

Uniform estimation in stochastic block models is slow*

Ismaël Castillo

*Laboratoire Probabilités, Statistique et Modélisation,
Sorbonne Université*

e-mail: ismael.castillo@sorbonne-universite.fr

and

Peter Orbanz

*Gatsby Computational Neuroscience Unit,
University College London*

e-mail: porbanz@gatsby.ucl.ac.uk

Abstract: We explicitly quantify the empirically observed phenomenon that estimation under a stochastic block model (SBM) is hard if the model contains classes that are similar. More precisely, we consider estimation of certain functionals of random graphs generated by a SBM. The SBM may or may not be sparse, and the number of classes may be fixed or grow with the number of vertices. Minimax lower and upper bounds of estimation along specific submodels are derived. The results are nonasymptotic and imply that uniform estimation of a single connectivity parameter is much slower than the expected asymptotic pointwise rate. Specifically, the uniform quadratic rate does not scale as the number of edges, but only as the number of vertices. The lower bounds are local around any possible SBM. An analogous result is derived for functionals of a class of smooth graphons.

MSC2020 subject classifications: Primary 62G20.

Keywords and phrases: Stochastic blockmodel, semiparametric estimation of functionals, minimax rates, spectral clustering, graphon model.

Received September 2021.

Contents

1	Introduction	2948
2	Preliminaries and notation	2951
3	Main results: lower bounds	2954
4	Further results: upper bounds and smooth graphons	2958
5	Discussion	2965
6	Proofs of the lower bounds in SBMs	2965
7	Proofs for results on graphon functionals	2974
A	Upper bounds computable in polynomial-time	2979
B	Extension to sparse graphs	2982

*I. C.'s work is supported by ANR grant ANR-17-CE40-0001 (BASICS).

C Remaining proofs: likelihood-based upper bounds	2984
D Remaining proofs: lower and upper bounds in the sparse case	2992
Acknowledgments	2998
References	2998

1. Introduction

Network data occurs in a range of fields, and its analysis has become a highly interdisciplinary effort [14, 18, 24, 28]. In statistical network analysis, two classes of models have recently received particular attention: Graphon models [9, 21, 29], and the subclass of stochastic block models (SBMs) [1, 7, 17]. The results of this paper show, informally speaking, that estimation under a SBM becomes difficult if the parameters specifying two classes are close to each other.

SBM and graphon models parametrize a random graph by a symmetric measurable function w , which can be interpreted as representing an adjacency matrix in the limit of infinite graph size [9]. In a SBM, the function is in particular piece-wise constant. Examples of statistical problems arising in this field include *estimation problems* (see below), *class label recovery* [5, 25, 26, 31, 36], and *signal detection*, which refers to testing for the presence of a signal in settings where observed data constitutes a network or array [3, 4, 10, 33].

We consider estimation problems. SBMs label each vertex in a graph with a category (a “community”), and this labelling is typically not observed. There is a substantial body of work on rates of estimation in such models [1, 2, 7, 8, 11, 12]. This literature considers asymptotic pointwise rates, and shows that, informally speaking, estimators of finite-dimensional statistics can converge quickly even if the labelling of vertices is not observed. Estimation of the entire function w has also been studied [16, 23, 34]. In a case where this parameter has infinite dimension and is estimated in a uniform way, Gao, Lu, and Zhou [16] show that not observing labels slows the rate. Our results show that that is not a consequence of the nonparametric setting: The best uniform (rather than pointwise) rate for estimating a finite-dimensional statistic—even a very simple one, and under a very simple parametric SBM—is slow. The same holds for a simple, one-dimensional functional of a smooth, infinite-dimensional parameter function w .

1.1. An informal overview

The remainder of this section provides an informal overview of our results. Rigorous definitions and statements follow in the next sections. A SBM is defined by two parameters, a probability distribution π on k categories (which we regard as a vector in $[0, 1]^k$), and a matrix $M \in [0, 1]^{k \times k}$. The model generates undirected random graphs $G_n(\pi, M)$ of any size $n \in \mathbb{N}$: Each vertex $i \leq n$ is assigned a category $\varphi(i) \in \{1, \dots, k\}$ drawn randomly from π , and an edge between vertices i and j is then added with probability $M_{\varphi(i)\varphi(j)}$. (A proper definition follows in Section 2.)

The main results. SBMs pose an estimation problem: Given an observed graph G_n with n vertices, estimate π and M . The purpose of our work is to show that, loosely speaking, the estimation problem can be harder than it appears, or indeed than previous results may suggest at first glance. Our results are phrased as minimax lower bounds: Suppose \mathcal{S} is a set of parameter pairs (π, M) . A minimax bound specifies a decreasing function τ such that, informally,

$$\inf_{\text{all estimators}} \sup_{(\pi, M) \in \mathcal{S}} \mathbb{E}_{\pi, M} [\|\text{estimate of } (\pi, M) \text{ computed from } G_n(\pi, M) - (\pi, M)\|^2] \geq \tau(n).$$

That is, given an observed graph with n vertices, there exists no estimator whose quadratic risk is smaller than $\tau(n)$ for all parameter values in \mathcal{S} . Since the supremum means that shrinking \mathcal{S} will not increase the lower bound, it can suffice to consider a subclass $\mathcal{S}' \subset \mathcal{S}$ of parameters—any lower bound for \mathcal{S}' is also a lower bound for \mathcal{S} . Indeed, we will see that one can obtain a meaningful bound by choosing a very small subclass with one degree of freedom, where π is fixed to the uniform distribution, and the matrix M is a function $M(\theta)$ of a one-dimensional surrogate parameter $\theta \in [-1/2, 1/2]$. The statement above then takes the form

$$\inf_{\text{all estimators}} \sup_{\theta \in [-1/2, 1/2]} \mathbb{E}_{\theta} [\text{estimate of } \theta \text{ computed from } G_n(\pi, M(\theta)) - \theta]^2 \geq \tau(n). \quad (1)$$

Our main result shows that the relevant lower bound is

$$\tau(n) = \frac{\text{constant} \cdot k}{n},$$

where $k = k(n)$ is permitted to depend on n . This is Theorem 1 (for the simplest case $k = 2$) and Theorem 2 (for $k = k(n) \geq 2$). Indeed our proofs imply a stronger statement:

- The results are completely non-asymptotic, and it is possible to explicitly determine numerical values for all relevant constants. See Remark 2 for an example.
- The minimax bound holds *locally*, not just globally: In principle, a slow minimax rate may be caused by just a few “pathological” points in the set \mathcal{S} . One can ask whether the rate τ can be improved by removing a small part of \mathcal{S} . That is not the case here: Shrinking the set of all SBM parameter pairs (π, M) to any open Euclidean neighborhood of any specific pair still results in the same rate (see Section 3.3). Informally, every region of parameter space contains parameters that prevent the rate from improving.

SBMs are often used in “sparse” forms, and we verify in Appendix B that the result also applies in the sparse case. Since sparsification reduces the amount

of available data, it slows convergence. Theorems 8 and 9 show the rate in the sparse setting also scales linearly in the expected number of edges.

Interpretation. If we were to simplify the estimation problem artificially by assuming that the assignment variables $\varphi(i)$ are observed, π could be estimated at rate $1/\sqrt{n}$ and M at rate $1/n$, both by computing sample averages. (The rates differ since π is estimated from n vertices, whereas M is estimated from edges, and the expected number of edges grows quadratically with n .) Since our bounds are phrased in terms of a quadratic risk, both rates must be squared: in the sequel, a bound $\tau(n) \approx 1/n$ as above is referred to as *slow rate*; by contrast, a *fast rate* corresponds to $\tau(n) \approx 1/n^2$.

One must distinguish uniform rates (which hold uniformly over sets of parameters) and asymptotic pointwise rates (where asymptotics in n are considered at just a given point). Previous work has established that estimation of M can be fast even if the $\varphi(i)$ are not observed. For example, a remarkable result of [7] shows that, if $\hat{\pi}$ and \hat{M} are chosen as certain profile maximum likelihood estimators, then, as $n \rightarrow \infty$,

$$\sqrt{n}(\hat{\pi} - \pi) \rightarrow Z_\pi \quad \text{and} \quad n(\hat{M} - M) \rightarrow Z_M \quad \text{in distribution.} \quad (2)$$

where Z_π and Z_M are multivariate Gaussian variables. This holds up to label switching (see Section 2), and requires that the columns of M are “not too similar”. Related results can be found in [1, 11, 12]. Since this result does not use a quadratic risk, it can be paraphrased informally as:

- Under suitable conditions on the model, the matrix M can be estimated, at least asymptotically and pointwise, at a fast rate.

It has long been recognized in statistics that pointwise asymptotic rates can be hard to interpret: As (π, M) runs through some set \mathcal{S} , the constants implicit in the rate may change locally around a given parameter as well as with n , and if they do so quickly enough, that results in an effective change in the rate. The Hodges phenomenon, for example, illustrates that highly pathological behavior of an estimator may only be visible in its uniform rate, whereas the asymptotic pointwise rate suggests good performance [see e.g. Section 8 and Figure 8.1 in 32]. Our result says:

- If measured uniformly over any given neighborhood in parameter space, the best achievable rate for connectivity parameters (i.e. for M) is always slow.

In other words, the change of constants is indeed an issue here, and makes the rate drop from a fast to a slow one. Since the minimax bound is local, this problem cannot be avoided by removing some (fixed) parameters (π, M) .

Further results. Section 4 and the Appendix provide additional results on achievability, i.e. upper bounds to complement the lower ones, and on graphon models and sparse graphs.

Upper bounds. Like most lower bound results, Theorems 1 and 2 do not show whether the bound τ is achievable—that is, the convergence rate of any

actual estimator could be even slower than τ . To show that a rate is achievable uniformly, one has to specify an estimator whose uniform risk matches the lower bound up to constants. Estimators for SBMs and their convergence rates are subject of a substantial literature, but these rates are, once again, generally pointwise. Obtaining a tight uniform upper bound for arbitrary SBMs is beyond the scope of this work, but we do consider the “hard” one-dimensional model (1), and show the following:

- For estimation of θ in (1), the rate τ is achieved by a type of maximum likelihood estimator. That is shown in Theorem 3 (for $k = 2$) and in Theorem 4 (for general k), in Section 4.1.

However, this estimator is not generally computable in polynomial time, which raises the additional question whether the problem exhibits a computational gap—that is, whether this is a problem where a sample of size n contains enough information to achieve a given rate $\tau(n)$, but this information cannot be extracted in polynomial time, and every practically computable estimate converges at a slower rate. In this context, we show:

- Under additional conditions, a spectral estimator (based on work of Lei and Zhu [26]) achieves τ , and is computable in polynomial time. (See Theorem 5 in Section 4.2 for $k = 2$ classes, and the appendix for the general case $k \geq 2$.)

Thus, in the submodel specified by the conditions, there is no computational gap. We do not know at present whether the same holds for general SBMs.

Smooth graphons. SBMs are a special case of so-called graphon models, which parametrize a random graph $G_n(w)$ on n vertices by a function w of a certain form. In SBMs, w is piece-wise constant. Section 4.4 instead considers a class \mathcal{S} of smooth graphons. It is known that uniform estimation of such a graphon w from $G_n(w)$ is only possible at a slow rate [16, 23]. Theorem 6 considers a simple, real-valued statistic $\vartheta(w)$ that can be read as a form of standard deviation. It shows that

$$\inf_{\text{all estimators}} \sup_{w \in \mathcal{S}} \mathbb{E}_w [\text{estimate of } \vartheta(w) \text{ computed from } G_n(w) - \vartheta(w)]^2 \geq \text{constant} \cdot \frac{1}{n} .$$

In other words, even if the infinite-dimensional quantity w is substituted by the much simpler, one-dimensional quantity ϑ , the rate is still slow. In this sense, Theorem 6 can be seen as a semiparametric counterpart to the nonparametric results of [16].

2. Preliminaries and notation

This section defines the models we consider, and briefly reviews some related background.

Notation. We abbreviate $[k] := \{1, \dots, k\}$, so that $[k]^n$ is the set of all mappings $\{1, \dots, n\} \rightarrow \{1, \dots, k\}$. For a subset A of the integers, $|A|$ denotes cardinality. If M is a square matrix, $\|M\|_F$ is its Frobenius norm and $\|M\|_{S_p}$ its spectral norm. Let $\text{Be}(p)$ be a shorthand for the Bernoulli(p) distribution. By $ER(p)$, we denote the law of an Erdős-Renyi random graph edge probability p over n nodes, that is, $ER(p) = ER(p, n) = \otimes_{i < j \leq n} \text{Be}(p)$, where \otimes denotes a tensor product of distributions.

Stochastic block models. Consider sampling at random an undirected, simple graph G on the vertex set $\mathcal{V} = \{1, \dots, n\}$ as follows. Fix some $k \in \{1, \dots, n\}$. Let $\pi = (\pi_1, \dots, \pi_k)$ be a probability distribution on the set $\{1, \dots, k\}$, with π identified as a line vector of size k . Let $M := (M_{lm})$ be a symmetric $k \times k$ matrix with elements $M_{lm} \in [0, 1]$. To sample a graph G , we generate its adjacency matrix $X = (X_{ij})_{i, j \in \mathcal{V}}$. Since G is undirected, it suffices to sample entries with $i < j$,

1. For each vertex $i \in \mathcal{V}$, independently generate a label $\varphi(i) \sim \pi$.
2. For each pair $i < j$ in \mathcal{V} , independently sample from the distribution $X_{ij} \mid \varphi(i), \varphi(j) \sim \text{Be}(M_{\varphi(i)\varphi(j)})$.

In this notation, φ is a (random) mapping $\varphi : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ that attributes a label to each node of the graph. It is random because labels are by definition randomly sampled. The distribution $P_{\pi, M}$ so defined on the set of undirected, simple graphs is called a *stochastic blockmodel* of order k with parameters π and M . One can also write

$$\begin{aligned} (\varphi(1), \dots, \varphi(n)) &\sim \pi^{\otimes n} \\ (X_{ij})_{i < j} \mid \varphi &\sim \bigotimes_{i < j} \text{Be}(M_{\varphi(i)\varphi(j)}), \end{aligned} \quad (3)$$

where $\pi^{\otimes n} = \pi \otimes \dots \otimes \pi$, and here and in the sequel $i < j$ refers to all pairs of indices $(i, j) \in \mathcal{V}^2$ with $i < j$. Any given φ partitions the vertex set $\{1, \dots, n\}$ into k distinct classes. We call π the *proportions* vector and M the matrix of *connectivity* parameters.

Graphon models. SBMs can be regarded as a special case of a more general class of random graphs, parametrized by the set of all measurable functions $w : [0, 1]^2 \rightarrow [0, 1]$ that are symmetric, i.e. $w(x, y) = w(y, x)$. Any such w defines a random graph G : denoting by $\text{Unif}[0, 1]$ the uniform distribution on $[0, 1]$, and $(U_i)_i = (U_i)_{1 \leq i \leq n}$, set

$$\begin{aligned} (U_i)_i &\sim \text{Unif}[0, 1]^{\otimes n} \\ (X_{ij})_{i < j} \mid (U_i)_i &\sim \bigotimes_{i < j} \text{Be}(w(U_i, U_j)). \end{aligned} \quad (4)$$

The law P_w of the graph G defined by the random matrix X in (4) is called a *graphon model* [9]. SBMs are recovered by choosing w as a histogram: subdivide

the unit interval into k intervals $I_s := [\sum_{i<s} \pi_i, \sum_{i\leq s} \pi_i)$ of respective lengths π_s , and set

$$w(x, y) := M_{ij} \quad \text{for } x \in I_i, y \in I_j . \tag{5}$$

Then $P_w = P_{\pi, M}$. In a graphon model, the continuous vertex labels U_i are almost surely distinct; in a stochastic block model, labels coincide whenever two vertices belong to the same class. Thus, the SBM labels can be regarded as discretization of graphon labels. Conversely, any graphon can be approximated by a sequence of stochastic blockmodels of increasing order k ; indeed, the set of stochastic blockmodels—that is, of graphons of the form (5) for all k, π and M —is dense in the set of functions w endowed with its natural topology [see e.g. 22, for details]. This idea can be used to construct SBM-valued estimators for graphons [16, 34]. SBMs and graphon models both generalize to directed graphs, by dropping the symmetry constraints on π and w , and requiring only $i \neq j$ rather than $i < j$; in the following, we consider only the undirected case.

Label switching and identifiability. The distribution (4) remains invariant if w is replaced by $w \circ g$, for any measure-preserving transformation g of $[0, 1]$: $P_w = P_{\tilde{w}}$ for $\tilde{w}(x, y) = w(g(x), g(y))$. More generally, two graphons w and w' are considered equivalent if $P_w = P_{w'}$. The equivalence class $\langle w \rangle$ of w is called a graph limit. Similarly in (3), if σ is a fixed arbitrary permutation of $\{1, \dots, k\}$, with permutation matrix Σ , then $P_{\pi, M} = P_{\pi\Sigma, \Sigma M \Sigma^T}$. The parameters of the SBM can only be recovered up to label switching. We refer to [1] and [11] for detailed identifiability statements.

Fixed and random design. In models (3)-(4), the latent variables, respectively φ and U , are random. Sometimes, a slightly different version of the model is considered, where φ and U are still unobserved, but fixed, non-random quantities. For instance, under this setting (3) becomes

$$(X_{ij})_{i<j} \sim \bigotimes_{i<j} \text{Be}(M_{\varphi(i)\varphi(j)}),$$

for a given, unknown, $\varphi : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$, and the data distribution is denoted $P_{\varphi, M}$. Such models will be referred to as *fixed design* SBM and *random design* SBM respectively. The term SBM as used in the literature typically refers to a random design. Some theoretical arguments simplify in the fixed design case, for which the data distribution is a product measure, rather than a mixture of products measures. Most results below are obtained for both cases.

Mixture interpretation. The n -tuple (U_i) in a graphon model, or, equivalently, the mapping φ in a SBM, are in general not observed, and can hence be interpreted as latent variables. In other words, the distribution of the data $(X_{ij})_{i<j}$ is a *mixture*. The mixture representation is useful to relate fixed and random designs to each other. In the random design case, we have

$$P_{\pi, M} = \sum_{\varphi \in [k]^n} \mu_{\pi}[\varphi] \bigotimes_{i<j} \text{Be}(M_{\varphi(i)\varphi(j)}), \tag{6}$$

where $\mu_\pi[\varphi] = \prod_{l=1}^k \pi_l^{N_l(\varphi)}$ and $N_j(\varphi) = \sum_{i=1}^n \mathbb{1}_{\varphi(i)=j}$ is the number of times the label j is present. In the fixed design model, the labels given through φ are also unobserved, but fixed, so that the distribution is $P_{\varphi, M}$ given by

$$P_{\varphi, M} = \bigotimes_{i < j} \text{Be}(M_{\varphi(i)\varphi(j)}).$$

3. Main results: lower bounds

In this section, we first construct in Section 3.1 natural submodels of a SBM with $k = 2$ along which the two classes become close, and derive an estimation lower bound in terms of the quadratic risk for the submodel parameter. We then consider in Section 3.2 the more general setting of a SBM with k classes ‘containing’ the previously constructed difficult submodel and derive a minimax estimation lower bound in this setting, which is local around any possible SBM of this type, as we discuss in more detail in Section 3.3.

3.1. The case $k = 2$

Consider the set of distributions

$$\mathcal{M} = \{P_\theta := P_{e, Q^\theta}, \quad \theta \in [-1/2, 1/2]\}, \tag{7}$$

where e and Q^θ are given by

$$e = \left[\frac{1}{2}, \frac{1}{2} \right] \tag{8}$$

$$Q^\theta = \begin{bmatrix} \frac{1}{2} + \theta & \frac{1}{2} - \theta \\ \frac{1}{2} - \theta & \frac{1}{2} + \theta \end{bmatrix}. \tag{9}$$

The set \mathcal{M} is a 1-dimensional submodel of the set of all SBMs with at most two classes. For $\theta = 0$ the matrix Q^0 is degenerate and the model is simply an Erdős-Reyni graph model with parameter $1/2$, that is all edges are independent and have a probability $1/2$ of being present. SBMs with connectivity matrices that—like Q^θ above—specify only two, one for intra-group and one for between-group connections, are known as affiliation models [e.g. 2, 3, 27].

Theorem 1. *Consider a stochastic blockmodel (3) with $k = 2$ specified by \mathcal{M} , that is $P_\theta = P_{e, Q^\theta}$ with e, Q^θ given by (8)-(9). There exists a constant $c_1 > 0$ such that for all $n \geq 2$,*

$$\inf_T \sup_{\theta \in [-1/2, 1/2]} E_\theta [T(X) - \theta]^2 \geq \frac{c_1}{n},$$

where the infimum is taken over all estimators T of θ in the model \mathcal{M} .

Proof. See Section 6.1. □

Theorem 1 states that, even in a very simple SBM with $k = 2$ classes and only one unknown parameter in its connectivity matrix, the minimax estimation rate is no faster than $1/n$. This is no contradiction to the fast rate obtained by Bickel et al. [7] (meaning a $1/n$ rate for the convergence in distribution but a $1/n^2$ rate for the quadratic risk): the latter is a pointwise asymptotic result, and assumes that no two lines of the connectivity matrix are the same, whereas Theorem 1 is nonasymptotic and uniform. It shows that the rate in a two-class model changes for distributions close to an Erdős-Renyi model ($k = 1$); informally, models close to the ‘boundary’ are harder to estimate. We note the result does not require the sub-model \mathcal{M} to include the Erdős-Renyi model; see the remark below. The phenomenon is reminiscent of effects familiar from community detection, where matrices similar to (9) naturally arise as most difficult submodels. Community detection is a *testing* problem, though, as opposed to the *estimation* problem considered here. For a different but related result in the very sparse case, see [27].

Remark 1 (different parameter choices). One can easily check that the result of Theorem 1 remains unchanged if instead of $1/2$ in the matrix Q^θ in (9), another number $a_0 \in (0, 1)$ is used. If a_0 is bounded away from 0 and 1, assuming $\min(a_0, 1 - a_0) \geq \rho > 0$, then the result is only modified by constants. Also, if the proportions vector π is of the form $[b, 1 - b]$ with $b > 0$, similar results continue to hold, provided the matrix Q^θ is replaced by

$$Q_b^\theta = \begin{bmatrix} \frac{1}{2} + c_b\theta & \frac{1}{2} - d_b\theta \\ \frac{1}{2} - d_b\theta & \frac{1}{2} + c_b\theta \end{bmatrix}$$

for suitable constants c_b, d_b that depend on b (one can take e.g. $c_b = 1 - b$ and $d_b = b$).

Remark 2 (numerical constants). In Theorem 1, one can take $c_1 = 1/107$; additionally, the supremum can be restricted to $(-\theta_n, \theta_n)$ for

$$\theta_n = \frac{c_0}{\sqrt{n}} \quad \text{and} \quad c_0 = \frac{1}{3 \cdot 2^{3/4}} \approx 0.56.$$

Moreover, the proof implies that one can restrict the supremum to a set not actually containing $\theta = 0$, but rather two points close enough to $\theta = 0$, namely $\theta_1 = c_1/\sqrt{n}$, $\theta_2 = c_2/\sqrt{n}$ for suitably chosen, fixed constants $c_1, c_2 > 0$.

Fixed design. A result similar to Theorem 1 holds for fixed designs. In this case, the map φ is deterministic, and the model can be written as $\mathcal{M}_F = \{P_{\theta, \varphi} := P_{\varphi, Q^\theta}, \theta \in [-1/2, 1/2], \varphi \in [2]^n\}$. Expectations with respect to the measures P_θ and $P_{\theta, \varphi}$ are denoted respectively E_θ and $E_{\theta, \varphi}$. We then have

$$\inf_{T_f} \sup_{\theta \in [-1/2, 1/2], \varphi \in [2]^n} E_{\theta, \varphi} [T_f(X) - \theta]^2 \geq \frac{c_1}{n},$$

where the infimum is taken over all estimators T_f of θ in the fixed design model. The proof is the same as for Theorem 1, see Section 6.1.

3.2. The general case

We now consider an arbitrary number k of classes. Above, we have perturbed a SBM with $k = 2$ classes around an Erdős-Renyi model. We now similarly perturb a k -class SBM around one with $k - 1$ classes. The connectivity matrix of a SBM with at most $k - 1$ classes is of the form, for $a_0, a_i, b_{ij} \in [0, 1]$ for $i, j \in [k - 2]$,

$$M = \begin{bmatrix} a_0 & a_1 & \cdots & a_{k-2} \\ a_1 & b_{11} & \cdots & b_{1k-2} \\ \vdots & \vdots & & \vdots \\ a_{k-2} & b_{1k-2} & \cdots & b_{k-2k-2} \end{bmatrix}. \quad (10)$$

For simplicity of notation and easy comparison with Section 3.1, we assume $a_0 = 1/2$ throughout. Results are easily adapted to the case $a_0 \in (0, 1)$, requiring only that a_0 be bounded away from 0 and 1. If needed, one can ensure the number of classes is exactly $k - 1$ by requiring no two rows of M coincide, which we require only in Theorems 7 and 8, which describe the behavior of spectral estimators.

We consider 1-dimensional submodels in the parameter space of connectivity matrices: Set

$$e_k = \left[\frac{1}{k}, \dots, \frac{1}{k} \right], \quad (11)$$

and, for coefficients $\{a_i\}, \{b_{ij}\}$ as above, define

$$M^\theta = \begin{bmatrix} \frac{1}{2} + \theta & \frac{1}{2} - \theta & a_1 & \cdots & a_{k-2} \\ \frac{1}{2} - \theta & \frac{1}{2} + \theta & a_1 & \cdots & a_{k-2} \\ a_1 & a_1 & b_{11} & \cdots & b_{1k-2} \\ \vdots & \vdots & \vdots & & \vdots \\ a_{k-2} & a_{k-2} & b_{1k-2} & \cdots & b_{k-2k-2} \end{bmatrix} = \begin{bmatrix} Q^\theta & A \\ A^T & B \end{bmatrix}, \quad (12)$$

where

$$Q^\theta = \begin{bmatrix} \frac{1}{2} + \theta & \frac{1}{2} - \theta \\ \frac{1}{2} - \theta & \frac{1}{2} + \theta \end{bmatrix}$$

and

$$A = \begin{bmatrix} a_1 & a_2 & \cdots & a_{k-2} \\ a_1 & a_2 & \cdots & a_{k-2} \end{bmatrix}, \quad B = \begin{bmatrix} b_{11} & \cdots & b_{1k-2} \\ \vdots & & \vdots \\ b_{1k-2} & \cdots & b_{k-2k-2} \end{bmatrix}.$$

Thus, M^θ is a symmetric $k \times k$ matrix, obtained from M by replacing the scalar coefficient a_0 by the 2×2 matrix Q^θ , and repeating the vector $(a_i)_{1 \leq i \leq k-2}$.

The number of nodes in a given class will be specified as follows. For simplicity, we choose the proportions vector π in (3) equiproportional and equal to e_k in (11). (As in the case $k = 2$, analogous results can be obtained if the proportions are of similar sizes.) Consider the model defined by

$$\mathcal{M}_k = \{P_\theta := P_{e_k, M^\theta}, \quad \theta \in [-1/2, 1/2]\}, \quad (13)$$

for e_k, M^θ as in (11)-(12). This is a 1-dimensional submodel of the set of all SBMs with at most k classes. For $\theta = 0$, the matrix M^0 again has two identical rows, and the model becomes a SBM with at most $k-1$ classes, with connectivity matrix given by M defined in (10). By E_θ , we denote the expectation under P_θ in the model \mathcal{M}_k given by (13).

Theorem 2. Consider a stochastic blockmodel (3) with $k \geq 2$ classes specified by \mathcal{M}_k in (13), that is $P_\theta = P_{e_k, M^\theta}$ with e_k, M^θ given by (11)-(12), for fixed matrices A, B with arbitrary coefficients. There exists a constant $c_3 = c_3(\rho) > 0$, independent of A, B , such that, for all $n \geq 12k$,

$$\inf_T \sup_{\theta \in [-1/2, 1/2]} E_\theta [T(X) - \theta]^2 \geq c_3 \frac{k}{n},$$

where the infimum is taken over all estimators T of θ in the model \mathcal{M}_k .

Proof. See Section 6.2. □

Fixed design. A similar result holds for the fixed design case, assuming that classes, given by the mapping φ , are balanced in the following sense. Let Σ_e denote the set of maps $\varphi \in [k]^n$ such that, for some constants c_1, c_2 , for any $1 \leq j \leq k$,

$$c_1 \frac{n}{k} \leq |\varphi^{-1}(j)| \leq c_2 \frac{n}{k}.$$

The set Σ_e thus consists of those maps φ that produce k classes all of size of order n/k . Then the conclusion of Theorem 2 still holds, provided E_θ is replaced by $E_{\theta, \varphi}$, and the supremum taken over $\theta \in [-1/2, 1/2]$ and $\varphi \in \Sigma_e$ as defined just above.

3.3. Some comments on the results

Theorem 2 establishes that the minimax estimation rate of θ in model (13) is at best of the order k/n , uniformly over k and n . An intuitive explanation for this particularly slow rate is as follows: the phenomenon observed for $k = 2$ is still present but this time the part of the matrix Z^θ containing information about θ is smaller, as only of the order $2/k$ of the nodes will be assigned to classes 1 or 2, which are the elements of the connectivity matrix that depend on θ .

An important point is that this lower bound is minimax *local* (as opposed to more commonly proved minimax global results) that is, not only does this slow rate occur around one specific least-favorable point in the parameter space, it does occur around *any* point. More precisely: If we start with any $k \geq 2$, any proportions vector, and any connectivity matrix M as in (10) with $k-1$ classes, there exists at least one submodel around M , namely \mathcal{M}_k in (13), such that estimation of a connectivity parameter in M cannot be faster than k/n . In Theorem 1, the model given by $\theta = 0$ is an Erdős-Renyi graph, which raises the question whether the slow rate in Theorem 1 is a consequence of the distinguished properties of the Erdős-Renyi model. This is not the case. Proving

such a local bound makes the proof of Theorem 2 more involved in the random design case, as one has to quantify the L^1 -distance between two mixtures of probability measures, instead of between one fixed measure and a mixture as is often the case in proving minimax global bounds.

It is interesting to compare the rate in Theorem 2 to the one that would be obtained if the labels were observed. If k is fixed, Lemma 2 in Bickel et al. [7] gives a quadratic rate of order $1/n^2$ for connectivity parameters when labels are observed. This result can be easily adapted to the case where k possibly grows with n , say in an asymptotic setting with $n \rightarrow \infty$ and $k/n \rightarrow 0$, leading to a quadratic rate of order $(k/n)^2$. The *uniform* rate in Theorem 2 is the square-root of this rate and thus much slower.

4. Further results: upper bounds and smooth graphons

In this Section, we complement our main results by upper bounds (under some conditions when $k \geq 3$) and results for certain smooth graphons, which can be seen as a continuous analogue of the results for the SBM parameter θ .

We establish upper bounds that show that the lower bounds in the previous section can be matched for certain subsets of connectivity matrices. In the case of $k \geq 3$ classes, the conditions are arguably somewhat restrictive and can probably be improved. However, since the lower bounds are proved to be local around any possible SBM containing two classes that are close, the rate, if not matched, can only become worse. As we show below, some conditions are in fact necessary. Indeed, we give an example in Section 4.3 where the rate drops further, illustrating the difficulty of the estimation problem.

4.1. Upper bounds via maximum likelihood

Theorems 1 and 2 provide lower bounds. There are corresponding, matching upper-bound, which we obtain next.

The case $k = 2$. We define an estimator of θ as follows. For any σ an element of $[2]^n$, i.e. for any mapping $\{1, \dots, n\} \rightarrow \{1, 2\}$, define

$$2Z_n(\sigma, X) := - \sum_{i < j, \sigma(i) = \sigma(j)} (1 - 2X_{ij}) + \sum_{i < j, \sigma(i) \neq \sigma(j)} (1 - 2X_{ij}). \quad (14)$$

Maximising (14) in σ leads to set

$$\hat{\sigma} = \operatorname{argmax}_{\sigma \in [2]^n} |Z_n(\sigma, X)|$$

which leads to the profile maximum likelihood estimate

$$\hat{\theta} = \frac{Z_n(\hat{\sigma}, X)}{b_n} \quad \text{and} \quad b_n = \binom{n}{2} = \frac{n(n-1)}{2}. \quad (15)$$

This estimator can be seen as a (pseudo)-maximum likelihood estimate, see Appendix C.

Theorem 3. Consider a stochastic blockmodel (3) with $k = 2$ specified by \mathcal{M} , that is, $P_\theta = P_{e, Q^\theta}$ with e, Q^θ given by (8)-(9). Let $\hat{\theta}$ be the estimator defined by (15). There exists a constant $C_1 > 0$ such that for all $n \geq 2$,

$$\sup_{\theta \in [-1/2, 1/2]} E_\theta [\hat{\theta} - \theta]^2 \leq \frac{C_1}{n}.$$

The same risk bound holds for $\hat{\theta}$ in the fixed design model, uniformly over θ and $\varphi \in [2]^n$.

Proof. See Appendix C.2. □

The main takeaway from this result is that the *uniform* quadratic rate for estimating the connectivity parameter along the submodel \mathcal{M} is exactly of order n^{-1} , up to constants. That follows from combining Theorems 1 and 3. This ‘slow’ rate (as compared to the asymptotic pointwise quadratic rate n^{-2} of (2)) arises even if all other parameters—here, the vector of proportions π —are assumed known. The submodel built for $k = 2$ can be regarded as a local perturbation of an Erdős-Renyi graph model with connection probability $1/2$. The drop in the rate is already noteworthy, as the rate of estimation of p for a $ER(p)$ model is of the order n^{-2} .

The case $k \geq 2$. For this case, we make additional (but fairly mild) assumptions on the matrix M . These conditions are for simplicity of presentation and could, in some cases, be improved. Our main purpose here is to show that, for ‘typical’ matrices A and B in (12), the rate of estimation of θ in (12) is indeed exactly of the order k/n . In Section 4.3 below, we show that at least *some* conditions on possible matrices A, B are necessary: for certain unfavourable matrices, the rate drops below k/n . As was the case for Theorem 1, the result of Theorem 3 remains unchanged if the constant $1/2$ in Q^θ is replaced by any $a_0 \in (0, 1)$.

We modify the criterion function (14) by restricting it to a given subset $S \subset \{1, 2, \dots, n\}$ of indices,

$$2Z_n(\sigma, S, X) = - \sum_{i < j, i, j \in S, \sigma(i) = \sigma(j)} (1 - 2X_{ij}) + \sum_{i < j, i, j \in S, \sigma(i) \neq \sigma(j)} (1 - 2X_{ij}). \tag{16}$$

To avoid technicalities, we maximize over a grid, which constitutes no loss of generality. To this end, define the regular grid $\Theta_n = \{i/(2n^2), i = -n^2, \dots, n^2\}$ in $\Theta = [-1/2, 1/2]$, and

$$(\tilde{\sigma}, \tilde{\theta}) := \operatorname{argmin}_{\sigma \in \Sigma_e, \theta \in \Theta_n} \sum_{i < j} (X_{ij} - Z_{\sigma(i)\sigma(j)}^\theta)^2 \tag{17}$$

$$\tilde{S}_I := \tilde{\sigma}^{-1}(\{1, 2\}). \tag{18}$$

Equation (17) defines a global maximum-likelihood type estimator, which is then used to obtain an estimate \tilde{S}_I of the set of nodes labelled 1 or 2. Given

this estimate, one can apply the profile-type method already used in the case $k = 2$: For \tilde{S}_I as in (18), $\tilde{n}_k = \binom{|\tilde{S}_I|}{2}$, and Z_n as in (16), set

$$\hat{\sigma}_I = \operatorname{argmax}_{\sigma \in \Sigma_e} |Z_n(\sigma, \tilde{S}_I, X)| \quad (19)$$

$$\hat{\theta} = \frac{Z_n(\hat{\sigma}_I, \tilde{S}_I, X)}{\tilde{n}_k}. \quad (20)$$

We require the coefficient a_0 of the matrix M in (10) to be sufficiently distinct from the remaining entries: Let $\mathcal{C} = \{a_i, b_{ij}, 1 \leq i, j \leq k - 2\}$ be the set of coefficients of the matrices A and B in (12), with $a_0 = 1/2$,

$$\min_{c \in \mathcal{C}} \{|c - a_0|\} \geq 2\kappa > 0. \quad (21)$$

Theorem 4. Consider a stochastic blockmodel (3) with $k \geq 2$ classes specified by $P_\theta = P_{e_k, M^\theta}$ with e_k and M^θ given by (11)-(12), for fixed matrices A and B . Define $\hat{\theta} = \hat{\theta}(X)$ as in (20). Suppose (21) holds and that, for some small enough d and κ as in (21),

$$k^3 \log k \leq d\kappa^4 n. \quad (22)$$

Then there exists a universal constant $C_1 > 0$ such that for $n \geq 5$,

$$\sup_{|\theta| \leq \kappa} E_\theta \left[\hat{\theta} - \theta \right]^2 \leq C_1 \frac{k}{n}.$$

The same risk bound holds for $\hat{\theta}$ in the fixed design model, uniformly over $|\theta| \leq \kappa, \varphi \in \Sigma_e$.

Proof. See Appendix C.3. □

Note κ in (22) may depend on k and n , and may go to zero in a framework where k, n go to infinity. Below are two examples for the behaviour of κ . These examples illustrate that our conditions are indeed met in commonly encountered settings, in particular, with high probability, if M is a random matrix and k does not grow too rapidly with n .

Example 1 (well-separated block). If κ is a fixed positive constant e.g. $1/4$, then the submatrix Q^θ is well separated from the other coefficients of the matrix M . The procedure above then correctly picks up a sensible approximation of the true set $\sigma^{-1}(\{1, 2\})$ via \tilde{S}_I and the rate k/n is achieved, as long as k does not grow faster than $n^{1/3}/\log n$, an already fairly important number of classes.

Example 2 (randomly sampled matrix M). Suppose that the symmetric matrix $M = (c_{ij})$ in (10) is a random matrix whose upper triangular entries are drawn i.i.d. with uniform distribution $\mathcal{U}[0, 1]$, except $c_{11} = 1/2$. The distribution of $|c_{ij} - 1/2|$ except for $i = j = 1$ is then $\mathcal{U}[0, 1/2]$, and it is a standard fact that the first order statistic of a uniformly distributed sample of size N is Beta(1, N) distributed. That implies the random variable $2 \min_{c_{ij} \in \mathcal{C}} |c_{ij} - 1/2|$ has law

$\text{Beta}(1, k(k-1)/2 - 1)$. Therefore, κ in (21) is of order no less than $1/k^2$ with high probability. From (22) one deduces that for k of the form n^δ with $\delta < 1/11$ and n large enough, the rate k/n is achieved uniformly and locally, for typical matrices M . Inspection of the proof of Theorem 4 reveals that $k = o(n^\delta)$ with $\delta < 1/7$ in fact suffices for the rate k/n to be attained with high probability when M is random: this is achieved by distinguishing c_{ij} of the types a_i or b_{ij} in the proof and noting that the minimum of $|a_i - 1/2|$ over i will be of larger order k^{-1} , instead of k^{-2} for the minimum over i, j of $|b_{ij} - 1/2|$.

Remark 3 (conditions $|\theta| \leq \kappa$ and (21)). We slightly restrict the range of θ in the upper bound of Theorem 4. Formally, the matching upper bound is obtained for a somewhat smaller interval than $[-1/2, 1/2]$ when $k \geq 3$. (If k is fixed and $n \rightarrow \infty$, the restriction is only to $[-\delta, \delta]$ for a small enough constant $\delta > 0$.) The condition is needed to ensure, in combination with (21), that the block Q^θ in the matrix (12) is separated sufficiently from the other submatrices A and B . If this is not the case, the estimation problem can become more difficult, and the rate hence slower. This is formally shown in Section 4.3, where the extreme case of all coefficients of A, B being equal to $1/2$ is discussed. This phenomenon can also occur if only some parts of A and B are close to $1/2$, or to either $1/2 + \theta$ or $1/2 - \theta$ for some $\theta \in (0, 1/2)$.

We do not claim that the restriction to $[-\kappa, \kappa]$ and (21) are sharp conditions; they can probably be improved. However, the argument above shows some condition of this form is needed, although it may vary depending on the estimation procedure considered: for spectral estimators as considered in Appendix A, for example, we need a similar separation assumption, although it takes a slightly different form (see Theorem 7 in Appendix A and the comments below it). We also note that small values of θ are conceptually the most interesting case, since the 2×2 subproblem becomes easier the larger θ becomes.

4.2. Upper bounds via spectral estimates

Since the maximum likelihood estimator (17) has to optimize over the set $[k]^n$, it need not be computable in polynomial time. It hence seems natural to ask whether there is a “computational gap”, that is, whether the best estimator computable in polynomial time converges at a slower rate than predicted by the minimax bound. We do not have a complete answer to his question, but for a somewhat restricted model class, no such gap exists: The estimator described below for the case $k = 2$ uses a spectral method, see e.g. [25]. A generalization to $k \geq 3$ classes is discussed in Appendix A, which requires further conditions. Within the remit of these conditions, however, the minimax rate is achievable in polynomial time. An extension to sparse graphs is considered in Appendix B. A small simulation study in Section A.2 illustrates the behaviour of the estimator.

With the convention that $X_{ii} = 1/2$ and $X_{ji} = X_{ij}$, define the $n \times n$ matrix Δ by

$$\Delta := X - \frac{1}{2}J, \quad \text{where} \quad J := (1)_{i,j \leq n}.$$

Let $\lambda_1^q(\Delta)$ denote the largest eigenvalue in absolute value of Δ and set

$$\tilde{\theta} := \frac{\lambda_1^q(\Delta)}{n-1}. \quad (23)$$

We refer to this procedure as spectral algorithm for $k=2$ and denote it \mathcal{S}_2 . The intuition behind this estimator in the fixed design setting is the following. For $i \neq j$, we have

$$E[X_{ij} - \frac{1}{2}] = M_{\varphi^{(i)}\varphi^{(j)}}^\theta - \frac{1}{2} = (-1)^{\mathbb{1}_{\{\varphi^{(i)} \neq \varphi^{(j)}\}}} \theta.$$

Set $v = ((-1)^{\mathbb{1}_{\{\varphi^{(i)}=1\}}})_{i \leq n}$ and $V := vv^t = ((-1)^{\mathbb{1}_{\{\varphi^{(i)} \neq \varphi^{(j)}\}}})_{i,j \leq n}$. Then for non-random φ ,

$$E[\Delta] = \theta(V - I_n),$$

where I_n is the identity matrix of size n . As $E[\Delta]$ is a rank 1 matrix whose non-zero eigenvalue equals $(n-1)\theta$ (with v the corresponding eigenvector), this leads us to introduce $\tilde{\theta}$ as in (23).

Theorem 5. *In the same setting as in Theorem 3, let $\tilde{\theta}$ be the estimator defined by (23). There exists a constant $C > 0$ such that for all $n \geq 2$,*

$$\sup_{\theta \in [-1/2, 1/2]} E_\theta [\tilde{\theta} - \theta]^2 \leq \frac{C}{n}.$$

The same risk bound holds for $\hat{\theta}$ in the fixed design model, uniformly over θ and $\varphi \in [2]^n$.

Proof. This follows as a special case of Theorem 8, in Appendix B. \square

4.3. Necessity of conditions on M

What precedes shows that the rate k/n is achieved under conditions on M in (10) and/or k . In general, we expect the rate to depend on the matrices M . Although we do not investigate this point in full here, we discuss it briefly.

The estimation methods investigated in Section 4.1 (MLE) and Appendix A (spectral method) require the upper-left 2×2 block of M^θ to be sufficiently separated from at least part of the other entries of M^θ . Among those matrices M^θ whose upper-left corner equals Q^θ , a worst case scenario should correspond to a matrix whose coefficients in A and B all equal $1/2$. This leads to the matrix

$$\check{M}^\theta = \begin{bmatrix} \frac{1}{2} + \theta & \frac{1}{2} - \theta & \frac{1}{2} & \cdots & \frac{1}{2} \\ \frac{1}{2} - \theta & \frac{1}{2} + \theta & \frac{1}{2} & \cdots & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \cdots & \frac{1}{2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \cdots & \frac{1}{2} \end{bmatrix}, \quad (24)$$

which is of course heavily over-specified from the SBM perspective. Consider the SBM in a fixed design case, where $\varphi : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ is unobserved. Suppose all classes $\sigma^{-1}(i)$ are of cardinality of order n/k , and the connectivity matrix is given by (24). This specific model can be regarded as a special case of the setting considered from a testing perspective by Butucea and Ingster [10] and Arias-Castro and Verzelen [3]. From Theorem 4.3 of [10], one can deduce that the minimax rate for the quadratic risk when estimating θ is no better than $\rho_n = \min\left(\frac{k^2}{n}, \sqrt{\frac{k \log k}{n}}\right)$, for $k, n \rightarrow \infty$ and $\rho_n = o(1)$. The rate is therefore no better than k^2/n for $k \leq n^{1/3}$, and remains much slower than k/n even for $k > n^{1/3}$.

4.4. Minimax rates for a class of functionals of smooth graphons

Stochastic block models can be identified with piecewise constant graphons; we now consider the case where the graphon is a smooth function instead. Let $w : [0, 1]^2 \rightarrow [0, 1]$ a measurable function, let $\langle w \rangle$ be its graphon equivalence class, and denote by $P_{\langle w \rangle} = P_w$ the distribution of data X generated by the graphon model (4). Consider the problem of estimating the functional

$$\vartheta(\langle w \rangle) = \left[\int_{[0,1]^2} \left(w(x, y) - \int_{[0,1]^2} w \right)^2 dx dy \right]^{1/2}, \quad (25)$$

for any representer w of $\langle w \rangle$. This is well defined in terms of the graphon, as the integral is invariant under any simultaneous (Lebesgue-)measure-preserving transformation of x and y .

The statistic (25) can be interpreted as a ‘graphon-standard deviation’. Its estimation under a smooth graphon model is, in a sense, analogous to the problem of estimating the functional θ in the simple SBM with two classes discussed in Section 3.1: Let h_θ be the piece-wise constant graphon characterizing the SBM defined by (8)–(9). Since $\vartheta(\langle h_\theta \rangle) = |\theta|$, estimating θ is then indeed equivalent (for positive values) to estimation of $\vartheta(\langle h_\theta \rangle)$.

Under a 2-class SBM, the results of Section 3.1 show θ in (9) cannot be estimated faster than c/n . It is natural to ask whether the same still holds if one works with ‘smoother’ graphons instead of histograms (where we refer to $\langle w \rangle$ as smooth if at least one of its representers is a smooth function). The following result addresses this question for a simple class of smooth graphons, both for $\vartheta(\cdot)$ and for a larger class of functionals containing $\vartheta(\cdot)$.

Let \mathcal{P}_B be the collection of all graphons that admit a representer which is a polynomial in x, y , with degree bounded by some integer $D \geq 2$ and coefficients bounded by an arbitrary constant $M > 0$ (this boundedness restriction is only to ensure a –nearly, up to a log term– matching upper-bound in the next result). For any $0 \leq \theta \leq 1$, let us denote by w_θ the function from $[0, 1]^2$ to $[0, 1]$ given by

$$w_\theta(x, y) = \frac{1}{2} - \theta \left(x - \frac{1}{2}\right) \left(y - \frac{1}{2}\right) \quad (26)$$

and let w_0 denote the constant function equal to $1/2$. The function w_θ can be interpreted as a ‘smooth’ counterpart to the histogram graphon underlying the SBM (9).

Theorem 6. *Let X be data from the graphon model (4). Let $\vartheta(\cdot)$ be defined as in (25). There exist constants $c_1, c_2 > 0$ such that*

$$\frac{c_1}{n} \leq \inf_{\hat{\vartheta}} \sup_{w \in \mathcal{P}_B} E_{P_w} \left[\hat{\vartheta}(X) - \vartheta(\langle w \rangle) \right]^2 \leq \frac{c_2 \log n}{n},$$

where the infimum is taken over all possible estimators of $\vartheta(\langle w \rangle)$ in model (4). Let ψ be an arbitrary functional defined on graphon equivalence classes satisfying

$$|\psi(\langle w_\theta \rangle) - \psi(\langle w_0 \rangle)| \geq c|\theta| \quad (27)$$

for some $c > 0$, for any $0 \leq \theta \leq 1$, and for the function w_θ in (26). Then for some $d > 0$,

$$\inf_{\hat{\psi}} \sup_{w \in \mathcal{P}_B} E_{P_w} \left[\hat{\psi}(X) - \psi(\langle w \rangle) \right]^2 \geq \frac{d}{n}.$$

Proof. See Section 7. □

The first part of Theorem 6 asserts that the quadratic minimax rate for estimating (25) cannot be faster than c/n , even if one restricts the parameter set to a small class of smooth graphons w , namely graphons with a polynomial representer of bounded degree. This class can be seen as a smooth analogue of the histogram graphon underlying model (7), or more generally the model with k classes and connectivity (12). The degree of the polynomial can be seen as the analog of k . The rate of order $1/n$ is obtained because the degree of the polynomial is assumed bounded. Although we do not investigate this further here, one may conjecture that the rate would slow even further for a larger class (e.g. growing degree of polynomials, or a nonparametric class such as a Hölder ball).

The second part of Theorem 6 indicates that the specific form of the functional $\vartheta(\cdot)$ in (25) is not essential for the lower bound to hold. A given functional $\psi(\cdot)$ leads to a rate at least as slow of $\vartheta(\cdot)$ over the considered class of graphons as soon as (27) holds. This condition intuitively means that the functional $\psi(\cdot)$ is at least as hard as to estimate as the functional $\vartheta(\cdot)$, for which the difference on the left hand-side of (27) indeed behaves like $|\theta|$. By direct computation we see that an example of such a graphon functional is

$$\psi(P_w) = \int_{[0,1]^2} \left| f(x, y) - \int_{[0,1]^2} |f(x, y)| dx dy \right| dx dy.$$

Providing a unified theory with matching lower and upper bounds for graphon functionals is an interesting topic for future research.

5. Discussion

Gao et al. [16] show that, if one estimates the parameter function w of a graphon model, not observing the vertex labels—in this case, the variables U_i in (4)—does (in general) impact on the optimal rate. In the present paper, we have considered uniform estimation of certain functionals of graphon models (in particular, the loss function is quite different from theirs). For estimation of certain random graph functionals—including the connectivity parameters considered by Bickel et al. [7]—we have shown that the uniform, minimax rate does depend on whether the labels are observed, i.e. the phenomenon described by [16] persists even if one does not try to recover the entire function w , but only a specific 1-dimensional aspect of w . The fast quadratic rate $1/n^2$ is not achievable uniformly. If the number k of classes is known and fixed, the quadratic rate becomes $1/n$. If the number of classes k grows with n , the rate drops to k/n . We have used some mild assumptions on the part of the connectivity matrix other than the 2×2 submodel. If those assumptions are not satisfied, the rate may even drop further. Similar results also hold for sparse graphs.

Interestingly, for the functionals considered here, the uniform rate is always, regardless of the number of classes k , much below the rate in the case where labels would be observed. This is in contrast with the problem of recovery of the mean adjacency matrix considered in [16], where for k is larger than $\sqrt{n \log n}$, the (non-normalised) rate $k^2 + n \log k$ is dominated by the ‘parametric’ rate k^2 , the rate if labels are observed.

We claim no novelty regarding the algorithms—the MLE and spectral method—which we have adapted from existing work to the problem at hand. Their purpose is to verify that the lower bound is tight (both algorithms achieve it) under some mild conditions, and that there is no computational gap (the spectral method does so in polynomial time). Yet, we are not aware of other work providing uniform rates for SBM connectivity parameters for these or other algorithms, which constitutes another novelty of the paper.

Aspects of our proofs reflect the fact that graphon models constitute a specific type of mixture model, and estimation in mixtures can be difficult if mixture components are hard to distinguish; although no general theory of these phenomena seems to exist, we refer to the early work on estimation in finite mixture models by [19] and [6], and e.g. to [20] and [15] for more recent results.

6. Proofs of the lower bounds in SBMs

The proofs of Theorems 1 and 2 rely on variations of Le Cam’s ‘two-points’ method, which bounds the minimax risk from below by a quantity involving the L_1 distance between a distribution and a finite mixture. (Specifically, this is the ‘point versus mixture’ variant of the two points method, see e.g. [35].) This and other relevant technical lemmas are recalled in Section 6.3 below; the two points method is Lemma 3. For Theorem 2, for $k \geq 2$ classes, one main idea is to ‘isolate’ the part corresponding to the submatrix Q^θ . More details comments are given along the proof in Section 6.2 below.

Notation. Recall that a SBM with k classes, proportions vector π and connectivity matrix M has distribution $P_{\pi,M}$ as given in (6). For a $n \times n$ symmetric matrix A with zero diagonal, we write

$$P_A = \bigotimes_{i < j} \text{Be}(A_{i,j}).$$

If A is only given by $A_{i,j}$ for $i < j$, one extends it by symmetry and sets $A_{i,i} = 0$. The distribution of a SBM in the fixed design case with given k, M and labelling function φ is hence P_{A^φ} , where $A_{i,j}^\varphi = A^\varphi(M)_{i,j} = M_{\varphi(i)\varphi(j)}$. In the random design case, if π is the vector with equal proportions $e_k = [k^{-1}, \dots, k^{-1}]$, then from (6),

$$P_{\pi,M} = P_{e_k,M} = \frac{1}{k^n} \sum_{\varphi \in [k]^n} P_{A^\varphi} \quad \text{where} \quad A_{i,j}^\varphi = M_{\varphi(i)\varphi(j)}.$$

We generically denote universal constants by C , where the value may change from line to line.

6.1. Two classes

Proof of Theorem 1. Let $N = 2^n$ and let A_1, \dots, A_N be the collection of symmetric $n \times n$ matrices with general term $a_{ij}(\varphi) = a_{ij}(\theta, \varphi) = Q_{\varphi(i)\varphi(j)}^\theta$, $i < j$ and zero diagonal, for all possible $\varphi \in [2]^n$ and some $\theta \in \Theta$. Let A_0 be the $n \times n$ matrix with all elements equal to $1/2$ on the off-diagonal, that is the matrix with $\theta = 0$. By Lemma 3, applied with $\vartheta = 0$ and $\theta = \theta_n$ small to be chosen below, in order to get a lower bound for the minimax risk, it is enough to bound the L^1 -distance $\|\mathbb{P} - \mathbb{Q}\|_1$ between

$$\mathbb{P} = P_{A_0}, \quad \mathbb{Q} = \frac{1}{N} \sum_{k=1}^N P_{A_k}.$$

If $\lambda_1 = \text{Be}(q)$, $\lambda_2 = \text{Be}(r)$, $\mu = \text{Be}(s)$, a simple computation leads to

$$\int \left(\frac{d\lambda_1}{d\mu} - 1\right) \left(\frac{d\lambda_2}{d\mu} - 1\right) d\mu = \frac{(q-s)(r-s)}{s(1-s)}.$$

By Lemma 4 applied to \mathbb{P} and \mathbb{Q} , where $\vartheta_{i,j}(\varphi, \psi) = (2a_{ij}(\varphi) - 1)(2a_{ij}(\psi) - 1)$,

$$\begin{aligned} \|\mathbb{P} - \mathbb{Q}\|_1^2 &\leq \frac{1}{4^n} \sum_{\varphi, \psi \in [2]^n} \prod_{i < j} (1 + \vartheta_{i,j}(\varphi, \psi)) - 1 \\ &\leq \left[\frac{1}{4^n} \sum_{\varphi, \psi \in [2]^n} e^{\sum_{i < j} \vartheta_{i,j}(\varphi, \psi)} \right] - 1, \end{aligned}$$

Note that $2a_{ij}(\varphi) - 1 = 2\theta(-1)^{\mathbb{1}_{\varphi(i) \neq \varphi(j)}}$ for any $\varphi \in [2]^n$. Denote $\eta_i = \mathbb{1}_{\varphi(i)=1} - \mathbb{1}_{\varphi(i)=2}$ and $\eta'_i = \mathbb{1}_{\psi(i)=1} - \mathbb{1}_{\psi(i)=2}$, for any index i . We have $2a_{ij}(\varphi) - 1 = 2\theta\eta_i\eta_j$,

so that $\vartheta_{i,j}(\varphi, \psi) = 4\theta^2 \eta_i \eta_j \eta'_i \eta'_j$. The term under brackets in the last display can be interpreted as an expectation over φ, ψ , where both variables are sampled uniformly from the set of all mappings from $\{1, \dots, n\}$ to $\{1, 2\}$. Under this distribution, the variables η_i for $i = 1, \dots, n$ are independent Rademacher, as well as the variables η'_i , and both samples are independent. Further note that the variables $R_i := \eta_i \eta'_i$ for $i = 1, \dots, n$ form again a sample of independent Rademacher variables. It is thus enough to bound,

$$E \left[e^{4\theta_n^2 \sum_{i < j} R_i R_j} \right],$$

where E denotes expectation under the law of the R_i . The previous exponent is an instance of Rademacher chaos; its Laplace transform can be bounded using Lemma 1. If $Z_n := \sum_{i < j} R_i R_j$, we have that for any ε (say $\varepsilon = 1/2$), there exists $\lambda > 0$ such that for all $n \geq 2$,

$$E \left[e^{|Z_n|/(\lambda n)} \right] \leq 1 + \varepsilon.$$

Choosing $n\theta_n^2 := 1/(4\lambda)$ leads to $\|\mathbb{P} - \mathbb{Q}\|_1^2 \leq \varepsilon = 1/2$, so that the minimax risk is bounded below by $(32n\lambda)^{-1}$.

To obtain the constants as in the remark below the Theorem, using Lemma 2 in the final step of the proof with $\theta_n^2 = 1/(12s_n)$, $s_n^2 = n(n-1)/2$, $r(\cdot)$ as in Lemma 2, gives

$$\begin{aligned} R_M &\geq \frac{\theta_n^2}{4} \left\{ 1 - \frac{1}{2} \sqrt{r(4\theta_n^2 s_n)} \right\}, \\ &\geq \frac{\sqrt{2}}{48n} \left\{ 1 - \frac{1}{2} \sqrt{r\left(\frac{1}{3}\right)} \right\} \geq \frac{0.45}{48n} \geq \frac{1}{107n}. \quad \square \end{aligned}$$

6.2. Lower bounds for k classes

Here the problem is more delicate compared to $k = 2$, as the typical number of nodes per class now depends on k , and, in the random design case, the data distribution for $\theta = 0$, around which we build the lower bound, is itself a mixture. As a first step, we start by establishing a result in a fixed design setting, that is

$$\inf_T \sup_{\theta \in \Theta, \varphi \in [k]^n} E_{\theta, \varphi} [T - \theta]^2 \geq c \frac{k}{n} \quad \text{for some } c > 0 \text{ and } k \geq 3. \quad (28)$$

Proof of (28). Define $m = m_k = 2 \lfloor \frac{n}{k} \rfloor$. Set $S_1 = \{1, \dots, m\}$ and $S_2 = \{m + 1, \dots, n\}$. Let $\varphi_0 \in [k]^n$ be a mapping such that

$$\varphi_0(S_1) \subset \{1, 2\} \quad \text{and} \quad \varphi_0(S_2) \subset \{3, \dots, k\}. \quad (29)$$

Let $\varphi \in [k]^n$ be such that

$$\varphi(i) = \varphi_0(i) \quad \text{whenever } i \in S_2 \quad \text{and} \quad \varphi(S_1) \subset \{1, 2\}, \quad (30)$$

and denote by $\mathcal{F} = \mathcal{F}(\varphi_0, S_1)$ the set of all such φ 's. Then the restriction $\varphi|_{S_1} =: \varphi_1$ of $\varphi \in \mathcal{F}$ to S_1 can be identified to an element of $[2]^m$.

Let M^θ be the $k \times k$ matrix defined in (12). For $\varphi \in \mathcal{F}$, let R_φ denote the matrix with general term $r_{ij} = r_{ij}(\varphi)$ equal to $M_{\varphi^{(i)\varphi(j)}}^\theta$. There are as many such matrices as possible φ 1s, that is $|[2]^m| = 2^m$. As φ and φ_0 are identical by construction on S_2 ,

$$r_{ij} = \begin{cases} M_{\varphi_1^{(i)\varphi_1(j)}}^\theta & \text{if } (i, j) \in S_1 \times S_1 \\ M_{\varphi_1^{(i)\varphi_0(j)}}^\theta = a_{\varphi_0(j)-2} = M_{\varphi_0^{(i)\varphi_0(j)}}^0 & \text{if } (i, j) \in S_1 \times S_2 \\ M_{\varphi_0^{(i)\varphi_1(j)}}^\theta = a_{\varphi_0(i)-2} = M_{\varphi_0^{(i)\varphi_0(j)}}^0 & \text{if } (i, j) \in S_2 \times S_1 \\ M_{\varphi_0^{(i)\varphi_0(j)}}^0 & \text{if } (i, j) \in S_2 \times S_2 \end{cases}$$

where φ_1 belongs to $[2]^m$. Next set, with A_0 the matrix with general term $m_{ij}(\varphi_0) = M_{\varphi_0^{(i)\varphi_0(j)}}^0$,

$$\mathbb{P}' = P_{A_0} \quad \text{and} \quad \mathbb{Q}' = \frac{1}{2^m} \sum_{\varphi \in \mathcal{F}} P_{R_\varphi}.$$

Now we apply Lemma 4 to \mathbb{P}', \mathbb{Q}' . Both P_{A_0} and P_{R_φ} are product measures over all pairs of indices (i, j) with $1 \leq i < j \leq n$. By construction, the individual components of these products coincide as soon as either i or j does not belong to S_1 . We write

$$P_{A_0} = \bigotimes_{i < j} P_{A_0}(i, j) \quad \text{and} \quad P_{R_\varphi} = \bigotimes_{i < j} P_{R_\varphi}(i, j).$$

where, for any indices i, j with $i < j$,

$$P_{A_0}(i, j) = \text{Be}(M_{\varphi_0^{(i)\varphi_0(j)}}^0) \quad \text{and} \quad P_{R_\varphi}(i, j) = \text{Be}(M_{\varphi^{(i)\varphi(j)}}^\theta).$$

For $\varphi, \psi \in [4]^n$, we set

$$\tau_{i,j}(\varphi, \psi) = P_{A_0}(i, j) \left[\left(\frac{dP_{R_\varphi}(i, j)}{dP_{A_0}(i, j)} - 1 \right) \left(\frac{dP_{R_\psi}(i, j)}{dP_{A_0}(i, j)} - 1 \right) \right].$$

If i or j belongs to S_2 , then $r_{ij}(\varphi) = M_{\varphi_0^{(i)\varphi_0(j)}}^0 = A_0(i, j) = r_{ij}(\psi)$ by definition, in which case the last display equals 0. In Lemma 4, where φ, ψ play the role of the indices k, l . Identifying $\psi|_{S_1}$ with the corresponding mapping $\psi_1 \in [2]^m$, we have $\|\mathbb{P}' - \mathbb{Q}'\|_1^2 \leq \chi^2(\mathbb{Q}', \mathbb{P}')$ and

$$\chi^2(\mathbb{Q}', \mathbb{P}') \leq \frac{1}{2^{2m}} \sum_{\varphi_1, \psi_1 \in [2]^m} \prod_{1 \leq i < j \leq m} (1 + \tau_{i,j}(\varphi_1, \psi_1)) - 1. \tag{31}$$

The last expression coincides with the bound obtained in the proof of Theorem 1, with n replaced by $m = m_k$. As in that proof, there hence exist independent Rademacher variables R_1, \dots, R_m such that $Z_m = m^{-1} \sum_{1 \leq i < j \leq m} R_i R_j$ satisfies

$$\chi^2(\mathbb{Q}', \mathbb{P}') \leq \mathbb{E} \exp [4m\theta^2 |Z_m|] - 1.$$

Provided θ is defined as, for a a small enough constant,

$$\theta^2 = a(4m)^{-1} \sim \left(\frac{n}{k}\right)^{-1},$$

using Lemma 1 as in the proof of Theorem 1 leads to the bound $\|\mathbb{P}' - \mathbb{Q}'\|_1 \leq 1/2$ if θ^2 is a small enough multiple of k/n , which again leads to a lower bound for the minimax risk of a positive constant times k/n , which proves (28). \square

Proof of Theorem 2. For $e = e_k$ and M^θ as in (11)-(12), let

$$\mathbb{Q}^\theta = \frac{1}{k^n} \sum_{\varphi \in [k]^n} \bigotimes_{i < j} \text{Be}(M_{\varphi(i)\varphi(j)}^\theta),$$

and set $\mathbb{P} = \mathbb{Q}^0$ corresponding to $\theta = 0$. Our aim is to show that \mathbb{Q}^θ and \mathbb{P} are close in the sense $\|\mathbb{Q}^\theta - \mathbb{P}\|_1 \leq 1/2$ say, while θ is a fixed positive multiple of $\sqrt{k/n}$. For a given $\varphi \in [k]^n$, set

$$S_1 := \varphi^{-1}(\{1, 2\}) \quad \text{and} \quad S_2 := \varphi^{-1}(\{3, 4, \dots, k\}).$$

By definition we have $S_1 = S_2^c := \{1, \dots, n\} \setminus S_2$ and $|S_1| + |S_2| = n$.

The proof has two steps. First, one shows that with high probability one can restrict to designs (i.e. specific mapping φ 's) such that there are around $2n/k$ nodes that have label either 1 or 2. Second, we show that estimation with a random design is 'harder' than in the (easiest) typical fixed design case. This argument is reminiscent of 'information processing inequalities' encountered in information theory, although here a maximisation also takes place for not knowing the class labels. It is then important to maximise only over designs obtained from Step 1, in order for the lower bound rate to be k/n .

Step 1. One first shows that it is possible to restrict the sum in the definition of \mathbb{Q}^θ and \mathbb{P} to φ 's in the set

$$\begin{aligned} \mathcal{A}_n &= \left\{ \varphi \in [k]^n, \quad \left| |\varphi^{-1}(\{1, 2\})| - \frac{2n}{k} \right| \leq \frac{n}{k} \right\} \\ &= \left\{ \varphi \in [k]^n, \quad \left| |\varphi^{-1}(\{3, \dots, k\})| - \frac{(k-2)n}{k} \right| \leq \frac{n}{k} \right\}. \end{aligned}$$

The reason is that the large majority of sets S_1 have a cardinality of the order close to n/k . The proportion of φ 's not in \mathcal{A}_n among all possible φ 's is given by the probability of a binomial $Y \sim \text{Bin}(n, 2/k)$ variable being farther than n/k from its mean. By Bernstein's inequality, as $v := \text{Var}[Y] = n(2/k)(1 - 2/k)$, for any $t > 0$,

$$\mathbb{P} \left[\left| Y - \frac{2n}{k} \right| > t \right] \leq 2 \exp \left\{ -\frac{t^2}{2v + t} \right\}.$$

Taking $t = n/k$ and setting $R_n := |\mathcal{A}_n|$, we have just shown that

$$0 \leq 1 - \frac{R_n}{k^n} \leq 2e^{-\frac{n}{k} \left(5 - \frac{8}{k}\right)^{-1}}.$$

Now set

$$\tilde{\mathbb{Q}}^\theta = \frac{1}{R_n} \sum_{\varphi \in \mathcal{A}_n} \bigotimes_{i < j} \text{Be}(M_{\varphi^{(i)\varphi(j)}}^\theta) \quad \text{and} \quad \tilde{\mathbb{P}} = \tilde{\mathbb{Q}}^0.$$

By the triangle inequality,

$$\|\mathbb{Q}^\theta - \mathbb{P}\|_1 \leq \|\mathbb{Q}^\theta - \tilde{\mathbb{Q}}^\theta\|_1 + \|\tilde{\mathbb{Q}}^\theta - \tilde{\mathbb{P}}\|_1 + \|\tilde{\mathbb{P}} - \mathbb{P}\|_1.$$

By Lemma 5, $\|\mathbb{Q}^\theta - \tilde{\mathbb{Q}}^\theta\|_1 + \|\tilde{\mathbb{P}} - \mathbb{P}\|_1$ is bounded above by $4(1 - R_n/k^n) \leq 8e^{-\frac{n}{k}(5-\frac{8}{k})^{-1}}$.

Step 2. We now focus on bounding the middle term $\|\tilde{\mathbb{Q}}^\theta - \tilde{\mathbb{P}}\|_1$. Let Σ_n denote the collection of subsets of $\{1, 2, \dots, n\}$ with $|S_2 - (k - 2)n/k| \leq n/k$. For a given $S \in \Sigma$, let $\varphi_S = \varphi|_S$ denote the restriction of φ to S . Below we use the notation \sum_{φ_S} with the meaning that each term of the sum corresponds to a possible mapping φ_S , that is a given collection of values $(\varphi(i))_{i \in S} \in \{1, \dots, k\}^S$.

To do so, we rewrite $\tilde{\mathbb{Q}}^\theta$ and $\tilde{\mathbb{P}}$ as ‘mixtures of mixtures’, by splitting the sum over φ into a sum over S_2, φ_{S_2} and φ_{S_1} given S_2 . Specifying φ is equivalent to giving oneself S_2 (then $S_1 = S_2^c$), φ_{S_2} and $\varphi_{S_2^c} = \varphi_{S_1}$. Denote

$$P_{\varphi}^\theta(i, j) = \text{Be}(M_{\varphi^{(i)\varphi(j)}}^\theta) \quad \text{and} \quad P_{\varphi}^\theta = \bigotimes_{i < j} P_{\varphi}^\theta(i, j).$$

For given S_2 and φ_{S_2} , set

$$T_{\varphi, S_2}^\theta = \frac{1}{2^{n-|S_2|}} \sum_{\varphi_{S_1} | S_2, \varphi_{S_2}} P_{\varphi}^\theta,$$

where one sums over all possible mappings φ_{S_1} , while S_2 and φ_{S_2} are fixed. We have

$$\tilde{\mathbb{Q}}^\theta = \sum_{S_2 \in \Sigma_n, \varphi_{S_2}} \frac{2^{n-|S_2|}}{R_n} T_{\varphi, S_2}^\theta.$$

Note that the above measures are normalised to be probability measures. Indeed, given $S_2 \in \Sigma_n$, there are $2^{|S_1|} = 2^{n-|S_2|}$ possible choices for φ_{S_1} . As $\tilde{\mathbb{Q}}^\theta$ is of total mass one, we have

$$\sum_{S_2 \in \Sigma_n, \varphi_{S_2}} \lambda_{S_2} = 1 \quad \text{for} \quad \lambda_{S_2} := \frac{2^{n-|S_2|}}{R_n}.$$

Using the triangle inequality, one can bound

$$\begin{aligned} \|\tilde{\mathbb{Q}}^\theta - \tilde{\mathbb{P}}\|_1 &= \left\| \sum_{S_2 \in \Sigma_n, \varphi_{S_2}} \lambda_{S_2} T_{\varphi, S_2}^\theta - \sum_{S_2 \in \Sigma_n, \varphi_{S_2}} \lambda_{S_2} T_{\varphi, S_2}^0 \right\|_1 \\ &\leq \sum_{S_2 \in \Sigma_n, \varphi_{S_2}} \lambda_{S_2} \|T_{\varphi, S_2}^\theta - T_{\varphi, S_2}^0\|_1 \leq \max_{S_2 \in \Sigma_n, \varphi_{S_2}} \|T_{\varphi, S_2}^\theta - T_{\varphi, S_2}^0\|_1. \end{aligned}$$

It is now sufficient to bound uniformly the above L^1 -distance. For simplicity, we denote

$$T_{\varphi, S_2}^\theta = \frac{1}{2^{n-|S_2|}} \sum_{\varphi_{S_1} | S_2} P_\varphi^\theta := \frac{1}{2^{n-|S_2|}} \sum_{\varphi_1} P_{\varphi_1, \varphi_2}^\theta,$$

where $\varphi_2 = \varphi_{S_2}$ and $\varphi_1 = \varphi_{S_1}$, and φ is the pair (φ_1, φ_2) . Set

$$\lambda_{S_1} = 2^{-(n-|S_2|)} = 2^{-|S_1|}.$$

Using the definition of T_{φ, S_2}^θ above,

$$\begin{aligned} \|T_{\varphi, S_2}^\theta - T_{\varphi, S_2}^0\|_1 &= \left\| \sum_{\varphi_1} \lambda_{S_1} P_{\varphi_1, \varphi_2}^\theta - \sum_{\varphi_1} \lambda_{S_1} P_{\varphi_1, \varphi_2}^0 \right\|_1 \\ &= \left\| \sum_{\varphi_1} \lambda_{S_1} \{P_{\varphi_1, \varphi_2}^\theta - \sum_{\varphi'_1} \lambda_{S_1} P_{\varphi'_1, \varphi_2}^0\} \right\|_1 \\ &\leq \sum_{\varphi_1} \lambda_{S_1} \left\| \sum_{\varphi'_1} \lambda_{S_1} P_{\varphi'_1, \varphi_2}^\theta - P_{\varphi_1, \varphi_2}^0 \right\|_1 \leq \max_{\varphi_1} \left\| \sum_{\varphi'_1} \lambda_{S_1} P_{\varphi'_1, \varphi_2}^\theta - P_{\varphi_1, \varphi_2}^0 \right\|_1. \end{aligned}$$

Combining this with the previous bounds one deduces that

$$\begin{aligned} \|\mathbb{Q}^\theta - \mathbb{P}\|_1 &\leq \|\tilde{\mathbb{Q}}^\theta - \tilde{\mathbb{P}}\|_1 + 8e^{-\frac{n}{k}(5-\frac{8}{k})^{-1}} \\ &\leq \max_{S_2 \in \Sigma_n, \varphi(S_2)} \max_{\varphi_1} \left\| \sum_{\varphi'_1} \lambda_{S_1} P_{\varphi'_1, \varphi_2}^\theta - P_{\varphi_1, \varphi_2}^0 \right\|_1 + 8e^{-\frac{n}{k}(5-\frac{8}{k})^{-1}}. \end{aligned}$$

To conclude the proof, observe that the structure of the bound in the maximum in the last display is nearly identical to the quantities appearing in Equation (31) for the fixed-design case.

In the present case, we have a fixed mapping $\varphi : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$, with $\varphi_1 = \varphi|_{S_1}$ and $\varphi_2 = \varphi|_{S_2}$, that plays the role of φ_0 in the fixed-design case. On the other hand, we have a collection of other mappings, say $\bar{\varphi}$, that coincide with φ on S_2 , that is $\bar{\varphi}_2 = \bar{\varphi}|_{S_2} = \varphi|_{S_2} = \varphi_2$, and that cover all possible cases for the image of S_1 , namely $\bar{\varphi}_1 = \bar{\varphi}|_{S_1} = \varphi'_1$. The only difference to the fixed-design case is that $|S_1|$ belongs to $[n/k, 3n/k]$, instead of being exactly $2\lfloor n/k \rfloor$, as specified in the definition of Σ_n above. That is, denoting as above $Z_m = m^{-1} \sum_{1 \leq i < j \leq m} R_i R_j$, with $m_1 = |S_1|$,

$$\left\| \sum_{\varphi'_1} \lambda_{S_1} P_{\varphi'_1, \varphi_2}^\theta - P_{\varphi_1, \varphi_2}^0 \right\|_1 \leq \mathbb{E} \exp [4m_1 \theta^2 |Z_{m_1}|] - 1.$$

This bound is uniform over S_2, φ_2 . As $m_1 \leq 3n/k$, if one chooses $\theta^2 \leq 1/(12\lambda n/k)$, with $\lambda = \lambda(1 + \varepsilon)$ the constant in Lemma 2, then this Lemma implies that for any m_1 between n/k and $3n/k$, the L^1 -distance in the last display is bounded

by ε . Crucially, the constant λ in Lemma 1 is independent of the number of terms in the Rademacher chaos. Deduce

$$\|\mathbb{Q}^\theta - \mathbb{P}\|_1 \leq \varepsilon + 8e^{-\frac{n}{k}(5-\frac{8}{k})^{-1}}.$$

Choosing $n/k > 12$ makes this bound smaller than $\varepsilon + 4/5 < 1$ for $\varepsilon < 1/5$. \square

6.3. Useful lemmas

Let $\{Z_i, i \geq 1\}$ be i.i.d. Rademacher variables. For reals x_{ij} and $N \geq 2$, set

$$Y := Y_N = \sum_{i < j \leq N} y_{ij} Z_i Z_j,$$

$$s(Y)^2 = \sum_{i < j \leq N} y_{ij}^2.$$

Lemma 1 (Corollary 3.2.6 of de la Peña and Giné [13]). *Let $N \geq 2$, and $Y = Y_N$ and $s(Y)$ as above. For every $c > 1$, there exists $\lambda = \lambda(c) > 0$ independent of N such that*

$$\mathbb{E} \exp \left[\frac{|Y|}{\lambda s(Y)} \right] \leq c.$$

We repeatedly use Lemma 1 in the case where all y_{ij} are equal to 1, for various values of N . In such a setting, a reformulation is as follows. For any $c > 1$ and $N \geq 2$, one can find a constant $a = a(c)$ independent of N such that

$$\mathbb{E} \exp \left[a \frac{|Y_N|}{N} \right] \leq c. \tag{32}$$

Lemma 2 (Rademacher chaos with explicit constant). *Let $N \geq 2$, and $Y = Y_N$ and $s(Y)$ as above. For any $0 \leq \delta \leq 1$,*

$$\mathbb{E} \exp \left[\delta \frac{|Y|}{s(Y)} \right] - 1 \leq r(\delta), \quad \text{with} \quad r(\delta) = \delta + \frac{\delta^2}{2} + \frac{8\delta^3}{6} + \frac{1}{1 - e\delta} \frac{(e\delta)^4}{\sqrt{8\pi}}.$$

The lemma applied with $\delta = 1/3$ gives a bound 1.87 for the right hand side.

Proof. Theorem 3.2.2 in [13] gives, for any $k \geq 2$,

$$E|Y|^k \leq (k - 1)^k s(Y)^k.$$

For $k = 1$ one has $E[|Y|] \leq E[Y^2]^{1/2} = s(Y)$. From this one deduces that for any $0 \leq \delta \leq 1$,

$$E \exp \left[\delta \frac{|Y|}{s(Y)} \right] \leq 1 + \delta + \frac{\delta^2}{2} + \frac{8\delta^3}{6} + \sum_{k \geq 4} \frac{(k - 1)^k}{k!} \delta^k,$$

and the result follows from an application of the nonasymptotic Stirling bound $k! \geq e^{-k} k^{k+\frac{1}{2}} \sqrt{2\pi}$ valid for $k \geq 1$. \square

Lemma 3 (Le Cam’s method ‘point versus mixture’). *Let $\mathcal{P} = \{P_M, M \in \mathcal{M}\}$ be a collection of probability measures indexed by an arbitrary set $\mathcal{M} = \{M_0, M_1, \dots, M_N\}$, $N > 1$. Set*

$$\mathbb{P} = P_{M_0}, \quad \mathbb{Q} = \frac{1}{N} \sum_{k=1}^N P_{M_k}.$$

If ψ is a real-valued functional such that $\psi(P_{M_0}) = \tau$ and $\psi(P_{M_i}) = \theta$ for any $i = 1, \dots, N$, then

$$\inf_{\hat{\theta}} \sup_{M \in \mathcal{M}} E_{P_M} (\hat{\psi}(X) - \psi(P_M))^2 \geq \frac{1}{4} (\theta - \vartheta)^2 \left(1 - \frac{1}{2} \|\mathbb{P} - \mathbb{Q}\|_1\right),$$

where the infimum is over all estimators $\hat{\psi}(X)$ of $\psi(P_M)$ based on the observation of $X \sim P_M$.

Proof. This is a standard variation on the case where $N = 1$ stated in e.g. [35]. □

Lemma 4 (Bound on total variation distance). *For $n \geq 1$, let P_1, \dots, P_n and $Q_1(k), \dots, Q_n(k)$ for $1 \leq k \leq N$, for some $N \geq 1$, be probability measures. Set*

$$P = \bigotimes_{i=1}^n P_i, \quad Q(k) = \bigotimes_{i=1}^n Q_i(k) \quad \mathbb{P} = P, \quad \mathbb{Q} = \frac{1}{N} \sum_{k=1}^N Q(k).$$

Suppose that for any i , $Q_i(k)$ has density $1 + \Delta_i(k)$ with respect to P_i . Denote $\vartheta_i(k, l) = P_i \Delta_i(k) \Delta_i(l)$. Then, for $\chi^2(\mathbb{Q}, \mathbb{P}) = \int (d\mathbb{Q}/d\mathbb{P} - 1)^2 d\mathbb{P}$,

$$\|\mathbb{P} - \mathbb{Q}\|_1^2 \leq \chi^2(\mathbb{Q}, \mathbb{P}) = \frac{1}{N^2} \sum_{k,l} \prod_{i=1}^n \{1 + \vartheta_i(k, l)\} - 1.$$

Proof. The first bound on distances is standard, while the second bound follows from elementary calculations. □

Lemma 5. *Let N, R be two integers with $N \geq 2, 1 \leq R \leq N$, and $(P_i)_{i \in I}$ be an arbitrary collection of probability measures with $|I| = N$. If $J \subset I$ and $|J| = R$, we have*

$$\left\| \frac{1}{N} \sum_{i \in I} P_i - \frac{1}{R} \sum_{i \in J} P_i \right\|_1 \leq 2 \left(1 - \frac{R}{N}\right).$$

Proof. The result follows by splitting the sum over I in a sum over J and $I \setminus J$, applying the triangle inequality and using the fact that $\|\sum_{j \in J} P_j\|_1 = |J|$. □

7. Proofs for results on graphon functionals

To prove Theorem 6, we observe that polynomial graphons of bounded degree include the graphon w_θ in (26). The proof approximates this smooth graphon w_θ by a piecewise constant graphon, and then uses a lower-bound for such piecewise constants. We prove this lower bound, Lemma 6, first. Similar to the SBM case, this builds on Le Cam's point versus mixture method. We then proceed to prove Theorem 6.

7.1. Auxiliary lower bound

Assume the function w is piecewise constant, with different values taken along blocks corresponding to a regular partition of $[0, 1]^2$ in $k \times k = k^2$ blocks, and k an even integer $k = 2l$, with $l \geq 1$. That defines a law of the form

$$\frac{1}{k^n} \sum_{\varphi \in [k]^n} \bigotimes_{i < j} \text{Be}(Q_{\varphi(i)\varphi(j)}) = P_{e_k, Q},$$

where φ is an element of $[k]^n$ and $Q = Q^\theta$ a given $k \times k$ matrix defined below. In the next statement and proof, E_θ denotes the expectation under this distribution. Denote by O_k the $k \times k$ matrix with only ones as coefficients,

$$O_k = \begin{bmatrix} 1 & \cdots & 1 \\ \vdots & \vdots & \vdots \\ 1 & \cdots & 1 \end{bmatrix},$$

and, for a symmetric $l \times l$ matrix A with coefficients $A_{ij} \in [0, 1]$, define the $k \times k = (2l) \times (2l)$ matrix

$$B := \left[\begin{array}{c|c} A & -A \\ \hline -A & A \end{array} \right].$$

We define $Q = Q^\theta$ as the $k \times k = (2l) \times (2l)$ matrix

$$Q = Q^\theta = \frac{1}{2} \cdot O_k + \theta \cdot B = \frac{1}{2} \cdot O_k + \theta \left[\begin{array}{c|c} A & -A \\ \hline -A & A \end{array} \right]. \quad (33)$$

Lemma 6. *Let $k = 2l$ be an even integer and A an arbitrary symmetric $l \times l$ matrix. Let $Q = Q^\theta$ be the matrix defined in (33). There exists a constant $c_3 > 0$ such that*

$$\inf_{\hat{\theta}} \sup_{\theta \in (-1/2, 1/2)} E_\theta(\hat{\theta}(X) - \theta)^2 \geq \frac{c_3}{n},$$

where the infimum is over all estimators of θ valid under $E_\theta = E_{P_{e_k, Q^\theta}}$.

Proof of Lemma 6. Let $\mathbb{Q} = P_{e_k, Q^\theta}$ be as above. That is,

$$\mathbb{Q} = \frac{1}{k^n} \sum_{\varphi \in [k]^n} P_{Z_\varphi^\theta}, \quad P_{Z_\varphi^\theta} = \bigotimes_{i < j} \text{Be}(Q_{\varphi^{(i)}\varphi^{(j)}}^\theta),$$

with $\{Z_\varphi^\theta\}$ the matrix of general term $z_{ij}(\varphi, \theta) = Q_{\varphi^{(i)}\varphi^{(j)}}^\theta$, for φ ranging over the set $[k]^n$. Let \mathbb{P} denote the Erdős-Renyi $ER(1/2)$ distribution over n nodes, which also corresponds to P_{e_k, Q^θ} for $\theta = 0$. Consider the functional ψ defined as,

$$\psi(P_{e_k, Q^\theta}) = \theta.$$

By definition, for any $\varphi \in [k]^n$, we have $\psi(\mathbb{P}) = \psi(P_{e_k, Q^0}) = 0$ and $\psi(\mathbb{Q}) = \psi(P_{e_k, Q^\theta}) = \theta$. The same computation as in the proof of Theorem 1 now shows that, for B given in the display below,

$$\|\mathbb{P} - \mathbb{Q}\|_1^2 \leq \frac{1}{k^{2n}} \sum_{\varphi, \psi \in [k]^n} \exp \left\{ 4\theta^2 \sum_{i < j} B_{\varphi^{(i)}\varphi^{(j)}} B_{\psi^{(i)}\psi^{(j)}} \right\} - 1.$$

The last term in the bound can be interpreted as an expectation over φ, ψ , where both variables are sampled uniformly from the set of all mappings from $\{1, \dots, n\}$ to $\{1, \dots, k\}$. Recall that $l = k/2$ and for any integer s , denote by $[s]_l$ the integer in $\{1, \dots, l\}$ that equals s modulo l , plus 1. The variable $B_{\varphi^{(i)}\varphi^{(j)}}$ can be written

$$B_{\varphi^{(i)}\varphi^{(j)}} = (-1)^{\varphi^{(i)} > l} (-1)^{\varphi^{(j)} > l} A_{[\varphi^{(i)}]_l [\varphi^{(j)}]_l}.$$

When φ follows the uniform distribution over $[k]^n$, the variables $(\varphi^{(i)})_i$ are independent and are marginally uniform over $[k]$. Also, the variables $((-1)^{\varphi^{(i)} > l})_i$ and $([\varphi^{(i)}]_l)_i$ are independent under the uniform distribution for φ as we show next. If Pr denotes the corresponding distribution, then for any $i \leq n$ and any $s \leq l$

$$\text{Pr} \left[\varphi^{(i)} > l, [\varphi^{(i)}]_l = s \right] = \frac{1}{k} = \frac{1}{2} \cdot \frac{1}{l} = \text{Pr}[\varphi^{(i)} > l] \text{Pr}[\varphi^{(i)}]_l = s].$$

Note the identity holds both for $k \leq n$ and $k > n$. Set $R_i := (-1)^{\varphi^{(i)} > l}$ and $a_{ij} := A_{[\varphi^{(i)}]_l [\varphi^{(j)}]_l}$. Deduce from the previous reasoning that the variables $(R_i)_i$ and $(a_{ij})_{i < j}$ are independent. Now, denoting by E the expectation under Pr ,

$$\begin{aligned} \|\mathbb{P} - \mathbb{Q}\|_1^2 &\leq E \left[\exp \left\{ 4\theta^2 \sum_{i < j} B_{\varphi^{(i)}\varphi^{(j)}} B_{\psi^{(i)}\psi^{(j)}} \right\} \right] - 1 \\ &\leq E \left[E \left[\exp \left\{ 4\theta^2 \sum_{i < j} a_{ij} R_i R_j \right\} \middle| a_{ij} \right] \right] - 1. \end{aligned}$$

As $(R_i)_i$ and $(a_{ij})_{i < j}$ are independent, one can compute the inner expectation in the last display under the distribution of $(R_i)_i$, the a_{ij} 's being fixed. The $(R_i)_i$ form a sample of independent Rademacher variables, hence

$$Z_n := \sum_{i < j} a_{ij} R_i R_j$$

is a Rademacher chaos of order 2 with weights (a_{ij}) . Suppose the matrix A is not identically zero (otherwise the bound below holds trivially). By Lemma 1, for any $c > 1$ one can find $\lambda > 0$ with

$$E \exp\left[\frac{|Z_n|}{\lambda \|Z_n\|_2} \mid a_{ij}\right] \leq c \quad \text{where} \quad \|Z_n\|_2^2 = E[|Z_n|^2 \mid a_{ij}] = \sum_{i < j} a_{ij}^2.$$

Choose $c = 3/2$. By definition, all a_{ij} s are bounded by 1. There is hence a $\lambda > 0$ such that, if $\theta^2 = 2/(\lambda n)$,

$$E \exp[4\theta^2 |Z_n| \mid a_{ij}] - 1 \leq 3/2 - 1 = 1/2.$$

The result now follows from an application of Lemma 3 to the functional ψ . \square

7.2. Proof of the theorem

Proof of Theorem 6. Let us recall the definition, for any $0 \leq \theta \leq 1$, of the function $w = w_\theta$ in (26)

$$w_\theta(x, y) = \frac{1}{2} - \theta(x - \frac{1}{2})(y - \frac{1}{2})$$

and let $\langle w_\theta \rangle$ be its graphon equivalence class. By definition, $\langle w_\theta \rangle$ belongs to \mathcal{P} . One has

$$\int_{[0,1]^2} w_\theta(x, y)^2 dx dy = \frac{1}{4} + c^2 \theta^2,$$

for some constant $c > 0$, so that $\vartheta(\langle w_\theta \rangle) - \vartheta(\langle w_0 \rangle) = c\theta$. The function w_0 is the constant $1/2$, and the density of the data distribution $P_{\langle w \rangle}$ with respect to counting measure on $\{0, 1\}^{n(n-1)/2}$ is

$$p_w(\{x_{ij}\}_{i < j}) = \int \cdots \int \prod_{i < j} \text{Be}(w(u_i, u_j))(x_{ij}) du_1 \cdots du_n,$$

where, for any z in $[0, 1]$ and x_{ij} in $\{0, 1\}$, we have set

$$\text{Be}(z)(x_{ij}) = z^{x_{ij}}(1 - z)^{1-x_{ij}}.$$

Next one shows that $P_{\langle w \rangle}$ is close in the total variation sense to a discrete mixture of the previous Bernoulli-probability distributions, provided the number of points in the mixture is suitably large. To do so, we approximate the function P_n defined by

$$P_n : (u_1, \dots, u_n) \mapsto \prod_{i < j} \text{Be}(w(u_i, u_j))(x_{ij}),$$

by a piecewise constant function $h_{N,\theta} = h_N$, where $[0, 1]^n$ is split into N^n blocks, $N \geq 1$, using a regular grid of $[0, 1]^n$ with points $(i_1/N, \dots, i_n/N)$ and

$0 \leq i_j \leq N$ for all j . To do so, one just replaces $w(u_i, u_j)$ by, say, the value of w on the middle of the block the point (u_i, u_j) belongs to. This defines a function

$$Q_{n,N} : (u_1, \dots, u_n) \mapsto \prod_{i < j} \text{Be}(\bar{w}(u_i, u_j))(x_{ij}),$$

where \bar{w} is constant on every block of the subdivision. Let Q_w^N denote the corresponding measure, with density

$$q_w^N(\{x_{ij}\}_{i < j}) = \int \cdots \int \prod_{i < j} \text{Be}(\bar{w}(u_i, u_j))(x_{ij}) du_1 \cdots du_n.$$

Taking $w = w_\theta$ as above, the function P_n is a polynomial in u_1, \dots, u_n , and its degree with respect to each variable u_i is $n - 1$. The partial derivatives of P_n can be computed, and each of them can be seen to be bounded by $n - 1$: For each variable, only $n - 1$ non-zero terms appear when evaluating the partial derivative, and each term is uniformly bounded by 1. Consequently, if (u_1, \dots, u_n) and (u'_1, \dots, u'_n) belong to the same block,

$$|P_n(u_1, \dots, u_n) - P_n(u'_1, \dots, u'_n)| \leq (n - 1) \sum_{i=1}^n |u_i - u'_i| \leq n^2/N.$$

For $w = w_\theta$ as above, we can thus bound the total variation distance as

$$\begin{aligned} \|P_{(w)} - Q_w^N\|_1 &\leq \sum_{x \in \{0,1\}^{\frac{n(n-1)}{2}}} |p_w - q_w^N|(x) \\ &\leq n^2 \max_{x \in \{0,1\}^{\frac{n(n-1)}{2}}} |p_w - q_w^N|(x) \leq n^2(n^2/N) = n^4/N. \end{aligned}$$

Each probability measure Q_w^N is a mixture of N^n distributions, each of which in turn corresponds to a block in the subdivision of $[0, 1]^n$. One can rewrite

$$Q_w^N = \frac{1}{N^n} \sum_{\varphi \in [N]^n} \bigotimes_{i < j} \text{Be}(M_{\varphi(i)\varphi(j)}),$$

where the matrix $M = (M_{pq})_{1 \leq p, q \leq N}$ is the symmetric matrix with terms

$$M_{pq} = w_\theta \left(\frac{p - \frac{1}{2}}{N}, \frac{q - \frac{1}{2}}{N} \right).$$

If N is even, which one can assume without loss of generality, the matrix M is exactly of the same form as Q in (33), with elements in $(0, 1)$, so one can use the bound in $\|\cdot\|_1$ -distance between measures obtained in the proof of Lemma 6. Note that the argument remains valid even if the number of classes exceeds the number of observations n , which will be of importance below. For a small constant c and $\theta^2 = \kappa/n$, we obtain

$$\|Q_{w_\theta}^N - Q_{w_0}^N\|_1 \leq c$$

for κ sufficiently small. Choosing $N = Cn^4$, for $C > 0$ large enough, leads to

$$\begin{aligned} \|P_{\langle w_\theta \rangle} - P_{\langle w_0 \rangle}\|_1 &\leq \|P_{\langle w_\theta \rangle} - Q_{w_\theta}^N\|_1 + \|Q_{w_\theta}^N - Q_{w_0}^N\|_1 + \|Q_{w_0}^N - P_{\langle w_0 \rangle}\|_1 \\ &\leq n^4/N + c + 0 \leq c' < 1/2. \end{aligned}$$

An application of Lemma 3 with the functional $\psi(P_{\langle w \rangle}) := \psi(\langle w \rangle)$ concludes the proof of the lower bound in Theorem 6 in the case where $\psi(\cdot) = \vartheta(\cdot)$. The lower bound for a general ψ follows by the same proof, noting that the specific form of the functional only comes in through the difference $\psi(\langle w_\theta \rangle) - \psi(\langle w_0 \rangle)$, which behaves as for $\vartheta(\cdot)$ by assumption.

For the upper-bound, we first link the squared distance to the truth for the functional to the squared L^2 -distance of corresponding graphons. Let w, w_1 be two fixed graphon functions, and suppose that at least one of these is non constant (almost everywhere), which means that either $\vartheta(\langle w \rangle) > 0$ or $\vartheta(\langle w_1 \rangle) > 0$. Then, writing simply \int to denote the double integral on $[0, 1]^2$,

$$\begin{aligned} \vartheta(\langle w_1 \rangle) - \vartheta(\langle w \rangle) &= \left\{ \int (w_1 - \int w_1)^2 \right\}^{1/2} - \left\{ \int (w - \int w)^2 \right\}^{1/2} \\ &= \frac{\int (w_1 - \int w_1)^2 - \int (w - \int w)^2}{\left\{ \int (w_1 - \int w_1)^2 \right\}^{1/2} + \left\{ \int (w - \int w)^2 \right\}^{1/2}}, \end{aligned}$$

where the denominator is nonzero by assumption on w, w_1 ; we henceforth denote it c . Then

$$\begin{aligned} \vartheta(\langle w_1 \rangle) - \vartheta(\langle w \rangle) &\leq c^{-1} \int \left\{ w_1 - w - \int (w_1 - w) \right\} \left\{ w_1 + w - \int (w_1 + w) \right\} \\ &\leq c^{-1} \left[\int \left\{ w_1 - w - \int (w_1 - w) \right\}^2 \right]^{1/2} \left[\int \left\{ w_1 + w - \int (w_1 + w) \right\}^2 \right]^{1/2}. \end{aligned}$$

The two factors in brackets are bounded as follows: For the second term, apply the inequality $(a + b)^2 \leq 2a^2 + 2b^2$, followed by $\sqrt{u + v} \leq \sqrt{u} + \sqrt{v}$, which yields

$$\begin{aligned} &\left[\int \left\{ w_1 + w - \int (w_1 + w) \right\}^2 \right]^{1/2} \\ &\leq \sqrt{2} \left[\int (w_1 - \int w_1)^2 + \int (w - \int w)^2 \right]^{1/2} \leq \sqrt{2}c. \end{aligned}$$

For the first term, use $0 \leq \int (g - \int g)^2 \leq \int g^2$ for a bounded measurable g , as one integrates over $[0, 1]^2$. That yields

$$\{\vartheta(\langle w_1 \rangle) - \vartheta(\langle w \rangle)\}^2 \leq 2 \int (w_1 - w)^2,$$

and this inequality clearly still holds true in case $\vartheta(\langle w \rangle) = \vartheta(\langle w_1 \rangle) = 0$. One concludes that

$$\{\vartheta(\langle w_1 \rangle) - \vartheta(\langle w \rangle)\}^2$$

$$\leq 2 \inf_{T \in \mathcal{T}} \int \int_{[0,1]^2} |w_1(T(x), T(y)) - w(x, y)|^2 dx dy =: \delta^2(w_1, w),$$

where \mathcal{T} is the set of all measure-preserving bijections of $[0, 1]$. Indeed, the previous inequalities hold true for any choice of representer of the graphon w_1 , so one can take the infimum over \mathcal{T} in the previous bounds. By Corollary 3.6 of [23], for data X generated from P_w , there exists an estimator $\hat{w} = \hat{w}(X)$ that satisfies $E_{P_w}[\delta^2(\hat{w}, w)] \leq C(\log n/n)$. Since $\langle w \rangle$ has a representer that belongs to \mathcal{P}_B by assumption, it belongs in particular to the Hölder class $\Sigma(1, L)$, provided L is chosen large enough. For the plug-in estimator $\hat{\vartheta}(X) := \vartheta(\hat{w})$, combining the previous result with the last display implies

$$E \left[\{\hat{\vartheta}(X) - \vartheta(\langle w \rangle)\}^2 \right] \leq C \frac{\log n}{n},$$

for C large enough depending only on \mathcal{P}_B , which concludes the proof. □

Appendix A: Upper bounds computable in polynomial-time

This section generalizes the polynomial-time estimate in Section 4.2 to $k \geq 2$ classes, by combining the spectral clustering method of Lei and Rinaldo [25] with a refinement due to Lei and Zhu [26]. The latter is based on a sample splitting, and under appropriate conditions on the connectivity matrix recovers the labels *exactly*, with high probability. Theorem 7 below shows that, under additional conditions, this polynomial-time estimator achieves the minimax rate.

A.1. Spectral estimation for $k \geq 2$ classes

Recall the assumed form of the connectivity matrix M^θ in (12). The conditions of the next results are in terms of an ‘aggregated’ $(k - 1) \times (k - 1)$ matrix N obtained from M^θ by merging the first and second row/columns when $\theta = 0$, that is

$$N = \begin{bmatrix} 1/2 & a_1 & \cdots & a_{k-2} \\ a_1 & b_{11} & \cdots & b_{1k-2} \\ \vdots & \vdots & & \vdots \\ a_{k-2} & b_{1k-2} & \cdots & b_{k-2k-2} \end{bmatrix}.$$

Recall that φ denotes the true labelling map. Define a labelling $\psi : [n] \rightarrow [k - 1]$ by $\psi(v) = 1$ if $\varphi(v) \in \{1, 2\}$ and $\psi(v) = \varphi(v) - 1$ if $v \in \{3, \dots, k\}$. That is, we ‘aggregate’ nodes of label 1 or 2 in one class and renumber the remaining labels so that the label set is, now, $[k - 1]$. Following [26], we write $g_v = \psi(v)$ for the true (aggregated) label of node $v \in [n]$ and $\mathcal{I}^{(l)} = \{v \in [n] : g_v = l\}$.

The algorithm **Spec- θ** specified in the frame below has three steps. First, one runs the exact label recovery algorithm **V-Clust** of Lei and Zhu [26] for $K = k - 1$ classes. Under some conditions on the matrix N , see (A1)–(A2) below, this finds the ‘aggregate’ labels ψ above *up to label permutation* with high

probability. Then the aim is to recover the aggregated class with original labels 1 and 2. Due to the label switching issue, this requires some extra condition on N . For simplicity (see also comments below) we assume in (A3) that the diagonal terms b_{ii} are separated from $1/2$, which enables to estimate the aggregated class label 1 by comparing diagonal empirical connectivities to $1/2$. Finally, in a third step one can run the spectral algorithm \mathcal{S}_2 from Section 3.1 on the nodes found at the previous step.

We set $K = k - 1$ and assume that, for a large enough universal constant C :

- (A1) N is full rank and any two rows of N are separated by at least $\gamma = \gamma(K) > 0$ in ℓ_2 -norm.
 (A2) For $\lambda = \lambda(K)$ the smallest absolute eigenvalue of N ,

$$n\lambda(K)\gamma(K) \geq CK^{4.5}, \quad n\gamma(K)^2 \geq CK^3 \log n, \quad n \geq CK^3.$$

- (A3) For all $i \in \{1, \dots, k - 2\}$,

$$|b_{ii} - 1/2| \geq \kappa,$$

$$\text{where } \kappa = \kappa(K) \geq C\sqrt{K(\log n)/n}.$$

Algorithm: Spectral method for estimation of θ (Spec- θ)

Input: adjacency matrix X (where we set $X_{ii} = 0$), number of classes k

Subroutines: V-Clust (Lei-Zhu), Initial community recovery \mathcal{S} (Lei-Rinaldo), Spectral algorithm \mathcal{S}_2 for $k = 2$ (Section 3.1)

1. Apply V-Clust on adjacency matrix X using $k - 1$ classes, \mathcal{S} and $V = 2$

$$\hat{g} = \text{V-Clust}(X, k - 1, V, \mathcal{S}).$$

2. Set $\hat{\mathcal{I}}^{(1)} = \{v \in [n] : \hat{g}_v = \hat{\ell}\}$, where

$$\hat{\ell} = \operatorname{argmin}_{\ell \in [k-1]} \left| \frac{1}{\binom{|\hat{g}^{-1}(\ell)|}{2}} \sum_{i < j, i, j \in \hat{g}^{-1}(\ell)} X_{ij} - \frac{1}{2} \right|$$

3. Run spectral algorithm \mathcal{S}_2 for $k = 2$ on corresponding nodes and set

$$\hat{\theta} = \mathcal{S}_2(X^{\hat{\mathcal{I}}^{(1)}}),$$

where $X^{\hat{\mathcal{I}}^{(1)}}$ is the induced adjacency matrix over nodes in $\hat{\mathcal{I}}^{(1)}$.

Comments on (A1)–(A3) follow below. For a version for sparse graphs, see Appendix B.

Theorem 7. *In the fixed design SBM model with k classes, under the assumptions (A1)–(A3), let us set, for c a small enough universal constant and*

$$K = k - 1,$$

$$T_K := c \frac{\lambda(K)\gamma(K)^{1/2}}{K^{5/4}} \wedge \frac{\kappa}{4}. \quad (34)$$

Then the obtained $\hat{\theta}$ from algorithm **Spec- θ** satisfies, for C_3 a large enough constant,

$$\sup_{|\theta| \leq T_K, \varphi \in \Sigma_e} E_{\theta, \varphi} \left[\hat{\theta} - \theta \right]^2 \leq C_3 \frac{k}{n}.$$

Proof. This is a special case of Theorem 9, in Appendix B below. \square

The algorithm **Spec- θ** , unlike the likelihood method considered below, only uses the fact that the connectivity matrix is of the form M^θ , but does not use specific knowledge of the vector a and matrix B to compute θ .

Comments on the assumptions. Conditions (A1) and (A2) are typical for spectral methods; their specific form is that assumed by Lei and Zhu [26], with the initial recovery algorithm being that of Lei and Rinaldo [25]. If K is fixed independently of n , then (A2) follows from (A1) if n is large enough. Condition (A3) is specific to our problem, and assumed in this form only for simplicity of exposition: To identify the special cluster arising from the $1/2$ coefficient in the matrix N (step 2. in **Spec- θ**), some identifiability condition is needed, because even the refined spectral clustering algorithm of [26] can only recover the original labels up to a permutation. Condition (A3) is similar in spirit to condition (21), but weaker. It can be replaced with any other condition that ensures cluster 1 can be identified from a noisy, permuted version of N (with noise amplitude going to zero fast, as k/n). Note that, if k is fixed and n large enough, (A3) simply requires the diagonal terms of B to differ from $1/2$.

Finally, a comment on T_K in (34). The label recovery in Steps 1–2 is run with $k - 1$ classes, and hence joins two of the k classes in the sample. The restriction on the range of θ ensures the classes joined are the first two, with high probability. Indeed, here we are interested in the situation where θ may be small, which makes identification of labels difficult, and the rate slow; if θ is large, the problem becomes easier. Again, note that if k is fixed, the condition simply requires that $|\theta|$ is smaller than a given constant.

A.2. Simulation study

Those estimators described above that are computationally feasible—the spectral and sample splitting estimators for $k = 2$, and the **Spec- θ** estimator for $k > 2$ —can be tested in simulation: Draw n vertices from a stochastic block model as in (13) with a given value of θ , compute the respective estimate, and report the empirical quadratic risk. Figure 1 shows how the risk develops as a function of sample size for different values of θ , for the two-community model (9). For $k > 2$ communities, the model is given by the connectivity matrix (12). Simulation results for $k = 5$, with $a_1 = \frac{1}{12}$, $a_2 = \frac{11}{12}$ and $a_3 = 1$, are shown in Figure 2. As is visible in Figures 1 and 2, smaller values of θ correspond overall to a larger risk, and a much slower decay of the empirical risk curves.

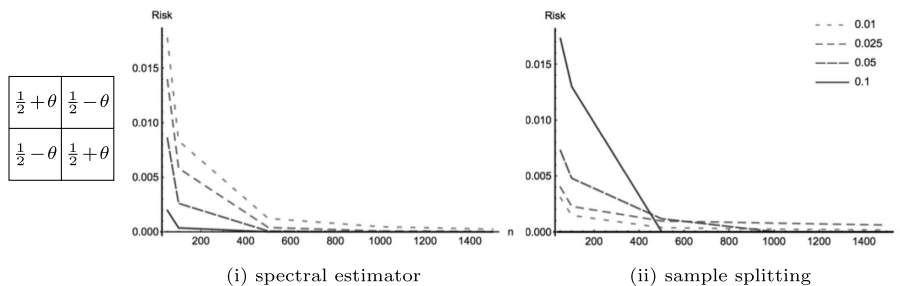


FIG 1. Estimation of θ in the two-class case, using (i) the spectral estimator (23) and (ii) sample splitting. Graphs of size $n = 50, 100, 500, 1000, 1500$ are generated from the graphon on the right, for $\theta = 0.01, 0.025, 0.05, 0.1$. Shown is the empirical risk (computed over 1000 experiments) as a function of the sample size n .

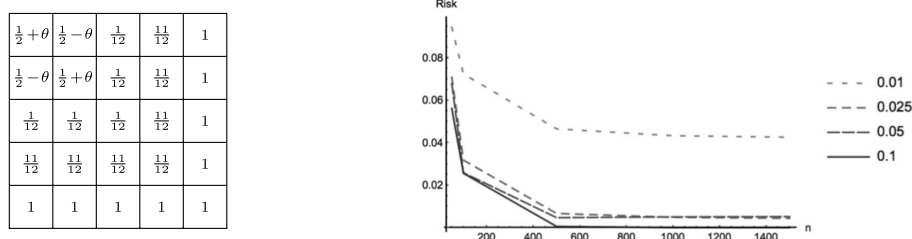


FIG 2. Estimation of θ in the five-class case, using the Spec- θ algorithm. Graphs of size $n = 50, 100, 500, 1000, 1500$ are generated from the graphon on the right, for $\theta = 0.01, 0.025, 0.05, 0.1$. Shown is the empirical risk (computed over 1000 experiments) as a function of the sample size n . For small values of θ , convergence slows visibly.

This illustrates our theoretical finding that there exists a range of parameters corresponding to two classes that become close where estimation is much slower.

Appendix B: Extension to sparse graphs

So far, we have for simplicity considered *dense* graphs, in the sense that at least some elements of the connectivity matrix (e.g. $1/2 + \theta$ or $1/2 - \theta$) are bounded away from zero.

B.1. Two classes

An α_n -sparse SBM model is generally defined as one in which the connectivity matrix M can be written, for α_n a sequence going to 0 with n , as $M = \alpha_n M_0$, for M_0 a nonnegative symmetric matrix with maximum entry 1 [e.g. 8, 25]. Here, we assume that the connectivity matrix is $M^\theta(\alpha_n)$ with

$$M^\theta(\alpha_n) = \alpha_n M^\theta, \quad (35)$$

and M^θ as in (12). Then the largest coefficient of M^θ is between $\alpha_n/2$ and α_n , as the coefficients of the upper 2×2 block are $\alpha_n(1/2 \pm \theta)$. We also set, for $\theta \in [-1/2, 1/2]$,

$$Q^\theta(\alpha_n) = \alpha_n Q^\theta, \quad Q^\theta = \begin{bmatrix} \frac{1}{2} + \theta & \frac{1}{2} - \theta \\ \frac{1}{2} - \theta & \frac{1}{2} + \theta \end{bmatrix}. \tag{36}$$

In constructing upper bounds below, we assume that for C_s a large enough constant,

$$(B0) \quad \alpha_n \geq C_s \frac{\log n}{n},$$

as up to a constant $\log n/n$ is the typical boundary between the moderately sparse and very sparse situations, the later requiring different tools, see [25]. For simplicity we also assume that α_n is known for the upper-bound results.

Theorem 8. *Consider a stochastic blockmodel (3) with $k = 2$ specified by $P_\theta = P_{e, Q^\theta(\alpha_n)}$ with $e, Q^\theta(\alpha_n)$ given by (8)-(36). There exists a constant $c_1 > 0$ such that for all $n \geq 2$,*

$$\inf_T \sup_{\theta \in [-1/2, 1/2]} E_\theta [T(X) - \theta]^2 \geq c_1 \left(1 \wedge \frac{1}{n\alpha_n} \right),$$

where the infimum is taken over all estimators T of θ in the model \mathcal{M} . Furthermore, if $\Delta_n = X - \alpha_n J/2$, and $\lambda_1^a(\Delta_n)$ the largest absolute eigenvalue of Δ_n , set $\tilde{\theta} := \lambda_1^a(\Delta_n)/\{(n-1)\alpha_n\}$. Then, under (B0), for some constant $C > 0$ and $n \geq 2$,

$$\sup_{\theta \in [-1/2, 1/2]} E_\theta [(\tilde{\theta} - \theta)^2] \leq \frac{C}{n\alpha_n}.$$

B.2. $k \geq 2$ classes

The case of k classes carries over to the sparse situation as follows. The lower bound result is only modified by a scaling factor $1/\alpha_n$. For upper bounds, considering the more easily computable spectral algorithm **Spec- θ** only, Assumption (A2) is replaced by (B2) below, where N has the same definition as in Appendix A.

(B2) For $\lambda = \lambda(K)$ the smallest absolute eigenvalue of N , there exists $C > 0$ such that

$$n\alpha_n \lambda(K) \gamma(K) \geq CK^{4.5}, \quad n\alpha_n \gamma(K)^2 \geq CK^3 \log n, \quad n \geq CK^3.$$

Theorem 9. *Consider a stochastic blockmodel (3) with $k \geq 2$ classes specified by \mathcal{M}_k in (13), that is $P_\theta = P_{e_k, M^\theta}$ with e_k, M^θ given by (11)-(35), for fixed matrices A, B with arbitrary coefficients. There exists a constant $c_3 = c_3(\rho) > 0$, independent of A, B , such that, for all $n \geq 12k$,*

$$\inf_T \sup_{\theta \in [-1/2, 1/2]} E_\theta [T(X) - \theta]^2 \geq c_3 \left(1 \wedge \frac{k}{n\alpha_n} \right),$$

where the infimum is taken over all estimators T of θ in the model \mathcal{M}_k . Let \mathcal{S}_{2,α_n} be the algorithm for $k = 2$ classes in the sparse case described in Theorem 8. Consider the fixed-design setting and suppose (B0), (B2), (A1) and (A3) are satisfied. Then the algorithm $\text{Spec-}\theta$ used with subroutine \mathcal{S}_{2,α_n} outputs an estimator $\hat{\theta}$ that satisfies, for T_K as in (34),

$$\sup_{|\theta| \leq T_K, \varphi \in \Sigma_e} E_{\theta, \varphi} \left(\hat{\theta} - \theta \right)^2 \leq C \frac{k}{n\alpha_n}.$$

Similar comments as for Theorems 2–7 can be made. Also, in the case that k does not grow with n , then (B2) follows from (B0) for n larger than a fixed constant. The proof of the lower bound in Theorem 9 is similar to that of Theorem 2 using the normalisation as in the proof of Theorem 8 and is omitted. The upper bound result includes that of Theorem 7 and is proved in Appendix D.

Appendix C: Remaining proofs: likelihood-based upper bounds

The proof for $k = 2$ below analyzes the least-squares criterion directly. For $k \geq 2$ classes, we ‘isolate’ the part corresponding to the 2×2 submodel, by controlling the number of errors in recovering the labels of the corresponding 2 classes. We then invoke the result for the case $k = 2$.

C.1. Interpretation as a pseudo-likelihood

We first justify the interpretation of the estimator $\hat{\theta}$ in (15) as a maximum (pseudo-)likelihood estimate. In the fixed design model, suppose the data is Gaussian $\mathcal{N}(\theta_{ij}, 1)$ instead of Bernoulli $\text{Be}(\theta_{ij})$. This suggests defining a (pseudo-)log-likelihood $\ell_n(\sigma, \theta)$ as follows, with $c_n = \binom{n}{2} \log(2\pi)$,

$$\begin{aligned} -2\ell_n(\sigma, \theta) &= \sum_{i < j} (X_{ij} - Q_{\sigma(i)\sigma(j)}^\theta)^2 + c_n \\ &= \sum_{i < j, \sigma(i)=\sigma(j)} (X_{ij} - (\frac{1}{2} + \theta))^2 + \sum_{i < j, \sigma(i) \neq \sigma(j)} (X_{ij} - (\frac{1}{2} - \theta))^2 + c_n \\ &= \binom{n}{2} \theta^2 + \theta \left(\sum_{i < j, \sigma(i)=\sigma(j)} (1 - 2X_{ij}) - \sum_{i < j, \sigma(i) \neq \sigma(j)} (1 - 2X_{ij}) \right) + C_n(X), \end{aligned}$$

for a constant $C_n(X)$ depending only on n and X . Setting $b_n = \binom{n}{2}$ and

$$2Z_n(\sigma, X) := - \sum_{i < j, \sigma(i)=\sigma(j)} (1 - 2X_{ij}) + \sum_{i < j, \sigma(i) \neq \sigma(j)} (1 - 2X_{ij}),$$

it is enough to study the function $g_n(\theta, \sigma) := b_n \theta^2 - 2Z_n(\sigma, X)\theta$, which satisfies

$$g_n(\theta, \sigma) = b_n \left(\theta - \frac{Z_n(\sigma, X)}{b_n} \right)^2 - \frac{Z_n(\sigma, X)^2}{b_n} \geq - \frac{Z_n(\sigma, X)^2}{b_n}.$$

Consequently, the pseudo maximum likelihood estimator $(\hat{\theta}, \hat{\sigma})$ is given by (15) as claimed.

C.2. Upper bound result, two classes

Proof of Theorem 3. We first prove the result in the fixed design case. Let θ_0, σ_0 denote the true values of θ, φ . The aim is to show that $E_{\theta_0, \sigma_0}(\hat{\theta} - \theta_0)^2 \leq C/n$ holds uniformly in θ_0, σ_0 . For a given $\sigma \in 2^{[n]}$,

$$r_{ij} := X_{ij} - 1/2 + (-1)^{\mathbb{1}_{\sigma_0(i)=\sigma_0(j)}}\theta_0 \quad \text{and} \quad R_n(\sigma) := \sum_{\sigma(i)=\sigma(j)} r_{ij} - \sum_{\sigma(i)\neq\sigma(j)} r_{ij}.$$

One can write, for any $\sigma \in 2^{[n]}$,

$$\begin{aligned} Z_n(\sigma_0, X) &= b_n\theta_0 + \sum_{i<j} (-1)^{\mathbb{1}_{\sigma_0(i)\neq\sigma_0(j)}} r_{ij} = b_n\theta_0 + R_n(\sigma_0), \\ Z_n(\sigma, X) &= \theta_0\delta(\sigma, \sigma_0) + R_n(\sigma), \end{aligned}$$

where we have set

$$\delta(\sigma, \sigma_0) = \sum_{i<j} (-1)^{\mathbb{1}_{\sigma_0(i)\neq\sigma_0(j)}} (-1)^{\mathbb{1}_{\sigma(i)\neq\sigma(j)}}. \tag{37}$$

For any $t > 0$ and $t_n = M_2/\sqrt{n}$, and for a large enough M_2 to be chosen below,

$$\begin{aligned} P_{\theta_0}[\sqrt{n}|\hat{\theta} - \theta_0| \geq t] &= P_{\theta_0}[\sqrt{n}|\hat{\theta} - \theta_0| \geq t] \mathbb{1}_{|\theta_0| \leq t_n} + P_{\theta_0}[\sqrt{n}|\hat{\theta} - \theta_0| \geq t] \mathbb{1}_{|\theta_0| > t_n} \\ &=: \mathcal{P}_1(t) + \mathcal{P}_2(t). \end{aligned}$$

By definition of $\hat{\theta}$, with $\delta(\sigma, \sigma_0)$ defined in (37),

$$b_n\hat{\theta} = Z_n(\hat{\sigma}, X) = \theta_0\delta(\hat{\sigma}, \sigma_0) + R_n(\hat{\sigma}) \quad \text{and} \quad Z_n(\sigma_0, X) = \theta_0b_n + R_n(\sigma_0).$$

For any $t \geq 4M_2$, using that $|\delta(\hat{\sigma}, \sigma_0)| \leq b_n$,

$$\begin{aligned} \mathcal{P}_1(t) &\leq P_{\theta_0} \left[\sqrt{n}|\theta_0| \left| \frac{\delta(\hat{\sigma}, \sigma_0)}{b_n} - 1 \right| + \frac{\sqrt{n}}{b_n} |R_n(\hat{\sigma})| \geq t \right] \mathbb{1}_{|\theta_0| \leq t_n} \\ &\leq P_{\theta_0} \left[\frac{\sqrt{n}}{b_n} |R_n(\hat{\sigma})| \geq t - 2M_2 \right] \\ &\leq P_{\theta_0} \left[\sup_{\sigma \in 2^{[n]}} |R_n(\sigma)| \geq \frac{b_n}{\sqrt{n}} \frac{t}{2} \right]. \end{aligned}$$

For $\mathcal{P}_2(t)$, there are two cases, depending on the sign of θ_0 ,

$$\mathcal{P}_2(t) \leq P_{\theta_0} \left[\sqrt{n}|\hat{\theta} - \theta_0| \geq t \right] \mathbb{1}_{\theta_0 > t_n} + P_{\theta_0} \left[\sqrt{n}|\hat{\theta} - \theta_0| \geq t \right] \mathbb{1}_{\theta_0 < -t_n}$$

$$\begin{aligned} &\leq \left\{ P_{\theta_0} \left[\sqrt{n}|\hat{\theta} - \theta_0| \geq t, Z_n(\hat{\sigma}, X) \geq 0 \right] \right. \\ &\quad \left. + P_{\theta_0} \left[\sqrt{n}|\hat{\theta} - \theta_0| \geq t, Z_n(\hat{\sigma}, X) < 0 \right] \right\} \mathbb{1}_{\theta_0 > t_n} \\ &\quad + \left\{ P_{\theta_0} \left[\sqrt{n}|\hat{\theta} - \theta_0| \geq t, Z_n(\hat{\sigma}, X) \geq 0 \right] \right. \\ &\quad \left. + P_{\theta_0} \left[\sqrt{n}|\hat{\theta} - \theta_0| \geq t, Z_n(\hat{\sigma}, X) < 0 \right] \right\} \mathbb{1}_{\theta_0 < -t_n} \end{aligned}$$

Let us discuss the term $\theta_0 > t_n$ first and note that if $Z_n(\hat{\sigma}, X) \geq 0$, then $Z_n(\hat{\sigma}, X) = |Z_n(\hat{\sigma}, X)| \geq |Z_n(\sigma_0, X)| \geq Z_n(\sigma_0, X)$ using the definition of $\hat{\sigma}$ as a maximum. First,

$$\begin{aligned} &P_{\theta_0} \left[\sqrt{n}|\hat{\theta} - \theta_0| \geq t, Z_n(\hat{\sigma}, X) \geq 0 \right] \mathbb{1}_{\theta_0 > t_n} \\ &\leq P_{\theta_0} \left[\frac{\sqrt{n}}{b_n} |Z_n(\hat{\sigma}, X) - Z_n(\sigma_0, X) + R_n(\sigma_0)| \geq t, Z_n(\hat{\sigma}, X) \geq 0 \right] \mathbb{1}_{\theta_0 > t_n} \\ &\leq P_{\theta_0} \left[\frac{\sqrt{n}}{b_n} (Z_n(\hat{\sigma}, X) - Z_n(\sigma_0, X)) + \frac{\sqrt{n}}{b_n} |R_n(\sigma_0)| \geq t \right] \mathbb{1}_{\theta_0 > t_n} \\ &\leq P_{\theta_0} \left[\frac{\sqrt{n}}{b_n} \theta_0 (\delta(\hat{\sigma}, \sigma_0) - b_n) + \frac{\sqrt{n}}{b_n} R_n(\hat{\sigma}) + 2 \frac{\sqrt{n}}{b_n} |R_n(\sigma_0)| \geq t \right] \mathbb{1}_{\theta_0 > t_n} \\ &\leq P_{\theta_0} \left[3 \frac{\sqrt{n}}{b_n} \sup_{\sigma \in 2^{[n]}} |R_n(\sigma)| \geq t \right], \end{aligned}$$

where the first three inequalities use identities obtained for $Z_n(\hat{\sigma}, X)$, $Z_n(\sigma_0, X)$ above and the inequality obtained before the display, and the last inequality uses $\delta(\hat{\sigma}, \sigma_0) - b_n \leq 0$ and $\theta_0 \geq 0$.

Second, as $Z_n(\hat{\sigma}, X) < 0$ implies $Z_n(\hat{\sigma}, X) < -|Z_n(\sigma_0, X)|$ by definition of the maximum,

$$\begin{aligned} &P_{\theta_0} \left[\sqrt{n}|\hat{\theta} - \theta_0| \geq t, Z_n(\hat{\sigma}, X) < 0 \right] \mathbb{1}_{\theta_0 > t_n} \\ &\leq P_{\theta_0} \left[|Z_n(\hat{\sigma}, X)| < -|Z_n(\sigma_0, X)| \right] \mathbb{1}_{\theta_0 > t_n} \\ &\leq P_{\theta_0} \left[|Z_n(\hat{\sigma}, X)| < -|\theta_0| b_n + |R_n(\sigma_0)| \right] \mathbb{1}_{\theta_0 > t_n} \\ &\leq P_{\theta_0} \left[\theta_0 (\delta(\hat{\sigma}, \sigma_0) + b_n) < 2 \sup_{\sigma \in 2^{[n]}} |R_n(\sigma)| \right] \mathbb{1}_{\theta_0 > t_n} \\ &\leq P_{\theta_0} \left[2 \sup_{\sigma \in 2^{[n]}} |R_n(\sigma)| > \frac{t_n b_n}{8} \right], \end{aligned}$$

where for the last inequality we have used the lower bound on δ obtained in Lemma 7.

The case $\theta_0 < -t_n$ is treated in a symmetric way, by distinguishing the two cases $Z_n(\hat{\sigma}, X) < 0$ and $Z_n(\hat{\sigma}, X) \geq 0$ respectively. To obtain a deviation bound for $\hat{\theta}$, it is enough to study the supremum of the process $|R_n(\sigma)|$. For any given σ and $y > 0$, by Hoeffding's inequality,

$$P_{\theta_0} [|R_n(\sigma)| > y] \leq 2 \exp\{-2y^2/b_n\}.$$

A union bound now leads to

$$P \left[\sup_{\sigma \in 2^{[n]}} |R_n(\sigma)| \geq y \right] \leq 2^n \exp \{ -2y^2/b_n \} \quad \text{for any } y > 0.$$

This bound is smaller than $\exp\{-2n\}$ if one chooses $y = n^{3/2}$. Combining the bounds obtained previously, and choosing M_2 above as $M_2 = 64$, one deduces

$$P[\sqrt{n}|\hat{\theta} - \theta_0| \geq t] \leq 6e^{-2n} \quad \text{for any } t \geq 4.$$

The deviation bound in turn implies the bound in expectation

$$\begin{aligned} E[n(\hat{\theta} - \theta_0)^2] &= E[n(\hat{\theta} - \theta_0)^2 \mathbb{1}_{\sqrt{n}|\hat{\theta} - \theta_0| \leq 4}] + E[n(\hat{\theta} - \theta_0)^2 \mathbb{1}_{\sqrt{n}|\hat{\theta} - \theta_0| > 4}], \\ &\leq 16 + 6ne^{-2n}, \end{aligned}$$

where for the second term we have used $|\hat{\theta} - \theta_0| \leq 1$, as Θ has diameter 1. This concludes the proof of Theorem 3 in the fixed design case.

In the random design case, one slightly updates the definition of r_{ij} . Here, the design is specified by φ , which is now random, but one can consider

$$\tilde{r}_{ij} = X_{ij} - E[X_{ij} | \varphi].$$

By definition, $E[\tilde{r}_{ij}] = E[E[\tilde{r}_{ij} | \varphi]] = 0$. Now one can follow the proof in the fixed design case by writing all statements conditionally on φ . As conditionally on φ the variables r_{ij} are independent and centered, the arguments leading to the various upper bounds remain unchanged. As the upper-bounds themselves do not depend on φ , the bounds also hold unconditionally. \square

What remains to be shown is the bound on δ used above:

Lemma 7. For any $\sigma_0, \sigma \in \Sigma$, with $\delta(\sigma, \sigma_0)$ defined in (37) and $b_n = \binom{n}{2}$,

$$-\frac{7}{8}b_n \leq \delta(\sigma, \sigma_0) \leq b_n \quad \text{for any } n \geq 5.$$

Proof. The upper bound corresponds to the number of terms in the sum. For the lower bound, denote $\mathcal{C}_1 = \sigma_0^{-1}(\{1\}) = \{i, \sigma_0(i) = 1\}$. By symmetry, one can always assume $|\mathcal{C}_1| \geq n/2$, otherwise one works with $\mathcal{C}_2 = \sigma_0^{-1}(\{2\})$. The number $T_\sigma(\mathcal{C}_1)$ of pairs $(i, j) \in \mathcal{C}_1 \times \mathcal{C}_1$ for which $\sigma(i) \neq \sigma(j)$ is at most $2N_1(\sigma)N_2(\sigma)$, if $N_i(\sigma) = |\sigma^{-1}(\{i\}) \cap \mathcal{C}_1|$, $i = 1, 2$. This implies $T_\sigma(\mathcal{C}_1) \leq |\mathcal{C}_1|^2/2$, using the inequality $p(q - p) \leq q^2/4$, for any $0 \leq p \leq q$. Thus the number of positive elements in the sum defining $\delta(\sigma, \sigma_0)$ is at least $(|\mathcal{C}_1|^2 - |\mathcal{C}_1|^2/2 - |\mathcal{C}_1|)/2$, where $|\mathcal{C}_1|$ corresponds to the diagonal terms and the division by 2 to the fact that the sum is restricted to $i < j$ only (note that the general term of the sum defining δ is symmetric in i, j). This is at least $|\mathcal{C}_1|^2/8$ if $|\mathcal{C}_1| \geq 3$. Hence,

$$\delta(\sigma, \sigma_0) \geq -b_n + 2\frac{|\mathcal{C}_1|^2}{8},$$

which is the desired bound in view of $|\mathcal{C}_1|^2 \geq b_n/2$. \square

C.3. Upper bound result, k classes

Proof of Theorem 4. First consider the fixed design case: As in the proof of Theorem 3, let θ_0, σ_0 denote the true values of θ, φ . The aim is to show that

$$\sup_{|\theta_0| \leq \kappa, \sigma_0 \in \Sigma_e} E_{\theta_0, \sigma_0}(\hat{\theta} - \theta)^2 \leq C \frac{k}{n}.$$

Let us denote by Z^0 and \tilde{Z} the matrices of general terms $Z_{i,j}^0 := M_{\sigma_0(i)\sigma_0(j)}^{\theta_0}$ and $\tilde{Z}_{ij} := M_{\tilde{\sigma}(i)\tilde{\sigma}(j)}^{\tilde{\theta}}$ respectively, with $\tilde{\sigma}, \tilde{\theta}$ given by (17) and $i \neq j$. One interpretation of (17) is that the matrix \tilde{Z} provides the best fit to the data X_{ij} with respect to the squared L^2 loss, when optimising over $\Sigma_e \times \Theta_n$.

As a first step, we show that \tilde{Z} and Z^0 are close with high probability, a result in the spirit of Gao et al [16], Theorem 2.1. This follows from Lemma 8 below, which states that $\|\tilde{Z} - Z^0\|^2 \leq Cn \log k$ with probability at least $1 - e^{-n \log k}$, where $\|\cdot\|$ is the Frobenius norm.

In a second step, denoting $S_I^0 := \sigma_0^{-1}(\{1, 2\})$ and recalling from (18) that $\tilde{S}_I = \tilde{\sigma}^{-1}(\{1, 2\})$, we show that \tilde{S}_I is close to S_I^0 . To do so, one separately bounds from below some terms from the quantity $\|Z^0 - \tilde{Z}\|^2 = \sum_{i,j} (Z_{i,j}^0 - \tilde{Z}_{i,j})^2$, recalling that one extends the Z matrices by symmetry and sets the diagonal to 0. First, using the definitions, (21), $|\theta| \leq \kappa$, and $n \geq 2$,

$$\sum_{i,j \in S_I^0 \setminus \tilde{S}_I} (Z_{ij}^0 - \tilde{Z}_{ij})^2 \geq \frac{\kappa^2}{2} |S_I^0 \setminus \tilde{S}_I|^2,$$

if $|S_I^0 \setminus \tilde{S}_I| \geq 2$ (otherwise the inequality below holds trivially), as well as

$$\sum_{i \in S_I^0 \cap \tilde{S}_I, j \in S_I^0 \setminus \tilde{S}_I} (Z_{ij}^0 - \tilde{Z}_{ij})^2 \geq \frac{\kappa^2}{2} |S_I^0 \cap \tilde{S}_I| |S_I^0 \setminus \tilde{S}_I|,$$

and, with $a \wedge b = \min(a, b)$, if $|S_I^0 \cap \tilde{S}_I| \geq 2$,

$$\sum_{i,j \in S_I^0 \cap \tilde{S}_I} (Z_{ij}^0 - \tilde{Z}_{ij})^2 \geq \frac{1}{2} [(\theta_0 - \tilde{\theta})^2 \wedge (\theta_0 + \tilde{\theta})^2] |S_I^0 \cap \tilde{S}_I|^2.$$

The previous bound on $\|\tilde{Z} - Z^0\|^2$ implies that

$$\max \left(|S_I^0 \setminus \tilde{S}_I|^2, |S_I^0 \cap \tilde{S}_I| |S_I^0 \setminus \tilde{S}_I| \right) \leq C \kappa^{-2} n \log k \quad \text{with high probability.}$$

It now follows from (22) that for any $\delta > 0$, one has $C \kappa^{-2} n \log k \leq \delta n^2 / k^2$, provided d in (22) is small enough. So for small d , as $|S_I^0| = |S_I^0 \cap \tilde{S}_I| + |S_I^0 \setminus \tilde{S}_I|$, and as by assumption $\sigma_0 \in \Sigma_e$ so that $|S_I^0| \asymp n/k$, one deduces $|S_I^0 \cap \tilde{S}_I| \gtrsim n/k$. From the bound on $\|\tilde{Z} - Z^0\|^2$, it follows that

$$(\theta_0 - \tilde{\theta})^2 \wedge (\theta_0 + \tilde{\theta})^2 \lesssim \frac{n \log k}{n^2 k^{-2}} \lesssim \frac{k^2 \log k}{n} \quad \text{with high probability.}$$

By (22) this shows that $\tilde{\theta}$ is close to either θ_0 or $-\theta_0$ up to $k(\log n/n)^{1/2} =: \rho \leq \kappa/2$. Now for any i, j in $\tilde{S}_I \setminus S_I^0$, if $\tilde{Z}_{ij} = \frac{1}{2} + \tilde{\theta}$ and $\tilde{\theta}$ is close to θ_0 (the cases where $\tilde{Z}_{ij} = \frac{1}{2} - \tilde{\theta}$ or $\tilde{\theta}$ is close to $-\theta_0$ are treated similarly), setting $Z_{ij}^0 = c_{ij}$ and using (22), with $a_0 = 1/2$,

$$\begin{aligned} |Z_{ij}^0 - \tilde{Z}_{ij}| &= |c_{ij} - (a_0 + \tilde{\theta})| = |c_{ij} - (a_0 + \theta^0) - (\tilde{\theta} - \theta^0)| \\ &\geq |c_{ij} - (a_0 + \theta^0)| - |\tilde{\theta} - \theta_0| \geq \kappa - \rho \geq \kappa/2. \end{aligned}$$

Therefore,

$$\kappa^2 |\tilde{S}_I \setminus S_I^0|^2 \lesssim \sum_{i,j \in \tilde{S}_I \setminus S_I^0} (Z_{ij}^0 - \tilde{Z}_{ij})^2 \lesssim n \log k \quad \text{with high probability}$$

and

$$\kappa^2 |S_I^0 \cap \tilde{S}_I| |\tilde{S}_I \setminus S_I^0| \lesssim \sum_{\substack{i \in S_I^0 \cap \tilde{S}_I \\ j \in \tilde{S}_I \setminus S_I^0}} (Z_{ij}^0 - \tilde{Z}_{ij})^2 \lesssim n \log k \quad \text{with high probability.}$$

Combining the previous bounds on cardinalities and denoting $A \Delta B := (A \setminus B) \cup (B \setminus A)$ for two sets A and B , one obtains

$$|\tilde{S}_I \Delta S_I^0|^2 + |\tilde{S}_I \Delta S_I^0| |S_I^0 \cap \tilde{S}_I| \lesssim \kappa^{-2} n \log k \quad \text{with high probability,}$$

which in turn implies that

$$|\{\tilde{S}_I \times \tilde{S}_I\} \Delta \{S_I^0 \times S_I^0\}| \lesssim \kappa^{-2} n \log k \quad \text{with high probability.} \quad (38)$$

In a third and last step, we follow the proof of Theorem 3. Let $\hat{\sigma} = \hat{\sigma}_I$ be the mapping in (19). It is a map $\tilde{S}_I \rightarrow \{1, 2\}$. Let $\bar{\sigma}$ be the mapping $S_I^0 \rightarrow \{1, 2\}$ that coincides with $\hat{\sigma}$ on $\tilde{S}_I \cap S_I^0$ and with σ_0 on $S_I^0 \setminus \tilde{S}_I$. By definition we have, with $\Delta_n := \kappa^{-2} n \log k$,

$$\begin{aligned} Z_n(\hat{\sigma}, \tilde{S}_I, X) &= \sum_{i < j, i, j \in \tilde{S}_I} (-1)^{\mathbb{1}_{\hat{\sigma}(i) \neq \hat{\sigma}(j)}} \left(\frac{1}{2} - X_{ij} \right) \\ &= \sum_{i < j, i, j \in S_I^0} (-1)^{\mathbb{1}_{\bar{\sigma}(i) \neq \bar{\sigma}(j)}} \left(\frac{1}{2} - X_{ij} \right) + O(\Delta_n) \\ &= \delta(\bar{\sigma}, \sigma_0) \theta_0 + R_n(\bar{\sigma}) + O(\Delta_n), \end{aligned}$$

where for the second identity we have used that the X_{ij} s are bounded by 1 and (38), and $R_n(\sigma)$ is defined as in the proof of Theorem 3. Similarly, denoting $n_k = \binom{|S_I^0|}{2}$ and $\tilde{n}_k = \binom{|\tilde{S}_I|}{2}$, we have

$$Z_n(\sigma_0, \tilde{S}_I, X) = n_k \theta_0 + O(\Delta_n) + R_n(\sigma_0) = \tilde{n}_k \theta_0 + O(\Delta_n) + R_n(\sigma_0),$$

with high probability, since (38) implies $|\tilde{n}_k - n_k| = O(\Delta_n)$ using that $\|A\| - \|B\| \leq |A\Delta B|$ for two sets A, B . Also, since $\hat{\theta} = Z_n(\hat{\sigma}, \tilde{S}_I, X)/\tilde{n}_k$ and $|\hat{\theta}| \leq 1/2$, by the same argument we have

$$\hat{\theta} = \frac{Z_n(\hat{\sigma}, \tilde{S}_I, X)}{n_k} + \frac{O(\Delta_n)}{n_k}. \quad (39)$$

Let v_k and t_k be two sequences depending on n and k whose specific values are determined below (see the last paragraph of the proof). If $Z_n(\hat{\sigma}, \tilde{S}_I, X) \geq 0$, then $Z_n(\hat{\sigma}, \tilde{S}_I, X) \geq |Z_n(\sigma_0, \tilde{S}_I, X)| \geq Z_n(\sigma_0, \tilde{S}_I, X)$. Let Σ_0 be the set of all maps $S_0 \rightarrow \{1, 2\}$. In the following inequalities we repeatedly use the fact that the normalisation \tilde{n}_k in the definition of $\hat{\theta}$ can be replaced by n_k up to a factor $O(\Delta_n)/n_k$, see (39),

$$\begin{aligned} & P_{\theta_0} \left[v_k |\hat{\theta} - \theta_0| \geq t, Z_n(\hat{\sigma}, \tilde{S}_I, X) \geq 0 \right] \mathbb{1}_{\theta_0 > t_k} \\ & \leq P_{\theta_0} \left[\frac{v_k}{n_k} |Z_n(\hat{\sigma}, \tilde{S}_I, X) - Z_n(\sigma_0, \tilde{S}_I, X) + O(\Delta_n) - R_n(\sigma_0)| \geq t, \right. \\ & \quad \left. Z_n(\hat{\sigma}, \tilde{S}_I, X) \geq 0 \right] \mathbb{1}_{\theta_0 > t_k} \\ & \leq P_{\theta_0} \left[\frac{v_k}{n_k} \left[Z_n(\hat{\sigma}, \tilde{S}_I, X) - Z_n(\sigma_0, \tilde{S}_I, X) \right] + \frac{v_k}{n_k} [|R_n(\sigma_0)| + O(\Delta_n)] \geq t \right] \mathbb{1}_{\theta_0 > t_k} \\ & \leq P_{\theta_0} \left[\frac{v_k}{n_k} \theta_0 (\delta(\bar{\sigma}, \sigma_0) - n_k) + \frac{v_k}{n_k} [R_n(\hat{\sigma}) + |R_n(\sigma_0)| + O(\Delta_n)] \geq t \right] \mathbb{1}_{\theta_0 > t_k} \\ & \leq P_{\theta_0} \left[2 \frac{v_k}{n_k} \sup_{\sigma \in \Sigma_0} |R_n(\sigma)| + \frac{v_k}{n_k} O(\Delta_n) \geq t \right], \end{aligned}$$

where the last inequality uses $\delta(\bar{\sigma}, \sigma_0) - n_k \leq 0$, see Lemma 7 with n_k in place of b_n , and $\theta_0 \geq 0$.

Second, as $Z_n(\hat{\sigma}, \tilde{S}_I, X) < 0$ implies $Z_n(\hat{\sigma}, \tilde{S}_I, X) < -|Z_n(\sigma_0, \tilde{S}_I, X)|$ by definition of $\hat{\sigma}$,

$$\begin{aligned} & P_{\theta_0} \left[Z_n(\hat{\sigma}, \tilde{S}_I, X) < 0 \right] \mathbb{1}_{\theta_0 > t_k} \\ & \leq P_{\theta_0} \left[Z_n(\hat{\sigma}, \tilde{S}_I, X) < -|Z_n(\sigma_0, \tilde{S}_I, X)| \right] \mathbb{1}_{\theta_0 > t_k} \\ & \leq P_{\theta_0} \left[Z_n(\hat{\sigma}, \tilde{S}_I, X) < -|\theta_0| n_k + O(\Delta_n) + |R_n(\sigma_0)| \right] \mathbb{1}_{\theta_0 > t_k} \\ & \leq P_{\theta_0} \left[\theta_0 (\delta(\bar{\sigma}, \sigma_0) + n_k) - 2O(\Delta_n) < 2 \sup_{\sigma \in \Sigma} |R_n(\sigma)| \right] \mathbb{1}_{\theta_0 > t_k} \\ & \leq P_{\theta_0} \left[2 \sup_{\sigma \in \Sigma_0} |R_n(\sigma)| > \frac{t_k n_k}{8} - O(\Delta_n) \right], \end{aligned}$$

where for the last inequality we have used the first inequality of Lemma 7. Also,

$$P_{\theta_0} \left[v_k |\hat{\theta} - \theta_0| \geq t \right] \mathbb{1}_{|\theta_0| \leq t_k}$$

$$\begin{aligned} &\leq P_{\theta_0} \left[\frac{v_k}{n_k} |\theta_0| |\delta(\bar{\sigma}, \sigma_0) - n_k| + \frac{v_k}{n_k} [O(\Delta_n) + |R_n(\bar{\sigma})| + |R_n(\sigma_0)|] \geq t \right] \mathbb{1}_{|\theta_0| \leq t_k} \\ &\leq P_{\theta_0} \left[2 \frac{v_k}{n_k} \sup_{\sigma \in \Sigma_0} |R_n(\sigma)| + \frac{v_k}{n_k} O(\Delta_n) \geq t - v_k t_k \right]. \end{aligned}$$

By the same argument as in the proof of Theorem 3, the supremum

$\sup_{\sigma \in \Sigma_0} |R_n(\sigma)|$ is of the order $|\Sigma_0|^{3/2} \asymp (n/k)^{3/2}$, by definition of Σ_e . Recall that $n_k \asymp (n/k)^2$. Set $v_k = \sqrt{n/k}$ and $t_k = Dv_k^{-1}$, with D a large enough constant. Assumption (22) ensures that $\Delta_n = \kappa^{-2} n \log k = O((n/k)^{3/2})$. Hence by taking t a large enough constant, one obtains that the last three displays are bounded above by $e^{-Cn/k}$, which concludes the proof in the fixed design case, proceeding as in the proof of Theorem 3 to get the final bound in expectation.

The proof in the random design case is obtained by first deriving the results conditionally on φ and then integrating out φ , as we did in the proof of Theorem 3. The first part is almost identical to the fixed design case: one only needs to note that one can restrict to mappings φ that belong to, essentially, Σ_e . Denote by Σ'_e the subset of those $\sigma \in \Sigma_e$ satisfying $|\sigma^{-1}(j) - \frac{n}{k}| < \frac{n}{2k}$. Then

$$\begin{aligned} E_{\theta_0}[(\hat{\theta} - \theta_0)^2] &= E_{\theta_0}[(\hat{\theta} - \theta_0)^2 \mathbb{1}\{\varphi \in \Sigma'_e\}] + E_{\theta_0}[(\hat{\theta} - \theta_0)^2 \mathbb{1}\{\varphi \in \Sigma'_e\}] \\ &= E_{\theta_0} \left[E_{\theta_0}[(\hat{\theta} - \theta_0)^2 | \varphi] \mathbb{1}\{\varphi \in \Sigma'_e\} \right] + P_{\theta_0}[\varphi \notin \Sigma'_e]. \end{aligned}$$

For the first term, one can apply the arguments above in fixed design, while for the second an application of Bernstein’s inequality gives

$$P_{\theta_0}[\varphi \notin \Sigma'_e] \leq kP \left[\left| \text{Bin}(n, k) - \frac{n}{k} \right| > \frac{n}{2k} \right] \leq 2ke^{-\frac{n}{10k}}.$$

By (22), $k^3 \leq Cn$ holds for C large enough—note that κ must be smaller than 1, as the entries of M are in $[0, 1]$. One deduces $ke^{-\frac{n}{10k}} \leq \frac{k}{n} ne^{-dn^{2/3}} \leq C\frac{k}{n}$, for d small enough and C large enough, so the quadratic risk is at most Ck/n in this case as well. \square

Lemma 8. Let $Z^0_{i,j} := M^{\theta_0}_{\sigma_0(i)\sigma_0(j)}$ and $\tilde{Z}_{ij} := M^{\tilde{\theta}}_{\tilde{\sigma}(i)\tilde{\sigma}(j)}$, with $\tilde{\sigma}, \tilde{\theta}$ given by (17). Let $\|\cdot\|$ denote the matrix Frobenius norm. With probability at least $1 - e^{-cn \log k}$,

$$\|\tilde{Z} - Z^0\|^2 \lesssim Cn \log k.$$

Proof of Lemma 8. Let θ_1 denote the element of Θ_n closest to θ_0 , so that $|\theta_0 - \theta_1| \leq n^{-2}$. Let Z^1 be the matrix given by $Z^1_{ij} := M^{\theta_1}_{\sigma_0(i)\sigma_0(j)}$. By definition, $\|X - \tilde{Z}\|^2 \leq \|X - Z^1\|^2$ and hence $\|Z^1 - \tilde{Z}\|^2 + 2\langle X - Z^1, Z^1 - \tilde{Z} \rangle \leq 0$, so

$$\begin{aligned} \|Z^1 - \tilde{Z}\|^2 &\leq 2\|Z^1 - \tilde{Z}\| \sup_{\theta \in \Theta_n, \sigma \in \Sigma_e} \left| \langle X - Z^1, \frac{M^{\theta}_{\sigma(i)\sigma(j)} - Z^1}{\|M^{\theta}_{\sigma(i)\sigma(j)} - Z^1\|} \rangle \right| \\ &\leq 2\|Z^1 - \tilde{Z}\| \left[\sup_{\theta \in \Theta_n, \sigma \in \Sigma_e} |T_n(\sigma, \theta)| + \|Z^0 - Z^1\| \right], \end{aligned}$$

where we denote $T_n(\sigma, \theta) := \langle X - Z^0, (M_{\sigma^{(i)}\sigma^{(j)}}^\theta - Z^1) / \|M_{\sigma^{(i)}\sigma^{(j)}}^\theta - Z^1\| \rangle$. As elements of the matrix $X - Z^0$ are between -1 and 1 , we note that $T_n(\sigma, \theta)$ is of the form $\sum_l \mu_l \varepsilon_l$, where $\varepsilon_l \in [-1, 1]$ are independent, and $\sum_l \mu_l^2 = 1$. So using Hoeffding's inequality, for any $t > 0$,

$$P[|T_n(\sigma, \theta)| > t] \leq 2 \exp\{-t^2/2\}.$$

The cardinality of the set $\Theta_n \times \Sigma_e$ is bounded above by $(2n^2 + 1)k^n \lesssim k^{Cn}$. A union bound then shows that, with probability at least $1 - e^{-cn \log k}$,

$$\sup_{\theta \in \Theta, \sigma \in \Sigma_e} |T_n(\sigma, \theta)| \leq C\sqrt{n \log k}.$$

Inserting this back into the previous inequality on $\|Z^1 - \tilde{Z}\|^2$ leads to $\|Z^1 - \tilde{Z}\| \leq C\sqrt{n \log k} + \|Z^0 - Z^1\|$ with probability at least $1 - e^{-cn \log k}$. As $\|Z^0 - Z^1\|^2 \leq Cn^2/n^2 \leq C$, the triangle inequality leads to the result. \square

Appendix D: Remaining proofs: lower and upper bounds in the sparse case

We begin with a brief overview of the proof techniques: For the lower bounds in the sparse setting (Theorems 8 and 9), proofs are very similar to the dense case, and it suffices to track the dependence on the sparsity parameter α_n . To upper-bound the convergence rate of spectral estimates (Theorems 5 and 8), we use the fact that θ can be estimated from the largest absolute eigenvalue of the (translated) adjacency matrix. The latter can in turn be estimated empirically. For the proofs of the upper bounds for k classes, we show that with high probability it is possible to recover the true aggregated labels, where aggregation means that classes 1 and 2, corresponding to the ‘hard submodel’ are merged (this is the ‘ g map’ introduced above in the second paragraph of Appendix A). To do so, one adapts techniques introduced by Lei and Rinaldo [25], and Lei and Zhu [26] and show that their results still hold ‘under small perturbations’, as explained in more details below. Once the true aggregated labels of classes 1 and 2 are obtained, it suffices to apply the (already derived) result for the case $k = 2$.

D.1. Proofs for the two-class case

Proof of the lower bound in Theorem 8. One proceeds in the same way as in the proof of Theorem 1 with $a_{ij}(\varphi)$ replaced by $b_{ij}(\varphi) := Q_{\varphi^{(i)}\varphi^{(j)}}^\theta(\alpha_n)$, $i < j$. If $\lambda_1 = \text{Be}(b_{ij}(\varphi))$, $\lambda_2 = \text{Be}(b_{ij}(\psi))$, $\mu = \text{Be}(\alpha_n/2)$, we now have

$$\vartheta_{ij}(\varphi, \psi) := \int \left(\frac{d\lambda_1}{d\mu} - 1\right) \left(\frac{d\lambda_2}{d\mu} - 1\right) d\mu = \frac{(2b_{ij}(\varphi) - \alpha_n)(2b_{ij}(\psi) - \alpha_n)}{\alpha_n(2 - \alpha_n)}.$$

This leads to, with $\eta_i = \mathbb{1}_{\varphi^{(i)}=1} - \mathbb{1}_{\varphi^{(i)}=2}$ and $\eta'_i = \mathbb{1}_{\psi^{(i)}=1} - \mathbb{1}_{\psi^{(i)}=2}$,

$$\vartheta_{i,j}(\varphi, \psi) = \frac{4\alpha_n \theta^2 \eta_i \eta_j \eta'_i \eta'_j}{2 - \alpha_n}.$$

By the same argument as in the proof of Theorem 1, it is enough to solve

$$\frac{n\alpha_n\theta^2}{2 - \alpha_n} = C$$

for θ , where C is a universal positive small enough constant, under the constraint that $|\theta| \leq 1/2$. This leads to take θ^2 equal up to a constant to $(n\alpha_n)^{-1} \wedge 1$ and the proof is complete. \square

Proof of the upper bound in Theorem 8. We write the proof directly in the possibly sparse setting. Let us first consider the fixed design case, where φ is non-random. Let $\|\cdot\|_{S_p}$ denote the spectral norm of a matrix (for a symmetric matrix Δ , $\|\cdot\|_{S_p} = \max(|\lambda_1(\cdot)|, |\lambda_n(\cdot)|)$), so $|\lambda_1^2(\Delta)| = \|\Delta\|_{S_p}$. By [25, Theorem 5.2], we have that for any $r > 0$, there exists a $C = C(r, c_0) > 0$ such that

$$\|X - E[X]\|_{S_p} \leq C\sqrt{n\alpha_n},$$

with probability at least $1 - n^{-r}$. From this one deduces that $\|\Delta_n - E[\Delta_n]\|_{S_p} \leq C\sqrt{n\alpha_n}$. The eigenvalues of Δ and those of $\Delta_n - E[\Delta_n]$ and $E[\Delta_n]$ can be related to each other by a Weyl-type inequality as

$$|\lambda_i(\Delta_n) - \lambda_i(E\Delta_n)| \leq \|\Delta_n - E\Delta_n\|_{S_p},$$

for any $1 \leq i \leq n$, see e.g. [30, eq. (1.64)]. Suppose for now that $\theta \geq 0$. In this case $\lambda_1(E\Delta_n) = (n-1)\alpha_n\theta$ and $\lambda_n(E\Delta_n) = 0$, which by the previous inequality implies, with high probability,

$$\left| \frac{\lambda_1(\Delta_n)}{\alpha_n(n-1)} - \theta \right| \leq \frac{C}{\sqrt{n\alpha_n}}, \quad \left| \frac{\lambda_n(\Delta_n)}{\alpha_n(n-1)} \right| \leq \frac{C}{\sqrt{n\alpha_n}}.$$

Now if $\theta > 2C/\sqrt{n\alpha_n}$, using the first inequality we have $\lambda_1(\Delta_n)/\{\alpha_n(n-1)\} > C/\sqrt{n\alpha_n}$ and $\hat{\lambda}_1 = \lambda_1(\Delta_n)$ follows from the second inequality, which means $|\hat{\theta}_n - \theta| \leq C/\sqrt{n\alpha_n}$. If $\theta \leq 2C/\sqrt{n\alpha_n}$, the triangle inequality and the second inequality imply $|\lambda_n(\Delta_n) - \theta| \leq 3C/\sqrt{n\alpha_n}$, which combined with the first inequality gives $|\hat{\theta}_n - \theta| \leq 3C/\sqrt{n\alpha_n}$. So, for $\theta \geq 0$, in all cases $|\hat{\theta}_n - \theta| \leq 3C/\sqrt{n\alpha_n}$ with high probability. The case $\theta < 0$ is treated similarly. In the random design setting, one can argue conditionally on φ , and then note that both the obtained bounds and the in-probability statements do not depend on φ , which gives the result in this setting as well. \square

D.2. Proofs for the general case

Proof of the lower bound in Theorem 9. The proof is similar to that of Theorem 2, where one now uses the sparse lower bound for two classes of Theorem 8 instead of Theorem 1, and is thus omitted. \square

Proof of the upper bound in Theorem 9. We show that the proof approach used by [26] to establish their Theorem 2 can be adapted to our problem. More precisely, it is amenable to a perturbation of the true matrix M of connection probabilities: We show that, for a graph generated by model (12) with a sufficiently small value of θ , the V-Clust algorithm with $K = k - 1$ classes recovers the aggregated labelling defined by g above with high probability. We do the proof in the possibly sparse situation, thereby also proving the upper-bound in Theorem 9.

There are three steps. First, we show that the initial label recovery algorithm \mathcal{S} of [25] recovers most of the labels correctly, and control the error. Second, we show that the scheme of proof of [26] carries over to the problem of recovering the aggregated clustering up to label permutation. Finally, using assumption (A3) one can recover the aggregated class 1 with high probability, and restricting to nodes with label in that class we can apply the spectral method \mathcal{S}_2 of the case $k = 2$.

First step (Perturbed spectral method of Lei and Rinaldo).

$$M^\theta = M^0 + \theta R, \quad \text{with} \quad R = \begin{bmatrix} 1 & -1 & 0_{2,k-2} \\ -1 & 1 & 0_{2,k-2} \\ 0_{k-2,1} & 0_{k-2,1} & 0_{k-2,k-2} \end{bmatrix}$$

The matrix M^0 (i.e. M^θ with $\theta = 0$) can be transformed into the matrix N above by removing the first line and then the first column.

Let X be the matrix (X_{ij}) . Since the relevant design is fixed, there exists a binary $n \times k$ matrix T , with a single 1 in each row, for which we have $E[X] = TM^\theta(\alpha_n)T^t + D$, where $D = -\text{Diag}(TM^\theta(\alpha_n)T^t)$ is a diagonal matrix with entries bounded by α_n . Lei and Rinaldo call T a membership matrix. It can be rewritten in terms of N , using the relation between M^0 and N noted above: for a $n \times K$ membership matrix S and $E[X]$ the expected value of X ,

$$E[X] = \alpha_n SNS^t + \alpha_n \theta TRT^t + D.$$

Now we can follow Lei and Rinaldo's analysis of simple spectral clustering with $K = k - 1$ and the expectation matrix SNS^t ; one only needs to show that, despite the perturbation θTRT^t , the argument still holds. Intuitively, this is guaranteed by the assumption that θ is small enough, which ensures that the spectrum of the perturbation θTRT^t does not interact much with that of SNS^t . More precisely, we decompose X as

$$X = P + W, \quad \text{with} \quad P := \alpha_n SNS^t, \quad W := \alpha_n \theta TRT^t + D + X - E[X].$$

Following the proof of Theorem 3.1 of [25], the pair (S, N) parametrises a SBM with $K = k - 1$ classes and N is full rank. By their Lemma 2.1, the eigendecomposition of $P = S(\alpha_n N)S^t$ can be written $P = UDU^t$, where U is the matrix of the K leading eigenvectors of P , and one can write $U = S\xi$, for some matrix $\xi \in \mathbb{R}^{K \times K}$ with orthogonal rows (and $\|\xi_{k*} - \xi_{l*}\|^2 = n_k^{-1} + n_l^{-1}$). It also follows from the proof of that Lemma that if γ_n denotes the smallest absolute nonzero

eigenvalue of P , we have $\gamma_n = n_{\min}\alpha_n\lambda(K)$, with n_{\min} the cardinality of the smallest class, here of order n/k using that classes are balanced, and $\lambda(K)$ the smallest absolute eigenvalue of N .

By Lemma 5.1 of [25], one can control the distance between the leading eigenspaces of X and P (for the first K non-zero eigenvalues) in terms of the spectral norm of W . The assumptions of that Lemma are fulfilled with P here of rank $K = k - 1$ and of smallest nonzero singular value γ_n . If $\hat{U} \in \mathbb{R}^{n \times K}$ is the matrix of the K leading eigenvectors of X (and U the one for P , as above), there exists a $K \times K$ orthogonal matrix Q such that, with $\|\cdot\|_K$ and $\|\cdot\|_{S_p}$ the Frobenius and spectral norms respectively,

$$\|\hat{U} - UQ\|_F \leq \frac{2\sqrt{2K}}{\gamma_n} \|X - P\|_{S_p}.$$

By the triangle inequality, the spectral norm $\|X - P\|_{S_p}$ is in turn bounded by

$$\|X - P\|_{S_p} \leq \|X - E[X]\|_{S_p} + \alpha_n|\theta|\|TRT^t\|_{S_p}.$$

The matrix R can be written $R = uu^t$, where u^t is the row $(1 - 1 \ 0 \dots 0)$ of length k . In particular, R is of rank 1, and $\|TRT^t\|_{S_p} = \|Tu\|_2^2$ (a nonzero eigenvector is Tu). By construction, $\|Tu\|_2^2 = n_1 + n_2$, the number of elements of classes 1 and 2, so that $\|TRT^t\|_{S_p} \leq Cn/K$. Also, $\|D\|_{S_p} \leq \alpha_n$ since D is diagonal with terms bounded by α_n . By Theorem 5.2 of [25], the norm $\|X - E[X]\|$ is, with probability at least $1 - 1/n^2$, no larger than $C\sqrt{n\alpha_n}$, for a sufficiently large constant C . Gathering the last bounds and using $\alpha_n \lesssim \sqrt{n\alpha_n}$, one obtains $\|X - P\|_{S_p} \leq C(\sqrt{n\alpha_n} + |\theta|n\alpha_n/K)$.

On the other hand, following Lei and Rinaldo [25], one can perform an $(1 + \varepsilon)$ -approximate k -means clustering on the rows of \hat{U} : Application of their Lemma 5.3 to the matrices \hat{U} and UQ shows the approximate k -means solution is a pair $(\hat{S}, \hat{\xi})$, where \hat{S} a membership matrix, $\hat{\xi}$ a $K \times K$ matrix, and $\hat{S}\hat{\xi}$ is an approximate least-squares fit to \hat{U} . Moreover, the estimated membership \hat{S} coincides with S up to label permutation, except on sets S_1, \dots, S_K that are characterized as follows: Recall that ψ is the ‘true’ labelling obtained by merging the original classes 1 and 2 of nodes. Each set $S_j \subset \psi^{-1}(j)$ satisfies

$$\frac{1}{n_{\min}} \sum_{j=1}^K |S_j| \leq 4(4 + 2\varepsilon)\|\hat{U} - UQ\|_F^2,$$

whenever

$$(16 + 8\varepsilon)\|\hat{U} - UQ\|_F^2 < 1. \tag{40}$$

This implies, using the previous bounds and $\gamma_n = n_{\min}\alpha_n\lambda(K) \gtrsim (n/K)\alpha_n\lambda(K)$, that

$$\frac{1}{n} \sum_{j=1}^K |S_j| \leq (16 + 8\varepsilon)\pi_0 \frac{\|\hat{U} - UQ\|_F^2}{K}$$

$$\begin{aligned} &\leq (16 + 8\varepsilon)\pi_0 \frac{8}{\gamma_n^2} 2C^2(n\alpha_n + \theta^2 \alpha_n^2 \frac{n^2}{K^2}) \\ &\leq \frac{CK^2}{\lambda(K)^2} \left(\frac{1}{n\alpha_n} + \frac{\theta^2}{K^2} \right), \end{aligned}$$

provided, for some suitably small constant $c > 0$, with $\lambda = \lambda(K)$,

$$K^3 \frac{1}{n\alpha_n \lambda^2} + K \frac{\theta^2}{\lambda^2} < c.$$

The first summand coincides with the condition in [25]. The second term accounts for the perturbation induced by $\alpha_n \theta R$. Provided that

$$\frac{K^3}{\alpha_n \lambda^2} < cn/2 \quad \text{and} \quad \theta^2 < \frac{c\lambda^2}{2K}, \quad (41)$$

the simple spectral clustering algorithm has recovery error at most $n/f(n\alpha_n, K)$, with

$$f(n\alpha_n, K) = C \left(\frac{n\alpha_n \lambda^2}{K^2} \wedge \frac{\lambda^2}{\theta^2} \right). \quad (42)$$

The conditions on $n, \alpha_n, \lambda, K, \theta$ permit this quantity to be chosen suitably large. This means that, with high probability, Step 1 of the algorithm with $K = k - 1$ recovers a sufficiently large proportions of the labels of N , up to label permutation.

Second step (Lei and Zhu's exact label recovery method via sample splitting). We can now use the method introduced by Lei and Zhu [26]: using a first rough estimate of the labels, one can refine it to an exact label recovery with high probability, provided $f(n\alpha_n, K)$ is large enough in terms of a certain function of K . The recovered labels are those of the original classes 3, 4, ..., k , and of the aggregated class containing classes 1 and 2. To verify that the proof of Lei and Zhu generalizes to the perturbed case, it suffices to note that the distortion of $E[X_{ij}]$ for $i, j \in \psi^{-1}(\{1, 2\})$ from $1/2$ to $1/2 \pm \theta$ does not interfere with the bounds of the proof of Theorem 2 in [26]. The sample splitting algorithm of Lei and Zhu involves two subroutines called **CrossClust** and **Merge**. The mean of X_{ij} enters in the proof of that result via two applications of Bernstein's inequality, in the proofs of Lemma 6 (which implies the consistency of **CrossClust** via Lemma 3) and Lemma 7 (consistency of **Merge**) of Lei and Zhu [26].

We impose the assumptions of Lei and Zhu [26], Theorem 2, on K and $f(\alpha_n n/2, K)$: one needs $f(n\alpha_n/2, K)\gamma(K) \geq CK^{2.5}$ and the inequalities $\alpha_n \geq CK^3 \log n/(\gamma^2(K)n)$ and $Cn \geq K^3$. The last two conditions are implied by (B2) (respectively (A2) in the dense case). By (42), the first one is satisfied if

$$C \frac{n\alpha_n \lambda^2}{K^2} \gamma(K) \geq CK^{2.5}, \quad \text{and} \quad \frac{\lambda^2}{\theta^2} \gamma(K) \geq CK^{2.5}.$$

Again, the first inequality holds by (B2). The second inequality asks for $\theta^2 < C\lambda^2\gamma(K)/K^{2.5}$, which is guaranteed by (34).

The parameter θ affects the proof of Lemma 3 of Lei and Zhu [26] as follows. The proof relies on bounding three terms T_1, T_2 and T_3 via Lemma 6 in the Appendix of their paper. To be able to apply Bernstein’s inequality on T_1 , one needs

$$(n_1 + n_2)\alpha_n|\theta| < \frac{1}{4}\pi_0\gamma(K)\frac{n\alpha_n}{30K^{3/2}},$$

where $\pi_0 = n_{\min}/(n/K) \gtrsim 1$ in the notation of [26], Definition 1. To bound T_2 , one needs

$$(n_1 + n_2)\alpha_n|\theta| < \frac{1}{4}\frac{\pi_0^2\gamma(K)n\alpha_n f(n\alpha_n/2, K)}{30K^{5/2}},$$

while, similarly, to bound T_3 one needs

$$(n_1 + n_2)\alpha_n|\theta| < \frac{1}{4}\pi_0\gamma(K)\frac{n\alpha_n}{15K^{3/2}}.$$

On the other hand, consistency of **Merge** with $V = 2$ requires

$$(n_1 + n_2)\alpha_n|\theta| < \frac{1}{4}\frac{\alpha_n\pi_0^2\gamma(K)n^2}{160K^3}.$$

The condition required for T_1 implies the remaining ones: The one required for T_2 is weaker, since $f(n\alpha_n, K) > C'K$ for some constant $C' > 0$, by (41)–(42). The condition for **Merge** is also weaker up to constants, provided that $n \gtrsim K^{3/2}$, which is satisfied under our conditions. To obtain the above Bernstein’s inequalities, one thus needs

$$|\theta| < c_1\frac{\gamma(K)}{\sqrt{K}}.$$

It follows from the spectral decomposition of N that $\lambda(K) \leq \gamma(K)/\sqrt{2}$. By combining with (41), we see that it is enough that θ satisfies $|\theta| \leq C\lambda/\sqrt{K}$, which was already required above. Finally, to see that the last inequality is satisfied, one notes that it is implied by (34), using that $\gamma(K) \leq \sqrt{K}$. We have just proved that the exact recovery from [26] also holds here.

Third step (Finding true cluster 1 and conclusion). The second step provides a labelling \hat{g} that coincides, up to permutation, with the aggregated labelling g with high probability. The assumed separation from $1/2$ allows us to identify cluster 1: For $l = 1, \dots, k - 1$, compute

$$\hat{N}_l := \frac{1}{\binom{\hat{g}^{-1}(l)}{2}} \sum_{i < j, i, j \in \hat{g}^{-1}(l)} X_{ij}.$$

Since class sizes are of order n/k , an application of Bernstein’s inequality gives

$$|\hat{N}_l - N_{\sigma(l)\sigma(l)}| \leq |\theta|\mathbb{1}_{\sigma(l)=1} + O\left(\sqrt{\frac{k}{n}\log n}\right)$$

for some permutation σ , with high probability. By (A3), if $l \neq 1$, we have $|N_l - 1/2| \geq \kappa$. So w.h.p. there is exactly one diagonal element $\hat{N}_l = \hat{N}_{\hat{\ell}\hat{\ell}}$

within $\kappa/2$ of $1/2$, since the conditions of the theorem imply

$$|\theta| + C\sqrt{\frac{k}{n} \log n} \leq \frac{\kappa}{2}.$$

The index $\hat{\ell}$ then identifies the first cluster of N —which is the aggregate cluster corresponding to clusters 1 and 2 defined by M^θ —with high probability. We can now apply the spectral algorithm for $k = 2$ to the induced submatrix $(X_{ij})_{i,j \in \hat{g}^{-1}(\hat{\ell})}$. Using the upper-bound part of Theorem 8 with a number of nodes $|\hat{g}^{-1}(\hat{\ell})| \asymp n/k$ leads to $E_\theta[(\hat{\theta} - \theta)^2] \leq Ck/(n\alpha_n)$, by observing that the event with high probability arising from the previous arguments (that is, the concentration result by [25] and Bernstein’s inequalities) holds with probability at least $1 - 1/n$. \square

Acknowledgments

I. C. is very grateful for the hospitality of Columbia’s statistics department, where parts of this work were carried out.

References

- [1] E. S. Allman, C. Matias, and J. A. Rhodes. Parameter identifiability in a class of random graph mixture models. *J. Statist. Plann. Inference*, 141(5):1719–1736, 2011. [MR2763202](#)
- [2] C. Ambroise and C. Matias. New consistent and asymptotically normal parameter estimates for random-graph mixture models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 74(1):3–35, 2012. [MR2885838](#)
- [3] E. Arias-Castro and N. Verzelen. Community detection in dense random networks. *Ann. Statist.*, 42(3):940–969, 2014. [MR3210992](#)
- [4] E. Arias-Castro, E. J. Candès, and A. Durand. Detection of an anomalous cluster in a network. *Ann. Statist.*, 39(1):278–304, 2011. [MR2797847](#)
- [5] P. Bickel and A. Chen. A nonparametric view of network models and Newman-Girvan and other modularities. *PNAS*, 106(50):21068–21073, 2009.
- [6] P. Bickel and H. Chernoff. Asymptotic distribution of the likelihood ratio statistic in a prototypical non regular problem. In *Statistics and Probability: A Raghu Raj Bahadur Festschrift*, pages 83–96. Wiley, New York, 1993.
- [7] P. Bickel, D. Choi, X. Chang, and H. Zhang. Asymptotic normality of maximum likelihood and its variational approximation for stochastic blockmodels. *Ann. Statist.*, 41(4):1922–1943, 2013. [MR3127853](#)
- [8] P. J. Bickel, A. Chen, and E. Levina. The method of moments and degree distributions for network models. *Ann. Statist.*, 39(5):2280–2301, 2011. [MR2906868](#)
- [9] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztegombi. Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing. *Adv. Math.*, 219(6):1801–1851, 2008. [MR2455626](#)

- [10] C. Butucea and Y. I. Ingster. Detection of a sparse submatrix of a high-dimensional noisy matrix. *Bernoulli*, 19(5B):2652–2688, 2013. [MR3160567](#)
- [11] A. Celisse, J.-J. Daudin, and L. Pierre. Consistency of maximum-likelihood and variational estimators in the stochastic block model. *Electron. J. Stat.*, 6:1847–1899, 2012. [MR2988467](#)
- [12] A. Channarond, J.-J. Daudin, and S. Robin. Classification and estimation in the stochastic blockmodel based on the empirical degrees. *Electron. J. Stat.*, 6:2574–2601, 2012. [MR3020277](#)
- [13] V. H. de la Peña and E. Giné. *Decoupling*. Probability and its Applications. Springer, 1999. [MR1666908](#)
- [14] D. Easley and J. Kleinberg. *Networks, crowds and markets*. Cambridge University Press, 2010. [MR2677125](#)
- [15] S. Gadat, J. Kahn, C. Marteau, and C. Maugis-Rabusseau. Parameter recovery in two-component contamination mixtures: the L^2 strategy. *Ann. Inst. Henri Poincaré Probab. Stat.*, 56(2):1391–1418, 2020. ISSN 0246–0203. [MR4076788](#)
- [16] C. Gao, Y. Lu, and H. H. Zhou. Rate-optimal graphon estimation. *Ann. Statist.*, 43(6):2624–2652, 2015. [MR3405606](#)
- [17] A. Goldenberg, A. X. Zheng, S. E. Fienberg, and E. M. Airoldi. A survey of statistical network models. *Foundations and Trends in Machine Learning*, 2(2):129–233, 2009.
- [18] M. S. Handcock and K. J. Gile. Modeling social networks from sampled data. *Ann. App. Statist.*, 4:5–25, 2010. [MR2758082](#)
- [19] J. A. Hartigan. A failure of likelihood asymptotics for normal mixtures. In *Proceedings of the Berkeley conference in honor of Jerzy Neyman and Jack Kiefer, Vol. II (Berkeley, Calif., 1983)*, pages 807–810. Wadsworth, 1985. [MR0822066](#)
- [20] P. Heinrich and J. Kahn. Strong identifiability and optimal minimax rates for finite mixture estimation. *Ann. Statist.*, 46(6A):2844–2870, 2018. ISSN 0090-5364. [MR3851757](#)
- [21] P. D. Hoff, A. E. Raftery, and M. S. Handcock. Latent space approaches to social network analysis. *J. Amer. Statist. Assoc.*, 97(460):1090–1098, 2002. ISSN 0162-1459. [MR1951262](#)
- [22] S. Janson. *Graphons, cut norm and distance, couplings and rearrangements*. New York Journal of Mathematics. NYJM Monographs. 2013. [MR3043217](#)
- [23] O. Klopp, A. B. Tsybakov, and N. Verzelen. Oracle inequalities for network models and sparse graphon estimation. *Ann. Statist.*, 45(1):316–354, 2017. [MR3611494](#)
- [24] E. D. Kolaczyk. *Statistical Analysis of Network Data*. Springer, 2009. [MR2724362](#)
- [25] J. Lei and A. Rinaldo. Consistency of spectral clustering in stochastic block models. *Ann. Statist.*, 43(1):215–237, 2015. ISSN 0090-5364. [MR3285605](#)
- [26] J. Lei and L. Zhu. Generic sample splitting for refined community recovery in degree corrected stochastic block models. *Statist. Sinica*, 27(4):1639–1659, 2017. ISSN 1017-0405. [MR3701502](#)
- [27] E. Mossel, J. Neeman, and A. Sly. Consistency thresholds for the planted

- bisection model. *Electron. J. Probab.*, 21:1–24, 2016. [MR3485363](#)
- [28] M. Newman. *Networks. An Introduction*. Oxford University Press, 2009. [MR2676073](#)
- [29] P. Orbanz and D. M. Roy. Bayesian models of graphs, arrays and other exchangeable random structures. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 37:437–461, 2015.
- [30] T. Tao. *Topics in random matrix theory*, volume 132 of *Graduate Studies in Mathematics*. American Mathematical Society, 2012. [MR2906465](#)
- [31] S. L. van der Pas and A. van der Vaart. Bayesian community detection. *Bayesian Anal.*, 13(3):767–796, 2018. ISSN 1936-0975. [MR3807866](#)
- [32] A. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998. [MR1652247](#)
- [33] N. Verzelen and E. Arias-Castro. Community detection in sparse random networks. *Ann. Appl. Probab.*, 25(6):3465–3510, 2015. ISSN 1050-5164. [MR3404642](#)
- [34] P. Wolfe and S. Olhede. Nonparametric graphon estimation. 2013. preprint.
- [35] B. Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, New York, 1997. [MR1462963](#)
- [36] A. Y. Zhang and H. H. Zhou. Minimax rates of community detection in stochastic block models. *Ann. Statist.*, 44(5):2252–2280, 2016. [MR3546450](#)