Empirical Reference Distributions for Networks of Different Size

ANNA SMITH^a, CATHERINE A. CALDER^{a,*}, CHRISTOPHER R. BROWNING^b

^a*Department of Statistics, The Ohio State University*

^b*Department of Sociology, The Ohio State University*

March 4, 2016

Abstract

Network analysis has become an increasingly prevalent research tool across a vast range of scientific fields. Here, we focus on the particular issue of comparing network statistics, i.e. graph-level measures of network structural features, across multiple networks that differ in size. Although “normalized” versions of some network statistics exist, we demonstrate via simulation why direct comparison is often inappropriate. We consider normalizing network statistics relative to a simple fully parameterized reference distribution and demonstrate via simulation how this is an improvement over direct comparison, but still sometimes problematic. We propose a new adjustment method based on a reference distribution constructed as a mixture model of random graphs which reflect the dependence structure exhibited in the observed networks. We show that using simple Bernoulli models as mixture components in this reference distribution can provide adjusted network statistics that are relatively comparable across different network sizes but still describe interesting features of networks, and that this can be accomplished at relatively low computational expense. Finally, we apply this methodology to a collection of ecological networks derived from the Los Angeles Family and Neighborhood Survey activity location data.

Keywords: *ERGM, L.A.FANS, mixture model, network comparison, normalized network statistics*

Acknowledgements

Support for this work was provided by grants from the National Science Foundation (NSF DMS-1209161), the National Institute of Health (NIH R01DA032371), the William T. Grant Foundation, and The Ohio State University Institute for Population Research (NIH P2CHD058484).

* Corresponding author: calder@stat.osu.edu

Empirical Reference Distributions for Networks of Different Size

March 4, 2016

Abstract

Network analysis has become an increasingly prevalent research tool across a vast range of scientific fields. Here, we focus on the particular issue of comparing network statistics, i.e. graph-level measures of network structural features, across multiple networks that differ in size. Although “normalized” versions of some network statistics exist, we demonstrate via simulation why direct comparison is often inappropriate. We consider normalizing network statistics relative to a simple fully parameterized reference distribution and demonstrate via simulation how this is an improvement over direct comparison, but still sometimes problematic. We propose a new adjustment method based on a reference distribution constructed as a mixture model of random graphs which reflect the dependence structure exhibited in the observed networks. We show that using simple Bernoulli models as mixture components in this reference distribution can provide adjusted network statistics that are relatively comparable across different network sizes but still describe interesting features of networks, and that this can be accomplished at relatively low computational expense. Finally, we apply this methodology to a collection of ecological networks derived from the Los Angeles Family and Neighborhood Survey activity location data.

Keywords: *ERGM, L.A.FANS, mixture model, network comparison, normalized network statistics*

1. Introduction

Networks have become ubiquitous as a tool for analysis in many areas of applied research. Here, we focus on the situation where the researcher is interested in considering and comparing multiple networks. Often the networks under consideration consist of differing numbers of nodes, and the comparison of statistics across these networks is not straightforward. That is, the distribution of a network statistic across networks may be very different, depending on the size of the network. This holds even when the networks are generated from a single model, as shown in Figure 1.

A tool for network comparisons would be useful in many fields - from traditional social network analysis to brain network studies in biomedical research. For example, many researchers have compared friendship networks within different classrooms in an effort to infer patterns of friendship formation and their role in academic performance (Lubbers, 2003; Lubbers and Snijders, 2007). Similarly, a growing area of study in biomedical research attempts to draw inference about patterns of disease, illness, and treatment effectiveness from brain networks by comparing networks across affected and unaffected patients (e.g. Bassett et al., 2008; He et al., 2008). Of course, many similar potential applications exist: comparing e-mail communication networks among employees in different departments of a company, comparing citation networks among researchers in different fields, comparing voting networks among politicians from different state congresses, etc.

Further, it would be advantageous to have measures of network structural features that can be included as covariates in typical regression models. For example, it might be interesting to relate the average shortest path length of an individual’s brain network to his/her stage of Alzheimer’s disease, or to relate the degree of transitivity (proportion of times the co-votee of my co-votee also co-votes with me) in a state congress’ voting network to the number of bipartisan legislative acts passed in that state. In order to include such network measures in a regression-style analysis, the measures must be on the same scale and have the same meaning across the various networks. However, when the networks being compared vary in size, often the values of these network statistics can be dominated by size effects.

We first demonstrate via simulation why direct comparison of network statistics across networks of different sizes is often inappropriate. We then investigate the use of an absolute reference distribution which is fully parameterized a priori, here choosing to compare to completely random graphs, or Erdős-Rényi graphs, of the appropriate size (e.g. Van Wijk et al., 2010; Stam, 2004; Schindler et al., 2008; Smit et al., 2008). We demonstrate via simulation why this is an improvement over direct comparison, but still sometimes problematic. We argue that a more appropriate method of comparison would both adjust for network size as well as incorporate some degree of appropriate network dependence structure. We propose utilizing

Figure 1: Realizations from a Single Model

The distribution of a network statistic (here, closeness centrality) across a range of network sizes (number of nodes, n).

a new reference distribution which comes from a mixture model, where we are mixing over the dependence structure demonstrated in a random subset of the empirically observed networks. To accomplish this, we consider three models for reference distribution simulations: the simple Bernoulli model, a mean degree preserving Bernoulli model (Krivitsky et al., 2011), and a Hierarchical Bernoulli model, where the latter is a special case of the HERGM, or Hierarchical ERGM, proposed by Schweinberger and Handcock (2015). We consider the performance of these methods in a series of simulation studies. Finally, we apply these techniques to neighborhood network data from the Los Angeles Family and Neighborhood Survey (L.A.FANS) to demonstrate the type of substantive statements that these methods can facilitate.

2. Background and Motivation

We consider networks or graphs, G , that consist of n nodes or vertices and represent them by their $n \times n$ adjacency matrix, $\mathbf{Y}$. Each element of this matrix, Y_{ij} , is a tie indicator variable such that

$$Y_{ij} = \begin{cases} 1 & \text{if nodes } i \text{ and } j \text{ are tied} \\ 0 & \text{otherwise} \end{cases} \quad \text{for } i, j = 1, \dots, n$$

For simplicity, we consider only one-mode, undirected and unweighted networks, and so we let $Y_{ii} = 0$ by definition. Certainly, the methodology proposed here could be extended to other types of networks. We refer to the number of nodes, n , as network size, but note that others denote this property as the *order* of the graph (e.g. Kolaczyk and Csárdi, 2014).

Generally, the issue of network comparisons has recieved little methodological attention, despite the popularity of network analysis. Further, much of researchers' attention has been focused on building methods to model multiple networks simultaneously and using this modeling framework as a way to proceed with any questions of network comparison, similar to a traditional meta-analysis approach (Snijders and Baerveldt, 2003; An, 2015). This style of analysis requires faith in the individual micro-level analyses themselves. However, working with existing network models can sometimes be difficult since these models (even in the simplest cases): often suffer from degeneracy (for some discussion, see Snijders et al., 2006; Chatterjee et al., 2013), can have dependence structures that are difficult to interpret (such as the ERGM terms suggested by Snijders et al., 2006), face computational issues, and, most importantly, can exhibit a rather striking lack of connection between model parameters and the specific structural features of networks generated from these models. In particular, many common network model parameters are often highly correlated so that individual parameters play different roles under different contexts. Consider the example highlighted by Snijders et al. (2006) of an ERGM with parameters for only edges and triangles: "for a fixed positive transitivity parameter, as the edge parameter becomes more negative there is a point at which... [simulated networks] change dramatically... from only complete graphs to only low density graphs". In this simple ERGM, the edges parameter does not seem to affect network density in the way that we might think, or at least the strength and smoothness of its effect appears to depend on the particular model specification itself. Thus, even in simple versions of existing network models, there are cases where model parameters do not appear to characterize the structural aspects of networks which they seem designed to describe. For this reason, when comparing multiple networks, we propose comparing direct measures of network structure itself - network or graph statistics - as a simple alternative to comparing fitted model parameter values whose interpretations are often not straight forward. The methodology we propose here could also be valuable as an additional viewpoint in combination with the existing meta-analysis style technique (e.g. as deccribed by Snijders and Baerveldt, 2003; An, 2015), where our proposed methodology could perhaps guide the specification of micro-level analyses and help to improve overall interpretability.

2.1. Common Network Statistics

We consider network statistics that are popular in the applied network science literature (Anderson et al., 1999; Smith and Moody, 2013). The statistics we consider here can be thought of as describing two different aspects of a network: centralization and topology. We consider degree, closeness, and betweenness as measures of centralization while our considered measures of topology are average path length and transitivity. We also consider network density.

All statistics are computed (and compared) at the graph level. For the centralization measures, we use Freeman’s formula (Freeman, 1979):

$$C(G) = \sum_{j=1}^n \left\{ \max_{1 \leq i \leq n} s(i) - s(j) \right\}$$

for any vertex-level statistic, $s(i)$. Note that at the graph level, $C(G)$ is a measure of the dispersion of the distribution of $s(i)$, across all vertices in the graph. We consider this “unnormalized” version of our centralization measures as well as the “normalized” version in which $C(G)$ is divided by the maximum possible value attainable by $C(G)$ for a graph of the same size. Unfortunately, the normalized version proposed by Butts (2006) which adjusts for both the size and density of G has not yet been implemented in standard software to our knowledge and so it will not be considered here.

Degree Centrality at the vertex level is simply a node’s degree, or the number of other nodes to which that node is tied:

$$s_{\text{deg}}(i) = \sum_{j \neq i} Y_{ij}.$$

Thus, at the vertex level, this statistic reports the size of the i th node’s personal or immediate network so that at the graph level, we measure the dispersion in personal network size.

Closeness Centrality at the vertex level is the reciprocal of the sum of the shortest path lengths from that node to all other nodes:

$$s_{\text{clo}}(i) = \frac{1}{\sum_{j \neq i} d_{ij}}$$

where d_{ij} is the length of the shortest path from node i to node j . Intuitively, $s_{\text{clo}}(i)$ captures the speed or efficiency with which information, disease, etc. could be spread from vertex i to the rest of the nodes in G . At the graph level, we measure the amount that this efficiency varies across all nodes.

Betweenness Centrality at the vertex level is the proportion of shortest paths between all pairs of nodes that pass through the node of interest:

$$s_{\text{btw}}(i) = \sum_{k \neq j; k, j \neq i} \frac{g_{kij}}{\tilde{g}_{kj}}$$

where g_{kij} is the number of paths from node k to node j that pass through node i , while $\tilde{g}_{kj}$ counts all paths from k to j . Adding an additional degree of complexity to closeness centrality, betweenness centrality at the vertex level measures the importance of the i th node as a central node in efficient communication or spread across G . Again, at the graph level, we get the degree of variation in this measure of a node’s importance.

Average Path Length is a graph-level measure and is defined as the average of the shortest path lengths between all pairs of nodes:

$$P(G) = \frac{1}{n-1} \sum_{i=1}^n \sum_{j \neq i} d_{ij}.$$

Similar to closeness centrality, we again are interested in the speed or efficiency with which nodes are able to transmit information across G . Recall that at the graph level, closeness centrality measures the dispersion of these efficiencies. Average path length, however, instead captures the “average efficiency”, where graphs with larger $P(G)$ are considered less efficient than graphs with smaller $P(G)$.

Transitivity, also called Clustering, at the vertex level is the proportion of 2-stars centered at node i which are closed triangles. At the graph level, we simply report the average:

$$T(G) = \frac{\# \text{ of closed triangles}}{\# \text{ of 2-stars}}$$

where triangles and 2-stars are the subgraph types depicted in Figure 2. This statistic is extremely popular in applied social network research, largely due to its relevance to much of social theory. Intuitively, $T(G)$ reflects the average amount of local

Figure 2: Subgraphs included in Transitivity Network Statistic

clustering in G .

Density is a graph-level measure of the proportion of all possible ties that are realized in the observed network:

$$D(G) = \frac{1}{\binom{n}{2}} \sum_{i < j} Y_{ij}.$$

Sometimes treated as a fixed attribute, similarly to the number of nodes, we instead prefer to treat density in the same manner as all other network statistics examined here. Note that other work sometimes uses the terms density and mean degree interchangeably (e.g Anderson et al., 1999; Van Wijk et al., 2010), where mean degree is $\frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} Y_{ij} = (n-1) D(G)$.

Despite the popularity of these common network statistics, there has been very little research that addresses their dependence on network size. Anderson et al. (1999) examine the distributions of a variety of graph-level network statistics (including degree and betweenness centrality) across a range of network sizes. However, they consider only networks generated from one particular model: the conditional uniform graph (CUG), where they allow both the network size and mean degree to vary. Generally, a CUG model is a uniform distribution across all possible graphs that share some common feature, such as degree distribution, density, etc. Using a CUG model conditional on shared network size and mean degree, Anderson et al. provide evidence that the distributions of these network statistics clearly depend on network size, and that they are also influenced by network density. Continuing in this line of work, we will examine distributions of these common network statistics across a range of generative network models, including some that better mimic real world network data than conditional uniform graphs.

Faust (2006) examines a wide range of real world networks (across a variety of applications and settings) to show that the triad census can be largely described by lower order graph properties, including size and density. The triad census is an example of a subgraph census, a commonly used measure of local graph structure. More specifically, the triad census is the collection of counts for each possible configuration of three (unlabeled) nodes across all triads in the graph of interest. Faust’s result that the triad census can be explained by lower order graph properties implies that other, higher-order network properties (perhaps such as the common network statistics mentioned here) might also be heavily influenced by network size.

Here, we will be primarily building off of the work of Van Wijk et al. (2010). The authors provide a detailed summary of the problems faced in comparing network statistics across multiple networks as well as a brief summary of a range of possible solutions, all through the lens of neuroscience research. They focus solely on the topological measures, average path length and transitivity, as those capture network features important in brain network theory. In fact, the authors provide closed form adjustments for these two statistics for networks simulated from particular network models (Van Wijk et al., 2010, Table 1). However, as Van Wijk et al. point out, the underlying generative model is typically unknown for empirical, real world networks.

Besides these few works, the issue of comparing network statistics across networks has recieved little methodological attention. Perhaps this is the result of simply ignoring a lesser-known issue. In fact, many of the network statistics examined here *appear* to be “normalized” in some way that incorporates network size, so that it is certainly plausible that a reasearcher might assume that these statistics ought to be comparable across sizes. In the case of the centralization measures, “normalized” versions are acheived by dividing by the largest possible value attainable on a graph *of that size*. The topological measures are versions of averages and are thus at least roughly adjusted for size effects. However, as we will see in the following section, these simple adjustments are not enough to account for what appears to be a complicated dependence between these common network statistics and network size.

2.2. Generative Models for Networks

One might hope that statistics computed on networks that come from the same generative model would have comparable values, regardless of network size. However, as has been shown for relatively simple classes of generative models (Anderson et al., 1999; Van Wijk et al., 2010) and as we will see for a range of other models, this is certainly not the case. To examine this issue, we simulate groups of networks that cover a range of sizes *from the same model* and examine the distribution of the common network statistics across the range of network sizes.

In an effort to capture the variety in structure, density, topology, etc. which is demonstrated in real world networks, we consider six network models of increasing complexity: the Erdős-Rényi model, a Bernoulli model, a mean degree preserving Bernoulli model (Krivitsky et al., 2011), a Markov ERGM, a Hierarchical Bernoulli model and a Hierarchical Markov ERGM. For simplicity, we will represent each model as a probability distribution for a random adjacency matrix, $\mathbf{Y}$. Alternatively, note that this gives us a joint distribution for the set of possible edges, $\{Y_{ij} : i < j; i, j = 1, \dots, n\}$. Since we are dealing exclusively with undirected graphs, we need only model $\frac{1}{2}$ of the $\binom{n}{2}$ possible edge variables, given that $Y_{ji} = Y_{ij} \forall i, j$ by definition in undirected networks.

Before describing these generative models in detail, we briefly outline how they will be used in the simulation studies in Section 2.3, where we monitor the performance of directly comparing statistics across networks of different size. For each generative model, we simulate $N_n = 200$ replicates of graphs with $n = 20, 30, \dots, 100$ nodes, so that each model gives rise to a total of $N = 1800$ graphs. This approach of using simulated network data allows us much greater levels of replication within network sizes than would working with any collection of real world networks.

Erdős-Rényi Model. This model is included as a simple baseline. Although there are a few formulations (Erdős and Rényi, 1960; Gilbert, 1959), we focus on the version that is equivalent to a Bernoulli graph with the probability of a tie, p , set to 0.50. In a Bernoulli graph, each tie indicator variable is modeled as an independent, identically distributed Bernoulli random variable, with $P(Y_{ij} = 1) = p$. This implies that

$$P(\mathbf{Y} = \mathbf{y} | \theta_{edge}) = \exp \left\{ \theta_{edge} h_{edge}(\mathbf{y}) - \psi(\theta_{edge}) \right\} \quad (1)$$

where $\theta_{edge} = \log \left(\frac{p}{1-p} \right)$, $h_{edge}(\mathbf{y})$ counts the number of edges in the observed network, and $\psi(\theta_{edge})$ is the normalizing term, which is calculated as follows:

$$\psi(\theta_{edge}) = \sum_{\mathbf{y}^* \in \mathcal{Y}} \exp \left\{ \theta_{edge} h_{edge}(\mathbf{y}^*) \right\}$$

where $\mathcal{Y}$ is the set of all possible graphs of size n . Although much is known about the properties and behavior of this model, it does not tend to reproduce the type of structure or topological features that we typically observe in real world networks.

Bernoulli Model. We can generalize the Erdős-Rényi model by changing the tie density, so that it better reflects what we often see in real world networks. Here, we set the probability of a tie, p , to be 0.20. Although we obtain reasonable density under this model, we would still not expect to see the type of topological structure that we associate with real world networks. This follows from the Bernoulli model's (unrealistic) assumption that all tie variables are independent.

Mean Degree Preserving Bernoulli Model. We implement the offset terms suggested by Krivitsky et al. (2011), which results in an reparameterized version of the traditional Bernoulli model:

$$P(\mathbf{Y} = \mathbf{y} | \theta_{deg}) = \exp \left\{ \theta_{deg} h_{edge}(\mathbf{y}) + \log \left(\frac{1}{n} \right) h_{edge}(\mathbf{y}) - \psi(\theta_{deg}) \right\} \quad (2)$$

where $h_{edge}(\mathbf{y})$ counts the number of edges in the observed network as before, and $\psi(\theta_{deg})$ is the normalizing constant:

$$\psi(\theta_{deg}) = \sum_{\mathbf{y}^* \in \mathcal{Y}} \exp \left\{ \theta_{deg} h_{edge}(\mathbf{y}^*) + \log \left(\frac{1}{n} \right) h_{edge}(\mathbf{y}^*) \right\}$$

The authors show that without an offset term, a Bernoulli model preserves density as network size increases and argue that this is not an appropriate effect for many real-world social networks (e.g. we would not expect the observed average number of friendships per person to increase at the same rate as the sample size). Instead, models that include the proposed offset are designed to maintain mean degree¹ as n increases. Krivitsky et al. (2011) show that under this model, the degree distribution converges to a Poisson distribution as $n \rightarrow \infty$, with average degree $\exp(\theta_{deg})$. We choose θ_{deg} such that the

¹Note that mean degree scales linearly with density by a factor of $n - 1$

Figure 3: Realizations from Simulated Datasets, $n = 20$

average degree is three.

Markov ERGM. Exponential random graph models (ERGMs) provide an easy way to introduce dependence between tie variables. Here, we use the Markov version of an ERGM which specifies that all tie variables which share a node are dependent (Frank and Strauss, 1986). The triad version of this model is formulated as follows:

$$P(Y = \mathbf{y}|\boldsymbol{\theta}) = \exp \{ \langle \boldsymbol{\theta}, \mathbf{h}(\mathbf{y}) \rangle - \psi(\boldsymbol{\theta}) \} \quad (3)$$

where $\langle \mathbf{x}_1, \mathbf{x}_2 \rangle$ denotes the inner product of $\mathbf{x}_1$ and $\mathbf{x}_2$, the elements of $\mathbf{h}(\mathbf{y})$ are counts for the number of edges, 2-stars, and triangles in the observed network $\mathbf{y}$, respectively, i.e. $\mathbf{h}(\mathbf{y}) = \{h_{\text{edge}}(\mathbf{y}), h_{2\text{-star}}(\mathbf{y}), h_{\text{triangle}}(\mathbf{y})\}$, $\boldsymbol{\theta}$ are the corresponding parameters, i.e. $\boldsymbol{\theta} = \{\theta_{\text{edge}}, \theta_{2\text{-star}}, \theta_{\text{triangle}}\}$, and $\psi(\boldsymbol{\theta})$ is the normalizing constant:

$$\psi(\boldsymbol{\theta}) = \sum_{\mathbf{y}^* \in \mathcal{Y}} \exp \{ \langle \boldsymbol{\theta}, \mathbf{h}(\mathbf{y}^*) \rangle \}.$$

We base parameter values for our generative Markov ERGM off of the well-known Florentine marriage network (Padgett, 1994)². Specifically, we fit the triad model to this network and use the resulting parameter estimates for our simulations: $\theta_{\text{edge}} = -1.55$, $\theta_{\text{two-star}} = -0.05$, and $\theta_{\text{triangle}} = 0.25$. This allows us to represent some of the structural and topological features that we often observe in real world networks. However, there are many well-known issues with ERG models, particularly the fact that all nontrivial specifications form non-projective families (Shalizi and Rinaldo, 2013). This indicates that parameter values do not have the same meaning across networks of different size. Yet we include the Markov ERGM here due to its widespread popularity in network analysis.

Note that we do not incorporate any of the latest modifications to ERGM statistics, such as the geometrically weighted terms developed by Snijders et al. (2006). These are largely solutions to model-fitting issues and not necessarily required for the generative models we specify here. Further, the relationship between these tuning parameter values and network size is unclear and potentially complicated.

Hierarchical Bernoulli Model. Both the Hierarchical Bernoulli Model and hierarchical Markov ERGM are special cases of the more general HERGM, or Hierarchical ERGM (Schweinberger and Handcock, 2015). Generally, this model utilizes a Bayesian approach to identify plausible partitions of the network into smaller sub-networks (or “neighborhoods”) and enforces strong dependence (e.g. a Markov ERGM) within neighborhoods with weak dependence (i.e. the Bernoulli model) between neighborhoods. In this way, this model avoids the ERGMs’ issue with non-projectability. Thus, we consider these generative models our best attempt to mimic real world networks that vary in size.

For both of our HERGM simulated datasets, we consider the number of neighborhoods, K , to be fixed (although neighborhood membership, $\mathbf{Z}$, is still assumed unknown) and we set $K = \frac{n}{5}$ so that neighborhood size remains relatively small while the number of neighborhoods varies smoothly with network size. Let $\mathbf{Y}_{k,k}$ denote the subgraph of $\mathbf{Y}$ corresponding to the k th neighborhood and let $\mathcal{A}_k$ represent the set of nodes belonging to the k th neighborhood. Then, the probability distribution for $\mathbf{Y}$ can be written as follows:

$$P(\mathbf{Y} = \mathbf{y}|\mathbf{Z}, \boldsymbol{\theta}) = \prod_{k=1}^K P(\mathbf{Y}_{k,k} = \mathbf{y}_{k,k}|\mathbf{Z}, \boldsymbol{\theta}) \times \prod_{l=1}^{k-1} \prod_{i \in \mathcal{A}_k, j \in \mathcal{A}_l} P(Y_{ij} = y_{ij}|\mathbf{Z}, \boldsymbol{\theta})$$

²This network was chosen arbitrarily, except for the requirement that the triad Markov ERGM could be successfully fit to it. Our primary goal is to create simulated networks that in some way resemble real world networks, but not to recreate the dynamics of this particular network per se.

where the node membership vectors are distributed as $\mathbf{Z}_i | \boldsymbol{\pi} \stackrel{iid}{\sim} \text{Multinomial}(1; \boldsymbol{\pi})$ for $i = 1, \dots, n$.

Notice that the probability distribution for $\mathbf{Y}$ is conditional on the model parameters as well as the membership vectors: the set of $\mathbf{Z}_i$'s. From (3), we can see that this model induces local dependence, since the probability mass function factorizes into within-neighborhood and between-neighborhood pieces. Further, as suggested by Schweinberger and Handcock, we make the following simplifying assumptions: First, we assume that the between-neighborhood probability mass function depends only on the presence of edges:

$$P(Y_{ij} = y_{ij} | \mathbf{Z}, \boldsymbol{\theta}) = \exp \{ \theta_{btw} y_{ij} - \psi(\theta_{btw}) \}$$

and secondly, that the within-neighborhood probability mass function resembles a generic ERGM:

$$P(\mathbf{Y}_{k,k} = \mathbf{y}_{k,k} | \mathbf{Z}, \boldsymbol{\theta}) = \exp \{ \langle \boldsymbol{\theta}_W, \mathbf{h}_W(\mathbf{y}_{k,k}) \rangle - \psi(\boldsymbol{\theta}_W) \} \quad (4)$$

where $\psi(\theta_{btw})$ and $\psi(\boldsymbol{\theta}_W)$ are normalizing constants, calculated similarly as for Equations (1) and (2).

As suggested by Schweinberger and Handcock, we specify a multivariate normal prior for $\boldsymbol{\theta}$ with parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, and chose a Dirichlet prior for $\boldsymbol{\pi}$, with parameter $\alpha \mathbf{1}_K$, where $\mathbf{1}_K$ is a $1 \times K$ vector of ones. For our simulations, we set $\alpha = 10$ and consider θ_{btw} fixed. For all models we consider elements of $\boldsymbol{\theta}$ uncorrelated and without loss of generality set the diagonal elements of $\boldsymbol{\Sigma}$ equal to one.

For the Hierarchical Bernoulli Model, the within-neighborhood ERGM depends only on tie density (i.e. $\mathbf{h}_W(\mathbf{y}_{k,k})$ in Equation (4) consists of an edge count for the k th neighborhood subgraph of $\mathbf{y}$). So, for this model, note that $\boldsymbol{\theta}_W$ is one-dimensional and we use the value corresponding to a within-neighborhood density of 0.20. We set θ_{btw} so that the probability of connections between neighborhoods is 0.10.

Hierarchical Markov ERGM. This is another special case of the HERGM, where $\mathbf{h}_W(\mathbf{y}_{k,k})$ is composed of subgraph counts for edges, 2-stars, and triangles, very similar to the setup for the Markov ERGM (see Equation (2)). We use the same parameter values as in our traditional Markov ERGM - those estimated from the Florentine marriage network. That is, we set $\boldsymbol{\theta}_W = \{\theta_{edge}, \theta_{2-star}, \theta_{triangle}\} = \{-1.55, -0.05, 0.25\}$ to govern the Markov ERGM within neighborhoods. We set θ_{btw} so that the probability of connections between neighborhoods is 0.25.

Examples of realizations from these generative network models are shown in Figure 3. We plot randomly selected networks of the smallest size examined here, $n = 20$, so that structural differences are more readily visible.

2.3. Direct Comparison

Recall that for each of these generative models, we simulate $N_n = 200$ replicates with $n = 20, 30, \dots, 100$ nodes. All networks are simulated using standard functions from the `statnet` (Handcock et al., 2015) and `hergm` (Schweinberger et al., 2015) packages for R. For each network in each of these simulated datasets, we compute each of the common network statistics described in Section 2.1. Then we consider the distribution of each of these common network statistics where the networks come from the same generative model, paying specific attention to how the distribution varies according to the size of the network. To examine this, consider the following three metrics.

First, perhaps the most intuitive method for examining differences in distributions is using histograms. In Figure 4, we have plotted the histograms for each common network statistic under each generative model. In each plot, the color of the histogram indicates the size of the network, with the smallest network size ($n = 20$) shown in white and the largest network size ($n = 100$) shown in dark blue. Ideally, we would like the histograms to stack up neatly on top of each other with little difference in overall shape.

Second, we attempt to quantify the difference in these distributions by considering 2-sample Kolmogorov-Smirnov statistics. Generically, let F_{M,η,n_i} denote the empirical distribution function (EDF) for a network statistic, $\eta(G)$, calculated on networks from Model M of size n_i , for $M = 1, \dots, 5$ denoting the generative models discussed in Section 2.2, $\eta(G)$ from the set of common network statistics described in Section 2.1, and $i = 1, \dots, 9$ corresponding to $n_i = 20, 30, \dots, 100$. Note that for all M, η , and n_i , the number of samples from F_{M,η,n_i} will be the same, precisely $N_n = 200$. We use the Kolmogorov-Smirnov statistic to compare empirical distributions of a statistic calculated on networks from the same model across two different network sizes:

$$\sup_{x \in \mathbb{R}} |F_{M,\eta,n_i}(x) - F_{M,\eta,n_j}(x)| \quad i \neq j.$$

Table 1: k-Sample Anderson-Darling Statistics for Unadjusted Statistics

		SIMULATED DATASET					
		Erdős-Rényi	Bernoulli	Mean Degree Bernoulli	Markov	Hierarchical Markov	Hierarchical Bernoulli
STATISTIC	Deg. Cent.	694.16	680.85	609.90	642.89	667.88	638.90
	Deg. Cent. (N)	209.65	194.63	391.44	329.30	266.07	226.86
	Betw. Cent.	633.16	580.17	694.14	639.21	591.44	600.95
	Betw. Cent. (N)	637.46	609.14	179.13	568.19	617.07	564.86
	Clos. Cent.	409.44	215.38	31.77	465.80	192.43	467.63
	Clos. Cent. (N)	229.20	452.93	304.60	146.13	439.05	130.31
	Transitivity	67.84	81.18	186.36	238.35	225.82	121.85
	Avg. Path Length	60.09	558.74	356.51	228.62	171.51	430.91
	Density	58.08	57.62	682.59	533.56	127.57	132.75

This method allows us to compare only two network sizes simultaneously, so we calculate this statistic pairwise for each possible pair of network sizes. The Kolmogorov-Smirnov statistic is naturally bounded between zero and one, and so we can easily compare across different models and network statistics. In Figure 5, we present heatmaps where the cells of each heatmap correspond to the Kolmogorov-Smirnov statistic computed as the difference between empirical distributions for networks of two different sizes. Dark blue cells correspond to distributions that are very similar, while violet cells correspond to distributions that are quite different. Note that these plots are symmetric, so it is sufficient to examine only the upper left or lower right triangle.

Third, we consider further reducing our summary information by computing k-Sample Anderson-Darling statistics (Scholz and Stephens, 1987). For a particular generative model and network statistic, we use this k-Sample statistic to summarize the difference across distributions of all network sizes simultaneously. The k-Sample Anderson-Darling statistic is formulated as follows:

$$\sum_{i=1}^9 N_n \int_{B_{M,\eta}} \frac{\{F_{M,\eta,n_i}(x) - H_{M,\eta}(x)\}^2}{H_{M,\eta}(x)\{1 - H_{M,\eta}(x)\}}$$

where $H_{M,\eta}(x) = \sum_{i=1}^9 F_{M,\eta,n_i}(x)$ is the empirical distribution function of the pooled sample across all network sizes and $B_{M,\eta} = \{x \in \mathbb{R} : H_{M,\eta}(x) < 1\}$. Note that the typical equations simplify here since, as mentioned previously, the number of samples from F_{M,η,n_i} is the same for all M, η , and n_i . The Anderson-Darling statistic is a quadratic empirical distribution function statistic (the Cramer-von Mises statistic is another special case) and measures the weighted distance between the empirical distributions under consideration. The k-Sample Anderson-Darling statistics for our network data are provided in Table 1.

We see that in most cases, statistics across networks of different sizes clearly have very different distributions and comparing them directly is generally not appropriate. Not surprisingly, in many of the heat maps and histogram plots, we notice that direct comparison is least appropriate when the difference between network sizes is large. Also notice that, given a fixed difference in network sizes, direct comparison appears to be less appropriate when the network sizes themselves are small. For example, comparing a network statistic across networks of sizes $n = 20$ and $n = 30$ appears to be more problematic than across networks of sizes $n = 90$ and $n = 100$. This is not surprising, since in the former case the difference in network size constitutes a much larger proportion of the absolute network sizes themselves.

However, there are cases where directly comparing network statistics is not terribly inappropriate. For example, the distributions of the topological measures appear to be rather comparable across different sizes. Yet, this statement does not appear to hold (nearly as well) if we are considering average path length in a group of Bernoulli graphs or density from a group of Markov graphs. Of course, in practice, the generative model is unknown. Thus, we would advise avoiding direct comparison if possible, even for these relatively stable topological measures.

In terms of the centralization measures, we notice that using the normalized versions appears to greatly lessen the issue, though it is not entirely solved. Notably, in the case of betweenness, normalization appears to be less effective (this is particularly evident in the heat maps in Figure 5 and the Anderson-Darling statistics in Table 1), except for networks from the mean degree preserving Bernoulli model. Further, in the case of closeness centrality, normalization appears to offer improvements *only* in combination with particular types of graphs. Particularly, if we examine the 2-Sample Kolmogorov-Smirnov statistics in Figure 5 for the two versions of the closeness centrality measure, we see that the normalized version is superior only in the cases where the simulated data are generated from an Erdős-Rényi model, Markov ERGM, or Hierarchical Bernoulli model. In the other simulated datasets, the unnormalized statistics are more comparable across network sizes. Perhaps the normalization for closeness centrality is sensitive to particular types of topological network structure. However, it is unclear from these preliminary results what the sensitivity might be and/or how to adjust for it.

Figure 4: Histograms of Unadjusted Statistics

The columns indicate different simulated datasets while the rows represent the different network statistics of interest.

In short, this simulation study has clearly demonstrated that across a range of popular and reasonable generative network models, network statistics are not directly comparable across networks of different size, even when the networks under consideration are generated from the same model. It follows then that statistics calculated on real world networks (say from the same data generating process) are not directly comparable either. Thus, we turn to a brief literature review of proposed adjustments to our common network statistics which should allow for comparison across networks of different size.

3. Comparing Network Statistics

3.1. Review of Existing Approaches

Anderson et al. (1999) suggest specifying a baseline network model and using the distribution of the desired statistic calculated on a set of networks simulated from this model as a reference distribution, which can be used to standardize the observed statistic. Based on the observation that many common network statistics depend on both the size and density of the underlying network, they suggest using a model which adjusts for these characteristics. However, when comparing networks that vary in both size and density, this means making comparisons across networks based on reference

Figure 5: Heat Maps of Kolmogorov-Smirnov Statistics for Unadjusted Statistics

The columns indicate different simulated datasets while the rows represent the different network statistics of interest. Each cell in a plot represents the 2-sample Kolmogorov-Smirnov statistic for the comparison of a network statistic's distribution across two different network sizes. The network sizes are indicated on both the x- and y-axes. Cells that are dark blue indicate that the distributions are very similar while those that are violet indicate that the distributions differ.

distributions that also vary across the networks. Instead, comparisons would be more straight forward and more readily interpretable if the comparisons across networks are made relative to a single reference distribution that is common across all the networks under consideration and represents some kind of absolute baseline model. In this sense, not only can one compare networks within a dataset but across datasets where the adjustments were perhaps performed by different researchers. Of course, choosing a realistic fully parameterized reference distribution for comparison a priori is certainly not a straight forward task. A simple choice might be a completely random graph, or the Erdős-Rényi model. We investigate this option as a special case of our more general method in Section 4. Intuitively, standardizing relative to the Erdős-Rényi model might be an improvement over direct comparison, but could remain problematic in some cases since Erdős-Rényi networks simply fail to exhibit the type of structural detail and topological features that we typically observe in real world networks. Of course, we could consider different versions of an a priori fully parameterized reference distribution, but as Van Wijk et al. (2010) point out, this requires complete trust in the validity of whichever baseline model is used.

3.2. A New Mixture Model Reference Distribution

We propose a new reference distribution which is formed by a mixture of random graph models, where we mix over the dependence and structural features actually manifested in the observed dataset. Specifically, we suggest mixing over the dependence structure exhibited in a subset of networks randomly selected from the entire collection of observed networks. In this way, the reference distribution is common across all networks being considered and will, by design, mimic the structure of the networks being compared as well as reflect the natural variability among the observed networks. Our proposed methodology hopes to draw on the historical success of mixture models as a class of flexible models which are particularly advantageous in situations where it is difficult or undesirable to fully specify a traditional model a priori, such as in Newcomb (1886)'s model for outliers or Pearson (1894)'s work with evolutionary populations. However, note that our proposed methodology will result in *relative* comparisons across networks, since the reference distribution is empirically-based rather than formed a priori.

This type of relative comparison is not unique to the methodology we propose here. Consider the familiar case in which a set of correlated variables are considered for inclusion as covariates in a regression analysis. Including the full set directly would lead to unstable coefficient estimates due to multicollinearity. Instead, a dimension-reduction technique like principal components analysis (PCA) can be used; PCA uses an orthogonal transformation to convert a set of possibly correlated variables into a set of linearly uncorrelated transformations of these variables, called principal components. And it does so in such a way that the first principal component accounts for the largest variance in the original set of variables, the second principal component accounts for the second largest amount of variance and so on. Although the value of the principal components themselves do not have the same physical meaning as the original set of variables which were first

considered, they still capture (most of) the variability and the underlying content of the original set. Similarly, the value of a network statistic adjusted via the mixture model adjustment will not have meaning in and of itself, but will be very meaningful relative to statistics calculated on other networks that are adjusted in the same way. Thus, similar to a PCA style analysis, results from the mixture model adjustment will pick out differences among the set of observed networks, rather than relative to some absolute scale.

The procedure for this method can be divided into two main components: Again, suppose we observe a collection of networks $\mathbf{Y} = \{\mathbf{Y}_1, \mathbf{Y}_2, \dots, \mathbf{Y}_N\}$ with corresponding sizes $\mathbf{n} = \{n_1, n_2, \dots, n_N\}$. Let $\eta(\cdot)$ be our network statistic of interest.

Component 1: Mixture Model Parameter Selection

1. Randomly select N_M networks from $\mathbf{Y}$
2. Fit a graph model separately to each of these networks. Let the obtained parameter estimates be denoted by

$$\hat{\boldsymbol{\theta}} = \{\hat{\boldsymbol{\theta}}_1, \hat{\boldsymbol{\theta}}_2, \dots, \hat{\boldsymbol{\theta}}_{N_M}\}$$

where $\hat{\boldsymbol{\theta}}_j$ is the vector of parameter estimates from fitting the model to the j th network selected in step one.

Component 2: Mixture Model Simulation and Adjustment

For $i = 1, \dots, N$,

1. For $j = 1, \dots, N_M$, simulate $\frac{N_S}{N_M}$ graphs of size n_i setting $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}_j$.
2. Call the entire collection of simulated networks

$$\tilde{\mathbf{Y}}^{(i)} = \{\tilde{\mathbf{Y}}_1^{(i)}, \tilde{\mathbf{Y}}_2^{(i)}, \dots, \tilde{\mathbf{Y}}_{N_S}^{(i)}\}.$$

3. Compute $\eta(\mathbf{Y}_k^{(i)})$ for $k = 1, \dots, N_S$ and find the sample mean, $m_\eta^{(i)}$, and sample standard deviation, $s_\eta^{(i)}$, of this distribution.
4. Finally compute the z-score:

$$z_i = z(\mathbf{Y}_i | \tilde{\mathbf{Y}}^{(i)}, N_S, N_M) = \frac{\eta(\mathbf{Y}_i) - m_\eta^{(i)}}{s_\eta^{(i)}}.$$

where $N_M \leq N$ and N_S is some large positive integer divisible by N_M . Note that the computational efficiency of this algorithm can be improved by letting i index only the unique elements in $\mathbf{n}$.

Here, we have chosen to use a very intuitive standardization measure, a z-score, though other measures have been suggested (Van Wijk et al., 2010). While distributions of network statistics for Erdős-Rényi graphs are known in some special cases, the distribution of a given statistic for some arbitrary network model is generally unknown. However, from past observations and Figure 4, a normal distribution might be a relatively well-fitting approximation in some cases, and so, a z-score is a plausible standardization measure here.

Notice that in this algorithm, we have intentionally left the form of the network model used for simulating the reference distribution in this procedure unspecified. There may be situations in which a researcher thinks one model is appropriate given the type of networks under consideration, or perhaps certain network statistics are more amenable to adjustments utilizing a particular class of network models. One particular special case is worth pointing out: If the Erdős-Rényi model is utilized, note that Component 1 of the above algorithm is no longer necessary since $\hat{\boldsymbol{\theta}}_j = \text{logit}(0.50) \forall j = 1, \dots, N_M$ and the reference distribution is not only common across all networks but also an absolute reference distribution (i.e. it is no longer a mixture distribution and does not depend on the data in any way).

Finally, note that this procedure requires fitting N_M separate network models and thus network models in which obtaining parameter estimates can be more easily automated will be preferred.

4. Simulation Study: n -dependence

4.1. Methods

To consider the efficacy of this method, we turn to a simulation study. Treating the raw statistics from Section 2.3 as our observed datasets, we apply this standardization method with $N_M = 30$ and $N_S = 1000$ to each network statistic and generative model pairing. Note that for each generative model considered, $N = 1800$ networks. In terms of the network model used for simulating the mixture reference distribution in this adjustment, we consider four options: the special case of a fully parameterized reference distribution utilizing the Erdős-Rényi model, the traditional Bernoulli model, the mean degree preserving Bernoulli model, and the Hierarchical Bernoulli model. For each resulting version of the Mixture Model Adjustment, the network models will be fit using standard functions in the *statnet* (Handcock et al., 2015) and *hergm* (Schweinberger et al., 2015) packages for R.

Recall that the Erdős-Rényi model is simply a special case of the Bernoulli model, where we fix the probability of a tie to be 0.50. Thus, we might expect to see improvements over results from an Erdős-Rényi adjustment, if we simply allow the density of the graphs in the reference distribution to vary, in a way that matches the empirical distribution of densities of our observed networks. This is precisely what we allow by utilizing the Bernoulli model in our Mixture Model Adjustment³

Similarly, using the mean degree preserving version of the Bernoulli model will mix over graphs that mimic the average degree of our observed networks. Since density and average degree are closely related, we might expect these two versions of the approach to perform similarly. However, recall that the main difference between these two models is their behavior as network size increases: the Bernoulli model maintains density, while adding Krivitsky et al. (2011)’s offset preserves mean degree. We might expect that whichever regime is present in the observed dataset, that it’s corresponding model might prove a better fit for the mixture model adjustment procedure.

Perhaps the next more complicated graph model would be some ERG-type model which includes not only a density effect but also some structural statistics, such as triangles or k-stars. However, due to Shalizi and Rinaldo’s result, we do not believe that ERGM parameters have the same meaning across different network sizes (2013). So, using an ERGM to simulate the mixture reference distribution would most likely be inappropriate. Further, ERGMs are notoriously difficult to fit, particularly when model terms are exclusively structural, and thus further unsuited for this methodology.

As mentioned earlier, the Hierarchical ERGMs developed by Schweinberger and Handcock do not face such issues. In fact, the authors show that HERGMs do form a type of projective family, although this type of consistency is a weaker condition than that considered by Shalizi and Rinaldo (Schweinberger and Handcock, 2015). However, HERGMs in their current implementation take much longer to fit than traditional ERGMs (hours rather than minutes, even for relatively small networks), though the resulting model fits are significantly more stable and reliable. Since our proposed method requires fitting N_M network models, we use only the Hierarchical Bernoulli version of the HERGM, rather than the more complex Hierarchical Markov model described in Section 2.2. In this implementation, when we fit the Hierarchical Bernoulli to the N_M randomly selected observed networks, we assume the number of neighborhoods is unknown and use a nonparametric stick-breaking prior for the neighborhood membership parameters, π . We set the maximum number of neighborhoods to be $K = \frac{n}{5}$. When we simulate the reference distribution formed from the parameter estimates obtained above, we follow the same model set up as for the simulated Hierarchical Bernoulli networks (described in Section 2.2), letting $\alpha = 10$ and setting K, θ, θ_{btw} , and the diagonal entries of Σ equal to the estimates from the model fitting step in the first component of the algorithm.

4.2. Results

As before, we use the three metrics from Section 2.3 to examine the results of these mixture model adjustments: histograms, heat maps of Kolmogorov-Smirnov statistics, and k-Sample Anderson-Darling statistics. Intuitively, if the method is effective, then, relative to the comparison metrics for the raw statistics, we expect the histograms to overlap each other, the heat maps to have more dark blue cells, and the Anderson-Darling statistics to be smaller. Note that the normalized versions of the centralization measures are excluded here since the normalization adjustment is linear for a fixed network size and thus z-scores for the normalized and unnormalized versions of these statistics are equivalent. These figures are available

³Note that the term “Mixture” in the name for our proposed methodology refers to the distribution of simulated networks in step one of Component 2 of our algorithm. These networks are simulated according to a mixture of different specifications of a particular network model where the components of this mixture are formed in step two of Component 1. For example, we might simulate networks from a Bernoulli model, mixing over the log odds of a tie between -2, 1.4, and 0.2. Note that our proposed methodology does not (currently) incorporate any of the so-called mixture models for random graphs (e.g. Daudin et al., 2008) in which a single network is modeled as a mixture of random graphs.

Figure 6: Performance of Mixture Model Adjustments

Each plot represents a different version of the mixture model approach, including no adjustment. The height of the bar corresponds to the inverse of the Anderson-Darling statistics, so that taller bars correspond to distributions that are more similar across varying network sizes. For simplicity, values above 0.10 are not shown here; the maximum is 0.13.

in the Supplementary Materials. As a summary of these results, we provide bar charts in Figure 6 of the reciprocal of the Anderson-Darling statistics across each adjustment method, so that taller bars correspond to better performance (i.e. adjusted distributions of network statistics that are more similar across network sizes).

Not surprisingly, the Erdős-Rényi adjustment works strikingly well when the data are actually generated from the Erdős-Rényi model. However, problems occur when the observed networks more strongly resemble real world network data, i.e. are simulated from the other generative models considered here. In particular, this adjustment performs very poorly for all of the topological measures computed on non-Erdős-Rényi simulated networks. However, we do see that using the Erdős-Rényi adjustment for the centralization measures appears to be an improvement over direct comparison. Thus, the Erdős-Rényi adjustment method described here is certainly an improvement over direct comparison, but is still sometimes problematic (especially in the case of topological measures). Note also that this adjustment method is very easily implemented, computationally inexpensive (for graphs similar in size to those examined here - the computational burden increases with network size), and is an example of a fully parameterized reference distribution.

Considering the first nontrivial version of the Mixture Model Adjustment, where the Bernoulli model is used, we see many improvements relative to adjustment using the Erdős-Rényi model. Recall that the Erdős-Rényi adjustment seemed most suited for centralization scores. However, we see even further improvements in the results based on this Bernoulli version of the Mixture Model Adjustment. The results from this mixture model adjustment are more comparable across *all* simulated datasets.

Recall also that the Erdős-Rényi adjustment performed quite poorly for the topological measures, so it is more appropriate to compare the results of the Mixture Model Adjustments to the raw topological network statistics themselves. Here we see significant improvements, particularly for the simulated datasets where the networks are generated from the Bernoulli model and both versions of the HERGMs implemented here.

However, the Bernoulli version of the Mixture Model Adjustment for the network statistics calculated on Markov graphs does not appear to offer any advantages over direct comparison of the statistics themselves. This is not too alarming, given the previously mentioned result which shows that ERG-type models like this one do not form projective families (Shalizi and Rinaldo, 2013). This result implies that parameters in such models are not comparable across different network sizes

and so, the dataset of simulated Markov graphs that we have used here might not be particularly meaningful, i.e. it might not exhibit the same type of structure across different network sizes. If this is the case, we would not expect any type of adjustment method to be able to “fix” such an issue.

The simulation results examined here demonstrate that our proposed mixture model adjustment, where we use the Bernoulli model to form the simulated mixture reference distribution, offers a computationally inexpensive way to make reliable comparisons of both centralization and topological network statistics across networks of different size (see Table 5 in the Appendix for a comparison of computation times across methods).

As expected, the version of the Mixture Model Adjustment which uses the mean degree preserving Bernoulli model performs well for datasets where network degree is preserved as network size grows, i.e. on networks generated from the mean degree preserving Bernoulli model. However, if this is not the case, the original Bernoulli version of the Mixture Model Adjustment appears to offer more comparable adjusted network statistics. Further, note that using the original Bernoulli version of the Mixture Model Adjustment for networks generated from the mean degree preserving version (and vice versa) appears to be a poor choice. Thus, in choosing a parametric model to implement in the Mixture Model Adjustment, it is very important to examine (both empirically and theoretically) whether density or average degree is preserved as network size grows.

Surprisingly, when we increase the complexity and flexibility of the reference model in the Mixture Model Adjustment algorithm to the Hierarchical Bernoulli model, we do not see much improvement and, in fact, the adjusted statistics look slightly less comparable across network sizes in many cases. Although the centrality measures look relatively comparable across network sizes, the topological network statistics’ distributions still show evidence of some dependence on n , particularly when calculated on non-Bernoulli graphs.

There are two plausible explanations for this behavior. First, perhaps this model is simply too complex to fit and simulate from in an automated way, as is required by the Mixture Model Adjustment. That is, perhaps the adjusted statistics would benefit from more model fine-tuning, such as trying different values of K_{max} or respecifying prior parameter values. Secondly, while the Hierarchical Bernoulli model is logically appealing and perhaps more suited to real-world network data than the simple Bernoulli model, it may be a poor model for the *simulated* datasets examined here. However, its mediocre performance for network statistics calculated on networks simulated from the Hierarchical Bernoulli distribution itself is curious.

5. Simulation Study: Feature Detection

We have demonstrated that the mixture model adjustment can be used to produce adjusted network statistics that are amenable to comparison across networks of different size. In this section, we design and implement a small simulation study to confirm that the adjusted network statistics can still be used in much the same way as the original, raw network statistics themselves: to identify relevant structural features of observed networks.

We simulate networks from two slightly different models, use the Mixture Model Adjustment to produce adjusted versions of our common network statistics, and examine these statistics to check that our adjusted statistics are able to detect the difference between the networks generated from different models, at least to the extent that the unadjusted statistics are capable of detecting the difference. We use our most realistic generative model, the Hierarchical Markov ERGM, under the same settings described in Section 2.2 but change θ_{btw} to create a second version of this model. More specifically, we consider two versions of the Hierarchical Markov ERGM here: one where the probability of connections between neighborhoods is 0.10 and one where this probability is 0.25. We simulate $N_n = 250$ replicates from these two Hierarchical Markov ERGMs with network sizes $n = 20, 30, \dots, 100$, for a total of $N = 2250$ simulated networks per generative model. Given its superior performance in the simulation studies in the previous section, we use the Bernoulli model to produce the mixture reference distribution and obtain adjusted network statistic values.

We provide boxplots of the results in Figure 7. We are primarily concerned with whether perceivable differences between the two generative models (with $\theta_{btw} = 0.10$ and $\theta_{btw} = 0.25$) which are evident in the boxplots of the raw, unadjusted statistics are either diminished, preserved, or amplified by the mixture model adjustment. While the empirical distribution of degree centrality appears to remain roughly the same under either model, we see differences in the empirical distributions of the other unadjusted/adjusted network statistics. Not surprisingly, in most cases, this difference is harder to distinguish among small networks but increasingly easier to distinguish as the network grows in size. Note that there is a difference between the empirical distributions of unadjusted transitivity and density statistics across the two generative models, and this distinction is preserved in the adjusted versions of these statistics. For the remaining network statistics - betweenness centrality, closeness centrality, and average path length - the difference between the empirical distributions of the unadjusted

Figure 7: Boxplots of Network Statistics from Two Hierarchical Markov ERGMs

The rows indicate the common network statistics, while the left column gives the unadjusted statistic and the right column gives those adjusted by a Bernoulli mixture model. In each plot, the x-axis gives the network sizes. The color indicates the two versions of the model, where the probability of connections between neighborhoods is given in the legend on the right.

statistics across the two generative models is actually amplified and made more noticeable by the adjustment.

Although this simulation study is rather limited in scope, we expect to see similar results under different settings. Once the network statistics' dependence on network size is removed, it is not surprising that we might be able to more easily pick up "true" differences between networks under comparison. Not only does the mixture model adjustment allow for comparison of common network statistics across networks of different size, but it also improves the ease of the comparative analysis itself, since differences among adjusted statistics are often more pronounced.

6. Application: L.A.FANS Ecological Networks

Finally, we apply the mixture model adjustment to neighborhood network data from the Los Angeles Family and Neighborhood Survey (L.A.FANS). Conducted in the early 2000s, L.A.FANS is a two-wave survey of roughly 3,000 families in 65 census tracts in Los Angeles County, California (Sastry et al., 2006). Here, we use activity data collected during the first wave (2000 and 2001). Survey participants were asked to report the locations where they conduct a few specific routine activities, such as where they go grocery shopping, where they go to the doctor and where they work.

Following earlier analysis with this dataset, we assign the reported activity locations to census block groups (Browning et al., 2014). Naturally, this data can be represented as a two-mode ecological network within each census tract, with the first mode being people who reside in a census tract and the second mode being the census block groups in which activity locations that these people reported visiting are physically located. We focus only on the one-mode projection of these networks, where people in a tract are connected if they report visiting activity locations that occur in the same census block group.

Note that we observe $N = 65$ networks, where the nodes in each network represent the survey participants from a particular census tract. The network sizes range from $n = 27$ to $n = 54$ people. Visually, we notice two structural patterns in these data. We give examples of these two patterns in Figure 8⁴. For the most part, these 65 observed networks look like either of the two networks in this figure. Either the network consists of a single group of very highly connected nodes and a few nodes with only a few connections (i.e. it exhibits some type of core-periphery structure) or is generally less dense, with

⁴The identity of the sampled census tracts in L.A.FANS is withheld to protect respondents' identity.

Figure 8: L.A.FANS Network Structure

The above networks are one-mode projections of activity pattern data from L.A.FANS. For both of these networks, $n = 34$

multiple groups of more highly connected nodes.

We analyze properties of these networks by mirroring the structure of this paper thus far: first, we examine the raw network statistics themselves; secondly, we standardize relative to a fully parameterized reference distribution, Erdős-Rényi graphs; finally, we apply the Mixture Model Adjustment where we use the Bernoulli model, the mean degree preserving Bernoulli model and the Hierarchical Bernoulli model to produce the mixture reference distributions. We use all of the same algorithm settings as described before. As was observed in the simulation study, we expect the performance of the Hierarchical Bernoulli mixture to be poor relative to its computational requirements. Further, recall that the simulation study results indicated that using the mean degree preserving version of the Bernoulli model would be superior to the original version if mean degree (rather than density) is preserved as network size increases. We consider a simple plot of density and mean degree against network size in Figure 9. Since the density of our observed networks appears to be preserved as network size changes, we expect the mixture model adjustment using Bernoulli model components to be superior to the version which includes Krivitsky et al. (2011)’s offset term.

Figure 9: Density and Mean Degree in L.A.FANS Networks

Density (left) and mean degree (right) are plotted against network size for the $N = 65$ observed L.A.FANS census tracts. The dotted lines represent fitted simple linear regressions.

Unlike in the simulation studies, here, we no longer have much replication across network size, and so we turn to new metrics for evaluating how comparable the network statistics are across different network sizes. In Figure 9 in the Supplementary Materials we provide simple plots of the network statistics by network size, n . If the network statistics are comparable across different network sizes, we would not expect to see dependence between network statistics and n . Or rather, we would expect to see the correlation between the raw statistics and n be reduced by our adjustment methods. We provide a table of such correlations in Table 2.

Although the plotting scale obviously differs across n , the adjustment methods do not seem to have a strong effect on the overall shape or pattern that we observe in any of these plots, with the exception of degree centrality. For the unadjusted degree centrality measures, we see a strong relationship with network size. However, applying any kind of adjustment method seems to ameliorate the issue. For all other network statistics, we do not observe such an obvious improvement.

However, when we examine the correlations in Table 2, we see that applying some type of an adjustment method does

Table 2: L.A.FANS Data: Correlation between Network Statistics and Network Size

		ADJUSTMENT MODEL				
		Unadjusted	Erdős-Rényi	Bernoulli	Mean Degree Bernoulli	Hierarchical Bernoulli
STATISTIC	Deg. Cent.	0.8228	0.3005	0.3026	0.3509	0.1877
	Betw. Cent.	0.6598	0.4670	0.4290	0.1351	0.3851
	Clos. Cent.	0.2228	0.0745	0.1060	0.1299	0.0417
	Transitivity	-0.0981	0.2543	-0.0926	0.2286	-0.1000
	Avg. Path	0.0081	0.1224	0.0651	-0.1659	0.0236
	Density	-0.0473	-0.1298	-0.0473	0.1878	-0.0522

reduce the correlation between network size and many of the network statistics. This holds for all of the centralization measures, but is violated for average path length as well as Erdős-Rényi adjusted transitivity and density. Recall that in our simulation study, the Erdős-Rényi adjustment performed rather poorly for topological measures, so it is not surprising that it appears to perform poorly for these same measures in this application as well. Further, although the correlation between average path length and network size is not reduced by the adjustment methods, it is very near zero to begin with and remains relatively near zero under both the Bernoulli and Hierarchical Bernoulli versions of the Mixture Model Adjustment.

Given these results and the conclusions from our simulation study, we recommend using network statistics adjusted by the Bernoulli version of the Mixture Model Adjustment in any further analysis involving these network statistics. Note that all versions of the Mixture Model Adjustment performed comparably, but that the Bernoulli version is much less computationally expensive (see Table 5 in the Supplementary Materials) than the Hierarchical Bernoulli version, and the assumptions of the Bernoulli model (constant density across network size) appear to be a good match for our observed data.

Moving forward with our analysis, we now use the adjusted network statistics to compare across networks of different size. Note that although the adjusted network statistics are z-scores, they should not be interpreted as conveying any type of statistical significance. Rather, they are measures of *relative* comparison (much like PCA component scores), that pick up differences among the networks being compared. Finally, we will briefly examine ways in which using these adjusted values can impact the types of conclusions we might draw about these data.

For example, if we examine the raw degree centralization values across our group of observed networks, we notice that the network for census tract #64 has a larger value than that for census tract #37 (see Figure 10). However, it turns out that this observation is entirely driven by network size. Once we perform the Bernoulli version of the Mixture Model Adjustment, we see the opposite: the network for census tract #37 has a much larger degree centrality value than we would expect for networks of that size while census tract #64's is closer to what we might expect for its size. And recall, that in this type of statement, we are comparing our observed networks to simulated networks of the same size which mimic the type of structure that we observe in our dataset (here, the only structural effect we are mimicking is density, since we are using the Bernoulli version of the Mixture Model Adjustment). Thus, we might conclude that, relative to the type of network data we observed, census tract #37 has higher degree centrality than census tract #64, net of network size effects. These types of rank changes are common throughout the various network statistics examined here.

In summary, using unadjusted or poorly adjusted (i.e. via the Erdős-Rényi adjustment) network statistics in a comparative analysis of network data can lead to very different conclusions than if one were to utilize our proposed adjusted statistics. In fact, we argue that these different conclusions could in fact be dangerously misleading because the statistics are not properly adjusted for network size. Thus, we encourage using a Bernoulli version (perhaps including Krivitsky et al. (2011)'s offset term) of the Mixture Model Adjustment to first adjust observed network statistics for size effects and then proceed with any desired analysis of these common network statistics per usual (e.g. as independent variables in a regression analysis). The decision to include Krivitsky et al. (2011)'s offset term could be guided by theoretical considerations of the expected behavior of density and mean degree for the particular type of network data under consideration, as network size increases. This decision should also be influenced by any patterns in the observed data, as was considered here. That is, simply plotting density and mean degree against network size for the observed networks should help inform whether or not an offset term is helpful.

7. Discussion

We have shown that network statistics are not amenable to direct comparison across networks of different size, even when the networks under consideration are generated from the same model. Although we can not make this statement absolutely, we have demonstrated this phenomenon across a range of popular and reasonable generative models for network data as well as across a variety of popular network statistics.

Figure 10: Degree Centrality in the L.A.FANS Neighborhood Data

The horizontal axis denotes network size, n , while the vertical axis gives the statistic value.

We previewed the advantages of the Mixture Model Adjustment and have shown that in many cases the adjusted statistics are more comparable across network sizes and are still capable of describing interesting features of networks. This provides researchers with a reliable way of making comparisons across networks of different sizes, without the difficulty of specifying a fully parameterized reference distribution a priori. The resulting adjusted network statistics can be interpreted as measures of relative comparison among the observed networks, similar to PCA component scores.

Due to its superior performance and low computational cost, we have recommended use of the Bernoulli model (perhaps including Krivitsky et al. (2011)'s offset term) for the components of the mixture reference distribution and provided a simple approach for deciding whether or not to include the offset term. Perhaps implementing network models other than those examined here as mixture components in this adjustment will offer even greater improvements, or perhaps certain network statistics are more amenable to adjustments utilizing a particular class of network models. We leave this topic for future investigation.

The strong performance of the Bernoulli models is not surprising. Anderson et al. (1999); Faust (2006) and Van Wijk et al. (2010) all document a relationship between network density (or, closely related, average degree, as examined by Anderson et al.; Van Wijk et al.) and a variety of graph-level network statistics. Thus, it is not surprising that a method which properly adjusts for network density would render the network statistics more comparable across network sizes.

Of course, the range of network statistics and generative network models examined here could be expanded as well. We have considered only statistical generative network models rather than mathematical models, such as the conditional uniform graphs examined by Anderson et al. (1999). Our intent was to create a mixture reference distribution which represents the type of structure that we *expect* to see in, say, another realization of our set of observed networks, rather than viewing the observed network structure as some *fixed* attribute and adjusting for it, i.e. rather than conditioning on *observed* network attributes as is done in CUG models.

Further examination of the performance of the Mixture Model Adjustment itself is also needed. In particular, a consideration of the effect of N_M , the number of mixture components, could prove to be a worthwhile investigation. Other measures of distributional comparison could be used, rather than the simple z-score which performs poorly for nonnormal distributions. Further, perhaps mixture weights could be incorporated within the Mixture Model Adjustment, proportional to some measure of model fit, to better adjust for observed variability in the network dataset.

References

- An, W. "Multilevel meta network analysis with application to studying network dynamics of network interventions." *Social Networks*, 43:48 – 56 (2015).
- Anderson, B. S., Butts, C., and Carley, K. "The interaction of size and density with graph-level indices." *Social Networks*, 21(3):239 – 267 (1999).

- Bassett, D. S., Bullmore, E., Verchinski, B. A., Mattay, V. S., Weinberger, D. R., and Meyer-Lindenberg, A. "Hierarchical organization of human cortical networks in health and schizophrenia." *The Journal of Neuroscience*, 28(37):9239–9248 (2008).
- Browning, C. R., Calder, C. A., Krivo, L. J., Mohr, A. L., and Boettner, B. "Socioeconomic segregation of activity spaces in urban neighborhoods: Does shared residence mean shared routines?" (2014). Manuscript submitted for publication.
- Butts, C. T. "Exact bounds for degree centralization." *Social Networks*, 28(4):283–296 (2006).
- Chatterjee, S., Diaconis, P., et al. "Estimating and understanding exponential random graph models." *The Annals of Statistics*, 41(5):2428–2461 (2013).
- Daudin, J.-J., Picard, F., and Robin, S. "A mixture model for random graphs." *Statistics and Computing*, 18(2):173–183 (2008).
- Erdős, P. and Rényi, A. "On the evolution of random graphs." *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17 – 61 (1960).
- Faust, K. "Comparing social networks: size, density, and local structure." *Metodološki zvezki*, 3(2):185–216 (2006).
- Frank, O. and Strauss, D. "Markov graphs." *Journal of the American Statistical Association*, 81(395):832–842 (1986).
- Freeman, L. C. "Centrality in social networks conceptual clarification." *Social Networks*, 1(3):215 – 239 (1979).
- Gilbert, E. N. "Random graphs." *The Annals of Mathematical Statistics*, 1141–1144 (1959).
- Handcock, M. S., Hunter, D. R., Butts, C. T., Goodreau, S. M., Krivitsky, P. N., Bender-deMoll, S., and Morris, M. *statnet: Software Tools for the Statistical Analysis of Network Data*. The Statnet Project, r package version 2015.11.0 edition (2015).
- He, Y., Chen, Z., and Evans, A. "Structural insights into aberrant topological patterns of large-scale cortical networks in alzheimer's disease." *The Journal of Neuroscience*, 28(18):4756–4766 (2008).
- Kolaczyk, E. D. and Csárdi, G. *Statistical analysis of network data with R*, volume 65. Springer (2014).
- Krivitsky, P. N., Handcock, M. S., and Morris, M. "Adjusting for network size and composition effects in exponential-family random graph models." *Statistical methodology*, 8(4):319–339 (2011).
- Lubbers, M. J. "Group composition and network structure in school classes: A multilevel application of the p^* model." *Social Networks*, 25(4):309 – 332 (2003).
- Lubbers, M. J. and Snijders, T. A. "A comparison of various approaches to the exponential random graph model: A reanalysis of 102 student networks in school classes." *Social Networks*, 29(4):489 – 507 (2007).
- Newcomb, S. "A generalized theory of the combination of observations so as to obtain the best result." *American Journal of Mathematics*, 8:343–66 (1886).
- Padgett, J. F. "Marriage and elite structure in Renaissance Florence, 1282-1500." *Social Science History Association* (1994).
- Pearson, K. "Contribution to the mathematical theory of evolution." *Philosophical Transactions of the Royal Society A*, 185:71–110 (1894).
- Sastry, N., Ghosh-Dastidar, B., Adams, J., and Pebley, A. R. "The design of a multilevel survey of children, families, and communities: The Los Angeles Family and Neighborhood Survey." *Social Science Research*, 35(4):1000 – 1024 (2006).
- Schindler, K. A., Bialonski, S., Horstmann, M.-T., Elger, C. E., and Lehnertz, K. "Evolving functional network properties and synchronizability during human epileptic seizures." *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 18(3):033119 (2008).
- Scholz, F. W. and Stephens, M. A. "K-Sample Anderson–Darling tests." *Journal of the American Statistical Association*, 82(399):918–924 (1987).
- Schweinberger, M., Handcock, M., and Luna, P. *hergm: Hierarchical Exponential-Family Random Graph Models with Local Dependence* (2015).
- Schweinberger, M. and Handcock, M. S. "Local dependence in random graph models: Characterization, properties and statistical inference." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(3):647–676 (2015).
- Shalizi, C. R. and Rinaldo, A. "Consistency under sampling of exponential random graph models." *The Annals of Statistics*, 41(2):508–535 (2013).

- Smit, D. J., Stam, C. J., Posthuma, D., Boomsma, D. I., and De Geus, E. J. "Heritability of "small-world" networks in the brain: a graph theoretical analysis of resting-state EEG functional connectivity." *Human brain mapping*, 29(12):1368–1378 (2008).
- Smith, J. A. and Moody, J. "Structural effects of network sampling coverage I: Nodes missing at random." *Social Networks*, 35(4):652 – 668 (2013).
- Snijders, T. A., Pattison, P. E., Robins, G. L., and Handcock, M. S. "New specifications for exponential random graph models." *Sociological methodology*, 36(1):99–153 (2006).
- Snijders, T. A. B. and Baerveldt, C. "A multilevel network study of the effects of delinquent behavior on friendship evolution." *The Journal of Mathematical Sociology*, 27(2-3):123–151 (2003).
- Stam, C. "Functional connectivity patterns of human magnetoencephalographic recordings: a 'small-world' network?" *Neuroscience Letters*, 355(1–2):25 – 28 (2004).
- Van Wijk, B. C., Stam, C. J., and Daffertshofer, A. "Comparing brain networks of different size and connectivity density using graph theory." *PloS One*, 5(10):e13701 (2010).