

ASYMPTOTIC NORMALITY OF MAXIMUM LIKELIHOOD AND ITS VARIATIONAL APPROXIMATION FOR STOCHASTIC BLOCKMODELS

BY PETER BICKEL , DAVID CHOI , XIANGYU CHANG AND HAI ZHANG

University of California, Berkeley

Carnegie Mellon University

Xi'an Jiaotong University

Northwest University

Variational methods for parameter estimation are an active research area, potentially offering computationally tractable heuristics with theoretical performance bounds. We build on recent work that applies such methods to network data, and establish asymptotic normality rates for parameter estimates of stochastic blockmodel data, by either maximum likelihood or variational estimation. The result also applies to various sub-models of the stochastic blockmodel found in the literature.

1. Introduction. The analysis of network data is an open statistical problem, with many potential applications in the social sciences (Lazer et al., 2009) and in biology (Proulx, Promislow and Phillips, 2005). In such applications, the models tend to pose both computational and statistical challenges, in that neither their fitting method nor their large sample properties are well-understood.

However, some results are becoming known for a model known as the stochastic blockmodel, which assumes that the network connections are explainable by a latent discrete class variable associated with each node. For this model, consistency has been shown for profile likelihood maximization (Bickel and Chen, 2009), a spectral-clustering based method (Rohe, Chatterjee and Yu, 2011), and other methods as well (Bickel, Chen and Levina, 2011; Channarond, Daudin and Robin, 2011; Choi, Wolfe and Airoldi, 2012; Coja-Oghlan and Lanka, 2008), under varying assumptions on the sparsity of the network and the number of classes. These results suggest that the model has reasonable statistical properties, and empirical experiments suggest that efficient approximate methods may suffice to find the parameter estimates. However, formally there is no satisfactory inference theory for the behavior of classical procedures such as maximum likelihood under the

AMS 2000 subject classifications: network statistics, stochastic blockmodeling, variational methods, maximum likelihood

model, nor for any procedure which is computationally not potentially NP under worst-case analysis.

In this note, we establish both consistency and asymptotic normality of maximum likelihood estimation, and also of a variational approximation method, considering sparse models and restricted sub-models. To some extent, we are following a pioneering paper of Celisse et. al. (Celisse, Daudin and Pierre, 2011), in which the dense model was considered, and consistency was established, but only for a subset of the parameters.

2. Preliminaries.

2.1. Stochastic Blockmodel. We consider a class of latent variable models considered by various authors (Karrer and Newman, 2011; Latouche, Birmelé and Ambroise, 2011; Snijders and Nowicki, 1997; Bickel and Chen, 2009), which we describe as follows. Let $Z = (Z_1, \dots, Z_n)$ be latent random variables corresponding to vertices $1, \dots, n$, taking values in $[K] \equiv \{1, \dots, K\}$. We will assume that K is fixed and does not increase with n . Let π be a distribution on $[K]$, and let H be a symmetric matrix in $[0, 1]^{K \times K}$. We define the *complete graph model* (CGM) for Z, A , where A is the $n \times n$ symmetric 0-1 adjacency matrix of a graph, by its distribution

$$(1) \quad f(Z, A) = \left(\prod_{i=1}^n \pi(Z_i) \right) \left(\prod_{i=1}^n \prod_{j=i+1}^n H(Z_i, Z_j)^{A_{ij}} (1 - H(Z_i, Z_j))^{1-A_{ij}} \right),$$

where we may interpret $H(Z_i, Z_j)$ as $\mathbb{P}(\text{edge}|Z_i, Z_j)$, and $\pi(a)$ as $\mathbb{P}(Z_i = a)$ for $a = 1, \dots, K$.

The *graph model* (GM) is defined by a distribution $g : \{0, 1\}^{n \times n} \rightarrow [0, 1]$, which satisfies $g(A) = \mathbb{P}(A; H, \pi)$ and is given by

$$g(A) = \sum_{z \in [K]^n} f(z, A)$$

It is data from GM which we assume we observe.

We will allow (H, π) to be parameterized by θ taking values in some restricted space Θ , so that parametric submodels of the blockmodel may be considered. We will consider parameterizations of the form $\theta = (\rho, \phi)$, in which

$$H_\theta \equiv \rho S_\phi, \quad \pi_\theta \equiv \pi_\phi, \quad \sum_{a,b=1}^K \pi_\phi(a) \pi_\phi(b) S_\phi(a, b) = 1$$

where $\rho > 0$ is a nonnegative scalar; ϕ is a Euclidean parameter ranging over an open set; S_ϕ is a symmetric matrix in $\mathbb{R}^{K \times K}$; and the map $\phi \mapsto (\pi_\phi, S_\phi)$ is assumed to be smooth. Let $\lambda = n\rho$. The interpretation of these parameters is that $\lambda = \mathbb{E}[\text{degree}]$ and $\rho = \mathbb{P}(A_{ij} = 1)$. The utility of this parameterization will be to analyze asymptotic behavior when $\rho \equiv \rho_n \rightarrow 0$ while ϕ is kept fixed, as seems reasonable for sparse network settings.

Identifiability of the model. We observe that f is symmetric under permutation of Z and θ ; that is, let $\sigma : [K] \rightarrow [K]$ denote a permutation of $[K]$, and let Π denote its permutation matrix. For $z \in [K]^n$, let $\sigma(z) = (\sigma(z_1), \dots, \sigma(z_n))$, and for $\theta \equiv (\pi, H)$, let $\sigma(\theta) = (\Pi\pi, \Pi H \Pi^T)$. It then holds for any permutation σ that

$$f(Z, A; \theta) = f(\sigma(Z), A; \sigma(\theta)),$$

and hence

$$g(A; \theta) = g(A; \sigma(\theta)),$$

showing that when Z is latent, the stochastic blockmodel is non-identifiable. Specifically, $\theta \equiv (\pi, H)$ is equivalent to $\sigma(\theta) \equiv (\pi\Pi, \Pi H \Pi^T)$ for any permutation σ . Let \mathcal{S}_θ denote this equivalence class, which corresponds to a relabeling of the latent classes $\{1, \dots, K\}$. By an estimate $\hat{\theta}$ under the GM blockmodel, we will mean the equivalence class $\mathcal{S}_{\hat{\theta}}$. By consistency and asymptotic normality of $\hat{\theta}$, we will mean that $\mathcal{S}_{\hat{\theta}}$ contains an element θ' that converges to the generative θ_0 , or has error $r_n(\theta' - \theta_0)$ that is asymptotically normal distributed for some rate $r_n \rightarrow 0$.

In our analysis, we will assume that the generative H has no identical rows, as we cannot expect to successfully distinguish classes which behave identically. If H did contain identical rows, then an additional source of non-identifiability would exist. Also, the generative model would be equivalent to a stochastic blockmodel of smaller order K . We do not treat such cases here.

We note that for some restricted submodels, identifiability can be restored by imposing a canonical ordering of the latent classes $1, \dots, K$. For example, the submodel may restrict H so that $H(Z_i, Z_j)$ depends only on whether $Z_i = Z_j$ or not; this assumption could reflect homogeneity of the classes, and is explored in (Rohe, Qin and Fan, 2012). This submodel is identifiable under ordering of π , and the latent structure might be more gracefully described as a partition, i.e., a variable $X \in \{0, 1\}^{n \times n}$ satisfying $X(i, j) = 1$ iff $Z_i = Z_j$. As a second example, the latent classes could be ordered by decreasing expected degree. If the submodel restricts the expected degrees to be unique, the submodel is identifiable; further discussion can be found in (Celisse, Daudin and Pierre, 2011; Bickel, Chen and Levina, 2011).

Degree-corrected blockmodels. An interesting class of submodels, discussed in (Karrer and Newman, 2011; Zhao, Levina and Zhu, 2013), are the “degree-corrected” blockmodels with UV -many classes obtained by considering $Z_i = (Z_{i1}, Z_{i2})$, for $i = 1, \dots, n$, which take values (u, v) ; where u takes values $1, \dots, U$ with probabilities $\alpha_1, \dots, \alpha_U$; and given parameters $\gamma_1, \dots, \gamma_V \in [0, 1]$, v takes values $\gamma_1, \dots, \gamma_V$ with probabilities β_1, \dots, β_V . We will assume Z_{i1} and Z_{i2} are independent. Additional parameters needed are a $U \times U$ symmetric matrix of probabilities G . We can now define

$$\mathbb{P}(Z_{i1} = a, Z_{i2} = \gamma_c, Z_{j1} = b, Z_{j2} = \gamma_d | A_{ij} = 1) = \alpha_a \alpha_b \beta_c \beta_d \gamma_c \gamma_d G(a, b).$$

So although this is a UV blockmodel, it has only $U(U+1)/2 + (U-1) + (2V-1)$ parameters. Its interpretation is that there are U subblocks, but within each subblock, vertices can hierarchically exhibit further affinities to vertices both within the same block and other blocks, thus enabling, for instance, distinction between vertices of high degree and low degree within each block. This distinction is not block-dependent, resulting in a reduction of parameters.

Many variants are of course possible; for example, one can choose to have more parameters by having the (u, v) block probabilities be free, so that the conditional distribution of Z_{i2} dependent on Z_{i1} , or fewer parameters by treating $\alpha(1), \dots, \alpha(U)$ as known.

More general models. The stochastic blockmodel is a special case of a more general latent variable model, considered by various authors (Hoff, Raftery and Handcock, 2002; Bickel and Chen, 2009; Bollobás, Janson and Riordan, 2007). In this model, the elements of Z take values in a general space \mathcal{Z} rather than $[K]$, π is a distribution on \mathcal{Z} , and H is replaced by a symmetric map $h : \mathcal{Z} \times \mathcal{Z} \rightarrow [0, 1]$. The CGM defines a density for (Z, A) , with respect to an appropriate reference measure, and GM satisfies the identity

$$(2) \quad \frac{g(A; \theta)}{g_0(A)} = \mathbb{E}_{\theta_0} \left[\frac{f(Z, A; \theta)}{f_0(Z, A)} \mid A \right],$$

where f_0 and g_0 denote the distribution under the generative θ_0 . This model is considered in (Bickel and Chen, 2009) with $\{Z_i\}_{i=1}^n$ assumed i.i.d Uniform $(0, 1)$. In (Handcock, Raftery and Tantrum, 2007), they are a multivariate mixture of Gaussians with unknown parameters. If we make no restrictions on h , these models are equivalent.

2.2. Maximum likelihood and variational estimates. For the complete graph blockmodel, maximum likelihood estimation of H and π (or of θ)

is basically understood. From Eq. (1) it can be seen that the log likelihood expression decomposes, so that π is estimated from Z independently of A , and H is estimated from A conditional on Z . We note that it is possible for the likelihood to have multiple local optima.

For the GM blockmodel, the maximum likelihood parameter estimate $\hat{\theta}^{ML}$ (i.e., the equivalence class $\mathcal{S}_{ML} \equiv \mathcal{S}_{\hat{\theta}^{ML}}$) is given by

$$\begin{aligned}\hat{\theta}^{ML} &= \arg \max_{\theta} g(A; \theta) \\ &= \arg \max_{\theta} \sum_{z \in [K]^n} f(z, A, \theta).\end{aligned}$$

Two difficulties in computing $\hat{\theta}^{ML}$ present themselves: first, multiple local optima in g may exist even if the CGM likelihood function f is concave in the appropriate parameterization, as we shall see for the ordinary unrestricted parameterization. Second, the maximum likelihood estimate involves a generally intractable marginalization over the latent variable Z .

Variational methods attempt to circumvent the second difficulty (while accepting the first) by introducing an approximate function J for which local optimization is computationally easier. For the GM blockmodel, the estimate $\hat{\theta}^{VAR}$ (i.e., the equivalence class $\mathcal{S}_{\hat{\theta}^{VAR}}$) is given by

$$\begin{aligned}\hat{\theta}^{VAR} &= \arg \max_{\theta} \max_{q \in \mathcal{D}} J(q, \theta; A) \\ &\triangleq \arg \max_{\theta} \max_{q \in \mathcal{D}} -D(q || f_{Z|A;\theta}) + \log g(A; \theta).\end{aligned}$$

Here \mathcal{D} is the set of all product distributions over \mathcal{Z}^n , with densities denoted by $\prod_{i=1}^n q_i(\cdot)$. The term $D(\cdot || \cdot)$ is the Kullback-Leibler divergence, and $f_{Z|A;\theta}$ is the conditional density of Z given A , i.e., $f_{Z|A;\theta}(Z) = \frac{f(Z, A; \theta)}{g(A; \theta)}$. The Kullback-Leibler divergence is given by

$$D(q || f_{Z|A;\theta}) = \sum_{z \in [K]^n} q(z) \log \frac{q(z)}{f_{Z|A;\theta}(z)}.$$

We note that J simplifies to

$$\begin{aligned}J(q, \theta; A) &= \sum_{i=1}^n \sum_{a=1}^K q_i(a) [-\log q_i(a) + \log \pi_{\theta}(a)] \\ &+ \sum_{i=1}^n \sum_{j=i+1}^n \sum_{a=1}^K \sum_{b=1}^K q_i(a) q_j(b) [A_{ij} \log H_{\theta}(a, b) + (1 - A_{ij}) \log(1 - H_{\theta}(a, b))].\end{aligned}$$

This formula indicates that, at least for the complete parameterization, a local optimum to J can be tractably computed for moderate n and K using the EM algorithm as in (Daudin, Picard and Robin, 2008). In contrast, optimization of g requires a summation over $[K]^n$ which is generally intractable. However, note that we have added $n(K - 1)$ new parameters.

Intuitively, we expect the variational estimate to approximate the maximum likelihood estimate when there exists $q \in \mathcal{D}$ which is close to $f_{Z|A;\theta}$.

We remark that $\max_q \exp(J(q, \theta; A))$ is upper and lower bounded by

$$(3) \quad f(z, A; \theta) \leq \max_q \exp(J(q, \theta; A)) \leq g(A; \theta),$$

for any $z \in [K]^n$. To see this, consider that the lower bound is an equality if $q = \delta_z$, while the upper bound holds due to nonnegativity of the Kullback-Leibler divergence.

Other estimation problems. Our focus here is on estimation of the generative θ_0 . In other papers, estimation of the latent Z is considered to be the primary inferential task (Rohe, Chatterjee and Yu, 2011; Choi, Wolfe and Airolidi, 2012). We feel that both tasks are of interest. For example, if the data A represents a network observed in its entirety, estimating Z and quantifying its uncertainty may give insight into the underlying network structure and the roles of its actors. On the other hand, if A is understood to be a representative sample of a larger population, whose overall structure is of interest, estimates of θ would be preferable.

3. Results.

3.1. *Asymptotic normality of maximum likelihood under CGM blockmodel.* We first review the asymptotics of the CGM blockmodel with complete parameterization.

Parameterize $\theta \equiv (\varpi, \nu)$, where $\varpi \in \mathbb{R}^K$ and $\nu \in \mathbb{R}^{K \times K}$ are the logit of π and H , given by

$$(4) \quad \begin{aligned} \varpi(a) &= \log \frac{\pi(a)}{1 - \sum_{b=1}^{K-1} \pi(b)} & a = 1, \dots, K-1 \\ \nu(a, b) &= \log \frac{H(a, b)}{1 - H(a, b)} & a, b = 1, \dots, K, \end{aligned}$$

and let \mathcal{T} denote the canonical parameter space $\{\theta : \varpi \in \mathbb{R}^{K-1}, \nu \in \mathbb{R}^{K(K+1)/2}\}$. Let (Z, A) denote data generated by the model, under the generative parameter θ_0 , and let f_0 denote f under θ_0 . For the CGM blockmodel,

the log likelihood ratio $\Lambda = \log \frac{f}{f_0}$ is given by

$$\begin{aligned} \Lambda(\theta, Z, A) &= \sum_{a=1}^{K-1} \left[(\varpi(a) - \varpi_0(a))n_a - n \log \frac{1 + \sum_{a=1}^{K-1} e^{\varpi(a)}}{1 + \sum_{a=1}^{K-1} e^{\varpi_0(a)}} \right] \\ &\quad + \frac{1}{2} \sum_{a=1}^K \sum_{b=1}^K \left[(\nu(a, b) - \nu_0(a, b))O_{ab} - n_{ab} \log \frac{1 + e^{\nu(a, b)}}{1 + e^{\nu_0(a, b)}} \right], \end{aligned}$$

where

$$n_a \equiv n_a(Z) = \sum_{i=1}^n 1\{Z_i = a\}, \quad n_{ab} \equiv n_{ab}(Z) = \sum_{i=1}^n \sum_{j \neq i}^n 1\{Z_i = a, Z_j = b\}$$

$$O_{ab} \equiv O_{ab}(A, Z) = \sum_{i=1}^n \sum_{j \neq i}^n 1\{Z_i = a, Z_j = b\} A_{ij}.$$

This is an exponential family in θ . The gradient of Λ conditioned on Z , evaluated at $\theta' \in \mathcal{T}$, is given by

$$\begin{aligned} \frac{\partial \Lambda}{\partial \varpi(a)}(\theta') &= n_a - n\pi'(a) & a = 1, \dots, K-1 \\ \frac{\partial \Lambda}{\partial \nu(a, b)}(\theta') &= O_{ab} - n_{ab}H'(a, b) & a, b = 1, \dots, K \end{aligned}$$

Using the parameterization (π, H) , the maximum likelihood estimates are given by

$$\hat{\pi}^{CGM}(a) = \frac{n_a}{n}, a = 1, \dots, K \quad \text{and} \quad \hat{H}^{CGM}(a, b) = \frac{O_{ab}}{n_{ab}}, a, b = 1, \dots, K.$$

We note that the parameterizations (ϖ, ν) and (π, ρ, S) are both identifiable under the CGM blockmodel.

LEMMA 1. *Assume the generative parameter $\theta_0 \in \mathcal{T}$ satisfies $(\log n)^{-1}\lambda_0 \rightarrow \infty$, with π_0 and S_0 constant in n . It holds that*

$$\begin{aligned} \sqrt{n}(\hat{\varpi}^{CGM} - \varpi_0) &\rightarrow N(0, \Sigma_1) \\ \sqrt{n\lambda_0}(\hat{\nu}^{CGM} - \nu_0) &\rightarrow N(0, \Sigma_2), \end{aligned}$$

where Σ_1 and Σ_2 are functions of θ_0 .

PROOF. The log likelihood ratio Λ can be decomposed into two terms which involve ϖ and ν separately. Asymptotic normality of $\hat{\pi}^{CGM}$ and $\hat{\varpi}^{CGM}$ follows from standard exponential family theory. It can be seen that

$$\begin{aligned} \sqrt{n\lambda_0} \left(\frac{\hat{H}^{CGM}(a,b)}{\rho_0} - \frac{H_0(a,b)}{\rho_0} \right) &= \sqrt{n^2/n_{ab}\sqrt{n_{ab}\rho_0}} \left(\frac{\hat{H}^{CGM}(a,b)}{\rho_0} - S_0 \right) \\ &= (\pi_0(a)\pi_0(b) + o_P(1))^{-1/2} \sqrt{n_{ab}\rho_0} \left(\frac{\hat{H}^{CGM}(a,b)}{\rho_0} - S_0 \right), \end{aligned}$$

which is asymptotically normal by a Lindeberg central limit theorem. Since $(\hat{\nu}^{CGM} - \nu_0) = \left(\log \frac{\hat{H}^{CGM}(a,b)}{\rho_0} - \log \frac{H_0(a,b)}{\rho_0} + o_P(1) \right)$, asymptotic normality of $\hat{\nu}^{CGM}$ follows by the delta method. \square

Let $\mathbf{H}(\theta') = D_{\theta'}^2 \Lambda(Z, A; \theta)|_{\theta=\theta'}$ denote the conditional hessian of Λ evaluated at θ' . For all $\theta' \in \mathcal{T}$, $\mathbf{H}(\theta')$ is given by

$$(5) \quad \frac{\partial^2 \Lambda}{\partial \varpi(a) \partial \varpi(a)}(\theta') = n\pi'(a)(1 - \pi'(a)) \quad a = 1, \dots, K-1$$

$$(6) \quad \frac{\partial^2 \Lambda}{\partial \varpi(a) \partial \varpi(b)}(\theta') = n\pi'(a)\pi'(b) \quad a, b = 1, \dots, K-1$$

$$(7) \quad \frac{\partial^2 \Lambda}{\partial \nu(a,b) \partial \nu(a,b)}(\theta') = n_{ab}H'(a,b)(1 - H'(a,b)) \quad a, b = 1, \dots, K$$

with all other terms equal to zero.

LEMMA 2 (Local asymptotic normality). *For the CGM blockmodel with parameter values $(\varpi_0, \nu_0) \equiv (\pi_0, \rho_n, S_0) \in \mathcal{T}$, it holds uniformly for any s, t in a compact set that*

$$(8) \quad \Lambda \left(\varpi_0 + \frac{s}{\sqrt{n}}, \nu_0 + \frac{t}{\sqrt{n^2 \rho_n}} \right) = s^T Y_1 + t^T Y_2 - \frac{1}{2} s^T \Sigma_1 s - \frac{1}{2} t^T \Sigma_2 t + o_P(1)$$

where Σ_1 and Σ_2 are functions of ϖ_0 and ν_0 , and Y_1, Y_2 are asymptotically normal distributed with zero mean and covariances Σ_1 and Σ_2 respectively.

PROOF. By Taylor expansion,

$$\begin{aligned} \Lambda \left(\varpi_0 + \frac{s}{\sqrt{n}}, \nu_0 + \frac{t}{\sqrt{n^2 \rho_n}} \right) &= \Lambda(\varpi_0, \nu_0) + \frac{1}{\sqrt{n}} s^T \nabla \Lambda_{\varpi}(\theta_0) + \frac{1}{\sqrt{n\lambda_0}} t^T \nabla \Lambda_{\nu}(\theta_0) \\ &\quad + \frac{1}{n} s^T \mathbf{H}_{\varpi}(\theta_0) s + \frac{1}{n^2 \rho_n} t^T \mathbf{H}_{\nu}(\theta_0) t + o_P(1), \end{aligned}$$

where $\nabla\Lambda_{\varpi}(\theta_0)$ and $\nabla\Lambda_{\nu}(\theta_0)$ denote the respective components of the gradient of Λ evaluated at θ_0 , and \mathbf{H}_{ϖ} and \mathbf{H}_{ν} are given by (5) - (7) which describe $\mathbf{H}(\theta_0)$. By inspection, $\Lambda(\varpi_0, \nu_0) = 0$; \mathbf{H}_{ϖ}/n and $\mathbf{H}_{\nu}/n^2\rho_n$ converge in probability to constant matrices; and the random vectors $n^{-1/2}\nabla\Lambda_{\varpi}$ and $(n\lambda_0)^{-1/2}\nabla\Lambda_{\nu}$ converge in distribution by central limit theorem. This establishes (8). \square

For submodels where $\theta \mapsto (\varpi, \nu)$ covers a restricted subset $\Theta \subset \mathcal{T}$, we generally have $\theta_0 - \hat{\theta}^{CGM} = O_P(n^{-1/2})$. However, if θ is separable into (θ_{π}, θ_S) such that $\pi = \pi_{\theta_{\pi}}$ and $S = S_{\theta_S}$, and θ_{π} and θ_S are allowed to vary freely, then θ_S has error that is asymptotically normal with the faster rate $\sqrt{n\lambda}$. Independence of the errors in θ_S and θ_{π} is then also valid as well.

3.2. Asymptotic normality of maximum likelihood under GM blockmodel. Our main result is that for graphs with poly-log expected degree, the likelihood ratios of the CGM and GM blockmodels are essentially equivalent with probability tending to 1, so that inference under the models is essentially equivalent up to the identifiability restrictions of the GM blockmodel.

THEOREM 1. *Let (Z, A) be generated from a blockmodel with $\theta_0 \in \mathcal{T}$, such that S_0 has no identical columns, and $\rho_0 = \rho_n$ satisfies $n\rho_n/\log n \rightarrow \infty$. Then for all $\theta \in \mathcal{T}$,*

$$(9) \quad \frac{g}{g_0}(A, \theta) = \max_{\theta' \in \mathcal{S}_{\theta}} \frac{f}{f_0}(Z, A, \theta')(1 + \epsilon_n(K, \theta')) + \epsilon_n(K, \theta'),$$

where $\sup_{\theta \in \mathcal{T}} \epsilon_n(K, \theta) = o_P(1)$.

Theorem 1 is proven in the appendix, and can be viewed as the sum of two parts.

1. In neighborhoods around (ϖ_0, ν_0) , of order $(n^{-1/2}, (n\lambda)^{-1/2})$, both f/f_0 and g/g_0 are of order 1 and their difference is $o_P(1)$. We show this using methods similar to (Bickel and Chen, 2009), but it may also be deduced from a general result in (Le Cam and Yang, 1988); in their terminology, the profile likelihood estimate is a distinguished statistic.
2. In the exterior of neighborhoods as above, both f/f_0 and g/g_0 are both $o_P(1)$ on complements of balls around \mathcal{S}_{θ_0} and converge uniformly to 0. Unlike the first, this part does not seem to follow from (Le Cam and Yang, 1988).

Asymptotic normality of $\hat{\theta}^{ML}$ follows from Theorem 1 and Lemma 1, as stated in the following theorem.

THEOREM 2. *Assuming the conditions of Theorem 1 and Lemma 2, let $\hat{\varpi}^{ML}, \hat{\nu}^{ML}$ and $\hat{\varpi}^{CGM}, \hat{\nu}^{CGM}$ be the corresponding maximum likelihood estimates over all $\theta \in \mathcal{T}$. It holds that \mathcal{S}_{ML} contains an element θ' satisfying*

$$\begin{aligned}\varpi' - \hat{\varpi}^{CGM} &= o_P(n^{-1/2}) \\ \nu' - \hat{\nu}^{CGM} &= o_P((n\lambda_0)^{-1/2}).\end{aligned}$$

PROOF. For each $\theta' \in \mathcal{S}_{ML}$ it holds that if either $|\varpi' - \hat{\varpi}^{CGM}| \neq o_P(n^{-1/2})$ or $|\nu' - \hat{\nu}^{CGM}| \neq o_P((n\lambda_n)^{-1/2})$, then by (8) and consistency of $\hat{\theta}^{CGM}$,

$$\Lambda(\hat{\theta}^{CGM}; A, Z) - \Lambda(\theta'; A, Z) = \Omega_P(1).$$

Thus, we may prove the Lemma by establishing the contrapositive. Since $\hat{\theta}^{ML}$ and $\hat{\theta}^{CGM}$ respectively maximize $\frac{g}{g_0}$ and $\frac{f}{f_0}$, it follows by Theorem 1 that for some $\theta' \in \mathcal{S}_{ML}$, $\left| \frac{f}{f_0}(Z, A, \theta^{CGM}) - \frac{f}{f_0}(Z, A, \theta') \right| = o_P(1)$, implying that $\Lambda(\hat{\theta}^{CGM}) - \Lambda(\theta') = o_P(1)$ for some $\theta' \in \mathcal{S}_{ML}$. \square

A parametrized submodel, such as the degree corrected block model as discussed earlier, has likelihood $g(A; \varpi(\theta), \nu(\theta))$. Theorem 1 applies, and if the mapping $\theta \mapsto (\varpi, \nu)$ is smooth, then if estimates for the CGM block submodel exist and are asymptotically normal, their equivalents in the corresponding GM model will have equivalent behavior, up to the identifiability issues that we have discussed. Of course, if CGM block submodel estimates do not exist or are not consistent, this will be inherited by the GM block submodel estimates as well.

3.3. *Asymptotic normality of variational estimates under GM blockmodel.* We show that same properties that we have established for maximum likelihood estimates under the GM blockmodel also hold for the more computable variational likelihood estimates.

Our proof will use a lemma which follows from the main result of (Bickel and Chen, 2009).

LEMMA 3. *Let (A, Z) be generated by $\theta_0 \equiv (\rho_n, \pi_0, S_0) \in \mathcal{T}$, such that $n\rho_n/\log n \rightarrow \infty$ and S_0 has no identical columns. It holds that*

$$(10) \quad f(A, Z; \theta_0)/g(A; \theta_0) = 1 + o_P(1).$$

PROOF. By exponential family theory, given a non-identity permutation σ , it holds that $\frac{f}{f_0}(A, Z; \sigma(\theta_0)) = o_P(1)$, and hence that $\frac{f}{f_0}(A, \sigma^{-1}(Z); \theta_0) =$

$o_P(1)$ as well. As a result,

$$(11) \quad \sum_{Z' \in \mathcal{S}_Z, Z' \neq Z} \frac{f(A, Z'; \theta_0)}{g(A; \theta_0)} \leq \sum_{Z' \in \mathcal{S}_Z, Z' \neq Z} \frac{f}{f_0}(A, Z'; \theta_0) = o_P(1).$$

Given (A, Z) generated under θ_0 , let $\hat{z}(A)$ denote the maximum profile likelihood estimate of Z , that is, the set $\arg \max_z \sup_{\theta} f(A, z; \theta)$. Let \mathcal{S}_Z denote the set of all labels Z' such that $Z' = \sigma(Z)$ for some permutation $\sigma : [K] \rightarrow [K]$. Theorem 1 from (Bickel and Chen, 2009) states that under the conditions of this lemma,

$$\lim \frac{\log \mathbb{P}_0(\mathcal{S}_Z \neq \hat{z}(A))}{\lambda_n} \leq -s_Q(\pi_0, S_0) < 0.$$

This implies that $\mathbb{P}(Z \notin \mathcal{S}_{\hat{z}(A)}) = o(1)$. By Markov's inequality, this implies that $\mathbb{P}(Z \notin \mathcal{S}_{\hat{z}(A)} | A) = o_P(1)$, which can be rewritten as

$$(12) \quad \sum_{Z' \notin \mathcal{S}_Z} f(A, Z'; \theta_0) / g(A; \theta_0) = o_P(1).$$

Combining (12) and (11) establishes (10). \square

Our result for the variational estimates is

THEOREM 3. *Let $J(\theta; A)$ denote $\max_{q \in \mathcal{D}} \exp[J(q, \theta; A)]$. Under the conditions of Theorem 1 and Lemma 2,*

$$(13) \quad \frac{J(\theta; A)}{g(A; \theta_0)} = \max_{\theta' \in \mathcal{S}_{\theta}} \frac{f}{f_0}(Z, A, \theta') (1 + \epsilon_n(K, \theta')) + \epsilon_n(K, \theta'),$$

where $\sup_{\theta \in \mathcal{T}} \epsilon_n(K, \theta) = o_P(1)$. Hence, the conclusions of Theorem 2 also apply to $(\hat{\pi}^{VAR}, \hat{S}^{VAR})$, the variational likelihood estimates.

PROOF. Recall (3) which states that for all z ,

$$f(z, A; \theta) \leq \max_q \exp(J(q, \theta; A)) \leq g(A; \theta).$$

Dividing the lower bound by $f(A, Z; \theta_0)$, which equals $g(A; \theta_0)(1 + o_P(1))$ by Lemma 3, yields

$$\max_{z \in \mathcal{S}_Z} \frac{f}{f_0}(z, A; \theta) \leq \frac{J(\theta; A)}{g(A; \theta_0)(1 + o_P(1))}.$$

The identity $\max_{z \in \mathcal{S}_Z} \frac{f}{f_0}(z, A; \theta) = \max_{\theta' \in \mathcal{S}_\theta} \frac{f}{f_0}(Z, A; \theta')$ thus implies

$$(14) \quad \max_{\theta' \in \mathcal{S}_\theta} \frac{f}{f_0}(z, A; \theta') \leq \frac{J(\theta; A)}{g(A; \theta_0)(1 + o_P(1))}.$$

Dividing the upper bound by $g(A; \theta_0)$, and applying Theorem 1 yields

$$(15) \quad \frac{J(\theta; A)}{g(A; \theta_0)} \leq \frac{g}{g_0}(A; \theta) \leq \max_{\theta' \in \mathcal{S}_\theta} \frac{f}{f_0}(Z, A; \theta')(1 + \epsilon_n(K, \theta')) + \epsilon_n(K, \theta').$$

Combining (14) and (15) to upper and lower bound $\frac{J(\theta; A)}{g(A; \theta_0)}$ proves the theorem. \square

4. Some statistical applications. With these results, we can show that some standard inference is valid using the likelihood or variational likelihood for blockmodels.

Confidence regions for θ . We have that $\hat{\theta}^{VAR}$ under P_{θ_0} is asymptotically normal with mean θ_0 and variance covariance matrices given by Theorem 2 and Lemma 1. Since $\theta \mapsto \Sigma(\theta)$ is continuous, we can evidently form tests and confidence regions based on $\sqrt{n}(\hat{\varpi}^{VAR} - \varpi_0)^T \hat{\Sigma}_1^{-1/2}$ and $\sqrt{n\hat{\lambda}}(\hat{\nu}^{VAR} - \nu_0)^T \hat{\Sigma}_2^{-1/2}$, where $\hat{\Sigma}_1$ and $\hat{\Sigma}_2$ are plug-in estimates of Σ_1 and Σ_2 using $\hat{\theta}^{VAR}$, and $\hat{\lambda}$ equals the average degree in the observed data. The same applies to $\hat{\theta}^{ML}$.

Wilks statistic for hypothesis testing. Under the CGM blockmodel with generative parameter θ_0 , the Wilks (or likelihood ratio) statistic is given by

$$\Lambda(Z, A; \hat{\theta}^{CGM}) \equiv 2 \log \frac{f}{f_0}(Z, A, \hat{\theta}^{CGM}) \rightarrow \chi_{\frac{K(K+3)}{2} - 1}^2.$$

This statistic can be used to test against a notional value for θ_0 .

A consequence of Theorem 1 is that

$$\sup_{\theta \in \mathcal{T}} \log \frac{g}{g_0}(A; \theta) = \sup_{\theta \in \mathcal{T}} \log \left(\frac{f}{f_0}(Z, A; \theta) \right) + o_P(1),$$

implying that

$$\Lambda_G(A; \hat{\theta}^{ML}) \equiv 2 \log \frac{g}{g_0}(A; \hat{\theta}^{ML}) = \Lambda(Z, A; \hat{\theta}^{CGM}) + o_P(1),$$

so that the Wilks statistic for the GM and CGM estimates have the same asymptotic distribution, enabling tests against a notional \mathcal{S}_{θ_0} using the GM likelihood ratio when Z is latent.

A similar result holds for the Wilks statistic of the variational estimate $\hat{\theta}^{VAR}$, which may be easier to compute. To see this, we observe that since $J(\theta; A) \equiv \max_{q \in \mathcal{D}} \exp[J(q, \theta; A)] \leq g(A; \theta)$, it holds that

$$\frac{J(\theta, A)}{J(\theta_0, A)} \geq \frac{J(\theta; A)}{g(A; \theta_0)},$$

so that Theorem 3 implies

$$\frac{J(\theta, A)}{J(\theta_0, A)} \geq \frac{f}{f_0}(Z, A; \theta)(1 + o_P(1)) + o_P(1).$$

To upper bound the same quantity, we observe that

$$\begin{aligned} \frac{J(\theta, A)}{J(\theta_0, A)} &\leq \frac{g(A; \theta)}{f(Z, A; \theta_0)} \\ &= \frac{g(A; \theta)}{g(A; \theta_0)f(Z, A; \theta_0)g(A; \theta_0)^{-1}} \\ &= \frac{g(A; \theta)}{g(A; \theta_0)(1 + o_P(1))}, \end{aligned}$$

using Lemma 3. Thus, the arguments used to bound Λ_G also imply

$$\Lambda_V(\hat{\theta}^{VAR}) \equiv 2 \log \frac{J(\hat{\theta}^{VAR}, A)}{J(\theta_0, A)} = \Lambda(Z, A; \hat{\theta}^{CGM}) + o_P(1).$$

Parametric bootstrap. The parametric bootstrap is also valid for $\hat{\theta}^{VAR}$. The algorithm is

1. Estimate θ by $\hat{\theta}^{VAR}$
2. Generate B graphs of size n according to the blockmodel with parameter $\hat{\theta}^{VAR}$, producing $(Z_1^*, A_1^*), \dots, (Z_B^*, A_B^*)$.
3. Fit A_1^*, \dots, A_B^* by variational likelihood to get $\hat{\theta}_1^{VAR*}, \dots, \hat{\theta}_B^{VAR*}$.
4. Compute the variance-covariance matrix of these B vectors and use it as an estimate of the truth, or similarly, use the empirical distribution function of the vectors

THEOREM 4. *Under the conditions of Theorem 2, the parametric bootstrap distribution of $\sqrt{n}(\hat{\varpi}^{VAR} - \varpi_0)$ and $\sqrt{n}\lambda(\hat{\nu}^{VAR} - \nu_0)$ converges to the Gaussian limits given by Lemma 2.*

PROOF. Without loss of generality we take $B = \infty$, so that we are asking that when the underlying parameter is $\hat{\theta}^{VAR}$, the random law of

$\sqrt{n}(\hat{\varpi}^{VAR*} - \hat{\varpi}^{VAR})$ and $\sqrt{n\lambda}(\hat{\nu}^{VAR*} - \hat{\nu}^{VAR})$ converges with P_{θ_0} probability tending to 1 to the Gaussian limits of $\sqrt{n}(\hat{\varpi}^{CGM} - \varpi_0)$ and $\sqrt{n\lambda}(\hat{\nu}^{CGM} - \nu_0)$ as generated under θ_0 .

Let $\hat{\varpi}^{CGM*}, \hat{\nu}^{CGM*}$ have the distribution of the CG MLE based on the data that we have generated from $P_{\hat{\theta}^{VAR}}$. By standard exponential theory such as our Lemma 2, we observe that

$$(16) \quad \sqrt{n}(\hat{\varpi}^{CGM*} - \hat{\varpi}^{VAR}) \xrightarrow{P_{\hat{\theta}^{VAR}}} N(0, \Sigma_1)$$

$$(17) \quad \sqrt{n\lambda}(\hat{\nu}^{CGM*} - \hat{\nu}^{VAR}) \xrightarrow{P_{\hat{\theta}^{VAR}}} N(0, \Sigma_2),$$

since the convergence is uniform on contiguous neighborhoods of θ_0 and the mapping $\theta \rightarrow (\Sigma_1(\theta), \Sigma_2(\theta))$ is smooth. As Theorem 3 implies local asymptotic normality, a theorem of Le Cam (Lehmann and Romano, 2005, Corollary 12.3.1) implies that $P_{\hat{\theta}^{VAR}} \triangleleft P_{\theta_0}$ with P_{θ_0} probability tending to 1, where \triangleleft denotes contiguity. As a result, Le Cam's first contiguity lemma (stated below) in conjunction with Theorem 3 implies that

$$\begin{aligned} \sqrt{n}(\hat{\varpi}^{CGM*} - \hat{\varpi}^{VAR*}) &= o_{P_{\hat{\theta}^{VAR}}}(1) \\ \sqrt{n\lambda}(\hat{\nu}^{CGM*} - \hat{\nu}^{VAR*}) &= o_{P_{\hat{\theta}^{VAR}}}(1). \end{aligned}$$

Using this result with Eqs. (16), it follows that

$$\begin{aligned} \sqrt{n}(\hat{\varpi}^{VAR*} - \hat{\varpi}^{VAR}) &\xrightarrow{P_{\hat{\theta}^{VAR}}} N(0, \Sigma_1) \\ \sqrt{n\lambda}(\hat{\nu}^{VAR*} - \hat{\nu}^{VAR}) &\xrightarrow{P_{\hat{\theta}^{VAR}}} N(0, \Sigma_2), \end{aligned}$$

establishing the theorem.

For completeness, we state Le Cam's first contiguity lemma as found in (Van der Vaart, 2000, Lemma 6.4).

LEMMA 4. *Let P_n and Q_n be sequences of probability measures on measurable spaces $(\Omega_n, \mathcal{A}_n)$. Then the following statements are equivalent:*

1. $Q_n \triangleleft P_n$.
2. *If dP_n/dQ_n converges in distribution under Q_n to U along a subsequence, then $P(U > 0) = 1$.*
3. *If dQ_n/dP_n converges in distribution under P_n to V along a subsequence, then $EV = 1$.*
4. *For any statistics $T_n : \Omega_n \mapsto \mathbb{R}^k$: If $T_n \xrightarrow{P_n} 0$, then $T_n \xrightarrow{Q_n} 0$.*

□

5. Conclusions. In this paper, we have studied stochastic block and extended blockmodels, such that the average degree tends to ∞ at least at a polylog rate, and the number of blocks K is fixed. We have shown:

1. Subject to identifiability restrictions, methods of estimation and parameter testing on maximum likelihood have exactly the same behavior as the same methods when the block identities are observed, such that an easily analyzed exponential family model is in force. The approach uses the methods of (Bickel and Chen, 2009) slightly corrected. Unfortunately, computation of the likelihood is as difficult as the NP-complete computation of modularities, which also yield parameter estimates that are usable in the same way.
2. We also show that the variational likelihood, introduced in this context by (Daudin, Picard and Robin, 2008), has the same properties as the ordinary likelihood under these conditions; hence, the procedures discussed above but applied to the variational likelihood behave in the same way. The variational likelihood can be computed in $\mathcal{O}(n^3)$ operations, making this a more attractive method.

These results easily imply that classical optimality properties of these procedures, such as achievement of the information bound, hold.

Discussion. A number of major issues still need to be resolved. Here are some:

1. Since the log likelihoods studied are highly non-concave, selection of starting points for optimization seems critical. The most promising approaches from both a theoretical and computational point of view are spectral clustering approaches (Rohe, Chatterjee and Yu, 2011; Chaudhuri, Chung and Tsias, 2012).
2. Blockmodels play the role of histogram approximations for more complex models of the type considered in (Bickel, Chen and Levina, 2011), and if observed covariates are added for models such as those of (Hoff, Raftery and Handcock, 2002). This implies permitting the number of blocks K to increase, which makes perfect classification and classical rates of parameter estimation unlikely. Issues of model selection and regularization come to the fore. Some work of this type has been done in (Rohe, Chatterjee and Yu, 2011; Choi, Wolfe and Airolidi, 2012; Chatterjee, 2012), but statistical approximation goals are unclear.
3. We have indicated that our results for (ϖ, ν) -parameterized blockmodels also apply to submodels which are sufficiently smoothly parameterizable. It seems likely that our methods can also apply to models where there are covariates associated to vertices or edges.

Appendix: Proof of Theorem 1. We adopt the convention of (Bickel and Chen, 2009) and let \mathbf{c} denote Z . Recall that $S = H/\rho_n$. Let $\mu_n = n^2\rho_n$. Let $L = \sum_{i \neq j} A_{ij}$. For any $\mathbf{e} \in [K]^n$, let

$$n_{ab}(\mathbf{e}) = \sum_{i=1}^n \sum_{j \neq i}^n 1\{\mathbf{e}_i = a, \mathbf{e}_j = b\}, \quad n_a(\mathbf{e}) = \sum_{i=1}^n 1\{\mathbf{e}_i = a\},$$

$$\pi_a(\mathbf{e}) = n_a(\mathbf{e})/n, \quad O_{ab}(A, \mathbf{e}) = \sum_{i=1}^n \sum_{j \neq i}^n 1\{\mathbf{e}_i = a, \mathbf{e}_j = b\} A_{ij}.$$

Let $|\mathbf{e} - \mathbf{c}|$ denote $\sum_{i=1}^n 1\{\mathbf{e}_i \neq \mathbf{c}_i\}$. Given \mathbf{e} , define $\bar{\mathbf{e}} = \arg \min_{\mathbf{e}' \in \mathcal{S}_{\mathbf{e}}} |\mathbf{e}' - \mathbf{c}|$. Define the confusion matrix $R \in [0, 1]^{K \times K}$ by

$$[R(\mathbf{e}, \mathbf{c})](a, a') = \frac{1}{n} \sum_i 1\{\mathbf{e}_i = a, \mathbf{c}_i = a'\}.$$

We observe that for fixed \mathbf{c} , R is constrained to the set $\mathcal{R} = \{R : R \geq 0, R^T \mathbf{1} = \pi(\mathbf{c})\}$. Let $RSR^T \equiv (RSR^T)(\mathbf{e})$ abbreviate $R(\mathbf{c}, \mathbf{e})SR^T(\mathbf{c}, \mathbf{e})$. Let $X(\mathbf{e}) = \mu_n^{-1}O(A, \mathbf{e}) - RSR^T$.

Let f_n denote the full data likelihood of the stochastic blockmodel,

$$f_n(A, \mathbf{e}; \theta) = \prod_{i=1}^n \pi_{\theta}(z_i) \prod_{i < j} H_{\theta}(\mathbf{e}_i, \mathbf{e}_j)^{A_{ij}} (1 - H_{\theta}(\mathbf{e}_i, \mathbf{e}_j))^{1 - A_{ij}}.$$

Let Q_n denote the likelihood modularity (Bickel and Chen, 2009), defined as $Q_n(A, \mathbf{e}) = \sup_{\theta} \log f_n(A, \mathbf{e}; \theta)$. We observe that Q_n equals

$$\begin{aligned} Q_n(A, \mathbf{e}) &= \sum_{i=1}^n n_a \log \frac{n_a}{n} \\ &\quad + \frac{1}{2} \sum_{a=1}^K \sum_{b=1}^K \left[O_{ab} \log \frac{O_{ab}}{n_{ab}} + \left(n_{ab} - \frac{O_{ab}}{n_{ab}} \right) \log \left(1 - \frac{O_{ab}}{n_{ab}} \right) \right]. \end{aligned}$$

For $\rho_n \rightarrow 0$, it is shown in (Bickel and Chen, 2009) that

$$Q_n(A, \mathbf{e}) = \mu_n \left(F \left(\frac{O(A, \mathbf{e})}{\mu_n}, \pi(\mathbf{e}) \right) + \frac{L}{\mu_n} \log \rho_n + o_P(1) \right),$$

where the function F is given by

$$\begin{aligned} F(M, t) &= \sum_{a,b} t_a t_b \tau \left(\frac{M_{ab}}{t_a t_b} \right) \\ \tau(x) &= x \log x - x \end{aligned}$$

for $M \in \mathbb{R}^{K \times K}$ and t in the K -simplex.

The result of (Bickel and Chen, 2009) establishes that the following properties hold for F_n (also see (Zhao, Levina and Zhu, 2013) for a reworked derivation):

1. The function $\mathbf{e} \mapsto F((RSR^T)(\mathbf{e}), \pi(\mathbf{e}))$ is maximized by any $\mathbf{e} \in \mathcal{S}_{\mathbf{c}}$.
2. The function F is uniformly continuous if M and t are restricted to any subset bounded away from 0.
3. Let $G(R, S) = F(RSR^T, R^T \mathbf{1})$. Given $(\pi, S) \in \mathcal{T}$, it holds for all $R \in \{R \geq 0, R^T \mathbf{1} = \pi\}$ that

$$\left. \frac{\partial G((1 - \epsilon) \text{diag}(\pi) + \epsilon R, S)}{\partial \epsilon} \right|_{\epsilon=0+} < -C < 0$$

4. The directional derivatives

$$\left. \frac{\partial^2 F}{\partial \epsilon^2} (M_0 + \epsilon(M_1 - M_0), t_0 + \epsilon(t_1 - t_0)) \right|_{\epsilon=0+}$$

are continuous in (M_1, t_1) for all (M_0, t_0) in a neighborhood of $(\text{diag}(\pi)S \text{diag}(\pi), \pi)$.

We will use an Bernstein inequality result, similar to that shown in (Bickel and Chen, 2009).

LEMMA 5. *Let $C_S = \max_{ab} S_{ab}$.*

$$(18) \quad \mathbb{P} \left(\max_{\mathbf{e}} \|X(\mathbf{e})\|_{\infty} \geq \epsilon \right) \leq 2K^{n+2} \exp \left(-\frac{1}{4} \epsilon^2 \mu_n \right)$$

for $\epsilon \leq 3$, and

$$(19) \quad \mathbb{P} \left(\max_{\mathbf{e}: |\mathbf{e}-\mathbf{c}| \leq m} \|X(\mathbf{e}) - X(\mathbf{c})\|_{\infty} \geq \epsilon \right) \leq 2 \binom{n}{m} K^{m+2} \exp \left(-\frac{n}{m(8C_S + 2)} \epsilon^2 \mu_n \right)$$

for $\epsilon \leq \frac{3m}{n}$.

PROOF OF LEMMA 5. $\mu_n X_{ab}$ is a sum of independent zero mean random variables bounded by 1. Thus by a Bernstein inequality,

$$\mathbb{P} (|\mu_n X_{ab}(\mathbf{e})| \geq \epsilon \mu_n) \leq \exp \left(\frac{-\epsilon^2 \mu_n^2}{2(\text{Var}(\mu_n X_{ab}) + \epsilon \mu_n / 3)} \right).$$

We may bound $\text{Var}(\mu_n X_{ab}) \leq \mu_n$ and $\epsilon \leq 3$ to yield for fixed a, b, \mathbf{e} that

$$\mathbb{P} (|X_{ab}(\mathbf{e})| \geq \epsilon \mu_n) \leq \exp \left(\frac{-\epsilon^2 \mu_n}{4} \right).$$

A union bound establishes (18).

Similarly, $\mu_n(X_{ab}(\mathbf{c}) - X_{ab}(\mathbf{e}))$ is a sum of independent zero mean random variables bounded by 1. Thus

$$\mathbb{P}(|\mu_n(X_{ab}(\mathbf{e}) - X_{ab}(\mathbf{c}))| \geq \epsilon\mu_n) \leq \exp\left(\frac{-\epsilon^2\mu_n^2}{2(\text{Var}\mu_n((X_{ab}(\mathbf{c}) - X_{ab}(\mathbf{e}))) + \epsilon\mu_n/3)}\right).$$

We may bound $\text{Var}(\mu_n(X_{ab}(\mathbf{e}) - X_{ab}(\mathbf{c}))) \leq 4mnC_S\rho_n = 4C_S\mu_n m/n$ and $\epsilon \leq 3m/n$ to yield for fixed a, b, \mathbf{e} that

$$\mathbb{P}(|\mu_n(X_{ab}(\mathbf{e}) - X_{ab}(\mathbf{c}))| \geq \epsilon\mu_n) \leq \exp\left(\frac{-\epsilon^2\mu_n}{(8C_S + 2)m/n}\right).$$

A union bound establishes (19), where we use that $|\{\mathbf{e} : |\mathbf{e} - \mathbf{c}| \leq m\}| \leq \binom{n}{m}K^m$ for fixed \mathbf{c} . \square

PROOF OF THEOREM 1. The proof can be separated into four parts.

Part 1: \mathbf{e} for which F is small. Here we show, for some $\delta_n \rightarrow 0$, that $F(O(\mathbf{e})/\mu, \pi(\mathbf{e}))$ is suboptimal by at least $\delta_n/2$ for all \mathbf{e} in a set E_{δ_n} . This will imply that $\sum_{\mathbf{e} \in E_{\delta_n}} \sup_{\theta} f(A, \mathbf{e}; \theta) = o_P(1) \sup_{\theta} f(A, \mathbf{c}; \theta)$.

By (18), $\mu_n^{-1}O(\mathbf{e}) \xrightarrow{P} RSR^T(\mathbf{e})$ uniformly over \mathbf{e} ; hence, by continuity of F there exists $\delta_n \rightarrow 0$ such that

$$\mathbb{P}\left(\max_{\mathbf{e}} \left| F\left(\frac{O(\mathbf{e})}{\mu_n}, \pi(\mathbf{e})\right) - F(RSR^T(\mathbf{e}), \pi(\mathbf{e})) \right| \geq \delta_n/2\right) = o(1).$$

As a result, given the sets

$$E_{\delta_n} = \{\mathbf{e} : |F((RSR^T)(\mathbf{e}), \pi(\mathbf{e})) - F((RSR^T)(\mathbf{c}), \pi(\mathbf{c}))| \geq \delta_n\},$$

it holds for all $\mathbf{e} \in E_{\delta_n}$ that $F\left(\frac{O(\mathbf{e})}{\mu_n}, \pi(\mathbf{e})\right) \leq F(RSR^T(\mathbf{c}), \pi(\mathbf{c})) - \delta_n/2 + o_P(\delta_n)$. We may choose δ_n to additionally satisfy

$$\begin{aligned} \sum_{\mathbf{e} \in E_{\delta_n}} e^{\mu_n F((RSR^T)(\mathbf{e}), \pi(\mathbf{e}))} &\leq \sum_{\mathbf{e} \in E_{\delta_n}} e^{\mu_n (F((RSR^T)(\mathbf{c}), \pi(\mathbf{c})) + o_P(\delta_n) - \delta_n/2)} \\ &\leq e^{\mu_n F((RSR^T)(\mathbf{c}), \pi(\mathbf{c}))} e^{-\mu_n(1+o_P(1))\delta_n/2} K^n \\ (20) \qquad \qquad \qquad &= e^{\mu_n F((RSR^T)(\mathbf{c}), \pi(\mathbf{c}))} o_P(1), \end{aligned}$$

where we require $\delta_n \rightarrow 0$ slowly enough that $\mu_n\delta_n \gg n$.

Part 2: A concentration inequality. We wish to show for $\mathbf{e} \notin E_{\delta_n}$ a result similar to Part 1. However, as some \mathbf{e} will be very close to \mathbf{c} , we must bound the suboptimality of $F(O(\mathbf{e})/\mu, \pi(\mathbf{e}))$ more carefully.

By (19), it holds that

$$\mathbb{P}\left(\max_{\mathbf{e}:|\mathbf{e}-\mathbf{c}|=m} \|X(\mathbf{e}) - X(\mathbf{c})\|_\infty \geq \epsilon \frac{m}{n}\right) \leq 2n^m K^{m+2} \exp\left(-\frac{m}{n(8C_S+2)} \epsilon^2 \mu_n\right)$$

It follows that we may choose $\epsilon \rightarrow 0$ such that

$$\begin{aligned} \mathbb{P}\left(\max_{\mathbf{e} \notin \mathcal{S}_c} \frac{\|X(\bar{\mathbf{e}}) - X(\mathbf{c})\|_\infty}{|\bar{\mathbf{e}} - \mathbf{c}|/n} \geq \epsilon\right) &\leq \sum_{m=1}^n \mathbb{P}\left(\max_{\mathbf{e}:|\mathbf{e}-\mathbf{c}|=m, \mathbf{e}=\bar{\mathbf{e}}} \frac{\|X(\mathbf{e}) - X(\mathbf{c})\|_\infty}{m/n} \geq \epsilon\right) \\ &\leq \sum_{m=1}^n 2K^K n^m K^{m+2} \exp\left(-\frac{m}{n(8C_S+2)} \epsilon^2 \mu_n\right) \\ &\leq \sum_{m=1}^n 2K^{K+2} e^{m\left(\log n + \log K - \frac{\epsilon^2 \mu_n}{n(8C_S+2)}\right)} \\ &= o(1), \end{aligned}$$

where the final equality holds because $\mu_n/n \gg \log n$, so that we may choose $\epsilon \rightarrow 0$ such that $\epsilon^2 \mu_n/n \gg \log n$. It follows that

$$(21) \quad \max_{\mathbf{e} \notin \mathcal{S}_c} \frac{\|X(\bar{\mathbf{e}}) - X(\mathbf{c})\|_\infty}{|\bar{\mathbf{e}} - \mathbf{c}|/n} = o_P(1).$$

Part 3: \mathbf{e} when F is large. Here we bound the suboptimality of $F(O(\mathbf{e})/\mu, \pi(\mathbf{e}))$ in similar fashion to Part 1.

Recall $F((RSR^T)(\mathbf{e}), \pi(\mathbf{e})) = G(R(\mathbf{e}), S)$ with $R(\mathbf{e}) \in \{R \geq 0, R^T \mathbf{1} = \pi(\mathbf{c})\}$. Let $h(\mathbf{e})$ abbreviate $RSR^T(\mathbf{e}) - RSR^T(\mathbf{c})$. Property 3 implies that for all \mathbf{e} ,

$$\left. \frac{\partial}{\partial \epsilon} F\left(RSR^T(\mathbf{c}) + \epsilon h(\mathbf{e}), (RSR^T(\mathbf{c}) + \epsilon h(\mathbf{e}))^T \mathbf{1}\right) \right|_{\epsilon=0^+} < -\Omega_P(1),$$

where $a_n = \Omega_P(b_n)$ denotes that a_n is bounded below (in probability) by b_n times a constant factor. As $\delta_n \rightarrow 0$, this implies for all $\mathbf{e} \notin E_{\delta_n}$,

$$F\left((RSR^T)(\mathbf{c}), \pi(\mathbf{c})\right) - F\left((RSR^T)(\mathbf{e}), \pi(\mathbf{e})\right) \geq \frac{1}{n} \Omega(|\bar{\mathbf{e}} - \mathbf{c}|).$$

As $(O(\mathbf{c})/\mu_n, \pi(\mathbf{c}))$ converges in probability to $(RSR^T(\mathbf{c}), \pi(\mathbf{c}))$, properties 3 and 4 together imply for all \mathbf{e} ,

$$\left. \frac{\partial}{\partial \epsilon} F\left(\frac{O(\mathbf{c})}{\mu_n} + \epsilon h(\mathbf{e}), \left(\frac{O(\mathbf{c})}{\mu_n} + \epsilon h(\mathbf{e})\right)^T \mathbf{1}\right) \right|_{\epsilon=0^+} < -\Omega_P(1),$$

and thus for $\mathbf{e} \notin E_{\delta_n}$,

$$F\left(\frac{O(\mathbf{c})}{\mu_n}, \pi(\mathbf{c})\right) - F\left(\frac{O(\mathbf{c})}{\mu_n} + h(\bar{\mathbf{e}}), \pi(\bar{\mathbf{e}})\right) \geq \frac{1}{n} \Omega_P(|\bar{\mathbf{e}} - \mathbf{c}|),$$

and hence also that

$$(22) \quad F\left(\frac{O(\mathbf{c})}{\mu_n}, \pi(\mathbf{c})\right) - F\left(\frac{O(\mathbf{c})}{\mu_n} + h(\bar{\mathbf{e}})(1 + o_P(1)), \pi(\bar{\mathbf{e}})\right) \geq \frac{1}{n} \Omega_P(|\bar{\mathbf{e}} - \mathbf{c}|).$$

It can be seen that $h(\bar{\mathbf{e}}) \equiv RSR^T(\bar{\mathbf{e}}) - RSR^T(\mathbf{c}) = \Omega(\|\bar{\mathbf{e}} - \mathbf{c}\|/n)$. As a result, by (21), for all $\mathbf{e} \notin E_{\delta_n}$,

$$\begin{aligned} \left\| \frac{O(\bar{\mathbf{e}})}{\mu_n} - \frac{O(\mathbf{c})}{\mu_n} - (RSR^T(\bar{\mathbf{e}}) - RSR^T(\mathbf{c})) \right\|_{\infty} &= o_P(|\bar{\mathbf{e}} - \bar{\mathbf{c}}|/n) \\ &= o_P(RSR^T(\bar{\mathbf{e}}) - RSR^T(\mathbf{c})), \end{aligned}$$

and hence manipulation yields for all $\mathbf{e} \notin E_{\delta_n}$,

$$\frac{O(\bar{\mathbf{e}})}{\mu_n} - \frac{O(\mathbf{c})}{\mu_n} = (RSR^T(\bar{\mathbf{e}}) - RSR^T(\mathbf{c})) (1 + o_P(1)),$$

where the $o_P(1)$ term is uniform over \mathbf{e} . As a result, it follows from (22) that for $\mathbf{e} \notin E_{\delta_n}$,

$$F\left(\frac{O(\mathbf{c})}{\mu_n}, \pi(\mathbf{c})\right) - F\left(\frac{O(\mathbf{e})}{\mu_n}, \pi(\bar{\mathbf{e}})\right) \geq \frac{1}{n} \Omega_P(|\bar{\mathbf{e}} - \mathbf{c}|),$$

where the $\Omega_P(|\bar{\mathbf{e}} - \mathbf{c}|)$ is uniform over \mathbf{e} . It follows that

$$\begin{aligned} \sum_{\mathbf{e} \notin E_{\delta_n}, \mathbf{e} \notin \mathcal{S}_{\mathbf{c}}} e^{\mu_n F\left(\frac{O(\mathbf{e})}{\mu_n}, \pi(\mathbf{e})\right)} &\leq \sum_{m=1}^n \sum_{\mathbf{e}: |\bar{\mathbf{e}} - \mathbf{c}| = m} e^{\mu_n F\left(\frac{O(\mathbf{e})}{\mu_n}, \pi(\mathbf{e})\right)} \\ &= \sum_{m=1}^n \sum_{\mathbf{e}: |\bar{\mathbf{e}} - \mathbf{c}| = m} e^{\mu_n \left[F\left(\frac{O(\mathbf{c})}{\mu_n}, \pi(\mathbf{c})\right) + F\left(\frac{O(\mathbf{e})}{\mu_n}, \pi(\mathbf{e})\right) - F\left(\frac{O(\mathbf{c})}{\mu_n}, \pi(\mathbf{c})\right) \right]} \\ &\leq \sum_{m=1}^n \sum_{\mathbf{e}: |\bar{\mathbf{e}} - \mathbf{c}| = m} e^{\mu_n F\left(\frac{O(\mathbf{c})}{\mu_n}, \pi(\mathbf{c})\right)} e^{-\frac{\mu_n}{n} \Omega_P(m)} \\ &\leq \sum_{m=1}^n e^{\mu_n F\left(\frac{O(\mathbf{c})}{\mu_n}, \pi(\mathbf{c})\right)} K^K n^m K^m e^{-\frac{\mu_n}{n} \Omega_P(m)} \\ &\leq \sum_{m=1}^n e^{\mu_n F\left(\frac{O(\mathbf{c})}{\mu_n}, \pi(\mathbf{c})\right)} K^K e^{m(\log n + \log K - \Omega_P(\mu_n/n))} \\ (23) \quad &= e^{\mu_n F\left(\frac{O(\mathbf{c})}{\mu_n}, \pi(\mathbf{c})\right)} o_P(1) \end{aligned}$$

Part 4: Putting the parts together. Combining (23) and (20) yields that

$$(24) \quad \sum_{\mathbf{e} \notin \mathcal{S}_{\mathbf{c}}} e^{\mu_n F\left(\frac{\rho(\mathbf{e})}{\mu}, \pi(\mathbf{e})\right)} \leq e^{\mu_n F((RSR^T)(\mathbf{c}), \pi(\mathbf{c}))} o_P(1).$$

Since $\frac{f}{f_0}(A, \mathbf{c}; \theta)$ is unimodal in θ , it holds that if $\frac{f}{f_0}(A, \mathbf{c}; \theta) \neq o_P(1)$, then $\theta \rightarrow \theta_0$, and hence by Lemma 1, $\frac{f}{f_0}(A, \mathbf{c}; \sigma(\theta)) = o_P(1)$ for any non-identity permutation σ . It follows that

$$(25) \quad \begin{aligned} \sum_{\mathbf{c}' \in \mathcal{S}_{\mathbf{c}}} f(A, \mathbf{c}'; \theta) &= \sum_{\theta' \in \mathcal{S}_{\theta}} f(A, \mathbf{c}; \theta') \\ &= \max_{\theta' \in \mathcal{S}_{\theta}} f(A, \mathbf{c}; \theta') (1 + o_P(1)) \end{aligned}$$

Combining (24) and (25) yields

$$\sum_{\mathbf{e} \neq \mathbf{c}} \sup_{\theta} f(A, \mathbf{e}; \theta) = \left(\sup_{\theta} f(A, \mathbf{c}; \theta) \right) o_P(1).$$

Letting F_0 abbreviate $\sup_{\theta} f(A, \mathbf{c}; \theta)$, and using $g(A; \theta) = \sum_{\mathbf{e}} f(A, \mathbf{e}; \theta)$,

$$\begin{aligned} \frac{g(A; \theta)}{g(A; \theta_0)} &= \frac{\sum_{\mathbf{e}} f(A, \mathbf{e}; \theta)}{\sum_{\mathbf{e}} f(A, \mathbf{e}; \theta_0)} \\ &= \frac{f(A, \mathbf{c}; \theta)}{f(A, \mathbf{c}; \theta_0) + \sum_{\mathbf{e} \neq \mathbf{c}} f(A, \mathbf{e}; \theta_0)} + \frac{\sum_{\mathbf{e} \neq \mathbf{c}} f(A, \mathbf{e}; \theta)}{f(A, \mathbf{c}; \theta_0) + \sum_{\mathbf{e} \neq \mathbf{c}} f(A, \mathbf{e}; \theta_0)} \\ &= \frac{f(A, \mathbf{c}; \theta)}{f(A, \mathbf{c}; \theta_0) + F_0 o_P(1)} + \frac{F_0 o_P(1)}{f(A, \mathbf{c}; \theta_0) + F_0 o_P(1)}, \end{aligned}$$

where in the last equality we have used the fact that for all θ

$$0 \leq \sum_{\mathbf{e} \neq \mathbf{c}} f(A, \mathbf{e}; \theta) \leq \sum_{\mathbf{e} \neq \mathbf{c}} \sup_{\theta} f(A, \mathbf{e}; \theta) = F_0 o_P(1).$$

Since F_0 equals the likelihood of the MLE under the CGM model, it holds that $\frac{F_0}{f(A, \mathbf{c}; \theta_0)}$ converges in distribution, and hence $\frac{F_0}{f(A, \mathbf{c}; \theta_0)} o_P(1) = o_P(1)$. We may therefore substitute $F_0 o_P(1) = f(A, \mathbf{c}; \theta_0) \frac{F_0}{f(A, \mathbf{c}; \theta_0)} o_P(1) = f(A, \mathbf{c}; \theta_0) o_P(1)$ to yield

$$\begin{aligned} \frac{g(A; \theta)}{g(A; \theta_0)} &= \frac{f(A, \mathbf{c}; \theta)}{f(A, \mathbf{c}; \theta_0)(1 + o_P(1))} + \frac{f(A, \mathbf{c}; \theta_0) o_P(1)}{f(A, \mathbf{c}; \theta_0)(1 + o_P(1))} \\ &= \frac{f}{f_0}(A, \mathbf{c}; \theta)(1 + o_P(1)) + o_P(1), \end{aligned}$$

which proves the theorem. □

Acknowledgements. We would like to thank the reviewers for their help in fixing an earlier version of the paper.

References.

- BICKEL, P. J. and CHEN, A. (2009). A nonparametric view of network models and Newman–Girvan and other modularities. *Proc. Natl. Acad. Sci. USA* **106** 21068.
- BICKEL, P. J., CHEN, A. and LEVINA, E. (2011). The method of moments and degree distributions for network models. *Ann. Statist.* **39** 38–59.
- BOLLOBÁS, B., JANSON, S. and RIORDAN, O. (2007). The phase transition in inhomogeneous random graphs. *Random Structures Algorithms* **31** 3–122.
- CELISSE, A., DAUDIN, J. J. and PIERRE, L. (2011). Consistency of maximum-likelihood and variational estimators in the Stochastic Block Model. *Arxiv preprint arXiv:1105.3288*.
- CHANNAROND, A., DAUDIN, J. J. and ROBIN, S. (2011). Classification and estimation in the Stochastic Block Model based on the empirical degrees. *Arxiv preprint arXiv:1110.6517*.
- CHATTERJEE, S. (2012). Matrix estimation by Universal Singular Value Thresholding. *arXiv preprint arXiv:1212.1247*.
- CHAUDHURI, K., CHUNG, F. and TSIATAS, A. (2012). Spectral Clustering of Graphs with General Degrees in the Extended Planted Partition Model. *J. Mach. Learn. Res.* 1–23.
- CHOI, D. S., WOLFE, P. J. and AIROLDI, E. M. (2012). Stochastic blockmodels with growing number of classes. *Biometrika* **99** 273–284.
- COJA-OGHLAN, A. and LANKA, A. (2008). Partitioning Random Graphs with General Degree Distributions. In *Fifth IFIP International Conference On Theoretical Computer Science–TCS 2008* 127–141.
- DAUDIN, J. J., PICARD, F. and ROBIN, S. (2008). A mixture model for random graphs. *Statist. Comput.* **18** 173–183.
- HANDCOCK, M. S., RAFTERY, A. E. and TANTRUM, J. M. (2007). Model-based clustering for social networks. *J. Roy. Statist. Soc. Ser. A* **170** 301–354.
- HOFF, P. D., RAFTERY, A. E. and HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *J. Am. Statist. Assoc.* **97** 1090–1098.
- KARRER, B. and NEWMAN, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Phys. Rev. E* **83** 016107.
- LATOUCHE, P., BIRMELÉ, E. and AMBROISE, C. (2011). Overlapping stochastic block models with application to the french political blogosphere. *Ann. of Appl. Statist.* **5** 309–336.
- LAZER, D., PENTLAND, A. S., ADAMIC, L., ARAL, S., BARABASI, A. L., BREWER, D., CHRISTAKIS, N., CONTRACTOR, N., FOWLER, J., GUTMANN, M. et al. (2009). Life in the network: the coming age of computational social science. *Science* **323** 721.
- LE CAM, L. and YANG, G. L. (1988). On the preservation of local asymptotic normality under information loss. *Ann. Statist.* **16** 483–520.
- LEHMANN, E. L. and ROMANO, J. P. (2005). *Testing statistical hypotheses*. Springer Verlag.
- PROULX, S. R., PROMISLOW, D. E. L. and PHILLIPS, P. C. (2005). Network thinking in ecology and evolution. *Trends in Ecology & Evolution* **20** 345–353.
- ROHE, K., CHATTERJEE, S. and YU, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.* **39** 1878–1915.
- ROHE, K., QIN, T. and FAN, H. (2012). The Highest Dimensional Stochastic Blockmodel with a Regularized Estimator. *arXiv preprint arXiv:1206.2380*.

- SNIJDERS, T. A. B. and NOWICKI, K. (1997). Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *Journal of Classification* **14** 75–100.
- VAN DER VAART, A. W. (2000). *Asymptotic statistics* **3**. Cambridge university press.
- ZHAO, Y., LEVINA, E. and ZHU, J. (2013). Consistency of community detection in networks under degree-corrected stochastic block models. *Ann. Statist.* To appear.

DEPARTMENT OF STATISTICS
 367 EVANS HALL
 BERKELEY, CA 94710
 E-MAIL: bickel@stat.berkeley.edu

HEINZ COLLEGE OF PUBLIC POLICY AND MANAGEMENT
 CARNEGIE MELLON UNIVERSITY
 PITTSBURGH, PA 15213
 E-MAIL: davidch@andrew.cmu.edu

DEPARTMENT OF APPLIED MATHEMATICS
 XI'AN JIAOTONG UNIVERSITY
 XI'AN 710049
 CHINA
 E-MAIL: xiangyuchang@gmail.com

DEPARTMENT OF MATHEMATICS
 NORTHWEST UNIVERSITY
 XI'AN 710069
 CHINA
 E-MAIL: zhanghai@nwu.edu.cn