

LIKELIHOOD-BASED MODEL SELECTION FOR STOCHASTIC BLOCK MODELS¹

BY Y. X. RACHEL WANG AND PETER J. BICKEL

Stanford University and University of California, Berkeley

The stochastic block model (SBM) provides a popular framework for modeling community structures in networks. However, more attention has been devoted to problems concerning estimating the latent node labels and the model parameters than the issue of choosing the number of blocks. We consider an approach based on the log likelihood ratio statistic and analyze its asymptotic properties under model misspecification. We show the limiting distribution of the statistic in the case of underfitting is normal and obtain its convergence rate in the case of overfitting. These conclusions remain valid when the average degree grows at a polylog rate. The results enable us to derive the correct order of the penalty term for model complexity and arrive at a likelihood-based model selection criterion that is asymptotically consistent. Our analysis can also be extended to a degree-corrected block model (DCSBM). In practice, the likelihood function can be estimated using more computationally efficient variational methods or consistent label estimation algorithms, allowing the criterion to be applied to large networks.

1. Introduction. Network modeling has attracted increasing research attention in the past few decades as the amount of data on complex systems accumulates at an unprecedented rate. Many complex systems in science and nature consist of interacting individual components which can be represented as nodes with connecting edges in a network. Network modeling has found numerous applications in studying friendship networks in sociology, Internet traffic in information technology, predator–prey interactions in ecology and protein–protein interactions and gene regulatory mechanisms in molecular biology.

One prominent feature of many of these networks is the presence of communities, where groups of nodes exhibit high internal connectivity. Communities provide a natural division of the network into subunits with certain traits. In social networks, they often arise based on people’s common interests and geographic locations. The World Wide Web forms communities or hubs based on the content of the web pages. In gene networks, communities correspond to genes with related functional groupings, many of which can act in the same biological pathway. Numerous heuristic algorithms have been proposed for detecting communities.

Received October 2015; revised February 2016.

¹Supported in part by NSF FRG Grant DMS-11-60319.

MSC2010 subject classifications. 62F05.

Key words and phrases. Stochastic block models, model misspecification, network communities, likelihood ratio statistic.

However, a generative model is needed to study the problem from a theoretical perspective.

The stochastic block model (SBM), proposed by Holland, Laskey and Leinhardt (1983) in social science, is one of the simplest random graph models incorporating community structures. It assigns each node a latent discrete block variable and the connectivity levels between nodes are determined by their block memberships. In practice, this model sometimes oversimplifies the structures of real networks and other variants have been proposed, including the degree-corrected SBM (DCSBM) [Karrer and Newman (2011)] relaxing the within-block degree homogeneity constraint and overlapping SBM [Airoldi et al. (2008)] allowing a node to be in multiple blocks. These models have been applied to model real networks in social science and biology [Bickel and Chen (2009), Daudin, Picard and Robin (2008), Airoldi et al. (2008), Karrer and Newman (2011)].

Much research effort has been devoted to the problems of estimating the latent block memberships and model parameters of a SBM, including modularity [Newman (2006a)] and likelihood maximization [Bickel and Chen (2009), Amini et al. (2013)], variational methods [Daudin, Picard and Robin (2008), Latouche, Birmelé and Ambroise (2012)], spectral clustering [Rohe, Chatterjee and Yu (2011), Fishkind et al. (2013)], belief propagation [Decelle et al. (2011)] to name but a few. The asymptotic properties of some of these methods have also been studied [Bickel and Chen (2009), Rohe, Chatterjee and Yu (2011), Celisse, Daudin and Pierre (2012), Bickel et al. (2013)]. However, these methods require knowing (or knowing at least a suitable range for) K , the number of blocks, a priori. Less attention has been paid to the problem of selecting K .

For general networks, this problem corresponds to the issue of determining the number of communities, which remains a challenging open problem. Recursive approaches have been adopted to extract [Zhao, Levina and Zhu (2011)] or divide [Bickel and Sarkar (2016)] one community sequentially, while using optimization strategies or hypothesis testing to decide whether the process should be stopped at one stage. A more general sequential test for comparing a fitted SBM against alternative models with finer structures is proposed in Lei (2014). Conceptually these approaches are more appealing for networks with a hierarchical structure. In other cases, it would be more desirable to be able to compare different community numbers directly. A few likelihood-based model selection criteria have been proposed [Daudin, Picard and Robin (2008), Latouche, Birmelé and Ambroise (2012), Saldana, Yu and Feng (2014)]. From an information-theoretic perspective, Peixoto (2013) proposed a criterion based on minimum length description. These approaches circumvent the difficulty of analyzing the likelihood directly by using variational approximations or assuming the node labels are fixed and using plug-in estimates obtained from other inference algorithms. Furthermore, the asymptotic studies of these criteria examining their large-sample performance remain incomplete. Empirically, a network cross-validation method has been investigated in Chen and Lei (2014). More recently, Le and Levina (2015) proposed a method

based on analyzing the spectral properties of graph operators, including the non-backtracking matrix and the Bethe Hessian matrix.

In this paper, we directly address the challenges involved in analyzing the asymptotic distribution of the maximum log likelihood function under model misspecification. We show the log likelihood ratio statistic is asymptotically normal in the case of underfitting. Although obtaining an explicit asymptotic distribution of the statistic in the case of overfitting is much more challenging, we have still derived its order of convergence and subsequently shown these two cases of misspecification can be separated with probability tending to one. We thus propose a model selection criterion taking the form of a penalized likelihood and show it is asymptotically consistent in the regime where network average degree grows at a polylog rate. In Section 2, we first derive our main results under the regular SBM assumptions and then outline how the arguments can be extended to a DCSBM. Computationally, the likelihood can be approximated with variational algorithms or consistent label estimation algorithms without affecting the asymptotic consistency of the criterion. We demonstrate the effectiveness of our method by comparing its performance with other model selection approaches on simulated and real networks in Sections 3 and 4.

2. Results.

2.1. Preliminaries. A SBM with K blocks on n nodes is defined as follows. A vector of latent labels $Z = (Z_1, \dots, Z_n)$ is generated with Z_i taking integer values from $[K] = \{1, \dots, K\}$ governed by a multinomial distribution with parameters $\pi = (\pi_1, \pi_2, \dots, \pi_K)$. Given $Z_i = a$, $Z_j = b$, an adjacency matrix A is generated with

$$A_{i,j} | (Z_i = a, Z_j = b) \sim \text{Bernoulli}(H_{a,b}), \quad i \neq j.$$

We consider a symmetric A with zero diagonal entries corresponding to an undirected graph, although our arguments generalize easily to directed graphs. H is a $K \times K$ symmetric matrix describing the connectivities within and between blocks. We denote the model parameters $\theta = (\pi, H)$ and let Θ_K be the parameter space of a K -block model,

$$\Theta_K = \left\{ \theta | \pi \in (0, 1)^K, \sum_{a=1}^K \pi_a = 1, H \in (0, 1)^{K \times K} \right\}.$$

Throughout the paper, $\theta^* = (\pi^*, H^*)$ will denote the true generative parameter giving rise to an observed A . We will further parametrize H^* by $H^* = \rho_n S^*$, where the degree density ρ_n may be $\Omega(1)$ or going to zero at a rate $n\rho_n / \log n \rightarrow \infty$. We assume $\theta^* \in \Theta_K$ and H^* has no identical columns, meaning the underlying model has K blocks and it is identifiable in the sense that it cannot be further collapsed to a smaller model. $z = (z_1, \dots, z_n) \in [K']^n$ represents another set of

labels under a K' -block model with K' not necessarily equaling K . $g(A; \theta)$ is the likelihood function describing the distribution of A with parameter $\theta \in \Theta_{K'}$ and can be written as the sum of the complete likelihood function $f(z, A; \theta)$ associated with the labels $z \in [K']^n$:

$$(2.1) \quad g(A; \theta) = \sum_{z \in [K']^n} f(z, A; \theta),$$

where $f(z, A; \theta)$ takes the form

$$f(z, A; \theta) = \left(\prod_{a=1}^{K'} \pi_a^{n_a(z)} \right) \left(\prod_{a=1}^{K'} \prod_{b=1}^{K'} H_{a,b}^{O_{a,b}(z)} (1 - H_{a,b})^{n_{a,b}(z) - O_{a,b}(z)} \right)^{1/2}$$

with count statistics

$$n_a(z) = \sum_{i=1}^n \mathbb{I}(z_i = a), \quad n_{a,b}(z) = \sum_{i=1}^n \sum_{j \neq i} \mathbb{I}(z_i = a, z_j = b),$$

$$O_{a,b}(z) = \sum_{i=1}^n \sum_{j \neq i} \mathbb{I}(z_i = a, z_j = b) A_{i,j}.$$

g and f are invariant with respect to a permutation on the block labels, $\tau : [K'] \rightarrow [K']$, and its corresponding permutations on the node labels z and the parameters θ . Furthermore, let $R(z)$ be the $K' \times K$ confusion matrix whose (k, a) th entry is

$$(2.2) \quad R_{k,a}(z, Z) = n^{-1} \sum_{i=1}^n \mathbb{I}(z_i = k, Z_i = a).$$

We take a likelihood-based approach toward model selection and first investigate whether different model choices can be separated using the log likelihood ratio

$$(2.3) \quad L_{K,K'} = \log \frac{\sup_{\theta \in \Theta_{K'}} g(A; \theta)}{\sup_{\theta \in \Theta_K} g(A; \theta)}.$$

Here, the comparison is made between the correct K -block model and fitting a misspecified K' -block model.

In the following sections, we analyze the asymptotic distribution of $L_{K,K'}$ for $K' \neq K$. The main focus of analysis lies in handling the sum in (2.1) which contains an exponential number of terms. It has been shown in [Bickel et al. \(2013\)](#) that when $\theta \in \Theta_K$, $\sup_{\theta \in \Theta_K} g(A; \theta)$ is essentially equivalent to maximizing the complete likelihood corresponding to the correct labels Z , $\sup_{\theta \in \Theta_K} f(Z, A; \theta)$. In the next section, we handle the case of underfitting and derive the asymptotic distribution of $L_{K,K'}$.




FIG. 1. A schematic representation of how H^* is merged to give $H = U_{K-1,K}(H^*, p)$. The green area contains unchanged parameters and the arrows indicate where mergings occur.

2.2. *Underfitting.* We start by considering $K' = K - 1$. Intuitively, a $(K - 1)$ -block model can be obtained by merging blocks in a K -block model. More specifically, given the correct labels $Z \in [K]^n$ and the corresponding block proportions $p = (p_1, \dots, p_K)$, $p_a = n_a(Z)/n$, we define a merging operation $U_{a,b}(H^*, p)$ which combines blocks a and b in H^* by taking weighted averages with proportions in p . For example, for $H = U_{K-1,K}(H^*, p)$,

$$\begin{aligned}
 H_{l,k} &= H_{l,k}^* \quad \text{for } 1 \leq l, k \leq K - 2; \\
 (2.4) \quad H_{l,K-1} &= \frac{p_l p_{K-1} H_{l,K-1}^* + p_l p_K H_{l,K}^*}{p_l p_{K-1} + p_l p_K} \quad \text{for } 1 \leq l \leq K - 2; \\
 H_{K-1,K-1} &= \frac{p_{K-1}^2 H_{K-1,K-1}^* + 2p_{K-1} p_K H_{K-1,K}^* + p_K^2 H_{K,K}^*}{p_{K-1}^2 + 2p_{K-1} p_K + p_K^2}.
 \end{aligned}$$

A schematic representation of H is given in Figure 1.

For consistency, when merging two blocks (a, b) with $b > a$, the new merged block will be relabeled a and all the blocks c with $c > b$ will be relabeled $c - 1$. Using this scheme, we also obtain the merged node labels $U_{a,b}(Z)$ and merged proportions $U_{a,b}(p)$ with $[U_{a,b}(p)]_a = p_a + p_b$.

Constraining the parameters to a smaller model results in a suboptimal likelihood and its distance from the likelihood associated with the correct model can be measured by the Kullback–Leibler divergence, denoted $D_{KL}(\cdot \| \cdot)$. Let

$$\begin{aligned}
 \gamma_1(x) &= x \log x + (1 - x) \log(1 - x), \\
 \gamma_2(x) &= x \log x - x
 \end{aligned}$$

and define

$$(2.5) \quad D_i(a, b) = \sum_{k,l=1}^{K-1} [U_{a,b}(\pi^*)]_k [U_{a,b}(\pi^*)]_l \gamma_i([U_{a,b}(H^*, \pi^*)]_{k,l}).$$

When $p = \pi^*$ and treating the labels Z as fixed parameters, denote $P_{A|Z, H^*}$ the probability distribution of A . Then the information loss incurred by the merging operation $U_{a,b}$ can be measured by

$$(2.6) \quad D_{KL}(P_{A|Z, H^*} \| P_{A|U_{a,b}(Z), U_{a,b}(H^*, \pi^*)}) = \begin{cases} \frac{n^2}{2} \left[\sum_{c,d=1}^K \pi_c^* \pi_d^* \gamma_1(H_{c,d}^*) - D_1(a, b) \right] + O(n), \\ \text{for } \rho_n = \Omega(1); \\ \frac{n^2 \rho_n}{2} \left[\sum_{c,d=1}^K \pi_c^* \pi_d^* \gamma_2(H_{c,d}^*) - D_2(a, b) \right] + O(n^2 \rho_n^2), \\ \text{for } \rho_n \rightarrow 0. \end{cases}$$

Thus, an optimal merging minimizing D_{KL} is essentially equivalent to maximizing $D_i(a, b)$.

We assume the following holds for θ^* .

ASSUMPTION 2.1. A unique maximum exists for $\max_{(a,b)} D_i(a, b)$.

This assumption is more of a notational convenience than necessity. From now on, without loss of generality assume the maximum is achieved at $a = K - 1$ and $b = K$, and denote $H' = U_{K-1, K}(H^*, \pi^*)$, $S' = H' / \rho_n$ and $Z' = U_{K-1, K}(Z)$. We also assume H' is identifiable in the sense that we have the following.

ASSUMPTION 2.2. S' has no identical columns.

Thus, the merged model cannot be collapsed further to a smaller model.

The next lemma argues $\sup_{\theta \in \Theta_{K-1}} g(A; \theta)$ is essentially dominated by the complete likelihood associated with the optimal merging.

LEMMA 2.3. Let $\mathcal{S}(z)$ be the set of labels which are equivalent up to a permutation τ , $\mathcal{S}(z) = \{\tau(z) | \tau : [K - 1] \rightarrow [K - 1]\}$. Then

$$(2.7) \quad \sum_{z \notin \mathcal{S}(Z')} \sup_{\theta \in \Theta_{K-1}} f(z, A; \theta) = \sup_{\theta \in \Theta_{K-1}} f(Z', A; \theta) o_P(1).$$

The proof is shown in the [Appendix](#).

This lemma provides a tractable bound on $\sup_{\theta \in \Theta_{K-1}} g(A; \theta)$, allowing the rest of the analysis to be carried out by usual Taylor expansion.

Define

$$\begin{aligned} \mu_1(\theta^*) &= \frac{1}{2} \left[D_1(K-1, K) - \sum_{c,d=1}^K \pi_c^* \pi_d^* \gamma_1(H_{c,d}^*) \right], \\ \mu_2(\theta^*) &= \mu_1(\theta^*) + \frac{1}{n} \{ (\pi_{K-1}^* + \pi_K^*) \log(\pi_{K-1}^* + \pi_K^*) \\ &\quad - \pi_{K-1}^* \log \pi_{K-1}^* - \pi_K^* \log \pi_K^* \}. \end{aligned}$$

Upon merging blocks $K-1$ and K , denote $u(a)$ as the new block label of block a , and define $d_i(a, b)$ such that

$$\begin{aligned} (2.8) \quad d_1(a, b) &= H_{a,b}^* \log \frac{H'_{u(a),u(b)}}{H_{a,b}^*} + (1 - H_{a,b}^*) \log \frac{1 - H'_{u(a),u(b)}}{1 - H_{a,b}^*}, \\ d_2(a, b) &= S_{a,b}^* \log \frac{S'_{u(a),u(b)}}{S_{a,b}^*} + (S'_{u(a),u(b)} - S_{a,b}^*). \end{aligned}$$

The following theorem gives the asymptotic distribution of $L_{K,K-1}$, the proof of which is shown in the [Appendix](#).

THEOREM 2.4. *Suppose the underlying model parameter generating A is $\theta^* = (\pi^*, H^*) \in \Theta_K$, then $L_{K,K-1}$ is asymptotically normal with*

$$(2.9) \quad \begin{aligned} n^{-3/2} L_{K,K-1} - \sqrt{n} \mu_1(\theta^*) &\xrightarrow{D} N(0, \sigma_1^2(\theta^*)) && \text{if } \rho_n = \Omega(1); \\ \rho_n^{-1} n^{-3/2} L_{K,K-1} - \rho_n^{-1} \sqrt{n} \mu_2(\theta^*) &\xrightarrow{D} N(0, \sigma_2^2(\theta^*)) && \text{if } \rho_n \rightarrow 0 \end{aligned}$$

when Assumptions 2.1 and 2.2 hold. Let $\Sigma(\pi^*)$ be the covariance matrix of a multinomial (π^*) distribution, that is, $\Sigma_{a,a}(\pi^*) = \pi_a^*(1 - \pi_a^*)$ and $\Sigma_{a,b}(\pi^*) = -\pi_a^* \pi_b^*$ for $a \neq b$. The variance $\sigma_i^2(\theta^*)$ is given by $J^T(\theta^*) \Sigma(\pi^*) J(\theta^*)$ for $i = 1, 2$, where $J(\theta^*) = (J_b(\theta^*))_{1 \leq b \leq K}$,

$$J_b(\theta^*) = \begin{cases} \pi_{K-1}^* d_i(b, K-1) + \pi_K^* d_i(b, K), & \text{for } 1 \leq b \leq K-2, \\ \sum_{a \neq b} \pi_a^* d_i(a, b) + \pi_b^* d_i(b, b), & \text{for } b = K-1, K. \end{cases}$$

REMARK 2.5. (i) A special case occurs for $K = 2$, $\pi_1^* = \pi_2^*$, $H_{1,1}^* = H_{2,2}^*$. In this case, $\sigma_i^2(\theta^*) = 0$ and $\rho_n^{-1} n^{-3/2} L_{K,K-1}$ converges to its asymptotic mean. In general, for homogeneous block models with $H_{a,a}^* = h$ and $H_{a,b}^* = g$ for $a \neq b$, $J(\theta^*)$ simplifies to $J_b(\theta^*) = 0$ for $b \leq K-2$, $J_{K-1}(\theta^*) = \pi_{K-1}^* d_i(K, K) + \pi_K^* d_i(K-1, K)$, and $J_K(\theta^*) = \pi_{K-1}^* d_i(K-1, K) + \pi_K^* d_i(K, K)$.

(ii) In general, underfitting a $K^- < K$ model will lead to the same type of limiting distribution under conditions similar to Assumptions 2.1 and 2.2, assuming the

uniqueness of the optimal merging scheme and identifiability after merging. That is,

$$(2.10) \quad \rho_n^{-1} n^{-3/2} L_{K,K^-} - \rho_n^{-1} \sqrt{n} \mu \xrightarrow{D} N(0, \sigma^2)$$

for some mean $\mu \sim C\rho_n$ and variance σ^2 . The proof will be similar but involve more tedious descriptions of how various merges can occur.

(iii) The asymptotic distributions derived under the null distribution of a K -block model suggest one might consider performing hypothesis testing directly to compare against an alternative simpler model. However, the asymptotic means depend on the true parameters, and its maximum likelihood estimate converges only at the rate \sqrt{n} [Bickel et al. (2013)].

(iv) Without Assumptions 2.1 and 2.2, it is easy to show

$$(2.11) \quad L_{K,K^-} \leq -\Omega_P(n^2 \rho_n),$$

where $\Omega(\cdot)$ denotes asymptotic lower bound, using the method in proving Theorem 2.7.

2.3. *Overfitting.* In the case of overfitting a K^+ -block model with $K^+ > K$, deriving the asymptotic distribution of L_{K,K^+} is much more challenging. Intuitively, embedding a K -block model in a larger model can be achieved by appropriately splitting the labels Z and there are an exponential number of possible splits. We first show a result analogous to Lemma 2.3. However, the number of summands involved in $\sup_{\theta \in \Theta_{K^+}} g(A; \theta)$ remains exponential this time.

Recall that for $z \in [K^+]^n$, $R(z, Z)$ is the $K^+ \times K$ confusion matrix. We first define a subset $\mathcal{V}_{K^+} \in [K^+]^n$ such that

$$\mathcal{V}_{K^+} = \{z \in [K^+]^n \mid \text{there is at most one nonzero entry in every row of } R(z, Z)\}.$$

\mathcal{V}_{K^+} is obtained by splitting of Z such that every block in z is always a subset of an existing block in Z . The next lemma shows it suffices to consider only the subclass of labels \mathcal{V}_{K^+} in the sum $g(A; \theta)$, the proof of which is given in the Appendix.

LEMMA 2.6. *Suppose $\theta^* \in \Theta_K$, then*

$$\sum_{z \in [K^+]^n} \sup_{\theta \in \Theta_{K^+}} f(z, A; \theta) = (1 + o_P(1)) \sum_{z \in \mathcal{V}_{K^+}} \sup_{\theta \in \Theta_{K^+}} f(z, A; \theta).$$

The lemma does not provide a direct simplification of the sum and suggests the reason why obtaining an asymptotic distribution for L_{K,K^+} is difficult. On the other hand, with appropriate concentration we can still derive the asymptotic order of the statistic.

THEOREM 2.7. *Suppose $\theta^* \in \Theta_K$, then overfitting by a K^+ -block model with $K^+ > K$ gives $L_{K,K^+} = O_P(n^{3/2} \rho_n^{1/2})$.*

The proof is provided in the Appendix.

2.4. *Model selection.* The results in the previous sections lead us to construct a penalized likelihood criterion for selecting the optimal block number. The criterion is consistent in the sense that asymptotically it chooses the correct K with probability one. Define

$$(2.12) \quad \beta(K') = \sup_{\theta \in \Theta_{K'}} \log g(A; \theta) - N_{K'} B_n,$$

where B_n gives the order of the penalty term, and $N_{K'}$ is a strictly increasing sequence indexed by K' describing the complexity of the model. The optimal K_0 is such that

$$(2.13) \quad K_0 = \arg \max_{K'} \beta(K').$$

COROLLARY 2.8. For $K' < K$, setting $B_n = o(n^2 \rho_n)$,

$$(2.14) \quad \mathbb{P}_{\theta^*}(\beta(K') < \beta(K)) \rightarrow 1.$$

For $K' > K$, setting B_n such that $B_n n^{-3/2} \rho_n^{-1/2} \rightarrow \infty$,

$$(2.15) \quad \mathbb{P}_{\theta^*}(\beta(K') < \beta(K)) \rightarrow 1.$$

PROOF. For $K' < K$, generalizing Theorem 2.4,

$$(2.16) \quad \begin{aligned} & \mathbb{P}_{\theta^*}(\beta(K') < \beta(K)) \\ &= \mathbb{P}_{\theta^*} \left(n^{-3/2} \rho_n^{-1} \log \frac{\sup_{\theta \in \Theta_{K'}} g(A; \theta)}{\sup_{\theta \in \Theta_K} g(A; \theta)} \right. \\ & \quad \left. - \sqrt{n} \rho_n^{-1} \mu < (N_{K'} - N_K) \frac{B_n}{n^{3/2} \rho_n} - \sqrt{n} \rho_n^{-1} \mu \right) \\ & \rightarrow 1, \end{aligned}$$

since $B_n = o(n^2 \rho_n)$ and $-\rho_n^{-1} \mu \geq C(\theta^*)$ for some positive constant depending on θ^* . In general, the same conclusion holds by Remark 2.5(iii).

For $K' > K$, using Theorem 2.7,

$$(2.17) \quad \begin{aligned} & \mathbb{P}_{\theta^*}(\beta(K') < \beta(K)) \\ &= \mathbb{P}_{\theta^*} \left(\frac{1}{n^{3/2} \rho_n^{1/2}} \log \frac{\sup_{\theta \in \Theta_{K'}} g(A; \theta)}{\sup_{\theta \in \Theta_K} g(A; \theta)} < (N_{K'} - N_K) \frac{B_n}{n^{3/2} \rho_n^{1/2}} \right) \\ & \rightarrow 1, \end{aligned}$$

when $B_n n^{-3/2} \rho_n^{-1/2} \rightarrow \infty$. \square

Since the ratio of the upper bound $n^2 \rho_n$ and the lower bound $n^{3/2} \rho_n^{1/2}$ tends to infinity, such a sequence B_n exists. Choosing B_n in this interval, we have $K_0 = K$

with probability tending to 1. However, we also note that for finite cases with moderate-sized n , $\sqrt{n}\rho_n^{-1}\mu$ in (2.16) is small, making it easy to over penalize with large B_n . At the same time, the lower bound in Theorem 2.7 is not tight and can be refined further.

We further assume the following holds for tractable approximation.

ASSUMPTION 2.9. The maximum is achieved in the set $\mathcal{N}_{K^+} = \{z \in \mathcal{V}_{K^+} | n_k(z) \geq \varepsilon n \text{ for all } k, \text{ for some } \varepsilon > 0, \}$.

Assumption 2.9 assumes the maximum can only be achieved on a loosely balanced block design. The assumption and Lemma 2.6 imply it remains to analyze the order of $\max_{z \in \mathcal{N}_{K^+}} \sup_{\theta \in \Theta_{K^+}} \log f(z, A; \theta)$. The following theorem shows the order of L_{K, K^+} can be refined to $O_P(n)$. The details can be found in the Appendix.

THEOREM 2.10. Under Assumption 2.9, L_{K, K^+} is of order $O_P(n)$ for $K^+ > K$.

It follows then choosing B_n growing slightly faster than n will ensure consistency in the sense described in Corollary 2.8. Thus, we choose a penalized likelihood of the following form:

$$(2.18) \quad \beta(K') = \sup_{\theta \in \Theta_{K'}} \log g(A; \theta) - \lambda \cdot \frac{K'(K' + 1)}{2} n \log n,$$

where the complexity term corresponds to the number of parameters in the edge probability matrix and the constant λ is a tuning parameter. Similar to many BIC-type criteria, choosing the tuning parameter is a challenging problem even though it does not affect the asymptotic properties. We discuss this problem in Section 3.

2.5. *Extension to a degree-corrected stochastic block model.* In practice, the SBM often oversimplifies the community structures by assuming all the nodes within a block have the same expected degree, thereby excluding networks with “hub” nodes and other possible degree variations within blocks. To address this limitation, Karrer and Newman (2011) proposed the degree-corrected stochastic block model (DCSBM) by setting

$$(2.19) \quad \mathbb{E}(A_{i,j} | Z, \omega) = \omega_i \omega_j H_{Z_i, Z_j}, \quad i \neq j,$$

where $\omega = (\omega_1, \dots, \omega_n)$ is the set of node degree parameters with some identifiability constraint.

As before, $Z \sim \text{Multinomial}(\pi)$. We also treat ω as a latent variable and assume $\omega_i | Z \sim n_k(Z) \cdot \text{Dirichlet}(\mathbf{1})$ for $Z_i = k$ so that ω satisfies the identifiability constraint $\sum_{i: Z_i=k} \omega_i = n_k$ for every k . Similar to Karrer and Newman (2011), we replace the Bernoulli likelihood by the Poisson likelihood and as-

sume $A_{i,i} \sim \text{Poisson}(\omega_i^2 H_{Z_i, Z_i} / 2)$ to simplify derivation. As noted in [Karrer and Newman \(2011\)](#) and [Zhao, Levina and Zhu \(2012\)](#), sparse networks are well approximated by the Poisson distribution and little difference was found in practice between the two choices. The assumption on the diagonal entries also does not change the asymptotic results. Therefore, given (Z, ω) , the log conditional likelihood of A is (up to a constant)

$$\begin{aligned}
 & \log f(A|Z, \omega; \theta) \\
 (2.20) \quad &= \frac{1}{2} \sum_{i,j} (A_{i,j} \log(\omega_i \omega_j H_{Z_i, Z_j}) - \omega_i \omega_j H_{Z_i, Z_j}) \\
 &= \sum_{i,j} A_{i,j} \log \omega_i + \frac{1}{2} \sum_{k,l} (O_{k,l}(Z) \log H_{k,l} - n_k(Z) n_l(Z) H_{k,l}).
 \end{aligned}$$

In this case, the likelihood function $f(Z, A; \theta)$ has a tractable form and one can show Lemma 2.3 holds provided $n^{1/2} \rho_n / \log n \rightarrow \infty$. The stricter condition on the degree density ensures even with node degree variations [in the worst case $\mathbb{E}(A_{i,j}|Z, \omega) \sim \rho_n/n$] there still exist enough edges for parameter estimation. Similar arguments apply to show that the criterion (2.18) is asymptotically consistent for this DCSBM. As the derivation is largely similar to the regular SBM case, we provide a proof sketch in the supplementary material [[Wang and Bickel \(2016\)](#)].

2.6. *Likelihood approximations.* In practice, direct computations of the likelihood function $g(A; \theta)$ and its supremum involve an exponential number of summands and quickly become intractable as n grows. In this section, we provide practical ways to approximate the likelihood and discuss conditions under which asymptotic consistency is preserved.

Variational likelihood for regular SBM. Using the EM algorithm to optimize over θ requires computing the conditional distribution of Z given A , which is not factorizable in this case. Variational methods tackle the true conditional distribution $f_{Z|A; \theta}$ with the mean field approximation, thus simplifying the local optimization at each iteration. Under the regular SBM, the variational log likelihood $J(q, \theta; A)$ for a K' -block model is defined as

$$(2.21) \quad J(q, \theta; A) = -D_{KL}(q \| f_{Z|A; \theta}) + \log g(A; \theta),$$

where $q \in \mathcal{D}_{K'}$ is any product distribution with $q(z) = \prod_{i=1}^n q_i(z_i)$, $1 \leq z_i \leq K'$. The variational estimates $\hat{\theta}_{K'}^{\text{VAR}}$ is given by

$$\hat{\theta}_{K'}^{\text{VAR}} = \arg \max_{\theta \in \Theta_{K'}} \max_{q \in \mathcal{D}_{K'}} J(q, \theta; A),$$

which can be optimized using the EM algorithm in [Daudin, Picard and Robin \(2008\)](#). Also we note that $J(q, \theta; A)$ simplifies to

$$J(q, \theta; A) = \sum_{i=1}^n \sum_{k=1}^{K'} q_i(k) (-\log q_i(k) + \log \pi(k)) + \sum_{i < j} \sum_{k,l=1}^{K'} q_i(k) q_j(l) (A_{ij} \log H_{k,l} + (1 - A_{ij}) \log(1 - H_{k,l}))$$

and hence can be easily evaluated.

We can replace the likelihood in (2.18) by the variational log likelihood J without changing its asymptotic performance. More precisely, the criterion with variational approximation

$$(2.22) \quad \beta^{\text{VAR}}(K') = \sup_{\theta \in \Theta_{K'}} \sup_{q \in \mathcal{D}_{K'}} J(q, \theta; A) - \lambda \cdot \frac{K'(K' + 1)}{2} n \log n$$

is still asymptotically consistent. Noting that:

- (i) $J(q, \theta; A) \leq \log g(A; \theta)$ for any $q \in \mathcal{D}_{K'}$;
- (ii) $\sup_{\theta \in \Theta_K} \sup_{q \in \mathcal{D}_K} J(q, \theta; A) - \sup_{\theta \in \Theta_K} \log g(A; \theta) = O_P(1)$ as shown in [Bickel et al. \(2013\)](#),

it can be easily verified that (2.16) and (2.17) still hold. Although (ii) applies to the global optimum of $J(q, \theta; A)$ which may not be achieved by the EM algorithm, we note that it can be relaxed to accommodate for the difference in practice. Provided the difference between the local optimum found by the algorithm and the global optimum is bounded by $o_P(n \log n)$, asymptotic consistency still holds.

Label estimation. The computation time of variational likelihood grows quickly with network size and becomes more complicated for degree corrected models. On the other hand, a number of algorithms are available for estimating the latent labels in a computationally efficient way under both regular SBM and DCSBM. Typically, these algorithms require specifying the block number, hence let $\hat{Z}(K')$ be the estimated labels corresponding to block number K' . Then $f_{\text{ML}}(\hat{Z}(K'), A) = \sup_{\theta \in \Theta_{K'}} f(\hat{Z}(K'), A; \theta)$ is the maximum complete likelihood by plugging in the estimated labels. We assume that the estimation algorithm is strongly consistent with the same convergence rate as in Theorem 1 of [Bickel and Chen \(2009\)](#).

ASSUMPTION 2.11. There exists a sequence $b_n \rightarrow \infty$ such that

$$(2.23) \quad \mathbb{P}(\hat{Z}(K) \neq \tau(Z)) = O(n^{-b_n}),$$

where τ is a permutation on $[K]$.

Such a convergence rate can be achieved by, for example, profile maximum likelihood [Bickel and Chen (2009)]. For computational efficiency, we will use the pseudo-likelihood algorithm developed by Amini et al. (2013), which is available for both regular SBM and DCSBM. Weak consistency for label estimation was shown in Amini et al. (2013). However, the plug-in estimates of the block model parameters are still consistent. Under Assumption 2.11, observing that:

1. $f_{\text{ML}}(\hat{Z}(K'), A) \leq \sup_{\theta \in \Theta_{K'}} \log g(A; \theta);$
2. $\log f_{\text{ML}}(\hat{Z}(K), A) = \sup_{\theta \in \Theta_K} \log g(A; \theta)(1 + o_P(1)),$

it is easy to see (2.16) and (2.17) still hold, and the criterion

$$(2.24) \quad \beta^{\text{ML}}(K') = \log f_{\text{ML}}(\hat{Z}(K'), A) - \lambda \cdot \frac{K'(K' + 1)}{2} n \log n$$

is asymptotically consistent.

In the next section, we use simulated data to demonstrate how the criterion performs in practice. We approximate the likelihood using the variational EM algorithm for regular SBM and the pseudo-likelihood algorithm for DCSBM.

3. Simulations.

3.1. *Goodness of fit.* We first examined how well the normal limiting distribution approximated the empirical distribution of the statistic in the case of underfitting. Figure 2 plots the distribution of $n^{-3/2}L_{K,K-1}$ for $n = 200$ and $n = 500$ obtained from 200 replications for the following two scenarios:

- (a) $K = 2, \pi^* = (0.4, 0.6), H^* = \begin{pmatrix} 0.15 & 0.05 \\ & 0.01 \end{pmatrix};$
- (b) $K = 3, \pi^* = (0.4, 0.3, 0.3), H^* = \begin{pmatrix} 0.2 & 0.1 & 0.1 \\ & 0.2 & 0.03 \\ & & 0.1 \end{pmatrix}.$

The log likelihoods are approximated by the variational EM algorithm initialized by regularized spectral clustering [Joseph and Yu (2013)]. The solid curves are normal densities with mean $\mu_2(\theta^*)$ and $\sigma(\theta^*)$ given in Theorem 2.4. Even though the $O(n)$ term in $\mu_2(\theta^*)$ diminishes asymptotically for ρ_n going to 0 slowly, we found it essential to correct for the bias in the finite sample regimes above. In both cases, the convergence to the Gaussian shape appears faster than the convergence to the mean, and a bias exists for $n = 200$. When the network size reaches 500, the empirical distributions are well approximated by their limiting distribution. We note that the bias should not have an adverse effect on model selection since it is in the direction away from zero, making it easier to separate the two models.

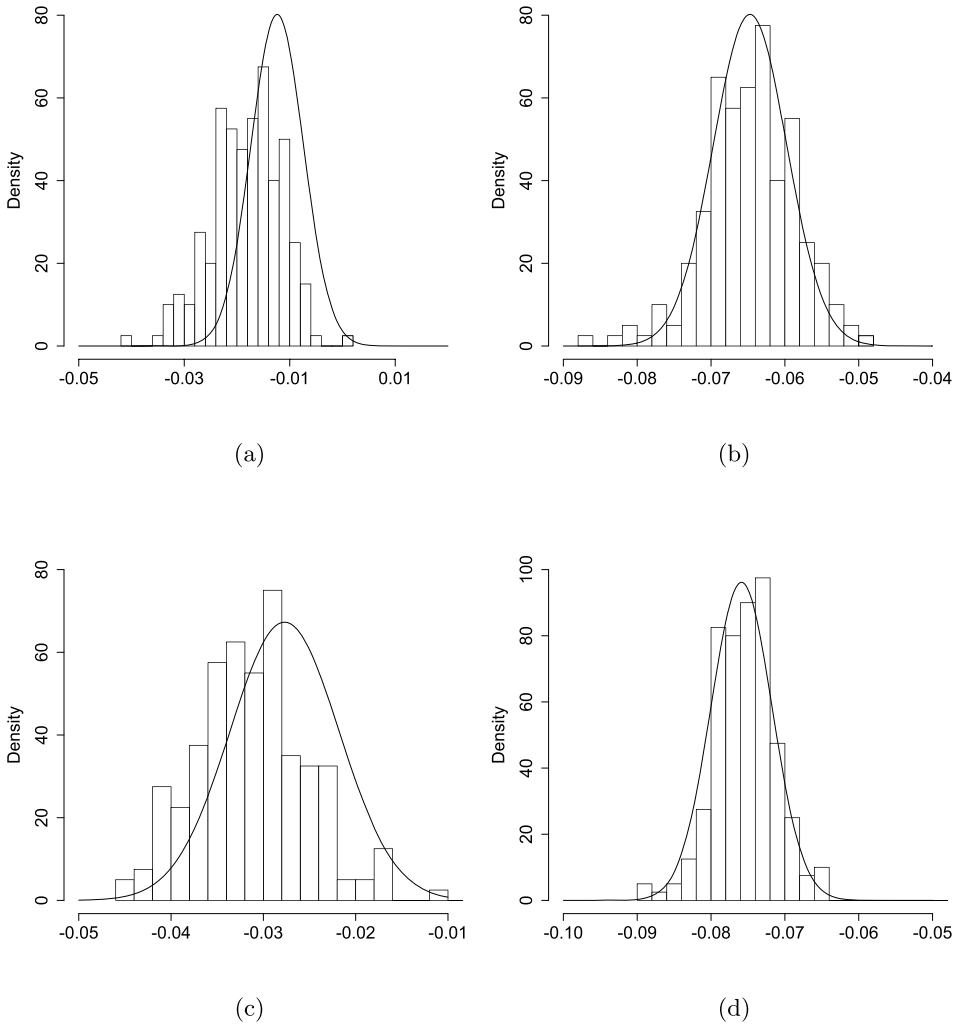


FIG. 2. Empirical distributions of $n^{-3/2}L_{K,K-1}$ for (a), (b) $K = 2$, π^* and H^* as described in scenario (a); (c), (d) $K = 3$, π^* and H^* as described in scenario (b). $n = 200$ in (a) and (c); $n = 500$ in (b) and (d). The solid curves are normal densities with mean $\mu_2(\theta^*)$ and $\sigma(\theta^*)$ as given in Theorem 2.4.

3.2. Selection of tuning parameter. Before we investigate the finite sample performance of our model selection criterion, we note that (2.18) involves a tuning parameter λ . We next propose a heuristic scheme for selecting λ . Similar to implementing BIC or AIC-type criteria, a maximum block number K_{\max} to be fitted needs to be chosen first. Then λ is chosen by the following algorithm:

1. For each choice of λ , compute $\beta(K')$ for $K' = 1, \dots, K_{\max}$.




FIG. 3. Comparison of the success rates of the penalized likelihood (plh) with variational Bayes (vb) and network cross validation (ncv). For every parameter setting, 50 networks were simulated from regular SBM with (a) $K = 2$, $\pi = (0.4, 0.6)$; (b) $K = 3$, $\pi = (0.3, 0.3, 0.4)$; (c) $K = 4$, $\pi_i = 0.25$ for all i . In all the cases, $H^* = \rho S^*$, where $\rho \in \{0.02, 0.04, \dots, 0.1\}$, the diagonal elements of S^* equal 2 and the off diagonal elements equal 1.

2. Normalize $(-\beta(1), \dots, -\beta(K_{\max}))$ into a probability vector $(w_1, \dots, w_{K_{\max}})$ so that they sum to 1.
3. Choose λ that maximizes the entropy $-\sum_k w_k \log(w_k)$.
4. If ties exist, choose the largest λ .

Here, $\beta(K')$ can be computed from either the variational likelihood (2.22) or the plug-in maximum likelihood (2.24). Heuristically, this algorithm chooses a λ that maximizes the “peakedness” of the profile of the penalized likelihoods $(\beta(1), \dots, \beta(K_{\max}))$, and hence the amount of signal contained in it. In the following sections, λ was chosen in the interval $[0, 0.3]$ with an increment of 1×10^{-3} ; $K_{\max} = 10$ for all simulated data.

3.3. Performance comparison with other methods. To see how our criterion (denoted plh for penalized likelihood) performs against other existing model selection methods, we compare its success rate with variational Bayes [Latouche, Birmelé and Ambroise (2012), denoted vb] and the 3-fold network cross validation method in Chen and Lei (2014) (denoted ncv). Since vb is only available for regular SBM, only ncv is included for DCSBM. plh is computed via either the variational EM for regular SBM or the pseudo-likelihood algorithm [Amini et al. (2013)] for DCSBM.

Figure 3 shows the average success rates of all three methods for data generated from regular SBM. 50 networks of size 500 were generated for each parameter set with $K = 2, 3, 4$, $H^* = \rho S^*$, and $\rho \in \{0.02, 0.04, \dots, 0.1\}$. The average degrees of these networks range from around 12 to 75. In general, the success rate of each method decreases as the networks become sparser and K increases, since the task of fitting also becomes harder. Overall, plh outperforms the other two methods.

TABLE 1

Comparison of the success rates of the penalized likelihood (plh) with network cross validation (ncv). For every parameter setting, 50 networks were simulated from the DCSBM with (a) $K = 2$, $\pi = (0.4, 0.6)$; (b) $K = 3$, $\pi = (0.3, 0.3, 0.4)$; (c) $K = 4$, $\pi_i = 0.25$ for all i . In all the cases, $H^* = \rho S^*$, where $\rho \in \{0.02, 0.04, \dots, 0.1\}$, the diagonal elements of S^* equal 2 and the off diagonal elements equal 1

ρ	$K = 2$			$K = 3$			$K = 4$		
	0.02	0.04	0.08	0.02	0.04	0.08	0.02	0.04	0.08
plh	0.88	0.96	1	0.08	0.66	1	0.12	0.64	0.98
ncv	0	0.26	1	0	0	0.54	0	0	0

Next, we simulated networks from DCSBM. To test if our method also works for more general DCSBM parameter settings and not limited by the specific Dirichlet prior assumption on ω , we generated degree parameters ω from a Unif(0.2, 1) distribution and further normalized them so that $\sum_{Z_i=k} \omega_i = n_k(Z)$. We set $H^* = \rho S^*$ as in the previous case with varying ρ . Binary adjacency matrices were generated even though our model in Section 2.5 is Poisson. As the simulation confirms, the approximation works well when networks are sparse. Table 1 shows the average success rates of plh (calculated using the pseudo-likelihood algorithm) and ncv for 50 networks of size 800 for each parameter set. Overall the problem is harder in this case than regular SBM as the inclusion of degree parameters induces more sparsity in some regions of the networks. plh shows a significant improvement over ncv in almost all cases.

4. Real world networks. In this section, we examine the performance of our method on real world networks. We set $K_{\max} = 30$ for the Facebook networks and $K_{\max} = 15$ for the others. We first implemented our method along with vb and ncv on nine Facebook ego networks, collected and labeled by Leskovec and McAuley (2012). An ego network is created by extracting subgraphs formed on the neighbors of a central (ego) node. Any isolated node was removed before analysis. Fitting regular SBM to these networks, variational EM was used to compute plh. The actual sizes of the networks and the number of communities selected by the three methods are shown in Table 2. The third row of the table shows the number of friend circles in every network with some individuals belonging to multiple circles, but not every individual possesses a circle label. These circle numbers give partial truth on how many communities there are in the networks. Overall, plh gives estimates closer to the circle numbers when the network is reasonably large and the number of circles is moderate. vb performs better on networks with a large number of circles but also overfits in a few cases. ncv tends to produce smaller community numbers.

We also implemented these methods on the political book network [Newman (2006b)], which consists of 105 books and their edges representing co-purchase

TABLE 2
Facebook ego networks and the number of communities selected by the three methods

# Nonisolated vertices	333	1034	224	150	168	61	786	534	52
Average degree	15	52	29	23	20	9	36	18	6
# Circles	24	9	14	7	13	13	17	32	17
plh	6	7	6	4	6	6	9	9	6
vb	11	24	16	9	11	6	25	23	6
ncv	3	6	4	2	4	2	2	2	3

information from Amazon. Again we treated this as a regular SBM and used variational EM to fit plh. Figure 4(a) shows the manual labeling of the books based on their political orientations being either “conservative,” “liberal” or “neutral.” As shown in (b), plh estimated $K = 6$ and essentially splits each of the communities in (a) into two, suggesting the presence of sub-communities with more uniform degree distributions. vb found 4 communities but merged two communities in (a) into one. ncv selected $K = 2$ and also merged two clusters in (a).

Finally, we used the political blog network [Adamic and Glance (2005)] as an example of DCSBM. The network consists of blogs on US politics and their web links as edges. Based on whether a blog is “liberal” or “conservative,” the network is divided into two communities. As is commonly done in the literature, we considered only the largest connected component containing a total of 1222 nodes. In this case, plh selected $K = 4$, splitting one of the communities into three as shown in Figure 5. ncv selected $K = 2$. On the other hand, we have also observed in simulations that ncv tends to have a bias toward lower K for DCSBM.

5. Discussion. In this paper, we have studied the problem of selecting the community number under both regular SBM and DCSBM, allowing the average degree to grow at a polylog rate and the true block number being fixed. We have shown the log likelihood ratio statistic has an asymptotic normal distribution when




FIG. 4. *Communities in 105 political books based on (a) manually curated ground truth; (b) penalized likelihood; (c) vb; (d) ncv.*




FIG. 5. *Political blog network (a) manually labeled ground truth; (b) penalized likelihood.*

a smaller model with fewer blocks is specified. In the case of misspecifying a larger model, we have obtained the convergence rate for the statistic. Combining these results, we arrive at a likelihood-based model selection criterion that is asymptotically consistent. For finite-sized networks, we have further refined the bound for the statistic in the overfitting case under reasonable assumptions to correct for the possibility of over-penalizing. Our method shows better performance than vb and ncv on simulated data and produces sensible results on a range of real world networks. We also note that vb is only available for regular SBM and ncv tends to be highly varying from run to run due to its use of random partitions.

There are a number of open problems for future work. (i) It would be interesting to investigate whether the results can be extended to other block model variants, such as overlapping SBM [Airoldi et al. (2008), Ball, Karrer and Newman (2011)]. (ii) We have performed our analysis with fixed block number as the number of nodes tends to infinity. However, in practice the number of communities is also likely to grow as a network expands [Choi, Wolfe and Airoldi (2012)], especially when we view block models as histogram approximations for more general models [Bickel and Chen (2009), Wolfe and Olhede (2013)]. Peixoto (2013) has provided some analysis on the maximum number of blocks detectable for a given SBM graph with fixed labels. In general, as more time-course network data become available in biology, social science and many other domains, incorporating dynamic features of community structures into network modeling will remain an interesting direction to explore.

APPENDIX: PROOFS OF LEMMAS AND THEOREMS

In this section, we prove all the lemmas and theorems in the main paper. Denote $\mu_n = n^2 \rho_n$, the total number of edges $L = \sum_{i=1}^n \sum_{j=i+1}^n A_{i,j}$, and $N(z) =$

$(n_{k,l}(z))_{1 \leq k,l \leq K'}$. For two sets of labels z and y , $|z - y| = \sum_{i=1}^n \mathbb{I}(z_i \neq y_i)$. $\|\cdot\|_\infty$ denotes the maximum norm of a matrix. We abbreviate $R(z, Z)S^*R(z, Z)^T$ as $RS^*R^T(z)$. C, C_1, \dots are constants which might be different at each occurrence. The following concentration inequalities bound the variations in A and will be used throughout the section.

LEMMA A.1. *Suppose $z \in [K']^n$ and define $X(z) = O(z)/\mu_n - RS^*R^T(z)$. For $\varepsilon \leq 3$,*

$$(A.1) \quad \mathbb{P}\left(\max_{z \in [K']^n} \|X(z)\|_\infty \geq \varepsilon\right) \leq 2(K')^{n+2} \exp(-C_1(S^*)\varepsilon^2\mu_n).$$

Let $y \in [K']^n$ be a fixed set of labels, then for $\varepsilon \leq 3m/n$,

$$(A.2) \quad \begin{aligned} & \mathbb{P}\left(\max_{z:|z-y|\leq m} \|X(z) - X(y)\|_\infty > \varepsilon\right) \\ & \leq 2 \binom{n}{m} (K')^{m+2} \exp\left(-C_2(S^*)\frac{n\varepsilon^2\mu_n}{m}\right). \end{aligned}$$

$C_1(S^*)$ and $C_2(S^*)$ are constants depending only on S^* .

PROOF. The proof follows from [Bickel and Chen \(2009\)](#) with minor modifications for general K' -block models and correcting for the zero diagonal in A . \square

Recall that

$$\begin{aligned} \gamma_1(x) &= x \log x + (1-x) \log(1-x), \\ \gamma_2(x) &= x \log x - x. \end{aligned}$$

Define $F_i(M, t)$, $i = 1, 2$, as

$$(A.3) \quad F_i(M, t) = \sum_{k,l=1}^{K'} t_{k,l} \gamma_i\left(\frac{M_{k,l}}{t_{k,l}}\right).$$

Then the log of the complete likelihood can be expressed as

$$(A.4) \quad \sup_{\theta \in \Theta_{K'}} \log f(z, A; \theta) = n \sum_{k=1}^{K'} \alpha(n_k(z)/n) + \frac{n^2}{2} F_1(O(z)/n^2, N(z)/n^2),$$

where $\alpha(x) = x \log(x)$. Noting the first term is of smaller order compared to the second term, and the conditional expectation of the argument in γ_1 given Z is $[RH^*R^T(z)]_{k,l}/[R\mathbf{1}\mathbf{1}^T R^T(z)]_{k,l}$ and $[RS^*R^T(z)]_{k,l}/[R\mathbf{1}\mathbf{1}^T R^T(z)]_{k,l}$ for γ_2 (up

to a diagonal difference) with fluctuation bounded by Lemma A.1, we will focus on analyzing the conditional expectation

$$G_1(R(z), S^*) = \sum_{k,l=1}^{K'} [R\mathbf{1}\mathbf{1}^T R^T(z)]_{k,l} \gamma_1 \left(\frac{[RH^* R^T(z)]_{k,l}}{[R\mathbf{1}\mathbf{1}^T R^T(z)]_{k,l}} \right) \tag{A.5}$$

for $\rho_n = \Omega(1)$,

$$G_2(R(z), H^*) = \sum_{k,l=1}^{K'} [R\mathbf{1}\mathbf{1}^T R^T(z)]_{k,l} \gamma_2 \left(\frac{[RS^* R^T(z)]_{k,l}}{[R\mathbf{1}\mathbf{1}^T R^T(z)]_{k,l}} \right) \tag{A.6}$$

for $\rho_n \rightarrow 0$.

The following lemma shows in the case of underfitting a $(K - 1)$ -block model, to maximize G_i over different configurations of $R(z, Z)$ with given Z , it suffices to consider the merging scheme described in Section 2.2 by combining two existing blocks in Z .

LEMMA A.2. *Given the true labels Z with block proportions $p = n(Z)/n$, maximizing the function $G_1(R(z), H^*)$ over R achieves its maximum in the label set*

$$\{z \in [K - 1]^n \mid \text{there exists } \tau \text{ such that } \tau(z) = U_{a,b}(Z), 1 \leq a < b \leq K\},$$

where $U_{a,b}$ merges Z_i with labels a and b .

Furthermore, suppose z_0 gives the unique maximum (up to permutation τ), for all R such that $R \geq 0, R^T \mathbf{1} = p$,

$$\frac{\partial G_1((1 - \varepsilon)R(z_0) + \varepsilon R, H^*)}{\partial \varepsilon} \Big|_{\varepsilon=0+} < -C < 0 \tag{A.7}$$

for $\rho_n = \Omega(1)$. The same conclusions hold for $G_2(R(z), S^*)$.

PROOF. Treating R as a $(K - 1) \times K$ -dimensional vector, it is easy to check $G_1(\cdot, H^*)$ is a convex function. Furthermore, since $R \geq 0, R^T \mathbf{1} = p$, the domain is part of a convex polyhedron $P_R = \{R \in \mathbb{R}^{K(K-1)} \mid R \geq 0, R^T \mathbf{1} = p\}$. Therefore, the maximum is attained at the vertices of P_R , that is R^{vert} such that for every a , exactly one $R_{k,a}^{\text{vert}}, (1 \leq k \leq K - 1)$ is nonzero. This is equivalent to assigning all $Z_i \in [K]$ with the same label into one group with a new label in $[K - 1]$. Let $u : [K] \rightarrow [K - 1]$ be the function specified by R^{vert} , then

$$G_1(R^{\text{vert}}, H^*) = \sum_{k,l} \sum_{\substack{a \in u^{-1}(k), \\ b \in u^{-1}(l)}} p_a p_b \gamma_1 \left(\frac{\sum_{a \in u^{-1}(k), b \in u^{-1}(l)} H_{a,b}^* p_a p_b}{\sum_{a \in u^{-1}(k), b \in u^{-1}(l)} p_a p_b} \right). \tag{A.8}$$

Note that there exists at least one $l \in [K - 1]$ such that $|\{u^{-1}(l)\}| > 1$, and $\{u^{-1}(k), k \in [K - 1]\}$ forms a partition on $[K]$. By strict convexity of γ_1 and

identifiability of H^* , to maximize G_1 it suffices to consider merging two of the labels in $[K]$ and mapping the other labels to the remaining labels in $[K - 1]$ in a one-to-one relationship.

The second part of the lemma holds since it is easy to see when the maximum is unique, the derivative of the G_1 at the optimal vertex is bounded away from 0 in all directions. The same arguments apply to G_2 . \square

Noting that when $p = \pi^*$, G_i evaluated at $R(U_{a,b}(Z))$ is equal to D_i defined in (2.5), it is easy to see Assumptions 2.1 and 2.2 guarantees the maximum is unique. We will now prove Lemma 2.3.

PROOF OF LEMMA 2.3. Taking the log of the complete likelihood,

$$\begin{aligned}
 \text{(A.9)} \quad & \sup_{\theta \in \Theta_{K-1}} \log f(z, A; \theta) \\
 &= n \sum_{k=1}^{K-1} \alpha(n_k(z)/n) + \frac{n^2}{2} F_1(O(z)/n^2, N(z)/n^2).
 \end{aligned}$$

By concentration of p_k , it suffices to consider $\{\|p - \pi^*\|_\infty < \eta\}$, where η is small enough that Z' remains the unique maximizer of $G_1(R(z), H^*)$ and $G_2(R(z), S^*)$, and distribution conditional on Z .

Using techniques similar to Bickel et al. (2013), we prove this by considering z far away from Z' and close to Z' (up to permutation τ). Let δ_n be a sequence converging to 0 slowly. Define

$$I_{\delta_n} = \{z \in [K - 1]^n : G_1(R(z), H^*) - G_1(R(Z'), H^*) < -\delta_n\}.$$

First, by (A.1) in Lemma A.1, for $\varepsilon_n \rightarrow 0$ slowly,

$$\begin{aligned}
 \text{(A.10)} \quad & |F_1(O(z)/n^2, N(z)/n^2) - G_1(R(z), H^*)| \\
 & \leq C \cdot \sum_{k,l} |O_{k,l}(z)/n^2 - (RH^*R^T(z))_{k,l}| + O(n^{-1}) \\
 & = o_P(\varepsilon_n)
 \end{aligned}$$

since γ_1 is Lipschitz on any interval bounded away from 0 and 1 and $\min H^* = \Omega(1)$. For $z \in I_{\delta_n}$ and $\rho_n = \Omega(1)$,

$$\begin{aligned}
 \text{(A.11)} \quad & \sum_{z \in I_{\delta_n}} \sup_{\theta \in \Theta_{K-1}} e^{\log f(z, A; \theta)} \\
 & \leq \sup_{\theta \in \Theta_{K-1}} f(Z', A; \theta) (K - 1)^n e^{O(n) + o_P(n^2 \varepsilon_n) - n^2 \delta_n} \\
 & = \sup_{\theta \in \Theta_{K-1}} f(Z', A; \theta) o_P(1)
 \end{aligned}$$

choosing $\delta_n \rightarrow 0$ slowly enough such that $\delta_n/\varepsilon_n \rightarrow \infty$. Similarly, for $\rho_n \rightarrow 0$, define

$$J_{\delta_n} = \{z \in [K - 1]^n : G_2(R(z), S^*) - G_2(R(Z'), S^*) < -\delta_n\}.$$

Note that in this case, for $\varepsilon_n \rightarrow 0$ slowly,

$$\begin{aligned} & F_1(O(z)/n^2, N(z)/n^2) \\ (A.12) \quad &= 2 \log \rho_n L/n^2 + \rho_n F_2(O(z)/\mu_n, N(z)/n^2) + O_P(\rho_n^2) \\ &= 2 \log \rho_n L/n^2 + \rho_n G_2(R(z), S^*) + o_P(\rho_n \varepsilon_n) + O_P(\rho_n^2), \end{aligned}$$

by (A.1) and the fact that γ_2 is Lipschitz on any interval bounded away from 0 and 1 and $\min S_{a,b}^* > 0$. Then for $z \in J_{\delta_n}$,

$$\begin{aligned} & \sum_{z \in J_{\delta_n}} \sup_{\theta \in \Theta_{K-1}} e^{\log f(z, A; \theta)} \\ (A.13) \quad & \leq \sup_{\theta \in \Theta_{K-1}} f(Z', A; \theta) (K - 1)^n e^{O(n) + O_P(\mu_n \rho_n) + o_P(\mu_n \varepsilon_n) - \mu_n \delta_n} \\ & = \sup_{\theta \in \Theta_{K-1}} f(Z', A; \theta) o_P(1) \end{aligned}$$

choosing $\varepsilon_n \rightarrow 0, \delta_n \rightarrow 0$ slowly enough.

For $z \notin J_{\delta_n}$, $|G_2(R(z), S^*) - G_2(R(Z'), S^*)| \rightarrow 0$. Let $\bar{z} = \min_{\tau} |\tau(z) - Z'|$. Since the maximum is unique up to τ , $\|R(\bar{z}) - R(Z')\|_{\infty} \rightarrow 0$ and $|\sum_k \alpha(n_k(\bar{z})/n) - \sum_k \alpha(n_k(Z')/n)| \rightarrow 0$.

By (A.2),

$$\begin{aligned} & \mathbb{P}\left(\max_{z \notin \mathcal{S}(Z')} \|X(\bar{z}) - X(Z')\|_{\infty} > \varepsilon |\bar{z} - Z'|/n\right) \\ (A.14) \quad & \leq \sum_{m=1}^n \mathbb{P}\left(\max_{z: z=\bar{z}, |\bar{z}-Z'|=m} \|X(z) - X(Z')\|_{\infty} > \varepsilon \frac{m}{n}\right) \\ & \leq \sum_{m=1}^n 2(K - 1)^{K-1} n^m (K - 1)^{m+2} \exp\left(-C \frac{m \mu_n}{n}\right) \rightarrow 0. \end{aligned}$$

It follows for $|\bar{z} - Z'| = m, z \notin J_{\delta_n}$,

$$\begin{aligned} & \left\| \frac{O(\bar{z})}{\mu_n} - \frac{O(Z')}{\mu_n} \right\|_{\infty} \\ (A.15) \quad &= o_P(1) \frac{|\bar{z} - Z'|}{n} + \|RS^*R^T(\bar{z}) - RS^*R^T(Z')\|_{\infty} \\ & \geq \frac{m}{n} (C + o_P(1)). \end{aligned}$$

Observe $\|O(Z')/\mu_n - RS^*R(Z')\|_\infty = o_P(1)$ by Lemma A.1, $N(Z')/n^2 = R\mathbf{1}^T R^T(Z') + o(1)$ on $\{\|p - \pi^*\|_\infty < \eta\}$, and $F_2(\cdot, \cdot)$ has continuous derivative in the neighborhood of $(O(Z')/\mu_n, N(Z')/n^2)$. Using (A.7) in Lemma A.2,

$$\frac{\partial F_2((1 - \varepsilon)\frac{O(Z')}{\mu_n} + \varepsilon M, (1 - \varepsilon)\frac{N(Z')}{n^2} + \varepsilon t)}{\partial \varepsilon} \Big|_{\varepsilon=0^+} < -\Omega_P(1) < 0$$

for (M, t) in the neighborhood of $(O(Z')/\mu_n, N(Z')/n^2)$. Hence,

$$\begin{aligned} (A.16) \quad & F_2(O(\bar{z})/\mu_n, N(\bar{z})/n^2) - F_2(O(Z')/\mu_n, N(Z')/n^2) \\ & \leq -\Omega_P(1)\frac{m}{n}. \end{aligned}$$

We have

$$\begin{aligned} (A.17) \quad & \sup_{\theta \in \Theta_{K-1}} \log f(z, A; \theta) - \sup_{\theta \in \Theta_{K-1}} \log f(Z', A; \theta) \\ & \leq n \left| \sum_{k=1}^{K-1} \alpha(n_k(\bar{z})/n) - \alpha(n_k(Z')/n) \right| \\ & \quad + n^2(F_1(O(\bar{z})/\mu_n, N(\bar{z})/n^2) \\ & \quad - F_1(O(Z')/\mu_n, N(Z')/n^2)) \\ & \leq (O(n) + o_P(\mu_n) - \Omega_P(\mu_n))\frac{m}{n} \\ & = -\Omega_P(\mu_n)\frac{m}{n} \end{aligned}$$

using (A.12) and (A.16). We can conclude

$$\begin{aligned} (A.18) \quad & \sum_{z \notin J_{\delta_n}, z \neq \tau(Z')} \sup_{\theta \in \Theta_{K-1}} e^{\log f(z, A; \theta)} \\ & \leq \sup_{\theta \in \Theta_{K-1}} f(Z', A; \theta) \sum_{m=1}^n (K-1)^{K-1} n^m (K-1)^m e^{-\Omega_P(\mu_n)m/n} \\ & = \sup_{\theta \in \Theta_{K-1}} f(Z', A; \theta) o_P(1). \end{aligned}$$

The bounds (A.13) and (A.18) yield (2.7). The case for $\rho_n = \Omega(1)$ can be shown in a similar way. \square

Now Theorem 2.4 follows by Taylor’s expansion.

PROOF OF THEOREM 2.4. First, note that

$$\begin{aligned}
 L_{K,K-1} &= \log \frac{\sup_{\theta \in \Theta_{K-1}} g(A; \theta)}{g(A; \theta^*)} - \log \frac{\sup_{\theta \in \Theta_K} g(A; \theta)}{g(A; \theta^*)} \\
 \text{(A.19)} \quad &= \sup_{\theta \in \Theta_{K-1}} \log \left[\frac{g(A; \theta)}{f(Z, A; \theta^*)} \cdot \frac{f(Z, A; \theta^*)}{g(A; \theta^*)} \right] + O_P(1) \\
 &= \sup_{\theta \in \Theta_{K-1}} \log \frac{g(A; \theta)}{f(Z, A; \theta^*)} + O_P(1)
 \end{aligned}$$

by a consequence of Theorem 1 and Lemma 3 in [Bickel et al. \(2013\)](#). Noting that $\sup_{\theta \in \Theta_{K-1}} f(Z', A; \theta)$ is uniquely maximized at (omitting the argument Z)

$$\begin{aligned}
 \hat{\pi}_a &= \frac{n_a}{n} = \pi_a^* + O_P(n^{-1/2}) \\
 &\qquad\qquad\qquad \text{for } 1 \leq a \leq K - 2, \hat{\pi}_{K-1} = \frac{n_{K-1} + n_K}{n}, \\
 \text{(A.20)} \quad \hat{H}_{a,b} &= \frac{O_{a,b}}{n_{a,b}} = H_{a,b}^* + O_P(\sqrt{\rho_n} n^{-1}) \quad \text{for } 1 \leq a \leq b \leq K - 2, \\
 \hat{H}_{a,K-1} &= \frac{O_{a,K-1} + O_{a,K}}{n_{a,K-1} + n_{a,K}} = H'_{a,K-1} + O_P(\rho_n n^{-1/2}) \\
 &\qquad\qquad\qquad \text{for } 1 \leq a \leq K - 2,
 \end{aligned}$$

$$\hat{H}_{K-1,K-1} = \frac{\sum_{a=K-1}^K \sum_{b=a}^K O_{a,b}}{\sum_{a=K-1}^K \sum_{b=a}^K n_{a,b}} = H'_{K-1,K-1} + O_P(\rho_n n^{-1/2}),$$

and Assumption 2.2 the merged S' is identifiable, we have

$$\frac{\sup_{\theta \in \Theta_{K-1}} \sum_{z \in \mathcal{S}(Z')} f(z, A; \theta)}{\sup_{\theta \in \Theta_{K-1}} f(Z', A; \theta)} = 1 + o_P(1).$$

Combined with Lemma 2.3

$$\begin{aligned}
 \text{(A.21)} \quad &\sup_{\theta \in \Theta_{K-1}} \log \frac{g(A; \theta)}{f(Z, A; \theta^*)} \\
 &= \sup_{\theta \in \Theta_{K-1}} \log \frac{f(Z', A; \theta)}{f(Z, A; \theta^*)} + o_P(1).
 \end{aligned}$$

We will check the expansion for the case $\rho_n \rightarrow 0$; the case $\rho_n = \Omega(1)$ can be shown in the same way.

$$\begin{aligned}
 &n^{-3/2} \rho_n^{-1} \sup_{\theta \in \Omega_{K-1}} \log \frac{g(A; \theta)}{f(Z, A; \theta^*)} \\
 &= n^{-3/2} \rho_n^{-1} \sup_{\theta \in \Omega_{K-1}} \log \frac{f(Z', A; \theta)}{f(Z, A; \theta^*)} + o_P(1)
 \end{aligned}$$

$$\begin{aligned}
 &= n^{-3/2} \rho_n^{-1} \left\{ n \sum_{a=1}^{K-1} \alpha(\hat{\pi}_a) \right. \\
 &\quad \left. + \frac{1}{2} \sum_{a=1}^{K-1} \sum_{b=1}^{K-1} \left(O_{a,b} \log \frac{\hat{H}_{a,b}}{1 - \hat{H}_{a,b}} + n_{a,b} \log(1 - \hat{H}_{a,b}) \right) \right. \\
 \text{(A.22)} \quad &\quad \left. - \sum_{a=1}^K n_a \log \pi_a^* \right. \\
 &\quad \left. - \frac{1}{2} \sum_{a=1}^K \sum_{b=1}^K \left(O_{a,b} \log \frac{H_{a,b}^*}{1 - H_{a,b}^*} + n_{a,b} \log(1 - H_{a,b}^*) \right) \right\} + o_P(1) \\
 &= n^{-1/2} \rho_n^{-1} [\alpha(\pi_{K-1}^* + \pi_K^*) - \alpha(\pi_{K-1}^*) - \alpha(\pi_K^*)] \\
 &\quad + n^{-3/2} \rho_n^{-1} \frac{1}{2} \sum_{(a,b) \in \mathcal{I}} \left(O_{a,b} \log \frac{H'_{u(a),u(b)}(1 - H_{a,b}^*)}{(1 - H'_{u(a),u(b)})H_{a,b}^*} \right. \\
 &\quad \left. + n_{a,b} \log \frac{1 - H'_{u(a),u(b)}}{1 - H_{a,b}^*} \right) + o_P(1),
 \end{aligned}$$

where \mathcal{I} be the set of indices affected by the merge,

$$\mathcal{I} = \{(a, b) \in [K]^2 \mid K - 1 \leq a \leq K \text{ or } K - 1 \leq b \leq K\}.$$

It is easy to see the expectation of this term is $\rho_n^{-1} \sqrt{n} \mu_2$, we have

$$\begin{aligned}
 &n^{-3/2} \rho_n^{-1} \sup_{\theta \in \Omega_{K-1}} \log \frac{g(A; \theta)}{f(Z, A; \theta^*)} - \sqrt{n} \rho_n^{-1} \mu_2 \\
 &= \frac{1}{2n^{3/2} \rho_n} \sum_{(a,b) \in \mathcal{I}} \left[(O_{a,b} - n^2 H_{a,b}^* \pi_a^* \pi_b^*) \log \frac{H'_{u(a),u(b)}(1 - H_{a,b}^*)}{(1 - H'_{u(a),u(b)})H_{a,b}^*} \right. \\
 \text{(A.23)} \quad &\quad \left. + (n_{a,b} - n^2 \pi_a^* \pi_b^*) \log \frac{1 - H'_{u(a),u(b)}}{1 - H_{a,b}^*} \right] + o_P(1) \\
 &= \frac{1}{2n^{3/2} \rho_n} \sum_{(a,b) \in \mathcal{I}} (n_{a,b} - n^2 \pi_a^* \pi_b^*) d_1(a, b) + o_P(1) \\
 &= \frac{\sqrt{n}}{2} \sum_{(a,b) \in \mathcal{I}} (n_{a,b}/n^2 - \pi_a^* \pi_b^*) d_2(a, b) + o_P(1),
 \end{aligned}$$

where $d_i(a, b)$ is defined in (2.8). (2.9) follows by the delta method. \square

PROOF OF LEMMA 2.6. The proof follows using arguments similar to Lemma 2.3. Note that in this case $G_1(R(z), H^*)$ is maximized at any $z \in \mathcal{V}_{K^+}$ with the value $\sum_{a,b} p_a p_b \gamma_1(H_{a,b}^*)$ [or $\sum_{a,b} p_a p_b \gamma_2(S_{a,b}^*)$ for $G_2(R(z), S^*)$].

It suffices to discuss the case $\rho_n \rightarrow 0$. Denote the optimal $G^* := \sum_{a,b} p_a p_b \times \gamma_2(S_{a,b}^*)$, define similarly to Lemma 2.3

$$J_{\delta_n} = \{z \in [K^+]^n : G_2(R(z), S^*) - G^* < -\delta_n\}$$

for $\delta_n \rightarrow 0$ slowly enough. It is easy to see

$$\sum_{z \in J_{\delta_n}} \sup_{\theta \in \Theta_{K^+}} f(z, A; \theta) \leq \sup_{\theta \in \Theta_{K^+}} f(z_0, A; \theta) o_P(1)$$

for any $z_0 \in \mathcal{V}_{K^+}$.

Next, note that treating $R(z)$ as a vector, $\{R(z) | z \in \mathcal{V}_{K^+}\}$ is a subset of the union of some of the $K^+ - K$ faces of the polyhedron P_R . For every $z \notin J_{\delta_n}, z \notin \mathcal{V}_{K^+}$, let z_\perp be such that $R(z_\perp) := \min_{R(z_0): z_0 \in \mathcal{V}_{K^+}} \|R(z) - R(z_0)\|_2$. $R(z) - R(z_\perp)$ is perpendicular to the corresponding $K^* - K$ face. Furthermore, this orthogonality implies the directional derivative of $G_2(\cdot, S^*)$ along the direction of $R(z) - R(z_\perp)$ is bounded away from 0. That is,

$$\left. \frac{\partial G_2((1 - \varepsilon)R(z_\perp) + \varepsilon R(z), S^*)}{\partial \varepsilon} \right|_{\varepsilon=0^+} < -C$$

for some universal positive constant C . Similar to (A.17),

$$\begin{aligned} \sup_{\theta \in \Theta_{K^+}} \log f(z, A; \theta) - \sup_{\theta \in \Theta_{K^+}} \log f(z_\perp, A; \theta) &\leq -\Omega_P(\mu_n) \frac{m}{n}, \\ \sup_{\theta \in \Theta_{K^+}} f(z, A; \theta) &\leq e^{-\Omega_P(\mu_n) \frac{m}{n}} \sup_{\theta \in \Theta_{K^+}} f(z_\perp, A; \theta), \end{aligned}$$

where $|z - z_\perp| = m$. We have

$$\begin{aligned} &\sum_{z \notin J_{\delta_n}, z \notin \mathcal{V}_{K^+}} \sup_{\theta \in \Theta_{K^+}} f(z, A; \theta) \\ &\leq \sum_{z \in \mathcal{V}_{K^+}} \sup_{\theta \in \Theta_{K^+}} f(z, A; \theta) \sum_{m=1}^n (K - 1)^m n^m e^{-\Omega_P(\mu_n) \frac{m}{n}} \\ &= o_P(1) \sum_{z \in \mathcal{V}_{K^+}} \sup_{\theta \in \Theta_{K^+}} f(z, A; \theta). \end{aligned}$$

Hence, the claim follows. \square

PROOF OF THEOREM 2.7. First, note

$$L_{K, K^+} = \log \frac{\sup_{\theta \in \Theta_{K^+}} g(A; \theta)}{f(Z, A; \theta^*)} + O_P(1),$$

where

$$\begin{aligned} \log \frac{\sup_{\theta \in \Theta_{K^+}} g(A; \theta)}{f(Z, A; \theta^*)} &\geq \log \frac{\sup_{\theta \in \Theta_{K^+}} f(Z, A; \theta)}{f(Z, A; \theta^*)} \\ \text{(A.24)} \qquad \qquad \qquad &= O_P(1). \end{aligned}$$

Let $D(\cdot)$ be a diagonal matrix, upper bounding by the maximum,

$$\begin{aligned}
 & \log \frac{\sup_{\theta \in \Theta_{K^+}} g(A; \theta)}{f(Z, A; \theta^*)} \\
 & \leq \max_z \sup_{\theta \in \Theta_{K^+}} \log \frac{f(z, A; \theta)}{f(Z, A; \theta^*)} + n \log K^+ \\
 & = \max_z \frac{n^2}{2} \{F_1(O(z)/n^2, N(z)/n^2) - F_1(D(p)H^*D(p), pp^T)\} + O_P(n) \\
 (A.25) \quad & \leq \max_z \frac{n^2}{2} |F_1(O(z)/n^2, N(z)/n^2) - F_1(RH^*R^T(z), R\mathbf{1}\mathbf{1}^T R^T(z))| \\
 & \quad + \max_z \frac{n^2}{2} [F_1(RH^*R^T(z), R\mathbf{1}\mathbf{1}^T R^T(z)) - F_1(D(p)H^*D(p), pp^T)] \\
 & \quad + O_P(n) \\
 & \leq C\mu_n \max_z \left\| \frac{O(z)}{\mu_n} - RS^*R^T \right\|_\infty + O_P(n) \\
 & = O_P(n^{3/2} \rho_n^{1/2})
 \end{aligned}$$

using (A.1) in Lemma A.1, and the fact that

$$\begin{aligned}
 & \max_{z \in [K^+]^n} F_1(RH^*R^T(z), R\mathbf{1}\mathbf{1}^T R^T(z)) \\
 & = F_1(D(p)H^*D(p), pp^T). \quad \square
 \end{aligned}$$

Next, we prove Theorem 2.10.

PROOF OF THEOREM 2.10. To upper bound L_{K, K^+} , by Lemma 2.6, it suffices to consider

$$\begin{aligned}
 & \max_{z \in \mathcal{V}_{K^+}} \sup_{\theta \in \Theta_{K^+}} \log f(z, A; \theta) - \sup_{\theta \in \Theta_K} \log g(A; \theta) + O(n) \\
 & = \max_{z \in \mathcal{V}_{K^+}} \sup_{\theta \in \Theta_{K^+}} \log f(z, A; \theta) - \log f(Z, A; \theta^*) + O(n).
 \end{aligned}$$

It follows from the definition of \mathcal{V}_{K^+} there exists a surjective function $h : [K^+] \rightarrow [K]$ describing the block assignments in $R(z, Z)$. We have

$$\begin{aligned}
 (A.26) \quad & \max_{z \in \mathcal{V}_{K^+}} \sup_{\theta \in \Theta_{K^+}} \log f(z, A; \theta) - \log f(Z, A; \theta^*) + O(n) \\
 & = \max_{z \in \mathcal{V}_{K^+}} n \sum_{k=1}^{K^+} \alpha(n_k(z)/n) - n \sum_{a=1}^K \alpha \left(\sum_{k \in h^{-1}(a)} n_k(z)/n \right)
 \end{aligned}$$

$$\begin{aligned}
& + \frac{1}{2} \sum_{k=1}^{K^+} \sum_{l=1}^{K^+} \left(O_{k,l} \log \frac{\hat{H}_{k,l}}{H_{h^{(k)},h^{(l)}}^*} + (n_{k,l} - O_{k,l}) \log \frac{1 - \hat{H}_{k,l}}{1 - H_{h^{(k)},h^{(l)}}^*} \right) \\
& + O_P(n),
\end{aligned}$$

where $\hat{H}_{k,l} = O_{k,l}(z)/n_{k,l}(z)$. The first part of the expression is nonpositive since α is superadditive.

Taylor expanding (A.26) and using the fact that $\hat{H}_{k,l} - H_{h^{(k)},h^{(l)}}^* = O_P(n^{-1/2} \rho_n^{1/2})$ for $z \in \mathcal{N}_{K^+}$ uniformly, (A.26) is upper bounded by $O_P(n)$. \square

Acknowledgments. The authors thank Haiyan Huang for helpful discussions and the reviewers for their valuable comments and suggestions.

SUPPLEMENTARY MATERIAL

Supplement to “Likelihood-based model selection for stochastic block models” (DOI: [10.1214/16-AOS1457SUPP](https://doi.org/10.1214/16-AOS1457SUPP); .pdf). A proof sketch of how the main results in the paper can be extended to the DCSBM described in Section 2.5 is provided in the supplement.

REFERENCES

- ADAMIC, L. A. and GLANCE, N. (2005). The political blogosphere and the 2004 US election: Divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery* 36–43. ACM, New York.
- AIROLDI, E. M., BLEI, D. M., FIENBERG, S. E. and XING, E. P. (2008). Mixed membership stochastic blockmodels. *J. Mach. Learn. Res.* **9** 1981–2014.
- AMINI, A. A., CHEN, A., BICKEL, P. J. and LEVINA, E. (2013). Pseudo-likelihood methods for community detection in large sparse networks. *Ann. Statist.* **41** 2097–2122. [MR3127859](#)
- BALL, B., KARRER, B. and NEWMAN, M. E. J. (2011). An efficient and principled method for detecting communities in networks. *Phys. Rev. E* (3) **84** 036103.
- BICKEL, P. and CHEN, A. (2009). A nonparametric view of network models and Newman–Girvan and other modularities. *Proc. Natl. Acad. Sci. USA* **106** 21068–73.
- BICKEL, P. J. and SARKAR, P. (2016). Hypothesis testing for automated community detection in networks. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **78** 253–273. [MR3453655](#)
- BICKEL, P., CHOI, D., CHANG, X. and ZHANG, H. (2013). Asymptotic normality of maximum likelihood and its variational approximation for stochastic blockmodels. *Ann. Statist.* **41** 1922–1943. [MR3127853](#)
- CELISSE, A., DAUDIN, J.-J. and PIERRE, L. (2012). Consistency of maximum-likelihood and variational estimators in the stochastic block model. *Electron. J. Stat.* **6** 1847–1899. [MR2988467](#)
- CHEN, K. and LEI, J. (2014). Network cross-validation for determining the number of communities in network data. Preprint. Available at [arXiv:1411.1715](https://arxiv.org/abs/1411.1715).
- CHOI, D. S., WOLFE, P. J. and AIROLDI, E. M. (2012). Stochastic blockmodels with a growing number of classes. *Biometrika* **99** 273–284. [MR2931253](#)
- DAUDIN, J.-J., PICARD, F. and ROBIN, S. (2008). A mixture model for random graphs. *Stat. Comput.* **18** 173–183. [MR2390817](#)
- DECELLE, A., KRZAKALA, F., MOORE, C. and ZDEBOROVÁ, L. (2011). Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Phys. Rev. E* (3) **84** 066106.

- FISHKIND, D. E., SUSSMAN, D. L., TANG, M., VOGELSTEIN, J. T. and PRIEBE, C. E. (2013). Consistent adjacency-spectral partitioning for the stochastic block model when the model parameters are unknown. *SIAM J. Matrix Anal. Appl.* **34** 23–39. [MR3032990](#)
- HOLLAND, P. W., LASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Social Networks* **5** 109–137. [MR0718088](#)
- JOSEPH, A. and YU, B. (2013). Impact of regularization on spectral clustering. Preprint. Available at [arXiv:1312.1733](#).
- KARRER, B. and NEWMAN, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Phys. Rev. E* (3) **83** 016107, 10. [MR2788206](#)
- LATOUCHE, P., BIRMELÉ, E. and AMBROISE, C. (2012). Variational Bayesian inference and complexity control for stochastic block models. *Stat. Model.* **12** 93–115. [MR2953099](#)
- LE, C. M. and LEVINA, E. (2015). Estimating the number of communities in networks by spectral methods. Preprint. Available at [arXiv:1507.00827](#).
- LEI, J. (2014). A goodness-of-fit test for stochastic block models. Preprint. Available at [arXiv:1412.4857](#).
- LESKOVEC, J. and MCAULEY, J. J. (2012). Learning to discover social circles in ego networks. In *Adv. Neural Inf. Process. Syst.* 539–547. Lake Tahoe.
- NEWMAN, M. E. J. (2006a). Modularity and community structure in networks. *Proc. Natl. Acad. Sci. USA* **103** 8577–8582.
- NEWMAN, M. E. J. (2006b). Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* (3) **74** 036104, 19. [MR2282139](#)
- PEIXOTO, T. P. (2013). Parsimonious module inference in large networks. *Phys. Rev. Lett.* **110** 148701.
- ROHE, K., CHATTERJEE, S. and YU, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.* **39** 1878–1915. [MR2893856](#)
- SALDANA, D. F., YU, Y. and FENG, Y. (2014). How many communities are there? Preprint. Available at [arXiv:1412.1684](#).
- WANG, Y. X. R. and BICKEL, P. J. (2016). Supplement to “Likelihood-based model selection for stochastic block models.” DOI:[10.1214/16-AOS1457SUPP](#).
- WOLFE, P. J. and OLHEDE, S. C. (2013). Nonparametric graphon estimation. Preprint. Available at [arXiv:1309.5936](#).
- ZHAO, Y., LEVINA, E. and ZHU, J. (2011). Community extraction for social networks. *Proc. Natl. Acad. Sci. USA* **108** 7321–7326.
- ZHAO, Y., LEVINA, E. and ZHU, J. (2012). Consistency of community detection in networks under degree-corrected stochastic block models. *Ann. Statist.* **40** 2266–2292. [MR3059083](#)

DEPARTMENT OF STATISTICS
SEQUOIA HALL
390 SERRA MALL
STANFORD UNIVERSITY
STANFORD, CALIFORNIA 94305-4020
USA
E-MAIL: ryxwang@stanford.edu

DEPARTMENT OF STATISTICS
UNIVERSITY OF CALIFORNIA, BERKELEY
367 EVANS HALL
BERKELEY, CALIFORNIA 94720
USA
E-MAIL: bickel@stat.berkeley.edu