

On Exchangeability in Network Models

Steffen L. Lauritzen*
 Department of Mathematical Sciences
 University of Copenhagen

Alessandro Rinaldo†
 Department of Statistics
 Carnegie Mellon University

Kayvan Sadeghi ‡
 Statistical Laboratory
 University of Cambridge

Dedicated to the memory of Steve Fienberg

Abstract

We derive representation theorems for exchangeable distributions on finite and infinite graphs using elementary arguments based on geometric and graph-theoretic concepts. Our results elucidate some of the key differences, and their implications, between statistical network models that are finitely exchangeable and models that define a consistent sequence of probability distributions on graphs of increasing size.

Keywords: deFinetti's theorem; graphons; Möbius simplex; finite exchangeability; positive semidefinite functions.

1 Introduction

Exchangeability is one of most basic forms of probabilistic invariance. When applied to probability distributions on graphs, it is equivalent to requiring that isomorphic graphs have the same probabilities. Indeed, exchangeability provides the probabilistic underpinning to the theory of dense graph limits developed recently in the graph-theoretic literature (see, e.g., [Diaconis and Janson, 2008](#); [Borgs et al., 2008](#); [Lovász and Szegedy, 2006](#)).

In statistical network modeling, exchangeability is a very common simplifying assumption. However, it is typically only enforced for models on graphs of a given size, and not simultaneously over sequences of models on graphs of increasing size. This practice is born out of convenience: it is much easier to formulate probability distributions on finite as opposed to infinite graphs. However, the consequences of relying on this weaker assumption of finite exchangeability can be detrimental to the validity and generalizability of any statistical analysis: the properties of probability distributions on graphs of different sizes that are finitely exchangeable need not be related to each other in any meaningful way (or in any way at all); see [Shalizi and Rinaldo \(2013\)](#); [Crane and Dempsey \(2015\)](#); [Snijders \(2010\)](#).

*Email: lauritzen@math.ku.dk

†Email: arinaldo@cmu.edu

‡Email: k.sadeghi@statslab.cam.ac.uk

In this article, we investigate the relationship between exchangeability of random finite graphs and exchangeability of random infinite graphs using a combination of simple geometric arguments and standard graph-theoretic concepts. Our work can be seen as an extension to the graph setting of the geometric representation of finite exchangeability for random binary sequences developed by [Diaconis \(1977\)](#). We make the following contributions: (1) we formulate a finite deFinetti's theorem for random graphs that is both elementary and rely on well known graph-theoretic quantities (namely, density homomorphisms) only; (2) we extend this result to obtain a simple derivation of the well-known deFinetti's representation theorem for exchangeable distributions on (infinite) graphs; (3) we provide novel geometric characterizations of all the finite marginals of exchangeable distributions on finite graphs and discuss the implications of our findings.

Related Work. There is a vast literature on exchangeability of random arrays, of which random graphs are a special case; see, e.g., [Aldous \(1981, 1985\)](#), [Eagleson and Weber \(1978\)](#), [Hoover \(1979\)](#), [Kerns and Székely \(2006\)](#), [Lauritzen \(2008\)](#), [Kallenberg \(2005\)](#) and [Silverman \(1976\)](#), to name a few. Of particular significance is [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 [Borgs et al. \(2008\)](#) and [Lovász and Szegedy \(2006\)](#) (see also the book [Lovász \(2012\)](#)). Similarly, finite exchangeability for sequences and arrays has been thoroughly investigated: see [Diaconis \(1977\)](#), [Diaconis and Freedman \(1981\)](#), [Aldous \(1981\)](#) and, in particular, [Matúš \(1995\)](#); see also [Volfovsky and Airoidi \(2016\)](#).

In the companion paper [Lauritzen et al. \(2017\)](#), we rely on the tools and formalism of graphical modeling to study the Markov properties of finitely exchangeable network models. The results derived there complement the ones we obtain in the present paper. We will discuss the connection between the two papers later in Section 5.

The article is organized in the following way. Section 2 describes the geometry of finitely exchangeable distributions on finite graphs and exchangeable distributions on infinite graphs and introduces the Möbius parametrization, which we will use throughout to represent probabilities on graphs. In Section 3 we provide definitions and basic results for homomorphism and isomorphism densities in order to derive a deFinetti theorem for finitely exchangeable probability distributions on graphs based on the Möbius parametrization in Section 4. In Section 5 we study the manifold of dissociated and exchangeable random graphs and show that there exists dissociated and exchangeable random graphs that are not infinitely extendable.

Notation. For any integer $n \geq 2$ we denote with \mathcal{L}_n and \mathcal{U}_n the set of simple labeled graphs and simple unlabeled graphs with node set $[n] := \{1, \dots, n\}$, respectively, and set $\mathcal{L} = \bigcup_{n=2}^{\infty} \mathcal{L}_n$ and $\mathcal{U} = \bigcup_{n=2}^{\infty} \mathcal{U}_n$. We let \mathcal{L}_{∞} be the set of infinite simple labeled graphs. For any two graphs G and G' in \mathcal{L} , we will write $G \sim G'$ to signify that they are isomorphic and $[G]$ for the equivalence class of all graphs isomorphic to G . With a slight abuse of notation, we will at times identify the class $[G]$ with the undirected graph representing it. We will also identify \mathcal{L}_n with the Boolean algebra of all subsets of the node pairs $\{\{i, j\}, i \neq j\}$ partially ordered by inclusion by identifying each graph in \mathcal{L}_n with the binary vector $\{0, 1\}^{\binom{n}{2}}$ representing its edges. With this identification, \mathcal{L}_n indexes the coordinates of vectors in $\mathbb{R}^{2^{\binom{n}{2}}}$. If G and H are in \mathcal{L} , we write $H \subseteq G$ if H is a sub-graph (not necessarily induced) of G . For integers $2 \leq m \leq n \leq \infty$ and a $G \in \mathcal{L}_n$, $G[m]$ is the sub-graph of G induced by the nodes $[m]$.

In our analysis, we will often identify a graph G with its set of edges, hence ignoring isolated nodes. The set of labeled graphs on $[n]$ without isolated nodes is denoted by \mathcal{I}_n and we let $\mathcal{I} = \bigcup_{n=2}^{\infty} \mathcal{I}_n$. Similarly, we let \mathcal{J}_n denote the set of unlabeled graphs without isolated nodes and $\mathcal{J} = \bigcup_{n=2}^{\infty} \mathcal{J}_n$.

For any $n \geq 2$, if $G \in \mathcal{L}_n$ and $\sigma \in \mathcal{S}_n$, the permutation group on $[n]$, we will let G_{σ} be the graph obtained from G by relabeling its nodes according to σ . Thus nodes i and j are connected in G if and only if $\sigma(i)$ and $\sigma(j)$ are connected in G_{σ} . Finally, we let $\mathcal{S} = \bigcup_n \mathcal{S}_n$ be the set of all finite permutations.

2 Probabilities on graphs: exchangeability and geometry

We begin by introducing the notions of exchangeable distributions on networks and illustrating their geometry properties.

The set of all probability distributions on \mathcal{L}_n can be represented geometrically as the probability simplex in $\mathbb{R}^{\mathcal{L}_n}$, denoted with Δ_n . The coordinates of each vector $p \in \Delta_n$ are indexed by the graphs in \mathcal{L}_n , and the $2^{\binom{n}{2}}$ vertices of Δ_n are the unit masses at each $G \in \mathcal{L}_n$.

Definition 1. A probability distribution p on \mathcal{L}_n is m -**exchangeable** when $p(G[m]) = p(G[m]_\sigma)$ for all $\sigma \in \mathcal{S}_m$. Equivalently, p is m -exchangeable when $p(G[m]) = p(G[m]')$ if $G[m] \sim G[m]'$ in \mathcal{L}_m . If n is finite and $m = n$ we say that p is (finitely) **exchangeable**.

We denote with $\mathcal{E}_n \subset \Delta_n$ the set of all exchangeable distributions on \mathcal{L}_n . It is easy to show that exchangeable distributions are mixtures of uniform distributions over isomorphic classes. In fact, \mathcal{E}_n is affinely isomorphic to the probability simplex in $\mathbb{R}^{\mathcal{U}_n}$.

Lemma 2. \mathcal{E}_n is a simplex whose vertices are the uniform probability distributions over isomorphic classes of \mathcal{L}_n . The dimension of \mathcal{E}_n is equal to $|\mathcal{U}_n| - 1$.

Proof. For a given class $[H] \in \mathcal{U}_n$ let $p_{[H]}$ be the probability distribution on \mathcal{L}_n corresponding to the uniform distribution over $[H]$. That is, for any $G \in \mathcal{L}_n$,

$$p_{n,[H]}(G) = \begin{cases} \frac{1}{|[H]|} & \text{if } G \in [H], \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Then, $p_{[H]} \in \mathcal{E}_n$ for all $[H]$. The vectors $\{p_{[H]}, [H] \in \mathcal{U}_n\}$ are affinely independent, because they are supported on incomparable subsets of \mathcal{L}_n , regarded as a poset with respect to the subset inclusion. Thus, their convex hull is a simplex inside \mathcal{E}_n . We will show that this simplex in fact coincides with \mathcal{E}_n . Let p be any point in \mathcal{E}_n . By exchangeability, for any $[H] \in \mathcal{U}_n$, the value of p at each of the coordinates indexed by the graphs in the isomorphism class $[H]$ is the same. Thus,

$$p = \sum_{[H] \in \mathcal{U}_n} v_{[H]} p_{[H]},$$

for some sequence $\{v_{[H]}, [H] \in \mathcal{U}_n\}$ of non-negative numbers. Since $\sum_{\{G \in \mathcal{L}_n\}} p(G) = 1$, it follows that $\sum_{[G]} v_{[G]} = 1$ and, therefore, that p is in the convex hull of $\{p_{[H]}, [H] \in \mathcal{U}_n\}$. Furthermore, since the convex hull of the vectors $\{p_{[H]}, [H] \in \mathcal{U}_n\}$ is a simplex, the sequence $\{v_{[H]}, [H] \in \mathcal{U}_n\}$ is unique. ■

Consistent Models and Exchangeability

A highly desirable property of a probability distribution for network data of a given size, say m , is that the distribution be realized as the marginal of probability distributions over networks of larger sizes n for all $n > m$. We refer to this property as probabilistic consistency. To formalize this point, we introduce the notion of **marginal mapping**, which, for any pair of integers $2 \leq m \leq n$, is defined to be the function $\Pi_n^m: \Delta_n \rightarrow \Delta_m$ mapping any probability distribution $p_n \in \Delta_n$ into the probability distribution $\Pi_n^m(p_n) = p_n^m \in \Delta_m$ given by

$$p_n^m(H) = \sum_{\{G \in \mathcal{L}_n: H=G[m]\}} p_n(G), \quad H \in \mathcal{L}_m. \quad (2)$$

The definition of marginal map can be extended in a straightforward manner to be defined also over distributions on \mathcal{L}_∞ : for any integer $m \geq 2$ and any probability distribution p_∞ over \mathcal{L}_∞ , Π_∞^m takes p_∞ into the distribution $\Pi_\infty^m(p_\infty) = p_\infty^m$ in Δ_m given by

$$p_\infty^m(H) = \sum_{\{G \in \mathcal{L}_\infty : H=G(m)\}} p_\infty(G), \quad H \in \mathcal{L}_m.$$

Definition 3. A sequence $\{p_n\}_{n=2}^\infty$ of probability distributions such that $p_n \in \Delta_n$ for all n is **consistent** if

$$p_m = \Pi_n^m p_n, \quad \forall 2 \leq m \leq n. \quad (3)$$

If a probability distribution over networks of a given size is not part of a consistent sequence, then its properties may not be related in any meaningful way to the properties of any probability distribution over networks of different sizes.

When applied to an exchangeable distribution in \mathcal{E}_n , the marginal mapping Π_n^m always yields an exchangeable distribution in \mathcal{E}_m . However, Π_n^m is not surjective: there are exchangeable distributions in \mathcal{E}_m that cannot be obtained as marginals of any exchangeable distribution on \mathcal{E}_n , for all $n > m > 3$. Before formalizing this in Lemma 5 and exemplifying it by an example, we need the following graph-theoretic result from [Akiyama et al. \(1979\)](#):

Lemma 4. *Let H be a graph with n nodes. We then have the following:*

1. *All the induced subgraphs of H with m nodes, $2 \leq m \leq n-2$, are isomorphic if and only if H is a complete or empty graph.*
2. *All the induced subgraphs of H with $n-1$ nodes are isomorphic if and only if H is a complete graph, an empty graph, or a **node-transitive graph**, that is for any two nodes v_1 and v_2 , there is some automorphism t such that $t(v_1) = v_2$.*

Lemma 5. *For all integers $4 \leq m < n_1 < n_2$, it holds that $\Pi_{n_2}^m(\mathcal{E}_{n_2}) \subsetneq \Pi_{n_1}^m(\mathcal{E}_{n_1}) \subsetneq \mathcal{E}_m$.*

Proof. The linearity of the marginal mapping Π_n^m , $n > m$, implies that $\Pi_n^m(\mathcal{E}_n)$ is a polytope whose vertices are contained in the image under Π_n^m of the vertices of \mathcal{E}_n (and no vertex of $\Pi_n^m(\mathcal{E}_n)$ can be the image of any point in the interior of \mathcal{E}_n). Thus, consider a vertex of \mathcal{E}_n , which, by Lemma 2, can be represented by an undirected graph on n nodes, say U . Such a vertex is mapped by Π_n^m into a distribution giving positive probabilities to only induced subgraphs of U of size m . This is a convex combination of uniform distributions over the labeled version of each of the induced subgraphs, which are vertices of \mathcal{E}_m . Hence

$$\Pi_n^m(\mathcal{E}_n) \subseteq \mathcal{E}_m. \quad (4)$$

Furthermore, a vertex of \mathcal{E}_n is mapped into a vertex of \mathcal{E}_m if and only if it corresponds to the uniform distribution over isomorphic graphs in \mathcal{G}_n such that all induced subgraphs obtained by removing any set of $n-m$ nodes are isomorphic. By Lemma 4, if $n-m=1$ this condition is verified by node-transitive graphs, the empty graph and the complete graph, and if $n-m > 2$ only by the empty and complete graphs. Therefore, the inclusion (4) is strict. To show that $\Pi_{n_2}^m(\mathcal{E}_{n_2}) \subsetneq \Pi_{n_1}^m(\mathcal{E}_{n_1}) \subsetneq \mathcal{E}_m$, for all integers $2 < m < n_1 < n_2$, we will need to show that $\Pi_{n_2}^m(\mathcal{E}_{n_2}) \subsetneq \Pi_{n_1}^m(\mathcal{E}_{n_1})$. We reason by contradiction. Indeed, if it were the case that $\Pi_{n_2}^m(\mathcal{E}_{n_2}) = \Pi_{n_1}^m(\mathcal{E}_{n_1})$ then all the vertices of $\Pi_{n_1}^m(\mathcal{E}_{n_1})$ would also be the vertices of $\Pi_{n_2}^m(\mathcal{E}_{n_2})$. Since $\Pi_{n_2}^m(\mathcal{E}_{n_2}) = \Pi_{n_1}^m(\Pi_{n_2}^{n_1}(\mathcal{E}_{n_2}))$ this would imply that all the vertices of $\Pi_{n_1}^m(\mathcal{E}_{n_1})$ belong to the image under $\Pi_{n_2}^{n_1}$ of the vertices of the set $\Pi_{n_2}^{n_1}(\mathcal{E}_{n_2})$. However, as we have just shown, with the exclusion of the empty graph, the complete graph and, if $n_2 - n_1 = 1$, the node transitive graphs in \mathcal{E}_{n_2} , the image under $\Pi_{n_2}^{n_1}$ of all the vertices of \mathcal{E}_{n_2} belongs to the interior of \mathcal{E}_{n_1} and would not be mapped by $\Pi_{n_1}^m$ to any of the vertices of $\Pi_{n_1}^m(\mathcal{E}_{n_1})$, yielding a contradiction. Thus, $\Pi_{n_2}^m(\mathcal{E}_{n_2}) \subsetneq \Pi_{n_1}^m(\mathcal{E}_{n_1})$. ■

Example 6. Let p_5 the uniform distribution on \mathcal{L}_5 that assigns probability $1/15$ over all graphs isomorphic to the union of a four cycle and an isolated node. Then, its image p_5^4 in \mathcal{E}_4 under the marginal mapping is the convex combinations of two vertices of \mathcal{E}_4 : the uniform distributions over the 3 graphs isomorphic to the 4-cycle and the uniform distribution over the 12 graphs isomorphic to the union of a 2-star and an isolated node. The weights of this mixtures are $3/15$ and $12/15$, respectively. On the other hand, $P_5^4(\mathcal{E}_5)$ does not contain those two vertices of \mathcal{E}_4 , verifying that $P_5^4(\mathcal{E}_5)$ is a strict subset of \mathcal{E}_4 . Furthermore, the point p_5^4 happens to be a vertex of $P_5^4(\mathcal{E}_5)$. To see this, it is enough to observe that, for each of the other 33 unlabeled graphs in \mathcal{U}_5 , the set of induced subgraphs obtained by removing any one node is different than (and is never contained in) the set consisting of the 4-cycle and the union of a 2-star and an isolated node. Thus, it is not possible to represent p_5^4 as a convex combination of the marginals of uniform distributions over isomorphic graphs on 5 nodes. Since each vertex of $P_5^4(\mathcal{E}_5)$ is the image of some vertex of \mathcal{E}_5 , the claim follows.

The previous example has led us to conjecture that, for $4 \leq m < n$, each vertex of \mathcal{E}_n is mapped into a vertex of $\Pi_n^m(\mathcal{E}_n)$. When $m = 3$ and $n = 4$ this is not true, since the uniform distribution over graphs in \mathcal{L}_n isomorphic to the 3-path is mapped into the interior of \mathcal{E}_3 . In this case, we also have that $P_4^3(\mathcal{E}_4) = \mathcal{E}_3$, which explains the requirement that $m \geq 4$ in Lemma 5.

Lemma 5 implies that a sequence of finitely exchangeable probability distributions on graphs need not be consistent. That is, if one poses a finitely exchangeable distribution p_n on \mathcal{L}_n , while all its marginals will be exchangeable, there is no guarantee *a priori* that p_n can be realized as the marginal of any exchangeable distribution on larger graphs.

In order for a sequence of finitely exchangeable distribution on graphs to be consistent, finite exchangeability needs to be replaced by the stronger notion of exchangeability, which we define next. We remark that, though our definition may appear different from the classic definition of row and column exchangeability of symmetric random binary arrays as in e.g. [Silverman \(1976\)](#); [Aldous \(1981\)](#); [Diaconis and Freedman \(1981\)](#); [Lauritzen \(2008\)](#), it is in fact equivalent.

Definition 7. A probability distribution p_∞ on \mathcal{L}_∞ is exchangeable when $p_\infty(G) = p_\infty(G_\sigma)$ for all $\sigma \in \mathcal{S}$. Equivalently, p_∞ is finitely exchangeable when any pair G and G' of isomorphic graphs in \mathcal{L} have the same marginal probabilities under p_∞ .

It follows that all the finite marginals $\{p_n\}_{n=1,2,\dots}$ of an exchangeable distribution define a consistent sequence of finitely exchangeable distributions. Conversely, by the Kolmogorov-Bochner extension theorem (see, e.g. [Rao, 1971](#)), the existence of a consistent family of finitely exchangeable distribution will guarantee the existence of an exchangeable distribution on \mathcal{L}_∞ .

We denote with \mathcal{E}_∞ the set of all exchangeable distributions on \mathcal{L}_∞ . \mathcal{E}_∞ can be identified with a compact subset of $[0, 1]^{\mathcal{L}}$ (see (26) below) and is a Bauer simplex; see Section 4.2.1 below. One of our goals in this article is to establish a relationship between finite exchangeability and exchangeability. In particular, letting \mathcal{E}_∞ denote exchangeable probabilities distributions over \mathcal{L}_∞ , we seek a geometric characterization of the subset of \mathcal{E}_m given by

$$\lim_{n \rightarrow \infty} \Pi_n^m(\mathcal{E}_n) = \bigcap_n \Pi_n^m(\mathcal{E}_n) = \Pi_\infty^m(\mathcal{E}_\infty),$$

which, in light of Lemma 5, is a well-defined closed set. We provide a partial solution in Section 5.

The Möbius parametrization

Though canonical, the parametrization corresponding to the set \mathcal{E}_n is not the most convenient. As we will see, exchangeable distributions on graphs are better represented using marginal, as opposed to joint,

probabilities. We will refer to this parametrization as the **Möbius parametrization**, which we describe next. We take note that this is not a novel parametrization: it is simply the adaptation to the network setting of the well-known representation of multivariate binary distributions by means of the Möbius inversion formula.

Let M_n be the square matrix of dimension $2^{\binom{n}{2}}$ with entries indexed by graphs in \mathcal{L}_n and given by

$$M_n(F, G) = 1(F \subseteq G), \quad F, G \in \mathcal{L}_n. \quad (5)$$

Then, M_n has full rank (see, e.g., [Stanley, 2011](#)) and its inverse has entries

$$M_n^{-1}(F, G) = (-1)^{|G \setminus F|} 1(F \subseteq G), \quad F, G \in \mathcal{L}_n, \quad (6)$$

where, for $F \subseteq G$, $|G \setminus F|$ is the number of edges G has in excess of F . Borrowing the terminology from [Drton and Richardson \(2008\)](#), we define the **Möbius simplex** to be the set

$$\Delta_n^M = \{M_n p, p \in \Delta_n\}.$$

As Δ_n^M and Δ_n are in one-to-one correspondence with each other, Δ_n^M is a valid parametrization of all the probability distribution on \mathcal{L}_n .

The form and probabilistic interpretation of Δ_n^M are of course quite different from those of Δ_n . Indeed, we will index the coordinates of the point in Δ_n^M by the elements of \mathcal{I}_n , which we recall is the set of labeled graphs on $[n]$ without isolated nodes. For a $p \in \Delta_n$ and a $z \in \Delta_n^M$ with $z = M_n p$, the value of z at any such graph F is just the **marginal probability** that F is a subgraph of a random graph drawn from p . That is,

$$z(F) = \sum_{\{H \in \mathcal{L}_n : F \subseteq H\}} p(H) = \mathbb{P}(F \subseteq G), \quad F \in \mathcal{I}_n, \quad (7)$$

where G is a random graph in \mathcal{L}_n with distribution p . In particular, the values of z conform to the partial order on \mathcal{I}_n : $z(F) \geq z(F')$ if $F \subseteq F'$. Since the Möbius transform (5) is linear and invertible, Δ_n^M is also a polytope (in fact a simplex), whose vertices are the image by M_n of the vertices of Δ_n : $2^{\binom{n}{2}} - 1$ vertices indexed by all non-empty graphs in \mathcal{I}_n and the vector $1 \in \mathbb{R}^{2^{\binom{n}{2}}}$, the Möbius transform of the point mass on the empty graph.

The Möbius parametrization enjoys the following property, referred to as **backward compatibility**. For $m \leq n$, let $p_n \in \Delta_n$ and $p_n^m = \Pi_n^m p_n \in \Delta_m$ be its marginal, and let z_n and z_n^m denote their Möbius transforms. Then,

$$z_n^m(F) = z_n(F), \quad \forall F \in \mathcal{I}_m. \quad (8)$$

Backward compatibility is a direct consequence of the fact the Möbius parameters are marginal probabilities.

The image of the simplex \mathcal{E}_n of finitely exchangeable probability distribution by the Möbius transform is also a polytope (in fact, a simplex) of the same dimension, denoted by \mathcal{E}_n^M , whose vertices are the Möbius transform of the vertices of \mathcal{E}_n . Clearly, the vertices of \mathcal{E}_n^M can also be indexed by \mathcal{I}_n . By exchangeability and Equation (7), for each $z \in \mathcal{E}_n^M$, $z(F) = z(F')$ whenever $F \sim F'$ in \mathcal{I}_n . In fact, using Equation (6), \mathcal{E}_n^M is defined geometrically by these linear constraints, the linear constraint that $z(\emptyset) = 1$ (where \emptyset signifies the empty graph), and the facet defining inequalities

$$\sum_{\{F' \in \mathcal{I}_n : F \subseteq F'\}} (-1)^{|F' \setminus F|} z(F') \geq 0, \quad \forall F \in \mathcal{I}_n, F \neq \emptyset. \quad (9)$$

3 Homomorphism and isomorphism densities

Next, we will recall some graph-theoretic quantities that play a key role in our derivations. It is not a coincidence that these very same quantities are also used in the theory of graph limits. See, e.g., [Lovász](#)

(2012).

Let $G \in \mathcal{L}_n$ and $F \in \mathcal{L}_m$, where $m \leq n$ (this last requirement is not technically necessary; however we will assume it throughout). The **homomorphism density** of H in G is

$$t_{\text{hom}}(F, G) = \frac{\text{hom}(F, G)}{n^m}, \quad (10)$$

and is equal to the fraction of all mappings from $[m]$ into $[n]$ that define a homomorphism (adjacency preserving mapping) between F and G . Density homomorphisms are **multiplicative**:

$$t_{\text{hom}}(F_1 F_2, G) = t_{\text{hom}}(F_1, G) t_{\text{hom}}(F_2, G), \quad (11)$$

where $F_1 F_2$ is the node-disjoint union of F_1 and F_2 . As we will see, this is the graph-theoretical counterpart to a fundamental probabilistic property known in the literature on exchangeable arrays as the dissociated property. Moreover, it asserts that the homomorphism density t_{hom} is a **character** (Berg et al., 1976, 1984) on the semigroup of graphs under node-disjoint union.

A related concept is that of the **injective homomorphism density** of F into G ,

$$t_{\text{inj}}(F, G) = \frac{\text{inj}(F, G)}{(n)_m}, \quad (12)$$

where $(n)_m = n!/(n-m)!$, and $\text{inj}(F, G)$ is the number of injective mappings from $[m]$ into $[n]$ that define a homomorphism between F and G .

Remark. If F has isolated nodes, the values of both $t_{\text{hom}}(F, G)$ and $t_{\text{inj}}(F, G)$ do not change if F is replaced by the smaller sub-graphs induced by the nodes of positive degree. Therefore, there is no loss of generality in assuming that $F \in \mathcal{I}_n$ when dealing with the quantities in (10) and (12).

In a similar manner, for $G \in \mathcal{L}_n$ and $F \in \mathcal{L}_m$, we define the **isomorphism density** of F into G as

$$t_{\text{iso}}(F, G) = \frac{\text{iso}(F, G)}{n^m},$$

where $\text{iso}(F, G)$ is the number of maps from $[m]$ into $[n]$ that preserve both adjacency and non-adjacency, i.e., such that the induced subgraph of G is isomorphic to F . Finally, let

$$t_{\text{ind}}(F, G) = \frac{\text{ind}(F, G)}{(n)_m} \quad (13)$$

be the injective isomorphism density, where $\text{ind}(F, G)$ is the number of injective mappings from $[m]$ into $[n]$ that preserve both adjacency and non-adjacency, i.e. the number of isomorphisms from F into induced subgraphs of G with m nodes.

The next two results, whose proofs are straightforward and therefore omitted, provide a more statistically transparent interpretation of homomorphism and isomorphism densities. The difference between injective and non-injective densities is precisely the difference between sampling with and without replacement.

Lemma 8. Fix a $G \in \mathcal{L}_n$. Let (U_1, \dots, U_m) be independent random variables uniformly distributed over $[n]$ and H be the random graph on $[m]$ where $i \sim j$ in H if and only if $U_i \sim U_j$ in G . Then, for any graph $F \in \mathcal{L}_m$,

$$t_{\text{hom}}(F, G) = \mathbb{P}(F \subseteq H) \quad \text{and} \quad t_{\text{iso}}(F, G) = \mathbb{P}(F = H).$$

Let (U'_1, \dots, U'_m) be the sequence of random variables describing the outcomes of m draws without replacements of n labelled equiprobable balls. Let H be a random graph on m nodes such that $i \sim j$ if and only if $U'_i \sim U'_j$ in G . Then, for any graph $F \in \mathcal{L}_m$,

$$t_{\text{inj}}(F, G) = \mathbb{P}'(F \subseteq H) \quad \text{and} \quad t_{\text{ind}}(F, G) = \mathbb{P}'(F = H).$$

Using the above representation we immediately obtain the following well-known bound on the difference between subgraph densities arising from injective and non-injective mappings. These bounds will also be used in the proof of Theorem 11.

Lemma 9. *Let $G \in \mathcal{L}_n$ and $m \leq n$. For any $A \subseteq \mathcal{L}_m$, set*

$$t_{\text{ind}}(A, G) = \sum_{F \in A} t_{\text{ind}}(F, G) \quad \text{and} \quad t_{\text{iso}}(A, G) = \sum_{F \in A} t_{\text{iso}}(F, G). \quad (14)$$

Then,

$$\sup_{A \subseteq \mathcal{L}_m} |t_{\text{iso}}(A, G) - t_{\text{ind}}(A, G)| \leq 1 - \frac{\binom{n}{m}}{n^m}. \quad (15)$$

As a result,

$$\sup_{F \in \mathcal{L}_m} |t_{\text{hom}}(F, G) - t_{\text{inj}}(F, G)| \leq 1 - \frac{\binom{n}{m}}{n^m}. \quad (16)$$

Proof. It is enough to prove (15), since (16) clearly follows from it. By Lemma 8, (15) can be established by a well known bound on the total variation distance between a sample with and without replacement: see, e.g., Freedman (1977). Here we give an alternative proof based on sub-graph densities, and inspired by the arguments used in Matúš (1995). Let $\text{notinjiso}(F, G)$ denote the number of non-injective mappings from $[k]$ into $[n]$ that define isomorphisms between F and G so that $\text{iso}(F, G) = \text{notinjiso}(F, G) + \text{ind}(F, G)$. For any $A \subseteq \mathcal{A}_m$, let $\text{notinjiso}(A, G) = \sum_{F \in A} \text{notinjiso}(F, G)$. By Lemma 8, both $t_{\text{ind}}(A, G)$ and $t_{\text{iso}}(A, G)$ are probabilities and therefore, take values in $[0, 1]$. Thus,

$$\begin{aligned} t_{\text{iso}}(A, G) - t_{\text{ind}}(A, G) &= \frac{\text{notinjiso}(A, G) + \text{ind}(A, G)}{n^k} - \frac{\text{ind}(A, G)}{\binom{n}{k}} \\ &= \frac{\text{notinjiso}(A, G)}{n^k} - \frac{\text{ind}(A, G)}{\binom{n}{k}} \left(1 - \frac{\binom{n}{k}}{n^k}\right). \end{aligned}$$

Since, trivially, $0 \leq \text{notinjiso}(A, G) \leq n^k - \binom{n}{k}$, we obtain that $0 \leq \frac{\text{notinjiso}(A, G)}{n^k} \leq \left(1 - \frac{\binom{n}{k}}{n^k}\right)$. Using both bounds in the previous display yields that

$$-t_{\text{ind}}(A, G) \left(1 - \frac{\binom{n}{k}}{n^k}\right) \leq t_{\text{iso}}(A, G) - t_{\text{ind}}(A, G) \leq (1 - t_{\text{ind}}(A, G)) \left(1 - \frac{\binom{n}{k}}{n^k}\right).$$

The claimed bound follows since $0 \leq t_{\text{ind}}(A, G) \leq 1$. ■

Remark. The above bound can be weakened to the simpler bound $\binom{m}{2}/n$. See also Lemma 2.1 in Lovász and Szegedy (2006).

The value of $t_{\circ}(F, G)$ remains unchanged if one or both of its arguments F and G are replaced by isomorphic graphs $F' \sim F$ and $G' \sim G$, where t_{\circ} is any of the densities introduced above. Thus, these graph densities remain well defined if one or both of their arguments belong to \mathcal{U} . The next result uses this fact to establish a correspondence between injective densities and the concepts introduced in Section 2. It will be used in the proof of Theorem 11.

Lemma 10. *Let p_n be the vertex of \mathcal{E}_n corresponding to the uniform distribution over the class $[G]$ and p_n^m its image under the marginal mapping Π_n^m , where $2 \leq m < n$. Let z_n^m be the Möbius transform of p_n^m . Then, for all $F \in \mathcal{L}_m$, $z_n^m(F) = t_{\text{inj}}(F, [G])$ and, for all $F \in \mathcal{L}_m$, $p_n^m(F) = t_{\text{ind}}(F, [G])$.*

Proof. We will give a proof only for the identity involving the injective homomorphism density, since the same arguments apply to the one involving the injective isomorphism density. Let $[H] \in \mathcal{U}_n$ be a given isomorphism class in \mathcal{L}_n and $p_{[H]}$ the point in \mathcal{E}_n corresponding to the uniform distribution over $[H]$. For a given $F \in \mathcal{I}_m$, let $B = \{G \in \mathcal{L}_n : F \subseteq G\}$. For any $\sigma \in \mathcal{S}_n$, let $\sigma^{-1}(B) = \{G \in \mathcal{L}_n : F \subseteq G_\sigma\}$, that is, the set of G such that $G_\sigma \in B$. With a slight abuse of notation we write $p_{[H]}(B) = \sum_{G \in B} p_{[H]}(G)$. Then,

$$\begin{aligned}
z_{[H]}^m(F) &= z_{[H]}(F) \\
&= \sum_{G \in \mathcal{L}_n} 1(F \subseteq G) p_{[H]}(G) \\
&= p_{[H]}(B) \\
&= \frac{1}{n!} \sum_{\sigma \in \mathcal{S}_n} p_{[H]}(\sigma^{-1}(B)) \\
&= \frac{1}{n!} \sum_{G \in [H]} \sum_{\sigma \in \mathcal{S}_n} 1(F \subseteq G_\sigma) p_{[H]}(G) \\
&= \sum_{G \in [H]} \frac{|\{\sigma \in \mathcal{S}_n : F \subseteq G_\sigma\}|}{n!} p_{[H]}(G) \\
&= t_{\text{inj}}(F, [H]) \sum_{G \in [H]} p_{[H]}(G) \\
&= t_{\text{inj}}(F, [H]).
\end{aligned}$$

The first identity follows from the backward compatibility of the Möbius transform, the fourth identity follows from exchangeability and the last identity uses the facts that, for any $\sigma \in \mathcal{S}_n$, $G_\sigma \in [H]$ if and only if $G \in [H]$ and that $t_{\text{inj}}(F, \cdot)$ is constant over $[H]$ (with the common value denoted as $t_{\text{inj}}(F, \cdot)$). ■

4 deFinetti theorems for exchangeable distributions on graphs

4.1 A finite deFinetti theorem

In this section we will use the sub-graph densities introduced in Section 3 to derive a deFinetti theorem for finitely exchangeable probability distributions on graphs based on the Möbius parametrization. The results show that the Möbius parameters of a finitely exchangeable distribution on \mathcal{L}_m that extends to a finitely exchangeable distribution on \mathcal{L}_n , where $2 \leq m < n$, are the expected values of the injective density homomorphisms. These are approximated uniformly well by the expected values of the density homomorphisms with the approximation error of order $O(m^2/n)$. The proof is a simple application of Lemma 10 and Lemma 9 and is the graph-theoretical counterpart of the proof of Theorem 1 in Matúš (1995).

Theorem 11 (deFinetti's theorem for finitely exchangeable distributions on graphs). *Let $p_n \in \mathcal{E}_n$ and $p_n^m = \Pi_n^m p_n$ where $m \leq n$. Let z_n^m be the corresponding Möbius transform of p_n^m . Then, for any subgraph $F \in \mathcal{I}_m$,*

$$z_n^m(F) = \sum_{G \in \mathcal{L}_n} t_{\text{inj}}(F, G) p_n(G) \quad (17)$$

and

$$\max_{F \in \mathcal{I}_m} \left| z_n^m(F) - \sum_{G \in \mathcal{L}_n} t_{\text{hom}}(F, G) p_n(G) \right| \leq 1 - \frac{\binom{n}{m}}{n^m}. \quad (18)$$

Proof. The proof relies on Lemma 9 and can be regarded as extension to the network setting of the geometric arguments used in Diaconis (1977). See also Diaconis and Freedman (1981) and Kerns and Székely (2006), and in particular, Matúš (1995).

For any isomorphism class $[H] \in \mathcal{U}_n$, and any $F \in \mathcal{I}_m$,

$$t_{\text{inj}}(F, G) = t_{\text{inj}}(F, G'), \quad \forall G, G' \in [H]. \quad (19)$$

We will denote by $t_{\text{inj}}(F, [H])$ the common values of all injective homomorphisms of F into the graphs in $[H]$.

Next, let $p_{n,[H]}$ be the finitely exchangeable probability distribution on \mathcal{L}_n corresponding to the uniform distribution over the isomorphic class $[H]$, as described in Equation (1). Since $p_{n,[H]}$ is a vertex of \mathcal{E}_n by Lemma 2, its Möbius transform $z_{n,[H]}$ is a vertex of \mathcal{E}_n^M . Then, because \mathcal{E}_n^M is a simplex, any point in $z_n \in \mathcal{E}_n^M$ can be written as

$$z_n = \sum_{[H] \in \mathcal{U}_n} w_{[H]} z_{n,[H]},$$

for a unique sequence of non-negative numbers $\{w_{[H]}, [H] \in \mathcal{U}_n\}$ such that $\sum_{[H] \in \mathcal{U}_n} w_{[H]} = 1$.

Let $p_{n,[H]}^m = \Pi_n^m p_{n,[H]}$ be the probability distribution over \mathcal{L}_m obtained by marginalizing over $p_{n,[H]}$ and $z_{n,[H]}^m$ be its Möbius transform. Then, for any $F \in \mathcal{I}_m$,

$$z_{n,[H]}^m(F) = t_{\text{inj}}(F, [H]) = \sum_{G \in [H]} \frac{1}{|[H]|} t_{\text{inj}}(F, G), \quad (20)$$

where the first identity follows from Lemma 10.

As a result, for any $F \in \mathcal{I}_m$,

$$\begin{aligned} z_n^m(F) &= \sum_{[H] \in \mathcal{U}_n} w_{[H]} z_{n,[H]}^m(F) \\ &= \sum_{[H] \in \mathcal{U}_n} w_{[H]} t_{\text{inj}}(F, [H]) \\ &= \sum_{[H] \in \mathcal{U}_n} w_{[H]} \left(\sum_{G \in [H]} \frac{1}{|[H]|} t_{\text{inj}}(F, G) \right) \\ &= \sum_{[H] \in \mathcal{U}_n} w_{[H]} \left(\sum_{G \in [H]} p_{n,[H]}(G) t_{\text{inj}}(F, G) \right) \\ &= \sum_{G \in \mathcal{G}_n} t_{\text{inj}}(F, G) \left(\sum_{[H] \in \mathcal{U}_n} w_{[H]} p_{n,[H]}(G) \right) \\ &= \sum_{G \in \mathcal{G}_n} t_{\text{inj}}(F, G) p_n(G), \end{aligned}$$

where the first, second and fourth identities follow from the linearity of marginal operation, Equation (20), and Equation (1), respectively. Thus, Equation (17) follows.

Using the previous identity, for a given $F \in \mathcal{I}_m$ of size, say, k ,

$$\begin{aligned} \left| z_n^m(F) - \sum_{G \in \mathcal{G}_n} t_{\text{hom}}(F, G) p_n(G) \right| &= \left| \sum_{G \in \mathcal{G}_n} t_{\text{inj}}(F, G) p_n(G) - \sum_{G \in \mathcal{G}_n} t_{\text{hom}}(F, G) p_n(G) \right| \\ &\leq \sum_{G \in \mathcal{G}_n} |t_{\text{inj}}(F, G) - t_{\text{hom}}(F, G)| p_n(G) \\ &\leq 1 - \frac{\binom{n}{k}}{n^k}, \end{aligned}$$

where the last inequality is due to Lemma 9. Equation (18) is established by noting that

$$1 - \frac{\binom{n}{k}}{n^k} \leq 1 - \frac{\binom{n}{m}}{n^m}, \quad (21)$$

for all integer $k < m$. ■

The theorem further implies that, for any finitely exchangeable distributions on \mathcal{L}_n , the marginal probabilities of all its small sub-graphs are well approximated by a certain mixture of densities homomorphisms of such sub-graphs, with the unique mixing measure defined over isomorphisms class in \mathcal{L}_n . Formally, we have the following:

Corollary 12. *Assume $2 \leq m < n$. Let $p_n \in \mathcal{E}_n$ and z_n be its Möbius transform. Then, there exists a unique probability distribution $\{w_U, U \in \mathcal{U}_n\}$ on \mathcal{U}_n such that*

$$\left| z_n(F) - \sum_{U \in \mathcal{U}_n} w_U t_{\text{hom}}(F, U) \right| \leq 1 - \frac{\binom{n}{m}}{n^m}, \quad \forall F \in \mathcal{I}_m.$$

Proof. By backward compatibility (8), $z_n^m(F) = z_n(F)$ for all $F \in \mathcal{I}_m$. Furthermore, by Lemma 2, each $p_n \in \mathcal{E}_n$ can be written as

$$p_n = \sum_{[G] \in \mathcal{U}_n} p_{n, [G]} w_{[G]},$$

for a unique probability distribution $\{w_{[G]}, [G] \in \mathcal{U}_n\}$ on \mathcal{U}_n , where $p_{n, [G]}$ is the uniform distribution over the class $[G] \in \mathcal{U}_n$. The claim follows from the fact that $t_{\text{hom}}(F, \cdot)$ takes on the same value $t_{\text{hom}}(F, [G])$ over $[G]$ and collecting terms. ■

By Lemma 2, p_n being extremal is equivalent to p_n being a uniform distribution over some isomorphism class, say $[G]$ in \mathcal{L}_n . If p_n is extremal then $w_{[G]} = 1$. Since homomorphism densities are multiplicative, we can use the fact that $\binom{n}{m}/n^m \leq \binom{m}{2}/n$ to conclude that, for n of larger order than m^2 and if p_n is an extremal distribution on \mathcal{L}_n ,

$$z_n(F) \approx z_n(F_1) z_n(F_2), \quad (22)$$

for each $F \in \mathcal{I}_m$ of the form $F = F_1 F_2$, where we recall that $F_1 F_2$ is the vertex-disjoint union of F_1 and F_2 . As we will show in the next section, the approximation in Equation (22) becomes an equality if p_n is embedded into a sequence of consistent finitely exchangeable distributions that extend to an extremal exchangeable distribution over \mathcal{L}_∞ . Furthermore, all such extremal distributions are defined by these identities.

A result analogous to Theorem 11 holds also for joint probabilities. We have chosen to focus on marginal probabilities since they are more natural in this context, as they directly lead to the key approximation property of Equation (22).

Corollary 13. Consider the setting of Theorem 11. Then, for any $F \in \mathcal{L}_m$,

$$p_n^m(F) = \sum_{G \in \mathcal{L}_n} t_{\text{ind}}(F, G) p_n(G) \quad (23)$$

and, as a result,

$$\max_{F \in \mathcal{L}_m} \left| p_n^m(F) - \sum_{G \in \mathcal{L}_n} t_{\text{iso}}(F, G) p_n(G) \right| \leq 1 - \frac{\binom{n}{m}}{n^m}. \quad (24)$$

Furthermore, letting \tilde{p}_n^m the probability distribution on \mathcal{L}_m given by

$$\tilde{p}_n^m(F) = \sum_{G \in \mathcal{L}_n} t_{\text{iso}}(F, G) p_n(G), \quad F \in \mathcal{L}_m,$$

we have

$$d_{\text{TV}}(p_n^m, \tilde{p}_n^m) \leq 1 - \frac{\binom{n}{m}}{n^m}, \quad (25)$$

where $d_{\text{TV}}(P, Q)$ denotes the total variation distance between the probability distributions P and Q .

Proof. We omit the proofs of (23) and (24), since they are nearly identical to the proofs of (17) and (18) given above. To prove (25), let $A \subset \mathcal{L}_m$ and recall the quantities defined in eq. (14):

$$t_{\text{ind}}(A, G) = \sum_{F \in A} t_{\text{ind}}(F, G) \quad \text{and} \quad t_{\text{iso}}(A, G) = \sum_{F \in A} t_{\text{iso}}(F, G).$$

Notice that by Lemma 9, $|t_{\text{ind}}(A, G) - t_{\text{iso}}(A, G)| \leq 1 - \frac{\binom{n}{m}}{n^m}$, for any $G \in \mathcal{L}_n$ and $A \subseteq \mathcal{L}_m$. Then,

$$\begin{aligned} \left| \sum_{F \in A} p_n^m(F) - \sum_{F \in A} \tilde{p}_n^m(F) \right| &= \left| \sum_{F \in A} \left(\sum_{G \in \mathcal{G}_n} t_{\text{ind}}(F, G) p_n(G) \right) - \sum_{F \in A} \left(\sum_{G \in \mathcal{G}_n} t_{\text{iso}}(F, G) p_n(G) \right) \right| \\ &= \left| \sum_{G \in \mathcal{L}_n} (t_{\text{ind}}(A, G) - t_{\text{iso}}(A, G)) p_n(G) \right| \\ &\leq \sum_{G \in \mathcal{L}_n} |t_{\text{ind}}(A, G) - t_{\text{iso}}(A, G)| p_n(G) \\ &\leq 1 - \frac{\binom{n}{m}}{n^m}. \end{aligned}$$

Inequality (25) now follows since, by definition,

$$d_{\text{TV}}(p_n^m, \tilde{p}_n^m) = \sup_{A \subseteq \mathcal{L}_m} \left| \sum_{F \in A} p_n^m(F) - \sum_{F \in A} \tilde{p}_n^m(F) \right|.$$

This completes the proof. ■

Just like in Corollary 12, we can equivalently express (23) as

$$p_n^m(F) = \sum_{U \in \mathcal{U}_n} t_{\text{ind}}(F, U) w_U,$$

for a unique probability distribution $\{w_U, U \in \mathcal{U}_n\}$.

4.2 From finite exchangeability to exchangeability

Below we will strengthen the conclusions of Theorem 11 by further assuming that each $p_n \in \mathcal{E}_n$ is an element of a sequence $\{p_n\}_{n=2}^\infty$ of finitely exchangeable distributions that are consistent, i.e. satisfy Equation (3). As noted above, each such sequence extends uniquely to one element in the simplex \mathcal{E}_∞ of exchangeable probability distribution on \mathcal{L}_∞ . Below, we will establish a deFinetti type of theorem for exchangeable distributions and, along the way, relate it to the theory of graph limit. This connection is well known and has been elucidated in [Diaconis and Janson \(2008\)](#).

We begin by introducing a few concepts that are necessary to represent distributions over infinite graphs and graph sequences. First, it is easy to see that any probability distributions in \mathcal{E}_∞ admits a Möbius parametrization that is completely analogous to the one given for distributions of finite random graphs and are based on marginal probabilities of finite subgraphs without isolated nodes. In detail, for a point $p_\infty \in \mathcal{E}_\infty$ we will write

$$z = z(p_\infty) = (z(F), F \in \mathcal{I}) \in [0, 1]^{\mathcal{I}}, \quad (26)$$

for the sequence of Möbius parameters given by

$$z(F) = \sum_{\{H \in \mathcal{L}_\infty : F \subseteq H\}} p_\infty(H) = \mathbb{P}(F \subset G) \quad F \in \mathcal{I}, \quad (27)$$

where G is an infinite random graph with distribution p_∞ and we recall that \mathcal{I} is the set of all finite sub-graphs without isolated nodes. In particular, the Möbius parametrization $(z(F), F \in \mathcal{I}) \in [0, 1]^{\mathcal{I}}$ of a distribution on \mathcal{L}_∞ satisfies the properties that $z(\emptyset) = 1$, $z(F) = z(F')$ if $F \sim F'$ and Equation (9) holds for all n .

Furthermore, for any $H \in \mathcal{I}$, the homomorphism density $t_{\text{hom}}(H, \cdot)$ is well defined on \mathcal{L}_∞ , and is given by

$$t_{\text{hom}}(H, G) = \lim_n t_{\text{hom}}(H, G(n)) = \lim_n t_{\text{hom}}(H, [G(n)]), \quad G \in \mathcal{L}_\infty. \quad (28)$$

Clearly, $t_{\text{inj}}(H, G) = t_{\text{hom}}(H, G)$, for any $H \in \mathcal{I}$, $F \in \mathcal{L}$ and $G \in \mathcal{L}_\infty$.

We will also require the notion of graph limits: see [Lovász and Szegedy \(2006\)](#), [Borgs et al. \(2008\)](#), [Lovász \(2012\)](#). Following [Diaconis and Janson \(2008\)](#), we let \mathcal{U}_∞ be the collection of all sequences

$$(x_F, F \in \mathcal{I}) \in [0, 1]^{\mathcal{I}} \quad (29)$$

of the form

$$x_F = \lim_n t_{\text{hom}}(F, U_n)$$

for some sequence $\{U_n\}_n$ of unlabeled graphs, with $U_n \in \mathcal{U}_n$ for all n . The set \mathcal{U}_∞ consists of all possible limits of sequences of unlabeled graphs, according to the definition of [Lovász and Szegedy \(2006\)](#) and [Borgs et al. \(2008\)](#). Intuitively, one can think of each $U \in \mathcal{U}_\infty$ as an “infinite unlabeled graph”. Indeed, notice the similarity between (29) and (26). In order to emphasize the role of density homomorphisms in this definition and also to draw a parallel with the limits of density homomorphisms for infinite labeled graphs (see 28), with a slight abuse of notation we will write $t_{\text{hom}}(F, U)$ for the element of the sequence $U \in \mathcal{U}_\infty$ corresponding to $F \in \mathcal{I}$. Notice that if F and H are isomorphic then $t_{\text{hom}}(F, U) = t_{\text{hom}}(H, U)$, for all $U \in \mathcal{U}_\infty$.

The set \mathcal{U}_∞ is a compact subset of the compact metric space $[0, 1]^{\mathcal{I}}$ when endowed with the metric

$$d(x, y) = \sum_{i=1}^{\infty} \frac{1}{2^i} |x_{F_i} - y_{F_i}|,$$

where F_1, F_2, \dots is an enumeration of all the graphs in \mathcal{I} .

The main result of this section provides a representation of the Möbius parameters of the probability distributions in \mathcal{E}_∞ as expected density homomorphism of all the graphs in \mathcal{I} . In addition, the Möbius parameters of the extremal distributions in \mathcal{E}_∞ satisfy a defining set by polynomial equations given below in (32).

Theorem 14 (deFinetti Theorem for exchangeable random networks). *The Möbius parameters corresponding to the probability distribution $p_\infty \in \mathcal{E}_\infty$ are given by*

$$z(F) = \sum_{\{G \in \mathcal{L}_\infty\}} t_{\text{hom}}(F, G) p_\infty(G), \quad F \in \mathcal{I}, \quad (30)$$

or, equivalently, by

$$z(F) = \mathbb{E}[t_{\text{hom}}(F, U)], \quad \forall F \in \mathcal{I}, \quad (31)$$

where the expectation is with respect to the distribution of a random variable U taking values in \mathcal{U}_∞ . A point p_∞ is extremal in \mathcal{E}_∞ if and only if its Möbius parameters satisfy the conditions

$$z(F) = z(F_1)z(F_2) \quad (32)$$

for all $F \in \mathcal{I}$ with $F = F_1 F_2$. Furthermore, in this case we have

$$z(F) = t_{\text{hom}}(F, U), \quad \forall F \in \mathcal{I}, \quad (33)$$

implying that p_∞ is concentrated on a single deterministic graph limit $U \in \mathcal{U}_\infty$

Proof. We will use throughout the notation in [Diaconis and Janson \(2008\)](#). Represent the set \mathcal{L}_∞ as the product set $\{0, 1\}^{E(K_\infty)}$, where $E(K_\infty)$ denotes all subsets of edges of an infinite complete graph. Endowed with the product topology, \mathcal{L}_∞ is a compact metric space. For any $n \in \mathbb{N}$ and $G \in \mathcal{L}_\infty$, write $G(n)$ for the subgraph of G induced by the vertex set $[n]$.

Let $G \in \mathcal{L}_\infty$ be a random infinite labeled graph. Then, since \mathcal{L}_∞ has the product topology, the sequence $\{G(n)\}_{n \geq 2}$ converges in distribution to G . If G is exchangeable, then, for a fixed $m \geq 2$ and each $n \geq m$, [Theorem 11](#) yields that

$$\left| \mathbb{P}(G(m) \supset F) - \mathbb{E}[t_{\text{hom}}(F, G(n))] \right| \leq 1 - \frac{\binom{n}{m}}{n^m},$$

uniformly over all $F \in \mathcal{I}_m$. Taking the limit in n , by weak convergence and boundedness of density homomorphism, we get that

$$\mathbb{P}(G(m) \supset F) = \mathbb{E}[t_{\text{hom}}(F, G)]$$

for all $F \in \mathcal{I}_m$ and all $m \in \mathbb{N}$. Thus,

$$\mathbb{P}(G \supset F) = \mathbb{E}[t_{\text{hom}}(F, G)], \quad \forall F \in \mathcal{I},$$

which is Equation (30).

To show (32), we will rely on the following result of [Diaconis and Janson \(2008\)](#).

Lemma 15 ([Diaconis and Janson \(2008\)](#), Corollary 5.4). *There is a one-to-one correspondence between the extreme points of the set \mathcal{E}_∞ and the set \mathcal{U}_∞ , given by*

$$t(F, U) = z(F), \quad \forall F \in \mathcal{I}, \quad (34)$$

where $U \in \mathcal{U}_\infty$ and $z(F)$ is the value of Möbius parameter at F for the corresponding p_∞ (see [27](#)).

We can now prove (32). Assume that p_∞ is extremal in \mathcal{E}_∞ with Möbius parameters $\{z(F), F \in \mathcal{I}\}$. Let $U \in \mathcal{U}_\infty$ its corresponding sequence. Then there exists a sequence $\{U_n\}_n \subset \mathcal{U}$ of unlabeled graphs with $U_n \in \mathcal{U}_n$ for all n such that $t(F, U) = \lim_n t_{\text{hom}}(F, U_n)$, for any $F \in \mathcal{I}$. Consider any pair of node disjoint graphs F_1 and F_2 in \mathcal{I} . Without loss of generality, we may take the nodesets of F_1 and F_2 to be $[m_1]$ and $\{m_1 + 1, \dots, m_1 + m_2\}$, respectively. Using lemma 15,

$$z(F_1 F_2) = t(F_1, F_2, U) = \lim_n t_{\text{hom}}(F_1 F_2, U_n) = \lim_n t_{\text{hom}}(F_1, U_n) t_{\text{hom}}(F_2, U_n) = t(F_1, U) t(F_2, U) = z(F_1) z(F_2),$$

where the third identity follows from the multiplicative property of density homomorphisms, which holds for all $n \geq m_1 + m_2$; see (11). The same argument applies to any pair of node-disjoint graphs F_1 and F_2 in \mathcal{I} , and (32) follows. Now suppose that (32) holds. Using Equation (30), for any pair of node-disjoint isomorphic graphs F_1 and F_2 in \mathcal{I} ,

$$z(F_1 F_2) = \mathbb{E}[t(F_1 F_2, G)] = \mathbb{E}[t^2(F_1, G)],$$

where G is the random graph in \mathcal{L}_∞ with distribution p_∞ . Using (32) and the fact that $z(F_1) = z(F_2)$ (since F_1 and F_2 are isomorphic), we have that

$$\mathbb{E}[t^2(F_1, G)] = z(F_1 F_2) = z(F_1) z(F_2) = z^2(F_1) = (\mathbb{E}[t(F_1, G)])^2,$$

and, therefore, that $t(F_1, G)$ is almost surely constant. Since the choice of F_1 is arbitrary, we conclude that the random variable $t_{\text{hom}}(F, G)$ is almost surely constant for each $F \in \mathcal{I}$. It then follows from Lemma 15 that the distribution of G is extremal. Indeed, let G be a realization from p_∞ and consider the sequence $\{G(n)\}_n$ of induced sub-graphs. Then,

$$\lim_n t_{\text{hom}}(F, G(n)) = t_{\text{hom}}(F, G) = z(F), \quad \forall F \in \mathcal{I}.$$

For each n let U_n the unlabeled graph representing the isomorphism class of $G(n)$. Then, $t_{\text{hom}}(F, G(n)) = t_{\text{hom}}(F, U_n)$ and, therefore, $\lim_n t_{\text{hom}}(F, U_n)$ exists for all $F \in \mathcal{I}$ (and, in fact, coincides with $z(F)$). Thus there exists a $U \in \mathcal{U}_\infty$ such that (34) holds. By Lemma 15, p_∞ is extremal.

Equations (31) and (33) are a direct consequence of Lemma 15. ■

The identity (31) signifies that the Möbius parameters of any $p_\infty \in \mathcal{E}_\infty$ can be expressed as an average of density homomorphisms over graph limits, while Equation (33) expresses the result that there is a one-to-one correspondence between graph limits and extremal distributions in \mathcal{E}_∞ (this formally stated in Lemma 15 above).

Theorem 14 gives a reformulation of well known results about symmetric binary exchangeable arrays (see, e.g. Aldous, 1981, 1985; Lauritzen, 2008; Silverman, 1976; Eagleson and Weber, 1978; Hoover, 1979; Kallenberg, 2005) and can also be directly linked to the theory of graph limits, as shown in particular by Diaconis and Janson (2008) (see also Lovász, 2012, Chapter 11). Our contribution is a relatively simple proof that combines the finite exchangeability bound from Theorem 11 with classic arguments from the theory of weak convergence of measure as detailed in Diaconis and Janson (2008).

An equivalent version of Theorem 14 can also be given for probability parameters as opposed to Möbius parameters. However, the parametrization of extremal distributions in \mathcal{E}_∞ by the induced probabilities of finite graphs does seem to satisfy any factorization properties, such as the ones expressed in (32) for the Möbius parameters. For this reason, we find the Möbius parametrization more convenient. We refrain from providing the details.

One of the main implications of Theorem 14 is that, for any integer $n \geq 2$, if $p_n \in \mathcal{E}_n$ is the marginal of an extremal exchangeable distribution on \mathcal{L}_∞ , then, by Equation (32), its Möbius parameters (marginal probabilities) satisfy the identities

$$z_n(F_1, F_2) = z_n(F_1) z_n(F_2), \quad \forall F_1 F_2 \in \mathcal{I}_n, \tag{35}$$

i.e. the approximation Equation (22) holds exactly. This property holding for all n is equivalent to a well-known measure-theoretical property of exchangeable distributions over binary arrays, known as *dissociatedness*; see, e.g., [Silverman \(1976\)](#). In fact, dissociatedness is a necessary and sufficient condition for an exchangeable distribution over arrays to be extremal ([Aldous, 1985](#)). In the graph limit literature (see, e.g., [Lovász, 2012](#), Chapter 11) an equivalent formulation of Equation (35) for all n is referred to as the *local* property of the associated sequence of distributions.

We will refer to distributions in \mathcal{E}_n satisfying Equation (35) as **dissociated**. Notice that, according to our definition, a dissociated distribution in \mathcal{E}_n needs not be the marginal of any dissociated or even finitely exchangeable distributions over larger graphs. This is in contrast with the classic notion of dissociatedness used in the probabilistic literature, which requires Equation (35) to hold for all n and therefore applies to all the marginals of an exchangeable distribution. Indeed, as we will show below, there exist dissociated distributions in \mathcal{E}_n , for all $n \geq 4$ that cannot be extended to any distributions on larger graphs.

Since dissociated distributions in \mathcal{E}_n contain the Möbius parameters of the marginals of all extremal distributions in \mathcal{E}_∞ , in order to understand the subset of \mathcal{E}_n corresponding to image under the marginal mapping of all exchangeable distributions it is crucial to study dissociated distributions, which we do next in the next Section.

4.2.1 Connection with harmonic analysis

There is an interesting connection between Theorem 14 and harmonic analysis on semigroups ([Berg et al., 1984](#); [Ressel, 2008](#)). More precisely, if we consider the semigroup $(\mathcal{J}, +)$ of unlabeled graphs without isolated nodes, where $+$ denotes node disjoint union, the Möbius parameters clearly satisfy

$$z(F) = \phi([F]),$$

for some function $\phi : \mathcal{J} \rightarrow \mathbb{R}_+$. As is shown in the lemma below, the function ϕ is **positive definite** on $(\mathcal{J}, +)$, meaning that any matrix of the form

$$m_{ij} = \phi([F_i] + [F_j]), \quad i, j = 1, \dots, n$$

is positive semidefinite.

Lemma 16. *Let G be a random exchangeable graph with Möbius parameters z given as above. Then the function ϕ is bounded and positive definite on $(\mathcal{J}, +)$.*

Proof. Clearly $\phi(\emptyset) = 1$ and ϕ is bounded. Introduce the binary random variables X_{ij} for $i \neq j \in \mathbb{N}$ where $X_{ij} = 1$ if $i \sim j$ in G and $X_{ij} = 0$ otherwise; X is the (random) adjacency matrix of G . Then, clearly

$$z(F) = \mathbb{E} \left(\prod_{ij: i \sim j \in F} X_{ij} \right).$$

So elementary calculations will verify that

$$\begin{aligned} \sum_{u,v=1}^n c_u c_v \phi([F_u] + [F_v]) &= \sum_{u,v=1}^n c_u c_v \mathbb{E} \left(\prod_{i \sim j \in F_u} X_{ij} \prod_{i \sim j \in F_v} X_{ij} \right) \\ &= \mathbb{E} \left\{ \sum_u c_u \prod_{i \sim j \in F_u} X_{ij} \right\}^2 \geq 0. \end{aligned}$$

This completes the proof. ■

We note that the property in Lemma 16 is referred to as *reflection positivity* in Lovász and Szegedy (2006).

Now Berg et al. (1976) show that the set of bounded positive definite functions on an Abelian semigroup is a Bauer simplex with the set of characters as extreme points; this is essentially equivalent to the statement in Theorem 14.

4.2.2 Connection with graphons

The conclusions of Theorem 11 can be equivalently expressed using graphons. Indeed, paraphrasing a deep result about exchangeable arrays established by Aldous (1981) and Hoover (1979) (see also Kallenberg, 2005), the Möbius sequence $(z(F), F \in \mathcal{I})$ corresponding to an extremal exchangeable distribution admits the representation

$$z(F) = \int_{[0,1]} \int_{[0,1]^n} \prod_{(j,j) \in E(F)} W(x_i, x_j) dx_1 \dots dx_n,$$

for all $F \in \mathcal{I}_n$ and all $n \geq 2$, and some symmetric measurable function $W: [0,1]^2 \rightarrow [0,1]$ (which is not uniquely defined). The same result was also established in the context of graph limits by Borgs et al. (2008) and Lovász and Szegedy (2006), who termed the function W a graphon. Furthermore, equation (31) takes the form

$$z(F) = \int_{[0,1]} \left\{ \int_{[0,1]^n} \prod_{(j,j) \in E(F)} \phi(\alpha, x_i, x_j) dx_1 \dots dx_n \right\} d\alpha,$$

for all $F \in \mathcal{I}_n$ and all $n \geq 2$, for some measurable function $\phi: [0,1]^3 \rightarrow [0,1]$ (not necessarily uniquely defined), symmetric in its last two arguments. See, e.g., Chapter 14 in Aldous (1981).

5 The manifold of dissociated exchangeable distributions

Let $\mathcal{D}_n \subset \mathcal{E}_n^M$ be the set of Möbius parameters of finitely exchangeable dissociated distributions on \mathcal{L}_n . By definition, \mathcal{D}_n is comprised of all the points in \mathcal{E}_n^M that satisfy the system of polynomial equations (35). Therefore, \mathcal{D}_n is the intersection of \mathcal{E}_n^M with a smooth manifold, in fact an affine variety in $(z(F), F \in \mathcal{I}_n)$. For this reason, we will refer to \mathcal{D}_n as the **dissociated manifold**.

Clearly, the image of \mathcal{D}_n under the inverse Möbius transform is a subset of \mathcal{E}_n that can be also defined by a system of polynomial equations in the probability parameters, though these relations are not as simple as the ones in (35).

The next result describes some of the properties of the set \mathcal{D}_n . In particular, it shows that if p_n is the marginal of an exchangeable, non-extremal distribution on \mathcal{L}_∞ , then its Möbius parameters (marginal probabilities) are mixtures of the Möbius parameters of dissociated distributions in \mathcal{E}_n . From this, we obtain a partial geometric characterization of the set $\Pi_\infty^n(\mathcal{E}_\infty)$. It should be apparent now why Möbius parameters (marginal probabilities) are better suited to describe exchangeability in our context.

Lemma 17. *The dimension of \mathcal{D}_n is the number of unlabeled connected graphs with at most n nodes. If p_∞ is a distribution in \mathcal{E}_∞ and $z = z(p_\infty)$ is as in Equation (27), then*

$$z_n \in \begin{cases} \mathcal{D}_n & \text{if } p_\infty \text{ is extremal} \\ \text{convhull}(\mathcal{D}_n) & \text{otherwise,} \end{cases}$$

where $z_n = (z(F), F \in \mathcal{I}_n)$. As a result, for each $n \geq 2$,

$$P_\infty^m(\mathcal{E}_\infty) \subset \{M_n^{-1}z_n : z_n \in \text{convhull}(\mathcal{D}_n)\}.$$

Proof. That claim about the dimension of \mathcal{D}_n follows from counting the number of polynomial equations in (35) (see, e.g., [Drton and Richardson \(2008\)](#) and references therein for a similar calculation) and taking into account the fact that $z_n(F) = z_n(F')$ for all $F \sim F'$. To show the second statement, let $p_\infty \in \mathcal{E}_\infty$ and $z = z(p_\infty)$ be as in Equation (27). Then, by (31),

$$z_n(F) = z(F) = \mathbb{E}[t_{\text{hom}}(F, U)], \quad \forall F \in \mathcal{I}_n,$$

where in the above display U is a random variable taking values in \mathcal{U}_∞ . The claim is then established after noting that for any deterministic $U \in \mathcal{U}_\infty$, it holds that $t_{\text{hom}}(F, U) \in \mathcal{D}_n$ for all $F \in \mathcal{I}_n$, by the second part of Theorem 14. ■

Lemma 17 should be compared with the results in [Erdős et al. \(1979\)](#), where it is shown, with a different language, that the set of Möbius parameters in \mathcal{E}_n^M arising from extremal distributions in \mathcal{E}_∞ belongs to \mathcal{D}_n and has a non-empty interior, of dimension equal to the number of unlabeled connected graphs with at most n nodes. The authors further remark that not much else is known about this set for any $n \geq 3$, including its topological properties (though they do show that it is path-connected).

In light of this, one may be led to conjecture that each point in the dissociated manifold \mathcal{D}_n arises as the Möbius parameter of an extremal exchangeable distribution on \mathcal{E}_∞ . However, quite surprisingly, this is not the case. The following example provides a family of strictly positive probability distributions in \mathcal{D}_4 that are not extendable. Geometrically, this set is a line segment in \mathcal{D}_4 . Other examples of dissociated distributions with zero entries in \mathcal{D}_4 that are not extendable are given in tables 1 and 2, which we discuss in the next section.

Example 18. Take $n = 4$. Consider the (strictly positive) probability distribution $p^* = \alpha p_1 + (1 - \alpha)p_2$ on \mathcal{L}_4 , where $\alpha \in (0, 1)$, p_1 is the Erdős-Renyi distribution on \mathcal{L}_4 with $p = 1/2$, and p_2 is the distribution on \mathcal{L}_4 corresponding to the mixture of the point mass at the empty graph, the uniform distribution over the 4 graphs isomorphic to the union of a triangle and an isolated node, and the uniform distribution over the 3 graphs isomorphic to the 4-cycle, with weights $1/8$, $1/2$ and $3/8$ respectively. Both p_1 and p_2 are finitely exchangeable, and, therefore, so is p^* . Furthermore, since under both p_1 and p_2 the Möbius parameters corresponding to graphs isomorphic to an edge and to the disjoint union of two edges are $1/2$ and $1/4$ respectively, $p^* \in \mathcal{D}_n$ for each $\alpha \in (0, 1)$. Yet, as p_2 is not contained in the image of the marginal mapping Π_5^4 over \mathcal{L}_5 , there is no distribution on \mathcal{L}_n , $n \geq 5$, whose marginal in \mathcal{L}_4 is p^* . Therefore, $z(p^*) \notin \Pi_\infty^5(\mathcal{E}_\infty)$.

It remains unknown whether there exist any simple criteria to determine whether the inverse Möbius transform of any point in \mathcal{D}_n is also in $\Pi_\infty^m(\mathcal{E}_\infty)$.

Connections with [Lauritzen et al. \(2017\)](#)

In [Lauritzen et al. \(2017\)](#), we have also investigated exchangeable network models as graphical models on binary data with symmetric restrictions. There we have shown that distributions in \mathcal{E}_n can only be compatible with few Markov properties, and we have identified all the possible conditional independence structures that such distributions may exhibit. Furthermore, we have proved that the only non-trivial conditional independence structure that yields a consistent sequence of finitely exchangeable probability distributions corresponds to a certain bi-directed graphical model for marginal independence. Such a model, which can be thought of as a canonical parametric model encompassing all finitely exchangeable networks of any given size, belongs to the class of marginal models for binary data studied by [Drton and Richardson \(2008\)](#) (see, also, [Roverato et al., 2013](#)). In particular, it is obtained from enforcing the dissociatedness constraints (35) in addition to exchangeability. One of the implications of these results is that the image under Π_∞^m of all extremal families in \mathcal{E}_∞ is a strict submodel of a graphical model for marginal independence. Finally, the

model can be parametrized as a curved exponential family on \mathcal{L}_n with natural sufficient statistics given by the injective density homomorphisms and dimension equal to the number of connected unlabeled graphs on n nodes.

5.1 Maximum likelihood estimation

In this section we further investigate some of the statistical properties of the models specified by the manifold \mathcal{D}_n of exchangeable and dissociated distributions. We will focus on the basic problem of estimating the Möbius parameters by maximizing the likelihood based on a sample of size one.

If $G \in \mathcal{L}_n$ is the observed network, a maximum likelihood estimator (MLE) of the Möbius parameters under the dissociated model is a point in the set

$$\operatorname{argmax}_{z \in \mathcal{D}_n} \ell(G, z),$$

where $\ell(\cdot, G)$ is the likelihood function, given by

$$z \in \mathcal{D}_n \mapsto \ell(z, G) = \sum_{\{F \in \mathcal{I}_n, : G \subseteq F\}} (-1)^{|F \setminus G|} z(F).$$

Using exchangeability, we can rewrite the likelihood function as

$$\ell(z, G) = \sum_{U \in \mathcal{U}_n} (-1)^{E(U) - E(G)} r_U(G) z(U),$$

where for a (labeled or unlabeled) graph G , $E(G)$ is the number of its edges, $r_U(G)$ is the number of graphs in \mathcal{L}_n containing G as a subgraph and belonging to the isomorphism class represented by U , and $z(U)$ is the common value of the coordinates of the Möbius parameters z corresponding to the graphs in the isomorphism class represented by U . See Examples 1 and 2 in [Lauritzen et al. \(2017\)](#).

As remarked in the previous Section, points in \mathcal{D}_n correspond to the closure of the mean-value space of a curved exponential family of probability distributions on \mathcal{L}_n . The MLE of the Möbius parameters may be on the boundary of \mathcal{D}_n and may not be unique. Both cases are problematic from a statistical standpoint: the former case implies that the probability distribution corresponding to the MLE assigns zero probability to some graphs in \mathcal{L}_n (when in fact all probabilities should be positive) and the latter case renders statistical inference based on such estimator ill-posed.

In order to study both issues, we have obtained numerically all the possible maximum likelihood estimates under the constraints of exchangeability and dissociatedness for all the realizations of one network on four nodes. We have carried the calculations in *Mathematica* using the built in optimization method. [Drton and Richardson \(2008\)](#) propose a general algorithm for computing the MLE of the Möbius parameters of marginal models for binary data that could in principle be used for our problem. While such algorithm is more efficient and presumably faster than the brute force optimization, its performance guarantees require positive counts, a condition that is never satisfied when the data take the form of one observed network.

When $n = 4$, there are 11 isomorphism classes, shown below in [Figure 1](#) as unlabeled graphs, along with their respective sizes.

[Table 1](#) and [Table 2](#) show the maximum likelihood estimates of the Möbius parameter and of the actual probabilities, respectively. An empty entry in the table signifies a value of zero. It is apparent from [Table 2](#) that all the estimates contain zero coordinates, a fact that implies that, with only one observed network, all the maximum likelihood estimates lie on the boundary of the parameter space. Furthermore, the MLE is not unique:




Figure 1: All non-isomorphic graphs $U \in \mathcal{U}_4$, along with the size m of the isomorphism class each represents, denoted by $\times m$.

it can be seen from Tables 1 and 2 that when the observed network consists of two parallel edges, or is a path or a cycle then the likelihood is maximized along line segments on the boundary of both the simplex and the Möbius simplex. (See also Example 7 in Lauritzen et al. (2017)). Finally, direct calculations reveal that, with the exception of the point masses at the empty and complete graphs, none of the maximum likelihood estimates of the probability distributions extend to exchangeable distributions over larger networks.

The fact that there are zeros in the MLEs of Table 2 means that it is relatively easy to check, for each case, that no exchangeable distribution on 5-node graphs can marginalize to that MLE. Consider for example the second row, corresponding to observing a 4-node graph with only one edge. The MLE is a mixture of a point mass at the complete graph and of the uniform distribution over graphs isomorphic to the observed one. In order for the MLE to be the marginal of some exchangeable distribution on 5-node graphs, that distribution must in turn be a mixture of uniform distributions over isomorphic 5-node graphs (and a point mass on the complete graph on 5 nodes) such that the removal of any one node will either be a 4-node graph with one edge or a complete graph. Such distribution does not exist (because there does not exist any 5-node unlabeled graph such that removing any one node will produce as an induced subgraph a 4-node graph with only one edge). Other cases can be checked by similar arguments.

Table 1: Maximum likelihood estimates of the Möbius parameters for all possible samples of size one from \mathcal{L}_4 . Along the rows of the table we report only the isomorphic classes in \mathcal{G}_4 , each represented as an undirected graph in \mathcal{U}_4 . Indeed, by exchangeability, isomorphic graphs in \mathcal{G}_4 yield the same maximum likelihood estimate. The columns of the table are indexed by all $U \in \mathcal{U}_4$ without isolated nodes. Thus the entry (U', U) represents the maximum likelihood estimate of $z(F')$, where $U = [F]$ if the observed network is a graph in the class represented by U' . An empty cell is equivalent to 0. In addition, $-(3\sqrt{2}/4) + 1 \leq b, d \leq (3\sqrt{2}/4) - 1$ and $-1/32 \leq c \leq 1/32$.

$x \setminus U$												
	1/5	1/25	1/25	1/25	1/25	1/25	1/25	1/25	1/25	1/25	1/25	1/25
	1/2	5/16	1/4	1/4	1/4	1/4	1/4	1/4	1/4	1/4	1/4	1/4
	$\sqrt{2} - 1$	$(3\sqrt{2}/2) - 2$	$3 - 2\sqrt{2}$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$
	1/2	1/4	1/4	1/8	1/8	1/8	1/8	1/8	1/8	1/8	1/8	1/8
	1/2	3/16	1/4	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c
	1/2	1/4	1/4	1/4	1/4	1/4	1/4	1/4	1/4	1/4	1/4	1/4
	$2 - \sqrt{2}$	$1 - \sqrt{2}/2$	$6 - 4\sqrt{2}$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$
	1/2	5/16	1/4	3/16	3/16	3/16	3/16	3/16	3/16	3/16	3/16	3/16
	4/5	16/25	16/25	12/25	12/25	12/25	12/25	12/25	12/25	12/25	12/25	12/25
	1	1	1	1	1	1	1	1	1	1	1	1

Table 2: Maximum likelihood estimates of the probability parameters for all possible samples of size one from \mathcal{L}_4 . Along the rows and columns of the Table we report only the isomorphic classes in \mathcal{G}_4 , each represented as an undirected graph in \mathcal{U}_4 . Thus, the entry corresponding to the pair (U', U) is the maximum likelihood estimate of the probability of observing a graph in the isomorphism class represented by U if the observed network is a graph in the class represented by U' . An empty cell is equivalent to 0. In addition, $-(3\sqrt{2}/4) + 1 \leq b, d \leq (3\sqrt{2}/4) - 1$ and $-1/32 \leq c \leq 1/32$.

$x \setminus U$												
	1	4/25	1/16	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$	$(3\sqrt{2}/4) - 1 - b$
	1/8	1/8	$3 - 2\sqrt{2}$	1/8	1/8	1/8	1/8	1/8	1/8	1/8	1/8	1/8
	1/8	1/8	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c	1/32 - c
	1/4	1/4	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$	$(3\sqrt{2}/4) - 1 - d$
	1/25	1/25	1/25	1/25	1/25	1/25	1/25	1/25	1/25	1/25	1/25	1/25

Table 3: Dimension of the sets \mathcal{E}_n and \mathcal{D}_n a function of the number of nodes n . The numbers are sourced from OEIS Foundation Inc. (2011), The On-Line Encyclopedia of Integer Sequences, see <http://oeis.org/A000088> and <http://oeis.org/A001349>.

	n									
	3	4	5	6	7	8	9	10	11	
$\dim(\mathcal{E}_n)$	3	10	33	155	1,043	12,345	274,667	12,005,167	1,018,997,863	
$\dim(\mathcal{D}_n)$	3	9	30	142	995	12,112	273,192	11,989,763	1,018,690,328	

6 Conclusions

It is worth commenting on the difference between Theorem 11 and 14 and analogous results for finitely exchangeable and exchangeable random binary sequences; see, e.g. Diaconis (1977); Diaconis and Freedman (1981); Kerns and Székely (2006). First, the n -dimensional marginal of any exchangeable binary sequence can be described geometrically as a uniquely determined point in the convex hull of the intersection of the one-dimensional variety corresponding to the surface of independence inside the $(2^n - 1)$ -dimensional simplex with the n dimensional affine subset of finitely exchangeable distributions. In our setting the manifold \mathcal{D}_n of dissociated distributions (actually, as we saw, a certain non-trivial subset of it) plays an analogous role, though in the Möbius parametrization. However, it is clear that \mathcal{D}_n is much more complex, and, unlike the surface of independence, which has fixed dimension 1, its dimension increases with n . Table 3 provides the dimension of \mathcal{E}_n and of \mathcal{D}_n (which are the number of unlabeled graphs minus 1 and the number of connected unlabeled graphs, respectively) for all nodes of size up to 11. As it can be seen and is also simple to show, the ratio of the dimension of \mathcal{D}_n over that of \mathcal{E}_n converges to 1 very rapidly as n grows. The other striking difference is the fact that not all points on the manifold of dissociated distributions correspond to extremal exchangeable distributions. This is in contrast with the sequence case, in which every points on the surface of independence corresponds to an extremal exchangeable distributions, for each n . Thus, exchangeability in graphs (a special case of exchangeability for binary 2-dimensional array) is considerably more subtle and complicated than the sequence case.

Acknowledgments

Alessandro Rinaldo and Kayvan Sadeghi were partially supported by AFOSR grant FA9550-14-1-014.

References

- Akiyama, J., G. Exoo, and F. Harary (1979). The graphs with all induced subgraphs isomorphic. *The Bulletin of the Malaysian Mathematical Society Series 2* 2, 43–44.
- 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.
- Berg, C., J. P. R. Christensen, and P. Ressel (1976). Positive definite functions on Abelian semigroups. *Mathematische Annalen* 259, 253–274.

- Berg, C., J. P. R. Christensen, and P. Ressel (1984). *Harmonic Analysis on Semigroups*. New York: Springer-Verlag.
- Borgs, C., J. Chayes, L. Lovász, V. Sós, and K. Vesztegombi (2008). Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing. *Advances in Mathematics* 219, 1801–1851.
- 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. 36, 271–281.
- 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(2), 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.
- Erdős, P., L. Lovász, and J. Spencer (1979). Strong independence of graphcopy functions. In J. Bondy and U. Murty (Eds.), *Graph Theory and Related Topics*, New York, pp. 165–172. Academic Press.
- Freedman, D. (1977). A remark on the difference between sampling with and without replacement. *Journal of the American Statistical Association* 72, 681–681.
- Hoover, D. N. (1979). Relations on probability spaces and arrays of random variables. Preprint, Institute of Advanced Study, Princeton.
- Kallenberg, O. (2005). *Probabilistic Symmetries and Invariance Principles*. New York: Springer-Verlag.
- Kerns, G. J. and G. J. Székely (2006). DeFinetti’s theorem for abstract finite exchangeable sequences. *Journal of Theoretical Probability* 19, 589–608.
- Lauritzen, S., A. Rinaldo, and K. Sadeghi (2017). Random networks, graphical models, and exchangeability. arXiv:1701.08420.
- 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, Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Prague.
- 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.
- Rao, M. M. (1971). Projective limits of probability spaces. *Journal of Multivariate Analysis* 1, 28–57.
- Ressel, P. (2008). Exchangeability and semigroups. *Rendiconti di Matematica, Serie VII* (28), 63–81.

- Roverato, A., M. Lupparelli, and L. La Rocca (2013). Log-mean linear models for binary data. *Biometrika* 100, 485–494.
- 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.
- Stanley, P. R. (2011). *Enumerative Combinatorics* (2nd ed.). Cambridge, UK: Cambridge University Press.
- Volfovsky, A. and E. M. Airoidi (2016). Sharp total variation bounds for finitely exchangeable arrays. *Statistics and Probability Letters* 114, 54 – 59.