

FACTOR MODELS ON LOCALLY TREE-LIKE GRAPHS

AMIR DEMBO*, ANDREA MONTANARI†, AND NIKE SUN‡

ABSTRACT. We consider homogeneous factor models on uniformly sparse graph sequences converging locally to a (unimodular) random tree T , and study the existence of the free energy density ϕ , the limit of the log-partition function divided by the number of vertices n as n tends to infinity. We provide a new interpolation scheme and use it to prove existence of, and to explicitly compute, the quantity ϕ subject to uniqueness of a relevant Gibbs measure for the factor model on T . By way of example we compute ϕ for the independent set (or hard-core) model at low fugacity, for the ferromagnetic Ising model at all parameter values, and for the ferromagnetic Potts model with both weak enough and strong enough interactions. Even beyond uniqueness our interpolation provides useful explicit bounds on ϕ .

In the regimes in which we establish existence of the limit, we show that it coincides with the Bethe free energy functional evaluated at a suitable fixed point of the belief propagation recursions on T . In the special case that T has a Galton-Watson law, this formula coincides with the non-rigorous "Bethe prediction" obtained by statistical physicists using the "replica" or "cavity" methods. Thus our work is a rigorous generalization of these heuristic calculations to the broader class of sparse graph sequences converging locally to trees. We also provide a variational characterization for the Bethe prediction in this general setting, which is of independent interest.

1. INTRODUCTION

Let $G = (V, E)$ be a finite undirected graph, and \mathcal{X} a finite alphabet of *spins*. A *factor model* on G is a probability measure on the space of (*spin*) *configurations* $\underline{\sigma} \in \mathcal{X}^V$ of form

$$\nu_{G, \underline{\psi}}^{\beta, B}(\underline{\sigma}) = \frac{1}{Z_{G, \underline{\psi}}(\beta, B)} \prod_{(ij) \in E} \psi^\beta(\sigma_i, \sigma_j) \prod_{i \in V} \bar{\psi}^B(\sigma_i), \quad (1.1)$$

where $\psi \equiv \psi^\beta$ is a symmetric function $\mathcal{X}^2 \rightarrow \mathbb{R}_{\geq 0}$ parametrized by $\beta \in \mathbb{R}$, $\bar{\psi} \equiv \bar{\psi}^B$ is a positive function $\mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ parametrized by $B \in \mathbb{R}$, and $Z_{G, \underline{\psi}}(\beta, B)$ is the normalizing constant, called the *partition function* (with its logarithm called the *free energy*). The pair $\underline{\psi} \equiv (\psi, \bar{\psi})$ is called a *specification* for the factor model (1.1).

In this paper we study the asymptotics of the free energy for sequences of (random) graphs $G_n = (V_n = [n], E_n)$ in the *thermodynamic limit* $n \rightarrow \infty$. More precisely, with $Z_n(\beta, B) \equiv Z_{G_n, \underline{\psi}}(\beta, B)$ and \mathbb{E}_n denoting expectation with respect to the law of G_n , we seek to establish the existence of the *free energy density*

$$\phi(\beta, B) \equiv \lim_{n \rightarrow \infty} \phi_n(\beta, B), \quad \text{where } \phi_n(\beta, B) \equiv \frac{1}{n} \mathbb{E}_n[\log Z_n(\beta, B)], \quad (1.2)$$

Date: October 24, 2011.

2010 Mathematics Subject Classification. 05C80, 60K35, 82B20, 82B23.

Key words and phrases. Factor models, random graphs, belief propagation, Bethe measures, Potts model, independent set, Gibbs measures, free energy density, local weak convergence.

Research partially supported by NSF grant DMS-1106627. †Research partially supported by Department of Defense (AFRL/AFOSR) NDSEG Fellowship.

and to determine its value. (In the literature, $\phi(\beta, B)$ is also referred to as the “free entropy density” or “pressure.”)

The primary example we consider is the q -Potts model on G with inverse temperature β and magnetic field B , the probability measure on \mathcal{X}^V (with $\mathcal{X} = [q] \equiv \{1, \dots, q\}$) given by

$$\nu_G^{\beta, B}(\underline{\sigma}) = \frac{1}{Z_G(\beta, B)} \exp \left\{ \beta \sum_{(ij) \in E} \delta_{\sigma_i, \sigma_j} + B \sum_{i \in V} \delta_{\sigma_i, 1} \right\}. \quad (1.3)$$

The model is said to be *ferromagnetic* if $\beta \geq 0$, *anti-ferromagnetic* otherwise. (In analogy with the Potts model, in the general factor model setting we continue to refer to β as the interaction or temperature parameter and to B as the magnetic field.)

Potts models have been intensively studied in statistical mechanics because of their key role in the theory of phase transitions [40], critical phenomena [43], and conformally invariant scaling limits [33]. As demonstrated for instance in [30] for the Ising model (i.e. the Potts model with $q = 2$), determining the limit (1.2) plays a key role in characterizing the asymptotic structure of the measures $\nu_{G_n}^{\beta, B}$ in the thermodynamic limit. Potts models are also of great interest in combinatorics: recall in fact that the partition function admits a random-cluster representation ([14, 19]; see also §4.2), which at $B = 0$ reads

$$Z_G(\beta, 0) = \sum_{F \subseteq E} (e^\beta - 1)^{|F|} q^{k(F)},$$

with $k(F)$ denoting the number of connected components induced by the subset of edges $F \subseteq E$ (cf. (4.2)). Up to a multiplicative constant this coincides with the Tutte polynomial $T_G(x, y)$ of G evaluated at $x = 1 + q(e^\beta - 1)^{-1}$, $y = e^\beta$ (see e.g. [37]).

Mathematical statistical mechanics has focused so far on specific graph sequences G_n , e.g. on finite exhaustions of the rectangular grid or other regular lattices in d dimensions with d fixed. Under mild conditions on the sequence, existence of the free energy density is a consequence of the following well-known argument (see e.g. [34, Propn. 2.3.2]): each graph G_n can be decomposed into smaller blocks by deleting a collection of edges whose number is negligible in comparison with the volume. Consequently the sequence $\log Z_{G_n}$ is approximately sub-additive in n , implying existence of the limit (see [22]).

In this paper we consider sparse graphs with a locally tree-like structure — formally, graph sequences G_n converging locally weakly to (random) trees; see Defn. 1.1 below. Although the study of statistical mechanics “beyond \mathbb{Z}^d ” is not directly motivated by physics considerations, physicists have been interested in models on alternative graph structures for a long time (an early example being [12]). Moreover, the study of factor models on *sparse* graphs has many motivations coming from computer science and statistical inference (see [8, 29]). Indeed, another example we will consider is the *independent set* or *hard-core* model on G with fugacity $\lambda > 0$, the probability measure on \mathcal{X}^V (with $\mathcal{X} = \{0, 1\}$) given by

$$\nu_G^\lambda(\underline{\sigma}) = \frac{1}{Z_G(\lambda)} \prod_{(ij) \in E} \mathbf{1}\{\sigma_i \sigma_j \neq 1\} \prod_{i \in V} \lambda^{\sigma_i}, \quad (1.4)$$

(where $B \equiv \log \lambda$ is the magnetic field). This model always has anti-ferromagnetic interactions, and is of significant interest in computer science. The independent set decision problem is NP-complete (via the clique decision problem [7, 24]). As λ increases the measure ν_G^λ becomes increasingly concentrated on the maximal independent sets; the optimization problem of finding such sets is NP-hard [26] and hard to approximate ([44] and references therein). The problem of counting independent sets (i.e. computing $Z_G(1)$) for graphs of maximum degree Δ is $\#P$ -complete for

$\Delta \geq 3$ ([18] and references therein). Although there exists a PTAS (polynomial-time approximation scheme) for $Z_G(\lambda)$ for λ below a certain “uniqueness threshold” [39], there is strong evidence that computation is hard above this threshold (see [31, 36, 39] and the references therein).

Since infinite trees are non-amenable, G_n cannot be decomposed by removing a vanishing fraction of edges, so the preceding argument no longer applies: in physics terms, *surface effects are non-negligible even in the thermodynamic limit*. Despite this, statistical physicists expect the free energy density (1.2) to exist on a large class of locally tree-like graphs. Even more surprisingly, employing non-rigorous but mathematically sophisticated heuristics such as the “replica” or “cavity” methods they derive exact formulae for this limit for a number of statistical mechanics models on locally tree-like graphs (see e.g. [29] and the references therein). The primary example considered in these works is the graph chosen uniformly at random from those with n vertices and $m = m(n)$ edges, with $m/n \rightarrow \gamma \in \mathbb{R}$; such graphs converge locally to the Galton-Watson tree with $\text{Pois}(2\gamma)$ offspring distribution. The Galton-Watson tree with general offspring distribution can be obtained as the local weak limit of random graphs with specified degree profile corresponding to the offspring distribution; the physics heuristics extend to this and even more general settings.

There is no good argument for why the limit (1.2) exists; the heuristic replica or cavity methods compute this limit starting from the postulate that it exists. A significant breakthrough was achieved by the interpolation method first developed by Guerra and Toninelli [21] for the Sherrington-Kirkpatrick model from spin-glass theory, and then generalized to a number of statistical physics models on sparse graphs [15, 16, 32] and related constraint satisfaction problems [4]. This method establishes super-additivity of $\log Z_{G_n}$ which implies existence of the limit (1.2). Unfortunately, this approach appears limited to models with *repulsive* interactions, i.e. in which higher weight is given to configurations in which neighboring vertices take different values. In particular, it does not apply to the ferromagnetic Potts model. This is especially puzzling because the heuristic physics predictions do not distinguish between the two cases, and there is no fundamental reason why the limit should be computable in one case and not in the other. Further, this interpolation method only applies to very restricted classes of graph sequences (typically, uniformly random given the degree sequence); notably, existence of the limit is not proved for *deterministic* graph sequences. Finally, the method gives no way to actually compute the limit, although interpolation has been used to prove upper bounds [15, 16, 32].

In this paper we follow a different approach relying only on local weak convergence of the graph sequence $(G_n)_{n \geq 1}$ to some limiting (random) tree. The general idea is that the corresponding factor models (1.1) must converge (passing to a subsequence as needed), to a Gibbs measure on the limiting tree; the task then “reduces” to the one of identifying the correct limit. This is still a substantial challenge because, in general, there is an uncountable number of “candidate” Gibbs measures for the limit. Nevertheless, this program was carried through in [9] for Ising models on graphs converging locally to Galton-Watson tree, under a “uniform sparsity” assumption (Defn. 1.3), on the degree distribution (it is further assumed in [9] that the distribution has finite second moment; this condition was relaxed in [11], thereby handling the case of power law graphs). The result of [9, 11] provides also a fairly explicit expression $\Phi(\beta, B)$ for the free energy density, defined solely in terms of the limiting tree. This expression coincides with the so-called “Bethe prediction” of statistical physics, derived earlier for random graphs with given degree distribution using the “replica” or “cavity” methods.

We develop this approach here in more generality. Rather than considering a specific model such as the Ising, we establish results for general abstract factor models satisfying mild regularity conditions (see (H1) below), covering in particular the Potts and independent set models. We also make no distributional assumptions on the graphs G_n or the limiting random tree, other

than some integrability conditions (see Defn. 1.3 and (H2) below). In this setting we develop a general interpolation scheme (Thm. 1.13) which, under appropriate assumptions, bounds differences $\phi_n(\beta, B)$ in the limit $n \rightarrow \infty$ by differences in a functional $\Phi(\beta, B)$ (see (1.10)) defined solely in terms of the limiting tree. Although we continue to refer to this $\Phi(\beta, B)$ as the ‘‘Bethe prediction,’’ we remark that it is a considerable generalization of earlier formulae obtained in the special case of Galton-Watson trees by statistical physics methods. It is defined as the evaluation of the ‘‘Bethe free energy functional’’ (1.9) at a specific Gibbs measure on the limiting tree, and corresponds to what physicists call the ‘‘replica symmetric solution’’: whereas it is expected to hold in the high-temperature regime (i.e. with small enough interactions), for many factor models it is incorrect at low temperature. However, we will show that in ‘‘uniqueness regimes,’’ where the set of Gibbs measures on the limiting tree corresponding to the factor model specification $\underline{\psi}$ is a singleton, the upper and lower bounds of Thm. 1.13 match to completely verify the Bethe prediction (Thm. 1.14).

We then apply our interpolation scheme to compute the free energy density in specific models. We verify the Bethe prediction for the independent set model with low fugacity (Thm. 1.11) as a consequence of Thm. 1.14. Further, by using monotonicity properties to restrict the set of relevant Gibbs measures, we obtain results for the Potts model going beyond the implications of Thm. 1.14: for $q = 2$ (Ising), we verify the Bethe prediction for all $\beta \geq 0$, $B \in \mathbb{R}$ (Thm. 1.8), extending the results of [9, 11] to general locally tree-like graph sequences. For general q , we verify the prediction in regimes of non-negative (β, B) in which two specific Gibbs measures on the limiting tree coincide, namely, the Gibbs measures arising from free and 1 boundary conditions coincide (see Defn. 1.7 below). This condition is satisfied throughout the range $\{\beta \geq 0, B > 0\}$ for $q = 2$; when $q \geq 3$ there are regimes of non-uniqueness in which it fails, but we will show that it is satisfied both at β sufficiently small and sufficiently large, i.e. at high and low temperatures.

Thm. 1.13 can give useful bounds even beyond uniqueness regimes. As an illustration, we study the Potts model in the case that G_n converges locally to the d -regular tree \mathbb{T}_d . In Thm. 1.10 we explicitly characterize the non-uniqueness regime of this model and use Thm. 1.13 to give bounds for $\phi_n(\beta, B)$ within this regime. In a forthcoming work with Allan Sly we prove that in this setting, $\phi(\beta, B)$ exists and matches the lower bound of Thm. 1.10. We also compute there the asymptotic free energy $\phi(\lambda)$ (all $\lambda \geq 0$) for the independent set model on d -regular bipartite graphs. In contrast, for generic non-bipartite G_n the consensus in physics is for a full replica symmetry breaking for large enough λ , and consequently there does not exist even a heuristic prediction for the free energy density in this regime.

As mentioned above, the Bethe prediction $\Phi(\beta, B)$ is the evaluation of the Bethe free energy functional at a specific Gibbs measure on the limiting tree. This Gibbs measure has a characterization in terms of ‘‘messages’’ $h_{x \rightarrow y} \equiv h_{(T, x \rightarrow y)}$ defined on the directed edges $x \rightarrow y$ of each tree T , such that the entire collection of messages is a fixed point of a certain ‘‘belief propagation recursion’’ (1.14). Motivated by the finite-graph optimization of [41], we provide a variational characterization of the Bethe prediction (Thm. 1.16) which is of independent interest. In particular, this formulation suggests non-trivial connections with large deviation principles.

1.1. Local weak convergence and the Bethe prediction. We study factor models on graphs which are ‘‘locally tree-like’’ in a sense which we now formalize, starting with a few notations and conventions. All graphs are taken to be undirected and locally finite. In a graph $G = (V, E)$, let d denote graph distance, and for $v \in V$ write $B_t(v)$ for the sub-graph of G induced by $\{w \in V : d(v, w) \leq t\}$. Write $v \sim w$ if v, w are neighbors in G , and write ∂v for the set of neighbors of v and $D_v \equiv |\partial v|$. Let \mathcal{G}_\bullet denote the space of isomorphism classes of (finite or infinite) rooted, connected graphs (G, o) . A metric on this space is given by defining the distance between (G_1, o_1) and (G_2, o_2)

in \mathcal{G}_\bullet to be $1/(1+R)$ where R is the maximal $r \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ such that $B_R(o_1) \cong B_R(o_2)$; with this definition \mathcal{G}_\bullet is a complete separable metric space (see e.g. [1]). Let $\mathcal{T}_\bullet \subset \mathcal{G}_\bullet$ denote the closed subspace of (rooted) trees $T \equiv (T, o)$, the acyclic elements of \mathcal{G}_\bullet . We write T^t for $B_t(o)$ in T , and in particular we use T^0 to denote the single-vertex tree. We now define the precise notion of graph limits considered throughout this paper.

Definition 1.1. Let $G_n = (V_n, E_n)$ ($n \geq 1$) be a sequence of random graphs, and let I_n be a vertex chosen uniformly at random from V_n . We say G_n converges locally (weakly) to the random tree T if for each $t \geq 0$, $B_t(I_n)$ converges in law to T^t in the space \mathcal{G}_\bullet . We say in this case that the G_n are locally tree-like.

We will make repeated use of the fact that local weak limits of graph sequences have the property of unimodularity, whose definition we recall here (for a detailed account see [1]). Let $\mathcal{G}_{\bullet\bullet}$ denote the space of isomorphism classes of bi-rooted, connected graphs with a distinguished ordered pair, denoted (G, i, j) (we do not require $i \sim j$); $\mathcal{G}_{\bullet\bullet}$ is metrizable in a similar manner as \mathcal{G}_\bullet .

Definition 1.2. A Borel probability measure μ on \mathcal{G}_\bullet is said to be unimodular if it obeys the mass-transport principle:

$$\mathbb{E}_\mu \left[\sum_{x \in V(G)} f(G, o, x) \right] = \mathbb{E}_\mu \left[\sum_{x \in V(G)} f(G, x, o) \right] \quad \forall f : \mathcal{G}_{\bullet\bullet} \rightarrow \mathbb{R}_{\geq 0} \text{ Borel.} \quad (1.5)$$

We say that μ is involution invariant if (1.5) holds when restricted to f supported only on those (G, x, y) with $x \sim y$.

A measure μ on \mathcal{G} is involution invariant if and only if it is unimodular [1, Propn. 2.2]. The concept of unimodularity first appeared in [5], where it was observed that local weak limits of graph sequences must be unimodular [5, §3.2]. The converse of this implication remains a well-known open question (see [1]).

Definition 1.3. The graph sequence G_n is uniformly sparse if the D_{I_n} are uniformly integrable, that is, if

$$\lim_{L \rightarrow \infty} \left(\limsup_{n \rightarrow \infty} \mathbb{E}_n [D_{I_n} \mathbf{1}\{D_{I_n} \geq L\}] \right) = 0$$

(where \mathbb{E}_n denotes expectation over the law of G_n and I_n).

We assume throughout that G_n ($n \geq 1$) is a uniformly sparse graph sequence converging locally weakly to the random tree T of (unimodular) law μ such that $\mu(D_o = 0) < 1$; this entire setting is hereafter denoted $G_n \rightarrow_{lwc} \mu$. In this setting we will describe general conditions under which the asymptotic free energy $\phi(\beta, B)$ for the factor model (1.1) exists and agrees with the ‘‘Bethe energy prediction,’’ which we now describe. (If the sequence of random graphs G_n is such that $G_n \rightarrow_{lwc} \mu$ for per realization of the sequence (a.s.), then our results apply instead to the a.s. limit of $n^{-1} \log Z_n(\beta, B)$.)

Let $\Delta_{\mathcal{X}}$ denote the $(|\mathcal{X}| - 1)$ -dimensional probability simplex. Let \mathcal{T}_\bullet^+ denote \mathcal{T}_\bullet without the single-vertex tree T^0 , and let $\mathcal{T}_e \subset \mathcal{G}_{\bullet\bullet}$ be the space of isomorphism classes of trees $T \in \mathcal{T}_\bullet^+$ rooted at a directed edge $x \rightarrow y$, written $(T, x \rightarrow y)$ or simply $x \rightarrow y$ for short. If T has law μ for μ a unimodular measure on \mathcal{T}_\bullet , we let μ^\uparrow and μ^\downarrow denote the laws of $(T, J \rightarrow o)$ and $(T, o \rightarrow J)$, respectively, for J chosen uniformly at random from ∂o conditioned on the event $\{T \in \mathcal{T}_\bullet^+\}$. Unimodularity of μ is then equivalent to

$$\mathbb{E}_{\mu^\downarrow} [D_x f(T, x \rightarrow y)] = \mathbb{E}_{\mu^\uparrow} [D_y f(T, x \rightarrow y)],$$

so in particular μ^\uparrow and μ^\downarrow are equivalent measures.

Definition 1.4. The *message space* is the space $\mathcal{H} \equiv \mathcal{H}_\mu$ of measurable functions

$$h : \mathcal{T}_e \times \mathbb{R}^2 \rightarrow \Delta_{\mathcal{X}}, \quad ((T, x \rightarrow y), \beta, B) \mapsto (h_{x \rightarrow y}^{\beta, B}(\sigma))_{\sigma \in \mathcal{X}},$$

taken up to μ^\uparrow -equivalence.

Remark 1.5. For $(T, x \rightarrow y) \in \mathcal{T}_e$ let $T_{x \rightarrow y}$ denote the component sub-tree rooted at x which results from deleting edge (x, y) from T . The interpretation of $h_{x \rightarrow y}$ is that it is a message from x to y on the tree T , giving the distribution of σ_x for the factor model (1.1) on $T_{x \rightarrow y}$. Indeed, although we do not require it in general, in our concrete examples $h_{x \rightarrow y}$ depends only on this component sub-tree.

For $T \in \mathcal{T}_\bullet$ and $h \in \mathcal{H}$, let

$$\Phi_T(\beta, B, h) \equiv \Phi_T^{\vee x}(\beta, B, h) - \Phi_T^{\circ}(\beta, B, h) \equiv \Phi_T^{\vee x}(\beta, B, h) - \frac{1}{2} \sum_{j \in \partial o} \Phi_T^{(oj)}(\beta, B, h) \quad (1.6)$$

where

$$\Phi_T^{\vee x}(\beta, B, h) \equiv \log \left\{ \sum_{\sigma} \bar{\psi}(\sigma) \prod_{j \in \partial o} \left(\sum_{\sigma_j} \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j) \right) \right\}, \quad (1.7)$$

$$\Phi_T^{(oj)}(\beta, B, h) \equiv \log \left\{ \sum_{\sigma, \sigma_j} \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j) h_{o \rightarrow j}(\sigma) \right\}. \quad (1.8)$$

We take the usual convention that the empty sum is zero and the empty product is one, so $\Phi_T = \log(\sum_{\sigma} \bar{\psi}(\sigma))$ in case $T = T^0$. Although we suppress it from the notation, in the above equations $\bar{\psi}$ and h are taken to be evaluated at (β, B) . The *Bethe free energy functional* on \mathcal{H} for the factor model (1.1) on $G_n \rightarrow_{lwc} \mu$ is defined by

$$\Phi_\mu(\beta, B, h) \equiv \mathbb{E}_\mu[\Phi_T(\beta, B, h)], \quad (1.9)$$

provided the integral is well-defined (see Lem. 2.2). The *Bethe prediction* is that the asymptotic free energy $\phi(\beta, B)$ of (1.2) exists and equals

$$\Phi_\mu^{\text{Bethe}}(\beta, B) \equiv \Phi_\mu(\beta, B, h^*) \quad (1.10)$$

where $h^* \in \mathcal{H}$ is a certain fixed point of the ‘‘belief propagation recursion’’ (1.15) below. We often drop the subscript μ when it is clear from context.

Remark 1.6. In the case that the recursion (1.15) has multiple fixed points, the *Bethe prediction* is defined to be the supremum of $\Phi(\beta, B, h)$ over admissible fixed points h . While in the abstract factor model setting all fixed points are admissible, in specific models typically there are ‘‘natural’’ criteria restricting the set of admissible fixed points. We will demonstrate this in the Ising and Potts models where restrictions are imposed by monotonicity and symmetry considerations.

The rationale for the Bethe prediction is explained in detail in [8, §3]. In brief, when G is a finite tree and μ_G is the law of (G, I) for I a uniform element of V (μ_G is a measure on \mathcal{T}_\bullet , but not necessarily unimodular), the equations (1.15) have a unique solution, given by the so-called ‘‘standard message set’’ (see [8, Rmk. 3.5]). In this setting it holds exactly that $|V|^{-1} \log Z_G(\beta, B) = \Phi_{\mu_G}(\beta, B)$ [8, Propn. 3.7]. The heuristic then is that for G_n locally tree-like, the (normalized) free energy ϕ_n is approximated by a similar expression for n large. We emphasize that *no* averaging over the vertices of the tree T takes place in the definition of Φ_T : instead, the averaging of $\Phi_{(G,v)}$ over the vertices $v \in G$ in the evaluation of $\Phi_{\mu_G}(\beta, B)$ corresponds to the averaging with respect

to the law μ in the evaluation of the Bethe prediction $\Phi_\mu(\beta, B)$. Indeed, for non-amenable $T \in \mathcal{T}_\bullet$ the sub-trees T^t typically do *not* converge locally weakly to T (see e.g. [8, Lem. 2.8]).

Here is a terminology which we adopt throughout the paper:

Definition 1.7. If G is any graph and U a sub-graph, the *external boundary* ∂U of U is the set of vertices of $G \setminus U$ adjacent to U . Let U^+ denote the sub-graph of G induced by the vertices in $V_U \cup \partial U$. For U finite (so U^+ is finite, since G is locally finite), and ν^\ddagger a measure on $\mathcal{X}^{\partial U}$, the factor model on U with ν^\ddagger *boundary conditions* is the probability measure on configurations $\underline{\sigma}_U \in \mathcal{X}^{V_U}$ given by

$$\nu_{U,G,\underline{\psi}}^\ddagger(\underline{\sigma}_U) \cong \int \prod_{(ij) \in E_{U^+}} \psi(\sigma_i, \sigma_j) \prod_{i \in U} \bar{\psi}(\sigma_i) d\nu^\ddagger(\underline{\sigma}_{\partial U}). \quad (1.11)$$

(Throughout, \cong indicates equivalence up to a positive normalizing constant.) The case in which ν^\ddagger gives probability one to the identically- σ_0 spin configuration on ∂U ($\sigma_0 \in \mathcal{X}$) is referred to as σ_0 *boundary conditions* and denoted $\nu^\ddagger = \nu^{\sigma_0}$, while the case in which ν^\ddagger is uniform measure on $\mathcal{X}^{\partial U}$ is referred to as *free boundary conditions* and denoted $\nu^\ddagger = \nu^f$.

1.2. Application to Ising, Potts, and independent set. Before formally stating our main theorem for general factor models, we mention its consequences in some models of interest: we verify the Bethe prediction for the ferromagnetic Ising model at all temperatures, the ferromagnetic Potts model with field $B \geq 0$ in uniqueness regimes, and the independent set model with low fugacity λ .

1.2.1. Ising model. The *Ising model* is the Potts model (1.3) with $q = 2$. For convenience we use the equivalent formulation which takes $\mathcal{X} = \{\pm 1\}$ and defines the probability measure on \mathcal{X}^V

$$\nu_G^{\beta,B}(\underline{\sigma}) = \frac{1}{Z_G(\beta, B)} \exp \left\{ \beta \sum_{(ij) \in E} \sigma_i \sigma_j + B \sum_{i \in V} \sigma_i \right\}. \quad (1.12)$$

For $T \in \mathcal{T}_\bullet$ let $\bar{h}_T^{t,+} \equiv \bar{h}_T^{t,+, \beta, B}$ denote the root marginal for the Ising model of parameters (β, B) on T^t with $+$ boundary conditions (i.e., with σ_v conditioned to be $+1$ for all v at level $t+1$), and similarly define \bar{h}_T^f corresponding to free boundary conditions. For $\ddagger \in \{f, +\}$ let $\bar{h}_T^\ddagger \equiv \bar{h}_T^{\ddagger, \beta, B} \equiv \lim_{t \rightarrow \infty} \bar{h}_T^{t, \ddagger, \beta, B}$. (Existence of the limits \bar{h}_T^f, \bar{h}_T^+ for the Ising model is an easy consequence of Griffiths's inequality; see Lem. 4.1.) We then define messages $h^\ddagger \in \mathcal{H}_\mu$ by

$$h_{x \rightarrow y}^\ddagger = \bar{h}_{T_{x \rightarrow y}}^\ddagger,$$

for $T_{x \rightarrow y}$ as defined in Rmk. 1.5. For $G_n \rightarrow_{lwc} \mu$, the Bethe free energy prediction for the Ising model with $\beta \geq 0$, $B > 0$ is $\phi(\beta, B) = \Phi_\mu(\beta, B, h^+)$. This prediction was verified in [9, Thm. 2.4] for uniformly sparse graph sequences converging locally weakly to Galton-Watson trees subject to the second-moment condition $\mathbb{E}_\mu[D_o^2] < \infty$, which was relaxed in [11] to a $(1+\epsilon)$ -moment condition. We have the following generalization of this result to an arbitrary limiting law.

Theorem 1.8. *For the Ising model (1.12) on $G_n \rightarrow_{lwc} \mu$,*

$$\phi(\beta, B) = \Phi_\mu(\beta, B, h^f) = \Phi_\mu(\beta, B, h^+)$$

for $\beta \geq 0$, $B > 0$. Also $\phi(\beta, B) = \phi(\beta, -B)$ and $\phi(\beta, 0) = \lim_{B \rightarrow 0} \phi(\beta, B)$.

1.2.2. *Potts model.* Throughout the remainder let $(\beta_0, B_0) \leq (\beta_1, B_1)$, where \leq means coordinate-wise less than or equal to. An *interpolation path* is a piecewise linear path, with each piece parallel to a coordinate axis, increasing from (β_0, B_0) to (β_1, B_1) with respect to the partial order \leq .

We restrict our attention to the Potts model with $\beta, B \geq 0$. In this regime we are able to use a random-cluster representation to extract important monotonicity properties. For $T \in \mathcal{T}_\bullet$ and $\ddagger \in \{f, 1\}$ let $\bar{h}_T^{\ddagger} \equiv \bar{h}_T^{\ddagger, \beta, B}$ denote the root marginal for the Potts model on T^t with \ddagger boundary conditions. Let $\bar{h}_T^{\ddagger} \equiv \bar{h}_T^{\ddagger, \beta, B} \equiv \lim_{t \rightarrow \infty} \bar{h}_T^{t, \ddagger, \beta, B}$ (existence of the limits \bar{h}_T^f, \bar{h}_T^1 for the Potts model follows from monotonicity properties of the random-cluster representation; see Cor. 4.4). We then define messages $h^\ddagger \in \mathcal{H}_\mu$ by $h_{x \rightarrow y}^\ddagger = \bar{h}_{T_{x \rightarrow y}}^\ddagger$, and let

$$\mathcal{R}_\mu \equiv \{(\beta, B) : h^f = h^1, \mu^\uparrow\text{-a.s.}\}.$$

We also define

$$\mathcal{R}_\infty \equiv (\{0\} \times \mathbb{R}_{\geq 0}) \cup (\mathbb{R}_{\geq 0} \times \{\infty\}) \cup (\{\infty\} \times \mathbb{R}_{> 0}).$$

Theorem 1.9. *For the Potts model (1.3) with $q > 2$ and $\beta, B \geq 0$ on $G_n \rightarrow_{lwc} \mu$, the following hold (with $\Phi \equiv \Phi_\mu$, $\mathcal{R} \equiv \mathcal{R}_\mu$):*

(a) *If there exists an interpolation path contained in \mathcal{R} joining (β, B) and \mathcal{R}_∞ , then*

$$\phi(\beta, B) = \Phi(\beta, B, h^f) = \Phi(\beta, B, h^1).$$

(b) *If there exists an interpolation path from (β_0, B_0) to (β_1, B_1) along which h^f is continuous (in the interpolation parameter), then*

$$\liminf_{n \rightarrow \infty} [\phi_n(\beta_1, B_1) - \phi_n(\beta_0, B_0)] \geq \Phi(\beta_1, B_1, h^f) - \Phi(\beta_0, B_0, h^f).$$

If h^f is replaced with h^1 then we have instead

$$\limsup_{n \rightarrow \infty} [\phi_n(\beta_1, B_1) - \phi_n(\beta_0, B_0)] \leq \Phi(\beta_1, B_1, h^1) - \Phi(\beta_0, B_0, h^1).$$

We obtain more explicit results when the limiting tree is the d -regular tree \mathbb{T}_d .

Theorem 1.10. *For the Potts model (1.3) with $q > 2$ and $\beta, B \geq 0$ on $G_n \rightarrow_{lwc} \mathbb{T}_d$, the following hold (with $\Phi \equiv \Phi_{\mathbb{T}_d}$, $\mathcal{R} \equiv \mathcal{R}_{\mathbb{T}_d}$, and $\mathcal{R}_\neq \equiv \{\beta, B \geq 0\} \setminus \mathcal{R}$):*

(a) *If $d = 2$, $\mathcal{R}_\neq = \emptyset$. If $d > 2$ and $q = 2$, there exists $0 < \beta_- < \infty$ such that $\mathcal{R}_\neq = \{B = 0, \beta > \beta_-\}$. If $d > 2$ and $q > 2$, there are smooth curves $\beta_f(B)$ and $\beta_+(B)$ defined on $[0, B_+]$ with $\beta_f(B_+) = \beta_+(B_+)$ such that*

$$\mathcal{R}_\neq = \{B = 0, \beta \geq \beta_f(0)\} \cup \{0 < B < B_+, \beta \in [\beta_f(B), \beta_+(B)]\}.$$

(b) *For $(\beta, B) \notin \mathcal{R}_\neq$, $\phi(\beta, B) = \Phi(\beta, B, h^f) = \Phi(\beta, B, h^1)$. If $(\beta, B) \in \partial \mathcal{R}_\neq$ with $\beta = \beta_f(B)$, then $\phi(\beta, B) = \Phi(\beta, B, h^f)$. If $(\beta, B) \in \partial \mathcal{R}_\neq$ with $\beta \geq \beta_+(B)$, then $\phi(\beta, B) = \Phi(\beta, B, h^1)$. For (β, B) in the interior \mathcal{R}_\neq° of \mathcal{R}_\neq ,*

$$\begin{aligned} \max\{\Phi(\beta, B, h^1), \Phi(\beta, B, h^f)\} &\leq \liminf_{n \rightarrow \infty} \phi_n(\beta, B) \\ &\leq \limsup_{n \rightarrow \infty} \phi_n(\beta, B) \leq \min\{\tilde{\Phi}^f(\beta, B), \tilde{\Phi}^1(\beta, B)\} \end{aligned}$$

where

$$\begin{aligned} \tilde{\Phi}^1(\beta, B) &\equiv \Phi(\beta_f(B), B, h^f) + [\Phi(\beta, B, h^1) - \Phi(\beta_f(B), B, h^1)], \\ \tilde{\Phi}^f(\beta, B) &\equiv \Phi(\beta_+(B), B, h^1) - [\Phi(\beta_+(B), B, h^f) - \Phi(\beta, B, h^f)]. \end{aligned}$$



 FIGURE 1. Potts model with $q = 30, d = 10$

Fig. 1 illustrates the situation of Thm. 1.10 with $q = 30, d = 10$. Fig. 1a shows the recursion (1.15) at $\beta = 0.7, B = 0.2$, restricted to those h which are symmetric among the spins $\neq 1$, parametrized by the log-likelihood ratio $r \equiv \log h(1) - \log h(2)$. The fixed point at $r = 0$ corresponds to h^f while the uppermost fixed point corresponds to h^1 , and both fixed points are stable. Fig. 1b shows the region \mathcal{R}_{\neq} in which this situation occurs.

1.2.3. *Independent set model.* We consider the independent set model (1.4) in the regime of low fugacity. For $\dagger \in \{0, 1\}$ let $\bar{h}_T^{t, \dagger} \equiv \bar{h}_T^{t, \dagger, \lambda}$ denote the root marginal on T^t with \dagger boundary conditions on ∂T^t (i.e. at level $t+1$). Let $\bar{h}_T^{\dagger} \equiv \lim_{t \rightarrow \infty} \bar{h}_T^{2t-1, \dagger}$ (existence of the limits \bar{h}_T^0, \bar{h}_T^1 for the independent set model follows from anti-monotonicity; see §2.4). We then define messages $h^{\dagger} \in \mathcal{H}_{\mu}$ by $h_{x \rightarrow y}^{\dagger} = \bar{h}_{T_{x \rightarrow y}}^{\dagger}$, and let

$$\lambda_c \equiv \lambda_{c, \mu} \equiv \inf\{\lambda \geq 0 : \mu^{\uparrow}(h_{x \rightarrow y}^{0, \lambda} = h_{x \rightarrow y}^{1, \lambda}) < 1\}$$

denote the uniqueness threshold. For $T \in \mathcal{T}_{\bullet}$, we use $\text{br } T$ to denote the branching number of T (see [28, §2]).

Theorem 1.11. *Consider the independent set model (1.4) on $G_n \rightarrow_{\text{hwc}} \mu$, and write $\lambda_c \equiv \lambda_{c, \mu}$.*

(a) *If $\lambda < \lambda_c$ and h has total variation bounded by a deterministic constant on $[0, \log \lambda]$, then*

$$\phi(\lambda) = \Phi_{\mu}(\lambda, h^0) = \Phi_{\mu}(\lambda, h^1), \quad (1.13)$$

which converges to $\phi(\lambda_c)$ as $\lambda \uparrow \lambda_c$.

(b) *If $\text{br } T_{x \rightarrow y} \leq \Delta - 1$ μ^{\uparrow} -a.s. for Δ a deterministic constant, then (1.13) holds for $\lambda < \lambda_c$ with $\lambda(\Delta - 2) < 1$.*

(c) *If $\mu = \delta_{\mathbb{T}_d}$, then (1.13) holds for $\lambda \leq \lambda_c$.*

For the d -regular tree \mathbb{T}_d , the uniqueness threshold $\lambda_c(d)$ is $(d-1)^{d-1}/(d-2)^d$ (see [25, §2]), and [39, Thm. 2.3] shows that \mathbb{T}_d has the lowest value of λ_c among trees with maximum degree at most d . The identity (1.13) has been proved in the case that the G_n are random d -regular graphs [2, 3]. It is also suggested by Weitz's PTAS for $Z_G(\lambda)$ on a finite graph G of maximum degree Δ and with $\lambda < \lambda_c(\Delta)$ [39, Cor. 2.8]. For μ a unimodular measure on \mathcal{T}_{\bullet} giving a local tree approximation to G (in the sense of Defn. 1.1), $\lambda_{c, \mu}$ is often an improvement over $\lambda_c(\Delta)$, making it possible to compute

$\phi(\lambda)$ above $\lambda_c(\Delta)$ provided (H3^B) can be verified. However, there is strong evidence that no PTAS exists for all graphs of maximum degree Δ when $\lambda > \lambda_c(\Delta)$ [31, 36].

1.3. Results for general factor models. We now state our results for the factor model (1.1). With the convention $\log 0 \equiv -\infty$, let $\log \psi \equiv \xi$ and $\log \bar{\psi} \equiv \bar{\xi}$, and impose the regularity condition

(H1) The specification is *permissive*, that is, $\bar{\psi}(\sigma) > 0$ for all $\sigma \in \mathcal{X}$, and there exists a “permitted state” $\sigma^p \in \mathcal{X}$ such that $\min_{\sigma} \psi(\sigma, \sigma^p) > 0$.

For any $\sigma \in \mathcal{X}$, $\bar{\xi}^B(\sigma)$ is continuously differentiable in B . For any $\sigma, \sigma' \in \mathcal{X}$, $\xi^\beta(\sigma, \sigma')$ is either identically $-\infty$ over all β , or finite and continuously differentiable in β .

Recall Defn. 1.4 of the message space \mathcal{H} , and consider the *belief propagation (BP) recursion* defined by $\text{BP} \equiv \text{BP}^{\beta, B} : \mathcal{H} \rightarrow \mathcal{H}$,

$$(\text{BP}^{\beta, B} h)_{x \rightarrow y}(\sigma) \equiv \frac{1}{z_{x \rightarrow y}(\beta, B)} \bar{\psi}^B(\sigma) \prod_{v \in \partial x \setminus y} \left(\sum_{\sigma_v} \psi^\beta(\sigma, \sigma_v) h_{v \rightarrow x}(\sigma_v) \right), \quad (1.14)$$

with $z_{x \rightarrow y}(\beta, B)$ normalizing constants. For μ a measure on \mathcal{T}_\bullet and fixed (β, B) , let $\mathcal{H}_\mu^*(\beta, B)$ denote the space of measurable functions $h : \mathcal{T}_e \rightarrow \Delta_{\mathcal{X}}$, again taken up to μ^\uparrow -equivalence, which are fixed points of the BP recursion: that is, satisfying

$$h = \text{BP}^{\beta, B} h, \quad \mu^\uparrow\text{-a.s.} \quad (1.15)$$

For $h \in \mathcal{H}_\mu$ we can define $\bar{h} : \mathcal{T}_\bullet \rightarrow \Delta_{\mathcal{X}}$ up to μ -equivalence by

$$\bar{h}_T(\sigma) \cong \bar{\psi}(\sigma) \prod_{j \in \partial o} \left(\sum_{\sigma_j} \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j) \right). \quad (1.16)$$

In particular, if $h \in \mathcal{H}_\mu^*(\beta, B)$ and $T \in \mathcal{T}_\bullet^+$, then clearly we have

$$\bar{h}_T(\sigma) \cong \sum_{\sigma_j} \psi(\sigma, \sigma_j) h_{o \rightarrow j}(\sigma) h_{j \rightarrow o}(\sigma_j) \quad (1.17)$$

independently of the choice of $j \in \partial o$. From now on, for $h \in \mathcal{H}_\mu$, we will write $h \in \mathcal{H}^*$ to indicate that $h^{\beta, B} \in \mathcal{H}_\mu^*(\beta, B)$ for (β, B) in the range being considered.

Remark 1.12. The elements of \mathcal{H}^* are consistent with the recursion structure of the tree in the following precise sense: for $T \in \mathcal{T}_\bullet$ and U a finite connected subgraph of T , consider the factor model $\nu_{U, T}^h$ on U with boundary conditions $\sigma_v \sim h_{v \rightarrow p(v)}$ independently for $v \in \partial U$, where $p(v)$ denotes the (necessarily unique) neighbor of v inside U . Then the marginal of $\nu_{T^t, T}^h$ on T^{t-1} is exactly the factor model $\nu_{T^{t-1}, T}^{\text{BP}^h}$ on T^{t-1} with boundary conditions $\sigma_u \sim (\text{BP}^h)_{u \rightarrow p(u)}$ independently for $u \in \partial T^{t-1}$, including any u which are leaves of T^t . This statement remains valid if ∂T^t or even ∂T^{t-1} is empty, since if $\partial T^t = \emptyset$ then $\nu_{T^t, T}^h$ is simply ν_T as defined by (1.1). Continuing the BP recursion up the tree, we see that $h \in \mathcal{H}^*$ implies that the marginal law of σ_o will be \bar{h}_T as defined by (1.16). From this it is easy to see that the measures $\nu_{U, T}^h$ form a consistent family of finite-dimensional marginals, so by the Kolmogorov consistency theorem they uniquely determine a probability measure $\nu_T \equiv \nu_T^h$ belonging to \mathcal{G}_T , the set of Gibbs measures (or Markov random fields) associated to the specification $\underline{\psi} \equiv (\psi, \bar{\psi})$ on T .¹ (In fact this mapping is one-to-one, e.g. by Rmk. 2.3 below.) Each ν_T belongs to a special class of measures in \mathcal{G}_T which are called *Markov chains* or *splitting Gibbs measures* in the literature, and the entire collection $(\nu_T)_{T \in \mathcal{T}_\bullet}$ arising from

¹Strictly speaking the term “Gibbs measures” refers to the case $\psi > 0$, but we will follow common practice and say Gibbs measures also for the general case. For the general theory of Gibbs measures see e.g. [17].

$h \in \mathcal{H}_\mu^*$ has a consistency property which leads us to term them “unimodular Markov chains” (see §2.3).

In this general setting, the *Bethe prediction* is the supremum of $\Phi_\mu(\beta, B, h)$ over $\mathcal{H}_\mu^*(\beta, B)$ (cf. Rmk. 1.6). (It will be shown in Lem. 2.2 that Φ_μ is uniformly bounded on $\mathcal{H}_\mu^*(\beta, B)$ subject to $\mathbb{E}_\mu[D_o^2] < \infty$; if further $\psi > 0$ then Φ_μ is in fact uniformly bounded on \mathcal{H} subject only to $\mathbb{E}_\mu[D_o] < \infty$.) We define the following integrability condition for unimodular measures μ on \mathcal{T}_\bullet (not necessarily arising from a graph sequence):

(H2) The probability measure μ on \mathcal{T}_\bullet satisfies $\mathbb{E}_\mu[D_o] < \infty$. If ψ is not everywhere positive, then furthermore $\mathbb{E}_\mu[e^{cD_o}] < \infty$ for all $c \in \mathbb{R}$.

Note that if $G_n \rightarrow_{lwc} \mu$ and $\psi > 0$ then (H2) holds trivially by the assumption of uniform sparsity. We will in fact justify our interpolation scheme under a weaker assumption than (H2); for the exact condition see (H2 $^\beta$), (H2 B) in §2.2.

1.3.1. *Bethe interpolation.* We will deduce the results of §1.2 from the abstract interpolation method given by Thm. 1.13 below, which bounds differences of $\phi(\beta, B)$ by differences of $\Phi(\beta, B, h)$ ($h \in \mathcal{H}^*$) when the limiting expectation of a certain edge or vertex functional in the finite graph (capturing respectively $\partial_\beta \phi_n$ or $\partial_B \phi_n$) is bounded by the expectation of an analogous functional on the infinite tree.

To be more precise, recall that I_n denotes a uniformly random vertex of V_n . Let $\langle \cdot \rangle_n^{\beta, B}$ denote expectation with respect to $\nu_{G_n, \underline{\psi}}$, conditioned on G_n . For $h \in \mathcal{H}_\mu^*(\beta, B)$ and $T \in \mathcal{T}_\bullet$, let $\llbracket \cdot \rrbracket_T^{h, \beta, B}$ denote expectation with respect to ν_T^h (as defined in Rmk. 1.12), conditioned on T , and define

$$\begin{aligned} a_n^e(\beta, B) &\equiv \frac{1}{2} \mathbb{E}_n \left[\sum_{j \in \partial I_n} \langle \partial_\beta \xi(\sigma_{I_n}, \sigma_j) \rangle_n^{\beta, B} \right], & a^e(\beta, B, h) &\equiv \frac{1}{2} \mathbb{E}_\mu \left[\sum_{j \in \partial o} \llbracket \partial_\beta \xi(\sigma_o, \sigma_j) \rrbracket_T^{h, \beta, B} \right], \\ a_n^{\text{vx}}(\beta, B) &\equiv \mathbb{E}_n [\langle \partial_B \bar{\xi}(\sigma_{I_n}) \rangle_n^{\beta, B}], & a^{\text{vx}}(\beta, B, h) &\equiv \mathbb{E}_\mu \left[\llbracket \partial_B \bar{\xi}(\sigma_o) \rrbracket_T^{h, \beta, B} \right]. \end{aligned}$$

The left-hand expressions are the derivatives $\partial_\beta \phi_n, \partial_B \phi_n$ (Lem. 2.1). The right-hand expressions are the infinite-tree analogues, which we will show in Propn. 2.4 may be thought of as derivatives in β and B of Φ_μ .

For interpolation in β on a compact interval $[\beta_0, \beta_1]$ using $h \in \mathcal{H}^*$, we require the following regularity condition:

(H3 $^\beta$) On $[\beta_0, \beta_1]$, μ^\uparrow -a.s. the functions $\beta \mapsto h_{j \rightarrow o}^\beta(\sigma)$ are continuous with total variation bounded by a deterministic constant depending only on β_0, β_1 .

We define the analogous condition (H3 B) for interpolation in B on a compact interval $[B_0, B_1]$. The condition of boundedness in total variation is implied for example whenever the functions h are (anti-)monotone in the interpolation parameter.

Theorem 1.13. *Let $\underline{\psi} \equiv (\psi, \bar{\psi})$ specify a factor model (1.1) on $G_n \rightarrow_{lwc} \mu$ such that (H1) and (H2) are satisfied.*

(a) *If on $[\beta_0, \beta_1]$ we have $h \in \mathcal{H}^*$ satisfying (H3 $^\beta$), and*

$$\limsup_{n \rightarrow \infty} a_n^e(\beta, B) \leq a^e(\beta, B, h), \tag{1.18}$$

then $\limsup_{n \rightarrow \infty} [\phi_n(\beta_1, B) - \phi_n(\beta_0, B)] \leq \Phi(\beta_1, B, h) - \Phi(\beta_0, B, h)$.

(b) *If on $[B_0, B_1]$ we have $h \in \mathcal{H}^*$ satisfying (H3 B), and*

$$\limsup_{n \rightarrow \infty} a_n^{\text{vx}}(\beta, B) \leq a^{\text{vx}}(\beta, B, h), \tag{1.19}$$

then $\limsup_{n \rightarrow \infty} [\phi_n(\beta, B_1) - \phi_n(\beta, B_0)] \leq \Phi(\beta, B_1, h) - \Phi(\beta, B_0, h)$.

The same results hold if all inequalities are reversed, replacing limit superior with inferior.

The conditions (1.18), (1.19) (and their reverses) are automatically verified in the following special case:

Theorem 1.14. *Let $\underline{\psi} \equiv (\psi, \bar{\psi})$ specify a factor model (1.1) on $G_n \rightarrow_{\text{luc}} \mu$ satisfying (H1) and (H2). We say that uniqueness holds if \mathcal{G}_T at (β, B) consists of a single measure ν_T , μ -a.s. In this case, $\mathcal{H}_\mu^*(\beta, B)$ is a singleton.*

(a) *If on $[\beta_0, \beta_1] \times \{B\}$ uniqueness holds and the unique element $h \in \mathcal{H}^*$ satisfies (H3 $^\beta$), then*

$$\lim_{n \rightarrow \infty} [\phi_n(\beta_1, B) - \phi_n(\beta_0, B)] = \Phi(\beta_1, B, h) - \Phi(\beta_0, B, h).$$

(b) *If on $\{\beta\} \times [B_0, B_1]$ uniqueness holds and the unique element $h \in \mathcal{H}^*$ satisfies (H3 B), then*

$$\lim_{n \rightarrow \infty} [\phi_n(\beta, B_1) - \phi_n(\beta, B_0)] = \Phi(\beta, B_1, h) - \Phi(\beta, B_0, h).$$

Uniqueness for \mathcal{G}_T corresponds to the vanishing effect of boundary conditions on ∂T^t as $t \rightarrow \infty$ [17, Ch. 7]. Dobrushin's uniqueness theorem (see e.g. [35]) gives a sufficient condition for uniqueness to hold, together with a bound on the rate of convergence of the root marginal in T^t to the limit as $t \rightarrow \infty$. Note that if the convergence rate is uniform in β, B then the continuity required in (H3 $^\beta$) and (H3 B) immediately follows. We will obtain continuity in uniqueness regimes via a different route, making use of certain monotonicity properties (see the proof of Thm. 1.8).

1.3.2. Variational principle. We further develop the theory by providing a variational principle for the Bethe prediction: we express $\Phi_\mu(\beta, B)$ as an optimum of a function $\Phi_\mu(\beta, B, \mathbf{h})$ defined for \mathbf{h} in a larger space \mathcal{H}_{loc} which, unlike $\mathcal{H}_\mu^*(\beta, B)$, is *independent* of β, B . This alternative characterization of Φ_μ is the infinite-tree analogue of the finite-graph optimization problem that is considered in [41]. Recall from §1.1 that \mathcal{T}_e denotes the space of trees rooted at a *directed* edge.

Definition 1.15. The *local polytope* $\mathcal{H}_{\text{loc}} \equiv \mathcal{H}_{\text{loc}, \mu}$ is the space of measurable functions

$$\mathbf{h} : \mathcal{T}_e \rightarrow \Delta_{\mathcal{X}^2}, \quad (T, x \rightarrow y) \mapsto \mathbf{h}_{(T, x \rightarrow y)} \equiv \mathbf{h}_{xy},$$

taken up to μ^\uparrow -equivalence, such that

- (i) $\mathbf{h}_{xy}(\sigma, \sigma') = \mathbf{h}_{yx}(\sigma', \sigma)$ for all $\sigma, \sigma' \in \mathcal{X}$, and
- (ii) for $T \in \mathcal{T}_\bullet^+$, the one-point marginal $\bar{h}_x(\sigma) \equiv \bar{h}_{(T, x)}(\sigma) \equiv \sum_{\sigma_y} \mathbf{h}_{xy}(\sigma_x, \sigma_y)$ is well-defined independently of the choice of $y \in \partial x$.

If $T = T^0$, we define $\bar{h}_T(\sigma) \cong \bar{\psi}(\sigma)$, in accordance with (1.16). Denote by $\mathcal{H}_{\text{loc}}^\circ$ the subset of $\mathbf{h} \in \mathcal{H}_{\text{loc}}$ for which μ^\uparrow -a.s. \mathbf{h}_{xy} is strictly positive on \mathcal{X}^2 .

For *fixed* (β, B) , by symmetry of ψ^β and (1.17), the space $\mathcal{H}_\mu^*(\beta, B)$ has a natural mapping into \mathcal{H}_{loc} given by

$$h \mapsto \mathbf{h}, \quad \mathbf{h}_{xy}(\sigma, \sigma') \cong \psi(\sigma, \sigma') h_{x \rightarrow y}(\sigma) h_{y \rightarrow x}(\sigma'). \quad (1.20)$$

With $\underline{\psi}$ permissive this is in fact an embedding; see Rmk. 2.3. We define the *Bethe free energy functional* on \mathcal{H}_{loc} by

$$\Phi_\mu(\mathbf{h}) \equiv \mathbb{E}_\mu \left[\langle \bar{\xi}(\sigma_o) \rangle_{\bar{h}_o} - (D_o - 1) H(\bar{h}_o) \right] + \frac{1}{2} \mathbb{E}_\mu \left[\sum_{j \in \partial o} \left\{ \langle \xi(\sigma_o, \sigma_j) \rangle_{\mathbf{h}_{oj}} + H(\mathbf{h}_{oj}) \right\} \right], \quad (1.21)$$

where $H(p)$ denotes the Shannon entropy $-\sum_k p_k \log p_k$ for p a probability measure on a finite space. This is an infinite-tree analogue of the definition of [41, (37)-(38)] for finite graphs. With

the usual conventions $\log 0 \equiv -\infty$, $0 \log 0 \equiv 0$ and $0 \log(0/0) \equiv 0$, Φ_μ is bounded above on \mathcal{H}_{loc} whenever $\mathbb{E}_\mu[D_o] < \infty$, and we show in Lem. 3.1 that for μ unimodular, this Φ_μ extends the previous definition (1.9) on \mathcal{H}^* (under the embedding (1.20)), provided the latter is finite. Furthermore, writing $H(q \parallel p)$ for the relative entropy $\sum_k q_k \log(q_k/p_k)$ between q and p (well-defined for any non-negative reference measure p), for μ unimodular we can alternatively express

$$\Phi_\mu(\mathbf{h}) = -\mathbb{E}_\mu[H(\bar{h}_o \parallel \bar{\psi})] - \frac{1}{2} \mathbb{E}_\mu \left[\sum_{j \in \partial o} H(\mathbf{h}_{oj} \parallel \psi) \right] - \mathbb{E}_\mu [D_o H(\bar{h}_o)] \quad (1.22)$$

$$\begin{aligned} &= -\mathbb{E}_\mu[H(\bar{h}_o \parallel \bar{\psi})] - \frac{1}{2} \mathbb{E}_\mu \left[\sum_{j \in \partial o} \left\{ H(\mathbf{h}_{oj} \parallel \psi) + H(\bar{h}_o) + H(\bar{h}_j) \right\} \right] \\ &= -\mathbb{E}_\mu[H(\bar{h}_o \parallel \bar{\psi})] - \frac{1}{2} \mathbb{E}_\mu \left[\sum_{j \in \partial o} H(\mathbf{h}_{oj} \parallel \bar{h}_o \times_\psi \bar{h}_j) \right], \end{aligned} \quad (1.23)$$

where $(\bar{h}_o \times_\psi \bar{h}_j)(\sigma_o, \sigma_j) \equiv \bar{h}_o(\sigma_o) \psi(\sigma_o, \sigma_j) \bar{h}_j(\sigma_j)$, and unimodularity is used in the second identity.

This extended definition of Φ_μ provides the following variational principle for the Bethe free energy:

Theorem 1.16. *Let $\underline{\psi} \equiv (\psi, \bar{\psi})$ specify a factor model (1.1) satisfying (H1), and let μ be a unimodular measure on $\overline{\mathcal{T}}_\bullet$ with $\mathbb{E}_\mu[D_o] < \infty$.*

- (a) $\tilde{\Phi}_\mu(\beta, B) \equiv \sup_{\mathbf{h} \in \mathcal{H}_{\text{loc}}} \Phi_\mu(\beta, B, \mathbf{h})$ is continuous in (β, B) .
- (b) If further $\psi > 0$, then any local maximizer of $\Phi_\mu(\beta, B)$ belongs to $\mathcal{H}_{\text{loc}}^\circ$. Any stationary point of $\Phi_\mu(\beta, B)$ belonging to $\mathcal{H}_{\text{loc}}^\circ$ is the image under (1.20) of an element of $\mathcal{H}_\mu^*(\beta, B)$. In particular, if Φ_μ attains its supremum on \mathcal{H}_{loc} , then

$$\tilde{\Phi}_\mu(\beta, B) = \max_{h \in \mathcal{H}_\mu^*(\beta, B)} \Phi_\mu(\beta, B, h) \equiv \Phi_\mu^{\text{Bethe}}(\beta, B),$$

so that the Bethe free energy is also continuous in (β, B) .

Although we do not pursue this point, we mention that even in specific models where the abstract definition of Φ^{Bethe} is supplanted by $\Phi(\beta, B, h)$ for some “naturally” distinguished h , an adaptation of Thm. 1.16 (involving a restricted subspace of \mathcal{H}_{loc} which is independent on (β, B)), may be relevant.

Remark 1.17. In case $G_n \rightarrow_{lwc} \mathbb{T}_d$ the d -regular tree, \mathcal{H}_{loc} is parametrized by a single measure \mathbf{h}_{xy} on \mathcal{X}^2 whose one-point marginals are required to agree, and the formula (1.23) simplifies to

$$-\Phi_\mu(\mathbf{h}) = H(\bar{h}_0 \parallel \bar{\psi}) + \frac{d}{2} H(\mathbf{h}_{01} \parallel \bar{h}_0 \times_\psi \bar{h}_1). \quad (1.24)$$

For $\underline{\sigma} \in \mathcal{X}^{V_n}$ let $L_n^{\text{vx}} \equiv n^{-1} \sum_{i \in V_n} \delta_{\sigma_i}$ and $L_n^e \equiv (2|E_n|)^{-1} \sum_{(ij) \in E_n} [\delta_{(\sigma_i, \sigma_j)} + \delta_{(\sigma_j, \sigma_i)}]$ denote the induced empirical and pair empirical measures respectively. If G_n is d -regular, then the one-point marginals of L_n^e coincide with L_n^{vx} , and

$$\phi_n = \log |\mathcal{X}| + \frac{1}{n} \mathbb{E}_n \left[\log \mathbb{E}_{\bar{u}_n} \exp \left\{ n \langle \bar{\xi} \rangle_{L_n^{\text{vx}}} + \frac{nd}{2} \langle \xi \rangle_{L_n^e} \right\} \right]$$

where the law of $\underline{\sigma}$ is the uniform measure \bar{u}_n on $\mathcal{X}^{[n]}$ and $\mathbb{E}_{\bar{u}_n}$ denotes expectation with respect to \bar{u}_n (with G_n fixed).

If (G_n) is an independent sequence of uniformly random d -regular graphs and $\underline{\sigma}_n \sim \bar{u}_n$, one might guess that for a.e. (G_n) the induced sequence L_n^e satisfies a large deviation principle (LDP)

with good rate function

$$I(\mathbf{h}_{01}) = H(\bar{h}_0 \parallel \bar{u}) + \frac{d}{2}H(\mathbf{h}_{01} \parallel \bar{h}_0 \times \bar{h}_1), \quad (1.25)$$

where $\bar{u} \equiv \bar{u}_1$. If this were the case, it would be an immediate consequence of Varadhan's lemma (see [10, §4.3.1]) that $\phi_n \rightarrow \tilde{\Phi}_\mu(\beta, B)$ (as defined in Thm. 1.16) for *any* factor model satisfying (H1). However, for many of these models the Bethe prediction is known to fail at low temperature for $d \geq 3$. So, while Thm. 1.16 suggests a potential connection to large deviations theory, such a connection would be highly non-trivial and applicable only in certain regimes of (β, B) .

One special case in which everything trivializes is the (rooted) infinite line \mathbb{T}_2 , the local weak limit of the simple path G_n on n vertices. In this case \bar{u}_n may be viewed as the law of a stationary reversible Markov chain on \mathcal{X} with transitions $q(\sigma, \sigma') = \bar{u}(\sigma')$ and reversing measure \bar{u} , and it is well-known (see e.g. [10, Thm. 3.1.13]) that the associated pair empirical measure L_n^e satisfies an LDP with good rate function $I(\mathbf{h}_{01}) = H(\mathbf{h}_{01}(\sigma, \sigma') \parallel \bar{h}_0(\sigma)q(\sigma, \sigma'))$ which matches (1.25). The implication of Varadhan's lemma is also easy to see: a factor model on the simple path G_n with general positive specification $\underline{\psi}$ corresponds in the limit $n \rightarrow \infty$ to a reversible Markov chain with transition kernel p and positive reversing measure π given by

$$p(\sigma, \sigma') = \frac{1}{\rho} \tilde{\psi}(\sigma, \sigma') \frac{m(\sigma')}{m(\sigma)}, \quad \pi(\sigma) = m(\sigma)^2,$$

where ρ and m are the Perron-Frobenius eigenvalue and eigenvector of the symmetric positive $|\mathcal{X}|$ -dimensional matrix with entries $\tilde{\psi}(\sigma, \sigma') \equiv \psi(\sigma, \sigma') \bar{\psi}(\sigma)^{1/2} \bar{\psi}(\sigma')^{1/2}$. The Bethe free energy functional (1.24) is then maximized at $\mathbf{h}_{01}(\sigma, \sigma') = \tilde{\psi}(\sigma, \sigma') m(\sigma) m(\sigma') / \rho$, where it takes the value $\Phi_\mu(\mathbf{h}) = \log \rho$ which coincides with ϕ by the Perron-Frobenius theorem (see e.g. [10, Thm. 3.1.1]).

Outline of the paper.

- In §2 we prove the abstract interpolation results. §2.1 presents some preliminary lemmas which will be useful in our proofs. Our main result for abstract factor models, Thm. 1.13, is proved in §2.2. §2.3 contains the specialization of this theorem to the uniqueness case (Thm. 1.14) and also contains discussion on unimodular Markov chains. §2.4 shows how to deduce our result for independent set (Thm. 1.11) from Thm. 1.13.
- In §3 we prove the variational characterization for the Bethe prediction (Thm. 1.16).
- §4 contains applications of our abstract results to the Ising and Potts models. In §4.1 we prove Thm. 1.8, generalizing the results of [9, 11]. In §4.2 we prove Thm. 1.9 by appealing to a random-cluster representation. Finally §4.3 analyzes the d -regular case and proves Thm. 1.10.

Acknowledgments. We thank Allan Sly and Ofer Zeitouni for many helpful conversations. A.D. and N.S. thank the Microsoft Research Theory Group for supporting a visit during which part of this work was completed.

2. BETHE INTERPOLATION FOR GENERAL FACTOR MODELS

2.1. Preliminaries. We begin with some straightforward observations on the boundedness of the free energy ϕ_n and the Bethe free energy Φ_μ as defined on \mathcal{H} , and we prove that the mapping (1.20) of \mathcal{H}^* into \mathcal{H}_{loc} is in fact an embedding for permissive specifications.

Lemma 2.1. *For the factor model (1.1) satisfying (H1) on $G_n \rightarrow_{\text{lwc}} \mu$, the functions $\phi_n(\beta, B)$ are uniformly bounded and equicontinuous on compact regions of (β, B) , with*

$$\partial_\beta \phi_n(\beta, B) = \frac{1}{n} \mathbb{E}_n[\partial_\beta \log Z_n(\beta, B)], \quad \partial_B \phi_n(\beta, B) = \frac{1}{n} \mathbb{E}_n[\partial_B \log Z_n(\beta, B)]. \quad (2.1)$$

Further

$$\frac{1}{n} \partial_\beta \log Z_n(\beta, B) = \frac{1}{2} \sum_{j \in \partial I_n} \langle \partial_\beta \xi(\sigma_{I_n}, \sigma_j) \rangle_n^{\beta, B}, \quad \frac{1}{n} \partial_B \log Z_n(\beta, B) = \langle \partial_B \bar{\xi}(\sigma_{I_n}) \rangle_n^{\beta, B},$$

with the convention $\partial_\beta \xi(\sigma, \sigma') \equiv 0$ in case $\xi(\sigma, \sigma') \equiv -\infty$.

Proof. The expressions for $n^{-1} \partial_\beta \log Z_n(\beta, B)$ and $n^{-1} \partial_B \log Z_n(\beta, B)$ are obtained by a straightforward computation. Now note that if $G_n \rightarrow_{lwc} \mu$ then the uniform sparsity assumption gives

$$\frac{1}{n} \mathbb{E}_n[|E_n|] = \frac{1}{2} \mathbb{E}_n[D_{I_n}] \rightarrow \frac{1}{2} \mathbb{E}_\mu[D_o] < \infty. \quad (2.2)$$

Let (β, B) vary within a given compact region. By (H1) we have $\xi, \bar{\xi} \leq \xi_{\max}$ as well as $\xi(\sigma^P, \cdot), \bar{\xi} \geq \xi_{\min}$. Therefore

$$(1 + |E_n|/n) \xi_{\min} \leq n^{-1} \log Z_n(\beta, B) \leq \log |\mathcal{X}| + (1 + |E_n|/n) \xi_{\max}$$

so $\phi_n = n^{-1} \mathbb{E}_n[\log Z_n(\beta, B)]$ is uniformly bounded by uniform sparsity. The exchange of differentiation and integration in (2.1) is justified by Vitali's convergence theorem, in view of the boundedness of $\partial_\beta \xi, \partial_B \bar{\xi}$ and the uniform integrability of $|E_n|/n$. It follows furthermore that $\partial_\beta \phi_n(\beta, B)$ and $\partial_B \phi_n(\beta, B)$ are bounded uniformly in n , from which equicontinuity follows. \square

Lemma 2.2. *Let $\psi \equiv (\psi, \bar{\psi})$ specify a factor model (1.1) satisfying (H1), and let μ be a unimodular measure on \mathcal{T}_\bullet . For any compact region of (β, B) there exists a deterministic constant $C < \infty$ such that*

- (a) $|\Phi_T(\beta, B, h)| \leq C(D_o^2 + 1)$ for any $h \in \mathcal{H}^*$, and
- (b) if further $\psi > 0$, then $|\Phi_T(\beta, B, h)| \leq C(D_o + 1)$ for any $h \in \mathcal{H}$.

Hence, on any compact region of (β, B) , Φ_μ is uniformly bounded on \mathcal{H}_μ^* provided $\mathbb{E}_\mu[D_o^2] < \infty$, and, if $\psi > 0$, uniformly bounded on \mathcal{H}_μ subject only to $\mathbb{E}_\mu[D_o] < \infty$.

Proof. Let ξ_{\min}, ξ_{\max} be as in the proof of Lem. 2.1. Then, for any $h \in \mathcal{H}$,

$$\log |\mathcal{X}| + (D_o + 1) \xi_{\max} \geq \Phi_T^{\text{vx}}(h) \geq (D_o + 1) \xi_{\min}.$$

If $\psi > 0$ then we also have

$$D_o \xi_{\max} \geq 2\Phi_T^e(h) \geq D_o \xi_{\min},$$

so $|\Phi_T(\beta, B, h)| \leq C(D_o + 1)$ on \mathcal{H} , which proves (b). For general permissive ψ , the preceding lower bound on $\Phi_T^e(h)$ may fail, but (1.15) implies that for $h \in \mathcal{H}^*$,

$$\log h_{o \rightarrow j}(\sigma^P) \geq D_o(\xi_{\min} - \xi_{\max}) - \log |\mathcal{X}|, \quad \forall j \in \partial o. \quad (2.3)$$

Therefore,

$$D_o \xi_{\max} \geq 2\Phi_T^e(h) \geq \sum_{j \in \partial o} \left(\xi_{\min} + \log h_{o \rightarrow j}(\sigma^P) \right) \geq D_o(\xi_{\min} - \log |\mathcal{X}|) + D_o^2(\xi_{\min} - \xi_{\max}),$$

which proves (a). \square

Remark 2.3. It is now easy to see that the mapping (1.20) of $\mathcal{H}_\mu^*(\beta, B)$ into \mathcal{H}_{loc} is injective: if $h, h' \in \mathcal{H}$ give rise to the same \mathbf{h} , then

$$h_{x \rightarrow y}(\sigma) h_{y \rightarrow x}(\sigma^P) = z_{x,y} h'_{x \rightarrow y}(\sigma) h'_{y \rightarrow x}(\sigma^P) \quad \forall \sigma \in \mathcal{X},$$

for $z_{x,y}$ a positive scaling factor. If $h, h' \in \mathcal{H}_\mu^*(\beta, B)$, then (2.3) implies that μ^\uparrow -a.s. both $h_{y \rightarrow x}$ and $h'_{y \rightarrow x}$ give positive measure to σ^P . Therefore, μ^\uparrow -a.s. the $|\mathcal{X}|$ -dimensional vectors $h_{x \rightarrow y}$ and $h'_{x \rightarrow y}$

are equivalent up to scaling, and since both are probability measures on \mathcal{X} , we must have $h = h'$ μ^\uparrow -a.s. as claimed.

2.2. Bethe interpolation. We now prove Thm. 1.13(a). The result is for fixed B so we suppress it from the notation. The proof of Thm. 1.13(b) is very similar and will be given in brief at the end of this section.

Our interpolation procedure relies on the proposition below which expresses Φ_μ as the integral of its partial derivative with respect to β only, ignoring the dependence on β through the function h . Recall that although it is suppressed from the notation, $\underline{\psi}$ and h depend on β , and are taken to be evaluated at β in expressions such as $\Phi_T(\beta, B)$. We will prove our result under the following integrability condition, which by (2.3) is a relaxation of (H2):

(H2 $^\beta$) The probability measure μ on \mathcal{T}_\bullet satisfies $\mathbb{E}_\mu[D_o] < \infty$. If ψ is not everywhere positive, then furthermore

$$\mathbb{E}_\mu \left[\sum_{j \in \partial o} \sup_{\beta \in [\beta_0, \beta_1]} \frac{1}{h_{j \rightarrow o}^\beta(\sigma^j)} \right] < \infty.$$

We define the analogous condition (H2 B) on an interval $[B_0, B_1]$.

Proposition 2.4. *Let $\underline{\psi} \equiv (\psi, \bar{\psi})$ be a specification satisfying (H1), and μ a unimodular measure on \mathcal{T}_\bullet . If on $[\beta_0, \beta_1]$ we have $h \in \mathcal{H}^*$ satisfying (H2 $^\beta$) and (H3 $^\beta$), then*

$$\Phi(\beta_1, h) - \Phi(\beta_0, h) = \int_{\beta_0}^{\beta_1} a^e(\beta, h) d\beta.$$

Proof. For fixed $T \in \mathcal{T}_\bullet$ we shall regard Φ_T simply as a function of a vector $(\beta, h_{x \rightarrow y})_{x \rightarrow y \in T^1}$ in $(1 + 2|\mathcal{X}|D_o)$ -dimensional euclidean space (with h depending on β). We begin by computing the partial derivatives of this function with respect to β and h . We abbreviate $\widehat{h}_{o \rightarrow j}^\beta(\sigma) \equiv (\text{BP}^\beta h)_{o \rightarrow j}(\sigma)$ for the belief propagation mapping of (1.14), which for fixed T and each $j \in \partial o$ is a well-defined function on the same euclidean space as Φ_T . Making use of (H1) we find

$$\frac{\partial \Phi_T^{\text{yx}}}{\partial \beta}(\beta, h) = \sum_{j \in \partial o} \frac{\sum_{\sigma, \sigma_j} \partial_\beta \xi(\sigma, \sigma_j) \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j) \widehat{h}_{o \rightarrow j}^\beta(\sigma)}{\sum_{\sigma, \sigma_j} \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j) \widehat{h}_{o \rightarrow j}^\beta(\sigma)}, \quad (2.4)$$

$$\frac{\partial \Phi_T^{(oj)}}{\partial \beta}(\beta, h) = \frac{\sum_{\sigma, \sigma_j} \partial_\beta \xi(\sigma, \sigma_j) \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j) h_{o \rightarrow j}(\sigma)}{\sum_{\sigma, \sigma_j} \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j) h_{o \rightarrow j}(\sigma)}. \quad (2.5)$$

If $h \in \mathcal{H}^*$ then $\widehat{h}^\beta = h$, therefore (recalling the notation $\llbracket \cdot \rrbracket_T^{h, \beta}$ from §1.3.1) we re-express the above as

$$\frac{\partial \Phi_T^{\text{yx}}}{\partial \beta}(\beta, h) = \sum_{j \in \partial o} \llbracket \partial_\beta \xi(\sigma_o, \sigma_j) \rrbracket_T^{h, \beta}, \quad \frac{\partial \Phi_T^{(oj)}}{\partial \beta}(\beta, h) = \llbracket \partial_\beta \xi(\sigma_o, \sigma_j) \rrbracket_T^{h, \beta},$$

and combining gives

$$\frac{\partial \Phi_T}{\partial \beta}(\beta, h) = \frac{1}{2} \sum_{j \in \partial o} \llbracket \partial_\beta \xi(\sigma_o, \sigma_j) \rrbracket_T^{h, \beta} \equiv a_T^e(\beta, h). \quad (2.6)$$

Likewise we compute that for $T \in \mathcal{T}_\bullet^+$,

$$\begin{aligned} \frac{\partial \Phi_T^{\text{vx}}(\beta, h)}{\partial h_{o \rightarrow j}(\sigma)} &= 0, & \frac{\partial \Phi_T^{\text{vx}}(\beta, h)}{\partial h_{j \rightarrow o}(\sigma_j)} &= \widehat{g}_{\sigma_j}^\beta(j \rightarrow o; h) \equiv \frac{\sum_\sigma \psi(\sigma, \sigma_j) \widehat{h}_{o \rightarrow j}^\beta(\sigma)}{\sum_{\sigma', \sigma'_j} \psi(\sigma', \sigma'_j) h_{j \rightarrow o}(\sigma'_j) \widehat{h}_{o \rightarrow j}^\beta(\sigma')}, \\ \frac{\partial \Phi_T^e(\beta, h)}{\partial h_{j \rightarrow o}(\sigma_j)} &= \frac{1}{2} g_{\sigma_j}^\beta(j \rightarrow o; h), & \frac{\partial \Phi_T^e(\beta, h)}{\partial h_{o \rightarrow j}(\sigma)} &= \frac{1}{2} g_\sigma^\beta(o \rightarrow j; h) \end{aligned}$$

where g_σ^β is the same as \widehat{g}_σ^β but with h in place of \widehat{h} . Note that for permissive ψ and any $\sigma \in \mathcal{X}$,

$$\widehat{g}_\sigma^\beta(x \rightarrow y; h) \leq \frac{\sum_{\sigma'} \psi(\sigma', \sigma) \widehat{h}_{y \rightarrow x}^\beta(\sigma')}{\sum_{\sigma'} \psi(\sigma', \sigma^{\text{p}}) h_{x \rightarrow y}(\sigma^{\text{p}}) \widehat{h}_{y \rightarrow x}^\beta(\sigma')} \leq \frac{\psi_{\max}^\beta}{\psi_{\min}^\beta h_{x \rightarrow y}(\sigma^{\text{p}})}. \quad (2.7)$$

If further $\psi > 0$ everywhere, then $\widehat{g}_\sigma^\beta(x \rightarrow y; h) \leq \psi_{\max}^\beta / \psi_{\min}^\beta$ is uniformly bounded on $[\beta_0, \beta_1]$.

Consider now a small sub-interval $[\beta, \beta + \delta]$ of $[\beta_0, \beta_1]$. Writing $\Delta_{\beta, \delta} h \equiv h^{\beta + \delta} - h^\beta$ and applying the mean value theorem to the differentiable function $t \mapsto \Phi_T(\beta + t\delta, h + t\Delta_{\beta, \delta} h)$ for $t \in [0, 1]$ gives

$$\Phi_T(\beta + \delta, h) - \Phi_T(\beta, h) = \frac{\partial \Phi_T}{\partial \beta}(\beta + t\delta, h^\beta + t\Delta_{\beta, \delta} h)\delta + \Gamma_T(\beta, \delta) + E_T(\beta, \delta) \quad (2.8)$$

for some $t \equiv t_{\beta, \delta} \in [0, 1]$, where

$$\begin{aligned} \Gamma_T(\beta, \delta) &\equiv \sum_\sigma \sum_{j \in \partial o} \left\{ \frac{\partial \Phi_T}{\partial h_{j \rightarrow o}(\sigma)}(\beta, h) \Delta_{\beta, \delta} h_{j \rightarrow o}(\sigma) + \frac{\partial \Phi_T}{\partial h_{o \rightarrow j}(\sigma)}(\beta, h) \Delta_{\beta, \delta} h_{o \rightarrow j}(\sigma) \right\}, \\ E_T(\beta, \delta) &\equiv \sum_\sigma \sum_{x \rightarrow y}^* \left\{ \frac{\partial \Phi_T}{\partial h_{x \rightarrow y}(\sigma)}(\beta + t\delta, h^\beta + t\Delta_{\beta, \delta} h) - \frac{\partial \Phi_T}{\partial h_{x \rightarrow y}(\sigma)}(\beta, h) \right\} \Delta_{\beta, \delta} h_{x \rightarrow y}(\sigma), \end{aligned}$$

and $\sum_{x \rightarrow y}^*$ indicates the sum over the $2D_o$ directed edges $x \rightarrow y$ within T^1 .

Setting $\delta \equiv \delta_m \equiv (\beta_1 - \beta_0)/m$, we now sum $\Phi(\beta + \delta_m, h) - \Phi(\beta, h)$ over $\beta \in \Pi_m \equiv \{\beta_0 + k\delta_m : 0 \leq k < m\}$ and analyze separately the contribution of each term on the right-hand side of (2.8):

(a) First we show that $\mathbb{E}_\mu[\Gamma_T(\beta, \delta)] = 0$ for any $[\beta, \beta + \delta] \subseteq [\beta_0, \beta_1]$. Indeed, since $h \in \mathcal{H}^*$ we have $\widehat{h}^\beta = h^\beta$ and $\widehat{g}^\beta = g^\beta$. Therefore,

$$\Gamma_T(\beta, \delta) = \frac{1}{2} \sum_\sigma \sum_{j \in \partial o} \left\{ g_\sigma^\beta(j \rightarrow o; h) \Delta_{\beta, \delta} h_{j \rightarrow o}(\sigma) - g_\sigma^\beta(o \rightarrow j; h) \Delta_{\beta, \delta} h_{o \rightarrow j}(\sigma) \right\}.$$

The result then follows from unimodularity of μ , subject to μ -integrability of

$$\sum_\sigma \sum_{j \in \partial o} |g_\sigma^\beta(j \rightarrow o; h) \Delta_{\beta, \delta} h_{j \rightarrow o}(\sigma)|.$$

Clearly $|\Delta_{\beta, \delta} h_{x \rightarrow y}(\sigma)| \leq 2$ so integrability certainly holds when $\psi > 0$, since $\mathbb{E}_\mu[D_o] < \infty$ and g_σ^β is deterministically uniformly bounded on $[\beta_0, \beta_1]$ as noted above. More generally, for permissive ψ the required μ -integrability follows from (2.7) and (H2 $^\beta$).

(b) The total contribution of the first term on the right-hand side of (2.8) is

$$A_m \equiv \delta \mathbb{E}_\mu \left[\sum_{\beta \in \Pi_m} \frac{\partial \Phi_T}{\partial \beta}(\beta + t_{\beta, \delta} \delta, h^\beta + t_{\beta, \delta} \Delta_{\beta, \delta} h) \right].$$

Observe that $A_m = \int Y_m d(\lambda \times \mu)$ where λ is Lebesgue measure on $[\beta_0, \beta_1]$ and

$$Y_m(\beta', T) \equiv \sum_{\beta \in \Pi_m} \mathbf{1}\{\beta \leq \beta' < \beta + \delta\} \frac{\partial \Phi_T}{\partial \beta}(\beta + t_{\beta, \delta} \delta, h^\beta + t_{\beta, \delta} \Delta_{\beta, \delta} h).$$

For $(\lambda \times \mu)$ -a.e. (β', T) , this sum has at most one non-zero term, in which the argument of $\partial_\beta \Phi_T$ converges by (H3 $^\beta$) to $(\beta', h^{\beta'})$ as $m \rightarrow \infty$. From (H1), (1.14), and the computation of $\partial_\beta \Phi_T$ in (2.4)-(2.5), we see that $\partial_\beta \Phi_T(\beta, h)$ is continuous in (β, h) . Therefore, $Y_m(\beta', T) \rightarrow a_T^e(\beta', h)$, $(\lambda \times \mu)$ -a.e. Furthermore, (H1) implies that $|\partial_\beta \xi| \leq C$ uniformly on $[\beta_0, \beta_1]$ for some deterministic constant C , so $|Y_m| \leq 2CD_o$ for all m , $(\lambda \times \mu)$ -a.e. (see (2.4) and (2.5)). Dominated convergence then gives

$$\lim_{m \rightarrow \infty} A_m = \int a_T^e(\beta', h) d(\lambda \times \mu) = \int_{\beta_0}^{\beta_1} a^e(\beta', h) d\beta'.$$

(c) The contribution of the final term in (2.8) is $\mathbb{E}_\mu[E_{T,m}]$ where

$$E_{T,m} \equiv \sum_{\beta \in \Pi_m} E_T(\beta, \delta),$$

and we conclude the proof by showing that $\lim_{m \rightarrow \infty} \mathbb{E}_\mu[E_{T,m}] = 0$.

Indeed, it is not hard to see that $\lim_{m \rightarrow \infty} E_{T,m} = 0$ μ -a.s.: by the uniform bound on total variation assumed in (H3 $^\beta$), there exists deterministic C such that

$$|E_{T,m}| \leq C \sum_{x \rightarrow y}^* \max_{\sigma} \sup_{\beta \in [\beta_0, \beta_1]} \sup_{t \in [0,1]} \left| \frac{\partial \Phi_T}{\partial h_{x \rightarrow y}(\sigma)}(\beta + t\delta_m, h^\beta + t\Delta_{\beta, \delta_m} h) - \frac{\partial \Phi_T}{\partial h_{x \rightarrow y}(\sigma)}(\beta, h) \right|$$

μ -a.s., uniformly in m . It also follows from (H3 $^\beta$) that μ -a.e. h^β is uniformly continuous on $[\beta_0, \beta_1]$. Using (H1), the partials $\partial_h \Phi_T$ computed above are uniformly continuous in (β, h) for $\beta \in [\beta_0, \beta_1]$ and $h_{j \rightarrow o}^\beta(\sigma^P)$ uniformly bounded away from zero. By (2.3) there exists deterministic c such that

$$\inf_{\beta \in [\beta_0, \beta_1]} h_{j \rightarrow o}^\beta(\sigma^P) \geq e^{-c(D_j+1)} \quad \forall j \in \partial o, \mu\text{-a.s..}$$

Combining these observations gives $\lim_{m \rightarrow \infty} E_{T,m} = 0$ μ -a.s.

To take the limit in μ -expectation, we argue similarly as in part (a): by (2.7) and (H1) there exists deterministic C' such that

$$\left| \frac{\partial \Phi_T}{\partial h_{x \rightarrow y}(\sigma)}(\beta + t\delta, h^\beta + t\Delta_{\beta, \delta} h) \right| \leq \frac{C'}{h_{x \rightarrow y}^\beta(\sigma^P) + t\Delta_{\beta, \delta} h_{x \rightarrow y}(\sigma^P)} \leq \sup_{\beta' \in [\beta_0, \beta_1]} \frac{C'}{h_{x \rightarrow y}^{\beta'}(\sigma^P)}$$

for all $\beta \in [\beta_0, \beta_1 - \delta]$, $x \rightarrow y \in T^1$, $\sigma \in \mathcal{X}$ and $t \in [0, 1]$, hence

$$|E_{T,m}| \leq CC' \sum_{x \rightarrow y}^* \sup_{\beta \in [\beta_0, \beta_1]} \frac{1}{h_{x \rightarrow y}^\beta(\sigma^P)}.$$

This is integrable by (H2 $^\beta$) and unimodularity of μ , so dominated convergence implies that $\lim_{m \rightarrow \infty} \mathbb{E}_\mu[E_{T,m}] = 0$ as claimed.

Combining (a)-(c) gives the result of the proposition. \square

Proof of Thm. 1.13(a). Recalling Lem. 2.1,

$$\limsup_{n \rightarrow \infty} [\phi_n(\beta_1) - \phi_n(\beta_0)] = \limsup_{n \rightarrow \infty} \int_{\beta_0}^{\beta_1} a_n^e(\beta, B) d\beta \leq \int_{\beta_0}^{\beta_1} \limsup_{n \rightarrow \infty} a_n^e(\beta, B) d\beta \leq \int_{\beta_0}^{\beta_1} a^e(\beta, h) d\beta,$$

where the first inequality follows by (the reversed) Fatou's lemma and the second one by the hypothesis (1.18). By Propn. 2.4 the right-most expression equals to $\Phi(\beta_1, h) - \Phi(\beta_0, h)$, so the theorem is proved. \square

The justification for interpolation in B is entirely similar:

Proof of Thm. 1.13(b). Now β is fixed so we suppress it from the notation. For $h \in \mathcal{H}$ and $T \in \mathcal{T}_\bullet^+$ then

$$\frac{\partial \Phi_T(B, h)}{\partial B} = \frac{\partial \Phi_T^{\text{vx}}(B, h)}{\partial B} = \frac{\sum_{\sigma, \sigma_j} \partial_B \bar{\xi}(\sigma) \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j) \widehat{h}_{o \rightarrow j}(\sigma)}{\sum_{\sigma, \sigma_j} \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j) \widehat{h}_{o \rightarrow j}(\sigma)} \quad \forall j \in \partial o,$$

while if $T = T^0$ then $\partial_B \Phi_T = \sum_{\sigma} \partial_B \bar{\xi}(\sigma) \bar{\psi}(\sigma) / \sum_{\sigma} \bar{\psi}(\sigma)$. If $h \in \mathcal{H}^*$ then $\widehat{h}^B = h^B$, so

$$\mathbb{E}_\mu[\partial_B \Phi_T(B, h)] = \mathbb{E}_\mu[\partial_B \bar{\xi}(\sigma_o)]_T^{h, B} \equiv a^{\text{vx}}(B, h).$$

The result now follows by adapting the proofs of Propn. 2.4 and Thm. 1.13(a). \square

2.3. Discussion and first consequences. We now prove Thm. 1.14 by considering an extended notion of local weak convergence. As discussed in [1], a graph $G = (V, E)$ together with a spin configuration $\underline{\sigma} \in \mathcal{X}^V$ on the graph can be regarded as a graph with *marks* in \mathcal{X} . Let $\mathcal{G}_\bullet^{\mathcal{X}}$ and $\mathcal{G}_{\bullet\bullet}^{\mathcal{X}}$ denote the spaces of marked isomorphism classes of connected, rooted and bi-rooted graphs, respectively, with marks in \mathcal{X} . These spaces are metrizable by the obvious generalizations of the metrics on $\mathcal{G}_\bullet, \mathcal{G}_{\bullet\bullet}$ defined in §2.1, giving rise to the notion of local weak convergence for pairs $(G_n, \underline{\sigma}_n)$ of graphs with spin configurations. Defn. 1.2 generalizes naturally to this setting, and we show next that if $\underline{\sigma}_n$ is a random configuration on G_n with law $\nu_{G_n, \underline{\psi}}$ (as defined in (1.1)), then a local weak limit of $(G_n, \underline{\sigma}_n)$, if it exists, must be unimodular.

Lemma 2.5. *If $G_n \rightarrow_{lwc} \mu$ and $\underline{\sigma}_n \sim \nu_{G_n, \underline{\psi}}$ then the laws of $(G_n, \underline{\sigma}_n)$ have subsequential local weak limits belonging to the space \mathcal{U} of unimodular measures on $\mathcal{G}_\bullet^{\mathcal{X}}$.*

Proof. For each fixed t , the laws of $B_t(I_n)$ are weakly convergent, hence by Prohorov's theorem form a uniformly tight sequence. Consequently, for each $\epsilon > 0$ there exists $K_\epsilon \subseteq \mathcal{G}_\bullet$ compact with $\sup_n \mathbb{P}_n(B_t(I_n) \notin K_\epsilon) \leq \epsilon$. Further, K_ϵ may be taken to contain only graphs of depth at most t , whereby the minimal distance between any two graphs in K_ϵ is uniformly bounded below (by $1/(1+t)$), hence the compactness of K_ϵ implies that it must be a finite set. The collection of all marked graphs in $\mathcal{G}_\bullet^{\mathcal{X}}$ whose underlying graph is in K_ϵ must therefore be finite, hence compact as well. Thus, by yet another application of Prohorov's theorem, the joint laws of $(B_t(I_n), \underline{\sigma}_{B_t(I_n)})$ are uniformly tight in $\mathcal{G}_\bullet^{\mathcal{X}}$ and consequently have subsequential weak limits. By extracting successive subsequences for increasing t and taking the diagonal subsequence it follows that the sequence $(G_n, \underline{\sigma}_n)$ admits subsequential local weak limits $\widehat{\mu} \in \mathcal{U}$. \square

For $\widehat{\mu} \in \mathcal{U}$, the marginal μ of $\widehat{\mu}$ is a unimodular measure on \mathcal{G}_\bullet . If it is supported on a single tree T as in the d -regular case, then clearly $\widehat{\mu}$ may be represented as $\delta_T \times \nu$ where $\nu \in \mathcal{G}_T$, the space of Gibbs measures on T corresponding to specification ψ . To make such a statement in the general setting, note that there is a continuous mapping π from \mathcal{G}_\bullet to the space \mathcal{N}_\bullet of graphs on $\mathbb{Z}_{\geq 0}$ rooted at 0, taking an isomorphism class to its canonical representative [1, p. 1461]. Thus $\widehat{\mu}$ may be regarded as a measure on the product space $\mathcal{N}_\bullet \times \mathcal{X}^{\mathbb{Z}_{\geq 0}}$, and consequently $\widehat{\mu}$ has a representation as the measure $\mu \otimes \nu$ on pairs $(T, \underline{\sigma})$ where T has law μ and $\underline{\sigma}$ given T has law $\nu_T \in \mathcal{G}_T$. In particular, if $|\mathcal{G}_T| = 1$ μ -a.s., then $\mu \otimes \nu$ is uniquely determined.

Let μ be a unimodular measure on \mathcal{T}_\bullet . It was noted in Rmk. 1.12 that there is a mapping from $\mathcal{H}_\mu^*(\beta, B)$ to collections $(\nu_T \in \mathcal{G}_T)_{T \in \mathcal{T}_\bullet}$. For such ν , $\mu \otimes \nu$ belongs to \mathcal{U} : if f is a non-negative Borel function on $\mathcal{G}_{\bullet\bullet}^{\mathcal{X}}$, it follows from the \mathcal{T}_e -measurability of elements of \mathcal{H}_{loc} that

$$\mathbb{E}_{\mu \otimes \nu} \left[\sum_{j \in \partial o} f((T, \underline{\sigma}), o, j) \right] = \mathbb{E}_\mu \left[\sum_{j \in \partial o} \bar{f}(T, o, j) \right],$$

where \bar{f} is a non-negative Borel function on $\mathcal{G}_{\bullet\bullet}$. The unimodularity of the underlying measure μ then gives

$$\mathbb{E}_{\mu \otimes \nu} \left[\sum_{j \in \partial o} f((T, \underline{\sigma}), o, j) \right] = \mathbb{E}_{\mu \otimes \nu} \left[\sum_{j \in \partial o} f((T, \underline{\sigma}), j, o) \right],$$

and therefore $\mu \otimes \nu \in \mathcal{U}$.

Remark 2.6. An element $\nu \in \mathcal{G}_T$ is called a *Markov chain* (or *splitting Gibbs measure*) if for any finite connected subgraph $U \subseteq T$, the marginal of ν on U is a Markov random field ([42], see also [17, Ch. 12] and [38]). A collection $\Lambda_T \equiv (\lambda_i^j)_{(ij) \in E_T}$ of probability measures on \mathcal{X} is called an *entrance law* (or *boundary law*) for the specification $\underline{\psi} \equiv (\psi, \bar{\psi})$ on T if it satisfies the consistency requirement ([42, (3.4)])

$$\lambda_i^j(\sigma_i) = \prod_{k \in \partial i \setminus j} \left(\sum_{\sigma_k} \phi_{ik}(\sigma_i, \sigma_k) \lambda_k^i(\sigma_k) \right),$$

where $\phi_{ij}(\sigma_i, \sigma_j) \equiv \bar{\psi}(\sigma)^{1/D_i} \psi(\sigma, \sigma') \bar{\psi}(\sigma')^{1/D_j}$, the pairwise interaction potential corresponding to $\underline{\psi}$. It is shown in [42, Thm. 3.2] that there is a one-to-one correspondence between Markov chains ν and entrance laws Λ_T , given by

$$\nu(\underline{\sigma}_U) \cong \prod_{(ij) \in E_U} \phi_{ij}(\sigma_i, \sigma_j) \prod_{i \in \partial U} \left(\sum_{\sigma_i} \phi_{ip(i)}(\sigma_i, \sigma_{p(i)}) \lambda_i^{p(i)}(\sigma_i) \right)$$

for U any finite connected subgraph of T , with $p(i)$ denoting the unique neighbor of i inside U for $i \in \partial U$. In particular, we see that the Gibbs measure ν_T arising from $h \in \mathcal{H}_\mu^*(\beta, B)$ is precisely the Markov chain with entrance law $\lambda_j^i(\sigma) \cong h_{j \rightarrow i}(\sigma) \bar{\psi}(\sigma)^{-1/D_j}$. Extremal elements of \mathcal{G}_T are Markov chains [42, Thm. 2.1] but the converse is false; for example, the free-boundary Ising Gibbs measure is non-extremal at low temperature (see [13, 23]). The measures $\mu \otimes \nu$ arising from elements of $\mathcal{H}_\mu^*(\beta, B)$ might naturally be termed “unimodular Markov chains” in the sense that the entrance laws for the entire collection $(\nu_T)_{T \in \mathcal{T}_\bullet}$ are specified by a single measurable function $h : \mathcal{T}_e \rightarrow \Delta_{\mathcal{X}}$. In the case $\mu = \delta_{\mathbb{T}_d}$ these correspond precisely to the *completely homogeneous Markov chains* studied in [42, §4].

Proof of Thm. 1.14. Suppose uniqueness holds at (β, B) , i.e. $\mathcal{G}_T = \{\nu_T\}$ μ -a.s. Then $\mathcal{H}_\mu^*(\beta, B)$ has size at most one by Rmk. 2.3. For μ -a.e. T , the measure ν_T is extremal, and so specifies a Markov chain on T with entrance law Λ_T (see Rmk. 2.6). If we define $h_{x \rightarrow y}(\sigma) \equiv h_{(T, x \rightarrow y)}(\sigma) \cong \lambda_x^y(\sigma) \bar{\psi}(\sigma)^{1/D_x}$, then $h \in \mathcal{H}_\mu^*(\beta, B)$, which proves that $\mathcal{H}_\mu^*(\beta, B)$ is a singleton.

Now consider interpolation in β or B . All the conditions of Thm. 1.13 are satisfied by assumption except (1.18) and (1.19). If uniqueness holds at (β, B) , it follows from the preceding discussion that there is a unique $\mu \otimes \nu \in \mathcal{U}$ corresponding to the specification $(\psi^\beta, \bar{\psi}^B)$. Any local weak limit of $(G_n, \underline{\sigma}_n)$ must be such a measure, so $(G_n, \underline{\sigma}_n) \rightarrow_{lwc} \mu \otimes \nu$; likewise, any element of $\mathcal{H}_\mu^*(\beta, B)$ gives rise to $\mu \otimes \nu$. Therefore,

$$\lim_{n \rightarrow \infty} a_n^e(\beta, B) = \frac{1}{2} \mathbb{E}_{\mu \otimes \nu} \left[\sum_{j \in \partial o} \partial_\beta \xi(\sigma_o, \sigma_j) \right] = a^e(\beta, B, h),$$

where the limit in expectation is justified by the boundedness of $\partial_\beta \xi$ on compacts and uniform sparsity (as in the proof of Lem. 2.1). This verifies (1.18), and the verification of (1.19) is entirely similar. The result therefore follows from Thm. 1.13. \square

Remark 2.7. If uniqueness of Gibbs measures does not hold, one may consider extremal decomposition of the subsequential local weak limits $\hat{\mu}$ of $(G_n, \underline{\sigma}_n)$, either in the spaces \mathcal{G}_T (possibly losing

unimodularity in the decomposition) or in the space \mathcal{U} . Extremal decomposition in \mathcal{U} is discussed in [1, §4] but it is unclear whether extremal elements would be unimodular Markov chains in the sense described here. A decomposition of $\hat{\mu} = \mu \otimes \nu$ into unimodular Markov chains $\mu \otimes \nu'$ would obviously yield a substantial generalization of Thm. 1.14.

2.4. Application to independent set. We now prove Thm. 1.11, our result for the independent set model (1.4), by verifying the conditions of Thm. 1.14 for the interpolation parameter $B \equiv \log \lambda$. In this setting a convenient parametrization for the messages $h \in \mathcal{H}$ is $u \equiv h(0)$, so that the BP mapping (1.14) becomes

$$(\text{BP}^\lambda u)_{x \rightarrow y} = \frac{1}{1 + \lambda \prod_{v \in \partial x \setminus y} u_{v \rightarrow x}}. \quad (2.9)$$

A single BP iteration is anti-monotone in the messages $u_{v \rightarrow x}$, so a double iteration is monotone. Since the root marginal for an independent set model in T^{2t-1} is obtained by an even number of BP iterations starting from level $2t$ (see Rmk. 1.12), it is monotone in the boundary conditions. Recalling from §1.2.3 the definition of $\bar{h}_T^{t, \ddagger} \equiv \bar{h}_T^{t, \ddagger, \lambda}$ for $\ddagger \in \{0, 1\}$ and writing $\bar{u}_T^{t, \ddagger} \equiv \bar{h}_T^{t, \ddagger}(0)$, the above implies that for $1 \leq s \leq t$,

$$\bar{u}_T^{2s-1, 0} \geq \bar{u}_T^{2t-1, 0} \geq \bar{u}_T^{2t-1, 1} \geq \bar{u}_T^{2s-1, 1} \geq \bar{u}_T^{1, 1} = \frac{1}{1 + \lambda}.$$

Thus the $t \rightarrow \infty$ limits \bar{h}_T^0, \bar{h}_T^1 are well-defined with $\bar{h}_T^0(1) \geq \bar{h}_T^1(1) \geq 1/(1 + \lambda)$, and using these we define messages $h^\ddagger \in \mathcal{H}$, $h_{x \rightarrow y}^\ddagger = \bar{h}_{T_{x \rightarrow y}}^\ddagger$. The next lemma gives the boundary values for the interpolation.

Lemma 2.8. *For the independent set model on $G_n \rightarrow_{lwc} \mu$,*

$$\lim_{\lambda \downarrow 0} \limsup_{n \rightarrow \infty} |\phi_n(\lambda)| = 0 = \lim_{\lambda \downarrow 0} \Phi(\lambda, h^\ddagger), \quad \ddagger \in \{0, 1\}.$$

Proof. The left limit follows from the trivial bounds $1 \leq Z_n \leq (1 + \lambda)^n$. Next, for any $h \in \mathcal{H}$,

$$\Phi_T^{yx}(\lambda, h) = \log \left\{ 1 + \lambda \prod_{j \in \partial o} h_{j \rightarrow o}(0) \right\},$$

so $\Phi_T^{yx}(\lambda, h) \rightarrow 0$ both μ -a.s. and in μ -expectation as $\lambda \downarrow 0$, by bounded convergence. The same holds for $\Phi_T^e(\lambda, h^\ddagger)$, $\ddagger \in \{0, 1\}$, using the bound $h_{x \rightarrow y}^\ddagger(0) \geq 1/(1 + \lambda)$. \square

Proof of Thm. 1.11. The independent set model (1.4) is of form (1.1) with $\mathcal{X} = \{0, 1\}$, $\psi(\sigma, \sigma') = \mathbf{1}\{\sigma\sigma' \neq 1\}$, and $\bar{\psi}(\sigma) = \lambda^\sigma \equiv e^{B\sigma}$, so (H1) is clearly satisfied with $\sigma^p = 0$ the permitted state. By definition of λ_c , if $\lambda < \lambda_c$ then $h^0 = h^1 \equiv h$ in \mathcal{H} , and it then follows from the recursive structure of the tree that $h \in \mathcal{H}_\mu^*(\lambda)$. Since $h_{x \rightarrow y}^\ddagger(0) \geq 1/(1 + \lambda)$ as noted above, (H2^B) is satisfied on any compact interval of λ .

For $T \in \mathcal{T}_\bullet$, as noted above the root occupation probability on T^s for $s \geq 2t - 1$ with any boundary conditions is sandwiched between $\bar{h}_T^{2t-1, 0}(1)$ and $\bar{h}_T^{2t-1, 1}(1)$, with the former increasing to $\bar{h}_T^0(1)$ and the latter decreasing to $\bar{h}_T^1(1)$. Since the $\bar{h}_T^{t, \ddagger}$ are clearly continuous in λ , it follows that $\bar{h}_T^0(1)$ and $\bar{h}_T^1(1)$ are respectively lower and upper semi-continuous in λ , so if they coincide then their common value $\bar{h}_T(1)$ is continuous in λ . Applying this with $T = T_{x \rightarrow y}$ gives the μ^\uparrow -a.s. continuity of $h_{x \rightarrow y}^\ddagger$ on $(0, \lambda_c)$.

For $T \in \mathcal{T}_\bullet$, \bar{h}_T^\ddagger for $\ddagger \in \{0, 1\}$ is a function of $(h_{j \rightarrow o}^{1-\ddagger})_{j \in \partial o}$, so for $\lambda < \lambda_c$ we have that $\bar{h}_T^0 = \bar{h}_T^1$, μ -a.s. It then follows from the preceding observations and Rmk. 1.12 that the boundary effect vanishes and $|\mathcal{G}_T| = 1$ μ -a.s. Thus, we are in the setting of Thm. 1.14(b), and it remains only to complete the verification of (H3^B), i.e. the boundedness in total variation of the messages $h_{x \rightarrow y}$:

- (a) No verification is needed since boundedness in total variation is simply assumed.
 (b) For $T \in \mathcal{T}_\bullet$, $\bar{u}_T^{t,\ddagger} \equiv \bar{h}_T^{t,\ddagger}(0)$ satisfies

$$\log \bar{u}_T^{2t+1,\ddagger} = -\log \left(1 + \lambda \prod_{j \in \partial o} \frac{1}{1 + \lambda \prod_{k \in \partial j \setminus o} \bar{u}_{T_{k \rightarrow j}}^{2t-1,\ddagger}} \right).$$

Differentiating with respect to λ , we find that $r_T^{t,\ddagger} \equiv (1 + \lambda) \partial_\lambda \log \bar{u}_T^{t,\ddagger}$ satisfies

$$|r_T^{2t+1,\ddagger}| \leq 1 + \frac{\lambda}{1 + \lambda} D_o + \left(\frac{\lambda}{1 + \lambda} \right)^2 \sum_{j \in \partial o} \sum_{k \in \partial j \setminus o} |r_{T_{k \rightarrow j}}^{2t-1,\ddagger}|. \quad (2.10)$$

Since $\bar{u}_T^{1,1} = 1/(1 + \lambda)$ for any $T \in \mathcal{T}_\bullet$, we find that

$$\sup_{t \geq 0} |r_T^{2t-1,1}| \leq 1 + \sum_{\ell \geq 1} \left(\frac{\lambda}{1 + \lambda} \right)^\ell |\partial T^{\ell-1}|.$$

If $\lambda_1/(1 + \lambda_1) < 1/\text{br } T$ then this is finite and uniformly bounded on $[\lambda_0, \lambda_1]$ (see [28, §2]), and consequently $\bar{u}_T^1 \equiv \lim_{t \rightarrow \infty} \bar{u}_T^{2t-1,1}$ has deterministically bounded total variation on $[\lambda_0, \lambda_1]$. If $\lambda_1 < \lambda_{c,\mu}$ then $h_{x \rightarrow y}^{2t-1,1} \rightarrow h_{x \rightarrow y}$ on $[\lambda_0, \lambda_1]$, so if $\text{br } T_{x \rightarrow y} \leq \Delta - 1$ μ^\uparrow -a.s. and $\lambda_1/(1 + \lambda_1) < 1/(\Delta - 1)$ (i.e. $\lambda_1(\Delta - 2) < 1$), then h has deterministically bounded total variation on $[\lambda_0, \lambda_1]$.

- (c) Since the limiting measure is supported on \mathbb{T}_d , only $h \equiv h_{(\mathbb{T}_d, x \rightarrow y)}$ is of relevance, and (2.9) reduces to $\text{BP}^\lambda u = (1 + \lambda u^{d-1})^{-1}$. For $\lambda \leq \lambda_c = \lambda_c(d)$ there is a unique fixed point (see [25, §2]), which is then easily seen to be monotone in λ .

Thus (H3^B) is verified in parts (a)-(c). Also, $\phi(\lambda_c) = \lim_{\lambda \uparrow \lambda_c} \phi(\lambda)$ as an immediate consequence of Lem. 2.1. The rest of the theorem follows by applying Thm. 1.14 and then taking $B_0 = \log \lambda_0 \rightarrow -\infty$, relying on the boundary value given by Lem. 2.8. \square

3. BETHE PREDICTION AS OPTIMIZATION OVER LOCAL POLYTOPE

Throughout this section we assume that $\psi \equiv (\psi, \bar{\psi})$ satisfying (H1) specifies a factor model of the form (1.1), with μ a unimodular measure on \mathcal{T}_\bullet with $\mathbb{E}_\mu[D_o] < \infty$. We study the Bethe prediction as the optimization of the Bethe free energy functional Φ_μ on \mathcal{H}_{loc} as defined by (1.21). We first verify that this agrees with the previous definition (1.9) of Φ_μ on $\mathcal{H}_\mu^*(\beta, B)$, which we always regard as being embedded into \mathcal{H}_{loc} via (1.20). Recall from Defn. 1.15 that for $\mathbf{h} \in \mathcal{H}_{\text{loc}}$, the one-point marginals of \mathbf{h}_{xy} are denoted \bar{h}_x and \bar{h}_y , and are measurable functions $\mathcal{T}_\bullet \rightarrow \Delta_{\mathcal{X}}$.

Lemma 3.1. *The functional Φ_μ on \mathcal{H}_{loc} given by (1.23) agrees with the previous definition (1.9) on $\mathcal{H}_\mu^*(\beta, B)$, subject to finiteness of $\mathbb{E}_\mu[\Phi_T^e]$.*

Proof. If \mathbf{h} corresponds to $h \in \mathcal{H}_\mu^*(\beta, B)$, then (1.20) and (1.15) imply that

$$\begin{aligned} \mathbf{h}_{xy}(\sigma, \sigma') \exp\{\Phi_T^{(xy)}(h)\} &= \psi(\sigma, \sigma') h_{x \rightarrow y}(\sigma) h_{y \rightarrow x}(\sigma'), \\ \bar{h}_o(\sigma) \exp\{\Phi_T^{\vee x}(h)\} &= \bar{\psi}(\sigma) \prod_{j \in \partial o} \left(\sum_{\sigma_j} \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j) \right). \end{aligned}$$

Letting $\Phi_i(\mathbf{h})$ ($1 \leq i \leq 3$) denote the three terms on the right-hand side of (1.22), it follows from the above that

$$\begin{aligned}\Phi_1(\mathbf{h}) &= \mathbb{E}_\mu[\Phi_T^{\vee\mathbf{x}}(h)] - \mathbb{E}_\mu\left[\sum_{j \in \partial o} \sum_{\sigma} \bar{h}_o(\sigma) \log\left(\sum_{\sigma_j} \psi(\sigma, \sigma_j) h_{j \rightarrow o}(\sigma_j)\right)\right], \\ \Phi_2(\mathbf{h}) &= \mathbb{E}_\mu[\Phi_T^e(h)] - \frac{1}{2} \mathbb{E}_\mu\left[\sum_{j \in \partial o} \sum_{\sigma, \sigma_j} \mathbf{h}_{oj}(\sigma, \sigma_j) \log(h_{o \rightarrow j}(\sigma) h_{j \rightarrow o}(\sigma_j))\right] \\ &= \mathbb{E}_\mu[\Phi_T^e(h)] - \mathbb{E}_\mu\left[\sum_{j \in \partial o} \sum_{\sigma} \bar{h}_o(\sigma) \log h_{o \rightarrow j}(\sigma)\right], \\ \Phi_3(\mathbf{h}) &= \mathbb{E}_\mu\left[\sum_{j \in \partial o} \sum_{\sigma} \bar{h}_o(\sigma) \log\left(\sum_{\sigma_j} \psi(\sigma, \sigma_j) h_{o \rightarrow j}(\sigma) h_{j \rightarrow o}(\sigma_j)\right)\right] - 2\mathbb{E}_\mu[\Phi_T^e(h)],\end{aligned}$$

where unimodularity was used in the simplification of Φ_2 . Adding these three identities gives $\Phi_\mu(\mathbf{h}) = \mathbb{E}_\mu[\Phi_T^{\vee\mathbf{x}}(h) - \Phi_T^e(h)]$ as claimed. \square

As mentioned in §1.3.2, our definition Φ_μ of the Bethe free energy functional on \mathcal{H}_{loc} is an infinite-tree analogue of the definition of [41] for finite graphs. It is proved in [41, Propn. 6] that when $\psi > 0$, all local maxima of the Bethe free energy lie in the *interior* of the local polytope. The following is the analogue of this result for infinite unimodular trees.

Proposition 3.2. *If $\psi > 0$ and \mathbf{h} is a local maximizer of Φ_μ over \mathcal{H}_{loc} , then*

$$\Omega_{\mathbf{h}} \equiv \{\mathbf{h}_{xy}(\sigma, \sigma') = 0 \text{ for some } \sigma, \sigma' \in \mathcal{X}\}$$

has zero μ^\uparrow -measure.

Proof. If $\mathbf{u} \in \mathcal{H}_{\text{loc}}^\circ$ and we set $\boldsymbol{\delta} \equiv \mathbf{u} - \mathbf{h}$, then it follows by convexity of \mathcal{H}_{loc} that $\mathbf{h}^\eta \equiv \mathbf{h} + \eta\boldsymbol{\delta} = (1 - \eta)\mathbf{h} + \eta\mathbf{u}$ belongs to $\mathcal{H}_{\text{loc}}^\circ$ for any $\eta \in (0, 1]$. Letting

$$R_\eta(\boldsymbol{\delta}) \equiv \frac{2}{\eta} [\Phi_\mu(\mathbf{h}^\eta) - \Phi_\mu(\mathbf{h})], \quad \widehat{R}_\eta(\boldsymbol{\delta}) \equiv \frac{R_\eta(\boldsymbol{\delta})}{\log(1/\eta)},$$

our claim will follow upon showing that if $\mu^\uparrow(\Omega_{\mathbf{h}}) > 0$, then there exists $\mathbf{u} \in \mathcal{H}_{\text{loc}}$ for which

$$\lim_{\eta \downarrow 0} \widehat{R}_\eta(\boldsymbol{\delta}) = \widehat{R}_0(\boldsymbol{\delta}) > 0.$$

To this end, note that by an easy computation $[H(\mathbf{h}^\eta) - H(\mathbf{h})]/\eta = -\langle \log \mathbf{h}^\eta \rangle_{\boldsymbol{\delta}} - \langle f_\eta(\boldsymbol{\delta}/\mathbf{h}) \rangle_{\mathbf{h}}$, where $f_\eta(r) \equiv \eta^{-1} \log(1 + \eta r)$ and $\boldsymbol{\delta}/\mathbf{h}$ is defined by coordinate-wise division in $\mathbb{R}^{\mathcal{X}^2}$. Thus from (1.21) we obtain

$$\begin{aligned}R_\eta(\boldsymbol{\delta}) &= \mathbb{E}_\mu\left[2\langle \bar{\xi} \rangle_{\bar{\delta}_o} + 2(D_o - 1)\langle \log \bar{h}_o^\eta \rangle_{\bar{\delta}_o} + 2(D_o - 1)\langle f_\eta(\bar{\delta}_o/\bar{h}_o) \rangle_{\bar{h}_o}\right] \\ &\quad + \mathbb{E}_\mu\left[\sum_{j \in \partial o} \left\{ \langle \xi \rangle_{\boldsymbol{\delta}_{oj}} - \langle \log \mathbf{h}_{oj}^\eta \rangle_{\boldsymbol{\delta}_{oj}} - \langle f_\eta(\boldsymbol{\delta}_{oj}/\mathbf{h}_{oj}) \rangle_{\mathbf{h}_{oj}} \right\}\right]\end{aligned}\tag{3.1}$$

Since $\lim_{\eta \downarrow 0} f_\eta(r) = r$ and $\sum_{\sigma, \sigma'} \boldsymbol{\delta}(\sigma, \sigma')$ is identically zero, the two terms above involving f_η vanish in the limit $\eta \downarrow 0$. Upon scaling by $\log(1/\eta)$, the terms involving $\bar{\psi}, \psi$ also tend to zero. It follows that

$$\widehat{R}_0(\boldsymbol{\delta}) = \mathbb{E}_\mu\left[(2 - 2D_o)\bar{u}_o(\{\sigma : \bar{h}_o(\sigma) = 0\}) + \sum_{j \in \partial o} \mathbf{u}_{oj}(\{\sigma, \sigma' : \mathbf{h}_{oj}(\sigma, \sigma') = 0\})\right],\tag{3.2}$$

using $\psi > 0$ and dominated convergence. Let $A_o \equiv A_{(T,o)} \equiv \{\sigma \in \mathcal{X} : \bar{h}_o(\sigma) = 0\}$ and suppose $\mathbf{u}_{xy}(\sigma, \sigma') = \bar{u}_x(\sigma)\bar{u}_y(\sigma')$: then, since $\mathbf{h}_{xy}(\sigma, \sigma') = 0$ whenever either $\sigma \in A_x$ or $\sigma' \in A_y$, inclusion-exclusion and the unimodularity of μ give

$$\begin{aligned} \widehat{R}_0(\boldsymbol{\delta}) &\geq \mathbb{E}_\mu \left[(2 - 2D_o)\bar{u}_o(A_o) + \sum_{j \in \partial o} \{ \bar{u}_o(A_o) + \bar{u}_j(A_j) - \bar{u}_o(A_o)\bar{u}_j(A_j) \} \right] \\ &= \mathbb{E}_\mu \left[2\bar{u}_o(A_o) - \sum_{j \in \partial o} \bar{u}_o(A_o)\bar{u}_j(A_j) \right] = \mathbb{E}_\mu \left[\sum_{j \in \partial o} \widehat{R}_{o \rightarrow j} \right], \end{aligned}$$

where $\widehat{R}_{o \rightarrow j} \equiv \mathbf{1}\{D_o > 0\}[D_o^{-1}\bar{u}_o(A_o) + D_j^{-1}\bar{u}_j(A_j) - \bar{u}_o(A_o)\bar{u}_j(A_j)]$ (note that necessarily $A_o = \emptyset$ when $D_o = 0$).

Now consider the measurable function $\bar{u} : \mathcal{T}_\mu^+ \rightarrow \Delta_{\mathcal{X}}$ defined (up to μ -equivalence) as follows: if $A_o = \emptyset$ then $\bar{u}_o \equiv \bar{u}_{(T,o)}$ is uniform measure on \mathcal{X} ; otherwise, \bar{u}_o assigns positive mass $1/(2D_o)$ uniformly over A_o and the remaining mass uniformly over A_o^c . Then, $\mathbf{u}_{xy}(\sigma, \sigma') \equiv \bar{u}_x(\sigma)\bar{u}_y(\sigma')$ belongs to $\mathcal{H}_{\text{loc}}^\circ$. With this choice of \mathbf{u} , if both A_o, A_j are non-empty then $\widehat{R}_{o \rightarrow j} = (D_o^{-2} + D_j^{-2} - D_o^{-1}D_j^{-1}/2)/2$, which is positive. Also, $\widehat{R}_{o \rightarrow j}$ is clearly positive when exactly one of A_o, A_j is non-empty, and zero when both are empty. Thus $\widehat{R}_0(\boldsymbol{\delta}) \geq 0$, with strict inequality if $\mu(A_o \neq \emptyset) > 0$. If $\mu(A_o = \emptyset) = 1$, then \mathbf{u} is uniform measure on \mathcal{X}^2 independently of T , so we have by (3.2) that

$$\widehat{R}_0(\boldsymbol{\delta}) = |\mathcal{X}|^{-2} \mathbb{E}_\mu \left[\sum_{j \in \partial o} \sum_{\sigma, \sigma'} \mathbf{1}\{\mathbf{h}_{oj}(\sigma, \sigma') = 0\} \right].$$

If $\mu^\uparrow(\Omega_{\mathbf{h}}) > 0$ then this is positive, completing the proof of our claim. \square

Our main result in this section is the following infinite-tree analogue of [41, Thm. 2], characterizing the interior stationary points of Φ_μ on \mathcal{H}_{loc} as fixed points of the BP recursion.

Proposition 3.3. *If $\psi > 0$, then any stationary point of Φ_μ inside $\mathcal{H}_{\text{loc}}^\circ$ belongs to \mathcal{H}^* .*

Proof. Let $\mathcal{H}_{\text{loc}}^\pm$ denote the space of measurable functions $\boldsymbol{\delta} : \mathcal{T}_e \rightarrow \mathbb{R}^{\mathcal{X}^2}$ (defined up to μ^\uparrow -equivalence), satisfying the symmetry relation $\boldsymbol{\delta}_{xy}(\sigma, \sigma') = \boldsymbol{\delta}_{yx}(\sigma', \sigma)$, with one-point marginals $\bar{\delta}_x(\sigma) \equiv \sum_{\sigma'} \boldsymbol{\delta}_{xy}(\sigma, \sigma')$ not depending on the choice of $y \in \partial x$, and with $\sum_\sigma \bar{\delta}(\sigma) = \sum_{\sigma, \sigma'} \boldsymbol{\delta}(\sigma, \sigma') \equiv 0$. Note that if $\mathbf{h} \in \mathcal{H}_{\text{loc}}^\circ$ and $\boldsymbol{\delta} \in \mathcal{H}_{\text{loc}}^\pm$ with $|\boldsymbol{\delta}| < \mathbf{h}$ in $\mathbb{R}^{\mathcal{X}^2}$ μ^\uparrow -a.s., then $\mathbf{h}^\eta \equiv \mathbf{h} + \eta\boldsymbol{\delta} \in \mathcal{H}_{\text{loc}}^\circ$ for all $|\eta| \leq 1$. If \mathbf{h} is an interior stationary point, then taking $\eta \rightarrow 0$ in (3.1) gives

$$0 = R_0(\boldsymbol{\delta}) \equiv \mathbb{E}_\mu \left[2 \sum_\sigma \bar{\delta}_o(\sigma) \bar{\kappa}'_o(\sigma) + \sum_{j \in \partial o} \sum_{\sigma, \sigma'} \boldsymbol{\delta}_{oj}(\sigma, \sigma') \boldsymbol{\kappa}'_{oj}(\sigma, \sigma') \right],$$

where $\bar{\kappa}'_x \equiv \bar{\xi} + (D_x - 1) \log \bar{h}_x$ and $\boldsymbol{\kappa}'_{xy} \equiv \xi - \log \mathbf{h}_{xy}$.

Consider now $\boldsymbol{\delta}$ with one-point marginals $\bar{\delta} \equiv 0$, so that the value of $\bar{\kappa}'$ becomes irrelevant. In this case the value of $R_0(\boldsymbol{\delta})$ is unchanged if we replace $\boldsymbol{\kappa}'$ by

$$\boldsymbol{\kappa}_{oj}(\sigma, \sigma') \equiv \boldsymbol{\kappa}'_{oj}(\sigma, \sigma') - \frac{1}{|\mathcal{X}|} \sum_\sigma \boldsymbol{\kappa}'_{oj}(\sigma, \sigma') - \frac{1}{|\mathcal{X}|} \sum_{\sigma'} \boldsymbol{\kappa}'_{oj}(\sigma, \sigma') + \frac{1}{|\mathcal{X}|^2} \sum_{\sigma, \sigma'} \boldsymbol{\kappa}'_{oj}(\sigma, \sigma').$$

By construction, $\boldsymbol{\kappa} \in \mathcal{H}_{\text{loc}}^\pm$ with one-point marginals $\bar{\kappa} \equiv 0$ μ -a.s. We claim that $\boldsymbol{\kappa} \equiv 0$ μ^\uparrow -a.s. Indeed, $\boldsymbol{\delta} = c\boldsymbol{\kappa}$ belongs to $\mathcal{H}_{\text{loc}}^\pm$ for any measurable $c : \mathcal{T}_e \rightarrow \mathbb{R}_{>0}$, $(T, x \rightarrow y) \mapsto c_{xy} = c_{yx}$, and we can choose c small enough so that $|\boldsymbol{\delta}| < |\mathbf{h}|$ μ^\uparrow -a.s. With this choice, $0 = R_0(\boldsymbol{\delta})$ is the μ -expectation

of a (weighted) sum of squares, so $\boldsymbol{\kappa} \equiv 0$, as claimed. By the definition of $\boldsymbol{\kappa}$ we find that there exists measurable $\lambda : \mathcal{T}_e \rightarrow \mathbb{R}^{\mathcal{X}}$ such that

$$\boldsymbol{\kappa}'_{oj}(\sigma, \sigma') = \xi(\sigma, \sigma') - \log \mathbf{h}_{oj}(\sigma, \sigma') = -\lambda_{o \rightarrow j}(\sigma) - \lambda_{j \rightarrow o}(\sigma'),$$

so the preceding expression for $R_0(\boldsymbol{\delta})$ (where $\boldsymbol{\delta} \in \mathcal{H}_{\text{loc}}^{\pm}$ with $|\boldsymbol{\delta}| < \mathbf{h}$), simplifies to

$$\begin{aligned} R_0(\boldsymbol{\delta}) &= \mathbb{E}_{\mu} \left[2 \sum_{\sigma} \bar{\delta}_o(\sigma) \bar{\kappa}'_o(\sigma) - \sum_{j \in \partial o} (\bar{\delta}_o(\sigma) \lambda_{o \rightarrow j}(\sigma) + \bar{\delta}_j(\sigma') \lambda_{j \rightarrow o}(\sigma')) \right] \\ &= 2 \mathbb{E}_{\mu} \left[\sum_{\sigma} \bar{\delta}_o(\sigma) \left(\bar{\kappa}'_o(\sigma) - \sum_{j \in \partial o} \lambda_{o \rightarrow j}(\sigma) \right) \right], \end{aligned}$$

using unimodularity of μ for the second identity. Now take $\bar{\delta} : \mathcal{T}_{\bullet} \rightarrow \mathbb{R}^{\mathcal{X}}$ measurable with $\sum_{\sigma} \bar{\delta}(\sigma) \equiv 0$ μ -a.s., and note that $\bar{\delta}_{xy}(\sigma, \sigma') \equiv \bar{\delta}_x(\sigma) \bar{\delta}_y(\sigma')$ is then in $\mathcal{H}_{\text{loc}}^{\pm}$. Considering all such $\boldsymbol{\delta}$ with $|\boldsymbol{\delta}| < \mathbf{h}$ and reasoning in the same way as before, we see that

$$\bar{\kappa}'_o(\sigma) - \sum_{j \in \partial o} \lambda_{o \rightarrow j}(\sigma) = \frac{1}{|\mathcal{X}|} \sum_{\sigma'} \left(\bar{\kappa}'_o(\sigma') - \sum_{j \in \partial o} \lambda_{o \rightarrow j}(\sigma') \right), \quad \mu^{\uparrow}\text{-a.e.}$$

Rearranging, we find that \mathbf{h} satisfies μ^{\uparrow} -a.e.

$$\mathbf{h}_{oj}(\sigma, \sigma') = \psi(\sigma, \sigma') \exp\{\lambda_{o \rightarrow j}(\sigma) + \lambda_{j \rightarrow o}(\sigma')\}, \quad (3.3)$$

$$\bar{h}_o(\sigma) \cong \exp \left\{ \frac{\sum_{j \in \partial o} \lambda_{o \rightarrow j}(\sigma) - \bar{\xi}(\sigma)}{D_o - 1} \right\}. \quad (3.4)$$

If we then reparametrize

$$\lambda_{o \rightarrow j} \equiv \bar{\xi} + \sum_{k \in \partial o \setminus j} \log \hat{m}_{k \rightarrow o} \quad \mu^{\uparrow}\text{-a.s.} \quad (3.5)$$

(well-defined, for each T and $\sigma \in \mathcal{X}$, by invertibility of the D_o -dimensional matrix $\mathbf{1}\mathbf{1}^t - I$), then the formula (3.4) for \bar{h}_o becomes

$$\bar{h}_o(\sigma) \cong \bar{\psi}(\sigma) \prod_{k \in \partial o} \hat{m}_{k \rightarrow o}(\sigma).$$

On the other hand, \bar{h}_o is the first marginal of \mathbf{h}_{oj} , and setting the above equal to the sum of (3.3) over σ' gives (making use of (3.5))

$$\hat{m}_{j \rightarrow o}(\sigma) \cong \sum_{\sigma'} \psi(\sigma, \sigma') e^{\lambda_{j \rightarrow o}(\sigma')}, \quad \mu^{\uparrow}\text{-a.s.}$$

Thus, if we define $m : \mathcal{T}_e \rightarrow \Delta_{\mathcal{X}}$, $m_{x \rightarrow y}(\sigma) \cong e^{\lambda_{x \rightarrow y}(\sigma)}$, then (3.5) can be written in terms of m as

$$m_{o \rightarrow j}(\sigma) \cong \bar{\psi}(\sigma) \prod_{k \in \partial o \setminus j} \left(\sum_{\sigma_k} \psi(\sigma, \sigma_k) m_{j \rightarrow o}(\sigma_k) \right), \quad \mu^{\uparrow}\text{-a.s.},$$

— that is, $m \in \mathcal{H}^*$. Then, (3.3) is precisely the statement that m maps to \mathbf{h} via (1.20), which concludes the proof. \square

Proof of Thm. 1.16. By (H1) the set $\mathcal{H}_{\text{loc}}^{\text{fin}}$ of $\mathbf{h} \in \mathcal{H}_{\text{loc}}$ for which $\Phi(\beta, B, \mathbf{h}) > -\infty$ is non-empty and does not depend on (β, B) , so without loss we will restrict to $\mathbf{h} \in \mathcal{H}_{\text{loc}}^{\text{fin}}$.

Again by (H1), the functions $(\beta, B) \mapsto \Phi_{\mu}(\beta, B, \mathbf{h})$ indexed by $\mathbf{h} \in \mathcal{H}_{\text{loc}}^{\text{fin}}$ are uniformly equicontinuous on compact regions of (β, B) : for any $\epsilon > 0$ there exists $\delta > 0$ sufficiently small so that if

(β, B) and (β', B') are within distance δ then $|\Phi_\mu(\beta, B, \mathbf{h}) - \Phi_\mu(\beta', B', \mathbf{h})| < \epsilon$ for all $\mathbf{h} \in \mathcal{H}_{\text{loc}}^{\text{fin}}$. Let $\mathbf{h} \in \mathcal{H}_{\text{loc}}^{\text{fin}}$ such that $\Phi_\mu(\beta, B, \mathbf{h}) \geq \tilde{\Phi}_\mu(\beta, B) - \epsilon$. Then,

$$\tilde{\Phi}_\mu(\beta', B') \geq \Phi_\mu(\beta', B', \mathbf{h}) \geq \tilde{\Phi}_\mu(\beta, B) - 2\epsilon$$

for all (β', B') within distance δ of (β, B) . Reversing the roles of (β, B) and (β', B') completes the proof of part (a). The statement of part (b) is a summary of the results of Lem. 3.1, Propn. 3.2 and 3.3. \square

4. APPLICATION TO ISING AND POTTS MODELS

In this section we apply Thm. 1.13 to prove our results for the ferromagnetic Ising and Potts models, Thms. 1.8–1.10. Although both models have regimes of multiple fixed points, monotonicity arguments allow us to restrict the space of fixed points. In the Ising model we can restrict to a unique fixed point and give a complete verification of the Bethe free energy prediction; in the Potts model with $q > 2$ there remain regimes of non-uniqueness where we can only provide bounds.

4.1. Ising model. We first prove Thm. 1.8. Recall the definition (1.12) for the Ising measure $\nu_G^{\beta, B}$ for a finite graph $G = (V, E)$, and more generally (from Defn. 1.7) the Ising measures $\nu_{U, G}^{\dagger, \beta, B}$ and $\nu_{U, G}^{+, \beta, B}$ for a finite sub-graph U of a (possibly infinite) graph G with free and $+$ boundary conditions. We will make use of the following direct consequence of the classical Griffiths's inequality (see e.g. [27, Thm. IV.1.21]).

Lemma 4.1. *For the Ising model with parameters $\beta, B \geq 0$ on U a finite sub-graph of a graph G with boundary conditions $\dagger \in \{\text{f}, +\}$, the magnetization $\langle \sigma_v \rangle_{U, G}^{\dagger, \beta, B}$ at vertex $v \in U$ is non-negative, non-decreasing in β, B , non-decreasing in U for $\dagger = \text{f}$, and non-increasing in U for $\dagger = +$.*

Recall from §1.2.1 the definitions of $\bar{h}_T^{t, \dagger}$ for $\dagger \in \{\text{f}, +\}$; the measure $\bar{h}_T^{t, \dagger}$ is parametrized by the corresponding magnetization $\bar{m}_T^{t, \dagger} \equiv \bar{h}_T^{t, \dagger}(+) - \bar{h}_T^{t, \dagger}(-)$. By Lem. 4.1, $\bar{m}_T^{t, \text{f}}$ is non-decreasing in t while $\bar{m}_T^{t, +}$ is non-increasing, so there exist well-defined limits $\bar{m}_T^\dagger(\beta, B) \equiv \lim_{t \rightarrow \infty} \bar{m}_T^{t, \dagger}(\beta, B)$. The following result from [11], an extension of [9, Lem. 4.3], shows that these limits agree on any $T \in \mathcal{T}_\bullet$.

Lemma 4.2 ([11, Lem. 3.1]). *For the Ising model (1.12) on an infinite tree T with $\beta, B > 0$, there exists a constant $C \equiv C(\beta, B)$ such that*

$$\bar{m}_T^{t, +} - \bar{m}_T^{t, \text{f}} \leq C/t \quad \forall t \geq 1.$$

By this result we can define $h \in \mathcal{H}$ by $h_{x \rightarrow y} = \bar{h}_{T_{x \rightarrow y}}^{\text{f}} = \bar{h}_{T_{x \rightarrow y}}^+$, and we now proceed to verify the Bethe prediction $\phi(\beta, B) = \Phi_\mu(\beta, B, h)$.

Proof of Thm. 1.8. The Ising model (1.12) is of form (1.1) with $\mathcal{X} = \{\pm 1\}$, $\xi(\sigma, \sigma') = \beta \sigma \sigma'$, and $\bar{\xi}(\sigma) = B \sigma$, so (H1) and (H2) are clearly satisfied (with no additional moment conditions on D_\circ , since $\psi > 0$). It follows directly from the recursive structure of the tree that $h \in \mathcal{H}^*$. It will be shown in Lem. 4.5 that for $\beta \geq 0$ fixed,

$$\lim_{B \rightarrow \infty} \limsup_{n \rightarrow \infty} |\phi_n(\beta, B) - \Phi_\mu(\beta, B, h)| = 0,$$

so to prove the theorem we will interpolate from (β, B) to (β, B_1) , then take $B_1 \rightarrow \infty$.

It follows from Lem. 4.1 and Lem. 4.2 that for $T \in \mathcal{T}_\bullet$, $\bar{m}_T^{\text{f}}(\beta, B) = \bar{m}_T^+(\beta, B) \equiv m_T(\beta, B)$ is the increasing limit of $\bar{m}_T^{t, \text{f}}(\beta, B)$ and the decreasing limit of $\bar{m}_T^{t, +}(\beta, B)$. The $\bar{m}_T^{t, \dagger}(\beta, B)$ are continuous and non-decreasing in β, B , so m inherits these properties by the same argument as in the proof

of Thm. 1.11, and so (since it takes values in $[-1, 1]$) is of uniformly bounded total variation. This verifies both $(H3^\beta)$ and $(H3^B)$ (though we will use only the latter).

We conclude by showing (cf. (1.19)), that

$$\lim_{n \rightarrow \infty} \mathbb{E}_n[\langle \partial \bar{\xi}(\sigma_{I_n}) \rangle_n^{\beta, B}] = a^{\text{vx}}(\beta, B) \equiv \mathbb{E}_\mu[\|\partial_B \bar{\xi}(\sigma_o)\|_T^{h, \beta, B}].$$

Here $\partial_B \bar{\xi}(\sigma) = \sigma$, and it follows from Lem. 4.1, our assumption of $G_n \rightarrow_{lwc} \mu$ and Fatou's lemma that

$$\begin{aligned} \mathbb{E}_\mu[\|\sigma_o\|_T^{h^f, \beta, B}] &\leq \liminf_{t \rightarrow \infty} \mathbb{E}_\mu[\langle \sigma_o \rangle_{T^t, T}^{f, \beta, B}] \leq \liminf_{n \rightarrow \infty} \mathbb{E}_n[\langle \sigma_{I_n} \rangle_n^{\beta, B}] \\ &\leq \limsup_{n \rightarrow \infty} \mathbb{E}_n[\langle \sigma_{I_n} \rangle_n^{\beta, B}] \leq \limsup_{t \rightarrow \infty} \mathbb{E}_\mu[\langle \sigma_o \rangle_{T^t, T}^{+, \beta, B}] \leq \mathbb{E}_\mu[\|\sigma_o\|_T^{h^+, \beta, B}]. \end{aligned}$$

The left-most and right-most expressions coincide by Lem. 4.2 so equality holds throughout.

By Thm. 1.13(b), $\phi(\beta, B) = \Phi(\beta, B, h^+) = \Phi(\beta, B, h^f)$ for $\beta \geq 0$, $B > 0$. Since ϕ_n is symmetric in B and continuous at $B = 0$ (uniformly in n), we have $\phi(\beta, B) = \phi(\beta, -B)$ and $\phi(\beta, 0) = \lim_{B \rightarrow 0} \phi(\beta, B)$. \square

4.2. Potts model. We now apply Thm. 1.13 to deduce our result Thm. 1.9 for the Potts model (1.3) with $\beta, B \geq 0$. From now on we let $\mathcal{X} \equiv [q]$ with $q \geq 2$. It will be convenient to generalize (1.3) to the *inhomogeneous* Potts model

$$\nu_G^{\beta, B}(\underline{\sigma}) \cong \exp \left\{ \sum_{(ij) \in E} \beta_{ij} \delta_{\sigma_i, \sigma_j} + \sum_{i \in V} B_i \delta_{\sigma_i, 1} \right\}, \quad \underline{\sigma} \in \mathcal{X}^V.$$

We now introduce the coupling of the Potts model with a random-cluster model which we use to obtain monotonicity properties. The following representation is as in [20] (see also [6]). If $G = (V, E)$ is a finite graph, let G^* be the graph formed by adding an edge from every $v \in V$ to a ‘‘ghost vertex’’ v^* , that is, $G^* = (V^*, E^*)$ where $V^* = V \cup \{v^*\}$ and $E^* = E \cup \{(v, v^*) : v \in V\}$. Writing $\underline{\sigma}$ for elements of \mathcal{X}^{V^*} and $\underline{\eta}$ for elements of $\{0, 1\}^{E^*}$ (*bond configurations*), consider the probability measure on pairs $(\underline{\sigma}, \underline{\eta})$ defined by

$$\varpi_G^{\beta, B}(\underline{\sigma}, \underline{\eta}) \cong \delta_{\sigma_{v^*}, 1} \prod_{\eta_{ij}=1} (e^{\beta_{ij}} - 1) \delta_{\sigma_i, \sigma_j} \prod_{\eta_i=1} (e^{B_i} - 1) \delta_{\sigma_i, \sigma_{v^*}}. \quad (4.1)$$

The marginal on $\underline{\sigma}_V$ is the inhomogeneous Potts measure $\nu_G^{\beta, B}$, while the marginal on $\underline{\eta}$ is the (*inhomogeneous*) *random-cluster measure*

$$\pi_G^{\beta, B}(\underline{\eta}) \cong \prod_{e \in E^*} p_e^{\eta_e} (1 - p_e)^{1 - \eta_e} \prod_{C \in \underline{\eta}} \Theta(C), \quad (4.2)$$

where $p_{ij} \equiv 1 - e^{-\beta_{ij}}$ for $(i, j) \in E$ and $p_{iv^*} \equiv 1 - e^{-B_i}$ for $i \in V$, and the last product is taken over connected components C of $\underline{\eta}$, with $\Theta(C) = q$ unless $v^* \in C$ in which case $\Theta(C) = 1$. Given a configuration $\underline{\eta}$, a realization of the conditional law $\varpi_G^{\beta, B}(\underline{\sigma} = \cdot | \underline{\eta})$ is obtained by choosing a constant spin on each connected component C of $\underline{\eta}$ independently and uniformly over $[q]$, except for C containing v^* which is given spin 1.

For a detailed account the random-cluster model see [19]; we will use only the following basic properties:

Proposition 4.3. *The random-cluster measure $\pi_G^{\beta, B}$ is FKG. It is also increasing, in the sense of stochastic domination, in (β, B) .*

Proof. The FKG property follows by a straightforward modification of the proof of [6, Thm. III.1(i)]. Monotonicity in $(\underline{\beta}, \underline{B})$ follows by modifying the proof of [19, Thm. 3.21]. \square

Recalling Defn. 1.7, for U a finite sub-graph of a graph G and $\ddagger \in \{f\} \cup [q]$ (with $f = \text{free}$), let $\nu_{U,G}^{\ddagger,\beta,B}$ denote the Potts model on U with \ddagger boundary conditions.

Corollary 4.4. *For the Potts model with parameters $\beta, B \geq 0$ on U a finite sub-graph of a graph G with boundary conditions $\ddagger \in \{f, 1\}$, and for any vertices $v, w \in U$, the quantities*

$$\nu_{U,G}^{\ddagger,\beta,B}(\sigma_v = 1), \quad \nu_{U,G}^{\ddagger,\beta,B}(\sigma_v = \sigma_w)$$

are non-decreasing in β and B , non-increasing in U for $\ddagger = 1$, and non-decreasing in U for $\ddagger = f$.

Proof. Note that $\nu_{U,G}^{f,\beta,B}$ is the marginal on $\underline{\sigma}_U$ of the measure $\varpi_G^{\beta,B}$ with

$$B_i = B \quad \forall i \in V, \quad \beta_e = \beta \mathbf{1}\{e \in E_U\}.$$

Similarly, $\nu_{U,G}^{1,\beta,B}$ is the marginal on $\underline{\sigma}_U$ of the measure $\varpi_G^{\beta',B'}$ with

$$\beta'_e = \beta \quad \forall e \in E, \quad B'_i = B \mathbf{1}\{i \in V_U\} + \infty \mathbf{1}\{i \notin V_U\}.$$

Clearly, $(\underline{\beta}, \underline{B})$ is non-decreasing in U while $(\underline{\beta}', \underline{B}')$ is non-increasing, and both are non-decreasing in β, B . The result therefore follows from Propn. 4.3 by showing that for any $(\underline{\beta}, \underline{B})$, the conditional probabilities $\varpi_G^{\beta,B}(\sigma_v = 1|\underline{\eta})$ and $\varpi_G^{\beta,B}(\sigma_v = \sigma_w|\underline{\eta})$ are monotone functions of $\underline{\eta}$. Indeed, letting $\varpi \equiv \varpi_G^{\beta,B}$ and writing $v \rightsquigarrow w$ to indicate that v, w belong to the same connected component of $\underline{\eta}$, we have

$$\begin{aligned} \varpi(\sigma_v = 1|\underline{\eta}) &= \mathbf{1}\{v \rightsquigarrow v^*\} + \frac{1 - \mathbf{1}\{v \rightsquigarrow v^*\}}{q}, \\ \varpi(\sigma_v = \sigma_w|\underline{\eta}) &= \mathbf{1}\{v \rightsquigarrow w\} + \frac{1 - \mathbf{1}\{v \rightsquigarrow w\}}{q}. \end{aligned}$$

These are increasing functions of $\underline{\eta}$ so the proof is complete. \square

Under the measures with $\ddagger \in \{f, 1\}$, any one-vertex marginal must be uniform on the spins $\neq 1$, and so is characterized by the probability given to spin 1. In particular, recall from §1.2.2 the definitions of $\bar{h}_T^{t,\ddagger}$ for $\ddagger \in \{f, 1\}$; existence of the $t \rightarrow \infty$ limits \bar{h}_T^\ddagger is now justified by Cor. 4.4, so we can define $h^\ddagger \in \mathcal{H}$ by $h_{x \rightarrow y}^\ddagger = \bar{h}_{T_{x \rightarrow y}}^\ddagger$. The following lemma gives the boundary values for the interpolation in (β, B) using h^\ddagger :

Lemma 4.5. *For the Potts model on $G_n \rightarrow_{lwc} \mu$, let*

$$\tilde{\Phi}_\mu(\beta, B) \equiv B + \beta \mathbb{E}_\mu[D_o]/2 + \mathbb{E}_\mu[\bar{\varphi}(|T|)], \quad \bar{\varphi}(n) \equiv \bar{\varphi}^B(n) \equiv n^{-1} \log(1 + (q-1)e^{-Bn}).$$

(a) *For all $B \in \mathbb{R}$ and any $h \in \mathcal{H}$, $\phi(0, B) = \log(e^B + q - 1) = \Phi_\mu(0, B, h)$.*

(b) *For $\beta \geq 0$ and $h \in \mathcal{H}^*$,*

$$\lim_{B \rightarrow \infty} \limsup_{n \rightarrow \infty} |\phi_n(\beta, B) - \tilde{\Phi}_\mu(\beta, B)| = 0 = \lim_{B \rightarrow \infty} |\Phi_\mu(\beta, B, h) - \tilde{\Phi}_\mu(\beta, B)|$$

(c) *For $B \geq 0$, $\lim_{\beta \rightarrow \infty} \limsup_{n \rightarrow \infty} |\phi_n(\beta, B) - \tilde{\Phi}_\mu(\beta, B)| = 0$.*

(d) *For $B > 0$ and $\ddagger \in \{f, 1\}$, $\lim_{\beta \rightarrow \infty} |\Phi_\mu(\beta, B, h^\ddagger) - \tilde{\Phi}_\mu(\beta, B)| = 0$.*

Proof. (a) At $\beta = 0$, $\psi \equiv 1$ so the spins are independent. Thus, for all $n \geq 1$, $h \in \mathcal{H}$, and $T \in \mathcal{T}_\bullet$,

$$\phi_n(0, B) = \log(e^B + q - 1) = \Phi_T^{\text{yx}}(0, B, h) = \Phi_T(0, B, h),$$

since $\Phi_T^{(oj)} \equiv 0$ for all $j \in \partial o$.

(b) The value of $Z_n(\beta, B)$ is bounded below by considering only the ground state $\underline{\sigma} \equiv 1$, and bounded above by decomposing \mathcal{X}^V according to the subset of k vertices where the spin is not 1. For $\beta \geq 0$ this gives

$$1 \leq Z_n(\beta, B)e^{-Bn - \beta|E_n|} \leq \sum_{k=0}^n \binom{n}{k} (q-1)^k e^{-Bk} = (1 + (q-1)e^{-B})^n,$$

so if we define $\bar{\phi}_n(\beta, B) \equiv \phi_n(\beta, B) - B - \beta \mathbb{E}_n[|E_n|]/n$, then $\lim_{B \rightarrow \infty} \limsup_{n \rightarrow \infty} |\bar{\phi}_n(\beta, B)| = 0$. Recalling (2.2), this proves the left identity in (b).

We next define

$$\bar{\Phi}_T^{\text{yx}} \equiv \Phi_T^{\text{yx}} - B - \beta D_o, \quad \bar{\Phi}_T^e \equiv \Phi_T^e - \beta D_o/2, \quad \bar{\Phi}_T \equiv \bar{\Phi}_T^{\text{yx}} - \bar{\Phi}_T^e, \quad \bar{\Phi}_\mu \equiv \mathbb{E}_\mu \bar{\Phi}_T,$$

so that to prove the right identity in (b) it suffices to show $\lim_{B \rightarrow \infty} \bar{\Phi}_\mu(\beta, B, h) = 0$ for any $h \in \mathcal{H}^*$. Indeed, (1.15) gives that μ -a.s., $\lim_{B \rightarrow \infty} h_{o \rightarrow j}^{\beta, B}(\sigma) = \delta_{\sigma, 1}$ for all $j \in \partial o$, hence also $\lim_{B \rightarrow \infty} h_{j \rightarrow o}^{\beta, B}(\sigma) = \delta_{\sigma, 1}$ for all $j \in \partial o$ by equivalence of μ^\uparrow and μ^\downarrow . Thus

$$\lim_{B \rightarrow \infty} \bar{\Phi}_T^{\text{yx}}(\beta, B, h) = 0 = \lim_{B \rightarrow \infty} \bar{\Phi}_T^e(\beta, B, h), \quad \mu\text{-a.s.}$$

It is easily verified that

$$-\beta D_o \leq \bar{\Phi}_T^{\text{yx}}(\beta, B, h) \leq \log q, \quad -\beta D_o/2 \leq \bar{\Phi}_T^e(\beta, B, h) \leq 0, \quad (4.3)$$

so $\bar{\Phi}_\mu(\beta, B, h) \rightarrow 0$ by dominated convergence.

(c) Suppose first that G_n is connected. Then $Z_n(\beta, B)$ is bounded below by considering only the q constant-spin configurations, and bounded above by decomposing \mathcal{X}^V according to the subset of ℓ edges across which the spins disagree. Since G_n is connected, removing ℓ edges leaves at most $\ell + 1$ connected components, of sizes k_0, \dots, k_ℓ summing to n . Therefore, with $\varphi(n) \equiv \varphi^B(n) \equiv n\bar{\varphi}^B(n)$, we have

$$e^{\varphi(n)} \leq Z_n(\beta, B)e^{-Bn - \beta|E_n|} \leq \sum_{\ell=0}^{|E_n|} \binom{|E_n|}{\ell} e^{-\beta\ell} \max_{k_0, \dots, k_\ell} \left\{ \exp \left\{ \sum_{r=0}^{\ell} \varphi(k_r) \right\} \right\},$$

where the maximum is taken over $k_0, \dots, k_\ell \in \mathbb{Z}_{\geq 0}$ summing to n . By convexity of φ this maximum is achieved with $k_r = n$ for some r , so

$$\varphi(n) \leq n\bar{\phi}_n(\beta, B) \leq \varphi(n) + \mathbb{E}_n \left[\log \left\{ \sum_{\ell=0}^{|E_n|} \binom{|E_n|}{\ell} e^{-\beta\ell} q^\ell \right\} \right] = \varphi(n) + \mathbb{E}_n[|E_n|] \log(1 + qe^{-\beta}). \quad (4.4)$$

If G_n has connected components $C^j = (V^j, E^j)$, $j \geq 1$, with $|V^j| = n^j$, then clearly $Z_n(\beta, B) = \prod_j Z_{C^j}(\beta, B)$, so

$$0 \leq \bar{\phi}_n(\beta, B) - \frac{1}{n} \mathbb{E}_n \left[\sum_j \varphi(n^j) \right] \leq \frac{1}{n} \mathbb{E}_n[|E_n|] \log(1 + qe^{-\beta}). \quad (4.5)$$

With $j(i)$ denoting the index of the connected component of G_n containing vertex i , we have $n^{-1}\mathbb{E}_n[\sum_j \varphi(n^j)] = \mathbb{E}_n[\bar{\varphi}(n^{j(I_n)})]$. Then, since $\bar{\varphi}'(n) \leq 0$,

$$\mathbb{E}_n[\bar{\varphi}(|B_t(I_n)|)\mathbf{1}\{B_t(I_n) = C^{j(I_n)}\}] \leq \mathbb{E}_n[\bar{\varphi}(n^{j(I_n)})] \leq \mathbb{E}_n[\bar{\varphi}(|B_t(I_n)|)].$$

Since $G_n \xrightarrow{law} \mu$, letting $n \rightarrow \infty$ followed by $t \rightarrow \infty$ in the above inequalities gives $\mathbb{E}_n[\bar{\varphi}(n^{j(I_n)})] \rightarrow \mathbb{E}_\mu[\bar{\varphi}(|T|)]$, and so (c) follows from (4.5) by taking first $n \rightarrow \infty$ and then $\beta \rightarrow \infty$.

(d) Clearly $h_T^f = h_T^1$ for any finite $T \in \mathcal{T}_\bullet$ (as $\partial T^t = \emptyset$ for large enough t). In the $\beta \rightarrow \infty$ limit only the constant-spin configurations contribute, so

$$\lim_{\beta \rightarrow \infty} h_T^{\ddagger, \beta, B}(\sigma) = e^{-\varphi(|T|) - B|T|(1 - \delta_{\sigma, 1})}, \quad \ddagger \in \{f, 1\}. \quad (4.6)$$

For T infinite, recall from Cor. 4.4 that $h_{T^t}^f(1) \leq h_T^f(1) \leq h_T^1(1)$, so if $B > 0$ then

$$1 = \lim_{t \rightarrow \infty} \lim_{\beta \rightarrow \infty} h_{T^t}^{f, \beta, B}(1) \leq \lim_{\beta \rightarrow \infty} h_T^{f, \beta, B}(1) \leq \lim_{\beta \rightarrow \infty} h_T^{1, \beta, B}(1),$$

so that (4.6) again holds for T infinite. We then compute

$$\begin{aligned} \lim_{\beta \rightarrow \infty} \bar{\Phi}_T^{\text{vx}}(\beta, B, h^\ddagger) &= - \sum_{j \in \partial o} \varphi(|T_{j \rightarrow o}|) + \varphi(|T|), \\ \lim_{\beta \rightarrow \infty} \bar{\Phi}_T^e(\beta, B, h^\ddagger) &= -\frac{1}{2} \sum_{j \in \partial o} \varphi(|T_{j \rightarrow o}|) - \frac{1}{2} \sum_{j \in \partial o} \varphi(|T_{o \rightarrow j}|) + \frac{D_o}{2} \varphi(|T|), \end{aligned}$$

μ -a.s., where the first identity uses $|T| = 1 + \sum_{j \in \partial o} |T_{j \rightarrow o}|$ and the second uses $|T| = |T_{o \rightarrow j}| + |T_{j \rightarrow o}|$. Convergence also holds in μ -expectation, using the upper bounds in (4.3) together with

$$\bar{\Phi}_T^{\text{vx}}(\beta, B, h^\ddagger) \geq \sum_{j \in \partial o} \log h_{j \rightarrow o}^{f, \beta, B}(1), \quad \bar{\Phi}_T^e(\beta, B, h^\ddagger) \geq \frac{1}{2} \sum_{j \in \partial o} \log h_{j \rightarrow o}^{f, \beta, B}(1) + \frac{1}{2} \sum_{j \in \partial o} \log h_{o \rightarrow j}^{f, \beta, B}(1),$$

and the fact that $h_{x \rightarrow y}^{f, \beta, B}(1) \geq 1/q$ for $\beta, B \geq 0$ (by Cor. 4.4). Thus, using unimodularity of μ , we have

$$\lim_{\beta \rightarrow \infty} \bar{\Phi}_\mu(\beta, B, h^\ddagger) = \mathbb{E}_\mu[(1 - D_o/2)\varphi(|T|)],$$

and we conclude by showing that this coincides with $\mathbb{E}_\mu[\bar{\varphi}(|T|)]$. The case $|T| = \infty$ is trivial; otherwise, another application of unimodularity gives

$$\begin{aligned} \frac{1}{2} \mathbb{E}_\mu[D_o \varphi(|T|)] &= \frac{1}{2} \mathbb{E}_\mu \left[D_o \sum_{x \in T} \bar{\varphi}(|T|) \right] = \frac{1}{2} \mathbb{E}_\mu \left[\sum_{x \in T} D_x \bar{\varphi}(|T|) \right] \\ &= \mathbb{E}_\mu[\bar{\varphi}(|T|)|E_T] = \mathbb{E}_\mu[\varphi(|T|)] - \mathbb{E}_\mu[\bar{\varphi}(|T|)]. \end{aligned}$$

Therefore, $\lim_{\beta \rightarrow \infty} \bar{\Phi}_\mu(\beta, B, h) = \mathbb{E}_\mu[\bar{\varphi}(|T|)]$ which concludes the proof. \square

Proof of Thm. 1.9. The Potts model (1.3) is of form (1.1) with $\mathcal{X} = [q]$, $\xi(\sigma, \sigma') = \beta \delta_{\sigma, \sigma'}$, and $\bar{\xi}(\sigma) = B \delta_{\sigma, 1}$, so (H1) and (H2) are clearly satisfied. It follows from the recursive structure of the tree that $h^\ddagger \in \mathcal{H}^*$ for $\ddagger \in \{f, 1\}$. For part (a), along any interpolation path contained in \mathcal{R}_μ , both (H3 $^\beta$) and (H3 B) are satisfied by Cor. 4.4 and the same argument used in the proof of Thm. 1.11. For part (b), (H3 $^\beta$) and (H3 B) are satisfied by the additional hypothesis of continuity.

The inequalities in part (b) then follow from Theorem 1.13 once we verify (cf. (1.18), (1.19))

$$\begin{aligned} a^{\text{vx}}(\beta, B, h^{\text{f}}) &\leq \liminf_{n \rightarrow \infty} a_n^{\text{vx}}(\beta, B) \leq \limsup_{n \rightarrow \infty} a_n^{\text{vx}}(\beta, B) \leq a^{\text{vx}}(\beta, B, h^1), \\ a^{\text{e}}(\beta, B, h^{\text{f}}) &\leq \liminf_{n \rightarrow \infty} a_n^{\text{e}}(\beta, B) \leq \limsup_{n \rightarrow \infty} a_n^{\text{e}}(\beta, B) \leq a^{\text{e}}(\beta, B, h^1), \end{aligned}$$

where $a_n^{\text{vx}}(\beta, B) = \mathbb{E}_n[\langle \delta_{\sigma_{I_n, 1}} \rangle_n^{\beta, B}]$ and $a_n^{\text{e}}(\beta, B) = \frac{1}{2} \mathbb{E}_n[\sum_{j \in \partial I_n} \langle \delta_{\sigma_{I_n, \sigma_j}} \rangle_n^{\beta, B}]$. Indeed, by Vitali's convergence theorem, the assumption $G_n \rightarrow_{\text{lwc}} \mu$ and Cor. 4.4 (with $U = B_t(I_n) \subseteq G_n$), we have

$$a^{\text{e}}(\beta, B, h^{\text{f}}) = \liminf_{t \rightarrow \infty} \frac{1}{2} \mathbb{E}_\mu \left[\sum_{j \in \partial o} \langle \delta_{\sigma_o, \sigma_j} \rangle_{T^t, T}^{\text{f}, \beta, B} \right] \leq \liminf_{n \rightarrow \infty} a_n^{\text{e}}(\beta, B),$$

and the other inequalities are proved similarly. Together these inequalities imply that

$$\lim_{n \rightarrow \infty} [\phi_n(\beta', B') - \phi_n(\beta, B)] = \Phi(\beta', B', h^{\ddagger}) - \Phi(\beta, B, h^{\ddagger})$$

for any (β, B) and (β', B') joined by an interpolation path contained in \mathcal{R}_μ . The result of part (a) then follows by letting (β', B') approach \mathcal{R}_∞ and applying Lem. 4.5. \square

4.3. Potts model with d -regular limiting tree. In this section we prove Thm. 1.10, which amounts to determining the shape of \mathcal{R}_\neq and establishing continuity of h^{f} and h^1 in certain regimes.

Since the limiting measure is supported on \mathbb{T}_d , only $h \equiv h_{(\mathbb{T}_d, x \rightarrow y)}$ is of relevance. Further, h^{\ddagger} is symmetric among the spins $\neq 1$ for $\ddagger \in \{\text{f}, 1\}$, so determination of h^{\ddagger} reduces to solving a univariate recursion for $h^{\ddagger}(1)$,

$$h \mapsto \frac{e^B [e^\beta h + (1-h)]^{d-1}}{e^B [e^\beta h + (1-h)]^{d-1} + (q-1) [h + \frac{1-h}{q-1} (e^\beta + q - 2)]^{d-1}}.$$

Our result follows from analysis of the fixed points of this mapping (similar computations have appeared e.g. in [38, 42] so some overlap among the analyses may occur).

A convenient parametrization is given by the log likelihood ratio $r \equiv \log h - \log[(1-h)/(q-1)]$, in terms of which the recursion becomes

$$r \mapsto f(r) \equiv f(r; \beta, B) = B + (d-1) \log \left(\frac{e^{\beta+r} + q - 1}{e^r + e^\beta + q - 2} \right).$$

With $f^{(t)}$ the t -fold iteration of f , let r^{f} denote the increasing limit of $f^{(t)}(0)$ and r^1 the decreasing limit of $f^{(t)}(\infty)$, as $t \rightarrow \infty$. The region \mathcal{R}_\neq corresponds to those $\beta, B \geq 0$ for which $r^{\text{f}} \neq r^1$.

Lemma 4.6. *There exists $\beta_- > 0$ such that for $\beta \leq \beta_-$ the map f has exactly one fixed point for any $B \in \mathbb{R}$. For $\beta > \beta_-$ there exist real-valued $B_-(\beta) < B_+(\beta)$ (smooth in β) such that f has one, two, or three fixed points depending on whether B is in $[B_-, B_+]^c$, $\{B_-, B_+\}$, or (B_-, B_+) . The curves extend continuously to $B_-(\beta_-) = B_+(\beta_-)$.*

Proof. We have

$$f'(r) = \frac{(d-1)e^r (e^\beta - 1) (q + e^\beta - 1)}{(q + e^r + e^\beta - 2) (q + e^{r+\beta} - 1)} \quad (4.7)$$

so f is increasing in r with $f'(r) \rightarrow 0$ as $r \rightarrow \pm\infty$. Since $f(r; \beta, B) = f(\beta; r, B)$, it easily follows from (4.7) that $\partial_\beta f(r)$ has the same sign as r while $\partial_\beta [f'(r)] > 0$. Further

$$f''(r) = - \frac{(d-1)e^{r+\beta} (e^\beta - 1) (q + e^\beta - 1) (e^{2r} - \alpha)}{(q + e^r + e^\beta - 2)^2 (q + e^{r+\beta} - 1)^2}, \quad \alpha \equiv (q-1)(1 + (q-2)e^{-\beta}),$$

with $\alpha > 0$ since $q > 1$. Notice that $f''(r) > 0$ for r sufficiently negative and $f''(r) < 0$ for r sufficiently positive, with a single sign change occurring at $(\log \alpha)/2$ which is zero for $q = 2$ and strictly positive for $q > 2$. This proves that f has between one and three fixed points. When $B = 0$, one fixed point is always given by $r^f(\beta, 0) = 0$. Further $f(r; 0, 0) \equiv 0$, so (by monotonicity of f' in β) there exists $\infty \geq \beta_- \geq 0$ such that $f' \leq 1$ everywhere for $\beta \leq \beta_-$ and f' exceeds 1 somewhere for $\beta > \beta_-$.

Solving the equation $f'(r) = 1$ in terms of $t \equiv e^r$ yields solutions

$$t_{\pm}(\beta) = -\gamma \pm \sqrt{\gamma^2 - \alpha}, \quad \gamma \equiv e^{\beta} + q - 2 - \frac{d}{2}(1 - e^{-\beta})(e^{\beta} + q - 1).$$

Since $\alpha > 0$, $t_{\pm}(\beta)$ are not positive if $\gamma > -\sqrt{\alpha}$, equal to $\sqrt{\alpha} > 0$ if $\gamma = -\sqrt{\alpha}$, and positive but not equal if $\gamma < -\sqrt{\alpha}$. If $d \geq 2$, it is easy to check that both α and γ decrease smoothly in β , starting at $\gamma|_{\beta=0} = q - 1$ and $\alpha|_{\beta=0} = (q - 1)^2$, so there is a unique value $\beta = \beta_- > 0$ at which $\gamma = -\sqrt{\alpha}$: if $d = 2$ then $\beta_- = \infty$, and if $d > 2$ then β_- is the logarithm of the unique finite positive root b_- of

$$(d - 2)^2 b^2 + (d - 2)^2 (q - 2)b - d^2 (q - 1) = 0. \quad (4.8)$$

Hence, the equation $f'(r) = 1$ has no solutions for $\beta < \beta_-$, and it has solutions $\rho_{\pm}(\beta) \equiv \log t_{\pm}(\beta)$ for $\beta \geq \beta_-$, with $\rho_-(\beta_-) = \rho_+(\beta_-)$ and $\rho_-(\beta) < \rho_+(\beta)$ for $\beta > \beta_-$. The values of $B_-(\beta)$, $B_+(\beta)$ are then given explicitly by

$$B_{\pm}(\beta) = \rho_{\mp}(\beta) - f(\rho_{\mp}(\beta); \beta, 0), \quad (4.9)$$

which clearly meet at $\beta = \beta_-$ and are smooth for $\beta > \beta_-$. \square

Considering hereafter only $d > 2$ (so that $\beta_- < \infty$), suppose $\beta > \beta_-$, so that the functions ρ_{\pm} are defined. Since $\partial_{\beta}[f'(r)] > 0$, ρ_- and ρ_+ must be, respectively, decreasing and increasing in β . Further, since f has a unique inflection point at $(\log \alpha)/2$, we must have $\rho_-(\beta) \leq (\log \alpha)/2 \leq \rho_+(\beta)$, with strict inequalities unless $\rho_-(\beta) = \rho_+(\beta)$. For $q = 2$ (Ising), this implies $\rho_- \leq 0 \leq \rho_+$ from which it is easy to see that whenever $B > 0$ we have $r^f(\beta, B) = r^1(\beta, B)$, which is then continuous in (β, B) by the same argument as in the proof of Thm. 1.11. When $B = 0$, $r^f(\beta, 0)$ is zero for all β , while $r^1(\beta, B)$ is zero for $\beta \leq \beta_-$ and strictly positive for $\beta > \beta_-$.

For $q > 2$ (Potts), this implies that $\rho_+(\beta, B) > 0$ while $\rho_-(\beta, B) \geq 0$ if and only if $f'(0; \beta, B) \leq 1$. From the calculations above, $f'(0)$ is zero at $\beta = 0$ and increases in β . We therefore define

$$\begin{aligned} \beta_f &\equiv \inf\{\beta \geq 0 : f(r; \beta, 0) = r \text{ for some } r > 0\}, \\ \beta_+ &\equiv \inf\{\beta \geq 0 : \rho_-(\beta) \leq 0\} = \inf\{\beta \geq 0 : f'(0; \beta, 0) \geq 1\} = \log\left(1 + \frac{q}{d-2}\right) \end{aligned} \quad (4.10)$$

(where the formula for β_+ comes from (4.7)). Clearly $\beta_- \leq \beta_f \leq \beta_+$, and in fact these inequalities are strict: at β_f , f' must exceed one between zero and the positive fixed point, so $\beta_- < \beta_f$.² Likewise, if $f'(0) \geq 1$ at $\beta = \beta_f$, the concavity of $f(r)$ at $r = 0$ would imply the existence of a positive fixed point at some β below β_f which is a contradiction, so $\beta_f < \beta_+$.

We refer again to Fig. 1 which illustrates this for $q = 30$, $d = 10$. Fig. 1a shows the map $f(r; 0.7, 0.2)$ with three positive fixed points. Fig. 1b shows the region of $(\beta, B) \geq (0, 0)$ delineated by the curves $B_{\pm}(\beta)$.

Proof of Thm. 1.10. (a) We found above that $\mathcal{R}_{\neq} = \emptyset$ for $d = 2$ and $\mathcal{R}_{\neq} = (\beta_-, \infty)$ for $q = 2$, so suppose $d, q > 2$. If $B > 0$, $r^f = r^1$ holds for all $\beta \geq 0$ with $\beta \notin (\beta_-, \beta_+)$. For $\beta \in (\beta_-, \beta_+)$ there is a closed interval $[B_-(\beta) \vee 0, B_+(\beta)]$ of B values for which $r^f < r^1$: this interval is strictly positive

²Note that $r^1(\beta_f, 0) > 0$, i.e., the 1-biased fixed point ‘‘arises discontinuously.’’

for $\beta < \beta_f$ and includes zero for $\beta \geq \beta_f$. If $B = 0$, $r^f = r^1$ for $0 \leq \beta < \beta_f$ and $r^f < r^1$ for $\beta \geq \beta_f$. Recalling (4.9),

$$\partial_\beta B_\pm(\beta) = \partial_\beta[\rho_\mp(\beta) - f(\rho_\mp(\beta))] = [1 - f'(\rho_\mp(\beta))]\partial_\beta \rho_\mp(\beta) - (\partial_\beta f)(\rho_\mp(\beta)) = -(\partial_\beta f)(\rho_\mp(\beta)).$$

This has the same sign as $-\rho_\mp(\beta)$, which are both negative for $0 \leq \beta < \beta_+$, so the curves $B_\pm(\beta)$ are decreasing. Inverting them gives the curves $\beta_f(B), \beta_+(B)$ which delineate the region \mathcal{R}_\neq as described in the theorem statement, with $\beta_f(0) = \beta_f$ and $\beta_+(0) = \beta_+$.

(b) Away from the boundary of \mathcal{R}_\neq , h^f and h^1 correspond to isolated zeroes of a smooth function, and so are continuous by the implicit function theorem. From part (a), any point of \mathcal{R} is connected to \mathcal{R}_∞ by an interpolation path contained in \mathcal{R} , so applying Thm. 1.9(a) verifies the Bethe prediction for $(\beta, B) \notin \mathcal{R}_\neq$.

Since changing B only translates $f(r; \beta, B)$, it is not difficult to see that when $\beta \in (\beta_-, \beta_+)$, the function $h^f(\beta, B)$ is continuous in B for $B \in [0, B_+(\beta)]$ while $h^1(\beta, B)$ is continuous for $B \in [B_-(\beta) \vee 0, \infty)$. It follows by Lem. 2.1 that for $(\beta, B) \in \partial\mathcal{R}_\neq$ with $\beta = \beta_f(B)$, $\phi(\beta, B) = \Phi(\beta, B, h^f)$, while for $(\beta, B) \in \partial\mathcal{R}_\neq$ with $\beta \geq \beta_+(B)$, $\phi(\beta, B) = \Phi(\beta, B, h^1)$.

Recall our convention that $\beta_0 \leq \beta_1, B_0 \leq B_1$: by Thm. 1.9(b) we may interpolate in B from $(\beta, B_0) \in \mathcal{R}_\neq^\circ$ to $(\beta, B_1) \in \mathcal{R}$ using the message h^1 , yielding $\liminf_{n \rightarrow \infty} \phi_n(\beta, B) \geq \Phi(\beta, B, h^1)$ for $(\beta, B) \in \mathcal{R}_\neq^\circ$. Likewise, we may interpolate in B from $(\beta, B_0) \in \mathcal{R}$ to $(\beta, B_1) \in \mathcal{R}_\neq$ using h^f (and once inside \mathcal{R}_\neq we may also interpolate in β using h^f), which gives $\liminf_{n \rightarrow \infty} \phi_n(\beta, B) \geq \Phi(\beta, B, h^f)$ for $(\beta, B) \in \mathcal{R}_\neq^\circ$.

Next, since $h^f(\beta, B)$ and $h^1(\beta, B)$ are lower and upper semi-continuous respectively in β , and both are non-decreasing in β , for $0 < B < B_+$ we have that $h^f(\beta, B) \uparrow h^f(\beta_+(B), B)$ as $\beta \uparrow \beta_+(B)$ and $h^1(\beta, B) \downarrow h^1(\beta_f(B), B)$ as $\beta \downarrow \beta_f(B)$. Again by Thm. 1.9(b), we may interpolate in β from $(\beta_0, B) = (\beta_f(B), B) \in \partial\mathcal{R}_\neq$ to $(\beta_1, B) \in \mathcal{R}_\neq^\circ$ using h^1 , and from $(\beta_0, B) \in \mathcal{R}_\neq^\circ$ to $(\beta_1, B) = (\beta_+(B), B) \in \partial\mathcal{R}_\neq$ using h^f , giving

$$\limsup_{n \rightarrow \infty} \phi_n(\beta, B) \leq \min\{\tilde{\Phi}^f(\beta, B), \tilde{\Phi}^1(\beta, B)\}, \quad (\beta, B) \in \mathcal{R}_\neq^\circ,$$

which completes the proof. \square

REFERENCES

- [1] D. Aldous and R. Lyons. Processes on unimodular random networks. *Electron. J. Probab.*, 12(54):1454–1508, 2007.
- [2] A. Bandyopadhyay and D. Gamarnik. Counting without sampling. New algorithms for enumeration problems using statistical physics. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 890–899, New York, 2006. ACM.
- [3] A. Bandyopadhyay and D. Gamarnik. Counting without sampling: asymptotics of the log-partition function for certain statistical physics models. *Random Structures Algorithms*, 33(4):452–479, 2008.
- [4] M. Bayati, D. Gamarnik, and P. Tetali. Combinatorial approach to the interpolation method and scaling limits in sparse random graphs, 2009.
- [5] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6(23):13 pp. (electronic), 2001.
- [6] M. Biskup, C. Borgs, J. T. Chayes, and R. Kotecký. Gibbs states of graphical representations of the Potts model with external fields. *J. Math. Phys.*, 41(3):1170–1210, 2000. Probabilistic techniques in equilibrium and nonequilibrium statistical physics.
- [7] S. A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of computing*, STOC '71, pages 151–158, New York, NY, USA, 1971. ACM.
- [8] A. Dembo and A. Montanari. Gibbs measures and phase transitions on sparse random graphs. *Braz. J. Probab. Stat.*, 24(2):137–211, 2010.
- [9] A. Dembo and A. Montanari. Ising models on locally tree-like graphs. *Ann. Appl. Probab.*, 20(2):565–592, 2010.

- [10] A. Dembo and O. Zeitouni. *Large deviations: techniques and applications*, volume 38 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, second edition, 1998.
- [11] S. Dommers, C. Giardinà, and R. van der Hofstad. Ising models on power-law random graphs. *J. Stat. Phys.*, 141(4):638–660, 2010.
- [12] F. J. Dyson. Existence of a phase-transition in a one-dimensional Ising ferromagnet. *Comm. Math. Phys.*, 12(2):91–107, 1969.
- [13] W. Evans, C. Kenyon, Y. Peres, and L. J. Schulman. Broadcasting on trees and the Ising model. *Ann. Appl. Probab.*, 10(2):410–433, 2000.
- [14] C. M. Fortuin and P. W. Kasteleyn. On the random-cluster model. I. Introduction and relation to other models. *Physica*, 57:536–564, 1972.
- [15] S. Franz and M. Leone. Replica bounds for optimization problems and diluted spin systems. *J. Statist. Phys.*, 111(3-4):535–564, 2003.
- [16] S. Franz, M. Leone, and F. Lucio Toninelli. Replica bounds for diluted non-Poissonian spin systems. *J. Phys. A*, 36(43):10967–10985, 2003. Special issue: Statistical physics of disordered systems: from real materials to optimization and codes.
- [17] H.-O. Georgii. *Gibbs measures and phase transitions*, volume 9 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, 1988.
- [18] C. Greenhill. The complexity of counting colourings and independent sets in sparse graphs and hypergraphs. *Comput. Complexity*, 9(1):52–72, 2000.
- [19] G. Grimmett. *The random-cluster model*, volume 333 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 2006.
- [20] G. Grimmett. Correlation inequalities of GKS type for the Potts model. ArXiv preprint, 2009.
- [21] F. Guerra and F. L. Toninelli. The thermodynamic limit in mean field spin glass models. *Comm. Math. Phys.*, 230(1):71–79, 2002.
- [22] J. M. Hammersley. Generalization of the fundamental theorem on sub-additive functions. *Proc. Cambridge Philos. Soc.*, 58:235–238, 1962.
- [23] Y. Higuchi. Remarks on the limiting Gibbs states on a $(d + 1)$ -tree. *Publ. Res. Inst. Math. Sci.*, 13(2):335–348, 1977/78.
- [24] R. M. Karp. Reducibility among combinatorial problems. In *Complexity of computer computations (Proc. Sympos., IBM Thomas J. Watson Res. Center, Yorktown Heights, N.Y., 1972)*, pages 85–103. Plenum, New York, 1972.
- [25] F. P. Kelly. Stochastic models of computer communication systems. *J. Roy. Statist. Soc. Ser. B*, 47(3):379–395, 415–428, 1985. With discussion.
- [26] E. L. Lawler, J. K. Lenstra, and A. H. G. Rinnooy Kan. Generating all maximal independent sets: NP-hardness and polynomial-time algorithms. *SIAM J. Comput.*, 9(3):558–565, 1980.
- [27] T. M. Liggett. *Interacting particle systems*. Classics in Mathematics. Springer-Verlag, Berlin, 2005. Reprint of the 1985 original.
- [28] R. Lyons. Random walks and percolation on trees. *Ann. Probab.*, 18(3):931–958, 1990.
- [29] M. Mézard and A. Montanari. *Information, physics, and computation*. Oxford Graduate Texts. Oxford University Press, Oxford, 2009.
- [30] A. Montanari, E. Mossel, and A. Sly. The weak limit of Ising models on locally tree-like graphs. *Probab. Theory Related Fields*, 2011 (in press).
- [31] E. Mossel, D. Weitz, and N. Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Probab. Theory Related Fields*, 143(3-4):401–439, 2009.
- [32] D. Panchenko and M. Talagrand. Bounds for diluted mean-fields spin glass models. *Probab. Theory Related Fields*, 130(3):319–336, 2004.
- [33] V. Riva and J. Cardy. Holomorphic parafermions in the Potts model and stochastic Loewner evolution. *J. Stat. Mech. Theory Exp.*, (12):P12001, 19 pp. (electronic), 2006.
- [34] D. Ruelle. *Statistical mechanics: Rigorous results*. W. A. Benjamin, Inc., New York-Amsterdam, 1969.
- [35] B. Simon. A remark on Dobrushin’s uniqueness theorem. *Comm. Math. Phys.*, 68(2):183–185, 1979.
- [36] A. Sly. Computational transition at the uniqueness threshold. *Foundations of Computer Science, Annual IEEE Symposium on*, 0:287–296, 2010.
- [37] A. D. Sokal. The multivariate Tutte polynomial (alias Potts model) for graphs and matroids. In *Surveys in combinatorics 2005*, volume 327 of *London Math. Soc. Lecture Note Ser.*, pages 173–226. Cambridge Univ. Press, Cambridge, 2005.

- [38] F. Spitzer. Markov random fields on an infinite tree. *Ann. Probab.*, 3(3):387–398, 1975.
- [39] D. Weitz. Counting independent sets up to the tree threshold. In *STOC'06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 140–149. ACM, New York, 2006.
- [40] F. Y. Wu. The Potts model. *Rev. Modern Phys.*, 54(1):235–268, 1982.
- [41] J. S. Yedidia, W. T. Freeman, and Y. Weiss. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Trans. Inform. Theory*, 51(7):2282–2312, 2005.
- [42] S. Zachary. Countable state space Markov random fields and Markov chains on trees. *Ann. Probab.*, 11(4):894–903, 1983.
- [43] J. Zinn-Justin. *Quantum field theory and critical phenomena*, volume 85 of *International Series of Monographs on Physics*. The Clarendon Press Oxford University Press, New York, second edition, 1993. Oxford Science Publications.
- [44] D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, STOC '06, pages 681–690, New York, NY, USA, 2006. ACM.

*DEPARTMENT OF MATHEMATICS, STANFORD UNIVERSITY
BUILDING 380, SLOAN HALL, STANFORD, CALIFORNIA 94305

†DEPARTMENT OF ELECTRICAL ENGINEERING, STANFORD UNIVERSITY
350 SERRA MALL, STANFORD, CALIFORNIA 94305

*††DEPARTMENT OF STATISTICS, STANFORD UNIVERSITY
SEQUOIA HALL, 390 SERRA MALL, STANFORD, CALIFORNIA 94305