# Bayesian Community Detection 

S. L. van der Pas ${ }^{\ddagger *}$ and A. W. van der Vaart ${ }^{\S \dagger}$


#### Abstract

We introduce a Bayesian estimator of the underlying class structure in the stochastic block model, when the number of classes is known. The estimator is the posterior mode corresponding to a Dirichlet prior on the class proportions, a generalized Bernoulli prior on the class labels, and a beta prior on the edge probabilities. We show that this estimator is strongly consistent when the expected degree is at least of order $\log ^{2} n$, where $n$ is the number of nodes in the network.


MSC 2010 subject classifications: Primary 62F15, 90B15.
Keywords: stochastic block model, community detection, networks, consistency, Bayesian inference, modularities, MAP estimation.

## 1 Introduction

The stochastic block model (SBM) (Holland et al., 1983) is a model for network data in which individual nodes are considered members of classes or communities, and the probability of a connection occurring between two individuals depends solely on their class membership. It has been applied to social, biological and communication networks, for example in Park and Bader (2012), Bickel and Chen (2009) and Snijders and Nowicki (1997) amongst many others. There are many extensions of the SBM for various applications, including the degree-corrected SBM (Karrer and Newman, 2011; Zhao et al., 2012) which accounts for possible heterogeneity among nodes within the same class, and the mixed-membership SBM (Airoldi et al., 2008), in which the assumption that the classes are disjoint is removed. These extensions allow for additional modelling flexibility.

Two main SBM research directions are the recovery of the class labels (community detection) and recovery of the remaining model parameters, consisting of the probability vector generating the class labels, and the class-dependent probabilities of creating an edge between nodes. In this paper, we focus on community detection, noting that once strong consistency of a community detection method has been established, consistency of the natural plug-in estimators for the remaining parameters follows directly by results in (Channarond et al., 2012).

A large number of methods for recovering the class labels has been proposed. Those most closely related to this work are the modularities. Newman and Girvan (2004) introduced the term modularity for 'a measure of the quality of a particular division of a network'. They described one such measure for models in which edges are more

[^0]likely to occur within classes than between classes, in which case there is a community structure in the colloquial sense, although the SBM does not require this assumption. Bickel and Chen (2009) studied more general modularities, defining them as functions of the number of connections between all combinations of classes and the proportion of nodes placed in each class. They introduced the likelihood modularity, and provided general conditions under which modularities are consistent. Their method and theory was extended to the degree-corrected SBM by Zhao et al. (2012).

Spectral methods for community detection have gained in popularity, and refined results on error bounds are now available for the SBM and extensions of the SBM, as evidenced in Rohe et al. (2011), Jin (2015), Sarkar and Bickel (2015) and Lei and Rinaldo (2015) for example. Many other algorithms have been introduced, most of them currently lacking formal proofs of consistency. A notable exception is the Largest Gaps algorithm (Channarond et al., 2012), which only takes the degree of each node as its input, and is strongly consistent under a separability condition.

A Bayesian approach towards recovering the class assignments in the SBM was first suggested by Snijders and Nowicki (1997), motivated by computational advantages of Gibbs sampling over maximum likelihood estimation. They considered two classes and proposed uniform priors on the class proportions and the edge probabilities. This approach was extended in (Nowicki and Snijders, 2001) to allow for more classes, with a Dirichlet prior on the class proportions and beta priors on the edge probabilities. Hofman and Wiggins (2008) described a similar Bayesian approach for a special case of the SBM and suggested a variational approach to overcome the computational issues associated with maximizing over all possible class assignments.

Bayesian methods for the SBM have barely been studied from a theoretical point of view, although recent results for parameter recovery by Pati and Bhattacharya (2015), for detecting the number of communites by Hayashi et al. (2016) and for an empirical Bayes approach to community detection by Suwan et al. (2016) are encouraging. In this work, we provide theoretical results on community detection, establishing that the Bayesian posterior mode is strongly consistent for the class labels if the expected degree is at least of order $\log ^{2} n$, where $n$ is the number of nodes. This is proven by relating the posterior mode to the maximizer of the likelihood modularity of Bickel and Chen (2009). The likelihood modularity has been claimed to be strongly consistent under the weaker assumption that the expected degree is of larger order than $\log n$ (Bickel and Chen, 2009; Zhao et al., 2012; Bickel et al., 2015). However, their proof assumes that the likelihood modularity is globally Lipschitz, while it is only locally so. The Bayesian method is based on a combination of likelihood and prior, and for this reason the proof of our main theorem, Theorem 1, runs into a similar problem. We were able to resolve this only under the slightly stronger assumption that the expected degree is of larger order than $(\log n)^{2}$. The literature on other methods for community detection shows that the order $\log n$ is sufficient for consistent detection. However, these results are usually obtained under additional assumptions such as a restriction to two classes or an ordering of the connection probabilities, and their implications for the likelihood or Bayesian modularities is unclear. We discuss this and the relevant literature further following the statement of our main result in Section 3.5.

The main result of the present paper is that the posterior mode is strongly consistent in the frequentist setup, a property that it shares with the maximizer of the likelihood modularity. As the number of parameters of the model ("labellings") increases rapidly with $n$, this result is certainly not covered by standard theory for parametric models, and in fact we shall see that the prior on the labellings plays a special role for consistency. That the posterior mode behaves well in terms of consistency is encouraging, and makes one hope that other aspects of the posterior distribution will also be useful for inference. The present paper may be considered a first step and further study of such aspects is desirable. One possible research direction would be to use the full posterior distribution on the labels to quantify uncertainty in the estimate of the class labels. A second issue that may be resolved by the Bayesian approach is the question of estimating the number of classes, $K$. This remains an important open question, as noted by Bickel and Chen (2009), despite recent attempts (e.g. Saldana et al. (2014), Chen and Lei (2014) and Wang and Bickel (2015)). By introducing a prior on $K$, such as the Poissonprior suggested by McDaid et al. (2013), the number of communities $K$ can be detected by the posterior. A third open question is whether the Bayesian estimator can be improved by incorporating prior knowledge of the community structure. Recent work on incorporating prior information in Gaussian graphical models (Kpogbezan et al., 2016) is encouraging, and has not been translated to the SBM yet.

This paper is organized as follows. We introduce the SBM and the associated notation in Section 2. Our main results are in Section 3, where we describe the prior and the link with the likelihood modularity, present the consistency results and discuss the underlying assumptions, especially those on the expected degree. After an illustration of the method on a data set in Section 4, we conclude with the proofs, first of weak consistency in Section 5 and finally of strong consistency of the Bayesian modularity in Section 6.

### 1.1 Notation

For a vector $v$ we denote by $\operatorname{Diag}(v)$ the diagonal matrix with diagonal $v$, and for a matrix $M$ we denote its diagonal by $\operatorname{diag}(M)$.

The $\|\cdot\|_{1}$-norm of a matrix $M$ is the sum of the absolute values of all entries of $M$.
We write $f(n)=\mathcal{O}(g(n))$ as $n \rightarrow \infty$ if there exist $C, n_{0}>0$ such that $|f(n)| \leq$ $C|g(n)|$ for all $n>n_{0}$.

## 2 The stochastic block model

We introduce the notation and generative model for the $\operatorname{SBM}$ with $K \in\{1,2, \ldots\}$ classes. Consider an undirected random graph with $n$ nodes, numbered $1,2, \ldots, n$, and edges encoded by the $n \times n$ symmetric adjacency matrix $\left(A_{i j}\right)$, with entries in $\{0,1\}$. Thus $A_{i j}=A_{j i}$ is equal to 1 or 0 if the nodes $i$ and $j$ are or are not connected by an edge, respectively. Self-loops are not allowed, so $A_{i i}=0$ for $i=1, \ldots, n$. The generative model for the random graph is:

1. The nodes are randomly labeled with i.i.d. variables $Z_{1}, \ldots, Z_{n}$, taking values in a finite set $\{1, \ldots, K\}$, according to probabilities $\pi=\left(\pi_{1}, \ldots, \pi_{K}\right)$.
2. Given $Z=\left(Z_{1}, \ldots, Z_{n}\right)^{T}$, the edges are independently generated as Bernoulli variables with $\mathbb{P}\left(A_{i j}=1 \mid Z\right)=P_{Z_{i}, Z_{j}}$, for $i<j$, for a given $K \times K$ symmetric matrix $P=\left(P_{a b}\right)$.

The probability vector $\pi$ is considered fixed, but unknown. Although this is not visible in the notation, the matrix $P$ may change with $n$, a case of particular interest being that $P$ tends to zero, which gives a sparse graph. The order of magnitude of $\|P\|_{\infty}=\max _{a, b} P_{a b}$ is the same as the order of magnitude of $\rho_{n}=\sum_{a, b} \pi_{a} \pi_{b} P_{a b}$, the probability of there being an edge between two randomly selected nodes. The expected degree of a randomly selected node is $\lambda_{n}=(n-1) \rho_{n}$, and twice the expected total number of edges in the network is $\mu_{n}=n(n-1) \rho_{n}$.

The likelihood for the model is given by

$$
\begin{equation*}
\prod_{i<j} P_{Z_{i} Z_{j}}^{A_{i j}}\left(1-P_{Z_{i} Z_{j}}\right)^{1-A_{i j}} \prod_{i} \pi_{Z_{i}}=\prod_{a \leq b} P_{a b}^{O_{a b}(Z)}\left(1-P_{a b}\right)^{n_{a b}(Z)-O_{a b}(Z)} \prod_{a} \pi_{a}^{n_{a}(Z)} \tag{1}
\end{equation*}
$$

where $O_{a b}(Z)$ is the number of edges between nodes labelled $a$ and $b$ by the labelling $Z$, $n_{a b}(Z)$ is the maximum number of edges that can be created between nodes labelled $a$ and $b$, and $n_{a}(Z)$ is the number of nodes labelled $a$, and $a$ and $b$ range over $\{1,2, \ldots, K\}$.

More formally, for a given labelling $e=\left(e_{1}, \ldots, e_{n}\right)^{T} \in\{1, \ldots, K\}^{n}$ of nodes, and class labels $a, b \in\{1, \ldots, K\}$, we define

$$
\begin{aligned}
O_{a b}(e) & = \begin{cases}\sum_{i, j} A_{i j} \mathbf{1}_{\left\{e_{i}=a, e_{j}=b\right\}}, & a \neq b, \\
\sum_{i<j} A_{i j} \mathbf{1}_{\left\{e_{i}=a, e_{j}=b\right\}}, & a=b,\end{cases} \\
n_{a b}(e) & = \begin{cases}n_{a}(e) n_{b}(e), & a \neq b, \\
\frac{1}{2} n_{a}(e)\left(n_{a}(e)-1\right), & a=b,\end{cases} \\
n_{a}(e) & =\sum_{i=1}^{n} \mathbf{1}_{\left\{e_{i}=a\right\}} .
\end{aligned}
$$

Since the matrix $A$ is symmetric with zero diagonal by assumption, for $a \neq b$ the variable $O_{a b}(e)$ can also be written as $\sum_{i<j} A_{i j}\left[\mathbf{1}_{\left\{e_{i}=a, e_{j}=b\right\}}+\mathbf{1}_{\left\{e_{j}=a, e_{i}=b\right\}}\right]$, which explains the different appearances of the diagonal and off-diagonal entries. The numbers $n_{a b}(e)$ are equal to the numbers $O_{a b}(e)$ when all $A_{i j}$ are equal to 1 . We collect the variables $O_{a b}(e)$ and $n_{a b}(e)$ in $K \times K$ matrices $O(e)$ and $n(e)$.

Now consider the $K \times K$ probability matrix $R(e, c)$ and $K$ probability vector $f(e)$ with entries

$$
\begin{equation*}
R_{a b}(e, c)=\frac{1}{n} \sum_{i=1}^{n} \mathbf{1}_{\left\{e_{i}=a, c_{i}=b\right\}}, \quad \quad f_{a}(e)=\frac{n_{a}(e)}{n} \tag{2}
\end{equation*}
$$

The row sums of $R(e, c)$ are equal to $R(e, c) \mathbf{1}=f(e)$, while the column sums are equal to $\mathbf{1}^{T} R(e, c)=f(c)^{T}$. Thus, the matrix $R(e, c)$ can be seen as a coupling of the marginal
probability vectors $f(e)$ and $f(c)$. If $e=c$, then it is diagonal with diagonal $f(c)=$ $f(e)$. More generally, the matrix can be viewed as measuring the discrepancy between labellings $e$ and $c$. This can be precisely measured as half the $L_{1}$-distance of $R(e, c)$ to its diagonal, as evidenced by Lemma 1, which is noted in Bickel and Chen (2009).

Recall that by $\|M\|_{1}$ we denote the sum of the absolute values of all entries of a matrix $M$.

Lemma 1. For every labelling $c, e$ in the $K$-class stochastic block model:

$$
\frac{1}{n} \sum_{i=1}^{n} \boldsymbol{1}_{\left\{c_{i} \neq e_{i}\right\}}=\frac{1}{2}\|\operatorname{Diag}(f(c))-R(e, c)\|_{1}
$$

Proof. The diagonal of $R(e, c)$ gives the fractions of labels on which $c$ and $e$ agree. Hence the left side of the lemma is $1-\sum_{a} R_{a a}(e, c)=\sum_{a}\left(f_{a}(c)-R_{a a}(e, c)\right)$. The elements of both $K \times K$ matrices $\operatorname{Diag}(f(c))$ and $R(e, c)$ can be viewed as probabilities that add up to 1 . Thus the sum of the differences of the diagonal elements is minus the sum of the differences of the off-diagonal elements. Because $f_{a}(c) \geq R_{a a}(e, c)$ for every $a$, we have $\sum_{a}\left(f_{a}(c)-R_{a a}(e, c)\right)=\sum_{a}\left|f_{a}(c)-R_{a a}(e, c)\right|$. Similarly the off-diagonal elements of $\operatorname{Diag}(f(c))$, which are zero, are smaller than the off-diagonal elements of $R(e, c)$ and hence we can add absolute values. Thus the sum over the diagonal is half the sum of the absolute values of all terms in $\operatorname{Diag}(f(c))-R(e, c)$.

## 3 Bayesian approach to community detection

Our main results are presented in this section. We first discuss the choice of prior in Section 3.1, and define the estimator, in Section 3.2. The resulting Bayesian modularity is closely related to the likelihood modularity of Bickel and Chen (2009). The relationship is clarified in Section 3.3. We briefly consider the issue of identifiability in the SBM in Section 3.4, and conclude with our main theorem on the strong consistency of the Bayesian modularity in Section 3.5.

### 3.1 The prior

We adopt the Bayesian approach of Nowicki and Snijders (2001). We put prior distributions on the parameters of the stochastic block model with $K$ known, the vector $\pi$ and the matrix $P$, yielding a joint probability distribution of $(A, Z, \pi, P)$. Next we marginalize over $\pi$ and $P$ as in McDaid et al. (2013), leading to a joint distribution of $(A, Z)$. Finally we "estimate" the unobserved vector $Z$ by the posterior mode of the conditional distribution of $Z$ given $A$. From a frequentist point of view this means that $Z$ is treated as a parameter of the problem, equipped with a hierarchical prior that chooses first $\pi$ and then $Z$. Accordingly we shall change notation from $Z$ to $e$, reserving $Z$ for the frequentist description of the stochastic block model in Section 2.

The prior on $\pi$ is a Dirichlet, and independently the $P_{a b}$ for $a \leq b$ receive independent beta priors:

$$
\begin{aligned}
\pi & \perp\left(P_{A b}\right) \\
\pi & \sim \operatorname{Dir}(\alpha, \ldots, \alpha) \\
P_{a b} & \stackrel{i . i . d .}{\sim} \operatorname{Beta}\left(\beta_{1}, \beta_{2}\right), \quad 1 \leq a \leq b \leq K .
\end{aligned}
$$

This is essentially the same set-up as in Nowicki and Snijders (2001) and McDaid et al. (2013), except that we use a more flexible $\operatorname{Beta}\left(\beta_{1}, \beta_{2}\right)$ instead of a uniform prior on the $P_{a b}$. We assume $\alpha, \beta_{1}, \beta_{2}>0$.

We complete the Bayesian model by specifying class labels $e=\left(e_{1}, \ldots, e_{n}\right)$ and edges $A=\left(A_{i j}: i<j\right)$ through

$$
\begin{aligned}
& \quad e_{i} \mid \pi, P \stackrel{i . i . d .}{\sim} \pi, \quad 1 \leq i \leq n, \\
& A_{i j} \mid \pi, P, e \stackrel{i n d .}{\sim} \operatorname{Bernoulli}\left(P_{e_{i}, e_{j}}\right), \quad 1 \leq i<j \leq n .
\end{aligned}
$$

Abusing notation we write $p(e), p(A \mid e)$ and $p(e \mid A)$ for marginal and conditional probability density functions.

### 3.2 The Bayesian modularity

The Bayesian estimator of the class labels will be the posterior mode, that is:

$$
\widehat{e}=\underset{e}{\operatorname{argmax}} p(e \mid A) .
$$

The posterior mode can be interpreted as a modularity-based estimator in the sense of Bickel and Chen (2009), in that it maximizes a function that only depends on the $O_{a b}(e)$ and the $n_{a}(e)$. This can be seen from the joint density of $(A, e)$, which is found by marginalizing the likelihood (1) over $\pi$ and $P$. The conjugacy between the multinomial and Dirichlet distributions gives the marginal density of the class assignment $e$ as:

$$
\begin{equation*}
p(e)=\int_{S_{K}} \prod_{a} \pi_{a}^{n_{a}(e)} \frac{\prod_{a} \pi_{a}^{\alpha-1}}{D(\alpha)} d \pi=\frac{\Gamma(\alpha K)}{\Gamma(\alpha)^{K} \Gamma(n+\alpha K)} \prod_{a} \Gamma\left(n_{a}(e)+\alpha\right) \tag{3}
\end{equation*}
$$

Here the integral is relative to the Lebesgue measure on the $K$-dimensional unit simplex and $D(\alpha)=\Gamma(\alpha)^{K} / \Gamma(K \alpha)$ is the norming constant for the Dirichlet density. Similarly the conjugacy between the Bernoulli and Beta distributions gives the marginal conditional density of $A$ given $e$ as:

$$
\begin{align*}
p(A \mid e) & =\int_{[0,1]^{K(K+1) / 2}} \prod_{a \leq b} P_{a b}^{O_{a b}(e)}\left(1-P_{a b}\right)^{n_{a b}(e)-O_{a b}(e)} \prod_{a \leq b} \frac{P_{a b}^{\beta_{1}-1}\left(1-P_{a b}\right)^{\beta_{2}-1}}{B\left(\beta_{1}, \beta_{2}\right)} d P \\
& =\prod_{a \leq b} \frac{1}{B\left(\beta_{1}, \beta_{2}\right)} B\left(O_{a b}(e)+\beta_{1}, n_{a b}(e)-O_{a b}(e)+\beta_{2}\right) \tag{4}
\end{align*}
$$

where $B(x, y)=\Gamma(x) \Gamma(y) / \Gamma(x+y)$ is the beta-function. The joint density of $A$ and $e$ is given by the product of (3) and (4), and $n^{-2}$ times its logarithm is up to a constant that is free of $e$ equal to
$Q_{B}(e)=\frac{1}{n^{2}} \sum_{1 \leq a \leq b \leq K} \log B\left(O_{a b}(e)+\beta_{1}, n_{a b}(e)-O_{a b}(e)+\beta_{2}\right)+\frac{1}{n^{2}} \sum_{a=1}^{K} \log \Gamma\left(n_{a}(e)+\alpha\right)$.

This is a modularity in the sense of Bickel and Chen (2009), which we define as the Bayesian modularity. As $p(e \mid A)$ is proportional to $p(e, A)$, the posterior mode is equal to the class assignment that maximizes the Bayesian modularity, so the Bayesian estimator is equal to:

$$
\begin{equation*}
\widehat{e}=\underset{e}{\operatorname{argmax}} Q_{B}(e) . \tag{6}
\end{equation*}
$$

### 3.3 Similarity to the likelihood modularity

The Bayesian modularity $Q_{B}(e)$ consists of two parts, originating from the likelihood and the prior on the classification, respectively. The first part is close to the likelihood modularity given by

$$
Q_{M L}(e)=\frac{1}{n^{2}} \sum_{1 \leq a \leq b \leq K} n_{a b}(e) \tau\left(\frac{O_{a b}(e)}{n_{a b}(e)}\right)
$$

where $\tau(x)=x \log x+(1-x) \log (1-x)$. This criterion, obtained in Bickel and Chen (2009), results from replacing in the log conditional likelihood of $A$ given $e$ (the logarithm of (1) with $Z$ replaced by $e$ and discarding the term involving the parameters $\pi_{a}$ ) the parameters $P_{a b}$ by their maximum likelihood estimators $\hat{P}_{a b}=O_{a b}(e) / n_{a b}(e)$. In other words, the parameters are profiled out rather than integrated out as for the Bayesian modularity. The corresponding estimator

$$
\widehat{e}_{M L}=\underset{e}{\operatorname{argmax}} Q_{M L}(e)
$$

is consistent, and hence one may hope that the Bayesian estimator can be proved consistent by showing that the Bayesian and likelihood modularities are close. This will indeed be our line of approach, but we shall see that the proximity of the two criteria is not close enough to explain the strong consistency of the two methods. In particular, the second, prior part of the Bayesian modularity, resulting from the prior density (3) over the labels, does play a role in the proof of strong consistency. We discuss this in more detail at the end of Section 3.5.

The following lemma links the Bayesian and likelihood modularities. The final assertion shows that they are at most of the order $\log n / n$ apart, which will be seen to be enough in the proof of weak consistency. For the proof of strong consistency we shall need the first assertion of the lemma, which makes the discrepancy between the two modularities explicit up to order $\log n / n^{2}$.

Lemma 2. There exists a constant $C$ such that, for $\mathcal{E}=\{1, \ldots, K\}^{n}$ the set of all possible labellings:

$$
\max _{e \in \mathcal{E}}\left|Q_{B}(e)-Q_{M L}(e)-Q_{P}(e)\right| \leq \frac{C \log n}{n^{2}}
$$

for

$$
Q_{P}(e)=\frac{1}{n^{2}} \sum_{a: n_{a}+\lfloor\alpha\rfloor \geq 2} n_{a}(e) \log \left(n_{a}(e)\right)-\frac{1}{n}
$$

Consequently $\max _{e \in \mathcal{E}}\left|Q_{B}(e)-Q_{M L}(e)\right|=\mathcal{O}(\log n / n)$ as $n \rightarrow \infty$.

### 3.4 Identifiability and consistency

A classification $\widehat{e}$ is said to be weakly consistent if the fraction of misclassified nodes tends to zero (partial recovery), and strongly consistent if the probability of misclassifying any of the nodes tends to zero (exact recovery). In defining consistency in a precise manner, the complication of the possible unidentifiability of the labels needs to be dealt with. From the observed data $A$ we can at best recover the partition of the $n$ nodes in the $K$ classes with equal labels $Z_{i}$, but not the values $Z_{1}, \ldots, Z_{n}$ of the labels, in the set $\{1,2, \ldots, K\}$, attached to the classes. Thus consistency will be up to a permutation of labels.

To make this precise define, for a given permutation $(1, \ldots, K) \rightarrow(\sigma(1), \ldots, \sigma(K))$, the permutation matrix $P_{\sigma}$ as the matrix with rows

$$
\begin{gathered}
e_{\sigma(1)}^{T} \\
\vdots \\
e_{\sigma(K)}^{T},
\end{gathered}
$$

for $e_{1}, \ldots, e_{K}$ the unit vectors in $\mathbb{R}^{K}$. Then pre-multiplication of a matrix by $P_{\sigma}$ permutes the rows, and post-multiplication by $P_{\sigma}^{T}$ the columns: $P_{\sigma} R$ is the matrix with $j$ th row equal to the $\sigma(j)$ th row of $R$, and $R P_{\sigma}^{T}$ is the matrix with $j$ th column the $\sigma(j)$ th column of $R$. Thus $P_{\sigma} R(e, Z)=R\left(P_{\sigma} e, Z\right)$ is the matrix that would result if we would permute the labels of the classes of the assignment $e$, and $P_{\sigma} P P_{\sigma}^{T}$ and $P_{\sigma} R(e, Z) P_{\sigma}^{T}=R\left(P_{\sigma} e, P_{\sigma} Z\right)$ are the matrices that would result if we would relabel the classes throughout. Since we cannot recover the labels, the matrix $P_{\sigma} R(e, Z)$ is just as good or bad as $R(e, Z)$ for measuring discrepancy between a labelling $e$ and the true labelling $Z$; furthermore, nothing should change if we choose different names for the classes.

Thus, taking into account the unidentifiability of the labels, by Lemma 1, we define an estimator $\widehat{e}$ to be weakly consistent if

$$
\left\|P_{\sigma} R(\widehat{e}, Z)-\operatorname{Diag}(f(Z))\right\|_{1} \rightarrow 0
$$

for some permutation matrix $P_{\sigma}$. We say the classification $\widehat{e}$ is strongly consistent if

$$
\mathbb{P}\left(P_{\sigma} R(\widehat{e}, Z)=\operatorname{Diag}(f(Z))\right) \rightarrow 1
$$

for some permutation matrix $P_{\sigma}$.
The following lemma shows that the permutation matrix $P_{\sigma}$ is for large $n$ uniquely defined, unless there are empty classes.

Lemma 3. If for a given vector $\pi$ and matrix $R$, there exist permutation matrices $P_{\sigma}$ and $Q_{\sigma}$ such that both $\left\|P_{\sigma} R-\operatorname{Diag}(\pi)\right\|_{1} \leq \min _{a} \pi_{a}$ and $\left\|Q_{\sigma} R-\operatorname{Diag}(\pi)\right\|_{1} \leq \min _{a} \pi_{a}$, then $P_{\sigma}=Q_{\sigma}$.

Proof. Because the $L_{1}$-norm is invariant under permutations, we have $\left\|R-P_{\sigma}{ }^{-1} \operatorname{Diag}(\pi)\right\|_{1} \leq \min _{a} \pi_{a}$, and similarly for $Q_{\sigma}$. Therefore $\| P_{\sigma}^{-1} \operatorname{Diag}(\pi)-$ $Q_{\sigma}^{-1} \operatorname{Diag}(\pi) \|_{1} \leq 2 \min _{a} \pi_{a}$, by the triangle inequality. Again by invariance, the left side of this inequality is equal to $\left\|\left(Q_{\sigma} P_{\sigma}^{-1}\right) \operatorname{Diag}(\pi)-\operatorname{Diag}(\pi)\right\|_{1}$, which is at least two times the sum of the two smallest coordinates of $\pi$ if $Q_{\sigma} P_{\sigma}^{-1}$ is not equal to the identity matrix.

A necessary requirement for consistency is that the classes can be recovered from the likelihood, i.e. the model parameters must be identifiable. If $\pi$ has strictly positive coordinates, so that all labels will appear in the data eventually, then as explained in Bickel and Chen (2009) an appropriate condition is that $P$ does not have two identical rows. If $\pi_{a}=0$ for some $a$, then class $a$ will never be consumed; the identifiability condition should then be imposed after deleting the $a$ th column from $P$. Thus, we call the pair $(P, \pi)$ identifiable if the rows of $P$ are different after removing the columns corresponding to zero coordinates of $\pi$. Throughout we assume that $P$ is symmetric.

### 3.5 Consistency results and assumptions

We are now ready to present our results on consistency for the Bayesian maximum a posteriori (MAP) estimator (6). Recall that $\rho_{n}=\sum_{a, b} \pi_{a} \pi_{b} P_{a b}$ is the probability of a new edge, and $\lambda_{n}=(n-1) \rho_{n}$ is the expected degree of a node. Theorem 1 shows strong consistency of the Bayesian estimator if $\lambda_{n} \gg(\log n)^{2}$. The proof rests on a proof of weak consistency under similar conditions, stated in Section 5 as Theorem 2.

Theorem 1 (strong consistency). If $P=\rho_{n} S$, where either $\rho_{n}=1$ is fixed or $\rho_{n} \rightarrow 0$, and $(S, \pi)$ is fixed and identifiable with all entries of $P$ strictly smaller than 1 and all entries of $S$ being strictly positive, then the MAP classifier $\widehat{e}=\arg \max _{e} Q_{B}(e)$ is strongly consistent if $\rho_{n} \gg(\log n)^{2} / n$.

The theorem is proven in two steps: first for the dense case, where $\rho_{n}$ is fixed, and then for the sparse case, where $\rho_{n}$ goes to zero. The second is the most interesting of the two, as it touches on the question how much information is required to recover the underlying community structure. Much recent research effort has gone into determining detection and computational boundaries, in particular for special cases of the SBM with $K=2$ (see e.g. Mossel et al. (2012), Chen and Xu (2014), Abbe et al. (2014) and Zhang and Zhou (2015)).

Weakly consistent estimation of the class labels for an arbitrary, but known, number of classes is possible by some method under the assumption $\lambda_{n} \gg \log n$, as this was shown to hold for spectral clustering by Lei and Rinaldo (2015). Strong consistency of maximum likelihood was shown to hold in the special cases of planted bisection ( $K=2$ and equal community sizes) and planted clustering (equal community sizes and
$P_{a b}$ can take two values) by Abbe et al. (2014); Chen and Xu (2014), again under the assumption $\lambda_{n} \gg \log n$. Gao et al. (2015) and Gao et al. (2016) achieve optimality in different senses, under assumptions on the average within-community and betweencommunity edge probabilities; Gao et al. (2015) introduce a two-stage procedure which achieves the optimal proportion of misclassified nodes in a special case where $P_{a b}$ can only take two values, while Gao et al. (2016) obtain minimax rates for the proportion of misclassified nodes in the degree corrected SBM.

Strong consistency of the likelihood modularity for an arbitrary number of classes $K$ has been claimed under the same assumption $\lambda_{n} \gg \log n$ (Bickel and Chen, 2009; Bickel et al., 2015), and those results have been extended to the degree-corrected SBM (Zhao et al., 2012). However, these results were obtained by application of an abstract theorem to the special case of the likelihood modularity, which would require the function $\tau(x)=$ $x \log x+(1-x) \log (1-x)$, or the function $\sigma(x)=x \log x$, to be globally Lipschitz. As $\tau$ and $\sigma$ are only locally Lipschitz, it is still unclear whether $\lambda_{n} \gg \log n$ is a sufficient condition for either weakly or strongly consistent estimation by maximum likelihood. From our proof of Theorem 1, which proceeds by comparing the Bayesian modularity and the likelihood modularity, it follows that $\lambda_{n} \gg(\log n)^{2}$ is certainly sufficient. Given weak consistency the problem can be reduced to a neighbourhood of the true parameter on which the Lipschitz condition is satisfied. However, it is precisely our proof of weak consistency that needs the additional $\log n$ factor.

The Largest Gaps algorithm of Channarond et al. (2012) is strongly consistent provided that $\min _{a \neq b}\left|\sum_{k=1}^{K} \alpha_{k}\left(P_{a k}-P_{b k}\right)\right|$ is at least of order $\sqrt{\log n / n}$, implying that at least one of the $P_{a b}$ is of the same order, and thus $\lambda_{n} \gg \sqrt{n \log n}$. This much stronger condition is not surprising, as the Largest Gaps algorithm only uses the degree of a node and does not take into account any finer information on the group structure, such as the information contained in the $O_{a b}$.

To the best of our knowledge, for $K>2$, it remains to be shown that $\lambda_{n} \gg \log n$ is sufficient for strong consistency of any community detection method for the general SBM. For the minimax rate for the proportion of misclustered nodes in community detection, when only classes of sizes proportional to $n$ are considered, a phase transition when going from the case $K=2$ to $K \geq 3$ was observed by Zhang and Zhou (2015). Their results show that if $K=2$, communities of the same size are most difficult to distinguish, while if $K \geq 3$, small communities are harder to discover. This shift in the nature of the communities that are harder to detect may be what has been preventing a general strong consistency result under the assumption $\lambda_{n} \gg \log n$ so far.

While the prior on the class assignment plays a negligible role in the proof of weak consistency, our argument for strong consistency requires that the prior does not vary too much in a neighborhood of the truth. To be precise, denote by $Q_{B, 2}(e)=$ $n^{-2} \sum_{a=1}^{K} \log \Gamma\left(n_{a}(e)+\alpha\right)$ the second part of the Bayesian modularity (5), and let $Z$ be the true labelling. Then we need that for any $e$ that differs from $Z$ by at most $m$ nodes, the distance $\left|Q_{B, 2}(e)-Q_{B, 2}(Z)\right|$ is of smaller order than $m / n$. We thus find that a variation on general posterior contraction results (e.g. Ghosal et al. (2000)) holds for the SBM as well, namely that the prior mass should be spread homogeneously in a neighborhood of the truth.

The number $K$ of classes is held fixed in the preceding theorem. Our proofs suggest that consistency is retained if $K=K_{n} \rightarrow \infty$ and $\rho_{n} \gg K_{n}^{4}\left(\log K_{n}\right) n^{-1}\left(\log \left(n / \log K_{n}\right)\right)^{2}$, provided the model is asymptotically identifiable in a suitable sense. In Theorem 1 identifiability in the case of fixed $K$ is described as a property of the pair $(S, \pi)$. If $K_{n} \rightarrow \infty$, then the dimensions of these objects tend to infinity and identifiability must be defined in a different way. Our proofs suggest that a crucial quality is $\sum_{a} \pi_{a} K_{0}\left(S_{a b^{\prime}} \| S_{a, b}\right)$, where $K_{0}\left(s \| s^{\prime}\right)$ is the Kullback-Leibler divergence between two Poisson distributions with means $s$ and $s^{\prime}$. As seen in the proof of Lemma 11, this quantity drives local identifiability. It seems a reasonable assumption that this number be bounded away from zero, but any type of behaviour is possible as the matrix $S$ will grow in dimension. A reasonable global identifiability condition might be that the left side of Lemma 11 is bounded below by this number, and then the preceding bound on $\rho_{n}$ is valid. See Remark 1 for further discussion.

## 4 Application

Some options for implementing the Bayesian modularity are given in Section 4.1, after which the results of applying the Bayesian and likelihood modularities to the wellstudied karate club data of Zachary (1977) are discussed in Section 4.2.

### 4.1 Implementation

The Bayesian modularity, like the likelihood modularity, requires maximization over all possible labellings. This is computationally feasible even in large networks, as shown in two recent works on implementing Bayesian methods for the SBM. McDaid et al. (2013) followed the approach of Nowicki and Snijders (2001) and added a Poisson prior on $K$. After marginalizing over $\pi$ and $P$, they employ an allocation sampler to sample from the joint density of $K$ and $z$ given $A$, and use the posterior mode to estimate $K$. Their algorithm gives access to the full posterior distribution on the node labels and can scale to networks with approximately ten thousand nodes and ten million edges. Côme and Latouche (2014), claiming that the algorithm of McDaid et al. (2013) suffers from poor mixing properties, propose a greedy inference algorithm for the same problem. They demonstrate their algorithm on networks ranging in size from one hundred to ten thousand nodes, and compare the results to a range of other methods, including spectral clustering.

For the karate club data in Section 4.2, the network was small enough that a tabu search (Glover, 1989), run for a number of different initial configurations, yielded good results. This takes a similar amount of time as a tabu search in combination with the likelihood modularity, as in Bickel and Chen (2009). Although tabu search has been implemented on large networks consisting of approximately 1000 nodes for the degreecorrected version of the likelihood modularity (Zhao et al., 2012), we recommend the use of the methods designed for the stochastic block model proposed by McDaid et al. (2013) or Côme and Latouche (2014) for networks of medium and large sizes.


Figure 1: Communities detected by the Bayesian modularity when $K=2$ (left) and $K=4$ (right), with $\alpha=\beta_{1}=\beta_{2}=1 / 2$. The polygons contain the two groups the karate club was split into; the left one is Mr. Hi's club, the right one is the Officers' club. The shapes of the nodes represent the communities selected by the modularities. Figure made using the igraph package (Csardi and Nepusz, 2006).

### 4.2 Karate club

Zachary (1977) described a karate club which split into two clubs after a conflict over the price of the karate lessons. The new club was led by Mr. Hi, the karate teacher of the original club, while the remainder of the old club stayed under the former Officers' rule. The data consists of an adjacency matrix for those 34 individuals who interacted with other club members outside club meetings and classes. Each of these individuals' affiliations after the conflict is known.

We used $\alpha=1 / 2$ for the Dirichlet prior, and $\beta_{1}=\beta_{2}=1 / 2$ for the beta prior. The communities selected by the Bayesian modularity for $K=2$ and $K=4$ are given in Figure 1. In both instances, the tabu search led to nearly the same solution for both the Bayesian and likelihood modularities, only differing at one node for $K=4$, which is not surprising in light of Lemma 2. For $K=2$, the results of Bickel and Chen (2009) for this data set are recovered. For $K=4$, the partition in Figure 1 yields a higher value of the likelihood modularity than the partition into four classes found by Bickel and Chen (2009), and an even higher value is obtained by switching club member 20 to the second-largest class. This discrepancy is likely due to the heuristic nature of the tabu search algorithm, and for the same reason, it may be the case that improvement over the partitions found by the Bayesian modularity in Figure 1 are possible.

For $K=2$, the communities found by the algorithms do not correspond in the slightest to the two karate clubs, instead grouping the nodes with the highest degrees, corresponding to Mr. Hi, the president of the original club, and their closest supporters,
together. Incidentally, this partition is the same as the one returned by the Largest Gaps algorithm of Channarond et al. (2012), which solely uses the degrees of the nodes and discards all other information.

These bad results are no reason to shelve the Bayesian and likelihood modularities, as there is no reason to believe that the two karate clubs form communities in the sense of the stochastic block model. Mr. Hi and the club's president are clear outliers within their groups, and neither of the algorithms were designed to be robust to such a phenomenon. The communities selected by the modularities are communities in the sense that they form connections within and between the groups in a similar fashion. This sense does not correspond to the social notion of a community in this setting.

The results for four classes unify the social and stochastic senses of community. The prominent members of each of the new clubs are placed into two separate, small, communities. The other members are classified nearly perfectly, with two exceptions. However, one of those exceptional individuals is the only person described by Zachary (1977) as being a supporter of the club's president before the split, who joined Mr. Hi's club, making this person's affiliation up for debate. The second is described as only a weak supporter of Mr. Hi. The increased number of communities allows for some outliers within the social communities, and leads to a more detailed understanding of the dynamics within both of the groups. We essentially recover the two communities, each with a core that is more connective than the remainder of the nodes.

## 5 Weak consistency

The proof of Theorem 1 is built on our proof of weak consistency of the Bayesian modularity, which we present here. The following quantities will be used in the course of multiple proofs. The function $H_{P}$, with domain $K \times K$ probability matrices, is given by, for $\tau(u)=u \log u+(1-u) \log (1-u)$,

$$
\begin{equation*}
H_{P}(R)=\frac{1}{2} \sum_{a, b}(R \mathbf{1})_{a}(R \mathbf{1})_{b} \tau\left(\frac{\left(R P R^{T}\right)_{a b}}{(R \mathbf{1})_{a}(R \mathbf{1})_{b}}\right) \tag{7}
\end{equation*}
$$

For $\tau_{0}(u)=u \log (u)-u$, define

$$
G_{P}(R)=\frac{1}{2} \sum_{a, b}(R \mathbf{1})_{a}(R \mathbf{1})_{b} \tau_{0}\left(\frac{\left(R P R^{T}\right)_{a b}}{(R \mathbf{1})_{a}(R \mathbf{1})_{b}}\right)
$$

The sums defining these functions are over all pairs $(a, b)$ with $1 \leq a, b \leq K$, unlike the sums defining the modularities $Q_{B}$ and $Q_{M L}$, which are restricted to $a \leq b$.

We write $\operatorname{diag}(P)$ for the diagonal of $P$ if $P$ is a matrix, and $\operatorname{Diag}(f)$ for the diagonal matrix with diagonal $f$ if $f$ is a vector.

Theorem 2 (weak consistency). If $P=\rho_{n} S$ where either $\rho_{n}=1$ is fixed or $\rho_{n} \rightarrow 0$, and $(S, \pi)$ is fixed and identifiable, then the MAP classifier $\widehat{e}=\arg \max _{z} Q_{B}(e)$ is weakly consistent provided $n \rho_{n} \gg(\log n)^{2}$.

Proof. By Lemma 2 the Bayesian modularity $Q_{B}$ is equivalent to the likelihood modularity $Q_{M L}$ up to order $(\log n) / n$. With the notation $\widetilde{O}_{a b}(e)=O_{a b}(e)$ if $a \neq b$, and $\widetilde{O}_{a b}(e)=2 O_{a b}(e)$ if $a=b$, the likelihood modularity is in turn equivalent up to the same order to

$$
\begin{equation*}
\mathbb{L}(e)=\frac{1}{2 n^{2}} \sum_{a, b} n_{a}(e) n_{b}(e) \tau\left(\frac{\widetilde{O}_{a b}(e)}{n_{a}(e) n_{b}(e)}\right) \tag{8}
\end{equation*}
$$

Indeed the terms of $Q_{M L}(e)$ for $a<b$ are identical to the sums of the terms of $\mathbb{L}(e)$ for $a<b$ and $a>b$, while for $a=b$ the terms of $Q_{M L}(e)$ and $\mathbb{L}(e)$ differ only subtly: the first uses $n_{a a}(e)=\frac{1}{2} n_{a}(e)\left(n_{a}(e)-1\right)$, where the second uses $\frac{1}{2} n_{a}(e)^{2}$. Thus the difference is bounded in absolute value by the sum over $a$ of (where $e$ is suppressed from the notation)

$$
\left|\frac{n_{a}^{2}}{2 n^{2}} \tau\left(\frac{\widetilde{O}_{a a}}{n_{a}^{2}}\right)-\frac{n_{a}\left(n_{a}-1\right)}{2 n^{2}} \tau\left(\frac{\widetilde{O}_{a a}}{n_{a}\left(n_{a}-1\right)}\right)\right| \leq \frac{1}{2 n}\|\tau\|_{\infty}+\frac{n_{a}^{2}}{2 n^{2}} l\left(\frac{\widetilde{O}_{a a}}{n_{a}^{2}\left(n_{a}-1\right)}\right)
$$

where $l(x)=2 x(1 \vee \log (1 / x))$, in view of Lemma 6 . We now use that $n_{a} l\left(u / n_{a}\right) \lesssim$ $\log n_{a} \leq \log n$, for $0 \leq u \leq 1$.

Combining the preceding, we conclude that

$$
\eta_{n, 1}:=\max _{e}\left|\mathbb{L}(e)-Q_{B}(e)\right|=\mathcal{O}\left(\frac{\log n}{n}\right)
$$

Since $Q_{B}(\widehat{e}) \geq Q_{B}(Z)$, by the definition of $\widehat{e}$, it follows that $\mathbb{L}(\widehat{e})-\mathbb{L}(Z) \geq-2 \eta_{n, 1}$. The next step is to replace $\mathbb{L}$ in this equality by an asymptotic value.

For $x$ equal to a big multiple of $\left(\|P\|_{\infty}^{1 / 2} \vee n^{-1 / 2}\right) / n^{1 / 2}$, the right side of Lemma 4 tends to zero and hence $\max _{e}\|\widetilde{O}(e)-\mathbb{E}(\widetilde{O}(e) \mid Z)\|_{\infty} / n^{2}$ is of this order in probability. We also have, by Lemma 5 :

$$
\begin{aligned}
\max _{e}\left\|\frac{1}{n^{2}} \mathbb{E}(\widetilde{O}(e) \mid Z)-R(e, Z) P R(e, Z)^{T}\right\|_{\infty} & =\max _{e} \frac{1}{n}\|\operatorname{Diag}(R(e, Z) \operatorname{diag}(P))\|_{\infty} \\
& =O\left(\frac{\rho_{n}}{n}\right)
\end{aligned}
$$

as the row sums of the matrix $R(e, Z)$ are bounded above by one. By Lemma 6 , $|v \tau(x / v)-v \tau(y / v)| \leq l(|x-y|)$, uniformly in $v \in[0,1]$, where $l(x)=2 x(1 \vee \log (1 / x))$. It follows that

$$
\eta_{n, 2}:=\max _{e}|\mathbb{L}(e)-L(e)|=o_{P}\left(l\left(\frac{\|P\|_{\infty}^{1 / 2} \vee n^{-1 / 2}}{n^{1 / 2}}\right)\right),
$$

for

$$
L(e)=\frac{1}{2} \sum_{a, b} f_{a}(e) f_{b}(e) \tau\left(\frac{\left(R(e, Z) P R(e, Z)^{T}\right)_{a b}}{f_{a}(e) f_{b}(e)}\right)
$$

Combining this with the preceding paragraph, we conclude that $L(\widehat{e}) \geq L(Z)-2\left(\eta_{n, 1}+\right.$ $\left.\eta_{n, 2}\right)$. Since $L(e)=H_{P}(R(e, Z))$ for every $e$ and $H_{P}$ as defined in $(7)$, and $R(Z, Z)=$ $\operatorname{Diag}(f(Z))=\operatorname{Diag}\left(R(\widehat{e}, Z)^{T} \mathbf{1}\right)$, this can be translated into

$$
\begin{equation*}
H_{P}\left(\operatorname{Diag}\left(R(\widehat{e}, Z)^{T} \mathbf{1}\right)\right)-H_{P}(R(\widehat{e}, Z)) \leq 2\left(\eta_{n, 1}+\eta_{n, 2}\right) \tag{9}
\end{equation*}
$$

We complete the proof separately for the cases that $\rho_{n}$ is fixed or tends to zero.

For given $\delta>0$, let $\mathcal{R}_{\delta}$ be the set of all probability matrices $R$ with

$$
\min _{P_{\sigma}}\left\|P_{\sigma} R-\operatorname{Diag}\left(R^{T} \mathbf{1}\right)\right\|_{1} \geq \delta, \quad \text { and } \quad \min _{a: \pi_{a}>0}\left(R^{T} \mathbf{1}\right)_{a} \geq \delta
$$

Here the minimum is taken over the (finite) set of all permutation matrices $P_{\sigma}$ on $K$ labels. Furthermore, set

$$
\eta:=\inf _{R \in \mathcal{R}_{\delta}}\left[H_{P}\left(\operatorname{Diag}\left(R^{T} \mathbf{1}\right)\right)-H_{P}(R)\right]
$$

Because $\mathcal{R}_{\delta}$ is compact and the maps $R \mapsto H_{P}(R)$ and $R \mapsto \operatorname{Diag}\left(R^{T} \mathbf{1}\right)$ are continuous, the infimum in the display is assumed for some $R \in \mathcal{R}_{\delta}$. Because no $R \in \mathcal{R}_{\delta}$ can be transformed into a diagonal element by permuting rows and every $R \in \mathcal{R}_{\delta}$ has a nonzero element in every column $a$ with $\pi_{a}>0$, Lemma 7 shows that $\eta>0$. If $2\left(\eta_{n, 1}+\eta_{n, 2}\right)$ is smaller than $\eta$, then it follows from (9) that $R(\widehat{e}, Z)$ cannot be contained in $\mathcal{R}_{\delta}$. Since $R(\widehat{e}, Z)^{T} \mathbf{1}=f(Z) \xrightarrow{P} \pi$, by the law of large numbers, for sufficiently small $\delta>0$ this must be because $R(\widehat{e}, Z)$ fails the first requirement defining $\mathcal{R}_{\delta}$. That is, $\left\|P_{\sigma} R(\widehat{e}, Z)-\operatorname{Diag}(f(Z))\right\|_{1} \leq \delta$ for some permutation matrix $P_{\sigma}$. As this is true eventually for any $\delta>0$, it follows that $\min _{P_{\sigma}}\left\|P_{\sigma} R(\widehat{e}, Z)-\operatorname{Diag}(\pi)\right\|_{1} \xrightarrow{P} 0$.

Finally we consider the case where $\rho_{n} \rightarrow 0$. In view of Lemma 8 , the number $\eta=\eta_{n}$, which now depends on $n$, is now bounded below by $\rho_{n}$ times a positive number that depends on $(S, \pi)$. The preceding argument goes through provided $\eta_{n, 1}+\eta_{n, 2}$ is of smaller order than $\eta_{n}$. This leads to $l\left(\sqrt{\rho_{n} / n}\right)+\log (n) / n \ll \rho_{n}$, or $\left(\rho_{n} / n\right) \log ^{2}\left(n /\left(\rho_{n}\|S\|_{\infty}\right)\right) \ll$ $\rho_{n}^{2}$.

Remark 1. If $K_{n} \rightarrow \infty$, then the numbers $\eta_{n, 2}$ in the preceding proof need to be adapted to $\eta_{n, 2} \asymp K_{n}^{2} l\left(x_{n}+\rho_{n} / n\right)$, for $x_{n}$ a big multiple of $\left(\log K_{n} / n\right)^{1 / 2}\left(\rho_{n}^{1 / 2} \vee\right.$ $\left.\left(\log K_{n} / n\right)^{1 / 2}\right)$. Equation (9) remains valid. Rather than referring to the identifiability lemma, Lemma 8, we would now wish to lower bound the left side of (9) by a multiple of $n^{-1} \sum_{i=1}^{n} \mathbf{1}_{\hat{e}_{i} \neq Z_{i}}$. The proof of Lemma 11 combined with Lemma 1 shows that locally the left side of (9) is bounded below by a multiple of $\rho_{n} \sum_{a} \pi_{a} K_{0}\left(S_{a b^{\prime}} \| S_{a b}\right) n^{-1} \times$ $\sum_{i=1}^{n} \mathbf{1}_{\hat{e}_{i} \neq Z_{i}}$. If this is also globally true, then we obtain consistency as announced at the end of Section 3.5.

Lemmas 4-8 are more precise, or, in case of Lemma 7, corrected versions of lemmas from Bickel and Chen (2009); Zhao et al. (2012); Bickel et al. (2015), supporting the weak consistency theorem.
Lemma 4. Let $\widetilde{O}_{a b}(e)=O_{a b}(e)$ if $a \neq b$, and $\widetilde{O}_{a b}(e)=2 O_{a b}(e)$ if $a=b$. For any $x>0$,

$$
\mathbb{P}\left(\max _{e}\|\widetilde{O}(e)-\mathbb{E}(\widetilde{O}(e) \mid Z)\|_{\infty}>x n^{2}\right) \leq 2 K^{n+2} e^{-x^{2} n^{2} /\left(8\|P\|_{\infty}+4 x / 3\right)}
$$

Proof. This Lemma is adapted from Lemma 1.1 in Bickel and Chen (2009). There are $K^{n}$ possible values of $e$ and $\|\cdot\|_{\infty}$ is the maximum of the $K^{2}$ entries in the matrix.

We use the union bound to pull these maxima out of the probability, giving the factor $K^{n+2}$ on the right. Next it suffices to bound the tail probability of each variable
$\widetilde{O}_{a b}(e)-\mathbb{E}\left(\widetilde{O}_{a b}(e) \mid Z\right)=\sum_{i, j}\left(A_{i j}-\mathbb{E}\left(A_{i j} \mid Z\right)\right)\left(\mathbf{1}\left\{e_{i}=a, e_{j}=b\right\}+\mathbf{1}\left\{e_{i}=b, e_{j}=a\right\}\right)$.
The $n_{a b}(e)$ variables in this sum are conditionally independent given $Z$, take values in $[-2,2]$, and have conditional mean zero given $Z$ and conditional variance bounded by $4 \operatorname{var}\left(A_{i j} \mid Z\right) \leq 4 P_{Z_{i} Z_{j}}\left(1-P_{Z_{i} Z_{j}}\right) \leq 4\|P\|_{\infty}$. Thus we can apply Bernstein's inequality to find that

$$
\mathbb{P}\left(\left|\widetilde{O}_{a b}(e)-\mathbb{E}\left(\widetilde{O}_{a b}(e) \mid Z\right)\right|>x n^{2}\right) \leq 2 e^{-x^{2} n^{4} /\left(8 n_{a b}(e)\|P\|_{\infty}+4 x n^{2} / 3\right)}
$$

Finally we use the crude bound $n_{a b}(e) \leq n^{2}$ and cancel one factor $n^{2}$.
Lemma 5. Define $\widetilde{O}_{a b}(e)=O_{a b}(e)$ if $a \neq b$, and $\widetilde{O}_{a b}(e)=2 O_{a b}(e)$ if $a=b$. Then, for $R(e, Z)$ as defined in (2),

$$
\mathbb{E}\left(\widetilde{O}_{a b} \mid Z\right)=n^{2} R(e, Z) P R(e, Z)^{T}-n \operatorname{Diag}(R(e, Z) \operatorname{diag}(P))
$$

Proof. A similar expression, not taking into account the absence of self-loops, appears in Bickel and Chen (2009). The relevant computation for our situation is as follows:

$$
\begin{aligned}
& \mathbb{E}\left(\widetilde{O}_{a b}(e) \mid Z=c\right)=\sum_{i \neq j} P_{c_{i} c_{j}} \mathbf{1}\left\{e_{i}=a, e_{j}=b\right\} \\
& =\sum_{a^{\prime}, b^{\prime}} P_{a^{\prime} b^{\prime}} \sum_{i \neq j} \mathbf{1}\left\{c_{i}=a^{\prime}, c_{j}=b^{\prime}\right\} \mathbf{1}\left\{e_{i}=a, e_{j}=b\right\} \\
& =\sum_{a^{\prime}, b^{\prime}} P_{a^{\prime} b^{\prime}} \sum_{i, j} \mathbf{1}\left\{c_{i}=a^{\prime}, c_{j}=b^{\prime}\right\} \mathbf{1}\left\{e_{i}=a, e_{j}=b\right\} \\
& \quad-\delta_{a b} \sum_{i} \sum_{a^{\prime}} P_{a^{\prime} a^{\prime}} \mathbf{1}\left\{c_{i}=a^{\prime}\right\} \mathbf{1}\left\{e_{i}=a\right\} \\
& =n^{2} \sum_{a^{\prime}, b^{\prime}} P_{a^{\prime} b^{\prime}} R_{a a^{\prime}}(e, c) R_{b b^{\prime}}(e, c)-\delta_{a b} n \sum_{a^{\prime}} P_{a^{\prime} a^{\prime}} R_{a a^{\prime}}(e, c) .
\end{aligned}
$$

Lemma 6. The function $\tau:[0,1] \rightarrow \mathbb{R}$ satisfies $|\tau(x)-\tau(y)| \leq l(|x-y|)$, for $l(x)=$ $2 x(1 \vee \log (1 / x))$.

Proof. Write the difference between $x \log x$ and $y \log y$ as $\left|\int_{x}^{y}(1+\log s) d s\right|$. The function $s \mapsto 1+\log s$ is strictly increasing on $[0,1]$ from $-\infty$ to 1 and changes sign at $s=e^{-1}$. Therefore the absolute integral is bounded above by the maximum of

$$
-\int_{0}^{|x-y| \wedge e^{-1}}(1+\log s) d s=-\left(|x-y| \wedge e^{-1}\right) \log \left(|x-y| \wedge e^{-1}\right)
$$

and

$$
\int_{1-|x-y| \vee e^{-1}}^{1}(1+\log s) d s \leq|x-y|
$$

Lemma 7. For any probability matrix $R$,

$$
\begin{equation*}
H_{P}(R) \leq H_{P}\left(\operatorname{Diag}\left(R^{T} \mathbf{1}\right)\right) \tag{10}
\end{equation*}
$$

Furthermore, if $(P, \pi)$ is identifiable and the columns of $R$ corresponding to positive coordinates of $\pi$ are not identically zero, then the inequality is strict unless $P_{\sigma} R$ is a diagonal matrix for some permutation matrix $P_{\sigma}$.

Proof. This Lemma is related to the proof that the likelihood modularity is consistent given in Bickel and Chen (2009). This proof however rests on their incorrect Lemma 3.1, and thus we provide full details on how the argument can be adapted to avoid the use of their Lemma 3.1 altogether.

For $R$ a diagonal matrix the numbers $\left(R P R^{T}\right)_{a b} /(R \mathbf{1})_{a}(R \mathbf{1})_{b}$ reduce to $P_{a b}$. Consequently, by the definition of $H_{P}$,

$$
\begin{equation*}
H_{P}(\operatorname{Diag}(f))=\sum_{a, b} f_{a} f_{b} \tau\left(P_{a b}\right) \tag{11}
\end{equation*}
$$

For a general matrix $R$, by inserting the definition of $\tau$,

$$
\begin{aligned}
H_{P}(R)=\sum_{a, b}(R P & \left.R^{T}\right)_{a b} \log \frac{\left(R P R^{T}\right)_{a b}}{(R \mathbf{1})_{a}(R \mathbf{1})_{b}} \\
& +\sum_{a, b}\left((R \mathbf{1})_{a}(R \mathbf{1})_{b}-\left(R P R^{T}\right)_{a b}\right) \log \left(1-\frac{\left(R P R^{T}\right)_{a b}}{(R \mathbf{1})_{a}(R \mathbf{1})_{b}}\right)
\end{aligned}
$$

Because $(R \mathbf{1})_{a}(R \mathbf{1})_{b}-\left(R P R^{T}\right)_{a b}=\left(R(1-P) R^{T}\right)_{a b}$, with 1 the $(K \times K)$-matrix with all coordinates equal to 1 , we can rewrite this as

$$
\sum_{a, b} \sum_{a^{\prime}, b^{\prime}} R_{a a^{\prime}} R_{b b^{\prime}}\left[P_{a^{\prime} b^{\prime}} \log \frac{\left(R P R^{T}\right)_{a b}}{(R \mathbf{1})_{a}(R \mathbf{1})_{b}}+\left(1-P_{a^{\prime} b^{\prime}}\right) \log \left(1-\frac{\left(R P R^{T}\right)_{a b}}{(R \mathbf{1})_{a}(R \mathbf{1})_{b}}\right)\right]
$$

By the information inequality for two-point measures, the expressions in square brackets become bigger when $\left(R P R^{T}\right)_{a b} /(R \mathbf{1})_{a}(R \mathbf{1})_{b}$ is replaced by $P_{a^{\prime} b^{\prime}}$, with a strict increase unless these two numbers are equal. After making this substitution the term in square brackets becomes $\tau\left(P_{a^{\prime} b^{\prime}}\right)$, and we can exchange the order of the two (double) sums and perform the sum on $(a, b)$ to write the resulting expression as

$$
\sum_{a^{\prime}, b^{\prime}}\left(R^{T} \mathbf{1}\right)_{a^{\prime}}\left(R^{T} \mathbf{1}\right)_{b^{\prime}} \tau\left(P_{a^{\prime} b^{\prime}}\right)=H_{P}\left(\operatorname{Diag}\left(R^{T} \mathbf{1}\right)\right)
$$

This proves the first assertion (10) of the lemma.
If $R$ attains equality, then also for every permutation matrix $P_{\sigma}$, by the equality $H_{P}\left(P_{\sigma} R\right)=H_{P}(R)$ and the fact that $\left(P_{\sigma} R\right)^{T} \mathbf{1}=R^{T} \mathbf{1}$, we have

$$
\begin{equation*}
H_{P}\left(P_{\sigma} R\right)=H_{P}\left(\operatorname{Diag}\left(\left(P_{\sigma} R\right)^{T} \mathbf{1}\right)\right) \tag{12}
\end{equation*}
$$

We shall show that if $R$ satisfies this equality and $P_{\sigma} R$ has a positive diagonal, then $P_{\sigma} R$ is in fact diagonal. Furthermore, we shall show that there exists $P_{\sigma}$ such that $P_{\sigma} R$ has a positive diagonal.

Fix some $\left(P_{\sigma}\right)_{m}$ that maximizes the number of positive diagonal elements of $P_{\sigma} R$ over all permutation matrices $P_{\sigma}$, and denote $\bar{R}=\left(P_{\sigma}\right)_{m} R$. Because the information inequality is strict, the preceding argument shows that (12) can be true for $P_{\sigma}=\left(P_{\sigma}\right)_{m}$ (giving $P_{\sigma} R=\bar{R}$ ) only if

$$
\begin{equation*}
P_{a^{\prime} b^{\prime}}=\frac{\left(\bar{R} P \bar{R}^{T}\right)_{a b}}{(\bar{R} \mathbf{1})_{a}(\bar{R} \mathbf{1})_{b}}, \quad \text { whenever } \bar{R}_{a a^{\prime}} \bar{R}_{b b^{\prime}}>0 \tag{13}
\end{equation*}
$$

Denote the matrix on the right of the equality by $Q$.
If $\bar{R}$ has a completely positive diagonal, then we can choose $a=a^{\prime}$ and $b=b^{\prime}$ and find from (13), that $P_{a b}=Q_{a b}$, for every $a, b$. If also $\bar{R}_{a a^{\prime}}>0$, then we can also choose $b=b^{\prime}$ and find that $P_{a^{\prime} b}=Q_{a b}$, for every $b$. Thus the $a$ th and $a^{\prime}$ th rows of $P$ are identical. Since all rows of $P$ are different by assumption, it follows that no $a \neq a^{\prime}$ with $\bar{R}_{a a^{\prime}}>0$ exists.

If $\bar{R}$ does not have a fully positive diagonal, then the submatrix of $\bar{R}$ obtained by deleting the rows and columns corresponding to positive diagonal elements must be the zero matrix, since otherwise we might permute the remaining rows and create an additional nonzero diagonal element, contradicting that $\left(P_{\sigma}\right)_{m}$ already maximized this number. If $I$ and $I^{c}$ are the sets of indices of zero and nonzero diagonal elements, then the preceding observation is that $\bar{R}_{i j}$ is zero for every $i, j \in I$. If $\pi>0$, then we need to consider only $R$ with nonzero columns. For $i \in I$ a nonzero element in the $i$ th column of $\bar{R}$ must be located in the rows with label in $I^{c}$ : for every $i \in I$ there exists $k_{i} \in I^{c}$ with $\bar{R}_{k_{i} i}>0$. Then, for $i, j \in I$,
(1) for $a=k_{i}, b=k_{j}, a^{\prime}=i, b^{\prime}=j$, (13) implies $Q_{k_{i} k_{j}}=P_{i j}$.
(2) for $a=k_{i}, b \in I^{c}, a^{\prime}=i, b^{\prime}=b$, (13) implies $Q_{k_{i} b}=P_{i b}$.
(3) for $a=k_{i}, b \in I^{c}, a^{\prime}=k_{i}, b^{\prime}=b$, (13) implies $Q_{k_{i} b}=P_{k_{i} b}$.

We combine these three assertions to conclude that, for $a, i \in I$ and $b \in I^{c}$,

$$
\begin{aligned}
& P_{a i}=P_{i a} \stackrel{(1)}{=} Q_{k_{i} k_{a}} \stackrel{(2)}{=} P_{i k_{a}}=P_{k_{a} i}, \\
& P_{a b} \stackrel{(2)}{=} Q_{k_{a} b} \stackrel{(3)}{=} P_{k_{a} b} .
\end{aligned}
$$

Together these imply that the $a$ th and the $k_{a}$ th row of $P$ are equal. Since by assumption they are not (if $\pi>0$ ), this case can actually not exist (i.e. $k=0$ ).

Finally if $\pi_{a}=0$ for some $a$, then we follow the same argument, but we match only every column $i \in I$ with $\pi_{i}>0$ to a row $k_{i} \in I^{c}$. By the assumption on $R$ such $k_{i}$ exist, and the construction results in two rows of $P$ that are identical in the coordinates with $\pi_{a}>0$.

Lemma 8. For any fixed $(K \times K)$-matrix $P$ with elements in $[0,1]$, uniformly in probability matrices $R$, as $\rho_{n} \rightarrow 0$,

$$
\begin{equation*}
\frac{1}{\rho_{n}}\left(H_{\rho_{n} P}\left(\operatorname{Diag}\left(R^{T} \mathbf{1}\right)\right)-H_{\rho_{n} P}(R)\right) \rightarrow G_{P}\left(\operatorname{Diag}\left(R^{T} \boldsymbol{1}\right)\right)-G_{P}(R) \tag{14}
\end{equation*}
$$

Furthermore, if $(P, \pi)$ is identifiable and the columns of $R$ corresponding to positive coordinates of $\pi$ are not identically zero, then the right side is strictly positive unless $S R$ is a diagonal matrix for some permutation matrix $S$.

Proof. From the fact that $|(1-u) \log (1-u)+u| \leq u^{2}$, for $0 \leq u \leq 1$, it can be verified that, $\left|\rho_{n}^{-1} \tau\left(\rho_{n} u\right)-\left(u \log \rho_{n}+\tau_{0}(u)\right)\right| \leq \rho_{n} \rightarrow 0$, uniformly in $0 \leq u \leq 1$. It follows that, uniformly in $R$,

$$
\frac{1}{\rho_{n}} H_{\rho_{n} P}(R)=\log \rho_{n} \sum_{a, b}\left(R P R^{T}\right)_{a b}+\sum_{a, b}(R \mathbf{1})_{a}(R \mathbf{1})_{b} \tau_{0}\left(\frac{\left(R P R^{T}\right)_{a b}}{(R \mathbf{1})_{a}(R \mathbf{1})_{b}}\right)+\mathcal{O}\left(\rho_{n}\right)
$$

The first term on the right is equal to $\log \rho_{n}\left(R^{T} \mathbf{1}\right)^{T} P\left(R^{T} \mathbf{1}\right)$, and hence is the same for $R$ and $\operatorname{Diag}\left(R^{T} \mathbf{1}\right)$. Thus this term cancels on taking the difference to form the left side of (14), and hence (14) follows.

The right side of (14) is nonnegative, because the left side is, by Lemma 7. This fact can also be proved directly along the lines of the proof of Lemma 7, as follows. Write

$$
G_{P}(R)=\sum_{a, b} \sum_{a^{\prime}, b^{\prime}} R_{a a^{\prime}} R_{b b^{\prime}}\left[P_{a^{\prime} b^{\prime}} \log \frac{\left(R P R^{T}\right)_{a b}}{(R \mathbf{1})_{a}(R \mathbf{1})_{b}}-\frac{\left(R P R^{T}\right)_{a b}}{(R \mathbf{1})_{a}(R \mathbf{1})_{b}}\right]
$$

By the information inequality for two Poisson distributions the term in square brackets becomes bigger if $\left(R P R^{T}\right)_{a b} /(R \mathbf{1})_{a}(R \mathbf{1})_{b}$ is replaced by $P_{a^{\prime} b^{\prime}}$. It then becomes $\tau_{0}\left(P_{a^{\prime} b^{\prime}}\right)$ and the double sum on $(a, b)$ can be executed to see that the resulting bound is $G_{P}\left(\operatorname{Diag}\left(R^{T} \mathbf{1}\right)\right)$. Furthermore, the inequality is strictly unless (13) holds, with $\bar{R}=R$. Since also $G_{P}\left(P_{\sigma} R\right)=G_{P}(R)$, for every permutation matrix $P_{\sigma}$, the final assertion of the lemma is proved by copying the proof of Lemma 7 .

## Proof of Lemma 2

Proof. The second assertion of the lemma follows from the first and the fact that $\max _{e} Q_{P}(e) \lesssim(\log n) / n$. It suffices to prove the first assertion.

Recall that the Bayesian modularity is given by $n^{-2}$ times

$$
\begin{equation*}
n^{2} Q_{B}(e)=\sum_{a \leq b} \log B\left(O_{a b}(e)+\frac{1}{2}, n_{a b}(e)-O_{a b}(e)+\frac{1}{2}\right)+\sum_{a} \log \Gamma\left(n_{a}(e)+\alpha\right) \tag{15}
\end{equation*}
$$

We shall show that the first sum on the right is equivalent to $Q_{M L}(e)$, and the second sum is equivalent to $Q_{P}(e)$. We show this by comparing the sums defining the various modularities term by term. For clarity we shall suppress the argument $e$. We will repeatedly use the following bound from (Robbins, 1955): for $n \in \mathbb{N}_{\geq 1}$,

$$
\begin{equation*}
\Gamma(n+1)=\sqrt{2 \pi} n^{n+1 / 2} e^{-n} e^{a_{n}} \tag{16}
\end{equation*}
$$

with $(12 n+1)^{-1} \leq a_{n} \leq(12 n)^{-1}$, as well as the fact that $\Gamma(s)$ is monotone increasing for $s \geq 3 / 2$. In addition, we will bound remainder terms by using the inequality $x \log ((x+$ $c) / x) \leq c$ for $c \geq 0$ and the fact that $x \log ((x-1) / x)$ is bounded for $x>1$.

## First sum of (15)

Upper bound, case 1: $O_{a b} \neq 0$ and $n_{a b} \neq O_{a b}$. We apply (16):

$$
\begin{aligned}
\log B & \left(O_{a b}+\beta_{1}, n_{a b}-O_{a b}+\beta_{2}\right) \leq \log \frac{\Gamma\left(O_{a b}+\left\lfloor\beta_{1}\right\rfloor+1\right) \Gamma\left(n_{a b}-O_{a b}+\left\lfloor\beta_{2}\right\rfloor+1\right)}{\Gamma\left(n_{a b}+\left\lfloor\beta_{1}+\beta_{2}\right\rfloor\right)} \\
= & O_{a b} \log \left(\frac{O_{a b}+\left\lfloor\beta_{1}\right\rfloor}{n_{a b}+\left\lfloor\beta_{1}+\beta_{2}\right\rfloor-1}\right)+\left(n_{a b}-O_{a b}\right) \log \left(\frac{n_{a b}-O_{a b}+\left\lfloor\beta_{2}\right\rfloor}{n_{a b}+\left\lfloor\beta_{1}+\beta_{2}\right\rfloor-1}\right) \\
& +\left(\left\lfloor\beta_{1}\right\rfloor+1 / 2\right) \log \left(O_{a b}+\left\lfloor\beta_{1}\right\rfloor\right)+\left(\left\lfloor\beta_{2}\right\rfloor+1 / 2\right) \log \left(n_{a b}-O_{a b}+\left\lfloor\beta_{2}\right\rfloor\right) \\
& -\left(\left\lfloor\beta_{1}+\beta_{2}\right\rfloor-1 / 2\right) \log \left(n_{a b}+\left\lfloor\beta_{1}+\beta_{2}\right\rfloor-1\right)+\log \sqrt{2 \pi}-\left\lfloor\beta_{1}\right\rfloor-\left\lfloor\beta_{2}\right\rfloor \\
& +\left\lfloor\beta_{1}+\beta_{2}\right\rfloor-1+\alpha_{a b}+\beta_{a b}-\gamma_{a b},
\end{aligned}
$$

where $\alpha_{a b}, \beta_{a b}$ and $\gamma_{a b}$ are bounded by constants. By the inequality $x \log ((x+c) / x) \leq c$ for $c \geq 0$, and the fact that $x \log ((x-1) / x)$ is bounded for $x>1$, we find the upper bound:

$$
\log B\left(O_{a b}+\beta_{1}, n_{a b}-O_{a b}+\beta_{2}\right) \leq n_{a b} \tau\left(\frac{O_{a b}}{n_{a b}}\right)+\mathcal{O}\left(\log n_{a b}\right)
$$

Upper bound, case 2: $n_{a b}=1$ and $O_{a b}=0$ or $n_{a b}=O_{a b}$, or $n_{a b}=0$. In both cases, the corresponding term of the likelihood modularity vanishes, whereas the contribution of the Bayesian modularity is either $\log B\left(1+\beta_{1}, \beta_{2}\right), \log \left(\beta_{1}, 1+\beta_{2}\right)$, or $\log B\left(\beta_{1}, \beta_{2}\right)$.

Upper bound, case 3: $n_{a b} \geq 2$ and $O_{a b}=0$ or $n_{a b}=O_{a b}$. Again, the corresponding term of the likelihood modularity vanishes. We show the computations for the case $n_{a b}=O_{a b} ;$ for the case $O_{a b}=0$, switch $\beta_{1}$ and $\beta_{2}$. By (16):

$$
\begin{aligned}
& \log B\left(O_{a b}+\beta_{1}, n_{a b}-O_{a b}+\beta_{2}\right)=\log B\left(n_{a b}+\beta_{1}, \beta_{2}\right) \leq \log \frac{\Gamma\left(n_{a b}+\left\lfloor\beta_{1}\right\rfloor+1\right) \Gamma\left(\beta_{2}\right)}{\Gamma\left(n_{a b}+\left\lfloor\beta_{1}+\beta_{2}\right\rfloor\right)} \\
& = \\
& \quad\left(n_{a b}+\left\lfloor\beta_{1}\right\rfloor\right) \log \left(\frac{n_{a b}+\left\lfloor\beta_{1}\right\rfloor}{n_{a b}+\left\lfloor\beta_{1}+\beta_{2}\right\rfloor}\right)+(1 / 2) \log \left(n_{a b}+\left\lfloor\beta_{1}\right\rfloor\right) \\
& \\
& \quad-\left(\left\lfloor\beta_{1}+\beta_{2}\right\rfloor+1 / 2\right) \log \left(n_{a b}+\left\lfloor\beta_{1}+\beta_{2}\right\rfloor\right)+\log \Gamma\left(\beta_{2}\right)+\left\lfloor\beta_{1}+\beta_{2}\right\rfloor-1+\delta_{a b}-\epsilon_{a b}
\end{aligned}
$$

where $\delta_{a b}$ and $\epsilon_{a b}$ are bounded by constants. Arguing as before, the first term is bounded, while the remainder is of order $\log \left(n_{a b}\right)$. A lower bound is found analogously.

Lower bound. The computations for the lower bound are completely analogous, except that we require $O_{a b}+\beta_{1} \geq 2$ and $n_{a b}-O_{a b}+\beta_{2} \geq 2$. We study four cases. The cases (1) $O_{a b} \geq 2$ and $n_{a b}-O_{a b} \geq 2$, (2) $n_{a b}=0$ and (3) $n_{a b}>0$ and $n_{a b}=O_{a b}$ or $O_{a b}=0$ are similar to cases 1,2 and 3 respectively of the upper bound. The fourth case is $n_{a b}-O_{a b}=1$ and $O_{a b} \geq 2$, or $O_{a b}=1$ and $n_{a b}-O_{a b} \geq 1$. In both instances, the likelihood modularity is equality to a bounded term minus $\log n_{a b}$. By similar calculations as before, the Bayesian modularity is of the order $\log n_{a b}$ as well.

Conclusion. We find:

$$
\sum_{a \leq b} \log B\left(O_{a b}+\beta_{1}, n_{a b}-O_{a b}+\beta_{2}\right)=\sum_{a \leq b} n_{a b} \tau\left(\frac{O_{a b}}{n_{a b}}\right)+\mathcal{O}(\log n)
$$

## Second sum of (15)

We consider three cases. If $n_{a}+\lfloor\alpha\rfloor=0$, then $\alpha>0$, implies $n_{a}=0$, in which case $\log \Gamma\left(n_{a}+\alpha\right)=\log \Gamma(\alpha)$, which is bounded. In case $n_{a}+\lfloor\alpha\rfloor=1$, the term $\log \Gamma\left(n_{a}+\alpha\right)$ is equal to either $\log \Gamma(1+\alpha)$ or $\log \Gamma(\alpha)$ and thus bounded as well. For the case $n_{a}+\lfloor\alpha\rfloor \geq 2$, we study the upper bound $\Gamma\left(n_{a}+\alpha\right) \leq \Gamma\left(n_{a}+\lfloor\alpha\rfloor+1\right)$ and the lower bound $\Gamma\left(n_{a}+\alpha\right) \geq \Gamma\left(n_{a}+\lfloor\alpha\rfloor\right)$. By applying (16) in both cases, we conclude:

$$
\sum_{a} \log \Gamma\left(n_{a}+\alpha\right)=\sum_{a: n_{a}+\lfloor\alpha\rfloor \geq 2} n_{a} \log n_{a}-n+\mathcal{O}(\log n)
$$

## 6 Strong consistency

We build upon the foundations from the previous Section to prove Theorem 1. We need slightly adapted versions of the function $H_{P}$, given by, with $\delta_{a b}$ equal to 1 or 0 if $a=b$ or not,

$$
\begin{equation*}
H_{P, n}(R)=\frac{1}{2} \sum_{a, b}(R \mathbf{1})_{a}\left((R \mathbf{1})_{b}-\delta_{a b} / n\right) \tau\left(\frac{\left(R P R^{T}\right)_{a b}-\delta_{a b} \sum_{k} P_{k k} R_{k a} / n}{(R \mathbf{1})_{a}\left((R \mathbf{1})_{b}-\delta_{a b} / n\right)}\right) \tag{17}
\end{equation*}
$$

For given functions $t_{a b}:[0,1] \rightarrow \mathbb{R}$, let $X(e)$ be the $K \times K$ matrix with entries

$$
\begin{equation*}
X_{a b}(e)=t_{a b}\left(\frac{\widetilde{O}_{a b}(e)}{n^{2}}\right)-t_{a b}\left(\frac{\mathbb{E}\left(\widetilde{O}_{a b}(e) \mid Z\right)}{n^{2}}\right) \tag{18}
\end{equation*}
$$

## Proof of Theorem 1 [strong consistency]

Proof. We first prove the statement in case $\rho_{n}$ is fixed. By Theorem $2, \widehat{e}$ is weakly consistent, and hence with probability tending to one it belongs to the set of classifications $e$ such that the fractions $f(e)$ are close to $\pi$, and the matrices $R(e, Z)$ are close to $\operatorname{Diag}(\pi)$ after the appropriate permutation of the labels (that is, of rows of $R(e, Z)$ ). Therefore, it is no loss of generality to assume that $\widehat{e}$ is restricted to this set. By Lemmas 4 and 5 , the matrices $\widetilde{O}(e) / n^{2}$ are then close to $R(e, Z) P R(e, Z)^{T} \rightarrow \operatorname{Diag}(\pi) P \operatorname{Diag}(\pi)$, and hence are bounded away from zero and one if $P$ has this property.

If $\widehat{e}$ and $Z$ differ at $m$ nodes, then $\widehat{e}$ belongs to the set of $e$ with $\|R(Z, Z)-R(e, Z)\|_{1}=$ $m(2 / n)$, by Lemma 1. In that case $Q_{B}(e) \geq Q_{B}(Z)$, for some $e$ in this set, and hence by Lemma $2 Q_{M L}(e)-Q_{M L}(Z)+Q_{P}(e)-Q_{P}(Z) \geq-\eta_{n}$, for some $\eta_{n}$ of order $(\log n) / n^{2}$. It follows that:

$$
\begin{align*}
{\left[Q_{M L}(e)\right.} & \left.-H_{P, n}(R(e, Z))\right]-\left[Q_{M L}(Z)-H_{P, n}(R(Z, Z))\right] \\
& \geq H_{P, n}(R(Z, Z))-H_{P, n}(R(e, Z))-\left|Q_{P}(e)-Q_{P}(Z)\right|-\eta_{n} . \tag{19}
\end{align*}
$$

The first term on the right is bounded below by a multiple of $m / n$, by Lemmas 9 and 1. Because $(x+\alpha) \log x-(y+\alpha) \log y=\int_{x}^{y}(\log s+(s+\alpha) / s) d s$ is bounded in absolute value by a multiple of $|x-y| \log (x \vee y)$, if $\alpha \geq 0$ and $x, y>0$, the second term $-\left|Q_{P}(e)-Q_{P}(Z)\right|$ is bounded below by a multiple of $m(\log n) / n^{2}$, which is of smaller order than $m / n$. We conclude that the left side of (19) is bounded below by $C_{1} m / n$. The left side is $\sum_{a, b}\left(X_{a b}(e)-X_{a b}(Z)\right)$, for $X$ defined in (18) and $t$ the function with coordinates $t_{a b}(o)=f_{a}(e)\left(f_{b}(e)-\delta_{a b} / n\right) \tau\left(o / f_{a}(e)\left(f_{b}(e)-\delta_{a b} / n\right)\right)$. Because we restrict $e$ to classifications such that $O_{a b}(e) / n_{a b}(e)$ and $f_{a}(e) f_{b}(e)$ are bounded away from zero and one, only the values of the function $\tau$ on an open interval strictly within $(0,1)$ matter. On any such interval $\tau$ has uniformly bounded derivatives, and hence the bound of Lemma 12 is valid. Thus we find that

$$
\begin{aligned}
\operatorname{Pr}\left(\#\left(i: \widehat{e}_{i} \neq Z_{i}\right)=m\right) & \leq \operatorname{Pr}\left(\sup _{e: \#\left(i: e_{i} \neq Z_{i}\right) \leq m}\|X(e)-X(Z)\|_{\infty} \geq \frac{C_{1} m}{n}\right) \\
& \leq C_{2} K^{m}\binom{n}{m} e^{-c m^{2} /\left(m\|P\|_{\infty} / n+m / n\right)} \\
& \leq C_{2} e^{m \log (\text { Kne/m)-c, } m n}
\end{aligned}
$$

The sum of the right side over $m=1, \ldots, n$ tends to zero.
In case $\rho_{n} \rightarrow 0$, we follow the some proof, but in (19) use that $H_{P, n}(R(Z, Z))-$ $H_{P, n}(R(e, Z)) \geq \rho_{n} C\|R(Z, Z)-R(e, Z)\|_{1} \geq \rho_{n} C 2 m / n$, by Lemma 11. Since $\rho_{n} \gg$ $(\log n) / n$ by assumption, we have that the contribution $m(\log n) / n^{2}$ of $Q_{P}(e)-Q_{P}(Z)$ is still negligible and hence $\rho_{n} C 2 m / n$ is a lower bound for the left side of (19). As a bound on the left side of the preceding display, we then obtain

$$
\sum_{m=1}^{n} K^{m}\binom{n}{m} e^{-c_{2} \rho_{n}^{2} m^{2} /\left(m \rho_{n} / n+\rho_{n} m / n\right)} \leq \sum_{m=1}^{n} e^{m \log (K n e / m)-c_{3} \rho_{n} m n}
$$

This sum tends to zero provided that $n \rho_{n} \gg \log n$.

Lemmas 9-11 are explicit verifications of versions of condition IIIc of Bickel and Chen (2009).
Lemma 9. If $P$ is fixed and symmetric, $(P, \pi)$ is identifiable and $0<P<1$, then, for sufficiently small $\delta>0$,

$$
\begin{equation*}
\liminf _{n \rightarrow \infty} \inf _{0<\|R-\operatorname{Diag}(\pi)\|<\delta} \frac{H_{P, n}\left(\operatorname{Diag}\left(R^{T} \boldsymbol{1}\right)\right)-H_{P, n}(R)}{\left\|\operatorname{Diag}\left(R^{T} \boldsymbol{1}\right)-R\right\|}>0 \tag{20}
\end{equation*}
$$

Proof. We can reparametrize the $K \times K$ matrices $R$ by the pairs ( $R^{T} \mathbf{1}, R-\operatorname{Diag}\left(R^{T} \mathbf{1}\right)$ ), consisting of the $K$ vector $f=R^{T} \mathbf{1}$ and the $K \times K$ matrix $R-\operatorname{Diag}\left(R^{T} \mathbf{1}\right)$. The latter matrix is characterized by having nonnegative off-diagonal elements and zero column
sums, and can be represented in the basis consisting of all $K \times K$ matrices $\Delta_{b b^{\prime}}$, for $b \neq b^{\prime}$, defined by: $\left(\Delta_{b b^{\prime}}\right)_{b^{\prime} b^{\prime}}=-1,\left(\Delta_{b b^{\prime}}\right)_{b b^{\prime}}=1$ and $\left(\Delta_{b b^{\prime}}\right)_{a a^{\prime}}=0$, for all other entries $\left(a, a^{\prime}\right)$, i.e. the $b^{\prime}$ th column of $\Delta_{b b^{\prime}}$ has a 1 in the $b$ th coordinate and a -1 on the $b^{\prime}$ th coordinate and all its other columns are zero. Given any matrix $R \geq 0$ the matrix $R-\operatorname{Diag}\left(R^{T} \mathbf{1}\right)$ can be decomposed as

$$
R-\operatorname{Diag}\left(R^{T} \mathbf{1}\right)=\sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}} \Delta_{b b^{\prime}}
$$

for $\lambda_{b b^{\prime}}=R_{b b^{\prime}} \geq 0$. Since every $\Delta_{b b^{\prime}}$ has exactly one nonzero off-diagonal element, which is equal to 1 , and in a different location for each $b \neq b$, the sum of the off-diagonal elements of the matrix on the right side is $\sum_{b, b^{\prime}} \lambda_{b b^{\prime}}$. Because the sum of all its elements is zero, it follows that its sum of absolute elements is given by $\left\|R-\operatorname{Diag}\left(R^{T} \mathbf{1}\right)\right\|_{1}=$ $2 \sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}}$.

Thus we obtain a further reparametrization $R \leftrightarrow(f, \lambda)$, in which $R=\operatorname{Diag}(f)+$ $\sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}} \Delta_{b b^{\prime}}$. Here the vector $f$ is a probability vector, and all $\lambda_{b b^{\prime}}$ are nonnegative (as $R_{b b^{\prime}} \geq 0$ ). The nonnegativity of the diagonal elements of $R$ gives the further restrictions that $\sum_{b \neq a} \lambda_{b a} \leq f_{a}$, for every $a$; in particular $\lambda_{b a}$ is zero for every $b$ and $a$ such that $f_{a}=0$. Other restrictions on the $\lambda_{b b^{\prime}}$ follow from the fact that $R \leq 1$, but since we shall be interested in $\lambda_{b b^{\prime}}$ close to zero, these restrictions will not be active.

For given $P, f$ and $n$, define the function

$$
G(\lambda)=H_{P, n}\left(\operatorname{Diag}(f)+\sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}} \Delta_{b b^{\prime}}\right)
$$

Then we would like to show that there exists $C$ such that

$$
\frac{H_{P, n}\left(\operatorname{Diag}\left(R^{T} \mathbf{1}\right)\right)-H_{P, n}(R)}{\left\|R-\operatorname{Diag}\left(R^{T} \mathbf{1}\right)\right\|_{1}}=\frac{G(0)-G(\lambda)}{2 \sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}}} \geq C>0
$$

for every $f$ in a neighbourhood of $\pi, \lambda$ in a neighbourhood of 0 intersected with $\{\lambda$ : $\lambda \geq 0\}$ and $\cap_{a}\left\{\lambda: \sum_{b \alpha} \lambda_{b a} \leq f_{a}\right\}$, and every sufficiently large $n$. The numerator in the quotient is $g(0)-g(1)$ for the function $g(s)=G(s \lambda)$. Writing this difference in the form $-g^{\prime}(0)-\int_{0}^{1}\left(g^{\prime}(s)-g^{\prime}(0)\right) d s$ gives that the numerator is equal to

$$
\begin{equation*}
-\nabla G(0)^{T} \lambda-\int_{0}^{1}(\nabla G(s \lambda)-\nabla G(0))^{T} d s \lambda \tag{21}
\end{equation*}
$$

Here $\nabla G$ is the gradient of $G$, where we only include partial derivatives with respect to coordinates $\lambda_{b b^{\prime}}$ that vary freely, i.e. not the coordinates $\lambda_{b a}$ for which $f_{a}=0$. It suffices to show that the first term is bounded below by a multiple of $\|\lambda\|_{1}$ and that the second is negligible relative to the first, as $n \rightarrow \infty$, uniformly in $f$ in a neighbourhood of $\pi$ and $\lambda$ in a neighbourhood of 0 intersected with $\{\lambda: \lambda \geq 0\}$. Thus it is sufficient to show first that for every coordinate $\lambda_{b b^{\prime}}$ of $\lambda$ minus the partial derivative of $G$ at $\lambda=0$ with respect to $\lambda_{b b^{\prime}}$ is bounded away from 0 , as $n \rightarrow \infty$ uniformly in $f$, and second that every partial derivative is equicontinuous at $\lambda=0$ uniformly in $f$ and large $n$.

We have

$$
\begin{equation*}
G(\lambda)=\frac{1}{2} \sum_{a, a^{\prime}} f_{a}(\lambda)\left(f_{a^{\prime}}(\lambda)-\delta_{a a^{\prime}} / n\right) \tau\left(\frac{\left(R(\lambda) P R(\lambda)^{T}\right)_{a a^{\prime}}-\delta_{a a^{\prime}} e_{a}(\lambda) / n}{f_{a}(\lambda)\left(f_{a^{\prime}}(\lambda)-\delta_{a a^{\prime}} / n\right)}\right) \tag{22}
\end{equation*}
$$

for

$$
\begin{aligned}
f(\lambda) & =f+\sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}}\left(\Delta_{b b^{\prime}} \mathbf{1}\right) \\
R(\lambda) & =\operatorname{Diag}(f)+\sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}} \Delta_{b b^{\prime}}, \\
e_{a}(\lambda) & =\sum_{k} P_{k k} R_{a k}(\lambda)=P_{a a} f_{a}+\sum_{b \neq b^{\prime}} P_{b^{\prime} b^{\prime}} \lambda_{b b^{\prime}}\left(\delta_{a b}-\delta_{a b^{\prime}}\right)
\end{aligned}
$$

By a lengthy calculation, given in Lemma 10,

$$
\begin{equation*}
\frac{\partial}{\partial \lambda_{b b^{\prime}}} G(\lambda)_{\mid \lambda=0}=-\sum_{a} f_{a} K\left(P_{a b^{\prime}} \| P_{a b}\right)+\frac{1}{2 n} K\left(P_{b^{\prime} b^{\prime}} \| P_{b b}\right) \tag{23}
\end{equation*}
$$

for $K(p \| q)=p \log (p / q)+(1-p) \log ((1-p) /(1-q))$ the Kullback-Leibler divergence between the Bernoulli distributions with success probabilities $p$ and $q$. For $f$ sufficiently close to $\pi$ the numbers $f_{a}$ such that $\pi_{a}>0$ are bounded away from zero, and hence $\sum_{a} f_{a} K\left(P_{a b^{\prime}} \| P_{a b}\right)>0$, by identifiability of $(P, \pi)$, since it suffices that just one of the terms of the sum is nonzero. The whole expression is bounded below by the minimum over $\left(b, b^{\prime}\right)$ of these numbers minus $(2 n)^{-1}$ times the maximum of the numbers $K\left(P_{b^{\prime} b^{\prime}} \| P_{b b}\right)$, and hence is positive and bounded away from zero for sufficiently large $n$.

To verify the equicontinuity in $f$ of the partial derivatives, we can compute these explicitly at $\lambda$ and take their limit as $n \rightarrow \infty$. We omit the details of this calculation. However, we note that every term of $G(\lambda)$ is a fixed function of the quadratic forms in $\lambda$

$$
\begin{align*}
\left(f_{a}+\sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}}\left(\Delta_{b b^{\prime}} \mathbf{1}\right)_{a}\right)\left(f_{a^{\prime}}\right. & \left.+\sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}}\left(\Delta_{b b^{\prime}} \mathbf{1}\right)_{a^{\prime}}-\delta_{a a^{\prime}} / n\right)  \tag{24}\\
\left(\left(\operatorname{Diag}(f)+\sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}} \Delta_{b b^{\prime}}\right) P\right. & \left.\left(\operatorname{Diag}(f)+\sum_{b \neq b^{\prime}} \lambda_{b b^{\prime}} \Delta_{b b^{\prime}}^{T}\right)\right)_{a a^{\prime}} \\
& -\frac{\delta_{a a^{\prime}}}{n}\left(P_{a a} f_{a}+\sum_{b \neq b^{\prime}} P_{b^{\prime} b^{\prime}} \lambda_{b b^{\prime}}\left(\delta_{a b}-\delta_{a b^{\prime}}\right)\right) . \tag{25}
\end{align*}
$$

These forms are obviously smooth in $\lambda$, and their dependence and that of their derivatives on $n$ is seen to vanish as $n \rightarrow \infty$. For $f$ and $\lambda$ restricted to neighbourhoods of $\pi$ and 0 , the values of the quadratic forms are restricted to a domain in which the transformation that maps them into $G(\lambda)$ is continuously differentiable. Thus the desired equicontinuity follows by the chain rule.

Lemma 10. The partial derivatives of the function $G$ at 0 defined by (22) are given by (23).

Proof. For given differentiable functions $u$ and $v$ the map $\epsilon \mapsto u(\epsilon) \tau(v(\epsilon) / u(\epsilon)$ ) has derivative $v^{\prime} \log (v /(u-v))-u^{\prime} \log (u /(u-v))$. We apply this for every given pair $\left(a, a^{\prime}\right)$
to the functions $u$ and $v$ obtained by taking $\lambda_{b b^{\prime}}$ in (24) and (25) equal to $\epsilon$ and all other coordinates of $\lambda$ equal to zero. Then

$$
\begin{aligned}
u(0) & =f_{a}\left(f_{a^{\prime}}-\delta_{a a^{\prime}} / n\right) \\
v(0) & =f_{a}\left(f_{a^{\prime}}-\delta_{a a^{\prime}} / n\right) P_{a a^{\prime}} \\
u^{\prime}(0) & =\left(\Delta_{b b^{\prime}} \mathbf{1}\right)_{a}\left(f_{a^{\prime}}-\delta_{a a^{\prime}} / n\right)+f_{a}\left(\Delta_{b b^{\prime}} \mathbf{1}\right)_{a^{\prime}} \\
v^{\prime}(0) & =\left(\Delta_{b b^{\prime}} P\right)_{a a^{\prime}} f_{a^{\prime}}+f_{a}\left(\Delta_{b b^{\prime}} P\right)_{a^{\prime} a}-\left(\delta_{a a^{\prime}} / n\right) P_{b^{\prime} b^{\prime}}\left(\delta_{a b}-\delta_{a b^{\prime}}\right)
\end{aligned}
$$

It follows that $v(0) /(u(0)-v(0))=P_{a a^{\prime}} /\left(1-P_{a a^{\prime}}\right)$, and $u(0) /(u(0)-v(0))=$ $1 /\left(1-P_{a a^{\prime}}\right)$. Hence in view of (17) the partial derivative in (23) is equal to

$$
\frac{1}{2} \sum_{a, a^{\prime}}\left[v^{\prime}(0) \log \frac{P_{a a^{\prime}}}{1-P_{a a^{\prime}}}-u^{\prime}(0) \log \frac{1}{1-P_{a a^{\prime}}}\right]
$$

We combine this with the equalities

$$
\left(\Delta_{b b^{\prime}} \mathbf{1}\right)_{a}=\left\{\begin{array}{ll}
0 & \text { if } a \notin\left\{b, b^{\prime}\right\}, \\
-1 & \text { if } a=b^{\prime}, \\
1 & \text { if } a=b,
\end{array} \quad\left(\Delta_{b b^{\prime}} P\right)_{a a^{\prime}}= \begin{cases}0 & \text { if } a \notin\left\{b, b^{\prime}\right\} \\
-P_{b^{\prime} a^{\prime}} & \text { if } a=b^{\prime} \\
P_{b^{\prime} a^{\prime}} & \text { if } a=b\end{cases}\right.
$$

If $f_{a}$ or $f_{a^{\prime}}$ are zero, then the method to obtain the values found for $v(0) /(u(0)-v(0))$ and $u(0) /(u(0)-v(0))$ in the preceding (substituting the given values of $v(0)$ and $u(0))$ breaks down as we obtain a quotient of zeros. However, the values obtained are still correct when interpreted as the limits from the right at 0 . In (21) and (23) the gradient $\nabla G(0)$ and derivative at $\lambda=0$ may also be interpreted as limits from the right as $\lambda \downarrow 0$ of the gradient. With this substitution the arguments go through. If both $f_{a}$ and $f_{a^{\prime}}$ are zero, the term involving ( $a, a^{\prime}$ ) disappears completely from the analysis.

Lemma 11. If $S$ is fixed and symmetric, $(S, \pi)$ is identifiable and $S>0$ coordinatewise, then there exists $C>0$ such that, for sufficiently small $\delta>0$ and any $\rho_{n} \downarrow 0$,

$$
\liminf _{n \rightarrow \infty} \inf _{0<\|R-\operatorname{Diag}(\pi)\|<\delta} \frac{H_{\rho_{n} S, n}\left(\operatorname{Diag}\left(R^{T} 1\right)\right)-H_{\rho_{n} S, n}(R)}{\rho_{n}\left\|\operatorname{Diag}\left(R^{T} 1\right)-R\right\|} \geq C
$$

Proof. In the notation of the proof of Lemma 9 we must now show that $G(0)-G(\lambda) \geq$ $C \rho_{n}\|\lambda\|_{1}$, as $n \rightarrow \infty$, uniformly in $f$ in a neighbourhood of $\pi$, and $\lambda$ in a positive neighbourhood of 0 . As in that proof we write $G(0)-G(\lambda)$ in the form (21) and see that it suffices that the partial derivatives of $G$ at 0 divided by $\rho_{n}$ tend to negative limits, and that $\|\nabla G(\lambda)-\nabla G(0)\| / \rho_{n}$ becomes uniformly small as $\lambda$ is close enough to zero.

The partial derivative at 0 with respect to $\lambda_{b b^{\prime}}$ is given in (23), where we must replace $P$ by $\rho_{n} S$. Since the scaled Kullback-Leibler divergence $\rho_{n}^{-1} K\left(\rho_{n} s \| \rho_{n} t\right)$ of two Bernoulli laws converges to the Kullback-Leibler divergence $K_{0}(s \| t)=s \log (s / t)+t-s$ between two Poisson laws of means $s$ and $t$, as $\rho_{n} \rightarrow 0$, it follows that for $\rho_{n} \rightarrow 0$, uniformly in $f$,

$$
\frac{1}{\rho_{n}} \frac{\partial}{\partial \lambda_{b b^{\prime}}} G(\lambda)_{\mid \lambda=0} \rightarrow-\sum_{a} f_{a} K_{0}\left(S_{a b^{\prime}} \| S_{a b}\right)
$$

The right side is strictly negative for $f$ close to $\pi$, by the assumption of identifiability of $(S, \pi)$.

If $P=\rho_{n} S$, then the function $\lambda \mapsto v(\lambda)$ given in (25) takes the form $v=\rho_{n} v_{S}$, for $v_{S}$ defined in the same way but with $S$ replacing $P$. The function $u$ given in (24) does not depend on $P$ or $S$. Using again that the derivative of the map $\epsilon \mapsto u(\epsilon) \tau(v(\epsilon) / u(\epsilon))$ is given by $v^{\prime} \log (v /(u-v))-u^{\prime} \log (u /(u-v))$, we see that the partial derivative with respect to $\lambda_{b b^{\prime}}$ of the $\left(a, a^{\prime}\right)$ term in the sum defining $G$ takes the form

$$
\begin{aligned}
& \rho_{n} v_{S}^{\prime} \log \frac{\rho_{n} v_{S}}{u-\rho v_{S}}-u^{\prime} \log \frac{u}{u-\rho_{n} v_{S}} \\
& \quad=\rho_{n} v_{S}^{\prime} \log \rho_{n}-\rho_{n} v_{S}^{\prime} \log \left(v_{S} / u\right)-\left(\rho_{n} v_{S}^{\prime}-u^{\prime}\right) \log \left(1-\rho_{n} v_{S} / u\right)
\end{aligned}
$$

Here $u$ and $v_{S}$ are as in (24) and (25) (with $P$ replaced by $S$ ), and depend on ( $a, a^{\prime}$ ). From the fact that the column sums of the matrices $R(\lambda)$ do not depend on $\lambda$, we have that

$$
\sum_{a, a^{\prime}}\left[\left(R(\lambda) S R(\lambda)^{T}\right)_{a a^{\prime}}-\frac{\delta_{a a^{\prime}}}{n} \sum_{k} P_{k k} R(\lambda)_{a k}\right]=R(\lambda)^{T} \mathbf{1} S R(\lambda)^{T} \mathbf{1}-\sum_{k} P_{k k} \sum_{a} R(\lambda)_{a k}
$$

is constant in $\lambda$. This shows that $\sum_{a, a^{\prime}} v_{S}^{\prime}=0$ and hence the contribution of the term $\rho_{n} v_{S}^{\prime} \log \rho_{n}$ to the partial derivatives of $G$ vanishes. The term $-\left(\rho_{n} v_{S}^{\prime}-u^{\prime}\right) \log (1-$ $\left.\rho_{n} v_{S} / u\right)$ can be expanded as $\left(\rho_{n} v_{S}^{\prime}-u^{\prime}\right) \rho_{n} v_{S} / u$ up to $\mathcal{O}\left(\rho_{n}^{2}\right)$, uniformly in $f$ and $\lambda$. Since these are equicontinuous functions of $\lambda$, it follows that $\rho_{n}^{-1}(\nabla G(\lambda)-\nabla G(0))$ becomes arbitrarily small if $\lambda$ varies in a sufficiently small neighbourhood of 0 .

Lemma 12 shows that in a neighborhood of the truth, there is not much variation in the differences between the observed modularity and the modularity evaluated on the expected number of connections between classes given the true labelling.
Lemma 12. There exists a constant $c>0$ such that for $X(e)$ as in (18), for every twice differentiable function $t_{a, b}:[0,1] \rightarrow \mathbb{R}$ with $\left\|t_{a, b}^{\prime}\right\|_{\infty} \vee\left\|t_{a, b}^{\prime \prime}\right\|_{\infty} \leq 1$, and every $x>0$,

$$
\operatorname{Pr}\left(\max _{e: \#\left(e_{i} \neq Z_{i}\right) \leq m}\|X(e)-X(Z)\|_{\infty}>x\right) \leq 6\binom{n}{m} K^{m+2} e^{-\frac{c x^{2} n^{2}}{m\|P\|_{\infty} / n+x}}
$$

Proof. Given $Z$ there are at most $\binom{n}{m}$ groups of $m$ candidate nodes that can be assigned to have $e_{i} \neq Z_{i}$, and the label of each node can be chosen in at most $K-1$ ways. Thus conditioning the probability on $Z$, we can use the union bound to pull out the maximum over $e$, giving a sum of fewer than $\binom{n}{m} K^{m}$ terms. Next we pull out the norm giving another factor $K^{2}$. It suffices to combine this with a tail bound for a single variable $X_{a, b}(e)-X_{a, b}(Z)$. Write $t$ for $t_{a, b}$.

Assume for simplicity of notation that $e_{i}=Z_{i}$, for $i>m$, and decompose

$$
\begin{aligned}
\frac{1}{n^{2}} O_{a b}(e) & =\frac{1}{n^{2}}\left[\sum_{i \leq m \text { or } j \leq m} A_{i j} 1_{e_{i}=a, e_{j}=b}+\sum_{i>m \text { and } j>m} A_{i j} 1_{e_{i}=a, e_{j}=b}\right] \\
& =: S_{1}+S_{2}
\end{aligned}
$$

Let $O_{a b}(Z) / n^{2}=: S_{1}^{\prime}+S_{2}$, with the same variable $S_{2}$, be the corresponding decomposition if $e$ is changed to $Z$, and then decompose, where the expectation signs $\mathbb{E}$ denote conditional expectations given $Z$,

$$
\begin{aligned}
X_{a b}(e) & -X_{a b}(Z) \\
= & \left(t\left(S_{1}+S_{2}\right)-t\left(\mathbb{E} S_{1}+\mathbb{E} S_{2}\right)\right)-\left(t\left(S_{1}^{\prime}+S_{2}\right)-t\left(\mathbb{E} S_{1}^{\prime}+\mathbb{E} S_{2}\right)\right) \\
= & t\left(S_{1}+S_{2}\right)-t\left(\mathbb{E} S_{1}+S_{2}\right) \\
& \quad+\left(t\left(\mathbb{E} S_{1}+S_{2}\right)-t\left(\mathbb{E} S_{1}+\mathbb{E} S_{2}\right)\right)-\left(t\left(\mathbb{E} S_{1}^{\prime}+S_{2}\right)-t\left(\mathbb{E} S_{1}^{\prime}+\mathbb{E} S_{2}\right)\right) \\
& +t\left(\mathbb{E} S_{1}^{\prime}+S_{2}\right)-t\left(S_{1}^{\prime}+S_{2}\right)
\end{aligned}
$$

The first and third terms on the far right can be bounded above in absolute value by $\left\|t^{\prime}\right\|_{\infty}$ times the increment. To estimate the second term we write it as

$$
\left(S_{2}-\mathbb{E} S_{2}\right)\left(\mathbb{E} S_{1}-\mathbb{E} S_{1}^{\prime}\right) \int_{0}^{1} \int_{0}^{1} t^{\prime \prime}\left(u S_{2}+(1-u) \mathbb{E} S_{2}+v \mathbb{E} S_{1}+(1-v) \mathbb{E} S_{1}^{\prime}\right) d u d v
$$

Since the first and second derivatives of $t$ are uniformly bounded by 1 , it follows that

$$
\left|X_{a b}(e)-X_{a b}(Z)\right| \leq\left|S_{1}-\mathbb{E} S_{1}\right|+\left|S_{2}-\mathbb{E} S_{2}\right|\left|\mathbb{E} S_{1}-\mathbb{E} S_{1}^{\prime}\right|+\left|S_{1}^{\prime}-\mathbb{E} S_{1}^{\prime}\right|
$$

The variable $S_{1}-\mathbb{E} S_{1}$ is a sum of fewer than $2 m n$ independent variables, each with conditional mean zero, bounded above by $1 / n^{2}$ and of variance bounded above by $\|P\|_{\infty} / n^{4}$. Therefore Bernstein's inequality gives that

$$
\mathbb{P}\left(\left|S_{1}-\mathbb{E} S_{1}\right|>x\right) \leq e^{-\frac{1}{2} x^{2} /\left(2 m n\|P\|_{\infty} / n^{4}+x /\left(3 n^{2}\right)\right)}
$$

This is as the exponential factor in the bound given by the lemma, for appropriate $c$. The variable $S_{1}^{\prime}-\mathbb{E} S_{1}^{\prime}$ can be bounded similarly. Furthermore $\left|\mathbb{E} S_{1}-\mathbb{E} S_{1}^{\prime}\right| \leq 4 \mathrm{mn} / n^{2}=$ $4 m / n$, and $S_{2}-\mathbb{E} S_{2}$ is the sum of fewer than $n^{2}$ variables as before, so that

$$
\mathbb{P}\left(\left|S_{2}-\mathbb{E} S_{2}\right|\left|\mathbb{E} S_{1}-\mathbb{E} S_{1}^{\prime}\right|>x\right) \leq e^{-\frac{1}{2}(x n /(4 m))^{2} /\left(n^{2}\|P\|_{\infty} / n^{4}+x n /\left(12 m n^{2}\right)\right)}
$$

The exponent has a similar form as before, except for an additional factor $n / m \geq 1$.

## References

Abbe, E., Bandeira, A. S., and Hall, G. (2014). "Exact Recovery in the Stochastic Block Model." ArXiv:1405.3267v4. MR3447993. doi: https://doi.org/10.1109/ TIT.2015.2490670. 775, 776
Airoldi, E. M., Blei, D. M., Fienberg, S. E., and Xing, E. P. (2008). "Mixed Membership Stochastic Blockmodels." Journal of Machine Learning Research, 9: 1981-2014. 767

Bickel, P. J. and Chen, A. (2009). "A Nonparametric View of Network Models and Newman-Girvan and Other Modularities." Proceedings of the National Academy of Sciences of the United States of America, 106(50): 21068-21073. 767, 768, 769, 771, $772,773,775,776,777,778,781,782,783,788$

Bickel, P. J., Chen, A., Zhao, Y., Levina, E., and Zhu, J. (2015). "Correction to the Proof of Consistency of Community Detection." The Annals of Statistics, 43(1): 462-466. MR3311866. doi: https://doi.org/10.1214/14-AOS1271. 768, 776, 781

Channarond, A., Daudin, J.-J., and Robin, S. (2012). "Classification and Estimation in the Stochastic Blockmodel Based on the Empirical Degrees." Electronic Journal of Statistics, 6: 2574-2601. MR3020277. doi: https://doi.org/10.1214/12-EJS753. 767, 768, 776, 779

Chen, K. and Lei, J. (2014). "Network Cross-Validation for Determining the Number of Communities in Network Data." ArXiv:1411.1715v1. 769

Chen, Y. and Xu, J. (2014). "Statistical-Computational Tradeoffs in Planted Problems and Submatrix Localization with a Growing Number of Clusters and Submatrices." ArXiv:1402.1267v2. MR3491121. 775, 776

Côme, E. and Latouche, P. (2014). "Model Selection and Clustering in Stochastic Block Models with the Exact Integrated Complete Data Likelihood." ArXiv:1303.2962. MR3441229. doi: https://doi.org/10.1177/1471082X15577017. 777

Csardi, G. and Nepusz, T. (2006). "The igraph Software Package for Complex Network Research." InterJournal Complex Systems, 1695. 778

Gao, C., Ma, Z., Zhang, A. Y., and Zhou, H. H. (2015). "Achieving Optimal Misclassification Proportion in Stochastic Block Model." ArXiv:1505.03772v5. 776

Gao, C., Ma, Z., Zhang, A. Y., and Zhou, H. H. (2016). "Community Detection in Degree-Corrected Block Models." ArXiv:1607.06993. 776

Ghosal, S., Ghosh, J. K., and van der Vaart, A. W. (2000). "Convergence rates of posterior distributions." The Annals of Statistics, 28(2): 500-531. MR1790007. doi: https://doi.org/10.1214/aos/1016218228. 776
Glover, F. (1989). "Tabu Search - Part I." ORSA Journal on Computing, 1(3): 190-206. 777

Hayashi, K., Konishi, T., and Kawamoto, T. (2016). "A Tractable Fully Bayesian Method for the Stochastic Block Model." ArXiv:1602.02256v1. 768

Hofman, J. M. and Wiggins, C. H. (2008). "Bayesian Approach to Network Modularity." Physical Review Letters, 100: 258701. 768

Holland, P. W., Laskey, K. B., and Leinhardt, S. (1983). "Stochastic Blockmodels: First Steps." Social Networks, 5: 109-137. MR0718088. doi: https://doi.org/ 10.1016/0378-8733(83)90021-7. 767

Jin, J. (2015). "Fast Community Detection by SCORE." The Annals of Statistics, 43(1): 57-89. MR3285600. doi: https://doi.org/10.1214/14-AOS1265. 768

Karrer, B. and Newman, M. E. J. (2011). "Stochastic Blockmodels and Community Structure in Networks." Physical Review E, 83: 016107. MR2788206. doi: https://doi.org/10.1103/PhysRevE.83.016107. 767

Kpogbezan, G. B., van der Vaart, A. W., van Wieringen, W. N., Leday, G. G. R., and van de Wiel, M. A. (2016). "An empirical Bayes approach to network recovery using external knowledge." ArXiv:1605.07514. 769

Lei, J. and Rinaldo, A. (2015). "Consistency of Spectral Clustering in Stochastic Block Models." The Annals of Statistics, 43(1): 215-237. MR3285605. doi: https://doi.org/10.1214/14-AOS1274. 768, 775
McDaid, A. F., Brendan Murphy, T., Friel, N., and Hurley, N. J. (2013). "Improved Bayesian Inference for the Stochastic Block Model with Application to Large Networks." Computational Statistics and Data Analysis, 60: 12-31. MR3007016. doi: https://doi.org/10.1016/j.csda.2012.10.021. 769, 771, 772, 777

Mossel, E., Neeman, J., and Sly, A. (2012). "Reconstruction and Estimation in the Planted Partition Model." ArXiv:11202.1499v4. MR3383334. doi: https://doi.org/10.1007/s00440-014-0576-6. 775

Newman, M. and Girvan, M. (2004). "Finding and Evaluating Community Structure in Networks." Physical Review E, 69: 026113. MR1975193. doi: https://doi.org/ 10.1103/PhysRevE.67.026126. 767

Nowicki, K. and Snijders, T. A. B. (2001). "Estimation and Prediction for Stochastic Blockstructures." Journal of the American Statistical Association, 96(455): 10771087. MR1947255. doi: https://doi.org/10.1198/016214501753208735. 768, 771, 772, 777

Park, Y. and Bader, J. S. (2012). "How Networks Change with Time." Bioinformatics, 28(12): i40-i48. 767
Pati, D. and Bhattacharya, A. (2015). "Optimal Bayesian Estimation in Stochastic Block Models." ArXiv:1505.06794. 768

Robbins, H. (1955). "A Remark on Stirling's Formula." The American Mathematical Monthly, 62(1): 26-29. MR0069328. doi: https://doi.org/10.2307/2308012. 785
Rohe, K., Chatterjee, S., and Yu, B. (2011). "Spectral Clustering and the HighDimensional Stochastic Blockmodel." The Annals of Statistics, 39(4): 1878-1915. MR2893856. doi: https://doi.org/10.1214/11-AOS887. 768

Saldana, D. F., Yu, Y., and Feng, Y. (2014). "How Many Communities Are There?" ArXiv:1412.1684v1. MR3610418. doi: https://doi.org/10.1080/10618600. 2015.1096790. 769

Sarkar, P. and Bickel, P. J. (2015). "Role of Normalization in Spectral Clustering for Stochastic Blockmodels." The Annals of Statistics, 43(3): 962-990. MR3346694. doi: https://doi.org/10.1214/14-AOS1285. 768

Snijders, T. A. and Nowicki, K. (1997). "Estimation and Prediction for Stochastic Blockmodels for Graphs with Latent Block Structure." Journal of Classification, 14: 75100. MR1449742. doi: https://doi.org/10.1007/s003579900004. 767, 768

Suwan, S., Lee, D. S., Tang, R., Sussman, D. L., Tang, M., and Priebe, C. E. (2016). "Empirical Bayes estimation for the stochastic blockmodel." Electronic Journal of Statistics, 10: 761-782. MR3477741. doi: https://doi.org/10.1214/16-EJS1115. 768

Wang, Y. X. R. and Bickel, P. J. (2015). "Likelihood-Based Model Selection for Stochastic Block Models." ArXiv:1502.02069v1. MR3196592. doi: https://doi . org/ 10.1080/15598608.2013.771546. 769

Zachary, W. W. (1977). "An Information Flow Model for Conflict and Fission in Small Groups." Journal of Anthropological Research, 33(4): 452-473. 777, 778, 779

Zhang, A. Y. and Zhou, H. H. (2015). "Minimax Rates of Community Detection in Stochastic Block Models." Preprint available at http://www.stat.yale.edu/~hz68/ CommunityDetection.pdf. MR3546450. doi: https://doi.org/10.1214/15AOS1428. 775, 776

Zhao, Y., Levina, E., and Zhu, J. (2012). "Consistency of Community Detection in Networks under Degree-Corrected Stochastic Block Models." The Annals of Statistics, 40(4): 2266-2292. MR3059083. doi: https://doi.org/10.1214/12-AOS1036. 767, $768,776,777,781$


[^0]:    *Research supported by Netherlands Organization for Scientific Research NWO.
    ${ }^{\dagger}$ The research leading to these results has received funding from the European Research Council under ERC Grant Agreement 320637.
    ${ }^{\ddagger}$ Mathematical Institute, Leiden University, svdpas@math.leidenuniv.nl
    ${ }^{\text {§ Mathematical Institute, Leiden University, avdvaart@math.leidenuniv.nl }}$

