

SPIN SYSTEMS ON BETHE LATTICES

AMIN COJA-OGLAN, WILL PERKINS

ABSTRACT. In an extremely influential paper Mézard and Parisi put forward an analytic but non-rigorous approach called the cavity method for studying spin systems on the Bethe lattice, i.e., the random d -regular graph [Eur. Phys. J. B **20** (2001) 217–233]. Their technique was based on certain hypotheses; most importantly, that the phase space decomposes into a number of Bethe states that are free from long-range correlations and whose marginals are given by a recurrence called Belief Propagation. In this paper we establish this decomposition rigorously for a very general family of spin systems. In addition, we show that the free energy can be computed from this decomposition. We also derive a variational formula for the free energy. The general results have interesting ramifications on several special cases.

MSC: 05C80

1. INTRODUCTION

1.1. Disordered systems and the Bethe lattice. In 2001 in a ground-breaking contribution Mézard and Parisi proposed an analytic but non-rigorous technique that they called the cavity method for the study of spin glasses on the ‘Bethe lattice’¹, known in combinatorics as the random d -regular graph [50]. Mézard and Parisi argued that the Bethe lattice constitutes an attractive halfway point between classical ‘mean-field’ models such as the Sherrington-Kirkpatrick model with complete interaction between all sites and spatial models such as the Edwards-Anderson model. Indeed, the Bethe lattice induces a non-trivial metric on the sites, each of which interacts with only a bounded number of others. But at the same time Mézard and Parisi showed that the model is amenable to analytic methods, even though matters are significantly more complicated than in the fully connected case. They went on to argue that the spin glass on the Bethe lattice exhibits many of the properties expected of real glassy systems, such as replica symmetry breaking and the proliferation of pure states.

From the original contribution [50] sprang a truly enormous body of work that has had a transformative impact on an astounding variety of subjects, ranging from physics to combinatorics to machine learning. Many of the applications may appear unexpected, even surprising. Almost all of them hinge on the cavity method. Prominent success stories include the development of ‘low-density parity check codes’, a rare example of a statistical physics idea leading directly to an eminently useful, and widely used, algorithm [59]. A further example is a new algorithm for the compressed sensing problem, a fundamental signal processing task [61]. Other important cavity method-based contributions pertain to classical problems in mathematics, such as phase transitions in random graphs and other random structures [43, 48, 49]. The cavity method has also been used to put forward predictions in machine learning, including the capacity of the Hopfield model or on restricted Boltzmann machines [47].

Due to these numerous ramifications, vindicating the cavity method rigorously has become an important research task at the junction of mathematical physics, combinatorics and computer science. There has been a lot of progress recently, e.g., [10, 34, 21, 54]; we shall review the literature in greater detail in Section 2.5. However, much of this work is concerned with special cases, mostly the ‘replica symmetric’ scenario where there is just a single pure state.

The aim of the present paper is to move past such assumptions and special cases. We confirm several of the key hypotheses of Mézard and Parisi, particularly the decomposition into pure states and the validity of the Belief Propagation recurrence, the mainstay of the cavity calculations. Further, we obtain a general variational formula for the free energy that is perfectly in line with the Mézard-Parisi ansatz. Additionally, we show that the free energy can be computed from the Belief Propagation representation of the pure states of the model. We obtain these results not merely for a specific model, but for a broad family of models on the Bethe lattice. The prime example is, of course, the diluted spin glass model. But in addition, since the proof techniques that we develop are generic,

¹Sometimes the d -regular infinite tree is referred to as the ‘Bethe lattice’. However, as Mézard and Parisi point out, the d -regular infinite tree does not provide a particularly useful framework for the study of spin interactions because almost all sites belong to the boundary of the tree. The random d -regular graph, which they and hence we call the Bethe lattice, provides a useful way out: while the local geometry around a given vertex is just a d -regular tree, at long distances this tree ‘wraps around’.

the results apply to models that are of eminent interest in other areas, particularly combinatorics, such as the Potts antiferromagnet or the hard-core model. Crucially, the results apply universally to all parameter values (such as degree, inverse temperature) of the respective models.

We should point out, however, that the present results fall short of fully corroborating the Mézard-Parisi ansatz. Most importantly, while we prove that general random factor graph models possess pure state decompositions represented by (approximate) Belief Propagation fixed points, an important prediction of the Mézard-Parisi ansatz pertains to the relative geometry of these fixed points. Roughly speaking, Mézard and Parisi predict three different scenarios, depending on the parameter values (such as temperature). In the *replica symmetric* case, there is just one single pure state (or possibly a small bounded number due to inherent symmetries). Moreover, in the *1-step replica symmetry breaking* scenario an unbounded number of pure states occur, each represented by a Belief Propagation fixed point. But these fixed points are predicted to exhibit a strong symmetry property that ensures, e.g., that the empirical distributions of the Belief Propagation messages in the different pure states are (nearly) identical. Finally, in the *full replica symmetry breaking* scenario we also expect an unbounded number of pure states and assorted Belief Propagation fixed points, which are arranged hierarchically in the fashion of an ultrametric tree. While in the present paper we prove the existence of a pure state decomposition and of associated Belief Propagation fixed points, our methods do not suffice to establish these more precise predictions as to the geometry and relative weights of the pure states. A full verification of these predictions remains an important open problem as it would, among other things, lead to a significantly simplified variational formula for the free energy. For a more detailed discussion of the Mézard and Parisi ansatz, replica symmetry breaking, and ultrametricity we refer to [48, 56, 57].

Technically the paper builds upon and continues two intertwined threads of prior work. First, we bring to bear a variant of the ‘regularity method’ from combinatorics that we developed recently [11, 16, 26] in order to establish the pure state decomposition and to vindicate the Belief Propagation equations. Second, we seize upon Panchenko’s work on asymptotic Gibbs measures and the interpolation method, particularly in order to derive the variational formula for the free energy [55, 56]. Both of these methods were previously applied with great success to random graphs of Erdős-Rényi type. This line of work crucially exploited the relative geometric flexibility of the Erdős-Rényi model, whose Poisson degree distribution facilitates coupling arguments. By contrast, the geometry of the Bethe lattice is rigid. While this entails that the specification of the model, the cavity equations and their solution are quite ‘clean’, the rigidity poses substantial technical challenges that the present paper resolves.

Before presenting the main results of the paper, which cover a broad family of problems that we call random factor graph models, in Section 2, we illustrate the results and the concepts around which they revolve with the spin glass model from the original contribution of Mézard and Parisi. We also work out an additional application to the hard-core model and the independence number of the random regular graph. Several further applications, including the Potts model and the MAX q -CUT problem, are worked out in Section 7.

1.2. The diluted spin glass. For integers $d \geq 3$, $n > 0$ such that dn is even, let $\mathbb{G} = \mathbb{G}(n, d)$ be the uniformly random d -regular graph on the vertex set $V_n = \{v_1, \dots, v_n\}$. With each edge $e \in E(\mathbb{G})$ comes a standard Gaussian J_e . The random variables $(J_e)_{e \in E(\mathbb{G})}$ are mutually independent. For a given inverse temperature $\beta > 0$, the diluted spin glass on \mathbb{G} is the probability distribution on $\{\pm 1\}^{V_n}$ defined by

$$\mu_{\mathbb{G}}(\sigma) = \frac{1}{Z(\mathbb{G})} \prod_{vw \in E(\mathbb{G})} \frac{1 + \tanh(\beta J_{vw})\sigma_v\sigma_w}{2}, \quad (1.1)$$

where the partition function $Z(\mathbb{G})$ ensures normalization.² Without the couplings J_e , this would just be the ferromagnetic Ising model on \mathbb{G} . But since the J_e are independent Gaussians, some will be positive and others negative. In effect, some edges induce ferromagnetic and others antiferromagnetic interactions, causing frustration. Thus, $\mu_{\mathbb{G}}$ is a spin glass model, the well-known diluted spin glass on the Bethe lattice.

There are two fundamental problems associated with this and numerous similar models: first, to characterize the structure of the Boltzmann distribution $\mu_{\mathbb{G}}$. Does it exhibit long-range correlations? Does it decompose into one or several ‘pure states’, and if so, how can we characterize them? Second, to calculate the quantity $\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbb{G})]$, which we call the *free energy density*. Its fundamental importance is due to the fact that other important observables derive from it. Moreover, the singularities of the function $\beta \mapsto \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbb{G})]$ constitute the phase transitions of the model.

²The expression (1.1) is equivalent to the possibly more familiar formula $\mu_{\mathbb{G}}(\sigma) \propto \exp(\beta \sum_{vw} \sigma_v \sigma_w)$.

Bethe states and the Boltzmann distribution. With respect to the first problem, Mézard and Parisi hypothesized that the Boltzmann distribution always decomposes into one or a moderate (albeit not necessarily bounded) number of pure states. Further, they hypothesized that these pure states are characterized by fixed points of a recurrence called Belief Propagation. Our first theorem confirms this hypothesis.

To be precise, writing ∂v for the set of neighbors of a vertex v , let $\mathcal{M}(\mathbb{G})$ be the set of all families $(v_{u \rightarrow v})_{v \in V_n, u \in \partial v}$ such that $v_{u \rightarrow v} \in [0, 1]$. We call $v_{u \rightarrow v}$ the *message* from u to v . The messages need not be symmetric, i.e., possibly $v_{u \rightarrow v} \neq v_{v \rightarrow u}$. Furthermore, *Belief Propagation* is the operator $\text{BP} : \mathcal{M}(\mathbb{G}) \rightarrow \mathcal{M}(\mathbb{G})$, $v \mapsto \hat{v}$, where

$$\hat{v}_{v \rightarrow u} = \frac{\prod_{w \in \partial v \setminus u} 1 + 2 \tanh(\beta J_{vw})(v_{w \rightarrow v} - 1/2)}{\sum_{\sigma \in \{\pm 1\}} \prod_{w \in \partial v \setminus u} 1 + 2\sigma \tanh(\beta J_{vw})(v_{w \rightarrow v} - 1/2)}. \quad (1.2)$$

The motivation behind this operator, and the origin of the name ‘cavity method’, is this. Suppose we fix a vertex v in a d -regular graph along with a neighbor u . Now suppose we remove the vertex u , thereby creating a ‘cavity’. Then the ‘ideal’ message $\mu_{\mathbb{G}, u \rightarrow v}$ that we would like to compute is just the marginal probability $\mu_{\mathbb{G} - v, u}(1)$ that u takes spin 1 in the subgraph obtained by removing v . If the Boltzmann distribution $\mu_{\mathbb{G}}$ is free from long-range correlations, then these ideal messages should plausibly be a fixed point of the BP operator. Indeed, if we remove v , then very likely its former neighbors will be mutually far apart in the resulting graph. In effect, the joint distribution of their spins should factorize. If so, then a straightforward calculation verifies that the ideal messages are a fixed point of BP. In fact this reasoning goes back to Bethe’s classical work [16].

However, generally spin glass models do exhibit long-range correlations, a phenomenon called *replica symmetry breaking* (see, e.g., [18, 27] for proofs that replica symmetry breaking occurs in certain models). Yet the fundamental hypothesis of Mézard and Parisi holds that the phase space $\{\pm 1\}^{V_n}$ always decomposes into *Bethe states* S_1, \dots, S_ℓ in such a way that the conditional distributions $\mu_{\mathbb{G}}[\cdot | S_h]$ are free from long-range correlations. Formally, this means that if we pick a pair of vertices (v_i, v_j) uniformly at random, then typically the conditional joint distribution $\mu_{\mathbb{G}, v_i, v_j}[\cdot | S_h]$ of the spins of v_i and v_j is close to the product distribution $\mu_{\mathbb{G}, v_i}(\cdot | S_h) \otimes \mu_{\mathbb{G}, v_j}(\cdot | S_h)$, i.e.,

$$\frac{1}{n^2} \sum_{1 \leq i < j \leq n} \left\| \mu_{\mathbb{G}, v_i, v_j}(\cdot | S_h) - \mu_{\mathbb{G}, v_i}(\cdot | S_h) \otimes \mu_{\mathbb{G}, v_j}(\cdot | S_h) \right\|_{\text{TV}} = o(1). \quad (1.3)$$

In effect, within each Bethe state the ‘ideal’ messages are predicted to be an approximate fixed point of the BP operator. To be precise, for adjacent vertices u, v we write $\mu_{\mathbb{G}, v \rightarrow u}[S_h] = \mu_{\mathbb{G} - u, v}(1 | S_h)$ for the conditional probability given S_h that v takes spin 1 in the subgraph of \mathbb{G} with u removed. Then we expect that

$$\frac{1}{n} \sum_{i=1}^n \sum_{u \in \partial v_i} \left\| \mu_{\mathbb{G}, v_i \rightarrow u}[S_h] - \hat{\mu}_{\mathbb{G}, v_i \rightarrow u}[S_h] \right\|_{\text{TV}} = o(1) \quad \text{where} \quad (\hat{\mu}_{\mathbb{G}, v \rightarrow u}[S_h])_{v \in V_n, u \in \partial v} = \text{BP}(\mu_{\mathbb{G}, v \rightarrow u}[S_h])_{v \in V_n, u \in \partial v}. \quad (1.4)$$

Further, the cavity method predicts that the Boltzmann marginals can be obtained from the messages by a formula quite similar to (1.2):

$$\frac{1}{n} \sum_{i=1}^n \left| \mu_{\mathbb{G}, v_i}(1 | S_h) - \frac{\prod_{w \in \partial v_i} 1 + 2 \tanh(\beta J)(\mu_{\mathbb{G}, w}[S_h] - 1/2)}{\sum_{\sigma \in \{\pm 1\}} \prod_{w \in \partial v_i} 1 + 2\sigma \tanh(\beta J)(\mu_{\mathbb{G}, w}[S_h] - 1/2)} \right| = o(1). \quad (1.5)$$

The following theorem establishes these conjectures rigorously. We say that \mathbb{G} enjoys a property *with high probability* (‘w.h.p.’) if the probability that the property holds tends to one as $n \rightarrow \infty$.

Theorem 1.1. *For any $d \geq 3$, $\beta > 0$ the following is true. Let $L = L(n) \rightarrow \infty$ be any integer sequence that tends to infinity. Then there exists a decomposition $S_0 = S_0(\mathbb{G}), S_1 = S_1(\mathbb{G}), \dots, S_\ell = S_\ell(\mathbb{G})$, $\ell = \ell(\mathbb{G}) \leq L$, of the phase space $\{\pm 1\}^n$ into non-empty sets such that $\mu_{\mathbb{G}}(S_0) = o(1)$ and such that with high probability (1.3)–(1.5) are satisfied for $h = 1, \dots, \ell$.*

Crucially, and in contrast to much prior work in this area, Theorem 1.1 applies indiscriminately to all d, β . While it is expected that in the ‘high-temperature’ regime (small β) there is just a single pure state, it is widely conjectured that for large d and β the number of pure states is unbounded. Thus, we do not expect that it will be possible to replace the unbounded L in Theorem 1.1 by a constant. Yet Theorem 1.1 shows that the number of states can be upper bounded by an *arbitrarily* slowly growing function $L(n)$.

The free energy. The Bethe states and their associated messages contain all the information needed to compute the free energy. To be precise, once more following the ideas of Mézard and Parisi, we can set up a recurrence for computing the difference $\mathbb{E}[\log Z(\mathbb{G}(n+1, d))] - \mathbb{E}[\log Z(\mathbb{G}(n, d))]$, which in turn enables us to write a formula for $\frac{1}{n}\mathbb{E}[\log Z(\mathbb{G}(n, d))]$ by telescoping. To set up such a recurrence it is necessary to crack the rigid geometry of the random regular graph open a little bit. To this end, we resort to the idea of creating a few ‘cavities’. Specifically, we delete a few random vertices and edges from $\mathbb{G}(n, d)$. Formally, let $\omega > 0$ and let \mathbf{X}, \mathbf{Y} be two independent Poisson variables with mean ω . Moreover, let $\mathbf{u}_1, \dots, \mathbf{u}_X$ and $\mathbf{v}_1 \mathbf{w}_1, \dots, \mathbf{v}_Y \mathbf{w}_Y$ be sequences of uniformly random vertices and edges of \mathbb{G} , chosen independently. With S_1, \dots, S_ℓ the decomposition from Theorem 1.1, we introduce weights

$$\begin{aligned} \mathbf{z}_{\mathbb{G}, h} = \mu_{\mathbb{G}}(S_h) \cdot \prod_{i=1}^X \left(\sum_{\sigma \in \{\pm 1\}} \prod_{v \in \partial \mathbf{u}_i} 1 + 2\sigma \tanh(\beta J_{v \mathbf{u}_i})(\mu_{\mathbb{G}, v \rightarrow \mathbf{u}_i}[\sigma | S_h] - 1/2) \right)^{-1} \\ \cdot \prod_{i=1}^Y (1 + 4 \tanh(\beta J_{\mathbf{v}_i \mathbf{w}_i})(\mu_{\mathbb{G}, \mathbf{v}_i \rightarrow \mathbf{w}_i}[1 | S_h] - 1/2)(\mu_{\mathbb{G}, \mathbf{w}_i \rightarrow \mathbf{v}_i}[1 | S_h] - 1/2))^{-1} \end{aligned}$$

and $\mathbf{z}_{\mathbb{G}} = \sum_{h=1}^\ell \mathbf{z}_{\mathbb{G}, h}$. Further, let $\mathcal{C}(\mathbb{G})$ be the set of all vertices of degree less than d in the graph $\mathbb{G}_{n, \omega}$ obtained from \mathbb{G} by removing $\mathbf{u}_1, \dots, \mathbf{u}_X$ and $\mathbf{v}_1 \mathbf{w}_1, \dots, \mathbf{v}_Y \mathbf{w}_Y$. Then with high probability each $c \in \mathcal{C}(\mathbb{G})$ has degree precisely $d-1$, and we write c' for the erstwhile d ’th neighbor of c . Further, with $\mathbf{c}_1, \mathbf{c}_2, \dots$ a sequence of uniformly and independently chosen elements of $\mathcal{C}(\mathbb{G})$ and $(\mathbf{J}_i)_{i \geq 1}$ a sequence of independent standard Gaussians, we let

$$\begin{aligned} \mathcal{B}(\mathbb{G}) = \mathbb{E} \left[\log \sum_{h=1}^\ell \frac{\mathbf{z}_{\mathbb{G}, h}}{\mathbf{z}_{\mathbb{G}}} \sum_{\sigma \in \{\pm 1\}} \prod_{i=1}^d 1 + 2\sigma \tanh(\beta \mathbf{J}_i)(\mu_{\mathbb{G}, \mathbf{c}_i \rightarrow \mathbf{c}'_i}[S_h] - 1/2) \middle| \mathbb{G} \right] \\ - \frac{d}{2} \mathbb{E} \left[\log 1 + 4 \tanh(\beta \mathbf{J}_1) \sum_{h=1}^\ell \frac{\mathbf{z}_{\mathbb{G}, h}}{\mathbf{z}_{\mathbb{G}}} (\mu_{\mathbb{G}, \mathbf{c}_1 \rightarrow \mathbf{c}'_1}[1 | S_h] - 1/2)(\mu_{\mathbb{G}, \mathbf{c}_2 \rightarrow \mathbf{c}'_2}[1 | S_h] - 1/2) \middle| \mathbb{G} \right] - \frac{d}{2} \log 2. \end{aligned}$$

The expression $\mathcal{B}(\mathbb{G})$ mirrors our recurrence for the difference $\mathbb{E}[\log Z(\mathbb{G}(n+1, d))] - \mathbb{E}[\log Z(\mathbb{G}(n, d))]$. Having created a moderate number of cavities, we insert a new $(n+1)$ st vertex, connected to d randomly chosen ‘cavities’. The first summand above represents the ensuing change in the free energy. But this operation adds d more edges, whereas a random regular graph with $n+1$ vertices only has $d/2$ more edges than one with n vertices. Therefore, a correction term is needed. Hence the second summand.

Crucially, the functional $\mathcal{B}(\mathbb{G})$ depends only on the pure state decomposition from Theorem 1.1 and the associated messages. The following theorem shows that this information suffices to compute the free energy.

Theorem 1.2. *For all $d \geq 3, \beta > 0$ we have*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbb{G})] = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathbb{E}[\mathcal{B}(\mathbb{G})].$$

Entirely in line with the ideas developed in [50], Theorem 1.2 establishes a direct conceptual link between Belief Propagation and the pure state decomposition from Theorem 1.1 and the free energy for *all* d, β . Of course, in order to evaluate $\mathcal{B}(\mathbb{G})$ it is necessary to actually determine the pure state decomposition along with the corresponding Belief Propagation messages. The shape of this decomposition, and the practical difficulty of computing it, will depend significantly on the parameters d, β . Alternatively, as we see next, it is possible to derive a variational formula for the free energy.

A variational formula. The variational formula comes in terms of an optimization problem on a space that resembles the graphon space from the theory of graph limits [46]. To be precise, let $v : [0, 1]^2 \rightarrow [0, 1]$, $(s, x) \mapsto v_{s, x}$ and $v' : [0, 1]^2 \rightarrow [0, 1]$, $(s, x) \mapsto v'_{s, x}$ be measurable maps. We define the *cut distance* between v, v' by

$$\mathcal{D}_\square(v, v') = \inf_{\varphi, \varphi'} \sup_{S, X \subset [0, 1]} \left| \int_S \int_X v_{s, x}(\omega) - v'_{\varphi(s), \varphi'(x)}(\omega) dx ds \right|,$$

where $\varphi, \varphi' : [0, 1] \rightarrow [0, 1]$ are measurable maps that preserve the Lebesgue measure and $S, X \subset [0, 1]$ are measurable. Obtain the space \mathfrak{K} by identifying any v, v' with $\mathcal{D}_\square(v, v') = 0$. Then \mathfrak{K} endowed with the cut distance is a compact metric space. In addition, write \mathfrak{D} for the space of probability measures on \mathfrak{K} .

The formula for the free energy comes as a variational problem on a subspace \mathfrak{D}^* of \mathfrak{D} . Let $N, M \geq 0$ be integers. For $\mu \in \mathfrak{K}$ we define a randomly perturbed $\mu^{*(N, M)} \in \mathfrak{K}$ as follows. Let $(\mathbf{x}_{i, j})_{i, j \geq 1}$ be a family of uniform random

variables on $[0, 1]$ and let $(J_{i,j})_{i,j \geq 1}$ be a family of standard Gaussians, all mutually independent. Then for $s \in [0, 1]$ we define

$$z_s = \prod_{i=1}^N \left(\sum_{\sigma \in \{\pm 1\}} \prod_{j=1}^d 1 + 2 \tanh(\beta J_{i,j})(\mu_{s,x_{i,j}} - 1/2) \right) \prod_{i=1}^M (1 + 4 \tanh(\beta J_{i+N,1})(\mu_{s,x_{i+N,1}} - 1/2)(\mu_{s,x_{i+N,2}} - 1/2)).$$

Further, let

$$t = t(s) = \inf \left\{ \theta \in [0, 1] : \int_0^\theta z_u du \geq s \int_0^1 z_u du \right\}, \quad \text{and} \quad \mu_{s,x}^{*(N,M)} = \mu_{t,x} \in \mathcal{K}.$$

Now, suppose that $\pi \in \mathcal{D}$ is a distribution, and write $\mu^\pi \in \mathcal{K}$ for a sample from π . Then we let \mathcal{D}^* be the set of all $\pi \in \mathcal{D}$ such that the perturbed $\mu^{\pi*(N,M)}$ has distribution π again for all $N, M \geq 0$.

The definition of \mathcal{D}^* , which is an adaptation of the one stated by Panchenko [55] in the case of models of Erdős-Rényi type, mirrors a natural combinatorial invariance property of the graph $\mathbb{G}_{n,\omega}$ with the random cavities. Indeed, because the numbers X, Y of deleted edges and vertices are Poisson with a large mean ω , for any fixed N, M the random graph $\mathbb{G}_{n,\omega}$ with X deleted vertices and Y deleted edges is close in total variation to the one with merely $X - N$ deleted vertices and $Y - M$ deleted edges. Furthermore, because adding or removing a small number of edges only affects the Boltzmann weights by a bounded factor, we should expect that the Bethe states of these two factor graphs remain the same. But, of course, the relative probability masses of the Bethe states will be different. Accordingly, the weights z_s mirror the changes in the weights of the Bethe states upon re-insertion of N vertices, each with d incident edges, and another M edges into $\mathbb{G}_{n,\omega}$. Once we take ω and n to infinity, the closeness of the two random factor graphs in total variation translates into the statement that the distribution of the messages emitted by the cavities of $\mathbb{G}_{n,\omega}$ belongs to \mathcal{D}^* .

Finally, define a functional $\mathcal{B} : \mathcal{K} \rightarrow \mathbb{R}$ by letting

$$\begin{aligned} \mathcal{B}(\mu) = & \mathbb{E} \left[\log \left(\sum_{\sigma \in \{\pm 1\}} \int_0^1 \prod_{i=1}^d 1 + 2\sigma \tanh(\beta J_{i,j})(\mu_{s,x_{i,j}} - 1/2) ds \right) \right. \\ & \left. - \frac{d}{2} \log \left(1 + 4 \tanh(\beta J_{1,1}) \int_0^1 (\mu_{s,x_{1,1}} - 1/2)(\mu_{s,x_{1,2}} - 1/2) ds \right) \right] - \frac{d}{2} \log 2. \end{aligned}$$

We are ready to state the variational formula for the free energy.

Theorem 1.3. *For all $d \geq 3$ and $\beta > 0$ we have $\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z_\beta(\mathbb{G})] = \min_{\pi \in \mathcal{D}^*} \mathbb{E}[\mathcal{B}(\mu^\pi)]$.*

Theorem 1.2 provides the combinatorial interpretation of the optimal π for Theorem 1.3: it is the kernel representing the messages $(\mu_{\mathbb{G}, c \rightarrow c'}[\cdot | S_h])_{c \in \mathcal{C}(\mathbb{G}), h=1, \dots, \ell}$ sent out by the cavities on the individual Bethe states.

1.3. The hard-core model. As a second application we discuss the hard-core model on the random regular graph $\mathbb{G} = \mathbb{G}(n, d)$. This is a probability distribution on the collection of independent sets of \mathbb{G} parametrized by $\lambda > 0$, the *fugacity*. Formally, encoding subsets of the vertex set by their indicator vectors, we define

$$\mu_{\mathbb{G}}(\sigma) = \frac{\lambda^{\sum_{i=1}^n \sigma_i}}{Z(\mathbb{G})} \prod_{1 \leq i < j \leq n} 1 - \mathbf{1}\{\nu_i \in \partial \nu_j\} \sigma_i \sigma_j \quad (\sigma \in \{0, 1\}^n),$$

with $Z(\mathbb{G})$ the partition function that turns $\mu_{\mathbb{G}}$ into a probability measure. Thus, $\mu_{\mathbb{G}}(\sigma) = 0$ unless the 1-entries of σ form an independent set in \mathbb{G} , in which case the weight of σ is proportional to λ taken to the power of the size of the independent set.

The hard-core model, of great prominence in statistical physics, is of eminent importance in combinatorics as well because it is closely related to the problem of finding the size of the largest independent set of the random regular graph. For d large, this problem was solved by Ding, Sly, and Sun [35] using an intricate version of the second-moment method guided by insights from the 1-step replica symmetry breaking (1RSB) version of the cavity method. But according to the physics predictions [14], the 1RSB method runs into an inherent obstacle for small d as the model exhibits a continuous phase transition to a more complicated ‘full replica symmetry breaking’ (full RSB) phase. In Corollary 1.5 below we will derive a formula for the largest independent set size that holds for all d and that accommodates the full RSB scenario.

But let us first deal with the free energy of the hard-core model, in and of itself a well-known problem. To derive a variational formula for the free energy, obtain \mathcal{K}_λ from the space of all measurable functions $[0, 1]^2 \rightarrow$

$[0, \lambda/(1+\lambda)]$ by identifying any v, v' with $\mathcal{D}_\square(v, v') = 0$. Then \mathfrak{K}_λ is a compact. In addition, we let \mathfrak{D}_λ be the space of probability measures on \mathfrak{K}_λ . Similarly to the spin glass problem, the formula for the free energy comes as a variational problem on a subspace \mathfrak{D}_λ^* of \mathfrak{D}_λ . This subspace is defined as follows. Let $(\mathbf{x}_{i,j})_{i,j \geq 1}$ be a family of independent random variables, uniformly distributed on $[0, 1]$, and let $N, M \geq 0$ be integers. Then for $\mu \in \mathfrak{K}_\lambda$ we define a random $\mu^{*(N,M)} \in \mathfrak{K}_\lambda$ as follows. For $s \in [0, 1]$ let

$$\mathbf{z}_s = \prod_{i=1}^N \left(1 + \lambda \prod_{j=1}^d 1 - \mu_{s,\mathbf{x}_{i,j}} \right) \prod_{i=1}^M \left(1 - \mu_{s,\mathbf{x}_{i+N,1}} \mu_{s,\mathbf{x}_{i+N,2}} \right) \quad \text{and} \quad \mathbf{t} = \mathbf{t}(s) = \inf \left\{ \theta \in [0, 1] : \int_0^\theta \mathbf{z}_{\mu,u}^{N,M} \, du \geq s \int_0^1 \mathbf{z}_{\mu,u}^{N,M} \, du \right\}$$

and set

$$\mu_{s,x}^{*(N,M)} = \mu_{\mathbf{t},x} \in \mathfrak{K}_\lambda.$$

Further, suppose that $\pi \in \mathfrak{D}_\lambda$ is a distribution, and write $\boldsymbol{\mu}^\pi \in \mathfrak{K}_\lambda$ for an element chosen from π . Then we let \mathfrak{D}_λ^* be the set of all $\pi \in \mathfrak{D}_\lambda$ such that $\boldsymbol{\mu}^\pi$ and $\boldsymbol{\mu}^{\pi*(N,M)}$ are identically distributed for all $N, M \geq 0$. Finally, let $\mathcal{B} : \mathfrak{K}_\lambda \rightarrow \mathbb{R}$ be the function defined by

$$\mathcal{B}(\mu) = \mathbb{E} \left[\log \left(1 + \lambda \int_0^1 \prod_{j=1}^d 1 - \mu_{s,\mathbf{x}_{1,j}} \, ds \right) - \frac{d}{2} \log \left(1 - \int_0^1 \mu_{s,\mathbf{x}_{1,1}} \mu_{s,\mathbf{x}_{1,2}} \, ds \right) \right].$$

The variational formula for the free energy reads as follows.

Theorem 1.4. *For all $d \geq 3$ and $\lambda > 0$ we have*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbb{G})] = \Phi_{d,\lambda}, \quad \text{with} \quad \Phi_{d,\lambda} = \min_{\pi \in \mathfrak{D}_\lambda^*} \mathbb{E}[\mathcal{B}(\boldsymbol{\mu}^\pi)].$$

In the limit $\lambda \rightarrow \infty$ the distribution $\mu_{\mathbb{G},\lambda}$ concentrates on the maximum independent sets of the random graph. As an application of Theorem 1.4 we therefore obtain the following result on the size of the largest independent set, i.e., the independence number $\alpha(\mathbb{G})$ of the random graph.

Corollary 1.5. *For all $d \geq 3$ we have $\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\alpha(\mathbb{G})] = \lim_{\lambda \rightarrow \infty} \lambda \cdot (\Phi_{d,\lambda+1} - \Phi_{d,\lambda})$.*

The formula in Corollary 1.5 may not be easy to evaluate; in particular, it may be difficult to obtain a numerical estimate for a given value of d . Nonetheless, since the proofs show that the optimal π in Theorem 1.4 is closely related to the Belief Propagation fixed points on \mathbb{G} , it should be possible to extract combinatorial information about the independent set problem on random graphs. In any case, Theorem 1.4 and Corollary 1.5 put a lid on the complexity of the problem.

1.4. Organization. In Section 2 we present the main results of the paper, which cover a broad family of random factor graph models. At the end of Section 2 we are in a position to discuss related work in detail. Sections 3–6 deliver the proofs of these general results. Finally, in Section 7 we show how Theorems 1.1–1.4 and Corollary 1.5 follow from the general results in Section 2. In addition, we work through several more applications that have each received considerable attention in their own right, such as the Potts antiferromagnet.

2. RANDOM FACTOR GRAPHS

In this section we present the main results of the paper, which cover a broad class of models called *random factor graphs*. The class encompasses many well-studied examples of problems on random regular graphs or hypergraphs, including the spin glass model from the previous section. Some other cases, such as the hard-core model or extremal cuts, can be dealt with by taking limits; we will come to that in Section 7.

2.1. Definitions. To define random factor graph models, we consider a finite set $\Omega \neq \emptyset$ whose elements we call *spins*. Moreover, for an integer $k \geq 2$ we let (Ψ, P) be a probability space of *weight functions* $\psi : \Omega^k \rightarrow (0, 1)$. We always denote by $\boldsymbol{\psi}$ an element of Ψ chosen from the distribution P . The space Ψ may be finite or infinite. In the latter case we assume that

$$\mathbb{E}[\exp(1/\min_{\sigma \in \Omega^k} \psi(\sigma))] < \infty. \tag{2.1}$$

Furthermore, we always assume that the distribution P is invariant under permutations of the coordinates. That is, for any $\psi \in \Psi$ and for any permutation κ of $\{1, \dots, k\}$ the function $\psi^\kappa : \sigma \mapsto \psi(\sigma_{\kappa_1}, \dots, \sigma_{\kappa_k})$ belongs to Ψ as well and ψ^κ has the same distribution as ψ . Additionally, let p be a probability distribution on Ω with $p(\omega) > 0$ for all

$\omega \in \Omega$. Further, let $d \geq 3, n > 0$ be integers and set $m = \lfloor dn/k \rfloor$. Let $V_n = \{v_1, \dots, v_n\}$ be a set of *variable nodes* and let $F_m = \{a_1, \dots, a_m\}$ be a set of *constraint nodes*.

Definition 2.1. Suppose that k divides dn . The random factor graph $\mathbf{G} = \mathbf{G}(n, d, p, P)$ consists of

- a weight function $\psi_{a_i} \in \Psi$ drawn from the distribution P independently for each $i = 1, \dots, m$ and
- an independent uniformly random bijection $\partial_{\mathbf{G}} : F_m \times \{1, \dots, k\} \rightarrow V_n \times \{1, \dots, d\}$.

The definition resembles the pairing model of random regular graphs [41]. Accordingly, we use standard graph-theoretic terminology. For instance, we call $x_i \in V_n$ and $a_j \in F_m$ adjacent if there exist $s \in [d]$ and $k \in [k]$ such that $\partial_{\mathbf{G}}(a_j, t) = (x_i, s)$. We also use the symbol $\partial_{\mathbf{G}}(a_j, t)$ for the variable node x_i such that $\partial_{\mathbf{G}}(a_j, t) = (x_i, s)$. Further, we write $\partial_{\mathbf{G}} x_i$ for the set of all $a_j \in F_m$ that x_i is adjacent to, and similarly for a_j . We omit the index and just write $\partial x_i, \partial a_j$ etc. where the reference to the random graph is apparent. In particular, \mathbf{G} induces a bipartite graph on the variable and constraint nodes, and thereby the shortest path metric on $V_n \cup F_m$. Hence, by extension of the above notation, we write $\partial_{\mathbf{G}}^\ell u$ for the set of all nodes at distance precisely ℓ from u and $\nabla_{\mathbf{G}}^\ell u$ for set of all variable nodes at distance at most ℓ from u .

We let \mathcal{S} be the event that \mathbf{G} is simple, i.e., that there do not occur multiple edges between any variable and constraint nodes. Moreover, we denote by \mathbb{G} the conditional distribution of \mathbf{G} given \mathcal{S} . Let us make a note of the following well known fact.

Fact 2.2 ([41]). We have $\mathbb{P}[\mathbf{G} \in \mathcal{S}] \sim \exp[-(d-1)(k-1)/2 - \mathbf{1}\{k=2\}(d-1)^2/4]$.

The random factor graph induces a probability distribution on Ω^{V_n} . To define it, we introduce the shorthand $\psi_{a_i}(\sigma) = \psi_{a_i}(\sigma(\partial(a_i, 1)), \dots, \sigma(\partial(a_i, k)))$ for $i \in [m]$ and $\sigma \in \Omega^{V_n}$. Thus, $\psi_{a_i}(\sigma)$ is the weight that constraint node a_i gives to σ . Further, we introduce the *total weight*

$$\psi_{\mathbf{G}}(\sigma) = \prod_{i=1}^m \psi_{a_i}(\sigma) \quad (\sigma \in \Omega^{V_n}).$$

by multiplying up all the weight functions of the constraint nodes. The total weights $\psi_{\mathbf{G}}(\sigma)$ give rise to the *partition function* and the *Boltzmann distribution*:

$$Z(\mathbf{G}) = \sum_{\tau \in \Omega^{V_n}} \psi_{\mathbf{G}}(\tau) \prod_{i=1}^n p(\sigma(x_i)), \quad \mu_{\mathbf{G}}(\sigma) = \frac{\psi_{\mathbf{G}}(\sigma)}{Z(\mathbf{G})} \prod_{i=1}^n p(\sigma(x_i)) \quad (\sigma \in \Omega^{V_n}) \quad (2.2)$$

Since all the weight functions $\psi \in \Psi$ are strictly positive, the Boltzmann distribution is a well-defined probability measure on the phase space Ω^{V_n} .

We set out to investigate the structure of the Boltzmann distribution $\mu_{\mathbf{G}}(\cdot)$ and to compute the partition function $Z(\mathbf{G})$ or, more specifically, its logarithm, which we call the *free energy*. In Section 2.2 we will prove the main result of the paper, which provides that the Boltzmann distribution decomposes into a convex combination of relatively simple distributions called Bethe states. But before we come to that, let us look at an example.

Example 2.3 (the k -spin model). Let $\Omega = \{\pm 1\}$, let $k \geq 2$ be an integer and let $\beta > 0$ be a real parameter. The k -spin model is a generalization of the spin glass model from the previous section, which corresponds to the special case $k = 2$. The weight functions of the k -spin model read

$$\psi_{\beta, J}(\sigma_1, \dots, \sigma_k) = \frac{1}{2} \left(1 + \tanh(\beta J) \prod_{i=1}^k \sigma_i \right) \quad (J \in \mathbb{R}).$$

Thus, $\Psi = \{\psi_{\beta, J} : J \in \mathbb{R}\}$, and the distribution P on Ψ is defined by choosing J from the standard Gaussian distribution. This distribution clearly satisfies (2.1). Geometrically, this model lives on a generalized Bethe lattice where all variable nodes, representing the sites, have degree d , while all constraint nodes, representing the interactions, have degree k .

The hard-core model from Section 1.3 cannot be expressed as a factor graph model directly because of the requirement that all weight functions be strictly positive. But it is possible to arrive at the hard-core model by taking suitable limits; see Section 7 for details.

2.2. Bethe states. The Belief Propagation message-passing scheme provides the mainstay of the physicists' non-rigorous cavity method. Our first main result vindicates its use by showing that the Boltzmann distribution of any random factor graph model can be described in terms of Belief Propagation fixed points.

To introduce Belief Propagation let $\mathcal{M}(\mathbf{G})$ be the *message space*, consisting of all families

$$\nu = (\nu_{v \rightarrow a}, \nu_{a \rightarrow v})_{v \in V_n, a \in F_n: v \in \partial_{\mathbf{G}} a}$$

of probability measures $\nu_{v \rightarrow a}, \nu_{a \rightarrow v}$ on Ω . For adjacent a, v we interpret $\nu_{v \rightarrow a}$ as a 'message' from v to a , and $\nu_{a \rightarrow v}$ as a message in the reverse direction. We equip $\mathcal{M}(\mathbf{G})$ with the metric

$$\mathcal{D}_1(\nu, \nu') = \frac{1}{n} \sum_{v, a: v \in \partial_{\mathbf{G}} a} \|\nu_{v \rightarrow a} - \nu'_{v \rightarrow a}\|_{\text{TV}} + \|\nu_{a \rightarrow v} - \nu'_{a \rightarrow v}\|_{\text{TV}}.$$

Belief Propagation is the operator $\text{BP}: \mathcal{M}(\mathbf{G}) \rightarrow \mathcal{M}(\mathbf{G})$ that maps ν to $\hat{\nu}$ defined by

$$\hat{\nu}_{v \rightarrow a}(\sigma) = \frac{p(\sigma) \prod_{b \in \partial v \setminus a} \nu_{b \rightarrow v}(\sigma)}{\sum_{\tau \in \Omega} p(\tau) \prod_{b \in \partial v \setminus a} \nu_{b \rightarrow v}(\tau)}, \quad \hat{\nu}_{a \rightarrow v}(\sigma) = \frac{\sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\tau_v = \sigma\} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \nu_{w \rightarrow a}(\tau_w)}{\sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \nu_{w \rightarrow a}(\tau_w)}. \quad (2.3)$$

Further, a point $\nu \in \mathcal{M}(\mathbf{G})$ is an ε -Belief Propagation fixed point if $\mathcal{D}_1(\nu, \text{BP}(\nu)) < \varepsilon$.

For a thorough discussion and motivation of Belief Propagation we refer to [48]. The punch line is that on acyclic factor graphs a Belief Propagation fixed point computation provably yields the marginals of the Boltzmann distribution as well as the free energy. Since the random graph \mathbf{G} contains only very few short cycles, one may therefore expect that Belief Propagation renders meaningful information on random factor graphs as well, provided that the Boltzmann distribution is free of long-range correlations.

Alas, in general long-range correlations do occur. Nevertheless, we will prove that the Boltzmann distribution still decomposes into a convex combination of relatively few 'Bethe states', characterized by Belief Propagation fixed points. To be precise, suppose that $\emptyset \neq S \subset \Omega^{V_n}$ is an event. Let v be a variable node and let $a \in \partial_{\mathbf{G}} v$. Then we define $\mu_{\mathbf{G}, v \rightarrow a}(\cdot | S)$ as the conditional marginal of v given S under the Boltzmann distribution of the factor graph $\mathbf{G} - a$ obtained from \mathbf{G} by removing the constraint node a . In formulas, with $\langle \cdot, \mu_{\mathbf{G}}(\cdot | S) \rangle$ denoting the expectation with respect to σ drawn from $\mu_{\mathbf{G}}(\cdot | S)$, we have

$$\mu_{\mathbf{G}, v \rightarrow a}(\sigma | S) = \frac{\langle \mathbf{1}\{\sigma_v = \sigma\} / \psi_a(\sigma), \mu_{\mathbf{G}}(\cdot | S) \rangle}{\langle 1 / \psi_a(\sigma), \mu_{\mathbf{G}}(\cdot | S) \rangle} \quad (\sigma \in \Omega). \quad (2.4)$$

Similarly, we let $\mu_{\mathbf{G}, a \rightarrow v}(\cdot | S)$ be the conditional marginal of v under the Boltzmann distribution of the factor graph obtained from \mathbf{G} by removing all constraint nodes $b \in \partial_{\mathbf{G}} v \setminus a$ and disregarding the prior of v :

$$\mu_{\mathbf{G}, a \rightarrow v}(\sigma | S) = \frac{\langle \mathbf{1}\{\sigma_v = \sigma\} / (p(\sigma) \prod_{b \in \partial v \setminus a} \psi_b(\sigma)), \mu_{\mathbf{G}}(\cdot | S) \rangle}{\langle 1 / (p(\sigma_v) \prod_{b \in \partial v \setminus a} \psi_b(\sigma)), \mu_{\mathbf{G}}(\cdot | S) \rangle} \quad (\sigma \in \Omega). \quad (2.5)$$

We refer to $\mu_{\mathbf{G}, v \rightarrow a}(\cdot | S), \mu_{\mathbf{G}, a \rightarrow v}(\cdot | S)$ as the *standard messages given S*.

Definition 2.4. Let $\varepsilon > 0$. An event $S \subset \Omega^n$ is an ε -Bethe state of \mathbf{G} if the following two conditions hold.

BS1: the standard messages given S are an ε -Belief Propagation fixed point.

BS2: if $\ell, \ell' \leq 1/\varepsilon$ and if $I \subset V_n, J \subset F_m$ are independent uniformly random sets of sizes $|I| = \ell, |J| = \ell'$, then for every $\sigma \in \Omega^{V_n}$ we have

$$\mathbb{E} \left| \langle \mathbf{1}\{\forall v \in I \cup \partial J \cup \partial^2 I: \sigma_v = \sigma_v\}, \mu_{\mathbf{G}}(\cdot | S) \rangle \right. \\ \left. - \prod_{v \in I} \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a \setminus v} \mu_{w \rightarrow a}(\sigma_w | S)}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{w \rightarrow a}(\tau_w | S)} \cdot \prod_{a \in \partial J} \frac{\psi_a(\sigma) \prod_{w \in \partial a} \mu_{w \rightarrow a}(\sigma_w | S)}{\sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a} \mu_{w \rightarrow a}(\tau_w | S)} \right| < \varepsilon. \quad (2.6)$$

Thus, on a Bethe state the standard messages form an approximate Belief Propagation fixed point. Furthermore, locally around a bunch of randomly chosen variable and constraint nodes the Boltzmann distribution is characterized by the standard messages. In particular, setting $\ell = 0$ and $\ell' = 1$ in **BS2**, we see that the conditional joint distribution $\mu_{\mathbf{G}, \partial a}(\cdot | S)$ of the variables around a typical random constraint node a reads

$$\mu_{\mathbf{G}, \partial a}(\sigma | S) = \frac{\psi_a(\sigma) \prod_{w \in \partial a} \mu_{w \rightarrow a}(\sigma_w | S)}{\sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a} \mu_{w \rightarrow a}(\tau_w | S)} + O(\varepsilon) \quad (\sigma \in \Omega^{\partial a}). \quad (2.7)$$

Additionally, setting $\ell = 1$ and $\ell' = 0$, we find that the local distribution around a typical variable node v , i.e., the distribution $\mu_{\mathbf{G}, v \cup \partial^2 v}(\cdot | S)$ induced on the second neighborhood of v , reads

$$\mu_{\mathbf{G}, v \cup \partial^2 v}(\sigma | S) = \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a} \mu_{w \rightarrow a}(\sigma_w | S)}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a} \mu_{w \rightarrow a}(\tau_w | S)} + O(\varepsilon) \quad (\sigma \in \Omega^{v \cup \partial^2 v}). \quad (2.8)$$

Thus, for most variable nodes v the conditional Boltzmann marginal $\mu_{\mathbf{G}, v}(\cdot | S)$ satisfies

$$\mu_{\mathbf{G}, v}(\sigma | S) = \frac{p(\sigma) \prod_{a \in \partial v} \mu_{a \rightarrow v}(\sigma | S)}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \mu_{a \rightarrow v}(\chi | S)} + O(\varepsilon) \quad (\sigma \in \Omega). \quad (2.9)$$

Apart from the conditioning on S , the formulas (2.7)–(2.9) coincide with the ones known in the acyclic case [48].

In addition, (2.6) implies that if we pick a few variable and/or constraint nodes randomly, then the joint distribution of their neighborhoods approximately factorizes. Applied to $\ell = 2, \ell' = 0$, this means that once we condition on S , the joint distribution of two randomly chosen variable nodes is close to a product distribution:

$$\frac{1}{n^2} \sum_{1 \leq i < j \leq n} \mathbb{E} \left\| \mu_{\mathbf{G}, v_i, v_j}(\cdot | S) - \mu_{\mathbf{G}, v_i}(\cdot | S) \otimes \mu_{\mathbf{G}, v_j}(\cdot | S) \right\|_{\text{TV}} = O(\varepsilon); \quad (2.10)$$

in statistical physics jargon, the conditional distribution $\mu_{\mathbf{G}}(\cdot | S)$ is *replica symmetric*.

Confirming the picture sketched by the cavity method and vindicating the use of Belief Propagation for the study of the Boltzmann distribution, the following theorem shows that w.h.p. the Boltzmann distribution of a random factor graph decomposes into a relatively small number of Bethe states.

Theorem 2.5. *For any function $L = L(n) \rightarrow \infty$ there exists $\varepsilon = \varepsilon(n) \rightarrow 0$ such that the following is true. There exists a decomposition $S_0 = S_0(\mathbf{G}), S_1 = S_1(\mathbf{G}), \dots, S_\ell = S_\ell(\mathbf{G})$, $\ell = \ell(\mathbf{G}) \leq L$, of Ω^n into non-empty sets such that $\mu_{\mathbf{G}}(S_0) \leq \varepsilon$ such that with high probability $S_1, \dots, S_\ell \subset \Omega^n$ are ε -Bethe states. The same statement holds with \mathbf{G} replaced by \mathbb{G} .*

An important feature of Theorem 2.5 is that the upper bound L on the size of the Bethe state decomposition can be an arbitrarily slowly growing function of n . Thus, the Gibbs measure can generally be decomposed into relatively few Bethe states, within which long-range correlations are negligible and where short-range correlations are characterized by Belief Propagation.

2.3. The free energy. Apart from the structure of the Boltzmann distribution, a second key challenge is the computation of the free energy. More specifically, arguably the single most important quantity associated with a random factor graph model is the *free energy density*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\log Z(\mathbf{G})]. \quad (2.11)$$

Of course, it comes as no surprise that computing (2.11) generally poses a formidable challenge. In fact, even the existence of the limit remains an unresolved problem in several interesting cases.

The next theorem provides a formula for (and *en passant* establishes the existence of) the limit (2.11) in terms of the Bethe state decomposition from Theorem 2.5 for a broad class of models. We merely require a certain ‘convexity condition’. This condition can be stated neatly in terms of a space that resembles the graphon space from combinatorics [46]. Specifically, let \mathcal{K} be the space of all measurable maps $[0, 1]^2 \rightarrow \mathcal{P}(\Omega)$ modulo equality (Lebesgue-)almost everywhere. We call these maps *strong kernels*. For $(s, x) \in [0, 1]^2$ and $\mu \in \mathcal{K}$ we let $\mu_{s, x} \in \mathcal{P}(\Omega)$ denote the function value of μ at (s, x) . Further, for $\mu, \mu' \in \mathcal{K}$ we define the *cut distance*

$$\mathcal{D}_\square(\mu, \mu') = \inf_{\varphi, \varphi'} \sup_{\substack{S, X \subset [0, 1] \\ \omega \in \Omega}} \left| \int_S \int_X \mu_{s, x}(\omega) - \mu'_{\varphi(s), \varphi'(x)}(\omega) dx ds \right|, \quad (2.12)$$

where the infimum is over all measurable $\varphi, \varphi' : [0, 1] \rightarrow [0, 1]$ that preserve the Lebesgue measure and where the supremum runs over all measurable $S, X \subset [0, 1]$. Strictly speaking, $\mathcal{D}_\square(\cdot, \cdot)$ is a pre-metric (as possibly $\mathcal{D}_\square(\mu, \nu) = 0$ even though $\mu \neq \nu$). We therefore let \mathfrak{K} be the metric space where any two μ, ν with $\mathcal{D}_\square(\mu, \nu) = 0$ are identified. Then \mathfrak{K} is a compact Polish space [40]. Additionally, we write \mathfrak{D} for the space of all probability distributions on \mathfrak{K} .

Crucially, the convexity assumption that we require comes solely in terms of the distribution P on the set Ψ of weight functions. Namely, let $\mathbf{x} = (\mathbf{x}_i)_{i \geq 1}$ be a sequence of independent uniformly random points in $[0, 1]$, chosen independently of $\boldsymbol{\psi} \in \Psi$. Writing $\mathbb{E}[\cdot]$ for the expectation on $\mathbf{x}, \boldsymbol{\psi}$, we make the following assumption.

For all $\mu, \mu' \in \mathcal{K}$ and for every integer $\ell \geq 1$,

$$\begin{aligned} & \mathbb{E} \left[\left(1 - \sum_{\sigma \in \Omega^k} \boldsymbol{\psi}(\sigma) \int_0^1 \prod_{i=1}^k \mu_{s, \mathbf{x}_i}(\sigma_i) ds \right)^\ell \right] + (k-1) \mathbb{E} \left[\left(1 - \sum_{\sigma \in \Omega^k} \boldsymbol{\psi}(\sigma) \int_0^1 \prod_{i=1}^k \mu'_{s, \mathbf{x}_i}(\sigma_i) ds \right)^\ell \right] \\ & \geq \sum_{h=1}^k \mathbb{E} \left[\left(1 - \sum_{\sigma \in \Omega^k} \boldsymbol{\psi}(\sigma) \int_0^1 \mu_{s, \mathbf{x}_h}(\sigma_h) \prod_{i \in [k] \setminus \{h\}} \mu'_{s, \mathbf{x}_i}(\sigma_i) ds \right)^\ell \right]. \end{aligned} \quad (\text{POS})$$

We will see in Section 7 that **POS** is easily verified for several interesting models, including the spin glass model from Section 1.

To obtain the formula for the free energy, we will represent the Bethe state decomposition of the random factor graph by a point in \mathfrak{K} . Specifically, let \mathbf{X}, \mathbf{Y} be random variables with distribution $\text{Po}(\omega)$ for an integer $\omega > 0$, mutually independent and independent of \mathbf{G} . Then with S_1, \dots, S_ℓ the decomposition promised by Theorem 2.5 we introduce for $i = 1, \dots, \ell$,

$$\begin{aligned} \check{z}_{\mathbf{G}, i} &= \mu_{\mathbf{G}}(S_i) \cdot \prod_{i=1}^X \left(\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v_i} \sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\tau_{v_i} = \chi\} \psi_a(\tau) \prod_{w \in \partial a \setminus v_i} \mu_{w \rightarrow a}(\tau_w | S_i) \right)^{-1} \\ & \quad \cdot \prod_{i=1}^Y \left(\sum_{\tau \in \Omega^{\partial a_i}} \psi_{a_i}(\tau) \prod_{w \in \partial a_i} \mu_{w \rightarrow a_i}(\tau_w | S_i) \right)^{-1}, \end{aligned} \quad (2.13)$$

and we let $\check{z}_{\mathbf{G}} = \sum_{i=1}^\ell \check{z}_{\mathbf{G}, i}$. It will emerge that combinatorially $\check{z}_{\mathbf{G}, i} / \check{z}_{\mathbf{G}}$ represents the probability mass of the Bethe state S_i in the factor graph \mathbf{G}' where we remove the first \mathbf{Y} constraint nodes a_1, \dots, a_Y as well as the first \mathbf{X} variable nodes v_1, \dots, v_X along with their adjacent constraint nodes. While this removal operation has no discernible impact on the free energy (so long as $\omega = o(n)$), it enables us to set up a recurrence for computing this quantity.

The recurrence comes in terms of the messages sent out by those variable nodes that are left with degree $d-1$ after the removal operation. We thus set up a kernel that captures these messages. Specifically, let v_{h_1}, \dots, v_{h_t} be the variable nodes of degree $d-1$ in the factor graph \mathbf{G}' and let b_1, \dots, b_t be their \mathbf{G} -neighbors that got deleted. Then we define the kernel $\check{\mu}_{\mathbf{G}, X, Y} : [0, 1]^2 \rightarrow \mathcal{P}(\Omega)$ by letting

$$\check{\mu}_{\mathbf{G}, X, Y} : (s, x) \mapsto \sum_{i=1}^t \sum_{j=1}^{\ell} \mathbf{1} \left\{ t-1 \leq x < t, \sum_{h < j} \check{z}_{\mathbf{G}, h} < s \check{z}_{\mathbf{G}} \leq \sum_{h \leq j} \check{z}_{\mathbf{G}, h} \right\} \mu_{\mathbf{G}, v_{h_i} \rightarrow b_i}(\cdot | S_j). \quad (2.14)$$

Recalling that $\mathbf{G}, \mathbf{X}, \mathbf{Y}$ are random, we write $\check{\pi}_{n, \omega} \in \mathcal{D}$ for the distribution of $\check{\mu}_{\mathbf{G}, X, Y}$. Analogously, we write $\check{\pi}_{n, \omega, \mathcal{S}}$ for the distribution of $\check{\mu}_{\mathbf{G}, X, Y}$ defined for the simple random factor graph.

Finally, we introduce a functional on the space \mathcal{D} that encodes the recurrence for computing the free energy from the Bethe state decomposition. Namely, let $(\mathbf{x}_{i, j})_{i, j \geq 1}$ be a family of random variables that are uniform on $[0, 1]$, let $(\mathbf{h}_i)_{i \geq 1}$ be a family of random variables that are uniform on $\{1, \dots, k\}$, let $(\boldsymbol{\psi}_i)_{i \geq 1}$ be a sequence of samples from P , and let $\boldsymbol{\mu}^\pi \in \mathfrak{K}$ be a sample from $\pi \in \mathcal{D}$, all mutually independent; then

$$\mathcal{B}(\pi) = \mathbb{E} \left[\log \int_0^1 \sum_{\sigma \in \Omega} p(\sigma) \prod_{i=1}^d \sum_{\substack{\tau \in \Omega^k : \\ \tau_{h_i} = \sigma}} \boldsymbol{\psi}_i(\tau) \prod_{j \neq h_i} \boldsymbol{\mu}_{s, \mathbf{x}_{i, j}}^\pi(\tau_j) ds - d(1-k^{-1}) \log \int_0^1 \sum_{\tau \in \Omega^k} \boldsymbol{\psi}_1(\tau) \prod_{j=1}^k \boldsymbol{\mu}_{s, \mathbf{x}_{1, j}}^\pi(\tau_j) ds \right]. \quad (2.15)$$

We obtain the following expression for the free energy.

Theorem 2.6. *Assume that condition **POS** is satisfied. Then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\log Z(\mathbf{G})] = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\check{\pi}_{n, \omega}), \quad \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\log Z(\mathbb{G})] = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\check{\pi}_{n, \omega, \mathcal{S}}).$$

In particular, the limit on the left hand side exists, and it can be computed from the Bethe state decomposition.

2.4. A variational formula. We proceed to state a variational formula for the free energy of the random factor graph models akin to the one from Theorem 1.3 for the spin glass model. Namely, we express the limit (2.11) variationally as the infimum of $\mathcal{B}(\pi)$ over π chosen from a certain subspace $\mathcal{D}^* \subset \mathcal{D}$. The definition of \mathcal{D}^* is an adaptation to the Bethe lattice of the invariance property that Panchenko [55] put forward in the case of the Erdős-Rényi model.

To define the subspace \mathfrak{D}^* let $\mu \in \mathcal{K}$, let $s \in [0, 1]$ and let $N, M \geq 0$ be integers. We introduce the random variable

$$\begin{aligned} \mathbf{z}(s) = & \prod_{i=1}^N \left[\sum_{\sigma \in \Omega} p(\sigma) \prod_{j=1}^d \sum_{\tau \in \Omega^k} \mathbf{1}\{\tau_{\mathbf{h}_i} = \sigma\} \boldsymbol{\psi}_{di+j}(\tau) \prod_{h \neq \mathbf{h}_i} \mu_{s, \mathbf{x}_{k(di+j)+h}}(\tau_h) \right] \\ & \cdot \prod_{i=1}^M \left[\sum_{\tau \in \Omega^k} \boldsymbol{\psi}_{dN+i}(\tau) \prod_{j=1}^k \mu_{s, \mathbf{x}_{dk(N+1)+j}}(\tau_j) \right]. \end{aligned} \quad (2.16)$$

Further, let

$$\mathbf{t} = \mathbf{t}(s) = \inf \left\{ \theta \in [0, 1] : \int_0^\theta \mathbf{z}(u) du \geq s \int_0^1 \mathbf{z}(u) du \right\} \quad \text{and} \quad \mu_{s,x}^{*(N,M)} = \mu_{\mathbf{t},x}. \quad (2.17)$$

Thus, for each $\mu \in \mathcal{K}$ we obtain a random $\mu^{*(N,M)} \in \mathcal{K}$. Further, given $\pi \in \mathfrak{D}$ we can apply this operation to a randomly chosen kernel $\boldsymbol{\mu}^\pi \in \mathcal{K}$, thus obtaining a random kernel $\boldsymbol{\mu}^{\pi^{*(N,M)}}$. We denote the distribution of $\boldsymbol{\mu}^{\pi^{*(N,M)}}$ by $\pi^{*(N,M)}$. Now, let \mathfrak{D}^* be the set of all densities $\pi \in \mathfrak{D}$ such that $\pi^{*(N,M)} = \pi$ for all $N, M \geq 0$. Then we obtain the following self-contained formula for the free energy.

Theorem 2.7. *Assume that POS holds. Then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\log Z(\mathbf{G})] = \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\log Z(\mathbb{G})] = \min_{\pi \in \mathfrak{D}^*} \mathcal{B}(\pi).$$

Admittedly, the variational formula may not be easy to evaluate. But Theorem 2.7 places a lid on the complexity of the problem, and Theorem 2.6 provides an explicit combinatorial interpretation of the minimizer in terms of Belief Propagation fixed points and Bethe states.

2.5. Discussion and related work. Over the past two decades an enormous amount of research, based on both rigorous and non-rigorous techniques, has been devoted to random factor graph models. Much of this work has been sparked by the cavity method advanced in the original contribution of Mézard and Parisi [50]. A survey of this literature up until about 2008 can be found in [48]. More recently models of Bayesian inference problems such as the stochastic block model have received a great deal of attention as well; this literature is surveyed in [1, 53, 61].

Rigorous work on random factor graphs and the cavity method can broadly be split into two categories. First, contributions that investigate physics predictions on specific models. Many of these contributions, particularly the earlier ones, rely on ‘classical’ techniques such as the second moment method, albeit frequently with physics-inspired twists. Examples include work on the k -SAT threshold [4, 7, 22, 23, 33], which culminated in the proof of the k -SAT threshold conjecture for large k [34], the Potts model and the random graph coloring problem [5, 12, 17, 36] or the hard-core model [28, 35]. Some recent work is based on the powerful but technically demanding idea of ‘spatial coupling’, which has led to important results in, e.g., coding theory [39] and random constraint satisfaction problems [3]. A second line of work focused on the mathematical vindication of the cavity method in general, with applications to specific models of interest. Examples include work on the role of spatial mixing [8, 29, 30, 31], the use of the interpolation method [15, 45, 58], phase transitions in inference problems [13, 21], and contributions based on the asymptotic analysis of the Boltzmann distribution such as the influential work of Panchenko [55] as well as [11, 16]. The present paper belongs to this second category.

In the following we discuss the main results and methods of the paper and how they compare to prior mathematical research. Subsequently we compare the present work with the physics intuition and discuss directions for future research.

Mathematical work. We regard Theorem 2.5 as the main result of the paper. The theorem confirms in great generality one of the key assumptions behind the cavity method and explains the success of Belief Propagation as a device for analyzing random regular factor graph models. Indeed, the existence of a Bethe state decomposition has been conjectured explicitly, e.g., by Mézard and Montanari [48, Chapter 19]; see also Dembo and Montanari [29].

In a prior paper [25] we constructed a Bethe state decomposition for random factor graph models of Erdős-Rényi type, where the constraint nodes independently choose k -tuples of adjacent variable nodes. While we will be able to use some of the general tools developed in that work, the main argument breaks in the case of the Bethe lattice due to its rigid geometry. Indeed, the construction of the Bethe state decomposition hinges on coupling arguments involving, e.g., a coupling of a factor graph with n variable and m constraint nodes and another one with

parameters n' and m' such that $n = n' + O(1)$, $m = m' + O(1)$. Due to the Poisson degree distribution and the Stein-Chen property, such arguments are pretty straightforward in the Erdős-Rényi case. One might say that the Erdős-Rényi graph resembles a gentle climbing wall with footholds supplied by the irregularity of the Poisson degree distribution. By contrast, the Bethe lattice with its regular degree makes for a smooth cliff. As a consequence, the Bethe lattice requires new ideas, leading to a rather subtle but ultimately elegant argument. The upshot is that this proof, which we present in Section 4, can be expected to generalize to other random graph models with given degrees. Apart from the appeal of such lattice-like models from a physics perspective, these models play a vital role, e.g., in coding theory, where a suitably chosen degree sequence is apt to greatly boost performance [59].

Similarly, the variational formula for the free energy provided by Theorem 2.7 is a generalization and adaptation of the formula established by Panchenko [55] for models of Erdős-Rényi type with spins $\Omega = \{\pm 1\}$. Panchenko's proof relies on two ingredients: an interpolation argument and a coupling argument. So does ours. But while the interpolation argument, an adaptation of the technique of Franz and Leone [37], goes through without too much trouble, the coupling argument does not. Once more the rigidity of the Bethe lattice poses substantial challenges that require subtle new arguments. A further, albeit relatively minor extension is that the present work applies to relatively general models with two or more spins, subject only to the condition **POS**. A further similarity between Panchenko's work and ours is the embedding of discrete Boltzmann distributions into a compact metric space, which enables us to pick convergent subsequences. While Panchenko resorts to the Aldous-Hoover representation, here we use the cut metric and the associated kernel space, which is convenient to link the combinatorial representation of the measures in terms of messages directly with the free energy formula. That the Aldous-Hoover representation is closely related to graph limits is, of course, a well known fact [32].

Furthermore, Bayati, Gamarnik and Tetali [15] applied the interpolation method to factor graph models, including ones with regular degrees, to establish the *existence* of the limit $\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbf{G})]$ in certain cases via a super-additivity argument. In the process they also used arguments based on 'cavities', i.e., the removal of a small but linear number of vertices from the graph; a similar trick was used in [24] as well. But here, particularly in the construction of the Bethe state decomposition, we need to tread much more carefully. In particular, while removal of a small linear number of vertices does not shift the free energy too much, here we can only afford the creation of a very small number of cavities in order to avoid a distortion of the Boltzmann distribution, an extremely volatile object.

Theorem 2.6, which expresses the free energy density in terms of the Bethe state decomposition, is a synthesis of Theorems 2.5 and 2.7. The proof shows that the free energy can be expressed in terms of a particular distribution on kernels $[0, 1]^2 \rightarrow \mathcal{P}(\Omega)$, namely the one that encodes the Bethe state decomposition of the random factor graph or, more specifically, the associated Belief Propagation messages. No corresponding result was previously known even in the conceptually simpler Erdős-Rényi case.

Apart from the interpolation method and coupling arguments, the proofs of Theorems 2.5–2.7 rely on some of the techniques that we developed in [11, 21, 26, 16], particularly the cut metric and its ramifications. The cut metric, which we apply to kernel representations of probability distributions, was originally developed in the context of the regularity method [38] and the theory of graph limits in combinatorics [46]. Here we use the cut metric and certain assorted results, such as the 'pinning lemma' (Lemma 3.15 below) from [21] as tools, e.g., in the construction of the Bethe state decomposition.

While the present paper is concerned with diluted models where each node has a (fixed) bounded number of neighbors, there is also a substantial literature on fully connected models. The prime example, of course, is the Sherrington-Kirkpatrick model. The monographs of Panchenko [57] and Talagrand [60] provide an overview of this literature. In particular, the TAP equations, the (simplified) fixed point equations that correspond to the Belief Propagation equations in the fully connected case, have been established in several cases [9].

In Section 7 we work out several application of the general results to specific models, such as the spin glass model from Section 1. Pointers to related work on the specific problems can be found there.

The physics perspective. The seminal work of Mézard and Parisi [50] marks the starting point of a substantial body of physics work. Highlights include the Survey Propagation algorithm and precise predictions on phase transitions, including satisfiability thresholds in combinatorial problems [43, 51, 49].

The results provided by Theorem 2.5–2.7 are perfectly in line with the physics predictions. But we should comment on a subtle point that is apt to cause confusion. Namely, it has been pointed out that within the replica

symmetric phase of certain models the support of the Boltzmann distribution may decompose into an exponentially large number of tiny ‘clusters’ [49, 43], a phenomenon called ‘dynamic replica symmetry breaking’. Indeed, it has been conjectured that each of these tiny clusters induces a Bethe state [48]; for the special case of the random graph coloring problem, this can be verified rigorously [12]. At first glance this proliferation of Bethe states may appear to contradict Theorem 2.5, where the number of Bethe states is upper-bounded by an arbitrarily slowly growing function $L(n)$. Yet the Bethe state decomposition is not unique, and despite the abundance of tiny clusters, μ_G itself is replica symmetric (i.e., condition (2.10) holds for $S = \Omega^n$) throughout the dynamic RSB phase. In effect, Theorem 2.5 would render just a single Bethe state that comprises all of the tiny clusters. By contrast, beyond the dynamic RSB phase, within the so-called condensed phase, Theorem 2.5 would yield a non-trivial decomposition. The existence of a condensed phase has been established rigorously in several examples [18, 21].

The variational formula for the free energy furnished by Theorem 2.7 is in line with the physics work, which does, however, provide additional clues as to the structure of the minimizer of the functional $\mathcal{B}(\cdot)$. Specifically, three different scenarios are expected to occur, depending on the model and the choice of its parameters. First, the replica symmetric scenario with a single (or a bounded number of) Bethe states. Second, the so-called ‘one-step replica symmetry breaking’ scenario, where there are an unbounded number of ‘independent’ Bethe states. Third, the ‘full replica symmetry breaking’ scenario, where the Bethe states form a hierarchical structure; see [48] for a detailed discussion. Clearly, in order to better evaluate the variational formula it would be very valuable to establish this additional structural information rigorously; in the Erdős-Rényi case first attempts have been undertaken in [56].

2.6. Organization. In Section 3 we introduce the necessary pieces of notation and state some basic results that we will need. Then in Section 3.3 we revisit the cut metric. While much of what we need on this subject already appears in earlier papers, there are a few general preparations that we need to make and that we carry out in that section. Subsequently Section 4 deals with the proof of Theorem 2.5. In Sections 5 and 6 we then prove Theorem 2.7 about the variational formula for the free energy. Section 6 also contains the proof of Theorem 2.6. Finally, in Section 7 we work through a few applications, including the spin glass and hard-core models from Section 1.

3. PRELIMINARIES

3.1. Basics. For an integer $\ell \geq 1$ we use the shorthand $[\ell] = \{1, \dots, \ell\}$. Furthermore, the symbols $O(\cdot), \Omega(\cdot), \dots$ refer to the limit $n \rightarrow \infty$ by default. To indicate asymptotics with respect to another variable K tending to infinity, we write $O_K(\cdot), \Omega_K(\cdot)$, etc. Further, where set operations involve singletons, we usually omit braces. For instance, if $x \in X$, then we just write $X \setminus x$ rather than $X \setminus \{x\}$.

For a finite set \mathcal{X} we let $\mathcal{P}(\mathcal{X})$ be the set of all probability distributions on \mathcal{X} , endowed with the total variation distance. More generally, if $(\mathcal{X}, \mathcal{A})$ is a measurable space, then $\mathcal{P}(\mathcal{X}) = \mathcal{P}(\mathcal{X}, \mathcal{A})$ denotes the set of all probability measures on this space. Further, for probability measures $\pi, \pi' \in \mathcal{P}(\mathcal{X})$ we let $\Gamma(\pi, \pi')$ be the set of all couplings of π, π' . Thus, $\gamma \in \Gamma(\pi, \pi')$ is a probability distribution on $\mathcal{X} \times \mathcal{X}$ with marginals π, π' .

Suppose that \mathcal{X} is a finite set, that $n \geq 1$ is an integer and that $\mu \in \mathcal{P}(\mathcal{X}^n)$. Then we denote by $\sigma^\mu, \sigma^{1,\mu}, \sigma^{2,\mu}, \dots$ a sequence of independent samples from μ . We omit the superscript μ where it is evident from the context. Further, if $f : (\mathcal{X}^n)^\ell \rightarrow \mathbb{R}$ is a function, then we write $\langle f(\sigma^1, \dots, \sigma^\ell), \mu \rangle$ for the expectation of f with respect to independent samples from μ ; thus,

$$\langle f(\sigma^1, \dots, \sigma^\ell), \mu \rangle = \sum_{\sigma^1, \dots, \sigma^\ell \in \mathcal{X}^n} f(\sigma^1, \dots, \sigma^\ell) \prod_{i=1}^{\ell} \mu(\sigma^i).$$

Suppose that $\Omega, V \neq \emptyset$ are finite sets. For a distribution $\mu \in \mathcal{P}(\Omega^V)$ and a set $I \subset V$ we denote by μ_I the joint distribution of the coordinates I . That is,

$$\mu_I(\sigma) = \sum_{\tau \in \Omega^I} \mathbf{1}\{\forall i \in I : \tau_i = \sigma_i\} \mu(\tau) \quad (I \subset V, \sigma \in \Omega^I).$$

For $\sigma \in \Omega^I$ we use the shorthand $\mu(\sigma) = \mu_I(\sigma)$. Moreover, if $I = \{i_1, \dots, i_l\}$ we usually write μ_{i_1, \dots, i_l} instead of $\mu_{\{i_1, \dots, i_l\}}$. Additionally, if $I \subset V$ and $\tau \in \Omega^V$, then we let $\tau_I = (\tau_i)_{i \in I}$ be the restriction of τ to I .

We keep the notation from Section 2; in particular, Ω continues to denote a finite set of spins, p is a probability distribution on Ω , Ψ is a measurable space of functions $\Omega^k \rightarrow (0, 1)$, and P is a probability distribution on Ψ . In addition throughout the paper we denote by

$$\mathbf{x}_i, \hat{\mathbf{x}}_i, \mathbf{s}_i, \mathbf{x}_{i,j}, \mathbf{x}'_{i,j}, \mathbf{x}''_{i,j}, \hat{\mathbf{x}}_{i,j} \quad (i, j \geq 1)$$

uniformly distributed random variables with values in $[0, 1]$. Additionally,

$$\boldsymbol{\psi}, \boldsymbol{\psi}_i, \boldsymbol{\psi}_{i,j}, \boldsymbol{\psi}'_{i,j}, \boldsymbol{\psi}''_{i,j}, \hat{\boldsymbol{\psi}}_{i,j} \quad (i, j \geq 1)$$

denote elements of Ψ drawn from the distribution P . Further,

$$\mathbf{h}_i, \mathbf{h}_{i,j}, \mathbf{h}'_{i,j}, \hat{\mathbf{h}}_{i,j} \quad (i, j \geq 1)$$

are uniformly distributed random variables with values in $[k]$. All of the above random variables are mutually independent as well as independently of any other sources of randomness. These random variables yield random functions that will play an important role: for $i \geq 1$ we let

$$\begin{aligned} \boldsymbol{\varphi}_i : \Omega^{dk} &\rightarrow \mathbb{R}, & \sigma \mapsto \sum_{\chi \in \Omega} p(\chi) \prod_{j=1}^d \mathbf{1}\{\sigma_{k(j-1)+\mathbf{h}_{i,j}} = \chi\} \boldsymbol{\psi}_{i,j}(\sigma), \\ \hat{\boldsymbol{\varphi}}_i : \Omega^{dk} &\rightarrow \mathbb{R}, & \sigma \mapsto \sum_{\chi \in \Omega} p(\chi) \prod_{j=1}^d \mathbf{1}\{\sigma_{k(j-1)+\hat{\mathbf{h}}_{i,j}} = \chi\} \hat{\boldsymbol{\psi}}_{i,j}(\sigma). \end{aligned}$$

3.2. Factor graphs. In Section 2 we already introduced the random factor graph model $\mathbf{G}(n, d, p, P)$. To facilitate the proofs we need the following abstract definition.

Definition 3.1. Suppose that $(\mathcal{X}, \mathcal{A})$ is a measurable space. An \mathcal{X} -factor graph $G = (V, F, (\partial a)_{a \in F}, (\psi_a)_{a \in F}, \pi)$ consists of

- a finite set V of variable nodes,
- a finite set F of constraint nodes,
- a set $\partial a \subset V$ for each $a \in F$,
- a function $\psi_a : \mathcal{X}^{\partial a} \rightarrow [0, \infty)$ for each $a \in F$ and
- a probability measure \mathbf{p} on \mathcal{X}^V , called the prior.

A factor graph induces a bipartite graph on $V \cup F$, where $x \in V$ is adjacent to $a \in F$ iff $v \in \partial a$. Accordingly, for a variable node v we let $\partial v \subset F$ be the set of adjacent constraint nodes (i.e., $a \in \partial v$ iff $v \in \partial a$). The bipartite graph defines a metric on $V \cup F$, the shortest path metric. For a variable or constraint node u we let $\partial^\ell u = \partial_G^\ell u$ be the set of all nodes at distance precisely ℓ from u . Moreover, $\nabla^\ell u = V \cap (u \cup \bigcup_{i \leq \ell} \partial^i u)$ denotes the set of all variable nodes at distance no more than ℓ from u .

Further, for an assignment $\sigma \in \Omega^V$ and $a \in F$ we use the notation $\psi_a(\sigma) = \psi_a(\sigma_{\partial a})$ and we define

$$\psi_G(\sigma) = \prod_{a \in F} \psi_a(\sigma_{\partial a}) \quad \text{and} \quad Z(G) = \int_{\mathcal{X}^V} \psi_G(\sigma) d\mathbf{p}(\sigma).$$

Providing that $Z(G) > 0$, we introduce a probability measure μ_G on \mathcal{X}^V , the *Boltzmann distribution*, by letting

$$d\mu_G(\sigma) = \frac{\psi_G(\sigma)}{Z(G)} d\mathbf{p}(\sigma).$$

Mostly the factor graphs that we deal with will have a finite space $\mathcal{X} = \Omega$ and the prior \mathbf{p} will be the product measure $\mathbf{p} = \bigotimes_{v \in V} p_v$. In this case we introduce the standard messages given an event $S \subset \Omega^V$ as in Section 2: for a constraint node a and $v \in \partial a$ we let

$$\mu_{G, v \rightarrow a}(\sigma | S) = \frac{\langle \mathbf{1}\{\sigma_v = \sigma\} / \psi_a(\sigma), \mu_G(\cdot | S) \rangle}{\langle 1 / \psi_a(\sigma), \mu_G(\cdot | S) \rangle} \quad (\sigma \in \Omega), \quad (3.1)$$

$$\mu_{G, a \rightarrow v}(\sigma | S) = \frac{\langle \mathbf{1}\{\sigma_v = \sigma\} / (p_v(\sigma) \prod_{b \in \partial v \setminus a} \psi_b(\sigma)), \mu_G(\cdot | S) \rangle}{\langle 1 / (p_v(\sigma) \prod_{b \in \partial v \setminus a} \psi_b(\sigma)), \mu_G(\cdot | S) \rangle} \quad (\sigma \in \Omega). \quad (3.2)$$

In the case that $S = \Omega^V$ is the entire phase space, we omit the conditioning from the notation and just write $\mu_{G, v \rightarrow a}$ and $\mu_{G, a \rightarrow v}$, respectively.

3.3. The cut metric revisited. The cut metric, defined in (2.12), plays a key role in the proofs of the main results. In this section we summarize a few basic facts about the cut metric. Although some have been proved in prior work, we will need to provide a few extensions and adaptations for our purposes. In addition to the continuous version from (2.12), we also need a discrete version of the cut metric, which we present in Section 3.3.2.

3.3.1. *The continuous cut metric.* We remember that \mathcal{K} denotes the space of all measurable maps $[0, 1]^2 \rightarrow \mathcal{P}(\Omega)$, up to equality almost everywhere; we call such maps *strong kernels*. The cut distance (2.12) induces a pre-metric on this space [40]. Moreover, on the space \mathfrak{K} obtained by identifying points at cut distance 0 the cut distance yields a metric. The elements of \mathfrak{K} are called *weak kernels*. We drop the attribute and just speak of kernels where there is no danger of confusion.

Proposition 3.2 ([26]). *Endowed with the cut distance \mathfrak{K} is a compact Polish space.*

We continue to write \mathfrak{D} for the space of all probability measures on \mathfrak{K} . This space is endowed with the weak topology. Since \mathfrak{K} is a compact Polish space, so is \mathfrak{D} . Hence, there is a natural metric on \mathfrak{D} that induces the weak topology, the L_1 -Wasserstein metric. We take license to denote this metric by $\mathcal{D}_\square(\cdot, \cdot)$ as well. Thus, recalling that $\Gamma(\pi, \pi')$ is the set of all couplings of $\pi, \pi' \in \mathfrak{D}$, we have

$$\mathcal{D}_\square(\pi, \pi') = \inf \left\{ \int_{\mathfrak{K} \times \mathfrak{K}} \mathcal{D}_\square(\mu, \mu') d\gamma(\mu, \mu') : \gamma \in \Gamma(\pi, \pi') \right\}.$$

For $\pi \in \mathfrak{D}$ we let $\mu^\pi \in \mathfrak{K}$ denote a sample. We just write μ where π is apparent.

By comparison to other metrics on the space of measurable functions $[0, 1]^2 \rightarrow \mathcal{P}(\Omega)$ the cut metric is extremely weak; this is highlighted by the compactness of the space \mathfrak{K} provided by Proposition 3.2. Yet the cut metric is sufficiently strong to ensure that certain functions that will be of vital interest to us are continuous. Indeed, suppose that $m, n > 0$ are integers and that $f : \Omega^{m \times n} \rightarrow \mathbb{R}$ is a function. Then for $\mu \in \mathcal{K}$ we define the random variable

$$\langle f, \mu \rangle = \sum_{\sigma \in \Omega^{m \times n}} f(\sigma) \int_0^1 \cdots \int_0^1 \prod_{i=1}^m \prod_{j=1}^n \mu_{s_i, \mathbf{x}_j}(\sigma_{i,j}) ds_1 \cdots ds_m.$$

Because we average out the s_i and the \mathbf{x}_i are uniform, the random variables $\langle f, \mu \rangle$ and $\langle f, \nu \rangle$ are identically distributed if $\mathcal{D}_\square(\mu, \nu) = 0$. Thus, we may safely write $\langle f, \mu \rangle$ for $\mu \in \mathfrak{K}$.

Lemma 3.3. *For any $f : \Omega^{m \times n} \rightarrow \mathbb{R}$, $\ell \geq 1$ the map $\mu \in \mathfrak{K} \mapsto \mathbb{E}[\langle f, \mu \rangle^\ell]$ is continuous with respect to the cut metric.*

The proof of Lemma 3.3 can be found in the appendix. For a probability distribution $\pi \in \mathfrak{D}$ we let $\langle f, \pi \rangle$ be the random variable $\langle f, \mu^\pi \rangle$, with μ^π chosen independently of the \mathbf{x}_i . Since \mathfrak{D} carries the weak topology, Lemma 3.3 implies

Corollary 3.4. *For any $f : \Omega^{m \times n} \rightarrow \mathbb{R}$, $\ell \geq 1$ the map $\pi \in \mathfrak{D} \mapsto \mathbb{E}[\langle f, \pi \rangle^\ell]$ is continuous.*

We recall the functional $\mathcal{B}(\cdot)$ from (2.15).

Corollary 3.5. *The map $\pi \in \mathfrak{D} \mapsto \mathcal{B}(\pi)$ is continuous.*

Proof. Thanks to the tail bound (2.1), we can approximate the logarithms in (2.15) by polynomials. Therefore, the assertion follows from Corollary 3.4. \square

The set \mathfrak{D}^* of $\pi \in \mathfrak{D}$ that are invariant under the $*(N, M)$ -operation is a closed subset of \mathfrak{D} . To see this, and to interpret the $*(N, M)$ -operation nicely in terms of operations that are continuous under the cut metric, we introduce the following general transformation. Suppose that $f : \Omega^N \rightarrow (0, \infty)$ is a function and that $\mu \in \mathcal{K}$. Then we define a random $f * \mu \in \mathcal{K}$ as follows. Letting

$$z_s = z_s^{f, \mu} = \sum_{\sigma \in \Omega^N} f(\sigma) \prod_{i=1}^N \mu_{s, \hat{x}_i}(\sigma_i), \quad z = z^{f, \mu} = \int_0^1 z_s^{f, \mu} ds,$$

we introduce

$$t = t_s^{f, \mu} = \inf \left\{ \theta \in [0, 1] : \int_0^\theta z_u du \geq sz \right\}.$$

Now, $f * \mu_{s,x} = \mu_{t,x}$. We emphasize that $f * \mu \in \mathcal{K}$ is random, dependent on $\hat{x}_1, \dots, \hat{x}_N$. The kernel is characterized by the identity

$$\int_0^1 \int_0^1 g_{s,x} \cdot f * \mu_{s,x}(\omega) ds dx = \frac{1}{z} \sum_{\sigma \in \Omega^N} f(\sigma) \int_0^1 \int_0^1 g_{s,x} \cdot \mu_{s,x}(\omega) z_s ds dx \quad \text{for all } g : [0, 1]^2 \rightarrow [0, 1], \omega \in \Omega.$$

Further, since the \hat{x}_i are uniform, we have $\mathcal{D}_\square(f * \mu, f * \nu) = 0$ if $\mathcal{D}_\square(\mu, \nu) = 0$. Hence, the $*$ -operation extends to weak kernels. Furthermore, for a distribution π we let $f * \pi$ be the distribution of $f * \mu^\pi$.

Lemma 3.6. *For any function $f : \Omega^k \rightarrow (0, \infty)$ the map $\mathfrak{K} \rightarrow \mathfrak{D}$, $\mu \mapsto f * \mu$ is continuous.*

The proof of Lemma 3.6 can be found in the appendix.

The $*(N, M)$ -operation is an application of the above $*$ -operation to a particular random function f . To define this random function, we need one more piece of notation. Namely, for functions $f : \Omega^{M \times N} \rightarrow \mathbb{R}$, $g : \Omega^{M \times L} \rightarrow \mathbb{R}$ we define

$$f \oplus g : \Omega^{M \times (N+L)} \rightarrow \mathbb{R}, \quad \sigma \mapsto f((\sigma_{i,j})_{i \in [M], j \in [N]}) \cdot g((\sigma_{i,j+N})_{i \in [M], j \in [L]}).$$

In words, we stick the first N ‘columns’ of σ into f and the last L columns into g and multiply the results. Recalling the random functions $\hat{\psi}_i$, $\hat{\phi}_i$ from Section 3.1, we obtain the following.

Lemma 3.7. *For any $\mu \in \mathfrak{K}$ the random $\mu^{*(N, M)} \in \mathfrak{K}$ is distributed as $(\bigoplus_{i=1}^N \hat{\phi}_i \oplus \bigoplus_{i=1}^M \hat{\psi}_i) * \mu$.*

Proof. This is immediate from the construction of $\mu^{*(N, M)}$. \square

Corollary 3.8. *For any N, M the map $\mu \in \mathfrak{K} \mapsto \mu^{*(N, M)}$ is continuous with respect to the cut metric.*

Proof. Since \mathfrak{D} carries the weak topology, which is induced by the Wasserstein metric, the assertion follows from Lemmas 3.6–3.7 and (2.1). \square

As a further immediate consequence of Lemma 3.7 we obtain

Corollary 3.9. *A distribution $\pi \in \mathfrak{D}$ belongs to \mathfrak{D}^* if and only if $\pi = (\bigoplus_{i=1}^N \hat{\phi}_i \oplus \bigoplus_{i=1}^M \hat{\psi}_i) * \pi$ for all N, M .*

In particular, Corollaries 3.8 and 3.9 imply that the map $\pi \mapsto \pi^{*(N, M)}$ is continuous for all $N, M \geq 0$. Consequently, \mathfrak{D}^* is a closed subset of \mathfrak{D} .

3.3.2. The discrete version. Apart from the ‘continuous’ installment of the cut metric, defined on kernels, we also need a discrete variant, defined on probability measures on discrete sets. To be precise, with $\Omega \neq \emptyset$ our finite set of spins and V another finite set of size $n \geq 1$, we define a metric $\Delta_{\square}(\cdot, \cdot)$ on $\mathcal{P}(\Omega^V)$ as follows. Recalling that $\Gamma(\mu, \nu)$ is the set of all couplings of probability measures μ, ν on Ω^V , we let

$$\Delta_{\square}(\mu, \nu) = \frac{1}{n} \min_{\gamma \in \Gamma(\mu, \nu)} \max_{\substack{I \subseteq V \\ B \subseteq \Omega^V \times \Omega^V \\ \omega \in \Omega}} \left| \sum_{i \in I} \sum_{(\sigma, \tau) \in B} \gamma(\sigma, \tau) (\mathbf{1}\{\sigma_i = \omega\} - \mathbf{1}\{\tau_i = \omega\}) \right| \quad \text{for } \mu, \nu \in \mathcal{P}(\Omega^V). \quad (3.3)$$

Fact 3.10 ([25]). *$\Delta_{\square}(\cdot, \cdot)$ is a metric on $\mathcal{P}(\Omega^V)$.*

We refer to $\Delta_{\square}(\cdot, \cdot)$ as the *discrete cut metric*.

Suppose that V is a finite set. A measure $\mu \in \mathcal{P}(\Omega^V)$ can be represented by a point $\dot{\mu} \in \mathcal{K}$. Indeed, assume without loss that $V = [n]$ and that $\Omega = [q]$. Then the set Ω^V can be ordered lexicographically as $\sigma^{(1)}, \dots, \sigma^{(q^n)}$. We define $\dot{\mu} \in \mathcal{K}$ by letting

$$\dot{\mu}_{s,x} = \sum_{i=1}^n \sum_{j=1}^{q^n} \mathbf{1}\{(i-1)/n \leq x < i/n\} \mathbf{1}\left\{ \sum_{h < j} \mu(\sigma^{(h)}) \leq s < \sum_{h \leq j} \mu(\sigma^{(h)}) \right\} \delta_{\sigma_i^{(j)}}.$$

Comparing (3.3) with the definition (2.12) of the continuous cut metric, we see that

$$\mathcal{D}_{\square}(\dot{\mu}, \dot{\nu}) \leq \Delta_{\square}(\mu, \nu) \quad (\mu, \nu \in \Omega^V). \quad (3.4)$$

The discrete cut metric encodes a great deal of information about the discrete measures. A particularly important case occurs when a measure $\mu \in \mathcal{P}(\Omega^V)$ is close to a product measure. To be precise, we say that μ is ε -extremal if $\Delta_{\square}(\mu, \bigotimes_{v \in V} \mu_v) < \varepsilon$. In words, μ is close to the product measure with the same marginals. In addition, $\mu \in \mathcal{P}(\Omega^V)$ is (ε, ℓ) -symmetric if

$$\frac{1}{|V|^\ell} \sum_{v_1, \dots, v_\ell \in V} \|\mu_{v_1, \dots, v_\ell} - \mu_{v_1} \otimes \dots \otimes \mu_{v_\ell}\|_{\text{TV}} < \varepsilon. \quad (3.5)$$

Informally, if we choose ℓ coordinates randomly, then their joint distribution typically ‘nearly’ factorizes. The following statement shows that these concepts are essentially equivalent, up to a moderate loss in the parameters.

Proposition 3.11 ([25]). *For any Ω of size $1 < |\Omega| < \infty$, any $0 < \varepsilon < 1/2$ and any $\ell \geq 2$ there exists $n_0 > 0$ such that for all $n > n_0$ and all $\mu \in \mathcal{P}(\Omega^V)$ the following two statements hold.*

- (i) If μ is $(\varepsilon/9)^3$ -symmetric, then μ is ε -extremal.
- (ii) If μ is $\varepsilon^3/(128|\Omega|)^{4\ell}$ -extremal, then μ is (ε, ℓ) -symmetric.

It is an elementary observation that probability measures that are close in the discrete cut metric cannot have very different marginals. Formally, we have the following.

Lemma 3.12. *For any two probability measures μ, ν on Ω^V we have $\sum_{v \in V} \|\mu_v - \nu_v\|_{\text{TV}} \leq 2|\Omega| \Delta_{\square}(\mu, \nu)$.*

Proof. There exists $\omega \in \Omega$ such that

$$\sum_{v \in V} (\mu_v(\omega) - \nu_v(\omega)) \vee 0 \geq \frac{1}{2|\Omega|} \sum_{v \in V} \|\mu_v - \nu_v\|_{\text{TV}}. \quad (3.6)$$

Let $I = \{i \in V : \mu_i(\omega) \geq \nu_i(\omega)\}$ and $B = \Omega^V \times \Omega^V$. Then for any coupling $\gamma \in \Gamma(\mu, \nu)$,

$$\sum_{i \in I} \sum_{(\sigma, \tau) \in B} \gamma(\sigma, \tau) (\mathbf{1}\{\sigma_i = \omega\} - \mathbf{1}\{\tau_i = \omega\}) = \sum_{i \in V} (\mu_i(\omega) - \nu_i(\omega)) \vee 0. \quad (3.7)$$

Combining (3.6) and (3.7) completes the proof. \square

The converse bound, that close marginals imply closeness in the cut metric, holds for extremal measures.

Lemma 3.13. *For any two ε -extremal $\mu, \nu \in \mathcal{P}(\Omega^V)$ we have $\mathcal{D}_{\square}(\mu, \nu) \leq 2\varepsilon + \sum_{v \in V} \|\mu_v - \nu_v\|_{\text{TV}}$.*

Proof. Assume without loss that $V = [n]$ and let $\bar{\mu} = \bigotimes_{i=1}^n \mu_i$ and $\bar{\nu} = \bigotimes_{i=1}^n \nu_i$. Since μ, ν are ε -extremal, we have

$$\Delta_{\square}(\mu, \bar{\mu}) < \varepsilon, \quad \Delta_{\square}(\nu, \bar{\nu}) < \varepsilon. \quad (3.8)$$

Let $\gamma_i \in \mathcal{P}(\Omega \times \Omega)$ be an optimal coupling of μ_i, ν_i , i.e., $\|\mu_i - \nu_i\|_{\text{TV}} = \sum_{\sigma \neq \tau} \gamma_i(\sigma, \tau)$. Then $\gamma = \bigotimes_{i=1}^n \gamma_i$ is a coupling of $\bar{\mu}, \bar{\nu}$. Further, for any $I \subset [n], B \subset \Omega^n \times \Omega^n, \omega \in \Omega$ we have

$$\left| \sum_{i \in I} \sum_{(\sigma, \tau) \in B} \gamma(\sigma, \tau) (\mathbf{1}\{\sigma_i = \omega\} - \mathbf{1}\{\tau_i = \omega\}) \right| \leq \sum_{i \in I} \sum_{(\sigma, \tau) \in B} \gamma(\sigma, \tau) \mathbf{1}\{\sigma_i \neq \tau_i\} \leq \sum_{i=1}^n \|\mu_i - \nu_i\|_{\text{TV}}.$$

Hence, $\Delta_{\square}(\bar{\mu}, \bar{\nu}) \leq \sum_{i=1}^n \|\mu_i - \nu_i\|_{\text{TV}}$, and thus the assertion follows from (3.8) and the triangle inequality. \square

We also make a note of the following enhanced triangle inequality.

Lemma 3.14 ([25]). *Suppose that $\mu^{(1)}, \nu^{(1)}, \dots, \mu^{(\ell)}, \nu^{(\ell)}$ are probability measures on Ω^V and that $u_1, \dots, u_\ell \geq 0$ are numbers such that $\sum_{i=1}^\ell u_i = 1$. Then*

$$\Delta_{\square} \left(\sum_{i=1}^\ell u_i \mu^{(i)}, \sum_{i=1}^\ell u_i \nu^{(i)} \right) \leq \sum_{i=1}^\ell u_i \Delta_{\square}(\mu^{(i)}, \nu^{(i)}).$$

Finally, we come to an important fact, intimately related to the Szemerédi regularity lemma from combinatorics. Namely, any probability distribution $\mu \in \mathcal{P}(\Omega^V)$ is close in the cut metric to a mixture of a ‘small’ number of product measures. To state this results precisely, suppose that $I \subset V$ and that $\sigma \in \Omega^I$. Let

$$S^{I, \sigma} = \{\tau \in \Omega^V : \tau_I = \sigma\}$$

be the sub-cube of Ω^V where the entries of the coordinates in I coincide with σ . Further, assuming that $\mu \in \mathcal{P}(\Omega^V)$ and $\mu(S^{I, \sigma}) > 0$, we let

$$\mu^{I, \sigma} = \mu[\cdot | S^{I, \sigma}] \quad (3.9)$$

be the corresponding conditional distribution of μ . (If $\mu(S^{I, \sigma}) = 0$, then we agree that $\mu^{I, \sigma}$ is the uniform distribution on $S^{I, \sigma}$.) The following key lemma shows that $\mu^{I, \sigma}$ is likely ε -symmetric for suitably random I, σ .

Lemma 3.15 ([21, Lemma 3.5]). *For any set Ω of size $1 < |\Omega| < \infty$ and any $\varepsilon > 0$ there exist $n_0 > 0$ and a random variable $0 < \theta \leq 2\varepsilon^{-4} \log |\Omega|$ such that for all $n > n_0$ and all $\mu \in \mathcal{P}(\Omega^V)$ the following holds. Let $I \subset V$ be a uniformly random subset of size θ and choose $\sigma \in \Omega^I$ from μ_I . Then $\mathbb{P}[\mu^{I, \sigma} \text{ is } \varepsilon\text{-symmetric}] > 1 - \varepsilon$.*

We can apply Lemma 3.15 multiple times to obtain a decomposition of the set Ω^V into sub-cubes S_1, \dots, S_ℓ such that $\mu[\cdot | S_i]$ is ε -symmetric. To obtain these sub-cubes we just choose the set I randomly as in Lemma 3.15 and let σ range over all $|\Omega|^\theta$ possible assignments of I . We then obtain the following version of the regularity lemma.

Corollary 3.16 ([11]). *For any finite set $\Omega \neq \emptyset$ and any $\varepsilon > 0$ there exist L, n_0 such that for all $n > n_0$ the following is true. For any $\mu \in \mathcal{P}(\Omega^V)$ there exists a partition of Ω^V into pairwise disjoint sets S_0, \dots, S_ℓ , $\ell \leq L$, such that $\mu(S_0) < \varepsilon$ and such that $\mu(\cdot | S_i)$ is ε -symmetric for each $1 \leq i \leq \ell$.*

3.3.3. *Contiguity.* Suppose that $\Omega \neq \emptyset$ is a finite set and let $c \geq 1$. A probability distribution ν on Ω^n is c -contiguous with respect to another probability distribution μ if

$$\nu(\sigma) \leq c\mu(\sigma) \quad \text{for all } \sigma \in \Omega^n.$$

Moreover, μ, ν are *mutually c-contiguous* if each is c -contiguous with respect to the other.

Lemma 3.17. *For any $c \geq 1, \delta > 0$ there exists $\varepsilon > 0$ such that for all large enough n the following is true. Assume that $\mu \in \mathcal{P}(\Omega^n)$ is ε -extremal and that ν is c -contiguous with respect to μ . Then ν is δ -extremal and $\Delta_{\square}(\mu, \nu) < \delta$.*

Proof. Choose $0 < \varepsilon \ll \eta \ll \zeta \ll \delta$ and assume that $n > n_0(\varepsilon)$ is sufficiently large and that μ is ε -extremal. Applying Corollary 3.16 to the measure ν , we obtain a partition S_0, S_1, \dots, S_ℓ of the cube Ω^n into pairwise disjoint sets such that $\nu(S_0) < \eta$ and such that $\nu(\cdot | S_i)$ is η -symmetric for every $i = 1, \dots, \ell$. Moreover, ℓ is bounded by a number $L(\eta, \Omega) > 0$ that depends on η and $|\Omega|$ only.

Suppose that for every $1 \leq i \leq \ell$ with $\nu(S_i) \geq \eta/\ell$ we have

$$\sum_{j=1}^n \|\nu_j(\cdot | S_i) - \mu_j\|_{\text{TV}} \leq \zeta n. \quad (3.10)$$

Then Lemma 3.13 yields $\Delta_{\square}(\nu(\cdot | S_i), \mu) \leq 2\eta + \zeta$. Hence, Lemma 3.14 shows that

$$\Delta_{\square}(\nu, \mu) \leq 4\eta + \zeta < \delta. \quad (3.11)$$

Further, (3.10) implies that $\sum_{j=1}^n \|\nu_j - \mu_j\|_{\text{TV}} \leq \zeta + 2\eta$. Hence, letting $\bar{\mu} = \bigotimes_{i=1}^n \mu_i$, $\bar{\nu} = \bigotimes_{i=1}^n \nu_i$ and applying Lemma 3.13 a second time, we obtain $\Delta_{\square}(\bar{\mu}, \bar{\nu}) \leq \zeta + 2\eta$. Thus, invoking the ε -extremality of μ and (3.11), we conclude that

$$\Delta_{\square}(\nu, \bar{\nu}) \leq \Delta_{\square}(\nu, \mu) + \Delta_{\square}(\mu, \bar{\mu}) + \Delta_{\square}(\bar{\mu}, \bar{\nu}) \leq (4\eta + \zeta) + \varepsilon + (\zeta + 2\eta) < \delta. \quad (3.12)$$

In summary, if (3.10) is satisfied, then (3.11) and (3.12) yield $\Delta_{\square}(\nu, \mu) < \delta$ and $\Delta_{\square}(\nu, \bar{\nu}) < \delta$, as claimed.

Thus, we are left to establish (3.10). Assume for contradiction that there is $1 \leq i \leq \ell$ with $\nu(S_i) \geq \eta/\ell$ and $\sum_{j=1}^n \|\nu_j(\cdot | S_i) - \mu_j\|_{\text{TV}} > \zeta n$. Then there exist $J \subset [n]$ and $\omega \in \Omega$ such that $\sum_{j \in J} \nu_j(\omega | S_i) - \mu_j(\omega) > \zeta n / (2|\Omega|)$. In other words, the random variable $X(\sigma) = \sum_{j \in J} \mathbf{1}\{\sigma_j = \omega\}$ satisfies

$$\langle X, \nu(\cdot | S_i) \rangle - \langle X, \mu \rangle > \zeta n / (2|\Omega|). \quad (3.13)$$

Due to the η -symmetry of $\nu(\cdot | S_i)$ and the ε -symmetry of μ , the second moments work out as

$$\langle X(X-1), \nu(\cdot | S_i) \rangle = \sum_{j, j' \in J: j \neq j'} \langle \mathbf{1}\{\sigma_j = \sigma_{j'} = \omega\}, \nu(\cdot | S_i) \rangle \leq \eta n + \langle X, \nu(\cdot | S_i) \rangle^2, \quad \langle X(X-1), \mu \rangle \leq \varepsilon n + \langle X, \mu \rangle^2. \quad (3.14)$$

Combining (3.13) and (3.14) with Chebyshev's inequality and keeping in mind that $\varepsilon \ll \eta \ll \zeta$, we conclude that the event $B = \{X(\sigma) \geq \langle X, \mu \rangle + \zeta / (4|\Omega|)\}$ satisfies

$$\nu(B | S_i) \geq 3/4, \quad \mu(B) \leq \varepsilon^{1/4}. \quad (3.15)$$

However, if ν is c -contiguous with respect to μ , then (3.15) yields

$$c\varepsilon^{1/4} \geq c\mu(B) \geq \nu(B) = \nu(B | S_i) \nu(S_i) \geq 3\eta / (4\ell) \geq 3\eta / (4L(\eta, \Omega)),$$

which contradicts the choice of the parameters ε, η . \square

Corollary 3.18. *For any $\delta > 0$ there exists $\varepsilon > 0$ such that the following is true. Suppose that μ is ε -extremal and that $S \subset \Omega^n$ is an event such that $\mu(S) \geq \delta$. Then $\mu(\cdot | S)$ is δ -extremal and $\Delta_{\square}(\mu(\cdot | S), \mu) < \delta$.*

Proof. Since $\mu(\sigma | S) \leq \mu(\sigma) / \mu(S)$ for every σ , the conditional distribution $\mu(\cdot | S)$ is $1/\varepsilon$ -contiguous with respect to μ . Thus, the assertion follows from Lemma 3.17 immediately. \square

4. BETHE STATE DECOMPOSITIONS

In this section we prove Theorem 2.5. Before getting into the details, why might we expect the statement of the theorem to be correct? In order to construct the desired decomposition of the phase space Ω^{V_n} we could attempt to apply Lemma 3.15. The lemma shows that if we randomly pick a set $\mathbf{I} \subset V_n$ of a moderate size of, say, $\mathbf{I} = O(\log \log n)$, we will likely obtain a decomposition $(S^{I,\sigma})_{\sigma \in \Omega^I}$ such that (2.10) is satisfied for most of its parts. Formally, Lemma 3.15 guarantees that with high probability,

$$\frac{1}{n^2} \sum_{1 \leq i < j \leq n} \left\langle \left\| \mu_{\mathbf{G}, v_i, v_j}(\cdot | S^{I,\sigma}) - \mu_{\mathbf{G}, v_i}(\cdot | S^{I,\sigma}) \otimes \mu_{\mathbf{G}, v_j}(\cdot | S^{I,\sigma}) \right\|_{\text{TV}}, \mu_{\mathbf{G}} \right\rangle = o(1). \quad (4.1)$$

Proposition 3.11 extends (4.1) from pairwise independence to independence of bounded numbers of randomly chosen variable nodes. Thus, the pinning operation eliminates long-range correlations.

But why should short-range correlations be described by Belief Propagation? Recalling the standard messages from (2.4)–(2.5), Belief Propagation (2.3) asserts that with high probability,

$$\mu_{\mathbf{G}, v \rightarrow a}(\sigma | S^{I,\sigma}) \propto \prod_{b \in \partial v \setminus a} \sum_{\tau \in \Omega^{\partial b}} \mathbf{1}\{\tau_x = \sigma\} \psi_b(\tau) \prod_{y \in \partial} \mu_{\mathbf{G}, y \rightarrow b}(\tau_y | S^{I,\sigma}) + o(1). \quad (4.2)$$

This last formula expresses the notion that once we remove x and its adjacent constraints ∂x , the spins assigned to the variables y at distance two from x in \mathbf{G} are (essentially) stochastically independent. This is very much in line with the absence of long-range correlations: in $\mathbf{G} - (x \cup \partial x)$, the variables y likely are far apart from one another as \mathbf{G} contains only few short cycles. (For a more detailed discussion of the intuition behind Belief Propagation we refer to [48, Chapter 14].)

Yet (4.2) does not follow from (4.1) directly. Indeed, (4.1) merely states that the joint two randomly chosen variables of \mathbf{G} are likely approximately independent. But being second neighbors of x , the y variables in (4.2) are anything but a uniformly random family of variable nodes. For random factor graph models of Erdős-Rényi type, this difficulty is easily overcome [25] because the random factor graph obtained by removing $x \cup \partial x$ has essentially the same distribution as the original model (of order $n - 1$). In effect, the original factor graph is distributed nearly the same as the factor graph of order $n - 1$ plus a new variable plus adjacent constraints connected to a uniformly random family of variable nodes, and (4.1) applies to these random attachments. Of course, this trick does not work on the Bethe lattice, where the variables y stand out as they have degree less than d . To cope with this difficulty, we will consider an auxiliary model in which a random number of variable nodes along with their neighborhoods are removed. Moreover, we will apply Lemma 3.15 not merely to the entire original factor graph \mathbf{G} but also to the variable nodes that have degree $d - 1$ after the removal operation, which will enable us to establish the Belief Propagation equations on the reduced factor graph. Finally, we will use the decorrelation property (4.1) together with the tools from Section 3.3 to stitch the factor graph back up, i.e., to get back to the original model with no variables and constraints removed. Let us now carry out this strategy in detail.

4.1. The construction. We aim to show that the Boltzmann distribution $\mu_{\mathbf{G}}$ is well approximated by a collection of no more than L Belief Propagation fixed points for an arbitrarily slowly growing $L = L(n)$. As (4.2) shows, for a given variable node v the corresponding fixed point equations involve the messages sent by the constraint $a \in \partial v$, which in turn are determined by the messages sent out by the variables w at distance precisely two from v . Thus, to express a single application of the Belief Propagation operator we require information about the variable nodes at distance two from v . Therefore, in addition to the Boltzmann distribution $\mu_{\mathbf{G}}$ we will consider an enhanced measure $\hat{\mu}_{\mathbf{G}}$ that captures the joint distribution of the second neighborhoods.

To be precise, let $G = (V, F, (\partial a)_{a \in F}, (\psi_a)_{a \in F}, p^{\otimes n})$ be a factor graph. Then its Boltzmann distribution $\mu_{\mathbf{G}}$ ‘lives’ on the space $\Omega_G = \Omega^V$. In addition, recalling that $\nabla_G^2 v$ consists of all variable nodes at distance at most two from v , consider the space

$$\hat{\Omega}_G = \prod_{v \in V} \Omega^{\nabla_G^2 v}$$

of second neighborhood assignments, whose elements we denote as $\tau = (\tau(v, w))_{v \in V, w \in \nabla_G^2 v}$. The factor graph G induces an embedding

$$\Omega_G \rightarrow \hat{\Omega}_G, \quad \sigma \mapsto \hat{\sigma} = (\hat{\sigma}(v, w))_{v \in V, w \in \nabla_G^2 v}, \quad \text{where } \hat{\sigma}(x, y) = \sigma(y).$$

Thus, $\mu_{\mathbf{G}}$ induces a probability distribution $\hat{\mu}_{\mathbf{G}}$ on $\hat{\Omega}_G$. For a variable v we denote by $\hat{\mu}_{\mathbf{G}, v} \in \mathcal{P}(\Omega^{\nabla_G^2 v})$ the marginal distribution of $\hat{\mu}_{\mathbf{G}}$ on the v -factor of $\hat{\Omega}_G$.

The enhanced measure $\hat{\mu}_G$ will play a vital role in the construction of the Bethe state decomposition. Indeed, by comparison to the Erdős-Rényi case, the rigid geometry of the random regular graph causes significant difficulties. More precisely, while the Belief Propagation messages are defined in terms of removing one or a few constraints, such operations clearly destroy regularity. Hence, we need to create a bit of wiggling room. To this end, we remove some variable nodes along with their adjacent constraint nodes, thereby leaving a few variable nodes with degree $d - 1$ rather than d . We refer to these variables as ‘cavities’. Clearly, this operation loses some information and would therefore by itself not suffice to prove Theorem 2.5. However, what saves the day is that the enhanced measure $\hat{\mu}_G$ contains the extra information needed to stitch the graph back up without losing track of the Bethe decomposition.

Unsurprisingly, the construction is subtle and involves several steps. It requires a number of carefully chosen parameters. Specifically, given a slowly diverging monotonically increasing positive integer sequence $L = L(n) \rightarrow \infty$ as in Theorem 2.5, we choose a sequence $0 < \xi = \xi(L) = o(1)$ that tends to zero monotonically sufficiently slowly, a further sequence $\omega = \omega(\xi) \rightarrow \infty$ that tends to infinity monotonically sufficiently slowly, as well as sequences $0 < \vartheta = \vartheta(\omega) = o(1)$, $0 < \zeta = \zeta(\vartheta) = o(1)$, $0 < \beta = \beta(\vartheta) = o(1)$, $0 < \alpha = \alpha(\beta) = o(1)$, $0 < \eta = \eta(\alpha) = o(1)$ and $0 < \varepsilon = \varepsilon(\eta) = o(1)$ that tend monotonically to zero slowly enough. In summary, the pecking order reads

$$1 \ll 1/\varepsilon \ll 1/\eta \ll 1/\alpha \ll 1/\beta \ll 1/\zeta \ll 1/\vartheta \ll \omega \ll 1/\xi \ll L \ll \log \log n, \quad (4.3)$$

and we always assume tacitly that $n > n_0$ is sufficiently large.

We are ready to begin the construction. Let \mathbf{G}_* be the random factor graph obtained from \mathbf{G} as follows. Let $\boldsymbol{\theta}_*$ be a copy of the random variable $\boldsymbol{\theta}_\xi$ promised by Lemma 3.15; $\boldsymbol{\theta}_*$ is independent of \mathbf{G} . Further, let \mathbf{U}_* be a random set of $\boldsymbol{\theta}_*$ variable nodes of \mathbf{G} and draw $\boldsymbol{\sigma}_*$ from μ_G independently of $\boldsymbol{\theta}_*$ and \mathbf{U}_* . Now, obtain \mathbf{G}_* from \mathbf{G} by changing the prior distribution to

$$p_{\mathbf{G}_*}(\boldsymbol{\sigma}) = \prod_{u \in \mathbf{U}_* \cup \partial^2 \mathbf{U}_*} \mathbf{1}\{\boldsymbol{\sigma}_u = \boldsymbol{\sigma}_{*u}\} \prod_{v \notin \mathbf{U}_* \cup \partial^2 \mathbf{U}_*} p(\sigma_v). \quad (4.4)$$

Additionally, let $\boldsymbol{\omega}$ be a random variable with distribution $\text{Po}(\omega) \wedge 2\omega$, independent of everything else, and let $\mathbf{W} = \{v_{n-\omega+1}, \dots, v_n\}$. Finally, obtain \mathbf{G}'_* from \mathbf{G}_* by removing the variable nodes in \mathbf{W} along with their adjacent constraint nodes.

Thus, in \mathbf{G}'_* we pin the spins of the variable nodes in \mathbf{U}_* and their neighbors to the values observed under $\boldsymbol{\sigma}_*$, which is drawn from μ_G . Additionally, we create cavities by removing the last ω variable nodes along with their adjacent constraints. The following lemma shows that the removal of the variable nodes in \mathbf{W} does not shift the marginals of the enhanced Boltzmann distribution much.

Lemma 4.1. *With probability at least $1 - \omega^{-10}$ over the choice of $\boldsymbol{\theta}_*$, $\boldsymbol{\sigma}_*$ and \mathbf{G} the following statements are true.*

- (i) *both $\mu_{\mathbf{G}_*}$ and $\hat{\mu}_{\mathbf{G}_*}$ are $\xi^{1/4}$ -extremal.*
- (ii) *we have $\sum_{v \in V_n \setminus (\mathbf{W} \cup \partial^2 \mathbf{W})} \|\hat{\mu}_{\mathbf{G}_*,v} - \hat{\mu}_{\mathbf{G}'_*,v}\|_{\text{TV}} < \vartheta n$.*

Proof. By construction, $\hat{\mu}_{\mathbf{G}_*}$ is identical to the measure obtained through the pinning procedure of Lemma 3.15 applied to the \mathbf{U}_* -components of the space $\hat{\Omega}_{\mathbf{G}}$. Hence, Lemma 3.15 and Proposition 3.11 imply that $\hat{\mu}_{\mathbf{G}_*}$ is $\xi^{1/4}$ -extremal with probability at least $1 - \xi^{1/4}$. Since $\mu_{\mathbf{G}_*}$ is a projection of $\hat{\mu}_{\mathbf{G}_*}$, we obtain (i).

Further, let $\mathcal{V} = V_n \setminus (\mathbf{W} \cup \partial^2 \mathbf{W})$. If $\hat{\mu}_{\mathbf{G}_*}$ is $\xi^{1/4}$ -extremal, then by the definition of the cut metric the distribution $\hat{\mu}_{\mathbf{G}_*,\mathcal{V}}$ induced on the neighborhoods of \mathcal{V} is $2\xi^{1/4}$ -extremal, because $|\mathcal{V}| \geq n/2$. Additionally, there is $C = C(\omega)$ such that $\hat{\mu}_{\mathbf{G}_*,\mathcal{V}}$ is C -contiguous with respect to $\hat{\mu}_{\mathbf{G}'_*,\mathcal{V}}$ with probability at least $1 - \omega^{-11}$. This follows from (2.1), because \mathbf{G}'_* is obtained from \mathbf{G}_* by removing no more than $d\omega$ constraint nodes. Therefore, (ii) follows from (i) and Lemma 3.17, provided that ξ, ω, ϑ are chosen appropriately in accordance with (4.3). \square

The following proposition, which establishes the Belief Propagation equations on \mathbf{G}'_* , constitutes the main technical step of the proof.

Proposition 4.2. *With probability at least $1 - \alpha^9$, \mathbf{G}'_* enjoys the following properties.*

- (i) *the standard messages $(\mu_{\mathbf{G}'_*,v \rightarrow a}, \mu_{\mathbf{G}'_*,a \rightarrow v})_{v \in V(\mathbf{G}'_*) \setminus \mathbf{W}, a \in \partial v}$ form an α^{9d} -Belief Propagation fixed point.*
- (ii) *we have*

$$\sum_{v \in V(\mathbf{G}'_*)} \sum_{\sigma \in \Omega^{\nabla^2 v}} \left| \mu_{\mathbf{G}'_*}(\sigma) - \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a} \mu_{\mathbf{G}'_*,w \rightarrow a}(\sigma_w)}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}: \tau_v = \chi} \psi_a(\tau) \prod_{w \in \partial a} \mu_{\mathbf{G}'_*,w \rightarrow a}(\tau_w)} \right| < \alpha^{9d} n.$$

Before we prove Proposition 4.2 in Section 4.2, let us indicate how the theorem follows. As a final preparation we need the following basic fact.

Lemma 4.3. *For any factor graph G , for any variable node v , any $S \subset \partial v$ and any $\sigma \in \Omega$ we have*

$$\mu_{G-S,v}(\sigma) = \frac{\langle \mathbf{1}\{\sigma_v = \sigma\} / \prod_{a \in S} \psi_a(\sigma), \mu_{G,v \cup \partial^2 v} \rangle}{\langle 1 / \prod_{a \in S} \psi_a(\sigma), \mu_{G,v \cup \partial^2 v} \rangle}.$$

Proof. The partition function works out to be

$$\begin{aligned} Z(G-S) &= \sum_{\sigma \in \Omega^{V(G)}} \prod_{a \in F(G) \setminus S} \psi_a(\sigma_{\partial a}) \prod_{v \in V(G)} p_{G,v}(\sigma_v) \\ &= \sum_{\sigma \in \Omega^{V(G)}} \frac{\prod_{a \in F(G)} \psi_a(\sigma_{\partial a}) \prod_{v \in V(G)} p_{G,v}(\sigma_v)}{\prod_{a \in S} \psi_a(\sigma_{\partial a})} = Z(G) \left\langle 1 / \prod_{a \in S} \psi_a, \mu_G \right\rangle. \end{aligned}$$

Hence, for any $\tau \in \Omega^{V(G)}$,

$$\begin{aligned} \mu_{G-S}(\tau) &= \frac{1}{Z(G-S)} \prod_{a \in F(G) \setminus S} \psi_a(\tau_{\partial a}) \prod_{v \in V(G)} p_{G,v}(\sigma_v) \\ &= \frac{Z(G)}{Z(G-S)} \frac{1}{Z(G)} \frac{\prod_{a \in F(G)} \psi_a(\tau_{\partial a}) \prod_{v \in V(G)} p_{G,v}(\sigma_v)}{\prod_{a \in S} \psi_a(\tau_{\partial a})} = \frac{\mu_G(\tau)}{\langle 1 / \prod_{a \in S} \psi_a, \mu_G \rangle \prod_{a \in S} \psi_a(\tau_{\partial a})}, \end{aligned}$$

and the average in the denominator involves variables in $v \cup \partial^2 v$ only. \square

Proof of Theorem 2.5. For any assignment $\chi \in \mathcal{X} = \prod_{v \in \mathbf{U}_*} \Omega^{\nabla^2 v}$ of the variables in \mathbf{U}_* and their neighborhoods let

$$S(\chi) = \{\sigma \in \hat{\Omega}_G : \forall v \in \mathbf{U}_*, w \in \nabla^2 v : \sigma(v, w) = \chi(v, w)\}.$$

Then $(S(\chi))_{\chi \in \mathcal{X}}$ is a decomposition of $\hat{\Omega}_G$ into no more than $q^{(1+d(k-1))\theta_*}$ sub-cubes, corresponding to the neighborhood assignments of the θ_* variable nodes in \mathbf{U}_* . As Lemma 3.15 shows, by choosing the functions from (4.3) appropriately we can guarantee that $q^{(1+d(k-1))\theta_*} \leq L$. We are going to show that the decomposition $(S(\chi))_\chi$ meets the requirements of the theorem w.h.p.

For $\chi \in \mathcal{X}$ let $\mathbf{G}_*[\chi]$ be the random factor graph \mathbf{G}_* given that $\sigma_*(w) = \chi(v, w)$ for all $v \in \mathbf{U}_*$ and all $w \in \nabla^2 v$. Also let $\mathbf{G}'_*[\chi]$ be the factor graph obtained from $\mathbf{G}_*[\chi]$ by removing the variables in \mathbf{W} along with their adjacent constraint nodes. Further, let \mathcal{E}_χ be the event that the following four conditions are satisfied.

E1: Both $\mu_{\mathbf{G}_*[\chi]}$ and $\hat{\mu}_{\mathbf{G}_*[\chi]}$ are $\xi^{1/8}$ -extremal.

E2: We have

$$\sum_{v \in V_n \setminus (\mathbf{W} \cup \partial_2 \mathbf{W})} \left\| \hat{\mu}_{\mathbf{G}_*[\chi],v} - \hat{\mu}_{\mathbf{G}'_*[\chi],v} \right\|_{\text{TV}} < \vartheta n. \quad (4.5)$$

E3: On $\mathbf{G}'_*[\chi]$ the standard messages form an α^{9d} -BP fixed point and

$$\sum_{v \notin \mathbf{W} \cup \partial_2 \mathbf{W}} \sum_{\sigma \in \Omega^{\nabla^2 v}} \left| \mu_{\mathbf{G}'_*[\chi]}(\sigma) - \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a} \mu_{\mathbf{G}'_*[\chi],w \rightarrow a}(\sigma_w)}{\sum_{s \in \Omega} p(s) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}: \tau_v = s} \psi_a(\tau) \prod_{w \in \partial a} \mu_{\mathbf{G}'_*[\chi],w \rightarrow a}(\tau_w)} \right| < \alpha^{9d} n. \quad (4.6)$$

E4: There are no more than $\alpha^{10d} n$ constraint nodes a in \mathbf{G} such that $\min_{\sigma \in \Omega^k} \psi_a(\sigma) \leq \alpha^{1/4}$, nor are there more than $\varepsilon^{20} n$ constraint nodes a such that $\min_{\sigma \in \Omega^k} \psi_a(\sigma) \leq \varepsilon$.

Then (2.1), Lemma 4.1 and Proposition 4.2 yield $\mathbb{E}[\sum_\chi \mu_{\mathbf{G}}(\chi) (1 - \mathbf{1}_{\mathcal{E}_\chi}) | \mathbf{E4}] \leq \alpha^8$. Thus, Markov's inequality shows

$$\mathbb{P} \left[\sum_\chi \mu_{\mathbf{G}}(\chi) (1 - \mathbf{1}_{\mathcal{E}_\chi}) \geq \alpha^4 \right] \leq \varepsilon^4. \quad (4.7)$$

Hence, we are left to argue that $S(\chi)$ is an ε -Bethe state of \mathbf{G} if the event \mathcal{E}_χ occurs. As a first step, we are going to show that the standard messages of $\mathbf{G}'_*[\chi]$, $\mathbf{G}_*[\chi]$ are close: given \mathcal{E}_χ , we claim

$$\sum_{\sigma \in \Omega} \sum_{v \notin \mathbf{W} \cup \partial_2 \mathbf{W}} \sum_{a \in \partial v} \left| \mu_{\mathbf{G}_*[\chi],v \rightarrow a}(\sigma) - \mu_{\mathbf{G}'_*[\chi],v \rightarrow a}(\sigma) \right| + \left| \mu_{\mathbf{G}_*[\chi],a \rightarrow v}(\sigma) - \mu_{\mathbf{G}'_*[\chi],a \rightarrow v}(\sigma) \right| < \alpha^{8d} n. \quad (4.8)$$

To see this, recall that $\mu_{G_*[\chi], v \rightarrow a}$ is the marginal of v in the factor graph $\mathbf{G}_*[\chi] - a$ and that $\mu_{G_*[\chi], v \rightarrow a}$ is the marginal of v in the factor graph obtained from $\mathbf{G}_*[\chi]$ by removing all $b \in \partial v \setminus a$ and disregarding the prior. Hence, Lemma 4.3 shows that

$$\mu_{G_*[\chi], v \rightarrow a}(\sigma) = \frac{\langle \mathbf{1}\{\sigma_v = \sigma\} / \psi_a(\sigma), \mu_{G_*[\chi], \nabla^2 v} \rangle}{\langle 1 / \psi_a(\sigma), \mu_{G_*[\chi], \nabla^2 v} \rangle}, \quad (4.9)$$

$$\mu_{G_*[\chi], a \rightarrow v}(\sigma) = \frac{\langle \mathbf{1}\{\sigma_v = \sigma\} / (p(\sigma) \prod_{b \in \partial v \setminus a} \psi_b(\sigma)), \mu_{G_*[\chi], \nabla^2 v} \rangle}{\langle 1 / (p(\sigma) \prod_{b \in \partial v \setminus a} \psi_b(\sigma)), \mu_{G_*[\chi], \nabla^2 v} \rangle}, \quad (4.10)$$

and analogously for $\mathbf{G}'_*[\chi]$. Providing $\vartheta \ll \alpha^d$, we obtain (4.8) from (4.5), (4.9), (4.10) and **E2**, **E4**. Further, the estimate (4.8) and the fact that the standard messages of $\mathbf{G}'_*[\chi]$ are an α^{9d} -BP fixed point imply that the standard messages of $\mathbf{G}_*[\chi]$ are an η -BP fixed point. Thus, we have established **BS1**.

In order to prove **BS2**, we estimate the derivatives of a term like in (4.6) as follows:

$$\left| \frac{\partial}{\partial v_w(\sigma)} \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a} v_w(\sigma(w))}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a} v_w(\tau(w))} \right| \leq \frac{1}{\min_{a \in \partial v, \tau \in \Omega^{\partial a}} \psi_a(\tau)^{2d}}.$$

Hence, (4.5), (4.6), (4.8) and **E4** and the bound $|\mathbf{W} \cup \partial^2 \mathbf{W}| = O(\log n)$ yield

$$\sum_{v \in V_n} \sum_{\sigma \in \Omega^{\nabla^2 v}} \left| \mu_{G_*[\chi]}(\sigma) - \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a} \mu_{G_*[\chi], w \rightarrow a}(\sigma(w))}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a} \mu_{G_*[\chi], w \rightarrow a}(\tau(w))} \right| < \alpha^{2d} n. \quad (4.11)$$

Additionally, we claim that

$$\sum_{b \in F_m} \sum_{\sigma \in \Omega^{\partial b}} \left| \mu_{G_*[\chi]}(\sigma) - \frac{\psi_b(\sigma) \prod_{w \in \partial b} \mu_{G_*[\chi], w \rightarrow b}(\sigma(w))}{\sum_{\tau \in \Omega^{\partial b}} \psi_b(\tau) \prod_{w \in \partial b} \mu_{G_*[\chi], w \rightarrow b}(\tau(w))} \right| < \varepsilon^9 n. \quad (4.12)$$

To see this, suppose that b satisfies $\min_{\sigma \in \Omega^k} \psi_b(\sigma) \geq \varepsilon$, that $b \in \partial v$ for a variable node v such that

$$\sum_{\sigma \in \Omega^{\nabla^2 v}} \left| \mu_{G_*[\chi]}(\sigma) - \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a \setminus v} \mu_{G_*[\chi], w \rightarrow a}(\sigma_w)}{\sum_{\kappa \in \Omega} p(\kappa) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \mathbf{1}\{\tau_v = \kappa\} \prod_{w \in \partial a \setminus v} \mu_{G_*[\chi], w \rightarrow a}(\tau_w)} \right| < \alpha^d \quad (4.13)$$

and that

$$\sum_{\sigma \in \Omega} \left| \mu_{G_*[\chi], v \rightarrow b}(\sigma) - \frac{p(\sigma) \prod_{a \in \partial v \setminus b} \mu_{G_*[\chi], a \rightarrow v}(\sigma)}{\sum_{\kappa \in \Omega} p(\kappa) \prod_{a \in \partial v \setminus b} \mu_{G_*[\chi], a \rightarrow v}(\kappa)} \right| < \eta^{1/4}, \quad (4.14)$$

$$\sum_{a \in \partial v} \sum_{\sigma \in \Omega} \left| \mu_{G_*[\chi], a \rightarrow v}(\sigma) - \frac{\sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\tau_v = \sigma\} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{G_*[\chi], w \rightarrow a}(\tau_w)}{\sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{G_*[\chi], w \rightarrow a}(\tau_w)} \right| < \eta^{1/4}. \quad (4.15)$$

All but $\varepsilon^{10} n$ constraint nodes b enjoy these properties, due to **E4**, (4.11) and because the standard messages of $\mathbf{G}_*[\chi]$ form an η -BP fixed point. For any such b and any $\sigma \in \Omega^{\partial b}$ we obtain

$$\begin{aligned} \mu_{G_*[\chi], \partial b}(\sigma) &= \sum_{\tau \in \Omega^{\nabla^2 v}} \mathbf{1}\{\tau_{\partial b} = \sigma\} \mu_{G_*[\chi], \nabla^2 v}(\tau) \\ &\stackrel{(4.13)}{=} \sum_{\tau \in \Omega^{\nabla^2 v}} \frac{\mathbf{1}\{\tau_{\partial b} = \sigma\} p(\sigma_v) \prod_{a \in \partial v} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{G_*[\chi], w \rightarrow a}(\tau_w)}{\sum_{\kappa \in \Omega} \prod_{a \in \partial v} \sum_{\tau' \in \Omega^{\partial a}} \psi_a(\tau') \mathbf{1}\{\tau'_v = \kappa\} \prod_{w \in \partial a \setminus v} \mu_{G_*[\chi], w \rightarrow a}(\tau'_w)} + O(\alpha^d) \\ &\stackrel{(4.15)}{=} \frac{p(\sigma_v) \psi_b(\sigma) \prod_{w \in \partial b \setminus v} \mu_{G_*[\chi], w \rightarrow b}(\sigma_w) \prod_{a \in \partial v \setminus b} \mu_{G_*[\chi], a \rightarrow v}(\sigma_v) \sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{G_*[\chi], w \rightarrow a}(\tau_w)}{\sum_{\kappa \in \Omega^{\partial b}} p(\kappa) \psi_b(\kappa) \prod_{w \in \partial b \setminus v} \mu_{G_*[\chi], w \rightarrow b}(\kappa_w) \prod_{a \in \partial v \setminus b} \mu_{G_*[\chi], a \rightarrow v}(\kappa) \sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{G_*[\chi], w \rightarrow a}(\tau_w)} \\ &\quad + O(\alpha^d) \\ &= \frac{p(\sigma_v) \psi_b(\sigma) \prod_{w \in \partial b \setminus v} \mu_{G_*[\chi], w \rightarrow b}(\sigma_w) \prod_{a \in \partial v \setminus b} \mu_{G_*[\chi], a \rightarrow v}(\sigma_v)}{\sum_{\kappa \in \Omega^{\partial b}} p(\kappa) \psi_b(\kappa) \prod_{w \in \partial b \setminus v} \mu_{G_*[\chi], w \rightarrow b}(\kappa_w) \prod_{a \in \partial v \setminus b} \mu_{G_*[\chi], a \rightarrow v}(\kappa)} + O(\alpha^d) \\ &\stackrel{(4.14)}{=} \frac{\psi_b(\sigma) \prod_{w \in \partial b} \mu_{G_*[\chi], w \rightarrow b}(\sigma_w)}{\sum_{\kappa \in \Omega^{\partial b}} \psi_b(\kappa) \prod_{w \in \partial b} \mu_{G_*[\chi], w \rightarrow b}(\kappa_w)} + O(\alpha^d), \end{aligned}$$

whence (4.12) follows by averaging on b .

Finally, **BS2** follows from (4.11), (4.12), the $\xi^{1/8}$ -extremality of $\hat{\mu}_{G_*[\chi]}$. Indeed, let \mathbf{I}, \mathbf{J} be random sets of at most $1/\varepsilon$ variable/constraint nodes. For each $b \in \mathbf{J}$ pick a variable node $v_b \in \partial b$. Because $\hat{\mu}_{G_*[\chi]}$ is $\xi^{1/8}$ -extremal, Proposition 3.11 yields

$$\mathbb{E} \left\| \hat{\mu}_{G_*[\chi], \mathbf{I} \cup \{v_b: b \in \mathbf{J}\}} - \bigotimes_{v \in \mathbf{I} \cup \{v_b: b \in \mathbf{J}\}} \hat{\mu}_{G_*[\chi], v} \right\|_{\text{TV}} < \alpha^4. \quad (4.16)$$

Furthermore, (4.11) and (4.12) imply that with probability at least $1 - \varepsilon^2$ over the choice of \mathbf{I}, \mathbf{J} we have

$$\begin{aligned} \forall b \in \mathbf{J}: \sum_{\sigma \in \Omega^{\partial b}} \left| \mu_{G_*[\chi]}(\sigma) - \frac{\psi_b(\sigma) \prod_{w \in \partial b} \mu_{G_*[\chi], w \rightarrow b}(\sigma(w))}{\sum_{\tau \in \Omega^{\partial b}} \psi_b(\tau) \prod_{w \in \partial b} \mu_{G_*[\chi], w \rightarrow b}(\tau(w))} \right| &< \varepsilon^3, \\ \forall v \in \mathbf{I}: \sum_{\sigma \in \Omega^{\nabla^2 v}} \left| \mu_{G_*[\chi]}(\sigma) - \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a} \mu_{G_*[\chi], w \rightarrow a}(\sigma(w))}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a} \mu_{G_*[\chi], w \rightarrow a}(\tau(w))} \right| &< \varepsilon^3. \end{aligned}$$

If these estimates hold, then for any configuration $\sigma \in \Omega^{\mathbf{I} \cup \partial \mathbf{J} \cup \partial^2 \mathbf{I}}$ we obtain

$$\begin{aligned} & \left| \bigotimes_{v \in \mathbf{I} \cup \{v_b: b \in \mathbf{J}\}} \hat{\mu}_{G_*[\chi], v}(\sigma) \right. \\ & \left. - \prod_{v \in \mathbf{I}} \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a \setminus v} \mu_{G_*[\chi], w \rightarrow a}(\sigma_w)}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{G_*[\chi], w \rightarrow a}(\tau_w)} \cdot \prod_{a \in \mathbf{J}} \frac{\psi_a(\sigma) \prod_{w \in \partial a} \mu_{G_*[\chi], w \rightarrow a}(\sigma_w)}{\sum_{\tau \in \Omega^{\partial a}} \psi_a(\tau) \prod_{w \in \partial a} \mu_{G_*[\chi], w \rightarrow a}(\tau_w)} \right| &< \frac{\varepsilon}{2}. \end{aligned} \quad (4.17)$$

Thus, **BS2** follows from (4.16) and (4.17).

Finally, to obtain the Bethe state decomposition of the simple factor graph \mathbb{G} , we merely recall that $\mathbb{P}[\mathbf{G} \in \mathcal{S}] = \Omega(1)$ by Fact 2.2. Hence, the claim about \mathbb{G} follows immediately from the statement for \mathbf{G} and Bayes' rule. \square

4.2. Proof of Proposition 4.2. By construction, the random factor graph \mathbf{G}'_* comprises a pairing of variable clones $(v_i, h) \in V_n \times [d]$ and constraint clones $(a_j, h) \in F_m \times [d]$. But since we obtained \mathbf{G}'_* from \mathbf{G}_* by removing some variable nodes \mathbf{W} along with their adjacent constraint nodes, not all of the variable clones (v_i, h) with $i \leq n - \omega$ are paired. We call variables with at least one unpaired clone *cavities*. Let \mathcal{C} be the set of all cavities.

The basic idea behind the proof is as follows. We will add a new variable node v^+ along with new adjacent constraint nodes b_1, \dots, b_d to \mathbf{G}'_* . Apart from v^+ , these new constraint nodes are adjacent to some of the cavities. The fresh randomness afforded by this construction will facilitate the study of the standard messages from v^+ to the b_i as well as the reverse messages. Then we will argue that v^+ is essentially indistinguishable from a randomly chosen variable node of \mathbf{G}'_* , thereby extending the analysis to almost all the messages of \mathbf{G}'_* .

Formally, since ω is a Poisson variable with mean ω truncated at 2ω , w.h.p. we have $\omega/2 \leq |\mathcal{C}| \leq 2d(k-1)\omega$. Given that $|\mathcal{C}| \geq d(k-1)$, obtain \mathbf{G}'_* from \mathbf{G}_* by re-inserting one variable node $v_+ = v_{n-\omega+1}$ along with d new constraint nodes b_1, \dots, b_d . For each of these constraint nodes a random clone (b_i, h_i) , $h_i \in [k]$, is paired with a random clone v_+ . In addition, the b_i are paired randomly to $k-1$ cavities. The weight functions ψ_{b_i} are chosen independently from P . The following lemma shows that the distributions of \mathbf{G}'_* and \mathbf{G}_* are reasonably close.

Lemma 4.4. *For any event \mathcal{E} we have $\mathbb{P}[\mathbf{G}'_* \in \mathcal{E}] \leq \alpha^{-1} \mathbb{P}[\mathbf{G}_* \in \mathcal{E}] + O(\alpha^{100})$.*

Proof. We need to get a grip on the conditional distribution of the second neighborhood of v_+ in \mathbf{G} given \mathbf{G}'_* . This is non-trivial because of the revised prior of \mathbf{G}_* introduced by the pinning operation (4.4); for the assignment σ_* is correlated with the neighborhood of v_+ in \mathbf{G} . To begin, let \mathcal{A} be the event that no constraint node of \mathbf{G} is connected by two edges with the variable nodes $v_{n-\omega+1}, \dots, v_n$ and that all cavities have degree precisely $d-1$. Then (4.3) guarantees that

$$\mathbb{P}[\mathbf{G} \in \mathcal{A}] = 1 - O(\omega^2/n) = 1 - O(n^{-1/2}). \quad (4.18)$$

Further, given \mathcal{A} the total number of cavities of \mathbf{G}'_* is equal to $d(k-1)\omega$, and thus

$$\mathbb{P}[|\mathcal{C}| \geq \omega/2 \mid \mathcal{A}] = 1 - O(\omega^{-1}). \quad (4.19)$$

Let \mathcal{A}' be the event that \mathcal{A} occurs, that $|\mathcal{C}| \geq \omega/2$ and that the weight functions of all constraints adjacent to v_+ take a minimum value of at least $\alpha^{-1/(2dk)}$. We condition on the event \mathcal{A}' , which occurs with probability $1 + O(\alpha^{100})$ due to (2.1), (4.18) and (4.19).

Let $\mathcal{N}_1, \mathcal{N}_2$ be two possible outcomes of the depth-two neighborhoods of v_+ in \mathbf{G} given \mathbf{G}'_* . Thus, $\mathcal{N}_1, \mathcal{N}_2$ specify the weight functions of the d constraints adjacent to v_+ , the pairing of the clones of v_+ to those of these constraint

nodes, and the pairing of these d constraint nodes and the cavities \mathcal{C} . In addition, let \mathbf{G}' be the random factor graph obtained from \mathbf{G}'_* by restoring the prior to $p^{\otimes n}$. Then we can set up a coupling (Γ_1, Γ_2) of \mathbf{G} given $\mathbf{G}', \mathcal{N}_1, \mathcal{A}'$ and of \mathbf{G} given $\mathbf{G}', \mathcal{N}_2, \mathcal{A}'$ such that under Γ the two random factor graphs differ in no more than $2dk$ edges: the coupling simply switches the pairings occurring in \mathcal{N}_1 but not in \mathcal{N}_2 , and vice versa. In effect, on \mathcal{A}' the Boltzmann distributions $\mu_{\Gamma_1}, \mu_{\Gamma_2}$ are mutually α^{-1} -contiguous. Consequently, since the priors are amended according to samples from these respective Boltzmann distribution, we conclude that for any two outcomes $\mathcal{N}_1, \mathcal{N}_2$ of the second neighborhood of v^+ and for any possible outcome g of \mathbf{G}'_* ,

$$\mathbb{P}[\mathbf{G}'_* = g \mid \mathcal{N}_2, \mathbf{G}', \mathcal{A}'] \leq \alpha^{-1} \mathbb{P}[\mathbf{G}'_* = g \mid \mathcal{N}_1, \mathbf{G}', \mathcal{A}']. \quad (4.20)$$

Combining (4.18)–(4.20), we conclude that for any possible g and for any $2\omega/3 \leq \omega_0 \leq 3\omega/2$,

$$\mathbb{P}[\mathbf{G}'_* = g \mid \boldsymbol{\omega} = \omega_0] \leq \alpha^{-1} \mathbb{P}[\mathbf{G}'_* = g \mid \boldsymbol{\omega} = \omega_0 + 1] + O(\alpha^{100}). \quad (4.21)$$

Finally, the assertion follows from (4.21) because $d_{\text{TV}}(\boldsymbol{\omega}, \boldsymbol{\omega} + 1) = O(\omega^{-1/2}) = O(\alpha^{200})$ and $\mathbb{P}[2\omega/3 \leq \omega \leq 3\omega/2] = 1 - \exp(-\Omega(\omega)) = 1 - O(\alpha^{200})$. \square

Lemma 4.4 shows that studying the messages received by and emanating from v_+ is about as good as studying the messages of a random variable node of \mathbf{G}'_* . The randomness involved in the attachment process will help, but is not yet quite sufficient to actually verify the Belief Propagation equations. Namely, we also need to make sure that the Boltzmann distribution of the cavities is extremal in order to argue that typically the joint distribution of the variables where the new constraints b_1, \dots, b_d are anchored factorizes. Unfortunately, we do not know a priori that extremality holds. Indeed, while going from \mathbf{G} to \mathbf{G}'_* renders the Boltzmann distribution $\xi^{1/4}$ -extremal (by Lemma 4.1), the cavities are too few in number to conclude that their joint distribution is extremal.

Hence, we will apply a second round of pinning. But this time we will pin the cavities directly. To be precise, recalling the random variable $\boldsymbol{\theta}_+ = \boldsymbol{\theta}_\zeta$ from Lemma 3.15, let $\mathcal{C}_+ \subset \mathcal{C}$ be a random subset of size $\boldsymbol{\theta}_+ \wedge |\mathcal{C}|$. Further, draw a sample $\boldsymbol{\sigma}_+$ from $\mu_{\mathbf{G}'_*}$. The choice of $\mathcal{C}_+, \boldsymbol{\sigma}_+$ is independent of the choice of the constraints b_1, \dots, b_d , and $\boldsymbol{\sigma}_+$ is independent of \mathcal{C}_+ . Now, obtain \mathbf{G}''_* from \mathbf{G}'_* by changing the prior to

$$p_{\mathbf{G}''_*}(\sigma) \propto p_{\mathbf{G}'_*}(\sigma) \prod_{y \in \mathcal{C}_+} \mathbf{1}\{\sigma = \boldsymbol{\sigma}_+(y)\}. \quad (4.22)$$

Thus, we pin the cavities $y \in \mathcal{C}_+$ to the spins observed under $\boldsymbol{\sigma}_+$, which are independent of b_1, \dots, b_d .

Lemma 4.5. *The joint distribution $\mu_{\mathbf{G}''_*, \mathcal{C}}$ of the cavities is ζ -symmetric with probability at least $1 - \zeta$.*

Proof. Since $|\mathcal{C}| \geq \omega/2$ with probability $1 - \exp(-\Omega(\omega))$, the assertion follows immediately from Lemma 3.15 and the construction of \mathbf{G}''_* . \square

Additionally, obtain \mathbf{G}_*^+ from \mathbf{G}_*^- by changing the prior as per (4.22) as well, i.e.,

$$p_{\mathbf{G}_*^+}(\sigma) \propto p_{\mathbf{G}_*^-}(\sigma) \prod_{y \in \mathcal{C}_+} \mathbf{1}\{\sigma = \boldsymbol{\sigma}_+(y)\}. \quad (4.23)$$

We are ready to verify the Belief Propagation equations for v_+ on \mathbf{G}_*^+ .

Lemma 4.6. *With probability $1 - O(\alpha^{90})$ the random factor graph \mathbf{G}_*^+ has the following properties:*

$$\left| \mu_{\mathbf{G}_*^+, b_i \rightarrow v_+}(\sigma) - \frac{\sum_{\tau \in \Omega^{\partial b_i}} \mathbf{1}\{\tau_{v_+} = \sigma\} \psi_{b_i}(\tau) \prod_{w \in \partial b_i \setminus v_+} \mu_{\mathbf{G}_*^+, w \rightarrow b_i}(\tau_w)}{\sum_{\tau \in \Omega^{\partial b_i}} \psi_{b_i}(\tau) \prod_{w \in \partial b_i \setminus v_+} \mu_{\mathbf{G}_*^+, w \rightarrow b_i}(\tau_w)} \right| \leq \alpha^{70d} \quad \forall i \in [d], \sigma \in \Omega, \quad (4.24)$$

$$\left| \mu_{\mathbf{G}_*^+, v_+ \rightarrow b_i}(\sigma) - \frac{p(\sigma) \prod_{j \neq i} \sum_{\tau \in \Omega^{\partial b_j}} \mathbf{1}\{\tau_{v_+} = \sigma\} \psi_{b_j}(\tau) \prod_{w \in \partial b_j \setminus v_+} \mu_{\mathbf{G}_*^+, w \rightarrow b_j}(\tau_w)}{\sum_{\chi \in \Omega} p(\chi) \prod_{j \neq i} \sum_{\tau \in \Omega^{\partial b_j}} \psi_{b_j}(\tau) \mathbf{1}\{\tau_{v_+} = \chi\} \prod_{w \in \partial b_j \setminus v_+} \mu_{\mathbf{G}_*^+, w \rightarrow b_j}(\tau_w)} \right| \leq \alpha^{70d} \quad \forall i \in [d], \sigma \in \Omega, \quad (4.25)$$

$$\left| \mu_{\mathbf{G}_*^+}(\sigma) - \frac{p(\sigma) \prod_{i=1}^d \psi_{b_i}(\sigma) \prod_{w \in \partial b_i \setminus v_+} \mu_{\mathbf{G}_*^+, w \rightarrow b_i}(\sigma_w)}{\sum_{\chi \in \Omega} p(\chi) \prod_{i=1}^d \sum_{\tau \in \Omega^{\partial b_i}} \mathbf{1}\{\tau_{v_+} = \chi\} \psi_{b_i}(\tau) \prod_{w \in \partial b_i \setminus v_+} \mu_{\mathbf{G}_*^+, w \rightarrow b_i}(\tau_w)} \right| < \alpha^{70d} \quad \forall \sigma \in \Omega^{\nabla^2 v_+}. \quad (4.26)$$

Proof. Lemma 4.5 shows that $\mu_{G''_*, \mathcal{C}}$ is ζ -symmetric with probability at least $1 - \zeta$. Suppose it is. Then Proposition 3.11 shows that $\mu_{G''_*, \mathcal{C}}$ is $(\beta, d(k-1))$ -symmetric. We may also assume that $|\mathcal{C}| \geq \omega/2$, an event that occurs with probability at least $1 - \exp(-\Omega(\omega))$ by the construction of G'_* . Additionally, due to (2.1) we may assume that

$$\min_{\sigma \in \Omega^k} \psi_{b_i}(\sigma) \geq \alpha \quad \text{for all } i \in [d]. \quad (4.27)$$

According to (3.2), the standard message $\mu_{G'_*, b_1 \rightarrow v_+}$ is defined as the marginal of v_+ in the factor graph obtained from G'_* by removing b_2, \dots, b_d and replacing the prior of v_+ by the uniform distribution. By construction, this factor graph is obtained from G''_* by adding the variable node v_+ and constraint node b_1 and replacing the prior of v_+ by the uniform distribution. Therefore,

$$\mu_{G'_*, b_1 \rightarrow v_+}(\sigma) = \frac{\sum_{\tau \in \Omega^{\partial b_1}} \mathbf{1}\{\tau_{v_+} = \sigma\} \psi_{b_1}(\tau) \left\langle \mathbf{1}\{\forall w \in \partial b_1 \setminus v_+ : \sigma_w = \tau_w\}, \mu_{G''_*, \mathcal{C}} \right\rangle}{\sum_{\tau \in \Omega^{\partial b_1}} \psi_{b_1}(\tau) \left\langle \mathbf{1}\{\forall w \in \partial b_1 \setminus v_+ : \sigma_w = \tau_w\}, \mu_{G''_*, \mathcal{C}} \right\rangle} \quad (\sigma \in \Omega). \quad (4.28)$$

Further, the neighbors $\partial b_1 \setminus v_+$ are chosen uniformly from \mathcal{C} (without replacement). Because $|\mathcal{C}| \geq \omega/2$ and $\mu_{G''_*, \mathcal{C}}$ is $(\beta, d(k-1))$ -symmetric, we conclude that

$$\mathbb{P} \left[\left\| \mu_{G''_*, \partial b_1 \setminus v_+} - \bigotimes_{w \in \partial b_1 \setminus v_+} \mu_{G''_*, w} \right\|_{\text{TV}} \leq \beta^{1/3} \right] \geq 1 - \beta^{1/3}. \quad (4.29)$$

Combining (4.27), (4.28) and (4.29), we obtain the estimate

$$\mathbb{E} \left| \mu_{G'_*, b_1 \rightarrow v_+}(\sigma) - \frac{\sum_{\tau \in \Omega^{\partial b_1}} \mathbf{1}\{\tau_{v_+} = \sigma\} \psi_{b_1}(\tau) \prod_{w \in \partial b_1 \setminus v_+} \mu_{G''_*, w}(\tau_w)}{\sum_{\tau \in \Omega^{\partial b_1}} \psi_{b_1}(\tau) \prod_{w \in \partial b_1 \setminus v_+} \mu_{G''_*, w}(\tau_w)} \right| \leq \beta^{1/4}. \quad (4.30)$$

Moreover, the factor graph $G'_* - b_1$ is obtained from G''_* by adding v_+ and b_2, \dots, b_d . Hence, (4.27) implies that $\mu_{G'_* - b_1, \mathcal{C}}$ is $(2/\alpha)^{dk}$ -contiguous with respect to $\mu_{G''_*, \mathcal{C}}$. Since $\mu_{G''_*, \mathcal{C}}$ is ζ -symmetric, Proposition 3.11 and Lemma 3.17 yield $\Delta_{\square}(\mu_{G'_* - b_1, \mathcal{C}}, \mu_{G''_*, \mathcal{C}}) < \beta$. Because the neighborhood ∂b_1 is random, Lemma 3.12 therefore yields

$$\mathbb{E} \sum_{w \in \partial b_1 \setminus v_+} \left\| \mu_{G''_*, w} - \mu_{G'_*, w \rightarrow b_1} \right\|_{\text{TV}} \leq O(\beta). \quad (4.31)$$

Combining (4.27), (4.30) and (4.31), we obtain (4.24).

The proofs of (4.25) and (4.26) are similar. Indeed, $\mu_{G'_*, v_+ \rightarrow b_1}$ is the marginal of v_+ in $G'_* - b_1$, which is obtained from G''_* by adding b_2, \dots, b_d . Hence,

$$\mu_{G'_*, v_+ \rightarrow b_1}(\sigma) = \frac{\sum_{\tau \in \Omega^{\{v_+\} \cup \partial^2 v_+}} p(\sigma) \mathbf{1}\{\tau_{v_+} = \sigma\} \left\langle \mathbf{1}\{\forall w \in \partial^2 v_+ : \sigma_w = \tau_w\}, \mu_{G''_*, \mathcal{C}} \right\rangle \prod_{i=2}^d \psi_{b_i}(\tau)}{\sum_{\tau \in \Omega^{\{v_+\} \cup \partial^2 v_+}} p(\tau_{v_+}) \left\langle \mathbf{1}\{\forall w \in \partial^2 v_+ : \sigma_w = \tau_w\}, \mu_{G''_*, \mathcal{C}} \right\rangle \prod_{i=2}^d \psi_{b_i}(\tau)}.$$

Invoking the $(\beta, d(k-1))$ -symmetry of $\mu_{G''_*}$ and (4.27), we obtain

$$\mathbb{E} \left| \mu_{G'_*, v_+ \rightarrow b_1}(\sigma) - \frac{\sum_{\tau \in \Omega^{\{v_+\} \cup \partial^2 v_+}} p(\sigma) \mathbf{1}\{\tau_{v_+} = \sigma\} \prod_{i=2}^d \psi_{b_i}(\tau) \prod_{w \in \partial^2 v_+} \mu_{G''_*, w}(\tau_w)}{\sum_{\tau \in \Omega^{\{v_+\} \cup \partial^2 v_+}} p(\tau_{v_+}) \prod_{i=2}^d \psi_{b_i}(\tau) \prod_{w \in \partial^2 v_+} \mu_{G''_*, w}(\tau_w)} \right| \leq \beta^{1/4}.$$

Moreover, reordering the sums and products, we simplify the last expression and find

$$\mathbb{E} \left| \mu_{G'_*, v_+ \rightarrow b_1}(\sigma) - \frac{p(\sigma) \prod_{i=2}^d \sum_{\tau \in \Omega^{\partial b_i}} \mathbf{1}\{\tau_{v_+} = \sigma\} \psi_{b_i}(\tau) \prod_{w \in \partial b_i \setminus v_+} \mu_{G''_*, w}(\tau_w)}{\sum_{\chi \in \Omega} p(\chi) \prod_{i=2}^d \sum_{\tau \in \Omega^{\partial b_i}} \mathbf{1}\{\tau_{v_+} = \chi\} \psi_{b_i}(\tau) \prod_{w \in \partial b_i \setminus v_+} \mu_{G''_*, w}(\tau_w)} \right| \leq \beta^{1/4}. \quad (4.32)$$

Further, (4.27) ensures that for each $i \in [d]$ the distribution $\mu_{G'_* - b_i, \mathcal{C}}$ is $(2/\alpha)^{dk}$ -contiguous with respect to G''_* . Consequently, since the neighbors of b_i are chosen randomly from the set \mathcal{C} of cavities, Proposition 3.11 and Lemma 3.17 yield

$$\sum_{i=2}^d \mathbb{E} \left[\sum_{w \in \partial b_i \setminus v_+} \left\| \mu_{G''_*, w} - \mu_{G'_*, w \rightarrow b_i} \right\|_{\text{TV}} \right] \leq O(\beta). \quad (4.33)$$

Combining (4.32) and (4.33), we obtain (4.25).

Moving on to (4.26), we consider $\sigma \in \Omega^{\{v^+\} \cup \partial^2 v^+}$. Since \mathbf{G}_*^+ is obtained from \mathbf{G}_*'' by adding v_+ along with b_1, \dots, b_d , we have the exact formula

$$\mu_{\mathbf{G}_*^+}(\sigma) = \frac{p(\sigma) \left\langle \mathbf{1}\{\forall w \in \partial^2 v^+ : \sigma_w = \sigma_w, \mu_{\mathbf{G}_*''}\} \prod_{i=1}^d \psi_{b_i}(\sigma) \right\rangle}{\sum_{\chi \in \Omega} p(\chi) \left\langle \prod_{i=1}^d \sum_{\tau \in \Omega^{\partial b_i}} \mathbf{1}\{\tau_{v^+} = \chi\} \psi_{b_i}(\tau) \prod_{w \in \partial b_i \setminus v^+} \mathbf{1}\{\sigma_w = \tau_w\}, \mu_{\mathbf{G}_*''} \right\rangle} \quad (\sigma \in \Omega^{\nabla^2 v_+}).$$

Since $\partial^2 v^+$ is a random set of cavities, the $(\beta, d(k-1))$ -symmetry of $\mu_{\mathbf{G}_*''}$ and (4.27) ensure that

$$\mathbb{E} \left| \mu_{\mathbf{G}_*^+}(\sigma) - \frac{p(\sigma) \prod_{w \in \partial^2 v^+} \mu_{\mathbf{G}_*''(w)}(\sigma_w) \prod_{i=1}^d \psi_{b_i}(\sigma)}{\sum_{\chi \in \Omega} p(\chi) \prod_{i=1}^d \sum_{\tau \in \Omega^{\partial b_i}} \mathbf{1}\{\tau_{v^+} = \chi\} \psi_{b_i}(\tau) \prod_{w \in \partial b_i \setminus v^+} \mu_{\mathbf{G}_*''(w)}(\tau_w)} \right| \leq \beta^{1/4}. \quad (4.34)$$

Finally, to complete the proof we combine (4.33) and (4.34). \square

We set up the random factor graph \mathbf{G}_*^+ so as to facilitate the verification of the BP equations. But in a sense the model is a bit ‘out of whack’ because the prior is pinned according to a configuration σ_+ drawn from $\mu_{\mathbf{G}_*'}$ rather than $\mu_{\mathbf{G}_*^-}$; see (4.23). Thus, with $\theta_+ = \theta_\zeta$ and $\mathcal{C}_+ \subset \mathcal{C}$ as before, draw σ_{++} from $\mu_{\mathbf{G}_*^-}$ and let \mathbf{G}_*^{++} be the random factor graph obtained from \mathbf{G}_*^- by changing the prior to

$$p_{\mathbf{G}_*^{++}}(\sigma) \propto p_{\mathbf{G}_*'}(\sigma) \prod_{w \in \mathcal{C}_+} \mathbf{1}\{\sigma = \sigma_{++}(w)\}.$$

Hence, σ_{++} takes v_+, b_1, \dots, b_d into account.

Corollary 4.7. *With probability $1 - O(\alpha^{80})$ the bounds (4.24)–(4.26) hold with \mathbf{G}_*^+ replaced by \mathbf{G}_*^{++} .*

Proof. The only difference between \mathbf{G}_*^{++} and \mathbf{G}_*^+ lies in the choice of the configuration to which the variable nodes in \mathcal{C}_+ get pinned. But since \mathbf{G}_*^+ is obtained from \mathbf{G}_*'' by the mere addition of d constraint nodes b_1, \dots, b_d , (2.1) shows that σ_{++} is α^{-1} -contiguous with respect to the distribution of σ_+ with probability $1 - O(\alpha^{80})$. Thus, the assertion follows from Lemma 4.6. \square

We are finally ready to go back to the random factor graph \mathbf{G}_*'' . Indeed, basically the only difference between \mathbf{G}_*^{++} and \mathbf{G}_*'' is that the former has one more variable node, along with d adjacent constraint nodes. But since the number of variable nodes of \mathbf{G}_*'' is random, this difference should hardly be noticeable. Also \mathbf{G}_*'' is invariant under permutations of its variable nodes. Thus, whatever we can prove for the last variable node v_+ of \mathbf{G}_*^{++} carries over to a random variable node of \mathbf{G}_*'' . The following corollary makes this precise.

Corollary 4.8. *With probability $1 - O(\alpha^{70})$ we have*

$$\sum_{v \in V_n} \sum_{b \in \partial v} \sum_{\sigma \in \Omega} \left| \mu_{\mathbf{G}_*'', b \rightarrow v}(\sigma) - \frac{\sum_{\tau \in \Omega^{\partial b}} \mathbf{1}\{\tau_v = \sigma\} \psi_b(\tau) \prod_{w \in \partial b \setminus v} \mu_{\mathbf{G}_*'', w \rightarrow b}(\tau_w)}{\sum_{\tau \in \Omega^{\partial b}} \psi_b(\tau) \prod_{w \in \partial b \setminus v} \mu_{\mathbf{G}_*'', w \rightarrow b}(\tau_w)} \right| \leq n \alpha^{60d}, \quad (4.35)$$

$$\sum_{v \in V_n} \sum_{b \in \partial v} \sum_{\sigma \in \Omega} \left| \mu_{\mathbf{G}_*'', v \rightarrow b}(\sigma) - \frac{p(\sigma) \prod_{a \in \partial v \setminus b} \sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\tau_v = \sigma\} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{\mathbf{G}_*'', w \rightarrow a}(\tau_w)}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v \setminus b} \sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\tau_v = \chi\} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{\mathbf{G}_*'', w \rightarrow a}(\tau_w)} \right| \leq n \alpha^{60d} \quad (4.36)$$

$$\sum_{v \in V_n} \sum_{\sigma \in \Omega^{\nabla^2 v}} \left| \mu_{\mathbf{G}_*''}(\sigma) - \frac{p(\sigma) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a \setminus v} \mu_{\mathbf{G}_*'', w \rightarrow a}(\sigma_w)}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\tau_v = \chi\} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{\mathbf{G}_*'', w \rightarrow a}(\tau_w)} \right| < n \alpha^{60d}. \quad (4.37)$$

Proof. Consider the event \mathcal{E} that in \mathbf{G}_*'' , for the variable node $v_{n-\omega}$ with the largest index the estimate

$$\sum_{b \in \partial v_{n-\omega}} \sum_{\sigma \in \Omega} \left| \mu_{\mathbf{G}_*'', b \rightarrow v_{n-\omega}}(\sigma) - \frac{\sum_{\tau \in \Omega^{\partial b}} \mathbf{1}\{\tau(v_{n-\omega}) = \sigma\} \psi_b(\tau) \prod_{w \in \partial b \setminus v_{n-\omega}} \mu_{\mathbf{G}_*'', w \rightarrow b}(\tau(w))}{\sum_{\tau \in \Omega^{\partial b}} \psi_b(\tau) \prod_{w \in \partial b \setminus v_{n-\omega}} \mu_{\mathbf{G}_*'', w \rightarrow b}(\tau_w)} \right| \leq \alpha^{69} \quad (4.38)$$

holds. Since \mathbf{G}_*^{++} is obtained from \mathbf{G}_*^+ by the same process that produces \mathbf{G}_*'' from \mathbf{G}_*' , Lemma 4.4 and Corollary 4.7 show that $\mathbb{P}[\mathbf{G}_*'' \in \mathcal{E}] = 1 - O(\alpha^{70})$. But since the distribution of \mathbf{G}_*'' is invariant under permutations of the $n - O(\omega)$ variable nodes of degree d , we can replace $v_{n-\omega}$ in (4.38) by a random variable node of degree d . Thus, we obtain (4.35). The two bounds (4.36) and (4.37) follow analogously. \square

To complete the proof of Proposition 4.2, we finally need to get from \mathbf{G}_*'' back to \mathbf{G}_*' . Thus, we need to undo the additional pinning of the cavities that was required to verify the BP equations (4.35)–(4.37). The elegant insight that makes this possible is that (4.35)–(4.37) really just describe a property of the joint distribution of the second neighborhoods of the variable nodes v_i , $i \neq n - \omega$. Indeed, by Lemma 4.3 the standard messages, defined via the

removal of a few constraints adjacent to a single variable node v , can be expressed easily in terms of the joint distribution of the second neighborhood of v . Furthermore, Lemma 4.1 implies that the enhanced measure $\hat{\mu}_{G'_*}$ describing the second neighborhood distributions is $\xi^{1/4}$ -extremal, with ξ near the top of the pecking order (4.3). In effect, $\hat{\mu}_{G'_*}$ is impervious to the additional pinning required to go from G'_* to G''_* . Let us formalize this argument to finish the proof of Proposition 4.2.

Proof of Proposition 4.2. By Lemma 4.1 the measure $\hat{\mu}_{G'_*}$ is $\xi^{1/4}$ -extremal with probability $1 - \omega^{-1}$. Consequently, since G'_* is obtained by deleting $O(\omega)$ constraints, Lemma 3.17 and (2.1) ensure that $\hat{\mu}_{G'_*}$ is ϑ -extremal with probability $1 - \alpha^{10}$. Furthermore, $\hat{\mu}_{G''_*}$ is nothing but the conditional distribution $\hat{\mu}_{G'_*}$ given the event S_+ that the spins of the θ_ζ cavities \mathcal{C}_+ coincide with the ones of the reference configuration σ_+ . Since σ_+ is drawn from $\hat{\mu}_{G'_*}$, with probability at least $1 - \zeta$ we have

$$\hat{\mu}_{G'_*}(S_+) \geq \zeta q^{-\theta_\zeta}.$$

If so, and if $\hat{\mu}_{G'_*}$ is ϑ -extremal, then (4.3) and Corollary 3.18 imply that $\Delta_\square(\hat{\mu}_{G''_*}, \hat{\mu}_{G'_*}) \leq \beta$. In summary,

$$\mathbb{P} \left[\Delta_\square(\hat{\mu}_{G''_*}, \hat{\mu}_{G'_*}) \leq \beta \right] \geq 1 - 2\zeta.$$

In addition (2.1) ensures that with probability at least $1 - \alpha^{10}$,

$$\left| \left\{ a \in F(G'_*) : \min_{\sigma \in \Omega^k} \psi_a(\sigma) \leq \alpha \right\} \right| \leq n \alpha^{1000d}. \quad (4.39)$$

Further, by Corollary 4.8 the bounds (4.35)–(4.37) hold with probability $1 - O(\alpha^{70})$.

Thus, we are left to prove statements (i) and (ii) under the assumption that $\Delta_\square(\hat{\mu}_{G''_*}, \hat{\mu}_{G'_*}) \leq \beta$ and that (4.35)–(4.37) and (4.39) hold. Applying Lemma 3.12, we obtain

$$\sum_{v \in V} \left\| \hat{\mu}_{G''_*, v} - \hat{\mu}_{G'_*, v} \right\|_{\text{TV}} \leq O(\beta). \quad (4.40)$$

Further, Lemma 4.3 shows that the messages $\mu_{G''_*, v \rightarrow a}$, $\mu_{G''_*, a \rightarrow v}$ and $\mu_{G'_*, v \rightarrow a}$, $\mu_{G'_*, a \rightarrow v}$ can be expressed in terms of the marginal distributions $\hat{\mu}_{G''_*, v}$ and $\hat{\mu}_{G'_*, v}$ of the depth-two neighborhood. Indeed, according to (3.1)–(3.2), for any $a \in \partial v$ and $\sigma \in \Omega$,

$$\mu_{G''_*, v \rightarrow a}(\sigma) = \frac{\left\langle \mathbf{1}\{\sigma_v = \sigma\} / \psi_a(\sigma), \hat{\mu}_{G''_*, v} \right\rangle}{\left\langle 1 / \psi_a(\sigma), \hat{\mu}_{G''_*, v} \right\rangle}, \quad \mu_{G''_*, a \rightarrow v}(\sigma) = \frac{\left\langle \mathbf{1}\{\sigma_v = \sigma\} / (p(\sigma) \prod_{b \in \partial v \setminus a} \psi_b(\sigma)), \hat{\mu}_{G''_*, v} \right\rangle}{\left\langle 1 / (p(\sigma_v) \prod_{b \in \partial v \setminus a} \psi_b(\sigma)), \hat{\mu}_{G''_*, v} \right\rangle},$$

and analogously for G'_* . Hence, the total variation bound (4.40) and (4.39) imply that

$$\sum_{v \in V} \sum_{a \in \partial v} \left\| \mu_{G''_*, v \rightarrow a} - \mu_{G'_*, v \rightarrow a} \right\|_{\text{TV}} + \left\| \mu_{G''_*, a \rightarrow v} - \mu_{G'_*, a \rightarrow v} \right\|_{\text{TV}} = O(n \alpha^{900d}). \quad (4.41)$$

Combining (4.37) and (4.39)–(4.41), we obtain assertion (ii). Further, (4.35), (4.39) and (4.41) readily yield

$$\sum_{v \in V} \sum_{b \in \partial v} \sum_{\sigma \in \Omega} \left| \mu_{G'_*, b \rightarrow v}(\sigma) - \frac{\sum_{\tau \in \Omega^{\partial b}} \mathbf{1}\{\tau(v) = \sigma\} \psi_b(\tau) \prod_{w \in \partial b \setminus v} \mu_{G'_*, w \rightarrow b}(\tau(w))}{\sum_{\tau \in \Omega^{\partial b}} \psi_b(\tau) \prod_{w \in \partial b \setminus v} \mu_{G'_*, w \rightarrow b}(\tau(w))} \right| \leq O(n \alpha^{60d}). \quad (4.42)$$

Moreover, combining (4.36), (4.39), (4.40) and (4.42), we obtain

$$\sum_{v \in V} \sum_{b \in \partial v} \sum_{\sigma \in \Omega} \left| \mu_{G'_*, v \rightarrow b}(\sigma) - \frac{p(\sigma) \prod_{a \in \partial v \setminus b} \mu_{G'_*, a \rightarrow v}(\sigma)}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v \setminus b} \mu_{G'_*, a \rightarrow v}(\chi)} \right| \leq O(n \alpha^{50d}). \quad (4.43)$$

Finally, (4.42) and (4.43) show that the standard messages are a $O(\alpha^{40d})$ -BP fixed point. \square

5. THE FREE ENERGY: UPPER BOUND

5.1. Outline. In this section we derive the following upper bound on the free energy.

Proposition 5.1. *Assume that **POS** is satisfied. Then*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\log Z(G)] \leq \inf_{\pi \in \mathcal{D}^*} \mathcal{B}(\pi), \quad \limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\log Z(\mathbb{G})] \leq \inf_{\pi \in \mathcal{D}^*} \mathcal{B}(\pi).$$

The proof of Proposition 5.1 consists of two parts. First, we will prove that *any* $\mu \in \mathcal{K}$ yields an upper bound on $\mathbb{E}[\log Z(\mathbf{G})]$. Specifically, recalling the notation from Section 3.1, let

$$\mathcal{B}'(\mu) = \mathbb{E} \log \left\langle \bigoplus_{i=1}^n \boldsymbol{\varphi}_i, \mu \right\rangle, \quad \mathcal{B}''(\mu) = \mathbb{E} \log \left\langle \bigoplus_{1 \leq i \leq (k-1)dn/k} \boldsymbol{\psi}_{1,i}, \mu \right\rangle.$$

Then we have the following generic upper bound, which may be of interest in its own right.

Proposition 5.2. *Assume that **POS** is satisfied. Then $\mathbb{E}[\log Z(\mathbf{G})] \leq o(n) + \mathcal{B}'(\mu) - \mathcal{B}''(\mu)$ for any $\mu \in \mathcal{K}$.*

The proof of Proposition 5.2, based on the interpolation method, is relatively standard, although the fact that we deal with regular graphs requires a bit of care. The details are carried out in Section 5.2. This is the only place where condition **POS** is required.

The second step toward the proof of Proposition 5.1 is to show that for μ drawn from $\pi \in \mathfrak{D}^*$ the upper bound from Proposition 5.2 boils down to the expression $\mathcal{B}(\pi)$.

Proposition 5.3. *For any $\pi \in \mathfrak{D}^*$ we have $\mathcal{B}(\pi) = \mathbb{E}[\mathcal{B}'(\boldsymbol{\mu}^\pi) - \mathcal{B}''(\boldsymbol{\mu}^\pi)]$.*

We prove Proposition 5.3 in Section 5.3.

Proof of Proposition 5.1. The first assertion is immediate from Propositions 5.2 and 5.3. To obtain the second assertion, we apply Azuma's inequality and (2.1) to see that $n^{-0.51} |\log Z(\mathbf{G}) - \mathbb{E} \log Z(\mathbf{G})|$ converges to zero in probability. Hence, Fact 2.2 and Bayes' rule show that $\mathbb{E} \log Z(\mathbf{G}) = \mathbb{E} \log Z(\mathbf{G}) + o(n)$, and thus the second assertion follows from the first. \square

5.2. Proof of Proposition 5.2. We construct a family of random factor graph models parametrized by $t \in [0, 1]$. The free energy of the model at $t = 1$ will be easy to compute, and we will see that it is (nearly) equal to $\mathcal{B}'(\mu) - \mathcal{B}''(\mu)$. The model with $t = 0$ essentially coincides with \mathbf{G} . Furthermore, we will show that the derivative of the free energy is non-negative for all t , thus obtaining the desired upper bound on $\mathbb{E} \log Z(\mathbf{G})$.

To construct this interpolating family, fix $\mu \in \mathcal{K}$ and a small $\varepsilon > 0$. For $t \in [0, 1]$ let

$$\begin{aligned} \mathbf{m}_t &= \text{Po}((1-t) \exp(-\varepsilon) dn/k), \\ \mathbf{m}'_t &= \text{Po}(t \exp(-\varepsilon) dn), \\ \mathbf{m}''_t &= \text{Po}((1-t)(k-1) \exp(-\varepsilon) dn/k), \end{aligned}$$

all three mutually independent and independent of everything else. Given $k\mathbf{m}_t + \mathbf{m}'_t \leq dn$, we define the random factor graph \mathbf{G}_t as follows.

INT1: the set of variable nodes is $\mathcal{V} = \{s\} \cup V_n$, and the set of spins is $\mathcal{X} = \Omega \cup [0, 1]$.

INT2: the set of constraint nodes is

$$\mathcal{F}_t = \left\{ a_1, \dots, a_{\mathbf{m}_t}, a'_1, \dots, a'_{\mathbf{m}'_t}, a''_1, \dots, a''_{\mathbf{m}''_t} \right\}.$$

INT3: each constraint node a_i independently chooses a weight function ψ_{a_i} from P , and the a_i are joined to the variable nodes v_1, \dots, v_n by a random pairing of $V_n \times [d]$ and $\{a_1, \dots, a_{\mathbf{m}_t}\} \times [k]$.

INT4: each of the constraint nodes a'_i , $i \in [\mathbf{m}'_t]$, is adjacent to the variable node s and one further variable node from v_1, \dots, v_n ; the links between the a'_i and the v_j are constructed by choosing a random pairing between the \mathbf{h}'_i -clone of each a'_i and the clones in $V_n \times [d]$ that are not paired to a constraint node a_h . The weight function associated with a'_i reads

$$\psi_{a'_i}(s, \sigma) = \sum_{\tau \in \Omega^k} \mathbf{1}\{\tau_{\mathbf{h}'_i} = \sigma\} \boldsymbol{\psi}'_i(\tau) \prod_{h \neq \mathbf{h}'_i} \mu_{s, \mathbf{x}'_{i,h}}(\tau_h).$$

INT5: the constraint nodes a''_i , $i \in [\mathbf{m}''_t]$, are unary, adjacent to s only. Their weight functions read

$$\psi_{a''_i}(s) = \sum_{\tau \in \Omega^k} \boldsymbol{\psi}''_i(\tau) \prod_{h=1}^k \mu_{s, \mathbf{x}''_{i,h}}(\tau_h).$$

INT6: the prior \mathbf{p} is a product measure

$$d\mathbf{p}(\sigma) = \mathbf{1}\{\sigma_s \in [0, 1], \forall 1 \leq i \leq n : \sigma_{v_i} \in \Omega\} \prod_{i=1}^n p(\sigma_{x_i}) d\sigma_s;$$

thus, for each $v_i \in V_n$ a spin from Ω is chosen independently from p , and σ_s is uniform on $[0, 1]$.

Thus, the total weight, partition function and Boltzmann distribution of \mathbf{G}_t read

$$\begin{aligned} \psi_{\mathbf{G}_t}(\sigma) &= \prod_{i=1}^{\mathbf{m}_t} \psi_{a_i}(\sigma) \prod_{i=1}^{\mathbf{m}'_t} \psi_{a'_i}(\sigma) \prod_{i=1}^{\mathbf{m}''_t} \psi_{a''_i}(\sigma), & (\sigma_s \in [0, 1], \sigma_{x_i} \in \Omega), \\ Z(\mathbf{G}_t) &= \sum_{\sigma_{x_1}, \dots, \sigma_{x_n} \in \Omega} \int_0^1 \psi_{\mathbf{G}_t}(\sigma) d\sigma_s \prod_{i=1}^n p(\sigma_{x_i}), & d\mu_{\mathbf{G}_t}(\sigma) = \frac{\psi_{\mathbf{G}_t}(\sigma)}{Z(\mathbf{G}_t)} d\mathbb{P}(\sigma). \end{aligned} \quad (5.1)$$

The following lemma establishes the monotonicity of the free energy in t ; its proof is the only place where we use condition **POS**.

Lemma 5.4. *Suppose that **POS** is satisfied. Then uniformly for all $t \in (0, 1)$ we have*

$$\frac{1}{n} \frac{\partial}{\partial t} \mathbb{E} [\log Z(\mathbf{G}_t)] \geq o(1).$$

Proof. We recall the derivative of the Poisson density: for any $\lambda > 0$, $\ell \geq 1$,

$$\frac{\partial}{\partial \lambda} \mathbb{P} [\text{Po}(\lambda) = \ell] = \frac{\partial}{\partial \lambda} \frac{\lambda^\ell}{\ell!} \exp(-\lambda) = \frac{\lambda^{\ell-1}}{(\ell-1)!} \exp(-\lambda) - \frac{\lambda^\ell}{\ell!} \exp(-\lambda) = \mathbb{P} [\text{Po}(\lambda) = \ell-1] - \mathbb{P} [\text{Po}(\lambda) = \ell]. \quad (5.2)$$

The variable t affects the distribution of \mathbf{G}_t by way of the variables $\mathbf{m}_t, \mathbf{m}'_t, \mathbf{m}''_t$. Specifically, let

$$\lambda_t = (1-t) \exp(-\varepsilon) dn/k, \quad \lambda'_t = t \exp(-\varepsilon) dn, \quad \lambda''_t = (1-t)(k-1) \exp(-\varepsilon) dn/k.$$

Recall that $\mathbf{m}_t, \mathbf{m}'_t$ are conditional Poisson variables $\text{Po}(\lambda_t)$ and $\text{Po}(\lambda'_t)$, respectively, given that $k\mathbf{m}_t + \mathbf{m}'_t \leq dn$. Since $\varepsilon > 0$ is independent of n , (5.2) shows that for any two integers $m_t, m'_t \geq 1$,

$$\begin{aligned} \frac{1}{n} \frac{\partial}{\partial t} \mathbb{P} [\mathbf{m}_t = m_t, \mathbf{m}'_t = m'_t] &= \exp(-\Omega(n)) + \frac{1}{n} \frac{\partial}{\partial t} \mathbb{P} [\text{Po}(\lambda_t) = m_t] \mathbb{P} [\text{Po}(\lambda'_t) = m'_t] \\ &= \exp(-\Omega(n)) + \frac{1}{n} \frac{\partial}{\partial t} \mathbb{P} [\text{Po}(\lambda_t) = m_t] \mathbb{P} [\text{Po}(\lambda'_t) = m'_t] \\ &= \exp(-\Omega(n)) + (\mathbb{P} [\text{Po}(\lambda_t) = m_t] - \mathbb{P} [\text{Po}(\lambda_t) = m_t - 1]) \mathbb{P} [\text{Po}(\lambda'_t) = m'_t] \exp(-\varepsilon) d/k \\ &\quad + (\mathbb{P} [\text{Po}(\lambda'_t) = m'_t - 1] - \mathbb{P} [\text{Po}(\lambda'_t) = m'_t]) \cdot \mathbb{P} [\text{Po}(\lambda_t) = m_t] \exp(-\varepsilon) d. \end{aligned} \quad (5.3)$$

Further, given the event $k\mathbf{m}_t + \mathbf{m}''_t \leq dn - k$ let \mathbf{G}'_t be the random factor graph obtained from \mathbf{G}_t by adding one more constraint node $a_{\mathbf{m}_t+1}$ as per **INT3**. Similarly, given $k\mathbf{m}_t + \mathbf{m}''_t \leq dn - 1$ obtain \mathbf{G}''_t from \mathbf{G}_t by adding $a_{\mathbf{m}'_t+1}$ according to **INT4**. Additionally, obtain \mathbf{G}'''_t from \mathbf{G}_t by adding a unary $a_{\mathbf{m}''_t+1}$ as described in **INT5**. Since \mathbf{m}''_t is independent of $\mathbf{m}'_t, \mathbf{m}''_t$, (5.2) and (5.3) yield

$$\begin{aligned} \frac{1}{n} \frac{\partial}{\partial t} \mathbb{E} [\log Z(\mathbf{G}_t)] &= \exp(-\Omega(n)) + \sum_{\substack{m_t, m'_t, m''_t \geq 1 \\ km_t + m''_t \leq dn - k}} \mathbb{E} \left[\log Z(\mathbf{G}_t) \middle| \begin{pmatrix} \mathbf{m}_t \\ \mathbf{m}'_t \\ \mathbf{m}''_t \end{pmatrix} = \begin{pmatrix} m_t \\ m'_t \\ m''_t \end{pmatrix} \right] \cdot \frac{1}{n} \frac{\partial}{\partial t} \mathbb{P} \left[\begin{pmatrix} \mathbf{m}_t \\ \mathbf{m}'_t \\ \mathbf{m}''_t \end{pmatrix} = \begin{pmatrix} m_t \\ m'_t \\ m''_t \end{pmatrix} \right] \\ &= \exp(-\Omega(n)) - \exp(-\varepsilon) \frac{d}{k} \left[\mathbb{E} \log \frac{Z(\mathbf{G}'_t)}{Z(\mathbf{G}_t)} - k \mathbb{E} \log \frac{Z(\mathbf{G}''_t)}{Z(\mathbf{G}_t)} + (k-1) \mathbb{E} \log \frac{Z(\mathbf{G}'''_t)}{Z(\mathbf{G}_t)} \right]. \end{aligned} \quad (5.4)$$

Hence, it suffices to prove that for all $0 < t < 1$,

$$\mathbb{E} \left[\log \frac{Z(\mathbf{G}'_t)}{Z(\mathbf{G}_t)} \right] - k \mathbb{E} \left[\log \frac{Z(\mathbf{G}''_t)}{Z(\mathbf{G}_t)} \right] + (k-1) \mathbb{E} \left[\log \frac{Z(\mathbf{G}'''_t)}{Z(\mathbf{G}_t)} \right] \leq 0. \quad (5.5)$$

By the definition of the Boltzmann distribution (5.1),

$$\frac{Z(\mathbf{G}'_t)}{Z(\mathbf{G}_t)} = \left\langle \psi_{a_{\mathbf{m}_t+1}}, \mu_{\mathbf{G}_t} \right\rangle, \quad \frac{Z(\mathbf{G}''_t)}{Z(\mathbf{G}_t)} = \left\langle \psi_{a'_{\mathbf{m}'_t+1}}, \mu_{\mathbf{G}_t} \right\rangle, \quad \frac{Z(\mathbf{G}'''_t)}{Z(\mathbf{G}_t)} = \left\langle \psi_{a''_{\mathbf{m}''_t+1}}, \mu_{\mathbf{G}_t} \right\rangle.$$

Hence,

$$\log \frac{Z(\mathbf{G}'_t)}{Z(\mathbf{G}_t)} = \log \left\langle \psi_{a_{\mathbf{m}_t+1}}, \mu_{\mathbf{G}_t} \right\rangle = - \sum_{\ell \geq 1} \frac{1}{\ell} \left\langle 1 - \psi_{a_{\mathbf{m}_t+1}}, \mu_{\mathbf{G}_t} \right\rangle^\ell. \quad (5.6)$$

Further, in terms of the kernel representation $\dot{\mu}_{G_t}$ of the Boltzmann distribution we obtain

$$\mathbb{E} \left[\left\langle 1 - \psi_{a_{m_t+1}}, \mu_{G_t} \right\rangle^\ell \right] = \mathbb{E} \left[\left(1 - \sum_{\sigma \in \Omega^k} \psi(\sigma) \int_0^1 \prod_{i=1}^k \dot{\mu}_{G_t, z, x_i}(\sigma_i) dz \right)^\ell \right]. \quad (5.7)$$

Combining (5.6) and (5.7) yields

$$\mathbb{E} \left[\log \frac{Z(\mathbf{G}'_t)}{Z(\mathbf{G}_t)} \right] = -\mathbb{E} \left[\sum_{\ell \geq 1} \frac{1}{\ell} \left(1 - \sum_{\sigma \in \Omega^k} \psi(\sigma) \int_0^1 \prod_{i=1}^k \dot{\mu}_{G_t, z, x_i}(\sigma_i) dz \right)^\ell \right]. \quad (5.8)$$

Due to (2.1) and Fubini's theorem, we can exchange the sum and the expectation in (5.8); indeed, (2.1) yields

$$\sum_{\ell \geq 1} \mathbb{E} \left[\left(1 - \sum_{\sigma \in \Omega^k} \psi(\sigma) \int_0^1 \prod_{i=1}^k \dot{\mu}_{G_t, z, x_i}(\sigma_i) dz \right)^\ell \right] \leq \sum_{\ell \geq 1} \mathbb{E} \left[\max_{\sigma \in \Omega^k} |1 - \psi(\sigma)|^\ell \right] < \infty.$$

Thus, (5.8) becomes

$$\mathbb{E} \left[\log \frac{Z(\mathbf{G}'_t)}{Z(\mathbf{G}_t)} \right] = - \sum_{\ell \geq 1} \frac{1}{\ell} \mathbb{E} \left[\left(1 - \sum_{\sigma \in \Omega^k} \psi(\sigma) \int_0^1 \prod_{i=1}^k \dot{\mu}_{G_t, z, x_i}(\sigma_i) dz \right)^\ell \right]. \quad (5.9)$$

Following similar steps, we obtain expansions for the other two terms from (5.5) as well:

$$\mathbb{E} \left[\log \frac{Z(\mathbf{G}''_t)}{Z(\mathbf{G}_t)} \right] = -\frac{1}{k} \sum_{h=1}^k \sum_{\ell \geq 1} \frac{1}{\ell} \mathbb{E} \left[\left(1 - \sum_{\sigma \in \Omega^k} \psi(\sigma) \int_0^1 \dot{\mu}_{G_t, z, x_h}(\sigma_h) \prod_{i \neq h} \mu_{z, x_i}(\sigma_i) dz \right)^\ell \right], \quad (5.10)$$

$$\mathbb{E} \left[\log \frac{Z(\mathbf{G}'''_t)}{Z(\mathbf{G}_t)} \right] = - \sum_{\ell \geq 1} \frac{1}{\ell} \mathbb{E} \left[\left(1 - \sum_{\sigma \in \Omega^k} \psi(\sigma) \int_0^1 \prod_{i=1}^k \mu_{z, x_i}(\sigma_i) dz \right)^\ell \right]. \quad (5.11)$$

Finally, the assertion follows from **POS** and (5.5), (5.9), (5.10) and (5.11). \square

Proof of Proposition 5.2. Integrating t from 0 to 1 and applying Lemma 5.4, we obtain for any $\varepsilon > 0$,

$$\mathbb{E}[\log Z(\mathbf{G}_0)] \leq \mathbb{E}[\log Z(\mathbf{G}_1)] + o(n). \quad (5.12)$$

Letting

$$Y = \log \int_0^1 \prod_{i=1}^{\mathbf{m}_0''} \sum_{\sigma \in \Omega^k} \psi_i''(\sigma) \prod_{h=1}^k \mu_{z, x_{i,h}''}(\sigma_h) dz,$$

we claim that for a certain number $c = c(P) > 0$,

$$\mathbb{E} \log Z(\mathbf{G}) + \mathbb{E}[Y] \leq \mathbb{E} \log Z(\mathbf{G}_1) + \varepsilon cn. \quad (5.13)$$

Indeed, at $t = 0$ the variable node s is adjacent to the constraint nodes $a_i'', i \in [\mathbf{m}_t'']$, only. Hence, \mathbf{G}_0 decomposes into connected components, one of which comprises s and the a_i'' . Let \mathbf{G}_0'' be this component, and let \mathbf{G}_0' be the remainder of \mathbf{G}_0 . Then by construction we have $\mathbb{E} \log Z(\mathbf{G}_0'') = \mathbb{E}[Y]$. Thus, (5.12) yields

$$\mathbb{E}[\log Z(\mathbf{G}_0')] + \mathbb{E}[Y] = \mathbb{E}[\log Z(\mathbf{G}_0)] \leq \mathbb{E}[\log Z(\mathbf{G}_1)] + o(n). \quad (5.14)$$

Furthermore, \mathbf{G}_0' consists of the variable nodes v_1, \dots, v_n and the constraint nodes a_1, \dots, a_{m_1} , where \mathbf{m}_1 is a Poisson variable $\text{Po}(\exp(-\varepsilon)dn/k)$ conditioned on taking a value of at most dn/k . Thus, we can construct a random factor graph with the same distribution as \mathbf{G} from \mathbf{G}_0' by simply adding $dn/k - \mathbf{m}_1$ further random k -ary constraint nodes as per **INT3**. Since all weight functions $\psi \in \Psi$ take values in $(0, 2)$, we obtain $c = c(P) > 0$ such that

$$\mathbb{E} \log Z(\mathbf{G}) \leq \mathbb{E} \log Z(\mathbf{G}_0') + \varepsilon cn. \quad (5.15)$$

Combining (5.14) and (5.15), we obtain (5.13).

We further claim that there is a constant $c' = c'(P) > 0$ such that

$$\frac{1}{n} \mathbb{E}[Y] \leq \varepsilon c' + o(1) + \mathcal{B}''(\mu). \quad (5.16)$$

Indeed, $\mathcal{B}''(\mu) = \mathbb{E}[Y \mid \mathbf{m}_0'' = (k-1)dn/k]$. In other words, we can think of $\mathcal{B}''(\mu)$ as the free energy of \mathbf{G}_0'' given that $\mathbf{m}_0'' = (k-1)dn/k$. Thus, obtain \mathbf{G}_0''' from \mathbf{G}_0'' by adding $(k-1)dn/k - \mathbf{m}_0''$ more constraint nodes according to **INT5**, or by removing some random constraint nodes if $\mathbf{m}_0'' > (k-1)dn/k$. Then $\mathcal{B}''(\mu) = \mathbb{E} \log Z(\mathbf{G}_0''')$. Since \mathbf{m}_0'' is a Poisson variable with mean $\exp(-\varepsilon)(k-1)dn/k$, with probability $1 - \exp(-\Omega(n))$ we do not need to add or remove

more than $2\varepsilon(k-1)dn$ constraint nodes. The tail bound (2.1) therefore implies together with the Chernoff bound that (5.16) is satisfied for a certain $c' = c'(P)$.

By similar arguments, for a certain $c'' = c''(P)$ we have

$$\frac{1}{n}\mathbb{E}[\log Z(\mathbf{G}_1)] \leq \mathcal{B}'(\mu) + \varepsilon c'' + o(1). \quad (5.17)$$

Indeed, $\mathcal{B}'(\mu)$ is nothing but the conditional expectation of $\log Z(\mathbf{G}_1)$ given that $\mathbf{m}'_1 = dn$. Hence, if we pad \mathbf{G}_1 by adding the missing $dn - \mathbf{m}'_1$ constraint nodes a'_i according to **INT4**, then the total number of constraints added does not exceed $2\varepsilon dn$ with probability $1 - \exp(-\Omega(n))$. Hence, (5.17) follows from (2.1) and the Chernoff bound.

Finally, combining (5.12)–(5.17), we conclude that

$$\frac{1}{n}\mathbb{E}\log Z(\mathbf{G}) \leq \mathcal{B}'(\mu) - \mathcal{B}''(\mu) + \varepsilon c''' + o(1)$$

for a certain $c''' = c'''(P) > 0$. Since this is true for any fixed $\varepsilon > 0$, the assertion follows. \square

5.3. Proof of Proposition 5.3. Following Panchenko [55], who worked with factor graphs of Erdős-Rényi type, we are going to use the invariance property of $\pi \in \mathfrak{D}^*$ under the $*(N, M)$ -operation to simplify $\mathcal{B}', \mathcal{B}''$ separately.

Lemma 5.5. *Suppose that $\pi \in \mathfrak{D}^*$. Then*

$$\mathbb{E}[\mathcal{B}''(\boldsymbol{\mu}^\pi)] = \frac{d(k-1)n}{k}\mathbb{E}[\log \langle \boldsymbol{\psi}_1, \pi \rangle]. \quad (5.18)$$

Proof. Let $\phi = \mathbb{E}[\log \langle \boldsymbol{\psi}_1, \pi \rangle]$ for brevity. We claim that for any integer $m \geq 0$,

$$\mathbb{E}\left[\log \frac{\langle \bigoplus_{i=1}^{m+1} \boldsymbol{\psi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^m \boldsymbol{\psi}_i, \pi \rangle}\right] = \phi. \quad (5.19)$$

Then (5.18) follows by summing (5.19) on $0 \leq m < d(k-1)n/k$.

Thus, we are left to prove (5.19). Since $\pi \in \mathfrak{D}^*$, Corollaries 3.4 and 3.9 imply that for any integer $\ell \geq 1$,

$$\begin{aligned} \mathbb{E}\left[\left(\frac{\langle \bigoplus_{i=1}^{m+1} \boldsymbol{\psi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^m \boldsymbol{\psi}_i, \pi \rangle}\right)^\ell\right] &= \mathbb{E}\left[\left\langle \frac{\bigoplus_{i=1}^m \boldsymbol{\psi}_i}{\langle \bigoplus_{j=1}^m \boldsymbol{\psi}_j, \pi \rangle} \oplus \boldsymbol{\psi}_{m+1}, \pi \right\rangle^\ell\right] \\ &= \mathbb{E}\left[\left\langle \boldsymbol{\psi}_{m+1}, \bigoplus_{i=1}^m \boldsymbol{\psi}_i * \pi \right\rangle^\ell\right] = \mathbb{E}[\langle \boldsymbol{\psi}_{m+1}, \pi \rangle^\ell] = \mathbb{E}[\langle \boldsymbol{\psi}_1, \pi \rangle^\ell]. \end{aligned}$$

Consequently, for all $\ell \geq 1$ we have

$$\mathbb{E}\left[\left(1 - \frac{\langle \bigoplus_{i=1}^{m+1} \boldsymbol{\psi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^m \boldsymbol{\psi}_i, \pi \rangle}\right)^\ell\right] = \mathbb{E}[(1 - \langle \boldsymbol{\psi}_1, \pi \rangle)^\ell]. \quad (5.20)$$

Further, because the continuous function $z \in [-1, 1] \mapsto |z|$ is a uniform limit of polynomials, (5.20) yields

$$\mathbb{E}\left|\left(1 - \frac{\langle \bigoplus_{i=1}^{m+1} \boldsymbol{\psi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^m \boldsymbol{\psi}_i, \pi \rangle}\right)^\ell\right| = \mathbb{E}|(1 - \langle \boldsymbol{\psi}_1, \pi \rangle)^\ell|$$

Therefore, invoking (2.1), we obtain

$$\sum_{\ell \geq 1} \frac{1}{\ell} \mathbb{E}\left|\left(1 - \frac{\langle \bigoplus_{i=1}^{m+1} \boldsymbol{\psi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^m \boldsymbol{\psi}_i, \pi \rangle}\right)^\ell\right| = \sum_{\ell \geq 1} \frac{1}{\ell} \mathbb{E}|(1 - \langle \boldsymbol{\psi}_1, \pi \rangle)^\ell| \leq \sum_{\ell \geq 1} \mathbb{E}\left[\max_{\sigma \in \Omega^k} |1 - \boldsymbol{\psi}_1(\sigma)|^\ell\right] < \infty.$$

Hence, by (5.20) and Fubini's theorem,

$$\mathbb{E}\left[\log \frac{\langle \bigoplus_{i=1}^{m+1} \boldsymbol{\psi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^m \boldsymbol{\psi}_i, \pi \rangle}\right] = - \sum_{\ell \geq 1} \frac{1}{\ell} \mathbb{E}\left[\left(1 - \frac{\langle \bigoplus_{i=1}^{m+1} \boldsymbol{\psi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^m \boldsymbol{\psi}_i, \pi \rangle}\right)^\ell\right] = - \sum_{\ell \geq 1} \frac{1}{\ell} \mathbb{E}|(1 - \langle \boldsymbol{\psi}_1, \pi \rangle)^\ell| = \phi,$$

which is (5.19). \square

Lemma 5.6. *Suppose that $\pi \in \mathfrak{D}^*$. Then $\mathbb{E}[\mathcal{B}'(\boldsymbol{\mu}^\pi)] = \mathbb{E}\log \langle \boldsymbol{\psi}_1, \pi \rangle$.*

Proof. We use a similar argument as in the proof of Lemma 5.5. This time we set $\phi = \mathbb{E} \log \langle \boldsymbol{\varphi}_1, \pi \rangle$. It suffices to show that for every $n \geq 0$,

$$\mathbb{E} \left[\log \frac{\langle \bigoplus_{i=1}^{n+1} \boldsymbol{\varphi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^n \boldsymbol{\varphi}_i, \pi \rangle} \right] = \phi. \quad (5.21)$$

As in the proof of Lemma 5.5, we use that $\pi \in \mathfrak{D}^*$ and apply Corollaries 3.4 and 3.9 to obtain for any $\ell \geq 1$,

$$\mathbb{E} \left[\left(\frac{\langle \bigoplus_{i=1}^{n+1} \boldsymbol{\varphi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^n \boldsymbol{\varphi}_i, \pi \rangle} \right)^\ell \right] = \mathbb{E} \left[\left(\left\langle \frac{\bigoplus_{i=1}^n \boldsymbol{\varphi}_i}{\langle \bigoplus_{j=1}^n \boldsymbol{\varphi}_j, \pi \rangle} \oplus \boldsymbol{\varphi}_{n+1}, \pi \right\rangle \right)^\ell \right] = \mathbb{E} \left[\left\langle \boldsymbol{\varphi}_{n+1}, \bigoplus_{i=1}^n \boldsymbol{\varphi}_i * \pi \right\rangle^\ell \right] = \mathbb{E} \left[\langle \boldsymbol{\varphi}_1, \pi \rangle^\ell \right]. \quad (5.22)$$

Hence, for any $\ell \geq 1$,

$$\mathbb{E} \left[\left(1 - \frac{\langle \bigoplus_{i=1}^{n+1} \boldsymbol{\varphi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^n \boldsymbol{\varphi}_i, \pi \rangle} \right)^\ell \right] = \mathbb{E} \left[(1 - \langle \boldsymbol{\varphi}_1, \pi \rangle)^\ell \right]. \quad (5.23)$$

Further, approximating the absolute value by polynomials, we obtain from (5.22) that

$$\mathbb{E} \left| \left(1 - \frac{\langle \bigoplus_{i=1}^{n+1} \boldsymbol{\varphi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^n \boldsymbol{\varphi}_i, \pi \rangle} \right)^\ell \right| = \mathbb{E} \left| (1 - \langle \boldsymbol{\varphi}_1, \pi \rangle)^\ell \right|.$$

Thus, (5.21) follows from (5.23) and Fubini's theorem. \square

Finally, Proposition 5.3 is immediate from Lemmas 5.5 and 5.6.

6. THE FREE ENERGY: LOWER BOUND

6.1. Outline. In this section we prove the following lower bound on the free energy that matches the upper bound from Proposition 5.1. The lower bound does not require the assumption **POS**.

Proposition 6.1. *We have*

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\log Z(\mathbf{G})] \geq \inf_{\pi \in \mathfrak{D}^*} \mathcal{B}(\pi), \quad \liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\log Z(\mathbb{G})] \geq \inf_{\pi \in \mathfrak{D}^*} \mathcal{B}(\pi).$$

Theorem 2.7 follows immediately from Propositions 5.1 and 6.1.

The proof of Proposition 6.1 is based on a kind of coupling argument that is colloquially referred to as the 'Aizenman–Sims–Starr' scheme. This technique has been applied with great success to random factor graphs of Erdős–Rényi type, where the degree distribution is approximately Poisson [24, 16, 55]. The basic idea is to couple a random factor graph with n variable nodes with a random factor graph with $n+1$ variable nodes and to calculate the difference of their free energies very precisely. This coupling is very easy to set up in the Erdős–Rényi case due to the Stein–Chen property of the Poisson distribution.

However, in the case of random regular graphs matters are more intricate. Due to the rigid local structure there is no obvious way of coupling random regular factor graphs with n and $n+1$ variable nodes. As in Section 4, we therefore resort to the idea of creating a bit of wiggling room by carving out a few cavities, in such a way that the free energy does not change significantly. But the details of the construction are delicate.

Let n, ω be integers and let X, Y be two independent Poisson variables with mean ω . The protagonist of the proof is the random factor graph $\mathbf{G}_{n, \omega}$ defined as follows. Let

$$N_{n, \omega} = k \vee (n - X) \quad \text{and let} \quad \Delta_{n, \omega} \stackrel{d}{=} \text{Be}(dN_{n, \omega}/k - \lfloor dN_{n, \omega}/k \rfloor)$$

be independent of Y . Further, set

$$M_{n, \omega} = d \vee (\lfloor dN_{n, \omega}/k \rfloor \wedge (\lfloor dN_{n, \omega}/k \rfloor + \Delta_{n, \omega} - dX - Y)).$$

Then $\mathbf{G}_{n, \omega}$ has $N_{n, \omega}$ variable nodes v_i , $i \in [N_{n, \omega}]$, and $M_{n, \omega}$ constraint nodes a_i , $i \in [M_{n, \omega}]$. The weight functions ψ_{a_i} are chosen independently from P . Furthermore, the variable and constraint nodes are linked through a random (one-to-one) pairing

$$F_{M_{n, \omega}} \times [k] \rightarrow V_{N_{n, \omega}} \times [d].$$

Since $kM_{n, \omega} \leq dN_{n, \omega}$ by construction, such a pairing exists, but some variable clones may go unpaired. We are going to harness these unpaired 'cavities' to set up a coupling of $\mathbf{G}_{n, \omega}$ and $\mathbf{G}_{n+1, \omega}$.

To this end, consider a further random factor graph $\hat{\mathbf{G}}_{n,\omega}$ with $N_{n,\omega}$ variable nodes and $\hat{M}_{n,\omega} = M_{n+1,\omega} - d$ constraint nodes. The weight functions are chosen independently from P , and the connections between the constraint and variable nodes are induced by a random pairing

$$F_{\hat{M}_{n,\omega}} \times [k] \rightarrow V_{N_{n,\omega}} \times [d].$$

Rather than coupling $\mathbf{G}_{n,\omega}$ and $\mathbf{G}_{n+1,\omega}$ directly, we will couple $\mathbf{G}_{n,\omega}$ and $\hat{\mathbf{G}}_{n,\omega}$ as well as $\mathbf{G}_{n+1,\omega}$ and $\hat{\mathbf{G}}_{n,\omega}$.

This construction leads to an approximate formula for the free energy of $\mathbf{G}_{n,\omega}$ that comes in terms of the kernel representation of the Boltzmann distribution of $\hat{\mathbf{G}}_{n,\omega}$. To be precise, let $\hat{\mathcal{C}}$ be the set of variables $v_i \in V_{N_{n,\omega}}$ with at least one unpaired clone in $\hat{\mathbf{G}}_{n,\omega}$. Consider the random kernel

$$\hat{\rho}_{n,\omega} = \hat{\mu}_{\hat{\mathbf{G}}_{n,\omega}, \hat{\mathcal{C}}} \in \mathfrak{K}$$

representing the joint distribution of the cavities $\hat{\mathcal{C}}$. Further, let $\hat{\pi}_{n,\omega} \in \mathfrak{D}$ be the distribution of $\hat{\rho}_{n,\omega}$. To deal with the conditioning on the event \mathcal{S} , we also introduce versions $\mathbb{G}_{n,\omega}$, $\hat{\mathbb{G}}_{n,\omega}$ of the above random factor graphs conditional on \mathcal{S} . Let $\tilde{\rho}_{n,\omega} = \hat{\mu}_{\hat{\mathbb{G}}_{n,\omega}, \hat{\mathcal{C}}} \in \mathfrak{K}$ be the kernel representation of the corresponding Boltzmann distribution, and let $\tilde{\pi}_{n,\omega} \in \mathfrak{D}$ be the law of $\tilde{\rho}_{n,\omega}$. In Section 6.2 we will derive the following formula.

Proposition 6.2. *For any $\varepsilon > 0$ there exists $\omega > 0$ such that*

$$\begin{aligned} \liminf_{n \rightarrow \infty} \mathbb{E} \left[\log \frac{Z(\mathbf{G}_{n+1,\omega})}{Z(\mathbf{G}_{n,\omega})} \right] - \mathbb{E} \left[\log \langle \boldsymbol{\varphi}_1, \hat{\rho}_{n,\omega} \rangle - \log \left\langle \bigoplus_{\hat{M}_{n,\omega} < i \leq M_{n,\omega}} \boldsymbol{\psi}_{1,i}, \hat{\rho}_{n,\omega} \right\rangle \right] &\geq -\varepsilon, \\ \liminf_{n \rightarrow \infty} \mathbb{E} \left[\log \frac{Z(\mathbb{G}_{n+1,\omega})}{Z(\mathbb{G}_{n,\omega})} \right] - \mathbb{E} \left[\log \langle \boldsymbol{\varphi}_1, \tilde{\rho}_{n,\omega} \rangle - \log \left\langle \bigoplus_{\hat{M}_{n,\omega} < i \leq M_{n,\omega}} \boldsymbol{\psi}_{1,i}, \tilde{\rho}_{n,\omega} \right\rangle \right] &\geq -\varepsilon. \end{aligned}$$

There are still two gaps to fill toward the proof of Proposition 6.1. First, the estimate of the free energy provided by Proposition 6.2 does not quite match the functional $\mathcal{B}(\hat{\pi}_{n,\omega})$. Second, the distribution $\hat{\pi}_{n,\omega} \in \mathfrak{D}$ does not generally belong to the subspace \mathfrak{D}^* . The following proposition deals with the second issue, which holds the key to resolving the first. Recall that the topology of \mathfrak{D} is induced by the Wasserstein metric $\mathcal{D}_{\square}(\cdot, \cdot)$. We introduce a relaxed version of \mathfrak{D}^* by letting

$$\mathfrak{D}_{\varepsilon,N,M}^* = \{ \pi \in \mathfrak{D} : \mathcal{D}_{\square}(\pi, \pi^{*(u,w)}) \leq \varepsilon \text{ for all } u \leq N \text{ and } w \leq M \}.$$

Since (2.1) and Lemma 3.6 show that the map $\pi \mapsto \pi^{*(u,w)}$ is continuous, $\mathfrak{D}_{\varepsilon,N,M}^*$ is a closed subspace of the compact Polish space \mathfrak{D} .

Proposition 6.3. *For any $\varepsilon, L > 0$ there is $\omega_0 > 0$ such that for every $\omega > \omega_0$ for large enough n we have*

$$\hat{\pi}_{n,\omega}, \tilde{\pi}_{n,\omega} \in \mathfrak{D}_{\varepsilon,L,L}^*.$$

The proof of Proposition 6.3 can be found in Section 6.3. Finally, in Section 6.4 we derive Proposition 6.1 from Propositions 6.2 and 6.3.

6.2. Proof of Proposition 6.2. We assume throughout that $\omega > \omega_0$ for a big enough $\omega_0 = \omega_0(d, P)$ and that n sufficiently large.

Obtain the random factor graph $\mathbf{G}'_{n,\omega}$ from $\hat{\mathbf{G}}_{n,\omega}$ by adding $M_n - \hat{M}_{n,\omega}$ new random constraint nodes a_i , $\hat{M}_{n,\omega} < i \leq M_n$, whose weight functions are drawn from P independently and that are linked with the variable nodes via a random pairing with the cavities $\hat{\mathcal{C}}$ of $\hat{\mathbf{G}}_{n,\omega}$.

Further, if $k\hat{M}_{n,\omega} \leq dN_{n,\omega} - d(k-1)$, then obtain $\mathbf{G}''_{n,\omega}$ from $\hat{\mathbf{G}}_{n,\omega}$ by adding one new variable node $\hat{v} = v_{N_{n,\omega}+1}$ along with d random constraint nodes $\hat{a}_1, \dots, \hat{a}_d$ adjacent to \hat{v} whose weight functions are drawn independently from P . To be precise, the clones of \hat{v} are paired each with a uniformly random clone \hat{h}_i of \hat{a}_i for $i = 1, \dots, d$, and the remaining $d(k-1)$ clones of the \hat{a}_i are paired with randomly chosen cavities of $\hat{\mathbf{G}}_{n,\omega}$. If $k\hat{M}_{n,\omega} > dN_{n,\omega} - d(k-1)$, then obtain $\mathbf{G}''_{n,\omega}$ from $\hat{\mathbf{G}}_{n,\omega}$ by just adding a new isolated variable node \hat{v} .

Obtain $\mathbb{G}'_{n,\omega}, \mathbb{G}''_{n,\omega}$ analogously from $\hat{\mathbb{G}}_{n,\omega}$ while conditioning on the event that the outcome is simple. If it is impossible to add the required number of constraint nodes in such a way that the resulting factor graph is simple, then do not add any.

Lemma 6.4. For $\omega > 0$ we have

$$\begin{aligned} |\mathbb{E} \log Z(\mathbf{G}_{n,\omega}) - \mathbb{E} \log Z(\mathbf{G}'_{n,\omega})| &= o(1), & |\mathbb{E} \log Z(\mathbf{G}_{n+1,\omega}) - \mathbb{E} \log Z(\mathbf{G}''_{n,\omega})| &= o_\omega(1), \\ |\mathbb{E} \log Z(\mathbb{G}_{n,\omega}) - \mathbb{E} \log Z(\mathbb{G}'_{n,\omega})| &= o_\omega(1), & |\mathbb{E} \log Z(\mathbb{G}_{n+1,\omega}) - \mathbb{E} \log Z(\mathbb{G}''_{n,\omega})| &= o_\omega(1). \end{aligned} \quad (6.1)$$

Proof. Since $d \geq 3$ and $k \geq 2$,

$$\hat{M}_{n,\omega} - \lfloor dN_{n,\omega}/k \rfloor \leq dN_{n,\omega}/k + 1 + d/k - d - \lfloor dN_{n,\omega}/k \rfloor \leq 2 - d(1 - 1/k) \leq 1/2.$$

Because the left-hand side is an integer, we conclude that $\hat{M}_{n,\omega} \leq \lfloor dN_{n,\omega}/k \rfloor$. Similarly,

$$\hat{M}_{n,\omega} - (\lfloor dN_{n,\omega}/k \rfloor + \Delta_{n,\omega} - dX - Y) \leq 2 - d(1 - 1/k) \leq 1/2.$$

Thus, $M_{n,\omega} \geq \hat{M}_{n,\omega}$. Hence, $\mathbf{G}_{n,\omega}$ and \mathbf{G}' are identically distributed.

Moving on to the second claim, we consider the event \mathcal{A} that the last variable node is adjacent to precisely d distinct constraint nodes. Then

$$\mathbb{P}[\mathbf{G}_{n+1,\omega} \in \mathcal{A}] = 1 - O(\omega/n), \quad (6.2)$$

while $\mathbf{G}''_{n,\omega} \in \mathcal{A}$ with certainty. Given \mathcal{A} and given that $X + Y \leq \sqrt{n}$, say, the subgraph $\tilde{\mathbf{G}}_{n+1,\omega}$ obtained from $\mathbf{G}_{n+1,\omega}$ by deleting \tilde{v} along with its adjacent constraint nodes is distributed precisely as $\hat{\mathbf{G}}_{n,\omega}$, and therefore $\mathbf{G}''_{n,\omega}$ and $\mathbf{G}_{n+1,\omega}$ can be coupled identically. Hence,

$$\mathbb{E}[\log Z(\mathbf{G}''_{n,\omega}) | X + Y \leq \sqrt{n}] = \mathbb{E}[\log Z(\mathbf{G}_{n+1,\omega}) | \mathcal{A}, X + Y \leq \sqrt{n}]. \quad (6.3)$$

If, on the other hand, $X + Y \leq \sqrt{n}$ but \mathcal{A} does not occur, then we can couple $\mathbf{G}_{n+1,\omega}$ and $\mathbf{G}''_{n,\omega}$ such that both disagree on at most $2d$ constraint nodes. Indeed, suppose that $v_{N_{n+1,\omega}}$ has $\tilde{d} < d$ adjacent constraints in $\mathbf{G}_{n+1,\omega}$. Then the subgraph obtained by removing $v_{N_{n+1,\omega}}$, its \tilde{d} neighbors and another $d - \tilde{d}$ random constraint nodes is distributed precisely as $\hat{\mathbf{G}}_{n,\omega}$. Hence, we can obtain both $\mathbf{G}_{n+1,\omega}$ and $\mathbf{G}''_{n,\omega}$ from $\hat{\mathbf{G}}_{n,\omega}$ by adding d (possibly distinct) constraint nodes. Thus, (2.1) ensures that

$$\mathbb{E}[\log Z(\mathbf{G}''_{n,\omega}) | X + Y \leq \sqrt{n}] = \mathbb{E}[\log Z(\mathbf{G}_{n+1,\omega}) | \overline{\mathcal{A}}, X + Y \leq \sqrt{n}] + O(1). \quad (6.4)$$

Furthermore, (2.1) ensures that

$$\mathbb{E}[\log Z(\mathbf{G}''_{n,\omega}) | X, Y], \mathbb{E}[\log Z(\mathbf{G}_{n+1,\omega}) | X, Y] = O(n). \quad (6.5)$$

Since $\mathbb{P}[X + Y > \sqrt{n}] = o(n^{-2})$, (6.2)–(6.5) yield the second assertion.

Matters get slightly more complicated once we condition on \mathcal{S} . Since $X + Y \leq \log n$ with probability $1 - O(n^{-k})$, due to (2.1) the event $X + Y > \log n$ contributes no more than an additive $o(1)$ to the difference of the free energies. Hence, we may condition on $X + Y \leq \log n$. Let \mathbf{d}' be the vector comprising the variable degrees in $\hat{\mathbf{G}}_{n,\omega}$. Let \mathcal{D} be the set of all such sequences with entries either d or $d - 1$. A standard moment calculation shows that given any possible \mathbf{d}' , the event $\hat{\mathbf{G}}_{n,\omega} \in \mathcal{S}$ has probability $(1 + o(1)) \exp[-(d-1)(k-1)/2 - \mathbf{1}\{k=2\}(d-1)^2/4]$ (cf. Fact 2.2). Therefore, with $\tilde{O}(\cdot)$ hiding poly-logarithmic terms,

$$\mathbb{P}[\mathbf{d}' \in \mathcal{D} | X + Y \leq \log n] = 1 - \tilde{O}(1/n). \quad (6.6)$$

Similarly, let \mathbf{d} comprise the variable degrees of the factor graph $\mathbb{G}_{n,\omega}^-$ obtained from $\mathbb{G}_{n,\omega}$ by deleting the last d constraint nodes. Then

$$\mathbb{P}[\mathbf{d} \in \mathcal{D} | X + Y \leq \log n] = 1 - \tilde{O}(1/n). \quad (6.7)$$

Additionally, let \mathcal{E} be the set of all factor graphs that have a constraint node that is adjacent to variable nodes of degree less than d only. Then

$$\mathbb{P}[\hat{\mathbf{G}}_{n,d} \in \mathcal{E} | X + Y \leq \log n], \mathbb{P}[\mathbb{G}_{n,\omega}^- \in \mathcal{E} | X + Y \leq \log n] = \tilde{O}(n^{1-k}). \quad (6.8)$$

Further, on the event $\mathcal{D} \setminus \mathcal{E}$ we can couple $\mathbb{G}'_{n,\omega}$ and $\mathbb{G}_{n,\omega}$ identically, because there is no way of adding the missing constraint nodes to $\hat{\mathbf{G}}_{n,\omega}$ without obtaining a simple factor graph. Hence,

$$\mathbb{E}[\log Z(\mathbb{G}_{n,\omega}) | \mathbf{d} \in \mathcal{D}, \mathbb{G}_{n,\omega}^- \notin \mathcal{E}, X + Y \leq \log n] = \mathbb{E}[\log Z(\mathbb{G}'_{n,\omega}) | \hat{\mathbf{d}} \in \mathcal{D}, \hat{\mathbf{G}}_{n,\omega} \notin \mathcal{E}, X + Y \leq \log n]. \quad (6.9)$$

But (6.9) does not yet suffice to prove (6.1) because outside the event $\mathcal{D} \setminus \mathcal{E}$ the free energies of the two factor graphs may differ by $\Omega(n)$. Hence, we also need to consider the event \mathcal{D}' that \mathbf{d} has a single $d-2$ entry; this suffices because

$$\mathbb{P}[\mathbf{d} \notin \mathcal{D} \cup \mathcal{D}' \mid X + Y \leq \log n], \mathbb{P}[\hat{\mathbf{d}} \notin \mathcal{D} \cup \mathcal{D}' \mid X + Y \leq \log n] = 1 - \tilde{O}(n^{-2}) \quad (6.10)$$

and thus the contribution of the complement of $\mathcal{D} \cup \mathcal{D}'$ to the free energy difference is $o(1)$ due to (2.1). Considering the event \mathcal{D}' is indeed necessary because $\mathbb{P}[\hat{\mathbf{d}} \in \mathcal{D}' \mid X + Y \leq \log n] > \mathbb{P}[\mathbf{d} \in \mathcal{D}' \mid X + Y \leq \log n]$. Indeed, while $\hat{\mathbb{G}}_{n,\omega}$ is just a uniformly random simple factor graph with $N_{n,\omega}$ variable and $\hat{M}_{n,\omega}$ constraint nodes, $\mathbb{G}_{n,\omega}^-$ has a tilted distribution, with each possible simple graph being weighed according to the number of extensions into a simple graph with $M_{n,\omega}$ constraints. In effect, since variable nodes of degree less than $d-1$ leave us with fewer extensions, the event \mathcal{D}' is less likely in $\mathbb{G}_{n,\omega}^-$. Yet because $M_{n,\omega} - \hat{M}_{n,\omega}$ is bounded, on the event \mathcal{D}' we can couple $\mathbb{G}_{n,\omega}$ and $\mathbb{G}'_{n,\omega}$ such that both differ only in a bounded number of constraint nodes. As a consequence,

$$\mathbb{E}[\log Z(\mathbb{G}_{n,\omega}) \mid \mathbf{d} \in \mathcal{D}', \mathbb{G}_{n,\omega}^- \notin \mathcal{E}, X + Y \leq \log n] = \mathbb{E}[\log Z(\mathbb{G}'_{n,\omega}) \mid \hat{\mathbf{d}} \in \mathcal{D}', \hat{\mathbb{G}}_{n,\omega} \notin \mathcal{E}, X + Y \leq \log n] + O(1). \quad (6.11)$$

Additionally, we claim that also $\mathbb{G}_{n,\omega}$ given $\mathbf{d} \in \mathcal{D}$ and $\mathbb{G}'_{n,\omega}$ given $\hat{\mathbf{d}} \in \mathcal{D}'$ can be coupled such that with probability $1 - \tilde{O}(1/n)$ both differ only in $\tilde{O}(1)$ constraint nodes and that, in effect,

$$\mathbb{E}[\log Z(\mathbb{G}_{n,\omega}) \mid \mathbf{d} \in \mathcal{D}, \mathbb{G}_{n,\omega}^- \notin \mathcal{E}, X + Y \leq \log n] = \mathbb{E}[\log Z(\mathbb{G}'_{n,\omega}) \mid \hat{\mathbf{d}} \in \mathcal{D}', \hat{\mathbb{G}}_{n,\omega} \notin \mathcal{E}, X + Y \leq \log n] + \tilde{O}(1). \quad (6.12)$$

To see this, let u_1, \dots, u_ℓ be the variables nodes of degree less than d in $\mathbb{G}_{n,\omega}^-$; suppose, indeed, that all of them have degree $d-1$. Similarly, let $u'_1, \dots, u'_{\ell-1}$ be the cavities of $\hat{\mathbb{G}}_{n,\omega}$, all of degree $d-1$ except for $u'_{\ell-1}$, which has degree $d-2$. Pick a further variable node u'_ℓ of degree d randomly. Then with probability $1 - \tilde{O}(n^{-1})$ the second neighborhoods $\partial^2\{u_1, \dots, u_\ell\}, \partial^2\{u'_1, \dots, u'_\ell\}$ both have size $\ell(k-1)(d-1)$. Consequently, the subgraphs of $\mathbb{G}_{n,\omega}^-$ and $\hat{\mathbb{G}}_{n,\omega}$ obtained by removing u_1, \dots, u_ℓ and u'_1, \dots, u'_ℓ along with their neighbors, respectively, can be coupled such that both coincide with probability $1 - \tilde{O}(1/n)$. Thus, $\mathbb{G}_{n,\omega}$ and $\mathbb{G}'_{n,\omega}$ can be coupled such that the expected number of constraint nodes on which the two factor graphs differ is $\tilde{O}(1)$, whence we obtain (6.12).

To deal with the event \mathcal{E} , we may assume that $k=2$ due to (6.6). Furthermore, because of (6.8) and as

$$\mathbb{P}[\mathbf{d}' \notin \mathcal{D} \mid \hat{\mathbb{G}}_{n,\omega} \in \mathcal{E}, X + Y \leq \log n] = \tilde{O}(1/n), \quad \mathbb{P}[\mathbf{d} \notin \mathcal{D} \mid \mathbb{G}_{n,\omega}^- \in \mathcal{E}, X + Y \leq \log n] = \tilde{O}(1/n), \quad (6.13)$$

we may assume that $\mathbf{d}, \mathbf{d}' \in \mathcal{D}$. Since

$$\mathbb{P}[\hat{\mathbb{G}}_{n,\omega} \in \mathcal{E} \mid \mathbf{d}' \in \mathcal{D}, X + Y \leq \log n] \geq \mathbb{P}[\mathbb{G}_{n,\omega}^- \in \mathcal{E} \mid \mathbf{d} \in \mathcal{D}, X + Y \leq \log n] \quad (6.14)$$

because the event \mathcal{E} precludes certain extensions into a simple factor graph with $M_{n,\omega}$ constraints, we just need to consider the case that $\hat{\mathbb{G}}_{n,\omega} \in \mathcal{E}$ and $\mathbb{G}_{n,\omega}^- \notin \mathcal{E}$ given that $\mathbf{d}, \mathbf{d}' \in \mathcal{D}$. Let u_1, \dots, u_ℓ and u'_1, \dots, u'_ℓ be the cavities of $\mathbb{G}_{n,\omega}^-$ and $\hat{\mathbb{G}}_{n,\omega}$, respectively. Pick one further constraint node b of $\hat{\mathbb{G}}_{n,\omega}$. Then with probability $1 - \tilde{O}(1/n)$ the set $\partial^2\{u_1, \dots, u_\ell\}$ has size $\ell(d-1)$, and all variable nodes in this set have pairwise distance at least four. The same is true of the set $\partial^2\{u'_1, \dots, u'_\ell\} \cup \partial b$ with probability $1 - \tilde{O}(1/n)$. If these two events occur, then $\mathbb{G}_{n,\omega}$ and $\mathbb{G}_{n,\omega}'$ can be coupled such that they only differ on the $\tilde{O}(1)$ constraint nodes that are adjacent to u_1, \dots, u_ℓ and u'_1, \dots, u'_ℓ and b . Hence, we obtain a coupling such that $\mathbb{G}_{n,\omega}$ and $\mathbb{G}_{n,\omega}'$ only differ on $\tilde{O}(1/n)$ variable nodes in expectation, and thus

$$\mathbb{E}[\log Z(\mathbb{G}'_{n,\omega}) \mid \mathbf{d}' \in \mathcal{D}, \hat{\mathbb{G}}_{n,\omega} \in \mathcal{E}, X + Y \leq \log n] = \mathbb{E}[\log Z(\mathbb{G}_{n,\omega}) \mid \mathbf{d} \in \mathcal{D}, \mathbb{G}_{n,\omega}^- \notin \mathcal{E}, X + Y \leq \log n] + \tilde{O}(1). \quad (6.15)$$

Moreover, because given \mathcal{E} there is precisely one constraint involving variables of degree $d-1$ only with probability $1 - \tilde{O}(1/n)$, we obtain

$$\mathbb{E}[\log Z(\mathbb{G}'_{n,\omega}) \mid \mathbf{d}' \in \mathcal{D}, \hat{\mathbb{G}}_{n,\omega} \in \mathcal{E}, X + Y \leq \log n] = \mathbb{E}[\log Z(\mathbb{G}_{n,\omega}) \mid \mathbf{d} \in \mathcal{D}, \mathbb{G}_{n,\omega}^- \in \mathcal{E}, X + Y \leq \log n] + \tilde{O}(1). \quad (6.16)$$

Combining (6.15)–(6.16), we obtain the left bound stated in (6.1).

We proceed similarly to derive the right bound in (6.1). Indeed, in this case we do not need to consider the event \mathcal{E} separately, because all additional constraint nodes are connected with a variable node that does not belong to $\hat{\mathbb{G}}_{n,\omega}$ or $\mathbb{G}_{n,\omega}^-$, respectively. Hence, on the event \mathcal{D} we can couple $\mathbb{G}_{n,\omega}''$ and $\mathbb{G}_{n+1,\omega}$ identically, and thus

$$\mathbb{E}[\log Z(\mathbb{G}_{n+1,\omega}) \mid \mathbf{d} \in \mathcal{D}, X + Y \leq \log n] = \mathbb{E}[\log Z(\mathbb{G}_{n,\omega}'') \mid \mathbf{d}' \in \mathcal{D}, X + Y \leq \log n]. \quad (6.17)$$

In effect, due to (6.10) we just need to construct a coupling in the event that $\hat{\mathbb{G}}_{n,\omega} \in \mathcal{D}'$ and $\hat{\mathbb{G}}_{n,\omega}^- \in \mathcal{D}$. To this end, we proceed as above by coupling $\hat{\mathbb{G}}_{n,\omega}$, $\mathbb{G}_{n,\omega}^-$ given the second neighborhoods of the cavities such that both only

differ in an expected $\tilde{O}(1/n)$ constraint nodes. Since $\mathbb{P}[\mathbf{d}' \in \mathcal{D}' | \mathbf{X} + \mathbf{Y} \leq \log n], \mathbb{P}[\mathbf{d} \in \mathcal{D}' | \mathbf{X} + \mathbf{Y} \leq \log n] = \tilde{O}(1/n)$ and $\mathbb{P}[\mathbf{d} \in \mathcal{D}' | \mathbf{X} + \mathbf{Y} \leq \log n] \geq \mathbb{P}[\mathbf{d}' \in \mathcal{D}' | \mathbf{X} + \mathbf{Y} \leq \log n]$, the second part of (6.1) follows from (6.17). \square

We are ready to compare the free energies of $Z(\mathbf{G}'')$, $Z(\hat{\mathbf{G}})$ and $Z(\mathbf{G}')$, $Z(\hat{\mathbf{G}})$ and of the corresponding simple graphs. We will carry the proofs out for the case of the simple random factor graph $\hat{\mathbf{G}}$; the other case is simply obtained by skipping any deliberations pertinent to the event \mathcal{S} .

Lemma 6.5. *We have*

$$\mathbb{E} \left[\log \frac{Z(\mathbf{G}_{n,\omega}'')}{Z(\hat{\mathbf{G}}_{n,\omega})} \right] = \mathbb{E} [\log \langle \boldsymbol{\varphi}_1, \hat{\rho}_{n,\omega} \rangle] + o_\omega(1), \quad \mathbb{E} \left[\log \frac{Z(\mathbf{G}_{n,\omega}'')}{Z(\hat{\mathbf{G}}_{n,\omega})} \right] = \mathbb{E} [\log \langle \boldsymbol{\varphi}_1, \tilde{\rho}_{n,\omega} \rangle] + o_\omega(1).$$

Proof. Let \mathcal{A} be the event that $\hat{\mathbf{G}}_{n,\omega}$ has at least $\omega/2$ cavities, that all variable nodes have degree either d or $d-1$ and that no two variable nodes of degree $d-1$ are adjacent to the same constraint node. Then $\mathbb{P}[\mathcal{A}] = 1 - \exp(-\Omega_\omega(\omega))$. Hence, (2.1) ensures that

$$\mathbb{E} [\log(Z(\mathbf{G}_{n,\omega}'')/Z(\hat{\mathbf{G}}_{n,\omega}))] = \mathbb{E} [\log(Z(\mathbf{G}_{n,\omega}'')/Z(\hat{\mathbf{G}}_{n,\omega})) | \mathcal{A}] + o_\omega(1). \quad (6.18)$$

Moreover, on \mathcal{A} the random factor graph $\mathbf{G}_{n,\omega}''$ is obtained from $\mathbf{G}_{n,\omega}$ by adding one variable node \hat{v} along with d constraint nodes $\hat{a}_1, \dots, \hat{a}_d$, whose weight functions are drawn from P independently. Further, on the event \mathcal{A} all neighbors of the \hat{a}_i except \hat{v} belong to the set $\hat{\mathcal{C}}$ of cavities. Therefore, we have the exact formula

$$\frac{Z(\mathbf{G}_{n,\omega}'')}{Z(\hat{\mathbf{G}}_{n,\omega})} = \left\langle \sum_{\chi \in \Omega} p(\chi) \prod_{i=1}^d \sum_{\tau \in \Omega^{\partial \hat{a}_i}} \boldsymbol{\psi}_{\hat{a}_i}(\tau) \mathbf{1}\{\tau_{\hat{x}} = \chi, \forall y \in \partial \hat{a}_i \setminus \hat{x} : \boldsymbol{\sigma}_y = \tau_y\}, \mu_{\hat{\mathbf{G}}_{n,\omega}, \hat{\mathcal{C}}} \right\rangle. \quad (6.19)$$

To proceed, let $(\mathbf{v}_{i,j})_{i,j}$ be a sequence of uniformly and independently chosen cavities $\mathbf{v}_{i,j} \in \hat{\mathcal{C}}$. We claim that on the event \mathcal{A} ,

$$\begin{aligned} \mathbb{E} \left[\log \frac{Z(\mathbf{G}_{n,\omega}'')}{Z(\hat{\mathbf{G}}_{n,\omega})} \middle| \hat{\mathbf{G}}_{n,\omega} \right] \\ = o_\omega(1) + \mathbb{E} \left[\log \left\langle \sum_{\chi \in \Omega} p(\chi) \prod_{i=1}^d \sum_{\tau \in \Omega^k} \boldsymbol{\psi}_i(\tau) \mathbf{1}\{\tau_{\mathbf{h}_i} = \chi, \forall h \in [k] \setminus \mathbf{h}_i : \boldsymbol{\sigma}_{\mathbf{v}_{i,h}} = \tau_h\}, \mu_{\hat{\mathbf{G}}_{n,\omega}, \hat{\mathcal{C}}} \right\rangle \middle| \hat{\mathbf{G}}_{n,\omega} \right]. \end{aligned} \quad (6.20)$$

Indeed, the only difference between (6.19) and (6.20) is that in the former the neighbours $\partial \hat{a}_i \setminus \hat{v}$ are chosen from $\hat{\mathcal{C}}$ without replacement, whereas the $\mathbf{v}_{i,j}$ are chosen independently, i.e., with replacement. But since we choose a mere dk cavities $(\mathbf{v}_{i,j})_{i \in [d], j \in [k]}$ out of a total of at least $\omega/2$, the probability of hitting the same cavity twice is $o_\omega(1)$, and thus (6.20) follows from (2.1). Further, unravelling the definitions of $\tilde{\rho}_{n,\omega}$ and $\boldsymbol{\varphi}_1$, we see that

$$\begin{aligned} \mathbb{E} \left[\log \left\langle \sum_{\chi \in \Omega} p(\chi) \prod_{i=1}^d \sum_{\tau \in \Omega^k} \boldsymbol{\psi}_i(\tau) \mathbf{1}\{\tau_{\mathbf{h}_i} = \chi, \forall h \in [k] \setminus \mathbf{h}_i : \boldsymbol{\sigma}_{\mathbf{v}_{i,h}} = \tau_h\}, \mu_{\hat{\mathbf{G}}_{n,\omega}, \hat{\mathcal{C}}} \right\rangle \middle| \hat{\mathbf{G}}_{n,\omega} \right] \\ = \mathbb{E} [\log \langle \boldsymbol{\varphi}_1, \tilde{\rho}_{n,\omega} \rangle | \hat{\mathbf{G}}_{n,\omega}]. \end{aligned} \quad (6.21)$$

Finally, the assertion follows from (6.18)–(6.21) by taking the expectation on $\hat{\mathbf{G}}_{n,\omega}$. \square

Lemma 6.6. *We have*

$$\mathbb{E} \left[\log \frac{Z(\mathbf{G}'_{n,\omega})}{Z(\hat{\mathbf{G}}_{n,\omega})} \right] = \mathbb{E} \left\langle \log \bigoplus_{\hat{M}_{n,\omega} < i \leq M_{n,\omega}} \boldsymbol{\psi}_{1,i}, \hat{\rho}_{n,\omega} \right\rangle + o_\omega(1), \quad \mathbb{E} \left[\log \frac{Z(\mathbf{G}'_{n,\omega})}{Z(\hat{\mathbf{G}}_{n,\omega})} \right] = \mathbb{E} \left\langle \log \bigoplus_{\hat{M}_{n,\omega} < i \leq M_{n,\omega}} \boldsymbol{\psi}_{1,i}, \tilde{\rho}_{n,\omega} \right\rangle + o_\omega(1).$$

Proof. The proof is similar in spirit to the previous one. Once more we consider the event \mathcal{A} that $\hat{\mathbf{G}}_{n,\omega}$ has at least $\omega/2$ cavities, that all variable nodes have degree either d or $d-1$ and that no two variable nodes of degree $d-1$ are adjacent to a common constraint node. Then $\mathbb{P}[\mathcal{A}] = 1 - \exp(-\Omega_\omega(\omega))$ and

$$\mathbb{E} [\log(Z(\mathbf{G}'_{n,\omega})/Z(\hat{\mathbf{G}}_{n,\omega}))] = \mathbb{E} [\log(Z(\mathbf{G}_{n,\omega}'')/Z(\hat{\mathbf{G}}_{n,\omega})) | \mathcal{A}] + o_\omega(1). \quad (6.22)$$

Moreover, we have the pointwise exact formula

$$\frac{Z(\mathbf{G}'_{n,\omega})}{Z(\hat{\mathbf{G}}_{n,\omega})} = \left\langle \prod_{\hat{M}_{n,\omega} < i \leq M_{n,\omega}} \sum_{\tau \in \Omega^{\partial a_i}} \psi_{a_i}(\tau) \mathbf{1}\{\boldsymbol{\sigma}_{\partial a_i} = \tau\}, \mu_{\hat{\mathbf{G}}_{n,\omega}, \hat{\mathcal{C}}} \right\rangle. \quad (6.23)$$

With $(\mathbf{v}_{i,j})_{i,j}$ a sequence of independently chosen cavities $\mathbf{v}_{i,j} \in \hat{\mathcal{C}}$, we claim that on \mathcal{A} ,

$$\mathbb{E} \left[\log \frac{Z(\mathbb{G}'_{n,\omega})}{Z(\hat{\mathbb{G}}_{n,\omega})} \mid \hat{\mathbb{G}}_{n,\omega} \right] = o_\omega(1) + \mathbb{E} \left[\log \left\langle \prod_{i=1}^{M_{n,\omega} - \hat{M}_{n,\omega}} \sum_{\tau \in \Omega^k} \boldsymbol{\psi}_i(\tau) \prod_{j=1}^k \mathbf{1}\{\boldsymbol{\sigma}_{\mathbf{v}_{i,j}} = \tau_j\}, \mu_{\hat{\mathbb{G}}_{n,\omega}, \hat{\mathcal{C}}} \right\rangle \mid \hat{\mathbb{G}}_{n,\omega} \right]. \quad (6.24)$$

Indeed, the only difference is that in (6.24) the $\mathbf{v}_{i,j}$ are chosen independently, whereas in (6.23) the neighbors of the a_i are chosen without replacement. But since $M_{n,\omega} - \hat{M}_{n,\omega}$ is bounded while there are at least $\omega/2$ cavities, the two terms coincide up to $o_\omega(1)$. Finally, the construction of $\tilde{\rho}_{n,\omega}$ ensures that

$$\mathbb{E} \left[\log \left\langle \prod_{i=1}^{M_{n,\omega} - \hat{M}_{n,\omega}} \sum_{\tau \in \Omega^k} \boldsymbol{\psi}_i(\tau) \prod_{j=1}^k \mathbf{1}\{\boldsymbol{\sigma}_{\mathbf{v}_{i,j}} = \tau_j\}, \mu_{\hat{\mathbb{G}}_{n,\omega}, \hat{\mathcal{C}}} \right\rangle \mid \hat{\mathbb{G}}_{n,\omega} \right] = \mathbb{E} \left[\log \left\langle \bigoplus_{i=1}^{M_{n,\omega} - \hat{M}_{n,\omega}} \boldsymbol{\psi}_{1,i}, \hat{\rho}_{n,\omega} \right\rangle \mid \hat{\mathbb{G}}_{n,\omega} \right], \quad (6.25)$$

and thus the assertion follows from (6.22)–(6.25) by taking the expectation. \square

Finally, Proposition 6.2 is an immediate consequence of Lemmas 6.4, 6.5 and 6.6.

6.3. Proof of Proposition 6.3. Once more we will carry the proof out for the simple random factor graph, which is the (slightly) more intricate case; the unconditional case follows by skipping any considerations pertaining to the conditioning. The basic idea behind the proof of Proposition 6.3 is quite simple. With probability $1 - o_\omega(1)$ the random graph $\hat{\mathbb{G}}_{n,\omega}$ consists of $N_{n,\omega}$ variable and $\hat{M}_{n,\omega}$ constraint nodes and we have $N_{n,\omega} = n - \mathbf{X}$ and

$$\hat{M}_{n,\omega} = M_{n+1,\omega} - d = \lfloor dN_{n,\omega}/k \rfloor \wedge (\lfloor dN_{n,\omega}/k \rfloor + \Delta_{n,\omega} - d\mathbf{X} - \mathbf{Y}) - d$$

with independent $\text{Po}(\omega)$ variables \mathbf{X}, \mathbf{Y} . Fix two integers $\ell, \ell' \geq 0$. Given that $\mathbf{X} \geq \ell$ and $k\hat{M}_{n,\omega} \leq dN_{n,\omega} - k\ell' - d(k-1)\ell$, let $\hat{\mathbb{G}}_{n,\omega}[\ell, \ell']$ be the random factor graph obtained from $\hat{\mathbb{G}}_{n,\omega}$ by adding

- ℓ more variable nodes $\hat{v}_1 = v_{N_{n,\omega}+1}, \dots, \hat{v}_\ell = v_{N_{n,\omega}+\ell}$ along with $d\ell$ new constraint nodes $\hat{a}_{i,j}$, $i \in [\ell], j \in [d]$, each with a weight function chosen from P independently; connect a random clone $\hat{\mathbf{h}}_{i,j}$ of each $\hat{a}_{i,j}$ with a random clone of \hat{v}_i and pair the other $k-1$ clones of $\hat{a}_{i,j}$ with random cavities of $\hat{\mathbb{G}}_{n,\omega}$ left pending by the previous additions.
- ℓ' more constraint nodes $\hat{a}_1, \dots, \hat{a}_{\ell'}$, each endowed with a weight function chosen from P independently and each connected with k random cavities of $\hat{\mathbb{G}}_{n,\omega}$ left vacant by the previous operations.

The resulting random factor graph $\hat{\mathbb{G}}_{n,\omega}[\ell, \ell']$ is not necessarily simple. Yet the key insight behind Proposition 6.3 is that for any ℓ, ℓ' the distribution of $\hat{\mathbb{G}}_{n,\omega}[\ell, \ell']$ is close to that of the original graph $\hat{\mathbb{G}}_{n,\omega}$, provided that ω is big enough. Moreover, the perturbation of the Boltzmann distribution that ensues upon going from $\hat{\mathbb{G}}_{n,\omega}$ to $\hat{\mathbb{G}}_{n,\omega}[\ell, \ell']$ is close to the perturbation induced by the $\ast(\ell, \ell')$ -operation. We introduce similar notation $\hat{\mathbb{G}}_{n,\omega}[\ell, \ell']$ for the random graph without the conditioning on \mathcal{S} .

To formalize this idea, we first compare the distributions of $\hat{\mathbb{G}}_{n,\omega}$ and $\hat{\mathbb{G}}_{n,\omega}[\ell, \ell']$. For integers x, y we denote by $\hat{\mathbb{G}}_{n,\omega,x,y}$ the conditional $\hat{\mathbb{G}}_{n,\omega}$ given that $\mathbf{X} = x$ and $\mathbf{Y} = y$.

Lemma 6.7. *For any $\ell, \ell' \geq 0$ we have $d_{\text{TV}}(\hat{\mathbb{G}}_{n,\omega}, \hat{\mathbb{G}}_{n,\omega}[\ell, \ell']) = o_\omega(1)$ and analogously $d_{\text{TV}}(\hat{\mathbb{G}}_{n,\omega}, \hat{\mathbb{G}}_{n,\omega}[\ell, \ell']) = o_\omega(1)$.*

Proof. The event $\mathcal{A} = \{\omega/2 \leq \mathbf{X} \leq 2\omega, \omega/2 \leq \mathbf{Y} \leq 2\omega\}$ has probability $1 - o_\omega(1)$. Further, because \mathbf{X}, \mathbf{Y} are independent Poisson variables with a large mean ω while ℓ, ℓ' are fixed, the total variation distance of the pairs (\mathbf{X}, \mathbf{Y}) and $(\mathbf{X} - \ell, \mathbf{Y} - \ell')$ is of order $O_\omega(\omega^{-1/2})$. Hence, given \mathcal{A} the total variation distance of $\hat{\mathbb{G}}_{n,\omega}$ and $\hat{\mathbb{G}}_{n,\omega, \mathbf{X} - \ell, \mathbf{Y} - \ell'}$ is $o_\omega(1)$; in symbols,

$$d_{\text{TV}}(\hat{\mathbb{G}}_{n,\omega} \mid \mathcal{A}, \hat{\mathbb{G}}_{n,\omega, \mathbf{X} - \ell, \mathbf{Y} - \ell'} \mid \mathcal{A}) = o_\omega(1). \quad (6.26)$$

Further, let \mathcal{A}' be the event that $\hat{\mathbb{G}}_{n,\omega}$ enjoys the following additional properties.

- (i') The last ℓ variable nodes of $\hat{\mathbb{G}}_{n,\omega}$ satisfy $|\partial^2\{v_{N_{n,\omega}-\ell+1}, \dots, v_{N_{n,\omega}}\} \setminus \hat{\mathcal{C}}| = \ell d(k-1)$. Hence, there are $\ell d(k-1)$ distinct second neighbors, none of which is a cavity.
- (ii') The last ℓ' constraint nodes of $\hat{\mathbb{G}}_{n,\omega}$ satisfy $|\partial\{a_{\hat{M}_{n,\omega}-\ell'+1}, \dots, a_{\hat{M}_{n,\omega}}\} \setminus \hat{\mathcal{C}}| = k\ell'$. Hence, there are $k\ell'$ distinct second neighbors, none of them a cavity.
- (iii') We have $\partial\{v_{N_{n,\omega}-\ell+1}, \dots, v_{N_{n,\omega}}\} \cap \{a_{\hat{M}_{n,\omega}-\ell'+1}, \dots, a_{\hat{M}_{n,\omega}}\} = \emptyset$.
- (iv') Let

$$\mathcal{U} = \hat{\mathcal{C}} \cup \partial^2\{v_{N_{n,\omega}-\ell+1}, \dots, v_{N_{n,\omega}}\} \cup \partial\{a_{\hat{M}_{n,\omega}-\ell'+1}, \dots, a_{\hat{M}_{n,\omega}}\}.$$

Then for any constraint node $a \notin \partial\{v_{N_{n,\omega}-\ell+1}, \dots, v_{N_{n,\omega}}\} \cup a_{\hat{M}_{n,\omega}-\ell'+1}, \dots, a_{\hat{M}_{n,\omega}}$ we have $|\partial a \cap \mathcal{U}| \leq 1$. Thus, only the constraint nodes adjacent to the last ℓ variable nodes or the a_i with $i > \hat{M}_{n,\omega} - \ell'$ may be adjacent to more than one variable node in \mathcal{U} .

(v') All variable nodes $u \in \mathcal{U}$ have degree d or $d-1$.

Additionally, let \mathcal{A}'' be the event that $\hat{\mathbb{G}}_{n,\omega, X-\ell, Y-\ell'}$ has the following properties.

(i'') all variable nodes have degree either d or $d-1$.

(ii'') no two variable nodes of degree $d-1$ are adjacent to the same constraint node.

Then

$$\mathbb{P}[\hat{\mathbb{G}}_{n,\omega} \in \mathcal{A}' | \mathcal{A}] = 1 - o_\omega(1), \quad \mathbb{P}[\hat{\mathbb{G}}_{n,\omega, X-\ell, Y-\ell'} \in \mathcal{A}'' | \mathcal{A}] = 1 - o_\omega(1). \quad (6.27)$$

Furthermore, given $\mathcal{A}'' \cap \mathcal{A}$, the random factor graph $\hat{\mathbb{G}}_{n,\omega, X-\ell, Y-\ell'}[\ell, \ell']$ obtained by attaching ℓ new variable nodes and ℓ' new constraint nodes is distributed precisely as $\hat{\mathbb{G}}_{n,\omega}$ given $\mathcal{A}' \cap \mathcal{A}$. Indeed, the construction of the enhanced factor graph $\hat{\mathbb{G}}_{n,\omega, X-\ell, Y-\ell'}[\ell, \ell']$ expressly ensures that (i'')–(iii'') are satisfied, and (iv')–(v') follow from (i'')–(ii''). Hence, (6.27) yields

$$d_{\text{TV}}(\hat{\mathbb{G}}_{n,\omega} | \mathcal{A}, \hat{\mathbb{G}}_{n,\omega, X-\ell, Y-\ell'}[\ell, \ell'] | \mathcal{A}) = o_\omega(1). \quad (6.28)$$

Finally, since $\mathbb{P}[\mathcal{A}] = 1 - o_\omega(1)$, the assertion follows from (6.26) and (6.28). \square

Let $\tilde{\mathcal{C}}[\ell, \ell']$ be the set of cavities of $\hat{\mathbb{G}}_{n,\omega}[\ell, \ell']$ and let $\hat{\rho}_{n,\omega}[\ell, \ell'] \in \mathcal{K}$ be the kernel representing $\mu_{\hat{\mathbb{G}}_{n,\omega}[\ell, \ell'], \tilde{\mathcal{C}}[\ell, \ell']}$. Let $\hat{\pi}_{n,\omega}[\ell, \ell'] \in \mathcal{D}$ be the distribution of $\hat{\rho}_{n,\omega}[\ell, \ell']$. Define $\tilde{\mathcal{C}}[\ell, \ell']$, $\tilde{\rho}_{n,\omega}[\ell, \ell']$ analogously for $\hat{\mathbb{G}}_{n,\omega}$.

Lemma 6.8. *For any $\ell, \ell' \geq 0$ we have*

$$\mathcal{D}_{\square}(\hat{\pi}_{n,\omega}^{*(\ell, \ell')}, \hat{\pi}_{n,\omega}[\ell, \ell']) = o_\omega(1), \quad \mathcal{D}_{\square}(\tilde{\pi}_{n,\omega}^{*(\ell, \ell')}, \tilde{\pi}_{n,\omega}[\ell, \ell']) = o_\omega(1).$$

Proof. The event $\mathcal{A}_1 = \{\omega/2 \leq X, Y \leq 2\omega\}$ occurs with probability $1 - o_\omega(1)$. So does the event \mathcal{A}_2 that all variable nodes of $\hat{\mathbb{G}}_{n,\omega}$ have degree either d or $d-1$, and thus the same is true of $\mathcal{A} = \mathcal{A}_1 \cap \mathcal{A}_2$. Moreover, the construction of $\hat{\mathbb{G}}_{n,\omega}[\ell, \ell']$ is such that on the event \mathcal{A} we have the exact formula

$$\begin{aligned} \frac{Z(\hat{\mathbb{G}}_{n,\omega}[\ell, \ell'])}{Z(\hat{\mathbb{G}}_{n,\omega})} &= \left\langle \prod_{i=1}^{\ell} \varphi_i(\sigma) \prod_{i=1}^{\ell'} \psi_{\hat{a}_i}(\sigma), \mu_{\hat{\mathbb{G}}_{n,\omega}, \tilde{\mathcal{C}}} \right\rangle, \quad \text{where} \\ \varphi_i(\sigma) &= \sum_{\chi \in \Omega} p(\chi) \prod_{j=1}^d \sum_{\tau \in \Omega^{\partial \hat{a}_i}} \psi_{\hat{a}_{i,j}}(\tau) \mathbf{1}\{\tau_{\hat{v}_i} = \chi, \forall w \in \partial \hat{a}_{i,j} \setminus \hat{v}_i : \sigma_w = \tau_w\}. \end{aligned}$$

Consequently, the joint distribution of the cavities $\tilde{\mathcal{C}}[\ell, \ell']$ of $\hat{\mathbb{G}}_{n,\omega}[\ell, \ell']$ reads

$$\mu_{\hat{\mathbb{G}}_{n,\omega}[\ell, \ell'], \tilde{\mathcal{C}}[\ell, \ell']}(\sigma) = \frac{\left\langle \mathbf{1}\{\forall u \in \tilde{\mathcal{C}}[\ell, \ell'] : \sigma_u = \sigma_u\} \prod_{i=1}^{\ell} \varphi_i(\sigma) \prod_{i=1}^{\ell'} \psi_{\hat{a}_i}(\sigma), \mu_{\hat{\mathbb{G}}_{n,\omega}}(\sigma) \right\rangle}{\left\langle \prod_{i=1}^{\ell} \varphi_i(\sigma) \prod_{i=1}^{\ell'} \psi_{\hat{a}_i}(\sigma), \mu_{\hat{\mathbb{G}}_{n,\omega}} \right\rangle} \quad (\sigma \in \Omega^{\tilde{\mathcal{C}}[\ell, \ell']}). \quad (6.29)$$

Thus, with probability $1 - o_\omega(1)$, namely on the event \mathcal{A} , $\tilde{\rho}_{n,\omega}[\ell, \ell']$ is just the kernel representing the right hand side of (6.29). We claim that in this case $\tilde{\rho}_{n,\omega}[\ell, \ell']$ and $\tilde{\rho}_{n,\omega}^{*(\ell, \ell')}$ can be coupled to coincide with probability $1 - o_\omega(1)$. Indeed, the weight functions associated with the \hat{a}_i and the $\hat{a}_{i,j}$ are chosen from P independently, and they are connected to the cavities of $\hat{\mathbb{G}}_{n,\omega}$ by a random pairing. By comparison, we construct $\tilde{\rho}_{n,\omega}^{*(\ell, \ell')}$ by adjoining $\varphi_1, \dots, \varphi_\ell$ and $\psi_1, \dots, \psi_{\ell'}$ that evaluate the kernel $\tilde{\rho}_{n,\omega}$ at independent uniformly random points of the unit interval. Combinatorially, this is equivalent to attaching the new variable and constraint nodes to random cavities chosen with replacement, rather than without replacement as in the construction of $\hat{\mathbb{G}}_{n,\omega}[\ell, \ell']$. But since the number of cavities of $\hat{\mathbb{G}}$ is $\Omega_\omega(\omega)$, the two constructions have total variation distance $o_\omega(1)$. \square

Proof of Proposition 6.3. The proposition is immediate from Lemmas 6.7 and 6.8. \square

6.4. **Proof of Proposition 6.1.** We begin with the following lemma, whose proof is similar to the proof of Lemma 5.5.

Lemma 6.9. *We have*

$$\begin{aligned}\mathbb{E} \left[\log \langle \boldsymbol{\varphi}_1, \hat{\pi}_{n,\omega} \rangle - \log \left\langle \bigoplus_{\hat{M}_{n,\omega} < i \leq M_{n,\omega}} \boldsymbol{\psi}_{1,i}, \hat{\pi}_{n,\omega} \right\rangle \right] &= \mathcal{B}(\hat{\pi}_{n,\omega}) + o_\omega(1), \\ \mathbb{E} \left[\log \langle \boldsymbol{\varphi}_1, \tilde{\pi}_{n,\omega} \rangle - \log \left\langle \bigoplus_{\hat{M}_{n,\omega} < i \leq M_{n,\omega}} \boldsymbol{\psi}_{1,i}, \tilde{\pi}_{n,\omega} \right\rangle \right] &= \mathcal{B}(\tilde{\pi}_{n,\omega}) + o_\omega(1).\end{aligned}$$

Proof. Let $\pi = \hat{\pi}_{n,\omega}$ or $\pi = \tilde{\pi}_{n,\omega}$. Since $\mathbb{E}[M_{n,\omega} - \hat{M}_{n,\omega}] = d(k-1)/k + o_\omega(1)$, due to (2.1) it suffices to show that

$$\mathbb{E} \left[\log \left\langle \bigoplus_{\hat{M}_{n,\omega} < i \leq M_{n,\omega}} \boldsymbol{\psi}_{1,i}, \pi \right\rangle \right] = o_\omega(1) + \mathbb{E}[M_{n,\omega} - \hat{M}_{n,\omega}] \cdot \mathbb{E}[\log \langle \boldsymbol{\psi}_{1,1}, \pi \rangle]. \quad (6.30)$$

Thus, we need to cope with the correlations between $M_{n,\omega} - \hat{M}_{n,\omega}$ and $\hat{\mathbf{G}}_{n,\omega}$ or $\hat{\mathbf{G}}_{n,\omega}$, respectively. In other words, we need to assess the correlations between $M_{n,\omega} - \hat{M}_{n,\omega}$ and $\hat{N}_{n,\omega}, \hat{M}_{n,\omega}$. With probability $1 - \exp(-\Omega_\omega(\omega))$ we have

$$M_{n,\omega} - \hat{M}_{n,\omega} = W, \quad \text{where} \quad W = \lfloor dN_{n,\omega}/k \rfloor - \lfloor d(1+N_{n,\omega})/k \rfloor + \Delta_{n,\omega} - \Delta_{n+1,\omega} + d, \quad (6.31)$$

with independent Bernoulli variables $\Delta_{n,\omega}, \Delta_{n+1,\omega}$. Thus, $W \leq d+1$ and (2.1) ensures that

$$\begin{aligned}\mathbb{E} \left[\log \left\langle \bigoplus_{\hat{M}_{n,\omega} < i \leq M_{n,\omega}} \boldsymbol{\psi}_{1,i}, \pi \right\rangle \right] &= \mathbb{E} \left[\log \left\langle \bigoplus_{i=1}^W \boldsymbol{\psi}_{1,i}, \pi \right\rangle \right] + o_\omega(1) \\ &= \mathbb{E} \left[\mathbf{1} \left\{ \left| \hat{N}_{n,\omega} - \mathbb{E}[\hat{N}_{n,\omega}] \right|, \left| \hat{M}_{n,\omega} - \mathbb{E}[\hat{M}_{n,\omega}] \right| \leq \sqrt{\omega \log \omega} \right\} \log \left\langle \bigoplus_{i=1}^W \boldsymbol{\psi}_{1,i}, \pi \right\rangle \right] + o_\omega(1).\end{aligned} \quad (6.32)$$

Furthermore, since X, Y are independent Poisson variables with mean ω while W is bounded, for any \hat{n}, \hat{m} such that $|\hat{n} - \mathbb{E}[\hat{N}_{n,\omega}]|, |\hat{m} - \mathbb{E}[\hat{M}_{n,\omega}]| \leq \sqrt{\omega \log \omega}$ we obtain from (6.31) that

$$\mathbb{P}[\hat{M}_{n,\omega} = \hat{m} \mid \hat{N}_{n,\omega} = \hat{n}] = (1 + o_\omega(1)) \mathbb{P}[\hat{M}_{n,\omega} = \hat{m} \mid N_{n,\omega} = \hat{n}, W = h] \quad \text{for any } 0 \leq h \leq d+1.$$

Hence, introducing an independent copy W' of W , we obtain from (2.1) and (6.32) that

$$\mathbb{E} \left[\log \left\langle \bigoplus_{\hat{M}_{n,\omega} < i \leq M_{n,\omega}} \boldsymbol{\psi}_{1,i}, \pi \right\rangle \right] = \mathbb{E} \left[\log \left\langle \bigoplus_{i=1}^{W'} \boldsymbol{\psi}_{1,i}, \pi \right\rangle \right] + o_\omega(1). \quad (6.33)$$

Additionally, we claim that for any $0 \leq \omega \leq d+1$,

$$\mathbb{E} \left[\log \frac{\langle \bigoplus_{i=1}^{w+1} \boldsymbol{\psi}_{1,i}, \pi \rangle}{\langle \bigoplus_{i=1}^w \boldsymbol{\psi}_{1,i}, \pi \rangle} \right] = \mathbb{E}[\log \langle \boldsymbol{\psi}_{1,1}, \pi \rangle] + o_\omega(1). \quad (6.34)$$

Indeed, as in the proof of Lemma 5.5 we obtain

$$\mathbb{E} \left[\left(\frac{\langle \bigoplus_{i=1}^{w+1} \boldsymbol{\psi}_i, \pi \rangle}{\langle \bigoplus_{i=1}^w \boldsymbol{\psi}_i, \pi \rangle} \right)^\ell \right] = \mathbb{E} \left[\left\langle \boldsymbol{\psi}_{w+1}, \bigoplus_{i=1}^w \boldsymbol{\psi}_i * \pi \right\rangle^\ell \right]. \quad (6.35)$$

Further, (2.1), Corollary 3.4 and Proposition 6.3 yield

$$\mathbb{E} \left[\left\langle \boldsymbol{\psi}_{w+1}, \bigoplus_{i=1}^w \boldsymbol{\psi}_i * \pi \right\rangle^\ell \right] = \mathbb{E}[\langle \boldsymbol{\psi}_{w+1}, \pi \rangle^\ell] + o_\omega(1). \quad (6.36)$$

As the logarithm can be approximated arbitrarily well by polynomials due to (2.1), (6.34) follows from (6.35)–(6.36). Finally, (6.30) follows from (6.33) and (6.34). \square

Proof of Proposition 6.1. Proposition 6.3 and Lemma 6.9 show that for any $\ell \geq 1$ there exists $\omega_\ell > \omega_{\ell-1}$ such that for all sufficiently large n we have $\hat{\pi}_{n,\omega_\ell} \in \mathfrak{D}_{1/\ell, \ell, \ell}^*$ and

$$\frac{1}{n} \mathbb{E} \log Z(\mathbf{G}) \geq \mathcal{B}(\hat{\pi}_{n,\omega_\ell}) - 1/\ell. \quad (6.37)$$

Since \mathfrak{D} is compact, the sequence $(\hat{\pi}_{n,\omega_\ell})_n$ has a convergent subsequence, whose limit $\hat{\pi}^{(\ell)}$ lies in the closed set $\mathfrak{D}_{1/\ell,\ell,\ell}^*$. Furthermore, because Lemma 3.3 shows that $\mathcal{B}(\cdot)$ is continuous, (6.37) yields

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbf{G}) \geq \mathcal{B}(\hat{\pi}^{(\ell)}) - 1/\ell. \quad (6.38)$$

Additionally, $(\hat{\pi}^{(\ell)})_\ell$ has a subsequence that converges to $\tilde{\pi} \in \mathfrak{D}^* = \bigcap_{\ell} \mathfrak{D}_{1/\ell,\ell,\ell}^*$. Thus, the first assertion follows from (6.38) and the continuity of $\mathcal{B}(\cdot)$ established by Corollary 3.5.

The second assertion concerning the simple random factor graph \mathbb{G} is immediate from the first. Indeed, due to (2.1) a standard application of Azuma's inequality shows that $n^{-0.51} |\log Z(\mathbf{G}) - \mathbb{E} \log Z(\mathbf{G})| \rightarrow 0$ in probability. Hence, Fact 2.2 and Bayes' rule imply that $\mathbb{E} \log Z(\mathbb{G}) - \mathbb{E} \log Z(\mathbf{G}) = o(n)$. \square

6.5. Proof of Theorem 2.6. We begin by showing that the free energy of \mathbf{G} can be expressed in terms of the functional \mathcal{B} applied to $\hat{\pi}_{n,\omega}$ or $\tilde{\pi}_{n,\omega}$, respectively. Once more we will carry the details out for \mathbb{G} ; the unconditioned random factor graph \mathbf{G} is easier to deal with, and the proofs are just obtained from the \mathbb{G} case by dropping any considerations regarding multiple edges.

Lemma 6.10. *If POS is satisfied, then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbf{G}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\hat{\pi}_{n,\omega}), \quad \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbb{G}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\tilde{\pi}_{n,\omega}).$$

Proof. Proposition 6.2 and Lemma 6.9 show that for any $\varepsilon > 0$ there exists ω_0 such that for all $\omega > \omega_0$ there exists n_0 such that for all $n > n_0$ we have $n^{-1} \mathbb{E} \log Z(\mathbb{G}) \geq \mathcal{B}(\tilde{\pi}_{n,\omega}) - \varepsilon$. Hence, for any $\varepsilon > 0$ there is $\omega_0 > 0$ such that for all $\omega > \omega_0$ we have

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbb{G}) \geq \liminf_{n \rightarrow \infty} \mathcal{B}(\tilde{\pi}_{n,\omega}) - \varepsilon. \quad (6.39)$$

Indeed, since Propositions 5.1 and 6.1 show that $(\frac{1}{n} \mathbb{E} \log Z(\mathbb{G}))_n$ converges, (6.39) yields

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbb{G}) = \liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbb{G}) \geq \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\tilde{\pi}_{n,\omega}). \quad (6.40)$$

We are left to prove the converse inequality. The space \mathfrak{D} is compact and separable. Therefore, for any ω the sequence $(\tilde{\pi}_{n,\omega})_n$ has a subsequence that converges to $\tilde{\pi}^{(\omega)} \in \mathfrak{D}$ such that $\liminf_{n \rightarrow \infty} \mathcal{B}(\tilde{\pi}_{n,\omega}) = \mathcal{B}(\tilde{\pi}^{(\omega)})$. Further, $(\tilde{\pi}^{(\omega)})_\omega$ has a subsequence that converges to $\tilde{\pi}^*$ such that

$$\liminf_{\omega \rightarrow \infty} \mathcal{B}(\tilde{\pi}^{(\omega)}) = \mathcal{B}(\tilde{\pi}^*). \quad (6.41)$$

Proposition 6.3 shows that $\tilde{\pi}^* \in \mathfrak{D}^*$. Hence, Proposition 5.1 implies that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbb{G}) \leq \mathcal{B}(\tilde{\pi}^*) = \liminf_{\omega \rightarrow \infty} \mathcal{B}(\tilde{\pi}^{(\omega)}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\tilde{\pi}_{n,\omega}). \quad (6.42)$$

Thus, the assertion follows from (6.40)–(6.42). \square

To proceed we need a small twist on Lemma 6.10. Namely, instead of using $\hat{\mathbb{G}}_{n,\omega}$ as our reference point, we are going to work with $\mathbb{G}_{n,\omega}$. Thus, let \mathcal{C} be the set of cavities of $\mathbb{G}_{n,\omega}$ and let $\rho_{n,\omega,\mathcal{S}} \in \mathfrak{K}$ be the kernel representing $\mu_{\mathbb{G}_{n,\omega},\mathcal{C}}$. Further, let $\pi_{n,\omega,\mathcal{S}} \in \mathfrak{D}$ be the distribution of $\rho_{n,\omega,\mathcal{S}}$. Define $\rho_{n,\omega}, \pi_{n,\omega}$ analogously with respect to $\mathbf{G}_{n,\omega}$.

Corollary 6.11. *If POS is satisfied, then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbf{G}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\pi_{n,\omega}), \quad \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbb{G}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\pi_{n,\omega,\mathcal{S}}).$$

Proof. Since X, Y are Poisson variables with a large mean ω , $\hat{M}_{n,\omega}$ and $M_{n,\omega}$ can be coupled so that both coincide with probability $1 - o_\omega(1)$. This coupling naturally extends to a coupling of $\hat{\mathbb{G}}_{n,\omega}$ and $\mathbb{G}_{n,\omega}$ under which $\mathbb{G}_{n,\omega} = \hat{\mathbb{G}}_{n,\omega}$ with probability $1 - o_\omega(1)$. Consequently, recalling that $\mathcal{D}_\square(\cdot, \cdot)$ stands for the Wasserstein metric on \mathfrak{D} , we have $\mathcal{D}_\square(\pi_{n,\omega,\mathcal{S}}, \tilde{\pi}_{n,\omega}) = o_\omega(1)$. Thus, the assertion follows from the Corollary 3.5. \square

We recall the construction of the kernel $\check{\mu}_{\mathbf{G},X,Y} \in \mathfrak{K}$ from (2.14). Let $\check{\pi}_{n,\omega} \in \mathfrak{D}$ be the distribution of $\check{\mu}_{\mathbf{G},X,Y}$, and define $\check{\mu}_{\mathbf{G},X,Y} \in \mathfrak{K}$, $\check{\pi}_{n,\omega,\mathcal{S}} \in \mathfrak{D}$ analogously with respect to \mathbb{G} . Due to the inevitable divisibility condition required to construct a regular factor graph, these kernels are defined whenever $k|dn$. The following proposition summarizes the main step toward the proof of Theorem 2.6.

Proposition 6.12. *For any $\alpha > 0$ there is $\omega_0 > 0$ such that for every $\omega > \omega_0$ there exists $n_0 > 0$ such that for all $n > n_0$ with $k|dn$ we have*

$$\mathcal{D}_\square(\check{\pi}_{n,\omega}, \pi_{n,\omega}) < \alpha, \quad \mathcal{D}_\square(\check{\pi}_{n,\omega,\mathcal{A}}, \pi_{n,\omega,\mathcal{A}}) < \alpha.$$

To prove Proposition 6.12 we let $\mathcal{V} = \{v_i : i > N_{n,\omega}\}$ and $\mathcal{A} = \{a_i : i > M_{n,\omega}\} \cup \bigcup_{v \in \mathcal{V}} \partial v$ be the sets of variable and constraint nodes, respectively, that are present in \mathbb{G} but not in $\mathbb{G}_{n,\omega}$. Similarly as in Section 6.3, conditioning on the event that $kM_{n,\omega} \leq dN_{n,\omega} - d(k-1)\mathbf{X} - k\mathbf{Y}$, we define an enhanced random factor graph $\mathbb{G}^\#$ by

- adding the variable nodes \mathcal{V} to $\mathbb{G}_{n,\omega}$ along with $d\mathbf{X}$ new constraint nodes $a_{v,j}^\#$, $v \in \mathcal{V}$, $j \in [d]$. Each $a_{v,j}^\#$ is adjacent to v and $k-1$ random cavities of $\mathbb{G}_{n,\omega}$,
- adding \mathbf{Y} more constraint nodes $a_1^\#, \dots, a_{\mathbf{Y}}^\#$, each connected with k random cavities of $\mathbb{G}_{n,\omega}$.

Of course, the cavities in the above construction are drawn without replacement and all weight functions are chosen from P independently. We do *not* require that the outcome $\mathbb{G}^\#$ be simple. Let

$$\mathcal{A}^\# = \{a_{v,j}^\# : v \in \mathcal{V}, j \in [d]\} \cup \{a_i^\# : i \leq \mathbf{Y}\}$$

comprise the new constraint nodes.

Lemma 6.13. *We have $d_{\text{TV}}(\mathbb{G}, \mathbb{G}^\#) = o_\omega(1)$.*

Proof. Similarly as in the proof of Lemma 6.7, we consider the event $\mathcal{E} = \{\omega/2 \leq \mathbf{X} \leq 2\omega, \omega/2 \leq \mathbf{Y} \leq 2\omega\}$, which has probability $1 - o_\omega(1)$. Further, let \mathcal{E}' be the event that \mathbb{G} enjoys the following additional properties.

- (i') We have $|\partial^2 \mathcal{V}| = |\mathcal{V}|d(k-1)$.
- (ii') $|\partial\{a_{M_{n,\omega}+1}, \dots, a_m\}| = k(m - M_{n,\omega})$ and $\partial \mathcal{V} \cap \{a_{M_{n,\omega}}, \dots, a_m\} = \emptyset$.
- (iii') If $a \notin \mathcal{A}$, then a is connected to the set $\partial \mathcal{A}$ by at most one edge.

Additionally, let \mathcal{E}'' be the event that $\mathbb{G}_{n,\omega}$ has the following properties.

- (i'') all variable nodes have degree either d or $d-1$.
- (ii'') no two cavities are adjacent to the same constraint node.

We have

$$\mathbb{P}[\mathcal{E}] = 1 - o_\omega(1), \quad \mathbb{P}[\mathbb{G} \in \mathcal{E}'] = 1 - o(1), \quad \mathbb{P}[\mathbb{G}_{n,\omega} \in \mathcal{E}''] = 1 - o(1). \quad (6.43)$$

Moreover, \mathbb{G} given \mathcal{E}' is distributed precisely as $\mathbb{G}^\#$ given \mathcal{E}'' . Thus, the assertion follows from (6.43). \square

Due to Lemma 6.13 we can apply Theorem 2.5 to $\mathbb{G}^\#$. Let $S_1^\#, \dots, S_\ell^\#$ denote the resulting Bethe state decomposition of $\mathbb{G}^\#$. Let $T_i^\# = S_i^\# \cap \Omega^{V_n \setminus \mathcal{V}}$ for $i \in [\ell]$. Further, we introduce

$$\begin{aligned} \mathbf{z}_{\mathbb{G}^\#, i} &= \left\langle \mathbf{1}\{\boldsymbol{\sigma} \in S_i^\#\} / \sum_{\tau \in \Omega^{\mathcal{V}}} \prod_{v \in \mathcal{V}} p(\tau_v) \prod_{a \in \mathcal{A}^\#} \psi_a(\boldsymbol{\sigma}_{V_n \setminus \mathcal{V}}, \tau), \mu_{\mathbb{G}^\#} \right\rangle, \\ \mu_{\mathbb{G}^\#, i}(\sigma) &= \frac{\mu_{\mathbb{G}^\#}(\sigma) \mathbf{1}\{\sigma \in T_i^\#\}}{\mathbf{z}_{\mathbb{G}^\#, i} \sum_{\tau \in \Omega^{\mathcal{V}}} \prod_{v \in \mathcal{V}} p(\tau_v) \prod_{a \in \mathcal{A}^\#} \psi_a(\sigma, \tau)} \quad (\sigma \in \Omega^{V_n \setminus \mathcal{V}}). \end{aligned} \quad (6.44)$$

Thus, $\mu_{\mathbb{G}^\#, i} \in \mathcal{P}(\Omega^{V_n \setminus \mathcal{V}})$.

Lemma 6.14. *With probability $1 - o_\omega(1)$ the sets $T_1^\#, \dots, T_\ell^\#$ are pairwise disjoint and we have*

$$\mu_{\mathbb{G}^\#, i}(\tau) = \mu_{\mathbb{G}_{n,\omega}}(\tau | T_i^\#) \quad \text{for all } \tau \in T_i^\# \text{ and} \quad \mu_{\mathbb{G}_{n,\omega}}(T_i^\#) = \mathbf{z}_{\mathbb{G}^\#, i} / \sum_{j=1}^{\ell} \mathbf{z}_{\mathbb{G}^\#, j}. \quad (6.45)$$

Proof. We recall from Section 4.1 that the decomposition $S_1^\#, \dots, S_\ell^\#$ is constructed by pinning the values of a random set \mathbf{U}_* of variables to specific spins. Since the size of this set is bounded, with high probability we have $(\mathcal{C} \cup \mathcal{V}) \cap \mathbf{U}_* = \emptyset$. We will prove that in this case, $\mu_{\mathbb{G}^\#, i}(\sigma) = \mu_{\mathbb{G}_{n,\omega}}(\sigma | T_i^\#)$ for all i, σ .

If $(\mathcal{C} \cup \mathcal{V}) \cap \mathbf{U}_* = \emptyset$, then $T_1^\#, \dots, T_\ell^\#$ are pairwise disjoint. Thus, fix $i \in [\ell]$ and $\sigma \in T_i^\#$. Then by the construction of $\mathbb{G}^\#$,

$$\mu_{\mathbb{G}_{n,\omega}}(\sigma) = \frac{Z(\mathbb{G}^\#)}{Z(\mathbb{G}_{n,\omega})} \cdot \frac{\mu_{\mathbb{G}^\#}(\sigma)}{\sum_{\tau \in \Omega^\mathcal{V}} \prod_{v \in \mathcal{V}} p(\tau_v) \prod_{a \in \mathcal{A}^\#} \psi_a(\sigma, \tau)}, \quad (6.46)$$

$$\frac{Z(\mathbb{G}_{n,\omega})}{Z(\mathbb{G}^\#)} = \left\langle 1 / \sum_{\tau \in \Omega^\mathcal{V}} \prod_{v \in \mathcal{V}} p(\tau_v) \prod_{a \in \mathcal{A}^\#} \psi_a(\sigma, \tau), \mu_{\mathbb{G}} \right\rangle = \sum_{j=1}^{\ell} \mathbf{z}_{\mathbb{G}^\#, j}. \quad (6.47)$$

Combining (6.46) and (6.47), we obtain the second identity in (6.45). Further, combining the second part of (6.45) with (6.46) and (6.47), we find

$$\mu_{\mathbb{G}_{n,\omega}}(\sigma \mid T_i^\#) = \frac{\mu_{\mathbb{G}_{n,\omega}}(\sigma)}{\mu_{\mathbb{G}_{n,\omega}}(T_i^\#)} = \frac{\sum_{j=1}^{\ell} \mathbf{z}_{\mathbb{G}^\#, j}}{\mathbf{z}_{\mathbb{G}^\#, i}} \cdot \frac{Z(\mathbb{G}^\#)}{Z(\mathbb{G}_{n,\omega})} \cdot \mathbf{z}_{\mathbb{G}^\#, i} \cdot \mu_{\mathbb{G}^\#, i}(\sigma) = \mu_{\mathbb{G}^\#, i}(\sigma),$$

thereby establishing the first part of (6.45). \square

W.h.p. each cavity $v \in \mathcal{C}$ of $\mathbb{G}_{n,\omega}$ has degree $d - 1$. In this case, we denote by b_v the unique neighbour of v in $\mathbb{G}^\#$ that is not present in $\mathbb{G}_{n,\omega}$. Further, for $i \in [\ell]$ let $\nu_{\mathbb{G},i} \in \mathcal{P}(\Omega^\mathcal{C})$ be the product measure

$$\nu_{\mathbb{G}^\#, i} = \bigotimes_{v \in \mathcal{C}} \mu_{\mathbb{G}^\#, v \rightarrow b_v}(\cdot \mid S_i^\#).$$

In close analogy to the weights introduced in (2.13), we also define

$$\begin{aligned} \check{\mathbf{z}}_i^\# &= \mu_{\mathbb{G}^\#}(S_i^\#) \cdot \prod_{v \in \mathcal{V}} \left(\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\tau_v = \chi\} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{\mathbb{G}^\#, w \rightarrow a}(\tau_w \mid S_i^\#) \right)^{-1} \\ &\quad \cdot \prod_{i=1}^{\ell} \left(\sum_{\tau \in \Omega^{\partial a_i^\#}} \psi_{a_i^\#}(\tau) \prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\tau_w \mid S_i^\#) \right)^{-1}. \end{aligned} \quad (6.48)$$

Lemma 6.15. *With probability $1 - o_\omega(1)$ we have $\sum_{h=1}^{\ell} |\mathbf{z}_{\mathbb{G}^\#, h} - \check{\mathbf{z}}_h^\#| = o(1)$ and*

$$\sum_{h=1}^{\ell} \mu_{\mathbb{G}_{n,\omega}}(T_h^\#) \|\mu_{\mathbb{G}_{n,\omega}, \mathcal{C}}(\cdot \mid T_h^\#) - \nu_{\mathbb{G}^\#, h}\|_{\text{TV}} = o(1).$$

Proof. Fix $h \in [\ell]$ and suppose that $S_h^\#$ is an $o(1)$ -Bethe state, which occurs with probability $1 - o_\omega(1)$ due to Theorem 2.5 and Lemma 6.13. Then by **BS2** w.h.p. we have for any $\sigma \in \Omega^{\mathcal{V} \cup \mathcal{C}}$,

$$\begin{aligned} \mu_{\mathbb{G}^\#}(\sigma \mid S_h^\#) &\sim \prod_{v \in \mathcal{V}} \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a} \mu_{\mathbb{G}^\#, w \rightarrow a}(\sigma_w \mid S_h^\#)}{\sum_{\chi \in \Omega} \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\tau_v = \chi\} \psi_a(\tau) \prod_{w \in \partial a} \mu_{\mathbb{G}^\#, w \rightarrow a}(\tau_w \mid S_h^\#)} \\ &\quad \cdot \prod_{i=1}^{\ell} \frac{\psi_{a_i^\#}(\sigma) \prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\sigma_w \mid S_h^\#)}{\sum_{\tau \in \Omega^{\partial a_i^\#}} \psi_{a_i^\#}(\tau) \prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\tau_w \mid S_h^\#)}. \end{aligned} \quad (6.49)$$

Further, w.h.p. each cavity of $\mathbb{G}_{n,\omega}$ has degree $d - 1$; in this case, denote by c_v the unique neighbor of v in $\mathbb{G}^\#$ that is absent in $\mathbb{G}_{n,\omega}$. Then by (6.49) w.h.p. we have

$$\begin{aligned}
z_{\mathbb{G}^\#, h} &= \mu_{\mathbb{G}^\#}(S_h^\#) \left\langle 1 / \sum_{\tau \in \Omega^\mathcal{V}} \prod_{v \in \mathcal{V}} p(\tau_v) \prod_{a \in \mathcal{A}^\#} \psi_a(\sigma_{V \setminus \mathcal{V}}, \tau), \mu_{\mathbb{G}^\#}(\cdot | S_h^\#) \right\rangle \\
&\sim \mu_{\mathbb{G}^\#}(S_h^\#) \sum_{\sigma \in \Omega^\mathcal{V} \cup \mathcal{C}} \prod_{v \in \mathcal{V}} \frac{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a \setminus v} \mu_{\mathbb{G}^\#, w \rightarrow a}(\sigma_w | S_h^\#)}{p(\sigma_v) \prod_{a \in \partial v} \psi_a(\sigma) \cdot \sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\chi = \tau_v\} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{\mathbb{G}^\#, w \rightarrow a}(\tau_w | S_h^\#)} \\
&\quad \cdot \prod_{i=1}^Y \frac{\psi_{a_i^\#}(\sigma) \prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\sigma_w | S_h^\#)}{\psi_{a_i^\#}(\sigma) \cdot \sum_{\tau \in \Omega^{\partial a_i^\#}} \psi_{a_i^\#}(\tau) \prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\tau_w | S_h^\#)} \\
&= \mu_{\mathbb{G}^\#}(S_h^\#) \prod_{v \in \mathcal{V}} \frac{\sum_{\sigma \in \Omega^{\partial^2 v}} \prod_{a \in \partial v} \prod_{w \in \partial a \setminus v} \mu_{\mathbb{G}^\#, w \rightarrow a}(\sigma_w | S_h^\#)}{\sum_{\chi \in \Omega} p(\chi) \prod_{a \in \partial v} \sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\chi = \tau_v\} \psi_a(\tau) \prod_{w \in \partial a \setminus v} \mu_{\mathbb{G}^\#, w \rightarrow a}(\tau_w | S_h^\#)} \\
&\quad \cdot \prod_{i=1}^Y \frac{\sum_{\sigma \in \Omega^{\partial a_i^\#}} \prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\sigma_w | S_h^\#)}{\psi_{a_i^\#}(\sigma) \cdot \sum_{\tau \in \Omega^{\partial a_i^\#}} \psi_{a_i^\#}(\tau) \prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\tau_w | S_h^\#)} \\
&= \check{z}_h^\#.
\end{aligned} \tag{6.50}$$

Summing on h completes the proof of the first assertion.

With respect to the second assertion, for $\sigma \in \Omega^\mathcal{C}$ we have w.h.p.

$$\begin{aligned}
\mu_{\mathbb{G}_{n,\omega}, \mathcal{C}}(\sigma | T_h^\#) &= \mu_{\mathbb{G}^\#, h}(\sigma) && \text{[by Lemma 6.14]} \\
&= \mu_{\mathbb{G}^\#}(S_h^\#) \cdot \frac{\mu_{\mathbb{G}^\#}(\sigma | S_h^\#)}{z_{\mathbb{G}^\#, h} \sum_{\tau \in \Omega^\mathcal{V}} \prod_{v \in \mathcal{V}} p(\tau_v) \prod_{a \in \mathcal{A}^\#} \psi_a(\sigma, \tau)} && \text{[by (6.44)]} \\
&\sim \frac{\mu_{\mathbb{G}^\#}(S_h^\#)}{z_{\mathbb{G}^\#, h}} \left(\sum_{\tau \in \Omega^\mathcal{V}} \prod_{v \in \mathcal{V}} p(\tau_v) \prod_{a \in \mathcal{A}^\#} \psi_a(\sigma, \tau) \right)^{-1} \\
&\quad \cdot \prod_{i=1}^Y \frac{\psi_{a_i^\#}(\sigma) \prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\sigma_w | S_h^\#)}{\sum_{\tau \in \Omega^{\partial a_i^\#}} \psi_{a_i^\#}(\tau) \prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\tau_w | S_h^\#)} \\
&\quad \cdot \prod_{v \in \mathcal{V}} \frac{p(\sigma_{v_i}) \prod_{a \in \partial v} \psi_a(\sigma) \prod_{w \in \partial a \setminus v} \mu_{\mathbb{G}^\#, w \rightarrow a}(\sigma_w | S_h^\#)}{\sum_{\chi \in \Omega} \prod_{a \in \partial v_i} \sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\tau_v = \chi\} \psi_a(\tau) \prod_{w \in \partial a} \mu_{\mathbb{G}^\#, w \rightarrow a}(\tau_w | S_h^\#)} && \text{[by (6.49)]} \\
&= \frac{\mu_{\mathbb{G}^\#}(S_h^\#)}{z_{\mathbb{G}^\#, h}} \prod_{i=1}^Y \frac{\prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\sigma_w | S_h^\#)}{\sum_{\tau \in \Omega^{\partial a_i^\#}} \psi_{a_i^\#}(\tau) \prod_{w \in \partial a_i^\#} \mu_{\mathbb{G}^\#, w \rightarrow a_i^\#}(\tau_w | S_h^\#)} \\
&\quad \cdot \prod_{v \in \mathcal{V}} \frac{\prod_{a \in \partial v} \prod_{w \in \partial a} \mu_{\mathbb{G}^\#, w \rightarrow a}(\sigma_w | S_h^\#)}{\sum_{\chi \in \Omega} \prod_{a \in \partial v_i} \sum_{\tau \in \Omega^{\partial a}} \mathbf{1}\{\tau_v = \chi\} \psi_a(\tau) \prod_{w \in \partial a} \mu_{\mathbb{G}^\#, w \rightarrow a}(\tau_w | S_h^\#)} \\
&= v_{\mathbb{G}^\#, h}(\sigma), && \text{[by (6.50)]}
\end{aligned}$$

as claimed. \square

Proof of Proposition 6.12. Let $v_{\mathbb{G}^\#} = \sum_{i=1}^\ell \check{z}_i^\# v_{\mathbb{G}^\#, i} / \sum_{i=1}^\ell \check{z}_{\mathbb{G}^\#, i}$ and let $\pi_{n,\omega, \mathcal{S}}^\#$ be the distribution of the kernel representation $\check{v}_{\mathbb{G}^\#} \in \mathfrak{K}$. Then up to a renumbering of the variable and constraint nodes, $\check{\mu}_{\mathbb{G}^\#, X, Y} \in \mathfrak{K}$ is distributed as the representation of $v_{\mathbb{G}^\#}$. Specifically, in (2.14) we renumbered the nodes such that \mathcal{V} comprises the first X variable nodes and such that the $a_i^\#$, $i \in [Y]$, are the first Y constraint nodes. Due to Lemma 6.13 and because \mathbb{G} and \mathbf{G} are invariant under node permutations, we conclude that $\mathcal{D}_\square(\pi_{n,\omega, \mathcal{S}}^\#, \check{\pi}_{n,\omega, \mathcal{S}}) = o_\omega(1)$. Furthermore, combining Lemmas 3.14, 6.14 and 6.15, we see that $\mathbb{E}[\Delta_\square(\mu_{\mathbb{G}_{n,\omega}, \mathcal{C}}, v_{\mathbb{G}^\#})] = o_\omega(1)$. Hence, invoking (3.4), we conclude that $\mathbb{E}[\mathcal{D}_\square(\rho_{n,\omega, \mathcal{S}}, \check{v}_{\mathbb{G}^\#})] = o_\omega(1)$. Thus, the triangle inequality yields $\mathcal{D}_\square(\pi_{n,\omega, \mathcal{S}}, \check{\pi}_{n,\omega, \mathcal{S}}) = o_\omega(1)$. The same argument applies to $\pi_{n,\omega}$ and $\check{\pi}_{n,\omega}$. \square

As a final preparation toward the proof of Theorem 2.6, we need the following simple lemma.

Lemma 6.16. *For any fixed integer ℓ we have*

$$\mathcal{D}_\square(\pi_{n,\omega}, \pi_{n+\ell,\omega}) = o_\omega(1), \quad \mathcal{D}_\square(\pi_{n,\omega,\mathcal{S}}, \pi_{n+\ell,\omega,\mathcal{S}}) = o_\omega(1).$$

Proof. The random factor graph $\mathbb{G}_{n,\omega}$ or $\mathbf{G}_{n,\omega}$, respectively, has $n - \mathbf{X}$ variable nodes with probability $1 - o_\omega(1)$. Similarly, the number of variable nodes of $\mathbb{G}_{n+\ell,\omega}$ or $\mathbf{G}_{n+\ell,\omega}$ is $n + \ell - \mathbf{X}$ with probability $1 - o_\omega(1)$. Since \mathbf{X} is a Poisson variable with mean ω , we have $d_{\text{TV}}(n + \ell - \mathbf{X}, n - \mathbf{X}) = o_\omega(1)$. Hence, we can couple $\mathbb{G}_{n+\ell,\omega}$ and $\mathbb{G}_{n,\omega}$ as well as $\mathbf{G}_{n+\ell,\omega}$ and $\mathbf{G}_{n,\omega}$ in such a way that both coincide w.h.p. This coupling extends to the distributions $\rho_{n,\omega,\mathcal{S}}, \rho_{n+\ell,\omega,\mathcal{S}}$ and $\rho_{n,\omega}, \rho_{n+\ell,\omega}$. \square

Proof of Theorem 2.6. Corollary 6.11 yields the free energy formula in terms of the distributions $\pi_{n,\omega}$ and $\pi_{n,\omega,\mathcal{S}}$, respectively. Furthermore, Proposition 6.12 implies together with Corollary 3.5 that

$$\liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty, k \mid dn} \mathcal{B}(\pi_{n,\omega}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty, k \mid dn} \mathcal{B}(\check{\pi}_{n,\omega}), \quad (6.51)$$

$$\liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty, k \mid dn} \mathcal{B}(\pi_{n,\omega,\mathcal{S}}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty, k \mid dn} \mathcal{B}(\check{\pi}_{n,\omega,\mathcal{S}}), \quad (6.52)$$

with the limit on n confined to integers such that $k \mid dn$ each time. But Lemma 6.16 implies with Corollary 3.5 that this divisibility condition does not alter the limits on the left hand side of these equations, i.e.,

$$\liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty, k \mid dn} \mathcal{B}(\pi_{n,\omega}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\pi_{n,\omega}), \quad (6.53)$$

$$\liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty, k \mid dn} \mathcal{B}(\pi_{n,\omega,\mathcal{S}}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\pi_{n,\omega,\mathcal{S}}). \quad (6.54)$$

Thus, combining (6.51)–(6.54) and invoking Corollary 6.11, we obtain

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbf{G}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\check{\pi}_{n,\omega}), \quad \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \log Z(\mathbb{G}) = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathcal{B}(\check{\pi}_{n,\omega,\mathcal{S}}),$$

where, of course, the limit is confined to n such that $k \mid dn$ because \mathbb{G} , \mathbf{G} and $\check{\pi}_{n,\omega}, \check{\pi}_{n,\omega,\mathcal{S}}$ are defined only for such n ; this is the assertion. \square

7. APPLICATIONS

In Section 7.1 we prove that the spin glass model from Section 1.2 satisfies the condition **POS**; the results stated in Section 1.2 are then immediate from those in Section 2. Further, in Sections 7.2 and 7.3 we apply the results from Section 2 to two further models, the Potts antiferromagnet and the random regular k -SAT model. Finally, in Section 7.4 we show how the theorems from Section 2 can be brought to bear on the hard-core model, thereby proving the results stated in Section 1.3.

7.1. The spin glass. To derive the results on the spin glass model stated in Section 1 from the general theorems in Section 2, we just need to verify the condition **POS** for the spin glass model. In Example 2.3 we introduced the relevant weight function even in the more general case of the k -spin model; the case $k = 2$ corresponds to the spin glass on the Bethe lattice.

Lemma 7.1. *The k -spin model satisfies **POS** for all $d \geq 3, \beta > 0$ and all even $k \geq 2$.*

Proof. The lemma is already implicit in [37, 58]; but let us carry the simple proof out for completeness. Let \mathbf{J} be a standard Gaussian. Upon substituting the weight functions from Example 2.3 into **POS** and multiplying by 2^ℓ , **POS** reads

$$\begin{aligned} & \mathbb{E} \left[\left(1 - \tanh(\beta \mathbf{J}) \int_0^1 \prod_{i=1}^k (2\mu_{s,\mathbf{x}_i} - 1) ds \right)^\ell \right] + (k-1) \mathbb{E} \left[\left(1 - \tanh(\beta \mathbf{J}) \int_0^1 \prod_{i=1}^k (2\mu_{s,\mathbf{x}_i} - 1) ds \right)^\ell \right] \\ & - k \mathbb{E} \left[\left(1 - \tanh(\beta \mathbf{J}) \int_0^1 (2\mu_{s,\mathbf{x}_1} - 1) \prod_{i=2}^k (2\mu'_{s,\mathbf{x}_i} - 1) ds \right)^\ell \right] \geq 0. \end{aligned} \quad (7.1)$$

for all measurable $\mu, \mu' : [0, 1]^2 \rightarrow [0, 1]$. Expanding the first expectation yields

$$\mathbb{E} \left[\left(1 - \tanh(\beta \mathbf{J}) \int_0^1 \prod_{i=1}^k (2\mu_{s,\mathbf{x}_i} - 1) ds \right)^\ell \right] = \sum_{j=0}^{\ell} \binom{\ell}{j} (-1)^j \mathbb{E} \left[\tanh(\beta \mathbf{J})^j \left(\int_0^1 \prod_{i=1}^k (2\mu_{s,\mathbf{x}_i} - 1) ds \right)^j \right]$$

Since \mathbf{J} is independent of the \mathbf{x}_i , the last expectation vanishes if j is odd, while $\tanh(\beta \mathbf{J})^j \geq 0$ if j is even. Thus, in order to establish (7.1) it suffices to show that for any even $j \geq 2$,

$$\mathbb{E} \left[\left(\int_0^1 \prod_{i=1}^k (2\mu_{s,\mathbf{x}_i} - 1) ds \right)^j + (k-1) \left(\int_0^1 \prod_{i=1}^k (2\mu'_{s,\mathbf{x}_i} - 1) ds \right)^j - k \left(\int_0^1 (2\mu_{s,\mathbf{x}_1} - 1) \prod_{i=2}^k (2\mu'_{s,\mathbf{x}_i} - 1) ds \right)^j \right] \geq 0. \quad (7.2)$$

Let $\mathbf{s}_1, \dots, \mathbf{s}_j \in [0, 1]$ be uniformly distribution and mutually independent as well as independent of the \mathbf{x}_i . Then Fubini's theorem yields

$$\mathbb{E} \left[\left(\int_0^1 \prod_{i=1}^k (2\mu_{s,\mathbf{x}_i} - 1) ds \right)^j \right] = \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^j (2\mu_{s_h,\mathbf{x}_1} - 1) \mid \mathbf{s}_1, \dots, \mathbf{s}_j \right]^k \right], \quad (7.3)$$

$$\mathbb{E} \left[\left(\int_0^1 \prod_{i=1}^k (2\mu'_{s,\mathbf{x}_i} - 1) ds \right)^j \right] = \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^j (2\mu'_{s_h,\mathbf{x}_1} - 1) \mid \mathbf{s}_1, \dots, \mathbf{s}_j \right]^k \right], \quad (7.4)$$

$$\mathbb{E} \left[\left(\int_0^1 (2\mu_{s,\mathbf{x}_1} - 1) \prod_{i=2}^k (2\mu'_{s,\mathbf{x}_i} - 1) ds \right)^j \right] = \mathbb{E} \left[\prod_{h=1}^j (2\mu_{s_h,\mathbf{x}_1} - 1) \mid \mathbf{s}_1, \dots, \mathbf{s}_\ell \right] \mathbb{E} \left[\prod_{h=1}^j (2\mu'_{s_h,\mathbf{x}_1} - 1) \mid \mathbf{s}_1, \dots, \mathbf{s}_\ell \right]^{k-1}. \quad (7.5)$$

Since for even k we have $X^k + (k-1)Y^k - kXY^{k-1} \geq 0$ for all $X, Y \in \mathbb{R}$, (7.3)–(7.5) yield (7.2). \square

Due to Lemma 7.1, Theorem 1.1 follows from Theorem 2.5, Theorem 1.2 follows from Theorem 2.6 and Theorem 1.3 follows from Theorem 2.7.

Remark 7.2. *Indeed, together with Lemma 7.1 the results from Section 2 yield the Bethe state decomposition and the corresponding formulas for the free energy for the k -spin model for any even $k \geq 2$.*

7.2. The Potts model. For an integer $q \geq 2$ let $\Omega = \{1, \dots, q\}$ be a set of q distinct colors. Also let $\beta > 0$ be a real parameter, the *inverse temperature*. The *Potts antiferromagnet* on \mathbb{G} is the distribution on Ω^{V_n} defined by

$$\mu_{\mathbb{G}, \beta}(\sigma) = \frac{1}{Z_\beta(\mathbb{G})} \exp \left[-\beta \sum_{1 \leq i < j \leq n} \mathbf{1}\{\nu_i \in \partial \nu_j, \sigma(\nu_i) = \sigma(\nu_j)\} \right], \quad (\sigma \in \Omega^{V_n}),$$

where the partition function $Z_\beta(\mathbb{G})$ provides normalization; we omit the reference to β where possible. Thus, for a given $\sigma \in \Omega^{V_n}$ each monochromatic edge of \mathbb{G} incurs an $\exp(-\beta)$ penalty factor.

The Potts antiferromagnet and the associated optimization problems, the MAX q -CUT problem, are of fundamental importance in combinatorics. Krzakala and Zdeborová [44] brought the cavity method to bear on this model. In the following we show how the main results of the present paper apply to this model to underpin the predictions from [44] rigorously. In particular, we specialize the Belief Propagation equations to the Potts model, work out the variational formula for the free energy and apply this formula to the MAX q -CUT problem on the random regular graph.

The Potts model on $\mathbb{G}(n, d)$ can be cast as a random factor graph model with a single weight function

$$\psi_\beta : \Omega^2 \rightarrow (0, 1), \quad (\sigma, \tau) \mapsto \exp(-\beta \mathbf{1}\{\sigma = \tau\}).$$

Thus, $k = 2$, $\Psi = \{\psi_\beta\}$ and $P(\psi_\beta) = 1$ and the prior distribution p is uniform on Ω . Since the constraints are binary, the random regular factor graph \mathbf{G} can be identified with the usual random d -regular graph \mathbb{G} , with the edges representing the factor nodes.

Lemma 7.3. *The Potts model satisfies condition **POS** for all $\beta > 0$.*

Proof. We plug the definition of ψ_β into **POS** and notice that the $1 - e^{-\beta}$ factors cancel. Hence, the desired inequality reads

$$\mathbb{E} \left[\left(\sum_{\sigma=1}^q \int_0^1 \mu_{s,\mathbf{x}_1}(\sigma) \mu_{s,\mathbf{x}_2}(\sigma) \right)^\ell + \left(\sum_{\sigma=1}^q \int_0^1 \mu'_{s,\mathbf{x}_1}(\sigma) \mu'_{s,\mathbf{x}_2}(\sigma) \right)^\ell - 2 \left(\sum_{\sigma=1}^q \int_0^1 \mu_{s,\mathbf{x}_1}(\sigma) \mu'_{s,\mathbf{x}_2}(\sigma) \right)^\ell \right] \geq 0 \quad (\mu, \mu' \in \mathcal{R}). \quad (7.6)$$

Applying Fubini's theorem to take the expectation on $\mathbf{x}_1, \mathbf{x}_2$ inside, we find

$$\begin{aligned} \mathbb{E} \left[\left(\sum_{\sigma=1}^q \int_0^1 \mu_{s, \mathbf{x}_1}(\sigma) \mu_{s, \mathbf{x}_2}(\sigma) \right)^\ell \right] &= \sum_{\sigma_1, \dots, \sigma_\ell=1}^q \mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, \mathbf{x}_1}(\sigma_h) \mu_{s_h, \mathbf{x}_2}(\sigma_h) \right] \\ &= \sum_{\sigma_1, \dots, \sigma_\ell=1}^q \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, \mathbf{x}_1}(\sigma_h) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right] \mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, \mathbf{x}_2}(\sigma_h) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right] \right] \\ &= \sum_{\sigma_1, \dots, \sigma_\ell=1}^q \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, \mathbf{x}_1}(\sigma_h) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right]^2 \right]. \end{aligned} \quad (7.7)$$

Similar manipulations yield

$$\mathbb{E} \left[\left(\sum_{\sigma=1}^q \int_0^1 \mu'_{s, \mathbf{x}_1}(\sigma) \mu'_{s, \mathbf{x}_2}(\sigma) \right)^\ell \right] = \sum_{\sigma_1, \dots, \sigma_\ell=1}^q \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu'_{s_h, \mathbf{x}_1}(\sigma_h) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right]^2 \right], \quad (7.8)$$

$$\mathbb{E} \left[\left(\sum_{\sigma=1}^q \int_0^1 \mu_{s, \mathbf{x}_1}(1) \mu'_{s, \mathbf{x}_2}(1) \right)^\ell \right] = \sum_{\sigma_1, \dots, \sigma_\ell=1}^q \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, \mathbf{x}_1}(\sigma_h) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right] \mathbb{E} \left[\prod_{h=1}^{\ell} \mu'_{s_h, \mathbf{x}_1}(\sigma_h) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right] \right]. \quad (7.9)$$

Combining (7.7)–(7.9), we conclude that the l.h.s. of (7.6) is the expectation of a sum of squares, and thus non-negative. \square

The message space $\mathcal{S}(\mathbb{G})$ of the Potts model boils down to the set of all families $(\mu_{v \rightarrow w})_{v \in V_n, w \in \partial v}$, with $\mu_{v \rightarrow w} \in \mathcal{P}(\Omega)$. With this simplification the Belief Propagation operator $\text{BP} : \mathcal{S}(\mathbb{G}) \rightarrow \mathcal{S}(\mathbb{G})$, $v \mapsto \hat{v}$ of the Potts model reads

$$\hat{v}_{v \rightarrow u}(\sigma) = \frac{\prod_{w \in \partial v \setminus u} 1 - (1 - e^{-\beta}) \mu_{w \rightarrow v}(\sigma)}{\sum_{\tau \in \Omega} \prod_{w \in \partial v \setminus u} 1 - (1 - e^{-\beta}) \mu_{w \rightarrow v}(\tau)} \quad (\sigma \in \Omega). \quad (7.10)$$

With respect to Bethe states, we expect that the phase space Ω^n decomposes into S_1, \dots, S_ℓ such that the conditional distribution $\mu_{\mathbb{G}}[\cdot | S_i]$ are free of long-range correlations, that their standard messages form an approximate fixed point of BP and that the conditional marginals derive from the messages. In formulas, with high probability over the choice of the graph and with $(\hat{\mu}_{\mathbb{G}, v \rightarrow u}[\cdot | S_h])_{u \in \partial v} = \text{BP}(\mu_{\mathbb{G}, v \rightarrow u}[\cdot | S_h])_{u \in \partial v}$, we aim to show that

$$\frac{1}{n^2} \sum_{1 \leq i < j \leq n} \left\| \mu_{\mathbb{G}, v_i, v_j}[\cdot | S_h] - \mu_{\mathbb{G}, v_i}[\cdot | S_h] \otimes \mu_{\mathbb{G}, v_j}[\cdot | S_h] \right\|_{\text{TV}} = o(1), \quad (7.11)$$

$$\frac{1}{n} \sum_{i=1}^n \sum_{u \in \partial v_i} \left\| \mu_{\mathbb{G}, v_i \rightarrow u}[\cdot | S_h] - \hat{\mu}_{\mathbb{G}, v_i \rightarrow u}[\cdot | S_h] \right\|_{\text{TV}} = o(1), \quad (7.12)$$

$$\frac{1}{n} \sum_{i=1}^n \sum_{\sigma \in \Omega} \left| \mu_{\mathbb{G}, v_i}[\sigma | S_h] - \frac{\prod_{w \in \partial v_i} 1 - (1 - e^{-\beta}) \mu_{\mathbb{G}, v_i \rightarrow w}(\sigma)}{\sum_{\tau \in \Omega} \prod_{w \in \partial v_i} 1 - (1 - e^{-\beta}) \mu_{\mathbb{G}, v_i \rightarrow w}(\tau)} \right| = o(1). \quad (7.13)$$

The following theorem establishes these facts.

Theorem 7.4. *For any sequence $L = L(n) \rightarrow \infty$ and all $d \geq 3$, $\beta > 0$ the following is true. With high probability \mathbb{G} admits a decomposition S_0, S_1, \dots, S_ℓ , $\ell \leq L$, of the phase space Ω^n such that $\mu_{\mathbb{G}}(S_0) = o(1)$ and such that (7.11)–(7.13) are satisfied for $h = 1, \dots, \ell$.*

Proof. This is immediate from Theorem 2.5 applied to the factor graph representation of the Potts model. \square

With respect to the free energy, let \mathbf{X}, \mathbf{Y} be two independent Poisson variables with mean ω . Let $\mathbf{u}_1, \dots, \mathbf{u}_X$ and $\mathbf{v}_1 \mathbf{w}_1, \dots, \mathbf{v}_Y \mathbf{w}_Y$ be uniformly random vertices and edges of \mathbb{G} , chosen independently. With S_1, \dots, S_ℓ the decomposition from Theorem 7.4, we introduce the weights

$$\mathbf{z}_{\mathbb{G}, h} = \mu_{\mathbb{G}}(S_h) \prod_{i=1}^X \left(\sum_{\sigma \in \Omega} \prod_{v \in \partial \mathbf{u}_i} 1 - (1 - e^{-\beta}) \mu_{\mathbb{G}, v \rightarrow \mathbf{u}_i}(\sigma | S_h) \right)^{-1} \prod_{i=1}^Y \left(1 - (1 - e^{-\beta}) \sum_{\sigma \in \Omega} \mu_{\mathbb{G}, \mathbf{v}_i \rightarrow \mathbf{w}_i}(\sigma | S_h) \mu_{\mathbb{G}, \mathbf{w}_i \rightarrow \mathbf{v}_i}(\sigma | S_h) \right)^{-1}$$

and $\mathbf{z}_{\mathbb{G}} = \sum_{h=1}^\ell \mathbf{z}_{\mathbb{G}, h}$. Further, let $\mathcal{C}(\mathbb{G})$ be the set of all vertices of degree less than d in the graph obtained from \mathbb{G} by removing $\mathbf{v}_1, \dots, \mathbf{v}_X$ and $\mathbf{v}_1 \mathbf{w}_1, \dots, \mathbf{v}_Y \mathbf{w}_Y$. Then with high probability each $c \in \mathcal{C}(\mathbb{G})$ has degree precisely $d-1$,

and we write c' for the missing d 'th neighbor of c . Then with $\mathbf{c}_1, \mathbf{c}_2, \dots$ a sequence of uniformly and independently chosen elements of $\mathcal{C}(\mathbb{G})$, we let

$$\mathcal{B}(\mathbb{G}) = \mathbb{E} \left[\log \sum_{h=1}^{\ell} \frac{\mathbf{z}_{\mathbb{G},h}}{\mathbf{z}_{\mathbb{G}}} \sum_{\sigma \in \Omega} \prod_{i=1}^d 1 - (1 - e^{-\beta}) \mu_{\mathbb{G}, \mathbf{c}_i \rightarrow \mathbf{c}'_i}(\sigma) - \frac{d}{2} \log \sum_{h=1}^{\ell} \frac{\mathbf{z}_{\mathbb{G},h}}{\mathbf{z}_{\mathbb{G}}} 1 - (1 - e^{-\beta}) \sum_{\sigma \in \Omega} \mu_{\mathbb{G}, \mathbf{c}_1 \rightarrow \mathbf{c}'_1}(\sigma) \mu_{\mathbb{G}, \mathbf{c}_2 \rightarrow \mathbf{c}'_2}(\sigma) \middle| \mathbb{G} \right].$$

Theorem 7.5. *For all $d \geq 3, \beta > 0$ we have $\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbb{G})] = \liminf_{\omega \rightarrow \infty} \liminf_{n \rightarrow \infty} \mathbb{E}[\mathcal{B}(\mathbb{G})]$.*

Proof. This is an immediate consequence of Theorem 2.6 and Lemma 7.3. \square

Additionally, Theorem 2.7 yields a variational formula for the free energy. Writing out the specifics of the Potts case, we see that \mathfrak{D}^* consists of all $\pi \in \mathfrak{D}$ that satisfy the following property. For a measurable $\mu : [0, 1]^2 \rightarrow \mathcal{P}(\Omega)$ with $\Omega = [q]$ and integers $N, M \geq 0$ let

$$\mathbf{z}_{\mu, s}^{N, M} = \prod_{i=1}^N \left(\sum_{\sigma=1}^q \prod_{j=1}^d 1 - (1 - e^{-\beta}) \mu_{s, \mathbf{x}_{i,j}}(\sigma) \right) \prod_{i=1}^M \left(1 - (1 - e^{-\beta}) \sum_{\sigma=1}^q \mu_{s, \mathbf{x}_{i+N,1}}(\sigma) \mu_{s, \mathbf{x}_{i+N,2}}(\sigma) \right), \quad \text{and set}$$

$$\mathbf{t} = \mathbf{t}(s) = \inf \left\{ u \in [0, 1] : \int_0^u \mathbf{z}_{\mu, u}^{N, M} ds \geq s \int_0^1 \mathbf{z}_{\mu, u}^{N, M} du \right\}.$$

Then we let $\mu_{s, x}^{*(N, M)} = \mu_{\mathbf{t}, x}$. Now \mathfrak{D}_β^* is the set of all $\pi \in \mathfrak{D}$ such that for a random $\mu^\pi \in \mathfrak{K}$ drawn from π , the perturbed $\mu^{\pi*(N, M)} \in \mathfrak{K}$ again has distribution π . Furthermore, in the Potts model the functional $\mathcal{B}(\cdot)$ reads

$$\mathcal{B}_\beta(\pi) = \mathbb{E} \left[\log \left(\sum_{\sigma=1}^q \int_0^1 \prod_{j=1}^d 1 - (1 - e^{-\beta}) \mu_{s, \mathbf{x}_{1,j}}(\sigma) ds \right) - \frac{d}{2} \log \left(1 - (1 - e^{-\beta}) \sum_{\sigma=1}^q \int_0^1 \mu_{s, \mathbf{x}_{1,1}}(\sigma) \mu_{s, \mathbf{x}_{1,2}}(\sigma) ds \right) \right].$$

Theorem 7.6. *For all $d \geq 3, \beta > 0$ we have*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbb{G})] = \Phi_{d, \beta} \quad \text{with} \quad \Phi_{d, \beta} = \inf_{\pi \in \mathfrak{D}_\beta^*} \mathcal{B}_\beta(\pi).$$

As a further application we obtain a variational formula for the MAX q -CUT of the random regular graph, which is defined as

$$\text{MC}_q(\mathbb{G}) = \frac{dn}{2} - \frac{1}{2} \min_{\sigma : [n] \rightarrow [q]} \sum_{v, w=1}^n \mathbf{1}\{w \in \partial v, \sigma(v) = \sigma(w)\}. \quad (7.14)$$

Thus, $\text{MC}_q(\mathbb{G})$ equals the total number of edges of \mathbb{G} minus the ground state energy of the Potts model. In other words, $\text{MC}_q(\mathbb{G})$ is the maximum, over the choice of $\sigma : [n] \rightarrow [q]$, of the number of edges that link vertices of different colors. The MAX q -CUT problem is well-studied in combinatorics and computer science. In particular, the problem is well known to be NP-hard on worst-case instances.

Corollary 7.7. *For all $d \geq 3$ we have $\text{MC}_q(\mathbb{G})/n \xrightarrow{n \rightarrow \infty} \frac{d}{2} + \lim_{\beta \rightarrow \infty} \Phi_{d, \beta+1} - \Phi_{d, \beta}$ in probability.*

Proof. Since Azuma's inequality shows that $\text{MC}_q(\mathbb{G})$ is concentrated within $O(\sqrt{n \log n})$ about its mean, it suffices to prove that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\text{MC}_q(\mathbb{G})] = \frac{d}{2} + \lim_{\beta \rightarrow \infty} \Phi_{d, \beta+1} - \Phi_{d, \beta}. \quad (7.15)$$

Further, introducing $\mathcal{H}_{\mathbb{G}}(\sigma) = \frac{1}{2} \sum_{v, w=1}^n \mathbf{1}\{w \in \partial v, \sigma(v) = \sigma(w)\}$ and recalling (7.14), we can rewrite (7.15) as

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\min_{\sigma : [n] \rightarrow [q]} \mathcal{H}_{\mathbb{G}}(\sigma) \right] = \lim_{\beta \rightarrow \infty} \Phi_{d, \beta} - \Phi_{d, \beta+1}. \quad (7.16)$$

To prove (7.16) we write $\mu_{G, \beta} \in \mathcal{P}([q]^{V(G)})$ for the Potts distribution induced by a d -regular graph $G = (V(G), E(G))$. Moreover, let us denote the Potts Hamiltonian by \mathcal{H}_G and the partition function by $Z_\beta(G)$. It is well known that for any $\varepsilon > 0$ there exists $\beta_0(\varepsilon) > 0$ such that for all $\beta > \beta_0(\varepsilon)$ and all d -regular graphs G we have

$$\langle \mathcal{H}_G, \mu_{G, \beta} \rangle - \varepsilon |V(G)| \leq \min_{\sigma : V(G) \rightarrow [q]} \mathcal{H}_G(\sigma) \leq \langle \mathcal{H}_G, \mu_{G, \beta} \rangle. \quad (7.17)$$

Consequently, for all $\beta > \beta_0(\varepsilon)$ we have

$$\int_{\beta}^{\beta+1} \langle \mathcal{H}_G, \mu_{G,b} \rangle db - \varepsilon |V(G)| \leq \min_{\sigma:V(G) \rightarrow [q]} \mathcal{H}_G(\sigma) \leq \int_{\beta}^{\beta+1} \langle \mathcal{H}_G, \mu_{G,b} \rangle db. \quad (7.18)$$

Since $\langle \mathcal{H}_G, \mu_{G,\beta} \rangle = -\frac{\partial}{\partial \beta} \log Z_\beta(G)$, (7.18) yields

$$\log Z_\beta(G) - \log Z_{\beta+1}(G) - \varepsilon |V(G)| \leq \min_{\sigma:V(G) \rightarrow [q]} \mathcal{H}_G(\sigma) \leq \log Z_\beta(G) - \log Z_{\beta+1}(G). \quad (7.19)$$

Applying (7.19) to the random regular graph \mathbb{G} and taking expectations, we obtain

$$\frac{1}{n} \mathbb{E}[\log Z_\beta(\mathbb{G})] - \frac{1}{n} \mathbb{E}[\log Z_{\beta+1}(\mathbb{G})] - \varepsilon \leq \frac{1}{n} \mathbb{E} \left[\min_{\sigma:[n] \rightarrow [q]} \mathcal{H}_{\mathbb{G}}(\sigma) \right] \leq \frac{1}{n} \mathbb{E}[\log Z_\beta(\mathbb{G})] - \frac{1}{n} \mathbb{E}[\log Z_{\beta+1}(\mathbb{G})]. \quad (7.20)$$

Hence, taking $n \rightarrow \infty$, we obtain for all $\beta > \beta_0(\varepsilon)$,

$$\Phi_{d,\beta} - \Phi_{d,\beta+1} - \varepsilon \leq \liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\min_{\sigma:[n] \rightarrow [q]} \mathcal{H}_{\mathbb{G}}(\sigma) \right] \leq \limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\min_{\sigma:[n] \rightarrow [q]} \mathcal{H}_{\mathbb{G}}(\sigma) \right] \leq \Phi_{d,\beta} - \Phi_{d,\beta+1}. \quad (7.21)$$

Finally, there exists a subsequence (n_l) along which $\mathbb{E}[\min_{\sigma:[n_l] \rightarrow [q]} \mathcal{H}_{\mathbb{G}(n_l,d)}(\sigma)] / n_l$ converges to a number $\xi \geq 0$. Taking the limit of (7.20) along this subsequence, we obtain $\xi \leq \Phi_{d,\beta} - \Phi_{d,\beta+1} \leq \xi + \varepsilon$ for all $\beta > \beta_0(\varepsilon)$. Consequently, the limit $\lim_{\beta \rightarrow \infty} \Phi_{d,\beta} - \Phi_{d,\beta+1}$ exists. Therefore, taking $\beta \rightarrow \infty$ in (7.21), we conclude that

$$\lim_{n \rightarrow \infty} n^{-1} \mathbb{E} \left[\min_{\sigma:[n] \rightarrow [q]} \mathcal{H}_{\mathbb{G}}(\sigma) \right]$$

exists as well and that (7.16) is satisfied. \square

7.3. The regular k -SAT model. The k -SAT problem plays a major role in computer science, particularly in computational complexity theory. In its optimization version, known as the MAX k -SAT problem asks for the largest number of clauses of a propositional formula in conjunctive normal form with clauses of length k that can be satisfied simultaneously. Random instances of k -SAT and MAX k -SAT have been studied extensively as instructive benchmarks [6].

We can express the MAX k -SAT problem as a factor graph model with spins $\Omega = \{-1, 1\}$ corresponding to the Boolean values ‘true’ and ‘false’ as follows. With $k \geq 2$ an integer and $\beta > 0$ be a real parameter, we introduce the weight functions

$$\psi_{\beta,\chi} : \{\pm 1\}^k \rightarrow (0, 1), \quad \sigma \mapsto \frac{1 - \tanh(\beta \prod_{i=1}^k \chi_i \sigma_i)}{2} \quad (\chi \in \{\pm 1\}^k).$$

Let p be the uniform distribution on Ω and let P be uniform on $\Psi_\beta = \{\psi_{\beta,\chi} : \chi \in \Omega^k\}$. In terms of propositional formulas, the semantics is that $\psi_{\beta,\chi}$ encodes a k -clause whose i th literal is negated if $\chi_i = 1$ and positive if $\chi_i = -1$. Thus, $\prod_{i=1}^k \chi_i \sigma_i = 1$ if the truth assignment σ fails to satisfy the clause, and $\prod_{i=1}^k \chi_i \sigma_i = -1$ otherwise. In effect, $\psi_{\beta,s}(\sigma) = (1 - \tanh \beta)/2 \rightarrow 0$ as $\beta \rightarrow \infty$ if σ fails to satisfy the clause, whereas $\psi_{\beta,s}(\sigma) = (1 + \tanh \beta)/2 \rightarrow 1$ if σ is satisfying. Hence, the random factor graph \mathbf{G} models a random k -SAT formula in which every variable appears precisely d times, the *regular k -SAT model*. We are going to derive variational formulas for its free energy and its ground state energy.

Lemma 7.8. *The regular k -SAT model satisfies **POS** for all $d, k \geq 3$ and all $\beta > 0$.*

Proof. Once more this is already implicit in [37, 58], but we carry out the argument here for completeness. Let us write χ for a uniformly random element of $\{\pm 1\}^k$. Substituting $\psi_{\beta,\chi}$ into **POS** and cancelling positive constants, we are left to verify the inequality

$$\mathbb{E} \left[\left(\int_0^1 \prod_{i=1}^k \mu_{s,\chi_i}(\chi_i) ds \right)^\ell + \left(\int_0^1 \prod_{i=1}^k \mu'_{s,\chi_i}(\chi_i) ds \right)^\ell - (k-1) \left(\int_0^1 \mu_{s,\chi_1}(\chi_1) \prod_{i=2}^k \mu'_{s,\chi_2}(\chi_i) ds \right)^\ell \right] \geq 0 \quad (\mu, \mu' \in \mathfrak{K}). \quad (7.22)$$

Fubini's theorem yields

$$\mathbb{E} \left[\left(\int_0^1 \prod_{i=1}^k \mu_{s, \mathbf{x}_i}(\chi_i) ds \right)^\ell \right] = \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, \mathbf{x}_1}(\chi_1) \middle| s_1, \dots, s_\ell \right]^k \right], \quad (7.23)$$

$$\mathbb{E} \left[\left(\int_0^1 \prod_{i=1}^k \mu'_{s, \mathbf{x}_i}(\chi_i) ds \right)^\ell \right] = \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu'_{s_h, \mathbf{x}_1}(\chi_1) \middle| s_1, \dots, s_\ell \right]^k \right], \quad (7.24)$$

$$\mathbb{E} \left[\left(\int_0^1 \mu_{s, \mathbf{x}_1}(\chi_2) \prod_{i=2}^k \mu'_{s, \mathbf{x}_2}(\chi_i) ds \right)^\ell \right] = \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, \mathbf{x}_1}(\chi_1) \middle| s_1, \dots, s_\ell \right] \mathbb{E} \left[\prod_{h=1}^{\ell} \mu'_{s_h, \mathbf{x}_1}(\chi_2) \middle| s_1, \dots, s_\ell \right]^{k-1} \right]. \quad (7.25)$$

Since $X^k + (k-1)Y^k - kXY^{k-1} \geq 0$ for $X, Y \geq 0$, (7.23)–(7.25) yield (7.22). \square

Due to Lemma 7.8 we can bring the results from Section 2 to bear on the random regular k -SAT model. Specifically, for a measurable $\mu : [0, 1]^2 \rightarrow \mathcal{P}(\Omega)$ with $\Omega = \{\pm 1\}$ and integers $N, M \geq 0$ let $(\chi_{i,j})_{i,j \geq 1}$ be independent uniformly random elements of Ω and let

$$z_{\mu, s}^{N, M} = \prod_{i=1}^N \left(\sum_{\sigma \in \Omega} \prod_{j=1}^d 1 - \tanh(\beta) \sum_{\tau \in \Omega^{k-1}} \chi_{i,k} \sigma \prod_{j=1}^{k-1} \chi_{i,j} \tau_j \mu_{s, \mathbf{x}_{i,j}}(\tau_j) \right) \prod_{i=1}^M \left(1 - \tanh(\beta) \sum_{\tau \in \Omega^k} \prod_{j=1}^k \chi_{i+N,j} \tau_j \mu_{s, \mathbf{x}_{i+N,j}}(\tau_j) \right).$$

Further, let

$$\mathbf{t} = \mathbf{t}(s) = \inf \left\{ u \in [0, 1] : \int_0^u z_{\mu, u}^{N, M} ds \geq s \int_0^1 z_{\mu, u}^{N, M} du \right\}.$$

and $\mu_{s, x}^{N, M} = \mu_{\mathbf{t}, x}$. Then \mathcal{D}_β^* consists of all $\pi \in \mathcal{D}$ such that $\mu^{\pi, N, M}$ has distribution π . Furthermore, the functional $\mathcal{B}(\cdot)$ reads

$$\begin{aligned} \mathcal{B}_\beta(\pi) &= \mathbb{E} \left[\log \left(\sum_{\sigma \in \Omega} \int_0^1 \sum_{\sigma \in \Omega} \prod_{j=1}^d 1 - \tanh(\beta) \sum_{\tau \in \Omega^{k-1}} \chi_{i,k} \sigma \prod_{j=1}^{k-1} \chi_{i,j} \tau_j \mu_{s, \mathbf{x}_{i,j}}(\tau_j) ds \right) \right] \\ &\quad + \frac{d(k-1)}{k} \mathbb{E} \left[\left(\sum_{\tau \in \Omega^k} 1 - \tanh \beta \int_0^1 \prod_{j=1}^k \chi_{1,j} \tau_j \mu_{s, \mathbf{x}_{1,j}}(\tau_j) ds \right) \right] - dk \log 2. \end{aligned}$$

Let

$$\Phi_{d, \beta} = \inf_{\pi \in \mathcal{D}_\beta^*} \mathcal{B}_\beta(\pi).$$

Theorem 7.9. For all $d, k \geq 3, \beta > 0$ we have $\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbf{G})] = \Phi_{d, \beta}$.

Proof. This follows immediately from Theorem 2.7 and Lemma 7.8. \square

As a further application we also obtain a variational formula for the MAX k -SAT problem. Specifically, with the interpretation of $\sigma \in \Omega^n$ as a truth assignment, define $\mathcal{H}_G(\sigma)$ as the number of propositional clauses of \mathbf{G} that σ fails to satisfy. Further, let $\text{OPT}(\mathbf{G}) = dn/k - \min_{\sigma \in \Omega^k} \mathcal{H}_G(\sigma)$ be the maximum number of clauses that can be satisfied simultaneously. Following the steps of the proof of Corollary 7.7 precisely, we obtain the following result.

Corollary 7.10. For all $d, k \geq 3$ we have

$$\frac{1}{n} \text{OPT}(\mathbf{G}) \xrightarrow{n \rightarrow \infty} \frac{d}{k} + \lim_{\beta \rightarrow \infty} \Phi_{d, \beta+1} - \Phi_{d, \beta} \quad \text{in probability.}$$

7.4. The hard-core model. The proofs of Theorem 1.4 and Corollary 1.5 are not entirely straightforward because the hard-core model cannot be cast directly as a factor graph model as in Section 2. This is because of the 'hard' constraint that $\sigma_v \sigma_w = 0$ for any adjacent v, w . We therefore prove Theorem 1.4 and Corollary 1.5 by way of a relaxed 'soft-core model' and taking two limits, first in the 'softness' and then in the fugacity. Specifically, we obtain a random factor graph model with $\Omega = \{0, 1\}$ and the prior $p(0) = 1/(1+\lambda)$ and $p(1) = \lambda/(1+\lambda)$. In addition, to mimic the hard-core constraints we would like to introduce a binary weight function that forbids its two adjacent variable nodes from both taking the spin 1. But since it would take values $\{0, 1\}$, we instead introduce

$$\psi_\beta : \Omega^2 \rightarrow (0, 1), \quad (\sigma_1, \sigma_2) \mapsto 1 - (1 - e^{-\beta}) \sigma_1 \sigma_2.$$

Thus, $\beta > 0$ is a ‘softness parameter’, and upon taking $\beta \rightarrow \infty$ we recover the hard-core constraint: $\psi_\infty(\sigma_1, \sigma_2) = 1 - \sigma_1 \sigma_2$. For any β, λ and $d \geq 3$ we obtain the random factor factor graph model $\mathbf{G}_{\lambda, \beta}$ with the single binary weight function ψ_β .

Lemma 7.11. *The model $\mathbf{G}_{\lambda, \beta}$ satisfies **POS** for all $d \geq 3, \lambda > 0, \beta \in (0, \infty]$.*

Proof. Substituting ψ_β into **POS** and noticing that $1 - e^{-\beta} > 0$, we see that it suffices to verify the inequality

$$\mathbb{E} \left[\left(\int_0^1 \mu_{s, x_1}(1) \mu_{s, x_2}(1) ds \right)^\ell + \left(\int_0^1 \mu'_{s, x_1}(1) \mu'_{s, x_2}(1) ds \right)^\ell - 2 \left(\int_0^1 \mu_{s, x_1}(1) \mu'_{s, x_2}(1) ds \right)^\ell \right] \geq 0 \quad (\mu, \mu' \in \mathcal{K}). \quad (7.26)$$

By Fubini’s theorem,

$$\begin{aligned} \mathbb{E} \left[\left(\int_0^1 \mu_{s, x_1}(1) \mu_{s, x_2}(1) ds \right)^\ell \right] &= \mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, x_1}(1) \mu_{s_h, x_2}(1) \right] = \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, x_1}(1) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right] \mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, x_2}(1) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, x_1}(1) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right]^2 \right], \end{aligned} \quad (7.27)$$

and analogously

$$\mathbb{E} \left[\left(\int_0^1 \mu'_{s, x_1}(1) \mu'_{s, x_2}(1) ds \right)^\ell \right] = \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu'_{s_h, x_1}(1) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right]^2 \right], \quad (7.28)$$

$$\mathbb{E} \left[\left(\int_0^1 \mu_{s, x_1}(1) \mu'_{s, x_2}(1) ds \right)^\ell \right] = \mathbb{E} \left[\mathbb{E} \left[\prod_{h=1}^{\ell} \mu_{s_h, x_1}(1) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right] \mathbb{E} \left[\prod_{h=1}^{\ell} \mu'_{s_h, x_1}(1) \middle| \mathbf{s}_1, \dots, \mathbf{s}_\ell \right] \right]. \quad (7.29)$$

Combining (7.27)–(7.29), we conclude that the l.h.s. of (7.26) is the expectation of a square. \square

We proceed to prove Theorem 1.4. In light of Lemma 7.11, Theorem 2.7 readily yields a variational formula for $\mathbf{G}_{\lambda, \beta}$. The main issue that we have to confront is that the resulting variational problem for given λ, β ranges over a spaces that depends on these parameters. In effect, it is not a priori clear that these variational problems bear any relationship to the one stated in Theorem 2.7. To deal with this issue, let \mathcal{D}_λ be the set of all $\pi \in \mathcal{D}$ that are supported on $\mu \in \mathcal{K}$ such that $\mu_{s, x}(1) \leq \lambda / (1 + \lambda)$ for all $s, x \in [0, 1]$. Further, for $\pi \in \mathcal{D}_\lambda$ we let $\pi^{*\beta(N, M)}$ be the distribution obtained by the adjoining operation with respect to the weight function ψ_β . Finally, let

$$\mathcal{D}_{\lambda, \beta}^* = \{\pi \in \mathcal{D}_\lambda : \text{for all } N, M \geq 0 \text{ we have } \pi^{*\beta(N, M)} = \pi\}.$$

Lemma 7.12. *For any $N, M \geq 0$ the map $\pi \in \mathcal{D}_\lambda \mapsto \pi^{*\infty(N, M)}$ is continuous.*

Like in the case of Lemma 3.7, the proof is based on arguments involving the cut metric. The details can be found in Appendix A.

Lemma 7.13. *Let $N, M \geq 0$ be integers. Uniformly for all $\pi \in \mathcal{D}_\lambda$ we have $\pi^{*\beta(N, M)} \rightarrow \pi^{*\infty(N, M)}$ as $\beta \rightarrow \infty$.*

Proof. Let $\varepsilon > 0$. For any $\mu \in \mathcal{K}$ let $Z_{\mu, \beta}^{N, M}(s)$ be the weight from (2.16) with respect to the weight function ψ_β . Then we see that, uniformly for all μ and s ,

$$Z_{\mu, \beta}^{N, M}(s) \rightarrow Z_{\mu, \infty}^{N, M}(s) \quad \text{as} \quad \beta \rightarrow \infty. \quad (7.30)$$

Furthermore, if $\mu_{s, x} \leq \lambda / (1 + \lambda)$ for all s, x , then for all $\beta \in (0, \infty]$ we have

$$Z_{\mu, \beta}^{N, M}(s) \geq \left(\frac{1}{1 + \lambda} \right)^N \left(1 - \left(\frac{\lambda}{1 + \lambda} \right)^2 \right)^M > 0. \quad (7.31)$$

Combining (7.30) and (7.31) and recalling the construction of $\mu^{*\beta(N, M)}$, we can construct a measurable map $\xi : [0, 1] \rightarrow [0, 1]$ that preserves the Lebesgue measure such that for large enough β for all $S, X \subset [0, 1]$,

$$\left| \int_S \int_X \mu_{s, x}^{*\beta(N, M)} - \mu_{\xi(s), x}^{*\infty(N, M)} dx ds \right| < \varepsilon.$$

Thus, $\mathcal{D}_\square(\mu^{*\beta(N, M)}, \mu^{*\infty(N, M)}) < \varepsilon$ for large β . Since \mathcal{D}_λ is endowed with the W_1 -metric, the assertion follows. \square

Lemma 7.14. *The set \mathcal{K}_λ is closed.*

Proof. We can view \mathfrak{K}_λ as a scaled version of the space of weak kernels. Therefore, since \mathfrak{K} is complete, so is \mathfrak{K}_λ is complete. Hence, any Cauchy sequence in \mathfrak{K}_λ has a limit within this set, and thus \mathfrak{K}_λ is a closed subspace of \mathfrak{K} . \square

Corollary 7.15. *The set \mathfrak{D}_λ is closed.*

Proof. By Lemma 7.14 there exists an increasing sequence of continuous functions $u_n : \mathfrak{K} \rightarrow [0, 1]$ that converges pointwise to $1 - \mathbf{1}_{\mathfrak{K}_\lambda}$. Thus, $\mathfrak{D}_\lambda = \bigcap_{n \geq 1} \{\pi \in \mathfrak{D} : \int u_n d\pi = 0\}$ is closed in the weak topology. \square

Corollary 7.16. *We have $\liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbf{G}_{\lambda, \beta})] \geq \inf_{\pi \in \mathfrak{D}_\lambda^*} \mathcal{B}(\pi)$.*

Proof. Since \mathfrak{D}^* is compact, Proposition 6.1 shows that there exists $\pi \in \mathfrak{D}^*$ such that

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbf{G}_{\lambda, \beta})] \geq \mathcal{B}(\pi). \quad (7.32)$$

The construction of the π for which the lower bound is attained is based on Proposition 6.2, whose proof shows that the measure $\pi_{\lambda, \beta}$ for which the lower bound is attained in the limit of a sequence of distributions $(\pi_{\lambda, \beta, n})_{n \geq 1}$ that come from random factor graphs with the weight function ψ_β . Specifically, we considered a random factor graph $\mathbf{G}_{\lambda, \beta, n, \omega}$ with a random number of ‘cavities’ for a slowly growing $\omega = \omega_n \rightarrow \infty$. With $\mu_n \in \mathcal{P}(\Omega^{\mathcal{C}})$ the joint Boltzmann distribution of the spins of the cavities \mathcal{C} , the measure $\pi_{\lambda, \beta, n}$ is defined as the distribution of the representation of μ_n as an element of \mathcal{M} . Thus, we just need to show that these representations converge to points in \mathfrak{K}_λ .

The proof of this fact is based on Corollary 3.16. Specifically, let $\varepsilon > 0$. We obtain a decomposition S_1, \dots, S_ℓ of $\Omega^{\mathcal{C}}$ into classes by pinning a random set Θ_ε of cavities. The size $|\Theta_\varepsilon|$ of this set depends on ε only and

$$\Delta_\square(\mu_n, \bar{\mu}_n) < \varepsilon, \quad \text{where} \quad \bar{\mu}_n = \sum_{i=1}^{\ell} \mu(S_i) \bigotimes_{v \in \mathcal{C}} \mu_v(\cdot | S_i). \quad (7.33)$$

Now, consider a cavity $v \in \mathcal{C} \setminus \Theta_\varepsilon$, let $1 \leq i \leq \ell$ and consider a configuration $\sigma \in S_i$ with $\sigma_v = 1$. Obtain σ' by setting $\sigma'_v = 0$ and $\sigma'_w = \sigma_w$ for all $w \neq v$. Then $\sigma' \in S_i$ and the construction of the Boltzmann distribution ensures that $\mu_n(\sigma | S_i) \leq \lambda \mu_n(\sigma' | S_i)$. Hence, $\mu_v(1 | S_i) \leq \lambda / (1 + \lambda)$. Since $|\Theta_\varepsilon|$ is bounded in terms of ε only, whereas $|\mathcal{C}| \geq \omega_n/2 \rightarrow \infty$ with high probability, we deduce from (7.33) that the representation $\check{\mu}_n \in \mathfrak{K}$ satisfies $\mathcal{D}_\square(\check{\mu}_n, \mathfrak{K}_\lambda) < \varepsilon$ with high probability. Since, furthermore, the Wasserstein metric induces the weak topology on \mathfrak{D} , we conclude that $\pi_{\lambda, \beta, n}$ converges to a point π on in the closure of \mathfrak{D}_λ ; but since \mathfrak{D}_λ is closed, we conclude that $\pi \in \mathfrak{D}_\lambda$. Finally, Corollary 7.15 implies that $\pi \in \mathfrak{D}_\lambda \cap \mathfrak{D}^* = \mathfrak{D}_\lambda^*$. Thus, the assertion follows from (7.32). \square

We are ready to establish the lower bound on the free energy.

Proposition 7.17. *For all $d \geq 3, \lambda > 0$ we have $\liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbf{G}_{\lambda, \infty})] \geq \Phi_{d, \lambda}$.*

Proof. For any $\beta, \lambda > 0$ Corollary 7.16 supplies $\pi_{\lambda, \beta} \in \mathfrak{D}_\lambda^*$ such that

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbf{G}_{\lambda, \beta})] \geq \mathcal{B}_{d, \lambda, \beta}(\pi_{\lambda, \beta}). \quad (7.34)$$

Now consider the sequence $(\pi_{\lambda, \beta})_{\beta=1, 2, \dots}$. Since \mathfrak{D}_λ is compact, a subsequence $(\pi_{\lambda, \beta_j})_j$ converges to $\pi_\lambda \in \mathfrak{D}_\lambda$, i.e.,

$$\lim_{j \rightarrow \infty} \mathcal{D}_\square(\pi_{\lambda, \beta_j}, \pi_\lambda) = 0. \quad (7.35)$$

Further, since $\pi_{\lambda, \beta_j}^{*\beta_j(N, M)} = \pi_{\lambda, \beta_j}$ for all j and $N, M \geq 0$, Lemma 7.13 implies that for all pairs $N, M \geq 0$,

$$\lim_{j \rightarrow \infty} \mathcal{D}_\square(\pi_{\lambda, \beta_j}, \pi_{\lambda, \beta_j}^{*\infty(N, M)}) = 0. \quad (7.36)$$

Combining (7.35) and (7.36) with Lemma 7.12, we conclude that $\pi_\lambda \in \mathfrak{D}_\lambda^*$. Finally, since for every $\beta > 0$ we have $\mathcal{B}_{d, \lambda, \beta}(\cdot) \geq \mathcal{B}_{d, \lambda, \infty}(\cdot)$ on \mathfrak{D}_λ , the assertion follows from (7.34) and the continuity of the functional $\mathcal{B}_{d, \lambda, \infty}(\cdot)$. \square

A separate argument is needed to derive the upper bound on the free energy. Basically, we will prove the following proposition by checking that the interpolation argument from Section 5 goes through for the hard-core model.

Proposition 7.18. *For all $d \geq 3, \lambda > 0$ we have $\limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\mathbf{G}_{\lambda, \infty})] \leq \Phi_{d, \lambda}$.*

With $\boldsymbol{\varphi}_i$ and $\boldsymbol{\psi}_{1,i}$ defined with respect to the hard-core weight function ψ_∞ , let

$$\mathcal{B}'(\mu) = \mathbb{E} \log \left\langle \bigoplus_{i=1}^n \boldsymbol{\varphi}_i, \mu \right\rangle, \quad \mathcal{B}''(\mu) = \mathbb{E} \log \left\langle \bigoplus_{1 \leq i \leq dn/2} \boldsymbol{\psi}_{1,i}, \mu \right\rangle.$$

Lemma 7.19. *For any $\lambda > 0$ and any $\mu \in \mathfrak{K}_\lambda$ we have $\mathbb{E} [\log Z(\mathbf{G}_{\lambda, \infty})] \leq \mathcal{B}'(\mu) - \mathcal{B}''(\mu) + o(n)$.*

Proof. This follows along the lines of the proof of Proposition 5.2. In that proof we required the assumption that all weight functions are strictly positive, but only in one place. Namely, we required positivity in order expand the logarithm into a power series in equations (5.9)–(5.11). Yet this approximation is still valid in the hardcore model. Indeed, the term $\langle \psi_{a_{m_t+1}}, \mu_{G_t} \rangle$, whose logarithm we calculate in (5.9), is lower-bounded by $1 - \lambda/(1 + \lambda)$, because in the hard-core model the marginal probability that a single variable node has spin one is upper-bounded by $\lambda/(1 + \lambda)$. Similarly, the arguments of the logarithms in (5.10) and (5.11) are lower-bounded by $1 - \lambda/(1 + \lambda)$ because $\mu \in \mathfrak{K}_\lambda$. \square

Proof of Proposition 7.18. Based on Lemma 7.19, we follow the proof of Proposition 5.3 to complete the proof of Proposition 7.18. Specifically, we claim that for any $\pi \in \mathfrak{D}_\lambda^*$,

$$\mathbb{E}[\mathcal{B}''(\mu^\pi)] = \frac{dn}{2} \mathbb{E}[\log \langle \psi_\infty, \pi \rangle], \quad \text{and} \quad \mathbb{E}[\mathcal{B}'(\mu^\pi)] = \mathbb{E} \log \langle \boldsymbol{\varphi}_1, \pi \rangle. \quad (7.37)$$

This follows along the lines of Lemmas 5.5 and 5.6. In both cases we assumed that the weight functions are strictly positive in order to ensure that the arguments of the logarithms on the l.h.s. are bounded away from zero so that the logarithmic series applies. But the condition $\pi \in \mathfrak{D}_\lambda^*$ guarantees that

$$\left\langle \bigoplus_{i=1}^n \boldsymbol{\varphi}_i, \mu \right\rangle \geq (1/(1 + \lambda))^n \quad \text{and} \quad \left\langle \bigoplus_{1 \leq i \leq dn/2} \boldsymbol{\psi}_{1,i}, \mu \right\rangle \geq (1/(1 + \lambda))^{dn/2}.$$

Thus, the same manipulations as before yield (7.37). Finally, the assertion follows from (7.37) and Lemma 7.19. \square

Proof of Theorem 1.4. The theorem is an immediate consequence of Propositions 7.17 and 7.18. \square

Proof of Corollary 1.5. For a graph $G = (V(G), E(G))$ let $\mu_{G, \lambda} \in \mathcal{P}(\{0, 1\}^{V(G)})$ denote the hard-core model on G with fugacity λ , and let $Z_\lambda(G)$ be the corresponding partition function. Further, let $\alpha_\lambda(G) = \sum_{v \in V(G)} \langle \boldsymbol{\sigma}_v, \mu_{G, \lambda} \rangle$ be the average size of an independent set drawn from $\mu_{G, \lambda}$. Additionally, we write $\alpha(G)$ for the maximum independent set size. It is well known that

$$\alpha_\lambda(G) = \lambda \frac{\partial}{\partial \lambda} \log Z_\lambda(G) \quad (7.38)$$

and that

$$\frac{\alpha_\lambda(G)}{|V(G)|} \xrightarrow{\lambda \rightarrow \infty} \frac{\alpha(G)}{|V(G)|} \quad \text{uniformly for all } G. \quad (7.39)$$

As an immediate consequence of (7.38) we obtain

$$\log Z_{\lambda+1}(G) - \log Z_\lambda(G) = \int_\lambda^{\lambda+1} \frac{\alpha_t(G)}{t} dt \begin{cases} \leq \alpha_{\lambda+1}(G)/\lambda, \\ \geq \alpha_\lambda(G)/(\lambda+1). \end{cases}$$

Hence, (7.39) shows that for any $\varepsilon > 0$ there exists $\lambda_0 > 0$ such that for all $\lambda \geq \lambda_0$ and all d -regular graphs G we have

$$(1 - \varepsilon) \alpha(G) \leq \frac{\lambda}{1 + \lambda} \alpha_\lambda(G) \leq \lambda (\log Z_{\lambda+1}(G) - \log Z_\lambda(G)) \leq \alpha_\lambda(G) \leq \alpha(G). \quad (7.40)$$

Applying (7.40) to the random graph \mathbf{G}_λ and taking expectations, we obtain

$$(1 - \varepsilon) \mathbb{E} \left[\frac{\alpha(\mathbb{G})}{n} \right] \leq \mathbb{E} [\lambda (\log Z_{\lambda+1}(\mathbb{G}) - \log Z_\lambda(\mathbb{G}))] \leq \mathbb{E} \left[\frac{\alpha(\mathbb{G})}{n} \right]. \quad (7.41)$$

Theorem 1.4 guarantees that the sequence $(\mathbb{E} [\lambda (\log Z_{\lambda+1}(\mathbb{G}) - \log Z_\lambda(\mathbb{G}))])_n$ converges, and thus (7.41) yields

$$(1 - \varepsilon) \limsup_{n \rightarrow \infty} \mathbb{E} \left[\frac{\alpha(\mathbb{G})}{n} \right] \leq \lambda (\Phi_{d, \lambda+1} - \Phi_{d, \lambda}) \leq \liminf_{n \rightarrow \infty} \mathbb{E} \left[\frac{\alpha(\mathbb{G})}{n} \right]. \quad (7.42)$$

Further, there exists a subsequence $(n_l)_{l \geq 1}$ along which $\mathbb{E}[\alpha(\mathbb{G})/n]$ converges to $\alpha_* \in [0, 1]$, whence (7.41) yields

$$(1 - \varepsilon)\alpha_* \leq \lambda(\Phi_{d,\lambda+1} - \Phi_{d,\lambda}) \leq \alpha_*. \quad (7.43)$$

Since (7.43) holds for every $\varepsilon > 0$ for large enough λ , we conclude that $\lim_{\lambda \rightarrow \infty} \lambda(\Phi_{d,\lambda+1} - \Phi_{d,\lambda})$ exists. Hence, taking the limit $\varepsilon \rightarrow 0$, and thus $\lambda \rightarrow \infty$, in (7.42) completes the proof. \square

Acknowledgement. The first author thanks Max Hahn-Klimroth for helpful discussions on the cut metric.

REFERENCES

- [1] E. Abbe: Community detection and stochastic block models: recent developments. arXiv:1703.10146 (2017).
- [2] E. Abbe, C. Sandon: Detection in the stochastic block model with multiple clusters: proof of the achievability conjecture, acyclic BP, and the information-computation gap. arXiv:1512.09080 (2015).
- [3] D. Achlioptas, S. Hassani, N. Macris, R. Urbanke: Bounds for random constraint satisfaction problems via spatial coupling. Proc. 27th SODA (2016) 469–479.
- [4] D. Achlioptas, C. Moore: Random k -SAT: two moments suffice to cross a sharp threshold. SIAM Journal on Computing **36** (2006) 740–762.
- [5] D. Achlioptas, A. Naor: The two possible values of the chromatic number of a random graph. Annals of Mathematics **162** (2005) 1333–1349.
- [6] D. Achlioptas, A. Naor, Y. Peres: Rigorous location of phase transitions in hard optimization problems. Nature **435** (2005) 759–764.
- [7] D. Achlioptas, Y. Peres: The threshold for random k -SAT is $2^k \log 2 - O(k)$. Journal of the AMS **17** (2004) 947–973.
- [8] D. Alberici, P. Contucci: Solution of the monomer-dimer model on locally tree-like graphs. Rigorous results. Communications in Mathematical Physics **331** (2014) 975–1003.
- [9] A. Auffinger, A. Jagannath: Thouless-Anderson-Palmer equations for conditional Gibbs measures in the generic p-spin glass model. arXiv:1612.06359 (2016).
- [10] J. Banks, C. Moore, J. Neeman, P. Netrapalli: Information-theoretic thresholds for community detection in sparse networks. Proc. 29th COLT (2016) 383–416.
- [11] V. Bapst, A. Coja-Oghlan: Harnessing the Bethe free energy. Random Structures and Algorithms **49** (2016) 694–741.
- [12] V. Bapst, A. Coja-Oghlan, S. Hetterich, F. Rassmann, D. Vilenchik: The condensation phase transition in random graph coloring. Communications in Mathematical Physics **341** (2016) 543–606.
- [13] J. Barbier, N. Macris: The adaptive interpolation method: a simple scheme to prove replica formulas in Bayesian inference. arXiv:1705.02780 (2017).
- [14] J. Barbier, F. Krzakala, L. Zdeborová, P. Zhang: The hard-core model on random graphs revisited. Journal of Physics: Conference Series **473** (2013) 12–21.
- [15] M. Bayati, D. Gamarnik, P. Tetali: Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. Annals of Probability **41** (2013) 4080–4115.
- [16] H. Bethe: Statistical theory of superlattices. Proc. Roy. Soc. London A **150** (1935) 552–558.
- [17] A. Coja-Oghlan, C. Efthymiou, S. Hetterich: On the chromatic number of random regular graphs. Journal of Combinatorial Theory, Series B **116** (2016) 367–439.
- [18] A. Coja-Oghlan, C. Efthymiou, N. Jaafari, M. Kang, T. Kapetanopoulos: Charting the Replica Symmetric Phase. Communications in Mathematical Physics **359** (2018) 603–698.
- [19] A. Coja-Oghlan, M. Hahn-Klimroth: The cut metric for probability distributions. arXiv:1905.13619.
- [20] A. Coja-Oghlan, N. Jaafari: On the Potts model on random graphs. Electronic Journal of Combinatorics **23** (2016) P4.3.
- [21] A. Coja-Oghlan, F. Krzakala, W. Perkins and L. Zdeborova: Information-theoretic thresholds from the cavity method. Advances in Mathematics **333** (2018) 694–795.
- [22] A. Coja-Oghlan, K. Panagiotou: Catching the k -NAESAT threshold. Proc. 44th STOC (2012) 899–908.
- [23] A. Coja-Oghlan, K. Panagiotou: The asymptotic k -SAT threshold. Advances in Mathematics **288** (2016) 985–1068.
- [24] A. Coja-Oghlan, W. Perkins: Belief Propagation on replica symmetric random factor graph models. Annales de l'institut Henri Poincaré D **5** (2018) 211–249.
- [25] A. Coja-Oghlan, W. Perkins: Bethe states of random factor graphs. Communications in Mathematical Physics **366** (2019) 173–201.
- [26] A. Coja-Oghlan, W. Perkins, K. Skubch: Limits of discrete distributions and Gibbs measures on random graphs. European Journal of Combinatorics **66** (2017) 37–59.
- [27] A. Coja-Oghlan, L. Zdeborová: The condensation transition in random hypergraph 2-coloring. Proc. 23rd SODA (2012) 241–250.
- [28] V. Dani, C. Moore: Independent sets in random graphs from the weighted second moment method. Proc. 15th RANDOM (2011) 472–482.
- [29] A. Dembo, A. Montanari: Gibbs measures and phase transitions on sparse random graphs. Brazilian Journal of Probability and Statistics **24** (2010) 137–211.
- [30] A. Dembo, A. Montanari: Ising models on locally tree-like graphs. Annals of Applied Probability **20** (2010) 565–592.
- [31] A. Dembo, A. Montanari, N. Sun: Factor models on locally tree-like graphs. Annals of Probability **41** (2013) 4162–4213.
- [32] P. Diaconis, S. Janson: Graph limits and exchangeable random graphs. Rend. Mat. Appl. **28** (2008) 33–61.
- [33] J. Ding, A. Sly, N. Sun: Satisfiability threshold for random regular NAE-SAT. Communications in Mathematical Physics **341** (2016) 435–489.
- [34] J. Ding, A. Sly, N. Sun: Proof of the satisfiability conjecture for large k . Proc. 47th STOC (2015) 59–68.
- [35] J. Ding, A. Sly, N. Sun: Maximum independent sets on random regular graphs. Acta Mathematica **217** (2016) 263–340.
- [36] P. Contucci, S. Dommers, C. Giardina, S. Starr: Antiferromagnetic Potts model on the Erdős-Rényi random graph. Communications in Mathematical Physics **323** (2013) 517–554.
- [37] S. Franz, M. Leone: Replica bounds for optimization problems and diluted spin systems. J. Stat. Phys. **111** (2003) 535–564.

[38] A. Frieze, R. Kannan: Quick approximation to matrices and applications. *Combinatorica* **19** (1999) 175–220.

[39] A. Giurgiu, N. Macris, R. Urbanke: Spatial coupling as a proof technique and three applications. *IEEE Transactions on Information Theory* **62** (2016) 5281–5295.

[40] S. Janson: Graphons, cut norm and distance, couplings and rearrangements. *NYJM Monographs*, Volume 4, 2013.

[41] S. Janson, T. Łuczak, A. Ruciński: *Random Graphs*. Wiley (2000)

[42] F. Guerra: Broken replica symmetry bounds in the mean field spin glass model. *Comm. Math. Phys.* **233** (2003) 1–12.

[43] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborová: Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc. National Academy of Sciences* **104** (2007) 10318–10323.

[44] F. Krzakala, L. Zdeborová: Potts glass on random graphs. *Europhysics Letters* **81** (2008) 57005.

[45] M. Lelarge, M. Oulamara: Replica bounds by combinatorial interpolation for diluted spin systems. *Journal of Statistical Physics* **173** (2018) 917–940.

[46] L. Lovász: *Large networks and graph limits*. AMS (2012).

[47] M. Mézard: Mean-field message-passing equations in the Hopfield model and its generalizations *Phys. Rev. E* **95** (2017) 022117.

[48] M. Mézard, A. Montanari: *Information, physics and computation*. Oxford University Press 2009.

[49] M. Mézard, G. Parisi, R. Zecchina: Analytic and algorithmic solution of random satisfiability problems. *Science* **297** (2002) 812–815.

[50] M. Mézard, G. Parisi: The Bethe lattice spin glass revisited. *Eur. Phys. J. B* **20** (2001) 217–233.

[51] M. Mézard, G. Parisi: The cavity method at zero temperature. *Journal of Statistical Physics* **111** (2003) 1–34.

[52] A. Montanari, R. Restrepo, P. Tetali: Reconstruction and clustering in random constraint satisfaction problems. *SIAM Journal on Discrete Mathematics* **25** (2011) 771–808.

[53] C. Moore: The computer science and physics of community detection: landscapes, phase transitions, and hardness. arXiv:1702.00467 (2017).

[54] E. Mossel, J. Neeman, A. Sly: Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields* (2014) 1–31.

[55] D. Panchenko: Spin glass models from the point of view of spin distributions. *Annals of Probability* **41** (2013) 1315–1361.

[56] D. Panchenko: Structure of finite-RSB asymptotic Gibbs measures in the diluted spin glass models. *Journal of Statistical Physics* **162** (2016) 1–42.

[57] D. Panchenko: *The Sherrington-Kirkpatrick model*. Springer 2013.

[58] D. Panchenko, M. Talagrand: Bounds for diluted mean-fields spin glass models. *Probab. Theory Relat. Fields* **130** (2004) 319–336.

[59] T. Richardson, R. Urbanke: *Modern coding theory*. Cambridge University Press (2008).

[60] M. Talagrand: Mean field models for spin glasses. Volumes I and II. Springer (2011).

[61] L. Zdeborová, F. Krzakala: Statistical physics of inference: thresholds and algorithms. *Advances in Physics* **65** (2016) 453–552.

APPENDIX A. PROOF OF LEMMAS 3.6 AND 7.12

The proof of Lemma 3.6 requires the regularity lemma for measures from [26].³ Let λ denote the Lebesgue measure. For $\mu \in \mathcal{K}$ and measurable $S, X \subset [0, 1]$ we write

$$\mu_{S,X} = \frac{1}{\lambda(S)\lambda(X)} \int_S \int_X \mu_{s,x} dx ds \in \mathcal{P}(\Omega),$$

with the convention that $\mu_{S,X}$ is uniform if $\lambda(S)\lambda(X) = 0$. Further, let $\mathbf{X} = (X_1, \dots, X_K)$, $\mathbf{S} = (S_1, \dots, S_L)$ be a partitions of $[0, 1]$ into pairwise disjoint measurable sets. We write $\#\mathbf{X}, \#\mathbf{S}$ for the number K, L of classes, respectively. Then μ is ε -regular with respect to (\mathbf{X}, \mathbf{S}) if there exists $R \subset [\#\mathbf{X}] \times [\#\mathbf{S}]$ such that the following conditions hold.

REG1: $\lambda(X_i) > 0$ and $\lambda(S_j) > 0$ for all $(i, j) \in R$.

REG2: $\sum_{(i,j) \in R} \lambda(X_i)\lambda(S_j) > 1 - \varepsilon$.

REG3: for all $(i, j) \in R$ and almost all $s, s' \in S_j$ we have $\|\int_{X_i} \mu_{s,x} - \mu'_{s',x} dx\|_{\text{TV}} < \varepsilon \lambda(X_i)$.

REG4: if $(i, j) \in R$, then for every $U \subset X_i$ with $\lambda(U) \geq \varepsilon \lambda(X_i)$ and every $T \subset S_j$ with $\lambda(T) \geq \varepsilon \lambda(S_j)$ we have

$$\|\mu_{S,X_i} - \mu_{T,U}\|_{\text{TV}} < \varepsilon.$$

A *refinement* of a partition (\mathbf{X}, \mathbf{S}) is a partition $(\mathbf{X}', \mathbf{S}')$ such that for every pair $(i', j') \in [\#\mathbf{X}'] \times [\#\mathbf{S}']$ there is a pair $(i, j) \in [\#\mathbf{X}] \times [\#\mathbf{S}]$ such that $(X'_{i'}, S'_{j'}) \subset (X_i, S_j)$.

Theorem A.1 ([26]). *For any $\varepsilon > 0$ there exists $N = N(\varepsilon, \Omega)$ such that for every $\mu \in \mathcal{K}$ the following is true. Every partition (X_0, S_0) with $\#X_0 + \#S_0 \leq 1/\varepsilon$ has a refinement (\mathbf{X}, \mathbf{S}) such that $\#\mathbf{X} + \#\mathbf{S} \leq N$ with respect to which μ is ε -regular.*

Additionally, we need the *strong cut metric*, defined by

$$D_{\square}(\mu, \nu) = \sup_{S, X, \omega} \left| \int_S \int_X \mu_{s,x}(\omega) - \nu_{s,x}(\omega) dx ds \right| \quad (\mu, \nu \in \mathcal{K}),$$

³The arguments in the appendix are special cases of more general results on the cut metric from [19].

where S, X range over measurable subsets of the unit interval and $\omega \in \Omega$. It is well known that $D_{\square}(\cdot, \cdot)$ induces a metric on \mathcal{K} .

For $\mu, \nu \in \mathcal{K}$ we define $\mu \oplus \nu : [0, 1]^3 \rightarrow \mathcal{P}(\Omega^2)$ by $\mu \oplus \nu_{s, x_1, x_2} = \mu_{s, x_1} \otimes \mu_{s, x_2}$. Since $[0, 1]^2$ with the Lebesgue measure is isomorphic as a measure space to $[0, 1]$ with the Lebesgue measure, we can view $\mu \oplus \nu$ as a strong $\mathcal{P}(\Omega^2)$ -valued kernel. In particular, it makes sense to apply the strong cut metric to these kernels.

Proposition A.2. *The map $(\mu, \nu) \mapsto \mu \oplus \nu$ is continuous with respect to the strong cut metric.*

Proof. Given $\varepsilon > 0$ pick a small enough $\delta > 0$ and assume that $D_{\square}(\mu, \mu') < \delta$. Due to the triangle inequality it suffices to prove that $D_{\square}(\mu \oplus \nu, \mu' \oplus \nu) < \varepsilon$ for every ν . Thus, we need to show that for any $X \subset [0, 1]^2$, $S \subset [0, 1]$ and $\sigma, \tau \in \Omega$,

$$\left| \int_X \int_S (\mu_{s, x_1}(\sigma) - \mu'_{s, x_1}(\sigma)) \nu_{s, x_2}(\tau) ds dx_1 dx_2 \right| < \varepsilon. \quad (\text{A.1})$$

To this end, we may assume that $\lambda(S) > \varepsilon^2$ and that $\int_S \nu_{s, x_2}(\tau) ds > \varepsilon^2$ for all $(x_1, x_2) \in X$. Further, with $z = \int_0^1 \nu_{s, x_2}(\tau) ds$ consider the variable transformation

$$dt = \frac{\nu_{s, x_2}(\tau) ds}{z}. \quad (\text{A.2})$$

Let T be the inverse image of S under the transformation (A.2). Then we obtain for any $X_1 \subset [0, 1]$,

$$\int_{X_1} \int_S (\mu_{s, x_1}(\sigma) - \mu'_{s, x_1}(\sigma)) \nu_{s, x_2}(\tau) ds dx_1 = z \int_{X_1} \int_T \mu_{t, x_1}(\sigma) - \mu'_{t, x_1}(\sigma) dt dx_1. \quad (\text{A.3})$$

But the assumption $D_{\square}(\mu, \mu') < \delta$ implies that the double integral on the r.h.s. of (A.3) is bounded by ε^4 in absolute value (providing δ is small enough). Thus, (A.1) follows. \square

Proof of Lemma 3.6. We may assume without loss that $f(\tau) = \mathbf{1}\{\tau = \sigma\}$ for some $\sigma \in \Omega^k$. Let $\varepsilon > 0$, pick $\alpha = \alpha(\varepsilon)$, $\xi = \xi(\alpha) > 0$ small enough and assume that $\mu, \nu \in \mathcal{K}$ are such that $D_{\square}(\mu, \nu) < \delta$ for a small enough $\delta = \delta(\xi) > 0$. Applying Theorem A.1 twice, we obtain (X, S) with respect to which both μ, ν are ξ -regular, and $L = \#X + \#S$ is bounded in terms of ξ only. Let R' be the set of all pairs for which **REG1–REG4** are satisfied for both μ, ν and that satisfy $\lambda(X_i, S_j) > \xi^8/L$. Assuming that δ is sufficiently small, we obtain

$$|\mu_{S_i, X_j} - \nu_{S_i, X_j}| < \xi^8 \quad \text{for all } (i, j) \in R'. \quad (\text{A.4})$$

Furthermore, consider the random variables

$$\begin{aligned} z_i &= \prod_{h=1}^k \mu_{S_i, x_h}(\sigma_h), & z &= \sum_{i \leq \#S} z_i, \\ z'_i &= \prod_{h=1}^k \nu_{S_i, x_h}(\sigma_h), & z' &= \sum_{i \leq \#S} z'_i \end{aligned}$$

and define $\mu', \nu' \in \mathcal{K}$ as follows. To construct μ' , partition the interval $[0, 1]$ into pairwise disjoint sets T_i , $i \in [\#S]$, of measure z_i/z and fill the strip $T_i \times [0, 1]$ with a suitably scaled copy of $(\mu_{s, x})_{s \in S_i, x \in [0, 1]}$. Construct ν' analogously from the z'_i . Then $\mathcal{D}_{\square}(\mu', f * \mu) = \mathcal{D}_{\square}(\nu', f * \nu) = 0$. Furthermore, Proposition A.2 shows that with probability at least $1 - \alpha$ we have

$$\sum_{i=1}^{\#S} \lambda(S_i) |z_i - z'_i| < \alpha^2,$$

provided that ξ, δ are chosen small enough. Since also $z \geq \alpha$ because the function f is strictly positive, we conclude that with probability at least $1 - \alpha$ we have $\mathcal{D}_{\square}(\mu', \nu') < \alpha$. We thus obtain a coupling of the random variables $f * \mu, f * \nu$ under which the expected cut distance is bounded by ε , as desired. \square

Proof of Lemma 7.12. We proceed precisely as in the proof of Lemma 3.6, up until the point where the positivity of f is used. In the setup of Lemma 7.12, the function f may take the value 0 on kernels that take the value 1 with positive probability; however, since we are assuming that the values of the kernels are bounded by $\lambda/(1 + \lambda)$. Therefore, the function f always attains values that are bounded away from 0. \square

APPENDIX B. PROOF OF LEMMA 3.3

The proof of Lemma 3.3 requires the following operation. For functions $f : \Omega^{M \times N} \rightarrow \mathbb{R}$, $g : \Omega^{L \times N} \rightarrow \mathbb{R}$ we define

$$f \otimes g : \Omega^{(M+L) \times N} \rightarrow \mathbb{R}, \quad \sigma \mapsto f((\sigma_{i,j})_{i \in [M], j \in [N]}) \cdot g((\sigma_{i+M, j+N})_{i \in [L], j \in [N]}).$$

Thus, the first M rows of σ go into f , the last L rows go into g and we multiply the results.

We define a corresponding operation on kernels. Namely, for $\mu, \nu \in \mathcal{K}$ we define $\mu \otimes \nu : [0, 1]^3 \rightarrow \mathcal{P}(\Omega^2)$ by $\mu \oplus \nu_{s,t,x} = \mu_{s,x} \otimes \nu_{t,x}$. Since $([0, 1]^2, \lambda \otimes \lambda)$ is isomorphic $([0, 1], \lambda)$, we can view $\mu \otimes \nu$ as a $\mathcal{P}(\Omega^2)$ -valued kernel, and the cut metric extends to these kernels. Since the cut metric is invariant under swapping the axes, Proposition A.2 readily yields the following.

Proposition B.1. *The map $(\mu, \nu) \mapsto \mu \otimes \nu$ is continuous with respect to the cut metric.*

As a final preparation toward the proof of Lemma 3.3 we need the following fact.

Lemma B.2. *For any $f : \Omega \rightarrow \mathbb{R}$ the map $\mu \in \mathcal{K} \mapsto \mathbb{E}\langle f, \mu \rangle$ is continuous.*

Proof. We may assume without loss that $f(\tau) = \mathbf{1}\{\sigma = \tau\}$ for some $\sigma \in \Omega$. Then

$$\mathbb{E}\langle f, \mu \rangle = \int_0^1 \int_0^1 \mu_{s,x}(\sigma) dx ds,$$

and it is immediate from the definition of the cut metric that the integral on the right hand side is a continuous function of μ . \square

Proof of Lemma 3.3. Let $f : \Omega^{m \times n} \rightarrow \mathbb{R}$ and let $\mu \in \mathcal{K}$. Define $\nu = (\mu^{\oplus n})^{\otimes m}$. Then ν is a kernel with values in Ω^{mn} and the definition of $\langle \cdot, \cdot \rangle$ ensures that $\mathbb{E}\langle f, \mu \rangle = \mathbb{E}\langle f, \nu \rangle$. This already shows that the map $\mu \mapsto \mathbb{E}\langle f, \mu \rangle$ is continuous, because the map $\mu \mapsto \nu$ is continuous by Proposition A.2 and B.1 and the map $\nu \mapsto \mathbb{E}\langle f, \nu \rangle$ is continuous by Lemma B.2. Now fix an integer $\ell \geq 2$ and let $\eta = \nu^{\otimes \ell}$. Then

$$\mathbb{E}[\langle f, \mu \rangle^\ell] = \mathbb{E}[\langle f, \eta \rangle]$$

and thus the continuity of the map $\mu \mapsto \mathbb{E}[\langle f, \mu \rangle^\ell]$ follows from Proposition B.1 and Lemma B.2. \square

AMIN COJA-OGLAN, acoghan@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

WILL PERKINS, math@willperkins.org, DEPARTMENT OF MATHEMATICS, STATISTICS, AND COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS AT CHICAGO, CHICAGO, ILLINOIS, USA.