}\boldsymbol{h}[\alpha]\}.$$

Thus $\boldsymbol{\Phi}_2$ achieves its maximum on $\boldsymbol{\Delta}_2[\alpha]$ at ${}^{\otimes}\boldsymbol{h}[\alpha]$ with value $2\boldsymbol{\phi}(\alpha)$, or at ${}^{\mathrm{id}}\boldsymbol{h}[\alpha]$ with value $\boldsymbol{\phi}(\alpha)$, depending on the sign of $\boldsymbol{\phi}(\alpha)$.

Definition 4.2. For $\boldsymbol{h} \in \boldsymbol{\Delta}_2[\alpha]$, we write $\varrho(\boldsymbol{h}) \equiv \dot{\boldsymbol{h}}(\eta^1 = \eta^2 = \mathtt{1})$, and define the *near-independent regime* and the *near-identical regime*, respectively, by

$$\begin{aligned} \mathsf{IND}[\alpha] &\equiv \{\boldsymbol{h} \in \boldsymbol{\Delta}_2[\alpha] : 0 \leqslant \varrho(\boldsymbol{h}) \leqslant d^{-1.1}\}, \\ \mathsf{EQ}[\alpha] &\equiv \left\{\boldsymbol{h} \in \boldsymbol{\Delta}_2[\alpha] : \alpha - \frac{\varepsilon_\circ \log d}{d} \leqslant \varrho(\boldsymbol{h}) \leqslant \alpha\right\}, \end{aligned}$$

where $\varepsilon_\circ > 0$ is a small absolute constant.

4.1. Intermediate overlap regime

Write $\boldsymbol{Z}^2_{n\alpha}(n\varrho)$ for the contribution to $(\boldsymbol{Z}_{n\alpha})^2$ from pairs $(\underline{\bar{\boldsymbol{\eta}}}^1, \underline{\bar{\boldsymbol{\eta}}}^2)$ in which exactly $n\varrho$ vertices take spin $\mathtt{1}$ in both $\underline{\bar{\boldsymbol{\eta}}}^i$, $i=1,2$. We call ϱ the *overlap*. In this subsection we show that the value of ϱ maximizing $\mathbb{E}[\boldsymbol{Z}^2_{n\alpha}(n\varrho)]$ must either be very small (indicating proximity to ${}^{\otimes}\boldsymbol{h}[\alpha]$), or very close to α (indicating proximity to ${}^{\mathrm{id}}\boldsymbol{h}[\alpha]$). This will be done by comparison between $\boldsymbol{Z}^2_{n\alpha}(n\varrho)$ and $Z^2_{n\alpha}(n\varrho)$, where the latter denotes the contribution to $(Z_{n\alpha})^2$ from pairs $(\underline{x}^1, \underline{x}^2)$ of independent sets with overlap ϱ.

Write $\mathscr{P} \equiv \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}^2 \cup \{\mathtt{f!f}\}$. Given a pair $(\underline{\bar{\boldsymbol{\eta}}}^1, \underline{\bar{\boldsymbol{\eta}}}^2)$, we associate a configuration

$$\underline{\omega} \equiv \underline{\omega}(\underline{\bar{\boldsymbol{\eta}}}^1, \underline{\bar{\boldsymbol{\eta}}}^2) \in \mathscr{P}^V \tag{57}$$

by setting $\omega_v = (\eta^1_v, \eta^2_v)$ *unless* $(\eta^1_v, \eta^2_v) = \mathtt{ff}$ 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 $S \subseteq V$ and $\omega \in \mathscr{P}$ let $S_\omega \equiv \{v \in S : \omega_s = \omega\}$. Recall that $\textvisiblespace$ is shorthand for $\{\mathtt{0}, \mathtt{1}, \mathtt{f}\}$; and for $\eta \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}$ write

$$S_{\eta\textvisiblespace} \equiv \{v \in S : \eta^1_v = \eta\}, \quad S_{\textvisiblespace\eta} \equiv \{v \in S : \eta^2_v = \eta\}, \quad S_\eta \equiv S_{\eta\textvisiblespace} \cup S_{\textvisiblespace\eta}. \tag{58}$$

In the case $S{=}V$ we further set the following notation:

$$n_\omega \equiv n\pi_\omega \equiv |V_\omega| = |\{v\in V: \omega_v=\omega\}|, \qquad \text{for } \omega\in\mathscr{P},$$
$$n_\eta \equiv n\pi_\eta \equiv |V_\eta| = |\{v\in V: \omega_v\in\mathscr{P}_\eta\}|, \qquad \text{for } \eta\in\{\mathtt{0},\mathtt{1},\mathtt{f}\}.$$

In particular, $(\pi_\omega)_\omega$ defines a probability measure on $\mathscr{P}$.

If we write $M(n\alpha, n\varrho)$ for the space of empirical measures $\bar\pi$ on $\{\mathtt{0},\mathtt{1}\}^2$ with $\bar\pi_{\mathtt{11}}{=}\varrho$ and $\bar\pi_{\mathtt{1}\llcorner}{=}\alpha{=}\bar\pi_{\llcorner\mathtt{1}}$, then we have the decomposition

$$\mathbb{E}[Z^2_{n\alpha}(n\varrho)] = \sum_{\bar\pi\in M(n\alpha,n\varrho)} \mathbb{E}[Z^2(\bar\pi)]. \tag{59}$$

Similarly, in the frozen model we have

$$\mathbb{E}[\boldsymbol{Z}^2_{n\alpha}(n\varrho)] = \sum_{\pi\in \boldsymbol{M}(n\alpha,n\varrho)} \mathbb{E}[\boldsymbol{Z}^2(\pi)], \tag{60}$$

where $\boldsymbol{M}(n\alpha, n\varrho)$ is the space of empirical measures π on $\mathscr{P}$ such that $\pi_{\mathtt{11}}{=}\varrho$, and the marginal distribution π^i on ω^i satisfies (for both $i{=}1,2$) that $\pi^i_{\mathtt{f}}{\leqslant}\beta_{\max}$, with normalized intensity $\mathbf{i}(\pi){=}\pi^i_{\mathtt{1}}{+}\frac{1}{2}\pi^i_{\mathtt{f}}{=}\alpha$ (cf. (34)). Note that π is simply a projection of the measure $\boldsymbol{h}$ appearing in (55).

We calculate $\mathbb{E}[\boldsymbol{Z}^2(\pi)]$ as follows. The pair of frozen configurations $\underline{\boldsymbol{\bar\eta}}^i{=}(\underline{\eta}^i, \underline{m}^i)$ is encoded by $(\underline{\omega}, \underline{m})$ where $\underline{\omega}$ is as above, and $\underline{m}{\equiv}(\underline{m}^1, \underline{m}^2)$. Then

$$\mathbb{E}[\boldsymbol{Z}^2(\pi)] = \sum_{\underline{\omega},\underline{m}} \mathbb{P}(\mathbf{1}\{(\underline{\omega},\underline{m}) \text{ is valid on } G\}) \tag{61}$$

where the sum is taken over all $(\underline{\omega}, \underline{m})$ consistent with π, and the probability is with respect to the randomness of G. In the literature, this representation sometimes goes by the name "planted model." In the original setup, we first sample the random graph and then consider the configurations valid on the graph. In the planted setup, we first plant a configuration $(\underline{\omega}, \underline{m})$, and then sample the random graph and consider the probability for the configuration to be valid.

Let us write $\underline{m}{\subseteq}G$ to indicate that the edges of $\underline{m}^1, \underline{m}^2$ are present in the graph G. Conditioning on this event gives

$$\mathbb{E}[\boldsymbol{Z}^2(\pi)] = \sum_{\underline{\omega},\underline{m}} \mathbb{P}(\underline{m}\subseteq G)\, \underbrace{\mathbb{P}(\mathbf{1}\{(\underline{\omega},\underline{m}) \text{ is valid on } G\} \mid \underline{m}\subseteq G)}_{\mathbf{G}(\pi)},$$

where it follows by symmetry considerations that the second factor is a function of π alone, so we denote it $\mathbf{G}(\pi)$. Therefore

$$\mathbb{E}[\boldsymbol{Z}^2(\pi)] = \mathbf{G}(\pi)\sum_{\underline{\omega},\underline{m}} \mathbb{P}(\underline{m}\subseteq G) = \mathbf{G}(\pi)\mathbf{c}(\pi)\mathbf{j}_{\mathtt{f}}(\pi), \tag{62}$$

where $\mathbf{c}(\pi)$ denotes the multinomial coefficient

$$\mathbf{c}(\pi)=\binom{n}{n\pi}=\frac{n!}{\prod_{\omega\in\mathscr{P}}(n\pi_\omega)!},$$

and $\mathbf{j}_{\mathtt{f}}(\pi)$ is the expected number of matchings on the free vertices in $\underline{\omega}$, where $\underline{\omega}\in\mathscr{P}^V$ is an arbitrary fixed configuration with empirical measure π.

We now estimate the terms appearing in (62). It will be useful to let $U_{\mathtt{0f}}\subseteq V_{\mathtt{0f}}$ denote the set of $\mathtt{0f}$-vertices whose matched partner under $\underline{m}^2$ does not have spin $\mathtt{1f}$. Symmetrically, let $U_{\mathtt{f0}}\subseteq V_{\mathtt{f0}}$ denote the set of $\mathtt{f0}$-vertices whose matched partner under $\underline{m}^1$ does not have spin $\mathtt{f1}$. Write

$$\begin{aligned}\mathsf{u}_{\mathtt{0f}}\equiv|U_{\mathtt{0f}}|=n_{\mathtt{0f}}-n_{\mathtt{1f}},\quad \mathsf{p}_{\mathtt{0f}}\equiv\frac{\mathsf{u}_{\mathtt{0f}}}{n}=\pi_{\mathtt{0f}}-\pi_{\mathtt{1f}},\\ \mathsf{u}_{\mathtt{f0}}\equiv|U_{\mathtt{f0}}|=n_{\mathtt{f0}}-n_{\mathtt{f1}},\quad \mathsf{p}_{\mathtt{f0}}\equiv\frac{\mathsf{u}_{\mathtt{f0}}}{n}=\pi_{\mathtt{f0}}-\pi_{\mathtt{f1}}.\end{aligned}$$

LEMMA 4.3. *Given $\pi\in\boldsymbol{M}(n\alpha,n\varrho)$, let $\mathbf{j}_{\mathtt{f}}(\pi)$ be the expected number of matchings on the free vertices in $\underline{\omega}$, with $\underline{\omega}\in\mathscr{P}^V$ being an arbitrary fixed configuration with empirical measure π. Then $(2n)^{-1}\mathbf{k}_{\mathtt{f}}(\pi)\leqslant\mathbf{j}_{\mathtt{f}}(\pi)\leqslant\mathbf{k}_{\mathtt{f}}(\pi)$, where*

$$\mathbf{k}_{\mathtt{f}}(\pi)\equiv\frac{d^{n_{\mathtt{f}}+n_{\mathtt{f!f}}}}{[d/(d-1)]^{n_{\mathtt{f!f}}}}\,\frac{(n_{\mathtt{ff}}-1)!!(n_{\mathtt{0f}})_{n_{\mathtt{1f}}}(n_{\mathtt{f0}})_{n_{\mathtt{f1}}}(\mathsf{u}_{\mathtt{0f}}+n_{\mathtt{f!f}}-1)!!(\mathsf{u}_{\mathtt{f0}}+n_{\mathtt{f!f}}-1)!!}{[nd]_{(n_{\mathtt{f}}+n_{\mathtt{f!f}})/2}}.$$

Proof. The upper bound is easy to see: for each vertex in $V_{\mathtt{f}}\setminus V_{\mathtt{f!f}}$ we distinguish one half-edge to participate in the matching. For each $\mathtt{f!f}$-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 $n_{\mathtt{f}}+n_{\mathtt{f!f}}$ distinguished half-edges in total. If $n_{\mathtt{f!f}}<2$ then $\mathbf{j}_{\mathtt{f}}(\pi)=\mathbf{k}_{\mathtt{f}}(\pi)$; in general, $\mathbf{k}_{\mathtt{f}}$ provides an upper bound for $\mathbf{j}_{\mathtt{f}}$ because it counts matchings without enforcing the constraint that two $\mathtt{f!f}$-vertices cannot be matched in both $\underline{m}^1$ and $\underline{m}^2$.

For the lower bound when $n_{\mathtt{f!f}}\geqslant 2$, suppose without loss of generality that $\mathsf{u}_{\mathtt{f0}}\geqslant\mathsf{u}_{\mathtt{0f}}$, and note that this implies that $\mathsf{a}\equiv\mathsf{u}_{\mathtt{f0}}+n_{\mathtt{f!f}}$ must be at least 4—otherwise $n_{\mathtt{f!f}}=2$ while $\mathsf{u}_{\mathtt{0f}}=\mathsf{u}_{\mathtt{f0}}=0$ so there is no valid matching on the $\mathtt{f!f}$-vertices. Of the half-edges chosen to participate in either $\underline{m}_1$ or $\underline{m}_2$, suppose that all have been matched *except* for the a half-edges incident to $U_{\mathtt{f0}}\cup V_{\mathtt{f!f}}$ 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 $(\mathsf{a}/2)-1$ pairs is $\geqslant(\mathsf{a}-2)(\mathsf{a}-4)\dots 2\geqslant(\mathsf{a}-3)!!\geqslant(\mathsf{a}-1)!!/n$.

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 $\frac{1}{2}\mathsf{a}-1$ pairs, and switch the half-edges in one of two ways to produce a valid matching (Figure 2). Thus each failed matching maps to $\mathsf{a}-2$ valid matchings.

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

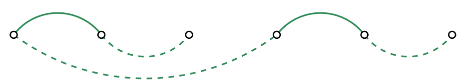

(b) One valid switching.

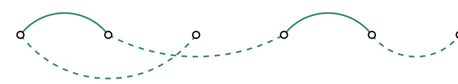

(c) A distinct valid switching.

Figure 2. Two valid switchings of a failed matching (dashed lines $\underline{m}^1$, solid lines $\underline{m}^2$).

In the reverse direction, given the final pair the failed matching can be uniquely recovered from the valid matching, so each valid matching has at most $\frac{1}{2}n_{\mathtt{f!f}}$ preimages. Thus the ratio of failed-to-valid matchings is (recalling that $\mathsf{a}\geqslant 4$) at most

$$\frac{\frac{1}{2}n_{\mathtt{f!f}}}{\mathsf{a}-2}\leqslant\left(2-\frac{4}{\mathsf{a}}\right)^{-1}\leqslant 1.$$

This proves that the matching procedure succeeds with probability at least $\frac{1}{2}$, and the claimed lower bound follows. ☐

COROLLARY 4.4. *For $\pi\in\boldsymbol{M}(n\alpha,n\varrho)$, let $\bar{\pi}\in M(n\alpha,n\varrho)$ be defined by*

$$\bar{\pi}_{\mathtt{11}}=\pi_{\mathtt{11}},\quad \bar{\pi}_{\mathtt{10}}=\pi_{\mathtt{10}}+\pi_{\mathtt{1f}}+\tfrac{1}{2}\pi_{\mathtt{f}\_},\quad \bar{\pi}_{\mathtt{01}}=\pi_{\mathtt{01}}+\pi_{\mathtt{f1}}+\tfrac{1}{2}\pi_{\_\mathtt{f}},$$

with the remaining probability going to $\bar{\pi}_{\mathtt{00}}$. Then $\mathbf{c}(\pi)\mathbf{j}_{\mathtt{f}}(\pi)=d^{O(n\beta_{\max})}\mathbf{c}(\bar{\pi})$.

Proof. Let us first compare $\mathbf{c}(\pi)\mathbf{j}_{\mathtt{f}}(\pi)$ against

$$\mathbf{c}''(\pi)\equiv\binom{n-n_{\mathtt{f}}}{n_{\mathtt{00}}!n_{\mathtt{01}}!n_{\mathtt{10}}!n_{\mathtt{11}}!}.$$

It follows from Lemma 4.3 and Stirling's formula that

$$\begin{aligned}\frac{\mathbf{c}(\pi)\mathbf{j}_{\mathtt{f}}(\pi)}{\mathbf{c}''(\pi)}&=n^{O(1)}\frac{d^{n_{\mathtt{f}}}(d-1)^{n_{\mathtt{f!f}}}(n)_{\mathtt{f}}}{[nd]_{(n_{\mathtt{f}}+n_{\mathtt{f!f}})/2}}\frac{(n_{\mathtt{ff}}-1)!!(\mathsf{u}_{\mathtt{0f}}+n_{\mathtt{f!f}}-1)!!(\mathsf{u}_{\mathtt{f0}}+n_{\mathtt{f!f}}-1)!!}{n_{\mathtt{1f}}!n_{\mathtt{f1}}!\mathsf{u}_{\mathtt{0f}}!\mathsf{u}_{\mathtt{f0}}!n_{\mathtt{ff}}!n_{\mathtt{f!f}}!}\\&=d^{O(n_{\mathtt{f}})}\frac{(\mathsf{p}_{\mathtt{0f}}+\pi_{\mathtt{f!f}})^{(\mathsf{u}_{\mathtt{0f}}+n_{\mathtt{f!f}})/2}(\mathsf{p}_{\mathtt{f0}}+\pi_{\mathtt{f!f}})^{(\mathsf{u}_{\mathtt{f0}}+n_{\mathtt{f!f}})/2}}{(\pi_{\mathtt{1f}})^{n_{\mathtt{1f}}}(\pi_{\mathtt{f1}})^{n_{\mathtt{f1}}}(\mathsf{p}_{\mathtt{0f}})^{\mathsf{u}_{\mathtt{0f}}}(\mathsf{p}_{\mathtt{f0}})^{\mathsf{u}_{\mathtt{f0}}}(\pi_{\mathtt{ff}})^{n_{\mathtt{ff}}/2}(\pi_{\mathtt{f!f}})^{n_{\mathtt{f!f}}}}=d^{O(n\beta_{\max})},\end{aligned}\tag{63}$$

where the last estimate uses that for $\delta\leqslant 1/e$, $\max\{|x\log x|:0\leqslant x\leqslant\delta\}=|\delta\log\delta|$. Note also that, for $0\leqslant x-\delta\leqslant x\leqslant 1/e$, we have the elementary inequality

$$|(x-\delta)\log(x-\delta)-x\log x|\leqslant\min\{|x\log x|,\delta|1+\log(x-\delta)|\}\lesssim|\delta\log\delta|.$$

From this it is easily seen that

$$\frac{\mathbf{c}''(\pi)}{\mathbf{c}(\bar{\pi})}=\exp(O(n\beta_{\max}\log d)),$$

concluding the proof. ☐

We now return to equation (62) and the calculation of $\mathbb{E}[\boldsymbol{Z}^2(\pi)]$. From now on, let $(\underline{\omega}, \underline{m})$ be a fixed configuration consistent with π. With respect to this fixed configuration, define $\boldsymbol{J}_1$ to be the indicator that $\mathtt{1}$-vertices neighbor only $\mathtt{0}$-vertices. Define $\boldsymbol{J}_{\mathtt{0}}$ to be the indicator that every $\mathtt{0}$-vertex is forced, that is, has at least two neighbors with spin $\mathtt{1}$. Then

$$\mathbf{G}(\pi) = \mathbb{E}[\boldsymbol{J}_1 \boldsymbol{J}_{\mathtt{0}}].$$

Now recall the notation of (58), and recall also that we use $\mathtt{Z}$ to denote $\{\mathtt{0}, \mathtt{f}\}$. For an integer K and a pair of integers $F \equiv (F_1, F_2)$, define

$$\boldsymbol{Z}^2(\pi, K, F) \equiv \boldsymbol{Z}^2(\pi)\chi_{K,F}, \quad \text{where } \chi_{K,F} \equiv \mathbf{1}\left\{\begin{array}{c} \mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{01}}) = K, \\ \mathtt{E}(V_{\mathtt{1Z}}, V_{\mathtt{0f}}) - n_{\mathtt{1f}} = F_1, \\ \mathtt{E}(V_{\mathtt{Z1}}, V_{\mathtt{f0}}) - n_{\mathtt{f1}} = F_2 \end{array}\right\}.$$

That is, K counts edges between $V_{\mathtt{10}}$ and $V_{\mathtt{01}}$, F_1 counts non-matching edges between $V_{\mathtt{10}} \cup V_{\mathtt{1f}}$ and $V_{\mathtt{0f}}$, and F_2 counts non-matching edges between $V_{\mathtt{01}} \cup V_{\mathtt{f1}}$ and $V_{\mathtt{f0}}$. Recall that each $\mathtt{0}$-variable must have at least two edges going to $\mathtt{1}$-variables, which implies that $K \geqslant 2\max\{n_{\mathtt{10}}, n_{\mathtt{01}}\}$, $F_1 \geqslant 2n_{\mathtt{0f}} - n_{\mathtt{1f}}$, and $F_2 \geqslant 2n_{\mathtt{f0}} - n_{\mathtt{f1}}$. We decompose $\mathbb{E}[\boldsymbol{Z}^2(\pi)]$ as the sum over K and F of

$$\mathbb{E}[\boldsymbol{Z}^2(\pi, K, F)] = \mathbf{c}(\pi)\mathbf{j}_{\mathtt{f}}(\pi) \overbrace{\underbrace{\mathbb{E}[\boldsymbol{J}_1 \chi_{K,F}]}_{\mathbf{j}_1(\pi,K,F)} \underbrace{\mathbb{E}[\boldsymbol{J}_{\mathtt{0}} \mid \boldsymbol{J}_1 = 1 = \chi_{K,F}]}_{\mathbf{j}_{\mathtt{0}}(\pi,K,F)}}^{\mathbf{G}(\pi,K,F)}. \tag{64}$$

For $\omega \in \mathscr{P}$ write E_ω for the number of unmatched half-edges incident to V_ω, not counting half-edges participating in $\underline{m}$. Likewise, for $\eta \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}$ write E_η for the number of unmatched half-edges incident to V_η. The total number of unmatched half-edges remaining after the placement of $\underline{m}$ is $\mathsf{E} = \sum_\omega \mathsf{E}_\omega = nd - n_{\mathtt{f}} - n_{\mathtt{f!f}}$. Note that if these half-edges are matched uniformly at random without regard to the model constraints, we would expect $K \approx K_\star$ and $F \approx F_\star$, where

$$K_\star \equiv \frac{\mathsf{E}_{\mathtt{10}} \mathsf{E}_{\mathtt{01}}}{\mathsf{E}} \quad \text{and} \quad F_\star \equiv (F_1, F_2)_\star \equiv \left(\frac{\mathsf{E}_{\mathtt{1Z}} \mathsf{E}_{\mathtt{0f}}}{\mathsf{E}}, \frac{\mathsf{E}_{\mathtt{Z1}} \mathsf{E}_{\mathtt{f0}}}{\mathsf{E}}\right).$$

Note $K_\star \asymp n\varepsilon^2 (\log d)^2/d$ while $F_{1,\star} \asymp n_{\mathtt{0f}} \varepsilon \log d$ and $F_{2,\star} \asymp n_{\mathtt{f0}} \varepsilon \log d$. In what follows we will show that the main contribution to $\mathbb{E}[\boldsymbol{Z}^2(\pi)]$ comes from (K, F) near $(K, F)_\star$.

LEMMA 4.5. *Let $\varrho = \alpha - \varepsilon(\log d)/d$, where $\varepsilon \geqslant \varepsilon_\circ$ for an absolute constant $\varepsilon_\circ > 0$.*(10) *Then, for all $\pi \in \boldsymbol{M}(n\alpha, n\varrho)$, we have*

$$\mathbb{E}[\boldsymbol{Z}^2(\pi)] = \sum_{K,F} \mathbb{E}[\boldsymbol{Z}^2(\pi, K, F)] \asymp \sum_{K,F}^{*} \mathbb{E}[\boldsymbol{Z}^2(\pi, K, F)], \tag{65}$$

(10) The lower bound d_0 may depend on $\varepsilon_\circ$, but we will ultimately take $\varepsilon_\circ$ to be an absolute constant, so that d_0 is also an absolute constant.

where $*$ indicates that the sum is restricted to

$$[1-(\log d)^{-1}]K_\star \equiv K_{\mathrm{lbd}} \leqslant K \leqslant K_{\mathrm{ubd}} \equiv [1+(\log d)^{-1}]K_\star,$$
$$(\log\log d)^{-1}F_\star \equiv F_{\mathrm{lbd}} \leqslant F \leqslant F_{\mathrm{ubd}} \equiv F_\star(\log\log d).$$

Further, for K and F in the restricted regime we have

$$\frac{\mathbb{E}[\boldsymbol{Z}^2(\pi,K,F)]}{\mathbb{E}[Z^2(\bar\pi,K)]} = \exp(O(nd^{-1}(d^{-0.49}+d^{-3\varepsilon/4}))), \tag{66}$$

with $\bar\pi \in M(n\alpha, n\varrho)$ as defined in Corollary 4.4.

Proof. Recalling (64), we now estimate $\mathbf{j}_1(\pi,K,F)$ and $\mathbf{j}_0(\pi,K,F)$.

Edges from $\mathtt{1}$*-vertices.* First assign the $\mathsf{E}_1 = n_1 d - n_{\mathtt{1f}} - n_{\mathtt{f1}}$ unmatched half-edges from V_1:

$$\mathbf{j}_1(\pi,K,F) = \frac{(\mathsf{E}_{10})_K(\mathsf{E}_{01})_K}{K!}\frac{(\mathsf{E}_{0\mathtt{f}})_{F_1}(\mathsf{E}_{1\mathtt{z}}-K)_{F_1}}{(F_1)!}\frac{(\mathsf{E}_{\mathtt{f}0})_{F_2}(\mathsf{E}_{\mathtt{z}1}-K)_{F_2}}{(F_2)!}\frac{(\mathsf{E}_{00})_{\mathsf{E}_1-2K-F_1-F_2}}{[\mathsf{E}]_{\mathsf{E}_1-K}}.$$

We then estimate

$$\mathbf{j}_1(\pi,K,F) = d^{O(n\beta_{\max})}\mathbf{g}_1(\mathsf{E},\mathsf{E}_1,\mathsf{E}_{00})\mathbf{a}_1(\pi,K)\mathbf{b}_1(\pi,K,F),$$

where $\mathbf{g}_1(\mathsf{E},\mathsf{E}_1,\mathsf{E}_{00}) = (\mathsf{E}_{00})_{\mathsf{E}_1}/(\mathsf{E})_{\mathsf{E}_1}$ as defined by (14), and

$$\mathbf{a}_1(\pi,K) \equiv \frac{(\mathsf{E}_{10})_K(\mathsf{E}_{01})_K/K!}{[\mathsf{E}_{00}-\mathsf{E}_1+2K]_K} = \frac{(\mathsf{E}_{10})_K(\mathsf{E}_{01})_K/K!}{(n[d-O(\log d)])^K}, \tag{67}$$

$$\mathbf{b}_1(\pi,K,F) \equiv \frac{(\mathsf{E}_{0\mathtt{f}})_{F_1}(\mathsf{E}_{1\mathtt{z}}-K)_{F_1}}{(F_1)!(\mathsf{E}_{00})^{F_1}}\frac{(\mathsf{E}_{\mathtt{f}0})_{F_2}(\mathsf{E}_{\mathtt{z}1}-K)_{F_2}}{(F_2)!(\mathsf{E}_{00})^{F_2}} \leqslant e^{O[(F_1+F_2)_\star]}. \tag{68}$$

Forcing of $\mathtt{0}$*-vertices.* Recall from (15) that $\mathbf{g}_0(n, nd\zeta_{\mathtt{1}\_})$ denotes the probability, with respect to a uniformly random assignment of $nd\zeta_{\mathtt{1}\_}$ half-edges to n degree-d vertices, that each vertex is forced (that is, receives at least two of the incoming half-edges). Similarly:

(1) Let $\tilde{\mathbf{g}}_0(n,\mathsf{u},nd\zeta_{\mathtt{1}\_})$ denote the probability, with respect to a uniformly random assignment of $nd\zeta_{\mathtt{1}\_}$ half-edges to n vertices of degree $d-1$, that each vertex $1 \leqslant i \leqslant \mathsf{u}$ receives at least two incoming half-edges, while each vertex $\mathsf{u} < i \leqslant n$ receives at least one.

(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 $\{\mathtt{11},\mathtt{10}\}$ and from $\{\mathtt{11},\mathtt{01}\}$).

Thus $\mathbf{j}_0(\pi, K, F)$ can be written as

$$\begin{aligned}\mathbf{j}_0(\pi,K,F) = &\overbrace{\mathbf{g}_0(n_{10},K)\,\mathbf{g}_0(n_{01},K)}^{(\mathrm{A})}\overbrace{\tilde{\mathbf{g}}_0(n_{0\mathtt{f}},\mathtt{u}_{0\mathtt{f}},F_1)\tilde{\mathbf{g}}_0(n_{\mathtt{f}0},\mathtt{u}_{\mathtt{f}0},F_2)}^{(\mathrm{B})}\\ &\times\underbrace{\mathbf{g}_{00}(n_{00},(n_{11}d,n_{1\mathtt{z}}d-K-n_{1\mathtt{f}}-F_1,n_{\mathtt{z}1}-n_{\mathtt{f}1}-K-F_2))}_{(\mathrm{C})}.\end{aligned}$$

The functions $\mathbf{g}_0$ and $\mathbf{g}_{00}$ were estimated in Proposition 2.10. It is clear that $\tilde{\mathbf{g}}_0$ satisfies the same estimates as those stated for $\mathbf{g}_0$ in Proposition 2.10 (a), so we conclude that

$$\begin{aligned}&(\mathrm{A}) = \exp(O((n_{10}+n_{01})d^{-4\varepsilon/5})) && \text{for } K \geqslant \tfrac{9}{10}K_\star;\\ &(\mathrm{B}) = \exp(O((n_{0\mathtt{f}}+n_{\mathtt{f}0})d^{-4\varepsilon/5})) && \text{for } F \geqslant \tfrac{9}{10}F_\star;\\ &(\mathrm{C}) = \exp(O(nd^{-1.65})) && \text{for } K \lesssim \min\{n_{10},n_{01}\}.\end{aligned} \tag{69}$$

Note also that, if we fix K and vary only F, Proposition 2.10 (a) gives

$$\frac{\mathbf{j}_0(\pi,K,F')}{\mathbf{j}_0(\pi,K,F)} = \exp(O(\|F'-F\|_1)). \tag{70}$$

Maximization over K. If $K \notin [K_{\mathrm{lbd}}, K_{\mathrm{ubd}}]$ then, recalling (67) and (68), we have

$$\begin{aligned}\frac{\mathbf{j}_1(\pi,K,F)}{\mathbf{j}_1(\pi,K_\star,F_\star)} &= d^{O(n\beta_{\max})}\frac{\mathbf{a}_1(\pi,K)}{\mathbf{a}_1(\pi,K_\star)}\frac{\mathbf{b}_1(\pi,K,F)}{\mathbf{b}_1(\pi,K_\star,F_\star)} \leqslant d^{O(n\beta_{\max})}\frac{\mathbf{a}_1(\pi,K)}{\mathbf{a}_1(\pi,K_\star)}\\ &\leqslant d^{O(n\beta_{\max})}\exp\left(-\Omega\left(\frac{K^\star}{(\log d)^2}\right)\right) \leqslant \exp\left(-\Omega\left(\frac{n\varepsilon^2}{d}\right)\right).\end{aligned}$$

Meanwhile, trivially $\mathbf{j}_0 \leqslant 1$. We have also seen above that $\mathbf{j}_0(\pi, K_\star, F_\star)$ is not much smaller than 1 (in the sense that there is a small exponential gap). Thus, for all $K \notin [K_{\mathrm{lbd}}, K_{\mathrm{ubd}}]$,

$$\begin{aligned}\frac{\mathbb{E}[\boldsymbol{Z}^2(\pi,K,F)]}{\mathbb{E}[\boldsymbol{Z}^2(\pi,K_\star,F_\star)]} &= \frac{\mathbf{j}_1(\pi,K,F)}{\mathbf{j}_1(\pi,K_\star,F_\star)}\frac{\mathbf{j}_0(\pi,K,F)}{\mathbf{j}_0(\pi,K_\star,F_\star)} \leqslant \frac{\mathbf{j}_1(\pi,K,F)}{\mathbf{j}_1(\pi,K_\star,F_\star)}\frac{1}{\mathbf{j}_0(\pi,K_\star,F_\star)}\\ &\leqslant \frac{\exp(O[nd^{-1}(d^{-0.65}+d^{-4\varepsilon/5}\log d)])}{\exp(\Omega(n\varepsilon^2/d))} \ll 1.\end{aligned}$$

Maximization over F. Now suppose $K \in [K_{\mathrm{lbd}}, K_{\mathrm{ubd}}]$ but $F_1 > F_{1,\mathrm{ubd}} = F_{1,\star}(\log\log d)$, and let $F' \equiv (F_{1,\star}, F_2)$. Then combining (68) with (70) gives

$$\frac{\mathbb{E}[\boldsymbol{Z}^2(\pi,K,F)]}{\mathbb{E}[\boldsymbol{Z}^2(\pi,K,F')]} = \frac{\mathbf{b}_1(\pi,K,F)}{\mathbf{b}_1(\pi,K,F')}\frac{\mathbf{j}_0(\pi,K,F)}{\mathbf{j}_0(\pi,K,F')} \leqslant \frac{e^{O(F_1)}}{(\log\log d)^{\Omega(F_1)}} \ll 1.$$

If instead $F_1 < F_{1,\mathrm{lbd}} = F_{1,\star}/(\log\log d)$ then, letting $F'' \equiv (2F_{1,\mathrm{lbd}}, F_2)$, and applying (68) and (70), again gives

$$\frac{\mathbb{E}[\boldsymbol{Z}^2(\pi,K,F)]}{\mathbb{E}[\boldsymbol{Z}^2(\pi,K,F'')]} \leqslant \frac{\exp(O(F_{1,\mathrm{lbd}}))}{\exp(\Omega[F_{1,\mathrm{lbd}}(\log\log\log d)])} \ll 1.$$

The same argument applies if $F_2 \notin [F_{2,\mathrm{lbd}}, F_{2,\mathrm{ubd}}]$, and this concludes the proof of (65).

Comparison with second moment of independent sets. We have

$$\frac{\mathbb{E}[\boldsymbol{Z}^2(\pi,K,F)]}{\mathbb{E}[Z^2(\bar\pi,K)]} = \frac{\mathbf{c}(\pi)\mathbf{j}_{\mathfrak{f}}(\pi)}{\mathbf{c}(\bar\pi)} \frac{\mathbf{j}_1(\pi,K,F)}{\mathbf{j}_1(\bar\pi,K,0)} \frac{\mathbf{j}_0(\pi,K,F)}{\mathbf{j}_0(\bar\pi,K,0)},$$

where Corollary 4.4 implies that the first ratio is $\exp(O(n\beta_{\max}\log d))$. Let $\bar{\mathsf{E}}_\eta$ and $\bar{\mathsf{E}}_\omega$ be defined as E_η and E_ω above, but with $\bar\pi$ in place of π. Then

$$\frac{\mathbf{j}_1(\pi,K,F)}{\mathbf{j}_1(\bar\pi,K,0)} = \underbrace{\frac{\mathbf{g}_1(\mathsf{E},\mathsf{E}_1,\mathsf{E}_{00})}{\mathbf{g}_1(\bar{\mathsf{E}},\bar{\mathsf{E}}_1,\bar{\mathsf{E}}_{00})} \frac{\mathbf{a}_1(\pi,K)}{\mathbf{a}_1(\bar\pi,K)}}_{\exp(O(n\beta_{\max}\log d))} \frac{\mathbf{b}_1(\pi,K,F)}{\mathbf{b}_1(\bar\pi,K,0)}.$$

Recall from (67) that

$$\mathbf{b}_1(\pi,K,F) \leqslant d^{O(n\beta_{\max})}$$

always. For $K \in [K_{\mathrm{lbd}}, K_{\mathrm{ubd}}]$ and $F \in [F_{\mathrm{lbd}}, F_{\mathrm{ubd}}]$,

$$\mathbf{b}_1(\pi,K,F) \geqslant (\log\log d)^{-O(F_{\mathrm{ubd}})} \geqslant \exp(-F_\star(\log\log d)^2).$$

Combining this estimate with (69) gives (66), as claimed. □

COROLLARY 4.6. *Let $\varrho = \alpha - \varepsilon(\log d)/d$, where $\varepsilon \geqslant \varepsilon_\circ$ for an absolute constant $\varepsilon_\circ > 0$. Then*

$$\mathbb{E}[\boldsymbol{Z}^2_{n\alpha}(n\varrho)] = \mathbb{E}[Z^2_{n\alpha}(n\varrho)] \exp(O[nd^{-1}(d^{-0.49} + d^{-3\varepsilon/4})]), \tag{71}$$

and the ratio $\mathbb{E}[\boldsymbol{Z}^2_{n\alpha}(n\varrho)]/\mathbb{E}[\boldsymbol{Z}^2_{n\alpha}]$ is exponentially small in n for

$$\varrho \in [d^{-1.1}, \alpha - \varepsilon_\circ(\log d)/d].$$

Proof. The first assertion (71) follows by combining (59), (60), and the estimate (66) from Lemma 4.5. We now show that the ratio of $\mathbb{E}[\boldsymbol{Z}^2_{n\alpha}(n\varrho)]$ to $\max_r \mathbb{E}[\boldsymbol{Z}^2_{n\alpha}(nr)]$ is exponentially small in n for all $d^{-1.45} \leqslant \varrho \leqslant \alpha - \varepsilon_\circ(\log d)/d$. Let us begin with an analogous calculation in the independent set model:

$$\begin{aligned}\frac{\mathbb{E}[Z^2_{n\alpha}(n\varrho)]}{\mathbb{E}[Z_{n\alpha}]} &= \sum_{\pi \in M(n\alpha,n\varrho)} \frac{\mathbf{c}(\pi)\mathbb{E}[\boldsymbol{J}_1]}{\binom{n}{n\alpha}\mathbf{g}_1(nd,nd\alpha,0)} \\ &= \sum_{\pi \in M(n\alpha,n\varrho)} \frac{\mathbf{c}(\pi)\mathbf{g}_1(nd,n_1d,0)\sum_K \mathbf{a}_1(\pi,K)}{\binom{n}{n\alpha}\mathbf{g}_1(nd,nd\alpha,0)},\end{aligned}$$

with $\mathbf{a}_1(\pi,K)$ as defined in (67). We then estimate

$$\begin{aligned}\mathbf{g}_1(nd,n_1d,0) &= \exp\left(-\frac{n_1^2 d}{2n} + O(nd^{-2}(\log d)^3)\right), \\ \sum_K \mathbf{a}_1(\pi,K) &= \exp\left(\frac{n_{10}n_{01}d}{n} + O(nd^{-2}(\log d)^3)\right),\end{aligned}$$

so altogether we find

$$\mathbb{E}[Z^2_{n\alpha}(n\varrho)] = \mathbb{E}[Z_{n\alpha}]\exp(ng_\alpha(\varrho)+O(nd^{-2}(\log d)^3)),$$

where

$$g_\alpha(\varrho)\equiv \alpha H\left(\frac{\varrho}{\alpha}\right)+(1-\alpha)H\left(\frac{\alpha-\varrho}{1-\alpha}\right)-\frac{d}{2}(\alpha^2-\varrho^2).$$

The function g_α has first and second derivatives

$$\begin{aligned} g'_\alpha(\varrho) &= 2\log(\alpha-\varrho)-\log\varrho-\log(1-2\alpha+\varrho)+d\varrho,\\ g''_\alpha(\varrho) &= d-2(\alpha-\varrho)^{-1}-\varrho^{-1}-(1-2\alpha+\varrho)^{-1}, \end{aligned}$$

so we see that g_α is strictly convex on the interval $2/d\leqslant\varrho\leqslant\alpha-3/d$. From the expression for g'_α we see that the (unique) minimizer $\varrho_\circ$ of g_α on this interval must lie near $(\log d)/d$, and so applying Corollary 4.6 gives

$$\frac{\sup_{2d^{-1.48}\leqslant\varrho\leqslant\varrho_\circ}\mathbb{E}[\boldsymbol{Z}^2_{n\alpha}(n\varrho)]}{\mathbb{E}[\boldsymbol{Z}^2_{n\alpha}(nd^{-1.48})]} = \exp(O(nd^{-1.49}))\frac{\sup_{2d^{-1.48}\leqslant\varrho\leqslant\varrho_\circ}\mathbb{E}[Z^2_{n\alpha}(n\varrho)]}{\mathbb{E}[Z^2_{n\alpha}(nd^{-1.48})]}, \tag{72}$$

$$\frac{\sup_{\varrho_\circ\leqslant\varrho\leqslant\alpha-\varepsilon_\circ(\log d)/d}\mathbb{E}[\boldsymbol{Z}^2_{n\alpha}(n\varrho)]}{\mathbb{E}[\boldsymbol{Z}^2_{n\alpha}(n\alpha-n\varepsilon_\circ(\log d)/2d)]} = \exp\left(O\left(\frac{n}{d}\right)\right)\frac{\sup_{\varrho_\circ\leqslant\varrho\leqslant\alpha-\varepsilon_\circ(\log d)/d}\mathbb{E}[Z^2_{n\alpha}(n\varrho)]}{\mathbb{E}[Z^2_{n\alpha}(n\alpha-n\varepsilon_\circ(\log d)/2d)]}. \tag{73}$$

We estimate $g'_\alpha(\varrho)\leqslant-\frac{1}{10}\log d$ for $d^{-1.85}\varrho\leqslant 100/d$, and similarly we have $g'_\alpha(\varrho)\geqslant\frac{1}{10}\log d$ for $1.6(\log d)/d\leqslant\varrho\leqslant\alpha-d^{-1.25}$. Thus the quantity in (72) is bounded from above by

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

and the expression in (73) is bounded from above by

$$\exp\left(n\left[g_\alpha\left(n\alpha-\frac{n\varepsilon_\circ\log d}{d}\right)-g_\alpha\left(n\alpha-\frac{n\varepsilon_\circ\log d}{2d}\right)\right]+O\left(\frac{n}{d}\right)\right)\leqslant\exp\left(-\frac{n\varepsilon_\circ\log d}{d}\right).$$

These estimates cover the full interval $d^{-1.45}\leqslant\varrho\leqslant\alpha-\varepsilon_\circ(\log d)/d$, implying the result. $\square$

4.2. Boundary estimates

Proposition 4.7. *There are constants $\delta_d,\bar\delta_d>0$, depending on d but not on n, such that*

(a) $\arg\max\{\mathbb{E}[\boldsymbol{Z}_{n(\alpha-\beta/2),n\beta}]:0\leqslant\beta\leqslant\beta_{\max}\}\geqslant\delta_d$;

(b) $\arg\max\{\boldsymbol{Z}^2_{n\alpha}(n\varrho):0\leqslant\varrho\leqslant\alpha-\varepsilon_\circ(\log d)/d\}\geqslant\delta_d$;

(c) *for $\delta_d\leqslant\varrho\leqslant\alpha-\varepsilon_\circ(\log d)/d$, any $\pi\in\boldsymbol{M}(n\alpha,n\varrho)$ maximizing $\mathbb{E}[\boldsymbol{Z}^2(\pi)]$ must have either* (i) $\frac{1}{2}\beta_{\max}\leqslant\max\{\pi_{\mathtt{f\_}},\pi_{\mathtt{\_f}}\}\leqslant\beta_{\max}$, *or* (ii) $\min\{\{\pi_\omega\}_{\omega\in\mathscr{P}}\cup\{\mathsf{p}_{\mathtt{0f}},\mathsf{p}_{\mathtt{f0}}\}\}\geqslant\bar\delta_d$, *where we recall that $\mathsf{p}_{\mathtt{0f}}=\mathsf{u}_{\mathtt{0f}}/n=\pi_{\mathtt{0f}}-\pi_{\mathtt{1f}}$, and $\mathsf{p}_{\mathtt{f0}}$ is defined symmetrically.*

Proof. The result follows by a series of comparison estimates. The general strategy is to show that, if some π_ω is too small, then making a tiny increase in π_ω gives an entropy gain which is much larger than any probability costs that may be incurred. We first address (c), which is the most difficult. Take $\pi \in \boldsymbol{M}(n\alpha, n\varrho)$, and assume that $\max\{\pi_{\texttt{f\textvisiblespace}}, \pi_{\texttt{\textvisiblespace f}}\} \leqslant \frac{1}{2}\beta_{\max}$, since otherwise (i) holds and we are done. We must then show that π satisfies all the properties listed under (ii). For each of these claimed properties, we will show that if the measure $\pi \in \boldsymbol{M}(n\alpha, n\varrho)$ fails to satisfy the property, then we can find another measure $\bar{\pi} \in \boldsymbol{M}(n\alpha, n\varrho)$ such that $\mathbb{E}[\boldsymbol{Z}^2(\bar{\pi})] \gg \mathbb{E}[\boldsymbol{Z}^2(\pi)]$, proving that π is not a maximizer. Take $\delta \equiv \delta_d \equiv 1/\exp(\exp(\delta))$; we will always arrange for $\|\bar{\pi} - \pi\|_1 \lesssim \delta$. Recall from (62) and (64) that

$$\mathbb{E}[\boldsymbol{Z}^2(\pi)] = \mathbf{c}_{\mathtt{f}}(\pi)\mathbf{G}(\pi), \quad \text{where } \begin{cases} \mathbf{c}_{\mathtt{f}}(\pi) \equiv \mathbf{c}(\pi)\mathbf{j}_{\mathtt{f}}(\pi), \\ \mathbf{G}(\pi) \equiv \sum_{K,F} \mathbf{j}_{\mathtt{1}}(\pi, K, F)\mathbf{j}_{\mathtt{0}}(\pi, K, F). \end{cases}$$

It follows by crude estimates that for $\|\bar{\pi} - \pi\| \lesssim \delta$ we have $\mathbf{G}(\bar{\pi})/\mathbf{G}(\pi) = \exp(O(nd^2\delta))$. In what follows we will, given π, find $\bar{\pi}$ such that the ratio $\mathbf{c}_{\mathtt{f}}(\bar{\pi})/\mathbf{c}_{\mathtt{f}}(\pi)$ is much larger—on the order of $(1/\delta)^{\Omega(n\delta)}$, reflecting the entropy gain. The estimates are derived from the expression (63) which we repeat here for convenience:

$$\frac{\mathbf{c}_{\mathtt{f}}(\pi)}{\mathbf{c}''(\pi)} = n^{O(1)} \frac{d^{n_{\mathtt{f}}}(d-1)^{n_{\mathtt{f!f}}}(n)_{n_{\mathtt{f}}}}{[nd]_{(n_{\mathtt{f}}+n_{\mathtt{f!f}})/2}} \frac{(n_{\mathtt{ff}}-1)!!(\mathsf{u}_{\mathtt{0f}}+n_{\mathtt{f!f}}-1)!!(\mathsf{u}_{\mathtt{f0}}+n_{\mathtt{f!f}}-1)!!}{n_{\mathtt{1f}}!\, n_{\mathtt{f1}}!\, \mathsf{u}_{\mathtt{0f}}!\, \mathsf{u}_{\mathtt{f0}}!\, n_{\mathtt{ff}}!\, n_{\mathtt{f!f}}!}. \tag{74}$$

Note that for $\|\bar{\pi} - \pi\|_1 \lesssim \delta$ we have $\mathbf{c}''(\bar{\pi})/\mathbf{c}''(\pi) = \exp(O(n\delta \log d))$.

(1) Suppose $\mathsf{p}_{\mathtt{0f}} = \mathsf{u}_{\mathtt{0f}}/n \leqslant \delta^2$, and consider the measure

$$\bar{\pi} = \pi + \delta\,(2 \cdot \mathbf{1}\{\mathtt{0f}\} - \mathbf{1}\{\mathtt{00}, \mathtt{01}\}). \tag{75}$$

(As we assumed that $\max\{\pi_{\texttt{f\textvisiblespace}}, \pi_{\texttt{\textvisiblespace f}}\} \leqslant \frac{1}{2}\beta_{\max}$, we certainly have $\max\{\bar{\pi}_{\texttt{f\textvisiblespace}}, \bar{\pi}_{\texttt{\textvisiblespace f}}\} \leqslant \beta_{\max}$ and therefore $\bar{\pi} \in \boldsymbol{M}(n\alpha, n\varrho)$.) Then from (74) we have

$$\frac{\mathbf{c}_{\mathtt{f}}(\bar{\pi})}{\mathbf{c}_{\mathtt{f}}(\pi)} = \frac{d^{O(n\delta)} n^{n\delta}(\mathsf{u}_{\mathtt{0f}}!)}{(\mathsf{u}_{\mathtt{0f}}+n\delta)!} \underbrace{\frac{(n_{\mathtt{0f}}+n_{\mathtt{f!f}}-1)!!}{(n_{\mathtt{f!f}}-1)!!}}_{\geqslant 1} \geqslant d^{O(n\delta)} \left(\frac{1}{\delta}\right)^{n\delta}.$$

Therefore $\mathbb{E}[\boldsymbol{Z}^2(\bar{\pi})] \gg \mathbb{E}[\boldsymbol{Z}^2(\pi)]$, so π is not a maximizer.

(2) If $\pi_{\mathtt{ff}} \leqslant \delta^2$, take $\bar{\pi} = \pi + \delta(2 \cdot \mathbf{1}\{\mathtt{ff}\} - \mathbf{1}\{\mathtt{10}, \mathtt{01}\})$. Then (74) gives

$$\frac{\mathbf{c}_{\mathtt{f}}(\bar{\pi})}{\mathbf{c}_{\mathtt{f}}(\pi)} = \frac{d^{O(n\delta)} n^{n\delta}(n_{\mathtt{ff}}!)}{(n_{\mathtt{ff}}+n\delta)!} \underbrace{\frac{(n_{\mathtt{ff}}+n\delta-1)!!}{(n_{\mathtt{ff}}-1)!!}}_{\geqslant 1} \geqslant d^{O(n\delta)} \left(\frac{1}{\delta}\right)^{n\delta}.$$

We hereafter assume that $\min\{\mathsf{p}_{\mathtt{0f}}, \mathsf{p}_{\mathtt{f0}}, \pi_{\mathtt{ff}}\} \geqslant \delta^2$.

(3) If $\pi_{\mathtt{f!f}} \leqslant \delta^4$, take $\bar\pi = \pi + \gamma(2\cdot\mathbf{1}\{\mathtt{f!f}\} - \mathbf{1}\{\mathtt{0f},\mathtt{f0}\})$ with $\gamma = \delta^3$. Then (74) gives

$$\begin{aligned}\frac{\mathbf{c}_{\mathtt{f}}(\bar\pi)}{\mathbf{c}_{\mathtt{f}}(\pi)} &= d^{O(n\gamma)} \frac{(\mathsf{u}_{\mathtt{0f}})^{n\gamma}(\mathsf{u}_{\mathtt{f0}})^{n\gamma} n_{\mathtt{f!f}}!}{(n_{\mathtt{f!f}}+2n\gamma)!} \frac{(\mathsf{u}_{\mathtt{0f}}+n_{\mathtt{f!f}})^{n\gamma/2}(\mathsf{u}_{\mathtt{0f}}+n_{\mathtt{f!f}})^{n\gamma/2}}{n^{n\gamma}} \\ &= d^{O(n\gamma)} \frac{(\mathsf{p}_{\mathtt{0f}})^{3n\gamma/2}(\mathsf{p}_{\mathtt{f0}})^{3n\gamma/2}}{(\gamma)^{2n\gamma}} \geqslant d^{O(n\gamma)}\left(\frac{\delta^3}{\gamma^2}\right)^{n\gamma} = d^{O(n\gamma)}\left(\frac{1}{\gamma}\right)^{n\gamma}.\end{aligned}$$

(4) If $\pi_{\mathtt{1f}} \leqslant \delta^4$, take $\bar\pi = \pi + \gamma(\mathbf{1}\{\mathtt{1f},\mathtt{0f}\} - 2\cdot\mathbf{1}\{\mathtt{ff}\})$ with $\gamma = \delta^3$. Then (74) gives

$$\frac{\mathbf{c}_{\mathtt{f}}(\bar\pi)}{\mathbf{c}_{\mathtt{f}}(\pi)} = d^{O(n\gamma)}\left(\frac{\pi_{\mathtt{ff}}}{\gamma}\right)^{n\gamma} \geqslant d^{O(n\gamma)}\left(\frac{\delta^2}{\gamma}\right)^{n\gamma} = d^{O(n\gamma)}\left(\frac{1}{\gamma}\right)^{n\gamma/3}.$$

This concludes the proof of (c). The proofs of (a) and (b) are similar but much simpler: for (a), going from $(\pi_{\mathtt{0}},\pi_{\mathtt{1}},\pi_{\mathtt{f}}) = (1-\alpha,\alpha,0)$ to $(\bar\pi_{\mathtt{0}},\bar\pi_{\mathtt{1}},\bar\pi_{\mathtt{f}}) = (1-\alpha-\delta,\alpha-\delta,2\delta)$ gives a large gain in the first moment for δ sufficiently small (depending on d). For (b), if $\pi_{\mathtt{00}} = 0$, then there is a large gain going from π to $\bar\pi = \pi + \delta(\mathbf{1}\{\mathtt{00},\mathtt{11}\} - \mathbf{1}\{\mathtt{01},\mathtt{10}\})$. $\square$

We now turn our attention to the maximizer(s) $\boldsymbol{h}$ for $\mathbf{\Phi}_2$ restricted to $\mathbf{\Delta}_2[\alpha]\setminus\mathsf{EQ}[\alpha]$:

$$\boldsymbol{h} \in \arg\max\{\mathbf{\Phi}_2(\boldsymbol{h}') : \boldsymbol{h}' \in \mathbf{\Delta}_2[\alpha]\setminus\mathsf{EQ}[\alpha]\}. \tag{76}$$

The tuple (π, K, F) introduced in (64) is simply a projection of $\boldsymbol{h}$, and up to now we have proved estimates concerning

$$\lim_{n\to\infty} n^{-1}\log\mathbb{E}[\boldsymbol{Z}^2(\pi,K,F)] = \max\{\mathbf{\Phi}_2(\boldsymbol{h}) : \boldsymbol{h} \text{ projects to } (\pi,K,F)\}.$$

Of course, $\boldsymbol{h}$ contains richer information than (π, K, F), and we now turn to estimates for this additional information. To this end, let us say that a $\mathtt{0f}$-variable is of type $\mathtt{0}^{\mathtt{1}}\mathtt{f}$ if the matched partner of v lies in $V_{\mathtt{1f}}$, and of type $\mathtt{0}^{\mathtt{Z}}\mathtt{f}$ otherwise. (Recall that we used $U_{\mathtt{0f}}$ to denote the subset of $\mathtt{0}_{\mathtt{Z}}\mathtt{f}$-variables.) Recall also from (32) that the $\mathtt{0}$ variable spin can be further subdivided into $\mathtt{0}^{\mathtt{r}}$ and $\mathtt{0}^{\mathtt{s}}$; and let us further subdivide $\mathtt{0}^{\mathtt{1}}\mathtt{f}$ into $\mathtt{0}^{\mathtt{r},\mathtt{1}}\mathtt{f}$ and $\mathtt{0}^{\mathtt{s},\mathtt{1}}\mathtt{f}$. Let Π be the measure induced by $\boldsymbol{h}$ on the expanded alphabet

$$\begin{aligned}\mathscr{Q} \equiv (\{\mathtt{0}^{\mathtt{r}},\mathtt{0}^{\mathtt{s}},\mathtt{1}\}\times\{\mathtt{0}^{\mathtt{r}},\mathtt{0}^{\mathtt{s}}\}) &\cup (\{\mathtt{0}^{\mathtt{r}},\mathtt{0}^{\mathtt{s}}\}\times\{\mathtt{0}^{\mathtt{r}},\mathtt{0}^{\mathtt{s}},\mathtt{1}\}) \cup \{\mathtt{ff},\mathtt{f!f}\} \\ &\cup (\{\mathtt{0}_{\mathtt{Z}},\mathtt{0}^{\mathtt{r},\mathtt{1}},\mathtt{0}^{\mathtt{s},\mathtt{1}}\}\times\{\mathtt{f}\}) \cup (\{\mathtt{f}\}\times\{\mathtt{0}_{\mathtt{Z}},\mathtt{0}^{\mathtt{r},\mathtt{1}},\mathtt{0}^{\mathtt{s},\mathtt{1}}\}).\end{aligned}$$

Thus $\boldsymbol{h}$ projects to Π, which in turn projects to a measure π on $\mathscr{P}$ as considered above.

LEMMA 4.8. *Let $\boldsymbol{h}$ be a maximizer as in (76), with edge marginal $\bar h$. Let Π be the measure on $\mathscr{Q}$ induced by $\boldsymbol{h}$. Then* $\operatorname{supp}\Pi = \mathscr{Q}$, *and*

$$\bar h(\sigma^1,\sigma^2) = \bar h(\sigma^2,\sigma^1) > 0 \quad \textit{for any } \sigma^1 \in \{\mathtt{10},\mathtt{1f},\mathtt{01},\mathtt{f1}\} \textit{ and } \sigma^2 = \mathtt{11}. \tag{77}$$

Proof. Recall from Lemma 4.5 that there are K edges between $V_{\texttt{01}}$ and $V_{\texttt{10}}$, where we have $K/|V_{\texttt{01}}| \asymp \varepsilon \log d \asymp K/|V_{\texttt{10}}|$ or else the contribution to $\mathbb{E}[\boldsymbol{Z}^2(\pi)]$ is negligible. Simply by the pigeonhole principle this implies that $V_{\texttt{01}}$ has a positive fraction of variables with more than two edges going to $V_{\texttt{10}}$, meaning $\Pi_{\texttt{0}^{\texttt{r}}\texttt{1}} \geqslant \delta$ for some constant $\delta \equiv \delta_d > 0$, which can depend on d but not on n.

On the other hand, conditional on each $v \in V_{\texttt{01}}$ receiving at least two edges from $V_{\texttt{10}}$, these edges are uniformly assigned. (As in (61), we are considering the distribution of the random graph G given a planted solution $(\underline{\omega}, \underline{m})$ consistent with π.) Thus, with this conditioning, for δ small enough it is exponentially unlikely for fewer than $n\delta$ vertices in $V_{\texttt{01}}$ to receive exactly two neighbors from $V_{\texttt{10}}$. We therefore conclude that at any maximizer of $\boldsymbol{\Phi}_2$ we must have $\Pi_{\texttt{0}^{\texttt{s}}\texttt{1}} \geqslant \delta$. By symmetry we also have $\min\{\Pi_{\texttt{10}^{\texttt{r}}}, \Pi_{\texttt{10}^{\texttt{s}}}\} \geqslant \delta$.

Similarly, given the matching $\underline{m}$, $V_{\texttt{0f}}$ is divided into variables of type $\texttt{0}_{\texttt{1}}\texttt{f}$ and $\texttt{0}_{\texttt{Z}}\texttt{f}$. We saw in Proposition 4.7 (c) that each of these types occupies a positive density of variables. Each $v \in V_{\texttt{0f}}$ has $d-1$ remaining half-edges that do not participate in the matching, and F_1 of these will match to half-edges coming from $V_{\texttt{1Z}}$. These edges are uniformly assigned, conditioned on the requirement that each $\texttt{0}$ is forced by at least two neighboring $\texttt{1}$'s. We saw in Lemma 4.5 that $F_1/|V_{\texttt{0f}}|$ equals $\varepsilon \log d$ up to a factor of $\log\log d$, or else the contribution to $\mathbb{E}[\boldsymbol{Z}^2(\pi)]$ is negligible. Thus it follows by the same considerations as above that $\Pi_\omega \geqslant \delta$ for all $\omega \in (\{\texttt{0}_{\texttt{Z}}, \texttt{0}^{\texttt{r},\texttt{1}}, \texttt{0}^{\texttt{s},\texttt{1}}\} \times \{\texttt{f}\}) \cup (\{\texttt{f}\} \times \{\texttt{0}_{\texttt{Z}}, \texttt{0}^{\texttt{r},\texttt{1}}, \texttt{0}^{\texttt{s},\texttt{1}}\})$. The assertion in (77) is an immediate consequence.

Lastly, recall that $V_{\texttt{00}}$ receives $n\varrho d$ edges from $V_{\texttt{11}}$, and $\asymp n\varepsilon \log d$ edges each from $V_{\texttt{1Z}}$ and $V_{\texttt{Z1}}$. These edges are uniformly assigned, conditioned on the requirement that each $\texttt{0}$ is forced by at least two neighboring $\texttt{1}$'s. Thus, with this conditioning, for δ small enough it is exponentially unlikely to have $|V_{\eta\eta'}| \leqslant n\delta$ for any $\eta, \eta' \in \{\texttt{0}^{\texttt{r}}, \texttt{0}^{\texttt{s}}\}$. This proves the assertion that Π has full support. $\square$

LEMMA 4.9. *Let $\boldsymbol{h}$ be a maximizer as in (76), and let $\bar{h}$ be the edge marginal. Let $\bar{f}$ be the projection of $\bar{h}$ under the mapping $\mathsf{proj}_{\mathrm{v}}$ of (30). Then*

$$\bar{f}(\sigma^1 \in R^1, \sigma^2 \in R^2) > 0$$

for all subsets R^1 and R^2 in the partition $\{\texttt{10}\}, \{\texttt{01}\}, \{\overline{\texttt{ff}}\}, \texttt{fZ} = \{\texttt{f0}, \texttt{ff}\}, \texttt{0Z} = \{\texttt{00}, \texttt{0f}\}$.

Proof. In the pair vertex-auxiliary model, the spin on a variable-clause edge can be written as $\tau \equiv \boldsymbol{oi}$, where $\boldsymbol{o} \equiv (\boldsymbol{o}^1, \boldsymbol{o}^2)$ is the pair of variable-to-clause messages, and $\boldsymbol{i} \equiv (\boldsymbol{i}^1, \boldsymbol{i}^2)$ is the pair of clause-to-variable messages. All messages $\boldsymbol{o}^i$ and $\boldsymbol{i}^i$ take values in $\{\texttt{0}, \texttt{1}, \texttt{f}, \bar{\texttt{f}}\}$, and we write $\sigma^i \equiv \boldsymbol{o}^i \boldsymbol{i}^i$. Abbreviate $\bar{f}(R^1, R^2) \equiv \bar{f}(\sigma^1 \in R^1, \sigma^2 \in R^2)$. From (32) and the definition (30) of $\mathsf{proj}_{\mathrm{v}}$, we see that if $\underline{\sigma}$ is a valid configuration of the vertex-auxiliary model, then the tuple of spins $\underline{\dot{\sigma}}_v$ around a variable $v \in V$ must satisfy one of the following:

$$\dot{\underline{\sigma}}_v \in \begin{cases} (\mathtt{10}^d), & \text{if } \eta_v = \mathtt{1}, \\ \mathrm{Per}[(\overline{\mathtt{ff}}, \mathtt{fZ}^{d-1})], & \text{if } \eta_v = \mathtt{f}, \\ \mathrm{Per}[(\mathtt{01}^k, \mathtt{0Z}^{d-k})_{2\leqslant k\leqslant d}], & \text{if } \eta_v = \mathtt{0}. \end{cases}$$

Note that $\bar{f}(\mathtt{0Z},\mathtt{0Z})\geqslant\bar{f}(\mathtt{00},\mathtt{00})>0$, simply because the vast majority of edges in the graph must be internal to $V_{\mathtt{00}}$—there are $nd/2$ edges total, and only $O(n\log d)$ can be incident to $\{V_\omega\}_{\omega\neq\mathtt{00}}$. We now finish the proof by making deductions from the conclusions of Lemma 4.5 and Proposition 4.7. The conclusions are summarized in the following table:

	$\mathtt{10}$	$\overline{\mathtt{ff}}$	$\mathtt{fZ}$	$\mathtt{01}$	$\mathtt{0Z}$
$\mathtt{10}$	$\pi_{\mathtt{11}}>0$	$\pi_{\mathtt{1f}}>0$	$\pi_{\mathtt{1f}}>0$	$K>0$	$K\leqslant K_{\mathrm{ubd}}$
$\overline{\mathtt{ff}}$	$\pi_{\mathtt{f1}}>0$	$\pi_{\mathtt{ff}}>0$	$\pi_{\mathtt{f!f}}>0$	$\pi_{\mathtt{f1}}>0$	$\mathsf{p}_{\mathtt{f0}}>0$
$\mathtt{fZ}$	$\pi_{\mathtt{f1}}>0$	$\pi_{\mathtt{f!f}}>0$	$\pi_{\mathtt{ff}}>0$	$\mathsf{p}_{\mathtt{f0}}>0$	$\mathsf{p}_{\mathtt{f0}}>0$
$\mathtt{01}$	$K>0$	$\pi_{\mathtt{1f}}>0$	$\mathsf{p}_{\mathtt{0f}}>0$	$\pi_{\mathtt{11}}>0$	$K\leqslant K_{\mathrm{ubd}}$
$\mathtt{0Z}$	$K\leqslant K_{\mathrm{ubd}}$	$\mathsf{p}_{\mathtt{0f}}>0$	$\mathsf{p}_{\mathtt{0f}}>0$	$K\leqslant K_{\mathrm{ubd}}$	(see above)

where the row index is R^1, the column index is R^2, and the (R^1,R^2) entry in the table gives the explanation for why $\bar{f}(R^1,R^2)$ is positive. Most of the entries are self-explanatory, but we supply some detail for the three entries labeled $\mathsf{p}_{\mathtt{0f}}>0$. Recall that $\mathsf{p}_{\mathtt{0f}}=\mathsf{u}_{\mathtt{0f}}/n$, where $\mathsf{u}_{\mathtt{0f}}=n_{\mathtt{0f}}-n_{\mathtt{1f}}$ counts the variables $v\in V_{\mathtt{0f}}$ whose partner w under $\underline{m}^2$ is not in $V_{\mathtt{1f}}$. Therefore, if a denotes the clause joining v to w, on the edge (av) we see the spins $\sigma^1\in\mathtt{0Z}$, $\sigma^2=\overline{\mathtt{ff}}$, proving the first assertion that $\bar{f}(\mathtt{0Z},\overline{\mathtt{ff}})>0$. Further, at least two non-matching edges leaving v must lie in $V_{\mathtt{1Z}}$, proving $\bar{f}(\mathtt{01},\mathtt{fZ})>0$. Finally, recalling from Lemma 4.5 (see (65)) that the contribution to $\mathbb{E}[\boldsymbol{Z}^2(\pi)]$ is negligible for $F_1\geqslant F_{1,\mathrm{ubd}}$, we see that most of the edges leaving $V_{\mathtt{0f}}$ must go to $V_{\mathtt{ZZ}}$, proving $\bar{f}(\mathtt{0Z},\mathtt{fZ})>0$. □

COROLLARY 4.10. *In the setting of Lemma* 4.9, $\operatorname{supp}\bar{f}=(\mathscr{M}_{\mathrm{v}})^2$.

Proof. As above, use $\tau\equiv(\sigma^1,\sigma^2)$ to denote an edge spin in the pair vertex-auxiliary model. We can express $\tau\equiv\boldsymbol{oi}$, where $\boldsymbol{o}\equiv(\boldsymbol{o}^1,\boldsymbol{o}^2)$, $\boldsymbol{i}\equiv(\boldsymbol{i}^1,\boldsymbol{i}^2)$, and $\sigma^i\equiv\boldsymbol{o}^i\boldsymbol{i}^i$. If $\bar{f}(\boldsymbol{oi})=0$, then Lemma 4.9 implies $\bar{f}(\boldsymbol{oi}')>0$ for some $\boldsymbol{i}'\neq\boldsymbol{i}$. Moreover $\bar{f}(\boldsymbol{io})=0$ by symmetry of $\bar{f}$ under R, so another application of Lemma 4.9 gives $\bar{f}(\boldsymbol{io}')>0$ for some $\boldsymbol{o}'\neq\boldsymbol{o}$, and again by symmetry under R we have $\bar{f}(\boldsymbol{o}'\boldsymbol{i})>0$.

This means that under the configuration $\underline{\tau}=(\underline{\sigma}^1,\underline{\sigma}^2)$ we have a positive density of variable-clause edges (av) with $\tau_{zv}=\boldsymbol{oi}'$, as well as of variable-clause edges (bw) with $\tau_{bw}=\boldsymbol{o}'\boldsymbol{i}$. If we cut the edges (av), (bw) and replace them by the switched edges (aw), (bv), then a valid configuration on the new graph is given by taking $\tau_{bv}=\boldsymbol{oi}$ and $\tau_{aw}=\boldsymbol{o}'\boldsymbol{i}'$, and keeping the same spins as before on all the other edges.

In preceding arguments (see (61) and the proof of Lemma 4.8), we drew conclusions about $\mathbb{E}[\boldsymbol{Z}^2(\pi)]$ by considering the random graph G given a planted configuration $(\underline{\omega},\underline{m})$

consistent with π. We can now draw further conclusions about $\mathbb{E}[\boldsymbol{Z}^2(\boldsymbol{h})]$ by considering a richer planted configuration, as follows. Suppose a variable v has incident spins $\underline{\dot{\sigma}}_v$, and take a half-edge incident to v with outgoing messages $\boldsymbol{o}^i$ and incoming messages $\boldsymbol{i}^i$. Define modified incoming messages $\boldsymbol{m}^i \in \{\mathtt{0}, \mathtt{1}, \bar{\mathtt{f}}, \mathtt{Z}\}$, $i=1,2$, by setting $\boldsymbol{m}^i = \boldsymbol{i}^i$ *unless* $\boldsymbol{i}^i$ is free to take either value in $\{\mathtt{0}, \mathtt{f}\}$, in which case we define $\boldsymbol{m}^i = \mathtt{Z}$.

Each $v \in V$ has d incident half-edges, and we now plant messages $\boldsymbol{o}$ and $\boldsymbol{m}$ on all half-edges in a manner consistent with $\boldsymbol{h}$. We can then consider sampling a random graph given the planted configuration: if v has an incident half-edge labeled $\boldsymbol{o}_v\boldsymbol{m}_v$ and w has an incident half-edge labeled $\boldsymbol{o}_w\boldsymbol{m}_w$, we can join up the half-edges provided $\boldsymbol{o}_v \in \boldsymbol{m}_w$ and $\boldsymbol{o}_w \in \boldsymbol{m}_v$. If a denotes the clause joining these half-edges in the resulting graph, then we will have $\tau_{av} = \boldsymbol{o}_v\boldsymbol{o}_w$ and $\tau_{aw} = \boldsymbol{o}_w\boldsymbol{o}_v$.

Now take $\boldsymbol{i}$, $\boldsymbol{i}'$, $\boldsymbol{o}$, and $\boldsymbol{o}'$ as above, let $\boldsymbol{O}$ be the $\boldsymbol{m}$-message associated with $\boldsymbol{o}$ and $\boldsymbol{o}'$, and let $\boldsymbol{I}$ be the $\boldsymbol{m}$-message associated with $\boldsymbol{i}$ and $\boldsymbol{i}'$. In the planted configuration we have a positive density of half-edges with each of the labels $\boldsymbol{oI}$, $\boldsymbol{o}'\boldsymbol{I}$, $\boldsymbol{iO}$, and $\boldsymbol{i}'\boldsymbol{O}$. If we take a uniformly random matching of half-edges that respects these labels, it is exponentially unlikely to have no matchings between $\boldsymbol{oI}$ and $\boldsymbol{iO}$. This contradicts the assumption $\bar{f}(\boldsymbol{oi})=0$, and the result follows. $\square$

COROLLARY 4.11. *In the setting of Lemma* 4.9, $\operatorname{supp} \bar{h} = \mathscr{M}^2$.

Proof. We will show that

$$\bar{h}(\sigma^1, \sigma^2) > 0 \quad \text{whenever } \sigma^1 \text{ or } \sigma^2 \text{ lies in } \{\mathtt{10}, \mathtt{1f}\}. \tag{78}$$

It then follows by symmetry of $\bar{h}$ under R (as in (32)) that

$$\bar{h}(\sigma^1, \sigma^2) > 0 \quad \text{whenever } \sigma^1 \text{ or } \sigma^2 \text{ lies in } \{\mathtt{01}, \mathtt{f1}\}.$$

Finally, if neither σ^i lies in $\{\mathtt{10}, \mathtt{1f}, \mathtt{01}, \mathtt{f1}\}$, then it follows directly from Corollary 4.10 that $\bar{h}(\sigma^1, \sigma^2) > 0$, which gives the claimed result. Thus it remains to establish (78). Recall that, by (77), we already know that $\bar{h}(\sigma^1, \sigma^2) > 0$ for $\sigma^1 \in \{\mathtt{10}, \mathtt{1f}\}$ and $\sigma^2 = \mathtt{11}$.

As before (cf. (61)), consider sampling the random graph G conditioned on the event that some planted configuration $(\underline{\omega}, \underline{m})$ is valid on G. We claim that, for any $\eta \in \{\mathtt{0}^{\mathtt{r}}, \mathtt{0}^{\mathtt{s}}\}$ and any $\eta', \eta'' \in \{\mathtt{0}, \mathtt{f}\}$, it is exponentially unlikely for the planted graph to have fewer than $n\delta$ non-matching edges between $V_{\mathtt{1}\eta'}$ and $V_{\eta\eta''}$. (As before, δ denotes a small positive constant, which can depend on d but not n.) To see this, suppose there are no edges between $V_{\mathtt{1}\eta'}$ and $V_{\mathtt{0}^{\mathtt{s}}\eta''}$, or equivalently that $\bar{h}(\mathtt{f1}, \eta'\eta'')=0$. Since $\bar{f}$ has full support by Corollary 4.10, we must have $\bar{h}(\mathtt{01}, \eta'\eta'') > 0$, which means that there is a positive density of edges between $v \in V_{\mathtt{1}\eta'}$ and $w \in V_{\mathtt{0}^{\mathtt{r}}\eta''}$. On the other hand, since $\Pi_{\mathtt{0}^{\mathtt{s}}\eta''} > 0$ by Lemma 4.8, there must be a positive density of edges between $v' \in V_{\mathtt{1}\eta'''}$ and $w' \in V_{\mathtt{0}^{\mathtt{s}}\eta''}$

for some $\eta'''\in\{\mathtt{0},\mathtt{1},\mathtt{f}\}$. If we cut the edges (vw) and $(v'w')$, and form the switched edges (vw') and $(v'w)$, then $(\underline{\omega},\underline{m})$ remains a valid configuration for the switched graph. It follows that in the randomly planted graph it is exponentially unlikely that there are fewer than $n\delta$ non-matching edges between $V_{\mathtt{1}\eta'}$ and $V_{\mathtt{0^s}\eta''}$. Exactly the same argument applies to show that we must also have $\geqslant n\delta$ non-matching edges between $V_{\mathtt{1}\eta'}$ and $V_{\mathtt{0^r}\eta''}$, and this proves that $\bar{h}(\sigma^1,\sigma^2)>0$ for $\sigma^1\in\{\mathtt{10},\mathtt{1f}\}$ and $\sigma^2\in\{\mathtt{00},\mathtt{0f},\mathtt{f0},\mathtt{ff}\}$.

If we condition on the half-edges incident to $V_{\mathtt{10}}$ and $V_{\mathtt{01}}$ that will participate in the K edges between them, then any matching of these half-edges is equally likely, so with high probability there is a positive density of edges between $V_{\mathtt{1}\eta}$ and $V_{\eta'\mathtt{1}}$ for all $\eta,\eta'\in\{\mathtt{0^r},\mathtt{0^s}\}$. This implies that $\bar{h}(\sigma^1,\sigma^2)>0$ for $\sigma^1\in\{\mathtt{10},\mathtt{1f}\}$ and $\sigma^2\in\{\mathtt{01},\mathtt{f1}\}$. Similarly, there is a positive density of edges between $V_{\mathtt{11}}$ and $V_{\eta\eta'}$ for all $\eta,\eta'\in\{\mathtt{0^r},\mathtt{0^s}\}$, which implies that $\bar{h}(\sigma^1,\sigma^2)>0$ if both $\sigma^i\in\{\mathtt{10},\mathtt{1f}\}$. This concludes the proof. □

PROPOSITION 4.12. *Let $\boldsymbol{h}$ be a maximizer as in (76). Then* $\operatorname{supp}\boldsymbol{h}=\operatorname{supp}\varphi_2$.

Proof. As we have noted before (Remark 3.8), the functional form of $\boldsymbol{\Phi}_2$ implies that at the optimizer $\boldsymbol{h}$ the functions $\dot{\boldsymbol{h}}$ and $\hat{\boldsymbol{h}}$ must be symmetric, with

$$\boldsymbol{\Phi}_2(\boldsymbol{h})=H(\dot{\boldsymbol{h}})-\tfrac{1}{2}dH(\bar{h}).$$

For $\boldsymbol{\delta}\equiv(\dot{\boldsymbol{\delta}},\hat{\boldsymbol{\delta}})$ such that $\boldsymbol{h}+t\boldsymbol{\delta}$ lies in $\boldsymbol{\Delta}_2[\alpha]$ for $t\geqslant 0$ small, consider

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

To show that $\boldsymbol{h}$ is not a maximizer it suffices to exhibit $\partial^{\log}(\boldsymbol{h};\boldsymbol{\delta})>0$ for some $\boldsymbol{\delta}$.

In particular, for any $\boldsymbol{h}\in\boldsymbol{\Delta}_2[\alpha]$ and with ${}^{\otimes}\boldsymbol{h}[\alpha]$ as in the statement of Theorem 4.1, it follows by convexity of the space $\boldsymbol{\Delta}_2[\alpha]$ that

$$\boldsymbol{h}+t({}^{\otimes}\boldsymbol{h}[\alpha]-\boldsymbol{h})\in\mathsf{IND}^{\circ}[\alpha]$$

for $t>0$ small. Therefore, if $\boldsymbol{h}$ is a maximizer such that the edge marginal $\bar{h}$ has full support $\operatorname{supp}\bar{h}=\mathscr{M}^2$, then necessarily $\operatorname{supp}\boldsymbol{h}=\operatorname{supp}\varphi_2$, since otherwise we would have $\partial^{\log}(\boldsymbol{h};{}^{\otimes}\boldsymbol{h}[\alpha]-\boldsymbol{h})>0$. □

The following proposition is established by similar (but much easier) arguments as for Proposition 4.12, so we omit the proof.

PROPOSITION 4.13. *Let $\boldsymbol{h}$ be a maximizer of $\boldsymbol{\Phi}$ on $\boldsymbol{\Delta}[\alpha]$. Then* $\operatorname{supp}\boldsymbol{h}=\operatorname{supp}\varphi$.

4.3. Near-independence regime

In this subsection we complete our analysis of the near-independence regime $\mathsf{IND}[\alpha]$ to prove the following result.

PROPOSITION 4.14. *The unique maximizer of the restriction of* $\mathbf{\Phi}_2$ *to* $\mathsf{IND}[\alpha]$ *is the product measure* ${}^{\otimes}\boldsymbol{h}[\alpha]\equiv{}^{\star}\boldsymbol{h}[\alpha]\otimes{}^{\star}\boldsymbol{h}[\alpha]$.

Recall that a pair frozen configuration is encoded by $(\underline{\omega},\underline{m})$, with $\underline{\omega}$ as defined in (57) and $\underline{m}\equiv(\underline{m}^1,\underline{m}^2)$. Given the graph G, the configuration $(\underline{\omega},\underline{m})$ corresponds bijectively to a pair message configuration $\underline{\tau}\equiv(\underline{\sigma}^1,\underline{\sigma}^2)$ on G. We write

$$(G,\underline{\omega},\underline{m})\in\boldsymbol{h}$$

to indicate that the corresponding $\underline{\tau}$ has pair empirical measure $\boldsymbol{h}$ (see (56)). We now study the maximization of $\mathbf{\Phi}_2(\boldsymbol{h})$ over $\boldsymbol{h}\in\mathsf{IND}[\alpha]\subset\mathbf{\Delta}_2[\alpha]$ (Definition 4.2). Recall that $\boldsymbol{h}$ induces a measure π on $\mathscr{P}$, with $\pi\in\boldsymbol{M}(n\alpha,n\varrho)$ for some $\varrho\leqslant d^{-1.1}$.

LEMMA 4.15. *Let* $\boldsymbol{h}$ *be any maximizer of the restriction of* $\mathbf{\Phi}_2$ *to* $\mathsf{IND}[\alpha]$, *and let* π *be the induced measure on* $\mathscr{P}$. *Then* π *must satisfy*

$$\frac{\pi_{\mathtt{1f}}}{\pi_{\mathtt{0f}}\pi_{\mathtt{10}}},\frac{\pi_{\mathtt{f1}}}{\pi_{\mathtt{f0}}\pi_{\mathtt{01}}},\frac{d\pi_{\mathtt{ff}}}{\pi_{\mathtt{0f}}\pi_{\mathtt{f0}}},\frac{\pi_{\mathtt{f!f}}}{\pi_{\mathtt{0f}}\pi_{\mathtt{f0}}}\lesssim 1$$

as well as $\pi_{\mathtt{0f}}\leqslant(\log d)^{O(1)}d^{-y}$.

Proof. Let $\boldsymbol{h}$ be the purported maximizer of $\mathbf{\Phi}_2$ on $\mathsf{IND}[\alpha]$. Recall from Proposition 4.12 that $\boldsymbol{h}$ has full support, so $\min\{\dot{\boldsymbol{h}},\hat{\boldsymbol{h}}\}\geqslant\delta$ for some small positive constant δ, which depends on d but not on n. We will define a switching operation ς that maps a valid tuple $(G,\underline{\omega},\underline{m})\in\boldsymbol{h}$ to a set of valid tuples $(G',\underline{\omega}',\underline{m}')$, with the property that for each such $(G',\underline{\omega}',\underline{m}')$ the corresponding $\boldsymbol{h}'$ still lies in $\mathsf{IND}[\alpha]$ and satisfies $\|\boldsymbol{h}'-\boldsymbol{h}\|_1\leqslant\delta'$, where $\delta'\leqslant\exp(-\exp(1/\delta))$. Let $\varsigma^{-1}(G',\underline{\omega}',\underline{m}')$ denote the preimage of $(G',\underline{\omega}',\underline{m}')$ under ς. Suppose we prove under these conditions that

$$|\varsigma(G,\underline{\omega},\underline{m})|\geqslant a\quad\text{and}\quad|\varsigma^{-1}(G,\underline{\omega},\underline{m})|\leqslant b.$$

Then, using $\mathbb{P}(G)=\mathbb{P}(G')$, we have

$$\begin{aligned}
a\mathbb{E}\boldsymbol{Z}^2(\boldsymbol{h})&\leqslant\sum_{(G,\underline{\omega},\underline{m})\in\boldsymbol{h}}\mathbb{P}(G)\sum_{(G',\underline{\omega}',\underline{m}')}\mathbf{1}\{(G',\underline{\omega}',\underline{m}')\in\varsigma(G,\underline{\omega},\underline{m})\}\\
&=\sum_{(G',\underline{\omega}',\underline{m}')}\mathbb{P}(G')\sum_{(G,\underline{\omega},\underline{m})\in\boldsymbol{h}}\mathbf{1}\{(G',\underline{\omega}',\underline{m}')\in\varsigma(G,\underline{\omega},\underline{m})\}\\
&\leqslant b\sum_{\boldsymbol{h}'\in\mathsf{IND}[\alpha]}\mathbb{E}\boldsymbol{Z}^2(\boldsymbol{h}')\leqslant n^{O(1)}b\,\mathbb{E}\boldsymbol{Z}^2(\boldsymbol{h}),
\end{aligned}\tag{79}$$

where the last step is by the assumed optimality of $\boldsymbol{h}$. This gives the bound $a\leqslant n^{O(1)}b$. To prove the stated bounds we will apply this argument with a few different choices of ς, which are explained below and illustrated in Figure 3. Recall the notation in (58).

(1) *Figure* 3 (a): *estimate for* $\mathtt{1f}$*-variables.*

For $\omega\in\mathscr{P}$ let W_ω denote the subset of variables $v\in V_\omega$ with no neighbors in $V_{\mathtt{11}}\cup V_{\mathtt{f1}}$. Let $W_\omega(l)$ be the subset of such variables with exactly l neighbors in $V_{\mathtt{1Z}}\cup V_{\mathtt{fZ}}$. For $\omega=\mathtt{00}$, note that $n^{-1}|W_{\mathtt{00}}(l)|\geqslant\min\dot{\boldsymbol{h}}\geqslant\delta$ for each $2\leqslant l\leqslant d-2$ by Proposition 4.12. Moreover, since $|V_{\mathtt{11}}\cup V_{\mathtt{f1}}|\leqslant n(\varrho+\beta_{\max})$, we clearly have $|W_{\mathtt{00}}|\geqslant\frac{3}{4}n_{\mathtt{00}}$.

Fix any $2\leqslant l\leqslant d-2$. Given $(G,\underline{\omega},\underline{m})$, choose a subset $S_{\mathtt{1f}}\subseteq V_{\mathtt{1f}}$ of size $n\delta'$. For each $v\in S_{\mathtt{1f}}$, select a corresponding $w\in W_{\mathtt{00}}(l)$, and write $S_{\mathtt{00}}(l)$ for the subset of chosen w. We require that for each $x\neq x'$ in $S_{\mathtt{1f}}\cup S_{\mathtt{00}}(l)$ the distance between x and x' is $\geqslant 5$. By definition, w has exactly l neighbors in $V_{\mathtt{1Z}}\cup V_{\mathtt{fZ}}$, which we denote $w_1,...,w_l$. Meanwhile v has $d-1$ neighbors in $V_{\mathtt{0Z}}$; choose l of these and denote them $v_1,...,v_l$.

Now for each $1\leqslant j\leqslant l$, cut the edges (vv_j) and (ww_j), and replace them by the switched edges (vw_j) and (wv_j). Then set $\omega'_v=\mathtt{0f}$ and $\omega'_w=\mathtt{10}$. Repeat this for each chosen pair (v,w), then set $\omega'_u=\omega_u$ for all $u\notin S_{\mathtt{1f}}\cup S_{\mathtt{00}}(l)$. This defines a map ς satisfying

$$|\varsigma(G,\underline{\omega},\underline{m})|=[(n\delta')!]^{-1}\exp\biggl(O\biggl(\frac{nd^5(\delta')^2}{\delta}\biggr)\biggr)((d-1)_l|V_{\mathtt{1f}}|\,|W_{\mathtt{00}}(l)|)^{n\delta'},$$

$$|\varsigma^{-1}(G,\underline{\omega},\underline{m})|\leqslant[(n\delta')!]^{-1}\exp\biggl(O\biggl(\frac{nd^5(\delta')^2}{\delta}\biggr)\biggr)((d-1)_l|V_{\mathtt{10}}|\,|W_{\mathtt{0f}}(l)|)^{n\delta'},$$

where the factor $\exp(O(nd^5(\delta')^2/\delta))$ accounts for the restriction that the $2n\delta'$ chosen vertices must lie at pairwise distance at least five. Applying (79) gives

$$(|V_{\mathtt{1f}}|\,|W_{\mathtt{00}}(l)|)^{n\delta'}\leqslant n^{O(1)}\exp\biggl(O\biggl(\frac{nd^5(\delta')^2}{\delta}\biggr)\biggr)(|V_{\mathtt{10}}|\,|W_{\mathtt{0f}}(l)|)^{n\delta'},$$

so we conclude $|V_{\mathtt{1f}}|\,|W_{\mathtt{00}}(l)|\lesssim|V_{\mathtt{10}}|\,|W_{\mathtt{0f}}(l)|$. Summing over l and rearranging gives

$$\pi_{\mathtt{1f}}\lesssim\frac{|V_{\mathtt{10}}|\,|V_{\mathtt{0f}}|}{|W_{\mathtt{00}}|}\lesssim\pi_{\mathtt{10}}\pi_{\mathtt{0f}}.$$

(2) *Figure* 3 (b): *estimate for* $\mathtt{ff}$*-variables.*

Let $F_{\mathtt{00}}$ denote the set of pairs $(x,y)\in(V_{\mathtt{00}})^2$ such that x neighbors y, and neither x nor y has any neighbors in $V_{\mathtt{11}}$. The number of pairs $(x,y)\in(V_{\mathtt{00}})^2$ such that either x or y neighbors $V_{\mathtt{11}}$ is $\leqslant 2nd^2\varrho$, so $|F_{\mathtt{00}}|\geqslant\frac{1}{2}nd$. Let $F_{\mathtt{00}}(k,l)$ denote the subset of such pairs such that x has exactly k neighbors in $V_{\mathtt{1Z}}$, and y has exactly l neighbors in $V_{\mathtt{Z1}}$. It is clear that we can adjust δ so that $|F_{\mathtt{00}}(k,l)|\geqslant n\delta$ for each $2\leqslant k,l\leqslant d-2$.

Now fix some $2\leqslant k,l\leqslant d-2$, and choose a subset $S_{\mathtt{00}}(k,l)\subseteq F_{\mathtt{00}}(k,l)$ of size $n\delta'$. For each $(x,y)\in S_{\mathtt{00}}(k,l)$ choose a corresponding pair $(v,w)\in(V_{\mathtt{ff}})^2$ that are matched under $\underline{m}$, and let $S_{\mathtt{ff}}$ denote the subset of chosen pairs (v,w). We require that for any two chosen pairs $(u,u')\neq(u'',u''')$ in $S_{\mathtt{00}}(k,l)\cup S_{\mathtt{ff}}$ the graph distance between u and u'' is at least five. By definition, x has exactly k neighbors $x_1,...,x_k$ with spin $\texttt{1Z}$, and y has exactly l neighbors $y_1,...,y_l$ with spin $\texttt{Z1}$. Choose k neighbors $v_1,...,v_k$ of v (none equal to w), and choose l neighbors $w_1,...,w_l$ of w (none equal to v).

Now perform the following operation: cut the edges (vw) and (xy), and form the switched edges (vy) and (wx). Then cut (vv_j) and (xx_j), and form (vx_j) and (vx_j) for each $1\leqslant j\leqslant k$; also cut (ww_j) and (yy_j), and form (wy_j) and (wy_j) for each $1\leqslant j\leqslant l$. Set $\omega'_v=\omega'_y=\texttt{0f}$, $\omega'_w=\omega'_w=\texttt{f0}$, and declare (vy) and (wx) to be matched edges (under both $\underline{m}^1$ and $\underline{m}^2$). Repeat this for each of the $n\delta'$ chosen 4-tuples (x,y,v,w), then set $\omega'_u=\omega_u$ for all $u\in V$ not appearing in one of the chosen 4-tuples. This defines a map ς satisfying

$$|\varsigma(G,\underline{\omega},\underline{m})|=[(n\delta')!]^{-1}\exp\left(O\left(\frac{nd^5(\delta')^2}{\delta}\right)\right)((d-1)_k(d-1)_l|V_{\texttt{ff}}|\,|F_{\texttt{00}}(k,l)|)^{n\delta'},$$
$$|\varsigma^{-1}(G,\underline{\omega},\underline{m})|\leqslant[(n\delta')!]^{-1}\exp\left(O\left(\frac{nd^5(\delta')^2}{\delta}\right)\right)((d-1)_k(d-1)_l|V_{\texttt{0f}}(k)|\,|V_{\texttt{f0}}(l)|)^{n\delta'},$$

where $V_{\texttt{0f}}(k)$ is the subset of variables in $V_{\texttt{0f}}$ with exactly k neighbors in $V_{\texttt{1Z}}$, and $V_{\texttt{f0}}(l)$ is symmetrically defined. Applying (79) gives

$$|V_{\texttt{ff}}|\,|F_{\texttt{00}}(k,l)|\lesssim|V_{\texttt{0f}}(k)|\,|V_{\texttt{f0}}(l)|.$$

Summing over k and l, and rearranging, gives $\pi_{\texttt{ff}}\lesssim\pi_{\texttt{0f}}\pi_{\texttt{f0}}/d$.

(3) *Figure 3(c): estimate for* $\texttt{f!f}$*-variables.*

Write $X_{\texttt{00}}$ for the subset of variables in $V_{\texttt{00}}$ that have no neighbors in $V_{\texttt{11}}$, and note that $|X_{\texttt{00}}|\geqslant\frac{1}{2}n$. Let $X_{\texttt{00}}(l)$ denote the subset of such variables that have exactly l neighbors in $V_{\texttt{1Z}}$. Then we have $|X_{\texttt{00}}(l)|\geqslant n\delta$ for all $2\leqslant l\leqslant d-2$.

Now fix some $2\leqslant l\leqslant d-2$. Choose a subset $S_{\texttt{f!f}}\subset V_{\texttt{f!f}}$ of size $n\delta'$. For each $v\in S_{\texttt{f!f}}$, choose a corresponding $x\in X_{\texttt{00}}(l)$, and let $S_{\texttt{00}}(l)$ denote the set of chosen x. We require that for all $u\neq u'$ in $S_{\texttt{f!f}}\cup S_{\texttt{00}}(l)$ the graph distance between u and u' is at least five. Let v_1 be the matched partner of v under $\underline{m}^1$. Let $x_1,...,x_l$ be the l neighbors of x in $V_{\texttt{1Z}}$. Choose $l-1$ neighbors $v_2,...,v_l$ of v (none equal to v_1). Cut the edges (vv_j) and (xx_j), and form the new edges (vx_j) and (xv_j). Then set $\omega'_v=\texttt{0f}$, $\omega'_x=\texttt{f0}$, and declare (xv_1) to be matched under $\underline{m}^1$. Repeat this for each of the $n\delta'$ chosen pairs (v,x), then set $\omega'_u=\omega_u$ for all u not appearing among any of the chosen pairs (v,x). This defines a map ς satisfying

$$|\varsigma(G,\underline{\omega},\underline{m})|=[(n\delta')!]^{-1}\exp\left(O\left(\frac{nd^5(\delta')^2}{\delta}\right)\right)((d-1)_{l-1}|V_{\texttt{f!f}}|\,|X_{\texttt{00}}(l)|)^{n\delta'},$$
$$|\varsigma^{-1}(G,\underline{\omega},\underline{m})|\leqslant[(n\delta')!]^{-1}\exp\left(O\left(\frac{nd^5(\delta')^2}{\delta}\right)\right)((d-1)_{l-1}|V_{\texttt{f0}}|\,|X_{\texttt{0f}}(l)|)^{n\delta'},$$

where $X_{\texttt{0f}}(l)$ denotes the subset of variables in $V_{\texttt{0f}}$ having exactly l neighbors in $V_{\texttt{1Z}}$. Applying (79) gives

$$|V_{\texttt{f!f}}|\,|X_{\texttt{00}}(l)| \lesssim |V_{\texttt{f0}}|\,|X_{\texttt{0f}}(l)|.$$

Summing over l and rearranging gives $\pi_{\texttt{f!f}} \lesssim \pi_{\texttt{f0}}\pi_{\texttt{0f}}$.

(4) *Figure* 3 (d): *estimate for* $\texttt{0f}$*-variables.*

Let $Y_{\texttt{0f}}$ be the subset of variables $u\in V_{\texttt{0f}}$ that have $d-1$ neighbors in $V_{\texttt{\_0}}$, and exactly one neighbor v in $V_{\texttt{0f}}$ (which is necessarily matched to u). By the preceding estimates, $V_{\texttt{0f}}$ accounts for most of $V_{\texttt{\_f}}$. The total size of $V_{\texttt{\_f}}$ is very small compared with $V_{\texttt{00}}$, so we have $|Y_{\texttt{0f}}|\geqslant\frac{1}{2}|V_{\texttt{0f}}|$. If we let $Y_{\texttt{01}}$ denote the number of pairs (x,y) such that $x\in V_{\texttt{01}}$ and $y\in V_{\texttt{Z}}\texttt{0}^{\texttt{r}}$, then we also have $|Y_{\texttt{01}}|\geqslant\frac{1}{2}d|V_{\texttt{01}}|$.

Choose a subset $S_{\texttt{0f}}$ of $n\delta'$ pairs (u,v) such that $u\in Y_{\texttt{0f}}$, and $v\in V_{\texttt{0f}}$ is the matched partner of u. Let $w\neq u$ be another neighbor of v. For each (u,v) choose a corresponding pair $(x,y)\in Y_{\texttt{01}}$, and let $S_{\texttt{00}}$ denote the subset of chosen pairs. We require that for any two chosen pairs $(u,u')\neq(u'',u''')$ in $S_{\texttt{0f}}\cup S_{\texttt{00}}$, the graph distance between u and u'' is at least five.

Cut the edges (vw) and (xy) and form the switched edges (vx) and (wy). Set $\omega'_u=\texttt{01}$, $\omega'_v=\texttt{00}$, $\omega'_x=\texttt{01}$, $\omega'_y=\texttt{00}$, and remove (uv) from the matching $\underline{m}^2$. Repeat this for each chosen 4-tuple (u,v,x,y), then set $\omega'_z=\omega_z$ for all z not appearing in any of the $n\delta'$ chosen 4-tuples. This defines a map ς satisfying

$$|\varsigma(G,\underline{\omega},\underline{m})| = [(n\delta')!]^{-1}\exp\biggl(O\biggl(\frac{nd^5(\delta')^2}{\delta}\biggr)\biggr)|Y_{\texttt{0f}}|\,|Y_{\texttt{01}}|(d-1),$$
$$|\varsigma^{-1}(G,\underline{\omega},\underline{m})| \leqslant [(n\delta')!]^{-1}\exp\biggl(O\biggl(\frac{nd^5(\delta')^2}{\delta}\biggr)\biggr)|V_{\texttt{00}^{\texttt{s}}}|\,|V_{\texttt{00}}|d.$$

Applying (79) and rearranging gives $\pi_{\texttt{0f}}\lesssim\pi_{\texttt{00}^{\texttt{s}}}/d\pi_{\texttt{01}}$. Applying the estimate (23) from Proposition 2.10, we conclude that $\pi_{\texttt{0f}}\leqslant(\log d)^{O(1)}d^{-y}$, as claimed.

This concludes the proof of the lemma. □

LEMMA 4.16. *Any maximizer of* $\mathbb{E}[\boldsymbol{Z}_{\alpha-\beta/2,\beta}]$ *over* $\beta\leqslant\beta_{\max}$ *satisfies*

$$\beta\leqslant\frac{(\log d)^{O(1)}}{d^y}.$$

Proof. This follows by essentially the same argument used to prove that

$$\pi_{\texttt{0f}}\leqslant\frac{(\log d)^{O(1)}}{d^y}$$

in Lemma 4.15. □

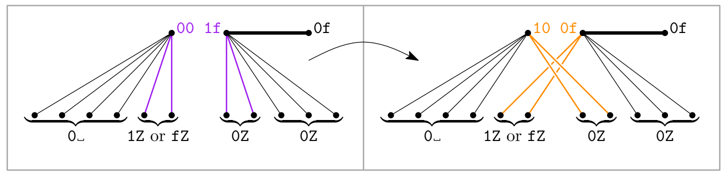

(a) Switching 1f, 00 for 1f, 0f.

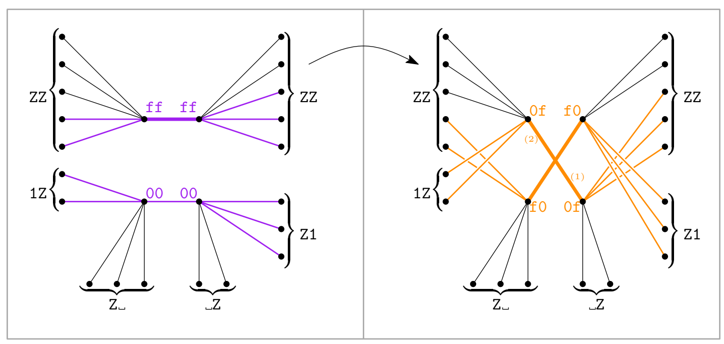

(b) Switching ff, 00 for 0f, f0.

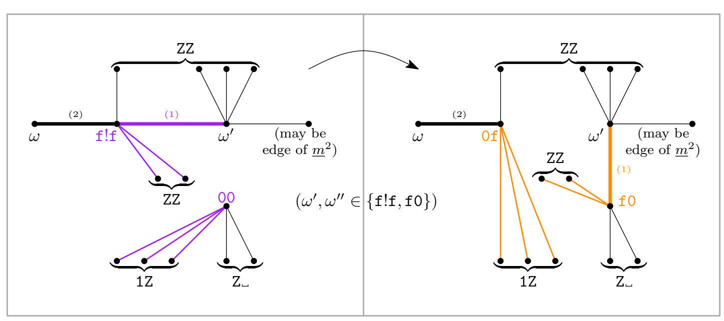

(c) Breaking an f!f chain.

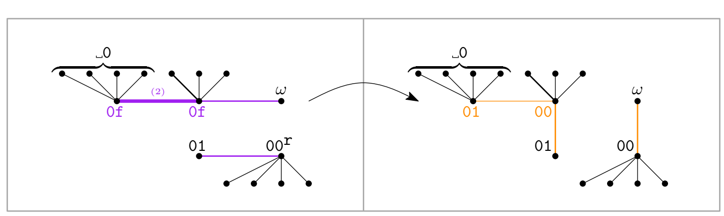

(d) Switching 0f, ηf for 01, η0.

Figure 3. Switching arguments (Lemma 4.15). In the original $(G, \underline{\omega}, \underline{m})$ we choose $n\delta'$ (mutually disjoint) copies of the left panel. The arrow indicates the operation that we apply in each of these copies to form the image $(G', \underline{\omega}', \underline{m}')$. Matching edges are indicated by thick lines; a thick line is marked (1) or (2) if the matching occurs under $\underline{m}^1$ or $\underline{m}^2$ only. Changes made by the switching operation are highlighted in color.

Corollary 4.17. *For $d \geqslant d_0$ and $\alpha_{\mathrm{lbd}} \leqslant \alpha \leqslant \alpha_{\mathrm{ubd}}$,*

(a) *any global maximizer of $\mathbf{\Phi}$ on $\mathbf{\Delta}[\alpha]$ lies in the interior $\mathbf{\Delta}^\circ[\alpha]$, and*

(b) *any global maximizer of $\mathbf{\Phi}_2$ on $\mathsf{IND}[\alpha]$ must be an interior stationary point.*

Proof. Proposition 4.13 and Lemma 4.16 combine to give (a), while (b) follows by combining Corollary 4.6, Proposition 4.12 and Lemma 4.15. □

Corollary 4.17 was required in the proof of Theorem 3.11; it also implies (with Lemma 3.9) that any maximizer $\boldsymbol{h}$ on $\mathbf{\Phi}_2$ on $\mathsf{IND}[\alpha]$ corresponds to a solution h of the pair Bethe recursions for some $\lambda \equiv (\lambda_1, \lambda_2)$ ((40) and (39) 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 ${}^\otimes\boldsymbol{h}[\alpha]$. For $\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d} \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}$ let us abbreviate

$$\hat{h}\binom{\mathtt{ab}}{\mathtt{cd}} \equiv \hat{h}(\boldsymbol{i} = \boldsymbol{i}^1\boldsymbol{i}^2 = \mathtt{ac}, \boldsymbol{o} = \boldsymbol{o}^1\boldsymbol{o}^2 = \mathtt{bd}) = \hat{h}(\sigma^1 = \boldsymbol{i}^1\boldsymbol{o}^1 = \mathtt{ab}, \sigma^2 = \boldsymbol{i}^2\boldsymbol{o}^2 = \mathtt{cd}).$$

Lemma 4.18. *In the setting of Proposition 4.14,*

$$\mathrm{B}_{\mathtt{cd}} \equiv \frac{\hat{h}\binom{\mathtt{0f}}{\mathtt{cd}}}{\hat{h}\binom{\mathtt{00}}{\mathtt{cd}}} \ll_d d^{-1} \quad \textit{for all } \mathtt{c}, \mathtt{d} \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}.$$

Proof. By the trivial symmetries (82) we have that $\mathrm{B}_{\mathtt{0}\eta} = \mathrm{B}_{\mathtt{f}\eta}$ and $\mathrm{B}_{\mathtt{00}} = \mathrm{B}_{\mathtt{10}} = \mathrm{B}_{\mathtt{f0}}$, so it suffices to show $\mathrm{B}_\sigma \ll_d d^{-1}$ for all $\sigma \in \{\mathtt{00}, \mathtt{01}, \mathtt{0f}, \mathtt{11}, \mathtt{1f}\}$. For any $\mathtt{a}, \mathtt{b} \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}$ we have

$$\mathrm{F}_{\mathtt{ab}} \equiv \frac{\hat{h}\binom{\mathtt{0a}}{\mathtt{0b}}}{\hat{h}\binom{\mathtt{00}}{\mathtt{00}}} = \frac{\hat{h}\binom{\mathtt{0a}}{\mathtt{0b}}\hat{h}\binom{\mathtt{a0}}{\mathtt{b0}}}{\hat{h}\binom{\mathtt{00}}{\mathtt{00}}\hat{h}\binom{\mathtt{00}}{\mathtt{00}}} = \frac{\hat{\boldsymbol{h}}\left(\binom{\mathtt{a0}}{\mathtt{b0}}, \binom{\mathtt{0a}}{\mathtt{0b}}\right)}{\hat{\boldsymbol{h}}\left(\binom{\mathtt{00}}{\mathtt{00}}, \binom{\mathtt{00}}{\mathtt{00}}\right)},$$

where the intermediate step is by (82) and the last step is by (39) and (40). Thus

$$\mathrm{F}_{\mathtt{10}} = \frac{\hat{h}\binom{\mathtt{01}}{\mathtt{00}}}{\hat{h}\binom{\mathtt{00}}{\mathtt{00}}} = \frac{\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}^{\mathtt{r}}\mathtt{0}})}{2\mathtt{E}(V_{\mathtt{00}}, V_{\mathtt{00}})} = \frac{nd\pi_{\mathtt{10}} - \mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{01}} \cup V_{\mathtt{0f}} \cup V_{\mathtt{0}^{\mathtt{s}}\mathtt{0}})}{nd[1 - O((\log d)/d)]} = \alpha + O(d^{-1}), \tag{80}$$

where the last step uses the estimate $K \leqslant K_{\mathrm{ubd}}$ from Lemma 4.5, together with the crude bounds $\mathtt{E}(V, V_{\mathtt{0f}}) \leqslant nd\beta_{\max}$ and $\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{0}}) \leqslant 2|V_{\mathtt{0}^{\mathtt{s}}\mathtt{0}}| \leqslant 2n$. We also have

$$\mathrm{F}_{\mathtt{11}} \lesssim \pi_{\mathtt{11}}, \quad \mathrm{F}_{\mathtt{ff}} \lesssim \pi_{\mathtt{ff}} + \pi_{\mathtt{f!f}}, \quad \mathrm{F}_{\mathtt{0f}} \asymp \pi_{\mathtt{0f}}, \quad \mathrm{F}_{\mathtt{f0}} \asymp \pi_{\mathtt{f0}}, \tag{81}$$

and so we deduce that

$$\mathrm{B}_{\mathtt{0f}} = \mathrm{F}_{\mathtt{ff}}/\mathrm{F}_{\mathtt{0f}} \ll_d d^{-1}, \quad \mathrm{B}_{\mathtt{00}} = \mathrm{F}_{\mathtt{f0}} \ll_d d^{-1}, \quad \mathrm{B}_{\mathtt{01}} = \mathrm{F}_{\mathtt{f1}}/\mathrm{F}_{\mathtt{01}} \lesssim \pi_{\mathtt{f1}}/\alpha \ll_d d^{-1}.$$

By a similar calculation as for $\mathrm{F}_{\mathtt{ab}}$, we have

$$\begin{aligned}
\mathrm{B}_{\mathtt{11}} &= \frac{\hat{h}\binom{\mathtt{0f}}{\mathtt{11}}}{\hat{h}\binom{\mathtt{00}}{\mathtt{11}}} = \frac{\hat{\boldsymbol{h}}\left(\binom{\mathtt{f0}}{\mathtt{11}}, \binom{\mathtt{0f}}{\mathtt{11}}\right)}{\hat{\boldsymbol{h}}\left(\binom{\mathtt{00}}{\mathtt{11}}, \binom{\mathtt{00}}{\mathtt{11}}\right)} \leqslant \frac{\pi_{\mathtt{f!f}}}{\pi_{\mathtt{0f}} - \pi_{\mathtt{1f}} - \pi_{\mathtt{f!f}}} \ll_d d^{-1}, \\
\mathrm{B}_{\mathtt{1f}} &= \frac{\hat{h}\binom{\mathtt{0f}}{\mathtt{1f}}}{\hat{h}\binom{\mathtt{00}}{\mathtt{1f}}} = \frac{\mathtt{E}(V_{\mathtt{f0}^{\mathtt{s}}}, V_{\mathtt{01}})}{\mathtt{E}(V_{\mathtt{00}^{\mathtt{s}}}, V_{\mathtt{01}})} \lesssim \frac{\mathtt{E}(V_{\mathtt{f0}^{\mathtt{s}}}, V_{\mathtt{Z1}})}{\mathtt{E}(V_{\mathtt{00}^{\mathtt{s}}}, V_{\mathtt{Z1}})} = \frac{|V_{\mathtt{f0}^{\mathtt{s}}}|}{|V_{\mathtt{00}^{\mathtt{s}}}|} \ll_d d^{-1},
\end{aligned}$$

where the last step uses the estimate (23) from Proposition 2.10 together with the estimates $K \leqslant K_{\mathrm{ubd}}$ and $F \leqslant F_{\mathrm{ubd}}$ from Lemma 4.5. □

LEMMA 4.19. *In the setting of Proposition* 4.14,

$$\mathrm{A}_{\mathtt{cd}} \equiv \frac{\hat{h}(\begin{smallmatrix}\mathtt{1f}\\ \mathtt{cd}\end{smallmatrix})}{\hat{h}(\begin{smallmatrix}\mathtt{10}\\ \mathtt{cd}\end{smallmatrix})} \ll_d d^{-1} \quad \text{for all } \mathtt{cd} \in \{\mathtt{0},\mathtt{1},\mathtt{f}\}.$$

Proof. By the trivial symmetries (82) we have that $\mathrm{A}_{\mathtt{0}\eta} = \mathrm{A}_{\mathtt{f}\eta}$ and $\mathrm{A}_{\mathtt{00}} = \mathrm{A}_{\mathtt{10}} = \mathrm{A}_{\mathtt{f0}}$, so it suffices to show $\mathrm{A}_\sigma \ll_d d^{-1}$ for all $\sigma \in \{\mathtt{00},\mathtt{01},\mathtt{0f},\mathtt{11},\mathtt{1f}\}$. For $\sigma = \mathtt{00}$ we have

$$\mathrm{A}_{\mathtt{00}} = \frac{\hat{h}\binom{\mathtt{1f}}{\mathtt{00}}}{\hat{h}\binom{\mathtt{10}}{\mathtt{00}}} = \frac{\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{0}})}{\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}^{\mathtt{r}}\mathtt{0}})} = \frac{\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{0}})}{\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}\_}) - \mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{01}} \cup V_{\mathtt{0f}}) - \mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{0}})}.$$

By its definition, $\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}\_}) = nd\alpha$. Next note that $\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{01}} \cup V_{\mathtt{0f}}) \leqslant K + F_1$, which by Lemma 4.5 is $\lesssim nd\alpha^2$. Lastly $\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{0}}) \leqslant 2|V_{\mathtt{0}^{\mathtt{s}}\mathtt{0}}|$, and it follows by the estimate (23) from Proposition 2.10 that at the maximizer we must have $|V_{\mathtt{0}^{\mathtt{s}}\mathtt{0}}|/n \ll_d 1/d$. Altogether this proves $\mathrm{A}_{\mathtt{00}} \ll_d 1/d$. Next,

$$\begin{aligned}\mathrm{A}_{\mathtt{01}} = \frac{\hat{h}\binom{\mathtt{1f}}{\mathtt{01}}}{\hat{h}\binom{\mathtt{10}}{\mathtt{01}}} &= \frac{\mathtt{E}(V_{\mathtt{10}^{\mathtt{r}}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{1}})}{\mathtt{E}(V_{\mathtt{10}^{\mathtt{r}}}, V_{\mathtt{0}^{\mathtt{r}}\mathtt{1}})} \leqslant \frac{\mathtt{E}(V_{\mathtt{10}^{\mathtt{r}}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{1}})}{\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{01}}) - \mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{1}}) - \mathtt{E}(V_{\mathtt{10}^{\mathtt{s}}}, V_{\mathtt{01}})} \\ &\leqslant \frac{2|V_{\mathtt{0}^{\mathtt{s}}\mathtt{1}}|}{K - 2|V_{\mathtt{0}^{\mathtt{s}}\mathtt{1}} \cup V_{\mathtt{10}^{\mathtt{s}}}|} \ll_d \frac{1}{d},\end{aligned}$$

where the last estimate uses that $K \lesssim nd\alpha^2$ by Lemma 4.5 and $|V_{\mathtt{0}^{\mathtt{s}}\mathtt{1}}|/|V_{\mathtt{01}}| \ll_d 1/d$ by (23). By similar considerations (and recalling $\mathtt{Z} \equiv \{\mathtt{0},\mathtt{f}\}$) we have

$$\mathrm{A}_{\mathtt{0f}} = \frac{\hat{h}\binom{\mathtt{1f}}{\mathtt{0f}}}{\hat{h}\binom{\mathtt{10}}{\mathtt{0f}}} = \frac{\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{f}})}{\mathtt{E}(V_{\mathtt{10}}, V_{\mathtt{0}^{\mathtt{r}}\mathtt{f}})} \leqslant \frac{2|V_{\mathtt{0}^{\mathtt{s}}\mathtt{f}}|}{\mathtt{E}(V_{\mathtt{1Z}}, V_{\mathtt{0f}}) - \mathtt{E}(V_{\mathtt{1f}}, V_{\mathtt{0f}}) - 2|V_{\mathtt{0}^{\mathtt{s}}\mathtt{f}}|} \ll_d \frac{1}{d}.$$

For subsets $S, S' \subseteq V$ let us write $\mathtt{E}(S, S'; \underline{m}^i)$ for the number of edges between S and S' that participate in the matching $\underline{m}^i$. It follows from (23) that

$$\mathrm{A}_{\mathtt{11}} = \frac{\mathtt{E}(V_{\mathtt{1f}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{f}}; \underline{m}^2)}{\mathtt{E}(V_{\mathtt{1f}}, V_{\mathtt{0}^{\mathtt{r}}\mathtt{f}}; \underline{m}^2)} \ll_d \frac{1}{d} \quad \text{and} \quad \mathrm{A}_{\mathtt{1f}} = \frac{\mathtt{E}(V_{\mathtt{11}}, V_{\mathtt{0}^{\mathtt{s}}\mathtt{0}^{\mathtt{s}}})}{\mathtt{E}(V_{\mathtt{11}}, V_{\mathtt{0}^{\mathtt{r}}\mathtt{0}^{\mathtt{s}}})} \ll_d \frac{1}{d},$$

and this concludes the proof of the lemma. □

As already noted above, Corollary 4.17 (b) implies that any maximizer $\boldsymbol{h}$ of $\boldsymbol{\Phi}_2$ on $\mathsf{IND}[\alpha]$ corresponds, via (40), to a solution $h \equiv (\dot{h}, \hat{h})$ of the Bethe recursions (39) with respect to the factors $\varphi^{\lambda_1} \otimes \varphi^{\lambda_2}$, for some parameters (λ_1, λ_2). We now show that h satisfies the analogue of (50) for the pair model.

PROPOSITION 4.20. *Let $\boldsymbol{h}$ be a maximizer of of $\boldsymbol{\Phi}_2$ on $\mathsf{IND}[\alpha]$, and let $h \equiv (\dot{h}, \hat{h})$ be the corresponding solution of the Bethe recursions. Then $\hat{h}$ is invariant with respect to changes in the incoming variable-to-clause message*:

$$\hat{h}(\boldsymbol{i}\boldsymbol{o}) = \hat{h}(\boldsymbol{i}'\boldsymbol{o}),$$

where $\boldsymbol{i}, \boldsymbol{i}' \in \{\mathtt{0},\mathtt{1},\mathtt{f}\}^2$ are variable-to-clause messages, and $\boldsymbol{o} \in \{\mathtt{0},\mathtt{1},\mathtt{f}\}^2$ is the clause-to-variable message.

Proof. We apply the same argument as in the proof of Proposition 3.10, that is, we shall argue that the effect of changing the message $\boldsymbol{i}$ incoming to $\acute{T}$ does not percolate down the tree. As in the first moment we have some trivial symmetries:

$$\begin{aligned}&\text{(i) } \hat{h}(\boldsymbol{io}) \text{ is invariant under changing } \boldsymbol{i}^i \text{ between } \mathtt{0} \text{ and } \mathtt{f}\text{; and}\\&\text{(ii) if } \boldsymbol{o}^i=\mathtt{0} \text{ then } \hat{h}(\boldsymbol{io}) \text{ is invariant under any change in } \boldsymbol{i}^i.\end{aligned} \tag{82}$$

It remains now to show that $\hat{h}(\boldsymbol{io})$ is invariant under changing $\boldsymbol{i}^i$ between $\mathtt{0}$ and $\mathtt{1}$. As in the proof of Proposition 3.10, this follows by showing that a certain propagation is subcritical, for which it suffices to have

$$\mathrm{A}_\sigma \equiv \frac{\hat{h}\binom{\mathtt{1f}}{\sigma}}{\hat{h}\binom{\mathtt{10}}{\sigma}} \ll_d d^{-1} \quad\text{and}\quad \mathrm{B}_\sigma \equiv \frac{\hat{h}\binom{\mathtt{0f}}{\sigma}}{\hat{h}\binom{\mathtt{00}}{\sigma}} \ll_d d^{-1} \quad \text{for all } \sigma\in\mathscr{M}.$$

This was done in Lemmas 4.18 and 4.19, so the result follows. □

It follows from Proposition 4.20 that we have (cf. (50) and (51))

$$9\hat{h}_{\boldsymbol{io}} = \boldsymbol{q}_{\boldsymbol{o}} \quad \text{for a measure } \boldsymbol{q} \text{ on } \{\mathtt{0},\mathtt{1},\mathtt{f}\}^2, \tag{83}$$

where $\boldsymbol{q}$ solves the recursions

$$\begin{aligned}z\boldsymbol{q}_{\mathtt{11}}&=\lambda_1\lambda_2(\boldsymbol{q}_{\mathtt{ZZ}})^{d-1}, & z\boldsymbol{q}_{\mathtt{1Z}}&=\lambda_1[(\boldsymbol{q}_{\mathtt{Z\_}})^{d-1}-(\boldsymbol{q}_{\mathtt{ZZ}})^{d-1}],\\ z\boldsymbol{q}_{\mathtt{Z1}}&=\lambda_2[(\boldsymbol{q}_{\mathtt{\_Z}})^{d-1}-(\boldsymbol{q}_{\mathtt{ZZ}})^{d-1}], & z\boldsymbol{q}_{\mathtt{ZZ}}&=1-(\boldsymbol{q}_{\mathtt{Z\_}})^{d-1}-(\boldsymbol{q}_{\mathtt{\_Z}})^{d-1}+(\boldsymbol{q}_{\mathtt{ZZ}})^{d-1}.\end{aligned} \tag{84}$$

Proof of Proposition 4.14. Write $\alpha\equiv y(\log d)/d$. By Proposition 4.20,

$$\frac{\boldsymbol{q}_{\mathtt{Z1}}}{\boldsymbol{q}_{\mathtt{ZZ}}} = \frac{\hat{h}\binom{\mathtt{00}}{\mathtt{01}}+\hat{h}\binom{\mathtt{0f}}{\mathtt{01}}}{\hat{h}\binom{\mathtt{00}}{\mathtt{00}}+\hat{h}\binom{\mathtt{00}}{\mathtt{0f}}+\hat{h}\binom{\mathtt{0f}}{\mathtt{00}}+\hat{h}\binom{\mathtt{0f}}{\mathtt{0f}}}.$$

By Lemma 4.18 and (80), this equals $\alpha+O(d^{-1})$. The same argument applies for $\boldsymbol{q}_{\mathtt{1Z}}$, so

$$\frac{\boldsymbol{q}_{\mathtt{Z1}}}{\boldsymbol{q}_{\mathtt{ZZ}}} = \alpha+O(d^{-1}) = \frac{\boldsymbol{q}_{\mathtt{1Z}}}{\boldsymbol{q}_{\mathtt{ZZ}}}. \tag{85}$$

Substituting (85) into (40) gives

$$\frac{\varrho}{\alpha-\varrho} = \frac{|V_{\mathtt{11}}|}{|V_{\mathtt{10}}|} = \frac{\lambda_2(\boldsymbol{q}_{\mathtt{ZZ}})^d}{(\boldsymbol{q}_{\mathtt{Z\_}})^d[1+O(d^{-y})]}.$$

On the other hand, substituting (85) into (84) gives

$$\frac{\boldsymbol{q}_{\mathtt{11}}}{\boldsymbol{q}_{\mathtt{1Z}}} = \frac{\lambda_2(\boldsymbol{q}_{\mathtt{ZZ}})^{d-1}}{(\boldsymbol{q}_{\mathtt{Z\_}})^{d-1}[1+O(d^{-y})]} = \frac{1+O(d^{-y})}{\boldsymbol{q}_{\mathtt{ZZ}}/\boldsymbol{q}_{\mathtt{Z\_}}}\frac{|V_{\mathtt{11}}|}{|V_{\mathtt{10}}|} \asymp \frac{\varrho}{\alpha},$$

and in combination with (85) this shows that the probability measure $\boldsymbol{q}$ is mostly supported on $\mathtt{ZZ}$, with $\boldsymbol{q}_{1\mathtt{Z}}=\alpha+O(d^{-1})=\boldsymbol{q}_{\mathtt{Z}1}$ and $\boldsymbol{q}_{11}\asymp\varrho$. The normalization z in (84) thus satisfies

$$z=1+z\boldsymbol{q}_{1\mathtt{Z}}+z\boldsymbol{q}_{\mathtt{Z}1}+O(d^{-y}+\varrho).$$

Consequently, if we let ε be defined by $\boldsymbol{q}_{\mathtt{ZZ}}=(1+\varepsilon)\boldsymbol{q}_{\mathtt{Z}\_}\boldsymbol{q}_{\_\mathtt{Z}}$, we have

$$|\varepsilon|=\left|\frac{(z\boldsymbol{q}_{\mathtt{ZZ}})z}{(z\boldsymbol{q}_{\mathtt{Z}\_})(z\boldsymbol{q}_{\_\mathtt{Z}})}-1\right|=\left|\frac{[1+O(d^{-y})][1+z\boldsymbol{q}_{1\mathtt{Z}}+z\boldsymbol{q}_{\mathtt{Z}1}+O(d^{-y}+\varrho)]}{(1+z\boldsymbol{q}_{\mathtt{Z}1})(1+z\boldsymbol{q}_{1\mathtt{Z}})}-1\right|\ll_d\frac{1}{d}.$$

Substituting into (84) gives, with the abbreviations $X_1\equiv(\boldsymbol{q}_{\mathtt{Z}\_})^{d-1}$ and $X_2\equiv(\boldsymbol{q}_{\_\mathtt{Z}})^{d-1}$,

$$\begin{aligned}
z\boldsymbol{q}_{11}&=(\lambda_1X_1)(\lambda_2X_2)+O(\varrho d\varepsilon), &\qquad z\boldsymbol{q}_{1\mathtt{Z}}&=\lambda_1X_1(1-X_2)+O\left(\frac{\alpha d\varepsilon}{d^y}\right),\\
z\boldsymbol{q}_{\mathtt{Z}1}&=(1-X_1)\lambda_2X_2+O\left(\frac{\alpha d\varepsilon}{d^y}\right), &\qquad z\boldsymbol{q}_{\mathtt{ZZ}}&=(1-X_1)(1-X_2)+O\left(\frac{d\varepsilon}{d^{2y}}\right).
\end{aligned}$$

All the additive errors above are $O(d^{-0.05}\varepsilon)$, so the recursion gives

$$\frac{\boldsymbol{q}_{\mathtt{ZZ}}}{\boldsymbol{q}_{\mathtt{Z}\_}}=\frac{[1+O(d^{-0.05}\varepsilon)](1-X_2)}{1-X_2+\lambda_2X_2}\quad\text{and}\quad\frac{1}{\boldsymbol{q}_{\_\mathtt{Z}}}=\frac{1-X_2+\lambda_2X_2}{[1+O(d^{-0.05}\varepsilon)](1-X_2)}.$$

Recalling the definition of ε this implies $\varepsilon=O(d^{-0.05}\varepsilon)$, so the only possibility is that $\varepsilon=0$, meaning $\boldsymbol{q}_{\mathtt{ZZ}}=\boldsymbol{q}_{\mathtt{Z}\_}\boldsymbol{q}_{\_\mathtt{Z}}$. Substituting into the recursion (84) again proves that $\boldsymbol{q}$ must be a product measure over $\{\mathtt{0},\mathtt{1},\mathtt{f}\}^2$, and its marginal on the ith coordinate (for $i=1,2$) must satisfy the recursion (48) with respect to λ_i. By the one-to-one correspondence between α and q in the regime considered (see Theorem 3.11), we conclude $\boldsymbol{h}={}^{\otimes}\boldsymbol{h}[\alpha]$ as claimed. □

4.4. A-priori rigidity estimate

In this section we analyze near-identical frozen model configurations to prove the following result.

PROPOSITION 4.21. *The contribution to $\mathbb{E}[(\boldsymbol{Z}_{n\alpha})^2]$ from $\mathsf{EQ}[\alpha]$ is $\asymp_d\mathbb{E}\boldsymbol{Z}_{n\alpha}$ as $n\to\infty$.*

The proof of Proposition 4.21 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 MIS_n (Remark 1.1), we shall prove this estimate for graphs drawn from the following slight generalization of the configuration model which allows for some "dangling" half-edges (see Figure 4).

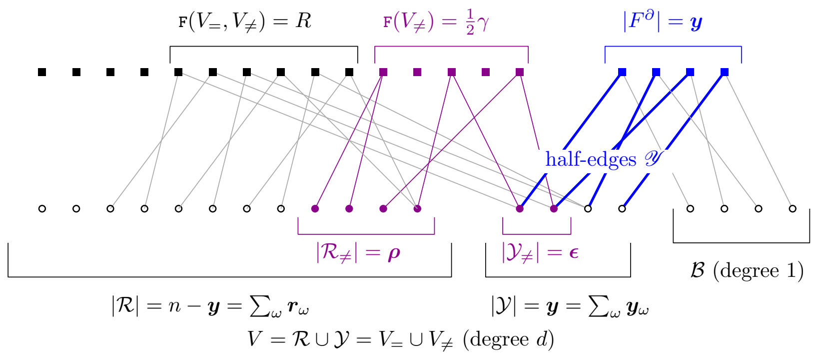

Figure 4. Graph G^∂ with dangling half-edges.

Let V be a set of n vertices, each incident to d half-edges. Let $\mathcal{Y}^{\rm ext}$ be a disjoint set of $\boldsymbol{y} \lesssim_d \log n$ vertices, each incident to a single half-edge. Let $V^\partial \equiv V \cup \mathcal{Y}^{\rm ext}$, and let G^∂ be the graph formed on V^∂ by taking a random matching of the $nd+\boldsymbol{y}$ available half-edges. We will equivalently regard G^∂ as a bipartite graph: $G^\partial = (V^\partial, F^\partial, \boldsymbol{E}^\partial)$ where F^∂ is a set of $\frac{1}{2}(nd+\boldsymbol{y})$ clauses, and $\boldsymbol{E}^\partial$ denotes the variable-clause edges. We write $F^{\rm ext}$ for the clauses incident to $\mathcal{Y}^{\rm ext}$, $\mathcal{Y}^{\rm int}$ for the variables in V that are incident to $F^{\rm ext}$, and $\mathcal{R} \equiv V \setminus \mathcal{Y}^{\rm int}$.

Definition 4.22. An auxiliary configuration on G^∂ is a message configuration $\underline{\sigma} \in \mathscr{M}^{\boldsymbol{E}^\partial}$ such that the configuration of messages incident to any variable or clause in $G^\partial \setminus \mathcal{Y}^{\rm ext}$ is valid (appears in $\operatorname{supp} \varphi$), and the total density of $\mathtt{f}$-variables is $\leqslant \beta_{\max}$. A frozen model configuration $\underline{\bar{\boldsymbol{\eta}}}$ on G^∂ is a spin configuration $\underline{\eta} \in \{\mathtt{0}, \mathtt{1}, \mathtt{f}\}^{V^\partial}$ together with a subset $\underline{m} \subseteq F$, obtained by applying $\mathsf{proj}_{\rm v}$ (see Proposition 3.5) to an auxiliary configuration on G^∂. Thus there is a one-to-one correspondence $\underline{\sigma} \leftrightarrow \underline{\bar{\boldsymbol{\eta}}}$ between the auxiliary and frozen models on G^∂.

As in (57), we associate with a frozen configuration pair $\underline{\bar{\boldsymbol{\eta}}}^i \equiv (\underline{\eta}^i, \underline{m}^i)$ $(i=1,2)$ on G^∂ a spin configuration $\underline{\omega} \equiv \underline{\omega}(\underline{\bar{\boldsymbol{\eta}}}^1, \underline{\bar{\boldsymbol{\eta}}}^2) \in \mathscr{P}^{V^\partial}$. As in (58), for $\omega \in \mathscr{P}$ and $S \subseteq V^\partial$ we write S_ω for the subset of vertices in S with spin ω. We will also write $S_{\neq}$ for the subset of vertices in S with spin in $\mathscr{P}_{\neq} \equiv \mathscr{P} \setminus \{\mathtt{00}, \mathtt{11}, \mathtt{ff}\}$, and $S_= \equiv S \setminus S_{\neq}$. Let

$$\boldsymbol{\varrho} \equiv |\mathcal{R}_{\neq}|, \quad \boldsymbol{\varepsilon} \equiv |\mathcal{Y}^{\rm ext}_{\neq}|, \quad \Delta \equiv |V_{\neq}| = |\mathcal{R}_{\neq} \cup \mathcal{Y}^{\rm ext}_{\neq}| = \boldsymbol{\varrho} + \boldsymbol{\varepsilon}.$$

As before, for $S, S' \subseteq V^\partial$ write $\mathtt{E}(S, S')$ for the number of edges joining S to S'. We shall abbreviate $\mathtt{E}(S) \equiv \mathtt{E}(S, S)$.

Write $\mathscr{Y} \subseteq \boldsymbol{E}^\partial$ for the variable-clause edges joining $\mathcal{Y}^{\rm int}$ to $F^{\rm ext}$. In the following, we fix boundary conditions $\underline{\tau}_{\mathscr{Y}} \equiv (\underline{\sigma}^1_{\mathscr{Y}}, \underline{\sigma}^2_{\mathscr{Y}}) \in \mathscr{M}^{2\mathscr{Y}}$. Let $\boldsymbol{O}^{2,\partial}[\pi | \underline{\tau}_{\mathscr{Y}}]$ denote the set of pair frozen model configurations on G^∂ which are consistent with $\underline{\tau}_{\mathscr{Y}}$ and have empirical measure $\pi_\omega = |\mathcal{R}_\omega|/|\mathcal{R}|$ when restricted to $\mathcal{R}$. Let

$$\boldsymbol{Z}^{2,\partial}[\pi\,|\,\underline{\tau}_{\mathscr{Y}}] \equiv \big|\boldsymbol{O}^{2,\partial}[\pi\,|\,\underline{\tau}_{\mathscr{Y}}]\big|.$$

We decompose

$$\boldsymbol{Z}^{2,\partial}[\pi\,|\,\underline{\tau}_{\mathscr{Y}}] = \sum_{A\geqslant 0} \boldsymbol{Z}^{2,\partial}_A[\pi\,|\,\underline{\tau}_{\mathscr{Y}}] = \sum_{A\geqslant 0} \big|\boldsymbol{O}^{2,\partial}_A[\pi\,|\,\underline{\tau}_{\mathscr{Y}}]\big| \tag{86}$$

where $\boldsymbol{O}^{2,\partial}_A[\pi\,|\,\underline{\tau}_{\mathscr{Y}}] \subseteq \boldsymbol{O}^{2,\partial}[\pi|\underline{\tau}_{\mathscr{Y}}]$ is the subset of configurations having

$$\tfrac{1}{2}\gamma \equiv \mathtt{E}(V_{\neq}) = \boldsymbol{\varrho} + A$$

internal edges among the unequal spins.[11] We shall compare below the expectation of $\boldsymbol{Z}^{2,\partial}_A[\pi\,|\,\underline{\tau}_{\mathscr{Y}}]$ 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. Note that $\boldsymbol{Z}^{\partial}[\pi^1\,|\,\underline{\sigma}^1_{\mathscr{Y}}] = \big|\mathsf{proj}^1(\boldsymbol{O}^{2,\partial}[\pi\,|\,\underline{\tau}_{\mathscr{Y}}])\big|$, where proj^1 is the projection mapping $(\underline{\bar{\boldsymbol{\eta}}}^1, \underline{\bar{\boldsymbol{\eta}}}^2) \mapsto \underline{\bar{\boldsymbol{\eta}}}^1$.

PROPOSITION 4.23. *There exists a small absolute constant $\varepsilon > 0$ such that for $d \geqslant d_0$, and any empirical measure π on $\mathscr{P}$ whose projections π^1 and π^2 have normalized intensities* (*cf.* (34)) $\mathbf{i}(\pi^i) = \pi^i_1 + \frac{1}{2}\pi^i_{\mathtt{f}} \in [\alpha_{\mathrm{lbd}}, \alpha_{\mathrm{ubd}}]$ (*though not necessarily equal*),

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

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

Proof. We will compute the ratio

$$R^1_A[\pi\,|\,\underline{\tau}_{\mathscr{Y}}] \equiv \frac{\mathbb{E}[\boldsymbol{Z}^{2,\partial}_A[\pi\,|\,\underline{\tau}_{\mathscr{Y}}]]}{\mathbb{E}[\boldsymbol{Z}^{\partial}[\pi^1\,|\,\underline{\sigma}^1_{\mathscr{Y}}]]}.$$

In what follows we will refer to the original d-regular graph model, although we keep the (equivalent) bipartite version in mind. Assume that the edges between V and $\mathcal{Y}^{\mathrm{ext}}$ are assigned, and that we are given $\underline{\omega} \in \mathscr{P}^{V^\partial}$ which is consistent with $\underline{\tau}_{\mathscr{Y}}$, and whose restriction to $\mathcal{R}$ has empirical measure π. Write E_ω for the number of half-edges incident to V_ω that are *not* matched to a half-edge from $\mathcal{Y}^{\mathrm{ext}}$; the total number of such half-edges is $\mathsf{E} \equiv \sum_\omega \mathsf{E}_\omega = nd - \boldsymbol{y}$.

Half-edges from $V_{\neq}$. We first estimate the probability that $\mathtt{E}(V_{\neq}, V_{\mathtt{11}}) = 0$. To this end, define

$$\mathscr{D}_{A,B}(\underline{\omega}) \equiv \big\{\mathtt{E}(V_{\neq}) = \tfrac{1}{2}\gamma \equiv \boldsymbol{\varrho} + A,\ \mathtt{E}(V_{\neq}, V_{\mathtt{ff}}) = B,\ \text{and } \mathtt{E}(V_{\neq}, V_{\mathtt{11}}) = 0\big\}.$$

---

[11] It is straightforward to see that $A \geqslant 0$.

Let $R\equiv\mathtt{E}(V_=,V_{\neq})=\mathsf{E}_{\neq}-\gamma$. Then, recalling the double factorial notation,

$$\mathbb{P}[\mathscr{D}_{A,B}(\underline{\omega})]\leqslant \underbrace{\frac{(\mathsf{E}_{\neq})_\gamma\,(\gamma-1)!!}{\gamma!\,[\mathsf{E}]_{\gamma/2}}}_{\leqslant e^{O(\gamma)}(\Delta^2 d/n\gamma)^{\gamma/2}}\cdot\underbrace{\frac{(\mathsf{E}_{\mathtt{ff}}+\mathsf{E}_{\mathtt{00}})_R}{[\mathsf{E}-\gamma]_R}}_{\leqslant d^{O(\Delta\varepsilon)}e^{-\mathsf{E}_{\mathtt{11}}R/\mathsf{E}}}\cdot\underbrace{\frac{\binom{R}{B}(\mathsf{E}_{\mathtt{ff}})_B(\mathsf{E}_{\mathtt{00}})_{R-B}}{(\mathsf{E}_{\mathtt{ff}}+\mathsf{E}_{\mathtt{00}})_R}}_{\leqslant d^{O(\Delta\varepsilon)}\binom{R}{B}\beta^B},\tag{87}$$

where $\beta\equiv\pi_{\mathtt{f}\_}$. Altogether we find

$$\mathbb{P}[\mathscr{D}_{A,B}(\underline{\omega})]\leqslant\frac{e^{O(A)}d^{O(\Delta\varepsilon)}}{\exp(\mathsf{E}_{\mathtt{11}}R/nd)}\left(\frac{\Delta^2 d}{n\gamma}\right)^{\gamma/2}\binom{R}{B}\beta^B.$$

Matchings on $\mathtt{f}$*-vertices.* Let $R'\equiv\mathsf{E}_{\neq}+R=2\mathsf{E}_{\neq}-\gamma\leqslant 2\Delta d$ denote the number of half-edges which are either incident to $V_{\neq}$ or matched with a half-edge from $V_{\neq}$. Let $U_{\mathtt{ff}}$ denote the subset of vertices in $V_{\mathtt{ff}}$ not matched to $\mathcal{Y}^{\mathrm{ext}}$ under $\underline{\tau}_{\mathscr{Y}}$. Conditioned on $\mathscr{D}_{A,B}(\underline{\omega})$, the expected number of matchings on $U_{\mathtt{ff}}$ is

$$\left(\prod_{v\in U_{\mathtt{ff}}}|\partial v\setminus(\mathcal{Y}^{\mathrm{ext}}\cup V_{\neq})|\right)\frac{(|U_{\mathtt{ff}}|-1)!!}{[\mathsf{E}-R']_{|U_{\mathtt{ff}}|/2}}.$$

In contrast, suppose we consider the first coordinate only: let $U_{\mathtt{f}\_}$ be the subset of vertices in $V_{\mathtt{f}\_}$ not matched to $\mathcal{Y}^{\mathrm{ext}}$ under $\underline{\sigma}^1_{\mathscr{Y}}$. The expected number of matchings on $U_{\mathtt{f}\_}$ is

$$\left(\prod_{v\in U_{\mathtt{f}\_}}|(\partial v)\setminus\mathcal{Y}^{\mathrm{ext}}|\right)\frac{(|U_{\mathtt{f}\_}|-1)!!}{[\mathsf{E}]_{|U_{\mathtt{f}\_}|/2}}.$$

Let $\mathscr{I}_{\mathtt{f}}^B$ denote the ratio between these two. Then for all B we have

$$\mathscr{I}_{\mathtt{f}}^B\leqslant e^{O(\Delta)}\left(\frac{1}{d\beta}\right)^{\Xi_{\mathtt{f}}/2}\quad\text{where }\ \Xi_{\mathtt{f}}\equiv|U_{\mathtt{f}\_}|-|U_{\mathtt{ff}}|.\tag{88}$$

Edges from $\mathtt{1}$*-vertices.* Suppose we are given $\underline{\omega}$ and a valid matching on $U_{\mathtt{ff}}$. Conditioning on $\mathscr{D}_{A,B}(\underline{\omega})$, the probability to correctly assign all half-edges leaving $V_{\mathtt{11}}$ is

$$\prod_{i=0}^{\mathsf{E}_{\mathtt{11}}-1}\frac{\mathsf{E}_{\mathtt{00}}-(R-B)-i}{\mathsf{E}_{=}-R-|U_{\mathtt{ff}}|-1-2i}.$$

In contrast, if we consider the first coordinate only, the probability to correctly assign all half-edges leaving $V_{\mathtt{1}\_}$ is

$$\prod_{i=0}^{\mathsf{E}_{\mathtt{1}\_}-1}\frac{\mathsf{E}_{\mathtt{0}\_}-i}{\mathsf{E}-|U_{\mathtt{f}\_}|-1-2i}.$$

Let $\mathscr{I}_{\mathtt{1}}^B$ be the ratio between these two; then

$$\mathscr{I}_{\mathtt{1}}^B\leqslant d^{O(\Delta\varepsilon)}\frac{\exp(\mathsf{E}_{\mathtt{11}}(\mathsf{E}_{\neq}+B)/nd)}{\exp(\mathsf{E}_{\mathtt{11}}(\mathsf{E}_{\mathtt{0}\_}-\mathsf{E}_{\mathtt{00}})/nd)}.\tag{89}$$

Forcing of $\mathtt{0}$*-vertices.* It remains to address the constraint that each $\mathtt{0}$ is forced by two neighboring $\mathtt{1}$'s. Let $\mathscr{J}_{\mathtt{0}}$ be the ratio of the conditional probability that all vertices in $V_{\mathtt{00}}$ are forced, with the conditional probability that all vertices in $V_{\mathtt{0}\llcorner}$ are forced in the first coordinate. Fix $0<\vartheta<1$ and define independent random variables $X^i\sim\mathrm{Bin}(|\partial i\setminus\mathcal{Y}^{\mathrm{ext}}|,\vartheta)$ with joint law $\mathbb{P}_\vartheta$. Then

$$\mathscr{J}_{\mathtt{0}}^B\leqslant\frac{\mathbb{P}_\vartheta(X^i+|\partial i\cap\mathcal{Y}^{\mathrm{ext}}_{\mathtt{1}\llcorner}|\geqslant 2 \text{ for all } i\in V_{\mathtt{00}}\setminus\partial V_{\neq}\,|\,\sum_{i\in V_{\mathtt{00}}\setminus\partial V_{\neq}}X^i=\mathsf{E}_{\mathtt{11}})}{\mathbb{P}_\vartheta(X^i+|\partial i\cap\mathcal{Y}^{\mathrm{ext}}_{\mathtt{1}\llcorner}|\geqslant 2 \text{ for all } i\in V_{\mathtt{0}\llcorner}\,|\,\sum_{i\in V_{\mathtt{0}\llcorner}}X^i=\mathsf{E}_{\mathtt{1}\llcorner})},$$

where we have ignored the forcing constraint on the vertices in $V_{\mathtt{00}}$ neighboring $V_{\neq}$. A small variation on the proof of Proposition 2.10 gives $\mathscr{J}_{\mathtt{0}}=e^{O(\Delta)}$. Combining with (87), (88) and (89), and recalling that $\mathsf{E}_{\neq}-R=\gamma=2(\boldsymbol{\varrho}+A)$, we see that

$$\sum_B\mathbb{P}(\mathscr{D}_{A,B}(\underline{\omega}))\mathscr{J}_{\mathtt{f}}^B\mathscr{J}_{\mathtt{1}}^B\mathscr{J}_{\mathtt{0}}^B\leqslant\frac{e^{O(A)}d^{O(\Delta\varepsilon)}}{\exp(\Delta_{\mathtt{0}}(d\alpha))}\left(\frac{1}{d\beta}\right)^{\Xi_{\mathtt{f}}/2}\left(\frac{\Delta^2 d}{n\gamma}\right)^{\gamma/2},\tag{90}$$

where $\alpha\equiv\pi_{\mathtt{1}\llcorner}$ and $\Delta_{\mathtt{0}}\equiv|V_{\mathtt{0}\llcorner}|-|V_{\mathtt{00}}|$.

Ratio of combinatorial factors. Finally, accounting for the permutations of $\underline{\omega}$, versus the permutations in the first-coordinate projection, we have

$$\mathscr{C}\equiv\binom{|\mathcal{R}|}{\boldsymbol{\varrho},|\mathcal{R}_{\mathtt{00}}|,|\mathcal{R}_{\mathtt{11}}|,|\mathcal{R}_{\mathtt{ff}}|}\Big/\binom{|\mathcal{R}|}{|\mathcal{R}_{\mathtt{0}\llcorner}|,|\mathcal{R}_{\mathtt{1}\llcorner}|,|\mathcal{R}_{\mathtt{f}\llcorner}|}\leqslant e^{O(\Delta)}\left(\frac{n}{\boldsymbol{\varrho}}\right)^{\boldsymbol{\varrho}}\alpha^{\boldsymbol{\varrho}_{\mathtt{1}}}\beta^{\boldsymbol{\varrho}_{\mathtt{f}}},$$

where $\boldsymbol{\varrho}_\eta\equiv|\mathcal{R}_{\eta\llcorner}|-|\mathcal{R}_{\eta\eta}|$. Combining with (90) (and recalling $\frac{1}{2}\gamma=\boldsymbol{\varrho}+A$) gives

$$\begin{aligned}R_A^1[\pi\,|\,\underline{\tau}_{\mathscr{Y}}]&\leqslant\mathscr{C}\sum_B\mathbb{P}(\mathscr{D}_{A,B}(\underline{\omega}))\mathscr{J}_{\mathtt{f}}^B\mathscr{J}_{\mathtt{1}}^B\mathscr{J}_{\mathtt{0}}^B\\&\leqslant e^{O(A)}d^{O(\Delta\varepsilon)}\left(\frac{\Delta^2}{\boldsymbol{\varrho}\gamma}\right)^{\boldsymbol{\varrho}}\left(\frac{\Delta^2 d}{n\gamma}\right)^{A}\frac{d^{\boldsymbol{\varrho}}\alpha^{\boldsymbol{\varrho}_{\mathtt{1}}}\beta^{\boldsymbol{\varrho}_{\mathtt{f}}}}{\exp(\Delta_{\mathtt{0}}(d\alpha))(d\beta)^{\Xi_{\mathtt{f}}/2}}.\end{aligned}\tag{91}$$

Recalling that $\Delta=\boldsymbol{\varrho}+\boldsymbol{\varepsilon}$, we can bound

$$\left(\frac{\Delta^2}{\boldsymbol{\varrho}\gamma}\right)^{\boldsymbol{\varrho}}=\left(\frac{(\boldsymbol{\varrho}+\boldsymbol{\varepsilon})^2}{\boldsymbol{\varrho}\gamma}\right)^{\boldsymbol{\varrho}}\leqslant e^{O(\Delta)}\left(\frac{\boldsymbol{\varrho}}{\gamma}\right)^{\boldsymbol{\varrho}}.$$

Next, observe that $\Xi_{\mathtt{f}}\leqslant\boldsymbol{\varrho}_{\mathtt{f}}+\boldsymbol{\delta}_{\mathtt{f}}$ for $\boldsymbol{\delta}_\eta\equiv|\mathcal{Y}^{\mathrm{int}}_{\eta\llcorner}|-|\mathcal{Y}^{\mathrm{int}}_{\eta\eta}|$. Therefore, using $\alpha\geqslant\frac{3}{2}(\log d)/d$ and $\beta\leqslant\beta_{\max}$, we have

$$\frac{d^{\boldsymbol{\varrho}}\alpha^{\boldsymbol{\varrho}_{\mathtt{1}}}\beta^{\boldsymbol{\varrho}_{\mathtt{f}}}}{\exp(\Delta_{\mathtt{0}}(d\alpha))(d\beta)^{\Xi_{\mathtt{f}}/2}}\leqslant\frac{(\log d)^{O(\Delta)}}{\exp\big(\big[\big(\frac{3}{2}\Delta_{\mathtt{0}}-\boldsymbol{\varrho}_{\mathtt{0}}\big)+\frac{1}{4}(\boldsymbol{\varrho}_{\mathtt{f}}-\boldsymbol{\delta}_{\mathtt{f}})\big]\log d\big)}.$$

Recall that $|\mathcal{R}_{\mathtt{1f}}|\leqslant|V_{\mathtt{0f}}|$ and $|\mathcal{R}_{\mathtt{f1}}|\leqslant|V_{\mathtt{f0}}|$. By symmetry, assume that $|V_{\mathtt{01}}\cup V_{\mathtt{0f}}|\geqslant|V_{\mathtt{10}}\cup V_{\mathtt{f0}}|$. Then

$$\begin{aligned}\tfrac{3}{2}\Delta_{\mathtt{0}}-\boldsymbol{\varrho}_{\mathtt{0}}&=\tfrac{1}{2}\Delta_{\mathtt{0}}+(\Delta_{\mathtt{0}}-\boldsymbol{\varrho}_{\mathtt{0}})\geqslant\tfrac{1}{2}\Delta_{\mathtt{0}}=\tfrac{1}{2}|V_{\mathtt{01}}\cup V_{\mathtt{0f}}|\geqslant\tfrac{1}{4}|V_{\mathtt{01}}\cup V_{\mathtt{0f}}\cup V_{\mathtt{10}}\cup V_{\mathtt{f0}}|\\&\geqslant\tfrac{1}{8}|V_{\mathtt{01}}\cup V_{\mathtt{0f}}\cup V_{\mathtt{10}}\cup V_{\mathtt{f0}}|+\tfrac{1}{8}(\boldsymbol{\varrho}-|\mathcal{R}_{\mathtt{1f}}|-|\mathcal{R}_{\mathtt{f1}}|-|\mathcal{R}_{\mathtt{f!f}}|)\\&\geqslant\tfrac{1}{8}(\boldsymbol{\varrho}+|V_{\mathtt{0f}}|-|\mathcal{R}_{\mathtt{1f}}|+|V_{\mathtt{f0}}|-|\mathcal{R}_{\mathtt{f1}}|-|\mathcal{R}_{\mathtt{f!f}}|)\geqslant\tfrac{1}{8}(\boldsymbol{\varrho}-|\mathcal{R}_{\mathtt{f!f}}|),\end{aligned}$$

so combining with $\boldsymbol{\varrho}_{\mathtt{f}}-\boldsymbol{\delta}_{\mathtt{f}}\geqslant\boldsymbol{\varrho}_{\mathtt{f}}-\boldsymbol{\varepsilon}$ gives, for some absolute constant C,

$$\sum_{A\geqslant 0} R^1_A[\pi\,|\,\underline{T}_{\mathscr{Y}}]\leqslant\frac{d^{O(\Delta\varepsilon)}(\log d)^{O(\Delta)}}{d^{(\boldsymbol{\varrho}-\boldsymbol{\varepsilon})/8}}\sum_{A\geqslant 0}\left(\frac{\boldsymbol{\varrho}}{\gamma}\right)^{\boldsymbol{\varrho}}\left(\frac{C\Delta^2 d}{n\gamma}\right)^{A}\leqslant\frac{d^{O(\Delta\varepsilon)}(\log d)^{O(\Delta)}}{d^{(\boldsymbol{\varrho}-\boldsymbol{\varepsilon})/8}}.$$

To see the last inequality, note that the sum over $A\geqslant 0$ is clearly $\lesssim 1$ if $C\Delta^2 d\leqslant 2n\gamma$. For $C\Delta^2 d>2n\gamma$, recalling that $\boldsymbol{\varrho}\leqslant\frac{1}{2}\gamma$ and optimizing over γ gives

$$\left(\frac{\boldsymbol{\varrho}}{\gamma}\right)^{\boldsymbol{\varrho}}\left(\frac{C\Delta^2 d}{n\gamma}\right)^{A}\leqslant d^{O(\Delta\varepsilon)}.$$

It follows that there exists a small absolute constant ε such that

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

implying the result. □

Proposition 4.21 now follows from Proposition 4.23 applied to our original random graph $\mathscr{G}_{d,n}$ with no dangling half-edges. Theorem 4.1 then follows by combining Propositions 4.14 and 4.21.

5. Negative-definiteness of free energy Hessians

THEOREM 5.1. *For $d\geqslant d_0$, the following hold uniformly over $\alpha_{\mathrm{lbd}}\leqslant\alpha\leqslant\alpha_{\mathrm{ubd}}$:*
(a) $\mathbb{E}\boldsymbol{Z}_{n\alpha}=[1+o_n(1)]\,\mathscr{C}(\alpha)\,n^{-1/2}\exp(n\boldsymbol{\phi}(\alpha))$, *for a smooth function $\mathscr{C}(\alpha)$, and*
(b) $\mathbb{E}[(\boldsymbol{Z}_{n\alpha})^2]\asymp_d(\mathbb{E}\boldsymbol{Z}_{n\alpha})^2+\mathbb{E}\boldsymbol{Z}_{n\alpha}$.

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

The calculation of this section is similar to that of [24, §7]. Let $\boldsymbol{h}\in\boldsymbol{\Delta}^{\circ}[\alpha]$ 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 $\boldsymbol{h}+s\boldsymbol{\delta}\in\boldsymbol{\Delta}^{\circ}[\alpha]$ for sufficiently small $|s|$. Then

$$(\partial_s)^2\boldsymbol{\Phi}(\boldsymbol{h}+s\boldsymbol{\delta})|_{s=0}=-\langle(\dot{\boldsymbol{\delta}}/\dot{\boldsymbol{h}})^2\rangle_{\dot{\boldsymbol{h}}}-\tfrac{1}{2}d\langle(\hat{\boldsymbol{\delta}}/\hat{\boldsymbol{h}})^2\rangle_{\hat{\boldsymbol{h}}}+d\langle(\bar{\delta}/\bar{h})^2\rangle_{\bar{h}},\tag{92}$$

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 (92) 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 the 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

$$\overline{H}^{-1}\bar\delta = d^{-1}\,[I+(d-1)\dot M]\,\dot\chi, \tag{93}$$

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

$$\dot M_{\sigma,\sigma'} \equiv \frac{1}{\bar h(\sigma)}\sum_{\underline{\dot\sigma}} \dot{\boldsymbol{h}}(\underline{\dot\sigma})\mathbf{1}\{(\sigma_1,\sigma_2)=(\sigma,\sigma')\}. \tag{94}$$

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 ${}^\otimes\boldsymbol{h}[\alpha]$ (second moment).

LEMMA 5.3. *For $d\geqslant d_0$, the eigenvalues of $\dot M$ counted with geometric multiplicity are*

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

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 the eigenvalue $-(d-1)^{-1}$ is given by $\bar x_\sigma=(d-1)\mathbf{1}\{\sigma=\texttt{11}\}-\mathbf{1}\{\sigma=\texttt{f0}$ or $\texttt{ff}\}$.

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

$$\dot M_{\texttt{1}} = \begin{array}{c c} & \begin{array}{cc} \texttt{10} & \texttt{1f} \end{array} \\ \begin{array}{c} \texttt{10} \\ \texttt{1f} \end{array} & \begin{pmatrix} \dfrac{q_{\texttt{0}}}{1-q_{\texttt{1}}} & \dfrac{q_{\texttt{f}}}{1-q_{\texttt{1}}} \\ \dfrac{q_{\texttt{0}}}{1-q_{\texttt{1}}} & \dfrac{q_{\texttt{f}}}{1-q_{\texttt{1}}} \end{pmatrix} \end{array} \quad\text{and}\quad \dot M_{\texttt{f}} = \begin{array}{c c} & \begin{array}{ccc} \texttt{11} & \texttt{f0} & \texttt{ff} \end{array} \\ \begin{array}{c} \texttt{11} \\ \texttt{f0} \\ \texttt{ff} \end{array} & \begin{pmatrix} 0 & \dfrac{q_{\texttt{0}}}{1-q_{\texttt{1}}} & \dfrac{q_{\texttt{f}}}{1-q_{\texttt{1}}} \\ \dfrac{1}{d-1} & \dfrac{d-2}{d-1}\dfrac{q_{\texttt{0}}}{1-q_{\texttt{1}}} & \dfrac{d-2}{d-1}\dfrac{q_{\texttt{f}}}{1-q_{\texttt{1}}} \\ \dfrac{1}{d-1} & \dfrac{d-2}{d-1}\dfrac{q_{\texttt{0}}}{1-q_{\texttt{1}}} & \dfrac{d-2}{d-1}\dfrac{q_{\texttt{f}}}{1-q_{\texttt{1}}} \end{pmatrix} \end{array},$$

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

$$\dot M_{\texttt{0}} = \begin{array}{c c} & \begin{array}{cccc} \texttt{f1} & \texttt{01} & \texttt{00} & \texttt{0f} \end{array} \\ \begin{array}{c} \texttt{f1} \\ \texttt{01} \\ \texttt{00} \\ \texttt{0f} \end{array} & \begin{pmatrix} \dfrac{1}{d-1} & 0 & \dfrac{d-2}{d-1}\dfrac{q_{\texttt{0}}}{1-q_{\texttt{1}}} & \dfrac{d-2}{d-1}\dfrac{q_{\texttt{f}}}{1-q_{\texttt{1}}} \\ 0 & \varepsilon+q_{\texttt{1}}(1-\varepsilon) & q_{\texttt{0}}(1-\varepsilon) & q_{\texttt{f}}(1-\varepsilon) \\ \varepsilon & q_{\texttt{1}}(1-\varepsilon) & q_{\texttt{0}}(1-\varepsilon) & q_{\texttt{f}}(1-\varepsilon) \\ \varepsilon & q_{\texttt{1}}(1-\varepsilon) & q_{\texttt{0}}(1-\varepsilon) & q_{\texttt{f}}(1-\varepsilon) \end{pmatrix} \end{array},$$

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

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

Write $\overline{H}_0$ for $\overline{H}$ restricted to $\{\mathtt{f1}, \mathtt{01}, \mathtt{00}, \mathtt{0f}\}$. Then the symmetric matrix $\overline{H}_0^{1/2}\dot{M}_0\overline{H}_0^{-1/2}$ 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

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

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 $\langle \overline{H}_0^{1/2} v_1, e_4\rangle=0$, so $\overline{H}_0^{1/2} v_1=\sum_i a_i e_i$ with $a_4=0$. We then calculate

$$\begin{aligned}\frac{\sum_i (\lambda_i a_i)^2}{\|a\|^2} &= \frac{\|\overline{H}_0^{1/2}\dot{M}_0 v_1\|^2}{\|\overline{H}_0^{1/2} v_1\|^2} = \frac{\varepsilon^2 q_0 q_1 (1-q_1)[q_1 + (1-q_1)/(d-2)^2]}{q_1[q_{\mathtt{f}}(q_1)^2 + q_0(1-q_1)/(d-2)^2]} \\ &\asymp q_1 (d\varepsilon)^2 \asymp q_1 (q_{\mathtt{f}} \log d)^2.\end{aligned}$$

From Theorem 3.1 we have $\lambda \geqslant d^{0.59}$, and substituting into (48) gives

$$q_{\mathtt{f}} \asymp \frac{d(q_1)^2}{\lambda} \leqslant d^{-1.58}.$$

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 $\overline{H}_0^{1/2} v_2=\sum_i b_i e_i$ gives

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

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

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

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

Proof of Proposition 5.2. Consider the first moment Hessian $H\mathbf{\Phi}({}^{\star}\boldsymbol{h}[\alpha])$ on the space $\mathbf{\Delta}[\alpha]$. For convenience, rewrite (93) as $\overline{H}^{-1/2}\bar{\delta}=\dot{S}\overline{H}^{1/2}\dot{\chi}$ with $\dot{S}$ being the symmetric matrix

$$\dot{S}\equiv d^{-1}\overline{H}^{1/2}[I+(d-1)\dot{M}]\overline{H}^{-1/2}.$$

By the requirement that both $\boldsymbol{h}$ and $\boldsymbol{h}+\eta\boldsymbol{\delta}$ lie in $\mathbf{\Delta}^{\circ}[\alpha]$ 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 Lemma 5 to span the kernel of $\dot{L}\equiv d^{-1}[I+(d-1)\dot{M}]$. For such $\bar{\delta}$, (93) is easily solved by considering the invariant action of $\dot{S}$ on the subspace $(\overline{H}^{1/2}\bar{x})^{\perp}$: 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 $\overline{H}^{1/2}\dot{\chi}=U(U^t\dot{S}U)^{-1}U^t\overline{H}^{-1/2}\bar{\delta}$, consequently the maximal value of $(\partial_\eta)^2\mathbf{\Phi}(\boldsymbol{h}+\eta\boldsymbol{\delta})|_{\eta=0}$ subject to marginals $\bar{\delta}$ is

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

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

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

so $\dot{Q}$ is non-singular since we saw in Lemma 5 that $(d-1)^{-1}\notin\mathsf{eigen}(\dot{M})$. Since we proved in Theorem 3.11 that ${}^{\star}\boldsymbol{h}[\alpha]\in\mathbf{\Delta}^{\circ}[\alpha]$ is the global maximizer of $\mathbf{\Phi}$ on $\mathbf{\Delta}[\alpha]$, the restriction of the above quadratic form to the space of permissible $\bar{\delta}$ (formally, to $U^t\overline{H}^{-1/2}\overline{W}$) 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 $H(\mathbf{\Phi}_2)({}^{\otimes}\boldsymbol{h}[\alpha])$ on $\mathbf{\Delta}_2[\alpha]$ is similar: first observe that ${}^{\otimes}\boldsymbol{h}[\alpha]={}^{\star}\boldsymbol{h}[\alpha]\otimes{}^{\star}\boldsymbol{h}[\alpha]$ 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 $\dot{M}_2$ 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 $(d-1)^{-1}\notin\mathsf{eigen}(\dot{M}_2)$. $\square$

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

$$\dot{H}\dot{\boldsymbol{\delta}}-\widehat{H}\hat{\boldsymbol{\delta}}=\mathbf{1}_\sigma-\mathbf{1}_{\sigma'}.$$

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

Proof. Suppose $\tau{=}\boldsymbol{io}$ and $\tau'{=}\boldsymbol{io}'$ where $\boldsymbol{i}$ is the variable-to-clause (pair) message and $\boldsymbol{o},\boldsymbol{o}'$ are the clause-to-variable (pair) messages. By the subcriticality established in the proof of Proposition 4.20, there exist configurations $\underline{\tau}$ and $\underline{\tau}'$ on $\acute{T}$, with spins τ and τ', respectively, on the root half-edge $\acute{e}$, such that the configurations $\underline{\tau}$ and $\underline{\tau}'$ concide below some depth l. Take $U{=}(V_U,F_U,E_U)$ to be the depth-l subtree of $\acute{T}$. Assume that the depth-l vertices of U are variables, and let δU denote the half-edges incident to these leaf variables. Use $\underline{\tau}$ to define a signed measure $\boldsymbol{a}{\equiv}(\dot{\boldsymbol{a}},\hat{\boldsymbol{a}})$:

$$\dot{\boldsymbol{a}}(\dot{\underline{\tau}})=\sum_{v\in V_U}\mathbf{1}\{\dot{\underline{\tau}}_v=\dot{\underline{\tau}}\}\quad\text{and}\quad\hat{\boldsymbol{a}}(\hat{\underline{\tau}})=\sum_{a\in F_U}\mathbf{1}\{\hat{\underline{\tau}}_a=\hat{\underline{\tau}}\}.$$

If we take the difference of $\dot{H}\dot{\boldsymbol{a}}$ and $\widehat{H}\hat{\boldsymbol{a}}$, we see that internal edges of U will cancel, so

$$\dot{H}\dot{\boldsymbol{a}}-\widehat{H}\hat{\boldsymbol{a}}=\mathbf{1}_\tau-\sum_{u\in\delta U}\mathbf{1}_{\tau_u}.$$

Similarly, use $\underline{\tau}'$ to define a signed measure $\boldsymbol{a}'{\equiv}(\dot{\boldsymbol{a}}',\hat{\boldsymbol{a}}')$. Since $\underline{\tau}$ and $\underline{\tau}'$ agree on δU, we have

$$\dot{H}\dot{\boldsymbol{a}}'-\widehat{H}\hat{\boldsymbol{a}}'=\mathbf{1}_{\tau'}-\sum_{u\in\delta U}\mathbf{1}_{\tau_u}.$$

It follows that $\boldsymbol{\delta}{\equiv}\boldsymbol{a}-\boldsymbol{a}'$ satisfies the requirements for $\boldsymbol{\delta}_{\tau-\tau'}$.

Now 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 required properties. The above proves that τ and τ' are connected if they differ only in the clause-to-variable message. By the same considerations of subcriticality, τ and τ' are also connected if they differ only in the variable-to-clause message. Thus $\tau{=}\boldsymbol{io}$ is connected to $\tau'{=}\boldsymbol{io}'$ which in turn is connected to $\tau''{=}\boldsymbol{i}'\boldsymbol{o}'$, so the entire graph is connected (hence complete), as required. □

Note that Lemma 5.4 implies that the matrix $H_{\boldsymbol{\Delta}}$ of Definition 3.7 is indeed surjective.

Proof of Theorem 5.1. Recalling Definition 3.7, $\mathbb{E}\boldsymbol{Z}_{n\alpha}$ 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 $\boldsymbol{g}{\equiv}(\dot{\boldsymbol{g}},\hat{\boldsymbol{g}}){\equiv}\big(n\dot{\boldsymbol{h}},\tfrac{1}{2}nd\hat{\boldsymbol{h}}\big)$ is integer-valued, and lies in the kernel of $H_{\boldsymbol{\Delta}}{\equiv}(\,\dot{H}\;-\widehat{H}\,)$ with (non-normalized) intensity $\mathbf{i}(\boldsymbol{g}){=}n\alpha$. Let $\mathbb{E}[\boldsymbol{Z}^\star_{n\alpha}]$ 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 Proposition 5.2 implies $\mathbb{E}\boldsymbol{Z}_{n\alpha}{=}[1{+}o_n(1)]\,\mathbb{E}[\boldsymbol{Z}^\star_{n\alpha}]$.

Lemma 5.4 shows that the integer matrix $H_{\boldsymbol{\Delta}}$ defines a surjection

$$H_{\boldsymbol{\Delta}}\colon L'\longrightarrow\{\bar{\delta}\in\mathbb{R}^{\mathscr{M}}:\langle\bar{\delta},1\rangle=0\},$$

where

$$L'\equiv\{\boldsymbol{\delta}\in\mathbb{R}^{\operatorname{supp}\psi}:\langle\dot{\boldsymbol{\delta}},1\rangle=\langle\hat{\boldsymbol{\delta}},1\rangle=\mathbf{i}(\boldsymbol{\delta})=0\}.$$

Thus we can conclude 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 $\mathbb{E}[\boldsymbol{Z}^{\star}_{n\alpha}]$ 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 (36) then shows that $n^{1/2}\,\mathbb{E}[\boldsymbol{Z}^{\star}_{n\alpha}]$ defines a convergent Riemann sum, implying the result for $\mathbb{E}\boldsymbol{Z}_{n\alpha}$. The same argument implies that the contribution to the second moment from measures in $\mathsf{IND}[\alpha]$ is $\asymp_d n^{-1}\exp(n\,[\boldsymbol{\Phi}_2({}^{\otimes}\boldsymbol{h}[\alpha])]) \asymp_d (\mathbb{E}\boldsymbol{Z}_{n\alpha})^2$, so combining with Theorem 4.1 and Proposition 4.21 gives the result for $\mathbb{E}[(\boldsymbol{Z}_{n\alpha})^2]$. $\square$

Proof of Theorem 2. Recall that part (a) follows directly from Theorem 3.1. The upper bound in part (b) follows from Theorem 5.1 (a) together with Markov's inequality. The lower bound in part (b) follows from Theorem 5.1 (b) together with (5). $\square$

References

[1] Achlioptas, D., Chtcherba, A., Istrate, G. & Moore, C., The phase transition in 1-in-k-SAT and NAE-3-SAT, in *Proceedings of the 12th ACM–SIAM Symposium on Discrete Algorithms* (Washington, DC, 2001), pp. 721–722. SIAM, Philadelphia, PA, 2001.

[2] Achlioptas, D. & Naor, A., The two possible values of the chromatic number of a random graph. *Ann. of Math.*, 162 (2005), 1335–1351.

[3] Achlioptas, D., Naor, A. & Peres, Y., Rigorous location of phase transitions in hard optimization problems. *Nature*, 435 (2005), 759–764.

[4] Achlioptas, D. & Peres, Y., The threshold for random k-SAT is $2^k(\ln 2 - O(k))$, in *Proceedings of the 35th Annual ACM Symposium on Theory of Computing* (New York, 2013), pp. 223–231. ACM, New York, 2003.

[5] Achlioptas, D. & Ricci-Tersenghi, F., Random formulas have frozen variables. *SIAM J. Comput.*, 39 (2009), 260–280.

[6] Aldous, D. J., The $\zeta(2)$ limit in the random assignment problem. *Random Structures Algorithms*, 18 (2001), 381–418.

[7] — Open problems. Posted online, 2003.

[8] Aldous, D. J. & Lyons, R., Processes on unimodular random networks. *Electron. J. Probab.*, 12 (2007), 1454–1508.

[9] Aldous, D. J. & Steele, J. M., The objective method: probabilistic combinatorial optimization and local weak convergence, in *Probability on Discrete Structures*, Encyclopaedia Math. Sci., 110, pp. 1–72. Springer, Berlin, 2004.

[10] Barbier, J., Krząkała, F., Zdeborová, L. & Zhang, P., The hard-core model on random graphs revisited. *J. Phys. Conf. Ser.*, 473 (2013), 012021, 9 pp.

[11] Bayati, M., Gamarnik, D. & Tetali, P., Combinatorial approach to the interpolation method and scaling limits in sparse random graphs, in *Proceedings of the* 2010 *ACM International Symposium on Theory of Computing* (New York, 2010), pp. 105–114. ACM, New York, 2010.

[12] Benjamini, I. & Schramm, O., Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6 (2001), 1–13.

[13] Bollobás, B., The independence ratio of regular graphs. *Proc. Amer. Math. Soc.*, 83 (1981), 433–436.

[14] Bollobás, B., Borgs, C., Chayes, J. T., Kim, J. H. & Wilson, D. B., The scaling window of the 2-SAT transition. *Random Structures Algorithms*, 18 (2001), 201–256.

[15] Braunstein, A., Mézard, M. & Zecchina, R., Survey propagation: an algorithm for satisfiability. *Random Structures Algorithms*, 27 (2005), 201–226.

[16] Chvatal, V. & Reed, B., Mick gets some (the odds are on his side), in *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science* (Pittsburgh, 1992), pp. 620–627. IEEE Computer Society, Washington, DC, 1992.

[17] Coja-Oghlan, A., Efthymiou, C. & Hetterich, S., On the chromatic number of random regular graphs. *J. Combin. Theory Ser. B*, 116 (2016), 367–439.

[18] Coja-Oghlan, A. & Panagiotou, K., The asymptotic k-SAT threshold. *Adv. Math.*, 288 (2016), 985–1068.

[19] Dembo, A. & Montanari, A., Gibbs measures and phase transitions on sparse random graphs. *Braz. J. Probab. Stat.*, 24 (2010), 137–211.

[20] Dembo, A., Montanari, A., Sly, A. & Sun, N., The replica symmetric solution for Potts models on d-regular graphs. *Comm. Math. Phys.*, 327 (2014), 551–575.

[21] Dembo, A., Montanari, A. & Sun, N., Factor models on locally tree-like graphs. *Ann. Probab.*, 41 (2013), 4162–4213.

[22] Dembo, A. & Zeitouni, O., *Large Deviations Techniques and Applications*. Stochastic Modelling and Applied Probability, 38. Springer, Berlin–Heidelberg, 2010.

[23] Dietzfelbinger, M., Goerdt, A., Mitzenmacher, M., Montanari, A., Pagh, R. & Rink, M., Tight thresholds for cuckoo hashing via XORSAT, in *Automata, Languages and Programming* (Bordeaux, 2010), pp. 213–225. Springer, Berlin–Heidelberg, 2010.

[24] Ding, J. & Sly, N. A. S, Satisfiability threshold for random regular NAE-SAT, in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing* (New York, 2014), pp. 814–822. ACM, New York, 2014.

[25] Frieze, A. & Łuczak, T., On the independence and chromatic numbers of random regular graphs. *J. Combin. Theory Ser. B*, 54 (1992), 123–132.

[26] Frieze, A. & Suen, S., Analysis of two simple heuristics on a random instance of k-SAT. *J. Algorithms*, 20 (1996), 312–355.

[27] Fu, Y. & Anderson, P. W., Application of statistical mechanics to NP-complete problems in combinatorial optimisation. *J. Phys. A*, 19 (1986), 1605–1620.

[28] Grimmett, G. R. & McDiarmid, C. J. H., On colouring random graphs. *Math. Proc. Cambridge Philos. Soc.*, 77 (1975), 313–324.

[29] Hartmann, A. K. & Weigt, M., *Phase Transitions in Combinatorial Optimization Problems*. Wiley-VCH Verlag, Weinheim, 2005.

[30] Janson, S., Łuczak, T. & Rucinski, A., *Random Graphs*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.

[31] Karp, R. M., Reducibility among combinatorial problems, in *Complexity of Computer Computations*, pp. 85–103. Plenum, New York, 1972.

[32] Kirousis, L. M., Kranakis, E., Krizanc, D. & Stamatiou, Y. C., Approximating the unsatisfiability threshold of random formulas. *Random Structures Algorithms*, 12 (1998), 253–269.

[33] Krząkała, F., Montanari, A., Ricci-Tersenghi, F., Semerjian, G. & Zdeborová, L., Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc. Natl. Acad. Sci. USA*, 104 (2007), 10318–10323.

[34] Maneva, E., Mossel, E. & Wainwright, M. J., A new look at survey propagation and its generalizations. *J. ACM*, 54 (2007), Art. 17, 41 pp.

[35] McKay, B. D., Independent sets in regular graphs of high girth. *Ars Combin.*, 23 (1987), 179–185.

[36] Mézard, M. & Montanari, A., *Information, Physics, and Computation*. Oxford Graduate Texts. Oxford Univ. Press, Oxford, 2009.
[37] Mézard, M. & Parisi, G., Replicas and optimization. *J. Physique Lett.*, 46 (1985), 771–778.
[38] — The Bethe lattice spin glass revisited. *Eur. Phys. J. B*, 20 (2001), 217–233.
[39] Panchenko, D., *The Sherrington–Kirkpatrick Model*. Springer Monographs in Mathematics. Springer, New York, 2013.
[40] Pittel, B. & Sorkin, G. B., The satisfiability threshold for k-XORSAT. *Combin. Probab. Comput.*, 25 (2016), 236–268.
[41] Rivoire, O., *Phases vitreuses, optimisation et grandes déviations*. Ph.D. Thesis, Université Paris-Sud, Orsay, 2005.
[42] Talagrand, M., The Parisi formula. *Ann. of Math.*, 163 (2006), 221–263.
[43] Wormald, N. C., Differential equations for random processes and random graphs. *Ann. Appl. Probab.*, 5 (1995), 1217–1235.
[44] — Models of random regular graphs, in *Surveys in Combinatorics* (Canterbury, 1999), London Math. Soc. Lecture Note Ser., 267, pp. 239–298. Cambridge Univ. Press, Cambridge, 1999.

Jian Ding
Statistics Department
University of Chicago
5747 South Ellis Avenue
Chicago, IL 60637
U.S.A.
jianding@galton.uchicago.edu

Nike Sun
Department of Statistics
Stanford University
Sequoia Hall, 390 Serra Mall
Stanford, CA 94305
U.S.A.
nikesun@berkeley.edu

Allan Sly
Department of Statistics
University of California, Berkeley
367 Evans Hall
Berkeley, CA 94720
U.S.A.
sly@stat.berkeley.edu

and

Mathematical Sciences Institute
Australian National University
John Dedman Building 27
Union Lane, Canberra, ACT 0200
Australia

Received March 23, 2014
Received in revised form January 11, 2016