

ROLE OF NORMALIZATION IN SPECTRAL CLUSTERING FOR STOCHASTIC BLOCKMODELS

BY PURNAMRITA SARKAR¹ AND PETER J. BICKEL

University of Texas, Austin and University of California, Berkeley

Spectral clustering is a technique that clusters elements using the top few eigenvectors of their (possibly normalized) similarity matrix. The quality of spectral clustering is closely tied to the convergence properties of these principal eigenvectors. This rate of convergence has been shown to be identical for both the normalized and unnormalized variants in recent random matrix theory literature. However, normalization for spectral clustering is commonly believed to be beneficial [*Stat. Comput.* **17** (2007) 395–416]. Indeed, our experiments show that normalization improves prediction accuracy. In this paper, for the popular stochastic blockmodel, we theoretically show that normalization shrinks the spread of points in a class by a constant fraction under a broad parameter regime. As a byproduct of our work, we also obtain sharp deviation bounds of empirical principal eigenvalues of graphs generated from a stochastic blockmodel.

1. Introduction. Networks appear in many real-world problems. Any dataset of co-occurrences or relationships between pairs of entities can be represented as a network. For example, the Netflix data can be thought of as a giant bipartite network between customers and movies, where edges are formed via ratings. Facebook is a network of friends, where edges represent who knows whom. Weblogs link to other blogs and give rise to blog networks. Networks can also be implicit; for example, in machine learning they are often built by computing pairwise similarities between entities.

Many problems in machine learning and statistics are centered around community detection. Viral marketing functions by understanding how information propagates through friendship networks, and community detection is key to this. Link farms in the World Wide Web are basically malicious tightly connected clusters of webpages which exploit web-search algorithms to increase their rank. These need to be identified and removed so that search results are authentic and do not mislead users.

Spectral clustering [7, 9] is a widely used network clustering algorithm. The main idea is to first represent the i th entity by a k -dimensional vector obtained by

Received October 2013; revised November 2014.

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

MSC2010 subject classifications. Primary 62H30; secondary 60B20.

Key words and phrases. Stochastic blockmodel, spectral clustering, networks, normalization, asymptotic analysis.

concatenating the i th elements of the top k eigenvectors of a graph, and then cluster this lower-dimensional representation. We will refer to this as the spectral representation. Due to its computational ease and competitive performance, emerging application areas of spectral clustering range widely from parallel computing [12], CAD (computer aided design) [11], parallel sparse matrix factorization [19] to image segmentation [23], general clustering problems in machine learning [17] and most recently, to fitting and classification using network blockmodels [21, 24].

A stochastic blockmodel is a widely used generative model for networks with labeled nodes [13, 21]. It assigns nodes to k different classes and forces all nodes in the same class to be stochastically equivalent. For example, in a two-class stochastic blockmodel, any pair of nodes belonging to different classes link with probability γ_n (a deterministic quantity possibly dependent on the size of the graph, i.e., n), whereas any pair belonging to class one (two) link with probability α_n (β_n).

Recently the consistency properties of spectral clustering in the context of stochastic blockmodels have attracted a significant amount of attention. Rohe, Chatterjee and Yu [21] showed that, under general conditions, for a sequence of normalized graphs with growing size generated from a stochastic blockmodels, spectral clustering yields the correct clustering in the limit. In a subsequent paper, Sussman et al. [24] showed that an analogous statement holds for an unnormalized sequence of graphs. For finite k , the above results can also be obtained using direct applications of results from [18].

This prior theoretical work does not distinguish between normalized and unnormalized spectral clustering, and hence cannot be used to support the common practice of normalizing matrices for spectral clustering. In this paper, we present both theoretical arguments and empirical results to give a quantitative argument showing that normalization improves the quality of clustering. While existing work [21, 24] bounds the classification accuracy, we do not take this route, since upper bounds can not be used to compare two methods. Instead, we focus on the variance within a class under the spectral representation using the top k eigenvectors. In this representation, by virtue of stochastic equivalence, points are identically distributed around their respective class centers. Hence the empirical variance can be computed using the average squared distance of points from their class center.

In this setting, the distance between the class centers can be thought of as bias; we show that this distance approaches the same deterministic quantity with or without normalization. Surprisingly, we also prove that normalization reduces the variance of points in a class by a constant fraction for a large parameter regime. So normalization does not change the bias, but shrinks the variance asymptotically. However, our results also indicate that the variance of points in a class increases as the graph gets sparser; hence methods which reduce the within-class variance are desired.

A simple consequence of our result is that in the completely associative case ($\gamma_n = 0$) as well as the completely dissociative case ($\alpha_n = \beta_n = 0$), the variance of the spectral embedding within a class is asymptotically four times less when the

matrix is normalized. While the completely associative case is on its own uninteresting, we build the proof of the general case using similar ideas and techniques.

Our results indicate that normalization has a clear edge when the parameters are close to the completely associative or completely dissociative settings. These seemingly easy to cluster regimes can be relatively difficult in sparse networks. Of course, as n grows, both methods have enough data to distinguish between the clusters and behave similarly. But for small and sparse graphs, it is indeed an important regime.

Sussman et al. [24] present a parameter setting where normalization is shown to hurt classification accuracy empirically. We show that this is but a partial picture; and in fact there is a large parameter regime where spectral clustering with normalized matrices yields tighter and hence better clusters.

Using quantifiable link prediction experiments on real world graphs and classifications tasks in labeled simulated graphs, we show that normalization leads to better classification accuracies for the regime dictated by our theory, and yields higher link prediction accuracy on sparse real world graphs.

We conclude the introduction with a word of caution. Our asymptotic theory is valid in the degree regime where networks are connected with probability approaching one. However, finite sparse networks can have disconnected or weakly connected small components, in the presence of which, the normalized method returns uninformative principal eigenvectors with support on the small components. This makes classification worse compared to the unnormalized method, whose principal eigenvectors are informative in spite of having high variance. Hence, our asymptotic results should be used only as a guidance for finite n , not as a hard rule. We deal with this problem by removing low degree nodes and performing experiments on the giant component of the remaining network.

2. Preliminaries and import of previous work. In this paper we will only work with two class blockmodels. Given a binary $n \times 2$ class membership matrix Z , the edges of the network are simply outcomes of $\binom{n}{2}$ independent Bernoulli coin flips. The stochastic blockmodel ensures stochastic equivalence of nodes within the same block; that is, all nodes within the same block have identical probability of linking with other nodes in the graph.

Thus the conditional expectation matrix P of the adjacency matrix A can be described by three probabilities, namely $\alpha_n, \beta_n, \gamma_n$; where α_n and β_n denote the probabilities of connecting within the first and second classes (C_1 and C_2), respectively, and γ_n denotes the probability of connecting across two classes. All statements in this paper are conditioned on $\alpha_n, \beta_n, \gamma_n$ and Z .

DEFINITION 2.1 (A stochastic blockmodel). Let $Z \in \{0, 1\}^{n \times 2}$ be a fixed and unknown matrix of class memberships such that every row has exactly one 1, and the first and second columns have $n\pi$ and $n(1 - \pi)$ ones, respectively. A stochastic blockmodel with parameters $(\alpha_n, \beta_n, \gamma_n, Z)$ generates symmetric graphs with

adjacency matrix A such that, $P(A(i, i) = 1) = 0, \forall i$. For $i > j, A_{ij} = A_{ji}$ are independent with $P(A_{ij} = 1|Z) = P_{ij}$, where P is symmetric with $P_{ij} = \alpha_n$ for $i, j \in C_1, \gamma_n$ for $i \in C_1, j \in C_2$ and β_n for $i, j \in C_2$.

For ease of exposition we will assume that the rows and columns of A are permuted such that all elements of the same class are grouped together. We have $P := E[A|Z]$. Clearly, P is a blockwise constant matrix with zero diagonal by construction.

We use a parametrization similar to that in [3] to allow for decaying edge probabilities as n grows. Formally α_n, β_n and γ_n are proportional to a common rate variable ρ_n where $\rho_n \rightarrow 0$ as $n \rightarrow \infty$, forcing all edge probabilities to decay at the same rate. Thus it suffices to replace α_n, β_n or γ_n by ρ_n in orders of magnitude; for example, the expected degree of nodes in either class is $C_0 n \rho_n$. We use “ C_0 ” to denote a generic positive constant. All expectations are conditioned on Z ; for notational convenience we write $E[\cdot]$ instead of $E[\cdot|Z]$.

First we consider the eigenvalues and eigenvectors of P without the constraint of zero diagonals. If $\alpha_n \beta_n \neq \gamma_n^2$, then this matrix (denoted by P_B) will have two eigenvalues with magnitude $O(n\rho_n)$ and $n - 2$ zero eigenvalues. Since $\|P - P_B\| = O(\rho_n)$, using Weyl’s inequality we see that the principal eigenvalues of P are $O(n\rho_n)$, whereas all other eigenvalues are $O(\rho_n)$.

Let $\mathbf{v}_i (\lambda_i)$ denote the i th eigenvector (eigenvalue) of matrix P . The ordering is in decreasing order of absolute value of the eigenvalues. We will denote the i th empirical eigenvector (eigenvalue) by $\hat{\mathbf{v}}_i (\hat{\lambda}_i)$. $\mathbf{v}_1, \mathbf{v}_2$ are piecewise constant.

Now we will define the normalized counterparts of the above quantities. Let $\tilde{A} := D^{-1/2}AD^{-1/2}$, and also let $\tilde{P} := \mathcal{D}^{-1/2}P\mathcal{D}^{-1/2}$, where D and \mathcal{D} are the diagonal matrices of degrees and expected degrees, respectively. We denote the first two eigenvectors by \mathbf{u}_1 and \mathbf{u}_2 , and the first two eigenvalues by ν_1 and ν_2 . Similar to \mathbf{v}_1 and \mathbf{v}_2 , \mathbf{u}_1 and \mathbf{u}_2 also are piecewise constant vectors. The empirical counterparts of the eigenvectors and values are denoted by $\tilde{\mathbf{u}}_i, \tilde{\nu}_i$. One interesting fact about $\tilde{\mathbf{u}}_1$ is that the i th entry is proportional to $\sqrt{d_i}$, where d_i is the degree of node i . However, one cannot explicitly obtain the form of $\hat{\mathbf{u}}_1(i)$.

Among the many variants of spectral clustering, we consider the algorithm used in [21]. The idea is to compute an $n \times k$ matrix \hat{Q} with the top k eigenvectors of A along its columns, and apply the `kmeans` algorithm on the rows of \hat{Q} . The `kmeans` algorithm searches over different clusterings and returns a local optima of an objective function that minimizes the squared Euclidean distance of points from their respective cluster centers. The clusters are now identified as estimates of the k blocks.

Probabilistic bounds on misclassification errors of spectral clustering under the stochastic blockmodel has been obtained in previous work [21, 24]. However, upper bounds cannot be used for comparing two algorithms. Instead, we define a sim-

ple clustering quality metric computable in terms of an appropriately defined deviation of empirical eigenvectors from their population counterparts, and we show that these are improved by normalization.

2.1. *Quality metrics.* The quality metrics are defined as follows: the algorithm passes the empirical eigenvectors to an oracle who knows the cluster memberships. The oracle computes cluster centers $K_k := \sum_{i \in C_k} \hat{Q}_i / |C_k|$, for us $k \in \{1, 2\}$. Let d_{11}^2 denote the mean squared distance of points in C_1 from K_1 , and let d_{12}^2 denote the mean squared distance of points in C_1 from K_2 . From now on we will denote by \hat{d}_{11}^2 and \tilde{d}_{11}^2 the distances obtained from the unnormalized and normalized methods, respectively.

To be concrete, we can write $\hat{d}_{11}^2 = \sum_{i \in C_1} \|\hat{Q}_i - K_1\|^2 / n\pi$. Similarly, define \hat{d}_{12}^2 as the mean square distance of points in C_1 from K_2 , that is, $\hat{d}_{12}^2 = \sum_{i \in C_1} \|\hat{Q}_i - K_2\|^2 / n\pi$. One can analogously define \hat{d}_{22}^2 and \hat{d}_{21}^2 . We will use the notation d_{11}^2 (or d_{12}^2) when we refer to the corresponding quantities in general, that is, without any particular reference to a specific method.

Although \hat{d}_{11}^2 seems like a simple average of squared distances, it actually has useful information about the quality of clustering. For definiteness, let us take the unnormalized case and examine points in C_1 . By stochastic equivalence, $\forall i \in C_1$, $\{\hat{v}_1(i), \hat{v}_2(i)\}$ are identically distributed (albeit dependent) random variables. Now \hat{d}_{11}^2 essentially is the trace of the 2×2 sample variance matrix, and hence measures the variance of these random variables.

Ideally a good clustering algorithm with or without normalization should satisfy $d_{11}^2 / d_{12}^2 \xrightarrow{P} 0$, but we will show that this ratio converges to zero at the same rate, with or without normalization, in consistence with previous work [18, 21] and [24]. Furthermore, we will show that $\tilde{d}_{12}^2 / \hat{d}_{12}^2 \xrightarrow{P} 1$; that is, the two methods do not distinguish between d_{12}^2 .

Interestingly, our results also imply that d_{11}^2 increases as the graphs become sparser; that is, ρ_n decreases. Hence, if a method can be shown to reduce the variance of points in a class by a constant fraction, it would be preferable for sparse graphs. Indeed we show that $\tilde{d}_{11}^2 / \hat{d}_{11}^2$ converges to a constant which is *less than* 1 for a broad range of parameter settings of α_n, β_n and γ_n . In the simple disconnected case with $\gamma_n = 0$, this constant is $1/4$.

Another advantage of d_{11}^2 is that it can be conveniently expressed in terms of an appropriately defined deviation of empirical eigenvectors from their population counterpart. For any population and empirical eigenvector pair $\{\mathbf{v}, \hat{\mathbf{v}}\}$, we consider the following orthogonal decomposition: $\mathbf{v} = c\hat{\mathbf{v}} + \mathbf{r}$, where $c := \mathbf{v}^T \hat{\mathbf{v}}$. The norm of residual \mathbf{r} will measure the deviation of $\hat{\mathbf{v}}$ from \mathbf{v} . The deviation of $\hat{\mathbf{u}}$ from \mathbf{u} can be measured similarly.

Since we are interested in two class blockmodels, we will mostly use $\mathbf{r}_i, i \in \{1, 2\}$ as the residual of the i th empirical eigenvector from its population counterpart, and $c_{jj} := \mathbf{v}_j^T \hat{\mathbf{v}}_j$. We denote by $v(C_1) := \sum_{i \in C_1} v(i) / n\pi$ the average of

entries of vector \mathbf{v} restricted to class C_1 . A key fact is that $\mathbf{v}_1, \mathbf{v}_2$ (or $\mathbf{u}_1, \mathbf{u}_2$) are both piecewise constant:

$$(2.1) \quad d_{11}^2 = \frac{1}{c_{11}^2} \left(\sum_{i \in C_1} \frac{r_1(i)^2}{n\pi} - r_1(C_1)^2 \right) + \frac{1}{c_{22}^2} \left(\sum_{i \in C_1} \frac{r_2(i)^2}{n\pi} - r_2(C_1)^2 \right),$$

$$(2.2) \quad d_{12}^2 = \frac{1}{n\pi} \sum_{i \in C_1} \|\hat{Q}_i - K_2\|^2 = d_{11}^2 + \|K_1 - K_2\|^2.$$

We will denote the distances obtained from A by $\hat{d}..$ and from \tilde{A} by $\tilde{d}..$. Even though K_i is defined in terms of $\hat{\mathbf{v}}_1$ and $\hat{\mathbf{v}}_2$, we will abuse this notation somewhat to use the above expressions for calculating $\tilde{d}..$, where K_i will be defined identically in terms of $\tilde{\mathbf{u}}_1$ and $\tilde{\mathbf{u}}_2$. For a wide regime of $(\alpha_n, \beta_n, \gamma_n)$, we prove that \tilde{d}_{11}^2 is asymptotically a constant factor smaller and hence better than \hat{d}_{11}^2 . First, using results from [10] we will prove that for $\gamma_n = 0$, $\tilde{d}_{11}^2 = 1/4\hat{d}_{11}^2(1 + o_P(1))$. In this case, the result can be proven using existing results on Erdős–Rényi graphs [10] and a simple application of Taylor’s theorem. In order to generalize the result to $\gamma_n \neq 0$, we would need new convergence results for A and \tilde{A} generated from a stochastic blockmodel. All results rely on the following assumption on ρ_n :

ASSUMPTION 2.1. We assume $\log n/n\rho_n \rightarrow 0$, as $n \rightarrow \infty$.

This assumption ascertains with high probability that the sequence of growing graphs are not too sparse. The expected degree is $np = O(n\rho_n)$, and this is the most commonly used regime where norm convergence of matrices can be shown [5, 6, 18]. Note that this is also the sharp threshold for connectivity of Erdős–Rényi graphs [4]. We will now formally define the sparsity regime in which we derive our results.

DEFINITION 2.2 (A semi-sparse stochastic blockmodel). Define a stochastic blockmodel with parameters $\alpha_n, \beta_n, \gamma_n$ and Z ; see Definition 2.1. Let α_n, β_n and γ_n be deterministic quantities of the form $C_0\rho_n$. If ρ_n satisfies Assumption 2.1, we call the stochastic blockmodel $(\alpha_n, \beta_n, \gamma_n, Z)$ a semi-sparse stochastic blockmodel.

The paper is organized as follows: we present the main results in Section 3. The proof of the simple $\gamma_n = 0$ case is in Section 4, whereas the expressions of \hat{d}_{11}^2 and \hat{d}_{12}^2 in the general case appear in Section 5. We derive the expressions of \tilde{d}_{11}^2 and \tilde{d}_{12}^2 in Section 6. Experiments on simulated and real data appear in Section 7. The proofs of some accompanying lemmas and ancillary results are omitted from the main manuscript for ease of exposition and are deferred to the Supplement [22].

2.2. *Import of previous work.* By virtue of stochastic equivalence of points belonging to the same class, eigenvectors of P map the data to k distinct points. This is why consistency of spectral clustering is closely tied to consistency properties of empirical eigenvalues and eigenvectors. We will show that current theoretical work on eigenvector consistency does not distinguish between the use of normalized or unnormalized A .

One of the earlier results on the consistency of spectral clustering can be found in [26], where weighted graphs generated from a geometric generative model are considered. While this is an important work, this does not apply to our random network model.

For any symmetric adjacency matrix A with independent entries, one can use results on random matrix theory from Oliveira [18] to show that the empirical eigenvectors of a semi-sparse stochastic blockmodel converge to their population counterpart at the same rate with or without normalization. If $\mathbf{p} : [n]^2 \rightarrow [0, 1]$ denotes the probability function $P(A_{ij} = 1) = 1 - P(A_{ij} = 0) = \mathbf{p}(i, j)$, and $d_{\mathbf{p}}$ denotes the expected degree, then:

THEOREM 2.1 (Theorem 3.1 of [18]). *For any constant $c > 0$, there exists another constant $C = C(c) > 0$, independent of n or \mathbf{p} , such that the following holds. Let $d := \min_{i \in [n]} d_{\mathbf{p}}(i)$, $\Delta := \max_{i \in [n]} d_{\mathbf{p}}(i)$. If $\Delta > C \log n$, then for all $n^{-c} \leq \delta \leq 1/2$,*

$$P(\|A - P\| \leq 4\sqrt{\Delta \log(n/\delta)}) \geq 1 - \delta.$$

Moreover, if $d \geq C \log n$, then for the same range of δ ,

$$P(\|\tilde{A} - \tilde{P}\| \leq 14\sqrt{\log(4n/\delta)/d}) \geq 1 - \delta.$$

Let $\Pi_{a,b}(A)$ denote the orthogonal projector onto the space spanned by the eigenvectors of A corresponding to eigenvalues in $[a, b]$. A simple consequence of Theorem 2.1 is that for suitably separated population eigenvalues, the operator norm of the difference of the eigenspaces also converges to zero.

COROLLARY 2.1 (Corollary 3.2 of [18]). *Given some $x > 0$, let $N_x(P)$ be the set of all pairs $a < b$ such that $a + x < b - x$, and P has no eigenvalues in $(a - x, a + x) \cup (b - x, b + x)$. Then for $x > 4\sqrt{\Delta \log(n/\delta)}$,*

$$\begin{aligned} \|A - P\| &\leq 4\sqrt{\Delta \log(n/\delta)} \\ \implies \forall (a, b) \in N_x(P), \\ \|\Pi_{a,b}(A) - \Pi_{a,b}(P)\| &\leq \left(\frac{4(b - a + 2x)}{\pi(x^2 - x\sqrt{\Delta \log(n/\delta)})} \right) \sqrt{\Delta \log(n/\delta)}. \end{aligned}$$

Similarly define $N_x(\tilde{P})$. Then for $x > 14\sqrt{\log(4n/\delta)/d}$,

$$\begin{aligned} \|\tilde{A} - \tilde{P}\| &\leq 14\sqrt{\frac{\log(4n/\delta)}{d}} \\ \implies \forall(a, b) &\in N_x(\tilde{P}), \\ \|\Pi_{a,b}(\tilde{A}) - \Pi_{a,b}(\tilde{P})\| &\leq \left(\frac{4(b-a+2x)}{\pi(x^2 - x\sqrt{\Delta \log(n/\delta)})}\right)\sqrt{\log(n/\delta)/d}. \end{aligned}$$

In particular the right-hand sides of the above equations hold with probability $\geq 1 - \delta$ for any $n^{-c} < \delta < 1/2$.

A straightforward application of this corollary yields that spectral clustering for a stochastic blockmodel with A and \tilde{A} lead to $O_P(\sqrt{\log n/n\rho_n})$ convergence of empirical eigenvectors to their population counterparts. Further analysis shows that the fraction of misclassified nodes go to zero at the same rate for A and \tilde{A} . We defer the proof to Section B of the Supplement [22].

COROLLARY 2.2. *Let A be generated from a semi-sparse stochastic blockmodel (Definition 2.2) with $\gamma_n > 0$ and $\alpha_n\beta_n \neq \gamma_n^2$. Then, for $i \in \{1, 2\}$, $\|\mathbf{v}_i\mathbf{v}_i^T - \hat{\mathbf{v}}_i\hat{\mathbf{v}}_i^T\| = O_P(\sqrt{\log n/n\rho_n})$. Furthermore the fraction of misclassified nodes can be bounded by $O_P(\log n/n\rho_n)$ for both methods.*

Spectral clustering with \tilde{A} derived from a stochastic blockmodel with growing number of blocks has been shown to be asymptotically consistent [21]. Further, the fraction of mis-clustered nodes is shown to converge to zero under general conditions. These results are extended to show that spectral clustering on unnormalized A also enjoys similar asymptotic properties [24]. Sussman et al. [24] also give an example of parameter setting for a stochastic blockmodel where spectral clustering using unnormalized A outperforms that using \tilde{A} . We, however, demonstrate using theory and experiments that this is only a partial picture, and there is a large regime of parameters where normalization indeed improves performance.

For ease of exposition, we list the different variables and their orders of magnitude in Table 1. For deterministic quantities x_n and c_n , $x_n \asymp c_n$, denotes that x_n/c_n converges to some constant as $n \rightarrow \infty$. For two random variables X_n and Y_n , we use $X_n \sim Y_n$ to denote $X_n = Y_n(1 + o_P(1))$. For the scope of this paper $\|\cdot\|$ denotes the L_2 norm, unless otherwise specified.

3. Main results. For the general case we derive the following asymptotic expressions of d_{11}^2 and d_{12}^2 . We recall that d_{11}^2 measures the variance of points in class one under the spectral representation, whereas d_{12}^2 basically measures the distance between the class centers, which can also be thought of as bias. We will show that

TABLE 1
Table of notation

ρ_n	Edge probability	I	The $n \times n$ identity matrix
n	Number of nodes in the network	Z	$n \times 2$ binary matrix of class memberships
C_i	The i th group, $i \in \{1, 2\}$	π	$ C_1 /n$
D	Diagonal matrix of degrees	\mathcal{D}	Diagonal matrix of expected degrees, conditioned on Z
A	Adjacency matrix	\tilde{A}	$D^{-1/2}AD^{-1/2}$
P	$E[A Z]$	\tilde{P}	$\mathcal{D}^{-1/2}P\mathcal{D}^{-1/2}$
α_n	$P[A_{ij} = 1 i \in C_1, j \in C_1] \asymp \rho_n$	μ_1	$E[D_{ii}/n i \in C_1] = \pi\alpha_n + (1 - \pi)\gamma_n - \alpha_n/n \asymp \rho_n$
β_n	$P[A_{ij} = 1 i \in C_2, j \in C_2] \asymp \rho_n$	μ_2	$E[D_{ii}/n i \in C_2] = \pi\gamma_n + (1 - \pi)\beta_n - \beta_n/n \asymp \rho_n$
γ_n	$P[A_{ij} = 1 i \in C_1, j \in C_2] \asymp \rho_n$	μ	$\sum_j D_{ij}/n^2 = \pi\mu_1 + (1 - \pi)\mu_2 \asymp \rho_n$
d_i	$D_{ii}, i \in \{1, \dots, n\} = O_P(n\rho_n)$	\tilde{d}_i	$\sum_j [A_{ij} - E[A_{ij} Z]] = O_P(\sqrt{n\rho_n})$
$\tilde{d}_i^{(1)}$	$\sum_{j \in C_1} [A_{ij} - E[A_{ij} Z]] = O_P(\sqrt{n\rho_n})$	$\tilde{d}_i^{(2)}$	$\sum_{j \in C_2} [A_{ij} - E[A_{ij} Z]] = O_P(\sqrt{n\rho_n})$
E_1	$\sum_{i \in C_1} d_i$	E_2	$\sum_{i \in C_2} d_i$
E	$\sum_i d_i$	$x(C_1)$	The average of \mathbf{x} restricted to C_1 , that is, $\sum_{i \in C_1} \mathbf{x}(i)$
λ_i	i th largest eigenvalue of P in magnitude $\asymp n\rho_n$, for $i \in \{1, 2\}$	v_i	i th largest eigenvalue of \tilde{P} in magnitude $\asymp 1$, for $i \in \{1, 2\}$
\mathbf{v}_i	i th eigenvector of P	\mathbf{u}_i	i th eigenvector of \tilde{P}
x_k	$\mathbf{v}_k(i) (k \in \{1, 2\}, i \in C_1) \asymp 1/\sqrt{n}$	\tilde{x}_k	$\mathbf{u}_k(i) (k \in \{1, 2\}, i \in C_1) \asymp 1/\sqrt{n}$
y_k	$\mathbf{v}_k(i) (k \in \{1, 2\}, i \in C_2) \asymp 1/\sqrt{n}$	\tilde{y}_k	$\mathbf{u}_k(i) (k \in \{1, 2\}, i \in C_2) \asymp 1/\sqrt{n}$
$\hat{\lambda}_i$	i th largest eigenvalue of A in magnitude;	\tilde{v}_i	i th largest eigenvalue of \tilde{A} in magnitude
$\hat{\mathbf{v}}_i$	i th eigenvector of A	$\tilde{\mathbf{u}}_i$	i th eigenvector of \tilde{A}
K_1	$\sum_{j \in C_1} \hat{Q}_j/n\pi$	K_2	$\sum_{j \in C_2} \hat{Q}_j/n\pi$
\hat{Q}	$n \times 2$ matrix of top two empirical eigenvectors (of A) along the columns	\hat{Q}	The population variant of \hat{Q}
$\hat{d}_{k\ell}^2$	$\sum_{i \in C_k} \ \hat{Q}_i - K_\ell\ ^2/n\pi$	$\tilde{d}_{k\ell}^2$	Variant of $\hat{d}_{k\ell}^2$ using eigenvectors of \tilde{A}
C	$\hat{Q}^T \hat{Q}$	c_{ij}	$C_{ij} := \mathbf{v}_i^T \hat{\mathbf{v}}_j$
$\hat{\mathbf{r}}_i$	$\mathbf{v}_i - (\mathbf{v}_i^T \hat{\mathbf{v}}_i) \hat{\mathbf{v}}_i, i \in \{1, 2\}$	$\tilde{\mathbf{r}}_i$	$\mathbf{u}_i - (\mathbf{u}_i^T \tilde{\mathbf{u}}_i) \tilde{\mathbf{u}}_i, i \in \{1, 2\}$

normalizing A asymptotically reduces the variance without affecting the bias. The proofs can be found in Sections 5 and 6.

THEOREM 3.1. *Let A be the adjacency matrix generated from a semi-sparse stochastic blockmodel $(\alpha_n, \beta_n, \gamma_n, Z)$ where $\gamma_n > 0$ and $\alpha_n\beta_n \neq \gamma_n^2$. We define λ_i, x_i, y_i , for $i \in \{1, 2\}$ and π as in Table 1:*

$$(3.1) \quad \hat{d}_{11}^2 \sim \left[\left(\frac{x_1^2}{\lambda_1^2} + \frac{x_2^2}{\lambda_2^2} \right) n\pi\alpha_n(1 - \alpha_n) + \left(\frac{y_1^2}{\lambda_1^2} + \frac{y_2^2}{\lambda_2^2} \right) n(1 - \pi)\gamma_n(1 - \gamma_n) \right],$$

$$(3.2) \quad \hat{d}_{12}^2 \sim 1/n\pi(1 - \pi).$$

THEOREM 3.2. *Let A be the adjacency matrix generated from a semi-sparse stochastic blockmodel $(\alpha_n, \beta_n, \gamma_n, Z)$ where $\gamma_n > 0$ and $\alpha_n\beta_n \neq \gamma_n^2$. We define μ_1 ,*

μ_2, v_2 and π as in Table 1:

$$(3.3) \quad \tilde{d}_{11}^2 \sim \left[\frac{n\pi\alpha_n(1-\alpha_n)}{n^3\pi\mu_1^2} \left(\frac{1}{4} + \frac{(1-\pi)\gamma_n}{\mu_1 v_2^2} \right) + \frac{n(1-\pi)\gamma_n(1-\gamma_n)}{n^3\mu_1^2} \left(\frac{1}{4\pi} + \frac{\pi\alpha_n}{(1-\pi)\mu_2 v_2^2} \right) \right],$$

$$(3.4) \quad \tilde{d}_{12}^2 \sim \frac{1}{n\pi(1-\pi)}.$$

Before explaining the above theorems, we present a special case for clarity.

REMARK (A special case). When $\gamma_n = 0$, we have $\lambda_1 = n\mu_1$, and $x_1 = 1/\sqrt{n\pi}$ and $y_1 = 0$, which immediately shows that normalization shrinks the variance of the spectral embedding within a class (d_{11}^2) by a factor of four. This is the completely associative case. Now consider the completely dissociative case, that is, $\gamma_n = 0$. It is easy to see that then $\lambda_1 = -\lambda_2 = n\sqrt{\pi(1-\pi)}\gamma_n$, and $y_1 = -y_2 = 1/\sqrt{2n(1-\pi)}$. Substituting these values into the distance formulas again shows that normalization shrinks d_{11}^2 by a factor of four in the completely dissociative case.

We call the completely associative case the zero communication case, which can be thought of as two disconnected Erdős–Rényi graphs. Under Assumption 2.1 each of the smaller graphs will be connected with probability tending to one. von Luxburg [25] already established that spectral clustering achieves perfect classification in this scenario. We merely present this simple setting because the ideas and proof techniques used for this case will be carried over to the general case with $\gamma_n \neq 0$. In particular, our results indicate that in the general case ($\alpha_n, \gamma_n > 0$), for parameter regimes close to the completely associative or completely dissociative models, the normalized method has a clear edge.

COROLLARY 3.1. Let A be the adjacency matrix generated from a semi-sparse stochastic blockmodel $(\alpha_n, \beta_n, \gamma_n, Z)$ (see Definition 2.2) where $\gamma_n = 0$ and $\alpha_n\beta_n \neq \gamma_n^2$. We have

$$(3.5) \quad \frac{\hat{d}_{11}^2}{\tilde{d}_{11}^2} \sim 4,$$

$$(3.6) \quad \frac{\hat{d}_{12}^2}{\tilde{d}_{12}^2} \sim 1,$$

$$(3.7) \quad \frac{\hat{d}_{11}^2}{\hat{d}_{12}^2} = O_P\left(\frac{1}{n\rho_n}\right),$$

$$(3.8) \quad \frac{\tilde{d}_{11}^2}{\tilde{d}_{12}^2} = O_P\left(\frac{1}{n\rho_n}\right).$$

The same holds for normalized and unnormalized versions of d_{22}^2 and d_{21}^2 .

While both of d_{11}^2 and d_{12}^2 (derived the unnormalized and normalized methods) are approaching zero in probability, d_{11}^2/d_{12}^2 is $C_0/n\rho_n$ for both the normalized and unnormalized cases. In our regime of ρ_n this translates to perfect classification as $n \rightarrow \infty$. This is not unexpected because existing literature has established that spectral clustering with both A and \tilde{A} are consistent in the semi-sparse regime. Also, $\tilde{d}_{12}^2/\hat{d}_{12}^2$ approaches one; thus if the limiting ratio $\tilde{d}_{11}^2/\hat{d}_{11}^2$ is smaller (larger) than one, then there is some indication that the normalized (unnormalized) method is to be preferred.

For simplicity, we consider a stochastic blockmodel with two equal sized classes and $\beta_n = \alpha_n$. We will now show that for this simple model, in the semi-sparse regime, our quality metric indicates that normalization would always improve performance. In the dense regime, that is, when degree grows linearly with n , there are parts of the parameter space where our quality metric prefers the unnormalized method. However, the network is so dense that the two class centers are well separated, leading to equally good performance of both methods.

For this simple model, the limiting $\tilde{d}_{11}^2/\hat{d}_{11}^2$ ratio has a concise form in the semi-sparse case, which is presented in Corollary 3.2, and plotted in Figure 1(A). Figure 1(B) shows the contour plot of the limiting ratio in the dense case. We also highlight the parameter regime where the ratio is close to or larger than one. Finally Figure 1(C) focuses on this highlighted area.

COROLLARY 3.2. *Let A be the adjacency matrix generated from the stochastic blockmodel $(\alpha_n, \alpha_n, \gamma_n, Z)$ where $\gamma_n = x\alpha_n > 0$ and $\pi = 1/2$. When $\rho_n \rightarrow 0$, we have the following limit, which is always smaller than one:*

$$(3.9) \quad \frac{\tilde{d}_{11}^2}{\hat{d}_{11}^2} \sim \frac{1}{4} + \frac{3}{2} \frac{x}{1+x^2}.$$

On the other hand, when ρ_n is a constant w.r.t. n , the above ratio is smaller than one, unless $x \geq 1$ or $\alpha_n \geq 1/2$. The universal upper bound is 1.31.

REMARK. Here we summarize the result in the above corollary.

(1) In the semi-sparse regime [Figure 1(A)], the limiting ratio is always less than one, thus favoring the normalized method.

(2) In the dense case [Figure 1(B) and (C)], where ρ_n is a constant w.r.t. n , this ratio can be larger than one when $x \geq 1$ or $\alpha_n \geq 1/2$, with an upper bound of 1.31. The upper bound is achieved for large α_n, γ_n pairs, for example, $\alpha_n = 1/3, \gamma_n = 1$ and $\alpha_n = 1, \gamma_n = 0.24$. In this dense regime, both methods perform equally well on any reasonably sized network. Using simulations on small networks (twenty nodes), we found that in terms of misclassification error, the methods perform comparably.

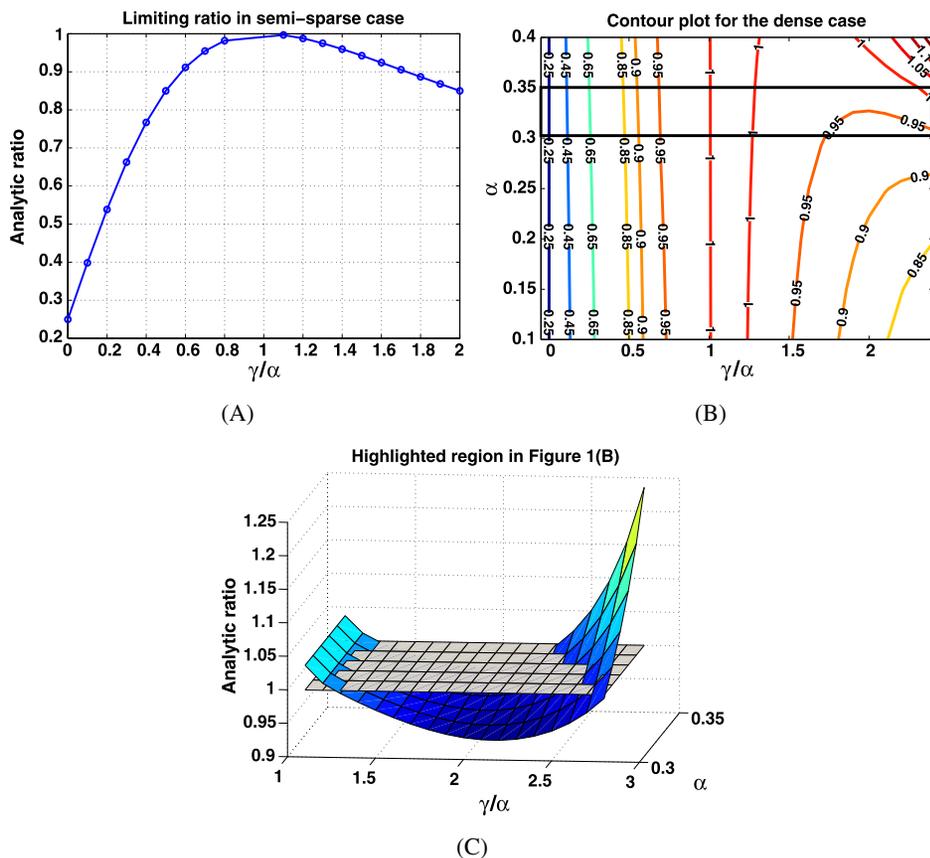


FIG. 1. Simple blockmodel with two equal sized classes and $\alpha_n = \beta_n$: (A) Limiting ratio of $\tilde{d}_{11}^2/\hat{d}_{11}^2$ in the semi-sparse case. (B) Contour plot for the ratio in the dense case. The rectangular area consists of parameters settings leading to a ratio bigger than one. This is highlighted in (C), which shows the surface plot for $\tilde{d}_{11}^2/\hat{d}_{11}^2$ along the Z axis in the regime where ratio is close to or larger than one. Y axis has varying α_n , X axis has varying γ_n/α_n . For reference we also plot the plane $Z = 1$.

(3) Because of the inherent symmetry of the simple model, for $y := \alpha_n/\gamma_n$, the ratio $\tilde{d}_{11}^2/\hat{d}_{11}^2 \rightarrow 1/4 + 3y/2(1 + y^2)$, in the semi-sparse regime. This again shows that normalization provides a clear edge close to the completely associative ($\gamma_n = 0$) or completely dissociative ($\alpha_n = 0$) cases.

We want to point out that for the simulated experiment with $\alpha_n = 0.42, \beta_n = 0.50, \gamma_n = 0.42$ (and $\pi = 0.60$), the unnormalized method performs better than the normalized method in Sussman et al. [24]. In this case the analytic ratio also is larger than one, and the graph is very close to an Erdős–Rényi graph.

3.1. A shortcoming of asymptotic analysis. While Corollary 3.2 suggests that normalization always reduces within class variance in the semi-sparse degree

regime, there is one caveat to this asymptotic result. In the semi-sparse regime, the network is connected with high probability as $n \rightarrow \infty$. However, finite sparse networks may have disconnected components consisting of a few nodes. In such scenarios, by construction the normalized method assigns eigenvalue one to eigenvectors with support on nodes in each of the connected components. As a result, the leading eigenvectors are uninformative, leading to poor performance. The unnormalized method, however, does not suffer from this problem and has the informative eigenvectors as the leading eigenvectors, albeit with a high variance. We get around this problem by removing small degree nodes and then working on the largest connected component. We also point out this problem in the discussion section.

3.2. Accuracy of the analytic ratio. Finally, we also use simulations to see how accurate the analytic ratio is. For $n = 1000$, we vary $\alpha_n \in [0.4, 0.6]$, $\beta_n \in [0.5, 0.9]$ and γ_n/α_n between zero and two such that $\forall \alpha_n, \gamma_n \leq 1$, and $\gamma_n^2 \neq \alpha_n \beta_n$. We note that the ratio increases for large (α_n, γ_n) pairs. The mean, median and maximum absolute relative error for $\tilde{d}_{11}^2/\hat{d}_{11}^2$ ($\tilde{d}_{12}^2/\hat{d}_{12}^2$) from their analytic counterparts is 0.02, 0.02 and 0.1 (0.001, 0.001 and 0.03), respectively. In both cases the maximum happens for the $\{\alpha_n, \beta_n, \gamma_n\}$ combination where $|\alpha_n \beta_n - \gamma_n^2|$ is the smallest, leading to most instability. Since all our $o_P(1)$ terms are $O_P(1/\sqrt{n\rho_n})$, for this experiment these errors are indeed justifiable.

4. The zero communication case. We will now present our result for two class blockmodels (see Definition 2.2) with $\gamma_n = 0$. We will heavily use the following orthogonal decomposition of the population eigenvectors:

$$\mathbf{v}_k := c_{kk} \hat{\mathbf{v}}_k + \mathbf{r}_k \quad \text{for } k \in \{1, 2\}, \text{ where } c_{kk} = \mathbf{v}_k^T \hat{\mathbf{v}}_k.$$

Since $\gamma_n = 0$, A can be thought of as two disconnected Erdős–Rényi graphs of size $n\pi$ and $n(1 - \pi)$ (let the two adjacency matrices be denoted by A_1 and A_2 , resp.). We assume WLOG $\pi\alpha_n > (1 - \pi)\beta_n$ so that $\lambda_1 = n\pi\alpha_n + O(\rho_n)$ and $\lambda_2 = n(1 - \pi)\beta_n + O(\rho_n)$. We also assume that rows and columns of A are permuted so that the first $n\pi$ entries are from C_1 . (We will not use this in our proofs; it only helps the exposition.)

Füredi and Komlós [10] show that for $i \in \{1, 2\}$, $\hat{\lambda}_i = \lambda_i + O_P(1)$ and $\max_{i>2} |\lambda_i| = O_P(\sqrt{n\rho_n})$. Hence for large n , the second largest eigenvalue will come from A_2 , and will have zeros along the first class, similar to the second population eigenvector. Thus $\hat{r}_1(i) = 0$ for $i \in C_2$, and vice versa.

Further, some algebra reveals that $K_1 = \{c_{11}/\sqrt{n\pi}, 0\}$ and $K_2 = \{0, c_{22}/\sqrt{n(1 - \pi)}\}$. Computing \hat{d}_{11}^2 or \tilde{d}_{11}^2 requires one to compute the norm and average of $\hat{\mathbf{r}}_k$ and $\tilde{\mathbf{r}}_k$ restricted to C_1 ; see equation (2.1). For $\gamma_n = 0$, this reduces to examining $\hat{\mathbf{r}}$ and $\tilde{\mathbf{r}}$ for two Erdős–Rényi graphs.

Let us consider an Erdős–Rényi graph $G_{n,p}$. Since self-loops are prohibited, the conditional expectation matrix P is simply $p(\mathbf{1}\mathbf{1}^T - I)$, which has n eigenvalues,

the largest of which is $\lambda := (n - 1)p$, and the rest are all $-p$. We denote by d_i the degree of node i , and $\bar{d}_i := d_i - (n - 1)p$.

Let λ, \mathbf{v} (ν, \mathbf{u}) be respectively the principal eigenvalue and eigenvector pair of P (\tilde{P}), whose empirical counterparts are given by $\hat{\lambda}$ and $\hat{\mathbf{v}}$ ($\hat{\nu}$ and $\hat{\mathbf{u}}$) respectively. In this simple case, \mathbf{v} and \mathbf{u} are the same. We require that all eigenvectors are unit-length. We denote by $\langle x_i \rangle$ the a n length vector with the i th entry equaling x_i . We note that $\mathbf{v} = \langle 1/\sqrt{n} \rangle$, and $\tilde{\mathbf{v}} = \langle \sqrt{d_i / \sum_j d_j} \rangle$. Let $\hat{c} := \hat{\mathbf{v}}^T \mathbf{v}$ and $\tilde{c} := \hat{\mathbf{u}}^T \mathbf{u}$. We will prove that $\|\hat{\mathbf{r}}\|^2 \sim 4\|\tilde{\mathbf{r}}\|^2$, which will help us prove Corollary 3.1.

Before proceeding with the result, for ease of exposition we recall the orders of magnitudes of some random variables used in the proof. Let E denote $\sum_i d_i$. We have $\sum_i \bar{d}_i = O_P(n\sqrt{\rho_n})$ [this is simply twice the sum of $\binom{n}{2}$ centered Bernoulli(p) variables and $\sum_i \bar{d}_i^2 = n^2 p(1 + o_P(1))$]. The later result can be obtained by showing that the expectation is $n(n - 1)p$, and the standard deviation is of a smaller order. A detailed proof can be found in [10].

LEMMA 4.1. *Write the first population eigenvector \mathbf{v} of an Erdős–Rényi (n, p) graph adjacency matrix A as $\mathbf{v} = \hat{c}\hat{\mathbf{v}} + \hat{\mathbf{r}}$. If $p = C_0\rho_n$ satisfies Assumption 2.1, we have*

$$\|\hat{\mathbf{r}}\|^2 \sim \frac{1}{((n - 1)p)^2} \frac{\sum_i \bar{d}_i^2}{n}.$$

PROOF. Before delving into the proof, we state the main result from [10]. For an Erdős–Rényi graph, $\hat{\lambda}_1 = \frac{\mathbf{1}^T A \mathbf{1}}{n} + (1 - p) + O_P(1/\sqrt{n})$. Since $\frac{\mathbf{1}^T A \mathbf{1}}{n} - (n - 1)p = O_P(\sqrt{p(1 - p)})$, we have $\hat{\lambda}_1 - (n - 1)p = O_P(1)$. As the explicit form of $\hat{\mathbf{v}}$ is not known, the following step is used to compute the norm of $\hat{\mathbf{r}}$:

$$(4.1) \quad (A - \hat{\lambda}_1 I)\hat{\mathbf{r}} = \langle (d_i - \hat{\lambda}_1)/\sqrt{n} \rangle.$$

The proof is straightforward. First we see that

$$A\mathbf{v} - \hat{\lambda}_1\mathbf{v} = A(\hat{c}\hat{\mathbf{v}} + \hat{\mathbf{r}}) - \hat{\lambda}_1\mathbf{v} = (A - \hat{\lambda}_1 I)\hat{\mathbf{r}}.$$

Using $\mathbf{v} = \langle 1/\sqrt{n} \rangle$, $A\mathbf{v} - \hat{\lambda}_1\mathbf{v} = \langle (d_i - \hat{\lambda}_1)/\sqrt{n} \rangle$, thus proving equation (4.1). Now equation (4.1) and standard norm-inequalities yield $\|A - \hat{\lambda}_1 I\| \leq (\hat{\lambda}_1 + \max(\hat{\lambda}_2, |\hat{\lambda}_n|))$, where $\hat{\lambda}_i$ is the i th largest eigenvalue of A .

Now, using results from [8] we have $\max(\hat{\lambda}_2, |\hat{\lambda}_n|) = O_P(\sqrt{np})$, and hence $\|A - \hat{\lambda}_1 I\| \sim np$. Interestingly, note that $\hat{\mathbf{r}} \perp \hat{\mathbf{v}}$, and hence $\|A\hat{\mathbf{r}}\|/\|\hat{\mathbf{r}}\| = O_P(\sqrt{np})$. Hence $\|(A - \hat{\lambda}_1 I)\hat{\mathbf{r}}\| \geq \hat{\lambda}_1(1 + o_P(1))\|\hat{\mathbf{r}}\|$. Combining this with the former upper bound, we have

$$\|(A - \hat{\lambda}_1 I)\hat{\mathbf{r}}\| \sim \hat{\lambda}_1\|\hat{\mathbf{r}}\|.$$

Since $\bar{d}_i/\sqrt{n} = O_P(\sqrt{p(1 - p)})$ and $E[d_i] = (n - 1)p$, we have

$$\sum_i \frac{(d_i - \hat{\lambda}_1)^2}{n} = \sum_i \frac{\bar{d}_i^2}{n} + (\hat{\lambda}_1 - (n - 1)p)^2 - 2(\hat{\lambda}_1 - (n - 1)p) \frac{\sum_i \bar{d}_i}{n} \sim \sum_i \frac{\bar{d}_i^2}{n}.$$

The last step is true because $\sum_i \frac{\bar{d}_i^2}{n} = O_P(np(1-p))$, whereas both $\hat{\lambda}_1 - (n-1)p$ and $\sum_i \bar{d}_i/n$ are $O_P(1)$. Simple application of the Cauchy–Schwarz inequality shows that the cross term is also $O_P(1)$. Now we have

$$(4.2) \quad \hat{\mathbf{r}}^T \hat{\mathbf{r}} = \frac{1}{\hat{\lambda}_1^2} \sum_i \frac{(d_i - \hat{\lambda}_1)^2}{n} \sim \frac{1}{((n-1)p)^2} \frac{\sum_i \bar{d}_i^2}{n}. \quad \square$$

Since the form of $\tilde{\mathbf{u}}$ is known, $\tilde{\mathbf{r}}^T \tilde{\mathbf{r}}$ can be obtained by using element-wise Taylor expansion. The only complication arises because we often approximate the norm of a length n vector by the norm of its first or second order Taylor expansion, where n is growing to infinity. Hence we present the following helping lemma, where we formalize sufficient conditions for neglecting lower order terms in such an expansion.

LEMMA 4.2. Consider length n vector $\mathbf{x}_n := c_n + \mathbf{x}_n^1 + \mathbf{R}_n$ where c_n is a vector of constants c . If both $\|\mathbf{R}_n\| = o_P(\|\mathbf{x}_n^1\|)$ and $|\sum_i x_n^1(i)/n| = o_P(\|\mathbf{x}_n^1\|/\sqrt{n})$, as $n \rightarrow \infty$, $\sum_i (x_n(i) - \sum_i x_n(i)/n)^2 \sim \|\mathbf{x}_n^1\|^2$.

The following lemma has the asymptotic form of $\|\tilde{\mathbf{r}}\|^2$.

LEMMA 4.3. Write the first population eigenvector \mathbf{u} of an Erdős–Rényi (n, p) graph normalized adjacency matrix \tilde{A} as $\mathbf{u} = \tilde{c}\tilde{\mathbf{u}} + \tilde{\mathbf{r}}$. If $p = O(\rho_n)$ satisfies Assumption 2.1,

$$\|\tilde{\mathbf{r}}\|^2 \sim \frac{1}{4n(n-1)p^2} \sum_i \frac{\bar{d}_i^2}{n}.$$

PROOF SKETCH. We will use the fact that $\|\tilde{\mathbf{r}}\|^2 = 1 - \tilde{c}^2 = \sum_i (\tilde{u}_i - \sum_i \tilde{u}_i/n)^2$. Since one can explicitly obtain the expression of u_i , the basic idea is to use Taylor approximation term by term to obtain the norm. However, the issue is that we are summing over n elements where n is going to infinity, and extra care is required for the remainder terms; in particular, we will bound them uniformly over n .

It is easy to check that the vector $\langle \sqrt{\bar{d}_i/E} \rangle$ is an eigenvector of \tilde{A} with eigenvalue one. By virtue of Assumption 2.1 we know that A is connected with high probability, and so the principal eigenvalue has multiplicity one. Thus $\tilde{\mathbf{u}}(i) = \sqrt{\bar{d}_i/E}$. Now termwise Taylor approximation gives

$$(4.3) \quad \tilde{u}_i = \frac{1}{\sqrt{n}} + \frac{\bar{d}_i}{2\sqrt{n(n-1)^2 p^2}} + R,$$

where R is a length n vector of remainder terms. We will now invoke Lemma 4.2. Let c_n be the vector of constants $1/\sqrt{n}$, and $x_n^1 := \frac{\bar{d}_i}{2\sqrt{d_0 E_0}}$, where d_0 and E_0

are respectively the expectation of d_i and E . In particular $d_0 = (n - 1)p$ and $E_0 = n(n - 1)p$. Hence $\|x_n^1\| \sim C_0/\sqrt{n\rho_n}$, and the mean of x_n^1 is $O_P(1/\sqrt{n^3\rho_n}) = o_P(\|x_n^1\|/\sqrt{n})$. Using standard probabilistic arguments and the form of R , we show that $\|R\| = o_P(\|1/\sqrt{n\rho_n}\|)$; for details, see Section D of the Supplement [22]. Hence we have

$$\|\tilde{\mathbf{r}}\|^2 = \sum_i \left(\tilde{u}_i - \sum_i \tilde{u}_i/n \right)^2 \sim \frac{1}{4n(n-1)p^2} \sum_i \frac{\tilde{d}_i^2}{n}. \quad \square$$

PROOF OF COROLLARY 3.1. In order to compute \tilde{d}_{11}^2 and \hat{d}_{11}^2 , we need to compute the norms and averages of $\hat{\mathbf{r}}_k$ and $\tilde{\mathbf{r}}_k, k \in \{1, 2\}$ restricted to class C_1 . First note that $\hat{\mathbf{r}}_1(C_1) = \hat{\mathbf{r}}_1^T \mathbf{v}_1/\sqrt{n\pi} = \|\hat{\mathbf{r}}_1\|^2/\sqrt{n\pi}$ by construction, and $\hat{r}_2(i) = 0$, for $i \in C_1$. Hence $\sum_{i \in C_1} \hat{r}_1(i)^2 = \|\hat{\mathbf{r}}_1\|^2$.

Using equation (2.1), $\hat{d}_{11}^2 = (\|\hat{\mathbf{r}}_1\|^2/n\pi - \|\hat{\mathbf{r}}_1\|^4/n\pi)/\hat{c}_{11}^2 = \|\hat{\mathbf{r}}_1\|^2/n\pi$. But $\|\hat{\mathbf{r}}_1\|^2$ is the norm-square of the residual of the principal eigenvector from A_1 which is an Erdős–Rényi $(n\pi, \alpha_n)$ graph; see Lemma 4.1.

Now we consider the corresponding quantities from \tilde{A} . The only issue is that A has two disconnected components (each of which is connected w.h.p., via Assumption 2.1), and hence \tilde{A} will have two eigenvalues equal to one; hence the first two eigenvectors can be any two orthogonal vectors spanning this eigenspace. Since Euclidean distances (e.g., $\tilde{d}_{11}^2, \tilde{d}_{12}^2$ etc.) are preserved under rotation, in this simple setting, any such pair of vectors can be shown to yield the same answer.

We will construct \mathbf{u}_1 and $\tilde{\mathbf{u}}_1$ as follows. \mathbf{u}_2 and $\tilde{\mathbf{u}}_2$ are constructed analogously.

$$\mathbf{u}_1(i) = \begin{cases} 1/\sqrt{n\pi}, & \text{for } i \in C_1, \\ 0, & \text{otherwise,} \end{cases} \quad \tilde{\mathbf{u}}_1(i) = \begin{cases} \sqrt{d_i/E}, & \text{for } i \in C_1, \\ 0, & \text{otherwise.} \end{cases}$$

Since \mathbf{u} and \mathbf{v} are identical in the zero communication case, we have $\tilde{d}_{11}^2 = \|\tilde{\mathbf{r}}_1\|^2/n\pi$. However, $\tilde{\mathbf{r}}_1$ is simply the residual of the principal eigenvector from \tilde{A}_1 . Now an application of Lemmas 4.1 and 4.3 proves equation (3.5).

As for \hat{d}_{12}^2 , note that $\hat{\mathbf{v}}(C_1) = \mathbf{v}_1^T \hat{\mathbf{v}}/\sqrt{n\pi} = c_{11}/\sqrt{n\pi}$. Hence, $K_1 = \{c_{11}/\sqrt{n\pi}, 0\}$ and $K_2 = \{c_{22}/\sqrt{n(1-\pi)}, 0\}$. Thus $\|K_1 - K_2\|^2 \sim 1/n\pi(1-\pi)$, since both $c_{11}^2 = 1 - \hat{\mathbf{r}}_1^T \hat{\mathbf{r}}_1$ and $c_{22}^2 = 1 - \hat{\mathbf{r}}_2^T \hat{\mathbf{r}}_2$ are $1 - o_P(1)$ (Lemma 4.1). Since $\hat{d}_{11}^2 = O_P(1/n^2\rho_n)$ is of smaller order than $\|K_1 - K_2\|^2$, using equation (2.2) we see that $\hat{d}_{12}^2 \sim 1/n\pi(1-\pi)$. An identical argument shows that $\tilde{d}_{12}^2 \sim 1/n\pi(1-\pi)$, yielding equation (3.6). With or without normalization, we have $d_{11}^2 = O_P(1/n^2\rho_n)$, whereas $d_{12}^2 \sim 1/n\pi(1-\pi)$; this yields equations (3.7) and (3.8). Finally, an identical argument proves the result for the normalized and unnormalized versions of d_{22}^2 and d_{21}^2 . \square

5. Analysis of the unnormalized method. In this section we obtain expressions for d_{11}^2 and d_{12}^2 when $\gamma_n \neq 0$ for A . First we give a simple lemma describing

the eigen-structure of the conditional probability matrix P . The proof is simple and is deferred to the Supplement.

LEMMA 5.1. *Define a stochastic blockmodel (see Definition 2.2) with parameters $(\alpha_n, \beta_n, \gamma_n, Z)$, where $\gamma_n > 0$ and $\alpha_n\beta_n \neq \gamma_n^2$. The two population eigenvectors of P are piecewise constant with first $n\pi$ elements x_1 and x_2 , respectively, and the second $n(1 - \pi)$ elements y_1 and y_2 , respectively. These elements are of the form C_0/\sqrt{n} , and they satisfy the following:*

$$(5.1) \quad x_1^2 + x_2^2 = 1/n\pi; \quad y_1^2 + y_2^2 = 1/n(1 - \pi); \quad x_1y_1 + x_2y_2 = 0.$$

The two principal population eigenvalues λ_1 and λ_2 are of the form $C'n\rho_n$ and $C''n\rho_n$, where C' and C'' are deterministic constants asymptotically independent of n ; also, $|\lambda_1 - \lambda_2|$ is of the form $C'''n\rho_n$ for some arbitrary constant C''' , when $\gamma_n > 0$. All other eigenvalues of P are $O(\rho_n)$.

We will now lay the groundwork for our result on \hat{d}_{11}^2 and \hat{d}_{12}^2 . In order to extend the simple zero-communication case to the general case, we will need some key results, which are listed below. Recall the following decomposition of the population eigenvector:

$$(5.2) \quad \mathbf{v}_k = c_{kk}\hat{\mathbf{v}}_k + \hat{\mathbf{r}}_k.$$

We will need the following three key components in order to use the same technique as in Lemma 4.1:

- (1) Sharp deviation of empirical eigenvalues. For $\gamma_n = 0$, we have $\hat{\lambda}_k = \lambda_k + O_P(1)$.
- (2) Upper bound on $\|A\hat{\mathbf{r}}_k\|$. For $\gamma_n = 0$, we have $\|A\hat{\mathbf{r}}_k\| = O_P(1)$.
- (3) Bound on the average of $\hat{\mathbf{r}}_k$ restricted to C_1 . For $\gamma_n = 0$ we have $\hat{\mathbf{r}}_k(C_1) = O_P(1/n^{3/2}\rho_n)$.

In Section E of the Supplement [22] we will provide detailed proofs of the following theorems, which show that the above results are also true when $\gamma_n \neq 0$.

In the following lemma we establish a sharp eigenvalue deviation result for blockmodels similar to the one for Erdős–Rényi graphs presented in [10]. Füredi and Komlós [10] use the von Mises iteration (also popularly known as power iteration), which intuitively returns a good approximation of the principal eigenvalue in a few iterations if the second eigenvalue is much smaller than the first. In [10] the second largest eigenvalue of the adjacency matrix is shown to be an order smaller than the first; hence two steps of power iteration can be shown to give a $O_P(1)$ close approximation of $\hat{\lambda}_1$. On the other hand, this approximation can also be shown to be $O_P(1)$ close to the population eigenvalue λ_1 , thus giving the sharp deviation bound.

In a stochastic blockmodel the second largest eigenvalue is of the same order as the first, which is problematic. However, the third largest eigenvalue can be

shown to be $O_P(\sqrt{n\rho_n \log n})$. Therefore we design a two-dimensional von Mises-style iteration argument, so that at any step, the residual vector is orthogonal to the first two empirical eigenvectors, and thus a $O_P(1)$ deviation of the empirical eigenvalues from their population counterparts can be proved. While we prove this result only for the two class blockmodels, the proof can be extended easily to k -class blockmodels, as long as $\lambda_i, i \in \{1, \dots, k\}$ are distinct.

LEMMA 5.2. *Consider an n node network generated from the semi-sparse stochastic blockmodel $(\alpha_n, \beta_n, \gamma_n, Z)$ with $\gamma_n > 0$. We have*

$$\text{For } i \in \{1, 2\}, \quad \hat{\lambda}_i = \lambda_i + O_P(1).$$

Next we need to show that $\|A\hat{\mathbf{r}}_k\| = O_P(1), k \in \{1, 2\}$, even when $\gamma_n \neq 0$. For definiteness let $k = 1$. We want to emphasize that proving $\|\hat{\mathbf{r}}_1\| = O_P(1/\sqrt{n\rho_n})$ is not enough to get the above. By construction $\hat{\mathbf{r}}_1$ is orthogonal to $\hat{\mathbf{v}}_1$, and hence $\|A\hat{\mathbf{r}}_1\|$ can be upper bounded by $|\hat{\lambda}_2|\|\hat{\mathbf{r}}_1\|$. For an Erdős–Rényi graph, $\hat{\lambda}_2 = O_P(\sqrt{n\rho_n})$ leading to an $O_P(1)$ bound, whereas for a stochastic blockmodel, $\hat{\lambda}_2 = O_P(n\rho_n)$ leading to a $O_P(\sqrt{n\rho_n})$ bound. We show the required result by proving that $\hat{\mathbf{v}}_2^T \hat{\mathbf{r}}_1 = O_P(1/n\rho_n)$. Since $\hat{\mathbf{v}}_1$ is orthogonal to $\hat{\mathbf{v}}_2$, $\hat{\mathbf{v}}_2^T \hat{\mathbf{r}}_1 = \hat{\mathbf{v}}_2^T \mathbf{v}_1$, which we prove to be $O_P(1/n\rho_n)$ in the following lemma.

LEMMA 5.3. *For the stochastic blockmodel $(\alpha_n, \beta_n, \gamma_n, Z)$ with $\gamma_n > 0$, define $c_{12} := \mathbf{v}_1^T \hat{\mathbf{v}}_2$ and $c_{21} := \mathbf{v}_2^T \hat{\mathbf{v}}_1$. We have*

$$\begin{aligned} \|\hat{\mathbf{r}}_1\|^2 &= 1 - c_{11}^2 = O_P(1/n\rho_n), & \|\hat{\mathbf{r}}_2\|^2 &= 1 - c_{22}^2 = O_P(1/n\rho_n), \\ c_{12} &:= \mathbf{v}_1^T \hat{\mathbf{v}}_2 = O_P(1/n\rho_n), & c_{21} &:= \mathbf{v}_2^T \hat{\mathbf{v}}_1 = O_P(1/n\rho_n). \end{aligned}$$

The final task is to show that $\hat{r}_k(C_1)$ and $\hat{r}_k(C_2)$ are small. The Cauchy–Schwarz inequality gives $|\hat{r}_k(C_1)| \leq \|\hat{\mathbf{r}}_k\|/\sqrt{n} = O_P(1/n\sqrt{\rho_n})$. However, for a stochastic blockmodel, by virtue of stochastic equivalence, \mathbf{v}_k for $k \in \{1, 2\}$ is piecewise constant; that is, all entries in C_1 have value x_k , whereas all in C_2 have value y_k . Now entries of $\hat{\mathbf{v}}_k$ in C_1 (C_2) constitute a noisy estimate of x_k (y_k). However, one should be able to get an even better estimate by considering $\hat{v}_k(C_1)$ and $\hat{v}_k(C_2)$. Since $\hat{r}_k(C_1)$ reflects the error of $\hat{v}_k(C_1)$ around x_k , it is plausible that $\hat{r}_k(C_1)$ is an order smaller than $\|\hat{\mathbf{r}}_k\|$, which is what we prove in the following lemma.

LEMMA 5.4. *Write $\mathbf{v}_i := c_{i1}\hat{\mathbf{v}}_1 + \hat{\mathbf{r}}_i$ for $i \in \{1, 2\}$. Now we have*

$$\text{For } i, j \in \{1, 2\}, \quad \hat{r}_i(C_j) = O_P(1/n^{3/2}\rho_n).$$

Before proceeding to prove our main result, we present the following simple concentration results, which are derived in the Supplement [22].

LEMMA 5.5. Denote $\bar{d}_i^{(1)}$ and $\bar{d}_i^{(2)}$ as the centered degree of node i restricted to blocks C_1 and C_2 , respectively. In particular, for $k \in \{1, 2\}$, $\bar{d}_i^{(k)} = \sum_{j \in C_k} (A_{ij} - E[A_{ij}|Z])$. We have

$$(5.3) \quad \sum_{i \in C_1} (\bar{d}_i^{(1)})^2 \sim (n\pi)^2 \alpha_n (1 - \alpha_n),$$

$$\sum_{i \in C_1} (\bar{d}_i^{(2)})^2 \sim n^2 \pi (1 - \pi) \gamma_n (1 - \gamma_n),$$

$$(5.4) \quad \sum_{i \in C_1} (x_1 \bar{d}_i^{(1)} + y_1 \bar{d}_i^{(2)})^2 \sim \left(x_1^2 \sum_{i \in C_1} (\bar{d}_i^{(1)})^2 + y_1^2 \sum_{i \in C_1} (\bar{d}_i^{(2)})^2 \right).$$

Now we prove Theorem 3.1. Surprisingly, \hat{d}_{12}^2 can be shown to be $(1 + o_P(1))/n\pi(1 - \pi)$, which does not depend on the parameters α_n, β_n or γ_n .

5.1. Derivation of the distance formulas for the unnormalized method. We now prove Theorem 3.1.

PROOF OF THEOREM 3.1. We will first prove equation (3.1) and then equation (3.2).

Derivation of \hat{d}_{11}^2 [equation (3.1)]. Define $\hat{\mathbf{r}}_i$ as in equation (5.2). First note that $\|\hat{\mathbf{r}}_i\|^2 = O_P(1/n\rho_n)$ by Lemma 5.2. An argument similar to Lemma 4.1 gives

$$(5.5) \quad \text{For } i \in \{1, 2\}, \quad (A - \hat{\lambda}_i I)\hat{\mathbf{r}}_i = (A - P)\mathbf{v}_i + (\lambda_i - \hat{\lambda}_i)\mathbf{v}_i.$$

As discussed earlier, we have $\hat{\mathbf{r}}_1^T \hat{\mathbf{v}}_2 = \mathbf{v}_1^T \hat{\mathbf{v}}_2$ since $\hat{\mathbf{v}}_1 \perp \hat{\mathbf{v}}_2$. But from Lemma 5.2, we know that $c_{12} = O_P(1/n\rho_n)$, and hence the projection of $\hat{\mathbf{r}}_1$ on the second eigen-space $\hat{\mathbf{v}}_2 \hat{\mathbf{v}}_2^T$ only contributes $\|\hat{\lambda}_2 c_{12} \hat{\mathbf{v}}_2\| = O_P(1)$. As $\hat{\lambda}_3 = O_P(\sqrt{n\rho_n \log n})$, $\|A\hat{\mathbf{r}}_1\| = O_P(1)$.

We compute \hat{d}_{11}^2 by deriving asymptotic expressions of $1/n\pi \sum_{i \in C_1} \hat{\mathbf{r}}_k(i)^2 - \hat{\mathbf{r}}_k(C_1)^2, k \in \{1, 2\}$. First we show that the second term is of lower order than the first. This is because $\sum_{i \in C_1} \hat{\mathbf{r}}_1(i)^2/n\pi \leq \|\hat{\mathbf{r}}_1\|^2/n\pi = O_P(1/n^2\rho_n)$, but $\hat{\mathbf{r}}_1(C_1)^2 = O_P(1/n^3\rho_n^2)$ using Lemma 5.4. We will now focus on the elements of $\hat{\mathbf{r}}_1$ belonging to C_1 . We also denote by $\hat{\mathbf{r}}_1(1)$ the subset of $\hat{\mathbf{r}}_1$ indexed by nodes in C_1 , and thus by $[A\hat{\mathbf{r}}_1](1)$ the subset of vector $A\hat{\mathbf{r}}_1$ indexed by C_1 . Also note that $\|[A\hat{\mathbf{r}}_1](1)\|^2 \leq \|A\hat{\mathbf{r}}_1\|^2 = O_P(1)$

$$[A\hat{\mathbf{r}}_1](1) - \hat{\lambda}_1 \hat{\mathbf{r}}_1(1) = [(A - P)v_1](1) + (\lambda_1 - \hat{\lambda}_1)v_1(1),$$

$$\sum_{i \in C_1} \hat{\mathbf{r}}_1(i)^2 \sim \frac{\sum_{i \in C_1} (x_1 \bar{d}_i^{(1)} + y_1 \bar{d}_i^{(2)})^2}{\lambda_1^2}.$$

The last step is valid because $\|(A - P)v_1\|$ can be shown to be $O_P(\sqrt{n\rho_n})$ (see Section A of the Supplement [22]) whereas $\|[A\hat{\mathbf{r}}_1](1)\| = O_P(1)$ and $\|(\lambda_1 -$

$\hat{\lambda}_1)v_1(1)\| = O_P(1)$ using Lemma 5.2. Similarly, $\sum_{i \in C_1} \hat{r}_2(i)^2 \sim \sum_{i \in C_1} (x_2 \bar{d}_i^{(1)} + y_2 \bar{d}_i^{(2)})^2 / \lambda_2^2$. Hence using Lemma 5.5, equation (5.4), we have

$$\hat{d}_{11}^2 \sim \frac{1}{n\pi} \left[\left(\frac{x_1^2}{\lambda_1^2} + \frac{x_2^2}{\lambda_2^2} \right) \sum_{i \in C_1} (\bar{d}_i^{(1)})^2 + \left(\frac{y_1^2}{\lambda_1^2} + \frac{y_2^2}{\lambda_2^2} \right) \sum_{i \in C_1} (\bar{d}_i^{(2)})^2 \right].$$

Now Lemma 5.5, equation (5.3) yields equation (3.1).

Derivation of \hat{d}_{12}^2 [equation (3.2)]. We recall that equation (2.2) gives $\hat{d}_{12}^2 = \hat{d}_{11}^2 + \|K_1 - K_2\|^2$, where $K_k = \{\hat{v}_1(C_k), \hat{v}_2(C_k)\}$, $k \in \{1, 2\}$. From equation (5.2), we see that $\hat{v}_i(C_1) = (v_i(C_1) - \hat{r}_i(C_1))/c_{ii}$, and hence we have

$$\hat{v}_i(C_1) - \hat{v}_i(C_2) = \left(\frac{x_i - y_i}{c_{ii}} \right) - \left(\frac{\hat{r}_i(C_1) - \hat{r}_i(C_2)}{c_{ii}} \right), \quad i \in \{1, 2\}.$$

We will now show that $\|K_1 - K_2\|^2 = \sum_{i=1}^2 (\hat{v}_i(C_1) - \hat{v}_i(C_2))^2 \sim ((x_1 - y_1)^2 + (x_2 - y_2)^2)$, which is $\sim 1/n\pi(1 - \pi)$, using Lemma 5.1 [equation (5.1)]. Since x_1, y_1, x_2, y_2 are of the form C_0/\sqrt{n} and $c_{ii}^2 = 1 - O_P(1/n\rho_n)$ (Lemma 5.3), we can show that

$$\sum_{i=1}^2 \left(\frac{x_i - y_i}{c_{ii}} \right)^2 = \frac{1 + o_P(1)}{n\pi(1 - \pi)}.$$

Also, for $i \in \{1, 2\}$, $\hat{r}_i(C_1) = O_P(1/n^{3/2}\rho_n)$ (Lemma 5.4), and $c_{ii}^2 = 1 - \|\hat{r}_i\|^2 = O_P(1/n\rho_n)$ (Lemma 5.3), and hence we have equation (3.2). \square

6. Analysis of the normalized method. As discussed in Section 4, both v_1 and \tilde{v}_1 (see Table 1) equal one, and $\tilde{\mathbf{u}}_1(i) = \sqrt{d_i/E}$. In our analysis what naturally appears is the following notion of density, defined by the expected degree over n . All expectations are conditioned on Z . Let μ_1 and μ_2 the $E[d_i|Z]/n$ for i in C_1 and C_2 , respectively. Also let $\mu = \sum_{ij} P_{ij}/n^2$. Hence $\mu_1 := \pi\alpha_n + (1 - \pi)\gamma_n - \alpha_n/n$, and $\mu_2 = (1 - \pi)\beta_n + \pi\gamma_n - \beta_n/n$, and $\mu = \pi\mu_1 + (1 - \pi)\mu_2$. Also, we recall that $\bar{d}_i^{(1)}$ is the centered $d_i^{(1)}$, that is, $d_i^{(1)} - (n\pi - 1)\alpha_n$ when $i \in C_1$ and $d_i^{(1)} - n\pi\gamma_n$ when $i \in C_2$. The properties of the eigen-spectrum of \tilde{P} are stated in the following lemma. Its proof is deferred to Section F of the Supplement [22].

LEMMA 6.1. *Define a semi-sparse stochastic blockmodel (see Definition 2.2) with parameters $(\alpha_n, \beta_n, \gamma_n, Z)$, where $\alpha_n\beta_n \neq \gamma_n^2$. The principal eigenvalues v_1 and v_2 , and the blockwise entries $\tilde{x}_1, \tilde{y}_1, \tilde{x}_2$ and \tilde{y}_2 of the principal eigenvectors of \tilde{P} are given by*

$$\begin{aligned} v_1 &= 1, & \tilde{x}_1 &= \sqrt{\frac{\mu_1}{n\mu}}, & \tilde{y}_1 &= \sqrt{\frac{\mu_2}{n\mu}}, \\ v_2 &= 1 - \frac{\gamma_n\mu}{\mu_1\mu_2}, & \tilde{x}_2 &= \sqrt{\frac{(1 - \pi)\mu_2}{n\pi\mu}}, & \tilde{y}_2 &= -\sqrt{\frac{\pi\mu_1}{n(1 - \pi)\mu}}. \end{aligned}$$

All other eigenvalues of \tilde{P} are $O(1/n)$.

In order to obtain \tilde{d}_{11}^2 [equation (2.1)], we need $\sum_{i \in C_1} (\tilde{u}_1(i) - \tilde{u}_1(C_1))^2$. Using $\tilde{u}(C_1) = \sum_{i \in C_1} \tilde{u}(i)/n\pi$ and arguing as in Lemma 4.3, we see that

$$(6.1) \quad \sum_{i \in C_1} (\tilde{u}_1(i) - \tilde{u}_1(C_1))^2 \sim \frac{1}{4n^3\mu\mu_1} \sum_{i \in C_1} \bar{d}_i^2.$$

Computing $\sum_{i \in C_1} (\tilde{u}_2(i) - \tilde{u}_2(C_1))^2$ requires more in-depth analysis, since $\tilde{\mathbf{u}}_2$ cannot be expressed in closed form as $\tilde{\mathbf{u}}_1$. Instead we look at a “good” approximation of $\tilde{\mathbf{u}}_2$, such that the approximation error cannot mask its $O_P(1/\sqrt{n\rho_n})$ deviation from the population counterpart \mathbf{u}_2 . The very first guess is to construct a vector orthogonal to $\tilde{\mathbf{u}}_1$. In this case, we present \mathbf{u}_g^0 as in equation (6.2). Define $E_1 := \sum_{i \in C_1} d_i$ and $E_2 := \sum_{i \in C_2} d_i$

$$(6.2) \quad \mathbf{u}_g^0(i) = \begin{cases} \frac{\sqrt{d_i}}{E_1}, & \text{for } i \in C_1, \\ -\frac{\sqrt{d_i}}{E_2}, & \text{for } i \in C_2. \end{cases}$$

In spite of being a fair guess, $\mathbf{u}_g^0/\|\mathbf{u}_g^0\|$ masks the $O_P(1/\sqrt{n\rho_n})$ error. So we take a von Mises iteration step starting with \mathbf{u}_g^0 , and get a finer approximation, namely \mathbf{u}_g . We now present element-wise Taylor expansions of \mathbf{u}_g .

LEMMA 6.2. *Define \mathbf{u}_g^0 as in equation (6.2). We have*

$$[\tilde{A}\mathbf{u}_g^0]_i = \begin{cases} \frac{v_2}{n\pi\sqrt{n\mu_1}} \left(1 - \frac{\bar{d}_i}{2n\mu_1} + \frac{\bar{d}_i^{(1)}}{n\mu_1 v_2} - \frac{\bar{d}_i^{(2)}}{n\mu_2 v_2} \frac{\pi}{1-\pi} + M_i \right), & i \in C_1, \\ -\frac{v_2}{n(1-\pi)\sqrt{n\mu_2}} \left(1 - \frac{\bar{d}_i}{2n\mu_2} - \frac{\bar{d}_i^{(1)}}{n\mu_1 v_2} \frac{1-\pi}{\pi} + \frac{\bar{d}_i^{(2)}}{n\mu_2 v_2} + M'_i \right), & i \in C_2. \end{cases}$$

The remainder vectors M and M' are of norm $o_P(C_0/\sqrt{\rho_n})$

$$\|\tilde{A}\mathbf{u}_g^0\| \sim v_2 \sqrt{\frac{\mu}{n^2\pi(1-\pi)\mu_1\mu_2}}.$$

The next lemma shows that \mathbf{u}_g has an approximation error of $O_P(\sqrt{\log n/n^2\rho_n^2})$. The proof again is deferred to Section F of the Supplement [22].

LEMMA 6.3. *Define $\mathbf{u}_g := \tilde{A}\mathbf{u}_g^0/\|\tilde{A}\mathbf{u}_g^0\|$. Let $c_g := (\tilde{\mathbf{u}}_2)^T \mathbf{u}_g$, that is, the projection of \mathbf{u}_g on $\tilde{\mathbf{u}}_2$ and $\mathbf{r}_g := \mathbf{u}_g - c_g \tilde{\mathbf{u}}_2$. We have*

$$\|\mathbf{r}_g\| = O_P\left(\sqrt{\frac{\log n}{n^2\rho_n^2}}\right); \quad c_g = 1 - o_P(1).$$

Now we are ready to derive the expressions of \tilde{d}_{11}^2 and \tilde{d}_{12}^2 (Theorem 3.2).

6.1. *Derivation of the distance formulas for the normalized method.* We now prove Theorem 3.2.

PROOF. We will first prove equation (3.3) and then equation (3.4).

Derivation of \tilde{d}_{11}^2 [equation (3.3)]. Computing \tilde{d}_{11}^2 only involves the entries of $\tilde{\mathbf{u}}_2$ indexed by nodes in C_1 ; hence we will apply Lemma 4.2 on $\tilde{u}_2(i), i \in C_1$. Using our construction,

$$(6.3) \quad \tilde{\mathbf{u}}_2 = (\mathbf{u}_g - \mathbf{r}_g)/c_g \quad \text{where } c_g = 1 - o_P(1).$$

Using Lemma 6.2, for $i \in C_1$, we can write each term of \mathbf{u}_g as

$$u_g(i) = \chi_n(1 + x_n^1(i) + M_i),$$

where x_n^1 and M are the first and remainder terms in the Taylor expansion of $u_g(i)/\chi_n$. We have

$$\begin{aligned} \chi_n &:= \frac{v_2}{n\pi \sqrt{n\mu_1} \|\tilde{A}\mathbf{u}_g^0\|} \sim \sqrt{\frac{(1-\pi)\mu_2}{n\pi\mu}}, \\ x_n^1(i) &:= -\frac{\bar{d}_i}{2n\mu_1} + \frac{\bar{d}_i^{(1)}}{n\mu_1 v_2} - \frac{\bar{d}_i^{(2)}}{n\mu_2 v_2} \frac{\pi}{1-\pi}, \quad i \in C_1. \end{aligned}$$

We have

$$\sum_{i \in C_1} \left(\frac{u_g(i) - u_g(C_1)}{\chi_n} \right)^2 = \sum_{i \in C_1} ((x_n^1(i) - x_n^1(C_1)) + (M_i - M(C_1)))^2.$$

While $\|x_n^1\| = C_0/\sqrt{\rho_n}$ [Lemma 5.5, equation (5.4)], $x_n^1(C_1) = O_P(1/\sqrt{n^2\rho_n})$, since it involves averages of $O(n^2)$ independent Bernoulli random variables. Also $\|M\| = o_P(1/\sqrt{\rho_n})$, and hence using a simple application of the Cauchy–Schwarz inequality, one has

$$(6.4) \quad \sum_{i \in C_1} (u_g(i) - u_g(C_1))^2 \sim \chi_n^2 \sum_{i \in C_1} x_n^1(i)^2.$$

Finally, since $\|\mathbf{r}_g\|^2 = O_P(\log n/(n\rho_n)^2)$ and $\sum_{i \in C_1} (u_g(i) - u_g(C_1))^2 = C_0/n\rho_n$, from equations (6.3) and (6.4), we have

$$(6.5) \quad \frac{1}{n\pi} \sum_{i \in C_1} (\tilde{u}_2(i) - \tilde{u}_2(C_1))^2 \sim \frac{\chi_n^2}{n\pi} \sum_{i \in C_1} x_n^1(i)^2.$$

With a little algebra, equations (6.1) and (6.5) give

$$\begin{aligned} \tilde{d}_{11}^2 &\sim \frac{1}{n\pi} \sum_{i \in C_1} \left(\frac{\mu_1}{n\mu} \frac{\bar{d}_i^2}{4n^2\mu_1^2} + \frac{(1-\pi)\mu_2}{n\pi\mu} \left(-\frac{\bar{d}_i}{2n\mu_1} + \frac{\bar{d}_i^{(1)}}{n\mu_1\nu_2} - \frac{\bar{d}_i^{(2)}}{n\mu_2\nu_2} \frac{\pi}{1-\pi} \right)^2 \right) \\ &\sim \frac{1}{n\pi} \sum_{i \in C_1} \left[\frac{(\bar{d}_i^{(1)})^2}{n^3\pi\mu_1^2} \left(\frac{1}{4} + \frac{(1-\pi)\gamma_n}{\mu_1\nu_2^2} \right) + \frac{(\bar{d}_i^{(2)})^2}{n^3\mu_1^2} \left(\frac{1}{4\pi} + \frac{\pi\alpha_n - \alpha_n/n}{(1-\pi)\mu_2\nu_2^2} \right) \right]. \end{aligned}$$

The last step uses Lemma 5.5 [equation (5.4)].

Derivation of \tilde{d}_{12}^2 [equation (3.4)]. Equation (2.2) gives $\tilde{d}_{12}^2 = \tilde{d}_{11}^2 + \|K_1 - K_2\|$. $K_i := \{\tilde{u}_1(C_i), \tilde{u}_2(C_i)\}$ for $i \in \{1, 2\}$. The Taylor expansion used in Lemma 4.3 shows that the second-order terms are $o_P(1/n)$ whereas the first is of the form C_0/n . For $\mu_1 \neq \mu_2$, neglecting second-order terms gives

$$(6.6) \quad (\tilde{u}_1(C_1) - \tilde{u}_1(C_2))^2 \sim \frac{(\sqrt{\mu_1} - \sqrt{\mu_2})^2}{n\mu}.$$

For the second part, equation (6.3) and an argument shown earlier gives

$$(6.7) \quad \begin{aligned} (\tilde{u}_2(C_1) - \tilde{u}_2(C_2))^2 &\sim (u_g(C_1) - u_g(C_2))^2 \\ &\sim \left(\sqrt{\frac{(1-\pi)\mu_2}{n\pi\mu}} + \sqrt{\frac{\pi\mu_1}{n(1-\pi)\mu}} \right)^2. \end{aligned}$$

Putting equations (6.6) and (6.7) together yields equation (3.4). When $\mu_1 = \mu_2$, the whole contribution comes from the second eigenvector, and \tilde{u}_1 only contributes $o_P(1/n)$ terms. \square

7. Experiments. We demonstrate the benefit of using normalization via classification tasks for simulated networks and link prediction experiments for real world co-authorship networks. For simulations, we investigate the behavior of misclassification error with: (a) a fixed parameter setting with increasing n and (b) changing parameter settings for a fixed n . For all simulations, a pair of training and test graphs are generated from a stochastic blockmodel with a given parameter setting. The model is fitted using spectral clustering (with or without normalization) using the training graph whereas misclassification error is computed using the test graph.

7.1. Simulated networks. For a stochastic blockmodel with $n = 1000$, $\beta_n = \alpha_n$ and $\pi = 1/2$, we focus on the semi-sparse regime, where expected degree is varied from 10–20. We vary $\alpha_n \in [0.01, 0.018]$ (y axis) and $\gamma_n/\alpha_n \in [0.005, 1.2] \setminus \{1\}$ (x axis). The $\gamma_n/\alpha_n = 1$ case causes instability because it reduces the stochastic blockmodel to an Erdős–Rényi graph and hence is excluded. Since `kmeans` can return a local optimum, we run `kmeans` five times and pick the most balanced

clustering, in particular the one whose smallest cluster size is largest among the five runs.

For each of the parameter settings average results from twenty random runs are reported with error bars. In order to ensure that our parameter settings reflect the regime of sparseness required for our theory to hold, we find the connected components of the graph, and only work with those settings where the size of the largest connected component is at least 95% of the size of the whole graph. All computations are carried out on the largest connected component. Therefore we never consider the simple case of disconnected clusters. We also assume that $k = 2$ is known.

In each subfigure of Figure 2 we hold γ_n/α_n fixed and plot the classification errors of the two algorithms along the Y axis against increasing α_n values on the X axis. Across the subfigures γ_n/α_n is increased. Our goal is to turn two knobs to adjust the hardness of the classification problem. If one increases α_n for a fixed value

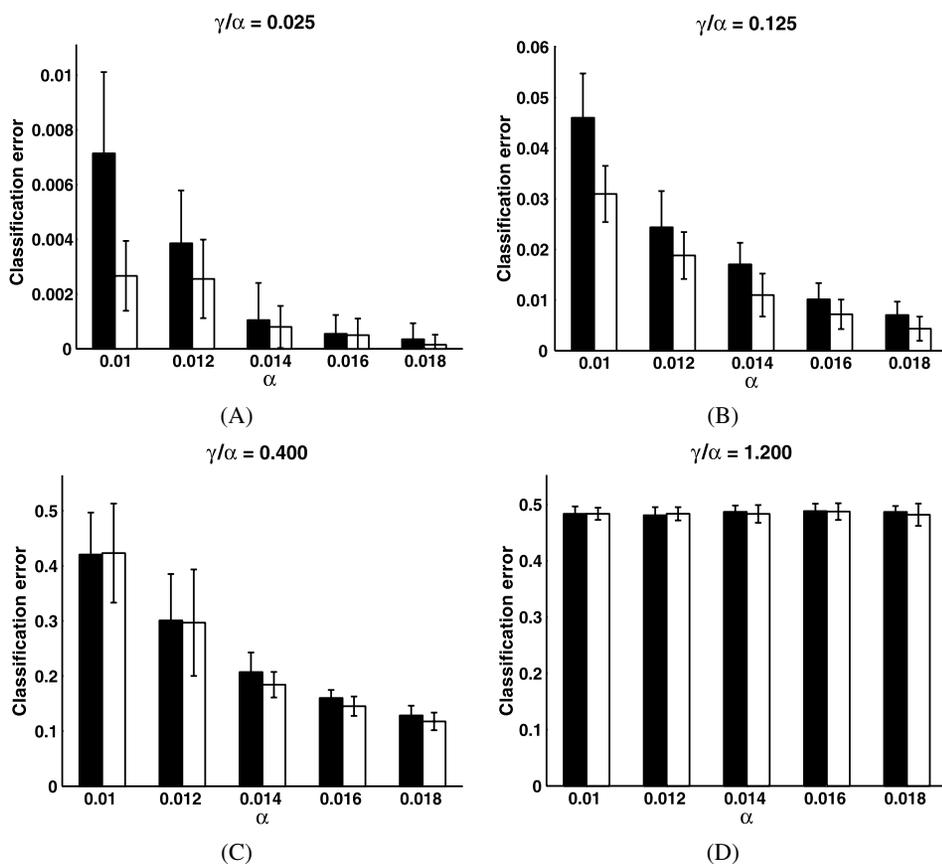


FIG. 2. For a fixed γ_n/α_n ratio miss-classification error is plotted on the Y axis with increasing α_n on the X axis. (A) $\gamma_n/\alpha_n = 0.025$, (B) $\gamma_n/\alpha_n = 0.125$, (C) $\gamma_n/\alpha_n = 0.4$ and (D) $\gamma_n/\alpha_n = 1.2$.

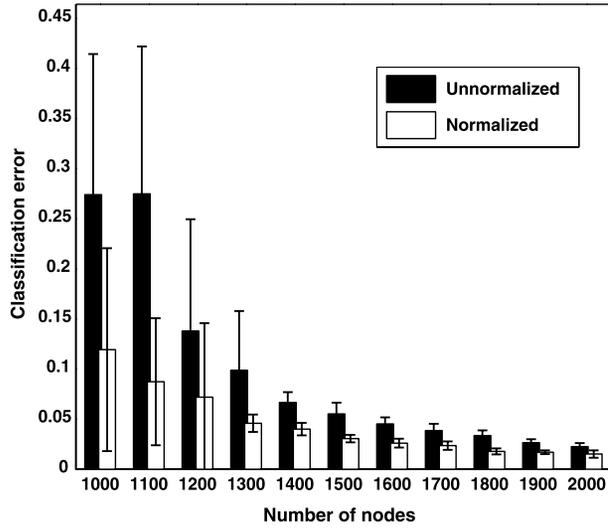


FIG. 3. Miss-classification error on the y axis and increasing n on the x axis.

of γ_n/α_n , then the problem becomes easier as the expected degree increases with increasing α_n . On the other hand, increasing γ_n/α_n makes it hard to distinguish between clusters.

According to our theoretical results, for small γ_n/α_n ratios, normalization performs better clustering under sparsity. In Figure 2(A) and (B), we see that normalization always has a smaller average error, although the difference is more striking for small α_n (average degree about 10). As α_n is increased, both methods start to perform equally well. In Figure 2(C) and (D), γ_n/α_n is larger, and thus the error rates are also larger. In Figure 2(C), both methods behave similarly and show improvement with increasing α_n . Finally in Figure 2(D), both misclassify about half of the nodes since the networks become close to Erdős–Rényi graphs; possibly with more data both methods would perform better. For the second simulation we fix $\alpha_n = \beta_n = 0.01$, $\gamma_n = 0.002$, $\pi = 0.40$. Now in Figure 3 we plot the error bars on classification error from twenty random runs along the Y axis, and n is varied from 1000 to 2000 in the X axis. One can see that the normalized method consistently outperforms the unnormalized method, the margin of improvement being smaller for n (smaller average degree and hence sparser graphs).

7.2. Real world networks. For real world datasets we use co-authorship networks over T timesteps. The nodes represent authors, and edges arise if two authors have co-authored a paper together. Since these networks are unlabeled, we cannot use classification accuracy to measure the quality of spectral clustering. Instead, we choose the task of link prediction to quantitatively assess the goodness of clustering. Since the number of clusters is unknown, we learn k via cross validation. We obtain the training graph (A_1) by merging the first $T-2$ datasets, use

TABLE 2
Table of AUC scores for real data

Dataset	n	Avg. degree	AUC scores with training links included			AUC scores with training links excluded		
			Unnorm.	Norm.	Katz	Unnorm.	Norm.	Katz
HepTH	4795	4.6	0.67	0.82	0.87	0.59	0.79	0.79
NIPS	986	4.4	0.75	0.89	0.75	0.71	0.90	0.69
Citeseer	3857	5.6	0.79	0.96	0.97	0.65	0.93	0.90

the $T-1$ th step (A_2) for cross-validating k and use the last timestep (A_3) as the test graph.

We use a subset of the high energy physics (HepTH) co-authorship dataset ($T = 6$), the NIPS data ($T = 9$) and the Citeseer data ($T = 11$). Each timestep considers 1–2 years of papers (so that the median degree of the test graph is at least 1). In order to match the degree regime of our theory, we remove all nodes with only one neighbor from the training graph, and work with the largest connected component of the resulting network. Cross validation and testing are done on the corresponding subgraphs of $T-1$ and T th timesteps, respectively. The number of nodes and average degrees are reported in Table 2.

In Section 8 we present the misclassification error on the political blogs network. This is possible because the entities are labeled as democratic and republican. We preprocess the network as discussed above, and use $k = 2$.

7.2.1. Link prediction task. First, we learn the $k \times k$ matrix \hat{P} of within and across class probabilities by counting edges between (or across) two clusters. For testing we pick a hundred nodes at random from nodes with at least one neighbor in the test graph. For node i we construct a prediction vector of length n , whose j th entry is the linkage probability \hat{P}_{ab} learned using spectral clustering; here node i belongs to the a th cluster, and node j belongs to the b th cluster. For ground truth we compute the zero one vector representing presence or absence of an edge between nodes i and j from A_3 . These vectors are concatenated to give one prediction vector and the corresponding ground truth.

Now the AUC score of the prediction vector is computed using the ground truth. This is simply the area under the ROC curve obtained by plotting the false positive rate along the x axis and the true positive rate along the y axis. In order to learn k , we vary k from ten to a hundred. For each value of k we estimate \hat{Z} using spectral clustering with k eigenvectors of A_1 (or its normalized counterpart) and then estimate the $k \times k$ conditional probability matrix; now AUC scores are computed using these estimated quantities from A_2 . The k with the largest AUC score is picked and mean AUC scores of five random runs on the test graph using this k is reported.

Since in a co-authorship network, the same edges tend to reappear over time, it is often possible to achieve high scores simply by predicting the edges which are already present in the training data. This is why we examine AUC scores from two experiments:

- (1) training links included in the test graph and
- (2) training links excluded from the test graph.

The second task is harder. We compare our methods with the Katz similarity measure between pairs of nodes [15]. This measure simply computes a weighted sum of number of paths between two nodes, the weights decreasing exponentially as the length of the path grows. It has been shown to give competitive prediction accuracy for link prediction tasks [16]. In both panels, the normalized score performs close to or better than the Katz score, and it outperforms the unnormalized score consistently.

8. Summary and discussion. Normalizing data matrices prior to spectral clustering is a common practice. In this paper we propose a theoretical framework to justify this seemingly heuristic choice. With a series of theoretical arguments, we show that for a large parameter regime, in the context of network blockmodels, normalization reduces the variance of points in a given class under the spectral representation. We also present quantifiable classification tasks on simulated networks and link prediction tasks on real networks to demonstrate that normalization improves prediction accuracy.

While we have only considered two class blockmodels, it should be possible to generalize our proof techniques to a constant number of classes if the population eigenvalues for the normalized or the unnormalized setting are distinct. In order to handle identical eigenvalues, one would need to update the proof techniques so as to argue with eigenspaces instead of individual eigenvectors. However, our simulations (omitted for brevity) show that a result similar in flavor to Figure 1 would hold. However, as γ_n/α_n increases, the ratio grows to one faster compared to that of the two-class blockmodel. In fact, for the real world graphs we learn k by cross validation, and it often exceeds two; our results show that normalization improves link prediction accuracy in these cases as well.

We conclude this paper with a discussion of some practical disadvantages of normalization. Unlike A , all disconnected components contribute eigenvalue one to the eigen-spectrum of \tilde{A} . Thus some of the top eigenvectors of \tilde{A} may have support on a small disconnected component and may be uninteresting. Another problem appears in the presence of small subgraphs weakly connected to the rest of the graph. Here the entries of \tilde{A} corresponding to edges in the subgraph may end up having relatively larger values than the rest of the elements. Hence the top eigenvectors may have high values for nodes in this subgraph leading to poor clustering.

For concreteness let us consider the political blogs network [1], which is a directed network of hyperlinks connecting nodes representing weblogs about US politics. The nodes are labeled as “liberal” and “conservative” blogs. We sym-

metrize the network, remove degree one nodes and find the largest connected component of the remaining network. On this preprocessed network, misclassification rates using spectral clustering for the political blogs dataset are 4% for normalized versus 37% for unnormalized.

If the degree one nodes are not removed prior to finding the largest connected component, then the misclassification error rate is 50% for normalized and 40% for unnormalized. On the other hand, removing degree-one nodes drastically improves the error rate of the normalized method to 4%, while not affecting the unnormalized method's performance significantly. We have also carried out the link prediction experiments without removing the degree-one nodes; the relative behavior of the different algorithms remained essentially unchanged.

In order to alleviate this problem, many regularization approaches [2, 5] have been proposed. These approaches ensure that, with high probability, the eigenvalues corresponding to the discriminating eigenvectors are of larger order than those corresponding to the uninteresting eigenvectors. Further analysis of regularization can be found in [14] and [20].

We want to point out that our analysis is not for a regularized variant of spectral clustering; however, our experiments do have a preprocessing step of operating on the largest connected component after removing low-degree nodes. This can be thought of as a regularizing procedure since this often removes small and weakly connected components and ranks the “useful eigenvectors” higher. In a nutshell, for the normalized method, sparse data artifacts may rank uninteresting eigenvectors high. In this paper we provide theoretical justification for the fact that the discriminating eigenvectors of \tilde{A} are often more useful than those of A .

SUPPLEMENTARY MATERIAL

Supplement to “Role of normalization in spectral clustering for stochastic blockmodels” (DOI: [10.1214/14-AOS1285SUPP](https://doi.org/10.1214/14-AOS1285SUPP); .pdf). Because of space constraints we have moved some of the technical details to the supplementary material [22].

REFERENCES

- [1] ADAMIC, L. A. and GLANCE, N. (2005). The political blogosphere and the 2004 U.S. election: Divided they blog. In *Proceedings of the 3rd Intl. Workshop on Link Discovery*. ACM, New York.
- [2] 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](#)
- [3] 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** 21068–21073.
- [4] BOLLOBÁS, B. (1998). *Modern Graph Theory. Graduate Texts in Mathematics* **184**. Springer, New York. [MR1633290](#)
- [5] CHAUDHURI, K., GRAHAM, F. C. and TSIATAS, A. (2012). Spectral clustering of graphs with general degrees in the extended planted partition model. *Journal of Machine Learning Research—Proceedings Track* **23** 35.1–35.23.

- [6] CHUNG, F. and RADCLIFFE, M. (2011). On the spectra of general random graphs. *Electron. J. Combin.* **18** Paper 215, 14. [MR2853072](#)
- [7] DONATH, W. E. and HOFFMAN, A. J. (1973). Lower bounds for the partitioning of graphs. *IBM J. Res. Develop.* **17** 420–425. [MR0329965](#)
- [8] FEIGE, U. and OFEK, E. (2005). Spectral techniques applied to sparse random graphs. *Random Structures Algorithms* **27** 251–275. [MR2155709](#)
- [9] FIEDLER, M. (1973). Algebraic connectivity of graphs. *Czechoslovak Math. J.* **23** 298–305. [MR0318007](#)
- [10] FÜREDI, Z. and KOMLÓS, J. (1981). The eigenvalues of random symmetric matrices. *Combinatorica* **1** 233–241. [MR0637828](#)
- [11] HAGEN, L. W. and KAHNG, A. B. (1992). New spectral methods for ratio cut partitioning and clustering. *IEEE Trans. on CAD of Integrated Circuits and Systems* **11** 1074–1085.
- [12] HENDRICKSON, B. and LELAND, R. (1995). An improved spectral graph partitioning algorithm for mapping parallel computations. *SIAM J. Sci. Comput.* **16** 452–469. [MR1317066](#)
- [13] HOLLAND, P. W., LASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Social Networks* **5** 109–137. [MR0718088](#)
- [14] JOSEPH, A. and YU, B. (2013). Impact of regularization on spectral clustering. *CoRR*.
- [15] KATZ, L. (1953). A new status index derived from sociometric analysis. *Psychometrika* **18** 39–43.
- [16] LIBEN-NOWELL, D. and KLEINBERG, J. (2003). The link prediction problem for social networks. In *Conference on Information and Knowledge Management ACM*, New York.
- [17] NG, A. Y., JORDAN, M. I. and WEISS, Y. (2001). On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems, Vancouver, British Columbia, Canada*. MIT Press, Cambridge, MA.
- [18] OLIVEIRA, R. I. (2009). Concentration of the adjacency matrix and of the Laplacian in random graphs with independent edges. Preprint.
- [19] POTHEN, A., SIMON, H. D. and LIOU, K.-P. (1990). Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Matrix Anal. Appl.* **11** 430–452. [MR1054210](#)
- [20] QIN, T. and ROHE, K. (2013). Regularized spectral clustering under the degree-corrected stochastic blockmodel. In *Advances in Neural Information Processing Systems, Lake Tahoe, Nevada, USA*. MIT Press, Cambridge, MA.
- [21] ROHE, K., CHATTERJEE, S. and YU, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.* **39** 1878–1915. [MR2893856](#)
- [22] SARKAR, P. and BICKEL, P. J. (2015). Supplement to “Role of normalization in spectral clustering for stochastic blockmodels.” DOI:[10.1214/14-AOS1285SUPP](#).
- [23] SHI, J. and MALIK, J. (2000). Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **22** 888–905.
- [24] SUSSMAN, D. L., TANG, M., FISHKIND, D. E. and PRIEBE, C. E. (2012). A consistent adjacency spectral embedding for stochastic blockmodel graphs. *J. Amer. Statist. Assoc.* **107** 1119–1128. [MR3010899](#)
- [25] VON LUXBURG, U. (2007). A tutorial on spectral clustering. *Stat. Comput.* **17** 395–416. [MR2409803](#)
- [26] VON LUXBURG, U., BELKIN, M. and BOUSQUET, O. (2008). Consistency of spectral clustering. *Ann. Statist.* **36** 555–586. [MR2396807](#)

DEPARTMENT OF STATISTICS AND DATA SCIENCE
 UNIVERSITY OF TEXAS, AUSTIN
 AUSTIN, TEXAS 78712
 USA
 E-MAIL: purna.sarkar@austin.utexas.edu

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