



HAL
open science

Maximum Likelihood Estimation of Sparse Networks with Missing Observations

Solenne Gaucher, Olga Klopp

► **To cite this version:**

Solenne Gaucher, Olga Klopp. Maximum Likelihood Estimation of Sparse Networks with Missing Observations. 2021. hal-02050003v2

HAL Id: hal-02050003

<https://hal.science/hal-02050003v2>

Preprint submitted on 27 Apr 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Maximum Likelihood Estimation of Sparse Networks with Missing Observations

Solenne Gaucher ^{*1} and Olga Klopp ^{†2,1}

¹CREST, ENSAE

²ESSEC Business School

April 27, 2021

Abstract

Estimating the matrix of connections probabilities is one of the key questions when studying sparse networks. In this work, we consider networks generated under the sparse graphon model and the inhomogeneous random graph model with missing observations. Using the Stochastic Block Model as a parametric proxy, we bound the risk of the maximum likelihood estimator of network connections probabilities, and show that it is minimax optimal. Moreover, we show that our estimator can be efficiently approximated using tractable variational methods, and thus used in practice.

Keywords— Missing observations, network models, sparse estimation, graphon model

1 Introduction

In the past two decades, networks have attracted considerable attention, as many scientific fields are concerned by the advances made in the understanding of these complex systems. In social sciences [53] as in physics [3] and biology [55], networks are used to represent a great variety of systems of interactions between social agents, particles, proteins or neurons. These networks are often modeled as an observation drawn from a random graph.

Missing observations is a common problem when studying real life networks. In social sciences, data coming from sample surveys are likely to be incomplete, especially, when dealing with large or hard-to-find populations. While biologists often use graphs to model interactions between proteins, experimental discovery of these interactions can require substantial time and investment from the scientific community [10]. In many cases, collecting complete information on relations between actors can be difficult, expensive and time-consuming [34, 56, 27, 24]. On the other hand, the emergence of detailed data sets coming, for example, from social networks or genome sequencing has fostered new challenges, as their large size makes using the full data computationally unattractive. This has lead scientists to consider only sub-samples of the available data [7]. However, incomplete observation of the network structure may considerably affect the accuracy of inference methods [33].

Our work focuses on the study of the inhomogeneous random graph model with *missing observations*. In this setting, the problem of estimating the matrix of connections probabilities is of primary interest. Minimax optimal convergence rates for this problem have been shown to be attained by the least square estimator under full observation of the network for dense graphs in [21] and for sparse graphs in [30]. In [20], the authors extended these results to the setting in which observations about the presence or absence of an edge are missing independently at random with the same probability p . However their estimator requires the knowledge of p , and cannot be extended to non-uniform sampling schemes. Unfortunately, least square estimation is too costly to be used in practice. Many other approaches have been proposed, for example, spectral clustering [43, 25, 49], modularity maximization [46, 9], belief propagation [19], neighborhood smoothing [57], convex relaxation of k-means clustering [23] and of likelihood maximization [5], and universal singular value thresholding [14, 31, 54]. It was conjectured the existence of possible computational gap when no polynomial time algorithm can achieve minimax optimal rate of convergence. The present work shows that this is not the case.

*solenne.gaucher@ensae.fr

†kloppolga@math.cnrs.fr

In this work, we consider the maximum likelihood estimator. This estimator is also NP-hard but its computationally efficient approximations (under some additional conditions) have been proposed in the literature (see, e.g., [42] for a detailed review of these methods). For example, the authors of [4] suggest to use pseudo-likelihood methods, as it leads to computationally tractable estimators. Alternatively, in [17] a tractable variational approximation of the maximum likelihood estimator is proposed. This method has been applied successfully to study biological networks, political blogosphere networks and seeds exchange networks [48, 50, 35]. The authors of [8] show asymptotic normality of the maximum likelihood estimate and of its variational approximation for sparse graphs generated by stochastic block models when the connections probabilities of the different communities are well separated. In [50], these results are extended to the case of missing observations. These methods suffer from a lack of theoretical guarantees when the model is misspecified or non-identifiable. On the other hand, to the best of our knowledge, no non-asymptotic bound has been established for the risk of the maximum likelihood estimator. In this work, we close this gap and show that the maximum likelihood estimator is minimax optimal in a number of scenarios, and adaptive to non-uniform missing data schemes. Moreover we show that it can be efficiently approximated using tractable variational methods.

Our results find a natural application in predicting the existence of non-observed edges, a commonly encountered problem called *link prediction* [39, 58]. Interaction networks are often incomplete, as detecting interactions can require significant experimental effort. Instead of exhaustively testing for every connection, one might be interested in deducing the pairs of agents which are most likely to interact based on the relations already recorded and on available covariates. If these estimations are precise enough, testing for these interactions would enable scientists to establish the network topology while substantially reducing the costs [16]. In this context, estimating the probabilities of connections through likelihood maximization enables to accordingly rank unobserved pairs of nodes. Link prediction also finds applications in recommender systems for social networks. The missing observation scheme studied in this work is motivated by the above examples, and generalizes the model described in [20].

1.1 Inhomogeneous random graph model

We consider an undirected, unweighted graph with n nodes indexed from 1 to n . Its connectivity can be encoded by its *adjacency matrix* \mathbf{A} , defined as follows: set \mathbf{A} a $n \times n$ symmetric matrix such that for any $i < j$, $\mathbf{A}_{ij} = 1$ if there exists an edge between node i and node j , $\mathbf{A}_{ij} = 0$ otherwise. In our model, we consider that there is no edge linking a node to itself, so $\mathbf{A}_{ii} = 0$ for any i . We assume that the variables $(\mathbf{A}_{ij})_{1 \leq i < j \leq n}$ are independent Bernoulli random variables of parameter Θ_{ij}^* , where Θ^* is a $n \times n$ symmetric matrix with zero diagonal entries. This matrix Θ^* corresponds to the matrix of probabilities of observing an edge between nodes i and j . This model is known as the *inhomogeneous random graph* model:

$$\forall 1 \leq i < j \leq n, \mathbf{A}_{ij} | \Theta_{ij}^* \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(\Theta_{ij}^*). \quad (1)$$

In the present paper we consider the following problem: from a single partial observation of the graph, that is, given a sample of entries of the adjacency matrix \mathbf{A} , we want to estimate the matrix of connections probabilities Θ^* .

The problem of estimating Θ^* when some entries of the adjacency matrix are not observed is closely related to the 1-bit matrix completion problem. The matrix completion problem [12, 32, 45] aims at recovering a matrix which is only partially observed. More precisely, we observe a random sample of its entries, which may be corrupted by some noise, and we wish to infer the rest of the matrix. In 1-bit matrix completion, first introduced in [18], the entries (i, j) of the observed matrix can only take two values $\{0, 1\}$ with probabilities given respectively by $f(\mathbf{M}_{ij})$ and $1 - f(\mathbf{M}_{ij})$. Here, the matrix \mathbf{M} corresponds to the real quantity of interest that one would like to infer, and the function f can be seen as the cumulative distribution function of the noise. A typical assumption in this setting is that the matrix \mathbf{M} is low-rank. In [29], the authors show that for 1-bit matrix completion the restricted penalized maximum likelihood estimator is minimax optimal up to a log factor. The methods used in our proofs are, to some extent, inspired by the methods developed for the framework of matrix completion. However, the problem we have in hand is in many aspects different from the 1-bit matrix completion problem. The structure of the connections probabilities matrix Θ^* and the sparsity of the network allow for faster rates of convergence, and the techniques of proof required to match the minimax optimal convergence rate are more involved.

Our approach for estimating the matrix of connections probabilities is based on the celebrated Regularity Lemma by Szemerédi [38], which implies that any graph can be well approximated by a stochastic block model (SBM). We refer to [38] for a more detailed presentation of this result. In the SBM, each node i is associated with a community $z^*(i)$, where $z^* : [n] \rightarrow [k]$ is called the index function. This index function can either be treated as a parameter to estimate (this model is sometimes called the conditional stochastic block model), or as a latent variable. In this case, the indexes follow a multinomial distribution: $\forall i, z^*(i) \stackrel{i.i.d.}{\sim} \mathcal{M}(1; \alpha^*)$ where $\forall l \in [k]$, α_l is the probability that node i belongs to the community l . Given this index function, the probability that there exists an edge between nodes i and j depends only on the communities of i and j . For example, when considering citations networks, where two articles are linked if one is cited by the other, it amounts to saying that the probability that two articles are linked

only depends on their topic. Similarly, if one considers students of a school in a social network, it is a reasonable assumption to say that the probability that two students are linked only depends on their cohorts. This implies that the matrix of connections probabilities Θ^* can be factorized as follows: $\Theta_{ij}^* = \mathbf{Q}_{z^*(i)z^*(j)}^*$, with \mathbf{Q}^* a $k \times k$ symmetric matrix such that \mathbf{Q}_{ab}^* is the probability that there exists an edge between a given member of the community a and a given member of the community b , so we have that the conditional SBM can be written as:

$$\begin{aligned} & \exists \mathbf{Q}^* \in [0, 1]_{\text{sym}}^{k \times k}, \exists z^* : [n] \rightarrow [k] \\ & \forall 1 \leq i < j \leq n, \mathbf{A}_{ij} | \mathbf{Q}^*, z^* \stackrel{\text{ind.}}{\sim} \text{Bernoulli} \left(\mathbf{Q}_{z^*(i)z^*(j)}^* \right), \mathbf{A}_{ii} = 0. \end{aligned} \quad (2)$$

While considering the SBM, the problem of estimating the matrix of connections probabilities reduces to estimating the label function z^* and the matrix of probabilities of connections between communities \mathbf{Q}^* .

In the past decade, the stochastic block model has known a growing interest from the statistical community and an important part of the work has focused on the problem of community recovery (i.e., the recovery of the vector of communities populations α^* , or of the label function z^* in the conditional model). Theoretical guarantees for this problem were established under quite strong assumptions on the matrix of probabilities of connections between communities, \mathbf{Q}^* , see, for example, [41, 11, 1, 44].

Note that our results hold without assuming the existence of the true community structure, that is, without assuming that the matrix Θ^* is block constant. With this in mind, we will focus on estimating the distribution giving rise to the adjacency matrix, i.e., on estimating Θ^* , rather than on estimating the label function or the populations of the communities. One important question in this setting is how to choose the number of communities for our estimator, as more communities implies a smaller bias and a greater variance. Optimizing this trade-off requires, first, establishing a non-asymptotic bound on the risk of our estimator for a number of communities that may depend on the number of nodes, and, in a second time, bounding the bias of an oracle block constant estimator.

Our work focuses on relevant in applications setting of partial observations of the network. We consider the following missing value setting. Let $\mathbf{X} \in \{0, 1\}_{\text{sym}}^{n \times n}$ denote the sampling matrix given by $\mathbf{X}_{ij} = 1$ if we observe \mathbf{A}_{ij} and $\mathbf{X}_{ij} = 0$ otherwise. We assume that the sampling matrix \mathbf{X} is random and, conditionally on Θ^* , independent from the adjacency matrix \mathbf{A} . For any $1 \leq i < j \leq n$, its entries \mathbf{X}_{ij} are mutually independent. Finally, we denote by $\mathbf{\Pi} \in [0, 1]_{\text{sym}}^{n \times n}$ the unknown matrix of sampling probabilities such that $\mathbf{X}_{ij} \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(\mathbf{\Pi}_{ij})$. This sampling scheme includes for instance node-based sampling schemes such as the exo-centered design described in [27], where we observe \mathbf{A}_{ij} if i or j belongs to the set of sampled nodes. It also covers random dyad sampling schemes (described, e.g., in [50]). In this case, the probability of observing the entry \mathbf{A}_{ij} is allowed to depend on the communities of i and j .

1.2 Graphon model

While studying exchangeable random graphs, important questions such as how to compare two graphs with different numbers of nodes or how to study graphs with an increasing number of nodes call for a more general, non-parametric model. One of such models that has attracted a lot of attention recently is the *graphon* model [47, 21, 30, 54]. In this model, the connections probabilities Θ_{ij}^* are the following random variables

$$\Theta_{ij}^* = W^*(\zeta_i, \zeta_j) \quad (3)$$

where ζ_1, \dots, ζ_n are unobserved (latent) independent random variables sampled uniformly in $[0, 1]$. The graph is then sampled according to the inhomogeneous random graph model (1). The function $W^* : [0, 1]^2 \rightarrow [0, 1]$ is measurable, symmetric and is called a graphon. Graphs encountered in practice are usually *sparse*: the expected number of edges grows as $\rho_n n^2$ where ρ_n is a decreasing sequence of sparsity inducing parameters. The dense graphon model can be modified in order to account for this sparsity:

$$\Theta_{ij}^* = \rho_n W^*(\zeta_i, \zeta_j). \quad (4)$$

Since the law of the graph is invariant under any change of labelling of its nodes, different graphons can give rise to the same distribution on the space of graphs of size n . More precisely, let W be a graphon and $\tau : [0, 1] \rightarrow [0, 1]$ be a measure-preserving function. We write $W_\tau(x, y) = W(\tau(x), \tau(y))$ and say that two graphons U and V are *weakly isomorphic* if there exists measure-preserving maps τ, ϕ such that $U_\tau = W_\phi$ almost everywhere. It is established in Section 10, [38] that two graphons define the same probability measure on graphs if and only if they are weakly isomorphic.

In the present paper we also consider the setting when the matrix of connections probabilities is generated following the sparse graphon model (4). We deal with two classes of graphon functions previously studied in the literature, step-function graphons and smooth graphons, under the scenario of partial observations of the network.

1.3 Outline of the paper

The present paper is devoted to the theoretical study of the maximum likelihood estimator in sparse network models with missing observations. First, we provide an oracle bound for the risk of the maximum likelihood estimator of the matrix of connections probabilities from a partial observation of the adjacency matrix \mathbf{A} . Our results hold under fairly general assumptions on the missing observations scheme and we show that the maximum likelihood estimator matches the minimax optimal rates of convergence in a variety of scenarios, while being fully adaptive to the missing data scheme. Second, we provide a parameter-free version of our estimator which, in particular, does not require the knowledge of the sparsity parameter ρ_n . We also bound the Kullback-Leibler divergence between the true matrix of connections probabilities and its block constant approximation, and derive an optimal choice for the number of communities defining the maximum likelihood estimator. We show that the maximum likelihood estimator can be approximated using tractable variational methods, and we provide a bound on the risk of the variational estimator.

This manuscript is organized as follows. In Section 2.1, we introduce the maximum likelihood estimator for the matrix of connections probabilities Θ^* from partial observation of the adjacency matrix \mathbf{A} . Then, Theorem 1 in Section 2.2 provides a non-asymptotic oracle bound on the risk of this estimator. As a consequence, we show that our estimator is minimax optimal in a number of scenarios and derive the corresponding bound for estimating Θ^* in the case of full observation of the adjacency matrix \mathbf{A} . Our estimation method requires bounds on the entries of Θ^* . In Section 2.3, we first propose a method to choose these bounds under fairly general assumptions and, in Section 2.4, we specify it to the case of sparse graphon model (4). We show that the resulting adaptive estimator is minimax optimal up to a log factor. In Section 2.5, Theorem 4, we provide the choice for the number of communities that achieves the best trade off between the variability of our estimate and the fit of the oracle model. Finally, in Section 3, we show that the maximum likelihood estimator can be consistently approximated using a tractable variational estimator.

1.4 Notations

We provide here a summary of the notations used throughout this paper.

- For any positive integer d , we denote by $[d]$ the set $\{1, \dots, d\}$.
- For any set \mathcal{S} , we denote by $|\mathcal{S}|$ its cardinality.
- For any matrix \mathbf{A} , we denote by A_{ij} its entry on row i and column j . If $\mathbf{A} \in [0, 1]^{n \times n}$ and \mathbf{A} is symmetric, we write $\mathbf{A} \in [0, 1]_{\text{sym}}^{n \times n}$.
- Let $\mathcal{K}(q, q') = q \log\left(\frac{q}{q'}\right) + (1-q) \log\left(\frac{1-q}{1-q'}\right)$ denote the Kullback-Leibler divergence of a Bernoulli distribution with parameter q from a Bernoulli distribution with parameter q' . For any three symmetric matrices with zero diagonal entries $\mathbf{A}, \mathbf{B}, \mathbf{X} \in [0, 1]_{\text{sym}}^{n \times n}$ we set

$$\mathcal{K}_{\mathbf{X}}(\mathbf{A}, \mathbf{B}) = \sum_{i < j} \mathbf{X}_{ij} \mathcal{K}(A_{ij}, B_{ij}) \quad \text{and} \quad \mathcal{K}(\mathbf{A}, \mathbf{B}) = \sum_{i < j} \mathcal{K}(A_{ij}, B_{ij}).$$

- For any three symmetric matrices with zero diagonal entries $\mathbf{A}, \mathbf{B}, \mathbf{C} \in \mathbb{R}_{\text{sym}}^{n \times n}$, let $\langle \mathbf{A} | \mathbf{B} \rangle = \sum_{i < j} \mathbf{A}_{ij} \mathbf{B}_{ij}$, $\langle \mathbf{A} | \mathbf{B} \rangle_{\mathbf{C}} = \sum_{i < j} \mathbf{C}_{ij} \mathbf{A}_{ij} \mathbf{B}_{ij}$, $\|\mathbf{A}\|_2 = \sqrt{\langle \mathbf{A} | \mathbf{A} \rangle}$, $\|\mathbf{A}\|_{2, \mathbf{C}} = \sqrt{\langle \mathbf{A} | \mathbf{A} \rangle_{\mathbf{C}}}$, and $\|\mathbf{A}\|_{\infty} = \max_{i, j} |A_{ij}|$. With these notations, $\|\mathbf{A}\|_{2, \mathbf{\Pi}} = \sqrt{\mathbb{E} \left[\|\mathbf{A}\|_{2, \mathbf{X}}^2 \right]}$ corresponds to the L_2 -norm of the matrix \mathbf{A} with respect to the sampling probabilities $\mathbf{\Pi}$.
- We denote by $\mathcal{Z}_{n, k}$ the label functions $z : [n] \rightarrow [k]$. For any $z \in \mathcal{Z}_{n, k}$, we denote by \mathcal{T}_z the set of block constant matrices corresponding to the label z :

$$\mathcal{T}_z \triangleq \left\{ \mathbf{A} : \forall i \in [n], A_{ii} = 0 \ \& \ \exists \mathbf{Q} \in [0, 1]^{k \times k}, \forall 1 \leq i < j \leq n, A_{ij} = A_{ji} = Q_{z(i)z(j)} \right\}.$$
- To ease notations, for $\mathbf{A} \in \mathcal{T}_z$ and $(a, b) \in [k]^2$, we sometimes denote by $\mathbf{A}_{z^{-1}(a)z^{-1}(b)}$ any entry \mathbf{A}_{ij} such that $(i, j) \in (z^{-1}(a), z^{-1}(b))$ and $i \neq j$. We write $\mathcal{T}_k = \bigcup_{z \in \mathcal{Z}_{n, k}} \mathcal{T}_z$.
- We denote by C and C' positive constants that can vary from line to line. These are absolute constants unless otherwise mentioned. For any two positive sequences $(a_n)_{n \in \mathbb{N}}, (b_n)_{n \in \mathbb{N}}$, we write $a_n = \omega(b_n)$ if $a_n/b_n \rightarrow \infty$.
- We denote respectively by $\mathbb{E}^{\mathbf{X}}$ and $\mathbb{P}^{\mathbf{X}}$ the expectation and the probability conditionally on the random variable \mathbf{X} , and respectively by \mathbb{E} and \mathbb{P} the expectation and the probability over all random variables.

2 Convergence rate for the maximum likelihood estimator

2.1 Maximum likelihood estimator under missing observations

We start by introducing the conditional log-likelihood for the model (1). Conditionally on the probability matrix Θ^* , the entries $(\mathbf{A}_{ij})_{1 \leq i < j \leq n}$ of the adjacency matrix \mathbf{A} are independent Bernoulli variables with parameters $(\Theta_{ij}^*)_{1 \leq i < j \leq n}$. Therefore, for any $\Theta \in [0, 1]^{n \times n}$, the conditional log-likelihood of the parameter matrix Θ with respect to the observed entries of the adjacency matrix \mathbf{A} is given by

$$\mathcal{L}_{\mathbf{X}}(\mathbf{A}; \Theta) = \sum_{i < j} \mathbf{X}_{ij} (\mathbf{A}_{ij} \log(\Theta_{ij}) + (1 - \mathbf{A}_{ij}) \log(1 - \Theta_{ij})).$$

For any $z \in \mathcal{Z}_{n,k}$ and $\mathbf{Q} \in [0, 1]_{\text{sym}}^{k \times k}$, the matrix of connections probabilities corresponding to the block model (z, \mathbf{Q}) is given by $\Theta_{ij} = \mathbf{Q}_{z(i)z(j)}$ for $1 \leq i < j \leq n$ and $\Theta_{ii} = 0$ for $i \in [n]$. With these notations, the conditional log-likelihood of a block model (z, \mathbf{Q}) with respect to the observed entries of the adjacency matrix \mathbf{A} is

$$\begin{aligned} \mathcal{L}_{\mathbf{X}}(\mathbf{A}; z, \mathbf{Q}) &= \sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} (\mathbf{A}_{ij} \log(\mathbf{Q}_{z(i)z(j)}) + (1 - \mathbf{A}_{ij}) \log(1 - \mathbf{Q}_{z(i)z(j)})) \\ &= \sum_{1 \leq a \leq b \leq k} \sum_{\substack{i \in z^{-1}(a), j \in z^{-1}(b) \\ i \neq j}} \mathbf{X}_{ij} (\mathbf{A}_{ij} \log(\mathbf{Q}_{ab}) + (1 - \mathbf{A}_{ij}) \log(1 - \mathbf{Q}_{ab})) \\ &= \sum_{a \leq b} \log(\mathbf{Q}_{ab}) \sum_{\substack{i \in z^{-1}(a), j \in z^{-1}(b) \\ i \neq j}} \mathbf{X}_{ij} \mathbf{A}_{ij} + \sum_{a \leq b} \log(1 - \mathbf{Q}_{ab}) \sum_{\substack{i \in z^{-1}(a), j \in z^{-1}(b) \\ i \neq j}} \mathbf{X}_{ij} (1 - \mathbf{A}_{ij}). \end{aligned}$$

The maximum likelihood estimator for the stochastic block model is

$$(\widehat{\mathbf{Q}}, \widehat{z}) \in \arg \max_{\mathbf{Q} \in [0, 1]_{\text{sym}}^{k \times k}, z \in \mathcal{Z}_{n,k}} \mathcal{L}_{\mathbf{X}}(\mathbf{A}; z, \mathbf{Q}).$$

The block constant maximum likelihood estimator of Θ^* is defined as $\widehat{\Theta}_{ij} = \widehat{\mathbf{Q}}_{z(i)z(j)}$ for any $i < j$. Note that maximizing the log-likelihood is equivalent to minimizing a sum of Bernoulli Kullback-Leibler divergences. Indeed, an easy calculation leads

$$(\widehat{\mathbf{Q}}, \widehat{z}) \in \arg \max_{\mathbf{Q} \in [0, 1]_{\text{sym}}^{k \times k}, z \in \mathcal{Z}_{n,k}} \mathcal{L}_{\mathbf{X}}(\mathbf{A}; z, \mathbf{Q}) = \arg \min_{\mathbf{Q} \in [0, 1]_{\text{sym}}^{k \times k}, z \in \mathcal{Z}_{n,k}} \sum_{i < j} \mathbf{X}_{ij} \mathcal{K}(\mathbf{A}_{ij}, \mathbf{Q}_{z(i)z(j)}). \quad (5)$$

Moreover, for any fixed assignment $z \in \mathcal{Z}_{n,k}$ and any sampling matrix \mathbf{X} , the log-likelihood with regards to the observed entries of \mathbf{A} will be maximized by taking

$$\mathbf{Q}_{ab} = \overline{\mathbf{X} \mathbf{A}_{ab}}^z \triangleq \frac{\sum_{\substack{i \in z^{-1}(a), j \in z^{-1}(b) \\ i \neq j}} \mathbf{X}_{ij} \mathbf{A}_{ij}}{\sum_{\substack{i \in z^{-1}(a), j \in z^{-1}(b) \\ i \neq j}} \mathbf{X}_{ij}}.$$

Note that for any label function z , maximizing the likelihood or minimizing the least square criterion defined as $\mathcal{C}_{\mathbf{X}}(\mathbf{A}; z, \mathbf{Q}) = \sum_{i < j} \mathbf{X}_{ij} (\mathbf{A}_{ij} - \mathbf{Q}_{z(i)z(j)})^2$ with regards to \mathbf{Q} yields the same estimator. However, the label functions selected by the two criterions can be different, as is shown in figure 1.

In the rest of this work, we will denote by $\widehat{\Theta}$ the oracle probability matrix, i.e., the best approximation to Θ^* in the sense of the weighted Kullback Leibler divergence:

$$\begin{aligned} \widetilde{\Theta}_{i < j} &= \mathbf{Q}_{z^*(i)z^*(j)}^*, \widetilde{\Theta}_{ii} = 0 \\ (\mathbf{Q}^*, z^*) &\in \arg \min_{\mathbf{Q} \in [0, 1]_{\text{sym}}^{k \times k}, z \in \mathcal{Z}_{n,k}} \sum_{i < j} \mathcal{K}_{\Pi}(\Theta_{ij}^*, \mathbf{Q}_{z(i)z(j)}). \end{aligned} \quad (6)$$

2.2 Upper bound on the risk of the restricted maximum likelihood estimator

In this section, we establish an upper bound on the risk of the maximum likelihood estimator and show that it matches the minimax convergence rate obtained in [30, 20]. We will measure the risk of our estimator in Π -weighted Frobenius norm. To bound the risk of the maximum likelihood estimator, we assume that there exists sequences ρ_n and γ_n such that $\forall i < j$,

$$0 < \gamma_n \leq \Theta_{ij}^* \leq \rho_n < 1. \quad (7)$$



(a) Communities obtained using the least squares criterion.

(b) Communities obtained using the maximum likelihood criterion.

Figure 1: We fit a SBM with two communities to the graph above using the maximum likelihood criterion (left) and the least squares criterion (right). Four nodes are classified differently by the two methods.

Note that for sparse graphs, ρ_n corresponds to the sparsity inducing sequence in equation (4).

The assumption (7) is systematic in the literature studying the maximum likelihood estimator for the stochastic block model, as it guarantees that the loss associated to the maximum likelihood estimator is Lipschitz. See, for example, [8] and [52], where the authors assume that the adjacency matrix is generated by an homogeneous stochastic block model for which the matrix \mathbf{Q}^*/ρ_n has entries bounded away from 0. In our model, this corresponds to imposing that $\rho_n = O(\gamma_n)$, i.e., that all entries of the matrix of connections probabilities Θ^* are of the same order of magnitude. Our assumptions are more general than the one developed in these articles, as they also cover the case $\gamma_n = o(\rho_n)$.

In [13], the authors consider the dense SBM and assume that the entries of \mathbf{Q}^* belong to $\{0\} \cup [\zeta, 1 - \zeta] \cup \{1\}$ for some $\zeta > 0$. They prove the consistency of the maximum likelihood estimator constrained to a restricted subset of the parameters. However, the definition of this subset implies knowing the set $\Omega_0 = \{(i, j) : \Theta_{ij}^* \in \{0, 1\}\}$ prior to estimating the matrix of connections probabilities. Note that, if we assume that Ω_0 is known and that \mathbf{Q}^* belong to $\{0\} \cup [\zeta, 1 - \zeta] \cup \{1\}$, we can set $\hat{\Theta}_{ij} = 0$ for any $(i, j) \in \Omega_0$ and estimate the remaining entries (which are bounded away from 0 and 1) with our procedure.

On the other hand, cases where the entries of Θ^* are of different order of magnitude are common in the literature in the case of planted partition models and assortative and disassortative SBM. In the planted partition model, the matrix of connections probabilities between communities is given by $\mathbf{Q}^* = (p - q)\mathbf{I}_k + q\mathbf{1}_k\mathbf{1}_k^T$, where $p > q$, \mathbf{I}_k is the identity matrix and $\mathbf{1}_k\mathbf{1}_k^T$ the matrix whose entries are all equal to 1. This amounts to saying that the probability that two nodes are connected only depends on whether they belong to the same community or not. This model can be relaxed to give rise to the assortative model, where the within group probabilities of connection \mathbf{Q}_{aa}^* are larger than the between group probabilities of connection \mathbf{Q}_{bc}^* : there exists $p, q \in [0, 1]$ such that for any $a \neq b$, one has $\mathbf{Q}_{ab}^* \leq q < p \leq \mathbf{Q}_{aa}^*$. The disassortative model corresponds to the case where between communities connections are more likely than within community connections: one has for any $a \neq b$, $\mathbf{Q}_{aa}^* \leq q < p \leq \mathbf{Q}_{ab}^*$. The last two models are closely related. Indeed, if \mathbf{A} is drawn from an assortative SBM, $\mathbf{1}_n\mathbf{1}_n^T - \mathbf{I}_n - \mathbf{A}$ corresponds to a realization of a disassortative SBM.

In the planted partition model, maximizing the likelihood is equivalent to finding a partition maximizing the within group connectivity, i.e., maximizing $\sum_{i < j} \mathbf{A}_{ij} \mathbf{Z}_{ij}$ where $\mathbf{Z}_{ij} = \mathbb{1}\{z(i) = z(j)\}$. Convex relaxations of the constraints on \mathbf{Z} have been studied in the literature [26, 6, 2], and theoretical guarantees for these algorithms for the problem of communities recovery have been established under assumptions on the gap $p - q$. In these models, communities are characterized by higher (respectively lower) connectivity, and the assumption that $q \ll p$ actually makes the recovery problem easier. By contrast, the definition of a community in the SBM as a set of nodes with the same stochastic behaviour is far more general. It covers settings not suitably described by assortative or disassortative models, as, e.g., graphs with leaders and followers such as the well known example of Zachary's Karate Club (see, e.g., [36]). In these models, leaders are seldomly linked one to another, but are highly connected to their own set of followers. On the other hand, followers rarely connect one to another or to more than one leader. By comparison, our results hold without any assumption on the assortativity or the disassortativity of the model.

In a first time, we assume that we know γ_n and ρ_n . We will discuss how to estimate these values in Section 2.4. Let $\hat{\Theta}$ be the block constant estimator based on the maximization of the likelihood among block constant matrices

with entries in $[\gamma_n, \rho_n]$:

$$\begin{aligned} \widehat{\Theta}_{i < j} &= \widehat{Q}_{\widehat{z}(i)\widehat{z}(j)}, \quad \widehat{\Theta}_{ii} = 0 \\ (\widehat{Q}, \widehat{z}) &\in \arg \min_{\substack{Q \in [\gamma_n, \rho_n]_{\text{sym}}^{k \times k}, z \in \mathcal{Z}_{n,k} \\ i < j}} \sum_{i < j} \mathbf{X}_{ij} \mathcal{K}(\mathbf{A}_{ij}, \mathbf{Q}_{z(i)z(j)}). \end{aligned} \quad (8)$$

Here we assume that k is fixed and that it can depend on the number of nodes n . k can be chosen using a network cross-validation method [15] or, when the graphon is a step function, it can be chosen using a sequential goodness-of-fit testing procedure [37] or a likelihood-based model selection method [52]. When graphon is Hölder-continuous, we provide a choice of k to optimize the usual trade-off between bias and variance of our estimator in Section 2.4. Note that the maximum likelihood estimator does not require the knowledge of the matrix of sampling probabilities $\mathbf{\Pi}$. On the other hand, this matrix, which characterizes the difficulty induced by missing observations, appears when bounding the error of this estimator. In Theorem 1, we provide a general bound on the $\|\cdot\|_{2, \mathbf{\Pi}}$ -norm of the error, while in Corollaries 1 and 2 we derive specific bounds on its $\|\cdot\|_2$ -norm under further assumptions on the matrix $\mathbf{\Pi}$.

Theorem 1. *Assume that \mathbf{A} is drawn according to (1), and that $\rho_n = \omega(n^{-1})$. Then, there exists absolute constants $C, C' > 0$ such that with probability at least $1 - 9 \exp(-C\rho_n n \log(k))$*

$$\|\Theta^* - \widehat{\Theta}\|_{2, \mathbf{\Pi}}^2 \leq C' \rho_n \left(\mathcal{K}_{\mathbf{\Pi}}(\Theta^*, \widetilde{\Theta}) + \frac{\rho_n^2}{(1 - \rho_n)^2 \wedge \gamma_n^2} (k^2 + n \log(k)) \right). \quad (9)$$

Remark 1. *Note that we are not interested in regimes for which $\rho_n = O(n^{-1})$, as Theorem 2 implies that in this setting the constant estimator with all entries equal to the average node degree attains the minimax rate.*

Remark 2. *This bound is stated as a function of the of weighted Kullback-Leibler divergence $\mathcal{K}_{\mathbf{\Pi}}$ and of the oracle matrix $\widetilde{\Theta}$ defined in (6). Note that it implies the weaker bound*

$$\|\Theta^* - \widehat{\Theta}\|_{2, \mathbf{\Pi}}^2 \leq C' \rho_n \left(\mathcal{K}(\Theta^*, \widetilde{\Theta}^f) + \frac{\rho_n^2}{(1 - \rho_n)^2 \wedge \gamma_n^2} (k^2 + n \log(k)) \right)$$

where $\widetilde{\Theta}^f$ is the oracle matrix for the full Kullback-Leibler divergence \mathcal{K} . Indeed, one has $\mathcal{K}_{\mathbf{\Pi}}(\Theta^*, \widetilde{\Theta}) \leq \mathcal{K}_{\mathbf{\Pi}}(\Theta^*, \widetilde{\Theta}^f) \leq \mathcal{K}(\Theta^*, \widetilde{\Theta}^f)$.

In the case where all entries are observed, that is $\mathbf{\Pi}_{ij} = 1$ for any $i < j$, the rate attained by the maximum likelihood estimator is given by the following corollary.

Corollary 1. *Assume that \mathbf{A} is drawn according to (1), that $\forall 1 \leq i < j \leq n$, $\mathbf{\Pi}_{ij} = 1$ and that $\rho_n = \omega(n^{-1})$. Then, there exists positive constants $C, C' > 0$ such that with probability at least $1 - 9 \exp(-C\rho_n n \log(k))$*

$$\|\Theta^* - \widehat{\Theta}\|_2^2 \leq C' \rho_n \left(\mathcal{K}(\Theta^*, \widetilde{\Theta}) + \frac{\rho_n^2}{(1 - \rho_n)^2 \wedge \gamma_n^2} (k^2 + n \log(k)) \right).$$

If we assume that the probability of observing any entry of the adjacency matrix is bounded away from 0, Theorem 1 can be adapted to provide a bound on the risk of our estimator under the Frobenius norm. Indeed, if $\min_{1 \leq i < j \leq n} \{\mathbf{\Pi}_{ij}\} \geq p$, then $\|\Theta^* - \widehat{\Theta}\|_2^2 \leq \frac{1}{p} \|\Theta^* - \widehat{\Theta}\|_{2, \mathbf{\Pi}}^2$ and we get the following result.

Corollary 2. *Assume that \mathbf{A} is drawn according to (1), that $\min_{1 \leq i < j \leq n} \{\mathbf{\Pi}_{ij}\} \geq p$ and that $\rho_n = \omega(n^{-1})$. Then, there exists absolute constants $C, C' > 0$ such that with probability at least $1 - 9 \exp(-C\rho_n (k^2 + n \log(k)))$*

$$\|\Theta^* - \widehat{\Theta}\|_2^2 \leq C' \frac{\rho_n}{p} \left(\mathcal{K}_{\mathbf{\Pi}}(\Theta^*, \widetilde{\Theta}) + \frac{\rho_n^2}{((1 - \rho_n)^2 \wedge \gamma_n^2)} (k^2 + n \log(k)) \right).$$

Previously, the problem of estimation of connections probabilities matrix Θ^* from partial observations of the network was studied, in particular, by Gao et al. [20]. In this paper, the authors assume that any entry of the adjacency matrix \mathbf{A} is observed independently from the others with the same probability p , which is assumed to be known. They establish the following lower bound on the risk of any estimator for the stochastic block model.

Theorem 2 (Gao et al., 2017). *Assume that \mathbf{A} is drawn according to (2), and that each edge is observed independently from the others with probability p . There exists universal constants $C, C' > 0$ such that*

$$\inf_{\widehat{\Theta}} \sup_{\Theta^* \in \mathcal{T}_k, \|\Theta^*\|_{\infty} \leq \rho_n} \mathbb{P} \left[\|\Theta^* - \widehat{\Theta}\|_2^2 \geq C \left(\frac{\rho_n (n \log(k) + k^2)}{p} \wedge \rho_n^2 n^2 \right) \right] > C'.$$

The authors of [20] study the following weighted least square estimator :

$$\widehat{\Theta} \in \arg \min \|\Theta\|_2^2 - \frac{2}{p} \sum_{i < j} \mathbf{X}_{ij} \Theta_{ij} \mathbf{A}_{ij}$$

where p is the probability of observing any entry of the adjacency matrix. They prove that this estimator is minimax optimal in the uniform sampling setting. However, this estimator cannot handle more realistic, non-uniform sampling schemes. Our estimator, on the other hand, works in the non-uniform sampling setting, is adaptative to the sampling design, and does not require information on the probability of observing the entries of the adjacency matrix \mathbf{A} . In the particular case when $\mathbf{X}_{ij} \sim \text{Bernoulli}(p)$ and $\rho_n = O(\gamma_n)$, Corollary 2 and the lower bound in Theorem 2 ensures that the maximum likelihood estimator is minimax optimal. We underline that although the lower bound has been established in [20] for $\Theta \in \mathcal{T}_k$, $\|\Theta\|_\infty \leq \rho_n$, its proof can be adapted to provide a lower bound on the convergence rate for a smaller set of parameters $\left\{ \Theta \in \mathcal{T}_k, \|\Theta\|_\infty \leq \rho_n, \min_{i < j} \{\Theta_{ij}\} \geq \gamma_n \right\}$. Indeed, the ‘‘non parametric’’ as well as the ‘‘clustering’’ components of the rate are established using matrices with entries close to $\frac{\rho_n}{2}$.

2.3 Choice of γ_n under general assumptions

In this section, we deal with the setting when condition $\min_{i < j} \Theta_{ij}^* > \gamma_n$ is violated. In what follows we consider the sparse case, that is $\rho_n \rightarrow 0$, so $\gamma_n \leq 1 - \rho_n$ for n large enough. As discussed in [30], we can easily estimate ρ_n (see also Section 2.4). On the other hand, when some entries of the matrix of connections probabilities Θ^* can be 0 or arbitrarily close to 0, choosing the best sequence γ_n comes down to a trade-off between errors caused by estimating entries smaller than γ_n by γ_n , and the bound obtained in Theorem 1. We first consider the case when there exists a sequence γ_n such that number of small entries $n_s = \sum_{i < j} \mathbb{1}\{\widehat{\Theta}_{ij} < \gamma_n\}$ is small enough. Then, we have the following result:

Corollary 3. *Assume that \mathbf{A} is drawn according to (1), that $\rho_n = \omega(n^{-1})$ and that $n_s \leq \frac{k^2 \vee (n \log(k))}{\rho_n}$. Then, there exists absolute constants $C, C' > 0$ such that with probability at least $1 - 9 \exp(-C \rho_n (k^2 + n \log(k)))$*

$$\|\Theta^* - \widehat{\Theta}\|_{2, \Pi}^2 \leq C' \rho_n \left(\mathcal{K}_\Pi(\Theta^*, \widetilde{\Theta}) + \frac{\rho_n^2}{\gamma_n^2} (k^2 + n \log(k)) \right).$$

To see it, we define

$$\begin{aligned} \widetilde{\Theta}_{ij}^s &= \mathbf{Q}_{z^*(i)z^*(j)}^s, \widetilde{\Theta}_{ii}^s = 0 \\ \mathbf{Q}_{ab}^s &= \mathbf{Q}_{ab}^* \vee \gamma_n \end{aligned} \tag{10}$$

where \mathbf{Q}^* is given by (6). Note that $\widetilde{\Theta}^s$ and $\widehat{\Theta}$ are defined on the same set, and thus $\mathcal{K}_\mathbf{X}(\mathbf{A}, \widehat{\Theta}) \leq \mathcal{K}_\mathbf{X}(\mathbf{A}, \widetilde{\Theta}^s)$. Adapting the proof of Theorem 1 gives

$$\begin{aligned} \|\Theta^* - \widehat{\Theta}\|_{2, \Pi}^2 &\leq C' \rho_n \left(\mathcal{K}_\Pi(\Theta^*, \widetilde{\Theta}^s) + \frac{\rho_n^2}{\gamma_n^2} (k^2 + n \log(k)) \right) \\ &\leq C' \rho_n \left(\mathcal{K}_\Pi(\Theta^*, \widetilde{\Theta}) + \mathcal{K}_\Pi(\Theta^*, \widetilde{\Theta}^s) - \mathcal{K}_\Pi(\Theta^*, \widetilde{\Theta}) + \frac{\rho_n^2}{\gamma_n^2} (k^2 + n \log(k)) \right) \\ &\leq C' \rho_n \left(\mathcal{K}_\Pi(\Theta^*, \widetilde{\Theta}) + 2\gamma_n n_s + \frac{\rho_n^2}{\gamma_n^2} (k^2 + n \log(k)) \right) \end{aligned} \tag{11}$$

where (11) follows from Lemma 20. Note that, if there exists a sequence γ_n such that $\rho_n = O(\gamma_n)$ and $n_s \leq k^2 \vee (n \log(k))/\rho_n$, the upper bound on the risk obtained in (11) matches the bound of Theorem 2 and is minimax optimal.

Without any assumption on the number of small entries of the matrix of connections probabilities, we choose $\gamma_n = \gamma(\rho_n) \triangleq n^{-\frac{2}{3}} \rho_n^{\frac{2}{3}} (k^2 + n \log(k))^{\frac{1}{3}}$ and obtain the following bound.

Corollary 4. *Assume that \mathbf{A} is drawn according to (1), and that $\rho_n = \omega(n^{-1})$. Let*

$$\begin{aligned} \widehat{\Theta}_{i < j} &= \widehat{\mathbf{Q}}_{\widehat{z}(i)\widehat{z}(j)}, \widehat{\Theta}_{ii} = 0 \\ (\widehat{\mathbf{Q}}, \widehat{z}) &\in \arg \min_{\mathbf{Q} \in [\gamma(\rho_n), \rho_n]_{\text{sym}}^{k \times k}, z \in \mathcal{Z}_{n, k}} \sum_{i < j} \mathbf{X}_{ij} \mathcal{K}(\mathbf{A}_{ij}, \mathbf{Q}_{z(i)z(j)}). \end{aligned}$$

There exists absolute constants $C, C' > 0$ such that with probability at least $1 - 9 \exp(-C \rho_n (k^2 + n \log(k)))$

$$\|\Theta^* - \widehat{\Theta}\|_{2, \Pi}^2 \leq C' \rho_n \left(\mathcal{K}_\Pi(\Theta^*, \widetilde{\Theta}) + \rho_n^{\frac{2}{3}} n^{\frac{4}{3}} (k^2 + n \log(k))^{\frac{1}{3}} \right).$$

If k is not too large, the rate of convergence is essentially multiplied by $(n\rho_n)^{\frac{2}{3}}$.

2.4 Choice of γ_n for sparse positive graphons

In Theorem 1 we have established an oracle bound for the maximum likelihood estimator with entries belonging to $[\gamma_n, \rho_n]$. Defining our estimator requires us to estimate the values of these two sparsity parameters, which are usually unknown. When the matrix of connections probabilities Θ^* is generated according to the sparse graphon model (4) where W^* is bounded away from 0, these bounds will be of the same order of magnitude and decrease as the expected node degree. Under this assumption, we can use \widehat{d} , the average number of edges, to estimate γ_n and ρ_n . Indeed, it is easy to see that, with probability close to 1, \widehat{d} is close to $d = \rho_n \int_0^1 \int_0^1 W^*(x, y) dx dy$, the expected node degree. Note that if the graphon W^* is Hölder continuous or is a step function, assuming that $W^* > 0$ is enough to ensure that there exists a constant $C_{inf} > 0$ such that $W^* \geq C_{inf}$.

To simplify the exposition, we will assume that we observe all the entries of \mathbf{A} . Our results can be extended to the missing observations scheme described in Section 2.1 under the assumption that the entries of the sampling probability matrix $\mathbf{\Pi}$ are bounded away from 0. Let Ω be a subset of $\{(i, j) \in [n]^2, i < j\}$ of size n sampled independently of \mathbf{A} , and let

$$\widehat{d} = \frac{1}{n} \sum_{(i,j) \in \Omega} \mathbf{A}_{ij}$$

$$\widehat{\rho}_n = (\log(n))^{\frac{1}{5}} \widehat{d}, \quad \widehat{\gamma}_n = (\log(n))^{-\frac{1}{5}} \widehat{d}.$$

We use $\widehat{\rho}_n$ and $\widehat{\gamma}_n$ to build the restricted maximum likelihood estimator of the matrix of connections probabilities based on the observations of \mathbf{A}_{ij} with $\{(i, j) \in [n]^2, i < j\} \setminus \Omega$:

$$\widehat{\Theta}_{i < j} = \widehat{Q}_{\widehat{z}(i)\widehat{z}(j)}, \quad \widehat{\Theta}_{ii} = 0$$

$$(\widehat{Q}, \widehat{z}) \in \arg \min_{\mathbf{Q} \in [\widehat{\gamma}_n, \widehat{\rho}_n]_{\text{sym}}^{k \times k}, z \in \mathcal{Z}_{n,k}, (i,j) \notin \Omega} \sum_{(i,j) \notin \Omega} \mathcal{K}(\mathbf{A}_{ij}, \mathbf{Q}_{z(i)z(j)}).$$

We prove the following upper bound on the risk of this adaptive estimator:

Theorem 3. *Assume that \mathbf{A} is drawn according to the sparse graphon model and that $C_{inf} \triangleq \inf_{(x,y) \in [0,1]^2} W^*(x, y) > 0$, $\rho_n = o(\log(n)^{\frac{-1}{5}})$ and $\rho_n = \omega(n^{-1})$. Then, there exists positive constants N, C, C' depending only on C_{inf} , such that, for $n \geq N$, with probability at least $1 - 7 \exp(-Cn\rho_n)$, we have*

$$\left\| \Theta^* - \widehat{\Theta} \right\|_2^2 \leq C' \rho_n \log(n) \left(\mathcal{K}(\Theta^*, \widetilde{\Theta}) + (k^2 + n \log(k)) \right).$$

In the sparse graphon model, if the graphon W^* is bounded away from 0 and $n^{-1} \ll \rho_n \ll \log(n)^{\frac{-1}{5}}$, our adaptive estimator is optimal in the minimax sense up to a log factor. This setting includes the sparse stochastic block model considered for example in [8]: in this model, the matrix of probabilities of connection between communities \mathbf{Q}^* is such that $\mathbf{Q}^* = \rho_n \mathbf{Q}^0$, for some fixed matrix \mathbf{Q}^0 with entries in $(0, 1]$ (in this case, the graphon function W^* is a step function bounded away from 0). More generally, Theorem 3 covers the case of $\Theta^* = \rho_n \Theta^0$, and Θ^0 is a fixed matrix with entries in $(0, 1]$.

When we can not assume that the graphon W^* is bounded away from 0, we can use the same trade-off as in (10) and choose $\widehat{\gamma}_n = \gamma(\widehat{\rho}_n)$. Then, with high probability, we obtain the following bound on the risk of the adaptive estimator:

$$\left\| \Theta^* - \widehat{\Theta} \right\|_2^2 \leq C' \rho_n \log(n) \left(\mathcal{K}(\Theta^*, \widetilde{\Theta}) + \left(\log(n)^{\frac{1}{5}} \rho_n \right)^{\frac{2}{3}} n^{\frac{4}{3}} \left(\frac{k^2}{n^2} + n \log(k) \right)^{\frac{1}{3}} \right).$$

2.5 Smooth graphons

We have established a non-asymptotic bound on the risk of the maximum likelihood estimator depending on the Kullback-Leibler divergence between Θ^* and its oracle approximation by a block constant matrix corresponding to a SBM with k communities. While studying the graphon model (4), two classes of graphons are of particular interest: step function graphons and Hölder continuous graphons [47, 30, 21, 54]. A graphon W is called a *step function* if there exists a partition $S_1 \cup \dots \cup S_k$ of $[0, 1]$ into measurable sets such that the graphon W is constant on any product set $S_a \times S_b$. For step function graphons, the model corresponds to the stochastic block model described in (2): in this

case, the oracle matrix $\tilde{\Theta}$ is equal to the matrix of connections probabilities Θ^* . Next, we bound the Kullback-Leibler divergence between Θ^* and its oracle approximation by a block constant matrix for Hölder continuous graphons. We also provide the optimal choice for the number of communities k for our estimator.

We consider graphons that are weakly isomorphic to a smooth function. More precisely, for any $\alpha > 0$ and $M > 0$, let $\mathcal{F}_\alpha(M)$ be the class of Hölder functions, defined as follows:

$$\mathcal{F}_\alpha(M) = \left\{ W : [0, 1]^2 \rightarrow [0, 1], \forall (x, y), (x', y') \in [0, 1]^2, \right. \\ \left. |W(x', y') - \mathcal{P}_{[\alpha]}((x, y), (x' - x, y' - y))| \leq M \left(|x - x'|^{\alpha - [\alpha]} + |y - y'|^{\alpha - [\alpha]} \right) \right\}$$

where $\mathcal{P}_{[\alpha]}((x, y), \cdot)$ is the Taylor polynomial of W of degree $[\alpha]$ at point (x, y) . In particular, if $W \in \mathcal{F}_\alpha(M)$, $\forall (x, y), (x', y') \in [0, 1]^2$,

$$|W(x', y') - W(x, y)| \leq M \left(|x - x'|^{\alpha \wedge 1} + |y - y'|^{\alpha \wedge 1} \right). \quad (12)$$

When the graphon is Hölder continuous, the following proposition provides an upper bound on the Kullback-Leibler divergence between Θ^* and $\tilde{\Theta}$.

Proposition 1. *Consider the sparse graphon model (4) with $W^* \in \mathcal{F}_\alpha(M)$ where $\alpha, M > 0$ and we assume that $C_{inf} \triangleq \inf_{(x, y) \in [0, 1]^2} W^*(x, y) > 0$, $\rho_n \leq 1 - C_{inf}$ and $\rho_n = \omega(n^{-1})$. Then, almost surely, there exists a k -block constant matrix Θ^{bc} such that*

$$\mathcal{K}(\Theta^*, \Theta^{bc}) \leq \frac{4n^2 \rho_n M^2}{C_{inf}(1 - \rho_n)} \left(\frac{1}{k} \right)^{2(\alpha \wedge 1)}. \quad (13)$$

Proposition 1 enables us to bound the bias of estimating Θ^* by an oracle SBM with k communities. On the other hand, the bound given in Theorem 1 can be considered as the variance term of a block constant estimator with k blocks. To optimize the trade-off between these two terms, we choose k as follows

$$k = \left\lceil n^{\frac{1}{1+(\alpha \wedge 1)}} \rho_n^{\frac{1}{2+2(\alpha \wedge 1)}} \right\rceil \quad (14)$$

and obtain the following result:

Theorem 4. *Consider the sparse graphon model (4) with $W^* \in \mathcal{F}_\alpha(M)$ where $\alpha, M > 0$ and we assume that $C_{inf} \triangleq \inf_{(x, y) \in [0, 1]^2} W^*(x, y) > 0$, $\rho_n \leq 1 - C_{inf}$ and that $\rho_n = \omega(n^{-1})$. Then, there exists constants $C, C' > 0$, depending only on M, α and C_{inf} , such that, the restricted maximum likelihood estimator defined by (6) constructed with k defined by (14) satisfies*

$$\left\| \Theta_{ij}^* - \hat{\Theta}_{ij} \right\|_2^2 \leq C \rho_n \left(n^{\frac{2}{1+(\alpha \wedge 1)}} \rho_n^{\frac{1}{1+(\alpha \wedge 1)}} + n \log(\rho_n n) \right)$$

with probability larger than $1 - 9 \exp(-C' \rho_n n \log(\rho_n n))$.

The bound obtained in Theorem 4 matches the minimax optimal rate established in [30] and proves that the maximum likelihood estimator is optimal for estimating the matrix of connections probabilities in graphon model for graphons W^* in the Hölder class.

3 Variational approximation to the maximum likelihood estimator

Due to the existence of local minima in the likelihood function and the subsequent necessary optimisation over the set of k^n labels, the maximum likelihood estimator defined in Equation (8) cannot be computed in polynomial time. A tractable variational approximation to this estimator has been introduced in [17] to study dense and fully-observed stochastic block models. It has been extended to sparse stochastic block models in [8]. More recently, the authors of [50] used variational methods to approximate the maximum likelihood estimator in networks with missing observations. This method estimates the matrix \mathbf{Q}^* of probabilities of connections between communities, and the proportions of the different communities, and can be used to estimate the community label function z^* and the matrix Θ^* of probabilities of connections between nodes.

Note that variational approximations to the maximum likelihood estimator have been introduced in a model closely related to the conditional stochastic block model described in Equation (2). In this model, the labels are assumed to be latent random variables drawn according to the following distribution

$$\forall 1 \leq i \leq n, z(i) \sim \mathcal{M}(\alpha^*)$$

where $\alpha^* = (\alpha_1^*, \dots, \alpha_k^*)$, and $\alpha_a^* n$ is the expected size of community a . Conditionally on $z = z^*$, the edges are then drawn according to the conditional stochastic block model described in Equation (2). Note that the stochastic block model with random labels is a special case of the graphon model described in Equation (3), in which the graphon W^* is a step function.

The likelihood of the observed adjacency matrix under a stochastic block model with random labels with parameters (α, \mathbf{Q}) is given by

$$l_{\mathbf{X}}(\mathbf{A}; \alpha, \mathbf{Q}) = \sum_{z \in \mathcal{Z}_{n,k}} \left(\prod_{i \leq n} \alpha_{z(i)} \right) \exp(\mathcal{L}_{\mathbf{X}}(\mathbf{A}; z, \mathbf{Q})).$$

The Expectation-Maximization (EM) algorithm cannot be used to maximise $l_{\mathbf{X}}$, as it would require to evaluate the expectation of the label function z for given parameters (α, \mathbf{Q}) by summing over k^n possible labels. To circumvent this problem, variational methods approximate the posterior distribution of z given the observed entries $\mathbf{X} \odot \mathbf{A}$ of the adjacency matrix \mathbf{A} , denoted $\mathbb{P}(\cdot | \mathbf{X} \odot \mathbf{A}, \alpha, \mathbf{Q})$, by a simpler distribution chosen so that the expectation step of the EM procedure becomes tractable. More precisely, the posterior distribution $\mathbb{P}(\cdot | \mathbf{X} \odot \mathbf{A}, \alpha, \mathbf{Q})$ is approximated by a multinomial distribution denoted \mathbb{P}_{τ} , such that $\mathbb{P}_{\tau}(z) = \prod_{1 \leq i \leq n} m(z | \tau^i)$, where $m(\cdot | \tau^i)$ is the density of the multinomial distribution with parameter $\tau^i = (\tau_1^i, \dots, \tau_k^i)$, and $\tau = (\tau^1, \dots, \tau^n)$. Then, the variational estimator is defined as

$$\begin{aligned} (\hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}, \hat{\tau}^{VAR}) &= \arg \max_{(\alpha, \mathbf{Q}) \in \mathcal{Q}, \tau \in \mathcal{T}} \mathcal{J}_{\mathbf{X}}(\mathbf{A}; \tau, \alpha, \mathbf{Q}) \\ \text{for } \mathcal{J}_{\mathbf{X}}(\mathbf{A}; \tau, \alpha, \mathbf{Q}) &= \log(l_{\mathbf{X}}(\mathbf{A}; \alpha, \mathbf{Q})) - KL(\mathbb{P}_{\tau}(\cdot) || \mathbb{P}(\cdot | \mathbf{X} \odot \mathbf{A}, \alpha, \mathbf{Q})) \end{aligned}$$

where \mathcal{Q} and \mathcal{T} are the respective parameter spaces for the parameters (α, \mathbf{Q}) and τ , and KL denotes the Kullback-Leibler divergence between two distribution. Note that $\exp(\mathcal{J}_{\mathbf{X}}(\mathbf{A}; \tau, \alpha, \mathbf{Q}))$ provides a lower bound on $l_{\mathbf{X}}(\mathbf{A}; \alpha, \mathbf{Q})$, as $KL(\mathbb{P}_{\tau}(\cdot) || \mathbb{P}(\cdot | \mathbf{X} \odot \mathbf{A}, \alpha, \mathbf{Q})) \geq 0$ for any parameter (α, \mathbf{Q}) . The variational estimator is obtained using the EM algorithm derived in [50], which alternates between the following two steps :

- Estimation Step : for given parameters (α, \mathbf{Q}) , the variational parameter τ maximizing $\mathcal{J}_{\mathbf{X}}(\mathbf{A}; \tau, \alpha, \mathbf{Q})$ is given by the following fixed point equation :

$$\tau_a^i = c_i \alpha_a \prod_{j \neq i: \mathbf{X}_{ij}=1} \prod_{b \leq k} \left(\mathbf{Q}_{ab}^{A_{ij}} (1 - \mathbf{Q}_{ab})^{1-A_{ij}} \right)^{\tau_b^j} \quad \text{where } c_i \text{ is a normalising constant;}$$

- Maximisation Step : for a given parameter τ , the parameters (α, \mathbf{Q}) maximizing $\mathcal{J}_{\mathbf{X}}(\mathbf{A}; \tau, \alpha, \mathbf{Q})$ are given by

$$\alpha_a = \frac{\sum_i \tau_a^i}{n}, \quad \mathbf{Q}_{ab} = \frac{\sum_{i \neq j} \mathbf{X}_{ij} \tau_a^i \tau_b^j A_{ij}}{\sum_{i \neq j} \mathbf{X}_{ij} \tau_a^i \tau_b^j}.$$

This algorithm is implemented in the package `missSBM`. Since the EM algorithm is not guaranteed to converge to a global maximum, a clever initialization of the parameter τ using a clustering step is performed.

Theoretical properties of the variational estimator have been first studied in [17] for the dense stochastic block model, in [8] for the sparse stochastic block model, and more recently in [40] for a dense stochastic block model with observations missing uniformly at random. These results establish that maximizing $\max_{\tau \in \mathcal{T}} \mathcal{J}_{\mathbf{X}}(\mathbf{A}; \tau, \alpha, \mathbf{Q})$ is equivalent to maximizing $l_{\mathbf{X}}(\mathbf{A}; \alpha, \mathbf{Q})$. They also establish that the estimator obtained by maximizing $l_{\mathbf{X}}(\mathbf{A}; \alpha, \mathbf{Q})$ converges to the true parameters (α^*, \mathbf{Q}^*) , and conclude that $(\hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})$ also converges to (α^*, \mathbf{Q}^*) .

We build on this approximation and use the following estimator for the label function z^* :

$$\forall i \leq n, \hat{z}^{VAR}(i) \triangleq \arg \max_{a \leq k} \left(\hat{\tau}^{VAR} \right)_a^i$$

Next, we compare the asymptotic behavior of \hat{z}^{VAR} with that of the profile maximum likelihood estimator of the label function \hat{z} , defined as

$$(\hat{\mathbf{Q}}, \hat{z}) \in \arg \min_{\mathbf{Q} \in \mathcal{Q}, z \in \mathcal{Z}_{n,k}} \sum_{i < j} \mathbf{X}_{ij} \mathcal{K}(\mathbf{A}_{ij}, \mathbf{Q}_{z(i)z(j)})$$

where we use the convention that $\mathbf{Q} \in \mathcal{Q}$ if there exists α such that $(\alpha, \mathbf{Q}) \in \mathcal{Q}$.

In Proposition 2, we establish the asymptotic equivalence of \hat{z}^{VAR} and of \hat{z} in the case of sparse stochastic block model with full observations. To prove it, we will assume that the proportions of different communities are held constant, while the probabilities of connections between communities may decrease at rate ρ_n . That is, the parameters (α^*, \mathbf{Q}^*) verify

A1 $\alpha^* = \alpha^0$ for some fixed α^0 such that $\alpha_a^0 > 0$ for any $a \in \{1, \dots, k\}$

A2 $\mathbf{Q}^* = \rho_n \mathbf{Q}^0$ for some fixed $\mathbf{Q}^0 \in (0, 1)^{k \times k}$ such that $\sum_{a,b=1}^k \alpha_a^0 \alpha_b^0 \mathbf{Q}_{ab}^0 = 1$

The normalisation constraint $\sum_{a,b=1}^k \alpha_a^0 \alpha_b^0 \mathbf{Q}_{ab}^0 = 1$ ensure the identifiability of the parameters (\mathbf{Q}^0, ρ_n) (see [8]). In the following, we denote by \mathcal{Q} the set of parameters (α, \mathbf{Q}) verifying Assumptions A1 and A2. Note that the stochastic block model is identifiable up to a simultaneous permutation of communities and of rows and columns of the parameters \mathbf{Q}^* . For any two label functions z, z' , we write $z \sim z'$ if there exists a permutation σ of $\{1, \dots, k\}$ such that $(z(\sigma(a)))_{a \leq k} = (z(a))_{a \leq k}$.

Proposition 2. *Assume that (z^*, \mathbf{A}) is generated from a stochastic block model with parameters $(\alpha^*, \mathbf{Q}^*) \in \mathcal{Q}$ such that \mathbf{Q}^0 has no identical columns and the sparsity inducing sequence ρ_n satisfies $\rho_n \gg \log(n)/n$. Then, under Assumptions A1 and A2, $\mathbb{P}(\widehat{z}^{VAR} \sim \widehat{z}) \rightarrow 1$ when $n \rightarrow \infty$.*

The proof of Proposition 2 relies on results established in [8] and [9], and is given in Section 5.9. These results have been extended in [40] to the case of dense stochastic block model with incomplete observation of the network, where entries of the adjacency matrix are observed uniformly at random. Building on their results, Proposition 2 can be extended to the case of missing observations.

Once we have estimated the community labels, we can derive an estimator of the matrix of connection probabilities. For all $(a, b) \in \{1, \dots, k\}$, define

$$n_{ab}(\widehat{z}^{VAR}) = \begin{cases} |(\widehat{z}^{VAR})^{-1}(a)| \times |(\widehat{z}^{VAR})^{-1}(b)| & \text{if } a \neq b \\ |(\widehat{z}^{VAR})^{-1}(a)| \times (|(\widehat{z}^{VAR})^{-1}(a)| - 1) & \text{otherwise} \end{cases}$$

and

$$\widehat{\mathbf{Q}}_{ab}^{ML-VAR} \triangleq \frac{\sum_{i \in (\widehat{z}^{VAR})^{-1}(a), j \in (\widehat{z}^{VAR})^{-1}(b)} \mathbf{A}_{ij}}{n_{ab}(\widehat{z}^{VAR})}.$$

Proposition 2 implies that with large probability, there exists a permutation σ of $\{1, \dots, k\}$ such that $(\widehat{z}^{VAR}(\sigma(a)))_{a \leq k} = (\widehat{z}(a))_{a \leq k}$. When this hold, $(\widehat{\mathbf{Q}}_{\sigma(a), \sigma(b)}^{ML-VAR})_{a,b \leq k} = (\widehat{\mathbf{Q}}_{a,b})_{a,b \leq k}$, and the risk bounds obtained in Theorem 1 for the estimator $\widehat{\Theta}$ hold for the tractable estimator $\widehat{\Theta}^{VAR}$, defined as

$$\widehat{\Theta}_{i \neq j}^{VAR} = \widehat{\mathbf{Q}}_{\widehat{z}^{VAR}(i), \widehat{z}^{VAR}(j)}^{ML-VAR}, \quad \widehat{\Theta}_{ii}^{VAR} = 0.$$

Proposition 3. *Assume that (z^*, \mathbf{A}) is generated from a stochastic block model with parameters $(\alpha^*, \mathbf{Q}^*) \in \mathcal{Q}$ such that \mathbf{Q}^0 has no identical columns and the sparsity inducing sequence ρ_n satisfies $\rho_n \gg \log(n)/n$. Then, under Assumptions A1 and A2, there exists a constant $C_{\mathbf{Q}^0} > 0$ depending on \mathbf{Q}^0 such that*

$$\mathbb{P}\left(\left\|\Theta^* - \widehat{\Theta}^{VAR}\right\|_2^2 \leq C_{\mathbf{Q}^0} \rho_n (k^2 + n \log(k))\right) \rightarrow 1$$

when $n \rightarrow \infty$.

4 Conclusion

We have studied the problem of estimating the matrix of connections probabilities for the inhomogeneous random graph model and the graphon model with missing link. We have established a non-asymptotic bound on the risk of the maximum likelihood estimator. In particular, we have shown that, if the entries of the probability matrix decrease at the same rate, our estimator achieves the minimax convergence rate. This result is adaptative to the unknown sampling scheme. Moreover, we show that our estimator can be efficiently approximated using variational methods, and thus used in practice.

Acknowledgments

The work of O. Klopp was conducted as part of the project Labex MME-DII (ANR11-LBX-0023-01). The authors want to thank Catherine Matias and Nicolas Verzelen for extremely valuable suggestions and discussions.

The authors declare that there is no conflict of interest.

5 Proofs

The proof of Theorem 1 requires bounding the domain of definition of our estimator away from 0 and 1 in order to ensure that the loss function associated with the maximum likelihood estimator is Lipschitz. The Lipschitz constant here is equal to $\frac{1}{(1-\rho_n)\wedge\gamma_n}$. We balance this term by ρ_n by taking advantage of the sparsity of the graph, which implies, in particular, the low variance of \mathbf{A} . For ease of notations, we will assume that $1-\rho_n \geq \gamma_n$. This is the case when the graph is sparse, and our results still hold in the dense case if we replace γ_n by $\gamma_n \wedge (1-\rho_n)$ in our bounds.

5.1 Proof of Theorem 1

In previous works on sparse network estimation, the authors used the definition of the least square estimator to bound the norm of the error of this estimator by the scalar product between this error and a noise term. The main difficulty in their proofs consists in bounding this scalar product. Unfortunately, the definition of the maximum likelihood estimator does not yield such a bound, and these technics cannot be used to control its error. Instead, we make use of more refined peeling arguments, in order to take advantage of the Lipschitz continuity of the loss.

Let $\epsilon_n = C \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2)$ for some absolute constant C defined as the maximum of the absolute constants appearing in Lemma 6, Lemma 11 and Lemma 14, and let $\epsilon^0 \triangleq \rho_n \epsilon_n$. We start by considering the following two cases:

Case 1: $\|\tilde{\Theta} - \hat{\Theta}\|_{2,\Pi}^2 \leq 2\epsilon^0$. Then the statement of Theorem 1 follows from Lemma 19:

$$\|\Theta^* - \hat{\Theta}\|_{2,\Pi}^2 \leq 2\|\tilde{\Theta} - \hat{\Theta}\|_{2,\Pi}^2 + 2\|\Theta^* - \tilde{\Theta}\|_{2,\Pi}^2 \leq 4\rho_n \epsilon_n + 16\rho_n \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}).$$

Case 2: $\|\tilde{\Theta} - \hat{\Theta}\|_{2,\Pi}^2 > 2\epsilon^0$. Then $\hat{\Theta}$ belongs to the set

$$\mathcal{S}_\Pi = \left\{ \Theta \in \bigcup_{z \in \mathcal{Z}_{n,k}} \mathcal{T}_z : \|\tilde{\Theta} - \Theta\|_{2,\Pi}^2 \geq 2\epsilon^0, \|\Theta\|_\infty \leq \rho_n, \min_{i < j} \{\Theta_{ij}\} \geq \gamma_n \right\}$$

The following lemma helps us bound the error in $\|\cdot\|_\Pi$ -norm by the error in $\|\cdot\|_X$ -norm. Its proof departs from previous technics used to study network with missing observations. It allows for non-uniform sampling design, and is one of the main technical contribution of this work.

Lemma 1. *There exists an absolute constant $C > 0$ such that for all $\Theta \in \mathcal{S}_\Pi$ simultaneously we have*

$$\left| \|\Theta - \tilde{\Theta}\|_{2,\Pi}^2 - \|\Theta - \tilde{\Theta}\|_{2,X}^2 \right| \leq \frac{1}{2} \|\Theta - \tilde{\Theta}\|_{2,\Pi}^2$$

with probability greater than $1 - 2\exp(-Cn \log(k))$.

Lemma 1 implies that with large probability, $\hat{\Theta}$ belongs to the set \mathcal{S}_X defined as

$$\mathcal{S}_X = \left\{ \Theta \in \bigcup_{z \in \mathcal{Z}_{n,k}} \mathcal{T}_z : \|\tilde{\Theta} - \Theta\|_{2,X}^2 \geq \epsilon^0, \|\Theta\|_\infty \leq \rho_n, \min_{i < j} \{\Theta_{ij}\} \geq \gamma_n \right\}.$$

To bound $\|\tilde{\Theta} - \hat{\Theta}\|_{2,\Pi}^2$ when $\hat{\Theta} \in \mathcal{S}_X \cap \mathcal{S}_\Pi$, we introduce the following notation. For $\Theta, \Theta' \in (0,1)_{sym}^{n \times n}$ and $\mathbf{B}, \mathbf{C} \in [0,1]_{sym}^{n \times n}$ we set $\Delta \mathcal{K}_B^C(\Theta, \Theta') = \mathcal{K}_B(\mathbf{C}, \Theta) - \mathcal{K}_B(\mathbf{C}, \Theta')$. Using Lemma 19 we get

$$\begin{aligned} \|\Theta^* - \hat{\Theta}\|_{2,\Pi}^2 &\leq 8\rho_n \mathcal{K}_\Pi(\Theta^*, \hat{\Theta}) \\ &\leq 8\rho_n \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}) + 8\rho_n \Delta \mathcal{K}_\Pi^{\Theta^*}(\hat{\Theta}, \tilde{\Theta}). \end{aligned}$$

On the other hand, the definition of $\hat{\Theta}$ implies that $\Delta \mathcal{K}_X^A(\hat{\Theta}, \tilde{\Theta}) \leq 0$ so

$$\begin{aligned} \|\Theta^* - \hat{\Theta}\|_{2,\Pi}^2 &\leq 8\rho_n \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}) + 8\rho_n \Delta \mathcal{K}_\Pi^{\Theta^*}(\hat{\Theta}, \tilde{\Theta}) - 8\rho_n \Delta \mathcal{K}_X^A(\hat{\Theta}, \tilde{\Theta}) \\ &\leq 8\rho_n \left(\mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}) + \left(\Delta \mathcal{K}_\Pi^{\Theta^*}(\hat{\Theta}, \tilde{\Theta}) - \Delta \mathcal{K}_X^A(\hat{\Theta}, \tilde{\Theta}) \right) + \left(\Delta \mathcal{K}_X^{\Theta^*}(\hat{\Theta}, \tilde{\Theta}) - \Delta \mathcal{K}_X^A(\hat{\Theta}, \tilde{\Theta}) \right) \right). \end{aligned} \quad (15)$$

To bound the terms involved in equation (15), we control $\sup_{\Theta \in \mathcal{S}_\Pi} \left| \Delta \mathcal{K}_\Pi^{\Theta^*}(\Theta, \tilde{\Theta}) - \Delta \mathcal{K}_X^{\Theta^*}(\Theta, \tilde{\Theta}) \right|$ using the concentration of \mathbf{X} around its expectation Π , and we control $\sup_{\Theta \in \mathcal{S}_X} \left| \Delta \mathcal{K}_X^{\Theta^*}(\Theta, \tilde{\Theta}) - \Delta \mathcal{K}_X^A(\Theta, \tilde{\Theta}) \right|$ conditionally on \mathbf{X} using the concentration of \mathbf{A} around its expectation Θ^* .

Lemma 2. *There exists absolute constants $C, C' > 0$ such that for all $\Theta \in \mathcal{S}_\Pi$ simultaneously we have*

$$\left| \Delta \mathcal{K}_\Pi^{\Theta^*}(\Theta, \tilde{\Theta}) - \Delta \mathcal{K}_X^{\Theta^*}(\Theta, \tilde{\Theta}) \right| \leq \frac{1}{2 \times 32 \rho_n} \left\| \Theta - \tilde{\Theta} \right\|_{2, \Pi}^2 + C \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2)$$

with probability greater than $1 - 2 \exp(-C' \rho_n n \log(k))$.

Lemma 3. *There exists absolute constants $C, C' > 0$ such that conditionally on \mathbf{X} , for any $\Theta \in \mathcal{S}_X$ simultaneously we have*

$$\left| \Delta \mathcal{K}_X^{\Theta^*}(\Theta, \tilde{\Theta}) - \Delta \mathcal{K}_X^A(\Theta, \tilde{\Theta}) \right| \leq \frac{1}{4 \times 32 \rho_n} \left\| \Theta - \tilde{\Theta} \right\|_{2, X}^2 + C \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2)$$

with probability greater than $1 - 5 \exp(-C' \rho_n n \log(k))$.

Combining Lemma 1, Lemma 2, Lemma 3 and (15) yields that there exists two absolute constants $C, C' > 0$ such that with probability greater than $1 - 9 \exp(-C' \rho_n n \log(k))$

$$\begin{aligned} \left\| \Theta^* - \hat{\Theta} \right\|_{2, \Pi}^2 &\leq 8 \rho_n \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}) + 8 \rho_n \times \frac{1}{2 \times 32 \rho_n} \left\| \hat{\Theta} - \tilde{\Theta} \right\|_{2, \Pi}^2 + 8 \rho_n \times \frac{1}{4 \times 32 \rho_n} \left\| \hat{\Theta} - \tilde{\Theta} \right\|_{2, X}^2 \\ &\quad + C \rho_n \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2) \\ &\leq 8 \rho_n \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}) + \frac{1}{8} \left\| \hat{\Theta} - \tilde{\Theta} \right\|_{2, \Pi}^2 + \frac{1}{16} \times \frac{3}{2} \left\| \hat{\Theta} - \tilde{\Theta} \right\|_{2, \Pi}^2 + C \rho_n \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2) \\ &\leq 8 \rho_n \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}) + \frac{1}{2} \left\| \Theta^* - \tilde{\Theta} \right\|_{2, \Pi}^2 + \frac{1}{2} \left\| \Theta^* - \hat{\Theta} \right\|_{2, \Pi}^2 + C \rho_n \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2). \end{aligned} \quad (16)$$

Lemma 19 and (16) imply that there exists two absolute constants $C, C' > 0$ such that with probability larger than $1 - 9 \exp(-C' \rho_n n \log(k))$,

$$\frac{1}{2} \left\| \Theta^* - \hat{\Theta} \right\|_{2, \Pi}^2 \leq 8 \rho_n \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}) + \frac{1}{2} \times 8 \rho_n \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}) + C \rho_n \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2).$$

This completes the proof of Theorem 1.

5.2 Proof of Lemma 1

To prove Lemma 1, we show that the probability of the following “bad” event is small:

$$\mathcal{E} \triangleq \left\{ \exists \Theta \in \mathcal{S}_\Pi : \left| \left\| \Theta - \tilde{\Theta} \right\|_{2, \Pi}^2 - \left\| \Theta - \tilde{\Theta} \right\|_{2, X}^2 \right| > \frac{1}{2} \left\| \Theta - \tilde{\Theta} \right\|_{2, \Pi}^2 \right\}.$$

We use a standard peeling argument (see, e.g., [28]): we slice \mathcal{S}_Π in different sets, on which we control $\left\| \Theta - \tilde{\Theta} \right\|_{2, \Pi}^2$.

Recall that $\epsilon_n \triangleq C \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2)$ where the absolute constant C is larger than the constant appearing in Lemma 6, and that $\epsilon^0 \triangleq \rho_n \epsilon_n$. For $l \in \mathbb{N}^*$, we set

$$\mathcal{S}_{l, \Pi} \triangleq \left\{ \Theta \in \mathcal{S}_\Pi : 2^{l-1} (2\epsilon^0) \leq \left\| \Theta - \tilde{\Theta} \right\|_{2, \Pi}^2 \leq 2^l (2\epsilon^0) \right\}.$$

If the event \mathcal{E} holds, there exists $l \in \mathbb{N}^*$ such that $\Theta \in \mathcal{S}_{l, \Pi}$ and

$$\left| \left\| \Theta - \tilde{\Theta} \right\|_{2, \Pi}^2 - \left\| \Theta - \tilde{\Theta} \right\|_{2, X}^2 \right| > \frac{1}{2} \left\| \Theta - \tilde{\Theta} \right\|_{2, \Pi}^2.$$

Note that $\mathbb{E} \left[\left\| \Theta - \tilde{\Theta} \right\|_{2, X}^2 \right] = \left\| \Theta - \tilde{\Theta} \right\|_{2, \Pi}^2$. The events that we need to control are the following:

$$\mathcal{E}_l \triangleq \left\{ \exists \Theta \in \mathcal{S}_{l, \Pi} : \left| \left\| \Theta - \tilde{\Theta} \right\|_{2, X}^2 - \mathbb{E} \left[\left\| \Theta - \tilde{\Theta} \right\|_{2, X}^2 \right] \right| > \frac{2^{l-1} (2\epsilon^0)}{2} \right\}.$$

If \mathcal{E} holds for some $\Theta \in \mathcal{S}_\Pi$, there exists $l \in \mathbb{N}^*$ such that $\Theta \in \mathcal{S}_{l, \Pi}$, thus there exists $l \in \mathbb{N}^*$ such that \mathcal{E}_l holds, i.e., $\mathcal{E} \subset \bigcup_{l \in \mathbb{N}^*} \mathcal{E}_{l, \Pi}$. For $T > 0$, let $\mathcal{S}_\Pi(T)$ be defined as follows:

$$\mathcal{S}_\Pi(T) = \left\{ \Theta \in \bigcup_{z \in \mathcal{Z}_{n, k}} \mathcal{T}_z : \left\| \Theta \right\|_\infty \leq \rho_n, \min_{i < j} \{\Theta_{ij}\} \geq \gamma_n, \left\| \Theta - \tilde{\Theta} \right\|_{2, \Pi}^2 \leq T \right\}.$$

We see that $\mathcal{S}_{l,\Pi} \subset \mathcal{S}_{\Pi}(2^l \epsilon^0)$, so we only need to control the probability of the events

$$\mathcal{E}(T) = \left\{ \exists \Theta \in \mathcal{S}_{\Pi}(T) : \left| \left\| \Theta - \tilde{\Theta} \right\|_{2,\mathbf{X}}^2 - \mathbb{E} \left[\left\| \Theta - \tilde{\Theta} \right\|_{2,\mathbf{X}}^2 \right] \right| > \frac{T}{4} \right\}.$$

The following lemma helps us bound the probability of the events $\mathcal{E}(T)$.

Lemma 4. *For $T > \epsilon^0$, let $Z_T = \sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \left\| \Theta - \tilde{\Theta} \right\|_{2,\Pi}^2 - \left\| \Theta - \tilde{\Theta} \right\|_{2,\mathbf{X}}^2 \right|$. There exists an absolute constant $C > 0$ such that*

$$\mathbb{P} \left(Z_T \geq \frac{T}{4} \right) \leq \exp \left(-\frac{CT}{\rho_n} \right).$$

Proof. To prove Lemma 4, we first show that Z_T concentrates around its expectation and then bound this term.

Lemma 5. *Let Z_T be defined as in 4. Then*

$$\mathbb{P} \left(Z_T > 2\mathbb{E}[Z_T] + \frac{T}{16} \right) \leq \exp \left(-\frac{T}{64\rho_n} \right).$$

Lemma 6. *Let Z_T be as in Lemma 4, then there exists an absolute constant $C > 0$ such that*

$$\mathbb{E}[Z_T] \leq \frac{T}{16} + C\rho_n \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2).$$

Putting together Lemma 5 and Lemma 6, we get that there exists an absolute constant $C > 0$ such that

$$\mathbb{P} \left(Z_T \geq \frac{3T}{16} + \frac{C}{8} \rho_n \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2) \right) \leq \exp \left(-\frac{T}{64\rho_n} \right).$$

Our choice of ϵ_0 allows us to conclude that for $T \geq 2\epsilon_0$, $\frac{C}{8} \rho_n \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2) \leq \frac{T}{16}$ and

$$\mathbb{P} \left(Z_T \geq \frac{T}{4} \right) \leq \exp \left(-\frac{T}{64\rho_n} \right).$$

□

For this choice of ϵ_0 ,

$$\begin{aligned} \mathbb{P}(\mathcal{E}) &\leq \sum_{l=1}^{\infty} \mathbb{P} \left(\mathcal{E}(2^l(2\epsilon^0)) \right) \\ &\leq \sum_{l=1}^{\infty} \exp \left(-2C\epsilon^0 2^l / \rho_n \right) \\ &\leq \sum_{l=1}^{\infty} \exp \left(-2Cl \log(2)\epsilon^0 / \rho_n \right) \\ &\leq \frac{\exp \left(-2C \log(2)\epsilon^0 / \rho_n \right)}{1 - \exp \left(-2C \log(2)\epsilon^0 / \rho_n \right)} = \frac{1}{\exp \left(-2C \log(2)\epsilon^0 / \rho_n \right) - 1} \leq 2 \exp \left(-Cn \log(k) \right) \end{aligned}$$

for n large enough. This completes the proof of Lemma 1.

5.2.1 Proof of Lemma 5

To control the deviation of Z_T from its expectation, we apply the following theorem from Bousquet, as stated in [22], Theorem 3.3.16.

Theorem 5 (Bousquet). *Let X_i , $i \in \mathbb{N}$ be independent \mathcal{S} -valued random variables, and let \mathcal{F} be a countable class of functions $f = (f_1, \dots, f_n) : \mathcal{S} \rightarrow [-1, 1]^n$ such that $\mathbb{E}[f_i(X_i)] = 0$ for any $f \in \mathcal{F}$ and $i \in [n]$. Set $Z = \sup_{f \in \mathcal{F}} \left| \sum_{1 \leq i \leq n} f_i(X_i) \right|$ and $v = \sup_{f \in \mathcal{F}} \sum_{1 \leq i \leq n} \mathbb{E}[f_i(X_i)^2]$. Then, for any $x > 0$,*

$$\mathbb{P} \left(Z > \mathbb{E}[Z] + \frac{x}{3} + \sqrt{2x(2\mathbb{E}[Z] + v)} \right) \leq \exp(-x).$$

We apply Theorem 5 to the random variable

$$\begin{aligned}
Z_T &= \sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \left\| \Theta - \tilde{\Theta} \right\|_{2, \Pi}^2 - \left\| \Theta - \tilde{\Theta} \right\|_{2, \mathbf{X}}^2 \right| \\
&= \sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} (\Pi_{ij} - \mathbf{X}_{ij}) (\Theta_{ij} - \tilde{\Theta}_{ij})^2 \right| \\
&= \rho_n \sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} f_{ij}^{\Theta}(\mathbf{X}_{ij}) \right|
\end{aligned}$$

where we set $f_{ij}^{\Theta}(\mathbf{X}_{ij}) \triangleq \frac{(\mathbf{X}_{ij} - \Pi_{ij})(\Theta_{ij} - \tilde{\Theta}_{ij})^2}{\rho_n}$. The set of functions $\{f_{ij}^{\Theta}, \Theta \in \mathcal{S}_{\Pi}(T)\}$ is separable and we can apply Theorem 5 (see, e.g., [22], Section 2.1). Note that for any $1 \leq i < j \leq n$, $\mathbb{E}[f_{ij}^{\Theta}(\mathbf{X}_{ij})] = 0$, $|f_{ij}^{\Theta}(\mathbf{X}_{ij})| \leq 1$, $\mathbb{E}[(\mathbf{X}_{ij} - \Pi_{ij})^2] \leq \Pi_{ij}$ and $|\Theta_{ij} - \tilde{\Theta}_{ij}| \leq \rho_n$ so

$$\begin{aligned}
\sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \sum_{1 \leq i < j \leq n} \mathbb{E}[f_{ij}^{\Theta}(\mathbf{X}_{ij})^2] &\leq \frac{1}{\rho_n^2} \sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \sum_{1 \leq i < j \leq n} \Pi_{ij} (\Theta_{ij} - \tilde{\Theta}_{ij})^4 \\
&\leq \sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \sum_{1 \leq i < j \leq n} \Pi_{ij} (\Theta_{ij} - \tilde{\Theta}_{ij})^2 \\
&\leq T.
\end{aligned}$$

Theorem 5 implies that

$$\begin{aligned}
\mathbb{P}\left(\frac{Z_T}{\rho_n} > \frac{1}{\rho_n} \mathbb{E}[Z_T] + \frac{x}{3} + \sqrt{2x \left(\frac{2}{\rho_n} \mathbb{E}[Z_T] + T\right)}\right) &\leq \exp(-x) \\
\mathbb{P}\left(Z_T > \mathbb{E}[Z_T] + \frac{\rho_n x}{3} + \sqrt{2x \rho_n (2\mathbb{E}[Z_T] + \rho_n T)}\right) &\leq \exp(-x) \\
\mathbb{P}\left(Z_T > \mathbb{E}[Z_T] + \frac{\rho_n x}{3} + 2\rho_n x + \mathbb{E}[Z_T] + \rho_n \sqrt{2xT}\right) &\leq \exp(-x)
\end{aligned}$$

where we have used $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ and $2\sqrt{ab} \leq a+b$. Setting $x = \frac{T}{64\rho_n}$ and noticing that $\rho_n \leq 1$ leads to

$$\mathbb{P}\left(Z_T > 2\mathbb{E}[Z_T] + \frac{T}{16}\right) \leq \exp\left(-\frac{T}{64\rho_n}\right).$$

5.2.2 Proof of Lemma 6

Once we have bounded Z_T by its expectation, we bound $\mathbb{E}[Z_T]$. To do so, we use a symmetrization argument and Talagrand's contraction principle (see, e.g., [22] for a proof).

Lemma 7 (Symmetrization). *Let $\{\mathbf{Y}_i\}_{1 \leq i \leq n}$ be independent random variables, $\{\epsilon_i\}_{1 \leq i \leq n}$ be a Rademacher sequence, and \mathcal{A} be a subset of \mathbb{R}^n , then*

$$\mathbb{E} \left[\sup_{\mathbf{A} \in \mathcal{A}} \left| \sum_{1 \leq i \leq n} (\mathbf{Y}_i - \mathbb{E}[\mathbf{Y}_i]) \mathbf{A}_i \right| \right] \leq 2\mathbb{E} \left[\sup_{\mathbf{A} \in \mathcal{A}} \left| \sum_{1 \leq i \leq n} \epsilon_i \mathbf{Y}_i \mathbf{A}_i \right| \right].$$

Lemma 8 (Talagrand's contraction principle). *Let $\{\phi_i\}_{1 \leq i \leq n} : \mathbb{R} \rightarrow \mathbb{R}$ be 1-Lipshitz functions vanishing at 0, \mathcal{A} be a compact subset of \mathbb{R}^n and $\{\epsilon_i\}_{1 \leq i \leq n}$ be a Rademacher sequence, then*

$$\mathbb{E} \left[\sup_{\Theta \in \mathcal{A}} \left| \sum_{1 \leq i \leq n} \epsilon_i \phi_i(\Theta_i) \right| \right] \leq 2\mathbb{E} \left[\sup_{\Theta \in \mathcal{A}} \left| \sum_{1 \leq i \leq n} \epsilon_i \Theta_i \right| \right].$$

Recall that

$$\mathbb{E}[Z_T] = \mathbb{E} \left[\sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} (\Pi_{ij} - \mathbf{X}_{ij}) (\Theta_{ij} - \tilde{\Theta}_{ij})^2 \right| \right].$$

Let $(\epsilon_{ij})_{1 \leq i < j \leq n}$ be a Rademacher sequence. Lemma 7 implies

$$\mathbb{E}[Z_T] \leq 2\mathbb{E} \left[\sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} \epsilon_{ij} \mathbf{X}_{ij} \left(\Theta_{ij} - \tilde{\Theta}_{ij} \right)^2 \right| \right].$$

For any $1 \leq i < j \leq n$, let $\phi_{ij} : x \rightarrow \frac{x^2}{2\rho_n}$. Note that on $[-\rho_n, \rho_n]$, ϕ_{ij} is a 1-Lipschitz and vanishes at 0. Applying Lemma 8, we get that

$$\begin{aligned} \mathbb{E}[Z_T] &\leq 4\rho_n \mathbb{E} \left[\sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} \epsilon_{ij} \phi_{ij} \left(\mathbf{X}_{ij} \left(\Theta_{ij} - \tilde{\Theta}_{ij} \right) \right) \right| \right] \\ &\leq 8\rho_n \mathbb{E} \left[\sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} \epsilon_{ij} \mathbf{X}_{ij} \left(\Theta_{ij} - \tilde{\Theta}_{ij} \right) \right| \right]. \end{aligned} \quad (17)$$

We bound the term in using the following lemma.

Lemma 9. *Let $\mathbf{B} \in \{\Pi, \mathbf{X}\}$ and let Σ be a random matrix such that almost surely, $\|\Sigma\|_{\infty} \leq 1$ and that conditionally on \mathbf{B} , for any $1 \leq i < j \leq n$ the coefficients Σ_{ij} are independent and centered. Assume that there exists $\alpha > 0$ such that for any $\Theta \in \mathbb{R}_{sym}^{n \times n}$, $\sum_{i < j} \mathbb{E}^{\mathbf{B}} [\Sigma_{ij}^2 \Theta_{ij}^2] \leq \alpha \|\Theta\|_{2, \mathbf{B}}^2$. There exists an absolute constant C such that*

$$\mathbb{E}^{\mathbf{B}} \left[\sup_{\Theta \in \mathcal{S}_{\mathbf{B}}(T)} \left| \sum_{1 \leq i < j \leq n} \Sigma_{ij} \left(\Theta_{ij} - \tilde{\Theta}_{ij} \right) \right| \right] \leq \frac{\gamma_n \alpha T}{32 \times 64^2 \rho_n^2} + C \frac{\rho_n^2}{\gamma_n} (n \log(k) + k^2).$$

Note that for any $1 \leq i < j \leq n$, $\mathbb{E}[\mathbf{X}_{ij}^2] \leq \Pi_{ij}$, so for any $\Theta \in \mathbb{R}_{sym}^{n \times n}$, $\sum_{i < j} \mathbb{E}[\epsilon_{ij}^2 \mathbf{X}_{ij}^2 \Theta_{ij}^2] \leq \|\Theta\|_{2, \Pi}^2$. We apply Lemma 9 with $\mathbf{B} = \Pi$, $\alpha = 1$ and for any $1 \leq i < j \leq n$, $\Sigma_{ij} = \epsilon_{ij} \mathbf{X}_{ij}$ and combine it with (17) to get that for some absolute constant C

$$\begin{aligned} \mathbb{E}[Z_T] &\leq \frac{T \gamma_n \times 8\rho_n}{32 \times 64^2 \rho_n^2} + C \rho_n \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2) \\ &\leq \frac{T}{4 \times 64} + C \rho_n \frac{\rho_n^2}{\gamma_n} (n \log(k) + k^2). \end{aligned}$$

This concludes the proof of Lemma 6.

5.2.3 Proof of Lemma 9

To get an upper bound on $\mathbb{E}^{\mathbf{B}} \left[\sup_{\Theta \in \mathcal{S}_{\mathbf{B}}(T)} \left| \sum_{1 \leq i < j \leq n} \Sigma_{ij} \left(\Theta_{ij} - \tilde{\Theta}_{ij} \right) \right| \right]$, we use Bernstein's inequality, which we state here for the reader's convenience:

Theorem 6 (Bernstein's inequality). *Let X_1, \dots, X_n be independent centered random variables. Assume that for any $i \in [n]$, $|X_i| \leq M$ almost surely, then*

$$\mathbb{P} \left(\left| \sum_{1 \leq i \leq n} X_i \right| \geq \sqrt{2t \sum_{1 \leq i \leq n} \mathbb{E}[X_i^2]} + \frac{2M}{3} t \right) \leq 2e^{-t}.$$

Recall that for $\mathbf{B} \in \{\Pi, \mathbf{X}\}$,

$$\mathcal{S}_{\mathbf{B}}(T) = \left\{ \Theta \in \bigcup_{z \in \mathcal{Z}_{n,k}} \mathcal{T}_z : \|\Theta\|_{\infty} \leq \rho_n, \min_{i < j} \{\Theta_{ij}\} \geq \gamma_n, \|\Theta - \tilde{\Theta}\|_{2, \mathbf{B}}^2 \leq T \right\}$$

and let $\mathcal{S}_z(T) \triangleq \mathcal{T}_z \cap \mathcal{S}_{\mathbf{B}}(T)$ be the set of matrices in $\mathcal{S}_{\mathbf{B}}(T)$ that are block constant for the label z . Let $\tilde{\Theta}^z$ be the projection of $\tilde{\Theta}$ onto \mathcal{T}_z for the \mathbf{B} -weighted Frobenius norm:

$$\tilde{\Theta}^z \triangleq \arg \min_{\Theta \in \mathcal{T}_z} \|\Theta - \tilde{\Theta}\|_{2, \mathbf{B}}.$$

Note that if $\mathcal{S}_z(T) \neq \emptyset$, then $\tilde{\Theta}^z \in \mathcal{S}_z(T)$. If $\mathcal{S}_z(T) = \emptyset$, we set $\sup_{\Theta \in \mathcal{S}_z(T)} \left| \langle \Sigma | \tilde{\Theta}^z - \Theta \rangle \right| = 0$. We decompose the error in two terms.

$$\begin{aligned} \mathbb{E}^{\mathcal{B}} \left[\sup_{\Theta \in \mathcal{S}_{\mathcal{B}}(T)} \left| \sum_{1 \leq i < j \leq n} \Sigma_{ij} (\Theta_{ij} - \tilde{\Theta}_{ij}) \right| \right] &\leq \mathbb{E}^{\mathcal{B}} \left[\sup_{z \in \mathcal{Z}_{n,k}, \mathcal{S}_z(T) \neq \emptyset} \left| \langle \Sigma | \tilde{\Theta} - \tilde{\Theta}^z \rangle \right| \right] \\ &\quad + \mathbb{E}^{\mathcal{B}} \left[\sup_{z \in \mathcal{Z}_{n,k}} \sup_{\Theta \in \mathcal{S}_z(T)} \left| \langle \Sigma | \tilde{\Theta}^z - \Theta \rangle \right| \right] \\ &\leq (I) + (II). \end{aligned} \tag{18}$$

The term (I) denotes $\mathbb{E}^{\mathcal{B}} \left[\sup_{z \in \mathcal{Z}_{n,k}, \mathcal{S}_z(T) \neq \emptyset} \left| \langle \Sigma | \tilde{\Theta} - \tilde{\Theta}^z \rangle \right| \right]$ and corresponds to the error induced by an error on the label. The term (II) denotes $\mathbb{E}^{\mathcal{B}} \left[\sup_{z \in \mathcal{Z}_{n,k}} \sup_{\Theta \in \mathcal{S}_z(T)} \left| \langle \Sigma | \tilde{\Theta}^z - \Theta \rangle \right| \right]$ and corresponds to the error induced by a Bernoulli noise.

Control of (I): To control the first term of (18), recall that for any $z \in \mathcal{Z}_{n,k}$ such that $\mathcal{S}_z(T) \neq \emptyset$, $\tilde{\Theta}^z \in \mathcal{S}_z(T)$ and by hypothesis, $\sum_{i < j} \mathbb{E}^{\mathcal{B}} \left[\Sigma_{ij}^2 (\tilde{\Theta}_{ij} - \tilde{\Theta}_{ij}^z)^2 \right] \leq \alpha \left\| \tilde{\Theta} - \tilde{\Theta}^z \right\|_{2,\mathcal{B}}^2 \leq \alpha T$. Furthermore, $\|\Sigma\|_{\infty} \leq 1$ so $\left| \Sigma_{ij} (\tilde{\Theta}_{ij} - \tilde{\Theta}_{ij}^z) \right| \leq \rho_n$. Since $|\mathcal{Z}_{n,k}| \leq n \log(k)$, the union bound and Bernstein's inequality imply

$$\begin{aligned} \mathbb{P}^{\mathcal{B}} \left(\sup_{z \in \mathcal{Z}_{n,k}, \mathcal{S}_z(T) \neq \emptyset} \left| \langle \Sigma | \tilde{\Theta} - \tilde{\Theta}^z \rangle \right| \geq \sqrt{2\alpha T(t + n \log(k))} + \frac{2\rho_n}{3}(t + n \log(k)) \right) &\leq 2e^{-t} \\ \mathbb{P}^{\mathcal{B}} \left(\sup_{z \in \mathcal{Z}_{n,k}, \mathcal{S}_z(T) \neq \emptyset} \left| \langle \Sigma | \tilde{\Theta} - \tilde{\Theta}^z \rangle \right| \geq \frac{\gamma_n \alpha T}{64^3 \rho_n^2} + \left(\frac{2\rho_n}{3} + \frac{64^3 \rho_n^2}{\gamma_n} \right) (t + n \log(k)) \right) &\leq 2e^{-t}. \end{aligned}$$

Integrating the last inequality and using $\frac{\rho_n}{\gamma_n} \geq 1$, we get that for some absolute constant C

$$(I) \leq \frac{\alpha \gamma_n T}{64^3 \rho_n^2} + C \frac{\rho_n^2}{\gamma_n} n \log(k). \tag{19}$$

Control of (II): The control of the second term of (18) is more involved. We adapt the argument developed in [30] and consider only $z \in \mathcal{Z}_{n,k}$ such that $\mathcal{S}_z(T) \neq \emptyset$. By property of the projection, we have for any $\Theta \in \mathcal{S}_z(T)$, $\left\| \Theta - \tilde{\Theta}^z \right\|_{2,\mathcal{B}}^2 \leq \left\| \Theta - \tilde{\Theta} \right\|_{2,\mathcal{B}}^2 \leq T$. Thus $(\tilde{\Theta}^z - \Theta) \in \mathcal{A}_z(T)$, where

$$\mathcal{A}_z(T) \triangleq \left\{ \Theta \in \mathcal{T}_z : \|\Theta\|_{\infty} \leq \rho_n, \|\Theta\|_{2,\mathcal{B}}^2 \leq T \right\},$$

so $\sup_{\Theta \in \mathcal{S}_z(T)} \left| \langle \Sigma | \tilde{\Theta}^z - \Theta \rangle \right| \leq \sup_{\Theta \in \mathcal{A}_z(T)} |\langle \Sigma | \Theta \rangle|$. Let $\hat{T}^z \in \mathcal{A}_z(T)$ be such that

$$\left| \langle \Sigma | \hat{T}^z \rangle \right| \triangleq \sup_{\Theta \in \mathcal{A}_z(T)} |\langle \Sigma | \Theta \rangle|. \tag{20}$$

Note that $\Theta \rightarrow |\langle \Sigma | \Theta \rangle|$ is continuous and reaches its supremum on $\mathcal{A}_z(T)$. Indeed, either for any $1 \leq i < j \leq n$, $\mathbf{B}_{ij} > 0$ so $\|\cdot\|_{\mathcal{B}}$ is a norm and $\mathcal{A}_z(T)$ is compact, or we can find a subspace \mathcal{V} of $\mathbb{R}^{\frac{(n-1)(n-2)}{2}}$ of dimension $|\{1 \leq i < j \leq n : \mathbf{B}_{ij} > 0\}|$ such that for any $\Theta \in \mathbb{R}^{\frac{(n-1)(n-2)}{2}}$, $\langle \Sigma | \Theta \rangle = \langle \Sigma | \mathcal{P}_{\mathcal{V}}(\Theta) \rangle$ where $\mathcal{P}_{\mathcal{V}}$ denotes the projection onto \mathcal{V} , and $\mathcal{A}_z(T) \cap \mathcal{V}$ is compact.

To control $\left| \langle \Sigma | \hat{T}^z \rangle \right|$, we build a finite set with small cardinality that approximates \hat{T}^z well both in the weighted Frobenius norm and in the supremum norm. More precisely, our goal is to construct a finite set $\tilde{\mathcal{C}}_z(T)$ containing a matrix $\hat{\mathbf{V}}$ such that $2(\hat{T}^z - \hat{\mathbf{V}}) \in \mathcal{A}_z(T)$. To apply Bernstein's inequality, we also need to be able to control the supremum norm on this set. Our first step will be to construct such a set.

We denote by \mathcal{B}_r the ball centered at $\mathbf{0}$ and of radius r for the weighted Frobenius norm $\|\cdot\|_{2,\mathcal{B}}$. Let \mathcal{C}_z be a minimal $\sqrt{T}/2$ -net for the weighted Frobenius norm on $\mathcal{B}_{\sqrt{T}} \cap \mathcal{T}_z$. Note that $\mathcal{A}_z \subset \mathcal{B}_{\sqrt{T}} \cap \mathcal{T}_z$, so there exists $\hat{\mathbf{V}} \in \mathcal{C}_z(T)$ such that $\left\| \hat{\mathbf{V}} - \hat{T}^z \right\|_{2,\mathcal{B}} \leq \frac{\sqrt{T}}{2}$. Since our choice of net does not allow us to directly bound $\left\| \hat{\mathbf{V}} - \hat{T}^z \right\|_{\infty}$, we extend

this net using the following argument. For any $\mathbf{V} \in \mathcal{C}_z$ and any matrix $\mathbf{U} \in \{-1, 0, 1\}^{k \times k}$, let $\mathbf{V}^{\mathbf{U}} \in \mathbb{R}^{n \times n}$ be such that $\mathbf{V}_{ii}^{\mathbf{U}} = 0$ and for any $i < j$,

$$\mathbf{V}_{ij}^{\mathbf{U}} = \text{sign}(\mathbf{V}_{ij}) (|\mathbf{V}_{ij}| \wedge \rho_n) (1 - |\mathbf{U}_{z(i)z(j)}|) + \mathbf{U}_{z(i)z(j)} \frac{\rho_n}{2}.$$

Recall that $\|\widehat{\mathbf{T}}^z\|_{\infty} \leq \rho_n$ so for any $\mathbf{V} \in \mathcal{C}_z(T)$ we have $|\text{sign}(\mathbf{V}_{ij}) (|\mathbf{V}_{ij}| \wedge \rho_n) - \widehat{\mathbf{T}}_{ij}^z| \leq |\mathbf{V}_{ij} - \widehat{\mathbf{T}}_{ij}^z|$. This implies that $\|\mathbf{V}^{\mathbf{0}} - \widehat{\mathbf{T}}^z\|_{2, \mathbf{B}} \leq \|\mathbf{V} - \widehat{\mathbf{T}}^z\|_{2, \mathbf{B}}$.

Now, let $\widetilde{\mathcal{C}}_z(T) = \{\mathbf{V}^{\mathbf{U}} : \mathbf{V} \in \mathcal{C}_z(T), \mathbf{U} \in \{-1, 0, 1\}^{k \times k}\}$ and $\widehat{\mathbf{U}} = \arg \min_{\mathbf{U} \in \{-1, 0, 1\}^{k \times k}} \|\widehat{\mathbf{V}}^{\mathbf{U}} - \widehat{\mathbf{T}}^z\|_{\infty}$. By definition, for any $(a, b) \in k \times k$, $\widehat{\mathbf{U}}$ minimises $|\widehat{\mathbf{V}}_{z^{-1}(a)z^{-1}(b)}^{\widehat{\mathbf{U}}} - \widehat{\mathbf{T}}_{z^{-1}(a)z^{-1}(b)}^z|$, so it is also a minimizer of $\|\widehat{\mathbf{V}}^{\widehat{\mathbf{U}}} - \widehat{\mathbf{T}}^z\|_{2, \mathbf{B}} = \sum_{a, b \in [k]} \left(\sum_{(i, j) \in z^{-1}(a) \times z^{-1}(b), i \neq j} \mathbf{B}_{ij} \right) |\widehat{\mathbf{V}}_{z^{-1}(a)z^{-1}(b)}^{\widehat{\mathbf{U}}} - \widehat{\mathbf{T}}_{z^{-1}(a)z^{-1}(b)}^z|^2$. Therefore

$$\|\widehat{\mathbf{V}}^{\widehat{\mathbf{U}}} - \widehat{\mathbf{T}}^z\|_{2, \mathbf{B}} \leq \|\widehat{\mathbf{V}}^{\mathbf{0}} - \widehat{\mathbf{T}}^z\|_{2, \mathbf{B}} \leq \|\widehat{\mathbf{V}} - \widehat{\mathbf{T}}^z\|_{2, \mathbf{B}} \leq \frac{\sqrt{T}}{2}.$$

Furthermore $\|\widehat{\mathbf{V}}^{\widehat{\mathbf{U}}} - \widehat{\mathbf{T}}^z\|_{\infty} \leq \|\widehat{\mathbf{V}}^{\mathbf{U}^*} - \widehat{\mathbf{T}}^z\|_{\infty}$, where $\mathbf{U}_{ab}^* = \text{sign}(\widehat{\mathbf{T}}_{z^{-1}(a)z^{-1}(b)}^z)$. By construction,

$$\|\widehat{\mathbf{V}}^{\mathbf{U}^*} - \widehat{\mathbf{T}}^z\|_{\infty} = \sup_{i < j} |\widehat{\mathbf{T}}_{ij}^z - \text{sign}(\widehat{\mathbf{T}}_{ij}^z) \frac{\rho_n}{2}| = \sup_{i < j} \left| \widehat{\mathbf{T}}_{ij}^z - \frac{\rho_n}{2} \right| \leq \frac{\rho_n}{2}.$$

Hence, $2(\widehat{\mathbf{T}}^z - \widehat{\mathbf{V}}^{\widehat{\mathbf{U}}}) \in \mathcal{A}_z(T)$. Thus, we have shown that

$$\begin{aligned} 2 \left| \langle \Sigma | \widehat{\mathbf{T}}^z - \widehat{\mathbf{V}}^{\widehat{\mathbf{U}}} \rangle \right| &\leq \sup_{\Theta \in \mathcal{A}_z(T)} |\langle \Sigma | \Theta \rangle| \triangleq \left| \langle \Sigma | \widehat{\mathbf{T}}^z \rangle \right| \\ 2 \left| \langle \Sigma | \widehat{\mathbf{T}}^z \rangle \right| - 2 \left| \langle \Sigma | \widehat{\mathbf{V}}^{\widehat{\mathbf{U}}} \rangle \right| &\leq \left| \langle \Sigma | \widehat{\mathbf{T}}^z \rangle \right| \\ \left| \langle \Sigma | \widehat{\mathbf{T}}^z \rangle \right| &\leq 2 \left| \langle \Sigma | \widehat{\mathbf{V}}^{\widehat{\mathbf{U}}} \rangle \right|. \end{aligned}$$

This and (20) allows us to conclude that

$$\sup_{z \in \mathcal{Z}_{n, k}} \sup_{\Theta \in \mathcal{S}_z(T)} \left| \langle \Sigma | \widetilde{\Theta}^z - \Theta \rangle \right| \leq 2 \sup_{z \in \mathcal{Z}_{n, k}} \sup_{\mathbf{V} \in \widetilde{\mathcal{C}}_z(T)} |\langle \Sigma | \mathbf{V} \rangle|. \quad (21)$$

To bound the right hand side of (21), we recall that by hypothesis for any $\mathbf{V} \in \widetilde{\mathcal{C}}_z(T)$, $\sum_{i < j} \mathbb{E}^{\mathbf{B}} [\Sigma_{ij}^2 \mathbf{V}_{ij}^2] \leq \alpha \|\mathbf{V}\|_{2, \mathbf{B}}^2$ and note that $\|\mathbf{V}\|_{\infty} \leq \rho_n$ and $\|\mathbf{V}\|_{2, \mathbf{B}} \leq \sqrt{T}$. We use Bernstein's inequality and the union bound to obtain

$$\mathbb{P} \left(\sup_{z \in \mathcal{Z}_{n, k}} \sup_{\mathbf{V} \in \widetilde{\mathcal{C}}_z(T)} |\langle \Sigma | \mathbf{V} \rangle| \geq \sqrt{2\alpha T t} + \frac{2}{3} t \right) \leq 2e^{-t+n \log(k) + \sup_{\mathbf{V} \in \widetilde{\mathcal{C}}_z(T)} \log(|\widetilde{\mathcal{C}}_z(T)|)}. \quad (22)$$

By construction of $\widetilde{\mathcal{C}}_z(T)$, we have $|\widetilde{\mathcal{C}}_z(T)| = |\mathcal{C}_z(T)| \times 3^{k^2}$. The following classical result on the covering number of a ball will help us bound $|\mathcal{C}_z(T)|$ (see, e.g., Lemma 5.2 in [51]).

Lemma 10. *Let \mathcal{B}_r the ball of a subspace of \mathbb{R}^n of dimension d centered at $\mathbf{0}$ and of radius r for the euclidean norm, and $\mathcal{N}(\mathcal{B}_r, \epsilon)$ its ϵ -covering number, that is the minimal cardinality of a set \mathcal{C} such that for any $\mathbf{X} \in \mathcal{B}_r$, there exists $\mathbf{Y} \in \mathcal{C}$ such that $\|\mathbf{X} - \mathbf{Y}\| \leq \epsilon$. Then*

$$\mathcal{N}(\mathcal{B}_r, \epsilon) \leq \left(\frac{3r}{\epsilon} \right)^d.$$

Extending the proof Lemma 10 to a weighed euclidean norm is straightforward. Putting Lemma 10 into equation (22) and noting that \mathcal{T}_z spans a subspace of $\mathbb{R}^{\frac{(n-1)(n-2)}{2}}$ of dimension $\frac{k(k-1)}{2}$, we get that for some absolute constant

C

$$\begin{aligned} & \mathbb{P} \left(\sup_{z \in \mathcal{Z}_{n,k}} \sup_{\mathbf{V} \in \tilde{\mathcal{C}}_z(T)} |\langle \boldsymbol{\Sigma} | \mathbf{V} \rangle| \geq \sqrt{2\alpha T (t + n \log(k) + k^2 \log(C))} + \frac{2\rho_n}{3} (t + n \log(k) + k^2 \log(C)) \right) \leq 2e^{-t} \\ & \mathbb{P} \left(\sup_{z \in \mathcal{Z}_{n,k}} \sup_{\mathbf{V} \in \tilde{\mathcal{C}}_z(T)} |\langle \boldsymbol{\Sigma} | \mathbf{V} \rangle| \geq \frac{\alpha\gamma_n T}{2 \times 64^3 \rho_n^2} + \left(\frac{2\rho_n}{3} + \frac{2 \times 64^2 \rho_n^2}{\gamma_n} \right) (t + n \log(k) + k^2 \log(C)) \right) \leq 2e^{-t}. \end{aligned}$$

We integrate and find for some absolute constant $C > 0$

$$\mathbb{E} \left[\sup_{z \in \mathcal{Z}_{n,k}} \sup_{\boldsymbol{\Theta} \in \mathcal{S}_z(T)} \left| \langle \boldsymbol{\Sigma} | \tilde{\boldsymbol{\Theta}}^z - \boldsymbol{\Theta} \rangle \right| \right] \leq \frac{\alpha\gamma_n T}{64^3 \rho_n^2} + C \frac{\rho_n^2}{\gamma_n} (n \log(k) + k^2). \quad (23)$$

Combining the bounds (23) and (19) yields the desired result.

5.3 Proof of Lemma 2

The proof of Lemma 2 closely follows that of Lemma 1 and we only sketch it. Recall that $\epsilon_n \triangleq C \frac{\rho_n^2}{\gamma_n} (n \log(k) + k^2)$ where the absolute constant C is larger than the constant appearing in Lemma 11, and that $\epsilon^0 \triangleq \rho_n \epsilon_n$. We show that the probability of the following “bad” event is small:

$$\mathcal{E} \triangleq \left\{ \exists \boldsymbol{\Theta} \in \mathcal{S}_{\Pi} : \left| \Delta \mathcal{K}_{\Pi}^{\boldsymbol{\Theta}^*}(\boldsymbol{\Theta}, \tilde{\boldsymbol{\Theta}}) - \Delta \mathcal{K}_{\mathbf{X}}^{\boldsymbol{\Theta}^*}(\boldsymbol{\Theta}, \tilde{\boldsymbol{\Theta}}) \right| > \frac{1}{2 \times 32 \rho_n} \left\| \boldsymbol{\Theta} - \tilde{\boldsymbol{\Theta}} \right\|_{2, \Pi}^2 + \epsilon_n \right\}.$$

Again, we slice \mathcal{S}_{Π} in different sets $\mathcal{S}_{l, \Pi}$ defined as $\mathcal{S}_{l, \Pi} \triangleq \left\{ \boldsymbol{\Theta} \in \mathcal{S}_{\Pi} : 32^{l-1} (2\epsilon^0) \leq \left\| \boldsymbol{\Theta} - \tilde{\boldsymbol{\Theta}} \right\|_{2, \Pi}^2 \leq 32^l (2\epsilon^0) \right\}$ on which we control the events $\mathcal{E}_l \triangleq \left\{ \exists \boldsymbol{\Theta} \in \mathcal{S}_{l, \Pi} : \left| \Delta \mathcal{K}_{\Pi}^{\boldsymbol{\Theta}^*}(\boldsymbol{\Theta}, \tilde{\boldsymbol{\Theta}}) - \Delta \mathcal{K}_{\mathbf{X}}^{\boldsymbol{\Theta}^*}(\boldsymbol{\Theta}, \tilde{\boldsymbol{\Theta}}) \right| > \frac{32^{l-1} \times 2\epsilon^0}{4 \times 32 \rho_n} + \epsilon_n \right\}$. To do this, we set $\mathcal{S}_{\Pi}(T) \triangleq \left\{ \boldsymbol{\Theta} \in \mathcal{S}_{\Pi} : \left\| \boldsymbol{\Theta} - \tilde{\boldsymbol{\Theta}} \right\|_{2, \Pi}^2 \leq T \right\}$ and we control the probability of the events

$$\mathcal{E}(T) = \left\{ \exists \boldsymbol{\Theta} \in \mathcal{S}_{\Pi}(T) : \left| \Delta \mathcal{K}_{\Pi}^{\boldsymbol{\Theta}^*}(\boldsymbol{\Theta}, \tilde{\boldsymbol{\Theta}}) - \Delta \mathcal{K}_{\mathbf{X}}^{\boldsymbol{\Theta}^*}(\boldsymbol{\Theta}, \tilde{\boldsymbol{\Theta}}) \right| > \frac{T}{64^2 \rho_n} + \epsilon_n \right\}.$$

The following lemma helps us bound the probability of the events $\mathcal{E}(T)$.

Lemma 11. *Let $\tilde{Z}_T = \sup_{\boldsymbol{\Theta} \in \mathcal{S}_{\Pi}(T)} \left| \Delta \mathcal{K}_{\Pi}^{\boldsymbol{\Theta}^*}(\boldsymbol{\Theta}, \tilde{\boldsymbol{\Theta}}) - \Delta \mathcal{K}_{\mathbf{X}}^{\boldsymbol{\Theta}^*}(\boldsymbol{\Theta}, \tilde{\boldsymbol{\Theta}}) \right|$. There exists two absolute constants $C, C' > 0$ such that*

$$\mathbb{P} \left(\tilde{Z}_T \geq \frac{T}{64^2 \rho_n} + C \frac{\rho_n^2}{\gamma_n} (n \log(k) + k^2) \right) \leq \exp \left(-\frac{C' T \gamma_n^2}{\rho_n^2} \right).$$

Proof. To prove Lemma 11, we first show that Z_T concentrates around its expectation and then bound this term.

Lemma 12. *Let \tilde{Z}_T be defined as in Lemma 11. Then there exists an absolute constant $C > 0$ such that*

$$\mathbb{P} \left(\tilde{Z}_T > 2\mathbb{E}[\tilde{Z}_T] + \frac{T}{2 \times 64^2 \rho_n} \right) \leq \exp \left(-\frac{CT \gamma_n^2}{\rho_n^2} \right).$$

Lemma 13. *Let \tilde{Z}_T be as in Lemma 11, then there exists an absolute constant $C > 0$ such that*

$$\mathbb{E} \left[\tilde{Z}_T \right] \leq \frac{T}{4 \times 64^2 \rho_n} + C \frac{\rho_n^2}{\gamma_n} (n \log(k) + k^2). \quad (24)$$

Putting together Lemma 12 and Lemma 13, we get that there exists two absolute constants $C, C' > 0$ such that

$$\mathbb{P} \left(\tilde{Z}_T \geq \frac{T}{64^2 \rho_n} + C \frac{\rho_n^2}{\gamma_n} (n \log(k) + k^2) \right) \leq \exp \left(-\frac{C' T \gamma_n^2}{\rho_n^2} \right).$$

This concludes the proof of Lemma 11. □

Lemma 11 and the arguments developed to prove Lemma 1 help us conclude the proof of Lemma 2.

5.3.1 Proof of Lemma 12

Recall that by definition of \tilde{Z}_T ,

$$\begin{aligned}\tilde{Z}_T &= \sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} (\Pi_{ij} - \mathbf{X}_{ij}) \left(\Theta_{ij}^* \log \left(\frac{\tilde{\Theta}_{ij}}{\Theta_{ij}} \right) + (1 - \Theta_{ij}^*) \log \left(\frac{1 - \tilde{\Theta}_{ij}}{1 - \Theta_{ij}} \right) \right) \right| \\ &= \frac{1}{\gamma_n} \sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} f_{ij}^{\Theta}(\mathbf{X}_{ij}) \right|\end{aligned}$$

where we set $f_{ij}^{\Theta}(\mathbf{X}_{ij}) \triangleq \gamma_n (\Pi_{ij} - \mathbf{X}_{ij}) \left(\Theta_{ij}^* \log \left(\frac{\tilde{\Theta}_{ij}}{\Theta_{ij}} \right) + (1 - \Theta_{ij}^*) \log \left(\frac{1 - \tilde{\Theta}_{ij}}{1 - \Theta_{ij}} \right) \right)$. Assuming that $\gamma_n \leq 1 - \rho_n$, $x \rightarrow \log(x)$ and $x \rightarrow \log(1 - x)$ are $\frac{1}{\gamma_n}$ -Lipshitz on $[\gamma_n, \rho_n]$ so

$$\left| \Theta_{ij}^* \log \left(\frac{\tilde{\Theta}_{ij}}{\Theta_{ij}} \right) + (1 - \Theta_{ij}^*) \log \left(\frac{1 - \tilde{\Theta}_{ij}}{1 - \Theta_{ij}} \right) \right| \leq \Theta_{ij}^* \frac{|\Theta_{ij} - \tilde{\Theta}_{ij}|}{\gamma_n} + (1 - \Theta_{ij}^*) \frac{|\Theta_{ij} - \tilde{\Theta}_{ij}|}{\gamma_n} \leq \frac{|\Theta_{ij} - \tilde{\Theta}_{ij}|}{\gamma_n}$$

which implies that for any $1 \leq i < j \leq n$, $|f_{ij}^{\Theta}(\mathbf{X}_{ij})| \leq 1$. Moreover for any $1 \leq i < j \leq n$, $\mathbb{E}[f_{ij}^{\Theta}(\mathbf{X}_{ij})] = 0$ and $\mathbb{E}[(\mathbf{X}_{ij} - \Pi_{ij})^2] \leq \Pi_{ij}$, hence

$$\begin{aligned}\sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \sum_{1 \leq i < j \leq n} \mathbb{E}[f_{ij}^{\Theta}(\mathbf{X}_{ij})^2] &\leq \gamma_n^2 \sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \sum_{1 \leq i \leq n} \Pi_{ij} \left(\Theta_{ij}^* \log \left(\frac{\tilde{\Theta}_{ij}}{\Theta_{ij}} \right) + (1 - \Theta_{ij}^*) \log \left(\frac{1 - \tilde{\Theta}_{ij}}{1 - \Theta_{ij}} \right) \right)^2 \\ &\leq \gamma_n^2 \sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \sum_{1 \leq i \leq n} \Pi_{ij} \frac{1}{\gamma_n^2} (\Theta_{ij} - \tilde{\Theta}_{ij})^2 \\ &\leq T.\end{aligned}$$

Then, Theorem 5 implies

$$\begin{aligned}\mathbb{P}\left(\gamma_n \tilde{Z}_T > \gamma_n \mathbb{E}[\tilde{Z}_T] + \frac{x}{3} + \sqrt{2x(2\gamma_n \mathbb{E}[\tilde{Z}_T] + T)}\right) &\leq \exp(-x) \\ \mathbb{P}\left(\tilde{Z}_T > \mathbb{E}[\tilde{Z}_T] + \frac{x}{3\gamma_n} + \frac{4x}{\gamma_n} + \mathbb{E}[\tilde{Z}_T] + \frac{2x \times 4 \times 64^2 \rho_n}{\gamma_n^2} + \frac{T}{4 \times 64^2 \rho_n}\right) &\leq \exp(-x) \\ \mathbb{P}\left(\tilde{Z}_T > 2\mathbb{E}[\tilde{Z}_T] + \frac{9 \times 64^2 x \rho_n}{\gamma_n^2} + \frac{T}{4 \times 64^2 \rho_n}\right) &\leq \exp(-x)\end{aligned}$$

where we have used $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$, $\sqrt{ab} \leq a+b$ and $\frac{\rho_n}{\gamma_n} \geq 1$. Setting $x = \frac{T\gamma_n^2}{9 \times 64^2 \times 4 \times 64^2 \rho_n^2}$ yields the desired result.

5.3.2 Proof of Lemma 13

In Lemma 13, we bound

$$\mathbb{E}[\tilde{Z}_T] = \mathbb{E} \left[\sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} (\Pi_{ij} - \mathbf{X}_{ij}) \left(\Theta_{ij}^* \log \left(\frac{\tilde{\Theta}_{ij}}{\Theta_{ij}} \right) + (1 - \Theta_{ij}^*) \log \left(\frac{1 - \tilde{\Theta}_{ij}}{1 - \Theta_{ij}} \right) \right) \right| \right].$$

Let $(\epsilon_{ij})_{1 \leq i < j \leq n}$ be a Rademacher sequence. We apply Lemma 7 and get

$$\mathbb{E}[\tilde{Z}_T] \leq 2\mathbb{E} \left[\sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} \epsilon_{ij} \mathbf{X}_{ij} \left(\Theta_{ij}^* \log \left(\frac{\tilde{\Theta}_{ij}}{\Theta_{ij}} \right) + (1 - \Theta_{ij}^*) \log \left(\frac{1 - \tilde{\Theta}_{ij}}{1 - \Theta_{ij}} \right) \right) \right| \right].$$

For any $1 \leq i < j \leq n$, let $\phi_{ij} : x \rightarrow \frac{\gamma_n}{2\rho_n} \mathbf{X}_{ij} \left(\Theta_{ij}^* \log \left(\frac{\tilde{\Theta}_{ij} - x}{\Theta_{ij}} \right) + (1 - \Theta_{ij}^*) \log \left(\frac{1+x - \tilde{\Theta}_{ij}}{1 - \Theta_{ij}} \right) \right)$. Note that on $[\tilde{\Theta}_{ij} - \rho_n, \tilde{\Theta}_{ij} - \gamma_n]$, ϕ_{ij} is 1-Lipschitz and vanishes at 0. Then we apply Lemma 8 and compute

$$\begin{aligned}\mathbb{E}[\tilde{Z}_T] &\leq \frac{4\rho_n}{\gamma_n} \mathbb{E} \left[\sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} \epsilon_{ij} \phi_{ij}(\mathbf{X}_{ij}(\Theta_{ij} - \tilde{\Theta}_{ij})) \right| \right] \\ &\leq \frac{8\rho_n}{\gamma_n} \mathbb{E} \left[\sup_{\Theta \in \mathcal{S}_{\Pi}(T)} \left| \sum_{1 \leq i < j \leq n} \epsilon_{ij} \mathbf{X}_{ij}(\Theta_{ij} - \tilde{\Theta}_{ij}) \right| \right].\end{aligned}$$

Now, applying Lemma 9 with $\alpha = 1$ and $\mathbf{B} = \mathbf{\Pi}$ allows us to conclude that there exists an absolute constant $C > 0$ such that

$$\mathbb{E} \left[\tilde{Z}_T \right] \leq \frac{T}{8 \times 64^2 \rho_n} + C \frac{\rho_n^3}{\gamma_n^2} (n \log(n) + k^2) \leq \frac{T}{8 \times 64^2 \rho_n} + C \frac{\rho_n^2}{\gamma_n^2} (n \log(n) + k^2).$$

5.4 Proof of Lemma 3

The proof of Lemma 3 closely follows that of Lemma 2, and we only sketch it. Recall that $\epsilon_n \triangleq C \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2)$ where the absolute constant C is larger than the constant appearing in Lemma 14, and that $\epsilon^0 \triangleq \rho_n \epsilon_n$. We show that conditionally on \mathbf{X} , the probability of the following ‘‘bad’’ event is small and does not depend on \mathbf{X} :

$$\mathcal{E}_{\mathbf{X}} \triangleq \left\{ \exists \Theta \in \mathcal{S}_{\mathbf{X}} : \left| \Delta \mathcal{K}_{\mathbf{X}}^{\Theta^*}(\Theta, \tilde{\Theta}) - \Delta \mathcal{K}_{\mathbf{X}}^A(\Theta, \tilde{\Theta}) \right| > \frac{1}{2 \times 64 \rho_n} \|\Theta - \tilde{\Theta}\|_{2, \mathbf{X}}^2 + \epsilon_n \right\}.$$

We slice $\mathcal{S}_{\mathbf{X}}$ in the following sets $\mathcal{S}_{l, \mathbf{X}} \triangleq \left\{ \Theta \in \mathcal{S}_{\mathbf{X}} : 64^{l-1} \epsilon^0 \leq \|\Theta - \tilde{\Theta}\|_{2, \mathbf{X}}^2 \leq 64^l \epsilon^0 \right\}$ and control the probability of the events $\mathcal{E}_{l, \mathbf{X}} \triangleq \left\{ \exists \Theta \in \mathcal{S}_{l, \mathbf{X}} : \left| \Delta \mathcal{K}_{\mathbf{X}}^{\Theta^*}(\Theta, \tilde{\Theta}) - \Delta \mathcal{K}_{\mathbf{X}}^A(\Theta, \tilde{\Theta}) \right| > \frac{64^l \epsilon^0}{2 \times 64^2 \rho_n} + \epsilon_n \right\}$. To do this, we control the probability of the events $\mathcal{E}_{\mathbf{X}}(T) = \left\{ \exists \Theta \in \mathcal{S}_{\mathbf{X}}(T) : \left| \Delta \mathcal{K}_{\mathbf{X}}^{\Theta^*}(\Theta, \tilde{\Theta}) - \Delta \mathcal{K}_{\mathbf{X}}^A(\Theta, \tilde{\Theta}) \right| > \frac{T}{2 \times 64^2 \rho_n} + \epsilon_n \right\}$ where $\mathcal{S}_{\mathbf{X}}(T) = \left\{ \Theta \in \mathcal{S}_{\mathbf{X}} : \|\Theta - \tilde{\Theta}\|_{2, \mathbf{X}}^2 \leq T \right\}$.

Lemma 14. *Let $Z_{T, \mathbf{X}} = \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \Delta \mathcal{K}_{\mathbf{X}}^{\Theta^*}(\Theta, \tilde{\Theta}) - \Delta \mathcal{K}_{\mathbf{X}}^A(\Theta, \tilde{\Theta}) \right|$. There exists two absolute constants $C, C' > 0$ such that $\mathbb{P}^{\mathbf{X}} \left(Z_{T, \mathbf{X}} \geq \frac{T}{2 \times 64^2 \rho_n} + C \frac{\rho_n^2}{\gamma_n^2} (n \log(k) + k^2) \right) \leq 4 \exp \left(-\frac{C' \gamma_n^2 T}{\rho_n^2} \right)$.*

Proof. To prove Lemma 14, we first show that $Z_{T, \mathbf{X}}$ concentrates around its expectation and then bound this term.

Lemma 15. *Let $Z_{T, \mathbf{X}}$ be as in Lemma 14, then there exists two absolute constants $C, C' > 0$ such that*

$$\mathbb{P}^{\mathbf{X}} \left(\left| Z_{T, \mathbf{X}} - \mathbb{E}^{\mathbf{X}}(Z_{T, \mathbf{X}}) \right| > \frac{C \rho_n}{\gamma_n^2} + \frac{T}{4 \times 64^2 \rho_n} \right) \leq 4 \exp \left(-\frac{C' T \gamma_n^2}{\rho_n^2} \right).$$

Lemma 16. *Let $Z_{T, \mathbf{X}}$ be as in Lemma 14, then there exists an absolute constant $C > 0$ such that*

$$\mathbb{E}^{\mathbf{X}} [Z_{T, \mathbf{X}}] \leq \frac{T}{4 \times 64^2 \rho_n} + \frac{C \rho_n^2}{\gamma_n^2} (n \log(k) + k^2).$$

Putting together Lemma 15 and Lemma 16, we get that

$$\mathbb{P}^{\mathbf{X}} \left(Z_{T, \mathbf{X}} \geq \frac{T}{2 \times 64^2 \rho_n} + C \frac{\rho_n^2}{\gamma_n^2} \left(n \log(k) + \frac{1}{\rho_n} + k^2 \right) \right) \leq 4 \exp \left(-\frac{C' \gamma_n^2 T}{\rho_n^2} \right).$$

If $n \rho_n \rightarrow \infty$, for n large enough $n > \frac{1}{\rho_n}$. This yields the desired result. \square

We combine Lemma 14 and the arguments developed in Lemma 1, and note that $\mathbb{P}^{\mathbf{X}}(\mathcal{E}_{\mathbf{X}})$ does not depend on \mathbf{X} to conclude the proof of Lemma 3.

5.4.1 Proof of Lemma 15

In this Section, we prove the Lemma 15 that helps us bound $|Z_{T, \mathbf{X}} - \mathbb{E}^{\mathbf{X}}(Z_{T, \mathbf{X}})|$ with high probability. To prove that $Z_{T, \mathbf{X}}$ concentrates around its mean, we use the following version of Talagrand’s Theorem for Lipschitz convex functions (for a proof, see Theorem 3.3 of [14]).

Theorem 7. *Suppose that $f : [-1, 1]^N \rightarrow \mathbb{R}$ is a convex Lipschitz function with Lipschitz constant L . Let R_1, \dots, R_N be independent random variables taking value in $[-1, 1]$. Let $Z := f(R_1, \dots, R_N)$. Then for any $t \geq 0$,*

$$\mathbb{P}(|Z - \mathbb{E}(Z)| > 16L + t) \leq 4e^{\left(\frac{-t^2}{2L^2} \right)}.$$

Recall that

$$\begin{aligned} Z_{T,\mathbf{X}} &= \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \Delta \mathcal{K}_{\mathbf{X}}^{\Theta^*}(\Theta, \tilde{\Theta}) - \Delta \mathcal{K}_{\mathbf{X}}^{\mathbf{A}}(\Theta, \tilde{\Theta}) \right| \\ &= \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \sum_{1 \leq i < j \leq n} (\mathbf{X}_{ij} \mathbf{A}_{ij} - \mathbf{X}_{ij} \Theta_{ij}^*) \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) - \log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right) \right|. \end{aligned}$$

Note that $Z_{T,\mathbf{X}} = f(\mathbf{A})$ where $f(\mathbf{R})$ is defined for $\mathbf{R} \in [-1, 1]^{\frac{(n-1)(n-2)}{2}}$ by

$$f : \mathbf{R} \rightarrow \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} (\mathbf{R}_{ij} - \Theta_{ij}^*) \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) - \log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right) \right|.$$

It is easy to see that f is indeed convex. Our next step is to show that f is Lipschitz. Let $\mathbf{R}, \mathbf{S} \in [-1, 1]^{\frac{(n-1)(n-2)}{2}}$,

$$\begin{aligned} |f(\mathbf{R}) - f(\mathbf{S})| &= \left| \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} (\mathbf{R}_{ij} - \Theta_{ij}^*) \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) - \log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right) \right| \right. \\ &\quad \left. - \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} (\mathbf{S}_{ij} - \Theta_{ij}^*) \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) - \log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right) \right| \right| \\ &\leq \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} (\mathbf{R}_{ij} - \Theta_{ij}^*) \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) - \log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right) \right| \\ &\quad - \left| \sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} (\mathbf{S}_{ij} - \Theta_{ij}^*) \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) - \log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right) \right| \\ &\leq \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} (\mathbf{R}_{ij} - \mathbf{S}_{ij}) \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) - \log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right) \right| \\ &\leq \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left\{ \sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} \left| (\mathbf{R}_{ij} - \mathbf{S}_{ij}) \log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) \right| + \mathbf{X}_{ij} \left| (\mathbf{R}_{ij} - \mathbf{S}_{ij}) \log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right| \right\} \\ &\leq \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \|\mathbf{R} - \mathbf{S}\|_2 \left(\sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) \right)^2 + \mathbf{X}_{ij} \left(\log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right)^2 \right)^{\frac{1}{2}} \end{aligned}$$

where we have used that $\mathbf{X} \in \{0, 1\}^{n \times n}$. Thus f is Lipschitz with Lipschitz constant

$$\sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left(\sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) \right)^2 + \mathbf{X}_{ij} \left(\log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right)^2 \right)^{\frac{1}{2}}.$$

As stated before, assuming that $\gamma_n \leq 1 - \rho_n$, $x \rightarrow \log(x)$ and $x \rightarrow \log(1 - x)$ are Lipschitz functions on $[\gamma_n, \rho_n]$ with Lipschitz constant γ_n^{-1} . Thus f is Lipschitz with Lipschitz constant

$$\sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left(\sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} \left(\frac{|\Theta_{ij} - \tilde{\Theta}_{ij}|}{\gamma_n} \right)^2 + \mathbf{X}_{ij} \left(\frac{|\Theta_{ij} - \tilde{\Theta}_{ij}|}{\gamma_n} \right)^2 \right)^{\frac{1}{2}}.$$

This implies

$$|f(\mathbf{R}) - f(\mathbf{S})| \leq \|\mathbf{R} - \mathbf{S}\|_2 \sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \frac{\sqrt{2} \|\tilde{\Theta} - \Theta\|_{2,\mathbf{X}}}{\gamma_n} \leq \|\mathbf{R} - \mathbf{S}\|_2 \frac{\sqrt{2T}}{\gamma_n}.$$

We have shown that f has a Lipschitz constant $L = \frac{\sqrt{2T}}{\gamma_n}$. Applying Theorem 7 for $t = \frac{T}{8 \times 64^2 \rho_n}$, we get

$$\mathbb{P}^{\mathbf{X}} \left(\left| Z_{T,\mathbf{X}} - \mathbb{E}(Z_{T,\mathbf{X}}) \right| > \frac{16\sqrt{2T}}{\gamma_n} + \frac{T}{8 \times 64^2 \rho_n} \right) \leq 4 \exp \left(\frac{-T\gamma_n^2}{8^2 \times 2 \times 64^4 \rho_n^2} \right).$$

Using for $\beta > 0$, $2\sqrt{ab} \leq \beta a^2 + b^2/\beta$ yields

$$\mathbb{P}^{\mathbf{X}} \left(\left| Z_{T,\mathbf{X}} - \mathbb{E}(Z_{T,\mathbf{X}}) \right| > \frac{8^2 \times 64^2 \times 16\rho_n}{\gamma_n^2} + \frac{T}{8 \times 64^2 \rho_n} + \frac{T}{8 \times 64^2 \rho_n} \right) \leq 4 \exp \left(\frac{-T\gamma_n^2}{4 \times 4^2 \times 32^4 \rho_n^2} \right).$$

This concludes the proof of Lemma 15.

5.5 Proof of Lemma 16

Once we have shown that $Z_{T,\mathbf{X}}$ concentrates around its mean, we bound $\mathbb{E}[Z_{T,\mathbf{X}}]$. To do so, we follow the steps of Lemma 13. Let $\epsilon_{1 \leq i < j \leq n}$ a Rademacher sequence. Applying Lemma 7, we get

$$\begin{aligned} \mathbb{E}^{\mathbf{X}} [Z_{T,\mathbf{X}}] &= \mathbb{E}^{\mathbf{X}} \left[\sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} (\mathbf{A}_{ij} - \mathbb{E}[\mathbf{A}_{ij}]) \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) - \log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right) \right| \right] \\ &\leq 2 \mathbb{E}^{\mathbf{X}} \left[\sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \sum_{1 \leq i < j \leq n} \mathbf{X}_{ij} \epsilon_{ij} \mathbf{A}_{ij} \left(\log \left(\frac{\Theta_{ij}}{\tilde{\Theta}_{ij}} \right) - \log \left(\frac{1 - \Theta_{ij}}{1 - \tilde{\Theta}_{ij}} \right) \right) \right| \right]. \end{aligned}$$

For any $1 \leq i < j \leq n$, let $\phi_{ij} : x \rightarrow \frac{1}{2} \gamma_n \mathbf{X}_{ij} \mathbf{A}_{ij} \left(\log \left(\frac{\tilde{\Theta}_{ij} - x}{\Theta_{ij}} \right) - \log \left(\frac{1+x-\tilde{\Theta}_{ij}}{1-\tilde{\Theta}_{ij}} \right) \right)$. Note that ϕ_{ij} is 1-Lipschitz and vanishes at 0 on the interval $[\tilde{\Theta}_{ij} - \rho_n, \tilde{\Theta}_{ij} - \gamma_n]$. Indeed,

$$\begin{aligned} \phi_{ij}(x)' &= \frac{1}{2} \gamma_n \mathbf{X}_{ij} \mathbf{A}_{ij} \left(\frac{-1}{\tilde{\Theta}_{ij} - x} - \frac{1}{1+x-\tilde{\Theta}_{ij}} \right) \\ &\leq \mathbf{X}_{ij} \mathbf{A}_{ij}. \end{aligned}$$

By definition of the functions ϕ_{ij} ,

$$\mathbb{E}^{\mathbf{X}} [Z_{T,\mathbf{X}}] \leq \frac{4}{\gamma_n} \mathbb{E}^{\mathbf{X}} \left[\sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \sum_{i < j} \epsilon_{ij} \phi_{ij}(\mathbf{X}_{ij} \mathbf{A}_{ij} (\tilde{\Theta}_{ij} - \Theta_{ij})) \right| \right].$$

We apply Lemma 8 to get

$$\mathbb{E}^{\mathbf{X}} [Z_{T,\mathbf{X}}] \leq \frac{8}{\gamma_n} \mathbb{E}^{\mathbf{X}} \left[\sup_{\Theta \in \mathcal{S}_{\mathbf{X}}(T)} \left| \sum_{i < j} \mathbf{X}_{ij} \epsilon_{ij} \mathbf{A}_{ij} (\tilde{\Theta}_{ij} - \Theta_{ij}) \right| \right]. \quad (25)$$

Next, we apply Lemma 9 with $\mathbf{B} = \mathbf{X}$, $\Sigma_{ij} = \mathbf{X}_{ij} \mathbf{A}_{ij} \epsilon_{ij}$ and $\alpha = \rho_n$. Note that $\|\Sigma\|_{\infty} \leq 1$ and that for any matrix Θ , $\sum_{i < j} \mathbb{E}^{\mathbf{X}} [\mathbf{X}_{ij}^2 \epsilon_{ij}^2 \Theta_{ij}^2] \leq \rho_n \|\Theta\|_{2,\mathbf{X}}^2$. Combining Lemma 9 and (25) yields

$$\mathbb{E}^{\mathbf{X}} [Z_{T,\mathbf{X}}] \leq \frac{8}{\gamma_n} \times \left(\frac{T\gamma_n}{32 \times 64^2 \rho_n} + C \frac{\rho_n^2}{\gamma_n} (n \log(k) + k^2) \right).$$

This concludes the proof of Lemma 16.

5.6 Proof of Theorem 3

Our proof relies on two steps: first, we show that with high probability, \hat{d} is close to its expected value, which belongs to $[\gamma_n, \rho_n]$. More precisely, let $\underline{\gamma}_n = \frac{C_{inf}}{2} \rho_n \log(n)^{\frac{-1}{5}}$ and $\underline{\rho}_n = \left(1 + \frac{C_{inf}}{2}\right) \rho_n \log(n)^{\frac{1}{5}}$. We prove that with high probability, $\underline{\gamma}_n \leq \hat{\gamma}_n \leq \gamma_n$ and $\rho_n \leq \hat{\rho}_n \leq \underline{\rho}_n$. Then this implies that the oracle matrix $\tilde{\Theta}$ belongs to the set of definition of our estimator and its likelihood is greater than that of $\hat{\Theta}$. Then both $\tilde{\Theta}$ and $\hat{\Theta}$ belong to the set $[\underline{\gamma}_n, \underline{\rho}_n]^{n \times n}$ and we adapt the proof of Theorem 1 to get the desired result.

Lemma 17. *Let $\mathcal{E} = \{\hat{\gamma}_n \in [\underline{\gamma}_n, \gamma_n], \hat{\rho}_n \in [\rho_n, \underline{\rho}_n]\}$. There exists a positive constant C and an integer N , both depending only on C_{inf} , such that $\forall n \geq N$, $\mathbb{P}(\mathcal{E}) \geq 1 - \exp(-Cn\rho_n)$.*

Proof. Note that $\|\mathbf{A} - \Theta^*\|_\infty \leq 1$ almost surely, and that for any $1 \leq i < j \leq n$, $(\mathbf{A}_{ij} - \Theta_{ij}^*)$ is centered and has a variance smaller than ρ_n . Applying Bernstein's inequality 6 yields

$$\mathbb{P}\left(\left|\sum_{(i,j) \in \Omega} (\mathbf{A}_{ij} - \Theta_{ij}^*)\right| \geq \sqrt{2tn\rho_n} + \frac{3t}{2}\right) \leq 2e^{-t}, \quad \forall t > 0.$$

Choosing $t = \rho_n n C$ with $C > 0$ such that $\sqrt{2C} + \frac{3C}{2} \leq \frac{C_{inf}}{2}$ yields

$$\mathbb{P}\left(\left|\hat{d} - \frac{\sum_{(i,j) \in \Omega} \Theta_{ij}^*}{n}\right| \geq \frac{C_{inf}}{2} \rho_n\right) \leq 2e^{-Cn\rho_n}.$$

Note that in the sparse graphon model (4), when $0 < C_{inf} \triangleq \inf_{(x,y) \in [0,1]^2} W^*(x,y)$, we see that $\gamma_n = C_{inf} \rho_n$ and $\frac{\sum_{(i,j) \in \Omega} \Theta_{ij}^*}{n} \in [\gamma_n, \rho_n] = [C_{inf} \rho_n, \rho_n]$. So, with probability greater than $1 - 2e^{-Cn\rho_n}$, $\hat{A} \in \left[\frac{C_{inf} \rho_n}{2}, (1 + \frac{C_{inf}}{2}) \rho_n\right]$. Let N be such that $\log(N)^{-\frac{1}{5}} \leq \frac{C_{inf}}{1 + \frac{C_{inf}}{2}}$ and $\log(N)^{\frac{1}{5}} \geq 2C_{inf}^{-1}$. For any $n \geq N$, with probability greater than $1 - 2e^{-Cn\rho_n}$, $\widehat{\gamma}_n \in [\underline{\gamma}_n, \overline{\gamma}_n]$ and $\widehat{\rho}_n \in [\underline{\rho}_n, \overline{\rho}_n]$. \square

To prove Theorem 3, we work conditionally on the event \mathcal{E} . Note that in the model (1), the law of the remaining entries $(A_{i,j})_{(i,j) \notin \Omega}$ is independent of \mathcal{E} . Since on \mathcal{E} both $\widehat{\Theta}$ and Θ^* belong to the set $[\underline{\gamma}_n, \overline{\rho}_n]$, we have

$$\begin{aligned} \|\widehat{\Theta} - \Theta^*\|_2^2 &= \sum_{(i,j) \in \Omega} (\widehat{\Theta}_{ij} - \Theta_{ij}^*)^2 + \sum_{(i,j) \notin \Omega} (\widehat{\Theta}_{ij} - \Theta_{ij}^*)^2 \\ &\leq n \underline{\rho}_n^2 + \sum_{(i,j) \notin \Omega} (\widehat{\Theta}_{ij} - \Theta_{ij}^*)^2. \end{aligned} \quad (26)$$

We adapt the proof of Theorem 1 to bound the second term. Let $\underline{\epsilon}_n \triangleq C \left(\frac{\rho_n}{\gamma_n}\right)^2 (n \log(k) + k^2)$ where C is the same absolute constant as in Theorem 1, and let $\underline{\epsilon}^0 \triangleq \underline{\rho}_n \underline{\epsilon}_n$. We start by considering the following two cases:

Case 1: $\sum_{(i,j) \notin \Omega} (\widehat{\Theta}_{ij} - \widetilde{\Theta}_{ij})^2 \leq \underline{\epsilon}^0$. Then, the statement of Theorem 3 follows from (26) and Lemma 19:

$$\begin{aligned} \sum_{(i,j) \notin \Omega} (\widehat{\Theta}_{ij} - \Theta_{ij}^*)^2 &\leq 2 \sum_{(i,j) \notin \Omega} (\widehat{\Theta}_{ij} - \widetilde{\Theta}_{ij})^2 + 2 \sum_{(i,j) \notin \Omega} (\widetilde{\Theta}_{ij} - \Theta_{ij}^*)^2 \\ &\leq 2 \underline{\rho}_n \underline{\epsilon}_n + 16 \underline{\rho}_n \mathcal{K}(\Theta^*, \widetilde{\Theta}) \\ &\leq C \log(n) \rho_n \left(\mathcal{K}(\Theta^*, \widetilde{\Theta}) + n \log(k) + k^2\right). \end{aligned}$$

Case 2: $\sum_{(i,j) \notin \Omega} (\widehat{\Theta}_{ij} - \widetilde{\Theta}_{ij})^2 > \underline{\epsilon}^0$. Then $\widehat{\Theta}$ belongs to the set

$$\underline{\mathcal{S}} = \left\{ \Theta \in \bigcup_{z \in \mathcal{Z}_{n,k}} \mathcal{T}_z : \sum_{(i,j) \notin \Omega} (\widehat{\Theta}_{ij} - \widetilde{\Theta}_{ij})^2 > \underline{\epsilon}^0, \|\Theta\|_\infty \leq \underline{\rho}_n, \min_{i < j} \{\Theta_{ij}\} \geq \underline{\gamma}_n \right\}.$$

As before, Lemma 19 implies

$$\begin{aligned} \sum_{(i,j) \notin \Omega} (\widehat{\Theta}_{ij} - \Theta_{ij}^*)^2 &\leq 8 \underline{\rho}_n \sum_{(i,j) \notin \Omega} \mathcal{K}(\Theta_{ij}^*, \widehat{\Theta}_{ij}) \\ &\leq 8 \underline{\rho}_n \mathcal{K}(\Theta^*, \widetilde{\Theta}) + 8 \underline{\rho}_n \sum_{(i,j) \notin \Omega} \left(\mathcal{K}(\Theta_{ij}^*, \widehat{\Theta}_{ij}) - \mathcal{K}(\Theta_{ij}^*, \widetilde{\Theta}_{ij})\right). \end{aligned}$$

On the event \mathcal{E} , $\widetilde{\Theta}$ belongs to the set of matrices on which the maximum likelihood estimator is defined, thus the definition of $\widehat{\Theta}$ implies $\sum_{(i,j) \notin \Omega} \left(\mathcal{K}(\mathbf{A}_{ij}, \widehat{\Theta}_{ij}) - \mathcal{K}(\mathbf{A}_{ij}, \widetilde{\Theta}_{ij})\right) \leq 0$ and

$$\sum_{(i,j) \notin \Omega} (\widehat{\Theta}_{ij} - \Theta_{ij}^*)^2 \leq 8 \underline{\rho}_n \mathcal{K}(\Theta^*, \widetilde{\Theta}) + 8 \underline{\rho}_n \sum_{(i,j) \notin \Omega} \left(\mathcal{K}(\Theta_{ij}^*, \widehat{\Theta}_{ij}) - \mathcal{K}(\mathbf{A}_{ij}, \widehat{\Theta}_{ij}) - \left(\mathcal{K}(\Theta_{ij}^*, \widetilde{\Theta}_{ij}) - \mathcal{K}(\mathbf{A}_{ij}, \widetilde{\Theta}_{ij})\right)\right).$$

The proof of the following lemma follows the lines of the proof of Lemma 3, and we do not present it.

Lemma 18. *There exists a constant $C > 0$ depending only on C_{inf} such that for any $\Theta \in \underline{\mathcal{S}}$ simultaneously we have*

$$\left| \sum_{(i,j) \notin \Omega} \left(\mathcal{K}(\Theta_{ij}^*, \Theta_{ij}) - \mathcal{K}(\mathbf{A}_{ij}, \Theta_{ij}) - \left(\mathcal{K}(\Theta_{ij}^*, \tilde{\Theta}_{ij}) - \mathcal{K}(\mathbf{A}_{ij}, \tilde{\Theta}_{ij}) \right) \right) \right| \leq \frac{1}{32\underline{\rho}_n} \sum_{(i,j) \notin \Omega} \left(\Theta_{ij} - \tilde{\Theta}_{ij} \right)^2 + \underline{\epsilon}_n$$

with probability at least $1 - 5 \exp(-C\underline{\rho}_n(n \log(k) + k^2))$.

This implies that on the event \mathcal{E} , with large probability,

$$\begin{aligned} \sum_{(i,j) \notin \Omega} \left(\hat{\Theta}_{ij} - \Theta_{ij}^* \right)^2 &\leq 8\underline{\rho}_n \mathcal{K}(\Theta^*, \tilde{\Theta}) + \frac{1}{4} \sum_{(i,j) \notin \Omega} \left(\hat{\Theta}_{ij} - \tilde{\Theta}_{ij} \right)^2 + 8\underline{\epsilon}_n \rho_n \\ &\leq 8\underline{\rho}_n \mathcal{K}(\Theta^*, \tilde{\Theta}) + \frac{1}{2} \sum_{(i,j) \notin \Omega} \left(\hat{\Theta}_{ij} - \Theta_{ij}^* \right)^2 + \frac{1}{2} \sum_{(i,j) \notin \Omega} \left(\Theta_{ij}^* - \tilde{\Theta}_{ij} \right)^2 + 8\underline{\epsilon}_n \rho_n \\ \frac{1}{2} \sum_{(i,j) \notin \Omega} \left(\hat{\Theta}_{ij} - \Theta_{ij}^* \right)^2 &\leq (8\underline{\rho}_n + 4\underline{\rho}_n) \mathcal{K}(\Theta^*, \tilde{\Theta}) + 8\underline{\epsilon}_n \rho_n. \end{aligned} \quad (27)$$

Using (26) and (27), we have shown that for $n \geq N$ and some constants $C, C' > 0$ depending only on C_{inf} , with probability at least $1 - 5 \exp(-C\underline{\rho}_n(n \log(k) + k^2)) - 2 \exp(-Cn\rho_n)$,

$$\left\| \hat{\Theta} - \Theta^* \right\|_2^2 \leq C \left(\underline{\rho}_n^2 n + \underline{\rho}_n \mathcal{K}(\Theta^*, \tilde{\Theta}) + \underline{\rho}_n \underline{\epsilon}_n \right).$$

We conclude the proof of Theorem 3 by noticing that $n\underline{\rho}_n^2 \leq \underline{\rho}_n \underline{\epsilon}_n$ and using that $\underline{\rho}_n = C \log(n)^{\frac{1}{5}} \rho_n$.

5.7 Proof of Proposition 1

By definition,

$$\begin{aligned} \mathcal{K}(\Theta^*, \Theta^{bc}) &= \sum_{i < j} \left(\Theta_{ij}^* \log \left(\frac{\Theta_{ij}^*}{\Theta_{ij}^{bc}} \right) + (1 - \Theta_{ij}^*) \log \left(\frac{1 - \Theta_{ij}^*}{1 - \Theta_{ij}^{bc}} \right) \right) \\ &\leq \sum_{i < j} \left(\Theta_{ij}^* \frac{\Theta_{ij}^* - \Theta_{ij}^{bc}}{\Theta_{ij}^{bc}} + (1 - \Theta_{ij}^*) \frac{\Theta_{ij}^{bc} - \Theta_{ij}^*}{1 - \Theta_{ij}^{bc}} \right) \\ &= \sum_{i < j} \frac{(\Theta_{ij}^{bc} - \Theta_{ij}^*)^2}{(1 - \Theta_{ij}^{bc}) \Theta_{ij}^{bc}} \end{aligned}$$

where the second line follows from the fact that for any $x > 0$, $\log(x) \leq x - 1$. Since for any $1 \leq i < j \leq n$, Θ_{ij}^{bc} and Θ_{ij}^* belong to $[C_{inf}\rho_n, \rho_n]$, this yields

$$\mathcal{K}(\Theta^*, \Theta^{bc}) \leq \sum_{i < j} \frac{(\Theta_{ij}^{bc} - \Theta_{ij}^*)^2}{(1 - \rho_n) C_{inf} \rho_n}.$$

Now, recall that $\Theta_{ij}^* = \rho_n W(\zeta_i, \zeta_j)$ and define $z^* : [n] \rightarrow [k]$ by $z^*(i) = \sum_{1 \leq a \leq k} a \mathbf{1} \left\{ \zeta_i \in \left[\frac{a-1}{k}, \frac{a}{k} \right] \right\}$ for any $i \in [n]$.

Moreover, define $\Theta_{ij}^{bc} = \rho_n W\left(\frac{z^*(i)}{k}, \frac{z^*(j)}{k}\right)$. Note that by definition of z^* , for any i , $\left| \zeta_i - \frac{z^*(i)}{k} \right| \leq \frac{1}{k}$. Thus

$$\begin{aligned} \mathcal{K}(\Theta^*, \Theta^{bc}) &\leq \frac{\rho_n}{C_{inf}(1 - \rho_n)} \sum_{i < j} \left(W(\zeta_i, \zeta_j) - W\left(\frac{z^*(i)}{k}, \frac{z^*(j)}{k}\right) \right)^2 \\ &\leq \frac{4\rho_n M^2}{C_{inf}(1 - \rho_n)} \sum_{i < j} \left(\frac{1}{k} \right)^{2(\alpha \wedge 1)} \end{aligned}$$

where the last equation follows from (12).

5.8 Technical lemmas

Lemma 19. For any $\Theta, \Theta' \in \mathbb{R}^{n \times n}$ and $\Pi \in [0, 1]^{n \times n}$,

$$\|\Theta - \Theta'\|_{2, \Pi}^2 \leq 8 (\|\Theta\|_\infty \vee \|\Theta'\|_\infty) \mathcal{K}_\Pi(\Theta, \Theta').$$

Proof. By definition of Bernoulli Kullback-Leibler divergence for any $0 < q, q' < 1$ we have that

$$\begin{aligned} \mathcal{K}(q, q') &= q \log \left(\frac{q}{q'} \right) + (1-q) \log \left(\frac{1-q}{1-q'} \right) \geq (\sqrt{q} - \sqrt{q'})^2 + (\sqrt{1-q} - \sqrt{1-q'})^2 \\ &\geq \frac{1}{2} \left[(\sqrt{q} - \sqrt{q'}) - (\sqrt{1-q} - \sqrt{1-q'}) \right]^2. \end{aligned}$$

Using Taylor's Theorem for some η between q and q' we get

$$\begin{aligned} \mathcal{K}(q, q') &\geq \frac{1}{2} \left[\frac{1}{2\sqrt{\eta}} (q - q') + \frac{1}{2\sqrt{1-\eta}} (q - q') \right]^2 = \frac{(q - q')^2}{8} \left[\frac{1}{\sqrt{\eta}} + \frac{1}{\sqrt{1-\eta}} \right]^2 \\ &= \frac{(q - q')^2}{8} \left[\frac{1}{\eta} + \frac{1}{1-\eta} \right] = \frac{(q - q')^2}{8} \frac{1}{\eta(1-\eta)} \geq \frac{(q - q')^2}{8(q \vee q')}. \end{aligned} \tag{28}$$

Now Lemma 19 follows from (28) and

$$\mathcal{K}_\Pi(\Theta, \Theta') = \sum_{i < j} \Pi_{ij} \mathcal{K}(\Theta_{ij}, \Theta'_{ij}).$$

□

Lemma 20. Let $\tilde{\Theta}^s$ and n_s be defined as in (10), and assume that $\gamma_n \leq \frac{1}{2}$, then

$$\mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}^s) - \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}) \leq 2\gamma_n n_s.$$

Proof.

$$\begin{aligned} \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}^s) - \mathcal{K}_\Pi(\Theta^*, \tilde{\Theta}_{ij}) &= \sum_{i < j} \Pi_{ij} \left(\Theta_{ij}^* \log \left(\frac{\tilde{\Theta}_{ij}}{\tilde{\Theta}_{ij}^s} \right) + (1 - \Theta_{ij}^*) \log \left(\frac{1 - \tilde{\Theta}_{ij}}{1 - \tilde{\Theta}_{ij}^s} \right) \right) \\ &= \sum_{i < j} \Pi_{ij} \mathbb{1} \{ \tilde{\Theta}_{ij} < \gamma_n \} \left(\Theta_{ij}^* \log \left(\frac{\tilde{\Theta}_{ij}}{\gamma_n} \right) + (1 - \Theta_{ij}^*) \log \left(\frac{1 - \tilde{\Theta}_{ij}}{1 - \gamma_n} \right) \right) \\ &\leq \sum_{i < j} \Pi_{ij} \mathbb{1} \{ \tilde{\Theta}_{ij} < \gamma_n \} (1 - \Theta_{ij}^*) \log \left(1 + \frac{\gamma_n - \tilde{\Theta}_{ij}}{1 - \gamma_n} \right) \\ &\leq \sum_{i < j} \Pi_{ij} \mathbb{1} \{ \tilde{\Theta}_{ij} < \gamma_n \} \frac{\gamma_n - \tilde{\Theta}_{ij}}{1 - \gamma_n} \leq \sum_{i < j} \Pi_{ij} \mathbb{1} \{ \tilde{\Theta}_{ij} < \gamma_n \} 2\gamma_n \leq 2n_s \gamma_n. \end{aligned}$$

□

5.9 Proof of Proposition 2

In the case of fully observed network, we alleviate notations and write

$$\begin{aligned} \mathcal{L}(\mathbf{A}; z, \mathbf{Q}) &= \sum_{i \neq j} \mathbf{A}_{ij} \log \left(\mathbf{Q}_{z(i), z(j)} \right) + (1 - \mathbf{A}_{ij}) \log \left(1 - \mathbf{Q}_{z(i), z(j)} \right), \\ \mathfrak{l}(\mathbf{A}; \alpha, \mathbf{Q}) &= \sum_{z \in \mathcal{Z}_{n,k}} \left(\prod_i \alpha_{z(i)} \right) \exp(\mathcal{L}(\mathbf{A}; z, \mathbf{Q})), \\ \text{and } \mathcal{J}(\mathbf{A}; \tau, \alpha, \mathbf{Q}) &= \log(\mathfrak{l}(\mathbf{A}; \alpha, \mathbf{Q})) - KL(\mathbb{P}_\tau(\cdot) \parallel \mathbb{P}(\cdot | \mathbf{A}, \alpha, \mathbf{Q})). \end{aligned}$$

For any $z \in \mathcal{Z}_{n,k}$ and $(\alpha, \mathbf{Q}) \in \mathcal{Q}$, we denote

$$\mathfrak{l}'(\mathbf{A}, z; \alpha, \mathbf{Q}) = \left(\prod_{i \leq n} \alpha_{z(i)} \right) \exp(\mathcal{L}(\mathbf{A}; z, \mathbf{Q}))$$

the likelihood of the parameters (α, \mathbf{Q}) and the label function z . Then, the likelihood of the stochastic block model with parameters (α, \mathbf{Q}) is given by $l(\mathbf{A}; \alpha, \mathbf{Q}) = \sum_{z \in \mathcal{Z}_{n,k}} l'(\mathbf{A}, z; \alpha, \mathbf{Q})$. Note that the likelihood functions $l(\mathbf{A}; \alpha, \mathbf{Q})$ and $l'(\mathbf{A}, z; \alpha, \mathbf{Q})$ provide lower and upper bounds on the variational objective function $\mathcal{J}(\mathbf{A}; \tau, \alpha, \mathbf{Q})$: for any parameter (α, \mathbf{Q}) and any label function $z \in \mathcal{Z}_{n,k}$,

$$l'(\mathbf{A}, z; \alpha, \mathbf{Q}) \leq \sup_{\tau \in \mathcal{T}} \exp(\mathcal{J}(\mathbf{A}; \tau, \alpha, \mathbf{Q})) \leq l(\mathbf{A}; \alpha, \mathbf{Q}). \quad (29)$$

To prove Proposition 2, we first show that $\mathbb{P}(\cdot | \mathbf{A}, \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})$, i.e. the posterior distribution of z at the variational estimator $(\hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})$, concentrates around $\delta_{z'}$, the dirac distribution at the label function $z' = \arg \max_{z: z \sim z^*} l'(\mathbf{A}, z; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})$:

$$\mathbb{P}(z' | \mathbf{A}, \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var}) = 1 - o_p(1). \quad (30)$$

Then, we show that it implies the concentration of the estimator \hat{z}^{Var} :

$$\mathbb{P}(\hat{z}^{Var} = z' | \mathbf{A}) = 1 - o_p(1). \quad (31)$$

Together (30) and (31) imply $\mathbb{P}(\hat{z}^{Var} \sim z^* | \mathbf{A}) = 1 - o_p(1)$. Since the random variable $\mathbb{P}(\hat{z}^{Var} \sim z^* | \mathbf{A})$ is bounded, Equation (31) also implies that it converges to 1 in expectation. Finally, we show that with probability going to one, the maximum likelihood estimator of the label function is equal to the true label function (up to permutation):

$$\mathbb{P}(\hat{z} \sim z^*) = 1 - o_p(1) \quad (32)$$

which concludes the proof of Proposition 2.

Proof of Equation (30)

The proof of Equation (30) relies on results proven in [8], which we recall for the sake of completeness. For any two parameters (α, \mathbf{Q}) and (α', \mathbf{Q}') in \mathcal{Q} , we say that $(\alpha', \mathbf{Q}') \in \mathcal{S}_{\alpha, \mathbf{Q}}$ if there exists a permutation σ of $\{1, \dots, k\}$ such that for any $(a, b) \in \{1, \dots, k\}^2$, $\mathbf{Q}'_{\sigma(a), \sigma(b)} = \mathbf{Q}_{a,b}$ and $\alpha'_{\sigma(a)} = \alpha_a$.

Theorem 8 (Theorem 1 in [8]). *Let (z^*, A) be generated from a stochastic block model with parameters $(\alpha^*, \mathbf{Q}^*) \in \mathcal{Q}$ such that \mathbf{Q}^0 has no identical columns and $\rho_n \gg \log(n)/n$. Then, for any $(\alpha, \mathbf{Q}) \in \mathcal{Q}$,*

$$\frac{l(\mathbf{A}; \alpha, \mathbf{Q})}{l(\mathbf{A}; \alpha^*, \mathbf{Q}^*)} = \max_{(\alpha', \mathbf{Q}') \in \mathcal{S}_{\alpha, \mathbf{Q}}} \frac{l'(\mathbf{A}, z^*; \alpha', \mathbf{Q}')}{l'(\mathbf{A}, z^*; \alpha^*, \mathbf{Q}^*)} (1 + \epsilon_n((\alpha', \mathbf{Q}'), k)) + \epsilon_n((\alpha', \mathbf{Q}'), k)$$

where $\sup_{(\alpha, \mathbf{Q}) \in \mathcal{Q}} \epsilon_n((\alpha, \mathbf{Q}), k) = o_p(1)$.

Lemma 21 (Lemma 3 in [8]). *Let (z^*, A) be generated from a stochastic block model with parameters $(\alpha^*, \mathbf{Q}^*) \in \mathcal{Q}$ such that \mathbf{Q}^0 has no identical columns and $\rho_n \gg \log(n)/n$. Then,*

$$\frac{l'(\mathbf{A}, z^*; \alpha^*, \mathbf{Q}^*)}{l(\mathbf{A}; \alpha^*, \mathbf{Q}^*)} = 1 + o_p(1).$$

Recall that $z' = \arg \max_{z: z \sim z^*} l'(\mathbf{A}, z; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})$. By definition of l and l' ,

$$\sum_{z \neq z'} l'(\mathbf{A}, z; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var}) = l(\mathbf{A}; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var}) - l'(\mathbf{A}, z'; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var}).$$

Thus

$$\frac{\sum_{z \neq z'} l'(\mathbf{A}, z; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})}{l'(\mathbf{A}, z'; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})} = \frac{l(\mathbf{A}; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})}{l'(\mathbf{A}, z'; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})} - \frac{l'(\mathbf{A}, z'; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})}{l'(\mathbf{A}, z'; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})}. \quad (33)$$

Using Lemma 21, we have that

$$\frac{l(\mathbf{A}; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})}{l'(\mathbf{A}, z'; \hat{\alpha}^{Var}, \hat{\mathbf{Q}}^{Var})} = 1 + o_p(1). \quad (34)$$

Moreover, we note that

$$\begin{aligned} \max_{(\alpha', \mathbf{Q}') \in \mathcal{S}_{\hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}}} l'(\mathbf{A}, z^*; \alpha', \mathbf{Q}') &= \max_{z \sim z^*} l'(\mathbf{A}, z; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}) \\ &= l'(\mathbf{A}, z'; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}) \end{aligned}$$

by the definition of z' . Then, applying Theorem 8, we get that

$$\frac{l(\mathbf{A}; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})}{l(\mathbf{A}; \alpha^*, \mathbf{Q}^*)} = \frac{l'(\mathbf{A}, z'; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})}{l'(\mathbf{A}, z^*; \alpha^*, \mathbf{Q}^*)} (1 + o_p(1)) + o_p(1). \quad (35)$$

Combining Equations (33), (34) and (35), we obtain that

$$\frac{\sum_{z \neq z'} l'(\mathbf{A}, z; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})}{l'(\mathbf{A}, z^*; \alpha^*, \mathbf{Q}^*)} = \frac{l'(\mathbf{A}, z'; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})}{l'(\mathbf{A}, z^*; \alpha^*, \mathbf{Q}^*)} o_p(1) + o_p(1).$$

Thus,

$$\sum_{z \neq z'} l'(\mathbf{A}, z; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}) = \max \left\{ l'(\mathbf{A}, z^*; \alpha^*, \mathbf{Q}^*), l'(\mathbf{A}, z'; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}) \right\} o_p(1). \quad (36)$$

On the one hand, using Equation (29) and the definition of $(\hat{\tau}^{VAR}, \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})$, we find that

$$\begin{aligned} l'(\mathbf{A}, z^*; \alpha^*, \mathbf{Q}^*) &\leq \sup_{\tau \in \mathcal{T}} \exp(\mathcal{J}(\mathbf{A}; \tau, \alpha^*, \mathbf{Q}^*)) \\ &\leq \exp\left(\mathcal{J}(\mathbf{A}; \hat{\tau}^{VAR}, \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})\right) \\ &\leq l(\mathbf{A}; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}). \end{aligned}$$

Also, by the definition of l and l' , we have that $l'(\mathbf{A}, z'; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}) \leq l(\mathbf{A}; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})$. Thus, Equation (36) implies

$$\sum_{z \neq z'} l'(\mathbf{A}, z; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}) = l(\mathbf{A}; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}) o_p(1). \quad (37)$$

Now, we can conclude the proof of Equation (30) by noticing that

$$\begin{aligned} \mathbb{P}(z' | \mathbf{A}, \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}) &= \frac{l'(\mathbf{A}, z'; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})}{l(\mathbf{A}; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})} \\ &= 1 - \frac{\sum_{z \neq z'} l'(\mathbf{A}, z; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})}{l(\mathbf{A}; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})} \end{aligned}$$

and using Equation (37).

Proof of Equation (31) By the definition of $\mathcal{J}(\mathbf{A}; \tau, \alpha, \mathbf{Q})$, we have that

$$KL\left(\mathbb{P}_{\hat{\tau}^{VAR}}(\cdot) \parallel \mathbb{P}(\cdot | \mathbf{A}, \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})\right) = \log\left(l(\mathbf{A}; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})\right) - \mathcal{J}(\mathbf{A}; \hat{\tau}^{VAR}, \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}).$$

Equation (29) implies that $\mathcal{J}(\mathbf{A}; \hat{\tau}^{VAR}, \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}) \geq \log\left(l(\mathbf{A}, z'; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})\right)$, so

$$KL\left(\mathbb{P}_{\hat{\tau}^{VAR}}(\cdot) \parallel \mathbb{P}(\cdot | \mathbf{A}, \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})\right) \leq \log\left(l(\mathbf{A}; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})\right) - \log\left(l(\mathbf{A}, z'; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})\right).$$

Note that Equation (30) implies

$$\log\left(l(\mathbf{A}; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})\right) - \log\left(l(\mathbf{A}, z'; \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR})\right) = o_p(1).$$

Now, using Pinsker's inequality, we see that

$$\left| \mathbb{P}_{\hat{\tau}^{VAR}}(z') - \mathbb{P}(z' | \mathbf{A}, \hat{\alpha}^{VAR}, \hat{\mathbf{Q}}^{VAR}) \right| = o_p(1).$$

We use Equation (30) and the definition of $\hat{z}^{(VAR)}$ to conclude the proof of Equation (31).

Proof of Equation (32)

Equation (32) is proven in [9]. In this work, the authors define the profile likelihood modularity $\mathcal{Q}_{LM}(A, z)$ of a label function $z \in \mathcal{Z}_{n,k}$ as

$$\mathcal{Q}_{LM}(A, z) = \frac{1}{2} \sum_{a,b} n_{ab} \left(\frac{\mathbf{O}_{ab}}{n_{ab}} \log \left(\frac{\mathbf{O}_{ab}}{n_{ab}} \right) + \left(1 - \frac{\mathbf{O}_{ab}}{n_{ab}} \right) \log \left(1 - \frac{\mathbf{O}_{ab}}{n_{ab}} \right) \right).$$

for $\mathbf{O}_{ab} = \sum_{i \in z^{-1}(a), j \in z^{-1}(b)} \mathbf{A}_{ij}$ and

$$n_{ab} = \begin{cases} |z^{-1}(a)| \times |z^{-1}(b)| & \text{if } a \neq b \\ |z^{-1}(a)| \times (|z^{-1}(a)| - 1) & \text{otherwise} \end{cases}$$

For $\hat{z}^{LM} = \arg \max_{z \in \mathcal{Z}_{n,k}} \mathcal{Q}_{LM}(A, z)$, the authors of [9] prove that under the assumptions of Proposition 2, with probability going to 1, $\hat{z}^{LM} \sim z^*$. Since maximizing $\mathcal{Q}(A, z)$ is equivalent to maximizing $\max_{\mathbf{Q}} \mathcal{L}(\mathbf{A}; \mathbf{Q}, z)$, this implies that $\hat{z} \sim z^*$ with probability going to 1.

5.10 Proof of Proposition 3

To prove Proposition 3, we show that the maximum likelihood estimator studied in Proposition 2 is equal to the restricted maximum estimator studied in Corollary 1. Proposition 3 then follows from Corollary 1.

Define $c_{min} = \min_{a,b} \mathbf{Q}_{a,b}^0$ and $c_{max} = \max_{a,b} \mathbf{Q}_{a,b}^0$. Corollary 1 implies that for some absolute constant $C > 0$,

$$\mathbb{P} \left(\left\| \hat{\Theta}^* - \hat{\Theta}^r \right\|_2^2 \leq C(c_{max}/c_{min})^2 \rho_n (k^2 + n \log(k)) \right) \rightarrow 1,$$

where the restricted maximum likelihood estimator $\hat{\Theta}^r$ is defined as

$$\begin{aligned} \hat{\Theta}_{i < j}^r &= \hat{\mathbf{Q}}_{\hat{z}^r(i)\hat{z}^r(j)}^r, \quad \hat{\Theta}_{ii}^r = 0 \\ (\hat{\mathbf{Q}}^r, \hat{z}^r) &\in \arg \min_{\mathbf{Q} \in [c_{min}\rho_n/2, 2c_{max}\rho_n]_{\text{sym}}^{k \times k}, z \in \mathcal{Z}_{n,k}} \sum_{i \neq j} \mathcal{K}(\mathbf{A}_{ij}, \mathbf{Q}_{z(i)z(j)}). \end{aligned}$$

One the other hand, Proposition 2 implies that with probability going to one, the variational estimator of the probabilities of connections $\hat{\Theta}^{VAR}$ is equal to the maximum likelihood estimator $\hat{\Theta}$ given by

$$\begin{aligned} \hat{\Theta}_{i < j} &= \hat{\mathbf{Q}}_{\hat{z}(i)\hat{z}(j)}, \quad \hat{\Theta}_{ii} = 0 \\ \text{for } (\hat{\mathbf{Q}}, \hat{z}) &\in \arg \min_{\mathbf{Q} \in \mathcal{Q}, z \in \mathcal{Z}_{n,k}} \sum_{i \neq j} \mathcal{K}(\mathbf{A}_{ij}, \mathbf{Q}_{z(i)z(j)}). \end{aligned}$$

Thus, it is enough to show that $\hat{\Theta} = \hat{\Theta}^r$ with large probability to prove Proposition 3. To do so, we show that with large probability, $\mathbf{Q}(\hat{z}) \in [c_{min}\rho_n/2, 2c_{max}\rho_n]^{k \times k}$. We define

$$n_{ab}(z) = \begin{cases} |z^{-1}(a)| \times |z^{-1}(b)| & \text{if } a \neq b \\ |z^{-1}(a)| \times (|z^{-1}(a)| - 1) & \text{otherwise} \end{cases}$$

for $z \in \mathcal{Z}_{n,k}$, and $\mathbf{Q}(z) = (\mathbf{Q}(z)_{ab})$ such that $\mathbf{Q}(z)_{ab} = \left(\sum_{i \in z^{-1}(a), j \in z^{-1}(b)} \mathbf{A}_{ij} \right) / n_{ab}(z)$. With these notations, we note that $\hat{\mathbf{Q}} = \mathbf{Q}(\hat{z})$.

Recall that $|(z^*)^{-1}(a)|$ is a sum of n independent Bernoulli random variables with mean α_a^0 . Using Bernstein's inequality 6, we find that for any a ,

$$\mathbb{P}(n\alpha_a^0 - |(z^*)^{-1}(a)| \geq 0.5n\alpha_a^0) \leq 2e^{-n\alpha_a^0/16}.$$

Thus,

$$\mathbb{P}\left(\min_a |(z^*)^{-1}(a)| \leq 0.5n \min_a \alpha_a^0\right) \leq 2ke^{-n \min_a \alpha_a^0/16}.$$

Therefore, the event $\Omega = \{\min_{a,b} n_{a,b}(z^*) \geq n^2 \min_a (\alpha_a^0)^2/5\}$ holds with probability going to 1.

Now, we show that on the event Ω , with large probability, $\mathbf{Q}(z^*) \in [c_{\min}\rho_n/2, 2c_{\max}\rho_n]^{k \times k}$. Recall that for any a, b , conditionally on z^* , $n_{ab}(z^*)\mathbf{Q}(z^*)_{ab}$ is a sum of $n_{ab}(z^*)$ independent Bernoulli random variables with mean $\rho_n \mathbf{Q}_{ab}^0$. Then, Bernstein's inequality 6 implies that for any $t > 0$

$$\mathbb{P}\left(|n_{ab}(z^*)\mathbf{Q}(z^*)_{ab} - n_{ab}(z^*)\rho_n \mathbf{Q}_{ab}^0| \geq \sqrt{2tn_{ab}(z^*)\rho_n \mathbf{Q}_{ab}^0} + \frac{2t}{3}\right) \leq 2e^{-t}.$$

Choosing $t = n_{ab}(z^*)\rho_n \mathbf{Q}_{ab}^0/16$ yields

$$\mathbb{P}\left(|n_{ab}(z^*)\mathbf{Q}(z^*)_{ab} - n_{ab}(z^*)\rho_n \mathbf{Q}_{ab}^0| \geq 0.5n_{ab}(z^*)\rho_n \mathbf{Q}_{ab}^0\right) \leq 2e^{-n_{ab}(z^*)\rho_n \mathbf{Q}_{ab}^0/16}.$$

On the event Ω , this implies that

$$\mathbb{P}\left(|n_{ab}(z^*)\mathbf{Q}(z^*)_{ab} - n_{ab}(z^*)\rho_n \mathbf{Q}_{ab}^0| \geq 0.5n_{ab}(z^*)\rho_n \mathbf{Q}_{ab}^0\right) \leq 2e^{-n^2 \rho_n \mathbf{Q}_{ab}^0 (\min_a \alpha_a^0)^2/80}.$$

A union bound yields

$$\mathbb{P}\left(\mathbf{Q}(z^*) \notin [c_{\min}\rho_n/2, 2c_{\max}\rho_n]^{k \times k}\right) \leq 2k^2 e^{-n^2 \rho_n \min_{a,b} \mathbf{Q}_{ab}^0 (\min_a \alpha_a^0)^2/80}$$

on the event Ω . Since $\mathbb{P}(\Omega) \rightarrow 1$ and $n^2 \rho_n \rightarrow +\infty$, this shows that

$$\mathbb{P}\left(\mathbf{Q}(z^*) \in [c_{\min}\rho_n/2, 2c_{\max}\rho_n]^{k \times k}\right) \rightarrow 1.$$

Now, Equation (32) shows that with probability going to 1, $\hat{z} \sim z^*$. Thus, $\mathbf{Q}(\hat{z}) \in [c_{\min}\rho_n/2, 2c_{\max}\rho_n]^{k \times k}$ with probability going to one, and the maximum likelihood estimator of the probabilities of connections between nodes coincides with the restricted maximum likelihood estimator. This concludes the proof of Proposition 3.

References

- [1] Abbe, E. and Sandon, C. (2015). Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery. In 2015 IEEE 56th Annual Symposium on Foundations of Computer Science, pages 670–688.
- [2] Agarwal, N., Bandeira, A. S., Koiliaris, K., and Kolla, A. (2017). Multisection in the Stochastic Block Model Using Semidefinite Programming, pages 125–162. Springer International Publishing, Cham.
- [3] Albert, R. and Barabási, A.-L. (2002). Statistical mechanics of complex networks. Reviews of Modern Physics, 74:47–97.
- [4] Amini, A. A., Chen, A., Bickel, P. J., and Levina, E. (2013). Pseudo-likelihood methods for community detection in large sparse networks. The Annals of Statistics, 41(4):2097 – 2122.
- [5] Amini, A. A. and Levina, E. (2018). On semidefinite relaxations for the block model. The Annals of Statistics, 46(1):149 – 179.
- [6] Bandeira, A. (2015). Random laplacian matrices and convex relaxations. Foundations of Computational Mathematics, 18.
- [7] Benyahia, O., Lameron, C., and Jeudy, B. (2017). Community detection in dynamic graphs with missing edges. In 2017 11th International Conference on Research Challenges in Information Science (RCIS), pages 372–381.
- [8] Bickel, P., Choi, D., Chang, X., and Zhang, H. (2013). Asymptotic normality of maximum likelihood and its variational approximation for stochastic blockmodels. The Annals of Statistics, 41(4):1922 – 1943.
- [9] Bickel, P. J. and Chen, A. (2009). A nonparametric view of network models and newman–girvan and other modularities. Proceedings of the National Academy of Sciences, 106(50):21068–21073.
- [10] Bleakley, K., Biau, G., and Vert, J.-P. (2007). Supervised reconstruction of biological networks with local models. Bioinformatics, 23(13):i57–i65.

- [11] Bordenave, C., Lelarge, M., and Massoulié, L. (2018). Nonbacktracking spectrum of random graphs: Community detection and nonregular Ramanujan graphs. The Annals of Probability, 46(1):1 – 71.
- [12] Candès, E. and Recht, B. (2012). Exact matrix completion via convex optimization. Foundations of Computational Mathematics, 55(6):111–119.
- [13] Celisse, A., Daudin, J.-J., and Pierre, L. (2012). Consistency of maximum-likelihood and variational estimators in the stochastic block model. Electronic Journal of Statistics, 6(none):1847 – 1899.
- [14] Chatterjee, S. (2015). Matrix estimation by Universal Singular Value Thresholding. The Annals of Statistics, 43(1):177 – 214.
- [15] Chen, K. and Lei, J. (2014). Network cross-validation for determining the number of communities in network data. Journal of the American Statistical Association, 113:241 – 251.
- [16] Clauset, A., Moore, C., and Newman, M. (2008). Hierarchical structure and the prediction of missing links in networks. Nature, 453:98–101.
- [17] Daudin, J.-J., Picard, F., and Robin, S. (2008). A mixture model for random graph. Statistics and Computing, 18:173–183.
- [18] Davenport, M. A., Plan, Y., van den Berg, E., and Wootters, M. (2014). 1-Bit matrix completion. Information and Inference: A Journal of the IMA, 3(3):189–223.
- [19] Decelle, A., Krzakala, F., Moore, C., and Zdeborová, L. (2011). Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. Physical review. E, 84:066106.
- [20] Gao, C., Lu, Y., Ma, Z., and Zhou, H. H. (2016). Optimal estimation and completion of matrices with biclustering structures. Journal of Machine Learning Research, 17(1):5602–5630.
- [21] Gao, C., Lu, Y., and Zhou, H. H. (2015). Rate-optimal graphon estimation. The Annals of Statistics, 43(6):2624 – 2652.
- [22] Giné, E. and Nickl, R. (2015). Mathematical Foundations of Infinite-Dimensional Statistical Models. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- [23] Giraud, C. and Verzelen, N. (2018). Partial recovery bounds for clustering with the relaxed K -means. Mathematical Statistics and Learning, 1(3):317–374.
- [24] Guimerà, R. and Sales-Pardo, M. (2009). Missing and spurious interactions and the reconstruction of complex networks. Proceedings of the National Academy of Sciences, 106(52):22073–22078.
- [25] Hagen, L. and Kahng, A. B. (1992). New spectral methods for ratio cut partitioning and clustering. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 11(9):1074–1085.
- [26] Hajek, B., Wu, Y., and Xu, J. (2016). Achieving exact cluster recovery threshold via semidefinite programming: Extensions. IEEE Transactions on Information Theory, 62(10):5918–5937.
- [27] Handcock, M. S. and Gile, K. J. (2010). Modeling social networks from sampled data. The Annals of Applied Statistics, 4(1).
- [28] Klopp, O. (2014). Noisy low-rank matrix completion with general sampling distribution. Bernoulli, 20(1):282 – 303.
- [29] Klopp, O., Lafond, J., Moulines, E., and Salmon, J. (2015). Adaptive multinomial matrix completion. Electronic Journal of Statistics, 9(2):2950 – 2975.
- [30] Klopp, O., Tsybakov, A. B., and Verzelen, N. (2017). Oracle inequalities for network models and sparse graphon estimation. The Annals of Statistics, 45(1):316 – 354.
- [31] Klopp, O. and Verzelen, N. (2017). Optimal graphon estimation in cut distance. Working Papers 2017-42, Center for Research in Economics and Statistics.
- [32] Koltchinskii, V., Lounici, K., and Tsybakov, A. B. (2011). Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. The Annals of Statistics, 39(5):2302 – 2329.

- [33] Kossinets, G. (2006). Effects of missing data in social networks. Social Networks, 28(3):247–268.
- [34] Kshirsagar, M., Carbonell, J., and Klein-Seetharaman, J. (2012). Techniques to cope with missing data in host–pathogen protein interaction prediction. Bioinformatics, 28(18):i466–i472.
- [35] Latouche, P., Birmelé, E., and Ambroise, C. (2011). Overlapping stochastic block models with application to the french political blogosphere. The Annals of Applied Statistics, 5(1):309–336.
- [36] Leger, J.-B., Vacher, C., and Daudin, J.-J. (2014). Detection of structurally homogeneous subsets in graphs. Statistics and Computing, 24(5):675–692.
- [37] Lei, J. (2016). A goodness-of-fit test for stochastic block models. The Annals of Statistics, 44(1):401 – 424.
- [38] Lovász, L. (2012). Large Networks and Graph Limits. American Mathematical Society colloquium publications. American Mathematical Society.
- [39] Lü, L. and Zhou, T. (2011). Link prediction in complex networks: A survey. Physica A: Statistical Mechanics and its Applications, 390(6):1150–1170.
- [40] Mariadassou, M. and Tabouy, T. (2020). Consistency and asymptotic normality of stochastic block models estimators from sampled data. Electronic Journal of Statistics, 14(2):3672 – 3704.
- [41] Massoulié, L. (2014). Community detection thresholds and the weak ramanujan property. In Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing, STOC ’14, page 694–703, New York, NY, USA. Association for Computing Machinery.
- [42] Matias, Catherine and Robin, Stéphane (2014). Modeling heterogeneity in random graphs through latent space models: a selective review*. ESAIM: Proc., 47:55–74.
- [43] McSherry, F. (2001). Spectral partitioning of random graphs. In Proceedings 42nd IEEE Symposium on Foundations of Computer Science, pages 529–537.
- [44] Mossel, E., Neeman, J., and Sly, A. (2016). Consistency thresholds for the planted bisection model. Electronic Journal of Probability, 21(none):1 – 24.
- [45] Negahban, S. and Wainwright, M. J. (2011). Estimation of (near) low-rank matrices with noise and high-dimensional scaling. The Annals of Statistics, 39(2):1069 – 1097.
- [46] Newman, M. E. J. (2006). Modularity and community structure in networks. Proceedings of the National Academy of Sciences, 103(23):8577–8582.
- [47] Olhede, S. C. and Wolfe, P. J. (2014). Network histograms and universality of blockmodel approximation. Proceedings of the National Academy of Sciences, 111(41):14722–14727.
- [48] Picard, F., Miele, V., Daudin, J., Cottret, L., and Robin, S. (2009). Deciphering the connectivity structure of biological networks using MixNet. BMC Bioinformatics, 10:1–11.
- [49] Rohe, K., Chatterjee, S., and Yu, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. The Annals of Statistics, 39(4):1878 – 1915.
- [50] Tabouy, T., Barbillon, P., and Chiquet, J. (2020). Variational inference for stochastic block models from sampled data. Journal of the American Statistical Association, 115(529):455–466.
- [51] Vershynin, R. (2012). Introduction to the non-asymptotic analysis of random matrices, page 210–268. Cambridge University Press.
- [52] Wang, Y. X. R. and Bickel, P. J. (2017). Likelihood-based model selection for stochastic block models. The Annals of Statistics, 45(2):500 – 528.
- [53] Wasserman, S. and Faust, K. (1994). Social Network Analysis: Methods and Applications. Structural Analysis in the Social Sciences. Cambridge University Press.
- [54] Xu, J. (2018). Rates of convergence of spectral methods for graphon estimation. In Dy, J. and Krause, A., editors, Proceedings of the 35th International Conference on Machine Learning, volume 80 of Proceedings of Machine Learning Research, pages 5433–5442. PMLR.

- [55] Yamanishi, Y., Vert, J.-P., and Kanehisa, M. (2004). Protein network inference from multiple genomic data: a supervised approach. Bioinformatics, 20 Suppl 1:i363–70.
- [56] Yan, B. and Gregory, S. (2012). Finding missing edges in networks based on their community structure. Physical review. E, 85:056112.
- [57] Zhang, Y., Levina, E., and Zhu, J. (2017). Estimating network edge probabilities by neighbourhood smoothing. Biometrika, 104(4):771–783.
- [58] Zhao, Y., Wu, Y.-J., Levina, E., and Zhu, J. (2017). Link prediction for partially observed networks. Journal of Computational and Graphical Statistics, 26(3):725–733.