FERROMAGNETIC POTTS MODEL: REFINED #BIS-HARDNESS AND RELATED RESULTS*

ANDREAS GALANIS†, DANIEL ŠTEFANKOVIČ‡, ERIC VIGODA§, AND LINJI YANG¶

Abstract. Recent results establish for the hard-core model (and more generally for 2-spin antiferromagnetic systems) that the computational complexity of approximating the partition function on graphs of maximum degree Δ undergoes a phase transition that coincides with the uniqueness/non-uniqueness phase transition on the infinite Δ-regular tree. For the ferromagnetic Potts model we investigate whether analogous hardness results hold. Goldberg and Jerrum showed that approximating the partition function of the ferromagnetic Potts model is at least as hard as approximating the number of independent sets in bipartite graphs, so-called #BIS-hardness. We improve this hardness result by establishing it for bipartite graphs of maximum degree Δ. To this end, we first present a detailed picture for the phase diagram for the infinite Δ-regular tree, giving a refined picture of its first-order phase transition and establishing the critical temperature for the coexistence of the disordered and ordered phases. We then prove for all temperatures below this critical temperature (corresponding to the region where the ordered phase "dominates") that it is #BIS-hard to approximate the partition function on bipartite graphs of maximum degree Δ. As a simple corollary of this result, we obtain that it is #BIS-hard to approximate the number of k-colorings on bipartite graphs of maximum degree Δ whenever $k \leq \Delta/(2\ln\Delta)$. The #BIS-hardness result for the ferromagnetic Potts model uses random bipartite regular graphs as a gadget in the reduction. The analysis of these random graphs relies on recent results establishing connections between the maxima of the expectation of their partition function, attractive fixpoints of the associated tree recursions, and induced matrix norms. In this paper we extend these connections to random regular graphs for all ferromagnetic models. Using these connections, we establish the Bethe prediction for every ferromagnetic spin system on random regular graphs, which says roughly that the expectation of the log of the partition function Z is the same as the log of the expectation of Z. As a further consequence of our results, we prove for the ferromagnetic Potts model that the Swendsen–Wang algorithm is torpidly mixing (i.e., exponentially slow convergence to its stationary distribution) on random Δ-regular graphs at the critical temperature for sufficiently large q.

Key words. ferromagnetic Potts model, approximate counting, spin systems, phase transitions, random regular graphs

1. Background.

1.1. Spin systems. We study the ferromagnetic Potts model and present tools which are useful for any ferromagnetic spin system on random regular graphs. Hence we begin with a general definition of a spin system.

A spin system is defined, for an n-vertex graph $G = (V, E)$ and integer $q \geq 2$, on the space Ω of configurations σ which are assignments $\sigma : V \to [q]$. The model is characterized by its energy or Hamiltonian $H(\sigma)$ which is a function of the spin assignments to the vertices. In the classical examples of the Ising ($q = 2$) and Potts ($q \geq 3$) models without external field, the Hamiltonian $H(\sigma)$ is the number of monochromatic edges in σ. Each configuration has a weight $w(\sigma) = \exp(-\beta H(\sigma))$ for a parameter β corresponding to the "inverse temperature" which controls the strength of edge interactions.

In our general setup, a *specification* of a q-state spin model is defined by a symmetric $q \times q$ interaction matrix $\mathbf{B} = \{B_{ij}\}_{i,j\in[q]}$ with nonnegative entries. For a graph $G = (V, E)$, the weight of a configuration $\sigma : V \to [q]$ is given by

$$w_G(\sigma) = \prod_{\{u,v\}\in E} B_{\sigma(u),\sigma(v)}.$$

We will occasionally drop the subscript G when the graph under consideration is clear from context. The *Gibbs distribution* $\mu = \mu_G$ is defined as $\mu(\sigma) = w(\sigma)/Z$, where $Z = Z_G(\mathbf{B}) = \sum_\sigma w(\sigma)$ is the *partition function*. We remark here that many of our results also apply to models with arbitrary external fields since we will work with Δ-regular graphs and in this case the external field can be incorporated into the interaction matrix.

The Ising ($q = 2$) and Potts ($q > 2$) models have interaction matrices with diagonal entries $B := \exp(-\beta)$ and off-diagonal entries 1. The models are called ferromagnetic if $B > 1$ since then neighboring spins prefer to align and are antiferromagnetic if $B < 1$. The hard-core model is an example of a 2-spin antiferromagnetic system, its interaction matrix is defined so that Ω is the set of independent sets of G and, for activity (external field) $\lambda > 0$, a configuration $\sigma \in \Omega$ has weight $w(\sigma) = \lambda^{|\sigma|}$ (with $|\sigma|$ denoting the cardinality of the independent set σ).

1.2. Ferromagnetic models. In this paper, we will focus on ferromagnetic models and pay special attention to the ferromagnetic Potts model. We are not aware of a general definition of ferromagnetic and antiferromagnetic models. We use the following notions which generalize the analogous notions for 2-spin and for the Potts model. The ferromagnetic definition captures that neighboring spins prefer to align (see Observation 2 below).

To avoid degenerate cases,[1] we assume throughout this paper that the interaction matrix $\mathbf{B}$ is ergodic, that is, irreducible and aperiodic. Hence, by the Perron–Frobenius theorem (since $\mathbf{B}$ has nonnegative entries) the eigenvalue of $\mathbf{B}$ with the largest magnitude is positive.

[1] If $\mathbf{B}$ is reducible, by a suitable permutation of the labels of the spins, $\mathbf{B}$ can be put in a block diagonal form where each of the blocks is either irreducible or zero. Such a model can be studied by considering the induced submodels of each block corresponding to irreducible symmetric matrices, since the partition function for the original model is simply the sum of the partition functions of each of these submodels. If $\mathbf{B}$ is periodic, and since $\mathbf{B}$ is symmetric, its period must be two. Such a model is interesting only on bipartite graphs (otherwise the partition function is zero), and the focus of our general results is for random Δ-regular graphs which are nonbipartite with high probability.

Definition 1. *A model is called* ferromagnetic *if* $\mathbf{B}$ *is positive definite. Equivalently we have that all of its eigenvalues are positive and also that*

$$\mathbf{B} = \hat{\mathbf{B}}^{\intercal}\hat{\mathbf{B}}$$

for some $q \times q$ *matrix* $\hat{\mathbf{B}}$.

In contrast to the above notion of a ferromagnetic system, in [17] a model is called *antiferromagnetic* if all of the eigenvalues of $\mathbf{B}$ are negative except for the largest (which, as noted above, is positive). Note that when the number of spins is greater than 2, there are models which are neither ferromagnetic nor antiferromagnetic.

The most alluring aspect of this definition is that for ferromagnetic models, neighboring vertices prefer to have the same spin. To see this, the following more general inequality is proved in [17], which is a simple application of the Cauchy–Schwarz inequality.

Observation 2. Let $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^q_{\geq 0}$ with $\|\mathbf{z}_1\|_1 = \|\mathbf{z}_2\|_1 = 1$. For ferromagnetic $\mathbf{B}$, we have

$$(\mathbf{z}_1^{\intercal}\mathbf{B}\mathbf{z}_1)(\mathbf{z}_2^{\intercal}\mathbf{B}\mathbf{z}_2) \geq (\mathbf{z}_1^{\intercal}\mathbf{B}\mathbf{z}_2)^2.$$

Equality holds iff $\mathbf{z}_1 = \mathbf{z}_2$. For antiferromagnetic $\mathbf{B}$, the inequality is reversed.

Observe that if we plug in to the above inequality the vectors with a single 1 in the positions i and j, respectively, we obtain that any two spins i, j induce a ferromagnetic 2-spin system.

As observed in [17], an appealing aspect of defining ferromagnetism in terms of the signature of the interaction matrix is that the definition remains invariant in the presence of external fields. More precisely, for Δ-regular graphs, any external field can be incorporated into the interaction matrix by a congruence transformation of the matrix $\mathbf{B}$. The modified interaction matrix has the same number of positive, zero, and negative eigenvalues as the original (this follows from the Sylvester's law of inertia), and hence it remains ferromagnetic/antiferromagnetic.

1.3. Known connections to phase transitions. Exact computation of the partition function is #P-complete, even for very restricted classes of graphs [24]. Hence we focus on whether there is a fully polynomial (randomized or deterministic) approximation scheme, a so-called FPRAS or FPTAS.

One of our goals in this paper is to refine our understanding of connections between approximating the partition function on graphs of maximum degree Δ with phase transitions on the infinite Δ-regular tree $\mathbb{T}_\Delta$. A phase transition of particular interest in the infinite tree $\mathbb{T}_\Delta$ is the uniqueness/nonuniqueness threshold. Roughly speaking, in the uniqueness phase, if one fixes a so-called boundary condition which is a configuration σ_ℓ (for instance, an independent set in the hard-core model) on the vertices distance ℓ from the root, then in the Gibbs distribution conditioned on this configuration, is the root "unbiased"? Specifically, for all sequences (σ_ℓ) of boundary conditions, in the limit $\ell \to \infty$, does the root have the same marginal distribution? If so, there is a unique Gibbs measure on the infinite tree and hence we say the model is in the uniqueness region. If there are sequences of boundary conditions which influence the root in the limit, then we say the model is in the nonuniqueness region.

For 2-spin antiferromagnetic spin systems, it was shown that there is an FPTAS for estimating the partition function for graphs of maximum degree Δ when the infinite tree $\mathbb{T}_\Delta$ is in the uniqueness region [32]. On the other side, unless NP=RP, there is no FPRAS for the partition function for Δ-regular graphs when $\mathbb{T}_\Delta$ is in the nonuniqueness region [44] (see also [16]). Recently, an analogous NP-hardness result was shown for approximating the number of k-colorings on triangle-free Δ-regular graphs for even k when $k < \Delta$ [17]. In contrast to the above inapproximability results for antiferromagnetic systems, for the ferromagnetic Ising model with or without external field [29] and for 2-spin ferromagnetic spin systems without external field [23] there is an FPRAS for all graphs. The situation for ferromagnetic multispin models, the ferromagnetic Potts being the most prominent example, is more intricate.

#BIS refers to the problem of computing the number of independent sets in bipartite graphs. A series of results has presented evidence that there is unlikely to be a polynomial-time algorithm for #BIS, since a number of unsolved counting problems have been shown to be #BIS-easy (for example, see [14, 3, 8]). The growing anecdotal evidence for #BIS-hardness suggests that the problem is intractable, though weaker than NP-hardness. More recently, it was shown in [7] that for antiferromagnetic 2-spin models it is #BIS-hard to approximate the partition function on *bipartite* graphs of maximum degree Δ when the parameters of the model lie in the nonuniqueness region of the infinite Δ-regular tree $\mathbb{T}_\Delta$. Also, for ferromagnetic 2-spin models with external field, [33] shows #BIS-hardness for some region of the parameter space (note that the known regions of the parameter space where an FPRAS exists (see [23, 33, 20]) do not yet completely complement the #BIS-hardness result).

1.4. Outline of results. Our focus in this paper is on understanding the behavior of ferromagnetic spin systems. Our main tools are bipartite random regular graphs and random regular graphs. Whether we use bipartite or general graphs depends on the context, and we use whichever yields the strongest results in that context. For instance, we establish #BIS-hardness for the ferromagenetic Potts model in section 2.1; to obtain hardness results on the class of bipartite graphs we use bipartite random regular graphs as the core of the gadget. In section 2.3 we establish results for the Swendsen–Wang algorithm; such results are more interesting for general graphs and hence we prove this result for random regular graphs.

In [17] we established concentration of the partition function for general spin systems on bipartite random regular graphs. At first glance the picture for random regular graphs is more complicated than for their bipartite counterparts since the connection to trees is less clear for general models; however, for ferromagnetic models an analogous connection holds, as we will establish in section 3.2. For ferromagnetic systems, we establish concentration on random regular graphs as detailed in section 3.1. As a consequence we establish the so-called Bethe prediction for random regular graphs as discussed in section 3.1.

2. Results for the Potts model.

2.1. #BIS-hardness for the Potts model. Goldberg and Jerrum [21] showed that approximating the partition function of the ferromagnetic Potts model is #BIS-hard, hence it appears likely that the ferromagnetic Potts model is inapproximable for general graphs. We refine this #BIS-hardness result for the ferromagnetic Potts model. We prove that approximating the partition function for the ferromagnetic Potts model on *bipartite* graphs of maximum degree Δ is #BIS-hard for temperatures below the appropriate phase transition point in the infinite tree $\mathbb{T}_\Delta$. The appropriate phase transition in the Potts model is not the uniqueness/nonuniqueness threshold, but rather it is the ordered/disordered phase transition which occurs at $B = \mathfrak{B}_o$ as explained in the next section.

Formally, we study the following problem.

Name. #BipFerroPotts(q,B,Δ).
Instance. A bipartite graph G with maximum degree Δ.
Output. The partition function for the q-state Potts model on G.

We use the notion of approximation-preserving reductions, denoted as $\leq_{\mathrm{AP}}$, formally defined in [14] (roughly, for counting problems $\#\Pi_1$ and $\#\Pi_2$, $\#\Pi_1 \leq_{\mathrm{AP}} \#\Pi_2$ implies that the existence of an FPRAS for $\#\Pi_2$ implies the existence of an FPRAS for $\#\Pi_1$). We can now formally state our main result.

THEOREM 3. *For all $q \geq 3$, for all $\Delta \geq 3$, for the ferromagnetic q-state Potts model, for any $B > \mathfrak{B}_o$,*

$$\#\mathrm{BIS} \leq_{\mathrm{AP}} \#\textsc{BipFerroPotts}(q, B, \Delta),$$

where $\mathfrak{B}_o$ is given by (3).

Theorem 3 has a simple, yet interesting, consequence for the problem of approximately counting k-colorings on *bipartite* graphs of maximum degree Δ. Recall that for general graphs of maximum degree Δ, approximately counting k-colorings is NP-hard whenever $k < \Delta$ (and k is even) due to the result of [17]. Theorem 3 yields #BIS-hardness for bipartite k-colorings whenever $k \leq \Delta/(2\log\Delta)$. Formally, we are interested in the following problem.

Name. #BipColorings(k,Δ).
Instance. A bipartite graph G with maximum degree Δ.
Output. The number of proper k-colorings of G.

We use a relatively simple reduction from the ferromagnetic Potts model to bipartite colorings, first observed in [14], which works even for bounded-degree graphs. Theorem 3 then yields the following corollary (which is proved in section 9.2).

COROLLARY 4. *For all $k, \Delta \geq 3$, whenever $k \leq \Delta/(2\ln\Delta)$, it holds that*

$$\#\mathrm{BIS} \leq_{\mathrm{AP}} \#\textsc{BipColorings}(k, \Delta).$$

It would be interesting to extend Corollary 4 to all $k < \Delta$.

2.2. Potts model phase diagram. To understand the critical point $\mathfrak{B}_o$ we need to delve into the nature of the phase transition in the ferromagnetic Potts model on the infinite Δ-regular tree $\mathbb{T}_\Delta$. We focus on how the phase transition manifests on a random Δ-regular graph.

For a configuration $\sigma \in \Omega$, denote the set of vertices assigned spin i by $\sigma^{-1}(i)$. Let $\triangle_q$ denote the $(q-1)$-simplex, where we recall that

$$\triangle_t = \{(x_1, x_2, \ldots, x_t) \in \mathbb{R}^t \,|\, \textstyle\sum_{i=1}^t x_i = 1 \text{ and } x_i \geq 0 \text{ for } i = 1, \ldots, t\}.$$

We refer to $\boldsymbol{\alpha} \in \triangle_q$ as a *phase.* For a phase $\boldsymbol{\alpha}$, denote the set of configurations with frequencies of colors given by $\boldsymbol{\alpha}$ as[2]

$$\Sigma^{\boldsymbol{\alpha}} = \left\{\sigma : V \to \{1, \ldots, q\} \,\middle|\, |\sigma^{-1}(i)| = \alpha_i n \text{ for } i = 1, \ldots, q\right\},$$

and denote the partition function restricted to these configurations by

$$Z_G^{\boldsymbol{\alpha}} = \textstyle\sum_{\sigma \in \Sigma^{\boldsymbol{\alpha}}} w_G(\sigma).$$

[2]Technically we need to define $\Sigma^{\boldsymbol{\alpha}} = \left\{\sigma : V \to [q] \,\middle|\, |\sigma^{-1}(i) \cap V| = \hat{\alpha}_i\right\}$, where $\{\hat{\alpha}_i\}$ are $\{\alpha_i n\}$ rounded in a canonical fashion so that their sum is preserved (for example, using "cascade rounding").

Let $\mathcal{G}$ denote the uniform distribution over Δ-regular graphs with n vertices (for Δn even). Denote the exponent of the first moment as

$$\Psi_1(\boldsymbol{\alpha}) := \Psi_1^{\mathbf{B}}(\boldsymbol{\alpha}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha}}\big]. \tag{1}$$

We derive the expression for Ψ_1 in section 4. Those $\boldsymbol{\alpha}$ which are global maxima of Ψ_1 we refer to as *dominant* phases. We will see in section 3.2 that, for all ferromagnetic models, roughly speaking, the candidates for dominant phases correspond to stable fixpoints of the so-called tree recursions.

For the ferromagnetic Potts model, there will be two types of phases with particular interest; we refer to these two types as the disordered phase and the ordered phase. The disordered phase is the uniform vector $\boldsymbol{\alpha} = (1/q, \ldots, 1/q)$. The ordered phase refers to a phase with one color dominating in the following sense: one coordinate is equal to $a > 1/q$ and the other $q-1$ coordinates are equal to $(1-a)/(q-1)$. Due to the symmetry of the Potts model, when the ordered phase dominates, in fact, the q symmetric ordered phases dominate. These ordered phases have a specific $a = a(q, B, \Delta)$ which corresponds to a fixpoint of the tree recursions. The exact definition of this marginal a is not important at this stage, and hence we defer its definition to a more detailed discussion which takes place in section 8 (see (38)).

One of the difficulties for the Potts model is that the nature of the uniqueness/nonuniqueness phase transition on $\mathbb{T}_\Delta$ is inherently different from that of the Ising model. The ferromagnetic Ising model undergoes a second-order phase transition on $\mathbb{T}_\Delta$ which manifests itself on random Δ-regular graphs in the following manner. In the uniqueness region the disordered phase dominates, and in the nonuniqueness region the two ordered phases dominate.

In contrast, the ferromagnetic Potts model undergoes a first-order phase transition at the critical activity $\mathfrak{B}_u$. For $B < \mathfrak{B}_u$ there is a unique Gibbs measure on $\mathbb{T}_\Delta$. For $B \geq \mathfrak{B}_u$ there are multiple Gibbs measures on $\mathbb{T}_\Delta$; however, there is a second critical activity $\mathfrak{B}_o$ corresponding to the disordered/ordered phase transition: for $B \leq \mathfrak{B}_o$ the disordered phase dominates, and for $B \geq \mathfrak{B}_o$ the ordered phases dominate (and at the critical point $\mathfrak{B}_o$ all of these $q+1$ phases dominate).

We present a detailed picture of the phase diagram for the ferromagnetic Potts model. Previously, Häggström [26] established the uniqueness threshold $\mathfrak{B}_u$ by studying percolation in the random cluster representation. In addition, Dembo et al. [12, 13] studied the ferromagnetic Potts model (including the case with an external field) and proved that for $B > \mathfrak{B}_u$, either the disordered or the q ordered phases are dominant, but they did not establish the precise regions where each phase dominates. For the simpler case of the complete graph (known as the Curie–Weiss model), [10] detailed the phase diagram.

Häggström [26] established that the uniqueness/nonuniqueness threshold for the infinite tree $\mathbb{T}_\Delta$ occurs at $\mathfrak{B}_u$ which is the unique value of B for which the following polynomial has a double root in $(0, 1)$:

$$(q-1)x^{\Delta} + (2 - B - q)\, x^{\Delta-1} + Bx - 1. \tag{2}$$

The disordered phase is dominant in the uniqueness region and continues to dominate until the following activity (which was considered by Peruggi, Di Liberto, and Monroy [41]):

$$\mathfrak{B}_o := \frac{q-2}{(q-1)^{(1-2/\Delta)} - 1}. \tag{3}$$

Finally, Häggström [26] considers the following activity $\mathfrak{B}_{rc}$, which he conjectures is a (second) threshold for uniqueness of the random cluster model, defined as

$$\mathfrak{B}_{rc} := 1 + \frac{q}{\Delta - 2}.$$

Note $\mathfrak{B}_u < \mathfrak{B}_o < \mathfrak{B}_{rc}$.

We prove the following picture for the phase diagram for the ferromagnetic Potts model in section 8. Note that to prove that a function has a local maximum at a point, a standard approach is to show that its Hessian matrix is negative definite. We often need this stronger condition in our proofs, hence we use the following definition. Those dominant phases $\boldsymbol{\alpha}$ with negative definite Hessian are called *Hessian dominant* phases. Note that dominant phases always exist but a dominant phase can fail to be Hessian (when some eigenvalue of the underlying Hessian is equal to zero). In section 3.2, we give an alternative formulation of the Hessian condition in terms of the local stability of fixpoints of the tree recursions.

THEOREM 5. *For the ferromagnetic Potts model the following holds at activity B:*

$B < \mathfrak{B}_u$**.** *There is a unique infinite-volume Gibbs measure on $\mathbb{T}_\Delta$. The disordered phase is Hessian dominant, and there are no other local maxima of Ψ_1.*

$\mathfrak{B}_u < B < \mathfrak{B}_{rc}$. *The local maxima of Ψ_1 are the disordered phase $\mathbf{u}$ and the q ordered phases (the ordered phases are permutations of each other). All of these $q+1$ phases are Hessian local maxima. Moreover:*

$\mathfrak{B}_u < B < \mathfrak{B}_o$: *The disordered phase is Hessian dominant.*

$B = \mathfrak{B}_o$: *Both the disordered phase and the ordered phases are Hessian dominant.*

$\mathfrak{B}_o < B < \mathfrak{B}_{rc}$: *The ordered phases are Hessian dominant.*

$B \geq \mathfrak{B}_{rc}$. *The q ordered phases (which are permutations of each other) are Hessian dominant. For $B > \mathfrak{B}_{rc}$ there are no other local maxima of Ψ_1.*

2.3. Swendsen–Wang algorithm. An algorithm of particular interest for the ferromagnetic Potts model is the Swendsen–Wang algorithm. The Swendsen–Wang algorithm is an ergodic Markov chain whose stationarity distribution is the Gibbs distribution. It utilizes the random cluster representation to overcome potential "bottlenecks" for rapid mixing that are expected to arise in the nonuniqueness region. As a consequence of the above picture for the phase diagram on the infinite tree $\mathbb{T}_\Delta$ and our tools for analyzing random regular graphs, we can prove torpid mixing of the Swendsen–Wang algorithm at the disordered/ordered phase transition point $\mathfrak{B}_o$. (Torpid mixing means that the mixing time is exponentially slow.)

The Swendsen–Wang algorithm utilizes the random cluster representation (see [25]) of the Potts model to potentially overcome bottlenecks that obstruct the simpler Glauber dynamics. It is formally defined as follows. From a configuration $X_t \in \Omega$,

1. let M be the set of monochromatic edges in X_t;
2. for each edge $e \in M$, delete it with probability $1/B$; let M' denote the set of monochromatic edges that were not deleted;
3. in the graph (V, M'), for each connected component, choose a color uniformly at random from $[q]$ and assign all vertices in that component the chosen color. Let X_{t+1} denote the resulting spin configuration.

There are few results establishing rapid mixing of the Swendsen–Wang algorithm beyond what is known for the Glauber dynamics, see [47] for recent progress showing rapid mixing on the two-dimensional lattice. However, there are several results establishing torpid mixing of the Swendsen–Wang algorithm at a critical value for the

q-state ferromagnetic Potts model: on the complete graph ($q \geq 3$) [22], on Erdös–Rényi random graphs ($q \geq 3$) [9], and on the d-dimensional integer lattice $\mathbb{Z}^d$ (q sufficiently large) [4, 5].

Using our detailed picture of the phase diagram of the ferromagnetic Potts model and our generic second moment analysis for ferromagnetic models on random regular graphs which we explain in a moment, we establish torpid mixing on random Δ-regular graphs at the phase coexistence point $\mathfrak{B}_o$.

THEOREM 6. *For all $\Delta \geq 3$ and $q \geq 2\Delta/\log\Delta$, with probability $1 - o(1)$ over the choice of a random Δ-regular graph, for the ferromagnetic Potts model with $B = \mathfrak{B}_o$, the Swendsen–Wang algorithm has mixing time $\exp(\Omega(n))$.*

We believe that the lower bound on q in Theorem 6 is an artifact of our proof; see Remark 34 in section 10 for details.

3. Results for ferromagnetic models.

3.1. Second moment and Bethe prediction. We analyze the Gibbs distribution on random Δ-regular graphs using second moment arguments. The challenging aspect of the second moment is determining the phase that dominates, as we will describe more precisely momentarily. In a straightforward analysis of the second moment, this reduces to an optimization problem over q^4 variables for a complicated expression. Even for $q = 2$ tackling this requires significant effort (see, for example, [40] for the hard-core model).

In a recent paper [17] we analyzed antiferromagnetic systems on *bipartite* random Δ-regular graphs, to use as gadgets for inapproximability results. In that work we presented a new approach for simplifying the analysis of the second moment for antiferromagnetic models using the theory of matrix norms. In this paper we extend that approach using the theory of matrix norms to analyze the second moment for random Δ-regular graphs (nonbipartite) for ferromagnetic systems. We obtain a short, elegant proof that the exponential order of the second moment is twice the exponential order of the first moment.

Denote the leading term of the second moment as

$$\Psi_2(\boldsymbol{\alpha}) := \Psi_2^{\mathbf{B}}(\boldsymbol{\alpha}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[\, (Z_G^{\boldsymbol{\alpha}})^2 \,\big]. \tag{4}$$

Our main technical result is the analysis of the second moment for ferromagnetic models. We will relate the maximum of the second moment to the maximum of the first moment. To analyze the second moment we need to determine the phase $\boldsymbol{\alpha}$ that maximizes Ψ_2. We will first show how to reexpress the critical points of Ψ_1 in a form that can be readily expressed in terms of matrix norms (see section 5.1). Then, using the Cholesky decomposition of the interaction matrix $\mathbf{B}$ and properties of matrix norms we will show that the second moment is maximized at a phase which is a tensor product of the dominant phases of the first moment. This results in the following theorem, which is proved in section 5.2.

THEOREM 7. *For a ferromagnetic model with interaction matrix $\mathbf{B}$,*

$$\max_{\boldsymbol{\alpha}} \Psi_2(\boldsymbol{\alpha}) = 2 \max_{\boldsymbol{\alpha}} \Psi_1(\boldsymbol{\alpha}).$$

More specifically, for dominant $\boldsymbol{\alpha}$, $\Psi_2(\boldsymbol{\alpha}) = 2\Psi_1(\boldsymbol{\alpha})$.

Combining Theorem 7 with an elaborate variance analysis known as the small subgraph conditioning method allows us to obtain a lower bound on $Z_G^{\boldsymbol{\alpha}}$ which matches its expectation up to a polynomial factor (see Lemma 16). In particular, we verify the so-called *Bethe prediction* (see [13, 12]) for general ferromagnetic models on random Δ-regular graphs, which is captured in our setting by (5) in the following theorem (the proof is in section 7.2). Our interest in the quantity $\mathbf{E}_{\mathcal{G}}[\log Z_G]$ stems from the fact that it gives information about the typical configurations on a random regular graph and hence Theorem 8 gives its value in terms of a much simpler quantity, $\log \mathbf{E}_{\mathcal{G}}[Z_G]$, which can be calculated much more easily; see section 4. (The equality in (5) is closely related to the cavity method; see [36].)

THEOREM 8. *Let* $\mathbf{B}$ *specify a ferromagnetic model. Then, if there exists a Hessian dominant phase, it holds that*

$$\lim_{n\to\infty} \frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G] = \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[Z_G]. \tag{5}$$

Note that for a ferromagnetic model the interaction matrix $\mathbf{B}$ is positive definite and hence the entries on the diagonal are all positive. Thus Z_G is always positive for every graph G (and hence $\log Z_G$ in (5) is well-defined).

Theorem 8 holds for all ferromagnetic models at any temperature. Dembo, Montanari, and Sun [13] consider general factor models on graph sequences converging locally to trees and verify the Bethe prediction when the underlying tree is in the uniqueness regime. In [12], the case of the ferromagnetic Potts model (with external field) is considered for graph sequences converging locally to trees and they obtain a general formula for the logarithm of the partition function.

Perhaps the most important conceptual content of Theorem 8 is that it shows that *all* ferromagnetic models, at any temperature, do not exhibit the complex behavior that other spin models, such as colorings or the antiferromagnetic Potts model, exhibit on random (regular) graphs. In particular, when the equality in (5) fails, we have the so-called condensation regime, and in that case calculating $\frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G]$ is a far more intricate task (see the recent works [1, 45]).

Theorem 8 can be extended to general models (not necessarily ferromagnetic) on random Δ-regular graphs under the stronger assumption that there is a unique semitranslation invariant Gibbs measure on $\mathbb{T}_\Delta$. In this setting, one also obtains the analogue of Theorem 7 and as a consequence concentration for $Z_G^{\boldsymbol{\alpha}}$ for the (unique) dominant phase $\boldsymbol{\alpha}$, which can be used to verify (in complete analogy) the Bethe prediction; see section 11.3 for more details.

3.2. Connection to tree recursions. As a consequence of Theorem 7, to analyze ferromagnetic models on random regular graphs, one only needs to analyze the first moment. To simplify the analysis of the first moment, we establish the following connection to the so-called tree recursions. An analogous connection was established in [17] for antiferromagnetic models on random bipartite Δ-regular graphs.

A key concept is the following recursions corresponding to the partition function on trees, and hence we refer to them as the (depth one) *tree recursions*:

$$\widehat{R}_i \propto \Big(\sum_{j=1}^{q} B_{ij} R_j\Big)^{\Delta-1}. \tag{6}$$

The fixpoints of the tree recursions are those $\mathbf{R} = (R_1, \dots, R_q)$ such that $\widehat{R}_i \propto R_i$ for all $i \in [q]$. We refer to a fixpoint $\mathbf{R}$ of the tree recursions as *Jacobian attractive* if the Jacobian at $\mathbf{R}$ has spectral radius less than 1. We prove the following theorem detailing the connections between the tree recursions and the critical points of the partition function for random regular graphs.

THEOREM 9. *Assume that the model is ferromagnetic. Jacobian attractive fixpoints of the (depth one) tree recursions are in one-to-one correspondence with the Hessian local maxima of* Ψ_1.

The above connection fails for antiferromagnetic models, e.g., for the antiferromagnetic Potts model the uniform distribution is a global maximum but it is not a stable fixpoint of the tree recursions for small enough temperature. (In fact, for antiferromagnetic models every solution of the tree recursions is a local maximum; see Remark 14.)

Using the above connection we establish the detailed picture for the dominant phases of the ferromagnetic Potts model as stated in Theorem 5.

4. Expressions for $\boldsymbol{\Psi_1}$ and $\boldsymbol{\Psi_2}$. In this section, we derive expressions for the first and second moments of $Z_G^{\boldsymbol{\alpha}}$, which will allow us to derive explicit expressions for the functions $\Psi_1(\boldsymbol{\alpha})$ and $\Psi_2(\boldsymbol{\alpha})$. Similar expressions have appeared in [12, section 2.1] in a slightly different form. Our exposition here is such that it provides a straightforward alignment with the analogous expressions in [17]. The minor differences are due to the model of Δ-regular random graphs, which in this paper is the pairing model $\mathcal{G}(n, \Delta)$. We first specify the model of Δ-regular random graphs.

The distribution $\mathcal{G}(n, \Delta)$ on Δ-regular multigraphs is generated by the following random process. For Δn even, consider the set $[\Delta n]$. Elements of $[\Delta n]$ will be called points. First, a random perfect matching of the Δn points is sampled. Then, for $i = 1, \ldots, n$, we identify the points $\Delta(i-1)+1, \ldots, \Delta i$ as a single vertex of a graph G. The edges of G are naturally induced by the edges of the random matching and hence every vertex has degree Δ. Note that G may contain parallel edges or self-loops. It is well-known that any property which holds asymptotically almost surely for the pairing model (i.e., with probability $1 - o(1)$ as $n \to \infty$) holds asymptotically almost surely for the uniform distribution on Δ-regular (simple) graphs as well; see, for example, [28]. This is going to be the case for our results.

Recall $\triangle_t$ denotes the simplex

$$\triangle_t = \{(x_1, x_2, \ldots, x_t) \in \mathbb{R}^t \mid \textstyle\sum_{i=1}^t x_i = 1 \text{ and } x_i \geq 0 \text{ for } i = 1, \ldots, t\}. \tag{7}$$

Let $G \sim \mathcal{G}(n, \Delta)$ and denote by V the vertex set of G. For a configuration $\sigma : V \to \{1, \ldots, q\}$, we denote the set of vertices assigned color i by $\sigma^{-1}(i)$. For $\boldsymbol{\alpha} \in \triangle_q$ and $n\boldsymbol{\alpha} \in \mathbb{Z}^q$, let

$$\Sigma^{\boldsymbol{\alpha}} = \{\sigma : V \to \{1, \ldots, q\} \,\big|\, |\sigma^{-1}(i)| = \alpha_i n \ \text{ for } i = 1, \ldots, q\},$$

that is, $\Sigma^{\boldsymbol{\alpha}}$ is the set of configurations σ which assign $\alpha_i n$ vertices of V the color i, for each $i \in [q]$. We are interested in the total weight $Z_G^{\boldsymbol{\alpha}}$ of configurations in $\Sigma^{\boldsymbol{\alpha}}$, namely,

$$Z_G^{\boldsymbol{\alpha}} = \textstyle\sum_{\sigma \in \Sigma^{\boldsymbol{\alpha}}} w_G(\sigma).$$

Note that $Z_G^{\boldsymbol{\alpha}}$ is a random variable, and as indicated earlier, we will look at its moments $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$ and $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2]$.

We begin with the first moment. For $\sigma \in \Sigma^{\boldsymbol{\alpha}}$ and $i, j \in [q]$, let $e_{ij}n$ denote the number of edges between vertices in $\sigma^{-1}(i)$ and $\sigma^{-1}(j)$. Clearly, $e_{ij} = e_{ji}$. It will be notationally convenient to reparameterize the variables e_{ij} as follows: for $i \neq j$ we set

$e_{ij} = \Delta x_{ij}$ and for $i = j$ we set $e_{ii} = \Delta x_{ii}/2$. For future use, when $G \sim \mathcal{G}(n, \Delta)$, we denote by $\mathbf{x}_G(\sigma)$ the random vector $(x_{11}, \ldots, x_{qq})$.

The number of perfect matchings between $2n$ vertices is $(2n-1)!! = (2n)!/(n!2^n)$. Under the convention that $0^0 \equiv 1$, we then have

$$\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}] = \binom{n}{\alpha_1 n, \ldots, \alpha_q n} \sum_{\mathbf{x}} \Biggl\{ \prod_i \binom{\Delta\alpha_i n}{\Delta x_{i1} n, \ldots, \Delta x_{iq} n} \times \frac{\bigl[\prod_{i \neq j} (\Delta x_{ij} n)!\bigr]^{1/2} \prod_i (\Delta x_{ii} n - 1)!!}{(\Delta n - 1)!!} \prod_{i,j} B_{ij}^{\Delta x_{ij} n/2} \Biggr\}, \tag{8}$$

where the sum ranges over all the possible values of the random vector $\mathbf{x}_G(\sigma)$. In particular, $\mathbf{x} = (x_{11}, \ldots, x_{qq})$ satisfies

$$\begin{aligned} \textstyle\sum_j x_{ij} = \alpha_i &\quad \bigl(\forall i \in [q]\bigr), \\ x_{ij} = x_{ji} \geq 0 &\quad \bigl(\forall i, j \in [q]\bigr). \end{aligned} \tag{9}$$

The first line in (8) accounts for the cardinality of $\Sigma^{\boldsymbol{\alpha}}$, while the second line is $\mathbf{E}_{\mathcal{G}}[w_G(\sigma)]$ for a fixed $\sigma \in \Sigma^{\boldsymbol{\alpha}}$, since by symmetry we may focus on any fixed σ. The first product is the number of ways to choose a partition of the points which is consistent with the values prescribed by $\mathbf{x}$, the fraction is the probability that the random matching connects the points as prescribed, and the last product is the weight of the configuration σ conditioned on $\mathbf{x}$.

We next consider the second moment of $Z_G^{\boldsymbol{\alpha}}$. The desired expression may be derived analogously to (8). For $(\sigma_1, \sigma_2) \in \Sigma^{\boldsymbol{\alpha}} \times \Sigma^{\boldsymbol{\alpha}}$, we need to compute the quantity $\mathbf{E}_{\mathcal{G}}[w_G(\sigma_1) w_G(\sigma_2)]$. To do this, for $i, k \in [q]$, let $\gamma_{ik} n = |\sigma_1^{-1}(i) \cap \sigma_2^{-1}(k)|$. The vector $\boldsymbol{\gamma}$ captures the overlap of the configurations σ_1, σ_2. Denote by $e_{ikjl} n$ the number of edges matching vertices in $\sigma_1^{-1}(i) \cap \sigma_2^{-1}(k)$ and $\sigma_1^{-1}(j) \cap \sigma_2^{-1}(l)$. We reparameterize as follows: for $(i,k) \neq (j,l)$ we set $e_{ikjl} = \Delta y_{ikjl}$ and for $(i,k) = (j,l)$ we set $e_{ikjl} = \Delta y_{ikjl}/2$.

$$\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2] = \sum_{\boldsymbol{\gamma}} \binom{n}{\gamma_{11} n, \ldots, \gamma_{qq} n} \sum_{\mathbf{y}} \Biggl\{ \prod_{i,k} \binom{\Delta\gamma_{ik} n}{\Delta y_{ik11} n, \ldots, \Delta y_{ikqq} n} \times \frac{\bigl[\prod_{(i,k) \neq (j,l)} (\Delta y_{ikjl} n)!\bigr]^{1/2} \prod_{i,k} (\Delta y_{ikik} n - 1)!!}{(\Delta n - 1)!!} \prod_{i,j,k,l} (B_{ij} B_{kl})^{\Delta y_{ikjl} n/2} \Biggr\}, \tag{10}$$

where the sums range over $\boldsymbol{\gamma} = (\gamma_{11}, \ldots, \gamma_{qq})$, $\mathbf{y} = (y_{1111}, \ldots, y_{qqqq})$ satisfying

$$\begin{gathered} \textstyle\sum_k \gamma_{ik} = \alpha_i \quad \bigl(\forall i \in [q]\bigr), \\ \textstyle\sum_i \gamma_{ik} = \alpha_k \quad \bigl(\forall k \in [q]\bigr), \\ \textstyle\sum_{j,l} y_{ikjl} = \gamma_{ik} \quad \bigl(\forall (i,k) \in [q]^2\bigr), \\ \gamma_{ik} \geq 0 \quad \bigl(\forall (i,k) \in [q]^2\bigr), \quad y_{ikjl} = y_{jlik} \geq 0 \quad \bigl(\forall (i,k,j,l) \in [q]^4\bigr). \end{gathered} \tag{11}$$

The sums in (8) and (10) are typically exponential in n. The most critical component of our arguments is to find the quantitative structure of configurations which determine the exponential order of the moments. Formally, we study the limits of

$\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha}}\big]$ and $\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}\big[(Z_G^{\boldsymbol{\alpha}})^2\big]$ as $n \to \infty$. These limits can be derived from (8) and (10) using Stirling's approximation formula. In particular, we shall use that for a constant $c > 0$ with cn even, we have

$$(12)\qquad \frac{1}{n}\ln\big[(cn)!\big] \sim c\ln n + c\ln c - c \ \text{ and } \ \frac{1}{n}\ln\big[(cn-1)!!\big] \sim \frac{c}{2}\ln n + \frac{c}{2}\ln c - \frac{c}{2}.$$

Under the usual conventions that $\ln 0 \equiv -\infty$ and $0\ln 0 \equiv 0$, the above formulas are correct even in the degenerate case $c = 0$.

We now derive asymptotics for the first moment $\mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha}}\big]$ in order to obtain the function $\Psi_1(\boldsymbol{\alpha})$; see (1). Applying (12) yields

$$(13)\qquad \begin{aligned} \Psi_1(\boldsymbol{\alpha}) := \lim_{n\to\infty}\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha}}\big] &= \max_{\mathbf{x}} \Upsilon_1(\boldsymbol{\alpha},\mathbf{x}),\\ \text{where}\quad \Upsilon_1(\boldsymbol{\alpha},\mathbf{x}) &:= (\Delta-1)f_1(\boldsymbol{\alpha}) + \Delta g_1(\mathbf{x}),\\ f_1(\boldsymbol{\alpha}) &:= \textstyle\sum_i \alpha_i \ln \alpha_i,\\ g_1(\mathbf{x}) &:= \textstyle\frac{1}{2}\sum_{i,j} x_{ij}\ln B_{ij} - \frac{1}{2}\sum_{i,j} x_{ij}\ln x_{ij}, \end{aligned}$$

defined on the region (9).

Completely analogously, for the second moment we obtain

$$(14)\qquad \begin{aligned} \Psi_2(\boldsymbol{\alpha}) := \lim_{n\to\infty}\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}\big[(Z_G^{\boldsymbol{\alpha}})^2\big] &= \max_{\boldsymbol{\gamma}}\max_{\mathbf{y}} \Upsilon_2(\boldsymbol{\gamma},\mathbf{y}),\\ \text{where}\quad \Upsilon_2(\boldsymbol{\gamma},\mathbf{y}) &:= (\Delta-1)f_2(\boldsymbol{\gamma}) + \Delta g_2(\mathbf{y}),\\ f_2(\boldsymbol{\gamma}) &:= \textstyle\sum_{i,k} \gamma_{ik}\ln \gamma_{ik},\\ g_2(\mathbf{y}) &:= \textstyle\frac{1}{2}\sum_{i,k,j,l} y_{ikj\ell}\ln(B_{ij}B_{kl}) - \frac{1}{2}\sum_{i,k,j,l} y_{ikjl}\ln y_{ikjl}, \end{aligned}$$

defined on the region (11).

Remark 10. It is useful to think of the second moment as the first moment of a paired-spin model with interaction matrix $\mathbf{B}\otimes\mathbf{B}$. Indeed, from (14), we can interpret $B_{ij}B_{kl}$ as the activity between the paired spins (i,k) and (j,l), thus giving the desired alignment.

5. Second moment analysis using induced matrix norms.

5.1. Critical points and matrix norms. It will be useful to reformulate function Ψ_1 into the following version, which will preserve the critical points and readily yield a formulation in terms of matrix norms. Let

$$(15)\qquad \Phi_1(\mathbf{R}) = \frac{\Delta}{2}\ln\Big(\sum_{i=1}^{q}\sum_{j=1}^{q} B_{ij}R_iR_j\Big) - (\Delta-1)\ln\Big(\sum_{i=1}^{q} R_i^{\Delta/(\Delta-1)}\Big),$$

where $\mathbf{R} = (R_1,\ldots,R_q)^{\intercal} \geq 0$, i.e., $\mathbf{R}$ has nonnegative entries. Let $p := \Delta/(\Delta-1)$. Note that (15) has the following appealing form:

$$(16)\qquad \exp(2\Phi_1(\mathbf{R})/\Delta) = \frac{\mathbf{R}^{\intercal}\mathbf{B}\mathbf{R}}{\|\mathbf{R}\|_p^2},$$

where $\|\mathbf{R}\|_p = (\sum_{i=1}^{n} R_i^p)^{1/p}$. This will allow us to use the techniques from the area of matrix norms in our arguments, more specifically, results on induced matrix norms.

The induced matrix norms will be denoted $\|\cdot\|_{p\to q'}$:

$$\|\mathbf{B}\|_{p\to q'} := \max_{\|\mathbf{z}\|_p=1} \|\mathbf{B}\mathbf{z}\|_{q'}. \tag{17}$$

Since we assume that $\mathbf{B}$ is ferromagnetic we have $\mathbf{B} = \hat{\mathbf{B}}^\intercal\hat{\mathbf{B}}$ and hence we can write

$$\exp(\Phi_1(\mathbf{R})/\Delta) = \frac{\|\hat{\mathbf{B}}\mathbf{R}\|_2}{\|\mathbf{R}\|_p}. \tag{18}$$

The next lemma describes the connection between Φ_1 and Ψ_1. We note that Φ_1 is not a reparameterization of Ψ_1; however, they do agree at the critical points. This is sufficient for our purpose: to understand the maxima of Ψ_1 it is enough to understand the maxima of Φ_1. The maximization

$$\max_{\mathbf{R}\geq 0} \frac{\|\hat{\mathbf{B}}\mathbf{R}\|_2}{\|\mathbf{R}\|_p} = \max_{\mathbf{R}} \frac{\|\hat{\mathbf{B}}\mathbf{R}\|_2}{\|\mathbf{R}\|_p} = \|\hat{\mathbf{B}}\|_{p\to 2} \tag{19}$$

is the induced $p \to 2$ matrix norm of $\hat{\mathbf{B}}$. The first equality in (19) follows from the fact that the maximum on the right-hand side (r.h.s.) of (16) is achieved for nonnegative $\mathbf{R}$ (this follows from the fact that $\mathbf{B}$ has nonnegative entries).

LEMMA 11. *There is a one-to-one correspondence between the fixpoints of the tree recursions and the critical points of Φ_1 (both considered for $R_i \geq 0$ in the projective space, that is, up to scaling by a constant). The following transformation $\mathbf{R} \mapsto \boldsymbol{\alpha}$ given by*

$$\alpha_i = R_i^{\Delta/(\Delta-1)} / \textstyle\sum_i R_i^{\Delta/(\Delta-1)} \tag{20}$$

yields a one-to-one-to-one correspondence between the critical points of Φ_1 and the critical points of Ψ_1 (in the region defined by $\alpha_i \geq 0$ and $\sum_i \alpha_i = 1$). Moreover, for the corresponding critical points $\mathbf{R}$ and $\boldsymbol{\alpha}$ one has

$$\Phi_1(\mathbf{R}) = \Psi_1(\boldsymbol{\alpha}). \tag{21}$$

Finally, the local maxima of Φ_1 and Ψ_1 happen at the critical points (that is, there are no local maxima on the boundary).

We omit the proof of Lemma 11 since it follows the proof of Theorem 4.1 in [17, section 4]. In that paper we consider random Δ-regular bipartite graphs and in analogy to $\Psi_1(\boldsymbol{\alpha})$ and $\Phi_1(\mathbf{R})$ we define $\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta})$ and $\Phi_1(\mathbf{R},\mathbf{C})$, respectively, where $\boldsymbol{\alpha},\mathbf{R}$ now correspond to the left side of the bipartition and $\boldsymbol{\beta},\mathbf{C}$ to the right side of the bipartition. The expressions in our setting (random Δ-regular graphs) are identical to those in [17] (random Δ-regular bipartite graphs) after identifying $\boldsymbol{\alpha}$ with $\boldsymbol{\beta}$ and $\mathbf{R}$ with $\mathbf{C}$. In fact, the proof of Theorem 4.1 in [17] works almost verbatim in our case after this identification.

5.2. Second moment analysis. For ferromagnetic models, Lemma 11 allows us to reexpress the optimization problem associated with the first moment in terms of matrix norms.

LEMMA 12. *Let $\mathbf{B} = \hat{\mathbf{B}}^\intercal\hat{\mathbf{B}}$ be the interaction matrix of a ferromagnetic spin system. We have*

$$\max_{\boldsymbol{\alpha}} \Psi_1(\boldsymbol{\alpha}) = \Delta \ln \|\hat{\mathbf{B}}\|_{\frac{\Delta}{\Delta-1}\to 2}.$$

Proof. Using Lemma 11 and (18) and (19), we obtain

$$\max_{\boldsymbol{\alpha}} \exp(\Psi_1(\boldsymbol{\alpha})/\Delta) = \max_{\mathbf{R}} \exp(\Phi_1(\mathbf{R})/\Delta) = \|\hat{\mathbf{B}}\|_{p\to 2}. \qquad \square$$

Recall the definition of Ψ_2 (see (4)) corresponding to the leading term of the second moment. A key fact is that Ψ_2 is given by a constrained first moment calculation on a "paired-spin" model where the interaction matrix in this model is the tensor product of the original interaction matrix with itself (see Remark 10 in section 4). The second moment considers a pair of configurations, say, σ and σ', which are constrained to have a given phase $\boldsymbol{\alpha}$. We capture this constraint using a vector $\boldsymbol{\gamma}$ corresponding to the overlap between σ and σ'; in particular, γ_{ij} is the number of vertices with spin i in σ and spin j in σ'.

Recall $\Psi_1^{\mathbf{B}}$ indicates the dependence of the function Ψ_1 on the interaction matrix $\mathbf{B}$; to simplify the notation we will drop the exponent if it is $\mathbf{B}$. More precisely,

$$\Psi_2(\boldsymbol{\alpha}) = \max_{\boldsymbol{\gamma}} \Psi_1^{\mathbf{B}\otimes\mathbf{B}}(\boldsymbol{\gamma}), \tag{22}$$

where the optimization in (22) is constrained to $\boldsymbol{\gamma}$ such that

$$\textstyle\sum_i \gamma_{ik} = \alpha_k \quad \text{and} \quad \sum_k \gamma_{ik} = \alpha_i. \tag{23}$$

Ignoring the two constraints in (23) can only increase the value of (22) and hence

$$\max_{\boldsymbol{\alpha}} \exp(\Psi_2(\boldsymbol{\alpha})/\Delta) \le \max_{\boldsymbol{\gamma}} \exp(\Psi_1^{\mathbf{B}\otimes\mathbf{B}}(\boldsymbol{\gamma})/\Delta) = \|\hat{\mathbf{B}}\otimes\hat{\mathbf{B}}\|^2_{\frac{\Delta}{\Delta-1}\to 2}. \tag{24}$$

For induced norms $\|\cdot\|_{p\to q'}$ with $p \le q'$ it is known (Proposition 10.1 in [2]) that

$$\|\hat{\mathbf{B}}\otimes\hat{\mathbf{B}}\|_{p\to q'} = \|\hat{\mathbf{B}}\|_{p\to q'}\|\hat{\mathbf{B}}\|_{p\to q'}. \tag{25}$$

Now we are ready to prove Theorem 7.

Proof of Theorem 7. Combining Lemma 12 and (24), (25) we obtain

$$\exp(\Psi_2(\boldsymbol{\alpha})/\Delta) = \max_{\boldsymbol{\gamma}} \exp(\Psi_1^{\mathbf{B}\otimes\mathbf{B}}(\boldsymbol{\gamma})/\Delta) \le \|\hat{\mathbf{B}}\|^2_{\frac{\Delta}{\Delta-1}\to 2} = \max_{\boldsymbol{\alpha}} \exp(2\Psi_1(\boldsymbol{\alpha})/\Delta).$$

This proves that if $\boldsymbol{\alpha}$ maximizes Ψ_1, we have $\Psi_2(\boldsymbol{\alpha}) \le 2\Psi_1(\boldsymbol{\alpha})$. The reverse inequality is trivial, yielding Theorem 7. $\square$

Remark 13. We will illustrate the necessity of the ferromagnetism assumption in Theorem 7 by giving an example of an antiferromagnetic model for which the second moment method fails (i.e., the second moment is larger than the square of the first moment by an exponential factor and therefore does not provide any useful concentration). Consider proper 3-colorings of random 10-regular graphs. As the size of the graph goes to infinity the probability of it being 3-colorable goes to zero. The intuitive effect of this is that to achieve a large value in the "paired-spin" model it is better to correlate the coordinates to agree. In terms of Ψ_1 and Ψ_2 we have that the maximum in the first moment is achieved for $\alpha_1 = \alpha_2 = \alpha_3 = 1/3$ with $\Psi_1 = 5\ln 2 - 4\ln 3 < 0$. To obtain a lower bound on the maximum in the second moment we take $\gamma_{11} = \gamma_{22} = \gamma_{33} = 1/3$, which yields $\Psi_2 = \Psi_1 > 2\Psi_1$. The argument actually applies whenever $\Psi_1 < 0$ (for models whose interaction matrices have 0's and 1's). By continuity (taking small B in the antiferromagnetic Potts model) one can obtain an example of a model without hard constraints for which the second moment fails.

6. Connections. In this section we prove Theorem 9, which describes the connection between the stable fixpoints of the tree recursions and the local maxima of Ψ_1. Theorem 9 will then be a key tool in our proof of Theorem 5. The technical core of the technique relies on the arguments in [17], where an analogous connection has been established for random bipartite regular graphs. The arguments here are a minor modification of this approach, suitably modified to account for random regular graphs.

Our starting point is the one-to-one correspondence between fixpoints of the tree recursions and the critical points of Ψ_1 (see [40] and also [17]). We show, roughly, that the stability of a fixpoint is equivalent to the local maximality of the corresponding critical point. This will be done by relating the Jacobian of the tree recursions at a fixpoint with the Hessian of Ψ_1 at the corresponding critical point. More precisely, we show that the Jacobian has spectral radius less than 1 (a sufficient condition for stability) iff the Hessian is negative definite (a sufficient condition for local maximality). Both constraints on the matrices are independent of the choice of local coordinates (that is, they are invariant under similarity transformations); however, to make the connection between the Jacobian and the Hessian apparent we will have to choose the local coordinates very carefully. A further technical complication is that the tree recursions are in the projective space and that the optimization of Ψ_1 is constrained.

We give a high level overview of the Jacobian; the proofs for the Δ-regular case follow the same reasoning as for the bipartite Δ-regular case (see [17, section 4.2.2]) after simply changing C_j's to R_j's and β_j's to α_j's. Assume that $R_1, \dots, R_q$ is a fixpoint of the tree recursions. Now we consider an infinitesimal perturbation of the fixpoint $R_1 + \varepsilon R'_1, \dots, R_q + \varepsilon R'_q$ and see how it is mapped by the tree recursions. Let $\alpha_i := \sum_j B_{ij} R_i R_j$. The right parametrization (choice of local coordinates) is to take $R'_i = r_i R_i/\sqrt{\alpha_i}$, where $r_1, \dots, r_q$ determines the perturbation. Note that $R_i/\sqrt{\alpha_i}$ depends on the fixpoint. The tree recursions map (in the projective space) the perturbation as follows:

(26)
$$\left(R_1 + \varepsilon r_1 \frac{R_1}{\sqrt{\alpha_1}}, \dots, R_q + \varepsilon r_q \frac{R_q}{\sqrt{\alpha_q}}\right) \mapsto \left(R_1 + \varepsilon \hat{r}_1 \frac{R_1}{\sqrt{\alpha_1}}, \dots, R_q + \varepsilon \hat{r}_1 \frac{R_q}{\sqrt{\alpha_q}}\right) + O(\varepsilon^2),$$

where $\hat{r}_i$'s are given by the linear transformation

$$\hat{r}_i = (\Delta - 1) \sum_{j=1}^{q} \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} r_j \tag{27}$$

and where the r_i's are required to satisfy

$$\sum_{i=1}^{q} \sqrt{\alpha_i} r_i = 0. \tag{28}$$

The condition (28) is invariant under the map (27) and corresponds to choosing the representative of $R_1, \dots, R_q$ with $\sum_i \sum_j B_{ij} R_i R_j = 1$.

Next we give a high level description of the Hessian; again, this is almost identical to the one in [17, section 4.2.1] after identifying C_j's with R_j's and β_j's with α_j's. Recall that Ψ_1 is a function of $\alpha_1, \dots, \alpha_q$. There is an alternative parameterization of Ψ_1: instead of $\alpha_1, \dots, \alpha_q$ (restricted to $\sum \alpha_i = 1$) we use $R_1, \dots, R_q$ (restricted to

$\sum_i \sum_j B_{ij} R_i R_j = 1$) and use the following:

$$\alpha_i = \sum_j B_{ij} R_i R_j \quad \forall\, i \in [q]. \tag{29}$$

Every $\boldsymbol{\alpha}$ can be achieved using parameterization by $\mathbf{R}$. Let $\alpha_1, \ldots, \alpha_q$ be a critical point of Ψ_1 and let $R_1, \ldots, R_q$ satisfy (29). We are going to evaluate Ψ_1 in a small neighborhood around $\alpha_1, \ldots, \alpha_q$. It is equivalent (and easier to understand) to perturb the $R_1, \ldots, R_q$ to $R_1 + \varepsilon R'_1, \ldots, R_q + \varepsilon R'_q$ and evaluate at the point given by (29). Again, the correct parameterization is to take $R'_i = r_i R_i / \sqrt{\alpha_i}$. This yields the following expression for the value of Ψ_1 at the perturbed point:

$$\begin{aligned} &\Psi_1(\alpha_1, \ldots, \alpha_q) \\ &\quad + \varepsilon^2 \sum_{i=1}^{q} \left(r_i + \sum_{j=1}^{q} \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} r_j \right) \left(\sum_{j=1}^{q} (\Delta - 1) \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} r_j - r_i \right) + O(\varepsilon^3). \end{aligned} \tag{30}$$

Note that there is no linear term, since we are at a critical point. Recall that the α_i have to satisfy $\sum_i \alpha_i = 1$ which corresponds to the restriction (28).

Now we are ready to prove Theorem 9. Let L be a linear map such that the Jacobian of the map $r \mapsto \hat{r}$ represented by (27) is $(\Delta - 1)L$. The Hessian of Ψ_1 is then $(I + L)((\Delta - 1)L - I)$. Finally let S be the linear subspace defined by (28).

Proof of Theorem 9. We will use the correspondence between fixpoints of the tree recursions and critical points of Ψ_1 given by Lemma 11. The constraint for the fixpoint to be Jacobian attractive is that $(\Delta - 1)L$ on S has spectral radius less than 1; see (26). The constraint for the critical point to be Hessian maximum is that the eigenvalues of $(I + L)((\Delta - 1)L - I)$ on S are negative; see (30).

Note that L is symmetric and if B is positive semidefinite, then L is positive semidefinite (since L is congruent to B; L is obtained by multiplying B by a diagonal matrix on the left and on the right). Hence L has nonnegative real spectrum. Note that S is invariant under L and hence the spectrum of L on S is a subset of the spectrum of L (it is still nonnegative real; the restriction wiped out the eigenvalue 1).

The constraint for the fixpoint to be Jacobian attractive, in terms of eigenvalues, is that for each eigenvalue x of L on S

$$(\Delta - 1)|x| < 1. \tag{31}$$

The constraint for the critical point to be Hessian maximum, in terms of eigenvalues, is that for each eigenvalue x of L on S

$$(1 + x)((\Delta - 1)x - 1) < 0. \tag{32}$$

Note that conditions (31) and (32) are equivalent (since $x \geq 0$). ∎

Remark 14. For antiferromagnetic models every critical point of Ψ_1 is a local maximum. Indeed, we need only to prove that (32) is satisfied for every critical point. The matrix L has nonnegative entries; hence 1 is the largest eigenvalue and all the other eigenvalues have magnitude less than 1 (since $\mathbf{B}$ is ergodic). Moreover the matrix L has the same signature as $\mathbf{B}$ (since they are congruent) and hence the eigenvalues other than 1 are negative. These two facts imply (32).

Remark 15. Note that one direction of the implication in Theorem 9, namely, that a Jacobian attractive fixpoint is Hessian local maximum, holds for every model (without the ferromagnetism assumption), since (31) always implies (32). However, for the reverse implication, the ferromagnetic assumption is essential. For example, in an antiferromagnetic model, by Remark 14, every critical point is a local maximum. For the antiferromagnetic Potts model, the only critical point is the uniform vector and hence it is always a local maximum for every value of B. On the other hand, it is straightforward to check that for the antiferromagnetic Potts model the uniform fixpoint is Jacobian unstable when $B < \frac{\Delta-q}{\Delta}$.

7. Bethe prediction for ferromagnetic models.

7.1. Small subgraph conditioning method. By Theorem 7, we have that for the random variable $Z_G^{\boldsymbol{\alpha}}$, when $\boldsymbol{\alpha}$ is a global maximizer of Ψ_1, the exponential order of its second moment is twice the exponential order of its first moment. This is not sufficient, however, to obtain high probability results, since it turns out that, in the limit $n \to \infty$, the ratio of the second moment to the square of the first moment converges to a constant greater than 1. Hence, the second moment method fails to give statements that hold with high probability over a uniform random Δ-regular graph. More specifically, to obtain our results we need sharp lower bounds on the partition function which hold for almost all Δ-regular graphs. In the setting we described, the second moment method only implies the existence of a graph which satisfies the desired bounds and even there in a not sufficiently strong form.

For random Δ-regular graph ensembles, the standard way to circumvent this failure is to use the small subgraph conditioning method of Robinson and Wormald [42]. While the method is quite technical, its application is relatively streamlined when employed in the right framework. The method was first used for the analysis of spin systems in the work of [40] for the hard-core model and subsequently in [43], [16]. In [17], we extended the approach to q-spin models for all $q \geq 2$, where the major technical obstacle was the computation of certain determinants which arise in the computation of the moments' asymptotics. While the arguments there are for random *bipartite* Δ-regular graphs, the approach extends in a straightforward manner to random Δ-regular graphs.

We defer the details of the application of the method in the present setting to section 11.1. We state here the following lemma, which is the final outcome of the method.

LEMMA 16. *For every ferromagnetic model* $\mathbf{B}$*, if* $\boldsymbol{\alpha}$ *is a Hessian dominant phase (cf. section* 3.2*) with probability* $1 - o(1)$ *over the choice of the graph* $G \sim \mathcal{G}(n, \Delta)$*, it holds that* $Z_G^{\boldsymbol{\alpha}} \geq \frac{1}{n}\mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha}}\big]$.

7.2. Proof of Theorem 8. Using Lemma 16, the proof of Theorem 8 is straightforward.

Proof of Theorem 8. Let $\boldsymbol{\alpha}$ be a Hessian dominant phase, whose existence is guaranteed by the assumptions. By Lemma 16, with probability $1 - o(1)$ over the choice of the graph, we have $Z_G^{\boldsymbol{\alpha}} \geq \frac{1}{n}\mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha}}\big]$, which implies $\frac{1}{n}\log Z_G \geq \Psi_1(\boldsymbol{\alpha}) + o(1)$.

Moreover, since the model is ferromagnetic, for Δ-regular graphs G with n vertices, $\frac{1}{n}\log Z_G \geq C$ for some constant $C > -\infty$ (explicitly, one can take $C := \frac{\Delta}{2}\log\max_{i\in[q]} B_{ii}$; see the remarks after Theorem 8). We thus obtain

$$\liminf_{n\to\infty} \frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G] \geq \liminf_{n\to\infty} \big[(1-o(1))\Psi_1(\boldsymbol{\alpha}) + o(1)C\big] = \Psi_1(\boldsymbol{\alpha}).$$

By Jensen's inequality, we also have

$$\limsup_{n\to\infty} \frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G] \le \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[Z_G].$$

All that remains to show is that $\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}[Z_G] = \Psi_1(\boldsymbol{\alpha})+o(1)$. This is straightforward; if we decompose Z_G as $Z_G = \sum_{\boldsymbol{\alpha}'} Z_G^{\boldsymbol{\alpha}'}$, we obtain $\exp(o(n))\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}] \ge \mathbf{E}_{\mathcal{G}}[Z_G] \ge \mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$. Note the factor $\exp(o(n))$, which is there to account for dominant phases which are not Hessian.

This concludes the proof. ∎

8. Phase diagram for the ferromagnetic Potts model. In this section we prove Theorem 5 detailing the phase diagram for the Potts model.

To prove Theorem 5, we will use Theorem 9 and the results of section 6. We briefly overview the approach. In order to determine the local maxima for the Potts model (or, more generally, for any ferromagnetic model), we need to compute the spectral radius of the map $L : (r_1, \dots, r_q) \mapsto (\widehat{r}_1, \dots, \widehat{r}_q)$, given by[3]

$$\widehat{r}_i = \sum_{j=1}^{q} \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} r_j \tag{33}$$

in the subspace

$$\sum_{i=1}^{q} \sqrt{\alpha_i} r_i = 0, \tag{34}$$

where the R_i's specify a fixpoint of the tree recursions (6) and the α_i's are given by

$$\alpha_i = R_i \sum_{j=1}^{q} B_{ij} R_j \text{ for } i = 1, \dots, q.$$

Our goal is to determine the local maxima by verifying when the spectral radius of this map (in the subspace (34)) is less than $1/(\Delta - 1)$. Let

$$\mathbf{M} = \left\{ \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} \right\}_{i,j=1}^{q}$$

be the matrix of the linear map L. Note that $\mathbf{M}$ is symmetric and has an eigenvalue equal to 1 with eigenvector $e = \left[\sqrt{\alpha_1}, \dots, \sqrt{\alpha_q}\right]^{\intercal}$. It follows that the eigenvalues of L in the subspace (34) are precisely the eigenvalues different from 1 of the matrix $\mathbf{M}$.

To proceed, we need to restrict our attention to the ferromagnetic Potts model. First we argue that the fixpoints of the tree recursions (6) in the case of the ferromagnetic Potts model are simple—they are supported on only two values.

LEMMA 17. *Let $(R_1, \dots, R_q)$ be a fixpoint of the tree recursions* (6) *of the ferromagnetic Potts model. Then the R_i's have at most two distinct values.*

[3]Note that, relative to (27), there is a factor of $(\Delta - 1)$ "missing" in the r.h.s. of (33). This factor will be accounted for later by demanding that the eigenvalues of the map L in (33) are less (in absolute value) than $1/(\Delta - 1)$ (instead of 1).

Proof. Without loss of generality we may assume that the implicit constant in (6) is 1. Let $r_i = R_i^{1/d}$ and $r = \sum_{i=1}^{q} r_i^d$, where $d := \Delta - 1$. We have

$$r_i = r + (B-1)r_i^d.$$

The polynomial $f(x) = (B-1)x^d - x + r$ has at most two positive roots (counted with their multiplicities, by the Descartes' rule of signs) and hence there are at most two different values of the r_i's. ∎

LEMMA 18. *The fixpoints of the tree recursions, assuming $R_1 \geq R_2 \geq \dots R_q$, satisfy $R_1 = R_2 = \cdots = R_t$ and $R_{t+1} = \dots R_q$ for some $1 \leq t \leq q$. It follows that $\alpha_1 = \alpha_2 = \cdots = \alpha_t$ and $\alpha_{t+1} = \cdots = \alpha_q$.*

Proof. This follows from Lemma 17. ∎

Remark 19. Two settings for t in the setting of Lemma 18 will be of particular interest, namely, $t = 1$ and $t = q$. We shall refer to the latter as the uniform fixpoint, and this corresponds to the disordered phase. We shall refer to fixpoints with $t = 1$ as the "majority" fixpoints. This class includes either one or two (depending on the value of B; cf. Lemma 24) fixpoints where color 1 dominates and the remaining appear with equal probability. The ordered phases correspond to the majority fixpoint for which the ratio R_1/R_q is maximum (cf. the upcoming Lemma 26).

Remark 20. We note here that fixpoints with $t \neq q$ exist iff $B \geq \mathfrak{B}_u$. In Lemma 24, we show only one side of this equivalence. In particular, we show that $B \geq \mathfrak{B}_u$ implies the existence of the majority fixpoints, and this turns out to be the only "existential" fact that is needed for the proof of Theorem 5. More precisely, in Lemmas 21 and 23, we show that fixpoints with $t \neq 1, q$ are not attractive, which in turn implies that the corresponding phases are not local maxima of Ψ_1 and, thus, are not dominant as well.

Lemma 18 implies that in the case of the ferromagnetic Potts model, $\mathbf{M}$ has a very simple structure. The following simple lemma describes the eigenvalues of $\mathbf{M}$.

LEMMA 21. *In the setting of Lemma* 18, $\mathbf{M}$ *has the following eigenvalues for* $1 \leq t < q$:

- 1 *with multiplicity* 1,
- $(B-1)R_1^2/\alpha_1$ *with multiplicity* $t-1$ *(assuming $t > 1$)*,
- $(B-1)R_q^2/\alpha_q$ *with multiplicity* $q-t-1$ *(assuming $t < q-1$), and*
- $(B+t-1)R_1^2/\alpha_1 + (B+q-t-1)R_q^2/\alpha_q - 1$ *with multiplicity* 1.

For $t = q$ the eigenvalues of $\mathbf{M}$ *are*

- 1 *with multiplicity* 1,
- $(B-1)R_1^2/\alpha_1$ *with multiplicity* $q-1$.

Proof. We already described the eigenvector for eigenvalue 1. For every i such that $2 \leq i \leq t$, a vector with 1 at position 1 and -1 at a position i (and zeros elsewhere) yields eigenvalue $(B-1)R_1^2/\alpha_1$. Similarly, for every i such that $t+1 \leq i < q$, a vector with 1 at position q and -1 at position i (and zeros elsewhere) yields eigenvalue $(B-1)R_q^2/\alpha_q$. Note that in the case $t = q$ this accounts for all the eigenvalues. In the case $t < q$ we deduce the remaining eigenvalue by considering the trace of $\mathbf{M}$:

$$t\frac{BR_1^2}{\alpha_1} + (q-t)\frac{BR_q^2}{\alpha_q} - (t-1)\frac{(B-1)R_1^2}{\alpha_1} - (q-t-1)\frac{(B-1)R_q^2}{\alpha_q} - 1.$$ ∎

Lemma 22. *The uniform fixpoint is Jacobian attractive if* $(\Delta-2)(B-1)<q$. *The uniform fixpoint is not attractive if* $(\Delta-2)(B-1)>q$.

Proof. The uniform fixpoint of the tree recursions corresponds to $R_1=\cdots=R_q$ and hence $\alpha_1=\cdots=\alpha_q=(B+q-1)R_1^2$. By Lemma 21, the only relevant eigenvalue is $(B-1)/(B+q-1)$ (with multiplicity $q-1$), which we compare with $1/(\Delta-1)$ to obtain the lemma. ∎

Lemma 22 allows us to restrict our focus on $q-1\geq t\geq 1$. In this setting, the tree recursions, with $x:=y^d:=\frac{R_1}{R_q}$ and $d:=\Delta-1$, yield

$$x=\left(\frac{(B+t-1)x+(q-t)}{tx+(B+q-t-1)}\right)^d \quad \text{or} \quad B-1=\frac{(y-1)\left(ty^d+q-t\right)}{y^d-y}. \tag{35}$$

The following lemma implies that all fixpoints with $q-1\geq t\geq 2$ are unstable in the whole nonuniqueness regime, since the respective matrices $\mathbf{M}$ have an eigenvalue greater than $1/(\Delta-1)$.

Lemma 23. *When* $\frac{R_1}{R_q}>1$, *it holds that* $(B-1)\frac{R_1^2}{\alpha_1}>\frac{1}{\Delta-1}$.

Proof. The desired inequality is equivalent to

$$(\Delta-1)(B-1)R_1>(B+t-1)R_1+(q-t)R_q,$$

which after simple manipulations reduces into

$$\big((\Delta-2)(B-1)-t\big)\frac{R_1}{R_q}>q-t.$$

Substituting $\frac{R_1}{R_q}=y^d$ and $B-1$ from (35), the inequality becomes

$$\left(\frac{(d-1)(y-1)\left(ty^d+q-t\right)}{y^d-y}-t\right)y^d-(q-t)>0.$$

Doing the necessary simplifications, we obtain the following equivalent inequality:

$$\frac{\left((d-1)y^{1+d}-dy^d+y\right)\left(q+t\left(y^d-1\right)\right)}{y^d-y}>0.$$

Since $y>1$, the only nontrivial factor to prove positivity is $p(y):=(d-1)y^{d+1}-dy^d+y$. By the Descartes' rule of signs, $p(y)$ can have at most two positive roots. It holds that $p(1)=p'(1)=0$, so that $p(y)$ is always positive for $y>1$. ∎

In light of Lemmas 21 and 23, it remains to classify fixpoints with $t=1$, i.e., the majority fixpoints. The following lemma gives the number of the majority fixpoints in the regimes of interest.

Lemma 24. *When* $\mathfrak{B}_u<B<\mathfrak{B}_{rc}$, *there are exactly two distinct majority fixpoints. When* $B\geq\mathfrak{B}_{rc}$, *there is exactly one majority fixpoint.*

Remark 25. At $B=\mathfrak{B}_u$, it follows from the proof of Lemma 24 that there is a unique majority fixpoint which "bifurcates" in the regime $\mathfrak{B}_u<B<\mathfrak{B}_{rc}$.

Proof of Lemma 24. We need to look at (35) for $t=1$ and check how many values of $y>1$ satisfy the equation in the two regimes $\mathfrak{B}_u<B<\mathfrak{B}_{rc}$ and $B\geq\mathfrak{B}_{rc}$. For $t=1$, (35) reads as

$$B-1=f(y):=\frac{(y-1)(y^d+q-1)}{y^d-y} \quad \text{so that} \quad f'(y)=\frac{p(y)}{(y^d-y)^2}, \tag{36}$$

where $p(y)$ is the polynomial

$$p(y) := y^{2d} - dy^{d+1} - (d-1)(q-2)y^d + d(q-1)y^{d-1} - (q-1). \tag{37}$$

Employing the Descartes' rule of signs we see that $p(y)$ has one or three positive roots counted by multiplicities. It is easy to check that $p(1) = p'(1) = 0$, so that p has in fact three positive roots counted by multiplicities (since 1 is a double root), and let ρ denote the other positive root. We next prove that $\rho > 1$ so that $p(y) \le 0$ if $1 \le y \le \rho$ and $p(y) \ge 0$ if $y \ge \rho$. It follows that for positive y we have $p(y) > 0$ iff $y > \rho$.

To prove that $\rho > 1$, for the sake of contradiction assume that $0 < \rho \le 1$. If $\rho = 1$, then 1 is a root with multiplicity 3 of the polynomial $p(y)$ and hence $p''(1) = 0$. By straightforward calculations we see that $p''(1) = (q-2)(d-d^2)$, which is clearly nonzero for $q \ge 3$ and $d \ge 2$. Thus, we may assume that $0 < \rho < 1$. Since $p(1) = p(\rho) = 0$, by Rolle's theorem there is a root $\rho' \in (\rho, 1)$ of the polynomial $p'(y) = dy^{d-2}g(y)$ where

$$g(y) := 2y^{d+1} - (d+1)y^2 - (d-1)(q-2)y + (d-1)(q-1).$$

Since $g(1) = g(\rho') = 0$, by the same token there is a root $\rho'' \in (\rho', 1)$ of

$$g'(y) = 2(d+1)y^d - 2(d+1)y - (d-1)(q-2).$$

We thus obtain the desired contradiction since, for $q \ge 3$ and $d \ge 2$, $g'(y) < 0$ for all $y \in [0, 1]$.

From the above, it follows that $f(\rho) = \min_{y \ge 1}\{f(y)\}$. Observe also that $f(y) \to \infty$ as $y \to \infty$, while $f(y) \to \frac{q}{d-1}$ as $y \downarrow 1$. Thus, when $y \downarrow 1$, we have that $B \uparrow \mathfrak{B}_{rc}$ (viewing B as a function of y; see (36)).

To obtain the lemma, it thus suffices to show that $\mathfrak{B}_u = f(\rho) + 1$. Recall that $\mathfrak{B}_u$ is the unique value of B for which the polynomial $(q-1)z^{d+1} + (2-B-q)z^d + Bz - 1$ has a double root in $(0, 1)$. We reparameterize $z \to 1/z$, so that $\mathfrak{B}_u$ is the unique value of B for which the following polynomial has a double root in $(1, \infty)$:

$$r(z) = z^{d+1} - Bz^d - (2-B-q)z - (q-1).$$

Let z_c be the double root of this polynomial when $B = \mathfrak{B}_u$. Solving each of $r(z_c) = 0$ and $r'(z_c) = 0$ with respect to B and equating the expressions, we obtain that $p(z_c) = 0$. It follows that for $B = \mathfrak{B}_u$ the double root of the polynomial $r(z)$ is equal to ρ. Now, solving $r(\rho) = 0$ with respect to B gives us that $\mathfrak{B}_u = f(\rho) + 1$, as wanted. $\square$

We can now classify the stability of fixpoints with $t = 1$.

LEMMA 26. *For $B > \mathfrak{B}_u$, exactly one majority fixpoint is Jacobian attractive. More precisely, the only Jacobian attractive fixpoint with $t = 1$ is the one maximizing $x = R_1/R_q$ (among the solutions of* (35) *for $t = 1$).*

We can now explicitly specify the marginal a in the definition of the ordered phase stated in the introduction. With x as in Lemma 26, we apply (20), which yields

$$a = \frac{x^{\Delta/(\Delta-1)}}{x^{\Delta/(\Delta-1)} + q - 1}. \tag{38}$$

Next, we give the proof of Lemma 26.

Proof of Lemma 26. In the setting of Lemma 18 we have $t = 1$ and thus the interesting eigenvalues of $\mathbf{M}$ are $\lambda_1 := (B-1)R_q^2/\alpha_q$ and $\lambda_2 := BR_1^2/\alpha_1 + (q-2+B)R_q^2/\alpha_q - 1$. Using that $\alpha_1 = R_1(BR_1 + (q-1)R_q)$ and $\alpha_q = R_q(R_1 + (q-2+B)R_q)$, we obtain the following equivalent expressions for λ_1, λ_2:

$$\lambda_1 = \frac{(B-1)R_q}{R_1 + (q-2+B)R_q}, \quad \lambda_2 = \frac{BR_1}{BR_1 + (q-1)R_q} - \frac{R_1}{R_1 + (q-2+B)R_q}. \tag{39}$$

We will show that

1. $\lambda_1 < \lambda_2$,
2. if R_1, R_q correspond to the fixpoint which maximizes x, it holds that $\lambda_2 < 1/(\Delta - 1)$; otherwise (i.e., if R_1, R_q correspond to the other fixpoint with $t = 1$), it holds that $\lambda_2 > 1/(\Delta - 1)$.

From these two items, the lemma follows.

For item 1, let $W := \lambda_2 - \lambda_1$. Then, expanding everything out, we obtain that

$$W = \frac{(B-1)(q-1)R_q(R_1 - R_q)}{\big(BR_1 + (q-1)R_q\big)\big(R_1 + (q-2+B)R_q\big)}$$

and thus $W > 0$ (since $B > 1$ and $R_1 > R_q$).

For item 2, let $Q := \lambda_2 - \frac{1}{\Delta-1}$. Expanding everything out, we have

$$Q = \frac{R_1R_q\big((\Delta-2)B(q-2+B) - \Delta(q-1)\big) - BR_1^2 - (q-1)(q-2+B)R_q^2}{(\Delta-1)(BR_1 + (q-1)R_q)(R_1 + (q-2+B)R_q)}.$$

Thus, to check whether $Q < 0$, it is equivalent to check, with $x = \frac{R_1}{R_q}$, whether

$$\big((\Delta-2)B(q-2+B) - \Delta(q-1)\big)x < Bx^2 + (q-1)(q-2+B). \tag{40}$$

Substituting $x = y^d$ and $B - 1$ from (35), we obtain the equivalent inequality

$$0 < \frac{y\left(y^d - 1\right)\left(y^d + q - 1\right)p(y)}{\left(y^d - y\right)^2},$$

where $p(y)$ is the polynomial defined in (37). By the proof of Lemma 24, $p(y) > 0$ iff $y > \rho$. The proof of Lemma 24 further yields that the latter inequality, throughout the regime $B > \mathfrak{B}_u$, is satisfied only by the majority fixpoint with x maximum, thus yielding item 2.

This concludes the proof. $\square$

Having classified the fixpoints which are Jacobian attractive, we now need to see when these are dominant. This entails comparing the values of Ψ_1 for the respective phases. Rather than doing this directly, we use Lemma 11. In particular, it is equivalent to compare the values of Φ_1 at the fixpoints. Moreover, note that the expression (15) is invariant upon scaling R_i's by the same factor and hence we only need to compare $\Phi_1(x, 1, \ldots, 1)$ and $\Phi_1(1, \ldots, 1)$, where x is a solution of (35) for $t = 1$.

LEMMA 27. *Let $t = 1$ and x be the solution of* (35) *with x maximum. Then $\Phi_1(x, 1, \ldots, 1) \geq \Phi_1(1, 1, \ldots, 1)$ iff $B \geq \mathfrak{B}_o$. Equality holds iff $B = \mathfrak{B}_o$.*

Proof. By a direct calculation

$$\begin{aligned}\Phi_1(x, 1, \ldots, 1) &= \frac{\Delta}{2}\ln\big((x+q-1)^2 + (B-1)(x^2+q-1)\big) \\ &\quad - (\Delta-1)\ln\big(x^{\Delta/(\Delta-1)} + q - 1\big), \\ \Phi_1(1, 1, \ldots, 1) &= \frac{\Delta}{2}\ln\big(q^2 + (B-1)q\big) - (\Delta-1)\ln(q).\end{aligned}$$

Using the substitutions $d = \Delta - 1$, $x = y^d$ and the second equation in (35), after careful manipulations we obtain

$$DIF := \Phi_1(x, 1, \ldots, 1) - \Phi_1(1, 1, \ldots, 1) = \frac{1}{2} \ln\left(\frac{q^{d-1}\left(y^d + q - 1\right)^{d+1}}{(q + y - 1)^{d+1}\left(y^{d+1} + q - 1\right)^{d-1}}\right).$$

It is straightforward to check that for $y = (q-1)^{2/(d+1)}$, $DIF = 0$. To find the respective value of B for this value of y, we just need to plug the value $y = (q-1)^{2/(d+1)}$ into the second equation in (35) for $t = 1$. In particular, from (35) we have that

$$B = 1 + \frac{(y-1)(y^d + q - 1)}{(y^d - y)} = \frac{y^{d+1} + (q-2)y - (q-1)}{y^d - y}.$$

For $y = (q-1)^{2/(d+1)}$, we have

$$y^{d+1} + (q-2)y - (q-1) = (q-2)\big(q - 1 + (q-1)^{2/(d+1)}\big)$$

and

$$\begin{aligned} y^d - y &= (q-1)^{2d/(d+1)} - (q-1)^{2/(d+1)} \\ &= \big((q-1)^{(d-1)/(d+1)} - 1\big)\big(q - 1 + (q-1)^{2/(d+1)}\big). \end{aligned}$$

It follows that

$$B = \frac{q-2}{(q-1)^{(d-1)/(d+1)} - 1} = \frac{q-2}{(q-1)^{1-2/\Delta} - 1} = \mathfrak{B}_o.$$

To prove the lemma, it thus remains to show that y is an increasing function of B and DIF increases as y increases. This is indeed true. Using (35), one calculates (see the relevant (36))

$$\frac{\partial y}{\partial B} \cdot \frac{p(y)}{(y^d - y)^2} = 1$$

and

$$\begin{aligned} \frac{\partial DIF}{\partial y} &= \frac{1}{2}\Big(\frac{d(d+1)y^{d-1}}{y^d + q - 1} - \frac{d+1}{y + q - 1} - \frac{(d-1)(d+1)y^d}{y^{d+1} + q - 1}\Big) \\ &= \frac{(d+1)(q-1)p(y)}{2(y + q - 1)(y^d + q - 1)(y^{d+1} + q - 1)}, \end{aligned}$$

where $p(y)$ is the polynomial defined in (37), whose positivity has already been established for all $y > \rho$; see the proof of Lemma 24. The claim follows. ∎

Proof of Theorem 5. We first argue about the local maxima. By Theorem 9, we just need to check the stability of the corresponding fixpoints. By Lemmas 23 and 26, only the uniform and the q majority fixpoints can be Jacobian attractive. The uniform fixpoint, by Lemma 22, is Jacobian attractive when $1 < B < \mathfrak{B}_{rc}$. The q majority fixpoints, by Lemma 26, are Jacobian attractive when $B > \mathfrak{B}_u$. This proves the assertions in Theorem 5 about the local maxima of the function Ψ_1.

To argue about the Hessian dominant phases, it only remains to find the regimes where the disordered/ordered phases are dominant. From the first part of Lemma 27, the disordered phase is dominant iff $B \le \mathfrak{B}_o$, whereas the q ordered phases are dominant iff $B \ge \mathfrak{B}_o$. From the second part of Lemma 27, all of the $q + 1$ phases are dominant at $B = \mathfrak{B}_o$. This completes the proof of Theorem 5. ∎

9. #BIS-hardness for Potts. We first give a rough description of our reduction. We will construct a gadget G which is a balanced, bipartite graph on $(2+o(1))n$ vertices. There will be $n' = O(n^{1/8})$ vertices on each side of G which will have degree $\Delta - 1$, and the remainder have degree Δ. The key is that G behaves similarly to a random bipartite Δ-regular graph. Hence, for the ferromagnetic Potts model, when $B > \mathfrak{B}_o$, the q ordered phases will dominate (ferromagnetic models on random bipartite Δ-regular graphs have the same dominant phases as on random Δ-regular graphs; see footnote 6). We will take an instance H for $\#\textsc{FerroPotts}(q, B, \Delta)$, where H has n' vertices. We then replace each vertex in H by a gadget G. Then we will use the degree $\Delta - 1$ vertices in these gadgets to encode the edges of H, while preserving bipartiteness. The resulting graph H^G will have bounded degree Δ and the Potts model on H^G will "simulate" the Potts model on H.

The gadget G is identical to the one used by Sly [43]. Before giving the detailed description of the gadget, it would be instructive to explain at an intuitive level the basic construction of the gadget and the properties we are trying to ensure. Roughly, the gadget is a random bipartite graph $\overline{G}$ with $(\Delta - 1)$-ary trees of depth $\Theta(\log n)$ attached on a small set of vertices on each side of $\overline{G}$. We use the random bipartite graph to ensure that the final gadget has q ordered phases when $B > \mathfrak{B}_o$, which will be used to encode the q spins of the Potts model in the graph H. The trees will ensure, roughly, that the relative error we introduce by approximating the partition function in H with the partition function in H^G is polynomially small. To explain this further, the reader should keep in mind that the analysis of the Potts model on the random graph $\overline{G}$ will only give estimates of its partition function within a multiplicative factor $(1 \pm \varepsilon)$ for some small constant $\varepsilon > 0$. In turn, using this bound to analyze the Potts model on H^G would result in estimating the partition function of H within a multiplicative factor $(1 \pm \varepsilon)^{|H|}$, which is way bigger than the polynomial accuracy we seek in approximation-preserving reductions. This obstacle is precisely the reason why the trees are attached to $\overline{G}$: the trees will boost the constant ε to a much smaller quantity of order $n^{-O(1)}$; then, the final approximation can be made polynomially small as desired.[4] As in [43, 16, 7], this boosting is possible due to the fact that the ordered phases correspond to nonreconstructible Gibbs measures on the infinite Δ-regular tree (we will expand on this later in the proof of Lemma 28, item 2).

Next, we give the description of the gadget G. The gadget G is defined by two parameters θ, ψ, where $0 < \theta, \psi < 1/8$. The construction of the gadget G has two parts. First construct the following bipartite graph $\overline{G}$ with vertex set $V^+ \cup V^-$. For $s \in \{+, -\}$, $|V^s| = n + m'$, where m' will be defined precisely later. Take Δ random perfect matchings between V^+ and V^-. Then remove a matching of size m' from one of the Δ matchings. Call this graph $\overline{G}$. For later use, let $U := U^+ \cup U^-$ denote the vertices of degree Δ in $\overline{G}$ and $W := W^+ \cup W^-$ denote the vertices of degree $\Delta - 1$ in $\overline{G}$.

[4] Let us remark that sometimes even the $(1 \pm \varepsilon)$ factor estimates are sufficient to get strong inapproximability results and the attachment of the trees is not needed; see, for example, the reductions in [44, 17]. The difference in those settings is that the corresponding counting problems, i.e., counting independent sets or counting colorings, can be connected to NP-hard problems (such as $\textsc{Max-Cut}$), yielding that it is NP-hard to approximate the corresponding partition function even within an exponential factor. In contrast, for problems like #BIS or $\#\textsc{FerroPotts}$ (which correspond to easy decision problems, e.g., finding the max independent set on a bipartite graph or the maximum weight configuration in the ferromagnetic Potts model), such strong inapproximability results are not known; in fact, we only have evidence that an FPRAS is unlikely to exist. This necessitates the study of approximation-preserving reductions in our setting (as in, e.g., [14, 7, 33]) and thus our quest for the polynomial precision.

In the second stage, for each side of $\overline{G}$, partition the degree $\Delta - 1$ vertices into n^θ equal sized sets and attach to each set a $(\Delta-1)$-ary tree of depth ℓ, where $\ell = \lfloor \psi \log_{\Delta-1} n \rfloor$. (Use the vertices of $\overline{G}$ as the leaves of these trees.) Hence each side contains n^θ trees of size $O(n^\psi)$. (More precisely, $(\Delta-1)^{\lfloor \theta \log_{\Delta-1} n \rfloor}$ trees, each having $(\Delta-1)^{\lfloor \psi \log_{\Delta-1} n \rfloor}$ leaves.) This defines the gadget G. For $s \in \{+,-\}$, let R^s denote the roots of the trees on side s and $R := R^+ \cup R^-$. Note that each vertex in R has degree $\Delta - 1$ and these will be used to encode the edges of H (we give the reduction explicitly just after stating the relevant gadget lemma, Lemma 28). Note that $m' = (\Delta-1)^{\lfloor \theta \log_{\Delta-1} n \rfloor + \lfloor \psi \log_{\Delta-1} n \rfloor}$ and $m' = o(n^{1/4})$.

Denote by $G = (V, E)$ the final graph. Recall, for a configuration $\sigma \in \Omega$, that the set of vertices assigned spin i is denoted by $\sigma^{-1}(i)$. The phase of a configuration $\sigma : V \to [q]$ is defined as the dominant spin among vertices in $U = U^+ \cup U^-$ (the vertices of degree Δ in $\overline{G}$):

$$Y(\sigma) := \arg\max_{i \in [q]} |\sigma^{-1}(i) \cap U|,$$

where ties are broken with an arbitrary deterministic criterion (e.g., the lowest index).

The gadget G behaves like a random bipartite Δ-regular graph because $m' \ll n$, as we will detail in the upcoming Lemma 28. Hence, since $B > \mathfrak{B}_o$, Theorem 5 implies that the q ordered phases are dominant. Therefore, we will get that for a sample σ from the Gibbs distribution, the phase of σ will be (close to) uniformly distributed over these q ordered phases. Let phase i refer to the ordered phase where spin i is the majority. Once we condition on the phase for the vertices in U, say, it is phase i, then each of the roots of the trees appended to G, roughly independently, will have spin i with probability $\approx p$ and spin $j \neq i$ with probability $\approx (1-p)/(q-1)$, where p is the probability that the root of the infinite $(\Delta-1)$-ary tree has spin i in the Gibbs measure corresponding to the ordered phase i.[5] Hence, for each of the q possible phases, we define the following product distribution on the configurations $\sigma_R : R \to [q]$. For $i \in [q]$, let

$$Q_R^i(\sigma_R) = p^{|\sigma_R^{-1}(i)|} \left(\frac{1-p}{q-1} \right)^{|R \setminus \sigma_R^{-1}(i)|}. \tag{41}$$

For future use, one can define completely analogously the product measure $Q_W^i(\cdot)$ on configurations $\sigma_W : W \to [q]$ (recall that W is the set of vertices with degree $\Delta - 1$ in $\overline{G}$).

The following lemma is proved using methods in [43] and its proof is given in section 9.1. Roughly, the first item in the lemma follows from the symmetries of the Potts model. For the second item, the rough idea is that when the phase is i, the marginal spin distribution of vertices in W in the graph $\overline{G}$ is close to Q_W^i. The purpose of the trees is to boost this effect, more precisely, to make the distance between the marginal spin distribution of vertices in R and Q_R^i an inverse polynomial factor (see item 2 in Lemma 28). In turn, the reason that the trees can accomodate the "boosting" is that the marginal distribution on W corresponds to an extremal Gibbs measure on the tree, which results in the spins of the roots of the trees being strongly concentrated.

[5] The ordered phase $\boldsymbol{\alpha} = (a, (1-a)/(q-1), \ldots, (1-a)/(q-1))$ specifies the marginal probabilities for the root of the infinite Δ-regular tree. To account for the root having degree $\Delta - 1$ one obtains that $p = \frac{a^{(\Delta-1)/\Delta}}{(a/(1-a))^{(\Delta-1)/\Delta} + (q-1)^{1/\Delta}}$. Alternatively, $p = x/(x + q - 1)$, for the same x as in (38).

LEMMA 28. *For every $q, \Delta \geq 3$ and $B > \mathfrak{B}_o$, there exist constants $\theta, \psi > 0$ such that the graph G satisfies the following with probability $1 - o(1)$ over the choice of the graph:*

1. *The phases occur with roughly equal probability, so that for every phase $i \in [q]$, we have*
$$\Big|\mu_G\big(Y(\sigma) = i\big) - \frac{1}{q}\Big| \leq n^{-2\theta}.$$
2. *Conditioned on the phase i, the spins of vertices in R are approximately independent, that is,*
$$\max_{\sigma_R}\Big|\frac{\mu_G\big(\sigma_R \,|\, Y = i\big)}{Q_R^i(\sigma_R)} - 1\Big| \leq n^{-2\theta}.$$

With Lemma 28 at hand, we can now formally state the reduction that we sketched earlier. Let $B > \mathfrak{B}_o$. Let H be a graph on n' vertices, where $n' \leq n^{\theta/4}$ and θ is as in Lemma 28. Assuming an FPRAS for the ferromagnetic Potts model on max-degree Δ graphs and parameter B, we will show that we can approximate $Z_H(B^*)$, the partition function of H in the ferromagnetic Potts model with parameter B^*, where B^* will be determined shortly.

To do this, we first construct a graph H^G. First, take $|H|$ disconnected copies of the gadget G in Lemma 28 and identify each copy with a vertex $v \in H$. Denote by $\hat{H}^G$ the resulting graph, by G_v the copy of the gadget associated to the vertex v in H, and by R_v^+, R_v^-, R_v the images of R^+, R^-, R in the gadget G_v, respectively. We next add the edges of H in $\hat{H}^G$. To do this, fix an arbitrary orientation of the edges of H. For each oriented edge (u, v) of H, we add an edge between one vertex in R_u^+ and one vertex in R_v^-, using mutually distinct vertices for distinct edges of H. The resulting graph will be denoted by H^G. Note that H^G is bipartite and has maximum degree Δ.

For a graph H and activity $B \geq 1$, recall that $Z_H(B)$ is the partition function for the ferromagnetic Potts model at activity B on the graph H. We have the following connection.

LEMMA 29. *Let $\Delta, q \geq 3$ and $B > \mathfrak{B}_o$. There exists $B^* > 1$ (depending only on q, Δ, B) such that the following holds for every graph H with n' vertices:*

$$\big(1 - O(n^{-\theta})\big)\frac{q^{n'} Z_{H^G}(B)}{C_H\big(Z_G(B)\big)^{n'}} \leq Z_H(B^*) \leq \big(1 + O(n^{-\theta})\big)\frac{q^{n'} Z_{H^G}(B)}{C_H\big(Z_G(B)\big)^{n'}},$$

where $C_H = D^{|E(H)|}$ and $D = 1 + (B-1)\big(\frac{2p(1-p)}{(q-1)^2} + (q-2)\frac{(1-p)^2}{(q-1)^2}\big)$.

Using Lemma 29 we can now prove that for all $\Delta \geq 3$, all $B > \mathfrak{B}_o$, it is #BIS-hard to approximate the partition function of the ferromagnetic Potts model on bipartite graphs of maximum degree Δ.

Proof of Theorem 3. Goldberg and Jerrum [21] showed that for every $B > 1$ it is #BIS-hard to approximate the partition function of the ferromagnetic Potts model on all graphs. Fix $\Delta, q \geq 3$ and $B > \mathfrak{B}_o$ for which we intend to prove Theorem 3, and let $B^* = B^*(q, \Delta, B) > 1$ be specified as in Lemma 29. We first show that an FPRAS for approximating the partition function with activity B on graphs with maximum degree Δ implies an FPRAS for approximating the partition function with activity B^* on all graphs. It will then be clear that our reduction is in fact approximation-preserving and hence the theorem will be proven.

Suppose that there exists an FPRAS for approximating the partition function with activity B on bipartite graphs with maximum degree Δ. Take an input instance H for which we would like to estimate the partition function of the Potts model at activity B^*. First generate a random gadget G using the construction defined earlier. This graph G satisfies the properties in Lemma 28 with probability $1-o(1)$. Approximate the partition function of G at activity B within a multiplicative factor $1 \pm \varepsilon/10n'$ using our presumed FPRAS (where, recall, n' is the number of vertices in H). Also, using the presumed FPRAS, approximate the partition function of H^G at activity B within a multiplicative factor $1 \pm \varepsilon/5$. The bounds for $Z_H(B^*)$ in Lemma 29 are then within a factor $1 \pm \varepsilon$ for sufficiently large n, implying an FPRAS for approximating the partition function at activity B^*. This, together with the result of [21], implies an FPRAS for counting independent sets in bipartite graphs. ∎

Proof of Lemma 29. Recall that $\hat{H}^G$ are the disconnected copies of the gadgets, as defined in the construction of H^G. Note that $Z_{\hat{H}^G}(B) = \big(Z_G(B)\big)^{n'}$. Hence to prove the lemma it suffices to analyze $\frac{Z_{H^G}(B)}{Z_{\hat{H}^G}(B)}$.

For a configuration σ on H^G, for each $v \in H$, let $Y_v(\sigma)$ denote the phase of σ on G_v. Denote the vector of these phases by $\mathcal{Y}(\sigma) = (Y_v(\sigma))_{v\in H} \in [q]^H$; we refer to $\mathcal{Y}(\sigma)$ as the phase vector for σ.

For $\mathcal{U} \in [q]^H$, let $\Omega_{\mathcal{U}}$ denote the set of configurations σ on H^G, where $\mathcal{Y}(\sigma) = \mathcal{U}$. Let $Z_{H^G}(\mathcal{U})$ be the partition function of H^G restricted to configurations $\sigma \in \Omega_{\mathcal{U}}$, that is,

$$Z_{H^G}(\mathcal{U}) = \sum_{\sigma\in\Omega_{\mathcal{U}}} B^{m(\sigma)},$$

where for a configuration σ, $m(\sigma)$ is the number of monochromatic edges under σ. We may view $\mathcal{U}$ as an assignment $V(H) \to [q]$, where $V(H)$ are the vertices in the graph H. Hence, we can consider the number of monochromatic edges in the graph H under the assignment $\mathcal{U}$, which we denote by $m(\mathcal{U})$. Recall the goal is to analyze $\frac{Z_{H^G}(B)}{Z_{\hat{H}^G}(B)}$. To this end we will analyze $\frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})}$ for every $\mathcal{U}$ and then we will use that every $\mathcal{U}$ is (close to) equally likely in $\hat{H}^G$ which will follow from property 1 in Lemma 28.

Denote by R_H the set of vertices $\cup_v R_v$, i.e., the union of all the vertices of degree $\Delta - 1$ in $\hat{H}^G$. Notice that once we fix an assignment to all of the vertices in R_H, by the Markov property of the model, we have that

$$\frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})} = \sum_{\sigma_{R_H}} \mu_{\hat{H}^G}(\sigma_{R_H} \,|\, \mathcal{Y}(\sigma) = \mathcal{U}) \prod_{(u,v)\in E(H^G)\setminus E(\hat{H}^G)} B^{\mathbf{1}\{\sigma_{R_H}(u)=\sigma_{R_H}(v)\}}.$$

Note that $\mu_{\hat{H}^G}(\sigma_{R_H} \,|\, \mathcal{Y}(\sigma) = \mathcal{U}) = \big(1 + O(n^{-\theta})\big) \prod_{v\in V(H)} Q^{\mathcal{U}_v}_{R_v}(\sigma_{R_v})$ since $\hat{H}^G$ is a union of disconnected copies of G and in each copy of G we have property 2 of Lemma 28. It follows that

$$\begin{aligned}\frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})} &= \big(1 + O(n^{-\theta})\big) \sum_{\sigma_{R_H}} \prod_{v\in V(H)} Q^{\mathcal{U}_v}_{R_v}(\sigma_{R_v}) \prod_{(u,v)\in E(H^G)\setminus E(\hat{H}^G)} B^{\mathbf{1}\{\sigma_{R_H}(u)=\sigma_{R_H}(v)\}} \\ &= \big(1 + O(n^{-\theta})\big) A^{m(\mathcal{U})} D^{|E(H)|-m(\mathcal{U})},\end{aligned}$$

where A (resp., D) is the expected weight of an edge connecting two gadgets which have the same (resp., different) phases. Simple calculations show that

$$A = 1 + (B-1)\Big(p^2 + \frac{(1-p)^2}{q-1}\Big), \quad D = 1 + (B-1)\Big(\frac{2p(1-p)}{q-1} + (q-2)\frac{(1-p)^2}{(q-1)^2}\Big).$$

Letting $B^* = A/D$ and $C_H = D^{|E(H)|}$, we obtain

$$\frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})} = \big(1 + O(n^{-\theta})\big)(B^*)^{m(\mathcal{U})} C_H. \tag{42}$$

Property 1 in Lemma 28 gives that for every $\mathcal{U}$ it holds that

$$\big(1 - O(n^{-\theta})\big)q^{-n'} \le \Big(\frac{1}{q} - n^{-2\theta}\Big)^{n'} \le \frac{Z_{\hat{H}^G}(\mathcal{U})}{Z_{\hat{H}^G}} \le \Big(\frac{1}{q} + n^{-2\theta}\Big)^{n'} \le \big(1 + O(n^{-\theta})\big)q^{-n'}. \tag{43}$$

We also have

$$Z_{H^G}(B) = \sum_{\mathcal{U}} Z_{H^G}(\mathcal{U}) = \sum_{\mathcal{U}} \frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})} Z_{\hat{H}^G}(\mathcal{U}) = Z_{\hat{H}^G} \sum_{\mathcal{U}} \frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})} \frac{Z_{\hat{H}^G}(\mathcal{U})}{Z_{\hat{H}^G}}. \tag{44}$$

Using the estimates (42), (43), in (44), we obtain

$$\big(1 - O(n^{-\theta})\big)q^{-n'} C_H Z_H(B^*) \le \frac{Z_{H^G}(B)}{Z_{\hat{H}^G}(B)} \le \big(1 + O(n^{-\theta})\big)q^{-n'} C_H Z_H(B^*).$$

The result follows after observing that $Z_{\hat{H}^G}(B) = \big(Z_G(B)\big)^{n'}$ and rearranging the inequality. $\square$

9.1. Proving the properties of the gadget. In this section, we prove the properties of the gadget we use, as stated in Lemma 28. We outline the proof and introduce the relevant notation. The proof follows the same approach as in [43, Theorem 2.1] and uses nonreconstruction results in [35]. We argue, however, more thoroughly for item 1 in Lemma 28, since in [43] a cruder bound for the probability that a phase appears was sufficient. In our case, the more delicate bound will follow from the symmetries of the Potts model. We first illustrate how symmetry comes into play.

Let Σ^i_G be the set of configurations on G which have phase i, i.e., $\Sigma^i_G := \{\sigma : V \to [q] \,|\, Y(\sigma) = i\}$. Moreover, let Σ^o_G be the set of configurations σ which satisfy $\big|\arg\max_{i\in[q]} |\sigma^{-1}(i) \cap U|\big| \ge 2$, that is, Σ^o_G consists of these configurations whose phase was determined by breaking a tie. We first show that item 1 in Lemma 28 will follow from showing that Σ^o_G has exponentially smaller contribution to the partition function of G than Σ^i_G for every $i \in [q]$.

To capture this, for a subset $\Sigma \subseteq \Omega_G$ of the configuration space, denote by $Z_G(\Sigma)$ the partition function restricted to configurations in Σ, that is,

$$Z_G(\Sigma) = \sum_{\sigma\in\Sigma} w_G(\sigma).$$

Let π be a permutation of the colors $[q]$ which maps color i to color j. For a configuration σ, we denote by $\pi(\sigma)$ the configuration $\pi \circ \sigma$. Clearly, for every configuration $\sigma \in \Sigma_G^i \backslash \Sigma_G^o$ we have $\pi(\sigma) \in \Sigma_G^j \backslash \Sigma_G^o$. It follows that for every two colors i, j we have $Z_G(\Sigma_G^i \backslash \Sigma_G^o) = Z_G(\Sigma_G^j \backslash \Sigma_G^o)$. Since

$$Z_G = Z_G(\Sigma_G^o) + \sum_i Z_G(\Sigma_G^i \backslash \Sigma_G^o),$$

to get the inequality in item 1 of Lemma 28 it suffices to show that $Z_G(\Sigma_G^o)$ is smaller than Z_G by a sufficiently large polynomial factor.

We briefly outline the argument for proving that $Z_G(\Sigma_G^o)$ is smaller than Z_G by a sufficiently large polynomial factor, introducing at the same time some relevant notation. First, note that the definition of the phase of a configuration makes sense for configurations on $\overline{G}$ as well. For convenience, we will henceforth use Z_G^o, Z_G^i as shorthand for $Z_G(\Sigma_G^o), Z_G(\Sigma_G^i)$ and $Z_{\overline{G}}^o, Z_{\overline{G}}^i$ for their analogues in $\overline{G}$. Roughly, we will first show that $Z_{\overline{G}}^o$ is exponentially smaller than $Z_{\overline{G}}^i$ with probability $1 - o(1)$ (over the choice of $\overline{G}$). This part follows from the fact that the q ordered phases are dominant and the fact that $Z_{\overline{G}}^i$ matches its expectation (up to a polynomial factor) with probability $1-o(1)$ (this is the analogue of Lemma 16 for the graph distribution induced by $\overline{G}$). We will then show that Z_G^o is smaller than Z_G^i by a factor of $\exp(n^{1/4})$ by crudely accounting for the contribution of the trees attached in the second step of the construction of G. Summing over $i \in [q]$ then yields the desired bound and thus completes the "symmetry" argument.

To formalize the outline in the previous paragraph, we will have to capture how the partition functions $Z_{\overline{G}}$ and Z_G interplay. Due to the Markov property, this happens only through vertices in W (recall this is the set of vertices of degree $\Delta-1$ in the graph $\overline{G}$ on which the trees are attached). Thus, we will partition the sets $\Sigma_{\overline{G}}^o, \Sigma_{\overline{G}}^i$ according to the configuration η on W. In particular, $\Sigma_{\overline{G}}^o(\eta)$ will be those configurations σ in $\Sigma_{\overline{G}}^o$ such that $\sigma_W = \eta$ and $Z_{\overline{G}}^o(\eta)$ will be the contribution to the partition function of $\overline{G}$ from configurations in $\Sigma_{\overline{G}}^o(\eta)$. Define similarly $\Sigma_{\overline{G}}^i(\eta)$ and $Z_{\overline{G}}(\eta)$.

We need a final piece of notation. Let J be the union of the trees appended in the second step of the construction of the gadget G. Note that the only vertices of $\overline{G}$ included in J are vertices in W. Let $Z_J(\eta)$ be the contribution to the partition function of J from configurations σ (on J) such that $\sigma_W = \eta$. We are now able to put these definitions to work. In particular, we have that

$$Z_G^i = \sum_{\eta: W \to [q]} Z_{\overline{G}}^i(\eta) Z_J(\eta) \text{ and } Z_G^o = \sum_{\eta: W \to [q]} Z_{\overline{G}}^o(\eta) Z_J(\eta).$$

We will need the following lemma, which is proved combining techniques from [43], [17, Appendices A and B], and the phase diagram for the Potts model (note that for ferromagnetic models, the dominant phases have the same quantitative structure in bipartite graphs[6]).

LEMMA 30. *Let $\mathcal{G} := \mathcal{G}_n$ denote the distribution of the random bipartite graph $\overline{G}$. For $B > \mathfrak{B}_o$, it holds that*

[6] Dominant phases on random bipartite Δ-regular graphs correspond to the global maximizers of $\max_{\mathbf{R},\mathbf{C}} \frac{\mathbf{R}^{\intercal}\mathbf{B}\mathbf{C}}{\|\mathbf{R}\|_p \|\mathbf{C}\|_p}$, where $p = \Delta/(\Delta-1)$; see [17, Theorem 4.1]. As a consequence of Observation 2, for a ferromagnetic model, any such maximum must satisfy $\mathbf{R} = \mathbf{C}$ (up to a scaling factor). This yields that the maximizers are in one-to-one correspondence with the maximizers of the r.h.s. of (16).

(i) *There exist constants* C_1, C_2, Ψ *depending only on* q, B, Δ*, such that for every* $i \in [q]$ *and* $\eta : W \to [q]$,

$$\begin{aligned} \mathbf{E}_{\mathcal{G}}\left[Z_{\overline{G}}^{i}\right] &= \left(1+O(n^{-1/2})\right)C_1 C_2^{2m'}\exp(n\Psi), \\ \mathbf{E}_{\mathcal{G}}[Z_{\overline{G}}^{i}(\eta)] &= \left(1+O(n^{-1/2})\right)Q_W^i(\eta)\mathbf{E}_{\mathcal{G}}\left[Z_{\overline{G}}^{i}\right]. \end{aligned} \tag{45}$$

(ii) *For all sufficiently small* $\varepsilon > 0$ *and sufficiently large* n*, for* $i \in [q]$,

$$\mathbf{E}_{\mathcal{G}}\left[Z_{\overline{G}}^{o}\right] \le \exp(-\varepsilon n)\mathbf{E}_{\mathcal{G}}\left[Z_{\overline{G}}^{i}\right]. \tag{46}$$

(iii) $\max_{i\in[q],\, \eta\colon W\to[q]} \Pr_{\mathcal{G}}\left(Z_{\overline{G}}^{i}(\eta) < \frac{1}{n}\mathbf{E}_{\mathcal{G}}[Z_{\overline{G}}^{i}(\eta)]\right) \to 0$ *as* $n \to \infty$.

Proof of Lemma 30. The first equality in (45) is proved in [17, Lemma B.3]. The second equality in (45) is proved in [17, Lemma 6.11]: the lemma is stated in the case where m' is a constant, but the proof also holds when $m' = o(n^{1/2})$ whenever the dominant phases are Hessian, as was first illustrated in [43]. The explicit error factors $O(n^{-1/2})$ in (45) are a consequence of Stirling's approximation (see [43, Lemma 3.1] for the explicit derivation in the hard-core model which is straightforward to adapt to the present setting as well). Combining the above yields item (i).

Item (ii) is a consequence of the fact that for some small $\varepsilon' > 0$, for each ordered phase i there exists an ε'-ball around it consisting solely of configurations which are contained in $\Sigma_{\overline{G}}^{i}\backslash\Sigma_{\overline{G}}^{o}$. Since the ordered phase i is dominant, a standard compactness argument (see, for example, the upcoming proof of Lemma 36) yields that $\frac{1}{n}\log\mathbf{E}_{\mathcal{G}}\left[Z_{\overline{G}}^{o}\right]$ is strictly less than $\frac{1}{n}\log\mathbf{E}_{\mathcal{G}}\left[Z_{\overline{G}}^{i}\right]$.

Finally, item (iii) follows from the small subgraph conditioning method in [17, Appendix A] (see also Theorem 37). ∎

We conclude this section by giving the proof of Lemma 28.

Proof of Lemma 28. To get item 1, by the symmetry argument described in the beginning of the section, it suffices to show that for every $i \in [q]$ it holds that $Z_G^o \le \exp(-n^{1/4})Z_G^i$ with probability $1 - o(1)$ over the choice of the graph G. We use Lemma 30. In particular, Markov's inequality yields

$$\Pr_{\mathcal{G}}\left(\sum_{\eta:W\to[q]} Z_{\overline{G}}^{o}(\eta)Z_J(\eta) > n\sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z_{\overline{G}}^{o}(\eta)]\right) \to 0 \text{ as } n\to\infty. \tag{47}$$

Item (iii) of Lemma 30 yields for every $i \in [q]$

$$\Pr_{\mathcal{G}}\left(\sum_{\eta:W\to[q]} Z_{\overline{G}}^{i}(\eta)Z_J(\eta) < \frac{1}{2n}\sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z_{\overline{G}}^{i}(\eta)]\right) \to 0 \text{ as } n\to\infty. \tag{48}$$

From (45) and $\sum_\eta Q_W^i(\eta) = 1$, it follows that

$$\begin{aligned} \sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z_{\overline{G}}^{i}(\eta)] &= (1+o(1))\mathbf{E}_{\mathcal{G}}[Z_{\overline{G}}^{i}]\sum_{\eta:W\to[q]} Z_J(\eta)Q_W^i(\eta) \\ &\ge (1+o(1))\mathbf{E}_{\mathcal{G}}[Z_{\overline{G}}^{i}]\min_{\eta:W\to[q]} Z_J(\eta). \end{aligned} \tag{49}$$

From the crude bound $\max_\eta Z_J(\eta) \le \exp(o(n^{1/4}))\min_\eta Z_J(\eta)$ and (46), it follows that

$$\begin{aligned} \mathbf{E}_{\mathcal{G}}[Z_{\overline{G}}^{i}]\min_{\eta:W\to[q]} Z_J(\eta) &\ge \exp(n^{1/2})\mathbf{E}_{\mathcal{G}}[Z_{\overline{G}}^{o}]\max_{\eta:W\to[q]} Z_J(\eta) \\ &\ge \exp(n^{1/2})\sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z_{\overline{G}}^{o}(\eta)]. \end{aligned} \tag{50}$$

Combining (49) and (50) yields

$$\sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z^i_{\overline{G}}(\eta)] \geq \exp(n^{1/2}) \sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z^o_{\overline{G}}(\eta)], \tag{51}$$

Combining (47), (48), and (51) yields that $Z^o_G \leq n\exp(-n^{1/2})Z^i_G \leq \exp(-n^{1/4})Z^i_G$ with probability $1-o(1)$ over the choice of the graph G, as wanted. This proves the first item of the lemma.

Item 2 of the lemma follows exactly the approach in [43]. The required non-reconstruction results to push the approach in [43] are given in [35, Proof of Theorem 1.4] (ferromagnetic Potts model on the tree with constant boundary condition). Together with Lemma 30, the proof of [43, Theorem 2.1] extends almost verbatim to our case as well. We briefly outline the main ideas of the proof as carried out in [43].

Recall that the goal is to show that, conditioned on the phase i, the distribution of the spins in vertices in R is close to $Q^i_R(\cdot)$. For $i \in [q]$, let (analogously to [43])

$$\mathcal{B}_i := \Big\{\eta : W \to [q] \mid \max_{\tau:R\to[q]} \Big|\mu_G\big(\sigma_R = \tau \mid \sigma_W = \eta\big) - Q^i_R(\sigma_R = \tau)\Big| > n^{-3\theta}\Big\},$$

i.e., $\mathcal{B}_i$ is the set of "bad" configurations on W which exert large influence on vertices in R. Note that while we defined $\mathcal{B}_i$ using the Gibbs distribution of the graph G, we could have used instead the Gibbs distribution of J, since, by the Markov property, conditioned on the spins of vertices in W, the spins of the vertices in J are conditionally independent from the rest of the vertices in the graph G. It follows that

$$\mathcal{B}_i = \Big\{\eta : W \to [q] \mid \max_{\tau:R\to[q]} \big|\mu_J\big(\sigma_R = \tau \mid \sigma_W = \eta\big) - Q^i_R(\sigma_R = \tau)\big| > n^{-3\theta}\Big\}. \tag{52}$$

Back to the proof, the result will follow from $\mu_G(\sigma_W \in \mathcal{B}_i \mid Y(\sigma) = i) \leq \exp(-n^{2\theta})$; for the technical details see [43, Proof of Theorem 2.1].

Note that

$$\mu_G\big(\sigma_W = \eta \,|\, Y(\sigma) = i\big) = \frac{Z^i_{\overline{G}}(\eta)Z_J(\eta)}{Z^i_G} = \frac{Z^i_{\overline{G}}(\eta)Z_J(\eta)}{\sum_{\eta':W\to[q]} Z^i_{\overline{G}}(\eta')Z_J(\eta')}. \tag{53}$$

Using inequalities analogous to those we used to prove item 1, it can be proved that

$$\mu_G\big(\sigma_W \in \mathcal{B}_i \,|\, Y(\sigma) = i\big) \leq poly(n)\nu^i(\sigma_W \in \mathcal{B}_i),$$

where the measure ν^i is defined on the space of all configurations $\eta : W \to [q]$ given by

$$\nu^i(\eta) := \frac{Z_J(\eta)Q^i_W(\eta)}{\sum_{\eta':W\to[q]} Z_J(\eta')Q^i_W(\eta')} \propto \mu_J(\sigma_W = \eta)Q^i_W(\eta). \tag{54}$$

Our goal is thus to show that

$$\nu^i(\sigma_W \in \mathcal{B}_i) \leq \exp(-n^{2\theta}). \tag{55}$$

It is useful to note at this point that the bound in (55) is a property of the trees and, in particular, does not depend on the Gibbs distribution of the (random) graph $\overline{G}$. (Indeed, ν^i is specified by the Gibbs distribution on the graph J, which is a disjoint union of $(\Delta-1)$-ary trees, and the product measure $Q^i(\eta)$. Also, $\mathcal{B}_i$ is specified by the Gibbs distribution of the graph J; see (52).) Also, since J is the disjoint union of a polynomial number of identical trees, by a union bound, it suffices to show (55) when J consists of a single $(\Delta-1)$-ary tree with height $\ell = \Theta(\log n)$ and W denoting the leaves of the tree. In turn, this will follow from the following doubly exponential upper bound:

$$\nu^i(\sigma_W \in \mathcal{B}_i) \leq \exp(-\exp(C\ell)), \tag{56}$$

where $C > 0$ is a constant. (To recover (55) from (56), we just need to tune the parameter θ of the gadget to ensure that the trees have sufficiently large height ℓ relative to θ; recall that $\ell = \lfloor \psi \log_{\Delta-1} n \rfloor$, so we can choose any constant θ so that $0 < \theta < \frac{C\psi}{2\ln(\Delta-1)}$.)

We will conclude this proof by sketching the main idea behind the strong bound in (56); a detailed proof of the bound with all the relevant connections can be found in Appendix A. The bound in (56) goes back to the works of Martinelli, Sinclair, and Weitz [34, 35], who studied the mixing time of Glauber dynamics on trees with boundary conditions, and was first used for the construction of gadgets by Sly [43]. A key idea, captured in [43, Proof of Lemma 4.3], is that the measures ν^i on configurations $\eta : W \to [q]$ can be viewed as projections of Gibbs measures corresponding to the ordered phases on the infinite $(\Delta-1)$-ary tree. The Gibbs measure corresponding to the ordered phase i can be obtained by taking the weak limit of the Potts distribution of a finite tree with depth ℓ whose leaves are conditioned to have spin i as $\ell \to \infty$. In Appendix A.2, we give an alternative Markov chain construction of these Gibbs measures using a broadcasting process which is more convenient to work with. These Gibbs measures are well-known to be extremal, or, equivalently, the broadcasting process has the nonreconstruction property, which roughly says that the spin of the root cannot be reconstructed from a typical configuration on the leaves of the tree (asymptotically in ℓ); see Appendix A.1 for details that are relevant in our setting and the survey [38] for more details on broadcasting processes on trees. Then, the techniques of [34, 35] further show that a certain eigenvalue condition of the relevant broadcasting matrix allows us to quantify the dependence on ℓ and thus obtain the bound in (56). In Appendix A, we use a similarly flavored result from Sly and Zhang [46], which we can apply more directly in our setting.

This concludes the proof of Lemma 28. ∎

9.2. #BIS-hardness for bipartite colorings. Using our #BIS-hardness result for the ferromagnetic Potts model on bounded-degree graphs, we now prove our #BIS-hardness result for colorings on bounded-degree bipartite graphs (Corollary 4). The reduction between these two problems was first observed in [14]; here, we just have to work out the bound on k that the application of Theorem 3 yields.

Proof of Corollary 4. We will show that for all integer $k, \Delta \geq 3$, it holds that

$$\#\textsc{BipFerroPotts}\left(q = k, B = \frac{k-1}{k-2}, \Delta\right) \leq_{\mathrm{AP}} \#\textsc{BipColorings}(k, \Delta) \tag{57}$$

and that

$$\text{whenever } k \leq \Delta/(2\ln\Delta), \text{ it holds that } B = \tfrac{k-1}{k-2} > \tfrac{k-2}{(k-1)^{1-2/\Delta}-1} = \mathfrak{B}_o. \tag{58}$$

The corollary will then follow from Theorem 3.

To prove (57), let $G = (V, E)$ be an input graph to the problem #BipFerro-Potts(k, B, Δ) with $B = (k-1)/(k-2) > 1$. Construct an instance G' of #Bip-Colorings(k, Δ) by subdividing each edge of G, i.e., G' is a graph with vertex set $V' = V \cup E$ and edge set $E' = \bigcup_{e=(u,v)\in E}\{(u, e), (e, v)\}$. It is clear that G' is bipartite and every vertex has degree at most Δ.

We claim that the partition function for the k-state ferromagnetic Potts model on G with $B = (k-1)/(k-2)$ is equal to the number of proper k-colorings on G' (times an easily computable factor equal to $(k-2)^{|E|}$).

To see this, for a k-coloring σ' of G', map σ' to a configuration σ on G given by the restriction of σ' to vertices in V. Let $\sigma : V \to [k]$ be any configuration of the Potts model on G. The claim will follow by showing that the number of colorings of G' which map to σ is given by $(k-1)^{m(\sigma)}(k-2)^{|E|-m(\sigma)}$, where $m(\sigma)$ denotes the number of monochromatic edges in G under σ. Indeed, for a monochromatic edge $e = (u, v) \in E$ under σ, there are $k-1$ ways to choose the color of the vertex $e \in V'$ in the graph G'. In contrast, if $e = (u, v) \in E$ is not monochromatic under σ, there are $k-2$ ways to choose the color of the vertex $e \in V'$ in the graph G'. This completes the proof of (57).

We next show (58). For $\Delta \le 16$, we have $k \le \frac{\Delta}{2\ln\Delta} < 3$, so we may assume that $\Delta \ge 17$ (otherwise, there is nothing to prove). We first reduce (58) to the case $k = \frac{\Delta}{2\ln\Delta}$. Let $d := \Delta-1$, $z_o := (k-1)^{1/(d+1)}$. The inequality $\frac{k-1}{k-2} > \frac{k-2}{(k-1)^{(d-1)/(d+1)}-1}$ is equivalent to

$$h_d(z_o) > 0 \text{ where } h_d(z) = -z_o^{2(d+1)} + z_o^{2d} + z_o^{d+1} - 1.$$

Fix $d \ge 16$. The polynomial $h_d(z)$ has two changes of sign, so by the Descartes' rule of signs, it has at most two positive roots. Clearly, $z = 1$ is a root of $h_d(z)$ and since $h_d'(1) > 0$, there is one more root $z_d > 1$. Thus, $h_d(z_o) > 0$ iff $z_o < z_d$ (note that $z_o > 1$). Thus, to show (58), it suffices to consider the case where $k = \Delta/(2\ln\Delta)$.

Now, we prove the desired inequality for $k = \Delta/(2\ln\Delta)$. The inequality is equivalent to

$$(k-1)^{-2/\Delta} > 1 - \frac{k-2}{(k-1)^2}.$$

Now using the bound $x^{-2/\Delta} = \exp\big(-\frac{2\ln x}{\Delta}\big) \ge 1 - \frac{2\ln x}{\Delta}$ for $x = k-1$, we only need to prove that

$$\frac{2\ln(k-1)}{\Delta} < \frac{k-2}{(k-1)^2}.$$

Now $2(k-1)\ln(k-1) \le 2k\ln k = \Delta(1 - \frac{\ln(2\ln\Delta)}{\ln\Delta})$, so the inequality will follow from

$$k - 1 > \frac{\ln\Delta}{\ln(2\ln\Delta)},$$

which holds for $k = \Delta/(2\ln\Delta)$ and all $\Delta \ge 17$, as wanted. This completes the proof of (58).

This concludes the proof of Corollary 4. ∎

10. Torpid mixing of Swendsen–Wang. In this section, we prove Theorem 6 about torpid mixing of the Swendsen–Wang algorithm at the critical activity $B = \mathfrak{B}_o$. More precisely, we will show that, with probability $1-o(1)$ over the choice of a random Δ-regular graph with n vertices, the mixing time of the Swendsen–Wang algorithm is exponential in n. We will exploit Theorem 5 for $B = \mathfrak{B}_o$, which in combination with Lemma 16 essentially implies that for this value of B, we have *coexistence* of the ordered and disordered phases in a random Δ-regular graph (with probability $1-o(1)$).

Denote by $\mathbf{u}$ the disordered phase and by $\mathbf{m}_1,\dots,\mathbf{m}_q$ the q ordered phases of Theorem 5. Note that the q ordered phases are identical up to a permutations of the colors and for the purposes of this section we can treat them in a uniform manner. Thus, denote $\mathbf{m}=\{\mathbf{m}_1,\dots,\mathbf{m}_q\}$. We will say that a configuration σ is close to $\mathbf{u}$ (resp., $\mathbf{m}$) if the color frequencies in σ are close to those prescribed by $\mathbf{u}$ (resp., one of $\mathbf{m}_1,\dots,\mathbf{m}_q$).

Lemma 16 implies that, with probability $1-o(1)$ over the choice of the graph, the set of configurations near $\mathbf{u}$ and $\mathbf{m}$ dominate the Gibbs distribution, in the sense that these sets each have measure $\geq 1/poly(n)$ and the rest of the configurations have exponentially smaller mass. To analyze the Swendsen–Wang algorithm we need a more refined picture which includes the number of monochromatic edges in such configurations. To this end, we define the following quantities which roughly correspond to the expected number of monochromatic edges for configurations in these two sets scaled by a factor of n (see the upcoming equation (64) and the remarks thereafter for the derivations). Hence, let

$$(59)\qquad E_{\mathbf{m}} := \frac{\Delta}{2}\frac{B(x^2+q-1)}{(x+q-1)^2+(B-1)(x^2+q-1)} \text{ and } E_{\mathbf{u}} := \frac{\Delta}{2}\frac{B}{q+B-1},$$

where x, defined in Lemma 27, is a solution of the normalized tree recursions.

Now we can define the set of configurations with vertex marginals close to $\mathbf{u}$ and $\mathbf{m}$ and edge marginals close to $nE_{\mathbf{u}}$ and $nE_{\mathbf{m}}$, respectively. For a configuration $\sigma\in\Omega$, let $e_G(\sigma)$ denote the number of *monochromatic* edges in G under the spin configuration σ. Recall $\sigma^{-1}(i)$ is the set of vertices with spin i in σ. Let $\mathbf{c}(\sigma)$ denote the vector $(|\sigma^{-1}(1)|/n,\dots,|\sigma^{-1}(q)|/n)$. For $\varepsilon>0$, let

(60)
$$U=U(\varepsilon):=\left\{\sigma\in\Omega \;\middle|\; \|\mathbf{c}(\sigma)-\mathbf{u}\|_\infty\leq\varepsilon \text{ and } |e_G(\sigma)-E_{\mathbf{u}}n|<\varepsilon n\right\}.$$

(61)
$$M=M(\varepsilon):=\left\{\sigma\in\Omega \;\middle|\; \text{ for some } \mathbf{m}_j,\ \|\mathbf{c}(\sigma)-\mathbf{m}_j\|_\infty\leq\varepsilon \text{ and } |e_G(\sigma)-E_{\mathbf{m}}n|<\varepsilon n\right\},$$

(62)
$$T=T(\varepsilon):=\Omega\setminus(U(\varepsilon)\cup M(\varepsilon)).$$

The following lemma is proved in section 10.1 and is the main tool to obtain our torpid mixing results.

Lemma 31. *Let $B=\mathfrak{B}_o$. For all sufficiently small $\varepsilon>0$, there exists $C>0$, such that with probability $1-o(1)$ over the choice of the graph $G\sim\mathcal{G}(n,\Delta)$, it holds that*

$$(63)\qquad \mu_G(U)\geq 1/poly(n),\qquad \mu_G(M)\geq 1/poly(n),\qquad \mu_G(T)\leq\exp(-Cn).$$

For the rest of this section, we fix a graph G whose Gibbs distribution satisfies (63). By Lemma 31, this holds for asymptotically almost all Δ-regular graphs.

Now to prove that the chain is torpidly mixing we will bound its conductance defined as $\Phi_{SW}=\min_{S;\emptyset\subset S\subset\Omega}\Phi_{SW}(S)$, where

$$\Phi_{SW}(S)=\frac{\sum_{\sigma\in S}\mu(\sigma)P(\sigma,\overline{S})}{\mu(S)\mu(\overline{S})},$$

where P denotes the transition matrix for Swendsen–Wang. To bound the conductance of the set M, we prove that a configuration in M is unlikely to transition to U in one step.

Lemma 32. *For $\sigma \in M$, $P\big(\sigma, U\big) < \exp(-cn)$ for some positive constant $c > 0$.*

Proof. We are going to argue that, with probability $1 - \exp(-\Omega(n))$, the number of mononochromatic edges after one transition of Swendsen–Wang is too large to be in the set U. Note that the third step of Swendsen–Wang cannot decrease the number of monochromatic edges, so it suffices to analyze the first two steps.

Since $\sigma \in M$, by definition, the number of monochromatic edges under σ is $(E_{\mathbf{m}} \pm \varepsilon)n$. The expected number of edges left after the second step of Swendsen–Wang is thus $(1 - 1/B)(E_{\mathbf{m}} \pm \varepsilon)n$. The following claim implies that for sufficiently small ε, this is greater than $(E_{\mathbf{u}} \pm \varepsilon)n$, the number of monochromatic edges in a configuration from U. The proof is given in section 10.1.

Claim 33. *Let $\Delta \geq 3$ and $q \geq 2\Delta/\log\Delta$. For $B = \mathfrak{B}_o$, it holds that $E_{\mathbf{m}}/E_{\mathbf{u}} > 1/(1 - 1/B)$.*

By standard Chernoff bounds, we can thus conclude that for sufficiently large n the transition from σ to U happens with exponentially small probability. ◻

Remark 34. We believe that the lower bound on q in Theorem 6 arises from ignoring the effect of the third step of Swendsen–Wang.

We can now bound the conductance Φ_{SW} of Swendsen–Wang.

$$\begin{aligned}
\Phi_{SW} \leq \Phi_{SW}(M) &= \frac{\sum_{\sigma\in M}\mu(\sigma)P(\sigma,\overline{M})}{\mu(M)\mu(\overline{M})} \\
&\leq \mathit{poly}(n)\Big(\sum_{\sigma\in M}\mu(\sigma)P(\sigma,U) + \sum_{\sigma\in M}\mu(\sigma)P(\sigma,T)\Big) \quad \text{by (63)} \\
&\leq \mathit{poly}(n)\Big(\exp(-Cn)\mu(M) + \sum_{\tau\in T}\mu(\tau)P(\tau,M)\Big) \text{ by Lemma 32 and reversibility} \\
&\leq \mathit{poly}(n)\big(\exp(-Cn) + \mu(T)\big) \\
&\leq \exp(-C'n) \qquad \text{by (63).}
\end{aligned}$$

Proof of Theorem 6. Standard conductance results imply that the mixing time is $\Omega(1/\Phi_{SW})$, which proves the theorem based on the above bounds. (See [37, p. 255] for such a statement using the form of (normalized) conductance as used here.) ◻

10.1. Phase coexistence for random Δ-regular graphs. In this section, we prove Lemma 31. The lemma will mostly follow from Lemma 16. We will need though a more refined analysis of the partition function conditioned on configurations close to $\mathbf{u}$ and $\mathbf{m}$.

In analogy to (60), (61), define

$$\begin{aligned}
\widehat{U} = \widehat{U}(\varepsilon) &:= \Big\{\sigma \in \Omega \,\Big|\, \|\mathbf{c}(\sigma) - \mathbf{u}\|_\infty \leq \varepsilon\Big\}, \\
\widehat{M}_j = \widehat{M}_j(\varepsilon) &:= \Big\{\sigma \in \Omega \,\Big|\, \|\mathbf{c}(\sigma) - \mathbf{m}_j\|_\infty \leq \varepsilon\Big\}. \\
\widehat{T} = \widehat{T}(\varepsilon) &:= \Omega\backslash\big(\widehat{U} \cup \widehat{M}_1 \cup \ldots \cup \widehat{M}_q\big),
\end{aligned}$$

and for a subset $\Sigma \subseteq \Omega$ of the configuration space, denote by $Z_G(\Sigma)$ the partition function restricted to configurations in Σ, that is,

$$Z_G(\Sigma) = \sum_{\sigma\in\Sigma} w_G(\sigma), \text{ so that } \mu_G(\Sigma) = \frac{Z_G(\Sigma)}{Z_G}.$$

We first prove the following weaker version of Lemma 31.

LEMMA 35. *Let $B = \mathfrak{B}_o$. Let $\Psi = \max_{\boldsymbol{\alpha} \in \triangle_q} \Psi_1(\boldsymbol{\alpha})$, where $\Psi_1(\boldsymbol{\alpha})$ is given by* (13). *For all sufficiently small $\varepsilon > 0$, there exists $C > 0$, such that with probability $1 - o(1)$ over the choice of the graph $G \sim \mathcal{G}(n, \Delta)$, it holds that $Z_G(\widehat{T}) \leq \exp(-Cn)\exp(\Psi n)$ and*

$$\frac{\exp(n\Psi)}{poly(n)} \leq \min\left\{Z_G(\widehat{U}), Z_G(\widehat{M}_1), \ldots, Z_G(\widehat{M}_q)\right\}$$
$$\leq \max\left\{Z_G(\widehat{U}), Z_G(\widehat{M}_1), \ldots, Z_G(\widehat{M}_q)\right\} \leq poly(n)\exp(n\Psi).$$

Proof of Lemma 35. Define the region $\mathcal{T}$ by

$$\mathcal{T} = \{\boldsymbol{\alpha} \in \triangle_q \mid \|\boldsymbol{\alpha} - \mathbf{u}\|_\infty \geq \varepsilon, \|\boldsymbol{\alpha} - \mathbf{m}_1\|_\infty \geq \varepsilon, \ldots, \|\boldsymbol{\alpha} - \mathbf{m}_q\|_\infty \geq \varepsilon\}.$$

Recall from Theorem 5 that, for $B = \mathfrak{B}_o$, the global maximum of $\Psi_1(\boldsymbol{\alpha})$ occurs exactly when $\boldsymbol{\alpha}$ is equal to one of $\mathbf{u}, \mathbf{m}_1, \ldots, \mathbf{m}_q$. It follows that

$$\max_{\boldsymbol{\alpha} \in \mathcal{T}} \Psi_1(\boldsymbol{\alpha}) < \Psi - C'$$

for some constant $C' = C'(\varepsilon) > 0$. Note that for fixed n the possible values of $\boldsymbol{\alpha} \in \mathcal{T}$ are polynomially many. For all sufficiently large n, we thus have

$$\mathbf{E}_{\mathcal{G}}[Z_G(\widehat{T})] \leq poly(n) \max_{\boldsymbol{\alpha} \in \mathcal{T}} \mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}] \leq poly(n)\exp\left((\Psi - C')n\right).$$

Let C be such that $C' > C > 0$. By Markov's inequality, we obtain $Z_G(\widehat{T}) \leq \exp(-Cn)\exp(\Psi n)$ with probability $1 - o(1)$. This establishes the first part. For the second part, the upper bounds follow from Markov's inequality by the same token. The lower bounds follow from Lemma 16 and the observation that at the disordered/ordered critical activity $\mathfrak{B}_o$ the expectations of $Z_G^{\mathbf{u}}, Z_G^{\mathbf{m}_1}, \ldots, Z_G^{\mathbf{m}_q}$ are within a polynomial factor of $\exp(n\Psi)$. By a simple union bound we can thus ensure all the properties stated with probability $1 - o(1)$, as desired. ∎

Having established Lemma 35, we are ready to start arguing about the empirical distribution of edges, more precisely, the fraction of edges in G whose endpoints are assigned colors i, j for configurations $\sigma \in \widehat{U}, \widehat{M}$. The rough idea is as follows. Fix an arbitrary $\boldsymbol{\alpha} \in \triangle_q$. In section 4, we established that $\Psi_1(\boldsymbol{\alpha}) = \max_{\mathbf{x}} \Upsilon_1(\boldsymbol{\alpha}, \mathbf{x}) = (\Delta - 1)f_1(\boldsymbol{\alpha}) + \Delta \max_{\mathbf{x}} g_1(\mathbf{x})$, where the latter maximization is over $\mathbf{x}$ which satisfy the constraints (9). Since $g_1(\mathbf{x})$ is strictly concave in the convex region it is defined, this maximum is attained for a unique vector $\mathbf{x}$, which from here on we shall denote by $\mathbf{x}_{\boldsymbol{\alpha}}$. Essentially by the same line of arguments as in the proof of Lemma 35, all the contribution to the first moment $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$ comes from those $\mathbf{x}$ which are close to $\mathbf{x}_{\boldsymbol{\alpha}}$. Thus, by Markov's inequality and the lower bounds of Lemma 35, if $\mathbf{x}$ is sufficiently far away from $\mathbf{x}_{\boldsymbol{\alpha}}$, with probability $1 - o(1)$ over the choice of the graph G, the empirical edge distribution of a configuration $\sigma \sim \mu_G$ will equal $\mathbf{x}$ with exponentially small probability. We are thus left to argue that for those $\mathbf{x}$ close to $\mathbf{x}_{\boldsymbol{\alpha}}$, the actual contribution to $Z_G^{\boldsymbol{\alpha}}$ from configurations with edge empirical distribution $\mathbf{x}$ is close to its expectation. But for a graph satisfying Lemma 35 this is immediately guaranteed, since we know a lower bound on $Z_G^{\boldsymbol{\alpha}}$ which matches its expectation up to a polynomial factor.

It is useful at this point to give the expressions for the optimal vector $\mathbf{x}_{\boldsymbol{\alpha}}$, when $\boldsymbol{\alpha}$ is a dominant phase. Note that a dominant phase $\boldsymbol{\alpha}$ corresponds (via (20)) to a fixpoint $(R_1, \dots, R_q)$ of the tree recursions (6). The entries of the optimal vector $\mathbf{x}^* = \mathbf{x}_{\boldsymbol{\alpha}}$ are given by

$$x^*_{ij} = \frac{B_{ij} R_i R_j}{\sum_{i,j} B_{ij} R_i R_j} \text{ for } i, j \in [q]. \tag{64}$$

Indeed, using that $(R_1, \dots, R_q)$ specify a fixpoint of the tree recursions (6), it can be checked that $\sum_j x^*_{ij} = \alpha_i$ for $i \in [q]$, and also that $\mathbf{x}^*$ is a critical point of $g(\mathbf{x})$. By the strict concavity of the function g_1, it follows that $\mathbf{x}^*$ is the unique optimal vector, as claimed. Note that the expressions for $E_{\mathbf{m}}$ and $E_{\mathbf{u}}$ in (59) can easily be derived from (64) when adapted to the ferromagnetic Potts model.

We next introduce some relevant notation. We first want to capture the contribution of configurations with specific edge empirical distribution. To do this, we need the notation $\mathbf{x}_G(\sigma)$ of section 4 introduced just before the expression for the first moment (8). For $\mathbf{x}_0 \in \triangle_{q^2}$, $\Sigma \subseteq \Omega$, define $Z_G(\Sigma, \mathbf{x}_0)$ as

$$Z_G(\Sigma, \mathbf{x}_0) = \sum_{\sigma \in \Sigma} w_G(\sigma) \mathbf{1}\{\mathbf{x}_G(\sigma) = \mathbf{x}_0\}, \text{ so that } Z_G(\Sigma) = \sum_{\mathbf{x} \in \triangle_{q^2}} Z_G(\Sigma, \mathbf{x}).$$

Thus, the definition of $Z_G(\Sigma, \mathbf{x}_0)$ restricts the partition function not only to configurations belonging to Σ but also to those having edge empirical distribution equal to $\mathbf{x}_0$. Now, we further extend this definition to capture that the edge empirical distribution is far from a prescribed vector $\mathbf{x}_0$. For $\varepsilon > 0$, define $Z_G(\Sigma, \mathbf{x}_0, \varepsilon)$ as

$$Z_G(\Sigma, \mathbf{x}_0, \varepsilon) = \sum_{\sigma \in \Sigma} \sum_{\mathbf{x}; \|\mathbf{x} - \mathbf{x}_0\|_\infty \geq \varepsilon} w_G(\sigma) \mathbf{1}\{\mathbf{x}_G(\sigma) = \mathbf{x}\}.$$

We are ready to prove the following.

LEMMA 36. *Let* $\Psi = \max_{\boldsymbol{\alpha} \in \triangle_q} \Psi_1(\boldsymbol{\alpha})$, *where* $\Psi_1(\boldsymbol{\alpha})$ *is given by* (13). *For all sufficiently small* $\varepsilon > 0$, *there exists* $C > 0$ *such that with probability* $1 - o(1)$ *over the choice of the graph* $G \sim \mathcal{G}(n, \Delta)$, *it holds that*

$$\max\{Z_G(\widehat{U}, \mathbf{x}_{\mathbf{u}}, \varepsilon), Z_G(\widehat{M}_1, \mathbf{x}_{\mathbf{m}_1}, \varepsilon), \dots, Z_G(\widehat{M}_q, \mathbf{x}_{\mathbf{m}_q}, \varepsilon)\} \leq \exp(-Cn) \exp(\Psi n).$$

Proof. We prove that the upper bound $Z_G(\widehat{U}, \mathbf{x}_{\mathbf{u}}, \varepsilon) \leq \exp(-Cn) \exp(\Psi n)$ holds with probability $1 - o(1)$ over the choice of the graph. The remaining random variables may be treated similarly and thus the claim follows by a union bound.

Observe that

$$Z_G\big(\widehat{U}, \mathbf{x}_{\mathbf{u}}, \varepsilon\big) = \sum_{\boldsymbol{\alpha}; \|\boldsymbol{\alpha} - \mathbf{u}\|_\infty \leq \varepsilon} Z_G(\Sigma^{\boldsymbol{\alpha}}, \mathbf{x}_{\mathbf{u}}, \varepsilon).$$

It follows that

$$\mathbf{E}_{\mathcal{G}}\big[Z_G\big(\widehat{U}, \mathbf{x}_{\mathbf{u}}, \varepsilon\big)\big] = \sum_{\boldsymbol{\alpha}; \|\boldsymbol{\alpha} - \mathbf{u}\|_\infty \leq \varepsilon} \mathbf{E}_{\mathcal{G}}[Z_G(\Sigma^{\boldsymbol{\alpha}}, \mathbf{x}_{\mathbf{u}}, \varepsilon)]. \tag{65}$$

Note that the sum in (65) is over polynomially many vectors $\boldsymbol{\alpha}$ satisfying $\|\boldsymbol{\alpha} - \mathbf{u}\|_\infty \leq \varepsilon$. Further, for fixed $\boldsymbol{\alpha}$, $\mathbf{E}_{\mathcal{G}}[Z_G(\Sigma^{\boldsymbol{\alpha}}, \mathbf{x}_{\mathbf{u}}, \varepsilon)]$ is also a sum over polynomially many $\mathbf{x}$ satisfying $\|\mathbf{x} - \mathbf{x}_{\mathbf{u}}\|_\infty \geq \varepsilon$. For a fixed $\mathbf{x}$, the exponential order of the term in the latter sum corresponding to $\mathbf{x}$ is given by the function $\Upsilon_1(\boldsymbol{\alpha}, \mathbf{x})$. By approximating the sums with their maximum terms (and using the continuity of the function Υ_1), it is standard to conclude from here that

$$\mathbf{E}_{\mathcal{G}}\big[Z_G\big(\widehat{U}, \mathbf{x}_{\mathbf{u}}, \varepsilon\big)\big] = \exp(o(n)) \exp\Big(n \max_{\|\boldsymbol{\alpha}-\mathbf{u}\|_\infty \leq \varepsilon} \max_{\|\mathbf{x}-\mathbf{x}^*_{\mathbf{u}}\|_\infty \geq \varepsilon} \Upsilon_1(\boldsymbol{\alpha}, \mathbf{x})\Big). \tag{66}$$

Note that the maximum in (66) is justified by standard compactness arguments (which we give below since we will need it in the proof).

Consider the region

$$\mathcal{T}(\varepsilon) := \{(\boldsymbol{\alpha}, \mathbf{x}) \mid \boldsymbol{\alpha} \in \triangle_q, \mathbf{x} \in \triangle_{q^2}, \|\boldsymbol{\alpha} - \mathbf{u}\|_\infty \leq \varepsilon, \|\mathbf{x} - \mathbf{x}_{\mathbf{u}}\|_\infty \geq \varepsilon\},$$

i.e., the region $T(\varepsilon)$ consists of those pairs $(\boldsymbol{\alpha}, \mathbf{x})$ such that $\boldsymbol{\alpha}$ is ε-close to $\mathbf{u}$, but $\mathbf{x}$ is ε-far from the optimal vector $\mathbf{x}_{\mathbf{u}}$. Let Ψ' be the maximum of $\Upsilon_1(\boldsymbol{\alpha}, \mathbf{x})$ over the region $\mathcal{T}(\varepsilon)$ (this exists since $\mathcal{T}(\varepsilon)$ is compact and $\Upsilon_1(\boldsymbol{\alpha}, \mathbf{x})$ is continuous). Recall now that for fixed $\boldsymbol{\alpha}$ the function $\Upsilon_1(\boldsymbol{\alpha}, \mathbf{x})$ is strictly concave in $\mathbf{x}$. Since the maximizers of $\Psi_1(\boldsymbol{\alpha}) = \max_{\mathbf{x} \in \triangle_{q^2}} \Upsilon_1(\boldsymbol{\alpha}, \mathbf{x})$ are the vectors $\mathbf{u}, \mathbf{m}_1, \ldots, \mathbf{m}_q$, it follows that the maximizers of $\Upsilon_1(\boldsymbol{\alpha}, \mathbf{x})$ are $(\mathbf{u}, \mathbf{x}_{\mathbf{u}}), (\mathbf{m}_1, \mathbf{x}_{\mathbf{m}_1}), \ldots, (\mathbf{m}_q, \mathbf{x}_{\mathbf{m}_q})$. Since for all sufficiently small $\varepsilon > 0$ none of these maximizers lies in the region $\mathcal{T}(\varepsilon)$, we have that $\Psi' < \Psi$ and hence there exists $C'(\varepsilon) > 0$ such that

$$\max_{\|\boldsymbol{\alpha}-\mathbf{u}\|_\infty \leq \varepsilon} \max_{\|\mathbf{x}-\mathbf{x}_{\mathbf{u}}\|_\infty \geq \varepsilon} \Upsilon_1(\boldsymbol{\alpha}, \mathbf{x}) < \Upsilon_1(\mathbf{u}, \mathbf{x}_{\mathbf{u}}) - C'.$$

By choosing C so that $0 < C < C'$, the desired bound now follows from (65) by an application of Markov's inequality. ∎

We are now ready to give the proof of Lemma 31.

Proof of Lemma 31. By a union bound, a graph $G \sim \mathcal{G}(n, \Delta)$ satisfies with probability $1 - o(1)$ both Lemmas 35 and 36. We have $Z_G = Z_G\big(\widehat{U}\big) + Z_G\big(\widehat{M}_1\big) + \cdots + Z_G\big(\widehat{M}_q\big) + Z_G\big(\widehat{T}\big)$. By Lemma 35, we thus have

$$\frac{\exp(n\Psi)}{poly(n)} \leq Z_G \leq poly(n) \exp(n\Psi).$$

Now observe that the sets U, M defined in (60) and (61) satisfy

$$Z_G(U) \geq Z_G\big(\widehat{U}\big) - Z_G\big(\widehat{U}, \mathbf{x}^*_{\mathbf{u}}, \varepsilon\big) \geq \frac{\exp(n\Psi)}{poly(n)},$$

$$Z_G(M) \geq Z_G\big(\widehat{M}\big) - Z_G\big(\widehat{M}, \mathbf{x}^*_{\mathbf{m}}, \varepsilon\big) \geq \frac{\exp(n\Psi)}{poly(n)},$$

$$Z_G(T) \leq Z_G\big(\widehat{T}\big) \leq \exp(-Cn) \exp(n\Psi).$$

The conclusion follows. ∎

To complete the proofs for section 10, we now give the proof of Claim 33.

Proof of Claim 33. Using (59), we have that

$$\frac{E_{\mathbf{m}}}{E_{\mathbf{u}}} = \frac{(q+B-1)(x^2+q-1)}{(x+q-1)^2 + (B-1)(x^2+q-1)} = 1 + \frac{(q-1)(x-1)^2}{(x+q-1)^2 + (B-1)(x^2+q-1)}.$$

It follows that

$$\frac{E_{\mathbf{m}}}{E_{\mathbf{u}}} - \frac{1}{1-1/B} = \frac{(q-1)(x-1)^2}{(x+q-1)^2 + (B-1)(x^2+q-1)} - \frac{1}{B-1}.$$

Using the substitutions $d=\Delta-1$, $x=y^d$ and the second equation in (35) (for $t=1$), the r.h.s. can be rewritten as

$$\frac{E_{\mathbf{m}}}{E_{\mathbf{u}}}-\frac{1}{1-1/B}=\frac{y\left(y^d-y\right)\left((q-2)y^d-(q-1)y^{d-1}-(q-1)\right)}{(y-1)\left(y^d+q-1\right)\left(y^{d+1}+q-1\right)}.$$

Recall, from the proof of Lemma 27, that for $B=\mathfrak{B}_o$, it holds that $y=y_o$, where $y_o=(q-1)^{2/(d+1)}$. Since $y_o>1$, to prove the claim we only need to show that $p(y_o)>0$, where $p(y):=(q-2)y^d-(q-1)y^{d-1}-(q-1)$. Massaging, we obtain the equivalent inequality

$$h_d((q-1)^{1/(d+1)})>0, \text{ where } h_d(z):=z^{2d}-z^{2(d-1)}-z^{d-1}-1. \tag{67}$$

Fix $d\geq 2$. If z_d is such that $h_d(z_d)>0$, then $h_d(z)>0$ for all $z>z_d$. This is again a consequence of the Descartes' rule of signs: the polynomial $h_d(z)$ has exactly one positive root, say, ρ_d, and hence $h_d(z)>0$ for z positive is equivalent to $z>\rho_d$. It follows that to prove (67) for $q\geq 2(d+1)/\ln(d+1)$, we only need to argue for its validity when $q=q_o:=2(d+1)/\ln(d+1)$. In other words, we need to show that $h_d(z_o)>0$, where $z_o:=(q_o-1)^{1/(d+1)}$. By direct calculations, it can be checked that the inequality is true for $d=2,\dots,9$. We therefore assume that $d\geq 10$ in what follows.

For all $w>1$ it holds that $(w+3/4)^2>w^2+w+1$, which for $w=z_o^{d-1}$ gives $(z_o^{d-1}+3/4)^2>z_o^{2(d-1)}+z_o^{d-1}+1$. Thus, we only need to show that $z_o^d>z_o^{d-1}+3/4$, or

$$z_o^{d-1}(z_o-1)>3/4. \tag{68}$$

To handle the two factors in (68), we will use the following bounds on z_o:

$$z_o\geq\left(C_1\frac{d+1}{\ln(d+1)}\right)^{1/(d+1)}, \quad z_o\geq 1+C_2\frac{\ln(d+1)}{d+1}, \text{ where } C_1:=\frac{3}{2},\ C_2:=\frac{3}{4}. \tag{69}$$

The first bound follows from the inequality $q_o\geq 1+\frac{3}{2}\frac{d+1}{\ln(d+1)}$. The second bound follows from the first bound and $x^{1/(d+1)}=\exp\left(\frac{\ln x}{d+1}\right)\geq 1+\frac{\ln x}{d+1}$ together with $\ln(C_1\frac{d+1}{\ln(d+1)})\geq C_2\ln(d+1)$.

Plugging the bounds (69) into (68), it suffices to check whether

$$C_2\frac{\ln(d+1)}{d+1}\left(C_1\frac{d+1}{\ln(d+1)}\right)^{(d-1)/(d+1)}\geq\frac{3}{4}, \text{ or } \left(\frac{4C_1C_2}{3}\right)^{d+1}\geq\left(C_1\frac{d+1}{\ln(d+1)}\right)^2.$$

This holds for all $d\geq 10$, completing the proof (for all $d\geq 2$). ∎

11. Remaining proofs.

11.1. Small subgraph conditioning method. In this section, we give the outline for the proof of Lemma 16. The proof is a minor modification of the arguments in [17, Appendices A and B] which were carried out for random Δ-regular bipartite graphs. Here, we just need to account for the nonbipartite case which turns out to be completely analogous. For completeness, we give the adaptation of the calculations therein to account for the slightly different setting.

The main tool we are going to use is the following theorem, which is due to [42]. The notation $[X]_m$ refers to the mth order falling factorial of the variable X.

THEOREM 37. *For $i = 1, 2, \ldots$, let $\lambda_i > 0$ and $\delta_i > -1$ be constants and assume that for each n there are random variables X_{in}, $i = 1, 2, \ldots$, and Y_n, all defined on the same probability space $\mathcal{G} = \mathcal{G}_n$ such that X_{in} is nonnegative integer valued, $Y_n \geq 0$, and $\mathbf{E}\big[Y_n\big] > 0$ (for n sufficiently large). Furthermore, the following hold:*

(A1) *$X_{in} \xrightarrow{d} Z_i$ as $n \to \infty$, jointly for all i, where $Z_i \sim \mathrm{Po}(\lambda_i)$ are independent Poisson random variables;*

(A2) *for every finite sequence $j_1, \ldots, j_m$ of nonnegative integers,*

$$\frac{\mathbf{E}_{\mathcal{G}}\big[Y_n[X_{1n}]_{j_1} \cdots [X_{mn}]_{j_m}\big]}{\mathbf{E}_{\mathcal{G}}\big[Y_n\big]} \to \prod_{i=1}^{m} \big(\lambda_i(1+\delta_i)\big)^{j_i} \qquad \text{as } n \to \infty; \tag{70}$$

(A3) *$\sum_i \lambda_i \delta_i^2 < \infty$;*

(A4) *$\mathbf{E}_{\mathcal{G}}\big[Y_n^2\big] / \big(\mathbf{E}_{\mathcal{G}}[Y_n]\big)^2 \leq \exp\big(\sum_i \lambda_i \delta_i^2\big) + o(1)$ as $n \to \infty$.*

Let $r(n)$ be a function such that $r(n) \to 0$ as $n \to \infty$. It holds that $Y_n > r(n)\mathbf{E}_{\mathcal{G}}\big[Y_n\big]$ asymptotically almost surely.

To obtain Lemma 16, we verify the assumptions of Theorem 37 for the random variables $Z_G^{\boldsymbol{\alpha}}$. Recall that we restrict our attention to $\boldsymbol{\alpha}$ which are Hessian dominant. For $G \sim \mathcal{G}(n, \Delta)$, let $X_i = X_{in}$ be the number of cycles of length i in G, $i = 1, 2, \ldots$.

The most technical part of this verification is assumption (A4), which requires computing the precise asymptotics of the moments. This in turn reduces to certain determinants which are not completely trivial. Nevertheless, the arguments have been carried out in full generality in [17]. The only minor modification required in the present case is to account for random Δ-regular graphs instead of the bipartite random Δ-regular graphs studied in [17].

We obtain the following lemmas.

LEMMA 38. *Assumption* (A1) *holds with $\lambda_i = \frac{(\Delta-1)^i}{2i}$.*

Lemma 38 is well-known; see, for example, [27].

LEMMA 39. *Assumption* (A2) *holds with $\delta_i = \sum_{j=1}^{q-1} \mu_j^i$, where $\mu_1, \mu_2, \ldots, \mu_{q-1}$ are the eigenvalues different from* 1 *of the matrix* $\mathbf{M}$ *defined in section* 8*. For Hessian dominant $\boldsymbol{\alpha}$, it holds that the μ_i are positive and strictly smaller than $1/(\Delta-1)$.*

Proof of Lemma 39. The proof is close to [17, Proof of Lemma A.6], which is in turn close to [40, Proof of Lemma 7.4]. We just modify the approach to account for the distribution induced by the pairing model. We make the minor notation change from X_i to X_ℓ, i.e., for $\ell \geq 1$, X_ℓ denotes the number of cycles of length ℓ in G. We show that assumption (A2) in Theorem 37 holds when $m = 1$ and $j_1 = 1$, and the extension to $m > 1$ and arbitrary indices $j_1, \ldots, j_m$ follow by standard arguments; see, for example, [31, section 2] for an exposition of the argument in a very similar setting.

Let $\mathcal{S} = \{S_1, \ldots, S_q\}$ be a partition of V such that $|S_i| = \alpha_i n$ for all $i \in [q]$. Note that $\mathcal{S}$ induces a configuration $\sigma(\mathcal{S})$ by setting, for every vertex $v \in V$, $\sigma(v) = i$ iff $v \in S_i$. Denote by $Y_{\mathcal{S}}$ the weight of the configuration $\sigma(\mathcal{S})$.

Fix a specific partition $\mathcal{S}$. By symmetry,

$$\frac{\mathbf{E}[Z_G^{\boldsymbol{\alpha}} X_\ell]}{\mathbf{E}[Z_G^{\boldsymbol{\alpha}}]} = \frac{\mathbf{E}[Y_{\mathcal{S}} X_\ell]}{\mathbf{E}[Y_{\mathcal{S}}]}. \tag{71}$$

We decompose X_ℓ as follows:

- ξ will denote a rooted and oriented ℓ-cycle, whose vertices are colored with $\{1,\ldots,q\}$ (note that the coloring is *not* assumed to be proper). A vertex colored i in ξ signifies that it corresponds to a vertex in S_i.
- Once we have specified ξ, we use ζ to denote the ℓ *points* that the cycle traverses in order, such that the prescription of the vertex colors of ξ is satisfied. (Recall from section 4 that points are elements of $[\Delta n]$; ζ specifies the preimage of the cycle ξ in the pairing model, i.e., the matched points that respect the colors ξ of the cycle with respect to the partition $\mathcal{S}$).
- $\mathbf{1}_{\xi,\zeta}$ is the indicator function whether the cycle specified by ξ,ζ is present in the graph G generated by the pairing model.

Once the vertex colors of a cycle have been specified, note that each possible cycle corresponds to exactly 2ℓ different configurations ξ (the number of ways to root and orient the cycle). For each of those ξ, the respective sets of configurations ζ are the same. Hence, we may write

$$X_\ell = \frac{1}{2\ell}\sum_{\xi}\sum_{\zeta}\mathbf{1}_{\xi,\zeta}.$$

Let $p_1 := \Pr[\mathbf{1}_{\xi,\zeta}=1]$. It follows that

$$\mathbf{E}[Y_{\mathcal{S}}X_\ell] = \frac{1}{2\ell}\sum_{\xi}\sum_{\zeta}p_1\cdot\mathbf{E}[Y_{\mathcal{S}}|\mathbf{1}_{\xi,\zeta}=1].$$

In light of (71), we need to study the ratio $\mathbf{E}[Y_{\mathcal{S}}\mid\mathbf{1}_{\xi,\zeta}=1]/\mathbf{E}[Y_{\mathcal{S}}]$. At this point, to simplify notation, we may assume that ξ,ζ are fixed.

We have shown in section 4 that

(72)

$$\mathbf{E}[Y_{\mathcal{S}}] = \sum_{\mathbf{x}}\prod_{i}\binom{\Delta\alpha_i n}{\Delta x_{i1}n,\ldots,\Delta x_{iq}n}\frac{\left[\prod_{i\neq j}(\Delta x_{ij}n)!\right]^{1/2}\prod_i(\Delta x_{ii}n-1)!!}{(\Delta n-1)!!}\prod_{i,j}B_{ij}^{\Delta x_{ij}n/2},$$

where the variables $\mathbf{x}=(x_{11},\ldots,x_{qq})$ capture the number of edges between the different color classes in $\mathcal{S}$. In particular, for $i\neq j$, $\Delta x_{ij}n$ is the number of edges between the sets S_i and S_j, whereas $\Delta x_{ii}n/2$ is the number of edges within the set S_i (cf. section 4 for more details).

To calculate $\mathbf{E}[Y_{\mathcal{S}}\mid\mathbf{1}_{\xi,\zeta}=1]$, we need some notation. For colors $i,j\in\{1,\ldots,q\}$, let a'_{ij} be the number of edges in ξ whose one endpoint has color i and the other j. It will be convenient to denote $a_{ii} := 2a'_{ii}$ and $a_{ij} := a'_{ij}$ whenever $i\neq j$. Finally, let c_i denote the number of vertices in ξ colored with i. The following equalities are immediate:

$$\textstyle\sum_j a_{ij} = 2c_i,\ \sum_{i,j}a_{ij} = 2\ell. \tag{73}$$

We are almost set to compute $\mathbf{E}[Y_{\mathcal{S}}\mid\mathbf{1}_{\xi,\zeta}=1]$. We denote by $\mathbf{x}$ the same set of variables as in (72). This number includes the a_{ij} edges prescribed by ξ,ζ. To make the following formulas easier to digest let $n\Delta x'_{ij} = n\Delta x_{ij}-a_{ij}$. We have

$$\mathbf{E}[Y_{\mathcal{S}}\mid\mathbf{1}_{\xi,\zeta}=1] = \sum_{\mathbf{x}}\prod_{i}\binom{\Delta\alpha_i n-2c_i}{\Delta x'_{i1}n,\ldots,\Delta x'_{iq}n}\frac{\left[\prod_{i\neq j}(\Delta x'_{ij}n)!\right]^{1/2}\prod_i(\Delta x'_{ii}n-1)!!}{(\Delta n-2\ell-1)!!}\prod_{i,j}B_{ij}^{\Delta x_{ij}n/2}.$$

Using that for constants $c_1, c_2 > 0$, it holds that $(c_1 n - c_2)!/(c_1 n)! = (1 + o(1))(c_1 n)^{-c_2}$, we obtain

$$\frac{\binom{\Delta\alpha_i n - 2c_i}{\Delta x'_{i1} n, \ldots, \Delta x'_{iq} n}}{\binom{\Delta\alpha_i n}{\Delta x_{i1} n, \ldots, \Delta x_{iq} n}} \sim \frac{\prod_j \left(x_{ij}\right)^{a_{ij}}}{\alpha_i^{2c_i}}, \quad \frac{\frac{\left[\prod_{i\neq j}(\Delta x'_{ij} n)!\right]^{1/2} \prod_i (\Delta x'_{ii} n - 1)!!}{(\Delta n - 2\ell - 1)!!}}{\frac{\left[\prod_{i\neq j}(\Delta x_{ij} n)!\right]^{1/2} \prod_i (\Delta x_{ii} n - 1)!!}{(\Delta n - 1)!!}} \sim \frac{1}{\prod_{i,j}(x_{ij})^{a_{ij}/2}}.$$

The asymptotics of the ratio $\mathbf{E}[Y_{\mathcal{S}} \mid \mathbf{1}_{\xi,\zeta} = 1]/\mathbf{E}[Y_{\mathcal{S}}]$ are determined from those $\mathbf{x}^*$ which maximize $\Upsilon_1(\boldsymbol{\alpha}, \mathbf{x})$. Thus, we obtain

$$\frac{\mathbf{E}[Y_{\mathcal{S}} \mid \mathbf{1}_{\xi,\zeta} = 1]}{\mathbf{E}[Y_{\mathcal{S}}]} \sim \frac{\prod_{i,j}\left(x^*_{ij}\right)^{a_{ij}/2}}{\prod_i \alpha_i^{2c_i}}.$$

For given ξ, the number of possible ζ in the pairing model is asymptotic to $\prod_i \left[\Delta(\Delta - 1)\alpha_i n\right]^{c_i} = \left[\Delta(\Delta-1)n\right]^{\ell} \prod_i \alpha_i^{c_i}$. Since $p_1 = (\Delta n - 2\ell - 1)!!/(\Delta n - 1)!! \sim (\Delta n)^{-\ell}$, we have

$$\begin{aligned}\frac{\sum_\zeta p_1 \mathbf{E}[Y_{\mathcal{S}} \mid \mathbf{1}_{\xi,\zeta} = 1]}{\mathbf{E}[Y_{\mathcal{S}}]} \sim \frac{(\Delta-1)^{\ell} \prod_i \alpha_i^{c_i} \prod_{i,j}\left(x^*_{ij}\right)^{a_{ij}/2}}{\prod_i \alpha_i^{2c_i}} &= \frac{(\Delta-1)^{\ell} \prod_{i,j}\left(x^*_{ij}\right)^{a_{ij}/2}}{\prod_i \alpha_i^{\frac{1}{2}\sum_j a_{ij}}} \\ &= (\Delta-1)^{\ell} \prod_{i\le j}\left(\frac{x^*_{ij}}{\sqrt{\alpha_i\alpha_j}}\right)^{a'_{ij}},\end{aligned}$$

where in the last equality we used that $a_{ii} = 2a'_{ii}$, $a_{ij} = a'_{ij}$ for $i \neq j$, $a'_{ij} = a'_{ji}$, and $x^*_{ij} = x^*_{ji}$. Note that the r.h.s. evaluates to 0 whenever there exist i,j such that $B_{ij} = 0$ but $a_{ij} \neq 0$, since then we have $x^*_{ij} = 0$ (cf. (64). This is in complete accordance with the fact that the configuration induced by the partition $\mathcal{S}$ has zero weight. Thus, by (71), we may write

$$\frac{\mathbf{E}[Z_G^{\boldsymbol{\alpha}} X_\ell]}{\mathbf{E}[Z_G^{\boldsymbol{\alpha}}]} \sim \frac{(\Delta-1)^{\ell}}{2\ell} \cdot \sum_{\mathbf{a}'} N_{\mathbf{a}'} \prod_{i\le j}\left(\frac{x^*_{ij}}{\sqrt{\alpha_i\alpha_j}}\right)^{a'_{ij}},$$

where $\mathbf{a}' = \{a'_{11}, \ldots, a'_{qq}\}$ and $N_{\mathbf{a}'}$ is the number of possible ξ with a'_{ij} edges of type $\{i,j\}$. Using (64), we have that $x^*_{ij}/\sqrt{\alpha_i\alpha_j}$ is equal to the (i,j)-entry of the matrix $\mathbf{M}$. Thus, the sum can be reformulated as the (multiplicative) weight of walks in a weighted multigraph whose (weighted) adjacency matrix is given by $\mathbf{M}$ (for more details on the technique see [27]). It thus follows that the sum equals $\mathrm{Tr}(\mathbf{M}^{\ell}) = 1 + \sum_{j=1}^{q-1} \mu_j^{\ell}$. The fact that the μ_i's are positive follows from the fact that $\mathbf{B}$ is a positive definite matrix, while the fact that the μ_i's are less than $1/(\Delta - 1)$ follows from the results of section 6. ∎

Lemma 40. *Assumption* (A3) *holds with*

$$\exp\left(\sum_{i\ge1} \lambda_i \delta_i^2\right) = \prod_{i=1}^{q-1}\prod_{j=1}^{q-1}\left(1 - (\Delta-1)\mu_i\mu_j\right)^{-1/2},$$

where the μ_i*'s are as in Lemma* 39.

Proof of Lemma 40. We have

$$\sum_{i\geq 1}\lambda_i\delta_i^2=\sum_{i\geq 1}\frac{(\Delta-1)^i}{2i}\sum_{j=1}^{q-1}\sum_{k=1}^{q-1}\mu_j^i\mu_k^i=-\frac{1}{2}\sum_{j,k\in[q-1]}\ln\big(1-(\Delta-1)\mu_j\mu_k\big),$$

where we used that $\sum_{i\geq 1}\frac{x^i}{i}=-\ln(1-x)$ for $|x|<1$. ∎

LEMMA 41. *For a ferromagnetic model, for all Hessian dominant* $\boldsymbol{\alpha}$ *it holds that*

$$\frac{\mathbf{E}_{\mathcal{G}}\big[(Z_G^{\boldsymbol{\alpha}})^2\big]}{\big(\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]\big)^2}\to\prod_{i=1}^{q-1}\prod_{j=1}^{q-1}\big(1-(\Delta-1)\mu_i\mu_j\big)^{-1/2},$$

where the μ_i*'s are as in Lemma* 39.

Proof. Let $\mathbf{x}^*=\arg\max_{\mathbf{x}}\Upsilon_1(\boldsymbol{\alpha},\mathbf{x})$, $(\boldsymbol{\gamma}^*,\mathbf{y}^*)=\arg\max_{\boldsymbol{\gamma},\mathbf{y}}\Upsilon_2(\boldsymbol{\gamma},\mathbf{y})$. For $\boldsymbol{\alpha}$ which is Hessian dominant, Theorem 7 yields $\Upsilon_2(\boldsymbol{\gamma}^*,\mathbf{y}^*)=2\Upsilon_1(\boldsymbol{\alpha},\mathbf{x}^*)$.

Using methods in [17, Appendix B], we show in section 11.2 that

$$\lim_{n\to\infty}\frac{(2\pi n)^{(q-1)/2}\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]}{e^{n\Upsilon_1(\boldsymbol{\alpha},\mathbf{x}^*)}}=\Big(\prod_{i\in[q]}\alpha_i\prod_{i\in[q-1]}(1+\mu_i)\Big)^{-1/2} \tag{74}$$

and

$$\lim_{n\to\infty}\frac{(2\pi n)^{q-1}\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2]}{e^{n\Upsilon_2(\boldsymbol{\gamma}^*,\mathbf{y}^*)}}=\Big(\prod_{i\in[q]}\alpha_i\prod_{i\in[q-1]}(1+\mu_i)\Big)^{-1}\prod_{i=1}^{q-1}\prod_{j=1}^{q-1}\big(1-(\Delta-1)\mu_i\mu_j\big)^{-1/2}. \tag{75}$$

The statement of the lemma follows. ∎

Proof of Lemma 16. Lemmas 38–41 verify the assumptions of Theorem 37. The lemma thus follows by applying Theorem 37 for $r(n)=1/n$. ∎

11.2. The asymptotics of the moments. We follow closely the proof in [17, Appendix B], where very similar asymptotics are computed in detail. We first overview the approach in [17, Appendix B] in our setting.

The asymptotics of $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$ and $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]^2$ are derived by first rewriting the sums in (8) and (10) as integrals and approximating the latter with Gaussian integrals. The principle behind the technique is the negative-definiteness of the Hessian at the maximizers of the functions Υ_1 and Υ_2, which control the exponential order of the terms in the sums (8) and (10), respectively. This allows us to focus on terms within $O(1/\sqrt{n})$ distance around the term with the maximum contribution. A thorough exposition of the technical details can be found in [28, section 9.4].

Carrying out the above scheme in our setting is impeded by the fact that the sums in (8) and (10) are over variables which are linearly dependent. We will get rid of this linear dependence in the simplest way: for each of the two sums, we pick a subset S of the variables (with minimum cardinality) such that every variable is a (nontrivial) linear combination of variables in S. Variables in S span a full-dimensional space, inducing what we call a "full-dimensional representation" of the functions Υ_1,Υ_2 when these are viewed as functions of the variables in S. The inconvenience that this procedure causes is that in the calculation of the Gaussian integrals, (the determinant of) the Hessian matrices of the full-dimensional representations of Υ_1,Υ_2 come into play.

In [17, Appendix B.1.1], the above setting is abstracted as follows: given a linear subspace $\mathbf{A}\mathbf{z} = 0$, compute (the determinant of) the Hessian matrix of the full-dimensional representation of a function $\Upsilon(\mathbf{z})$. It is not hard to see that a full-dimensional representation of $\mathbf{A}\mathbf{z} = 0$, assuming that $\mathbf{A}$ has row rank r, is obtained by first picking a submatrix $\mathbf{A}_f$ induced by r linearly independent rows of $\mathbf{A}$ and then picking r columns of $\mathbf{A}_f$ to obtain an $r \times r$ invertible submatrix $\mathbf{A}_{fs}$ (the variables corresponding to the columns of $\mathbf{A}_{fs}$ can be written as nontrivial linear combinations of the remaining variables; the latter yield the full-dimensional representation). We denote by $\mathbf{H}^f$ the Hessian matrix of the full-dimensional representation of Υ induced by the matrices $\mathbf{A}_f, \mathbf{A}_{fs}$. We further denote by $\mathbf{H}$ the Hessian matrix of $\Upsilon(z)$ (where $\mathbf{z}$ is now assumed to be unconstrained); note that $\mathbf{H}$ is diagonal. The following is proved in [17].

For a polynomial $p(s)$, $[s^t]p(s)$ denotes the coefficient of s^t in $p(s)$.

LEMMA 42 (see [17, Lemma B.7]). *Suppose $\mathbf{A}$ consists of m rows and has rank r. Let $\mathbf{T}$ be a positive semidefinite diagonal matrix such that $[\mathbf{T}\ \mathbf{A}]$ has full row rank. If $\mathbf{H}$ is invertible, then*

$$\mathrm{Det}\big(-\mathbf{H}^f\big) = \frac{L\big(\mathbf{A}_f, \mathbf{A}, \mathbf{T}\big)}{\mathrm{Det}\big(\mathbf{A}_{fs}\big)^2}\,\mathrm{Det}(-\mathbf{H})\,[\varepsilon^{m-r}]\,\mathrm{Det}\big(\varepsilon\mathbf{T} - \mathbf{A}\mathbf{H}^{-1}\mathbf{A}^\intercal\big), \tag{76}$$

where $L\big(\mathbf{A}_f, \mathbf{A}, \mathbf{T}\big) = (-1)^r\,\mathrm{Det}\big(\mathbf{A}_f\mathbf{A}_f^\intercal\big)\Big/[\varepsilon^{m-r}]\,\mathrm{Det}\big(\varepsilon\mathbf{T} - \mathbf{A}\mathbf{A}^\intercal\big)$.

Remark 43. When $\mathbf{A}$ has full row rank, i.e., $r = m$, one can take $\mathbf{T}$ to be the identity matrix. Then, the r.h.s. in (76) simplifies into

$$\mathrm{Det}(-\mathbf{H})\,\mathrm{Det}\big(-\mathbf{A}\mathbf{H}^{-1}\mathbf{A}^\intercal\big)/\mathrm{Det}\big(\mathbf{A}_{fs}\big)^2.$$

We will apply Lemma 42 in sections 11.2.1 and 11.2.2 to calculate the asymptotics of the moments. To do this, we will need more information on the maximizers of the functions Υ_1 and Υ_2 for a Hessian dominant phase $\boldsymbol{\alpha}$. In particular, let $\mathbf{x}^*$ be the maximizer of $\Upsilon_1(\boldsymbol{\alpha}, \mathbf{x})$ (cf. (64)) and $(\boldsymbol{\gamma}^*, \mathbf{y}^*)$ be the maximizer of $\Upsilon_2(\boldsymbol{\gamma}, \mathbf{y})$. Adapting the proof of [17, Lemma 3.2], we have that $\gamma^*_{ij} = \alpha_i\alpha_j$ for all $i, j \in [q]$, from where it easily follows that $y^*_{ikjl} = x^*_{ij}x^*_{kl}$.

Following [17, Appendix B.1.2], we use the following notation in sections 11.2.1 and 11.2.2. For a vector $\mathbf{z} \in \mathbb{R}^n$ we denote by $\mathbf{z}^D$ the $n \times n$ diagonal matrix whose i-entry on the diagonal equals z_i for $i \in n$. For vectors $\mathbf{z}_j \in \mathbb{R}^{m_j}$, $j = 1, \dots, t$, we denote by $[\mathbf{z}_1, \dots, \mathbf{z}_t]^\intercal$ the $\mathbb{R}^{\sum_j m_j}$ vector which is the concatenation of the vectors $\mathbf{z}_1, \dots, \mathbf{z}_t$. $\mathbf{I}_n$ will denote the identity matrix with dimensions $n \times n$ and $\mathbf{0}$ will denote the all-zeros matrix whose dimensions will be inferred from context. For matrices $\mathbf{M}_1$ and $\mathbf{M}_2$, $\mathbf{M}_1 \otimes \mathbf{M}_2$ will denote the Kronecker product of $\mathbf{M}_1, \mathbf{M}_2$, while $\mathbf{M}_1 \oplus \mathbf{M}_2$ denotes the direct sum of $\mathbf{M}_1, \mathbf{M}_2$, that is, the block matrix $\left[\begin{smallmatrix}\mathbf{M}_1 & \mathbf{0}\\ \mathbf{0} & \mathbf{M}_2\end{smallmatrix}\right]$.

11.2.1. Proof of (74). The first moment $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$ is a sum over $\mathbf{x}$ (and $\boldsymbol{\alpha}$ is fixed). Note that if $B_{ij} = 0$, then we may restrict the sum in (8) to those $\mathbf{x}$ which satisfy $x_{ij} = 0$ without changing the sum's value. Further, since $x_{ij} = x_{ji}$ and $B_{ij} = B_{ji}$, we may write the sum in (8) in terms of those x_{ij} with $i \le j$. Let

$$P_1 = \big\{(i, j) \in [q]^2 \,\big|\, B_{ij} > 0, i \le j\big\}. \tag{77}$$

Henceforth $\mathbf{x}$ will denote $\{x_{ij}\}_{(i,j)\in P_1}$. The observations above imply that the sum in (8) can be written over the possible values of the vector $\mathbf{x}$. Note that for ferromagnetic models we have that $(i,i) \in P_1$ for every $i \in [q]$ (cf. the discussion after Theorem 8). We are left to account for the linear dependencies induced by the q constraints $a_i = \sum_j x_{ij}$. In matrix form, we can write those as

$$\mathbf{A}_1 \mathbf{x} = \boldsymbol{\alpha}, \tag{78}$$

where $\mathbf{A}_1$ is a $\{0,1\}$-matrix with dimension $q \times |P_1|$. For a ferromagnetic model, as we shall display shortly, we have that the rank of $\mathbf{A}_1$ is q (this holds more generally for matrices $\mathbf{B}$ which are irreducible and aperiodic; see, for example, footnote 7). To get a full-dimensional representation of the space (78), we will eliminate q variables from the vector $\mathbf{x}$. This corresponds to picking q columns of $\mathbf{A}_1$ which induce a $q \times q$ invertible submatrix of $\mathbf{A}_1$. We will denote this submatrix by $\mathbf{A}_{1,s}$. For ferromagnetic models, we can choose the columns corresponding to the variables x_{ii} for $i \in [q]$, in which case $\mathbf{A}_{1,s}$ is simply the identity matrix.[7]

Adapting the proof of [17, Lemma B.3] yields the following asymptotics for the first moment $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$. The details of the proof can be found in Appendix B.

LEMMA 44. *For a ferromagnetic model,*[8] *it holds that*

$$\lim_{n\to\infty} \frac{(2\pi n)^{(q-1)/2}\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]}{e^{n\Upsilon_1(\boldsymbol{\alpha},\mathbf{x}^*)}} = \Big(2^{q-1}\prod_{(i,j)\in P_1} x^*_{ij}\Big)^{-1/2} \mathrm{Det}(-\mathbf{H}^f_{1,\mathbf{x}})^{-1/2}, \tag{79}$$

where $\mathbf{H}^f_{1,\mathbf{x}}$ is the Hessian of the full-dimensional representation of $g_1(\mathbf{x})$ evaluated at $\mathbf{x} = \mathbf{x}^*$.

To expand the determinant in (79), we apply Lemma 42 and in particular Remark 43. This yields

$$\begin{aligned}\mathrm{Det}\big(-\mathbf{H}^f_{1,\mathbf{x}}\big) &= \frac{1}{\mathrm{Det}(\mathbf{A}_{1,s})^2}\,\mathrm{Det}(-\mathbf{H}_{1,\mathbf{x}})\,\mathrm{Det}\big(-\mathbf{A}_1\mathbf{H}^{-1}_{1,\mathbf{x}}\mathbf{A}_1^{\intercal}\big)\\ &= \mathrm{Det}(-\mathbf{H}_{1,\mathbf{x}})\,\mathrm{Det}\big(-\mathbf{A}_1\mathbf{H}^{-1}_{1,\mathbf{x}}\mathbf{A}_1^{\intercal}\big),\end{aligned} \tag{80}$$

where $\mathbf{H}_{1,\mathbf{x}}$ is the $|P_1| \times |P_1|$ diagonal matrix corresponding to the Hessian matrix of $g_1(\mathbf{x})$ (when $\mathbf{x}$ is unconstrained), and in the second equality in (80) we used that $\mathrm{Det}(\mathbf{A}_{1,s}) = 1$ (by our choice of $\mathbf{A}_{1,s}$).

Since $\mathbf{H}_{1,\mathbf{x}}$ is diagonal, we obtain that

$$\mathrm{Det}(-\mathbf{H}_{1,\mathbf{x}})^{-1} = 2^q \prod_{(i,j)\in P_1} x^*_{ij},$$

[7]In general, the $q \times q$ invertible submatrices $\mathbf{A}_{1,s}$ can be characterized as follows. First, view $\mathbf{A}_1$ as the (unsigned) incidence matrix of a graph H with vertex set $[q]$ and edge set P_1, where vertex i corresponds to the ith row of $\mathbf{A}_1$ and an edge labeled (i,j) corresponds to the column labeled (i,j) in $\mathbf{A}_1$ (note that H has a self-loop on vertex i iff $B_{ii} > 0$). Then $\mathbf{A}_{1,s}$ specifies a subgraph H' of H with exactly q edges. It can be shown that $\mathbf{A}_{1,s}$ is invertible if H' is spanning (i.e., every vertex in H' has nonzero degree) and all of the connected components of H' are unicyclic and nonbipartite (i.e., every connected component of H' has a unique cycle of odd length, where self-loops count as cycles of length 1).

[8]We briefly comment on how the choice of the full-dimensional representation (i.e., the choice of $\mathbf{A}_{1,s}$) has been used in the derivation of (79). Relative to footnote 7, if the invertible submatrix $\mathbf{A}_{1,s}$ corresponds to a subgraph with exactly c components which contain a nontrivial odd cycle (i.e., an odd cycle of length ≥ 3), there is a correction factor 2^{-c} in the r.h.s. of (80). The factor comes from $(\bmod 2)$ constraints imposed by considering the sum of constraints in (78) corresponding to vertices in each such unicyclic component (in the derivation below, this factor cancels with the factor $|\mathrm{Det}(\mathbf{A}_{1,s})|$ coming from $\mathrm{Det}(-\mathbf{H}^f_{1,\mathbf{x}})^{-1/2}$; it can be shown that $|\mathrm{Det}(\mathbf{A}_{1,s})| = 2^c$). Note that for our choice of $\mathbf{A}_{1,s}$, c equals zero, since the subgraph induced by the columns of $\mathbf{A}_{1,s}$ consists of q components, each of which is a single vertex with a self-loop.

so (74) will follow from

$$\mathrm{Det}\big(-\mathbf{A}_1(\mathbf{H}_{1,\mathbf{x}})^{-1}\mathbf{A}_1^{\intercal}\big) = 2\prod_{i\in[q]}\alpha_i \prod_{i\in[q-1]}\big(1+\mu_i\big). \tag{81}$$

To show (81), it can be checked that

$$-\mathbf{A}_1\mathbf{H}_{1,\mathbf{x}}^{-1}\mathbf{A}^{\intercal} = \boldsymbol{\alpha}^D + \mathbf{S}_{\mathbf{x}}, \tag{82}$$

where $\boldsymbol{\alpha}^D$ is the $q\times q$ diagonal matrix whose ith diagonal entry is α_i and $\mathbf{S}_{\mathbf{x}}$ is the $q\times q$ symmetric matrix whose (i,j) entry (when $i\le j$) is x^*_{ij} whenever $(i,j)\in P_1$ and 0 otherwise.

Observe that $(\boldsymbol{\alpha}^D)^{-1/2}\mathbf{S}_{\mathbf{x}}(\boldsymbol{\alpha}^D)^{-1/2} = \mathbf{M}$, where $\mathbf{M}$ is the matrix in Lemma 39. From this, we obtain

$$\mathrm{Det}(-\mathbf{A}_1\mathbf{H}_{1,\mathbf{x}}^{-1}\mathbf{A}_1^{\intercal}) = \Big(\prod_{i\in[q]}\alpha_i\Big)\,\mathrm{Det}\big(\mathbf{I}_q+\mathbf{M}\big).$$

Since the spectrum of $\mathbf{M}$ is $\{1,\mu_1,\dots,\mu_q\}$, it follows that the spectrum of the matrix $\mathbf{I}_q+\mathbf{M}$ is $\{2,1+\mu_1,\dots,1+\mu_{q-1}\}$. This yields (81), thus completing the proof of (74).

11.2.2. Proof of (75). For the second moment, $\mathbf{E}_{\mathcal{G}}\big[(Z_G^{\boldsymbol{\alpha}})^2\big]$ is a sum over $\boldsymbol{\gamma},\mathbf{y}$ while $\boldsymbol{\alpha}$ is fixed. Analogously to (77), let

$$P_2 = \big\{(i,k,j,l)\in[q]^4 \,\big|\, B_{ij}B_{kl}>0, i\le j, k\le l\big\}. \tag{83}$$

Henceforth, $\mathbf{y}$ will denote $\{y_{ikjl}\}_{(i,k,j,l)\in P_2}$. The constraints in (11) can now be written as

$$\mathbf{A}_2\,\big[\boldsymbol{\gamma},\,\mathbf{y}\big]^{\intercal} = \big[\boldsymbol{\alpha},\,\boldsymbol{\alpha},\,\mathbf{0}_{q^2}\big]^{\intercal},\ \text{where } \mathbf{A}_2 = \begin{bmatrix} \mathbf{A}_{2,\boldsymbol{\gamma}} & \mathbf{0} \\ -\mathbf{I}_{q^2} & \mathbf{A}_{2,\mathbf{y}} \end{bmatrix}, \tag{84}$$

and $\mathbf{A}_{2,\boldsymbol{\gamma}},\mathbf{A}_{2,\mathbf{y}}$ are $\{0,1\}$-matrices with dimensions $2q\times q^2$ and $q^2\times|P_2|$, respectively. It is easy to see that $\mathbf{A}_{2,\mathbf{y}}$ has full row rank $r_{\mathbf{y}}=q^2$, while $\mathbf{A}_{2,\boldsymbol{\gamma}}$ has rank $r_{\boldsymbol{\gamma}}=2q-1$, so that the rank of $\mathbf{A}_2$ is $r_2=r_{\mathbf{y}}+r_{\boldsymbol{\gamma}}=q^2+2q-1$. Thus, to specify a full-dimensional representation of (84), we need to specify an $r_2\times r_2$ invertible submatrix $\mathbf{A}_{2,s}$ of $\mathbf{A}_2$. It can be checked that any such submatrix $\mathbf{A}_{2,s}$ of $\mathbf{A}_2$ must have the form

$$\mathbf{A}_{2,s} = \begin{bmatrix} \mathbf{A}^s_{2,\boldsymbol{\gamma}} & \mathbf{0} \\ -\mathbf{I}_{q^2} & \mathbf{A}^s_{2,\mathbf{y}} \end{bmatrix},$$

where $\mathbf{A}^s_{2,\boldsymbol{\gamma}},\mathbf{A}^s_{2,\mathbf{y}}$ are $r_{\boldsymbol{\gamma}}\times r_{\boldsymbol{\gamma}}$ and $r_{\mathbf{y}}\times r_{\mathbf{y}}$ invertible submatrices of $\mathbf{A}^s_{2,\boldsymbol{\gamma}},\mathbf{A}^s_{2,\mathbf{y}}$, respectively. Thus, we only need to specify the matrices $\mathbf{A}^s_{2,\boldsymbol{\gamma}},\mathbf{A}^s_{2,\mathbf{y}}$. We will choose $\mathbf{A}^s_{2,\boldsymbol{\gamma}}$ to be an arbitrary invertible submatrix of $\mathbf{A}_{2,\boldsymbol{\gamma}}$; since $\mathbf{A}_{2,\boldsymbol{\gamma}}$ is totally unimodular (it corresponds to the incidence matrix of the complete bipartite graph with q vertices on each side), we have $\mathrm{Det}(\mathbf{A}^s_{2,\boldsymbol{\gamma}})^2=1$. For ferromagnetic models, we can choose $\mathbf{A}^s_{2,\mathbf{y}}$ to be the identity matrix using the columns corresponding to variables y_{ikik} with $i,k\in[q]$, so $\mathrm{Det}(\mathbf{A}^s_{2,\mathbf{y}})^2=1$. It follows that $\mathrm{Det}(\mathbf{A}_{2,s})^2=1$. For future use (with the scope of applying Lemma 42), let $\mathbf{A}_{2,f}$ be the submatrix of $\mathbf{A}_2$ induced by the rows corresponding to rows of $\mathbf{A}_{2,s}$.

We have the following analogue of Lemma 44. The proof is given in Appendix B.

LEMMA 45. *For a ferromagnetic model, it holds that*

$$\lim_{n\to\infty}\frac{(2\pi n)^{q-1}\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2]}{e^{n\Upsilon_2(\boldsymbol{\gamma},\mathbf{y}^*)}}=\Big(2^{q^2-1}\prod_{(i,k,j,l)\in P_2}y^*_{ikjl}\Big)^{-1/2}\Delta^{-(q-1)^2/2}\mathrm{Det}(-\mathbf{H}_2^f)^{-1/2}, \tag{85}$$

where $\mathbf{H}_2^f$ is the Hessian of the full-dimensional representation of $\Upsilon_2(\boldsymbol{\gamma},\mathbf{y})/\Delta$ evaluated at $(\boldsymbol{\gamma},\mathbf{y})=(\boldsymbol{\gamma}^*,\mathbf{y}^*)$.

Denote by $\mathbf{H}_2$ the diagonal matrix corresponding to the Hessian matrix of $\Upsilon_2(\boldsymbol{\gamma},\mathbf{y})/\Delta$ (when $\boldsymbol{\gamma},\mathbf{y}$ are unconstrained). Note that we may decompose $\mathbf{H}_2$ as $\left[\begin{smallmatrix}\mathbf{H}_{2,\boldsymbol{\gamma}} & \mathbf{0}\\ \mathbf{0} & \mathbf{H}_{2,\mathbf{y}}\end{smallmatrix}\right]$, where $\mathbf{H}_{2,\boldsymbol{\gamma}}$ is the $q^2\times q^2$ diagonal matrix corresponding to the Hessian matrix of $(\Delta-1)f_2(\boldsymbol{\gamma})/\Delta$ and $\mathbf{H}_{2,\mathbf{y}}$ is the $|P_2|\times|P_2|$ diagonal matrix corresponding to the Hessian matrix of $g_2(\mathbf{y})$ (see (14) for the specification of the functions f_2 and g_2).

We next apply Lemma 42 with the matrix (recall that $\boldsymbol{\alpha}^D$ is the $q\times q$ diagonal matrix whose ith diagonal entry is α_i and $\mathbf{0}_{q^2}$ is the $q^2\times q^2$ all-zeros matrix),

$$\mathbf{T}_2=\boldsymbol{\alpha}^D\oplus\boldsymbol{\alpha}^D\oplus\mathbf{0}_{q^2}, \tag{86}$$

to obtain the following equality:

$$\begin{aligned}\mathrm{Det}\big(-\mathbf{H}_2^f\big)&=\frac{L(\mathbf{A}_2,\mathbf{A}_{2,f},\mathbf{T}_2)}{\mathrm{Det}(\mathbf{A}_{2,s})^2}\mathrm{Det}(-\mathbf{H}_2)\,[\varepsilon]\,\mathrm{Det}\big(\varepsilon\mathbf{T}_2-\mathbf{A}_2\mathbf{H}_2^{-1}\mathbf{A}_2^{\intercal}\big)\\&=\frac{1}{2}\mathrm{Det}(-\mathbf{H}_2)\,[\varepsilon]\,\mathrm{Det}\big(\varepsilon\mathbf{T}_2-\mathbf{A}_2\mathbf{H}_2^{-1}\mathbf{A}_2^{\intercal}\big),\end{aligned}$$

where in the latter equality we used that $\mathrm{Det}(\mathbf{A}_{2,s})^2=1$ (which was proved earlier) and $L(\mathbf{A}_2,\mathbf{A}_{2,f},\mathbf{T}_2)=1/2$ (follows by [17, Proof of Lemma B.8]). We calculate

$$\begin{aligned}\mathrm{Det}(-\mathbf{H}_2)^{-1}&=\mathrm{Det}(-\mathbf{H}_{2,\boldsymbol{\gamma}})^{-1}\,\mathrm{Det}(-\mathbf{H}_{2,\mathbf{y}})^{-1}\\&=\Big[(-1)^{q^2}\Big(\frac{\Delta}{\Delta-1}\Big)^{q^2}\prod_{i,k\in[q]}\gamma^*_{ik}\Big]\Big[2^{q^2}\prod_{(i,k,j,l)\in P_2}y^*_{ikjl}\Big],\end{aligned}$$

so that (75) will follow from

$$\begin{aligned}&\frac{[\varepsilon]\mathrm{Det}\big(\varepsilon\mathbf{T}_2-\mathbf{A}_2(\mathbf{H}_2)^{-1}\mathbf{A}_2^{\intercal}\big)}{\prod_{i\in[q]}\alpha_i^2\prod_{i,k\in[q]}\gamma^*_{ik}}\\&\qquad=(-1)^{q^2}\frac{4\Delta^{2q-1}}{(\Delta-1)^{q^2}}\prod_{i\in[q-1]}\big(1+\mu_i\big)^2\prod_{i,j\in[q-1]}\big(1-(\Delta-1)\mu_i\mu_j\big).\end{aligned} \tag{87}$$

We first write out the block structure of $\varepsilon\mathbf{T}_2-\mathbf{A}_2(\mathbf{H}_2)^{-1}\mathbf{A}_2^{\intercal}$. First, we have the following analogue of (82):

$$-\mathbf{A}_{2,\mathbf{y}}\mathbf{H}_{2,\mathbf{y}}^{-1}\mathbf{A}_{2,\mathbf{y}}^{\intercal}=\boldsymbol{\gamma}^D+\mathbf{S}_{\mathbf{y}},\quad \mathbf{A}_{2,\boldsymbol{\gamma}}\boldsymbol{\gamma}^D\mathbf{A}_{2,\boldsymbol{\gamma}}^{\intercal}=\begin{bmatrix}\boldsymbol{\alpha}^D & \mathbf{S}_{\boldsymbol{\gamma}}\\ \mathbf{S}_{\boldsymbol{\gamma}} & \boldsymbol{\alpha}^D\end{bmatrix}, \tag{88}$$

where $\mathbf{S}_{\boldsymbol{\gamma}}$ is the $q \times q$ matrix whose (i,j) entry is γ^*_{ij} and $\mathbf{S}_{\mathbf{y}}$ is the $q^2 \times q^2$ matrix whose $((i,k),(j,l))$ entry is y^*_{ikjl}. From

$$\begin{aligned}
&\varepsilon\mathbf{T}_2 - \mathbf{A}_2(\mathbf{H}_2)^{-1}\mathbf{A}_2^{\intercal}\\
&= \varepsilon\mathbf{T}_2 + \begin{bmatrix} \mathbf{A}_{2,\boldsymbol{\gamma}} & \mathbf{0} \\ -\mathbf{I}_{q^2} & \mathbf{A}_{2,\mathbf{y}} \end{bmatrix}\begin{bmatrix} -\frac{\Delta}{\Delta-1}\boldsymbol{\gamma}^D & \mathbf{0} \\ \mathbf{0} & -\mathbf{H}_{2,\mathbf{y}}^{-1} \end{bmatrix}\begin{bmatrix} \mathbf{A}_{2,\boldsymbol{\gamma}}^{\intercal} & -\mathbf{I}_{q^2} \\ \mathbf{0} & \mathbf{A}_{2,\mathbf{y}}^{\intercal} \end{bmatrix}\\
&= \frac{\Delta}{\Delta-1}\begin{bmatrix} \varepsilon\frac{\Delta-1}{\Delta}(\boldsymbol{\alpha}^D \oplus \boldsymbol{\alpha}^D) - \mathbf{A}_{2,\boldsymbol{\gamma}}\boldsymbol{\gamma}^D\mathbf{A}_{2,\boldsymbol{\gamma}}^{\intercal} & \mathbf{A}_{2,\mathbf{y}}\boldsymbol{\gamma}^D \\ \boldsymbol{\gamma}^D\mathbf{A}_{2,\mathbf{y}}^{\intercal} & -\boldsymbol{\gamma}^D - \frac{\Delta-1}{\Delta}\mathbf{A}_{2,\mathbf{y}}\mathbf{H}_{2,\mathbf{y}}^{-1}\mathbf{A}_{2,\mathbf{y}}^{\intercal} \end{bmatrix}
\end{aligned}$$

we obtain that

$$\frac{[\varepsilon]\mathrm{Det}\big(\varepsilon\mathbf{T}_2 - \mathbf{A}_2(\mathbf{H}_2)^{-1}\mathbf{A}_2^{\intercal}\big)}{\prod_{i\in[q]}\alpha_i^2\prod_{i,k\in[q]}\gamma^*_{ik}} = [\varepsilon]\mathrm{Det}\big(\mathbf{H}_2'\big), \tag{89}$$

where

$$\mathbf{H}_2' = \frac{\Delta}{\Delta-1}\begin{bmatrix} \varepsilon\frac{\Delta-1}{\Delta}\mathbf{I}_{2q} - \mathbf{V}\mathbf{V}^{\intercal} & \mathbf{V} \\ \mathbf{V}^{\intercal} & -\frac{\Delta-1}{\Delta}\mathbf{W} \end{bmatrix},$$

and the matrices $\mathbf{W}, \mathbf{V}$ are given by (recall that $\mathbf{M}$ is the matrix in Lemma 39 whose eigenvalues are $1, \mu_1, \ldots, \mu_{q-1}$)

$$\mathbf{W} = \frac{1}{\Delta-1}\mathbf{I}_{q^2} - \mathbf{M}\otimes\mathbf{M}, \quad \mathbf{V} = (\boldsymbol{\alpha}^D \oplus \boldsymbol{\alpha}^D)^{-1/2}\mathbf{A}_{2,\boldsymbol{\gamma}}\,(\boldsymbol{\gamma}^D)^{1/2}.$$

In light of (89), it suffices to compute $\mathrm{Det}(\mathbf{H}_2')$. To do this, we proceed by taking the Schur complement of the matrix $\mathbf{W}$. It is easy to see that $\mathbf{W}$ is invertible, since its spectrum is given by

$$\begin{gathered} t-1, t-\mu_1, t-\mu_1, \ldots, t-\mu_{q-1}, t-\mu_{q-1},\\ t-\mu_1^2, t-\mu_1\mu_2, \ldots, t-\mu_1\mu_{q-1}, t-\mu_2\mu_1, \ldots, t-\mu_{q-1}^2, \end{gathered}$$

where $t := 1/(\Delta-1)$. We also have

$$\mathrm{Det}(\mathbf{W}) = -\frac{(\Delta-2)}{(\Delta-1)^{q^2}}\prod_{i\in[q-1]}\big(1-(\Delta-1)\mu_i\big)^2\prod_{i,j\in[q-1]}\big(1-(\Delta-1)\mu_i\mu_j\big). \tag{90}$$

Considering the Schur complement of the matrix $\mathbf{W}$, we obtain

$$\begin{gathered} \mathrm{Det}(\mathbf{H}_2') = (-1)^{q^2}\Big(\frac{\Delta}{\Delta-1}\Big)^{2q}\mathrm{Det}\big(\mathbf{W}\big)\,\mathrm{Det}\Big(\varepsilon\frac{\Delta-1}{\Delta}\mathbf{I}_{2q} + \mathbf{Z}\Big)\\ \text{with } \mathbf{Z} = \frac{\Delta}{\Delta-1}\mathbf{V}\mathbf{W}^{-1}\mathbf{V}^{\intercal} - \mathbf{V}\mathbf{V}^{\intercal}. \end{gathered} \tag{91}$$

To compute $\mathrm{Det}\big(\varepsilon\frac{\Delta-1}{\Delta}\mathbf{I}_{2q} + \mathbf{Z}\big)$, we need to obtain a simpler form for $\mathbf{Z}$. The following lemma, which is proved at the end of this section, will allow for such a simplication.

Lemma 46. *It holds that* $\mathbf{V}\mathbf{W} = \big(\frac{1}{\Delta-1}\mathbf{I}_{2q} - \mathbf{M}'\big)\mathbf{V}$, *where* $\mathbf{M}' := \mathbf{M}\oplus\mathbf{M}$.

It is standard to express the eigenvalues of $\frac{1}{\Delta-1}\mathbf{I}_{2q}-\mathbf{M}'$ in terms of the eigenvalues of $\mathbf{M}$ and hence obtain that the former matrix is invertible (since the eigenvalues of $\mathbf{M}$ other than 1 are less in absolute value than $1/(\Delta-1)$). Thus, Lemma 46 gives

$$\begin{aligned}\mathbf{Z} &= \Big[-\mathbf{I}_{2q}+\frac{\Delta}{\Delta-1}\Big(\frac{1}{\Delta-1}\mathbf{I}_{2q}-\mathbf{M}'\Big)^{-1}\Big]\mathbf{V}\mathbf{V}^{\intercal}\\ &= \big(\mathbf{I}_{2q}+\mathbf{M}'\big)\big(\frac{1}{\Delta-1}\mathbf{I}_{2q}-\mathbf{M}'\big)^{-1}\mathbf{V}\mathbf{V}^{\intercal}. \end{aligned} \tag{92}$$

By (91), $\mathbf{Z}$ is trivially symmetric. Using (92), we obtain the eigenvalues of $\mathbf{Z}$.

LEMMA 47. *The spectrum of* $\mathbf{Z}$ *is given by*

$$0, 2f(1), f(\mu_1), f(\mu_1), f(\mu_2), f(\mu_2), \dots, f(\mu_{q-1}), f(\mu_{q-1}),$$

where $f(x)=(1+x)(\frac{1}{\Delta-1}-x)^{-1}$.

Proof of Lemma 47. Let $\mathbf{u}_1=[\sqrt{\boldsymbol{\alpha}},\sqrt{\boldsymbol{\alpha}}]^{\intercal}$, $\mathbf{u}_2=[\sqrt{\boldsymbol{\alpha}},-\sqrt{\boldsymbol{\alpha}}]^{\intercal}$. Note that $\mathbf{u}_1,\mathbf{u}_2$ are linearly independent eigenvectors of $\mathbf{M}'$ corresponding to the eigenvalue 1.

Using (88), we have that $\mathbf{V}\mathbf{V}^{\intercal}=\left[\begin{smallmatrix}\mathbf{I}_q & \mathbf{S}'_{\boldsymbol{\gamma}}\\ \mathbf{S}'_{\boldsymbol{\gamma}} & \mathbf{I}_q\end{smallmatrix}\right]$, where $\mathbf{S}'_{\boldsymbol{\gamma}}$ is the $q\times q$ matrix whose (i,j) entry is $\sqrt{\alpha_i\alpha_j}$. It follows that $\mathbf{u}_1$ and $\mathbf{u}_2$ are eigenvectors of $\mathbf{V}\mathbf{V}^{\intercal}$ with eigenvalues 2 and 0, respectively, and hence $\mathbf{u}_1$ and $\mathbf{u}_2$ are eigenvectors of $\mathbf{Z}$ with eigenvalues $2f(1)$ and 0, respectively.

Let $\mathbf{u}$ be an eigenvector of $\mathbf{M}'$ corresponding to an eigenvalue $\mu\neq 1$. Note that $\mathbf{u}$ is perpendicular to both $\mathbf{u}_1$ and $\mathbf{u}_2$. It follows that $\mathbf{V}\mathbf{V}^{\intercal}\mathbf{u}=\mathbf{u}$, so that $\mathbf{Z}\mathbf{u}=f(\mu)\mathbf{u}$. Thus, $\mathbf{u}$ is also an eigenvector of $\mathbf{Z}$ with eigenvalue $f(\mu)$. ∎

To simplify the expressions, set $r=(\Delta-1)/\Delta$. The matrix $\varepsilon r\mathbf{I}_{2q}$ shifts the eigenvalues of $\mathbf{Z}$ by εr. Thus, Lemma 47 yields

$$\mathrm{Det}\big(\varepsilon r\mathbf{I}_{2q}+\mathbf{Z}\big)=\varepsilon r\big(\varepsilon r+2f(1)\big)\prod_{i\in[q-1]}\big(\varepsilon r+f(\mu_i)\big)^2.$$

We have $f(1), f(\mu_i)\neq 0$ for every $i\in[q-1]$, so that

$$\begin{aligned}[\varepsilon]\mathrm{Det}\big(\varepsilon r\mathbf{I}_{2q}+\mathbf{Z}\big) &= 2rf(1)\prod_{i\in[q-1]}\big(f(\mu_i)\big)^2\\ &= -\frac{4(\Delta-1)^{2q}}{\Delta(\Delta-2)}\prod_{i\in[q-1]}\left(\frac{1+\mu_i}{1-(\Delta-1)\mu_i}\right)^2.\end{aligned} \tag{93}$$

Plugging (90) and (93) into (91), we obtain

$$[\varepsilon]\mathrm{Det}(\mathbf{H}_2')=(-1)^{q^2}\frac{4\Delta^{2q-1}}{(\Delta-1)^{q^2}}\prod_{i,j\in[q-1]}\big(1-(\Delta-1)\mu_i\mu_j\big)\prod_{i\in[q-1]}\big(1+\mu_i\big)^2.$$

Using this and (89), we obtain (87) as wanted. We conclude by giving the deferred proof of Lemma 46.

Proof of Lemma 46. For notational convenience, let $\mathbf{N}:=\mathbf{M}\otimes\mathbf{M}$. The lemma clearly reduces to proving $\mathbf{V}\,\mathbf{N}=\mathbf{M}'\,\mathbf{V}$. Let $\mathbf{D}:=\mathbf{V}\,\mathbf{N}$, $\mathbf{E}:=\mathbf{M}'\,\mathbf{V}$.

The matrices $\mathbf{D},\mathbf{E}$ clearly have the same dimensions, since $\mathbf{V}$ has dimensions $2q\times q^2$, $\mathbf{N}$ has dimensions $q^2\times q^2$, and $\mathbf{M}'$ has dimensions $2q\times 2q$. It remains to check that the entries of $\mathbf{D}, \mathbf{E}$ are equal. First, we give explicit expressions for the entries of $\mathbf{V}, \mathbf{N}$. We have

$$V_{t,(i,k)} = \begin{cases} \sqrt{\frac{\gamma^*_{ik}}{\alpha_i}}\mathbf{1}\{i=t\}, & 1 \le t \le q, \\ \sqrt{\frac{\gamma^*_{ik}}{\alpha_k}}\mathbf{1}\{k=t-q\}, & q+1 \le t \le 2q, \end{cases} \qquad N_{(i,k),(j,l)} = \frac{x^*_{ij}x^*_{kl}}{\sqrt{\gamma^*_{ik}\gamma^*_{jl}}}.$$

We next consider the $\big(i,(j,l)\big)$ entries of the matrices $\mathbf{D}, \mathbf{E}$. Assume first that $i \le q$. We have

$$\begin{aligned} D_{i,(j,l)} &= \sum_{i',k} V_{i,(i',k)} N_{(i',k),(j,l)} = \sum_{i',k} \sqrt{\frac{\gamma^*_{i'k}}{\alpha_{i'}}}\mathbf{1}\{i'=i\} \frac{x^*_{i'j}x^*_{kl}}{\sqrt{\gamma^*_{i'k}\gamma^*_{jl}}} \\ &= \frac{x^*_{ij}}{\sqrt{\alpha_i}\sqrt{\gamma^*_{jl}}} \sum_k x^*_{kl} = \frac{\alpha_l x^*_{ij}}{\sqrt{\alpha_i}\sqrt{\gamma^*_{jl}}}, \\ E_{i,(j,l)} &= \sum_{j'} M'_{i,j'} V_{j',(j,l)} = \sum_{j'} M_{i,j'} \sqrt{\frac{\gamma^*_{jl}}{\alpha_j}}\mathbf{1}\{j=j'\} = M_{i,j}\sqrt{\frac{\gamma^*_{jl}}{\alpha_j}} \\ &= \frac{x^*_{ij}\sqrt{\gamma^*_{jl}}}{\alpha_j\sqrt{\alpha_i}} = \frac{\alpha_l x^*_{ij}}{\sqrt{\alpha_i}\sqrt{\gamma^*_{jl}}}. \end{aligned}$$

An analogous calculation for $q < i \le 2q$ yields that $D_{i,(j,l)} = E_{i,(j,l)}$ for every i, j, l. ∎

11.3. Bethe prediction for general models on random regular graphs. In this section, we show how to extend Theorem 8 for general models on random regular graphs as discussed in section 3.1. A more general result has been derived in [13, Theorem 1.16] for sequences of graphs converging locally to (random) trees, under the assumption of uniqueness of the Gibbs measure on the underlying tree. For the special case of random Δ-regular graphs, we show how to extend Theorem 8 when there is a unique *semitranslation invariant* Gibbs measure. Our proof has a different perspective and yields a slightly simpler condition for random Δ-regular graphs.

Semitranslation invariant Gibbs measures on $\mathbb{T}_\Delta$ are Gibbs measures that are invariant under any parity-preserving automorphisms of $\mathbb{T}_\Delta$ (cf. [6]). They can be specified by a pair of probability vectors $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ for the even and odd, respectively, vertices. Note that if there is a unique semitranslation invariant measure, then this measure is also translation invariant. Hence, it corresponds to a fixpoint of the tree recursions (6).

THEOREM 48. *Let* $\mathbf{B}$ *be a regular matrix which specifies a model such that for all* Δ*-regular graphs* $Z_G > 0$*. If there is a unique semitranslation invariant Gibbs measure on* $\mathbb{T}_\Delta$ *and the corresponding fixpoint is Jacobian attractive, then*

$$\lim_{n\to\infty} \frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G] = \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[Z_G],$$

where $\mathcal{G}$ *is the uniform distribution on* Δ*-regular graphs with* n *vertices.*

The first assumption in the theorem is mainly to avoid pathological cases where $\log Z_G \equiv -\infty$ in which case the quantities are not well-defined. It is satisfied by many classes of models, e.g., permissive models [13] such as the hard-core and antiferromagnetic Potts model, or even nonpermissive such as q-colorings when $q \ge \Delta + 1$.

The proof of Theorem 48 is analogous to that of Theorem 8, once we establish the analogue of Theorem 7 for general models. As we illustrated in Remark 13, this is hopeless to achieve in general and we must thus use the uniqueness assumption that Theorem 8 requires. Note that if there is a unique semitranslation invariant measure (which is the assumption in Theorem 8), then this measure is also translation invariant.

Proof of Theorem 48. Let $\boldsymbol{\alpha}^*$ be a dominant phase. By semitranslational uniqueness we have that $\boldsymbol{\alpha}^*$ is unique. We next describe how to obtain the analogue of Theorem 7 under the assumptions of Theorem 48. Let $p = \Delta/(\Delta-1)$. We show that whenever there is a unique semitranslation Gibbs measure on $\mathbb{T}_\Delta$, it holds that $\exp(2\Psi_1(\boldsymbol{\alpha}^*)/\Delta) = \|\mathbf{B}\|_{p\to\Delta}$.

From (16) and (21), we obtain

$$\exp(2\Psi_1(\boldsymbol{\alpha}^*)/\Delta) = \max_{\boldsymbol{\alpha}} \exp(2\Psi_1(\boldsymbol{\alpha})/\Delta) = \max_{\mathbf{R}} \frac{\mathbf{R}^\intercal\mathbf{B}\mathbf{R}}{\|\mathbf{R}\|_p^2} \le \max_{\mathbf{R},\mathbf{C}} \frac{\mathbf{R}^\intercal\mathbf{B}\mathbf{C}}{\|\mathbf{R}\|_p\|\mathbf{C}\|_p}. \tag{94}$$

Note that the last inequality is trivial; we just enlarged the maximization region we consider. It is proved in [17] that the maximum of the r.h.s. is achieved at a semitranslation invariant fixpoint. If there is a unique semitranslation invariant Gibbs measure on $\mathbb{T}_\Delta$, this must be translation invariant and hence the maximum in the r.h.s. of (94) must occur at $\mathbf{R}=\mathbf{C}$. We thus obtain that (94) is satisfied at equality. The r.h.s. in (94) is equal to $\|\mathbf{B}\|_{p\to\Delta}$ [17, section 3.1], proving the desired claim.

By the same token, one has the bound

$$\exp(2\Psi_2(\boldsymbol{\alpha}^*)/\Delta) = \max_{\boldsymbol{\alpha}} \exp(2\Psi_2(\boldsymbol{\alpha})/\Delta) \le \|\mathbf{B}\otimes\mathbf{B}\|_{p\to\Delta},$$

and since $\|\mathbf{B}\otimes\mathbf{B}\|_{p\to\Delta} = \|\mathbf{B}\|^2_{p\to\Delta}$ [2, Proposition 10.3], we obtain that $\Psi_2(\boldsymbol{\alpha}^*) = 2\Psi_1(\boldsymbol{\alpha}^*)$, as wanted.

Since the dominant phase $\boldsymbol{\alpha}^*$ corresponds to a Jacobian attractive fixpoint (by assumption), it is also Hessian dominant (see Remark 15). With minor modifications (see footnotes 7 and 8), the results of section 11.1 can be adapted to obtain a lower bound on $Z_G^{\boldsymbol{\alpha}}$ as in Lemma 16. Thus the proof of Theorem 8 in section 7.1 extends to the present setting as well. ∎

Appendix A. Nonreconstruction for the ordered phases on the tree. In this appendix, we give in detail the proof of the doubly exponential upper bound in (56). This appendix is organized as follows. In Appendix A.1, we review broadcasting processes on trees, the nonreconstruction property, and a concentration result from [46]. In Appendix A.2, we review relevant connections between broadcasting processes and Gibbs measures defined by fixpoints of the tree recursions (6), which will allow us to apply the result of [46]. Finally, in Appendix A.3, we apply these results to the ferromagnetic Potts model and obtain the bound in (56).

Let us fix some notation that will be used throughout this section. We will denote by $T=(V,E)$ the infinite $(\Delta-1)$-ary tree. The root of T will be denoted by ρ. Also, for an integer $\ell \ge 0$, T_ℓ will denote the subtree of T consisting of the first ℓ levels of T and W_ℓ will denote the set of the leaves of T_ℓ. Further, for a configuration $\sigma : V \to [q]$, we denote by σ_{W_ℓ} the restriction of σ on W_ℓ.

A.1. Nonreconstruction in broadcasting processes on trees. Let $q \ge 2$ be an integer and $\mathbf{M} = (M_{ij})_{i,j\in[q]}$ be a $q\times q$ stochastic matrix (i.e., the entries are nonnegative and the entries in each row have sum equal to 1). We will further assume that $\mathbf{M}$ is irreducible and aperiodic, so that there exists a unique q-dimensional probability vector $\boldsymbol{\pi}^* = (\pi_i^*)_{i\in[q]}$ so that $\boldsymbol{\pi}^*\mathbf{M} = \boldsymbol{\pi}^*$. Note that the entries of $\boldsymbol{\pi}^*$ are all positive. We will refer to $\boldsymbol{\pi}^*$ as the stationary distribution of $\mathbf{M}$. We will also assume that $\mathbf{M}$ is reversible with respect to $\boldsymbol{\pi}^*$, i.e., $\pi_i^* M_{ij} = \pi_j^* M_{ji}$ for all $i,j \in [q]$ (every such matrix is similar to a symmetric matrix and thus has real eigenvalues).

Let $\boldsymbol{\pi} = (\pi_i)_{i\in[q]}$ be a q-dimensional probability vector with positive entries (note that it may hold that $\boldsymbol{\pi} \neq \boldsymbol{\pi}^*$). The broadcasting process $\mathbf{M}$ on the tree $T = (V, E)$ with root ρ is a probability distribution ν on the set of assignments $\sigma : V \to [q]$ such that

$$\nu(\sigma) = \pi_{\sigma(\rho)} \prod_{(u,v)\in E} M_{\sigma(u),\sigma(v)}. \tag{95}$$

To generate σ with distribution ν, first pick randomly the spin of the root from the distribution $\boldsymbol{\pi}$ and then broadcast the spin down the tree, where each edge of the tree acts as a noisy channel. In particular, for an edge (u, v) of the tree where u is the parent of v, conditioned on the spin $\sigma(u)$, the spin $\sigma(v)$ is picked randomly from the distribution $(M_{\sigma(u),1}, \ldots, M_{\sigma(u),q})$.

We next define the nonreconstruction property, which roughly captures whether, as we go deeper into the tree, the information about the spin of the root vanishes. (For distributions μ_1, μ_2 defined on the same space Ω, we denote by $d_{TV}(\mu_1, \mu_2)$ the total variation distance between μ_1, μ_2.)

Definition 49 (nonreconstruction). *A broadcasting process* $\mathbf{M}$ *has the nonreconstruction property on the tree* T *if*

$$\lim_{\ell\to\infty} \max_{i,j\in[q]} d_{TV}\big(\nu(\sigma_{W_\ell} = \cdot \mid \sigma_\rho = i), \nu(\sigma_{W_\ell} = \cdot \mid \sigma_\rho = j)\big) = 0. \tag{96}$$

Nonreconstruction is often closely connected to the second largest eigenvalue of $\mathbf{M}$. We will use the following concentration result of [46], which can be interpreted as quantifying the rate of convergence to 0, when the second largest eigenvalue of $\mathbf{M}$ is small with respect to the branching factor of the tree.

Theorem 50 (see [46, Theorem 2.3]; see also [34, 35]). *Consider a broadcasting process* $\mathbf{M}$ *on the infinite* $(\Delta - 1)$*-ary tree with no hard constraints (i.e., all entries of* $\mathbf{M}$ *are positive), whose spin at the root is chosen according to some distribution* $\boldsymbol{\pi}$ *with positive entries. Let* λ *be the second largest eigenvalue of* $\mathbf{M}$ *in absolute value. Then, if* $\mathbf{M}$ *has nonreconstruction and* $(\Delta - 1)\lambda^2 < 1$*, there exist constants* $C > 0$ *and* $\ell_0 \geq 1$ *such that the following holds.*

Let $\mathcal{B}_\ell := \big\{\eta : W_\ell \to [q] \,\big|\, \big\|\nu(\sigma_\rho = \cdot \mid \sigma_{W_\ell} = \eta) - \boldsymbol{\pi}\big\|_\infty \geq \exp(-C\ell)\big\}$*. Then, for all* $\ell \geq \ell_0$*,*

$$\nu(\sigma_{W_\ell} \in \mathcal{B}_\ell) \leq \exp(-\exp(C\ell)).$$

We remark here that the restriction in Theorem 50 that $\mathbf{M}$ has no hard constraints is not needed and, in fact, in [46], the analogous statement is proved for general models $\mathbf{M}$ whose state space satisfies a general connectivity condition. Since we will only apply the result of [46] to the ferromagnetic Potts model (which has no hard constraints), such connectivity issues are not present in our setting and thus are out of our scope. In particular, in the language/notation of [46], all colors c, c' will be trivially compatible in our setting and thus the measure $\mu^c(\cdot)$ in [46, Theorem 2.3], which conditions the broadcasting process in the space of configurations where the "parent of the root" has color c, is identical to the unconditioned broadcasting process (denoted by ν in our setting). Further, [46, Theorem 2.3] is stated for the case where $\boldsymbol{\pi} = \boldsymbol{\pi}^*$, i.e., when the distribution of the spin of the root ρ is chosen according to the stationary distribution of $\mathbf{M}$. We next display how to derive from this the slightly more general version stated in Theorem 50.

In particular, suppose that Theorem 50 is true for some distribution $\boldsymbol{\pi}$. Our goal is to show that it also holds for some other distribution $\boldsymbol{\pi}'$ (we assume that both $\boldsymbol{\pi}$ and $\boldsymbol{\pi}'$ have positive entries) for some constants $C', \ell_0' > 0$. We will denote by ν the broadcasting process when the initial distribution is $\boldsymbol{\pi}$ and by ν' when the initial distribution is $\boldsymbol{\pi}'$. We will also use $\mathcal{B}_\ell$ and $\mathcal{B}_\ell'$ to denote the set of "bad" configurations on W_ℓ for the two processes ν, ν', respectively (see Theorem 50). Let $\eta \notin \mathcal{B}_\ell$ and, for $i \in [q]$, set $z_i(\eta) := \nu(\sigma_\rho = i \mid \sigma_{W_\ell} = \eta)/\nu(\sigma_\rho = i)$. Denote also by $z_i'(\eta)$ the respective quantity for the measure ν'. Since both processes have the same broadcasting matrix, observe that for any colors $i, j \in [q]$ it holds that

$$\frac{z_i(\eta)}{z_j(\eta)} = \frac{\nu(\sigma_{W_\ell} = \eta \mid \sigma_\rho = i)}{\nu(\sigma_{W_\ell} = \eta \mid \sigma_\rho = j)} = \frac{\nu'(\sigma_{W_\ell} = \eta \mid \sigma_\rho = i)}{\nu'(\sigma_{W_\ell} = \eta \mid \sigma_\rho = j)} = \frac{z_i'(\eta)}{z_j'(\eta)}.$$

Since $\eta \notin \mathcal{B}_\ell$, the ratio $z_i(\eta)/z_j(\eta)$ is bounded by $1 \pm O(\exp(-C\ell))$ and thus the same is true for $z_i'(\eta)/z_j'(\eta)$. This gives that $\eta \notin \mathcal{B}_\ell'$ (for any constant $0 < C' < C$ and sufficiently large ℓ_0'), i.e., $\mathcal{B}_\ell' \subseteq \mathcal{B}_\ell$. To obtain that $\nu'(\sigma_{W_\ell} \in \mathcal{B}_\ell') \leq \exp(-\exp(C'\ell))$ for all sufficiently large ℓ, observe that for any $\eta : W_\ell \to [q]$ and $i \in [q]$ it holds that $\nu(\sigma_{W_\ell} = \eta \mid \sigma_\rho = i) = \nu'(\sigma_{W_\ell} = \eta \mid \sigma_\rho = i)$, so that

$$\frac{\nu'(\sigma_{W_\ell} = \eta)}{\nu(\sigma_{W_\ell} = \eta)} = \frac{\sum_{i \in [q]} \pi_i' \nu'(\sigma_{W_\ell} = \eta \mid \sigma_\rho = i)}{\sum_{i \in [q]} \pi_i \nu(\sigma_{W_\ell} = \eta \mid \sigma_\rho = i)} \leq \max_{i \in [q]} \frac{\pi_i'}{\pi_i}.$$

Thus the desired bound on $\nu'(\sigma_{W_\ell} \in \mathcal{B}_\ell')$ follows from the bound on $\nu(\sigma_{W_\ell} \in \mathcal{B}_\ell)$.

A.2. Broadcasting processes and fixpoints of the tree recursions. In light of Theorem 50, our strategy for proving the bound in (56) will be to show that the measure ν^i (corresponding to the ith ordered phase in the Potts model) corresponds to a broadcasting process on the $(\Delta-1)$-ary tree (and then simply verify the assumptions of the theorem). The purpose of this section is to make this correspondence explicit. In fact, we will work out the relevant connections for general spin models.

Let $\mathbf{B}$ be the interaction matrix of a q-spin system. As in section 1.2, we assume that $\mathbf{B}$ is symmetric, irreducible, and aperiodic. For an integer $\Delta \geq 3$, recall that a fixpoint of the tree recursions is a vector $\mathbf{R} = (R_1, \dots, R_q)$ with positive entries such that

$$R_i \propto \left(\sum_j B_{ij} R_j \right)^{\Delta - 1} \forall i \in [q]. \tag{6}$$

For the purpose of this section, we assume that the normalization in (6) is such that $\sum_i R_i = 1$, i.e., $\mathbf{R}$ is a q-dimensional probability vector.

We next define the broadcasting process corresponding to the fixpoint $\mathbf{R}$ by first specifying an appropriate broadcasting matrix. In particular, let $\mathbf{M}$ be the $q \times q$ matrix whose (i,j)-entry is given by

$$M_{ij} = \frac{B_{ij} R_j}{\sum_{j'} B_{ij'} R_{j'}} \text{ for } i, j \in [q]. \tag{97}$$

We remark here that the normalization of the R_i's in (6) is not important for defining the matrix $\mathbf{M}$ (the entries remain unchanged if we scale the R_i's); we normalize

$\mathbf{R}$ to be a probability vector so that we can use it as the initial distribution $\boldsymbol{\pi}$ of the spin of the root in the broadcasting process. In particular, in the notation of Appendix A.1, we will set $\boldsymbol{\pi} = \mathbf{R}$. This completes the specification of the broadcasting process (cf. (95)). Note that $\mathbf{M}$ is stochastic, irreducible, and aperiodic. Further, its stationary distribution $\boldsymbol{\pi}^*$ is given by the probability vector whose entries satisfy $\pi_i^* \propto R_i \sum_j B_{ij} R_j$ for all $i \in [q]$. Finally, we have that $\mathbf{M}$ is reversible with respect to $\boldsymbol{\pi}^*$.

In the rest of this section, we state several results that eventually will allow us to apply Theorem 50. First, we connect the spectral properties of $\mathbf{M}$ with the attractiveness of the fixpoint $\mathbf{R}$ of the tree recursions (see section 3.2 for the relevant definitions).

LEMMA 51. *Let* $\mathbf{R}$ *be a Jacobian attractive fixpoint of the tree recursions and let* $\mathbf{M}$ *be the broadcasting matrix corresponding to* $\mathbf{R}$*. Let* λ *be the second largest eigenvalue of* $\mathbf{M}$ *in absolute value. Then* $(\Delta - 1)\lambda < 1$*.*

Proof. Recall from section 6 (see also the beginning of section 8) that $\mathbf{R}$ is a Jacobian attractive fixpoint of the tree recursions if every eigenvalue $x \neq 1$ of the matrix

$$\widetilde{\mathbf{M}} = \left\{ \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} \right\}_{i,j=1}^{q} \text{ with } \alpha_i = R_i \sum_j B_{ij} R_j \text{ for } i \in [q]$$

satisfies $(\Delta - 1)|x| < 1$. The result will thus follow by showing that the eigenvalues of $\mathbf{M}$ are identical to those of $\widetilde{\mathbf{M}}$.

We will show that $\mathbf{M}$ and $\widetilde{\mathbf{M}}$ are similar matrices, thus showing the result. Let $\mathbf{A}$ be the diagonal matrix whose ith diagonal entry is given by $\sqrt{\alpha_i}$. Note that $\mathbf{A}$ is invertible (since the R_i's are positive). By a direct calculation, it also holds that $\mathbf{A}\mathbf{M}\mathbf{A}^{-1} = \widetilde{\mathbf{M}}$, thus proving that $\mathbf{M}$ and $\widetilde{\mathbf{M}}$ are similar. This concludes the proof. $\square$

We now focus on connecting the Gibbs distribution of the spin model with interaction matrix $\mathbf{B}$ and the broadcasting process $\mathbf{M}$. As before, let T be the infinite $(\Delta - 1)$-ary tree with root ρ and denote by T_ℓ the subtree of T consisting of the first ℓ levels of T and by W_ℓ the set of leaves of T_ℓ. We will denote by μ_ℓ the Gibbs distribution on T_ℓ corresponding to the spin system with interaction matrix $\mathbf{B}$. We will use σ to denote configurations on T_ℓ and by σ_{W_ℓ} the restriction of σ to the leaves W_ℓ.

To connect μ_ℓ to the broadcasting process $\mathbf{M}$ on T, we will need just a few more definitions. Let $Q_{W_\ell}(\cdot)$ be the following product distribution on configurations on the leaves W_ℓ. For a configuration $\eta : W_\ell \to [q]$,

$$Q_{W_\ell}(\eta) := \prod_{i \in [q]} (R_i)^{|\eta^{-1}(i) \cap W_\ell|}. \tag{98}$$

Finally, consider the following distribution $\widehat{\nu}_\ell$, which is also defined on configurations on the leaves W_ℓ, given by

$$\widehat{\nu}_\ell(\eta) \propto \mu_\ell(\sigma_{W_\ell} = \eta)\, Q_{W_\ell}(\eta) \quad \forall \eta : W_\ell \to [q]. \tag{99}$$

It is instructive at this point to spell out the interplay of these definitions with the bound in (56). Namely, the product distribution $Q_{W_\ell}(\eta)$ is the generalization of the product distribution $Q^i_W(\eta)$ (defined just after (41)) and $\widehat{\nu}_\ell(\cdot)$ is the generalization of the distribution $\nu^i(\cdot)$ (defined in (54)).

We are now ready to state the desired connection.

LEMMA 52. *Let* $\mathbf{B},\mathbf{R},\mathbf{M}$ *be as above. Let* ν *denote the broadcasting measure* $\mathbf{M}$ *on* T *(defined in* (95)*) and, for integer* $\ell \geq 0$*, let* μ_ℓ *be the Gibbs distribution on* T_ℓ *corresponding to the spin system with interaction matrix* $\mathbf{B}$*, and let* $\widehat{\nu}_\ell(\cdot)$ *be the distribution in* (99) *corresponding to the fixpoint* $\mathbf{R}$ *of the tree recursions.*

Then, for all $\ell \geq 0$*, for all* $\eta : W_\ell \to [q]$ *and* $i \in [q]$*, it holds that*

$$\widehat{\nu}_\ell(\eta) = \nu(\sigma_{W_\ell} = \eta) \text{ and } \mu_\ell(\sigma_\rho = i \mid \sigma_{W_\ell} = \eta) = \nu(\sigma_\rho = i \mid \sigma_{W_\ell} = \eta).$$

Proof. Let $d := \Delta - 1$. The proof is by induction on ℓ. For $\ell = 0$, the lemma is trivial. Let us assume that the lemma holds for ℓ, we will prove it for $\ell+1$. For a vertex $v \in W_\ell$ denote by $v_1,\ldots,v_d$ the children of v in T (note that $v_1,\ldots,v_d \in W_{\ell+1}$).

We prove first that $\widehat{\nu}_{\ell+1}(\eta) = \nu(\sigma_{W_{\ell+1}} = \eta)$ for all $\eta : W_{\ell+1} \to [q]$. By definition of the broadcasting process, conditioned on the configuration τ on W_ℓ, the spins of the vertices in $W_{\ell+1}$ are independent. We thus have that, for all $\eta : W_{\ell+1} \to [q]$,

$$\nu(\sigma_{W_{\ell+1}} = \eta) = \sum_{\tau : W_\ell \to [q]} \nu(\sigma_{W_\ell} = \tau) \prod_{v \in W_\ell} \prod_{j=1}^{d} M_{\tau_v,\eta_{v_j}}.$$

Using the induction hypothesis we have that $\nu(\sigma_{W_\ell} = \tau) = \widehat{\nu}_\ell(\tau) \propto \mu_\ell(\sigma_{W_\ell} = \tau)Q_\ell(\tau)$ for all $\tau : W_\ell \to [q]$ and, substituting the value of the product measure $Q_{W_\ell}(\tau)$, we obtain

$$\nu(\sigma_{W_{\ell+1}} = \eta) \propto \sum_{\tau : W_\ell \to [q]} \mu_\ell(\sigma_{W_\ell} = \tau) \prod_{v \in W_\ell} R_{\tau_v} \prod_{j=1}^{d} M_{\tau_v,\eta_{v_j}} \quad \forall \eta : W_{\ell+1} \to [q]. \tag{100}$$

We also have that $\mu_{\ell+1}(\sigma_{W_{\ell+1}} = \eta) \propto \sum_{\tau : W_\ell \to [q]} \mu_\ell(\sigma_{W_\ell} = \tau) \prod_{v \in W_\ell} \prod_{j=1}^{d} B_{\tau_v,\eta_{v_j}}$ for all $\eta : W_{\ell+1} \to [q]$, so substituting the value of the product measure $Q_{W_{\ell+1}}(\eta)$ we obtain

$$\widehat{\nu}_{\ell+1}(\eta) \propto \sum_{\tau : W_\ell \to [q]} \mu_\ell(\sigma_{W_\ell} = \tau) \prod_{v \in W_\ell} \prod_{j=1}^{d} R_{\eta_{v_j}} B_{\tau_v,\eta_{v_j}} \quad \forall \eta : W_{\ell+1} \to [q]. \tag{101}$$

To complete the induction step, it thus remains to show that the r.h.s. in (100) and (101) are proportional by a factor that does not depend on η. This will follow from

$$R_i M_{i,j_1} \cdots M_{i,j_d} \propto (R_{j_1} B_{i,j_1}) \cdots (R_{j_d} B_{i,j_d}) \quad \forall i, j_1, \ldots, j_d \in [q]. \tag{102}$$

(Set $i = \tau_v$, $j_1 = \eta_{v_1}, \ldots, j_d = \eta_{v_d}$ and multiply over $v \in W_\ell$.) To see (102), note that by (97), we have that, for every $i, j \in [q]$,

$$M_{ij} = \frac{B_{ij} R_j}{\sum_{j'} B_{ij'} R_{j'}}, \tag{97}$$

so, to prove (102), it suffices to show that the quantities $\frac{R_i}{\left(\sum_{j' \in [q]} B_{ij'} R_{j'}\right)^d}$ do not depend on $i \in [q]$. This is a consequence of the fact that $\mathbf{R}$ is a fixpoint of the tree recursions, i.e., $R_1, \dots, R_q$ satisfy (6). This completes the induction step for the first equality in the lemma.

We next show the induction step for the second equality in the lemma. Fix $\eta : W_{\ell+1} \to [q]$. Denote by $\rho_1, \dots, \rho_d$ the children of the root ρ. For $k \in [d]$, denote by $W_{\ell+1,k}$ the set of vertices in $W_{\ell+1}$ which are in the subtree of T rooted at ρ_k and by η_k the restriction of η on $W_{\ell+1,k}$. By the induction hypothesis, we have, for every $k \in [d]$ and $i \in [q]$,

$$X_k(i) := \mu_\ell(\sigma_{\rho_k} = i \mid \sigma_{W_{\ell+1,k}} = \eta_k) = \nu(\sigma_\rho = i \mid \sigma_{W_\ell} = \eta_k).$$

Note that $\nu(\sigma_{W_{\ell+1,k}} = \eta_k \mid \sigma_{\rho_k} = i) = \nu(\sigma_{W_\ell} = \eta_k \mid \sigma_\rho = i)$ from where we obtain that

$$\frac{\nu(\sigma_{\rho_k} = i \mid \sigma_{W_{\ell+1,k}} = \eta_k)}{\nu(\sigma_{\rho_k} = i)} \propto \frac{\nu(\sigma_\rho = i \mid \sigma_{W_\ell} = \eta_k)}{\nu(\sigma_\rho = i)} = \frac{X_k(i)}{R_i} \quad \forall i \in [q].$$

(Note that the normalizing factor depends on η_k.)

Using that $T_{\ell+1}$ is a tree, we then calculate that

$$\mu_{\ell+1}(\sigma_\rho = i \mid \sigma_{W_{\ell+1}} = \eta) \propto \prod_{k\in[d]} \left(\sum_{j\in[q]} B_{ij} X_k(j) \right) \text{ for } i \in [q] \tag{103}$$

and (see [46, Lemma 3.1] for a thorough derivation)

$$\nu(\sigma_\rho = i \mid \sigma_{W_{\ell+1}} = \eta) \propto R_i \prod_{k\in[d]} \left(\sum_{j\in[q]} M_{ij} \frac{X_k(j)}{R_j} \right) \text{ for } i \in [q]. \tag{104}$$

By an argument completely analogous to the one we used for (100) and (101) (i.e., using (97) and the fact that $\mathbf{R}$ is a fixpoint of the tree recursions (6)), we obtain that the r.h.s. in (103) and (104) are proportional by a factor that does not depend on i, thus completing the induction step for the second equality in the lemma.

This concludes the proof of Lemma 52. $\square$

A.3. Application to the ferromagnetic Potts model — proof of (56). We are now able to apply the results of Appendices A.1 and A.2 to the ferromagnetic Potts model and prove the bound (56) for the ordered phases on the tree.

Recall that the interaction matrix $\mathbf{B}$ of the q-state ferromagnetic Potts model has diagonal entries equal to $B > 1$ and off-diagonal entries equal to 1. An ordered phase corresponds to a fixpoint $\mathbf{R} = (R_1, \dots, R_q)$ of the tree recursions. Thus, the R_i's satisfy

$$R_i \propto \left(B_{ii} R_i + \sum_{j\neq i} R_j \right)^{\Delta-1} \quad \forall i \in [q]. \tag{105}$$

Recall that there are q ordered phases which are symmetric, each corresponding to a color $i \in [q]$. Without loss of generality, we will focus on the ordered phase corresponding to the color $i = 1$. As we showed in section 8, the solution of (105) corresponding to the ordered phase $i = 1$ is given by the vector $\mathbf{R}$ which satisfies (105), $R_1 > R_2 = \cdots = R_q$, and R_1/R_q is maximum (see Remark 19). Such a solution exists in the nonuniqueness region, i.e., when $B > \mathfrak{B}_u$. The broadcasting matrix

$\mathbf{M}$ corresponding to the ordered phase $i = 1$ is given by (97):
(106)

$$
\begin{aligned}
M_{11} &= \frac{BR_1}{BR_1 + (q-1)R_q}, & M_{1j} &= \frac{R_q}{BR_1 + (q-1)R_q} \text{ for } j \neq 1,\\
M_{ii} &= \frac{BR_q}{R_1 + (q-2+B)R_q} \text{ for } i \neq 1, & M_{ij} &= \frac{R_q}{R_1 + (q-2+B)R_q} \text{ for } i \neq 1,\ i \neq j.
\end{aligned}
$$

We need the following lemma, which can be inferred from [35, Proof of Theorem 1.4]. For completeness, we give the proof.

LEMMA 53. *Let $\Delta \geq 3$ be an integer and $B > \mathfrak{B}_u$. Then, the broadcasting process* $\mathbf{M}$ *defined by* (106) *is nonreconstructible on the* $(\Delta-1)$*-ary tree.*

Proof. Let $i, j \in [q]$ be two arbitrary colors with $i \neq j$ and consider two copies X, Y of the broadcasting process on the $(\Delta-1)$-ary tree where the spins of the root ρ are conditioned to be i and j, respectively. To show that the total variation distance between the distributions $\nu(\sigma_{W_\ell} = \cdot \mid \sigma_\rho = i)$ and $\nu(\sigma_{W_\ell} = \cdot \mid \sigma_\rho = j)$ goes to 0 as $\ell \to \infty$, it suffices to couple X, Y so that the expected number of disagreements, i.e., vertices in W_ℓ whose spins are different, goes to 0 as $\ell \to \infty$. In turn, it suffices to couple one step of the broadcasting process so that the expected number of disagreements is bounded by some constant $\kappa < 1/(\Delta-1)$, since this yields that the expected number of disagreements at level ℓ decays exponentially with ℓ, at least as fast as $((\Delta-1)\kappa)^\ell$.

In particular, let (u, v) be an arbitrary edge in the tree, with u being the parent of v. By the coupling lemma, conditioned on the spin of u in X and Y, we can couple the spins of v in X and Y so that the probability that they are different is bounded by κ, where

$$
\kappa := \max_{i,j \in [q]} d_{TV}\big(\nu(\sigma_v = \cdot \mid \sigma_u = i), \nu(\sigma_v = \cdot \mid \sigma_u = j)\big) = \max_{i,j} \frac{1}{2} \sum_{k \in [q]} |M_{ik} - M_{jk}|.
$$

In the following, we justify that $\kappa < 1/(\Delta-1)$. We will see that, in the case of the ferromagnetic Potts model, κ is related to the eigenvalues of the Jacobian matrix of the tree recursions (evaluated at the fixpoint), which we have already studied in section 8. First, we find a simpler expression for κ. Consider colors $i, j \neq 1$. Then

$$
\frac{1}{2} \sum_{k \in [q]} |M_{ik} - M_{jk}| = M_{ii} - M_{ij} = \frac{(B-1)R_q}{R_1 + (q-2+B)R_q} = \lambda_1,
$$

where λ_1 is as in (39). Consider now the case that $i = 1$ and $j \neq 1$. Using that $B > 1$ and $R_1 > R_q$, we have $M_{11} > M_{j1}$ and $M_{1k} < M_{jk}$ for $k \neq 1$. It follows that

$$
\frac{1}{2} \sum_{k \in [q]} |M_{1k} - M_{jk}| = M_{11} - M_{j1} = \frac{BR_1}{BR_1 + (q-1)R_q} - \frac{R_q}{R_1 + (q-2+B)R_q} = \lambda_2,
$$

where λ_2 is as in (39). In the proof of Lemma 26, we showed that $\lambda_1, \lambda_2 < 1/(\Delta-1)$ for all $B > \mathfrak{B}_u$, which shows that $\kappa < 1/(\Delta-1)$, thus completing the proof of the lemma. ∎

We conclude this appendix by giving the proof of (56).

Proof of (56). Recall that we only need to consider the case where J consists of a single $(\Delta-1)$-ary tree of height $\ell = \lfloor \psi \log_{\Delta-1} n \rfloor$. Using the second equality in Lemma 52, we have that the set of "bad" configurations $\mathcal{B}_i$ defined in (52) is a subset of the set $\mathcal{B}_\ell$ defined in Theorem 50 (for all $0 < \theta < \frac{C\psi}{3\ln(\Delta-1)}$, it holds that $n^{-3\theta} > \exp(-C\ell)$). Further, using Lemmas 51 and 53, the assumptions of Theorem 50 are all satisfied for the broadcasting process $\mathbf{M}$ defined by (106). As we observed just after (99), $\widehat{\nu}_\ell(\cdot)$ is identical to the distribution $\nu^i(\cdot)$ (defined in (54)). Thus, the conclusion of Theorem 50 and the first equality in Lemma 52 yield

$$\nu^i(\sigma_W \in \mathcal{B}_i) \leq \exp(-\exp(C\ell)), \tag{56}$$

as wanted. $\square$

Appendix B. Moment asymptotics—Proofs of Lemmas 44 and 45. In this appendix, we give the proofs of Lemmas 44 and 45, which express the asymptotics of the moments in terms of certain determinants. The proofs of these lemmas are similar and closely follow [17, Lemma B.3].

Proof of Lemma 44. We have already seen that the ferromagnetism of the model implies a full-dimensional representation of $\mathbf{x}$ which consists of the variables x_{ij} with $(i,j) \in P_1^*$, where $P_1^* = P_1 \backslash \{(i,i) \mid i \in [q]\}$. For $(i,j) \notin P_1^*$, we will still use x_{ij} as shorthand for the appropriate linear combination of the variables inside the full-dimensional representation.

Recall that

$$\begin{aligned}\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}] = \binom{n}{\alpha_1 n, \ldots, \alpha_q n} \sum_{\mathbf{x}} \Bigg\{ \prod_i \binom{\Delta\alpha_i n}{\Delta x_{i1} n, \ldots, \Delta x_{iq} n} \\ \times \frac{\left[\prod_{i\neq j}(\Delta x_{ij} n)!\right]^{1/2} \prod_i (\Delta x_{ii} n - 1)!!}{(\Delta n - 1)!!} \prod_{i,j} B_{ij}^{\Delta x_{ij} n/2} \Bigg\}.\end{aligned} \tag{8}$$

Note that the range of $\mathbf{x}$ in the summation is over those vectors $\mathbf{x}$ such that, for each $i \in [q]$, $\Delta x_{ii} n = 2e_{ii} n$ is an even integer, which yields the constraint that $\sum_{j\neq i} \Delta x_{ij} n \equiv \Delta\alpha_i n (\bmod 2)$. Since Δn is even and $\sum_i \alpha_i = 1$, only $q-1$ of these constraints are linearly independent $(\bmod 2)$.

Since $\boldsymbol{\alpha}$ is a dominant phase, we have $\alpha_i > 0$ for all $i \in [q]$ (from the last part of Lemma 11). Also, for the maximizer $\mathbf{x}^*$ of $g_1(\mathbf{x})$, it holds that $x_{ij}^* > 0$ for all $(i,j) \in P_1$ (see (64)). Pick δ sufficiently small such that

$$\|\mathbf{x} - \mathbf{x}^*\|_2 \leq \delta \text{ implies } x_{ij} > 0 \ \forall (i,j) \in P_1.$$

Since $g_1(\mathbf{x})$ has the unique global maximum $\mathbf{x}^*$, standard compactness arguments imply that there exists $\varepsilon(\delta) > 0$ such that $\|\mathbf{x} - \mathbf{x}^*\|_2 \geq \delta$ implies $g(\mathbf{x}^*) - g(\mathbf{x}) \geq \varepsilon$. It follows that the contribution of terms with $\|\mathbf{x} - \mathbf{x}^*\|_2 \geq \delta$ to $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$ is exponentially small and may be ignored. Hence we may restrict our attention to $\mathbf{x}$ satisfying $\|\mathbf{x} - \mathbf{x}^*\|_2 < \delta$.

We now approximate the terms in (8) with $\|\mathbf{x} - \mathbf{x}^*\|_2 < \delta$ using Stirling's approximation. The only difference with the asymptotics in section 4 is that now the relative error of the approximation will be asymptotically $O(n^{-1})$. In particular, we will use the following asymptotics for factorials (which are a refinement of (12)). For any constant $c > 0$, it holds that

$$\begin{aligned}(cn)! &= \left(1 + O(n^{-1})\right)\sqrt{2\pi cn}\exp(cn\ln n + cn\ln c - cn),\\ (cn-1)!! &= \left(1 + O(n^{-1})\right)\sqrt{2}\exp\left(\frac{cn}{2}\ln n + \frac{cn}{2}\ln c - \frac{cn}{2}\right).\end{aligned} \tag{107}$$

Using these asymptotics to expand the terms in (8) (together with $\sum_j x_{ij} = \alpha_i$ and $\sum_i \alpha_i = 1$), we obtain

$$(108)\quad \frac{(2\pi n)^{(q-1)/2}\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]}{e^{n\Upsilon_1(\boldsymbol{\alpha},\mathbf{x}^*)}} = \Big(1+O\big(n^{-1}\big)\Big)2^{(q-1)/2} \times \sum_{\mathbf{x}} \big(\sqrt{2\pi\Delta n}\big)^q \Big(\prod_{(i,j)\in P_1} \frac{1}{\sqrt{2\pi\Delta n x_{ij}}}\Big) e^{n\Delta\big(g_1(\mathbf{x})-g_1(\mathbf{x}^*)\big)}.$$

In the l.h.s., the factor $(2\pi n)^{(q-1)/2}$ comes from the expansion of $\binom{n}{\alpha_1 n,\ldots,\alpha_q n}$; in the r.h.s., the factor $\big(\sqrt{2\pi\Delta n}\big)^q$ comes from the expansion of $(\Delta\alpha_i n)!$ for $i \in [q]$, the factor $1/\sqrt{2\pi\Delta n x_{ij}}$ comes from the expansion of $(\Delta x_{ij} n)!$ for $(i,j) \in P_1$, and the factor $2^{(q-1)/2}$ comes from the expansion of $(\Delta n - 1)!!$ and $(\Delta x_{ii} n - 1)!!$ for $i \in [q]$.

We are now ready to compute

$$L := \lim_{n\to\infty} \frac{(2\pi n)^{(q-1)/2}\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]}{e^{n\Upsilon_1(\boldsymbol{\alpha},\mathbf{x}^*)}}.$$

Since $\mathbf{x}^*$ is a critical point of $g_1(\mathbf{x})$, for all sufficiently small $\delta > 0$, we have the expansion

$$g_1(\mathbf{x}) - g_1(\mathbf{x}^*) = \frac{1}{2}(\mathbf{x}-\mathbf{x}^*)^{\intercal}\mathbf{H}(\mathbf{x}-\mathbf{x}^*) + O(\delta^3),$$

where $\mathbf{H} = \mathbf{H}^f_{1,\mathbf{x}}$ is the Hessian of the full-dimensional representation of $g_1(\mathbf{x})$ evaluated at $\mathbf{x} = \mathbf{x}^*$ (the matrix $\mathbf{H}$ has dimension $(|P_1| - q) \times (|P_1| - q)$). Note that g_1 is a strictly concave function, so $\mathbf{H}$ is negative definite. Using standard techniques of rewriting sums as integrals and the dominated convergence theorem (see [28, section 9.4]), we obtain

$$(109)\quad \lim_{n\to\infty} \sum_{\mathbf{x}} \big(\sqrt{2\pi\Delta n}\big)^q \Big(\prod_{(i,j)\in P_1} \frac{1}{\sqrt{2\pi\Delta n x_{ij}}}\Big) e^{n\Delta\big(g_1(\mathbf{x})-g_1(\mathbf{x}^*)\big)} = \frac{\big(\sqrt{\Delta}\big)^{|P_1|-q}}{2^{q-1}\big(\prod_{(i,j)\in P_1} x^*_{ij}\big)^{1/2}} \Big(\frac{1}{(\sqrt{2\pi})^{|P_1|-q}} \int_{-\infty}^{\infty}\cdots\int_{-\infty}^{\infty} \exp\Big(\frac{\Delta}{2}\mathbf{x}^{\intercal}\,\mathbf{H}\mathbf{x}\Big)d\mathbf{x}\Big).$$

Note that in the r.h.s. of (109), the factor $1/2^{q-1}$ comes from the $q-1$ constraints $(\mathrm{mod}\, 2)$ restricting the range of $\mathbf{x}$ (discussed just after the expression (8) for $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$ in the beginning of the proof).

Plugging into (109) the value of the Gaussian integral and substituting back into (108) yields that

$$L = \Big(2^{q-1}\prod_{(i,j)\in P_1} x^*_{ij}\Big)^{-1/2}\big(\sqrt{\Delta}\big)^{|P_1|-q}\mathrm{Det}(-\Delta\mathbf{H})^{-1/2} = \Big(2^{q-1}\prod_{(i,j)\in P_1} x^*_{ij}\Big)^{-1/2}\mathrm{Det}(-\mathbf{H})^{-1/2},$$

as wanted. This concludes the proof of Lemma 44. ∎

Proof of Lemma 45. We proceed analogously to the proof of Lemma 44. In particular, we will assume a full-dimensional representation of $(\boldsymbol{\gamma}, \mathbf{y})$ (see the beginning of section 11.2.2 for details on the choice of the representation). Note that $\boldsymbol{\gamma}, \mathbf{y}$ have $(q-1)^2$ and $|P_2| - q^2$ variables, respectively.

Recall that

$$\text{(10)}\quad \mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2] = \sum_{\boldsymbol{\gamma}} \binom{n}{\gamma_{11}n, \dots, \gamma_{qq}n} \sum_{\mathbf{y}} \Bigg\{ \prod_{i,k} \binom{\Delta\gamma_{ik}n}{\Delta y_{ik11}n, \dots, \Delta y_{ikqq}n}$$

$$\times \frac{\big[\prod_{(i,k)\neq(j,l)}(\Delta y_{ikjl}n)!\big]^{1/2} \prod_{i,k}(\Delta y_{ikik}n-1)!!}{(\Delta n-1)!!} \prod_{i,j,k,l} \big(B_{ij}B_{kl}\big)^{\Delta y_{ikjl}n/2} \Bigg\}.$$

We have that $\gamma^*_{ik} > 0$ for all $i,k \in [q]$ and $y^*_{ikjl} > 0$ for $(i,k,j,l) \in P_2$ (see the paragraph following Remark 43). Pick δ sufficiently small such that

$$\|(\boldsymbol{\gamma},\mathbf{y}) - (\boldsymbol{\gamma}^*,\mathbf{y}^*)\|_2 \leq \delta \text{ implies } \gamma_{ik} > 0 \text{ for } i,k \in [q] \;\&\; y_{ikjl} > 0 \text{ for } (i,k,j,l) \in P_2.$$

Using that $\Upsilon_2(\boldsymbol{\gamma},\mathbf{y})$ is maximized (uniquely) at $(\boldsymbol{\gamma},\mathbf{y}) = (\boldsymbol{\gamma}^*,\mathbf{y}^*)$, the same line of arguments as in the proof of Lemma 44 yields

$$\frac{\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2]}{e^{n\Upsilon_2(\boldsymbol{\gamma}^*,\mathbf{y}^*)}} = \Big(1+O\big(n^{-1}\big)\Big) \sum_{\boldsymbol{\gamma},\mathbf{y}} \Bigg\{ (\sqrt{2\pi n}) 2^{(q^2-1)/2} \Delta^{q^2/2}$$

$$\Big(\prod_{(i,k,j,l)\in P_2} \frac{1}{\sqrt{2\pi\Delta y_{ikjl}n}} \Big) e^{n\left(\Upsilon_2(\boldsymbol{\gamma},\mathbf{y}) - \Upsilon_2(\boldsymbol{\gamma}^*,\mathbf{y}^*)\right)} \Bigg\}.$$

In the r.h.s., the factor $\sqrt{2\pi n}$ comes from the expansion of $n!$, the factor $2^{(q^2-1)/2}$ comes from the expansion of $(\Delta n - 1)!!$ and $(\Delta y_{ikik}n-1)!!$ for $i,k \in [q]$, the factor $\Delta^{q^2/2}$ from the expansion of $(\Delta\gamma_{ik}n)!$ for $i,k \in [q]$, and the factor $1/\sqrt{2\pi\Delta y_{ikjl}n}$ from the expansion of $(\Delta y_{ikjl}n)!$ for $(i,k,j,l) \in P_2$.

We now compute

$$L = \lim_{n\to\infty} \frac{(2\pi n)^{(q-1)}\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2]}{e^{n\Upsilon_2(\boldsymbol{\gamma}^*,\mathbf{y}^*)}}.$$

Since $\boldsymbol{\gamma}^*, \mathbf{y}^*$ is a critical point of $\Upsilon_2(\boldsymbol{\gamma},\mathbf{y})$, for all sufficiently small $\delta > 0$, we have the expansion

$$\Upsilon_2(\boldsymbol{\gamma},\mathbf{y}) - \Upsilon_2(\boldsymbol{\gamma}^*,\mathbf{y}^*) = \frac{\Delta}{2}\big([\boldsymbol{\gamma},\mathbf{y}] - [\boldsymbol{\gamma}^*,\mathbf{y}^*]\big)^{\intercal}\mathbf{H}\big([\boldsymbol{\gamma},\mathbf{y}] - [\boldsymbol{\gamma}^*,\mathbf{y}^*]\big) + O(\delta^3),$$

where $\mathbf{H} = \mathbf{H}_2^f$ is the Hessian matrix of Υ_2 evaluated at $(\boldsymbol{\gamma}^*,\mathbf{y}^*)$ scaled by $1/\Delta$. Now, we may proceed analogously to the proof of Lemma 44 and obtain

$$\begin{aligned} L &= \Big(2^{q^2-1} \prod_{(i,k,j,l)\in P_2} y^*_{ikjl}\Big)^{-1/2} (\sqrt{\Delta})^{|P_2|-q^2} \\ &\qquad \Big(\frac{1}{\sqrt{2\pi}}\Big)^{|P_2|-2q+1} \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} \exp\Big(\frac{\Delta}{2}[\boldsymbol{\gamma},\mathbf{y}]^{\intercal}\mathbf{H}[\boldsymbol{\gamma},\mathbf{y}]\Big) d\mathbf{y}\, d\boldsymbol{\gamma}, \\ &= \Big(2^{q^2-1} \prod_{(i,k,j,l)\in P_2} y^*_{ikjl}\Big)^{-1/2} \Delta^{-(q-1)^2/2}\, \mathrm{Det}(-\mathbf{H})^{-1/2}, \end{aligned}$$

where in the last equality we substituted the value of the Gaussian integral.

This concludes the proof of Lemma 45. ∎

REFERENCES

[1] V. Bapst, A. Coja-Oghlan, S. Hetterich, F. Raßmann, and D. Vilenchik, *The condensation phase transition in random graph coloring*, Comm. Math. Phys., 341 (2016), pp. 543–606.

[2] G. Bennett, *Schur multipliers*, Duke Math. J., 44 (1977), pp. 603–639.

[3] A. A. Bulatov, M. Dyer, L. A. Goldberg, M. Jerrum, and C. McQuillan, *The expressibility of functions on the boolean domain, with applications to counting CSPs*, J. ACM, 60 (2013), 32.

[4] C. Borgs, J. T. Chayes, A. Frieze, J. H. Kim, P. Tetali, E. Vigoda, and V. H. Vu, *Torpid mixing of some Monte Carlo Markov chain algorithms in statistical physics*, in Proceedings of the 40th IEEE Symposium on Foundations of Computer Science (FOCS), 1999, pp. 218–229.

[5] C. Borgs, J. T. Chayes, and P. Tetali, *Tight bounds for mixing of the Swendsen-Wang algorithm at the Potts transition point*, Probab. Theory Related Fields, 152 (2012), pp. 509–557.

[6] G. R. Brightwell and P. Winkler, *Random colorings of a Cayley tree*, in Contemporary Combinatorics, Bolyai Soc. Math. Stud. 10, János Bolyai Mathematical Society, Budapest, 2002, pp. 247–276.

[7] J.-Y. Cai, A. Galanis, L. A. Goldberg, H. Guo, M. Jerrum, D. Štefankovič, and E. Vigoda, *#BIS-hardness for 2-spin systems on bipartite bounded degree graphs in the tree nonuniqueness region*, J. Comput. System Sci., 82 (2016), pp. 690–711.

[8] X. Chen, M. E. Dyer, L. A. Goldberg, M. Jerrum, P. Lu, C. McQuillan, and D. Richerby, *The complexity of approximating conservative counting CSPs*, in Proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS), 2013, pp. 148–159.

[9] C. Cooper and A. M. Frieze, *Mixing properties of the Swendsen-Wang process on classes of graphs*, Random Structures Algorithms, 15 (1999), pp. 242–261.

[10] M. Costeniuc, R. S. Ellis, and H. Touchette, *Complete analysis of phase transitions and ensemble equivalence for the Curie-Weiss-Potts model*, J. Math. Phys., 46 (2005), 063301.

[11] A. Dembo and A. Montanari, *Ising models on locally tree-like graphs*, Ann. Appl. Probab., 20 (2010), pp. 565–592.

[12] A. Dembo, A. Montanari, A. Sly, and N. Sun, *The replica symmetric solution for Potts models on d-regular graphs*, Comm. Math. Phys., 327 (2014), pp. 551–575.

[13] A. Dembo, A. Montanari, and N. Sun, *Factor models on locally tree-like graphs*, Ann. Probab., 41 (2013), pp. 4162–4213.

[14] M. E. Dyer, L. A. Goldberg, C. S. Greenhill, and M. Jerrum. *The relative complexity of approximate counting problems*, Algorithmica, 38 (2003), pp. 471–500.

[15] A. Galanis, Q. Ge, D. Štefankovič, E. Vigoda, and L. Yang, *Improved inapproximability results for counting independent sets in the hard-core model*, Random Struct. Algorithms, 45 (2014), pp. 78–110.

[16] A. Galanis, D. Štefankovič, and E. Vigoda, *Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models*, Combin. Probab. Comput., 25 (2016), pp. 500–559.

[17] A. Galanis, D. Štefankovič, and E. Vigoda, *Inapproximability for antiferromagnetic spin systems in the tree non-uniqueness region*, J. ACM, 62 (2015), 50.

[18] H.-O. Georgii, *Gibbs Measures and Phase Transitions*, 2nd ed., de Gruyter Stud. Math. 9, Walter de Gruyter, Berlin, 2011.

[19] A. Gerschenfeld and A. Montanari, *Reconstruction for models on random graphs*, in Proceedings of the 48th Annual Symposium on Foundations of Computer Science (FOCS), 2007, pp. 194–204.

[20] H. Guo and P. Lu. *Uniqueness, Spatial Mixing, and Approximation for Ferromagnetic 2-Spin Systems*, arXiv:1511.00493, 2015.

[21] L. A. Goldberg and M. Jerrum, *Approximating the partition function of the ferromagnetic Potts model*, J. ACM, 59 (2012), 25.

[22] V. K. Gore and M. R. Jerrum, *The Swendsen-Wang process does not always mix rapidly*, J. Stat. Phys., 97 (1999), pp. 67–86.

[23] L. A. Goldberg, M. Jerrum, and M. Paterson, *The computational complexity of two-state spin systems*, Random Structures Algorithms, 23 (2003), pp. 133–154.

[24] C. Greenhill, *The complexity of counting colourings and independent sets in sparse graphs and hypergraphs*, Comput. Complexity, 9 (2000), pp. 52–72.

[25] G. Grimmett, *The Random-Cluster Model*, Grundlehren Math. Wiss. 333, Springer-Verlag, Berlin, 2006.

[26] O. Häggström, *The random-cluster model on a homogeneous tree*, Probab. Theory Related Fields, 104 (1996), pp. 231–253.

[27] S. Janson, *Random regular graphs: Asymptotic distributions and contiguity*, Combin. Probab. Comput., 4 (1995), pp. 369–405.

[28] S. Janson, T. Łuczak, and A. Rucinski, *Random Graphs*. Wiley Ser. Discrete Math. Optim., Wiley, New York, 2000.

[29] M. Jerrum and A. Sinclair, *Polynomial-time approximation algorithms for the Ising model*, SIAM J. Comput., 22 (1993), pp. 1087–1116.

[30] F. P. Kelly, *Loss networks*, Ann. Appl. Probab., 1 (1991), pp. 319–378.

[31] G. Kemkes, X. Pérez-Giménez, and N. Wormald, *On the chromatic number of random d-regular graphs*, Adv. Math., 223 (2010), pp. 300–328.

[32] L. Li, P. Lu, and Y. Yin, *Correlation decay up to uniqueness in spin systems*, in Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2013, pp. 47–66.

[33] J. Liu, P. Lu, and C. Zhang, *The complexity of ferromagnetic two-spin systems with external fields*, in Proceedings of Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM), 2014, pp. 843–856.

[34] F. Martinelli, A. Sinclair, and D. Weitz, *Glauber dynamics on trees: Boundary conditions and mixing time*, Commun. Math. Phys., 250 (2004), pp. 301–334.

[35] F. Martinelli, A. Sinclair, and D. Weitz, *Fast mixing for independent sets, colorings and other models on trees*, Random Structures Algorithms, 31 (2007), pp. 134–172.

[36] M. Mezard and A. Montanari, *Information, Physics, and Computation*. Oxford University Press, New York, 2009.

[37] R. Montenegro and P. Tetali, *Mathematical aspects of mixing times in Markov chains*, Foundat. Trends Theoret. Comput. Sci. 1 (2006), pp. 237–354.

[38] E. Mossel, *Survey: Information flow on trees*, in Graphs, Morphisms and Statistical Physics, *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.* 63, AMS, Providence, RI, 2004, pp. 155–170.

[39] E. Mossel and A. Sly, *Exact thresholds for Ising-Gibbs samplers on general graphs*, Ann. Probab., 41 (2013), pp. 294–328.

[40] E. Mossel, D. Weitz, and N. Wormald, *On the hardness of sampling independent sets beyond the tree threshold*, Probab. Theory Related Fields, 143 (2009), pp. 401–439.

[41] F. Peruggi, F. Di Liberto, and G. Monroy, *Phase diagrams of the q-state Potts model on Bethe lattices*, Phys. A, 141 (1987), pp. 151–186.

[42] R. W. Robinson and N. C. Wormald, *Almost all regular graphs are Hamiltonian*, Random Structures Algorithms, 5 (1994), pp. 363–374.

[43] A. Sly, *Computational transition at the uniqueness threshold*, in Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2010, pp. 287–296.

[44] A. Sly and N. Sun, *The computational hardness of counting in two-spin models on d-regular graphs*, in Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2012, pp. 361–369.

[45] A. Sly, N. Sun, and Y. Zhang, *The Number of Solutions for Random Regular NAE-SAT*, arXiv:1604.08546, 2016.

[46] A. Sly and Y. Zhang, *Glauber Dynamics of Colorings on Trees*, arXiv:1412.3156, 2014.

[47] M. Ullrich, *Rapid Mixing of Swendsen-Wang Dynamics in Two Dimensions*, Ph.D. thesis, Universität Jena, Germany, 2012, arXiv:1212.4908.

[48] L. G. Valiant, *The complexity of enumeration and reliability problems*, SIAM J. Comput., 8 (1979), pp. 410–421.

[49] D. Weitz, *Counting independent sets up to the tree threshold*, in Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC), 2006, pp. 140–149.