Random Networks, Graphical Models, and Exchangeability

Steffen Lauritzen*
University of Copenhagen

Alessandro Rinaldo
Carnegie Mellon University

Kayvan Sadeghi
University of Cambridge

November 22, 2017

Abstract

We study conditional independence relationships for random networks and their interplay with exchangeability. We show that, for finitely exchangeable network models, the empirical subgraph densities are maximum likelihood estimates of their theoretical counterparts. We then characterize all possible Markov structures for finitely exchangeable random graphs, thereby identifying a new class of Markov network models corresponding to bidirected Kneser graphs. In particular, we demonstrate that the fundamental property of dissociatedness corresponds to a Markov property for exchangeable networks described by bidirected line graphs. Finally we study those exchangeable models that are also summarized in the sense that the probability of a network only depends on the degree distribution, and identify a class of models that is dual to the Markov graphs of Frank and Strauss (1986). Particular emphasis is placed on studying consistency properties of network models under the process of forming subnetworks and we show that the only consistent systems of Markov properties correspond to the empty graph, the bidirected line graph of the complete graph, and the complete graph.

Keywords: bidirected Markov property, de Finetti's theorem, exchangeable arrays, exponential random graph model, graph limit, graphon, Kneser graph, marginal beta model, Möbius parametrization, Petersen graph.

*Steffen Lauritzen, Department of Mathematical Sciences, University of Copenhagen, Universitetsparken 5, 2100, Copenhagen, Denmark. email: lauritzen@math.ku.dk

1 Introduction

Over the last decades, the collection and rapid diffusion of network data originating from a wide spectrum of scientific areas have created the need for new statistical theories and methodologies for modeling and analyzing large random graphs. There now exists a very large and rapidly growing body of literature on network analysis: see, for example, Kolaczyk (2009), Newman (2010), and references therein. Since the seminal contributions of Frank and Strauss (1986) and Holland and Leinhardt (1981), the field of statistical network modeling has advanced significantly and researchers have now gained a much broader understanding of the properties of network data and the many open challenges associated to network modeling.

A common feature shared by many network models is that of invariance to the relabeling of the network units, or (finite) exchangeability, whereby isomorphic graphs have the same probabilities, and are therefore regarded as statistically equivalent. Exchangeability is a basic form of probabilistic invariance, but also a natural and convenient simplifying assumption to impose when formalizing statistical models for random graphs. Examples of popular network models which rely on exchangeability include many exponential random graph models, the stochastic block model, graphon-based models, latent space models, to name a few.

In the probability literature, exchangeability of infinite binary arrays — which encompasses networks — is very well understood: see for example Aldous (1981, 1985); Lauritzen (2003, 2008); and Kallenberg (2005). Of particular relevance here is the work of Diaconis and Janson (2008) (but see also Orbantz and Roy, 2015), which details the connections between exchangeability of random graphs and the notion of graph limits developed in Lovász and Szegedy (2006). While the existing results provide a canonical, *analytic* representation of infinite exchangeable networks (see later in Section 2), they appear to be too implicit for statistical modeling purposes. Indeed, the properties implied by such representation do not directly yield simple parametric statistical models for finite networks. Furthermore, while an exchangeable probability distribution for an infinite networks by construction induce consistent families of exchangeable distributions for all its finite subnetworks, there is no guarantee that an exchangeable distribution for networks of a given size may be extendable to a probability distribution for larger (possibly infinite) networks, or that the induced distributions share the similar properties. Thus, *a priori* any statistical model for finite exchangeable networks that does not conform to the analytic representation of infinite exchangeable networks need not be compatible in any meaningful way with models of the same type over networks of different sizes. These issues are of great significance, as they render very difficult to determine the extent to which statistical inference based on a subnetwork applies to a larger network, as argued in Crane and Dempsey (2015), and Shalizi and Rinaldo (2013).

The main goal of our article is to reveal and point out similarities and connections between seemingly unrelated concepts from different parts of the literature — in particular graphical models, network models, and notions of exchangeability — and our analysis is attempting to enhance understanding of the models and their properties.

We rely on the theory of exponential families and, especially, of graphical models (Lauritzen, 1996) in order to categorize Markov properties implied by the assumption of exchangeability, and the statistical models that can be derived from it. The idea of modeling networks as Markov random fields on binary variables is not new: it was originally suggested by Frank and Strauss (1986), who argue that a particular class of exponential random graph models (to be discussed below in Section 4) is able to capture what the authors, somewhat arbitrarily, posit as a natural form of dependence structure for networks, as well as the symmetries implied by exchangeability. In contrast, our analysis of the Markov properties of exchangeable network models is both principled, as it only relies on the assumption of exchangeability, and exhaustive, as it produces a complete list of all Markov properties expressed by exchangeable network models. We make the following specific contributions:

- We describe an exponential family representation for exchangeable network models where the sufficient statistics are given by all subgraph counts. The subgraph frequencies are maximum likelihood estimators of the corresponding mean value parameters, which are the marginal probabilities of all subgraphs.

- We study the Markov properties of finitely exchangeable network models, and show that such models are compatible only with four types of non-trivial conditional independence structures.

- We demonstrate that statistical models for finite networks induced by extremal exchangeable distributions on an infinite network are Markov with respect to a distinguished bidirected graphical model for marginal independence.

- Finally, we propose a novel class of models for exchangeable networks that are also *summarized,* in the sense that their properties only depend on the degree distribution of the observed network.

We should emphasize that the scope of this paper has been strictly limited to the study of full exchangeablity and Markov structures in the sense used for standard graphical models and, as we show, the assumption of full exchangeabiliy, extendibility, and Markov properties in this sense is very restrictive.

There are important types of Markov properties for random networks that we have not considered in this paper, notably that of *partial conditional independence* (Snijders et al., 2006), also used in Hunter et al. (2008). This concept allows the conditional independence properties of specific ties in a network to depend on the configuration of the remaining network. This type of Markov property is common in the analysis of spatial point processes, where it is referred to as *data dependent Markov neighbourhoods* (Baddeley and Møller, 1989); indeed, a random network can be considered as a point process on the space of possible ties. Also in the literature on Bayesian networks and graphical models, an analogous Markov theory has evolved and here the term *context-specific independence* is often used, see e.g. Boutilier et al. (1996) or Nyman et al. (2014) for a recent discussion.

Similarly, it is of interest but outside the scope of the present paper to discuss other variants of exchangeability, where we in particular mention *partial exchangeability* or *block exchangeability* as used in Schweinberger and Handcock (2015) or models derived from exchangeable measures, as in Caron and Fox (2017).

The structure of our paper is as follows. In Section 2 we develop the necessary terminology and detail the concepts of exchangeable arrays, random networks, bidirected Markov properties, and important parametrizations of the models. In Section 3 we provide parametrizations for exchangeable networks and their maximum likelihood estimation. In Section 4 we investigate the consequences of exchangeability in its interplay with associated Markov properties. In Section 5 we study the consequences of adding the restrictions of summarizedness, and in Section 6 we study consistency properties under subsampling of various systems of network models.

2 Preliminaries

In this section we gather some background material on graph-theoretic concepts, parametrizations of distributions of binary arrays and graphical modeling.

2.1 Graph-theoretic concepts

Following West (2001), a *labeled graph* is determined by an ordered pair $G = (V, E)$ consisting of a *vertex* set V, an *edge* set E, and a relation that with each edge associates two vertices, called its *endpoints*. When vertices u and v are the endpoints of an edge, these are *adjacent* and we write $u \sim v$; we denote the corresponding edge as uv.

In this paper, we will restrict to *simple* graphs, i.e. graphs without loops (the endpoints of each edge are distinct) or multiple edges (each pair of vertices are the endpoints of at most one edge); simple graphs are determined by the pair $G = (V, E)$ alone. We will distinguish three *types* of edge, denoted by *arrows*, *arcs* (solid lines with two-headed arrows) and *lines* (solid lines). Arrows can be represented by ordered pairs of vertices, while arcs and lines by 2-subsets of the vertex set. The graphs in the present paper only contain one of these types, and are called, respectively, *directed*, *bidirected*, or *undirected*.

The labeled graphs $F = (V_F, E_F)$ and $G = (V_G, E_G)$ are considered equal if and only if $(V_F, E_F) = (V_G, E_G)$. We omit the term labeled when the context prevents ambiguity.

A *subgraph* of a graph $G = (V_G, E_G)$ is a graph $F = (V_F, E_F)$ such that $V_F \subseteq V_G$ and $E_F \subseteq E_G$ and the assignment of endpoints to edges in F is the same as in G. We will write $F \subset G$ to signify that F is a subgraph of G. For a graph $G = (V, E)$, any non-empty subset A of the vertices generates the subgraph $G(A)$ consisting of all and only vertices in A and edges between two vertices in A, called the *subgraph induced by* A. Similarly, a subset $B \subseteq E$ of edges induces a subgraph that contains the edges in F and all and only vertices that are endpoints of edges in B. We note that *edge induced subgraphs have no isolated vertices.*

A *complete graph* is a graph where all nodes are adjacent, and a *k-star* is a graph where one node (called the *hub*) is adjacent to all other nodes and there is no other edge in the graph. The *line graph* $L(G)$ of a graph $G = (V, E)$ is the intersection graph of the edge set E, i.e. its vertex set is E and $e_1 \sim e_2$ if and only if e_1 and e_2 have a common endpoint (West, 2001, p. 168). We will in particular be interested in the line graph of the complete graph, which we will refer to as the *incidence graph.*

2.2 Homomorphism densities

We will rely on the notion of subgraph homomorphism density, a graph-theoretic concept that is central to the theory of graph convergence (see, e.g., Lovász, 2012).

A map $\phi : V_F \to V_G$ between the vertex sets of two simple graphs $F = (V_F, E_F)$ and $G = (V_G, E_G)$ is a *homomorphism* if it preserves edges, i.e. if uv in E_F implies $\phi(u)\phi(v) \in E_G$. A homomorphism ϕ is an *isomorphism* if it is a bijection and its inverse is a homomorphism, i.e. ϕ also preserves non-edges. If there is an isomorphism between F and G, the graphs are *isomorphic*, and we write $F \cong G$. An isomorphism can also be thought of as a *relabelling* of the vertex set. The isomorphism relation is an equivalence relation among graphs with the same vertex set. Accordingly, for a graph $G = (V, E)$ we may identify the corresponding equivalence class as an *unlabeled graph* with $|V|$ vertices. An isomorphism ϕ between the graphs G and G is an *automorphism*. Then $\mathrm{aut}(G)$ denotes the size of the *automorphism group* , i.e. the group of automorphisms of G.

Following Lovász (2012, Ch. 5.2) and assuming that $|V_F| \leq |V_G|$, we let $\mathrm{inj}(F, G)$ denote the number of injective graph homomorphisms between F and G. Note that if G is a complete graph (i.e. all the vertices are adjacent), all maps are homomorphisms, and $\mathrm{inj}(F, G) = (|V_G|)_{|V_F|} = |V_G|!/(|V_G| - |V_F|)!$. Furthermore, we have that $\mathrm{inj}(G, G) = \mathrm{aut}(G)$. The *injective homomorphism density* of F in G is defined to be

$$t_{\mathrm{inj}}(F, G) = \frac{\mathrm{inj}(F, G)}{(|V_G|)_{|V_F|}},$$

the proportion of all injective mappings from V_F into V_G that are homomorphisms. It is immediate from the above definition that homomorphism densities are invariant under isomorphisms, in the sense that $t_{\text{inj}}(F,G) = t_{\text{inj}}(F',G')$, for any pairs $F \cong F'$ and $G \cong G'$. Thus, density homomorphisms are also well defined over unlabeled graphs. Indeed, the injective homomorphism density $t_{\text{inj}}(F,G)$ can be interpreted as the probability that F is realized as a random subgraph of G (see, e.g. Diaconis and Janson, 2008).

Finally, we will be interested in the number $\text{sub}(F,G)$ of (not necessarily induced) subgraphs of G that are isomorphic to F. This is given as

$$\text{sub}(F,G) = \text{inj}(F,G)/\text{inj}(F,F) = \text{inj}(F,G)/\text{aut}(F). \tag{1}$$

Clearly, also $\text{sub}(F,G)$ remains unchanged if F and G are replaced by isomorphic graphs F' and G' so that $\text{sub}(\cdot,\cdot)$ is well defined over unlabeled graphs.

2.3 The Möbius parametrization for binary arrays

As explained later in detail, network modeling can be considered equivalent to modeling the joint distribution of a collection of binary random variables. For this purpose we will find it convenient to rely on a particular parametrization of such distributions that is based on the Möbius transform (see, e.g. Lauritzen, 1996, Appendix A3) and that therefore, following Drton and Richardson (2008), we call the *Möbius parametrization.*

Let V be a finite set indexing a collection $\{X_v, v \in V\}$ of binary random variables taking value in $\{0,1\}^V$ and for a non-empty subset B of V let $X_B = (X_v, v \in B)$. The Möbius parametrization arises from taking the Möbius transform of the vector of joint probabilities of each point in $\{0,1\}^V$. This transformation, which is linear and invertible, in turns yields a vector of marginal, as opposed to joint, probabilities. In detail, let $\mathcal{B}$ denotes the set of all subsets of V and, for each non-empty $B \in \mathcal{B}$, set $z_B = \mathbb{P}(X_B = 1_B)$, where 1_B is the $|B|$-dimensional vector of all 1's. Further, set $z_\varnothing = 1$. Then, using Möbius inversion, we have that, for each $H \in \mathcal{B}$,

$$\mathbb{P}(X_H = 1_H, X_{H^c} = 0_{H^c}) = \sum_{B \in \mathcal{B}: H \subseteq B} (-1)^{|B \setminus H|} z_B, \tag{2}$$

where $H^c = V \setminus H$. Similarly, for each $B \in \mathcal{B}$,

$$z_B = \mathbb{P}(X_B = 1_B) = \sum_{H \in \mathcal{B}: B \subseteq H} \mathbb{P}(X_H = 1_H, X_{H^c} = 0_{H^c}).$$

The marginal probabilities $\{z_B, B \in \mathcal{B}\}$ are called *the Möbius parameters* of the distribution of X_V. The above formulae show that there is a one-to-one and linear relation between joint probabilities and their Möbius parameters and that the only restriction on the set of feasible Möbius parameters is that all expressions of the form (2) should be non-negative, and $z_\varnothing = 1$.

We can represent the set of all strictly positive probabilities on $\{0,1\}^V$ as an exponential family with canonical sufficient statistic given by

$$x \in \{0,1\}^V \mapsto s(x) \in \{0,1\}^{|\mathcal{B}|-1},$$

where $s(x) = (s_B(x), B \in \mathcal{B} \setminus \varnothing)$, with

$$s_B(x) = \prod_{b \in B} x_b, \quad B \in \mathcal{B};$$

see, for instance, Frank and Strauss (1986). The exponential family on $\{0,1\}^V$ generated by these sufficient statistics consists of all probability distributions of the form

$$\mathbb{P}(X = x) = P_\nu(x) = \exp\left\{ \sum_{B \in \mathcal{B}\setminus\varnothing} \nu_B s_B(x) - \psi(\nu) \right\}, \quad x \in \{0,1\}^V, \tag{3}$$

for any choice of canonical parameters $\nu = (\nu_B, B \in \mathcal{B}\setminus\varnothing) \in \mathbb{R}^{\mathcal{B}\setminus\varnothing}$, and where $\psi(\nu)$ is the *log-partition function*, ensuring that probabilities add to unity. The above exponential family is minimal, full and regular, and the *mean value parameters* (Barndorff-Nielsen, 1978) for this representation are precisely the Möbius parameters corresponding to non-empty subsets of V:

$$\mathrm{E}\{s_B(X)\} = z_B, \quad B \in \mathcal{B} \setminus \varnothing.$$

2.4 Undirected and bidirected graphical models

Graphical models (see, e.g. Lauritzen, 1996) are statistical models expressing conditional independence clauses among a collection of random variables $X_V = (X_v, v \in V)$ indexed by a finite set V. A graphical model is determined by a graph $G = (V, E)$ over the indexing set V, and the edge set E (which may include edges of undirected, directed or bidirected type) encodes conditional independence relations among the variables, or *Markov properties*. A joint probability distribution P for X_V is Markov with respect to G if it satisfies the Markov properties expressed by G.

For non-empty subsets A and B of V and a (possibly empty) subset S of V, we use the notation $A \perp\!\!\!\perp B \,|\, S$ as shorthand for the conditional independence relation $X_A \perp\!\!\!\perp X_B \,|\, X_S$, thus identifying random variables with their labels. When $S = \emptyset$ the independence relation is intended as marginal independence.

For an undirected graph G — where all edges are undirected — the (global) Markov property expresses that $A \perp\!\!\!\perp B \,|\, S$ when every path between A and B has a vertex in S or, in other words, S *separates* A from B in G.

We shall be specifically interested in graphical models given by a *bidirected graph* G where all edges are bidirected. For such graphs the (global) Markov property (Cox and Wermuth, 1993; Kauermann, 1996; Richardson and Spirtes, 2002; Richardson, 2003) expresses that

$A \perp\!\!\!\perp B \mid S$ when every path between A and B has a vertex outside $S \cup A \cup B$, i.e. $V \setminus (A \cup B \cup S)$ separates A from B. Note the obvious duality between this and the Markov property for undirected graphs.

For example, in the undirected graph of Figure 1(a), the global Markov property implies that $\{i, l\} \perp\!\!\!\perp k \mid j$, whereas in the bidirected graph of Figure 1(b), the global Markov property implies that $\{i, l\} \perp\!\!\!\perp k$. Notice that, for simplicity, we write k and j instead of $\{k\}$ and $\{j\}$ as these are single elements.

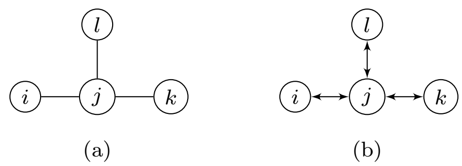

Figure 1: (a) An undirected dependence graph. (b) a bidirected dependence graph.

For a system of random variables X we define the *dependence skeleton* of X, denoted by $\mathrm{sk}(X)$ to be the undirected graph with vertex set V such that vertices u and v are not adjacent if and only if there is *some* subset S of V so that $u \perp\!\!\!\perp v \mid S$. Thus, if X is Markov with respect to an undirected graph G — such as the case considered by Frank and Strauss (1986), who chose S to be the set of remaining variables — $\mathrm{sk}(X)$ would be a subgraph of G; whereas if X is Markov with respect to a bidirected graph, we may choose $S = \emptyset$ and the corresponding skeleton will be a subgraph of the graph obtained by replacing all arcs with lines.

Bidirected graphical models for binary variables For the bidirected case, the exponential representation does not simplify. However, in the Möbius parametrization the bidirected Markov property takes a particular simple form. Indeed Drton and Richardson (2008) showed that a distribution P is Markov with respect to the bidirected graph $G = (V, E)$ if and only if for any $B \subset V$ that is disconnected in G, the corresponding Möbius parameter satisfies

$$z_B = z_{C_1} \cdots z_{C_k}, \tag{4}$$

where $C_1, \ldots, C_k$ partitions B into inclusion maximal connected sets (with respect to G). Thus, if we let $\mathcal{C} = \mathcal{C}(G)$ be the set of all non-empty connected subsets of V (with respect to G), the bidirected binary graphical model with graph G is completely, injectively, and smoothly parametrized by the Möbius parameters $(z_C, C \in \mathcal{C})$. More precisely we have that

$$\mathbb{P}(X_H = 1_H, X_{H^c} = 0_{H^c}) = \sum_{B \in \mathcal{B}: H \subseteq B} (-1)^{|B \setminus H|} \prod_{C \in \mathcal{C}_B} z_C, \tag{5}$$

where $\mathcal{C}_B$ denotes the collection of inclusion maximal connected subsets of B. Further, the Möbius parameters vary freely save for the restrictions that $z_\varnothing = 1$ and the right-hand expressions in (2) and (5) must be non-negative for all choices of $E \subseteq V$. For example, for the bidirected graph of Figure 2 we have that

$$\mathbb{P}(X_1 = 1, X_2 = 0, X_3 = 0) = z_1 - z_{12} - z_1 z_3 + z_{123}, \quad \mathbb{P}(X_1 = 1, X_3 = 1) = z_1 z_3.$$

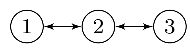

Figure 2: A dependence graph with vertex set $\{1, 2, 3\}$.

Drton and Richardson (2008) further show that the family of positive bidirected Markov distributions becomes *a curved exponential family* due to the restriction (4) which is non-linear in the canonical parameters ν_B and log-linear in the mean value parameters z_B (Roverato et al., 2013). Finally, the dimension of the corresponding model is the cardinality of $\mathcal{C}$, the number of connected induced subgraphs of G.

3 Network models and exchangeability

3.1 Some terminology

Given a finite or countably infinite node set $\mathcal{N}$ — representing *individuals* or *actors* in a given population of interest — we define a *random network* over $\mathcal{N}$ to be a collection $X = (X_d, d \in D(\mathcal{N}))$ of binary random variables taking values 0 and 1 indexed by a set $D(\mathcal{N})$ of *dyads.* Throughout the paper we take $D(\mathcal{N})$ to be the set of all unordered pairs ij of nodes in $\mathcal{N}$, and nodes i and j are said to have a *tie* if the random variable X_{ij} takes the value 1, and no tie otherwise. Thus, a network is a random variable taking value in $\{0,1\}^{\binom{\mathcal{N}}{2}}$ and can therefore be seen as a random *simple, undirected, labeled* graph with node set $\mathcal{N}$, whereby the ties form the random edges of the graphs. We will write $\mathcal{G}_\mathcal{N}$ for the set of all simple, labeled undirected graphs on $\mathcal{N}$.

We use the terms network, node, and tie rather than (random) graph, vertex, and (random) edge to differentiate from the terminology used in the graphical model sense. Indeed, as we shall discuss graphical models for networks, we will also consider each dyad d as a vertex in a graph $G = (D, E)$ representing the dependence structure of the random variables associated with the dyads.

3.2 Exponential random graph models

Exponential random graph models (Frank, 1991; Wasserman and Pattison, 1996), or ERGMs in short, are among the most important and popular statistical models for network data and, as we will see, the set of exchangeable distributions form an ERGM.

ERGMs are exponential families of distributions (Barndorff-Nielsen, 1978) on $\mathcal{G}_\mathcal{N}$ whereby the probability of observing a network $x \in \mathcal{G}_\mathcal{N}$ can be expressed as

$$P_\theta(x) = \exp\Big\{ \sum_{l=1}^{m} s_l(x)\theta_l - \psi(\theta) \Big\}, \quad \theta \in \Theta \subseteq \mathbb{R}^m, \tag{6}$$

where $s(x) = (s_1(x) \ldots, s_m(x)) \in \mathcal{R}^m$ are *canonical sufficient statistics* which capture some important feature of x, $\theta \in \mathbb{R}^m$ is a point in the canonical parameter space Θ, and $\psi\colon \Theta \to [0, \infty)$ is the log-partition function; in (6) we have thus chosen counting measure as the base measure of the representation.

The choice of the canonical sufficient statistics and any restriction imposed on Θ will determine the properties of the corresponding ERGM. The simplest ERGM is the *Erdös–Rényi* model (Erdös and Rényi, 1960), where there is only one parameter θ, and the canonical sufficient statistic is the number of ties in x. This is equivalent to ties occurring independently with probability p and $\theta = \log\{p/(1-p)\}$.

Beta models are also ERGMs and they can be considered a generalization of the Erdös-Rényi model: in these models it is also assumed that ties occur independently; the probability p_{ij} of a tie between nodes i and j is given as:

$$p_{ij} = \mathbb{P}(X_{ij} = 1) = \frac{e^{\beta_i + \beta_j}}{1 + e^{\beta_i + \beta_j}}, \quad \forall ij \in D, \tag{7}$$

where $(\beta_i, i \in \mathcal{N})$ can be interpreted as parameters that determine the propensity of node i to have edges. The probability of a network $x \in \mathcal{G}_\mathcal{N}$ is thus

$$P_\beta(x) = \prod_{ij \in D} p_{ij}^{x_{ij}} (1 - p_{ij})^{1 - x_{ij}} = \prod_{ij \in D} \frac{e^{(\beta_i + \beta_j) x_{ij}}}{1 + e^{(\beta_i + \beta_j)}} = \exp\left\{ \sum_{i \in \mathcal{N}} d_i(x)\beta_i - \psi(\beta) \right\}, \tag{8}$$

where $(d_i(x), i \in \mathcal{N})$ is the *degree sequence* of x, i.e. $d_i(x)$ is the number of ties in x involving node i.

Other ERGMs include, for example, the exponential family models of Holland and Leinhardt (1981) and their modifications (Fienberg and Wasserman, 1981), as well as the Markov graphs of Frank and Strauss (1986).

If $\mathcal{N}' \subseteq \mathcal{N}$, any probability distribution $P_\mathcal{N}$ for a network $X_\mathcal{N} = (X_d, d \in D(\mathcal{N}))$ induces a distribution $P_{\mathcal{N}'}$ for the subnetwork $X_{\mathcal{N}'} = (X_d, d \in D(\mathcal{N}'))$ corresponding to the dyads in $D(\mathcal{N}')$. In particular, $X_{\mathcal{N}'}$ is the subgraph of $X_\mathcal{N}$ induced by $\mathcal{N}'$.

One problem with ERGMs is that, typically, $P_{\mathcal{N}'}$ needs not be the related to $P_{\mathcal{N}}$ in a meaningful way, and needs not be an ERGM itself. In that sense, the ERGMs are not *marginalizable*; see also Snijders (2010) as well as Shalizi and Rinaldo (2013). The Erdös–Rényi models and the beta models are marginalizable in this sense. Many other network models suffer from this issue: see, e.g., Crane and Dempsey (2015). We shall return to this and related issues in more detail in Section 6.

3.3 Parametrizations of random network models

Since networks are arrays of binary variables, we may use the Möbius parametrization and the exponential form in Section 2.3 to represent their distributions.

For a given finite node set $\mathcal{N}$, we thus let $\mathcal{B}(\mathcal{N}) = \mathcal{B}$ be the collection of all subsets of $D(\mathcal{N})$. Now if X is a network on $\mathcal{N}$ and B a non-empty set in $\mathcal{B}$, then $X_B = (X_d, d \in B)$ is the subnetwork indexed by the dyads in B. It will be convenient to identify each non-empty set $B \in \mathcal{B}$ with the edge-induced subgraph of the complete graph on $\mathcal{N}$ comprised of the edges in B. Equivalently, $\mathcal{B} \setminus \varnothing$ represents all subgraphs of the complete graph on $\mathcal{N}$ without isolated nodes. In particular, for a non-empty $B \in \mathcal{B}$, we may equivalently write the event $\{X_B = 1_B\}$ as the event $\{B \subset X\}$ that the graph B is a subgraph of the random graph X.

The Möbius parameters corresponding to the distribution P of a network X are $(z_B, B \in \mathcal{B})$, where

$$z_B = \begin{cases} \mathbb{P}(B \subseteq X) & \text{if } B \neq \varnothing \\ 1 & \text{otherwise.} \end{cases}$$

Thus, the Möbius parameters for a network distribution are just the implied probabilities of observing all subgraphs of the complete graph on $\mathcal{N}$ without isolated nodes. In the exponential family parametrization (3), the canonical sufficient statistic s is defined over graphs $\mathcal{G}_{\mathcal{N}}$, and the coordinate indexed by any non-empty $B \in \mathcal{B}$ is simply the indicator function

$$s_B(x) = \begin{cases} 1 & \text{if } B \subseteq x \\ 0 & \text{otherwise.} \end{cases}$$

Example 1. Suppose that we observe the network x in Figure 3 (a), where $\mathcal{N} = \{1, 2, 3, 4\}$.

For any probability measure P on $\mathcal{G}_{\mathcal{N}}$, we have that, under both the Möbius parametrization and the canonical parametrization in (3),

$$\begin{aligned} P_\nu(x) &= z_{14,23,24,34} - z_{12,14,23,24,34} - z_{13,14,23,24,34} + z_{12,13,14,23,24,34} \\ &= \exp\{\nu_{14} + \nu_{23} + \nu_{24} + \nu_{34} + \nu_{14,23} + \nu_{14,24} + \nu_{14,34} + \nu_{23,24} + \nu_{23,34} + \\ &\quad \nu_{24,34} + \nu_{14,24,23} + \nu_{14,24,34} + \nu_{14,34,23} + \nu_{13,14,34} + \nu_{14,23,24,34} - \psi(\nu)\}, \end{aligned} \tag{9}$$

Figure 3: Two observed networks with four nodes.

where ν is the unique 63-dimensional vector parametrizing the distribution of X. Notice that the Möbius parametrization is considerably simpler than the exponential family parametrization in this instance, since x is a dense network and only parameters corresponding to supergraphs of the observed network enter into the calculation. However, if we instead observe the sparser network $y \in \mathcal{G}_{\mathcal{N}}$ in Figure 3 (b), we get

$$\begin{aligned} P_\nu(y) = z_{14,23} - z_{14,23,12} - z_{14,23,13} - z_{14,23,24} - z_{14,23,34} \\ +z_{14,23,12,13} + z_{14,23,12,24} + z_{14,23,12,34} + z_{14,23,13,34} + z_{14,23,13,24} + z_{14,23,24,34} \\ -z_{12,14,23,24,34} - z_{13,14,23,24,34} + z_{12,13,14,23,24,34} \\ = \exp\{\nu_{14} + \nu_{23} + \nu_{14,23} - \psi(\nu)\}. \end{aligned} \tag{10}$$

Thus the exponential form is simpler for the sparse case, save for the fact that the log-partition function $\psi(\nu)$ is complicated. □

3.4 Exchangeability

We are concerned with probability distributions on networks that are *exchangeable:* invariant under relabelings of the node set $\mathcal{N}$. Exchangeability is a well-known form of probabilistic invariance, and, from a modeling perspective, also a natural simplifying assumption to impose when formaliziing statistical models for networks. Below we summarize well-known facts on exchangeability.

A distribution P of a random array $X = (X_{ij})_{i,j \in \mathcal{N}}$ over a finite nodeset $\mathcal{N}$ is said to be *weakly exchangeable* (WE) (Silverman, 1976; Eagleson and Weber, 1978) if for all permutations $\pi \in S(\mathcal{N})$ we have that

$$\mathbb{P}\{(X_{ij} = x_{ij})_{i,j \in \mathcal{N}}\} = \mathbb{P}\{(X_{ij} = x_{\pi(i)\pi(j)})_{i,j \in \mathcal{N}}\}. \tag{11}$$

If the array X is symmetric — i.e. $X_{ij} = X_{ji}$, we say it is *symmetric weakly exchangeable* (SWE). In the following we shall for brevity say that X is WE, SWE, etc. in the meaning that its distribution is.

A symmetric binary array with zero diagonal can be interpreted as a matrix of ties (adjacency matrix) of a random network and thus the above concepts can be translated into networks. A random network is *exchangeable* if its adjacency matrix is SWE. Then a random network is exchangeable if and only if its distribution is invariant under relabeling of the nodes of the network. We will distinguish the case of a finite node set from that of a (countably) infinite one, and in the former case we will speak of *finite exchangeability.* A random network X over a countably infinite nodeset $\mathcal{N}$ is said to be exchangeable, if every finite induced subnetwork $X_{\mathcal{N}'}$ for $\mathcal{N}' \subset \mathcal{N}$ is.

When $\mathcal{N}$ is infinite, exchangeability is well-understood. Indeed, any doubly infinite random SWE array can be represented as a mixture of so-called *ϕ-matrices* (Diaconis and Freedman, 1981) in the following way. Let $\mathcal{P}_\infty$ be the convex set of all exchangeable distributions for networks with a countably infinite number of nodes and let $\mathcal{E}_\infty$ denote the extreme points of that set. Then we have:

Proposition 1. *$P \in \mathcal{E}_\infty$ if and only if for every finite $\mathcal{N}' \subset \mathcal{N}$,*

$$P_{\mathcal{N}'}(x) = \int_{[0,1]^{\mathcal{N}}} \prod_{ij \in D(\mathcal{N}')} \phi(u_i, u_j)^{x_{ij}} \{1 - \phi(u_i, u_j)\}^{1-x_{ij}}\, du, \quad \forall x \in \mathcal{G}_{\mathcal{N}'}, \tag{12}$$

where $P_{\mathcal{N}'}$ is the induced probability on $\mathcal{G}_{\mathcal{N}'}$ and ϕ is a measurable function from $[0,1]^2$ to $[0,1]$. The function ϕ is unique up to a measure-preserving transformation of the unit interval.

In a graph theoretic context, (equivalence classes of) ϕ-matrices are also known as *graphons* (Lovász and Szegedy, 2006).

Elements of $\mathcal{E}_\infty$ as given in (12) are all *dissociated* (Silverman, 1976) in the sense that $\{X_{ij}, i, j \in \mathcal{N}'\}$ are independent of $\{X_{kl}, k, l \in \mathcal{N}''\}$ whenever $\mathcal{N}' \cap \mathcal{N}'' = \emptyset$ in fact it holds (Aldous, 1981, 1985) that

Proposition 2. *$P \in \mathcal{E}_\infty$ if and only if P is exchangeable and dissociated.*

We emphasize that the representation of exchangeable distributions on networks by means of ϕ-matrices requires an infinite node set $\mathcal{N}$. If $\mathcal{N}$ is finite, a finitely exchangeable network on $\mathcal{N}$ needs not have a representation as a mixture of ϕ-matrices, and need not be *extendable*, i.e. the induced probability measure from a probability distribution over networks on a larger node set. In fact, the properties of finitely exchangeable networks are distinctively different; see for example Diaconis (1977); Diaconis and Freedman (1980); Matúš (1995); Kerns and Székely (2006); and Konstantopoulos and Yuan (2015). We will return to the issue of extendability later in Section 6 and focus instead on the properties of finitely exchangeable network distributions.

3.5 Exponential representation of exchangeable networks

For exchangeable random networks on a finite nodeset $\mathcal{N}$, both the Möbius and the exponential parameters simplify. More precisely, P is an exchangeable distribution on

$\mathcal{G}_{\mathcal{N}}$ if and only if $P(x) = P(x')$ whenever $x \cong x'$, i.e. whenever x and x' are identical save for a permutation of their labels. Thus, it follows that a random network X is exchangeable if and only if $z_B = z_{B'}$ whenever $B \cong B'$, where $z_B = \mathbb{P}(B \subseteq X)$.

Next we let $\mathcal{U}_{\mathcal{N}}$ denote the set of all unlabeled graphs on $\mathcal{N}$ and write $\varnothing$ for the empty graph. Each $U \in \mathcal{U}_{\mathcal{N}}$ can be identified with an isomorphism class in $\mathcal{G}_{\mathcal{N}}$, which we will denote with $[U]$. Then, any exchangeable distribution P on $\mathcal{G}_{\mathcal{N}}$ is parametrized by $\{z_U, U \in \mathcal{U}_{\mathcal{N}}\}$, with $z_\varnothing = 1$, where $z_U = \mathbb{P}(B \subseteq X)$, for any $B \in [U]$. Collecting identical terms in (2) we obtain

$$\mathbb{P}(X = x) = \sum_{U \in \mathcal{U}_{\mathcal{N}}} (-1)^{|U|-|x|} r_U(x) z_U, \quad x \in \mathcal{G}_{\mathcal{N}} \tag{13}$$

where $|U|$ is the number of edges in U and $r_U(x)$ is the number of graphs in $[U]$ that contain x as a subgraph.

Similarly, it holds for the exponential representation (3) that a probability distribution on $\mathcal{G}_{\mathcal{N}}$ is exchangeable if and only if for all B and $B = B'$ in $D(\mathcal{N})$ with $B \cong B'$, the corresponding canonical parameters satisfy $\nu_B = \nu_{B'}$. Thus we can represent the family of exchangeable network distributions on $\mathcal{N}$ as the exponential family of probability distributions of the form

$$P_\nu(x) = \exp\Big\{ \sum_{U \in \mathcal{U}_{\mathcal{N}} \setminus \varnothing} \sigma_U(x)\nu_U - \psi(\nu) \Big\}, \quad x \in \mathcal{G}_{\mathcal{N}}, \tag{14}$$

for any choice of the canonical parameters $\nu = (\nu_U, U \in \mathcal{U}_{\mathcal{N}} \setminus \varnothing) \in \mathbb{R}^{\mathcal{U}_{\mathcal{N}} \setminus \varnothing}$, where $\psi(\cdot)$ is the log-partition function and

$$\sigma_U(x) = \sum_{B \in [U]} s_B(x)$$

is the number of graphs in the isomorphism class corresponding to U that are subgraphs of x. Indeed from (1) we have that

$$\sigma_U(x) = \mathrm{inj}(U, x)/\mathrm{inj}(U, U). \tag{15}$$

Note that the set of exchangeable distributions again form a linear and regular exponential family with canonical sufficient statistics $(\sigma_U, U \in \mathcal{U}_{\mathcal{N}} \setminus \varnothing)$, canonical parameters $(\nu_U, U \in \mathcal{U}_{\mathcal{N}} \setminus \varnothing)$, and mean value parameters

$$\tau_U = \mathrm{E}\{\sigma_U(X)\} = \sum_{B \in [U]} z_U = \mathrm{sub}(U, D) z_U = \mathrm{inj}(U, D) z_U / \mathrm{inj}(U, U) \tag{16}$$

for $U \in \mathcal{U}_{\mathcal{N}} \setminus \varnothing$. In other words, *the family of finitely exchangeable network distributions is an ERGM.*

By standard theory of exponential families (Barndorff-Nielsen, 1978), upon observing a network $x \in \mathcal{G}_{\mathcal{N}}$, the maximum likelihood estimate of the Möbius parameters $(z_U, U \in \mathcal{U}_{\mathcal{N}} \setminus \varnothing)$ under the assumption of exchangeability is obtained by equating the observed canonical statistic to its expectation. As a result, the MLE of z_U, for any non-empty U is

$$\hat{z}_U = \sigma_U(x)/\mathrm{sub}(U, D) = \mathrm{inj}(U, x)/\mathrm{inj}(U, D) = t_{\mathrm{inj}}(U, x). \tag{17}$$

Strictly speaking, this is not the maximum likelihood estimate in the exponential family, but rather its *completion* (Barndorff-Nielsen, 1978) obtained by including all limit points of the set of mean value parameters. In fact, the canonical parameters ν_U are not estimable based on observation of a single network. Note that, interestingly, *the maximum likelihood estimates of the Möbius parameters are exactly the injective homomorphism densities introduced in Section 2.* For an extreme and infinitely exchangeable network (the dissociated case), these estimates are also well known to be consistent (Lovász and Szegedy, 2006, Theorem 2.5).

Example 2. Consider the networks x and y in Figure 3. Using the exponential representation (14) we see — with a notation that hopefully is transparent — that the probability of observing x is

$P_\nu(x) = \exp\Big\{4\nu$  $+ 5\nu$  $+ \nu$ 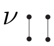 $+ \nu$ 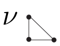 $+ \nu$ 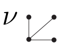 $+ 2\nu$ 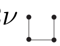 $+ \nu$ 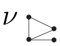 $- \psi(\nu)\Big\}.$ (18)

In terms of the Möbius parametrization the expression simplifies to

$P_\nu(x) = z$  $- 2z$ 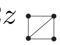 $+ z$ 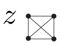. (19)

If x is observed, the non-zero estimates of Möbius parameters for $U \neq \varnothing$ are

$\hat{z}$  $= 2/3$, $\hat{z}$ 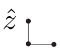 $= 5/12$, $\hat{z}$ 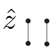 $= 1/3$, $\hat{z}$ 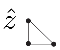 $= 1/4$,
$\hat{z}$ 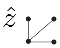 $= 1/4$, $\hat{z}$ 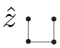 $= 1/6$, $\hat{z}$ 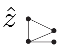 $= 1/12$; (20)

These numbers are obtained from (17) by calculating, for example, as

inj(, $x) = 3! = 6$, inj(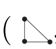, $D) = 4 \times 3 \times 2 = 24$.

Note that the canonical parameters in (18) are not estimable in this case as the set of sufficient statistics $(\sigma_U(X), U \in \mathcal{U}_{\mathcal{N}} \setminus \varnothing)$ are at the boundary of their convex support. $\square$

4 Independence structures for exchangeable networks

In this section we will investigate the Markov properties of exchangeable network distributions on a finite node set.

4.1 Possible independence structures

Recall that we use the term *incidence graph* $L(\mathcal{N})$ for the line graph of a complete graph on $\mathcal{N}$, whereby $L(\mathcal{N})$ has edges between dyads which are *incident* i.e. dyads having a node in common. Figure 4 displays the incidence undirected graph for $\mathcal{N} = \{1, 2, 3, 4\}$ and a bidirected version $L_{\leftrightarrow}(\mathcal{N})$ of the same. In fact, we show next that the skeleton

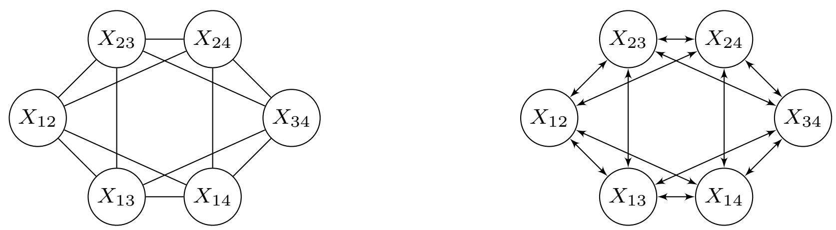

Figure 4: (a) The incidence graph for $\mathcal{N}\{1, 2, 3, 4\}$. (b) The bidirected incidence graph with 4 nodes.

graph of a finitely exchangeable network distribution can only be one of four possible types.

Proposition 3. *The dependence skeleton* $\mathrm{sk}(X)$ *of an exchangeable random network* X *is one of the following:*

(a) *the empty graph;*

(b) *the incidence graph* $L(\mathcal{N})$ *;*

(c) *the complement of the incidence graph* $L^c(\mathcal{N})$ *;*

(d) *the complete graph.*

Proof. Let $|\mathcal{N}| = n$. Exchangeability implies that one can relabel the $\binom{n}{2}$ vertices of the dependence skeleton corresponding to every permutation of $(1, \ldots, n)$ without changing independence statements of form $i \perp\!\!\!\perp j \,|\, S$ for some $S \subset D$. If the dependence skeleton is empty or complete it is clear that relabeling would not change such independence statements.

If there is an edge ij, kl, for mutually non-equal i, j, k, l, and a missing edge $i'j', k'l'$, for mutually non-equal i', j', k', l', in the dependence skeleton, then by the relabeling corresponding to sending (i', j', k', l') to (i, j, k, l) we obtain a missing edge between ij and kl. However, in the latter labeling, there is an independence of form $ij \perp\!\!\!\perp kl \,|\, S$ whereas in the former labeling there is no independence of form $ij \perp\!\!\!\perp kl \,|\, S$. Hence, such a graph cannot be the dependence skeleton of an exchangeable random network.

Similarly, if there is an edge ij, il, for mutually non-equal i, j, l, and a missing edge $i'j', i'l'$, for mutually non-equal i', j', l', in the dependence skeleton, then by the relabeling corresponding to sending (i', j', l') to (i, j, l) we obtain a missing edge between ij and il. With the same argument as in the previous case, such a graph cannot be the dependence skeleton of an exchangeable random network.

The remaining cases are when either all the pairs of vertices ij, il are adjacent and all the pairs of vertices ij, kl are not adjacent, which leads to the $L(\mathcal{N})$; or all the pairs of vertices ij, kl are adjacent and all the pairs of vertices ij, il are not adjacent, which leads to $L^c(\mathcal{N})$. □

□

To give a complete characterization of all possible graphical Markov structures with these skeletons, we need to give a precise definition of a generic graphical Markov structure. We abstain from doing so but just mention that it is clear that directed edges are not relevant for exchangeable random networks. If we vary the type of edge so that edges are either all bidirected or undirected, this leads to six different dependence structures since the undirected and bidirected graphs in cases (a) and (d) are Markov equivalent. If an exchangeable random network X satisfies the Markov property associated with an undirected or bidirected dependence graph G with vertex set $\{X_{ij}, ij \in D(\mathcal{N})\}$ then G is one of these cases. Dependence structures corresponding to the complete graph in (d) are uninteresting. The case (a) corresponds to the Erdös–Renyi model. Structures corresponding to the other four dependence graphs are classified below, where our main focus is on the bidirected incidence graph case as it corresponds to the dissociated model and has certain desired properties concerning extendability; see Section 6.

Undirected incidence graph – the Frank-Strauss model. Frank and Strauss (1986) showed the result below:

Proposition 4. *Every clique C of the incidence graph $L(\mathcal{N})$ of the complete network base $(\mathcal{N}, D)$ corresponds exactly to a subnetwork $(\mathcal{N}_C, C)$ of $(\mathcal{N}, D)$ that is either a triangle or a k-star.*

Based on this, they show that a random network X is Markov with respect to $L(\mathcal{N})$ if and only if the canonical parameters ν_B are zero unless B is a k-star or a triangle. Assuming also exchangeability, this model is known as the Frank-Strauss model and studied extensively in Frank and Strauss (1986).

Example 3. Consider again the network x in Figure 3 and any choice of the canonical parameters $\nu \in \mathbb{R}^{\mathcal{U}_\mathcal{N} \setminus \varnothing}$. Following Example 2, the probability of observing x under the Frank–Strauss model becomes

$P_\nu(x) = \exp\Big\{4\nu$  $+ 5\nu$  $+ \nu$  $+ \nu$ 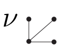 $- \psi(\nu)\Big\}, \quad (21)$

whereas no further simplification appears in terms of the mean-value parameters $\{z_U, U \in \mathcal{U}_\mathcal{N}\}$. □

Undirected complement of the incidence graph. Cliques in the complement of the incidence graph $L^c(\mathcal{N})$ correspond to collections of disjoint dyads. Therefore, the corresponding model is an ERGM with sufficient statistics corresponding to collections of disjoint dyads (including the subnetwork that is a single dyad). For the case of a network with five nodes, the dependence graph is the famous *Petersen graph* (Petersen, 1898) and for n vertices, the graph is known as the Kneser graph $KG_{n,2}$ (Godsil and Royle, 2001, Ch. 7); for this reason we shall refer to the model determined by all exchangeable random networks that are Markov w.r.t. the Kneser graph as the *Kneser model*.

Example 4. Consider again the network x in Figure 3. Still following Example 2 we get using the exponential representation (14) and the Markov property w.r.t. $L^c(\mathcal{N})$ that in the Kneser model, the probability of observing x is

$$P_\nu(x) = \exp\left\{4\nu_{\vert} + \nu_{\vert\,\vert} - \psi(\nu)\right\}, \tag{22}$$

for any $\nu \in \mathbb{R}^{\mathcal{U}_\mathcal{N} \setminus \varnothing}$. □

Bidirected complement of the incidence graph. Every connected subset of the complement of the incidence graph $L^c_\leftrightarrow(\mathcal{N})$ corresponds to a disconnected subnetwork of the of size $\mathcal{N}$ or the subnetwork that is a single tie (which corresponds to a single vertex in $L^c_\leftrightarrow(\mathcal{N})$). Therefore, the corresponding model has Möbius parameters corresponding to these. This model satisfies a property that is dual to dissociatedness so that, for example $X_{ij} \perp\!\!\!\perp X_{jk}$.

Example 5. If the random network X is Markov w.r.t. the complement of the bidirected incidence graph $L^c_\leftrightarrow(\mathcal{N})$, the Möbius parametrization of the probability of the network x in Figure 3 is

$$\mathbb{P}\left(X = x\right) = (z_{\vert})^2 z_{\vert\,\vert} - 2z_{\vert}(z_{\vert\,\vert})^2 + (z_{\vert\,\vert})^3. \tag{23}$$

This has been obtained, for example, as $z_{\triangleright} = (z_{\vert})^2 z_{\vert\,\vert}$ since, for an arbitrary labeling of $\triangleright$, two of the three disconnected induced subgraphs $L^c_\leftrightarrow(\mathcal{N})[12,34]$, $L^c_\leftrightarrow(\mathcal{N})[13,24]$, and $L^c_\leftrightarrow(\mathcal{N})[14,23]$ of $L^c_\leftrightarrow(\mathcal{N})$ correspond to the subnetwork $\vert$ and one to $\vert\,\vert$. □

Later in Section 6 we will argue that the only Markov structures that are also extendable to arbitrarily large networks are the complete graph, the empty graph, and the bidirected incidence graph; hence the latter of these structures is the main focus of the present paper.

4.2 Bidirected incidence graph – the dissociated model

We now consider in detail the class of finitely exchangeable network distributions that are Markov with respect to the bidirected incidence graph $G = L_\leftrightarrow(\mathcal{N})$, where $\mathcal{N}$ is finite.

This is equivalent to requiring an exchangeable network X on $\mathcal{N}$ to be dissociated: $\{X_{ij}, i, j \in \mathcal{N}'\}$ are independent of $\{X_{kl}, k, l \in \mathcal{N}''\}$ for each pair disjoint subsets $\mathcal{N}'$ and $\mathcal{N}''$ of $\mathcal{N}$.

We begin by stating a simple yet useful property of the incidence graph. For a given finite node set $\mathcal{N}$, let $C(\mathcal{N}) \subset D(\mathcal{N})$ denote the subsets of dyads corresponding to connected subgraphs of the complete graph on $\mathcal{N}$.

Lemma 1. *A random network X is dissociated if and only if for any non-empty sub-network B it holds that*

$$z_B = \prod_i z_{C_i},$$

where the C_i are the connected components of B.

Proof. The proof follows directly from (4) and the fact that there is a one-to-one correspondence between the set of all connected subgraphs of $L_{\leftrightarrow}(\mathcal{N})$ and $C(\mathcal{N})$. □
□

With some abuse of notation, for any non-empty $U \in \mathcal{U}_{\mathcal{N}}$, we will write $U = \dot{\cup} C$ for the decomposition of one arbitrary graph in the isomorphism class $[U]$ into the disjoint union of its connected components in $C(\mathcal{N})$. Note that each graph in $[U]$ has the same number of isomorphic connected components. By exchangeability and dissociatedness, for any $U \in \mathcal{U}_{\mathcal{N}}$,

$$z_U = \prod_{C \in C(\mathcal{N}): U = \dot{\cup} C} z_C,$$

regardless of the choice of the representing graph in $[U]$. From (4) and (13) we thus arrive at the following result.

Proposition 5. *A network X on $\mathcal{N}$ is exchangeable and dissociated if and only if*

$$\mathbb{P}(X = x) = \sum_{U \in \mathcal{U}_{\mathcal{N}}} (-1)^{|U| - |x|} r_U(x) \prod_{C \in C(\mathcal{N}): U = \dot{\cup} C} z_C, \quad \forall x \in \mathcal{G}_{\mathcal{N}}, \tag{24}$$

It is worth noting that the model combining dissociatedness and exhangeability remains a *log-mean linear model* in the sense of Roverato et al. (2013).

Example 6. The additional restriction of dissociatedness does not yield any helpful simplification of the exponential representations in (18) from Example 2, since the dissociatedness conditions entails restriction on the possible values of the canonical parameters and not on the form of the probabilities. For the observed network x, no further simplification of its Möbius parametrization in (19) is possible.

Continuing with Example 2, under dissociatedness, the likelihood function in (19) should now be maximized not only under the constraints that all expressions in (2) be non-negative, but also under the non-linear constraints described in (4), which in this case reduces to just one constraint $z = z^2$. Thus, although the likelihood function is linear in z, numerical maximization is in general necessary. In this particular case we find that

$$\hat{z} = 1/2,\ \hat{z} = 5/16,\ \hat{z} = 1/4\ \hat{z} = 3/16,$$
$$\hat{z} = 3/16,\ \hat{z} = 1/8,\ \hat{z} = 1/16,$$

and the remaining parameter estimates are zero, which should be compared to the results in the non-dissociated case. In fact, this estimate represents a mixture of the uniform distribution of all networks isomorphic to the observed network x, and the empty network, with weights 3/4 and 1/4, respectively. □

We emphasize that when dissociatedness is assumed, the maximum likelihood estimator may no longer be unique, as the example below shows.

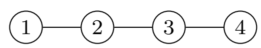

Figure 5: An observed networks with four nodes.

Example 7. Suppose that we observe the network in Figure 5. Then the likelihood function assuming dissociatedness is maximized at any value of λ satisfying $0 \le \lambda \le 1/16$ by the quantities

$$\hat{z} = 1/2,\ \hat{z} = 3/16,\ \hat{z} = 1/4,\ \hat{z} = 1/16 - \lambda,\ \hat{z} = \lambda,\ \hat{z} = 1/16,$$

and all other $\hat{z}$ equal to zero. Indeed this corresponds to a random network that has probability 3/4 of being isomorphic to the observation and the remaining probability mass of 1/4 is distributed arbitrarily between a triangle plus an isolated point, and a 3-star. □

5 Exchangeable and summarized random networks

An assumption commonly used in network modeling is that the probability of a network configuration x in $\mathcal{G}_\mathcal{N}$ be a function of its *degree distribution* $\{n_j(x), j \in \mathcal{N}\}$, where

$$n_j(x) = \sum_{i \in \mathcal{N}(x)} \delta_j\{d_i(x)\}, \quad j \in \mathcal{N},$$

with $d_i(x) = \sum_{j \in \mathcal{N}} X_{ij}$ the degree of node i and $\delta_j(k) = 1$ if $k = j$ and 0 otherwise. Note that here the degree distribution is not normalized and $n_j(x)$ is the number of nodes in x with degree j.

Lauritzen (2008) defines a random binary array $X = (X_{ij})_{i,j \in \mathcal{N}}$ to be *weakly summarized* (WS) if there is a function φ such that it holds that

$$\mathbb{P}\{(X_{ij} = x_{ij})_{i,j \in \mathcal{N}}\} = \varphi(R_i + C_i, i \in \mathcal{N}), \tag{25}$$

where $R_i = \sum_j X_{ij}$ and $C_j = \sum_i X_{ij}$ are the row and column sums. For symmetric arrays, it holds that $R_i = C_i$, and a distribution with the above property may be called *symmetric weakly summarized* (SWS). Again, this can be translated into statements about networks: A random network X on $\mathcal{N}$ is *summarized* if, for any $x \in \mathcal{G}_{\mathcal{N}}$, $\mathbb{P}(X = x) = \varphi(d_i(x), i \in \mathcal{N})$, where $d_i(x) = R_i = C_i$ is the degree of node i. SWE distributions are generally not SWS or vice versa (Lauritzen, 2008). If a distribution is both SWE and SWS then we write that the distribution is SWES. When $\mathcal{N}$ is infinite and X is a SWES array whose distribution is in $\mathcal{E}_\infty$, Lauritzen (2008) showed that Proposition 1 holds with ϕ having a specific form: $\phi(u, v) = a(u)a(v)/\{1 + a(u)a(v)\}$ (Lauritzen, 2008). The beta model in (7) then emerges when further conditioning on u and letting $\beta_i = \log a(u_i)$.

5.1 Characterizing summarizing statistics

We now study exchangeable and summarized distributions on a finite node set $\mathcal{N}$. Isomorphic graphs have the same degree distribution; but, as shown for instance in Example 8 below, the opposite is not true and thus the set of summarized and exchangeable network distributions on $\mathcal{N}$ is a subfamily of the set of exchangeable distribution on $\mathcal{N}$. In order to characterize the types of restriction that summarizedness entails, we begin with some preliminary lemmas which identify which subgraph counts are functions of the degree distribution only. Recall that subgraph counts are canonical sufficient statistics for the exponential family of exchangeable network distributions: see Section 3.5.

Lemma 2. *Given an unlabeled graph $U \in \mathcal{U}_{\mathcal{N}}$, it holds that $\sigma_U(x)$ is a function of the degree distribution of x, for all $x \in \mathcal{G}_{\mathcal{N}}$, if*

(a) *$U = S_k$, a k-star for any size $k \geq 0$ or $U =$ , the disjoint union of two independent edges, or*

(b) *$U = K_{|\mathcal{N}|-1}$, the complete graph of size $|\mathcal{N}| - 1$, or $U = K_{|\mathcal{N}|-1} \setminus \{e\}$, the complete graph of size $|\mathcal{N}| - 1$ with one edge removed.*

Proof. To prove (a), note that a 0-star is simply a vertex and a 1-star is an edge. If $U = S_k$ for $k \neq 1$ then

$$\sigma_{S_k}(x) = \sum_{j=0}^{\infty} \binom{j}{k} n_j(x),$$

where the sum only effectively extends at most from $j = k$ to $|\mathcal{N}| - 1$ as all other terms are zero. If $U = S_1$ we have to divide by two, and obtain

$$\sigma_{S_1}(x) = \sum_{j=0}^{\infty} j n_j(x)/2 = |E(x)| \tag{26}$$

where $|E(x)|$ is the number of edges in x. The formula (26) is also known as Euler's *handshaking lemma*; see e.g. (Wilson, 1996, p.12). Finally, if $U =$ we have

σ 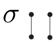 $= \binom{|E(x)|}{2} - \sigma_{S_2}(x).$

For (b), we have that

$$\sigma_{K_{|\mathcal{N}|-1}}(x) = \begin{cases} |\mathcal{N}|, & \text{if } n(x) = (0, \ldots, 0, |\mathcal{N}|); \\ 2, & \text{if } n(x) = (0, \ldots, 0, 2, |\mathcal{N}| - 2); \\ 1, & \text{if } n(x) = (0, \ldots, 0, 1_r, 0, \ldots, 0, |\mathcal{N}| - r - 1, r); \\ 0, & \text{otherwise,} \end{cases}$$

where 1_r is 1 at the rth entry. The top row corresponds to the complete graph, the second row to the complete graph with one edge removed, and the third row to a complete graph with a vertex from which $|\mathcal{N}| - r - 1$ edge is removed.

A similar function can be provided for the case of $K_{|\mathcal{N}|-1} \setminus \{e\}$, but we abstain from giving the somewhat cumbersome details.

□

□

The converse to the statement in Lemma 2 only holds with some modification:

Lemma 3. *Given an unlabeled graph $U \in \mathcal{U}_{\mathcal{N}}$, if $\sigma_U(x)$ is a function of the degree distribution of x, for all $x \in \mathcal{G}_{\mathcal{N}}$, then*

(a) *if U has at most $|\mathcal{N}| - 2$ vertices, U is a k-star or $U =$* *, the disjoint union of two independent edges;*

(b) *if U has at most $|\mathcal{N}| - 1$ vertices, U is a k-star, or $U =$* *, or $K_{|\mathcal{N}|-1}$, or $K_{|\mathcal{N}|-1} \setminus e$, a complete graph on $|\mathcal{N}| - 1$ vertices with one edge removed.*

Proof. It is enough to show that for any U that is not one of the above graphs, there exist two graphs x_1 and x_2 with the same degree sequence (and hence the same degree distribution) such that $\sigma_U(x_1) \neq \sigma_U(x_2)$.

(a): For U with at most $|\mathcal{N}| - 2$ vertices, if U contains two adjacent vertices i and j with degrees more then 1 then define x_1 and x_2 as follows: Let x_1 contain U and a disjoint edge kl. Let x_2 contain U with the ij edge removed and two non-adjacent vertices k and l adjacent to i and j respectively. These two graphs obviously have the same degree sequence and $\sigma_U(x_1) = 1$ whereas $\sigma_U(x_2) = 2$.

Hence, the cases remaining to be considered are such where U is a forest of stars. If there are at least two stars in U that are 2-stars or larger then call these S and T. Let x_1 be U and define x_2 as follows. Connect the hubs of S and T. In addition, take a leaf from each of S and T, connect them and remove the edge between them and their respective hubs. Now x_1 and x_2 have the same degree sequence but $\sigma_U(x_1) = 1$ and $\sigma_U(x_2) = 0$.

If U contains only one star S with hub s that is not an edge and some disjoint edges then take an edge ij in U which is disjoint from S. Define x_1 by adding a vertex l to U and connecting it to i and to a leaf k of S. Define x_2 by taking x_1, removing the sk edge and the ij edge, and adding the ki edge and the sj edge. Both graphs have the same degree sequence, but if S is a 3-star or larger we have $\sigma_U(x_1) = 2$ and $\sigma_U(x_2) = 3$; if S is a 2-star we have $\sigma_U(x_1) = 6$ and $\sigma_U(x_2) = 9$.

Finally, the only case that is left is when U is collection of three or more disjoint edges: ij, kl, hr, and a set $\mathcal{E}$ of disjoint edges. Let x_1 be the union of the cycle $i \sim k \sim l \sim j \sim i$, the hr edge, and $\mathcal{E}$; let x_2 be the union of the path $h \sim i \sim j \sim l \sim k \sim r$ and $\mathcal{E}$. Both graphs have the same degree sequence but $\sigma_U(x_1) = 2$ whereas $\sigma_U(x_2) = 1$.

(b): By (a) it is enough to work with U with $n-1$ vertices that is not one of those mentioned. Collections of disjoint edges and stars (including forest of edges) have been covered in the proof of (a). If U is disconnected and not one of these, then take two connected components with edges ij and kl in each component. Let x_1 be U, and define x_2 by removing ij and kl and connecting i to k and j to l. We have that $\sigma_U(x_1) = 1$ whereas $\sigma_U(x_2) = 0$.

Hence, henceforth, assume that U is connected. If there is an induced P_4 or C_4 (i, j, k, l) in U, i.e., an induced path or cycle with 4 nodes, then define x_1 and x_2 as follows: For x_1, add a vertex h and connect it to i. For x_2, add h and connect it to l, remove kl, and add ik. The graphs have the same degree sequence, but $\sigma_U(x_1) > 0$ whereas $\sigma_U(x_2) = 0$.

If there are no induced P_4 or C_4 the graph is trivially perfect (see Brandstadt et al. (1999) Section 7.1.2) and there is a vertex i that is adjacent to all other vertices in U.

If among the neighbours of i, there are two non-adjacent pairs of vertices and an edge we have the following two cases:

1) Among the neighbours of i, suppose that there are two non-adjacent pairs of vertices (j, k) and (k, l) and a pair of adjacent vertices (l, h), where h may be the same as j. In this case, define x_1 by adding u and connecting it to k, and define x_2 by adding

u and connecting it to h, removing lh, and connecting kl. The graphs have the same degree sequence, but $\sigma_U(x_1) > 0$ whereas $\sigma_U(x_2) = 0$.

2): Among the neighbours of i, suppose that there are two non-adjacent pairs of vertices (j,k) and (l,h) and a pair of adjacent vertices (k,l), where h may be the same as j. In this case, define x_1 by adding u and connecting it to h, and define x_2 by adding u and connecting it to k, adding lh, and removing kl. The graphs have the same degree sequence, but $\sigma_U(x_1) > 0$ whereas $\sigma_U(x_2) = 0$.

Therefore, the only cases that are left are stars, the complete graph, or the complete graph with one edge removed. This completes the proof. □

□

In general, if x_1 and x_2 have the same degree distributions, the fact that their probabilities should be identical induces a linear restriction on the parameters $(\nu_U, U \in \mathcal{U}_\mathcal{N})$ and hence the set of exchangeable and summarized distributions on a nodeset $\mathcal{N}$ is again a linear exponential family with dimension equal to the number of distinct degree distributions for simple graphs on $|\mathcal{N}|$ nodes; however, the restrictions induced on the parameters in (14) may be complicated, as illustrated in the example below.

Similarly, if x_1 and x_2 have the same degree distributions, a linear restriction is induced on the parameters $(z_U, U \in \mathcal{U}_\mathcal{N})$ in (13). This is also illustrated in the example below.

Example 8. For networks with at most four nodes, the isomorphism class of a network is uniquely determined by the degree distribution, so we consider networks with five nodes. For most such networks, the isomorphism class is also identified by the degree distribution, the six exceptions being displayed in Fig. 6, where the node degrees of x_1 and x_2 both are $(2,2,2,1,1)$, those of y_1 and y_2 $(3,2,2,2,1)$, and those of z_1 and z_2 $(3,3,2,2,2)$.

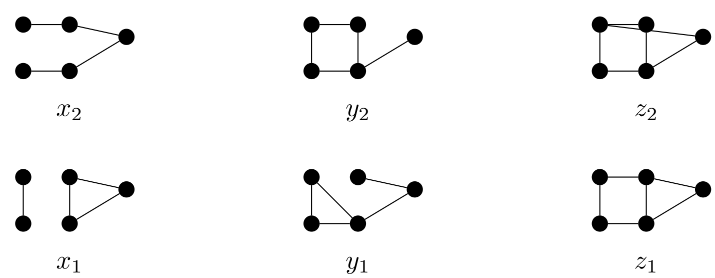

Figure 6: Three pairs of non-isomorphic graphs with identical degree distributions.

It thus follows that for any choice of canonical parameters $\nu \in \mathbb{R}^{\mathcal{U}_\mathcal{N} \setminus \varnothing}$, the distribution P_ν of an exchangeable network on five nodes is summarized if and only if $P_\nu(x_1) = P_\nu(x_2)$,

$P_\nu(y_1) = P_\nu(y_2)$, and $P_\nu(z_1) = P_\nu(z_2)$. Using the exponential representation (14) we obtain that the canonical parameters must satisfy three linear relations. From the relation $P_\nu(x_1) = P_\nu(x_2)$ we get

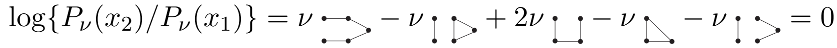

where we have used Lemma 2 to infer that all terms of the form $\nu_U \sigma_U(\cdot)$ must cancel if U is a k-star or $U =$ . Similarly, we get from $P_\nu(y_1) = P_\nu(y_2)$ that

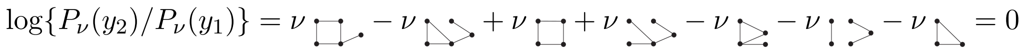

and a somewhat more involved linear relation appears from the fact that $P_\nu(z_1) = P_\nu(z_2)$. Similarly using (13) we obtain that the Möbius parameters must satisfy three linear relations. The simplest of the three relations with Möbius parameters for $P_\nu(z_1) = P_\nu(z_2)$ is

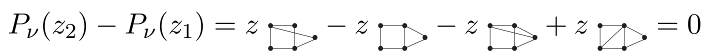

We refrain from giving details of the corresponding constraints with Möbius parameters obtained from $P_\nu(x_1) = P_\nu(x_2)$ and $P_\nu(y_1) = P_\nu(y_2)$ as well as with canonical parameters obtained from $P_\nu(z_1) = P_\nu(z_2)$. □

Corollary 1. *Let U be a finite unlabelled graph without isolated nodes. If there is a function φ_U so that $\sigma_U(x) = \varphi_U\{n_j(x), j = 0, 1, 2, \ldots\}$ for all networks x of any size, then $U = S_k$, a k-star, or $U =$* *, the disjoint union of two independent edges.*

Proof. If this is true for all networks, we can always choose $\mathcal{N}$ sufficiently large for Lemma 3, (a) to apply. □

□

5.2 The SE* model

Motivated by the results from the previous section, we can now simplify the exchangeable ERGM (14) and define an ERGM by the expression

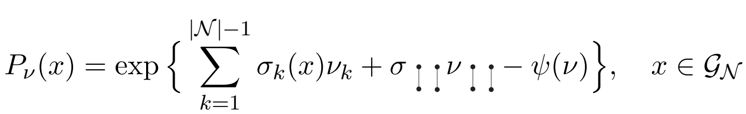

(27)

where $\sigma_k(x)$ is the number of k-stars in x, $1 \le k \le |\mathcal{N}| - 1$, and each of the $|\mathcal{N}|$ parameters ν 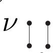 $, \nu_1, \ldots, \nu_{|\mathcal{N}|-1}$ vary independently on the real line.

We shall refer to this ERGM as the *SE*-model*, reflecting that all distributions within the model are both summarized and exchangeable, the star indicating that the model is not containing all such distributions. This is because neither the number of complete sub-graphs on $|\mathcal{N}|-1$ nor the number of complete sub-graphs on $|\mathcal{N}|-1$ nodes with one edge removed are sufficient statistics for the SE* model, so that, in virtue of Lemma 2 and Lemma 3, there exist summarized exchangeable network distributions that are not in the SE* model. The distributions implied by the marginal beta model, discussed below, are such an example. We reserve the term *SE-model* for the set of all distributions that are exchangeable and summarized and *DE-model* for the set of all distributions that are exchangeable and dissociated.

We note the similarities between the SE*-model (27) and the exchangeable Markov model (EM-model) of Frank and Strauss (1986), exponential family of the form

$$P_\nu(x) = \exp\left\{ \sum_{k=1}^{|\mathcal{N}|-1} \sigma_k(x)\nu_k + \sigma_{\triangle}\,\nu_{\triangle} - \psi(\nu) \right\}, \quad x \in \mathcal{G}_\mathcal{N}, \nu \in \mathbb{R}^{|\mathcal{N}|},$$

which differs from SE*-model only by $\sigma_{\triangle}\,\nu_{\triangle}$ replacing $\sigma_{\mid\,\mid}\,\nu_{\mid\,\mid}$. Though nearly identical, the justifications behind the EM and the SE* models are quite different. The EM model results from postulating certain symmetric Markov properties for the edges, covered by case (b) in Proposition 3. On the other hand, the SE* model arises as a convenient sub-model of the class of summarized and exchangeable network distributions. The intersection of the SE*-model and the EM-model only contains the counts for k-stars in the exponent, and so does the intersection of the SE-model and the EM model. The resulting *SEM-model* is the ERGM with sufficient statistics given by the entries degree distribution, studied in Sadeghi and Rinaldo (2014). This is because there is a one-to-one linear transformation between the number of k-stars of x and the degree distribution of x, as demonstrated in the proof of Lemma 2.

From the computational standpoint, parameter estimation in the SE* model and in the EM model can be carried out by the same procedures, such as pseudo-maximum likelihood estimation and MCMCMLE methods – see, e.g., Strauss and Ikeda (1990), Geyer and Thompson (1992) and Hunter et al. (2008) – and is likely to be just as problematic.

5.3 The marginal beta model

Here we consider *the marginal beta model,* an example of a family of network distributions that are dissociated, exchangeable, and summarized, and consequently satisfy all the conditions discussed above.

Consider the beta model defined in (7). Suppose that $(B_i, i \in \mathcal{N})$ are independent and identically distributed real-valued random variables with distribution F_B. Then the expression (7) can be understood as the conditional distribution of X_{ij} given B_i and B_j and thus the density satisfies the factorization $f(y) = \prod_i f(y_i \mid \mathrm{pa}(y_i))$ for the directed acyclic graph (DAG) D_n that consists of $(B_i, i \in \mathcal{N})$ and $(X_{ij}, ij \in D(\mathcal{N}))$ as vertices and arrows from each B_i to all the X_{ij}, $j \in \mathcal{N}$.

Figure 7 illustrates the case for $\mathcal{N} = \{1, 2, 3, 4\}$. Indeed the exact same graph also

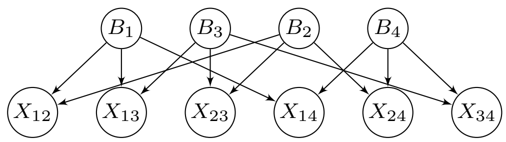

Figure 7: Directed acyclic graph D_4.

illustrates the Markov structure for a general ϕ-matrix, where just B_i should be replaced by uniformly distributed U_i and the arrows represent the ϕs.

By marginalization over and conditioning on the variables in the DAG model the resulting dependence structure can in general be captured by *mixed graphs*; for further details see, for example, Richardson and Spirtes (2002) and Sadeghi (2013). For example, by conditioning on the B_i we obtain the empty graph with all the X_{ij} as vertices. This implies the complete independence of X_{ij}, corresponding to the dyadic independence in the beta model (7). By marginalizing over B_i, the bidirected incidence graph appears, as studied extensively in this paper and exemplified in Figure 4. As discussed before, this implies that after marginalization over B_i, the network is dissociated. We refer to this model as the *marginal beta model*; the dissociatedness also follows immediately from the following "de Finetti type" representation for the distribution of X:

$$\mathbb{P}(X = x) = \int_{\mathbb{R}^n} \prod_{ij \in D(\mathcal{N})} \frac{e^{(\beta_i + \beta_j)x_{ij}}}{1 + e^{(\beta_i + \beta_j)}} dF_B^{\mathcal{N}}(\beta), \quad x \in \mathcal{G}_{\mathcal{N}}, \tag{28}$$

where $F_B^{\mathcal{N}}$ denotes the $|\mathcal{N}|$-fold product distribution from F_B. The function ϕ in Proposition 1 is in fact p_{ij} as described in (7). It clearly follows from (28) that the marginal beta distribution is exchangeable; in addition, (8) implies that the marginal beta model is summarized.

6 Consistent systems of network models

In this section we shall consider systems of network models and their relation to each other. We let $\mathcal{N}$ be the basic set of network nodes and consider models for finite subsets $\mathcal{N}' \subseteq \mathcal{N}$. The basic set $\mathcal{N}$ may be finite or countably infinite. If $P_{\mathcal{N}'}$ is a distribution of a random network $X_{\mathcal{N}'}$ with nodeset $\mathcal{N}'$ and $\mathcal{N}'' \subseteq \mathcal{N}'$ we let $\Pi^{\mathcal{N}''}_{\mathcal{N}'} P_{\mathcal{N}'}$ denote the distribution of the induced subnetwork $X_{\mathcal{N}''}$ under $P_{\mathcal{N}'}$. Thus the marginalization operator $\Pi^{\mathcal{N}'}_{\mathcal{N}''}$ takes as input a probability distribution on $\mathcal{G}_{\mathcal{N}'}$ and returns the induced or marginal probability distribution over $\mathcal{G}_{\mathcal{N}''}$ corresponding to the dyads with endpoints in $\mathcal{N}''$. By constructions, marginalization operators satisfy the property

$$\Pi^{\mathcal{N}'}_{\mathcal{N}'''} = \Pi^{\mathcal{N}''}_{\mathcal{N}'''} \circ \Pi^{\mathcal{N}'}_{\mathcal{N}''}.$$

for any $\mathcal{N}''' \subset \mathcal{N}'' \subset \mathcal{N}'$, so they form a *projective system*. A collection $\{\mathcal{P}_{\mathcal{N}'}, \mathcal{N}' \subseteq \mathcal{N}\}$ of random network models for all finite subsets $\mathcal{N}'$ of $\mathcal{N}$ is said to be (strongly) *consistent* if for every $\mathcal{N}'' \subseteq \mathcal{N}'$ in the collection it holds that

$$\mathcal{P}_{\mathcal{N}''} = \Pi^{\mathcal{N}'}_{\mathcal{N}''}(\mathcal{P}_{\mathcal{N}'}),$$

in other words it holds that the marginal to a subnetwork of any network model is identical to the corresponding network model for the subnetwork. Thus a strongly consistent system forms a *projective statistical field* in the sense of Lauritzen (1988). We say the system is *weakly consistent* if

$$\Pi^{\mathcal{N}'}_{\mathcal{N}''}(\mathcal{P}_{\mathcal{N}'}) \subseteq \mathcal{P}_{\mathcal{N}''}$$

i.e. marginalization of any network model for a larger nodeset is not in conflict with the network model for the subnetwork.

We first note that the bidirected Markov property itself is always *marginalizable* in the sense that if X is bidirected Markov w.r.t. a graph G, then the marginal distribution of X_A is automatically Markov w.r.t. the induced subgraph G_A, whereas the similar statement is generally not true for the undirected Markov property.

This holds in particular for the random network models when $G(\mathcal{N}') = L_{\leftrightarrow}(\mathcal{N}')$ is the bidirected incidence graph. In other words, *the system of dissociated random network models is weakly consistent.* The same holds for the system of exchangeable random networks $\{\mathcal{E}_{\mathcal{N}'}, \mathcal{N}' \subseteq \mathcal{N}\}$, and when both exchangeability and dissociatedness are combined to form the system of exchangeable and dissociated random networks $\{\mathcal{E}^{L_{\leftrightarrow}}_{\mathcal{N}'}, \mathcal{N}' \subseteq \mathcal{N}\}$, or when considering the system of models $\{\mathcal{E}^{L^c_{\leftrightarrow}}_{\mathcal{N}'}, \mathcal{N}' \subseteq \mathcal{N}\}$, which are all Markov w.r.t. the bidirected complements of the line graphs.

Correspondingly, the Möbius parametrization is *marginalizable* in the sense that if X is a random network with nodeset $\mathcal{N}$ and $\mathcal{N}' \subseteq \mathcal{N}$, the Möbius parameters $z^{\mathcal{N}'}_B$ for the induced subnetwork are identical to the corresponding parameters $z^{\mathcal{N}}_B$ for X:

$$z^{\mathcal{N}'}_B = z^{\mathcal{N}}_B = z_B, \text{ for all } B \in \mathcal{B}(\mathcal{N}'). \tag{29}$$

From any weakly consistent system $\{\mathcal{P}_{\mathcal{N}'}, \mathcal{N}' \subseteq \mathcal{N}\}$ we can construct a strongly consistent system $\{\mathcal{P}^*_{\mathcal{N}'}, \mathcal{N}' \subseteq \mathcal{N}\}$, by letting

$$\mathcal{P}^*_{\mathcal{N}'} = \bigcap_{\mathcal{N}'':\mathcal{N}''\supset\mathcal{N}'} \Pi^{\mathcal{N}''}_{\mathcal{N}'}(\mathcal{P}_{\mathcal{N}''})$$

provided that the intersection is non-empty. Thus each family in the system consists exactly of those distributions that can be obtained as marginals of distributions from models on arbitrarily large nodesets. We shall thus say that the elements in the strongly consistent variant of a weakly consistent system are *completely extendable.*

We note that, for example, the system of exchangeable random networks is not strongly consistent. In Example 2, the estimated distributions are easily seen not to be extendable to a network with more than four nodes. In general it is a difficult issue to determine whether a given system of Möbius parameters $\{z_B^{\mathcal{N}'}, B \in \mathcal{B}(\mathcal{N}')\}$ is extendable to a larger network with nodeset $\mathcal{N}$, i.e. if there exists a system of Möbius parameters $\{z_B^{\mathcal{N}}, B \in \mathcal{B}(\mathcal{N})\}$, such that (29) holds.

Nevertheless, we can characterize the strongly consistent sequences of exchangeable random networks as follows. Below, if $\mathcal{N}$ is infinite, we let $\mathcal{B}(\mathcal{N})$ be the set of all finite subsets of $D(\mathcal{N})$, which, as before, we represent as finite edge-induced subgraphs of the complete graph on $\mathcal{N}$.

Proposition 6. *All strongly consistent and infinitely extendable systems of exchangeable Möbius parameters $(z_B, B \in \mathcal{B}(\mathcal{N}))$ are those which satisfy the conditions*

(a) $z_\varnothing = 1$*;*

(b) $z_B = z_{B'} = z_U$ *if* B *and* B' *are in* $[U]$*;*

(c) *for all non-negative integers* n *and all* $\mathcal{N}' \subset \mathcal{N}$ *with* $|\mathcal{N}'| = n$

$$\sum_{U \in \mathcal{U}_{\mathcal{N}'}} (-1)^{|U|-|x|} r_U(x) z_U \geq 0, \quad \forall x \in \mathcal{G}_{\mathcal{N}'},$$

where we recall that $r_U(x)$ *is the number of graphs in* $[U]$ *that contain* x *as a subgraph. In addition, the sequences of Möbius parameters corresponding to the extremal distribution* $\mathcal{E}_\infty$ *are those satisfying the dissociated property that if* B *has connected components* $C_1, \ldots, C_k$ *then* $z_B = z_{C_1} \cdots z_{C_k}$.

Proof. From (29) it follows that any system of valid Möbius parameters is consistent and as previously argued, these are valid if and only if they satisfy $z_\varnothing = 1$ and the non-negativity restriction in equation (13) above. We have also previously argued that the additional product condition characterizes dissociated networks which are the extreme points according to de Finetti's theorem (Proposition 1). □

□

It is worth pointing out that, when $|\mathcal{N}| = \infty$, if X is a random network on $\mathcal{N}$ whose distribution is extremal (and therefore dissociated), then, for each non-empty $U \in \mathcal{U}_{\mathcal{N}'}$

with $|\mathcal{N}'| < \infty$,

$$z_U = \int_{[0,1]^{\mathcal{N}'}} \prod_{ij \in E(U)} \phi(u_i, u_j) du, \tag{30}$$

where $E(U)$ denotes the set of edges in U and ϕ is some measurable, symmetric function from $[0,1]^2$ into $[0,1]$.

Using the arguments in Chapter 14 of Aldous (1985) (see also Kallenberg, 2005, for a generalization), de Finetti's theorem (Proposition 1) and Proposition 6 yields the following result.

Corollary 2. *The convex set of feasible Möbius parameters are those which can be represented as*

$$z_B = \int_{[0,1]} \left\{ \int_{[0,1]^n} \prod_{ij \in B} \phi(u_i, u_j, \lambda)\, du \right\} d\lambda, \quad B \in \mathcal{B}(\mathcal{N}), \tag{31}$$

where $n = |B|$ and ϕ is a (not necessarily unique) measurable function from $[0,1]^3$ to $[0,1]$. The extreme points of this convex set correspond to the cases in which the function $\phi(u_i, u_j, \lambda)$ is a constant function of λ, for almost all $\lambda \in [0,1]$.

Still, these characterizations are quite implicit. It seems difficult to obtain more explicit characterizations. For example, Figure 16.1 on page 288 of Lovász (2012) displays the area of variation of $(z_{\mid}, z_{\triangle})$ as given in (30) for all exchangeable and dissociated infnite networks, and this is a very complicated region. The set of infinitely extendable pairs is given by the convex hull of the region (Bollobás, 1976).

Using Corollary 2 we can show that strongly consistent systems of exchangeable models which are Markov w.r.t. bidirected complement of the incidence graph must be Erdös–Renyi models. First, we need a preliminary lemma. The lemma and proof are analogues of Theorem 10 in Diaconis and Freedman (1981).

Lemma 4. *Assume X is an exchangeable random network with an infinitely extendable set of Möbius parameters satisfying, for some $\eta \in (0,1)$,*

$$z_{\mid} = \eta, \quad z_{\llcorner} = \eta^2, \quad z_{\square} = \eta^4.$$

Then, the edges in X are independent and identically distributed with $P(X_{ij} = 1) = \eta$, i.e. X is Erdös–Renyi.

Proof. Since all elements are exchangeable, we have from (31) that

$$\eta = z_{\mid} = \int_{[0,1]^3} \phi(u, v, \lambda)\, du\, dv\, d\lambda.$$

But also

$$z_{\llcorner} = \eta^2 = \int_{[0,1]^4} \phi(u,v,\lambda)\phi(u',v,\lambda)\,du\,du'\,dv\,d\lambda = \int_{[0,1]^2} \left\{ \int_{[0,1]} \phi(u,v,\lambda)\,du \right\}^2 dv\,d\lambda.$$

Thus,

$$\int_{[0,1]^2} \left\{ \int_{[0,1]} \phi(u,v,\lambda)\,du - \eta \right\}^2 dv\,d\lambda = 0,$$

and hence

$$\int \phi(u,v,\lambda)\,du = \eta \quad \text{for almost all } (v,\lambda),$$

which in turn implies that, if we assume $\mathcal{N} = \{1,2,3,\ldots\}$,

$$z_{\vert\,\vert} = P(X_{12}=1, X_{34}=1) = \int_{[0,1]^5} \phi(u,v,\lambda)\phi(u',v',\lambda)\,du\,du'\,dv\;dv'\,d\lambda = \eta^2.$$

Next, we have that

$$\begin{aligned} \eta^4 = z_{\square} &= \int_{[0,1]^5} \phi(u,v,\lambda)\phi(u',v,\lambda)\phi(u,v',\lambda)\phi(u',v',\lambda)\,du\,du'\,dv\,dv'\,d\lambda \\ &= \int_{[0,1]^3} \left\{ \int_{[0,1]} \phi(u,v,\lambda)\phi(u',v,\lambda)\,dv \right\}^2 du\,du'\,d\lambda, \end{aligned}$$

and hence we conclude as above that, for almost all (u,u',λ),

$$\int_{[0,1]} \phi(u,v,\lambda)\phi(u',v,\lambda)\,dv = \eta^2. \tag{32}$$

We would now wish to let $u = u'$ in the above equation but this might just be the exceptional set. However, we may extend ϕ periodically to be defined on $(-\infty,\infty) \times [0,1]^2$ by letting

$$\phi(u,v,\lambda) = \phi(u \bmod 1, v, \lambda)$$

so that (32) still holds almost surely. Then, using Lebesgue's theorem, we can write

$$\begin{aligned} \int_{[0,1]} \phi(u,v,\lambda)^2\,dv &= \int_{[0,1]} \lim_{\epsilon\to 0} \frac{1}{4\epsilon^2} \int_{-\epsilon}^{\epsilon} \int_{-\epsilon}^{\epsilon} \phi(u+s,v,\lambda)\phi(u+t,v,\lambda)\,ds\,dt\,dv \\ &= \lim_{\epsilon\to 0} \frac{1}{4\epsilon^2} \int_{-\epsilon}^{\epsilon} \int_{-\epsilon}^{\epsilon} \int_{[0,1]} \phi(u+s,v,\lambda)\phi(u+t,v,\lambda)\,dv\,ds\,dt = \eta^2. \end{aligned}$$

As above, we conclude, for almost all λ, u, v, that $\phi(u,v,\lambda) = \eta$ and hence we have $z_D = \eta^{|D|}$, implying that all edges are independent. Hence, the model is Erdös–Renyi. $\square$

$\square$

Proposition 7. *Consider a countable nodeset $\mathcal{N}$ and a strongly consistent family $\{\mathcal{P}_{\mathcal{N}'}, \mathcal{N}' \subseteq \mathcal{N}\}$ of random network models; if all elements of $\mathcal{P}_{\mathcal{N}'}$ are exchangeable and Markov w.r.t. the complement of the bidirected line graph $L^c_{\leftrightarrow}(\mathcal{N}')$, then they are all Erdös–Renyi models.*

Proof. We assume that $\mathcal{N} = \{1, 2, 3, \ldots\}$. By exchangeability, we have that, for some $\eta \in [0, 1]$, z  $= \mathbb{P}(X_{12} = 1) = \eta$. Since $X_{12} \perp\!\!\!\perp X_{23}$, we also have that

z  $= \mathbb{P}(X_{12} = X_{23} = 1) = \mathbb{P}(X_{12} = 1)\mathbb{P}(X_{23} = 1) = \eta^2 = z$ ,

where the last equality is obtained as in the proof of Lemma 4. The bidirected complement of the incidence graph for $\mathcal{N}' = \{1, 2, 3, 4\}$ contains the edges

$$\{12 \leftrightarrow 34, 13 \leftrightarrow 24, 14 \leftrightarrow 23\}.$$

Since all the distributions are Markov w.r.t. this graph we have

z  $= P(X_{12} = 1, X_{34} = 1)P(X_{14} = 1, X_{23} = 1) = (z$ $)^2 = \eta^4.$

and the conclusion now follows from Lemma 4. □

□

Analogous results hold true if the Markov assumption in Proposition 7 is replaced with Markov properties with respect to the undirected incidence graph as in the Frank–Strauss models or the undirected complement of the line graph, as in the Kneser models. The proof of these facts are quite elementary and we abstain from giving the details. In fact these sequences are not even weakly consistent, as marginals of undirected Markov distributions are not Markov with respect to the corresponding induced subgraph.

Finally we hasten to point out that indeed the *system of marginal beta models is strongly consistent* by construction and that this clearly remains true if the mixing distribution F is assumed to belong to any specific subset of distributions $\mathcal{F}$.

7 Discussion

We have derived a complete characterization of possible Markov properties of exchangeable network models, and studied some of the implied models. We also consider summarized exchangeable network models in which the probability of a network is only a function of its degree distributions, and relate them to the other models we discuss in the paper. Overall, our findings have unveiled several interesting properties of exchangeable network models and established connections among seemingly disparate concepts used in network analysis, the probability literature on exchangeable arrays and the discrete mathematics literature on graph limits.

While we have focused on exchangeable and extendable network models, we point out that these are by no means the only interesting network models. In fact, many other parametric models, not fulfilling those properties, can also be successfully deployed to represent and study networks. Examples include the beta model (Chatterjee et al., 2011; Yan and Xu, 2013; Rinaldo et al., 2013) and the sparse Bernoulli model considered in Krivitsky and Kolaczyk (2015); see also the remarks in the introduction on partial conditional independence and on weakened forms of exchangeability.

We conclude with a final remark on the computational difficulties associated with fitting the various models discussed in the paper. As it is very often the case with ERGMs, model fitting can be rather challenging, and our models are no exception. For the DE and DSE models, one could in principle rely on the algorithms proposed in Drton and Richardson (2008), appropriately modified to allow for zero counts and for the fact that the MLE of the model parameters is nearly all cases on the boundary of the parameter space. However such algorithms may not scale well with the network size. As for the model given in (14), and its submodels, the SE and SE* and the EM models, parameter estimation can be carried out using pseudo-maximum likelihood and MCMC-based methods, (see, e.g. Strauss and Ikeda, 1990; Geyer and Thompson, 1992), implemented, for instance, in the `R` package `ergm` (Hunter et al., 2008). In all these cases, however, we expect that fitting could be problematic, as these procedures do not scale well with the network size and MCMC convergence can be very slow or hard to assess. In particular, we expect that the issues of degeneracy (see Schweinberger, 2011; Rinaldo et al., 2009) plaguing many ERGMs will also impact estimation of these models.

Acknowledgements

The authors are grateful to Éva Czabarka for proving an earlier version of Lemma 3. Alessandro Rinaldo and Kayvan Sadeghi were partially supported by AFOSR grant FA9550-14-1-0141.

References

Aldous, D. (1981). Representations for partially exchangeable random variables. *Journal of Multivariate Analysis 11*, 581–598.

Aldous, D. (1985). Exchangeability and related topics. In P. Hennequin (Ed.), *École d'Été de Probabilités de Saint–Flour XIII — 1983*, pp. 1–198. Heidelberg: Springer-Verlag. Lecture Notes in Mathematics 1117.

Baddeley, A. and J. Møller (1989). Nearest-neighbour Markov point processes and random sets. *International Statistical Review 57*, 89–121.

Barndorff-Nielsen, O. E. (1978). *Information and Exponential Families in Statistical Theory.* Chichester, UK: John Wiley and Sons.

Bollobás, B. (1976). Relations between sets of complete subgraphs. In C. S. J. A. Nash-Williams and J. Sheehan (Eds.), *Proceedings of the 5th British Combinatorics Conference.*

Boutilier, C., N. Friedman, M. Goldszmidt, and D. Koller (1996). Context-specific independence in Bayesian networks. In *Proceedings of the Twelfth International Conference on Uncertainty in Artificial Intelligence*, UAI'96, San Francisco, CA, USA, pp. 115–123. Morgan Kaufmann Publishers Inc.

Brandstadt, A., V. B. Le, and J. Spinrad (1999). *Graph Classes: A Survey.* SIAM Monographs on Discrete Mathematics and Applications. Philadelphia: Society of Industrial and Applied Mathematics.

Caron, F. and E. Fox (2017). Sparse graphs using exchangeable random measures. *Journal of Royal Statistical Society, Series B 79*, 1295–1366.

Chatterjee, S., P. Diaconis, and A. Sly (2011). Random graphs with a given degree sequence. *Annals of Applied Probability 21*, 1400–1435.

Cox, D. R. and N. Wermuth (1993). Linear dependencies represented by chain graphs (with discussion). *Statistical Science 8*, 204–218; 247–277.

Crane, H. and W. Dempsey (2015). A framework for statistical network modeling. Available at https://arxiv.org/abs/1509.08185.

Diaconis, P. (1977). Finite forms of de Finetti's theorem on exchangeability. *Synthese 36*, 271–281.

Diaconis, P. and D. Freedman (1980). Finite exchangeable sequences. *Annals of Probability 8*, 745–764.

Diaconis, P. and D. Freedman (1981). On the statistics of vision: the Julesz conjecture. *Journal of Mathematical Psychology 24*, 112–138.

Diaconis, P. and S. Janson (2008). Graph limits and exchangeable random graphs. *Rendiconti di Matematica, Serie VII 28*, 33–61.

Drton, M. and T. S. Richardson (2008). Binary models for marginal independence. *Journal of the Royal Statistical Society Series B 70*, 287–309.

Eagleson, G. K. and N. C. Weber (1978). Limit theorems for weakly exchangeable arrays. *Mathematical Proceedings of the Cambridge Philosophical Society 84*, 123–130.

Erdös, P. and A. Rényi (1960). On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences 5*, 17–61.

Fienberg, S. E. and S. Wasserman (1981). Discussion of an exponential family of probability distributions for directed graphs by Holland and Leinhardt. *Journal of the American Statistical Association 76*, 54–57.

Frank, O. (1991). Statistical analysis of change in networks. *Statistica Neerlandica 45*, 283–293.

Frank, O. and D. Strauss (1986). Markov graphs. *Journal of the American Statistical Association 81*, 832–842.

Geyer, C. J. and E. A. Thompson (1992). Constrained Monte Carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society. Series B (Methodological) 54*, 657–699.

Godsil, C. and G. Royle (2001). *Algebraic Graph Theory.* New York: Springer.

Holland, P. and S. Leinhardt (1981). An exponential family of probability distributions for directed graphs. *Journal of the American Statistical Association 76*, 33–50.

Hunter, D. R., S. M. Goodreau, and M. S. Handcock (2008). Goodness of fit of social network models. *Journal of the American Statistical Association 103*, 248–258.

Hunter, D. R., M. S. Handcock, C. T. Butts, S. M. Goodreau, Morris, and Martina (2008). `ergm`: A package to fit, simulate and diagnose exponential family models for networks. *Journal of Statistical Software 24*, 1–29.

Kallenberg, O. (2005). *Probabilistic Symmetries and Invariance Principles.* New York: Springer–Verlag.

Kauermann, G. (1996). On a dualization of graphical Gaussian models. *Scandinavian Journal of Statistics 23*, 105–116.

Kerns, G. J. and G. J. Székely (2006). DeFinettis theorem for abstract finite exchangeable sequences. *Journal of Theoretical Probability 19*, 589–608.

Kolaczyk, E. D. (2009). *Statistical Analysis of Network Data: Methods and Models.* Springer Publishing Company, Incorporated.

Konstantopoulos, T. and L. Yuan (2015). On the extendibility of finitely exchangeable probability measures. Preprint, Department of Mathematics, Uppsala University. Available at https://arxiv.org/abs/1501.06188.

Krivitsky, P. N. and E. D. Kolaczyk (2015). On the question of effective sample size in network modeling: An asymptotic inquiry. *Statistical Science 30*, 184–198.

Lauritzen, S. L. (1988). *Extremal Families and Systems of Sufficient Statistics.* Heidelberg: Springer-Verlag. Lecture Notes in Statistics 49.

Lauritzen, S. L. (1996). *Graphical Models.* Oxford, United Kingdom: Clarendon Press.

Lauritzen, S. L. (2003). Rasch models with exchangeable rows and columns. In J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith, and M. West (Eds.), *Bayesian Statistics*, Volume 7, Oxford, UK, pp. 215–232. Oxford University Press.

Lauritzen, S. L. (2008). Exchangeable Rasch matrices. *Rendiconti di Matematica, Serie VII 28*, 83–95.

Lovász, L. (2012). *Large Networks and Graph Limits*, Volume 60 of *Colloquium Publications.* American Mathematical Society.

Lovász, L. and B. Szegedy (2006). Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B 96*, 933–957.

Matúš, F. (1995). Finite partially exchangeable arrays. Technical Report 1856, Inst. of Inf. Theory and Autom., http://staff.utia.cas.cz/matus/CsiMaGSI13.pdf.

Newman, M. (2010). *Networks: An Introduction.* Oxford University Press, Inc.

Nyman, H., J. Pensar, T. Koski, and J. Corander (2014). Stratified graphical models — context-specific independence in graphical models. *Bayesian Analysis 9*, 883–908.

Orbantz, P. and D. M. Roy (2015). Bayesian models of graphs, arrays, and other exchangeable structures. *IEEE Transactions on Pattern Analysis and Machine Intelligence 37*, 437–461.

Petersen, J. (1898). Sur le théorème de Tait. *L'Intermédiaire des Mathématiciens 5*, 225–227.

Richardson, T. (2003). Markov properties for acyclic directed mixed graphs. *Scandinavian Journal of Statistics 30*, 145–157.

Richardson, T. S. and P. Spirtes (2002). Ancestral graph Markov models. *Annals of Statistics 30*, 962–1030.

Rinaldo, A., S. E. Fienberg, and Y. Zhou (2009). On the geometry of discrete exponential families with application to exponential random graph models. *Electr. J. Statist. 3*, 446–484.

Rinaldo, A., S. Petrović, and S. E. Fienberg (2013). Maximum lilkelihood estimation in the beta model. *Annals of Statistics 41*, 1085–1110.

Roverato, A., M. Lupparelli, and L. La Rocca (2013). Log-mean linear models for binary data. *Biometrika 100*, 485–494.

Sadeghi, K. (2013). Stable mixed graphs. *Bernoulli 19*, 2330–2358.

Sadeghi, K. and A. Rinaldo (2014). Statistical models for degree distributions of networks. In *NIPS Workshop: From Graphs to Rich Data.* arXiv:1411.3825.

Schweinberger, M. (2011). Instability, sensitivity, and degeneracy of discrete exponential families. *Journal of the American Statistical Association 106*, 1361–1370.

Schweinberger, M. and M. S. Handcock (2015). Local dependence in random graph models: characterization, properties and statistical inference. *Journal of Royal Statistical Society, Series B 77*, 647–676.

Shalizi, C. and A. Rinaldo (2013). Consistency under sampling of exponential random graph models. *Annals of Statistics 41*, 508–535.

Silverman, B. W. (1976). Limit theorems for dissociated random variables. *Advances in Applied Probability 8*, 806–819.

Snijders, T. A. B. (2010). Conditional marginalization for exponential random graph models. *The Journal of Mathematical Sociology 34*, 239–252.

Snijders, T. A. B., P. E. Pattison, G. L. Robins, and M. S. Handcock (2006). New specifications for exponential random graph models. *Sociological Methodology 36*, 99–153.

Strauss, D. and M. Ikeda (1990). Pseudolikelihood estimation for social networks. *Journal of the American Statistical Association 85*, 204–212.

Wasserman, S. and P. Pattison (1996). Logit models and logistic regressions for social networks: I. an introduction to Markov graphs and p^*. *Psychometrika 61*, 401–425.

West, D. B. (2001). *Introduction to Graph Theory* (2nd ed.). Upper Saddle River, NJ, USA: Prentice Hall.

Wilson, R. J. (1996). *Introduction to Graph Theory* (4th ed.). Essex, UK: Longman.

Yan, T. and J. Xu (2013). A central limit theorem in the β-model for undirected random graphs with a diverging number of vertices. *Biometrika 100*, 519–524.