

Estimation in tensor Ising models

SOMABHA MUKHERJEE

Department of Statistics and Data Science, National Institute of Singapore, Singapore, Singapore

JAESUNG SON

Department of Statistics, Columbia University, New York, NY, USA

AND

BHASWAR B. BHATTACHARYA[†]

Department of Statistics and Data Science, University of Pennsylvania, Philadelphia, PA, USA

[†]Corresponding author; Email: bbhaswar@wharton.upenn.edu

[Received on 9 November 2022; revised on 13 September 2021; accepted on 9 February 2022]

The p -tensor Ising model is a one-parameter discrete exponential family for modeling dependent binary data, where the sufficient statistic is a multi-linear form of degree $p \geq 2$. This is a natural generalization of the matrix Ising model that provides a convenient mathematical framework for capturing, not just pairwise, but higher-order dependencies in complex relational data. In this paper, we consider the problem of estimating the natural parameter of the p -tensor Ising model given a single sample from the distribution on N nodes. Our estimate is based on the maximum pseudolikelihood (MPL) method, which provides a computationally efficient algorithm for estimating the parameter that avoids computing the intractable partition function. We derive general conditions under which the MPL estimate is $\sqrt{N}$ -consistent, that is, it converges to the true parameter at rate $1/\sqrt{N}$. Our conditions are robust enough to handle a variety of commonly used tensor Ising models, including spin glass models with random interactions and models where the rate of estimation undergoes a phase transition. In particular, this includes results on $\sqrt{N}$ -consistency of the MPL estimate in the well-known p -spin Sherrington–Kirkpatrick model, spin systems on general p -uniform hypergraphs and Ising models on the hypergraph stochastic block model (HSBM). In fact, for the HSBM we pin down the exact location of the phase transition threshold, which is determined by the positivity of a certain mean-field variational problem, such that above this threshold the MPL estimate is $\sqrt{N}$ -consistent, whereas below the threshold no estimator is consistent. Finally, we derive the precise fluctuations of the MPL estimate in the special case of the p -tensor Curie–Weiss model, which is the Ising model on the complete p -uniform hypergraph. An interesting consequence of our results is that the MPL estimate in the Curie–Weiss model saturates the Cramér–Rao lower bound at all points above the estimation threshold, that is, the MPL estimate incurs no loss in asymptotic statistical efficiency in the estimability regime, even though it is obtained by minimizing only an approximation of the true likelihood function for computational tractability.

Keywords: Markov random fields; parameter estimation; hypergraphs; phase transitions; pseudolikelihood; spin systems.

1. Introduction

The Ising model is a discrete exponential family with binary outcomes, where the sufficient statistic involves a quadratic term designed to capture correlations arising from pairwise interactions. This was originally developed in statistical physics to model ferromagnetism [37] and has since then found applications in various diverse fields, such as spatial statistics, social networks, computer vision, neural networks and computational biology [4, 26, 28, 32, 36, 47]. The increasing popularity of the Ising model as a foundational tool for understanding nearest-neighbor interactions in network data has made it imperative to develop computationally tractable algorithms for learning the model parameters and understanding their rates of convergence (statistical efficiencies). In particular, we are interested in estimating the parameters of the model given a single sample of binary outcomes from an underlying network. This problem was classically studied when the underlying network was a spatial lattice, where consistency and optimality of the maximum likelihood (ML) estimates were derived [20, 30, 33, 54]. However, for general networks, parameter estimation using the ML method turns out to be notoriously hard due to the appearance of an intractable normalizing constant in the likelihood. To circumvent this issue, Chatterjee [18] proposed using the maximum pseudolikelihood (MPL) estimator [8, 9], which is a computationally efficient algorithm for estimating the parameters of a Markov random field, that maximizes an approximation to the likelihood function (a ‘pseudolikelihood’) based on conditional distributions. This method and results in [18] were later generalized in [10] and [29] to obtain rates of estimation for Ising models on general weighted graphs and joint estimation of parameters, respectively. These techniques were recently used in Daskalakis et al. [24, 25] to obtain rates of convergence of the MPL estimate in general logistic regression models with dependent observations. Very recently, Dagan *et al.* [22] considered the problem of parameter estimation in a more general model where the binary outcomes can be influenced by various underlying networks and, as a consequence, improved some of the results in [10].

In many situations, both in modeling real-world network data and interacting spin systems, dependencies arise not just from pairs, but from interactions between groups of particles or individuals. This leads to the study of higher-order Ising models, specifically, the p -tensor Ising model, where the sufficient statistic is a multilinear polynomial of degree $p \geq 2$, for capturing higher-order interactions between the different particles. These models can be represented as a spin system on a p -uniform hypergraph, where the individual entities represent the vertices of the hypergraph and the p -tuples of interactions are indexed by the hyperedges. More formally, given a vector of binary outcomes $\mathbf{X} := (X_1, \dots, X_N) \in \mathcal{C}_N := \{-1, 1\}^N$ and a p -tensor $\mathbf{J}_N = ((J_{i_1 \dots i_p}))_{1 \leq i_1 \dots i_p \leq N}$, the p -tensor Ising model is a discrete exponential family with probability mass function:

$$\mathbb{P}_{\beta, p}(\mathbf{X}) = \frac{1}{2^N Z_N(\beta, p)} e^{\beta H_N(\mathbf{X})}, \quad (1.1)$$

where the sufficient statistic (Hamiltonian)

$$H_N(\mathbf{X}) := \sum_{1 \leq i_1, \dots, i_p \leq N} J_{i_1 \dots i_p} X_{i_1} \dots X_{i_p}, \quad (1.2)$$

and $\beta \geq 0$ is the *natural parameter* (referred to as the *inverse temperature* in statistical physics) of the model. The *normalizing constant* $Z_N(\beta, p)$ (also referred to as the *partition function*) is determined by

the condition $\sum_{X \in \mathcal{C}_N} \mathbb{P}_{\beta,p}(X) = 1$, that is,

$$Z_N(\beta, p) = \frac{1}{2^N} \sum_{X \in \mathcal{C}_N} \exp \left\{ \beta \sum_{1 \leq i_1, \dots, i_p \leq N} J_{i_1 \dots i_p} X_{i_1} \dots X_{i_p} \right\}$$

We will denote by $F_N(\beta, p) := \log Z_N(\beta, p)$ the *log-partition* function of the model. Higher-order Ising models arise naturally in the study of multi-atom interactions in lattice gas models, such as the square-lattice eight-vertex model, the Ashkin–Teller model, and Suzuki’s pseudo-three-dimensional anisotropic model (cf. [6, 35, 38, 40, 52, 58, 59, 64, 69] and the references therein). More recently, higher-order spin systems have also been proposed for modeling peer-group effects in social networks [24].

Hereafter, unless mentioned otherwise, we will assume that the tensor J_N satisfies the following two properties:

- (1) The tensor J_N is *symmetric*, that is, $J_{i_1 \dots i_p} = J_{i_{\sigma(1)} \dots i_{\sigma(p)}}$ for every $1 \leq i_1 < \dots < i_p \leq N$ and every permutation σ of $\{1, \dots, p\}$ and
- (2) The tensor J_N has zeros on the ‘diagonals’, that is, $J_{i_1 \dots i_p} = 0$, if $i_s = i_t$ for some $1 \leq s < t \leq p$.

In this paper, we consider the problem of estimating the parameter β given a single sample $X = (X_1, X_2, \dots, X_n)$ from the p -tensor Ising model (1.1). We establish general sufficient conditions under which the MPL estimate of β is $\sqrt{N}$ -consistent in the p -tensor Ising model, for any $p \geq 2$.

¹ The main bottleneck as one moves from the matrix ($p = 2$) to the tensor ($p \geq 3$) case is the lack of a natural spectral condition that is strong enough to control the fluctuations of the MPL function, but still verifiable in natural examples. To this end, we introduce the notion of a *local interaction matrix* that, given a configuration $x \in \{-1, 1\}^n$, measures the strength of the interaction between pairs of vertices (Definition 2.2). Our result shows that the MPL estimate is $\sqrt{N}$ -consistent, whenever we have an appropriate moment bound on the local interaction matrix, and if the normalized log-partition function stays bounded away from zero (Theorem 2.3). We illustrate the robustness and generality of our result by verifying the conditions of the theorem in various commonly studied tensor Ising models. This includes the $\sqrt{N}$ -consistency of the MPL estimate in the well-known p -spin Sherrington–Kirkpatrick (SK) model [11, 53] (Corollary 2.5), and in Ising models on p -uniform hypergraphs under appropriate conditions on the adjacency tensors (Corollary 2.6). The latter is also related to the recent work of Daskalakis *et al.* [24], where a general model for logistic regression with dependent observations using higher-order Ising models was proposed, which includes as a special case the model in (1.1). However, the conditions in [24] are based directly on the interaction tensor, hence, cannot handle models where the rate of estimation undergoes a phase transition. This is understandable because [24] considered the problem of jointly estimating multiple parameters in a more general model; hence, stronger assumptions were necessary for ensuring consistency. Our goal, on the other hand, is to pin down the precise conditions necessary for estimating the single parameter β and develop methods for verifying those conditions in natural examples. To this end, our general theorem recovers as a corollary, the results in [24] when specialized to the model (1.1). More importantly, our results can handle models where the

¹ A sequence of estimators $\{\hat{\beta}_N\}_{N \geq 1}$ is said to be *consistent* at β , if $\hat{\beta}_N \xrightarrow{P} \beta$ under $\mathbb{P}_\beta$, that is, for every $M > 0$, $\mathbb{P}_\beta(|\hat{\beta}_N(X) - \beta| \leq M) \rightarrow 1$ as $N \rightarrow \infty$. Moreover, a sequence of estimators $\{\hat{\beta}_N\}_{N \geq 1}$ is said to be *$\sqrt{N}$ -consistent* at β , if for every $\delta > 0$, there exists $M := M(\delta, \beta) > 0$ such that $\mathbb{P}_\beta(\sqrt{N}|\hat{\beta}_N(X) - \beta| \leq M) > 1 - \delta$, for all N .

rate of estimation has phase transitions, which happens whenever the underlying hypergraph becomes dense. To illustrate this phenomenon we consider the Ising model on a hypergraph stochastic block model (HSBM), a natural generalization of the widely studied (graph) stochastic block model, which serves as a natural model for capturing higher-order relational data [1, 41, 62, 63]. In this case, we show there is a critical value β_{HSBM}^* , such that if $\beta > \beta_{\text{HSBM}}^*$ then the MPL estimate is $\sqrt{N}$ -consistent, while if $\beta < \beta_{\text{HSBM}}^*$ there is no consistent estimator for β (Theorem 2.10). This phase transition stems from the fact that below the threshold the pairwise correlations are weak and the model is asymptotically indistinguishable from the uniform distribution on $\mathcal{C}_N$. While it is relatively straightforward to show the $\sqrt{N}$ -consistency of the MPL estimate above the threshold using our general theorem, proving that estimation is impossible below the threshold is more challenging. This is one of the technical highlights of the paper, which requires careful combinatorial estimates that go beyond the standard mean-field approximation techniques.

Next, we consider the special case of the p -tensor Curie–Weiss model, which is the Ising model on the complete p -uniform hypergraph. Here, using the special structure of the interaction tensor we are able to obtain the exact limiting distribution of the MPL estimate for all points above the estimation threshold (Theorem 2.15). In fact, in this regime the asymptotic variance of the MPL estimate saturates the Cramer–Rao lower bound, that is, the MPL estimate attains the best asymptotic variance among the class of consistent estimates. Finally, we derive the asymptotic distribution of the MPL estimate in the p -tensor Curie–Weiss model at the estimation threshold (this is the point below which consistent estimation is impossible and above which the MPL estimate is $\sqrt{N}$ -consistent). Interestingly, at the threshold the asymptotics of the MPL estimate depend on the value of p . In particular, the MPL estimate is $\sqrt{N}$ -consistent (with a non-Gaussian limiting distribution) when $p = 2$, but inconsistent for $p \geq 3$. The formal statements of the results and their various consequences are given below in Section 2.

1.1 Related work

The problem of structure learning in Markov random fields is another related area of active research, where the goal is to estimate the underlying graph structure given access to *multiple* i.i.d. samples from an Ising model, or a more general (undirected) graphical model. This begins with the seminal paper by Chow and Liu [19], where the problem was resolved for the special case of tree-structured graphical models using ML estimation, which in this special case is tractable. However, as discussed above, for undirected graphical models on general graphs the partition function, and hence, the ML estimators are intractable.

Over the years, various algorithms have been proposed for learning undirected graphical models on general graphs, which circumvent the intractability of the partition function. Towards this, an initial noteworthy result is due Ravikumar *et al.* [55] where a regularized pseudolikelihood-based estimator, with an additive L_1 penalty, was studied. This estimator offers guarantees for structure reconstruction under certain conditions. However, the conditions required for the result to hold are often hard to verify and the method performs poorly for graphs with long-range correlations [7]. Various other papers [2, 68] give alternative algorithms for estimating sparse high-dimensional Ising models, all of which require various restrictive assumptions to guarantee efficacy. Bresler *et al.* [13] was the first to show that it is possible to reconstruct any discrete graphical model with bounded degree with samples that is logarithmic in the number of variables and polynomial computational complexity (although the degree of the polynomial equals to the maximum degree of the underlying graph). The computational complexity was improved by Bresler [12] to quasi-quadratic for Ising models. However, the proposed algorithm had double-exponential sample dependencies on the model parameters, such as the maximum degree of the underlying graph and the largest coupling intensity, which is non-optimal with respect to

the information theoretic lower bound [57]. The first computationally efficient reconstruction algorithm for Ising models with a near-optimal sample complexity was the ‘interaction screening’ approach designed and analyzed in [66]. This has the structure of a M -estimator and can be solved via convex optimization combined with a thresholding procedure. Another related line of work has focused on testing properties of the underlying graph. In this direction, Daskalakis *et al.* [23] studied the problems of identity and independence testing, and Neykov *et al.* [15, 51] considered problems in graph property testing, given access to multiple samples from an Ising model.

Algorithms for learning discrete graphical models beyond pairwise and binary alphabets have been studied recently [34, 39, 65, 67]. The method in [34] works for arbitrary models with bounded degrees, however, as in [12], the sample complexity is double-exponential in the model parameters. On the other hand, the method in [39], which resembles a stochastic first order method with multiplicative updates, has a low computational complexity and is sample-efficient for structure recovery for multiwise graphical models over binary variables and pairwise models with general alphabets. The sample complexity in [39] was improved in subsequent work [67] in the case of pairwise models over non-binary variables. Very recently [65] generalized the interaction screening framework to obtain the first sample-efficient method based for learning general discrete factor models with node-specific alphabets and multi-body interactions.

All the aforementioned results, however, are in contrast with the present work, where the underlying graph structure is assumed to be known and the goal is to derive rates of estimation for the natural parameters given a *single* sample from the model. This is motivated by the applications mentioned above where it is often difficult, if not impossible, to generate many independent samples from the model within a reasonable amount of time. Although there are methodological similarities to some of the algorithms discussed above, such as the use of the pseudolikelihood method, the results obtained in this paper are fundamentally different. Here, unlike in the multiple samples case, one cannot treat the different neighborhoods as independent, which renders classical techniques for proving consistency inapplicable. Consequently, to handle the dependencies among the different neighborhoods in the pseudolikelihood function we need to use a different (non-classical) set of tools, such as the method of exchangeable pairs and non-linear large deviations. More closely related to the present work are the problems in hypothesis testing given a single sample from an Ising model [14, 49, 50]. In particular, [14] studied the problem of testing, based on a single sample, between the 2-spin Curie–Weiss model and an Ising model on a regular graph, where it was shown that testability depends on the magnitude of the inverse temperature β . In [49, 50] the authors considered the problem of global testing against sparse alternatives for the means of binary outcomes following an Ising model (with node-specific magnetic-field parameters).

2. Statements of the main results

In this section, we state our main results. We begin with a section on notations (Section 2.1). The general result about the $\sqrt{N}$ -consistency of the MPL estimate in tensor Ising models is discussed in Section 2.2. Applications of this result to the p -spin SK model and Ising models on various hypergraphs are discussed in Section 2.3. Finally, in Section 2.4 we obtain the limiting distribution of the MPL estimate in the p -spin Curie–Weiss model. Hereafter, we will often omit the dependence on p and abbreviate the probability measure $\mathbb{P}_{\beta,p}$, the partition function $Z_N(\beta,p)$ and log-partition function $F_N(\beta,p)$ by $\mathbb{P}_\beta$, $Z_N(\beta)$ and $F_N(\beta)$, respectively, when there is no scope of confusion.

2.1 Notations

For a vector $\mathbf{v} \in \mathbb{R}^N$, $\|\mathbf{v}\|$ will denote the Euclidean norm of $\mathbf{v}$. For a $N \times N$ matrix $\mathbf{A}$, $\|\mathbf{A}\| := \sup_{\|\mathbf{x}\|=1} \|\mathbf{A}\mathbf{x}\|$ will denote the operator norm of $\mathbf{A}$ and $\lambda_{\max}(\mathbf{A})$ and $\lambda_{\min}(\mathbf{A})$ the maximum and the

minimum eigenvalues of $\mathbf{A}$, respectively. Finally, for the set $[N] := \{1, 2, \dots, N\}$, $[N]^p$ will denote the p -fold Cartesian product $[N] \times [N] \times \dots \times [N]$, and $[N]_p$ is the collection of p -tuples in $[N]^p$ with distinct entries.

For positive sequences $\{a_n\}_{n \geq 1}$ and $\{b_n\}_{n \geq 1}$, $a_n = O(b_n)$ means $a_n \leq C_1 b_n$, $a_n = \Omega(b_n)$ means $a_n \geq C_2 b_n$, and $a_n = \Theta(b_n)$ means $C_1 b_n \leq a_n \leq C_2 b_n$, for all n large enough and positive constants C_1, C_2 . Similarly, for positive sequences $\{a_n\}_{n \geq 1}$ and $\{b_n\}_{n \geq 1}$, $a_n \lesssim b_n$ means $a_n \leq C_1 b_n$, and $a_n \gtrsim b_n$ means $a_n \geq C_2 b_n$, for all n large enough and positive constants C_1, C_2 . Moreover, subscripts in the above notation, for example $O_\square, \lesssim_\square$, denote that the hidden constants may depend on the subscripted parameters.

2.2 Rate of consistency of the MPL Estimator

The MPL method, introduced by Besag [8, 9], provides a way to conveniently approximate the joint distribution of $\mathbf{X} \sim \mathbb{P}_{\beta,p}$ that avoids calculations with the normalizing constant.

DEFINITION 2.1. ([8; 9]) Given a discrete random vector $\mathbf{X} = (X_1, X_2, \dots, X_N)$ whose joint distribution is parameterized by a parameter $\beta \in \mathbb{R}$, the MPL estimate of β is defined as

$$\hat{\beta}_N(\mathbf{X}) := \arg \max_{\beta \in \mathbb{R}} \prod_{i=1}^N f_i(\beta, \mathbf{X}),$$

where $f_i(\beta, \mathbf{X})$ is the conditional probability mass function of X_i given $(X_j)_{j \neq i}$.

To compute the MPL estimate in the p -tensor Ising model (1.1), fix $\beta > 0$ and consider $\mathbf{X} \sim \mathbb{P}_\beta$. Then from (1.1), the conditional distribution of X_i given $(X_j)_{j \neq i}$ can be easily computed as:

$$\mathbb{P}_\beta \left(X_i | (X_j)_{j \neq i} \right) = \frac{e^{p\beta X_i m_i(\mathbf{X})}}{e^{p\beta m_i(\mathbf{X})} + e^{-p\beta m_i(\mathbf{X})}}, \quad (2.1)$$

where $m_i(\mathbf{X}) := \sum_{1 \leq i_2, \dots, i_p \leq N} J_{ii_2 \dots i_p} X_{i_2} \cdots X_{i_p}$, is the local effect at the node $1 \leq i \leq N$ (often referred to as the local magnetization of the vertex i in the statistical physics literature). Then the pseudolikelihood estimate of β (as defined in (1.1)) in the p -tensor Ising model (2.1) is obtained by maximizing the function below, with respect to b ,

$$L(b|\mathbf{X}) := \prod_{i=1}^N \mathbb{P}_b \left(X_i | (X_j)_{j \neq i} \right) = \frac{1}{2^N} \exp \left\{ \sum_{i=1}^N \{pbX_i m_i(\mathbf{X}) - \log \cosh (pbm_i(\mathbf{X}))\} \right\}.$$

Now, since $\log L(b|\mathbf{X})$ is concave in b , the MPL estimator $\hat{\beta}_N(\mathbf{X})$ can be obtained by solving the gradient equation $\frac{\partial \log L(b|\mathbf{X})}{\partial b} = 0$, which simplifies to

$$H_N(\mathbf{X}) - \sum_{i=1}^N m_i(\mathbf{X}) \tanh (pbm_i(\mathbf{X})) = 0. \quad (2.2)$$

To ensure well-definedness, in case (2.2) does not have a solution or has more than one solution, the MPL estimate $\hat{\beta}_N(\mathbf{X})$ is more formally defined as:

$$\hat{\beta}_N(\mathbf{X}) := \inf \left\{ b \geq 0 : H_N(\mathbf{X}) = \sum_{i=1}^N m_i(\mathbf{X}) \tanh(pbm_i(\mathbf{X})) \right\}, \quad (2.3)$$

where the infimum of an empty-set is defined to be $+\infty$. Note that the expression in the right hand side (RHS) of the equality in (2.3) is an increasing function of b , hence $\hat{\beta}_N(\mathbf{X})$ can be very easily computed by the Newton–Raphson method or even a simple grid search.

Our first result is about the rate of consistency of the MPL estimate in general tensor Ising models. In particular, we show in the proposition below that the MPL estimate $\hat{\beta}_N(\mathbf{X})$, based on a single sample $\mathbf{X} \sim \mathbb{P}_\beta$ converges to the true parameter β at rate $1/\sqrt{N}$, whenever the interaction tensor $\mathbf{J}_N$ satisfies a certain spectral-type condition and the log-partition function is $\Omega(N)$ at the true parameter value. To state our result formally, we need the following definition:

DEFINITION 2.2. Given a p -tensor $\mathbf{J}_N = ((J_{i_1 i_2 \dots i_p}))_{1 \leq i_1, i_2, \dots, i_p \leq N}$ and $\mathbf{x} = (x_1, x_2, \dots, x_N) \in \mathcal{C}_N$, define the *local interaction matrix* of $\mathbf{J}_N$ at the point $\mathbf{x}$ as the $N \times N$ matrix $\mathbf{J}_N(\mathbf{x}) := ((J_{i_1 i_2}(\mathbf{x})))_{1 \leq i_1, i_2 \leq N}$, where the entries are given by:

$$J_{i_1 i_2}(\mathbf{x}) := \sum_{1 \leq i_3, \dots, i_p \leq N} J_{i_1 i_2 i_3 \dots i_p} x_{i_3} \cdots x_{i_p}. \quad (2.4)$$

(Note that in the case $p = 2$, $J_{i_1 i_2}(\mathbf{x}) = J_{i_1 i_2}$, that is, the local interaction matrix $\mathbf{J}_N(\mathbf{x})$ is same as the interaction matrix $\mathbf{J}_N$, for all $\mathbf{x} \in \mathcal{C}_N$.)

We are now ready to state our result on the convergence rate of the MPL estimate in a tensor Ising model.

THEOREM 2.3. Fix $p \geq 2$, $\beta > 0$ and a sequence of p -tensors $\{\mathbf{J}_N\}_{N \geq 1}$ such that the following two conditions hold:

- (1) $\sup_{N \geq 1} \mathbb{E}_\beta[\|\mathbf{J}_N(\mathbf{Z})\|^4] < \infty$, where the expectation is taken with respect to $\mathbf{Z} \sim \mathbb{P}_\beta$,
- (2) $\liminf_{N \rightarrow \infty} \frac{1}{N} F_N(\beta) > 0$.

Then given a single sample $\mathbf{X}$ from the model (1.1) with interaction tensor $\mathbf{J}_N$, the MPL estimate $\hat{\beta}_N(\mathbf{X})$, as defined in (2.3), is $\sqrt{N}$ -consistent for β , that is, for every $\delta > 0$, there exists $M := M(\delta, \beta) > 0$ such that

$$\mathbb{P}_\beta(\sqrt{N}|\hat{\beta}_N(\mathbf{X}) - \beta| \leq M) > 1 - \delta,$$

for all N large enough.

The proof of this theorem is given in Section 3. The proof has two main steps: in the first step we use the method of exchangeable pairs to show that the derivative of the log-pseudolikelihood (the left hand side (LHS) of (2.2)) is concentrated around zero at the true model parameter (see Lemma 3.1 for details). The proof adapts the method of exchangeable pairs introduced in [18] where a similar result was proved for matrix (2-spin) Ising models. The main technical challenge as one goes from the

matrix to the tensor case is the absence of a natural spectral condition in tensor models. To this end, we introduce condition (1), which requires that the fourth-moment of the spectral norm of the local interaction matrix is uniformly bounded. This condition allows us to prove the desired concentration of the log-pseudolikelihood and, as we will see below, can be easily verified for a large class of natural tensor models. The second step in the proof of Theorem 2.3 is to show that the log-pseudolikelihood is strongly concave, that is, its second derivative is strictly negative with high probability. Here, we use condition (2) to first show that the Hamiltonian is $\Omega(N)$ with high probability, which then implies the strong concavity of the log-pseudolikelihood by a truncated second-moment argument.

Remark 2.4 Regarding condition (1) in Theorem 2.3. The discussion in Definition 2.2 implies that when $p = 2$ the condition (1) simplifies to the boundedness of the spectral norm: $\sup_{N \geq 1} \|\mathbf{J}_N\| < \infty$. Hence, Theorem 2.3 recovers Chatterjee's result on $\sqrt{N}$ -consistency of MPL estimates in 2-spin Ising models [(18), Theorem 1.1]. Note that, since $\|\mathbf{J}_N\| \leq \max_{1 \leq i \leq N} \sum_{j=1}^N |J_{ij}|$, the boundedness of the spectral norm is implied by the bounded row sum ('degree') condition:

$$\sup_{N \geq 1} \max_{1 \leq i \leq N} \sum_{j=1}^N |J_{ij}| < \infty,$$

which, informally, ensures that every node has $O(1)$ contribution to the Hamiltonian. For tensor models ($p \geq 3$), as mentioned above, it is *a priori* unclear what would be the 'correct' analogue of the spectral condition. The L_4 moment condition (condition (1)) in Theorem 2.3 is one attempt at such a condition, which is satisfied in natural examples and is also sufficient to establish the concentration of the MPL function. Note that the L_4 moment condition on the local interaction matrix in Theorem 2.3 can be replaced by the following stronger L_∞ -condition, which is often easier to verify in examples:

$$\sup_{N \geq 1} \sup_{\mathbf{x} \in \mathcal{C}_N} \|\mathbf{J}_N(\mathbf{x})\| < \infty. \quad (2.5)$$

Moreover, similar to the 2-spin case, condition (2.5), hence condition (1) in Theorem 2.3, is weaker than the 'bounded-degree condition':

$$\sup_{1 \leq i_1 \leq N} \sum_{1 \leq i_2, i_3, \dots, i_p \leq N} |J_{i_1 i_2 i_3 \dots i_p}| = O(1). \quad (2.6)$$

As in the 2-spin case, this has the interpretation that every node has $O(1)$ contribution to the Hamiltonian. However, we do need the weaker condition (2.5) to handle the p -spin SK model, where the bounded-degree condition (2.6) fails to hold.

2.3 Applications

In this section we discuss the consequences of Theorem 2.3 to the p -spin SK model (Section 2.3.1), spin systems on general hypergraphs (Section 2.3.2), and the HSBM (Section 2.3.3).

2.3.1 The p -Spin SK Model In the p -spin SK model [11], the interaction tensor is of the form

$$J_{i_1 \dots i_p} = N^{\frac{1-p}{2}} g_{i_1 \dots i_p}, \quad (2.7)$$

where $(g_{i_1 \dots i_p})_{1 \leq i_1 < \dots < i_p < \infty}$ is a fixed realization of a collection of independent standard Gaussian random variables, and $g_{i_1 \dots i_p} = g_{\sigma(i_1) \dots \sigma(i_p)}$, for any permutation σ of $\{1, 2, \dots, p\}$. This is a canonical example of a spin glass model that has remarkable thermodynamic properties [45]. A whole new discipline has emerged from the study of this object, with many beautiful theorems that have unearthed deep connections between diverse areas in mathematics and statistical physics (cf. [11, 46, 53, 60] and the references therein). The problem of parameter estimation in the SK model was initiated by Chatterjee [18], where $\sqrt{N}$ -consistent of the MPL estimate for all $\beta > 0$ was proved for the 2-spin SK model. The following corollary extends this to all $p \geq 3$.

COROLLARY 2.5. In the p -spin SK model, the MPL estimate $\hat{\beta}_N(\mathbf{X})$ is $\sqrt{N}$ -consistent for all $\beta > 0$.

The proof of this result is given in Section 4. In this case, condition (2) in Theorem 2.3 can be easily verified using monotonicity and the well-known asymptotics of $F_N(\beta)$ in the high-temperature (small β) regime: in particular, we know from [(11), Theorem 1.1] that, almost surely, $\lim_{N \rightarrow \infty} \frac{1}{N} F_N(\beta) = \frac{\beta^2}{2}$, for $\beta > 0$ small enough. Hence, by the monotonicity of $F_N(\beta)$, we have $\lim_{N \rightarrow \infty} \frac{1}{N} F_N(\beta) > 0$ for all $\beta > 0$, which establishes (2). However, unlike when $p = 2$, verifying condition (1) in Theorem 2.3 when $p \geq 3$ requires more work.² To this end, note that for $p \geq 3$ and every fixed $\mathbf{x} \in \mathcal{C}_N$, the local interaction matrix $\mathbf{J}_N(\mathbf{x})$ is a Gaussian random matrix, but the elements are now dependent because of the symmetry of the tensor $\mathbf{J}_N$. This dependence, however, is relatively weak and using standard Gaussian process machinery we can show the validity of (2.5), and, hence, that of condition (1) in Theorem 2.3.

2.3.2 Ising models on hypergraphs The p -tensor model (1.1) can be interpreted as a spin system on a weighted p -uniform hypergraph, where the entries of the tensor correspond to the weights of the hyperedges. More precisely, given a symmetric tensor $\mathbf{J}_N = ((J_{i_1 i_2 \dots i_p}))_{1 \leq i_1, i_2, \dots, i_p \leq N}$, construct a weighted p -uniform hypergraph H_N with vertex set $[N] := \{1, 2, \dots, N\}$ and edge weights $w(\mathbf{e}) = J_{i_1 i_2 \dots i_p}$, for $\mathbf{e} = (i_1, i_2, \dots, i_p) \in [N]_p$. The model (1.1) is then a spin system on H_N where the Hamiltonian (1.2) can be rewritten as

$$H_N(\mathbf{X}) = \sum_{\mathbf{e} \in [N]_p} w(\mathbf{e}) X_{\mathbf{e}},$$

where $\mathbf{X} = (X_1, X_2, \dots, X_N) \in \mathcal{C}_N$ and $X_{\mathbf{e}} = X_{i_1} X_{i_2} \dots X_{i_p}$, for $\mathbf{e} = (i_1, i_2, \dots, i_p)$. For a tensor $\mathbf{J}_N = ((J_{i_1 i_2 \dots i_p}))$, define the (weighted) degree of the vertex i_1 as

$$d_{\mathbf{J}_N}(i_1) := \frac{1}{(p-1)!} \sum_{1 \leq i_2, i_3, \dots, i_p \leq N} |J_{i_1 i_2 i_3 \dots i_p}|,$$

² Note that when $p = 2$, $\mathbf{J}_N$ is a Wigner matrix, and hence, by ([5], Theorem 2.12) $\sup_{N \geq 1} \|\mathbf{J}_N\| < \infty$, thus verifying condition (1) of Theorem 2.3.

which is the sum of the absolute values weights of the hyperedges passing through the vertex i_1 . Similarly, define the weighted *co-degree* of the vertices i_1, i_2 as

$$d_{J_N}(i_1, i_2) := \frac{1}{(p-2)!} \sum_{1 \leq i_3, \dots, i_p \leq N} |J_{i_1 i_2 i_3 \dots i_p}|, \quad (2.8)$$

which is the sum of the absolute values of weights of the hyperedges incident on both i_1 and i_2 . Denote by $\mathbf{D}_{J_N} = ((d_{J_N}(i_1, i_2)))_{1 \leq i_1, i_2 \leq N}$, the *co-degree matrix* corresponding to the tensor J_N . The following corollary provides useful sufficient conditions under which the MPL estimate is $\sqrt{N}$ -consistent at all temperatures. The proof is given in Section 5.

COROLLARY 2.6. Suppose $\{J_N\}_{N \geq 1}$ is a sequence of p -tensors such that the following two conditions hold:

- (1) $\sup_{N \geq 1} \|\mathbf{D}_{J_N}\| < \infty$,
- (2) $\liminf_{N \rightarrow \infty} \frac{1}{N} \sum_{1 \leq i_1 < i_2 < \dots < i_p \leq N} J_{i_1 i_2 \dots i_p}^2 > 0$.

Then the MPL estimate $\hat{\beta}_N(X)$ is $\sqrt{N}$ -consistent for all $\beta > 0$.

The proof of this corollary entails verifying the conditions in Theorem 2.3. Towards this, it is easy to check that condition (1) in Corollary 2.6 implies the L_∞ -condition 2.6 and, hence condition (1) in Theorem 2.3. To next step is to show that condition (2) in Corollary (2.6) implies $F_N(\beta) = \Omega(N)$, for β in a neighborhood of zero. (Note that this implies condition (2) in Theorem 2.3 by monotonicity of the function $F_N(\cdot)$.) This follows by a Taylor expansion of $F_N(\beta)$ about zero and the observation $F_N''(0) = (p!)^2 \sum_{1 \leq i_1 < i_2 < \dots < i_p \leq N} J_{i_1 i_2 \dots i_p}^2$. The details are given in Section 5.

REMARK 2.7. Note that, since the L_2 -operator norm of a symmetric matrix is bounded by its L_∞ -operator norm,

$$\begin{aligned} \|\mathbf{D}_{J_N}\| &\leq \max_{1 \leq i_1 \leq N} \sum_{i_2=1}^N d_{J_N}(i_1, i_2) = \frac{1}{(p-2)!} \max_{1 \leq i_1 \leq N} \sum_{1 \leq i_2, i_3, \dots, i_p \leq N} |J_{i_1 i_2 \dots i_p}| \\ &\lesssim_p \max_{1 \leq i_1 \leq N} d_{J_N}(i_1), \end{aligned} \quad (2.9)$$

that is, if a tensor has bounded maximum degree, then condition (1) of Theorem 2.3 holds. This shows that Corollary 2.6 recovers the general theorem of [24], where $\sqrt{N}$ -consistency of the MPL was proved, albeit for a more general model, under condition (2) and condition (1) replaced by the bounded degree assumption $\max_{1 \leq i_1 \leq N} d_{i_1} = O(1)$.

As mentioned earlier, the conditions in Corollary 2.6, neither of which depend on the true parameter β , cannot hold for hypergraphs where the rate of estimation undergoes a phase transition. In fact, as explained in Remark 2.8, the scope of this corollary is really only restricted to Ising models on hypergraphs that are sparse. The importance of the second condition in Theorem 2.3 becomes evident when the hypergraph becomes dense, where $F_N(\beta)$ ceases to be $\Omega(N)$ for all β , and the rate of estimation changes as β varies. This is illustrated in Section 2.3.3 below, where the exact location of the phase transition is derived for Ising models on block hypergraphs.

REMARK 2.8. Suppose $H_N = (V(H_N), E(H_N))$ is a sequence of unweighted p -uniform hypergraphs with vertex set $V(H_N) = [N] = \{1, 2, \dots, N\}$ and edge set $E(H_N)$, with no isolated vertex. Denote by $\mathbf{A}_{H_N} = ((a_{i_1 i_2 \dots i_p}))_{1 \leq i_1, i_2, \dots, i_p \leq N}$ the adjacency tensor of H_N , that is, $a_{i_1 i_2 \dots i_p} = 1$ if $(i_1, i_2, \dots, i_p) \in E(H_N)$ and zero otherwise. Then in order to ensure that a p -spin-system on H_N , as in (1.1), has a non-trivial scaling limit, one needs to consider the scaled tensor,

$$\mathbf{J}_{H_N} = \frac{N}{|E(H_N)|} \mathbf{A}_{H_N}.$$

In this case, the Frobenius norm condition in Corollary 2.6 simplifies to,

$$\frac{1}{N} \|\mathbf{J}_{H_N}\|_F^2 = \frac{N}{|E(H_N)|^2} \sum_{1 \leq i_1, i_2, \dots, i_p \leq N} a_{i_1 i_2 \dots i_p} = \Theta\left(\frac{N}{|E(H_N)|}\right) = \Omega(1). \quad (2.10)$$

This implies, $|E(H_N)| = \Theta(N)$, since H_N has no isolated vertex. Moreover, condition (1) can be written as,

$$\|\mathbf{D}_{\mathbf{A}_{H_N}}\| = O\left(\frac{|E(H_N)|}{N}\right). \quad (2.11)$$

Therefore, combining (2.10), (2.11), and Corollary 2.6, shows that for any sequence of (unweighted) p -uniform hypergraphs $H_N = (V(H_N), E(H_N))$, such that $\|\mathbf{D}_{\mathbf{A}_{H_N}}\| = O(1)$ and $|E(H_N)| = O(N)$, the MPL estimate $\hat{\beta}_N(\mathbf{X})$ in the Ising model (1.1) with interaction tensor $\mathbf{J}_{H_N}$, is $\sqrt{N}$ -consistent for all $\beta > 0$. In particular, by the bound in (2.9) applied to the adjacency tensor $\mathbf{A}_{H_N}$, the MPL estimate $\hat{\beta}_N(\mathbf{X})$ is $\sqrt{N}$ -consistent for all $\beta > 0$, whenever H_N has bounded maximum degree and $O(N)$ edges.

2.3.3 HSBMs The HSBM is a random hypergraph model where each hyperedge is present independently with probability depending on the membership of the vertices to various blocks (see [3, 31, 43] and the references therein for more on the HSBM and its applications in higher-order community detection).

DEFINITION 2.9. (Hypergraph Stochastic Block Model) Fix $p \geq 2$, $K \geq 1$, a vector of community proportions $\boldsymbol{\lambda} := (\lambda_1, \dots, \lambda_K)$, such that $\sum_{j=1}^K \lambda_j = 1$, and a symmetric probability tensor $\boldsymbol{\Theta} := ((\theta_{j_1 \dots j_p}))_{1 \leq j_1, \dots, j_p \leq K}$, where $\theta_{j_1 \dots j_p} \in [0, 1]$, for $1 \leq i_1, \dots, i_p \leq K$. The HSBM with proportion vector $\boldsymbol{\lambda}$ and probability tensor $\boldsymbol{\Theta}$ is a p -uniform hypergraph H_N on $[N] = \{1, 2, \dots, N\}$ vertices with adjacency tensor $\mathbf{A}_{H_N} = ((a_{i_1 i_2 \dots i_p}))_{1 \leq i_1, i_2, \dots, i_p \leq N}$, where

$$a_{i_1 \dots i_p} \sim \text{Ber}\left(\theta_{j_1 \dots j_p}\right) \quad \text{for } i_1 < \dots < i_p \text{ and } (i_1, \dots, i_p) \in \mathcal{B}_{j_1} \times \dots \times \mathcal{B}_{j_p},$$

where $\mathcal{B}_j := (N \sum_{i=1}^{j-1} \lambda_i, N \sum_{i=j}^j \lambda_i] \cap [N]$, for $j \in \{1, \dots, K\}$, and $\{a_{i_1 \dots i_p}\}_{1 \leq i_1 < \dots < i_p \leq |V|}$ are independent. We denote this model by $\mathcal{H}_{p,K,N}(\boldsymbol{\lambda}, \boldsymbol{\Theta})$ and a realization from this model as $H_N \sim \mathcal{H}_{p,K,N}(\boldsymbol{\lambda}, \boldsymbol{\Theta})$.

In this section, we consider the problem of parameter estimation given a sample from an Ising model on a HSBM. The following theorem shows that for the p -tensor Ising models on a HSBM, there

is a critical value of β , below which estimation is impossible, and above which the MPL estimate is $\sqrt{N}$ -consistent. The location of the phase transition is determined by the first time the maximum of a certain variational problem, which arises from the mean-field approximation of the partition function, becomes non-zero. More formally, this is defined as,

$$\beta_{\text{HSBM}}^* := \sup \left\{ \beta \geq 0 : \sup_{(t_1, \dots, t_K) \in [0, 1]^K} \phi_\beta(t_1, \dots, t_K) = 0 \right\}, \quad (2.12)$$

where the function $\phi_\beta : [-1, 1]^K \mapsto \mathbb{R}$ is:

$$\phi_\beta(t_1, t_2, \dots, t_K) := \beta \left(\sum_{1 \leq j_1, \dots, j_p \leq K} \theta_{j_1 \dots j_p} \prod_{\ell=1}^p \lambda_{j_\ell} t_{j_\ell} \right) - \sum_{j=1}^K \lambda_j I(t_j), \quad (2.13)$$

and $I(t) := \frac{1}{2} \{(1+t) \log(1+t) + (1-t) \log(1-t)\}$ is the binary entropy function.

THEOREM 2.10. Fix $p \geq 2$ and let $\{H_N\}_{N \geq 1}$ be a sequence of realizations of HSBMs, with $H_N \sim \mathcal{H}_{p,K,N}(\lambda, \Theta)$, where λ is a proportion vector and Θ is a symmetric probability tensor as in Definition 2.9. Then for almost every (with probability 1) realization of the sequence $\{H_N\}_{N \geq 1}$, given a sample $X \sim \mathbb{P}_\beta$ from the model (1.1), with adjacency tensor $J_N = \frac{1}{N^{p-1}} A_{H_N}$, the following hold:

- (1) The MPL estimate $\hat{\beta}_N(X)$ is $\sqrt{N}$ -consistent for $\beta > \beta_{\text{HSBM}}^*$.
- (2) There does not exist any consistent sequence of estimators for any $\beta < \beta_{\text{HSBM}}^*$.

The proof of the above result is given in Section 6. To show the result in (1) we verify the conditions of Theorem 2.3. Here, we invoke the standard mean-field lower bound to the Gibbs variational representation of the partition function [16], from which it can be easily verified that $F_N(\beta) = \Omega(N)$, whenever $\beta > \beta_{\text{HSBM}}^*$. Perhaps the more interesting consequence of Theorem 2.10 is the result in (2), which shows that not only is the MPL estimate not $\sqrt{N}$ -consistent below the threshold, no estimator is consistent in this regime, let alone $\sqrt{N}$ -consistent. The main argument in this proof is to show that

$$F_N(\beta) = O(1), \quad \text{for } \beta < \beta_{\text{HSBM}}^*. \quad (2.14)$$

Once this is proved, then it can be easily verified that the Kullback–Leibler (KL) divergence between the measures $\mathbb{P}_{\beta_1, p}$ and $\mathbb{P}_{\beta_2, p}$, for any two $0 < \beta_1 < \beta_2 < \beta_{\text{HSBM}}^*$ remains bounded, which in turn implies that the measures $\mathbb{P}_{\beta_1, p}$ and $\mathbb{P}_{\beta_2, p}$ are untestable, and hence inestimable. The main technical difficulty in proving an estimate like (2.14) in tensor models is the absence of ‘Gaussian’ techniques [10, 21], which allows one to compare the partition function of Ising models with quadratic Hamiltonians with an appropriately chosen Gaussian model. This method, unfortunately, does not apply when $p \geq 3$, hence, to estimate the partition function we take the following more direct approach: we first consider the *averaged model* where the interaction tensor is replaced by the expected interaction tensor $\tilde{J}_N := \mathbb{E}J_N$. Using the block structure of the tensor $\tilde{J}_N$ the Hamiltonian in the averaged model can be written in terms of the average of the spins in the different blocks, and hence, the partition function in the averaged model can be accurately estimated using bare-hands combinatorics (Lemma 6.2). We then move from the averaged model to the actual model using standard concentration arguments (Lemma 6.4).

REMARK 2.11. Using the machinery of non-linear large deviations developed in [16], we can in fact show that for the HSBM,

$$\lim_{N \rightarrow \infty} \frac{1}{N} F_N(\beta) = \sup_{(t_1, \dots, t_K) \in [0,1]^K} \phi_\beta(t_1, \dots, t_K), \quad (2.15)$$

with probability 1. Although the proof of this result has not been included in the paper, because for proving Theorem 2.10 (1) we only need to establish a lower bound on $\frac{1}{N} F_N(\beta)$, this is worth mentioning as it motivates the definition of the threshold β_{HSBM}^* and corroborates the result in Theorem 2.10 (1). The result in (2.15) is, however, not strong enough to show that estimation is impossible below the threshold β_{HSBM}^* . Here, we need to understand the asymptotic behavior of $F_N(\beta)$ itself (without scaling by N), which is a more delicate matter that requires arguments beyond the purview of non-linear large deviations and mean-field approximations, as discussed above. In this case, the proof of Theorem 2.10 (2) shows that whenever the log-partition function is $o(N)$, which happens when $\beta < \beta_{\text{HSBM}}^*$, it is actually $O(1)$, and hence, there is a sharp transition from inestimability to $\sqrt{N}$ -consistency.

An important special case of the HSBM is the Erdős–Rényi random hypergraph model, where every hyperedge is present independently with the same fixed probability. More precisely, the HSBM reduces to the classical Erdős–Rényi random p -hypergraph model when the number of blocks $K = 1$. In this case, each hyperedge is present independently with probability $\theta \in (0, 1]$, and the variational problem (2.12) for the threshold simplifies to

$$\beta_{\text{ER}}^*(p, \theta) := \sup \left\{ \beta \geq 0 : \sup_{t \in [0,1]} \{ \beta \theta t^p - I(t) \} = 0 \right\}. \quad (2.16)$$

We will denote this hypergraph model by $\mathcal{G}_p(N, \theta)$. In this case, Theorem 2.10 implies the following:

Corollary 2.12 Erdős–Rényi random hypergraphs. Fix $p \geq 2$, $\theta \in (0, 1]$, and let $\{H_N\}_{N \geq 1}$, with $H_N \sim \mathcal{G}_p(N, \theta)$, be a sequence of realizations of Erdős–Rényi random p -hypergraphs. Then for almost every realization of the sequence $\{H_N\}_{N \geq 1}$ the following hold:

(1) The MPL estimate $\hat{\beta}_N(X)$ is $\sqrt{N}$ -consistent for all $\beta > \beta_{\text{ER}}^*(p, \theta)$.

(2) There does not exist any consistent sequence of estimators for any $\beta < \beta_{\text{ER}}^*(p, \theta)$.

REMARK 2.13. Note that by the change of variable $\kappa = \beta\theta$, it follows that $\beta_{\text{ER}}^*(p, \theta) = \beta_{\text{ER}}^*(p, 1)/\theta$. A simple analysis shows $\beta_{\text{ER}}^*(2, 1) = 0.5$, and hence, $\beta_{\text{ER}}^*(2, \theta) = 0.5/\theta$. For higher values of p , $\beta_{\text{ER}}^*(p, 1)$ can be easily computed numerically. In particular, we have $\beta_{\text{ER}}^*(3, 1) \approx 0.672$ and $\beta_{\text{ER}}^*(4, 1) \approx 0.689$. In fact, $\beta_{\text{ER}}^*(p, 1)$ is strictly increasing in p and $\lim_{p \rightarrow \infty} \beta_{\text{ER}}^*(p, 1) = \log 2$ (see Appendix A for a proof).

Another example is that of random p -partite p -uniform hypergraphs, which are natural extensions of random bipartite graphs.

Example 2.14 Random p -partite p -uniform hypergraphs. A p -uniform hypergraph is said to be p -partite if the vertex set of the hypergraph can be partitioned into p -nonempty sets in such a way that every edge intersects every set of the partition in exactly one vertex. A random p -partite p -uniform hypergraph, is a p -partite p -uniform hypergraph where each edge is present independently with some fixed probability $\theta \in (0, 1]$ [56]. More formally, given a vector $N = (N_1, N_2, \dots, N_p)$ of positive integers, such that

$\sum_{j=1}^p N_j = N$ and $\theta \in (0, 1]$, in the random p -partite p -uniform hypergraph $\mathcal{H}_p(N, \theta)$, the vertex set $[N] = \{1, 2, \dots, N\}$ is partitioned into p disjoint sets $S_1, \dots, S_p$, such that $|S_j| = N_j$ for $1 \leq j \leq p$, and each edge $e \in V_1 \times V_2 \times \dots \times V_p$ is present independently with probability θ . If N is such that $\frac{1}{N}N \rightarrow \lambda = (\lambda_1, \lambda_2, \dots, \lambda_p)$, as $N \rightarrow \infty$, then this is a special case of the HSBM and the threshold (2.12) simplifies to,

$$\beta_{\text{partite}}^*(p, \lambda, \theta) := \sup \left\{ \beta \geq 0 : \sup_{(t_1, \dots, t_p) \in [0, 1]^p} \left\{ \beta \theta \prod_{j=1}^p \lambda_j t_j - \sum_{j=1}^p \lambda_j I(t_j) \right\} = 0 \right\}. \quad (2.17)$$

Theorem 2.10 then implies that the MPL estimate is $\sqrt{N}$ -consistent for all $\beta > \beta_{\text{partite}}^*(p, \lambda, \theta)$, and consistent estimation is impossible for $\beta < \beta_{\text{partite}}^*(p, \lambda, \theta)$. In case the p partitioning sets have asymptotically equal size, that is, $\lambda_j = \frac{1}{p}$ for all $1 \leq j \leq p$, the threshold in (2.17) simplifies further to:

$$\beta_{\text{equipartite}}^*(p, \theta) := \sup \left\{ \beta \geq 0 : \sup_{(t_1, \dots, t_p) \in [0, 1]^p} \left\{ \beta \theta p^{-p} \prod_{j=1}^p t_j - \frac{1}{p} \sum_{j=1}^p I(t_j) \right\} = 0 \right\}. \quad (2.18)$$

Now, a simple analysis shows that $\beta_{\text{equipartite}}^*(p, \theta) = p^p \beta_{\text{ER}}^*(p, \theta)$. The upper bound $\beta_{\text{equipartite}}^*(p, \theta) \leq p^p \beta_{\text{ER}}^*(p, \theta)$ follows by substituting $t_1 = t_2 = \dots = t_p = t \in [0, 1]$ in (2.18) and relating it to (2.16). For the lower bound, note by the convexity of the function $I(x)$ and the AM–GM inequality, that

$$\beta \theta p^{-p} \prod_{j=1}^p t_j - \frac{1}{p} \sum_{j=1}^p I(t_j) \leq \beta \theta p^{-p} \left(\frac{1}{p} \sum_{j=1}^p t_j \right)^p - I \left(\frac{1}{p} \sum_{j=1}^p t_j \right).$$

Then, by the change of variable $\kappa = \beta p^{-p}$, it follows that $\beta_{\text{equipartite}}^*(p, \theta) \geq p^p \beta_{\text{ER}}^*(p, \theta)$.

2.4 Precise fluctuations in the Curie–Weiss model

The p -tensor Curie–Weiss model is the Ising model on the complete p -uniform hypergraph,³ where all the p -tuples of interactions are present [38]. In other words, this is the Ising model on the Erdős–Rényi p -hypergraph with $\theta = 1$. Denoting $\beta_{\text{CW}}^*(p) := \beta_{\text{ER}}^*(p, 1)$, we know from Corollary 2.12 that for $\beta < \beta_{\text{CW}}^*(p)$ consistent estimation is impossible, whereas for $\beta > \beta_{\text{CW}}^*(p)$ the MPL estimate $\hat{\beta}_N(X)$ is $\sqrt{N}$ -consistent. Given that we know the rate of consistency, the next natural question is to wonder whether anything can be said about the limiting distribution of the MPL estimate above the threshold. While tackling this question appears to be extremely difficult, if not impossible, for general models, the special structure of the Curie–Weiss model makes this tractable. This is formalized in the following theorem for $\beta > \beta_{\text{CW}}^*(p)$.

THEOREM 2.15. (Asymptotics of the MPL estimate above the threshold) Fix $p \geq 2$ and consider the p -spin Curie–Weiss model with interaction tensor $\mathbf{J}_N = ((J_{i_1 \dots i_p}))_{1 \leq i_1, \dots, i_p \leq N}$, where $J_{i_1 \dots i_p} = \frac{1}{N^{p-1}}$, for

³ In the complete p -uniform hypergraph with vertex set $[N] = \{1, 2, \dots, N\}$ the set of hyperedges is the collection of all the p -element subsets of $[N]$.




FIG. 1. The histogram $\sqrt{N}(\hat{\beta}_N(X) - \beta)$ in the 4-tensor Curie–Weiss model at $\beta = 0.75 > \beta_{\text{CW}}^*(4) \approx 0.689$ (above the estimation threshold).

all $1 \leq i_1, \dots, i_p \leq N$. Then for every $\beta > \beta_{\text{CW}}^*(p)$, as $N \rightarrow \infty$,

$$\sqrt{N}(\hat{\beta}_N(X) - \beta) \xrightarrow{D} N\left(0, -\frac{g''(m_*)}{p^2 m_*^{2p-2}}\right), \quad (2.19)$$

where $g(t) := \beta t^p - I(t)$, for $t \in [-1, 1]$, and $m_* = m_*(\beta, p)$ is the unique positive global maximizer of g .

The proof of this result is given in Section 7.1. This begins with the observation that in the Curie–Weiss model the MPL estimate can be written as a function of the sample mean $\bar{X}_N = \frac{1}{N} \sum_{i=1}^N X_i$ (see (7.3)). Then combining the recent results on the asymptotic distribution of $\bar{X}_N$ [48] and the delta theorem, we can get the precise fluctuations of the MPL estimate at all points above the estimation threshold $\beta_{\text{CW}}^*(p)$. Figure 1 shows the histogram (over 10^6 replications) of $\sqrt{N}(\hat{\beta}_N(X) - \beta)$ with $p = 4$, $\beta = 0.75$, and $N = 20000$. As predicted by the result above, we see a limiting Gaussian distribution, since $\beta = 0.75 > \beta_{\text{CW}}^*(4) \approx 0.689$ is above the estimation threshold.

The result in Theorem 2.15 can be used to construct a confidence interval for the parameter β for all points above the estimation threshold. Towards this, note, by [(48), Theorem 2.1], that $|\bar{X}_N| \xrightarrow{P} m_*$ under $P_{\beta,p}$, when $\beta > \beta_{\text{CW}}^*(p)$. The result in (2.19) then implies that

$$\left(\hat{\beta}_N(X) - \frac{|\bar{X}_N|^{1-p}}{p} \sqrt{\frac{-g''(|\bar{X}_N|)}{N}} z_{1-\frac{\alpha}{2}}, \hat{\beta}_N(X) + \frac{|\bar{X}_N|^{1-p}}{p} \sqrt{\frac{-g''(|\bar{X}_N|)}{N}} z_{1-\frac{\alpha}{2}} \right),$$

is an interval that contains β with asymptotic coverage probability $1 - \alpha$, whenever $\beta > \beta_{\text{CW}}^*(p)$. (For $\alpha \in (0, 1)$, z_α is the α -th quantile of the standard normal distribution, that is, $\mathbb{P}_\beta(N(0, 1) \leq z_\alpha) = \alpha$.)

REMARK 2.16. (Efficiency of the MPL estimate) An interesting consequence of Theorem 2.15 is that the limiting variance in (2.19) saturates the Cramer–Rao (information) lower bound of the model, when $\beta > \beta_{\text{CW}}^*(p)$. To see this, note that the (scaled) Fisher information in the model (1.1) (recall that the Cramer–Rao lower bound is the inverse of the Fisher information) is given by,

$$I_N(\beta) = \frac{1}{N} \mathbb{E}_\beta \left[\left(\frac{d}{d\beta} \log \mathbb{P}_\beta(X) \right)^2 \right] = \text{Var}_\beta(N^{\frac{1}{2}} \bar{X}_N^p) \rightarrow -\frac{p^2 m_*^{2p-2}}{g''(m_*)},$$

as $N \rightarrow \infty$, where the last step follows from the asymptotics of $\bar{X}_N$ derived in [48]. This implies, for $\beta > \beta_{\text{CW}}^*(p)$, the MPL estimate $\hat{\beta}_N(X)$ is *asymptotically efficient*, which means that no other consistent estimator can have lower asymptotic mean squared error than $\hat{\beta}_N(X)$ above the estimation threshold. While this has been shown for the ML estimate [21; 48], that the MPL estimate, which only maximizes an approximation to the true likelihood, also has this property, is particularly encouraging, as it showcases the effectiveness of the MPL method, both computationally as well as in terms of statistical efficiency.

The results above show that the MPL estimate is $\sqrt{N}$ -consistent and asymptotic efficient whenever $\beta > \beta_{\text{CW}}^*(p)$. On the other hand, for $\beta < \beta_{\text{CW}}^*(p)$, we know from Theorem 2.10 that consistent estimation is impossible. In particular, this means that the MPL estimate is inconsistent for $\beta < \beta_{\text{CW}}^*(p)$. Therefore, the only case that remains is at the threshold $\beta = \beta_{\text{CW}}^*(p)$. Here, the situation is much more delicate. We address this case in the theorem below, which shows that the MPL is $\sqrt{N}$ -consistent for $p = 2$ (with a non-Gaussian limiting distribution), but inconsistent for $p \geq 3$.

THEOREM 2.17. (Asymptotics of the MPL estimate at the threshold) Fix $p \geq 2$ and consider the p -spin Curie–Weiss model with interaction tensor $\mathbf{J}_N = ((J_{i_1 \dots i_p}))_{1 \leq i_1, \dots, i_p \leq N}$, where $J_{i_1 \dots i_p} = \frac{1}{N^{p-1}}$, for all $1 \leq i_1, \dots, i_p \leq N$. Suppose $\beta = \beta_{\text{CW}}^*(p)$. Denote by $m_* = m_*(\beta, p) \in (0, 1)$ the unique positive maximizer of the function $g := \beta t^p - I(t)$, for $t \in [-1, 1]$ and define

$$\alpha := \begin{cases} \frac{1}{1+2[(m_*^2-1)g''(m_*)]^{-\frac{1}{2}}} & \text{if } p \text{ is even,} \\ \frac{1}{1+[(m_*^2-1)g''(m_*)]^{-\frac{1}{2}}} & \text{if } p \text{ is odd.} \end{cases} \quad (2.20)$$

Then, the following hold as $N \rightarrow \infty$,

- (1) If $p = 2$ (recall $\beta_{\text{CW}}^*(2) = \frac{1}{2}$), then for every $t \in \mathbb{R}$,

$$\lim_{N \rightarrow \infty} \mathbb{P}_\beta \left(N^{\frac{1}{2}} \left(\hat{\beta}_N - \frac{1}{2} \right) \leq t \right) = \begin{cases} F(\sqrt{6}t) - F(-\sqrt{6}t) & \text{if } t \geq 0 \\ 0 & \text{if } t < 0 \end{cases} \quad (2.21)$$

where F is a probability distribution function with density given by $dF(t) \propto \exp\left(-\frac{t^4}{12}\right) dt$.

- (2) If $p \geq 3$, then

$$\sqrt{N}(\hat{\beta}_N(X) - \beta) \xrightarrow{D} (1 - \alpha)N \left(0, -\frac{g''(m_*)}{p^2 m_*^{2p-2}} \right) + \alpha \delta_\infty, \quad (2.22)$$

where $g(\cdot)$ is as defined Theorem 2.15 and δ_∞ denotes the point mass at ∞ .

(3) Moreover, at a finer scaling, the following hold:

(a) If $p \geq 4$ is even, then

$$N^{1-\frac{p}{2}} \hat{\beta}_N \xrightarrow{D} \alpha \left(\frac{1}{p|Z|^{p-2}} \right) + (1-\alpha)\delta_0, \quad (2.23)$$

where $Z \sim N(0, 1)$.

(b) If $p \geq 3$ is odd, then

$$N^{1-\frac{p}{2}} \hat{\beta}_N \xrightarrow{D} \frac{\alpha}{2} \left(\frac{1}{p|Z|^{p-2}} \right) + \frac{\alpha}{2}\delta_\infty + (1-\alpha)\delta_0. \quad (2.24)$$

The proof of this result is given in Section 7.1. As in the proof of Theorem 2.15, the main ingredient in the proof of the above result is the asymptotic distribution of the sample mean at the threshold derived in [21, 48]. The reason there is a change in the consistency rates of the MPL estimate as one moves from the 2-spin model to the p -spin model, for $p \geq 3$, is because the rate of convergence of the sample mean $\bar{X}_N$ in the Curie–Weiss model depends on the value of p at the threshold. More precisely, for $p = 2$ and $\beta = \beta_{\text{CW}}^*(2) = \frac{1}{2}$, $N^{\frac{1}{4}} \bar{X}_N \xrightarrow{D} F$, where F is as defined in Theorem 2.17 (1) (see [(21), Proposition 4.1]). On the other hand, when $p \geq 3$ and $\beta = \beta_{\text{CW}}^*(p)$, $N^{\frac{1}{2}} \bar{X}_N$ converges to a mixture of point masses with two or three components depending on whether p is odd or even, respectively (see [(48), Theorem 1.1]).

Taking derivatives in (2.21) shows that for $p = 2$ the MPL estimate has a limiting Gamma distribution with density $g(a) \propto \frac{1}{\sqrt{a}} e^{-3a^2} da$. Figure 2(a) shows the histogram of $\sqrt{N}(\hat{\beta}_N(X) - \beta)$ for $p = 2$ and $\beta = \beta_{\text{CW}}^*(2) = 0.5$, and the limiting density function (plotted in red). On the other hand, for $p \geq 3$, Theorem 2.17 (3) shows that the MPL estimate is inconsistent at the threshold (in fact, $\hat{\beta}_N(X) \xrightarrow{P} \infty$, for $p \geq 3$ and $\beta = \beta_{\text{CW}}^*(p)$). However, even though for $p \geq 3$ the MPL estimate is inconsistent when $\beta = \beta_{\text{CW}}^*(p)$, Theorem 2.17 (2) shows $\sqrt{N}(\hat{\beta}_N(X) - \beta)$ has a Gaussian limit with probability $1 - \alpha$, that is, MPL estimate is $\sqrt{N}$ -consistent at this point with probability $1 - \alpha$. In fact, the proof of Theorem 2.17 (2) shows that $\hat{\beta}_N(X)$ is not $\sqrt{N}$ -consistent at the threshold for $p \geq 3$, only when $\bar{X}_N$ is close to zero. More precisely, the proof shows that $\sqrt{N}(\hat{\beta}_N(X) - \beta) | \{\bar{X}_N \in A_\beta\} \xrightarrow{D} N(0, -\frac{g''(m_*)}{p^2 m_*^{2p-2}})$, if $A_\beta = [-1, 1] \setminus B_0$, where B_0 is a small neighborhood of zero. This is illustrated in Figure 2(b) that plots the histogram of this conditional distribution for $p = 4$ and $\beta = 0.6888 \approx \beta_{\text{CW}}^*(4)$.

2.5 Organization

The rest of the paper is organized as follows. In Section 3 we prove Theorem 3. The proofs of Corollary 2.5 and Corollary 2.6 are given in Section 4 and Section 5, respectively. The proof of Theorem 2.10 is given in Section 6. The proofs of Theorem 2.15 and Theorem 2.17 are given in Section 7. Additional properties of the Curie–Weiss threshold are given in Appendix A.




FIG. 2. (a) The histogram of $\sqrt{N}(\hat{\beta}_N(X) - \beta)$ in the 2-tensor Curie–Weiss model at the estimation threshold ($\beta = \frac{1}{2}$) and the limiting density function (in red) and (b) the histogram of the conditional distribution $\sqrt{N}(\hat{\beta}_N(X) - \beta) | \{X_N \in A_\beta\}$, where A_β is the interval $[-1, 1]$ minus a small neighborhood around zero, in the 4-tensor Curie–Weiss model at the estimation threshold, which has a limiting normal distribution.

3. Proof of Theorem 2.3

In this section, we prove Theorem 2.3. We first state the two main technical estimates required in the proof and show how these results can be used to complete the proof of Theorem 2.3. As mentioned before, the first step in the proof of Theorem 2.3 is to show that the (scaled) log-pseudolikelihood concentrates around zero at the true parameter value $\beta > 0$ at the desired rate. This is achieved by proving the following second-moment estimate on the scaled log-pseudolikelihood function. The proof of this lemma is given in Section 3.1.

LEMMA 3.1. Let $\beta > 0$ be such that assumption (1) of Theorem 2.3 holds. Then

$$\mathbb{E}_\beta \left[s_X^2(\beta) \right] = O_{\beta,p} \left(\frac{1}{N} \right),$$

where $s_X(b) = \frac{1}{pN} \frac{\partial \log L(b|X)}{\partial b} = \frac{1}{N} (H_N(X) - \sum_{i=1}^N m_i(X) \tanh(pbm_i(X)))$.

The next step of the proof is to show the strong concavity of the log-pseudolikelihood, that is, $-\frac{\partial}{\partial \beta} s_X(\beta)$ is strictly positive and bounded away from 0 with high probability. To this end, note that for any $M > 0$,

$$\begin{aligned} -\frac{\partial}{\partial \beta} s_X(\beta) &= \frac{p}{N} \sum_{i=1}^N m_i(X)^2 \operatorname{sech}^2(p\beta m_i(X)) \\ &\geq \frac{p}{N} \operatorname{sech}^2(p\beta M) \sum_{i=1}^N m_i(X)^2 \mathbf{1}_{\{|m_i(X)| \leq M\}}. \end{aligned} \quad (3.1)$$

Therefore, to show that $-\frac{\partial}{\partial\beta}s_X(\beta)$ is strictly positive, it suffices to show that $\sum_{i=1}^N m_i(X)^2 \mathbf{1}\{|m_i(X)| \leq M\}$ is $\Omega(N)$ with high probability. This is formalized in the following lemma, which is proved in Section 3.2.

LEMMA 3.2. Fix $0 < \delta < 1$. Then under the assumptions in Theorem 2.3, there exists $\varepsilon = \varepsilon(\delta, \beta) > 0$ and $M = M(\delta, \beta) < \infty$ such that

$$\mathbb{P}_\beta \left(\sum_{i=1}^N m_i(X)^2 \mathbf{1}\{|m_i(X)| \leq M\} \geq \varepsilon N \right) \geq 1 - \delta,$$

for all N large enough.

The proof of Theorem 2.3 can now be easily completed using the above lemmas. To this end, note that for any $M_1 > 0$, by Chebyshev's inequality and Lemma 3.1 we have,

$$\mathbb{P}_\beta \left(|s_X(\beta)| > \frac{M_1}{\sqrt{N}} \right) \lesssim_{\beta,p} \frac{1}{M_1^2}.$$

Now, fix $\delta > 0$. Therefore, it is possible to choose $M_1 = M_1(\delta, \beta)$ such that the RHS above is less than δ . Next, by Lemma 3.2 there exists $\varepsilon = \varepsilon(\delta, \beta) > 0$ and $M_2 = M_2(\varepsilon, \delta, \beta) < \infty$ such that

$$\mathbb{P}_\beta \left(\sum_{i=1}^N m_i(X)^2 \mathbf{1}\{|m_i(X)| \leq M_2\} \geq \varepsilon N \right) \geq 1 - \delta,$$

for N large enough. Thus, defining

$$T_N := \left\{ X \in \mathcal{C}_N : |s_X(\beta)| \leq \frac{M_1}{\sqrt{N}}, \sum_{i=1}^N m_i(X)^2 \mathbf{1}\{|m_i(X)| \leq M_2\} \geq \varepsilon N \right\},$$

gives $\mathbb{P}_\beta(T_N) \geq 1 - 2\delta$, for N large enough. For $X \in T_N$, recalling (3.1), gives

$$-\frac{\partial}{\partial\beta}s_X(\beta) \geq \frac{p}{N} \operatorname{sech}^2(p\beta M_2) \sum_{i=1}^N m_i(X)^2 \mathbf{1}\{|m_i(X)| \leq M_2\} \geq p\varepsilon \operatorname{sech}^2(p\beta M_2).$$

Therefore, for $X \in T_N$,

$$\begin{aligned} \frac{M_1}{\sqrt{N}} &\geq |s_X(\beta)| = |s_X(\beta) - s_X(\hat{\beta}_N(X))| \geq - \int_{\beta \wedge \hat{\beta}_N(X)}^{\beta \vee \hat{\beta}_N(X)} \frac{\partial}{\partial\beta} s_X(\beta) d\beta \\ &\geq \frac{\varepsilon}{M_2} |\tanh(pM_2 \hat{\beta}_N(X)) - \tanh(pM_2 \beta)|. \end{aligned}$$

Then, defining $M = M(\delta, \beta) := \frac{M_2}{M_1 \varepsilon}$, shows that

$$\mathbb{P}_\beta \left(\sqrt{N} |\tanh(pM_2 \hat{\beta}_N(X)) - \tanh(pM_2 \beta)| \leq R \right) \geq 1 - 2\delta.$$

The proof of Theorem 2.3 now follows by inverting the tanh function.

3.1 Proof of Lemma 3.1

For $\mathbf{x}, \mathbf{x}' \in \mathcal{C}_N$ define

$$F(\mathbf{x}, \mathbf{x}') = \frac{1}{2} \sum_{i=1}^N (m_i(\mathbf{x}) + m_i(\mathbf{x}')) (x_i - x'_i),$$

where m_i is as defined in (2.1). Note that F is antisymmetric, that is, $F(\mathbf{x}, \mathbf{x}') = -F(\mathbf{x}', \mathbf{x})$.

Now, choose a coordinate $I \in \{1, 2, \dots, N\}$ uniformly at random and replace the I -th coordinate of $\mathbf{X} \sim \mathbb{P}_\beta$ by a sample drawn from the conditional distribution of X_u given $(X_v)_{v \neq I}$. Denote the resulting vector by $\mathbf{X}'$. Note that $F(\mathbf{X}, \mathbf{X}') = m_I(\mathbf{X})(X_I - X'_I)$. Then

$$\begin{aligned} f(\mathbf{X}) &:= \mathbb{E}_\beta (F(\mathbf{X}, \mathbf{X}') | \mathbf{X}) = \frac{1}{N} \sum_{i=1}^N m_i(\mathbf{X}) \left\{ X_i - \mathbb{E}_\beta (X_i | (X_j)_{j \neq i}) \right\} \\ &= \frac{1}{N} \sum_{i=1}^N m_i(\mathbf{X}) (X_i - \tanh(p\beta m_i(\mathbf{X}))) \\ &= s_{\mathbf{X}}(\beta). \end{aligned} \tag{3.2}$$

Now, since $(\mathbf{X}, \mathbf{X}')$ is an exchangeable pair,

$$\mathbb{E}_\beta (f(\mathbf{X})F(\mathbf{X}, \mathbf{X}')) = \mathbb{E}_\beta (f(\mathbf{X}')F(\mathbf{X}', \mathbf{X})).$$

Again, because F is antisymmetric, we have $\mathbb{E}_\beta (f(\mathbf{X}')F(\mathbf{X}', \mathbf{X})) = -\mathbb{E}_\beta (f(\mathbf{X}')F(\mathbf{X}, \mathbf{X}'))$. Hence,

$$\begin{aligned} \mathbb{E}_\beta (f(\mathbf{X})^2) &= \mathbb{E}_\beta (f(\mathbf{X}) \mathbb{E}_\beta [F(\mathbf{X}, \mathbf{X}') | \mathbf{X}]) = \mathbb{E}_\beta (f(\mathbf{X})F(\mathbf{X}, \mathbf{X}')) \\ &= \frac{1}{2} \mathbb{E}_\beta ((f(\mathbf{X}) - f(\mathbf{X}'))F(\mathbf{X}, \mathbf{X}')). \end{aligned} \tag{3.3}$$

Now, for any $1 \leq t \leq N$ and $\mathbf{x} \in \mathcal{C}_N$, let

$$\mathbf{x}^{(t)} = (x_1, x_2, \dots, x_{t-1}, 1 - x_t, x_{t+1}, \dots, x_N),$$

and

$$p_t(\mathbf{x}) := \mathbb{P}_\beta(X'_t = -x_t | \mathbf{X} = \mathbf{x}, I = t) = \frac{e^{-p\beta x_t m_t(\mathbf{x})}}{e^{-p\beta m_t(\mathbf{x})} + e^{p\beta m_t(\mathbf{x})}}.$$

This implies

$$\begin{aligned} \mathbb{E}_\beta((f(\mathbf{X}) - f(\mathbf{X}'))F(\mathbf{X}, \mathbf{X}') | \mathbf{X}) &= \frac{1}{N} \sum_{t=1}^N (f(\mathbf{X}) - f(\mathbf{X}^{(t)}))F(\mathbf{X}, \mathbf{X}^{(t)})p_t(\mathbf{X}) \\ &= \frac{2}{N} \sum_{t=1}^N m_t(\mathbf{X})X_t p_t(\mathbf{X}) (f(\mathbf{X}) - f(\mathbf{X}^{(t)})). \end{aligned} \quad (3.4)$$

For $1 \leq s, t \leq N$, let $a_s(\mathbf{x}) := x_s - \tanh(p\beta m_s(\mathbf{x}))$ and $b_{st}(\mathbf{x}) := \tanh(p\beta m_s(\mathbf{x})) - \tanh(p\beta m_s(\mathbf{x}^{(t)}))$. Then, noting that $f(\mathbf{x}) = \frac{1}{N} \sum_{s=1}^N m_s(\mathbf{x})a_s(\mathbf{x})$ gives

$$\begin{aligned} f(\mathbf{X}) - f(\mathbf{X}^{(t)}) &= \frac{1}{N} \sum_{s=1}^N (m_s(\mathbf{X}) - m_s(\mathbf{X}^{(t)}))a_s(\mathbf{X}) + \frac{1}{N} \sum_{s=1}^N m_s(\mathbf{X}^{(t)})(a_s(\mathbf{X}) - a_s(\mathbf{X}^{(t)})) \\ &= A_t + B_t + C_t, \end{aligned} \quad (3.5)$$

where $A_t := \frac{2(p-1)X_t}{N} \sum_{s=1}^N J_{st}(\mathbf{X})a_s(\mathbf{X})$, $B_t := \frac{2m_t(\mathbf{X})X_t}{N}$, and $C_t := -\frac{1}{N} \sum_{s=1}^N m_s(\mathbf{X}^{(t)})b_{st}(\mathbf{X})$. Then using (3.3) and (3.4), we have

$$\mathbb{E}_\beta(f(\mathbf{X})^2) = \frac{1}{N} \sum_{t=1}^N \mathbb{E}_\beta[(A_t + B_t + C_t)m_t(\mathbf{X})X_t p_t(\mathbf{X})].$$

Now, define the following three quantities:

$$\mathbf{a}(\mathbf{X}) := (a_1(\mathbf{X}), \dots, a_N(\mathbf{X})), \quad \mathbf{m}(\mathbf{X}) := (m_1(\mathbf{X}), \dots, m_N(\mathbf{X})),$$

and $\mathbf{M}(\mathbf{X}) := (m_1(\mathbf{X})p_1(\mathbf{X}), \dots, m_N(\mathbf{X})p_N(\mathbf{X}))$. Note that $\mathbf{m}(\mathbf{X}) = \mathbf{X}\mathbf{J}_N(\mathbf{X})^\top$. Also, observe that each entry of $\mathbf{A}(\mathbf{X})$ is bounded in absolute value by 2, hence, $\|\mathbf{A}(\mathbf{X})\| \leq 2\sqrt{N}$. Moreover, using $p_t(\mathbf{X}) \leq 1$,

$$\|\mathbf{M}(\mathbf{X})\| = \left(\sum_{i=1}^N m_i(\mathbf{X})^2 p_i(\mathbf{X})^2 \right)^{\frac{1}{2}} \leq \|\mathbf{m}(\mathbf{X})\| = \|\mathbf{J}_N(\mathbf{X})\mathbf{X}^\top\| \leq \sqrt{N}\|\mathbf{J}_N(\mathbf{X})\|. \quad (3.6)$$

Hence, recalling the definition of A_t from (3.5) gives

$$\begin{aligned} \left| \frac{1}{N} \sum_{t=1}^N A_t m_t(\mathbf{X}) X_t p_t(\mathbf{X}) \right| &= \frac{2(p-1)}{N^2} \left| \mathbf{a}(\mathbf{X}) \mathbf{J}_N(\mathbf{X}) \mathbf{M}(\mathbf{X})^\top \right| \\ &\lesssim_p \frac{1}{N^2} \|\mathbf{J}_N(\mathbf{X})\| \|\mathbf{a}(\mathbf{X})\| \|\mathbf{M}(\mathbf{X})\| \\ &\lesssim_p \frac{\|\mathbf{J}_N(\mathbf{X})\|^2}{N}. \end{aligned} \quad (3.7)$$

Next, we consider the term corresponding to B_t :

$$\left| \frac{1}{N} \sum_{t=1}^N B_t m_t(\mathbf{X}) X_t p_t(\mathbf{X}) \right| = \frac{2}{N^2} \left| \sum_{t=1}^N m_t^2(\mathbf{X}) p_t(\mathbf{X}) \right| \lesssim \frac{1}{N^2} \|\mathbf{m}(\mathbf{X})\|^2 \leq \frac{\|\mathbf{J}_N(\mathbf{X})\|^2}{N}, \quad (3.8)$$

where the last step uses (3.6).

Finally, we consider the term corresponding to C_t . Let us define the matrix $\mathbf{J}_{N,2}(\mathbf{X}) := ((J_{ij}(\mathbf{X})^2))_{1 \leq i,j \leq N}$. Then, denoting by $\mathbf{e}_i$ the vector in $\mathbb{R}^N$ with the i -th entry 1 and 0 everywhere else, we get

$$\|\mathbf{J}_{N,2}(\mathbf{X})\| \leq \max_{1 \leq i \leq N} \sum_{j=1}^N J_{ij}(\mathbf{X})^2 \leq \max_{1 \leq i \leq N} \|\mathbf{e}_i^\top \mathbf{J}_N(\mathbf{X})\|^2 \leq \|\mathbf{J}_N(\mathbf{X})\|^2.$$

Let $h(x) := \tanh(p\beta x)$. It is easy to check that $\|h''\|_\infty \leq \beta^2$. Hence, by a Taylor expansion, for $1 \leq s \leq N$,

$$\left| h(m_s(\mathbf{X})) - h(m_s(\mathbf{X}^{(t)})) - (m_s(\mathbf{X}) - m_s(\mathbf{X}^{(t)})) h'(m_s(\mathbf{X})) \right| \lesssim_\beta \left(m_s(\mathbf{X}) - m_s(\mathbf{X}^{(t)}) \right)^2. \quad (3.9)$$

Note that $m_s(\mathbf{X}) - m_s(\mathbf{X}^{(t)}) = 2(p-1)J_{st}(\mathbf{X})X_t$ and $h(m_s(\mathbf{X})) - h(m_s(\mathbf{X}^{(t)})) = \tanh(p\beta m_s(\mathbf{X})) - \tanh(p\beta m_s(\mathbf{X}^{(t)})) = b_{st}(\mathbf{X})$, hence, (3.9) can be rewritten as:

$$\left| b_{st}(\mathbf{X}) - 2(p-1)J_{st}(\mathbf{X})X_t h'(m_s(\mathbf{X})) \right| \lesssim_{\beta,p} J_{st}(\mathbf{X})^2.$$

Using the above bounds, we have, for any $\mathbf{x} = (x_1, x_2, \dots, x_N), \mathbf{y} = (y_1, y_2, \dots, y_N) \in \mathbb{R}^N$,

$$\begin{aligned}
& \left| \sum_{1 \leq s, t \leq N} x_s y_t b_{st}(\mathbf{X}) \right| \\
& \leq \left| \sum_{1 \leq s, t \leq N} 2(p-1) x_s y_t J_{st}(\mathbf{X}) X_t h'(m_s(\mathbf{X})) \right| + \left| \sum_{1 \leq s, t \leq N} x_s y_t (b_{st}(\mathbf{X}) - 2(p-1) J_{st}(\mathbf{X}) X_t h'(m_s(\mathbf{X}))) \right| \\
& \lesssim_{\beta, p} \|J_N(\mathbf{X})\| \left(\sum_{s=1}^N (x_s h'(m_s(\mathbf{X})))^2 \right)^{\frac{1}{2}} \left(\sum_{t=1}^N (y_t X_t)^2 \right)^{\frac{1}{2}} + \sum_{1 \leq s, t \leq N} |x_s y_t| J_{st}(\mathbf{X})^2 \\
& \lesssim_{\beta, p} \|J_N(\mathbf{X})\| \|\mathbf{x}\| \|\mathbf{y}\| + \|J_{N,2}(\mathbf{X})\| \|\mathbf{x}\| \|\mathbf{y}\| \quad (\text{using } h'(m_s(\mathbf{X})) \lesssim_{\beta, p} 1) \\
& \lesssim_{\beta, p} \left(\|J_N(\mathbf{X})\| + \|J_N(\mathbf{X})\|^2 \right) \|\mathbf{x}\| \|\mathbf{y}\|. \tag{3.10}
\end{aligned}$$

Again, by a Taylor expansion and using the bound $\|h'\|_\infty \lesssim_{\beta, p} 1$, gives

$$|b_{st}(\mathbf{X})| \lesssim_{\beta, p} |m_s(\mathbf{X}) - m_s(\mathbf{X}^{(t)})| \lesssim_{\beta, p} |J_{st}(\mathbf{X})|.$$

Consequently,

$$\left| \sum_{1 \leq s, t \leq N} x_s y_t J_{st}(\mathbf{X}) b_{st}(\mathbf{X}) \right| \lesssim_{\beta, p} \sum_{1 \leq s, t \leq N} |x_s y_t| J_{st}(\mathbf{X})^2 \leq \|J_N(\mathbf{X})\|^2 \|\mathbf{x}\| \|\mathbf{y}\|. \tag{3.11}$$

Now, recalling the definition of C_t from (3.5) gives,

$$\begin{aligned}
\left| \frac{1}{N} \sum_{t=1}^N C_t m_t(\mathbf{X}) X_t p_t(\mathbf{X}) \right| &= \left| \frac{1}{N^2} \sum_{1 \leq s, t \leq N} m_s(\mathbf{X}^{(t)}) b_{st}(\mathbf{X}) m_t(\mathbf{X}) X_t p_t(\mathbf{X}) \right| \\
&= \left| \frac{1}{N^2} \sum_{1 \leq s, t \leq N} (m_s(\mathbf{X}) - 2(p-1) J_{st}(\mathbf{X}) X_t) b_{st}(\mathbf{X}) m_t(\mathbf{X}) X_t p_t(\mathbf{X}) \right| \\
&\lesssim_{\beta, p} \frac{\|J_N(\mathbf{X})\|^3 + \|J_N(\mathbf{X})\|^4}{N}, \tag{3.12}
\end{aligned}$$

where the last step uses (3.10), (3.11), and $\|\mathbf{m}(\mathbf{X})\| = \|J_N(\mathbf{X}) \mathbf{X}^\top\| \leq \sqrt{N} \|J_N(\mathbf{X})\|$.

Combining (3.7), (3.8) and (3.12), it follows that $\mathbb{E}_\beta(f(\mathbf{X})^2) \lesssim_{\beta, p} \frac{1}{N}$, since by condition of (1) of Theorem 2.3, $\mathbb{E}_\beta(\|J_N(\mathbf{X})\|^4)$ is uniformly bounded. This completes the proof of the lemma, since recalling (3.2), $f(\mathbf{X}) = s_{\mathbf{X}}(\beta)$.

3.2 Proof of Lemma 3.2

We begin with the following simple observation, which says that if $\liminf_{N \rightarrow \infty} \frac{1}{N} F_N(\beta) > 0$, we can find a γ small enough such that $\liminf_{N \rightarrow \infty} \frac{1}{N} F'_N(\beta - \gamma) > 0$.

Observation 3.3. Suppose $\beta > 0$ is such that $\liminf_{N \rightarrow \infty} \frac{1}{N} F_N(\beta) > 0$. Then

$$\lim_{\delta \rightarrow 0} \liminf_{N \rightarrow \infty} \frac{1}{N} F_N(\beta - \delta) > 0.$$

Proof. Denote $K := \sup_{N \geq 1} \mathbb{E}_\beta(|J_N(\mathbf{X})|) < \infty$ Then,

$$F'_N(\beta) := \frac{d}{d\beta} F_N(\beta) = \mathbb{E}_\beta(H_N(\mathbf{X})) = \mathbb{E}_\beta(\mathbf{X}' J_N(\mathbf{X}) \mathbf{X}) \leq N \mathbb{E}_\beta(|J_N(\mathbf{X})|) \leq KN.$$

Therefore, by a Taylor expansion

$$\lim_{\delta \rightarrow 0} \liminf_{N \rightarrow \infty} \frac{1}{N} F_N(\beta - \delta) \geq \lim_{\delta \rightarrow 0} \liminf_{N \rightarrow \infty} \left(\frac{1}{N} F_N(\beta) - \delta K \right) > 0,$$

as required. $\square$

Now, note that for any $\varepsilon, \gamma > 0$,

$$\mathbb{P}_\beta(H_N(\mathbf{X}) < \varepsilon N) = \mathbb{P}_\beta(e^{-\gamma H_N(\mathbf{X})} > e^{-\gamma \varepsilon N}) \leq e^{\gamma \varepsilon N + F_N(\beta - \gamma) - F_N(\beta)},$$

which, on taking logarithms, implies that

$$\log \mathbb{P}_\beta(H_N(\mathbf{X}) < \varepsilon N) \leq \varepsilon \gamma N - \int_{\beta - \gamma}^{\beta} F'_N(t) dt \leq \varepsilon \gamma N - F'_N(\beta - \gamma) \gamma,$$

by the monotonicity of $F'_N(\cdot)$. Dividing both sides by N and taking limits as $N \rightarrow \infty$ followed by $\varepsilon \rightarrow 0$ we have

$$\lim_{\varepsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{P}_\beta(H_N(\mathbf{X}) < \varepsilon N) \leq -\liminf_{N \rightarrow \infty} \frac{1}{N} F'_N(\beta - \gamma) \leq -\liminf_{N \rightarrow \infty} \frac{F_N(\beta - \gamma)}{N(\beta - \gamma)} < 0,$$

by choosing γ small enough (by Observation 3.3). This shows that, for every $0 < \delta < 1$ there exists $\varepsilon = \varepsilon(\delta) > 0$ such that, for N large enough,

$$\mathbb{P}_\beta(H_N(\mathbf{X}) < 2\varepsilon N) \leq \delta. \quad (3.13)$$

Next, by Lemma 3.1 and Chebyshev's inequality, there exists $M_1 = M_1(\delta) < \infty$ such that

$$\mathbb{P}_\beta \left(|s_{\mathbf{X}}(\beta)| > \frac{M_1}{\sqrt{N}} \right) \leq \delta. \quad (3.14)$$

Moreover, note that for any $M_2 > 0$,

$$\sum_{i=1}^N |m_i(\mathbf{X})| \mathbf{1}\{|m_i(\mathbf{X})| > M_2\} \leq \frac{1}{M_2} \sum_{i=1}^N m_i(\mathbf{X})^2 = \frac{1}{M_2} \|J_N(\mathbf{X})\mathbf{X}^\top\|^2 \leq \frac{N\|J_N(\mathbf{X})\|^2}{M_2}.$$

Therefore, using Markov's inequality and condition (1) of Theorem 2.3, we can choose $M_2 = M_2(\delta) < \infty$ such that for all N large enough,

$$\mathbb{P}_\beta \left(\sum_{i=1}^N |m_i(\mathbf{X})| \mathbf{1}\{|m_i(\mathbf{X})| > M_2\} > \varepsilon N \right) \leq \frac{\mathbb{E}_\beta(\|J_N(\mathbf{X})\|^2)}{\varepsilon M_2} \leq \delta. \quad (3.15)$$

Then, defining

$$T_N := \left\{ \mathbf{X} \in \mathcal{C}_N : H_N(\mathbf{X}) \geq 2\varepsilon N, |s_{\mathbf{X}}(\beta)| \leq \frac{M_1}{\sqrt{N}}, \sum_{i=1}^N |m_i(\mathbf{X})| \mathbf{1}\{|m_i(\mathbf{X})| > M_2\} \leq \varepsilon N \right\},$$

and combining (3.13), (3.14), and (3.15), gives $\mathbb{P}_\beta(T_N) \geq 1 - 3\delta$, for N large enough. Now, for $\mathbf{X} \in T_N$,

$$\begin{aligned} & \sum_{i=1}^N m_i(\mathbf{X})^2 \mathbf{1}\{|m_i(\mathbf{X})| \leq M_2\} + \varepsilon N \\ & \geq \frac{1}{p\beta} \sum_{i=1}^N |m_i(\mathbf{X})| \tanh(p\beta |m_i(\mathbf{X})|) \mathbf{1}\{|m_i(\mathbf{X})| \leq M_2\} + \sum_{i=1}^N |m_i(\mathbf{X})| \mathbf{1}\{|m_i(\mathbf{X})| > M_2\} \\ & \quad \text{(using } \tanh x \leq x) \\ & \gtrsim_{p,\beta} \sum_{i=1}^N |m_i(\mathbf{X})| \tanh(p\beta |m_i(\mathbf{X})|) \mathbf{1}\{|m_i(\mathbf{X})| \leq M_2\} + \sum_{i=1}^N |m_i(\mathbf{X})| \tanh(p\beta |m_i(\mathbf{X})|) \mathbf{1}\{|m_i(\mathbf{X})| > M_2\} \\ & \quad \text{(using } \tanh x \leq 1) \\ & \geq \sum_{i=1}^N m_i(\mathbf{X}) \tanh(p\beta m_i(\mathbf{X})) \\ & = H_N(\mathbf{X}) - N s_{\mathbf{X}}(\beta) \geq 2\varepsilon N - M_1 \sqrt{N}. \end{aligned}$$

Thus, on the set T_N ,

$$\sum_{i=1}^N m_i(\mathbf{X})^2 \mathbf{1}\{|m_i(\mathbf{X})| \leq M_2\} \gtrsim_{p,\beta} 2\varepsilon N - M_1 \sqrt{N} > \varepsilon N,$$

for all N large enough. This completes the proof of Lemma 3.2.

4. Proof of Corollary 2.5

To prove Corollary 2.5 we will verify that the conditions in Theorem 2.3 hold with probability 1. As mentioned before, in this case condition (2) is easy to verify. To this end, note that by [(11), Theorem 1.1], $\lim_{N \rightarrow \infty} \frac{1}{N} F_N(\beta) = \beta^2/2$ almost surely, for $\beta > 0$ small enough. This implies, since F_N on increasing on the positive half-line, $\lim_{N \rightarrow \infty} \frac{1}{N} F_N(\beta) > 0$ almost surely, for all $\beta > 0$. This establishes condition (2) in Theorem 2.3.

We now proceed to verify condition (1). To begin with, fix $\mathbf{x} \in \mathcal{C}_N$ and consider the Gaussian process

$$G_{\mathbf{u}}(\mathbf{x}) := \mathbf{u}^\top \mathbf{J}_N(\mathbf{x}) \mathbf{u} \quad (4.1)$$

indexed by $\mathbf{u} \in S^{N-1} := \{\mathbf{t} \in \mathbb{R}^N : \|\mathbf{t}\| = 1\}$. Here, $\mathbf{J}_N(\mathbf{x})$ is the local interaction matrix corresponding to the tensor (2.7) of the p -tensor SK model. Note that the maximum eigenvalue of $\mathbf{J}_N(\mathbf{x})$ can be expressed as $\lambda_{\max}(\mathbf{J}_N(\mathbf{x})) = \sup_{\mathbf{u} \in S^{N-1}} G_{\mathbf{u}}(\mathbf{x})$.

LEMMA 4.1. Fix $\mathbf{x} \in \mathcal{C}_N$ and consider the Gaussian process $\{G_{\mathbf{u}}(\mathbf{x}) : \mathbf{u} \in S^{N-1}\}$ as defined above in (4.1). Then, the following hold:

- (1) For every vector $\mathbf{u} \in S^{N-1}$, $\mathbb{E}[G_{\mathbf{u}}(\mathbf{x})^2] \lesssim_p \frac{1}{N}$;
- (2) For vectors $\mathbf{u}, \mathbf{v} \in S^{N-1}$, $\mathbb{E}[G_{\mathbf{u}}(\mathbf{x}) - G_{\mathbf{v}}(\mathbf{x})]^2 \lesssim_p \frac{1}{N} \sum_{i=1}^N (u_i - v_i)^2$.

Proof of (1): Fix $\mathbf{x} \in \mathcal{C}_N$ and $\mathbf{u} \in S^{N-1}$. Then

$$\begin{aligned} G_{\mathbf{u}}(\mathbf{x}) &= \sum_{1 \leq i_1, \dots, i_p \leq N} J_{i_1 \dots i_p} u_{i_1} u_{i_2} x_{i_3} \dots x_{i_p} \\ &= (p-2)! \sum_{1 \leq i_1 < \dots < i_p \leq N} J_{i_1 \dots i_p} \left(\sum_{1 \leq s \neq t \leq p} u_{i_s} u_{i_t} \prod_{a \in \{1, 2, \dots, p\} \setminus \{s, t\}} x_{i_a} \right). \end{aligned}$$

Hence,

$$\begin{aligned} \mathbb{E}[G_{\mathbf{u}}(\mathbf{x})^2] &\lesssim_p \frac{1}{N^{p-1}} \sum_{1 \leq i_1 < \dots < i_p \leq N} \left(\sum_{1 \leq s \neq t \leq p} u_{i_s} u_{i_t} \prod_{a \in \{1, 2, \dots, p\} \setminus \{s, t\}} x_{i_a} \right)^2 \\ &\lesssim_p \frac{1}{N^{p-1}} \sum_{1 \leq i_1 < \dots < i_p \leq N} \sum_{1 \leq s \neq t \leq p} u_{i_s}^2 u_{i_t}^2 \\ &\lesssim_p \frac{1}{N^{p-1}} \sum_{1 \leq s \neq t \leq p} \sum_{1 \leq i_1, \dots, i_p \leq N} u_{i_s}^2 u_{i_t}^2 \\ &= \frac{1}{N^{p-1}} \sum_{1 \leq s \neq t \leq p} N^{p-2} \|\mathbf{u}\|^4 \lesssim_p \frac{1}{N}. \end{aligned} \quad (4.2)$$

where in (4.2) we used the inequality $(\sum_{i=1}^n a_i)^2 \leq n \sum_{i=1}^n a_i^2$, for any sequence of real numbers $a_1, \dots, a_n$.

Proof of (2): Fix $\mathbf{x} \in \mathcal{C}_N$ and $\mathbf{u}, \mathbf{v} \in S^{N-1}$. Then,

$$\begin{aligned}
 & \mathbb{E} [G_{\mathbf{u}}(\mathbf{x}) - G_{\mathbf{v}}(\mathbf{x})]^2 \\
 &= ((p-2)!)^2 \mathbb{E} \left[\sum_{1 \leq i_1 < \dots < i_p \leq N} J_{i_1 \dots i_p} \left(\sum_{1 \leq s \neq t \leq p} (u_{i_s} u_{i_t} - v_{i_s} v_{i_t}) \prod_{a \in \{1, 2, \dots, p\} \setminus \{s, t\}} \tau_{i_a} \right) \right]^2 \\
 &\lesssim_p \frac{1}{N^{p-1}} \sum_{1 \leq i_1 < \dots < i_p \leq N} \left(\sum_{1 \leq s \neq t \leq p} (u_{i_s} u_{i_t} - v_{i_s} v_{i_t}) \prod_{a \in \{1, 2, \dots, p\} \setminus \{s, t\}} \tau_{i_a} \right)^2 \\
 &\lesssim_p \frac{1}{N^{p-1}} \sum_{1 \leq i_1 < \dots < i_p \leq N} \sum_{1 \leq s \neq t \leq p} (u_{i_s} u_{i_t} - v_{i_s} v_{i_t})^2 \\
 &\lesssim_p \frac{1}{N} \sum_{1 \leq i, j \leq N} (u_i u_j - v_i v_j)^2 \\
 &= \frac{2}{N} \left[1 - \left(\sum_{i=1}^N u_i v_i \right)^2 \right] \leq \frac{2}{N} \left[2 - 2 \sum_{i=1}^N u_i v_i \right] \lesssim \frac{1}{N} \sum_{i=1}^N (u_i - v_i)^2.
 \end{aligned}$$

This completes the proof of Lemma 4.1 (2).

Using the lemma above we first show that $\mathbb{E}[\sup_{\mathbf{u} \in S^{N-1}} G_{\mathbf{u}}(\mathbf{x})] \lesssim_p 1$. We do this comparing the supremum of the Gaussian process $\{G_{\mathbf{u}}(\mathbf{x}) : \mathbf{u} \in S^{N-1}\}$ with the supremum of the Gaussian process $\{H_{\mathbf{u}} : \mathbf{u} \in S^{N-1}\}$, where $H_{\mathbf{u}} = \sum_{i=1}^N g_i u_i$ and $g_1, \dots, g_N$ are independent standard Gaussians. Now, by Lemma 4.1 (2), there exists a constant $C := C(p) > 0$, such that for $\mathbf{u}, \mathbf{v} \in S^{N-1}$,

$$\mathbb{E} [G_{\mathbf{u}}(\mathbf{x}) - G_{\mathbf{v}}(\mathbf{x})]^2 \leq \frac{C}{N} \sum_{i=1}^N (u_i - v_i)^2 = \frac{C}{N} \mathbb{E} [H_{\mathbf{u}} - H_{\mathbf{v}}]^2.$$

Hence, by the Sudakov–Fernique inequality [(17), Theorem 1.1],

$$\begin{aligned}
 \mathbb{E} \left[\sup_{\mathbf{u} \in S^{N-1}} G_{\mathbf{u}}(\mathbf{x}) \right] &\leq \left(\frac{C}{N} \right)^{\frac{1}{2}} \mathbb{E} \left[\sup_{\mathbf{u} \in S^{N-1}} H_{\mathbf{u}} \right] = \left(\frac{C}{N} \right)^{\frac{1}{2}} \mathbb{E} \left[\left(\sum_{i=1}^N g_i^2 \right)^{\frac{1}{2}} \right] \\
 &\leq C^{\frac{1}{2}} := D.
 \end{aligned} \tag{4.3}$$

Now, since by Lemma 4.1 (1) there exists a constant $K := K(p) > 0$ such that $\sup_{\mathbf{u} \in S^{N-1}} \mathbb{E} [G_{\mathbf{u}}(\mathbf{x})^2] \leq \frac{K}{N}$, by the Borell–Tsirelson inequality [(27), Theorem 2.1.1], for any $t > 0$,

$$\mathbb{P} \left(\sup_{\mathbf{u} \in S^{N-1}} G_{\mathbf{u}}(\mathbf{x}) - \mathbb{E} \left[\sup_{\mathbf{u} \in S^{N-1}} G_{\mathbf{u}}(\mathbf{x}) \right] > t \right) \leq e^{-\frac{Nt^2}{2K}}.$$

This implies, by (4.3),

$$\mathbb{P} \left(\sup_{\mathbf{u} \in S^{N-1}} G_{\mathbf{u}}(\mathbf{x}) > D + t \right) \leq e^{-\frac{Nt^2}{2K}}.$$

Then, taking $t = \sqrt{2K}$ in the inequality above gives,

$$\mathbb{P} \left(\|J_N(\mathbf{x})\| > D + \sqrt{2K} \right) \leq 2e^{-N},$$

since we have $\lambda_{\max}(J_N(\mathbf{x})) \stackrel{D}{=} -\lambda_{\min}(J_N(\mathbf{x}))$, because $J_N(\mathbf{x}) \stackrel{D}{=} -J_N(\mathbf{x})$. Therefore, by an union bound,

$$\mathbb{P} \left(\sup_{\mathbf{x} \in \mathcal{C}_N} \|J_N(\mathbf{x})\| > D + \sqrt{2K} \right) \leq 2^{N+1} e^{-N} = 2(e/2)^{-N}.$$

Hence, by the Borel–Cantelli lemma, $\limsup_{N \rightarrow \infty} \sup_{\mathbf{x} \in \mathcal{C}_N} \|J_N(\mathbf{x})\| \leq D + \sqrt{2K}$ with probability 1, which establishes (2.5) and hence, condition (1) of Theorem 2.3.

5. Proof of Corollary 2.6

We begin by showing that condition (1) of Corollary 2.6 implies $\sup_{N \geq 1} \sup_{\mathbf{x} \in \mathcal{C}_N} \|J_N(\mathbf{x})\| < \infty$ and, hence, condition (1) of Theorem 2.3. To this end, fix $\mathbf{x} \in \mathcal{C}_N$ and take $\mathbf{u} \in S^{N-1} := \{\mathbf{t} \in \mathbb{R}^N : \|\mathbf{t}\| = 1\}$. Then,

$$\begin{aligned} |\mathbf{u}^\top J_N(\mathbf{x}) \mathbf{u}| &= \left| \sum_{1 \leq i_1, i_2, \dots, i_p \leq N} J_{i_1 i_2 \dots i_p} u_{i_1} u_{i_2} x_{i_3} \dots x_{i_p} \right| \leq \sum_{1 \leq i_1, i_2, \dots, i_p \leq N} |J_{i_1 i_2 \dots i_p}| |u_{i_1}| |u_{i_2}| \\ &= (p-2)! |\mathbf{u}|^\top \mathbf{D}_{J_N} |\mathbf{u}| \lesssim_p \|\mathbf{D}_{J_N}\|, \end{aligned} \quad (5.1)$$

since $|\mathbf{u}| := (|u_1|, \dots, |u_N|) \in S^{N-1}$. Taking supremum over all $\mathbf{u} \in S^{N-1}$ followed by the supremum over all $\mathbf{x} \in \mathcal{C}_N$ and further followed by the supremum over all $N \geq 1$ throughout (5.1), we have:

$$\sup_{N \geq 1} \sup_{\mathbf{x} \in \mathcal{C}_N} \|J_N(\mathbf{x})\| \lesssim_p \sup_{N \geq 1} \|\mathbf{D}_{J_N}\| < \infty.$$

Next, we verify condition (2) in Theorem 2.3. To this end, we need the following lemma:

LEMMA 5.1. For every $p \geq 2$, under the assumptions of Corollary 2.6, $|F_N^{(3)}(0)| = O(N)$, where $F^{(3)}(0)$ denotes the third derivative of $F_N(\beta)$ at $\beta = 0$.

The proof of the lemma is given below. First we show how it can be used to prove condition (2) in Theorem 2.3. To begin with, note that condition (2) of Corollary 2.6 implies that

$$F_N''(0) = \text{Var}_0(H_N(X)) = \mathbb{E}_0 H_N^2(X) = (p!)^2 \sum_{1 \leq i_1 < \dots < i_p \leq N} J_{i_1 \dots i_p}^2 = \Omega(N). \quad (5.2)$$

Hence, $\liminf_{N \rightarrow \infty} \frac{1}{N} F_N''(0) > 0$. Now, since by Lemma 5.1, $\limsup_{N \rightarrow \infty} \frac{1}{N} |F_N^{(3)}(0)| < \infty$, we can choose $\varepsilon > 0$ small enough, such that for all N large enough,

$$\frac{\varepsilon}{6N} |F_N^{(3)}(0)| < \frac{F_N''(0)}{4N}.$$

Therefore, because the fourth derivative $F_N^{(4)}(b) = \mathbb{E}_b H_N^4(X) \geq 0$ for all $b \geq 0$, a Taylor expansion gives the following for all $\beta \in (0, \varepsilon)$:

$$\frac{F_N(\beta)}{N} \geq \frac{\beta^2}{2N} F_N''(0) + \frac{\beta^3}{6N} F_N^{(3)}(0) \geq \frac{\beta^2}{2N} F_N''(0) - \frac{\beta^3}{6N} |F_N^{(3)}(0)| > \frac{\beta^2}{4N} F_N''(0) = \Omega(1),$$

where the last step uses (5.2). This verifies condition (2) of Theorem 2.3 for all $\beta > 0$, by the monotonicity of F_N .

Proof. Proof of Lemma 5.1: To begin with observe that $F_N^{(3)}(0) = \mathbb{E}_0 H_N^3(X)$. Now, the proof of the lemma for odd p is trivial. This is because, under $\mathbb{P}_0$, $X \stackrel{D}{=} -X$, and for odd p , $H_N(-X) = -H_N(X)$, which implies $\mathbb{E}_0 H_N^3(X) = 0$. Hence, we will assume that $p = 2q$, for $q \geq 1$, throughout the rest of the proof. Now, note that

$$\mathbb{E}_0 H_N^3(X) = \sum_{\substack{1 \leq i_1, \dots, i_p \leq N \\ \text{distinct}}} \sum_{\substack{1 \leq j_1, \dots, j_p \leq N \\ \text{distinct}}} \sum_{\substack{1 \leq k_1, \dots, k_p \leq N \\ \text{distinct}}} J_{i_1, \dots, i_p} J_{j_1, \dots, j_p} J_{k_1, \dots, k_p} \mathbb{E}_0 \left(\prod_{s=1}^p X_{i_s} X_{j_s} X_{k_s} \right). \quad (5.3)$$

Observe for each term in the sum above, the expectation is non-zero, if only if the multiplicity of each element in the multi-set $\{i_1, \dots, i_p\} \cup \{j_1, \dots, j_p\} \cup \{k_1, \dots, k_p\}$ is exactly 2. This implies, the number of distinct elements in $\{i_1, \dots, i_p\} \cup \{j_1, \dots, j_p\} \cup \{k_1, \dots, k_p\}$ is $3q$ and every pair of sets among $\{i_1, \dots, i_p\}$, $\{j_1, \dots, j_p\}$, $\{k_1, \dots, k_p\}$ must have exactly q elements in common. Therefore, from (5.3) and

recalling the definition of the matrix $\mathbf{D}_{J_N}$ from (2.8) we get,

$$\begin{aligned} |\mathbb{E}_0 H_N^3(\mathbf{X})| &\lesssim_p \sum_{\substack{1 \leq i_1, \dots, i_q, i_{q+1}, \dots, i_{2q}, i_{2q+1}, \dots, i_{3q} \leq N \\ \text{distinct}}} |J_{i_1, \dots, i_q, i_{q+1}, \dots, i_{2q}} J_{i_{q+1}, \dots, i_{2q}, i_{2q+1}, \dots, i_{3q}} J_{i_{2q+1}, \dots, i_{3q}, i_1, \dots, i_q}| \\ &\lesssim_p \sum_{\substack{1 \leq i_1, i_{q+1}, i_{2q+1} \leq N \\ \text{distinct}}} d_{J_N}(i_1, i_{q+1}) d_{J_N}(i_{q+1}, i_{2q+1}) d_{J_N}(i_1, i_{2q+1}) \\ &= \text{Trace}(\mathbf{D}_{J_N}^3) \leq N \|\mathbf{D}_{J_N}\| = O(N), \end{aligned}$$

where the last step uses the assumption that $\sup_{N \geq 1} \|\mathbf{D}_{J_N}\| < \infty$. $\square$

6. Proof of Theorem 2.10

We start by proving Theorem 2.10 (1). To this end, it suffices to verify Theorem 2.3. Note that condition (1) is easily satisfied because the bounded maximum degree condition (2.9) holds almost surely:

$$\max_{1 \leq i \leq N} d_{J_N}(i) = \frac{1}{N^{p-1}} \max_{1 \leq i \leq N} d_{A_{H_N}}(i) = O(1),$$

since $d_{A_{H_N}}(i) \leq N^{p-1}$, for all $1 \leq i \leq N$, in any p -uniform hypergraph H_N . Next, we verify condition (2) of Theorem 2.3. To this end, by the lower bound in [(16), Theorem 1.6] (which is the mean-field lower bound to the Gibbs variational representation of the log-partition function), we have

$$\begin{aligned} \mathbb{E} F_N(\beta) &\geq \mathbb{E} \left[\sup_{\mathbf{x} \in [-1, 1]^N} \left\{ \beta \sum_{1 \leq i_1, \dots, i_p \leq N} \mathbb{E}(J_{i_1 \dots i_p}) x_{i_1} \dots x_{i_p} - \sum_{i=1}^N I(x_i) \right\} \right] \\ &\geq \sup_{\mathbf{x} \in [-1, 1]^N} \left\{ \frac{\beta}{N^{p-1}} \sum_{1 \leq i_1, \dots, i_p \leq N} \mathbb{E}(a_{i_1 \dots i_p}) x_{i_1} \dots x_{i_p} - \sum_{i=1}^N I(x_i) \right\}. \end{aligned} \quad (6.1)$$

Now, take any $\mathbf{t} := (t_1, \dots, t_K) \in [-1, 1]^K$, and define $\mathbf{x} \in [-1, 1]^N$ by taking $x_i := t_j$, if $i \in \mathcal{B}_j$, where $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_K$ are as in Definition 2.9. Then, the term inside the supremum in the RHS of (6.1) equals $N\phi_\beta(t_1, \dots, t_K) + O(1)$ (recall the definition of the function $\phi_\beta(t_1, \dots, t_K)$ from (2.13)). Hence, (6.1) gives us,

$$\frac{\mathbb{E} F_N(\beta)}{N} \geq \sup_{(t_1, t_2, \dots, t_K) \in [0, 1]^K} \phi_\beta(t_1, \dots, t_K) + o(1). \quad (6.2)$$

The bound in (6.2) above combined with the definition of the threshold β_{HSBM}^* in (2.12) and now implies that for all $\beta > \beta_{\text{HSBM}}^*$, $\liminf_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} F_N(\beta) > 0$. Then by Lemma 6.1 below, it follows that $\liminf_{N \rightarrow \infty} \frac{1}{N} F_N(\beta) > 0$ with probability 1. This verifies condition (2) of Theorem 2.3 and shows that the MPL estimate $\hat{\beta}_N(\mathbf{X})$ is $\sqrt{N}$ -consistent for $\beta > \beta_{\text{HSBM}}^*$.

LEMMA 6.1. Let $F_N(\beta)$ denote the log-partition function of the p -tensor stochastic block model as in Theorem 2.10. Then, for every $\beta > 0$, the sequence $F_N(\beta) - \mathbb{E}F_N(\beta)$ is bounded in probability.

Proof. To start with, note that $F_N(\beta)$ is a function of the collection of i.i.d. random variables $\mathcal{A} := \{A_{i_1 \dots i_p}\}_{1 \leq i_1 < \dots < i_p \leq N}$, and so, it is convenient to denote $F_N(\beta)$ by $F_{N,\beta}(\mathcal{A})$. Let us take $\mathcal{A}' := \{A'_{i_1 \dots i_p}\}_{1 \leq i_1 < \dots < i_p \leq N}$, where $A'_{123 \dots p} = 1 - A_{123 \dots p}$ and $A'_{i_1 \dots i_p} = A_{i_1 \dots i_p}$ for all $(i_1, \dots, i_p) \neq (1, 2, 3, \dots, p)$. Note that

$$\left| \sum_{i_1 < \dots < i_p} A_{i_1 \dots i_p} X_{i_1} \cdots X_{i_p} - \sum_{i_1 < \dots < i_p} A'_{i_1 \dots i_p} X_{i_1} \cdots X_{i_p} \right| = |X_1 X_2 \cdots X_p| = 1.$$

Hence,

$$\exp \left\{ \frac{\beta p!}{N^{p-1}} \sum_{i_1 < \dots < i_p} A_{i_1 \dots i_p} X_{i_1} \cdots X_{i_p} \right\} \leq \exp \left\{ \frac{\beta p!}{N^{p-1}} \right\} \exp \left\{ \frac{\beta p!}{N^{p-1}} \sum_{i_1 < \dots < i_p} A'_{i_1 \dots i_p} X_{i_1} \cdots X_{i_p} \right\}.$$

The above inequality implies that $F_{N,\beta}(\mathcal{A}) \leq F_{N,\beta}(\mathcal{A}') + \beta p! N^{1-p}$. Similarly, we also have $F_{N,\beta}(\mathcal{A}') \leq F_{N,\beta}(\mathcal{A}) + \beta p! N^{1-p}$, and hence,

$$|F_{N,\beta}(\mathcal{A}) - F_{N,\beta}(\mathcal{A}')| \leq \beta p! N^{1-p}.$$

Of course, the above arguments hold if $\mathcal{A}'$ is obtained by flipping any arbitrary entry of $\mathcal{A}$ (not necessarily the $(1, 2, \dots, p)$ -th entry) and keeping all other entries unchanged. Hence, the assumption of McDiarmid's inequality [44] holds with bounding constants $c_{i_1 \dots i_p} = \beta p! N^{1-p}$. Therefore, for every $t > 0$:

$$\mathbb{P}(|F_N(\beta) - \mathbb{E}F_N(\beta)| \geq t) \leq 2 \exp \left\{ -\frac{2t^2}{\sum_{i_1 < \dots < i_p} c_{i_1 \dots i_p}^2} \right\} \leq 2 \exp \left\{ -\frac{2t^2 N^{p-2}}{\beta^2 (p!)^2} \right\},$$

which completes the proof of the lemma. $\square$

We will now use Lemma 6.1 to prove Theorem 2.10 (2). To this end, we will show that

$$\mathbb{E}F_N(\beta) = O(1), \quad \text{for } \beta < \beta_{\text{HSBM}}^*, \quad (6.3)$$

the expectation in (6.3) being taken with respect to the randomness of the HSBM. To see why this implies Theorem 2.10 (2), assume, on the contrary, that there is a sequence of estimates that is consistent for $\beta < \beta_{\text{HSBM}}^*$. Using this sequence of estimates, we can then construct a consistent sequence of tests

$\{\phi_N\}_{N \geq 1}$ for the following hypothesis testing problem:⁴

$$H_0 : \beta = \beta_1 \quad \text{versus} \quad H_1 : \beta = \beta_2, \quad (6.4)$$

if $\beta_1 < \beta_2 < \beta_{\text{HSBM}}^*$. To this end, denote by $\mathbb{Q}_{\beta,p}$ the joint distribution of the HSBM and the p -tensor Ising model with parameter β . Then a simple calculation shows that for any two positive real numbers $\beta_1 < \beta_2$, the KL divergence between the joint measures $\mathbb{Q}_{\beta_1,p}$ and $\mathbb{Q}_{\beta_2,p}$ is given by:

$$D_N(\mathbb{Q}_{\beta_1,p} \parallel \mathbb{Q}_{\beta_2,p}) = \mathbb{E} D_N(\mathbb{P}_{\beta_1,p} \parallel \mathbb{P}_{\beta_2,p}) = \mathbb{E} F_N(\beta_2) - \mathbb{E} F_N(\beta_1) - (\beta_2 - \beta_1) \mathbb{E} F'_N(\beta_1), \quad (6.5)$$

where, as before, the expectation in (6.5) is taken with respect to the randomness of the HSBM. Now, by the monotonicity of $F'_N(\cdot)$,

$$0 = (\beta_2 - \beta_1) F'_N(0) \leq (\beta_2 - \beta_1) F'_N(\beta_1) \leq \int_{\beta_1}^{\beta_2} F'_N(t) dt = F_N(\beta_2) - F_N(\beta_1).$$

Hence, by (6.3) and (6.5), $D_N(\mathbb{Q}_{\beta_1,p} \parallel \mathbb{Q}_{\beta_2,p}) = O(1)$. Then, by [(10), Proposition 6.1], there cannot exist any sequence of consistent tests for the hypothesis (6.4), which leads to a contradiction. This completes the proof of Theorem 2.10 (2).

6.1 Proof of (6.3)

The proof of (6.3) has the following two steps:

- (I) Define a new p -tensor Ising model on N nodes, with interaction tensor $\tilde{J}_N := \mathbb{E} J_N$. We will call this model $\mathcal{M}_0$. The first step in the proof of (6.3) is to show that the log-partition function $\tilde{F}_N$ of the model $\mathcal{M}_0$ is bounded, for every $\beta < \beta_{\text{HSBM}}^*$.
- (II) The second step is to show that the expected log-partition function $\mathbb{E} F_N(\beta)$ of the original model is bounded, for $\beta < \beta_{\text{HSBM}}^*$, by comparing it with the log-partition function $\tilde{F}_N$ of the model $\mathcal{M}_0$. The result in (6.3) then follows by an application of Lemma 6.1.

6.1.1 Proof of Step (I) Throughout this section we fix $\beta < \beta_{\text{HSBM}}^*$ and denote by $\mathbb{P}_{\beta, \mathcal{M}_0}$ the probability measure corresponding to the model $\mathcal{M}_0$ at the parameter β and $\mathbb{E}_{\beta, \mathcal{M}_0}$ the expectation with respect to the probability measure $\mathbb{P}_{\beta, \mathcal{M}_0}$.

LEMMA 6.2. Denote by $\tilde{F}_N(\beta)$ the log-partition function of the model $\mathcal{M}_0$. Then for $\beta < \beta_{\text{HSBM}}^*$, $\limsup_{N \rightarrow \infty} \tilde{F}_N(\beta) < \infty$.

Proof. *Proof of Lemma 6.2:* Denote the Hamiltonian of the model $\mathcal{M}_0$ by $\tilde{H}_N(\mathbf{X})$, that is,

$$\tilde{H}_N(\mathbf{X}) := \frac{1}{N^{p-1}} \sum_{1 \leq i_1, i_2, \dots, i_p \leq N} \mathbb{E}(a_{i_1 i_2 \dots i_p}) X_{i_1} X_{i_2} \dots X_{i_p}.$$

⁴ A sequence of tests $\{\phi_N\}_{N \geq 1}$ is said to be *consistent* if both its Type I and Type II errors converge to zero as $N \rightarrow \infty$, that is, $\lim_{N \rightarrow \infty} \mathbb{E}_{H_0} \phi_N = 0$, and the power $\lim_{N \rightarrow \infty} \mathbb{E}_{H_1} \phi_N = 1$.

For each $X \in \mathcal{C}_N$ and $1 \leq j \leq K$, define $S_j(X) := \sum_{i \in \mathcal{B}_j} X_i$. With these notations, we have $\tilde{H}_N(X) = \bar{H}_N(X) + O(1)$, where

$$\bar{H}_N(X) = \frac{1}{N^{p-1}} \sum_{1 \leq j_1, \dots, j_p \leq K} \theta_{j_1 \dots j_p} \prod_{\ell=1}^p S_{j_\ell}(X). \quad (6.6)$$

Let us define $\bar{Z}_N(\beta) := \frac{1}{2^N} \sum_{X \in \mathcal{C}_N} e^{\beta \bar{H}_N(X)}$. Since $\tilde{F}_N(\beta) = \log \bar{Z}_N(\beta) + O(1)$, it suffices to show that $\bar{Z}_N(\beta) = O(1)$. Towards this, for each $1 \leq j \leq K$, define the sets:

$$I_j := \left\{ -1, -1 + \frac{2}{|\mathcal{B}_j|}, -1 + \frac{4}{|\mathcal{B}_j|}, \dots, 1 - \frac{2}{|\mathcal{B}_j|}, 1 \right\} \quad \text{and} \quad A_s(j) := \{X \in \mathcal{C}_N : S_j(X) = s\}.$$

Recall, $\mathcal{B}_j = (N \sum_{i=1}^{j-1} \lambda_i, N \sum_{i=1}^j \lambda_i] \cap [N]$, hence $||\mathcal{B}_j| - N\lambda_j| \leq 2$. Now, note that

$$\begin{aligned} \bar{Z}_N(\beta) &= \frac{1}{2^N} \sum_{(\ell_1, \dots, \ell_K) \in I_1 \times \dots \times I_K} e^{\frac{\beta}{N^{p-1}} \sum_{1 \leq j_1, \dots, j_p \leq K} \theta_{j_1 \dots j_p} \prod_{m=1}^p \ell_{j_m} |\mathcal{B}_{j_m}|} \left| \bigcap_{j=1}^K A_{\ell_j |\mathcal{B}_j|}(j) \right| \\ &\lesssim \frac{1}{2^N} \sum_{(\ell_1, \dots, \ell_K) \in I_1 \times \dots \times I_K} e^{N\beta \sum_{1 \leq j_1, \dots, j_p \leq K} \theta_{j_1 \dots j_p} \prod_{m=1}^p \lambda_{j_m} \ell_{j_m}} \left| \bigcap_{j=1}^K A_{\ell_j |\mathcal{B}_j|}(j) \right| \end{aligned} \quad (6.7)$$

$$= T_1 + T_2, \quad (6.8)$$

where the term T_1 is obtained by restricting the sum in the RHS of (6.7) to the set $(I_1 \times \dots \times I_K) \cap [-\frac{1}{2}, \frac{1}{2}]^K$ and the term T_2 is the sum restricted to the set $(I_1 \times \dots \times I_K) \cap ([-\frac{1}{2}, \frac{1}{2}]^K)^c$. $\square$

Let us bound T_1 first. Note that

$$|A_{\ell_j |\mathcal{B}_j|}(j)| = \binom{|\mathcal{B}_j|}{\frac{|\mathcal{B}_j|(1+\ell_j)}{2}}.$$

Then by the Stirling's approximation of the binomial coefficient (see, for example, [(48), Lemma B.5]) and using the fact that the sets $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_K$ are disjoint, we have for all $(\ell_1, \dots, \ell_K) \in [-\frac{1}{2}, \frac{1}{2}]^K$,

$$\left| \bigcap_{j=1}^K A_{\ell_j |\mathcal{B}_j|}(j) \right| = 2^N \exp \left\{ - \sum_{j=1}^K |\mathcal{B}_j| I(\ell_j) \right\} O \left(\frac{1}{\sqrt{\prod_{j=1}^K |\mathcal{B}_j|}} \right).$$

Hence, denoting $\ell := (\ell_1, \dots, \ell_K)$ and $\mathcal{J} := (I_1 \times \dots \times I_K)$ gives,

$$\begin{aligned}
 T_1 &= \sum_{\ell \in \mathcal{J} \cap \left[-\frac{1}{2}, \frac{1}{2}\right]^K} e^{N\beta \sum_{1 \leq j_1, \dots, j_p \leq K} \theta_{j_1 \dots j_p} \prod_{m=1}^p \lambda_{jm} \ell_{jm} - \sum_{j=1}^K |\mathcal{B}_j| I(\ell_j)} O\left(\frac{1}{\sqrt{\prod_{j=1}^K |\mathcal{B}_j|}}\right) \\
 &\lesssim \sum_{\ell \in \mathcal{J} \cap \left[-\frac{1}{2}, \frac{1}{2}\right]^K} e^{N\beta \sum_{1 \leq j_1, \dots, j_p \leq K} \theta_{j_1 \dots j_p} \prod_{m=1}^p \lambda_{jm} \ell_{jm} - N \sum_{j=1}^K \lambda_j I(\ell_j)} O\left(\frac{1}{\sqrt{\prod_{j=1}^K |\mathcal{B}_j|}}\right) \\
 &= \sum_{\ell \in \mathcal{J} \cap \left[-\frac{1}{2}, \frac{1}{2}\right]^K} e^{N\phi_\beta(\ell_1, \dots, \ell_K)} O\left(\frac{1}{\sqrt{\prod_{j=1}^K |\mathcal{B}_j|}}\right)
 \end{aligned}$$

Now, since $\beta < \beta_{\text{HSBM}}^*$ we can choose $\varepsilon \in (0, 1)$, such that $\beta/(1 - \varepsilon) < \beta_{\text{HSBM}}^*$. Hence, recalling (2.12),

$$\phi_{\beta/(1-\varepsilon)}(\ell_1, \dots, \ell_K) \leq 0, \quad \text{that is,} \quad \phi_\beta(\ell_1, \dots, \ell_K) \leq -\varepsilon \sum_{j=1}^K \lambda_j I(\ell_j) \leq -\frac{\varepsilon}{2} \sum_{j=1}^K \lambda_j \ell_j^2,$$

where the last step uses $I(x) \geq x^2/2$. Then by a Riemann sum approximation (see [(48), Lemma B.2]),

$$T_1 \lesssim \sum_{\ell \in \mathcal{J} \cap \left[-\frac{1}{2}, \frac{1}{2}\right]^K} e^{-\frac{N\varepsilon}{2} \sum_{j=1}^K \lambda_j \ell_j^2} O\left(\frac{1}{\sqrt{\prod_{j=1}^K |\mathcal{B}_j|}}\right) \quad (6.9)$$

$$\begin{aligned}
 &\leq \frac{1}{2^K} O\left(\sqrt{\prod_{j=1}^K |\mathcal{B}_j|}\right) \int_{\mathbb{R}^K} \exp\left\{-\frac{N\varepsilon}{2} \sum_{j=1}^K \lambda_j x_j^2\right\} dx_1 \dots dx_K + O(1) \\
 &= \frac{1}{2^K} O\left(\sqrt{\prod_{j=1}^K \frac{|\mathcal{B}_j|}{N}}\right) \int_{\mathbb{R}^K} \exp\left\{-\frac{\varepsilon}{2} \sum_{j=1}^K \lambda_j y_j^2\right\} dy_1 \dots dy_K + O(1) \\
 &\lesssim \int_{\mathbb{R}^K} \exp\left\{-\frac{\varepsilon}{2} \sum_{j=1}^K \lambda_j y_j^2\right\} dy_1 \dots dy_K + O(1) = O(1). \quad (6.10)
 \end{aligned}$$

Now, we bound T_2 . For this we need the following combinatorial estimate:

Observation 6.3(61), Equation (5.4). For every integer s and positive integer m ,

$$\left| \left\{ \mathbf{x} \in \{-1, 1\}^m : \sum_{i=1}^m x_i = s \right\} \right| \leq 2^m e^{-mI(\frac{s}{m})}. \quad (6.11)$$

Using the bound in (6.11) and recalling that the sets $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_K$ are disjoint, we have for every $(\ell_1, \ell_2, \dots, \ell_K) \in (I_1, I_2, \dots, I_K)$,

$$\left| \bigcap_{j=1}^K A_{\ell_j|\mathcal{B}_j}(j) \right| \leq 2^N \exp \left\{ - \sum_{j=1}^K |\mathcal{B}_j| I(\ell_j) \right\}.$$

Hence, by following the arguments used to obtain (6.9) we get,

$$\begin{aligned} T_2 &\lesssim \sum_{\ell \in \mathcal{I} \cap \left(\left[-\frac{1}{2}, \frac{1}{2} \right]^K \right)^c} e^{-\frac{N\epsilon}{2} \sum_{j=1}^K \lambda_j \ell_j^2} \leq \left(\prod_{j=1}^K (|\mathcal{B}_j| + 1) \right) \max_{\ell \in \mathcal{I} \cap \left(\left[-\frac{1}{2}, \frac{1}{2} \right]^K \right)^c} e^{-\frac{N\epsilon}{2} \sum_{j=1}^K \lambda_j \ell_j^2} \\ &\leq \left(\prod_{j=1}^K (N\lambda_j + 1) \right) \max_{1 \leq j \leq K} e^{-\frac{N\epsilon \lambda_j}{8}} = o(1). \end{aligned} \quad (6.12)$$

This shows, combining (6.8), (6.10) and (6.12), that $\bar{Z}_N(\beta) = O(1)$ for all $\beta < \beta_{\text{HSBM}}^*$, completing the proof of Lemma 6.2.

6.1.2 Proof of Step (II) We now show that $\mathbb{E}F_N(\beta)$ is bounded, for $\beta < \beta_{\text{HSBM}}^*$, which will allow us to conclude, using Lemma 6.1, that $F_N(\beta) = O(1)$ with probability 1, for all $\beta < \beta_{\text{HSBM}}^*$, that is, (6.3) holds.

LEMMA 6.4. For every $\beta < \beta_{\text{HSBM}}^*$, $\limsup_{N \rightarrow \infty} \mathbb{E}F_N(\beta) < \infty$.

Proof. Fix $\beta < \beta_{\text{HSBM}}^*$. Then the partition function $Z_N(\beta)$ becomes:

$$\begin{aligned} Z_N(\beta) &= \frac{1}{2^N} \sum_{\mathbf{x} \in \mathcal{C}_N} e^{\frac{p\beta}{N^{p-1}} \sum_{i_1 < \dots < i_p} A_{i_1 \dots i_p} X_{i_1} \dots X_{i_p}} \\ &= \frac{1}{2^N} \sum_{\mathbf{x} \in \mathcal{C}_N} e^{\frac{p\beta}{N^{p-1}} \sum_{i_1 < \dots < i_p} (A_{i_1 \dots i_p} - \mathbb{E}A_{i_1 \dots i_p}) X_{i_1} \dots X_{i_p}} e^{\frac{p\beta}{N^{p-1}} \sum_{i_1 < \dots < i_p} \mathbb{E}A_{i_1 \dots i_p} X_{i_1} \dots X_{i_p}}. \end{aligned}$$

By Hoeffding's inequality, for each $X \in \mathcal{C}_N$ and $i_1 < \dots < i_p$,

$$\begin{aligned} \mathbb{E} \exp \left\{ \frac{p! \beta}{N^{p-1}} \sum_{i_1 < \dots < i_p} (A_{i_1 \dots i_p} - \mathbb{E} A_{i_1 \dots i_p}) X_{i_1} \dots X_{i_p} \right\} &\leq \prod_{i_1 < \dots < i_p} \exp \left\{ \frac{\beta^2 (p!)^2}{8N^{2p-2}} \right\} \\ &\leq \exp \left\{ \frac{\beta^2 p!}{8N^{p-2}} \right\}. \end{aligned}$$

This shows that

$$\mathbb{E} Z_N(\beta) \leq \exp \left\{ \frac{\beta^2 p!}{8N^{p-2}} \right\} \left[\frac{1}{2^N} \sum_{X \in \mathcal{C}_N} e^{\frac{p! \beta}{N^{p-1}} \sum_{i_1 < \dots < i_p} \mathbb{E} A_{i_1 \dots i_p} X_{i_1} \dots X_{i_p}} \right] = \exp \left\{ \frac{\beta^2 p!}{8N^{p-2}} \right\} \tilde{Z}_N(\beta),$$

where $\tilde{Z}_N$ is the partition function of the model $\mathcal{M}_0$. Now, taking logarithms and using Lemma 6.2 shows, $\limsup_{N \rightarrow \infty} \log \mathbb{E} Z_N(\beta) < \infty$, for $\beta < \beta_{\text{HSBM}}^*$. Then, by Jensen's inequality, we conclude that $\limsup_{N \rightarrow \infty} \mathbb{E} F_N(\beta) < \infty$, for $\beta < \beta_{\text{HSBM}}^*$, completing the proof of the lemma. $\square$

7. Proofs from Section 2.4

7.1 Proof of Theorem 2.15

Define the function $\phi_p : [-1, 1] \mapsto (-\infty, \infty]$ as:

$$\phi(t) = \phi_p(t) := \begin{cases} p^{-1} t^{1-p} \tanh^{-1}(t) & \text{if } p \text{ is even and } t \neq 0, \\ p^{-1} t^{1-p} \tanh^{-1}(t) & \text{if } p \text{ is odd and } t > 0. \\ \infty & \text{if } p \text{ is odd and } t < 0. \\ 0 & \text{if } t = 0. \end{cases} \quad (7.1)$$

Note that for every $t \neq 0$ when p is even, and every $t > 0$ when p is odd, the function ϕ is twice differentiable on some compact set containing t in its interior, and

$$\phi'(t) = -\frac{1}{p t^{p-1}} \left\{ \frac{(p-1) \tanh^{-1}(t)}{t} - \frac{1}{1-t^2} \right\}.$$

Hence, recalling the definition of the function g from the statement of the theorem we have,

$$\phi'(m_*) = -\frac{g''(m_*)}{p m_*^{p-1}}. \quad (7.2)$$

Moreover, from the definition of the MPL estimate in (2.3) it is easy to see that in the Curie–Weiss model $\hat{\beta}_N(\mathbf{X}) = \phi(\bar{X}_N)$. Note that $\phi(m_*) = \beta$, since $g'(m_*) = 0$. This implies,

$$\sqrt{N}(\hat{\beta}_N(\mathbf{X}) - \beta) = \sqrt{N}(\phi(\bar{X}_N) - \phi(m_*)). \quad (7.3)$$

We now consider the case p is even and p is odd separately. Throughout, we assume $\beta > \beta_{\text{CW}}^*$:

- $p \geq 3$ is odd: In this case, m_* is the unique global maximizer of the function g on $[-1, 1]$. This implies, by [(48), Theorem 2.1 (1)],

$$\sqrt{N}(\bar{X}_N - m_*) \xrightarrow{D} N\left(0, -\frac{1}{g''(m_*)}\right).$$

Hence, by (7.3) and the delta method [(42), Theorem 1.8.12],

$$\sqrt{N}(\hat{\beta}_N(\mathbf{X}) - \beta) \xrightarrow{D} N\left(0, -\frac{(\phi'(m_*))^2}{g''(m_*)}\right) \stackrel{D}{=} N\left(0, -\frac{g''(m_*)}{p^2 m_*^{2p-2}}\right),$$

where the last step uses (7.2). This completes the proof when p is odd.

- $p \geq 2$ is even: In this case, the function g has two (non-zero) global maximizers on $[-1, 1]$, which are given by m_* and $-m_*$ (as shown in [(48), Section C]). This implies, by [(48), Theorem 2.1 (2)] and the delta method,

$$\sqrt{N}(\hat{\beta}_N(\mathbf{X}) - \beta) \Big|_{\{\bar{X}_N \in (0, 1]\}} \xrightarrow{D} N\left(0, -\frac{g''(m_*)}{p^2 m_*^{2p-2}}\right). \quad (7.4)$$

Similarly, observing that $\phi(m_*) = -\phi(-m_*)$ and $g''(m_*) = g''(-m_*)$,

$$\sqrt{N}(\hat{\beta}_N(\mathbf{X}) - \beta) \Big|_{\{\bar{X}_N \in [-1, 0]\}} \xrightarrow{D} N\left(0, -\frac{g''(m_*)}{p^2 m_*^{2p-2}}\right). \quad (7.5)$$

Combining (7.4) and (7.5), gives the desired result when p is even.

7.2 Proof of Theorem 2.17

Fix $\beta = \beta_{\text{CW}}^*(p)$. We now consider the three cases in Theorem 2.17 separately.

Proof of (1): In this case, $p = 2$ and, hence, by (7.1), $\hat{\beta}_N = \frac{\tanh^{-1}(\bar{X}_N)}{2\bar{X}_N} \mathbf{1}_{\{\bar{X}_N \neq 0\}}$. Therefore, on the event $\mathcal{E}_N := \{\bar{X}_N \neq 0, |\bar{X}_N| < \frac{1}{2}\}$,

$$\begin{aligned} N^{\frac{1}{2}} \left(\hat{\beta}_N - \frac{1}{2} \right) &= N^{\frac{1}{2}} \left(\frac{\tanh^{-1}(\bar{X}_N)}{2\bar{X}_N} - \frac{1}{2} \right) \\ &= \frac{1}{2} N^{\frac{1}{2}} \sum_{s=1}^{\infty} \frac{\bar{X}_N^{2s}}{2s+1} = \frac{1}{6} N^{\frac{1}{2}} \bar{X}_N^2 + N^{\frac{1}{2}} O(\bar{X}_N^4). \end{aligned} \quad (7.6)$$

From [(21), Proposition 4.1] we know that $N^{\frac{1}{4}} \bar{X}_N \xrightarrow{D} F$, where F is as defined in the statement of Theorem 2.17. The result in (2.21) now follows from (7.6), and the observation that $\mathbb{P}_{\beta}(\mathcal{E}_N) = o(1)$, since $\bar{X}_N \xrightarrow{P} 0$ and $\mathbb{P}_{\beta}(\bar{X}_N = 0) = o(1)$ (from the proof of [(48), Lemma C.6]).

Proof of (2): Assume $p \geq 3$. To begin with, define the three intervals $A_1 := [-1, -\frac{m_*}{2}]$ and $A_2 := (-\frac{m_*}{2}, \frac{m_*}{2})$ and $A_3 := [\frac{m_*}{2}, 1]$ (recall that $m_* = m_*(p, \beta)$ is the unique positive maximizer of the function $g(t) := \beta t^p - I(t)$). We now consider the following two cases depending on whether $p \geq 3$ is odd or even:

- $p \geq 4$ is even: In this case, the function g has three global maximizers on $[-1, 1]$, which are given by $m_* > 0$ and $-m_*$, and 0. Then, by [(48), Theorem 2.1 (2)] and arguments as in (7.4), we have

$$\sqrt{N}(\hat{\beta}_N(\mathbf{X}) - \beta) | \{\bar{X}_N \in A_i\} \xrightarrow{D} N \left(0, -\frac{g''(m_*)}{p^2 m_*^{2p-2}} \right), \quad (7.7)$$

for $i \in \{1, 3\}$. Next, recalling $\beta_N(\mathbf{X}) = \phi(\bar{X}_N)$, where $\phi(\cdot)$ as in (7.1), note that

$$\hat{\beta}_N(\mathbf{X}) | \{\bar{X}_N \in A_2 \setminus \{0\}\} \xrightarrow{D} \infty, \quad (7.8)$$

since $\frac{\tanh^{-1}(t)}{t^{p-1}} \rightarrow \infty$, as $|t| \rightarrow 0$. Now, since $\mathbb{P}(\bar{X}_N = 0) = o(1)$ and $\mathbb{P}_{\beta}(\bar{X}_N \in A_2) \rightarrow \alpha$ (by [(48), Theorem 2.1 (2)]), combining (7.7) and (7.8) the result in (2.22) follows, if $p \geq 4$ is even.

- $p \geq 3$ is odd: In this case, the function g has two global maximizers on $[-1, 1]$ one of which is non-positive and the other is $m_* > 0$. Then, by similar arguments as above,

$$\sqrt{N}(\hat{\beta}_N(\mathbf{X}) - \beta) | \{\bar{X}_N \in A_3\} \xrightarrow{D} N \left(0, -\frac{g''(m_*)}{p^2 m_*^{2p-2}} \right). \quad (7.9)$$

Moreover, recalling (7.1),

$$\hat{\beta}_N(\mathbf{X}) | \{\bar{X}_N \in (A_1 \cup A_2) \setminus \{0\}\} \xrightarrow{D} \infty. \quad (7.10)$$

The result in (2.22) now follows from (7.9), (7.10), and the fact that $\mathbb{P}_\beta(\bar{X}_N \in A_1 \cup A_2) \rightarrow \alpha$, if p is odd (by [(48), Theorem 2.1 (2)]).

Proof of (3): Here, we prove the finer asymptotics of $\hat{\beta}_N(\mathbf{X})$. For this, note by [(48), Theorem 2.1 (2)] that $\sqrt{N}\bar{X}_N|\{\bar{X}_N \in \mathcal{B}\} \xrightarrow{D} N(0, 1)$ for any interval $\mathcal{B}$ containing 0, but no other maximizer of g . We now consider the following two cases:

- $p \geq 4$ is even: In this case, $\sqrt{N}\bar{X}_N|\{\bar{X}_N \in A_2\} \xrightarrow{D} N(0, 1)$. Then, (7.1) gives,

$$N^{1-\frac{p}{2}}\hat{\beta}_N(\mathbf{X}) = \frac{1}{p(\sqrt{N}\bar{X}_N)^{p-2}} \cdot \frac{\tanh^{-1}(\bar{X}_N)}{\bar{X}_N} \mathbf{1}_{\{\bar{X}_N \neq 0\}}.$$

Now, since $\frac{\tanh^{-1}(\bar{X}_N)}{\bar{X}_N} \mathbf{1}_{\{\bar{X}_N \neq 0\}}|\{\bar{X}_N \in A_2 \setminus \{0\}\} \xrightarrow{P} 1$, we have,

$$N^{1-\frac{p}{2}}\hat{\beta}_N(\mathbf{X})|\{\bar{X}_N \in A_2 \setminus \{0\}\} \xrightarrow{D} \frac{1}{pZ^{p-2}}, \quad (7.11)$$

where $Z \sim N(0, 1)$. The result in now (2.23) follows from (7.11), by noting $\mathbb{P}(\bar{X}_N = 0) = o(1)$, that $\mathbb{P}_\beta(\bar{X}_N \in A_2) \rightarrow \alpha$ and $N^{1-\frac{p}{2}}\hat{\beta}_N(\mathbf{X})|\{\bar{X}_N \in A_1 \cup A_3\} \xrightarrow{P} 0$ (by (7.7)).

- $p \geq 3$ is odd: In this case, since the function g has two global maximizers on $[-1, 1]$ one of which is non-positive and the other is $m_* > 0$, $\sqrt{N}\bar{X}_N|\{\bar{X}_N \in A_1 \cup A_2\} \xrightarrow{D} N(0, 1)$. Then, (7.1) gives,

$$N^{1-\frac{p}{2}}\hat{\beta}_N(\mathbf{X}) = \frac{1}{p(\sqrt{N}\bar{X}_N)^{p-2}} \cdot \frac{\tanh^{-1}(\bar{X}_N)}{\bar{X}_N} \mathbf{1}_{\{\bar{X}_N > 0\}} + \infty \mathbf{1}_{\{\bar{X}_N < 0\}}, \quad (7.12)$$

where we adopt the convention infinity times zero is zero. It follows from (7.12) that $N^{1-p/2}\hat{\beta}_N(\mathbf{X})$ is a non-negative random variable. Hence, denoting $\mathcal{A}_{12} := \{\bar{X}_N \in A_1 \cup A_2\}$ and taking $t \in [0, \infty)$ gives

$$\mathbb{P}\left(N^{1-\frac{p}{2}}\hat{\beta}_N(\mathbf{X}) \leq t | \mathcal{A}_{12}\right) = \mathbb{P}\left(N^{1-\frac{p}{2}}\hat{\beta}_N(\mathbf{X}) \leq t, \bar{X}_N > 0 | \mathcal{A}_{12}\right) + o(1),$$

since $\mathbb{P}_\beta(\bar{X}_N = 0 | \mathcal{A}_{12}) = o(1)$. Now, note that

$$\begin{aligned} & \mathbb{P}\left(N^{1-\frac{p}{2}}\hat{\beta}_N(\mathbf{X}) \leq t, \bar{X}_N > 0 | \mathcal{A}_{12}\right) \\ &= \mathbb{P}\left(\frac{1}{p(\sqrt{N}\bar{X}_N)^{p-2}} \cdot \frac{\tanh^{-1}(\bar{X}_N)}{\bar{X}_N} \leq t, \bar{X}_N \neq 0 | \mathcal{A}_{12}\right) - \mathbb{P}_\beta(\bar{X}_N < 0 | \mathcal{A}_{12}), \end{aligned}$$

since on the event $\{\bar{X}_N < 0\}$, $\frac{1}{p(\sqrt{N}\bar{X}_N)^{p-2}} \cdot \frac{\tanh^{-1}(\bar{X}_N)}{\bar{X}_N} < 0 \leq t$. Note that $\mathbb{P}_\beta(\bar{X}_N < 0 | \mathcal{A}_{12}) \rightarrow \frac{1}{2}$, since $\sqrt{N}\bar{X}_N | \mathcal{A}_{12} \xrightarrow{D} N(0, 1)$. Then, by arguments as in (7.11), it follows that

$$\mathbb{P}\left(N^{1-\frac{p}{2}}\hat{\beta}_N(\mathbf{X}) \leq t | \mathcal{A}_{12}\right) \rightarrow \mathbb{P}\left(\frac{1}{pZ^{p-2}} \leq t\right) - \frac{1}{2} = \frac{1}{2}\mathbb{P}\left(\frac{1}{p|Z|^{p-2}} \leq t\right).$$

Therefore, by (7.12),

$$N^{1-\frac{p}{2}}\hat{\beta}_N(\mathbf{X}) | \mathcal{A}_{12} \xrightarrow{D} \frac{1}{2}\left(\frac{1}{p|Z|^{p-2}}\right) + \frac{1}{2}\delta_\infty. \quad (7.13)$$

The result in (2.24) follows now from (7.13), by noting that $\mathbb{P}_\beta(\bar{X}_N \in A_1 \cup A_2) \rightarrow \alpha$ and $N^{1-\frac{p}{2}}\hat{\beta}_N(\mathbf{X}) | \{\bar{X}_N \in A_3\} \xrightarrow{P} 0$ (by (7.9)).

Acknowledgement

The authors thank the anonymous referees for their valuable comments that greatly improved the quality and the presentation of the paper.

Funding

BBB was supported by National Science Foundation (NSF) CAREER grant DMS 2046393 and a Sloan Research Fellowship.

REFERENCES

1. AGARWAL, S., BRANSON, K. & BELONGIE, S. (2006) Higher order learning with graphs. *Proceedings of the 23rd international conference on Machine learning*, 17–24.
2. ANIMASHREE ANANDKUMAR, VINCENT Y. F. TAN, FURONG HUANG, and ALLAN S. WILLSKY, High-dimensional structure estimation in Ising models: Local separation criterion, *Annals of Statistics*, Vol. **40** (3), 1346–1375 2012.
3. ANGELINI, M. C., CALTAGIRONE, F., KRZAKALA, F. & ZDEBOROVÁ, L. (2015) Spectral detection on sparse hypergraphs. *53rd IEEE Annual Allerton Conference*, pp. 66–73.
4. BANERJEE, S., CARLIN, B. P. & GELFAND, A. E. (2014) Hierarchical modeling and analysis for spatial data. *Chapman and Hall/CRC*.
5. BAI, Z. D. (1999) Methodologies in spectral analysis of large-dimensional random matrices, a review (with discussion). *Statist. Sinica*, **9**, 611–677.
6. ADRIANO BARRA, Notes on ferromagnetic p-spin and REM, *Math. Methods Appl. Sci.*, Vol. **32** (7), 783–797, 2009.
7. BENTO, J. & MONTANARI, A. (2009) Which graphical models are difficult to learn? *Neural Information Processing Systems (NeurIPS)*, 1303–1311.
8. BESAG, J. (1974) Spatial interaction and the statistical analysis of lattice systems. *J. Roy. Stat. Soc. B*, **36**, 192–236.
9. JULIAN BESAG, Statistical analysis of non-lattice data, *The Statistician*, Vol. **24** (3), 179–195, 1975.
10. BHASWAR B. BHATTACHARYA and Sumit Mukherjee, Inference in ising models, *Bernoulli*, Vol. **24** (1), 493–525, 2018.
11. BOVIER, A., KURKOVA, I. & LÖWE, M. (2002) Fluctuations of the free energy in the REM and the p-Spin SK models. *Annals of Probability*, **30**, 605–651.

12. BRESLER, G. (2015) Efficiently learning Ising models on arbitrary graphs. *Proceedings Symposium on Theory of Computing (STOC)*, 771–782.
13. GUY BRESLER, ELCHANAN MOSSEL, and ALLAN SLY, Reconstruction of Markov random fields from samples: Some observations and algorithms, *SIAM J. Comput.*, Vol. **42**(2), 563–578, 2013.
14. BRESLER, G. & NAGARAJ, D. (2018) Optimal single sample tests for structured versus unstructured network data. *Conference On Learning Theory (COLT)*, 1657–1690.
15. CAO, Y., NEYKOV, M. & LIU, H. (2021) High temperature structure detection in ferromagnets. *Information and Inference: A Journal of the IMA*, to appear.
16. CHATTERJEE, S. & DEMBO, A. (2016) Nonlinear large deviations. *Adv. Math.*, **299**, 396–450.
17. CHATTERJEE, S. (2008) An error bound in the Sudakov-Fernique inequality. arXiv:0510424.
18. SOURAV CHATTERJEE, Estimation in spin glasses: A first step, *Annals of Statistics*, Vol. **35** (5), 1931–1946, 2007.
19. C. CHOW and C. LIU, Approximating discrete probability distributions with dependence trees, *IEEE Transactions on Information Theory*, Vol. **14** (3), 462–467, 1968.
20. FRANCIS COMETS, On consistency of a class of estimators for exponential families of Markov random fields on the lattice, *Annals of Statistics*, Vol. **20** (1), 455–468, 1992.
21. FRANCIS COMETS and BASILIS GIDAS, Asymptotics of maximum likelihood estimators for the Curie–Weiss model, *Annals of Statistics*, **19**(2): 557–578, 1991.
22. DAGAN, Y., DASKALAKIS, C., DIKKALA, N. & KANDIROU, A. V. (2021) Learning Ising models from one or multiple samples. *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 161–168.
23. CONSTANTINOS DASKALAKIS, NISHANT DIKKALA, and GAUTAM KAMATH, Testing Ising models, *IEEE Transactions on Information Theory*, Vol. **65** (11), 6829–6852, 2019.
24. DASKALAKIS, C., DIKKALA, N. & PANAGEAS, I. (2020) Logistic regression with peer-group effects via inference in higher-order Ising models. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 3653–3663.
25. DASKALAKIS, C., DIKKALA, N. & PANAGEAS, I. (2019) Regression from dependent observations. *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 881–889.
26. CONSTANTINOS DASKALAKIS, ELCHANAN MOSSEL, and SEBASTIEN ROCH, Evolutionary trees and the Ising model on the Bethe lattice: A proof of Steel’s conjecture, *Probab. Theory Related Fields*, Vol. **149** (1), 149–189, 2011.
27. ADLER, R. J. & TAYLOR, J. E. (2007) Gaussian Inequalities. *Random Fields and Geometry. Springer Monographs in Mathematics*. New York, NY: Springer.
28. GEMAN, S. & GRAFFIGNE, C. (1986) Markov random field image models and their applications to computer vision. *Proceedings of the International Congress of Mathematicians*, 1496–1517.
29. PROMIT GHOSAL and SUMIT MUKHERJEE, Joint estimation of parameters in Ising model, *Annals of Statistics*, Vol. **48** (2), 785–810, 2020.
30. GIDAS, B. (1988) Consistency of maximum likelihood and pseudolikelihood estimators for Gibbs distributions. *Stochastic Differential Systems, Stochastic Control Theory and Applications*. (W. FLEMING & P.-L. LIONS eds). New York: Springer, pp. 129–145.
31. DEBARGHYA GHOSHDASTIDAR and AMBEDKAR DUKKIPATI, Consistency of spectral hypergraph partitioning under planted partition model, *Annals of Statistics*, Vol. **45** (1), 289–315, 2017.
32. PETER, J. (2002) Green and Sylvia Richardson, Hidden markov models and disease mapping. *J. Amer. Statist. Assoc.*, **97**, 1055–1070.
33. GUYON, X. & KÜNSCH, H. R. (1992) Asymptotic comparison of estimators in the Ising Model. *Stochastic Models, Statistical Methods, and Algorithms in Image Analysis, Lecture Notes in Statistics*, **74**, 177–198.

34. HAMILTON, L., KOEHLER, F. & MOITRA, A. (2017) Information theoretic properties of Markov Random Fields, and their algorithmic applications. *Advances in Neural Information Processing Systems (NIPS)*, 2463–2472.
35. HERINGA, J. R., BLÖTE, H. W. J. & HOOGLAND, A. (1989) Phase transitions in self-dual Ising models with multispin interactions and a field. *Phys. Rev. Lett.*, **63**, 1546–1549.
36. JOHN, J. (1982) Hopfield, Neural networks and physical systems with emergent collective computational abilities. *Proc. Nat. Acad. Sci. U.S.A.*, **79**, 2554–2558.
37. ISING, E. (1925) Beitrag zur theorie der ferromagnetismus. *Z. Phys. C*, **31**, 253–258.
38. THOMAS JÖRG, FLORENT KRZAKALA, JORGE KURCHAN, A. C. MAGGS, and JOSÉ PUJOS, Energy gaps in quantum first-order mean-field-like transitions: The problems that quantum annealing cannot solve, *EPL (Europhysics Letters)*, Vol. **89** (4), 40004, 2010.
39. KLIVANS, A. R. & MEKA, R. (2017) Learning Graphical Models Using Multiplicative Weights. *Foundations of Computer Science*, 343–354.
40. LIU, J., SINCLAIR, A. & SRIVASTAVA, P. (2019) The Ising partition function: zeros and deterministic approximation. *J. Statist. Phys.*, **174**, 287–315.
41. QINGSHAN LIU, YUCHI HUANG, and DIMITRIS N METAXAS. Hypergraph with sampling for image retrieval, *Pattern Recognition*, Vol. **44**(10–11), 2255–2262, 2011.
42. LEHMANN, E. L. & CASELLA, G. (2006) *Theory of point estimation*. Springer.
43. LESIEUR, T., MIOLANE, L., LELARGE, M., KRZAKALA, F. & ZDEBOROVÁ, L. (2017) Statistical and computational phase transitions in spiked tensor estimation. *IEEE International Symposium on Information Theory (ISIT)*, 511–515.
44. MCDIARMID, C. (1989) *On the method of bounded differences, Surveys in Combinatorics*. Cambridge University Press, pp. 148–188.
45. MÉZARD, M., PARISI, G. & VIRASORO, M. A. (1987) *Spin glass theory and beyond, World Scientific Lecture Notes in Physics*, vol. **9**. Teane: World Scientific Publishing Co. Inc.
46. MONTANARI, A. (2019) Optimization of the Sherrington-Kirkpatrick Hamiltonian. *IEEE Symposium on the Foundations of Computer Science (FOCS)*.
47. ANDREA MONTANARI and AMIN SABERI, The spread of innovations in social networks, *Proceedings of the National Academy of Sciences*, Vol. **107** (47), 20196–20201, 2010.
48. MUKHERJEE, S., SON, J. & BHATTACHARYA, B. B. (2020) Phase transitions of the maximum likelihood estimates in the p-spin Curie-Weiss model. arXiv:2005.03631.
49. MUKHERJEE, R. & RAY, G. (2021) On testing for parameters in Ising model, *Annales de l’Institut Henri Poincaré. Probabilités et Statistiques, to appear*.
50. RAJARSHI MUKHERJEE, SUMIT MUKHERJEE and MING YUAN, Global testing against sparse alternatives under Ising models, *Annals of Statistics*, Vol. **46** (5), 2062–2093, 2018.
51. MATEY NEYKOV and HAN LIU, Property testing in high-dimensional Ising models, *Annals of Statistics*, Vol. **47** (5), 2472–2503, 2019.
52. OHKUWA, M., NISHIMORI, H. & LIDAR, D. A. (2018) Reverse annealing for the fully connected p-spin model. *Phys. Rev. A*, **98**, 022314.
53. PANCHENKO, D. (2013) *The Sherrington-Kirkpatrick model*. Springer.
54. DAVID K. PICKARD, Inference for discrete Markov Fields: the simplest nontrivial case, *J. Amer. Statist. Assoc.*, Vol. **82** (397), 90–96, 1987.
55. PRADEEP RAVIKUMAR, MARTIN J. WAINWRIGHT and JOHN D. LAFFERTY, High-dimensional Ising model selection using ℓ_1 -regularized logistic regression, *Annals of Statistics*, Vol. **38** (3), 1287–1319, 2010.
56. VOJTĚCH ROŮDL, ANDRZEJ RUCIŃSKI, and MATHIAS SCHACHT, Ramsey properties of random k-partite, k-uniform hypergraphs, *SIAM J. Discrete Math.*, Vol. **21** (2), 442–460, 2007.
57. NARAYANA P. SANTHANAM and MARTIN J. WAINWRIGHT, Information-theoretic limits of selecting binary graphical models in high dimensions, *IEEE Transactions on Information Theory*, Vol. **58** (7), 4117–4134, 2012.

58. SUZUKI, M. (1972) Solution and critical behavior of some “Three-Dimensional” Ising Models with a four-spin interaction. *Phys. Rev. Lett.*, **28**, 507–510.
59. MASUO SUZUKI and MICHAEL E. FISHER, Zeros of the partition function for the Heisenberg, ferroelectric, and general Ising models, *J. Math. Phys.*, Vol. **12** (2), 235–246, 1971.
60. MICHAEL TALAGRAND, The Parisi formula, *Annals of Mathematics. Second Series* (**2**) Vol. 163, 221–263, 2006.
61. TALAGRAND, M. & GLASSES, S. (2003) *A Challenge for Mathematicians-Cavity and Mean Field Models*. Berlin: Springer.
62. SHULONG TAN, JIAJUN BU, CHUN CHEN, BIN XU, CAN WANG, and XIAOFEI HE, Using rich social media information for music recommendation via hypergraph model, *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, Vol. **7**(1), Vol. 22, 2011.
63. ZE TIAN, TAEHYUN HWANG, and RUI KUANG. A hypergraph-based learning algorithm for classifying gene expression and arrayCGH data with prior knowledge, *Bioinformatics*, Vol. **25** (21), 2831–2838, 2009.
64. LOÏC TURBAN, One-dimensional Ising model with multispin interactions, *Journal of Physics A: Mathematical and Theoretical*, **49**(35), 2016.
65. VUFFRAY, M., MISRA, S. & LOKHOV, A. Y. (2020) Efficient Learning of Discrete Graphical Models. *Neural Information Processing System (NeurIPS)*.
66. VUFFRAY, M., MISRA, S., LOKHOV, A. Y. & CHERTKOV, M. (2016) Interaction Screening: Efficient and Sample-Optimal Learning of Ising Models. *Neural Information Processing System (NeurIPS)*, 2595–2603.
67. WU, S., SANGHAVI, S. & DIMAKIS, A. G. (2019) Sparse Logistic Regression Learns All Discrete Pairwise Graphical Models. *Neural Information Processing System (NeurIPS)*, 8069–8079.
68. LINGZHOU XUE, HUI ZOU, TIANXI CAI, Nonconcave penalized composite conditional likelihood estimation of sparse Ising models, *Annals of Statistics*, Vol. **40** (3), 1403–1429, 2012.
69. YAMASHIRO, Y., OHKUWA, M., NISHIMORI, H. & LIDAR, D. A. (2019) Dynamics of reverse annealing for the fully-connected p-spin model. *Phys. Rev. A*, **100**, 052321.

A. Properties of the Curie–Weiss threshold

Here, we will prove various properties of the Curie–Weiss threshold $\beta_{\text{CW}}^*(p) = \beta_{\text{ER}}^*(p, 1)$ (recall (2.16)).

LEMMA A.1. The Curie–Weiss threshold $\beta_{\text{CW}}^*(p)$ has the following properties:

- (1) $\lim_{p \rightarrow \infty} \beta_{\text{CW}}^*(p) = \log 2$.
- (2) The sequence $\{\beta_{\text{CW}}^*(p)\}_{p \geq 2}$ is strictly increasing.
- (3) $\beta_{\text{CW}}^*(2) = 0.5$.

Proof. Define the function $g_{\beta,p}(t) := \beta t^p - I(t)$. Since $g_{\beta,p}(1) = \beta - \log 2$, recalling (2.16), it immediately follows that $\beta_{\text{CW}}^*(p) \leq \log 2$. Now, take any $\beta < \log 2$. Note that $g_{\beta,2}(1) < 0$ and the function $t \mapsto g_{\beta,2}(t)$ is continuous at 1. Therefore, there exists $r \in (0, 1)$, such that $g_{\beta,2}(t) < 0$ for all $t \in [r, 1]$. Clearly, $g_{\beta,p}(t) \leq g_{\beta,2}(t) < 0$ for all $p \geq 2$ and $t \in [r, 1]$. Now, note that for all $t \in [0, 1]$,

$$g''_{\beta,p}(t) = \beta p(p-1)t^{p-2} - \frac{1}{1-t^2} \leq \beta p(p-1)t^{p-2} - 1. \quad (\text{A.1})$$

Since $\lim_{p \rightarrow \infty} p(p-1)r^{p-2} = 0$, there exists $p(\beta) \geq 2$, such that $g''_{\beta,p}(r) < 0$ for all $p \geq p(\beta)$. Hence, $g''_{\beta,p}(t) < 0$ for all $t \in [0, r]$ and $p \geq p(\beta)$. This, together with the fact that $g'_{\beta,p}(0) = 0$, implies that $g_{\beta,p}$ is strictly decreasing on $[0, r]$ for all $p \geq p(\beta)$. Moreover, because $g_{\beta,p}(0) = 0$, it follows that $g_{\beta,p}(t) \leq 0$ for all $t \in [0, r]$ and $p \geq p(\beta)$. Hence, $g_{\beta,p}(t) \leq 0$ for all $t \in [0, 1]$ and $p \geq p(\beta)$, i.e. $\beta_{\text{CW}}^*(p) \geq \beta$ for all $p \geq p(\beta)$. This shows that $\beta_{\text{CW}}^*(p) \rightarrow \log 2$, as $p \rightarrow \infty$.

Next, we show that the sequence $\{\beta_{\text{CW}}^*(p)\}_{p \geq 2}$ is strictly increasing. Towards this, take any $p \geq 3$. It follows from [(48), Lemma F.1] that there exists $a \in (0, 1)$, such that 0 and a are both global maximizers of $g_{\beta_{\text{CW}}^*(p),p}$. In particular, $g_{\beta_{\text{CW}}^*(p),p}(a) = 0$, and hence, $g_{\beta_{\text{CW}}^*(p),p-1}(a) > 0$. The function $\beta \mapsto g_{\beta,p-1}(a)$ being continuous, there exists $\beta < \beta_{\text{CW}}^*(p)$, such that $g_{\beta,p-1}(a) > 0$. Hence, $\beta_{\text{CW}}^*(p-1) \leq \beta$, establishing that $\beta_{\text{CW}}^*(p-1) < \beta_{\text{CW}}^*(p)$.

Finally, we show that $\beta_{\text{CW}}^*(2) = 0.5$. By (A.1), $\sup_{t \in [0,1]} g''_{\beta,p}(t) < 0$ for all $\beta < 0.5$. This, coupled with the facts that $g'_{\beta,p}$ and $g_{\beta,p}$ vanish at 0, implies that $\sup_{t \in [0,1]} g_{\beta,p}(t) = 0$ for all $\beta < 0.5$. Hence, $\beta_{\text{CW}}^*(2) \geq 0.5$. On the other hand, for any $\beta > 0.5$, $g''_{\beta,p}(0) = 2\beta - 1 > 0$ and hence, by continuity of the function $g''_{\beta,p}$ at 0, there exists $\varepsilon > 0$, such that $\inf_{t \in [0,\varepsilon]} g''_{\beta,p}(t) > 0$. Once again, since $g'_{\beta,p}$ and $g_{\beta,p}$ vanish at 0, we have $g_{\beta,p}(t) > 0$ for all $t \in (0, \varepsilon]$. This shows that $\beta_{\text{CW}}^*(2) \leq 0.5$. $\square$