GRAPH LIMITS AND EXCHANGEABLE RANDOM GRAPHS

PERSI DIACONIS AND SVANTE JANSON

ABSTRACT. We develop a clear connection between deFinetti's theorem for exchangeable arrays (work of Aldous–Hoover–Kallenberg) and the emerging area of graph limits (work of Lovász and many coauthors). Along the way, we translate the graph theory into more classical probability.

1. INTRODUCTION

`S:intro`

DeFinetti's profound contributions are now woven into many parts of probability, statistics and philosophy. Here we show how developments from deFinetti's work on partial exchangeability have a direct link to the recent development of a limiting theory for large graphs. This introduction first recalls the theory of exchangeable arrays (Section 1.1). Then, the subject of graph limits is outlined (Section 1.2). Finally, the link between these ideas, which forms the bulk of this paper, is outlined (Section 1.3).

`secA`

1.1. Exchangeability, Partial Exchangeability and Exchangeable Arrays.

Let $\{X_i\}$ $1 \leq i < \infty$ be a sequence of binary random variables. They are *exchangeable* if

$$P(X_1 = e_1, \cdots, X_n = e_n) = P(X_1 = e_{\sigma(1)}, \cdots, X_n = e_{\sigma(n)})$$

for all n, permutations σ and all $e_i \in \{0, 1\}$. The celebrated representation theorem says

Theorem 1.1. (deFinetti) *If $\{X_i\}$ $1 \leq i < \infty$ is a binary exchangeable sequence, then*
  (i) *With probability 1, $X_\infty = \lim \frac{1}{n}(X_1 + \cdots + X_n)$ exists.*
  (ii) *If $\mu(A) = P\{X_\infty \in A\}$, then for all n, e_i $1 \leq i \leq n$*

$$P(X_1 = e_1, \cdots, X_n = e_n) = \int_0^1 x^s (1-x)^{n-s} \mu(dx) \qquad (1.1)$$

`eqP1`

for $s = e_1 + \cdots + e_n$.

It is natural to refine and extend deFinetti's theorem to allow more general observables (X_i with values in a Polish space) and other notions of symmetry (partial exchangeability). A definitive treatment of these developments is

---

given in [12]. Of interest here is the extension of deFinetti's theorem to two-dimensional arrays.

Definition. Let $\{X_{ij}\}$ $1 \leq i, j < \infty$ be binary random variables. They are *jointly exchangeable* if

$$P(X_{ij} = e_{ij}, 1 \leq i, j \leq n) = P(X_{ij} = e_{\sigma(i)\tau(j)} 1 \leq i, j \leq n)$$

for all n, all permutations σ, τ and all $e_{ij} \in \{0, 1\}$.

The question of two-dimensional versions of deFinetti's theorem under joint exchangeability arose from the statistical problems of two-way analysis of variance. Early workers expected a version of (1.1) with perhaps a two-dimensional integral. The probabilist David Aldous [1] and the logician Douglas Hoover [11] found that the answer is more complicated.

Define a random binary array $\{X_{ij}\}$ as follows: Let U_i, V_j $1 \leq i, j < \infty$ be independent and uniform in $[0, 1]$. Let $W(x, y)$ be a function from $[0, 1]^2$ to $[0, 1]$. Let X_{ij} be 1 or 0 as a $W(U_i, V_j)$ coin comes up heads or tails. Let P_W be the probability distribution of $\{X_{ij}\}$ $1 \leq i, j < \infty$. The family $\{X_{ij}\}$ is jointly exchangeable because of the symmetry of the construction. The Aldous–Hoover theorem says that any jointly exchangeable binary array is a mixture of such P_W:

Theorem 1.2. (Aldous–Hoover) *Let $X = \{X_{ij}\}$ $1 \leq i, j < \infty$ be a jointly exchangeable binary array. Then, there is a probability μ such that*

$$P\{X \in A\} = \int P_W(A)\mu(dW).$$

The uniqueness of μ resisted understanding; if $\widehat{W}$ is obtained from W by a measure-preserving change of each variable, clearly the associated process $\{\widehat{X}_{ij}\}$ has the same joint distribution as $\{X_{ij}\}$. Using model theory, Hoover was able to show that this was the only source of non-uniqueness. A 'probabilist's proof' was finally found by Kallenberg [12, Sect. 7.6 has details and references].

These results hold for higher dimensional arrays with X_{ij} taking values in a Polish space with minor change [12, Chap. 7]. The description above has not mentioned several elegant results of the theory. In particular, Kallenberg's 'spreadable' version of the theory replaces invariance under a group by invariance under subsequences. A variety of tail fields may be introduced to allow characterizing when W takes values in $\{0, 1\}$ [8, Sect. 4]. Much more general notions of partial exchangeability are studed in [9].

secB

1.2. **Graph Limits.** Large graphs, both random and deterministic, abound in applications. They arise from the internet, social networks, gene regulation, ecology and in mathematics. It is natural to seek an approximation theory: What does it mean for a sequence of graphs to converge? When can a large complex graph be approximated by a small graph?
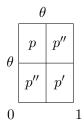
In a sequence of papers, Laszlo Lovász with coauthors (listed here in order of frequency) V. T. Sós, B. Szegedy, C. Borgs, J. Chayes, K. Vesztergombi, A. Schrijver, M. Freedman, have developed a beautiful, unifying limited theory. This sheds light on topics such as graph homomorphisms, Szemeredi's regularity lemma, quasi-random graphs, graph testing and external graph theory. Their theory has been developed for dense graphs (number of edges comparable with the square of number of vertices) but parallel theories for sparse graphs are beginning to emerge [4].

Roughly, a growing sequence of finite graphs G_n converges if, for any fixed graph F, the proportion of copies of F in G_n converges. Section 2 below has precise definitions.

Example 1.3. Define a probability distribution on graphs on n-vertices as follows. Flip a θ-coin for each vertex (dividing vertices into 'boys' and 'girls'). Connect two boys with probability p. Connect two girls with probability p'. Connect a boy and a girl with probability p''. Thus, if $p = p' = 1, p'' = 0$, we have a random bipartite graph. If $p = p' = 0, p'' = 1$, we have two disjoint complete graphs. If $p = p' = p''$, we have the Erdös–Renyi model. As n grows, these models generate a sequence of random graphs which converge almost surely to a limiting object described below.

If a sequence of graphs converges, what does it converge to? For exchangeable random graphs (defined below), there is a limiting object which may be thought of as a probability measure on infinite random graphs. Suppose $W(x, y) = W(y, x)$ is a function from $[0, 1]^2 \to [0, 1]$. Choose $\{U_i\}$ $1 \le i < \infty$ independent uniformly distributed random variables on $[0, 1]$. Form an infinite random graph by putting an edge from i to j with probability $W(U_i, U_j)$. This measure on graphs (or alternatively W) is the limiting object.

For the "boys and girls" example above, W may be pictured as

The theory developed shows that various properties of G_n can be well approximated by calculations with the limiting object. There is an elegant characterization of these 'continuous graph properties' with applications to algorithms for graph testing (Does this graph contain an Eulerian cycle?) or parameter estimation (What is an approximation to the size of the maximum cut?). There is a practical way to find useful approximations to a large graph by graphs of fixed size [5]. This paper contains a useful review of the current state of the theory with proofs and references.

We have sketched the theory for unweighted graphs. There are generalizations to graphs with weights on vertices and edges, to bipartite, directed and hypergraphs. The sketch leaves out many nice developments. For example, the useful cut metric between graphs [19] and connections to statistical physics.

secC

1.3. **Overview of the Present Paper.** There is an apparent similarity between the measure P_W of the Aldous–Hoover theorem and the limiting object W from graph limits. Roughly, working with symmetric W gives the graph limit theory; working with general W gives directed graphs. The main results of this paper make these connections precise.

Basic definitions are in Section 2 which introduces a probabilist's version of graph convergence equivalent to the definition using graph homomorphisms. Section 3 uses the well-established theory of weak convergence of a sequence of probability measures on a metric space to get properties of graph convergence. Section 4 carries things over to infinite graphs.

The main results appear in Section 5. This introduces exchangeable random graphs and gives a one-to-one correspondence between infinite exchangeable random graphs and the space of proper graph limits (Theorem 5.3). A useful characterization of the extreme points of the set of exchangeable random graphs is in Theorem 5.5. These results are translated to the equivalence between proper graph limits and the Aldous–Hoover theory in Section 6.

The equivalence involves symmetric $W(x,y)$ and a single permutation σ taking $W(U_i, U_j)$ to $W(U_{\sigma(i)}, U_{\sigma(j)})$. The original Aldous–Hoover theorem, with perhaps non-symmetric $W(x,y)$ and $W(U_i, V_j)$ to $W(U_{\sigma(i)}, V_{\tau(j)})$ translates to a limit theorem for bi-partite graphs. This is developed in Section 7. The extensions to weighted graphs are covered by allowing X_{ij} to take general values in the Aldous–Hoover theory. The extension to hypergraphs follows from the Aldous–Hoover theory for higher-dimensional arrays.

Despite these parallels, the theories have much to contribute to each other. The algorithmic, graph testing, Szemeredi partitioning perspective is new to exchangeability theory. Indeed, the "boys and girls" random graph was introduced to study the psychology of vision in Diaconis–Freedman (1981). As far as we know, its graph theoretic properties have not been studied. The various developments around shell-fields in exchangeability, which characterize $0/1$ $W(x,y)$, have yet to be translated into graph-theoretic terms.

2. Definitions and basic properties

Sdef

All graphs will be simple. Infinite graphs will be important in later sections, but will always be clearly stated to be infinite; otherwise, graphs will be finite. We denote the vertex and edge sets of a graph G by $V(G)$ and $E(G)$, and the numbers of vertices and edges by $v(G) := |V(G)|$ and $e(G) := |E(G)|$. We consider both labelled and unlabelled graphs; the labels will be the integers $1, \ldots, n$, where n is the number of vertices in the graph. A labelled graph is thus a graph with vertex set $[n] := \{1, \ldots, n\}$ for some $n \geq 1$; we let $\mathcal{L}_n$ denote the set of the $2^{\binom{n}{2}}$ labelled graphs on $[n]$ and let $\mathcal{L} := \bigcup_{n=1}^{\infty} \mathcal{L}_n$. An unlabelled graph can be regarded as a labelled graph where we ignore the labels; formally, we define $\mathcal{U}_n$, the set of unlabelled graphs of order n, as the quotient set $\mathcal{L}_n/\cong$ of labelled graphs modulo isomorphisms. We let $\mathcal{U} := \bigcup_{n=1}^{\infty} \mathcal{U}_n = \mathcal{L}/\cong$, the set of all unlabelled graphs.

Note that we can, and often will, regard a labelled graph as an unlabelled graph.

If G is an (unlabelled) graph and $v_1, \ldots, v_k$ is a sequence of vertices in G, then $G(v_1, \ldots, v_k)$ denotes the labelled graph with vertex set $[k]$ where we put an edge between i and j if v_i and v_j are adjacent in G. We allow the possibility that $v_i = v_j$ for some i and j. (In this case, there is no edge ij because there are no loops in G.)

We let $G[k]$, for $k \geq 1$, be the random graph $G(v_1, \ldots, v_k)$ obtained by sampling $v_1, \ldots, v_k$ uniformly at random among the vertices of G, with replacement. In other words, $v_1, \ldots, v_k$ are independent uniformly distributed random vertices of G.

For $k \leq v(G)$, we further let $G[k]'$ be the random graph $G(v'_1, \ldots, v'_k)$ where we sample $v'_1, \ldots, v'_k$ uniformly at random without replacement; the sequence $v'_1, \ldots, v'_k$ is thus a uniformly distributed random sequence of k distinct vertices.

If F and G are two graphs, we define, following [7] and [19], first assuming that F is labelled and with $k = v(F)$,

$$t(F, G) := \mathbb{P}\big(F \subseteq G[k]\big). \tag{2.1}$$ t

Note that both F and $G[k]$ are graphs on $[k]$, so the relation $F \subseteq G[k]$ is well-defined as containment of labelled graphs on the same vertex set, i.e. as $E(F) \subseteq E(G[k])$. Although the relation $F \subseteq G[k]$ may depend on the labelling of F, the probability in (2.1) does not, by symmetry, so $t(F, G)$ is really well defined for unlabelled F and G. (The definition by Lovász and Szegedy [19], and also in [7], is actually stated differently; they define $t(F, G)$ as the proportion of graph homomorphisms $F \to G$ among all mappings $V(F) \to V(G)$. This is evidently equivalent to (2.1).)

With F, G and k as in (2.1), we further define, again using the notation of [7] but stating the definitions in different but equivalent forms,

$$t_{\mathrm{inj}}(F, G) := \mathbb{P}\big(F \subseteq G[k]'\big) \tag{2.2}$$ tinj

and

$$t_{\mathrm{ind}}(F, G) := \mathbb{P}\big(F = G[k]'\big), \tag{2.3}$$ `tind`

provided F and G are (unlabelled) graphs with $v(F) \leq v(G)$. If $v(F) > v(G)$ we set $t_{\mathrm{inj}}(F, G) := t_{\mathrm{ind}}(F, G) := 0$.

Since the probability that a random sample $v_1, \ldots, v_k$ of vertices in G contains some repeated vertex is $\leq k^2/(2v(G))$, it follows that [19]

$$\big|t(F, G) - t_{\mathrm{inj}}(F, G)\big| \leq \frac{v(F)^2}{2v(G)}. \tag{2.4}$$ `a4`

Hence, when considering asymptotics with $v(G) \to \infty$, it does not matter whether we use t or t_{inj}. Moreover, if $F \in \mathcal{L}_k$, then, as pointed out in [7] and [19],

$$t_{\mathrm{inj}}(F, G) = \sum_{F' \in \mathcal{L}_k, \, F' \supseteq F} t_{\mathrm{ind}}(F, G) \tag{2.5}$$ `a4a`

and, by inclusion-exclusion,

$$t_{\mathrm{ind}}(F, G) = \sum_{F' \in \mathcal{L}_k, \, F' \supseteq F} (-1)^{e(F')-e(F)} t_{\mathrm{inj}}(F, G). \tag{2.6}$$ `a4b`

Hence, the two families $\{t_{\mathrm{inj}}(F, \cdot)\}_{F \in \mathcal{U}}$ and $\{t_{\mathrm{ind}}(F, \cdot)\}_{F \in \mathcal{U}}$ of graph functionals contain the same information and can replace each other.

The basic definition of Lovász and Szegedy [19] and Borgs, Chayes, Lovász, Sós and Vesztergombi [7] is that a sequence (G_n) of graphs converges if $t(F, G_n)$ converges for every graph F. We can express this by considering the map $\tau : \mathcal{U} \to [0, 1]^{\mathcal{U}}$ defined by

$$\tau(G) := (t(F, G))_{F \in \mathcal{U}} \in [0, 1]^{\mathcal{U}}. \tag{2.7}$$ `tau`

Then (G_n) converges if and only if $\tau(G_n)$ converges in $[0, 1]^{\mathcal{U}}$, equipped with the usual product topology. Note that $[0, 1]^{\mathcal{U}}$ is a compact metric space; as is well known, a metric can be defined by, for example,

$$d\big((x_F), (y_F)\big) := \sum_{i=0}^{\infty} 2^{-i} |x_{F_i} - y_{F_i}|, \tag{2.8}$$ `a5`

where $F_1, F_2, \ldots$ is some enumeration of all unlabelled graphs.

We define $\mathcal{U}^* := \tau(\mathcal{U}) \subseteq [0, 1]^{\mathcal{U}}$ to be the image of $\mathcal{U}$ under this mapping τ, and let $\overline{\mathcal{U}^*}$ be the closure of $\mathcal{U}^*$ in $[0, 1]^{\mathcal{U}}$. Thus $\overline{\mathcal{U}^*}$ is a compact metric space. (For explicit descriptions of the subset $\overline{\mathcal{U}^*}$ of $[0, 1]^{\mathcal{U}}$ as a set of graph functionals, see Lovász and Szegedy [19].)

As pointed out in [19] and [7] (in equivalent terminology), τ is not injective; for example, $\tau(K_{n,n})$ is the same for all complete bipartite graphs $K_{n,n}$. Nevertheless, as in [19] and [7], we can consider a graph G as an element of $\mathcal{U}^*$ by identifying G and $\tau(G)$ (thus identifying graphs with the same $\tau(G)$), and then convergence of (G_n) as defined above is equivalent to convergence in $\overline{\mathcal{U}^*}$. The limit is thus an element of $\overline{\mathcal{U}^*}$, but typically not a

graph in $\mathcal{U}^*$. The main result of Lovász and Szegedy [19] is a representation of the elements in $\overline{\mathcal{U}^*}$ to which we will return in Section 6.

Remark 2.1. As said above, $\overline{\mathcal{U}^*}$ is a compact metric space, and it can be given several equivalent metrics. One metric is the metric (2.8) inherited from $[0,1]^{\mathcal{U}}$, which for graphs becomes $d(G, G') = \sum_i 2^{-i}|t(F_i, G) - t(F_i, G')|$. Another metric, shown by Borgs, Chayes, Lovász, Sós and Vesztergombi [7] to be equivalent, is the cut-distance $\delta_\square$, see [7] for definitions.

The identification of graphs with the same image in $\mathcal{U}^*$ (i.e., with the same $t(F, \cdot)$ for all F) is sometimes elegant but at other times inconvenient. It can be avoided if we instead let $\mathcal{U}^+$ be the union of $\mathcal{U}$ and some one-point set $\{*\}$ and consider the mapping $\tau^+ : \mathcal{U} \to [0,1]^{\mathcal{U}^+} = [0,1]^{\mathcal{U}} \times [0,1]$ defined by

$$\tau^+(G) = \big(\tau(G),\, v(G)^{-1}\big). \tag{2.9}$$

Then τ^+ is injective, because if $\tau(G_1) = \tau(G_2)$ for two graphs G_1 and G_2 with the same number of vertices, then G_1 and G_2 are isomorphic and thus $G_1 = G_2$ as unlabelled graphs. (This can easily be shown directly: it follows from (2.1) that $G_1[k] \stackrel{\mathrm{d}}{=} G_2[k]$ for every k, which implies $G_1[k]' \stackrel{\mathrm{d}}{=} G_2[k]'$ for every $k \leq v(G_1) = v(G_2)$; now take $k = v(G_1)$. It is also a consequence of [7, Theorem 2.7 and Theorem 2.3 or Lemma 5.1].)

Consequently, we can identify $\mathcal{U}$ with its image $\tau^+(\mathcal{U}) \subseteq [0,1]^{\mathcal{U}^+}$ and define $\overline{\mathcal{U}} \subseteq [0,1]^{\mathcal{U}^+}$ as its closure. It is easily seen that a sequence (G_n) of graphs converges in $\overline{\mathcal{U}}$ if and only if either $v(G_n) \to \infty$ and (G_n) converges in $\overline{\mathcal{U}^*}$, or the sequence (G_n) is constant from some n_0 on. Hence, convergence in $\overline{\mathcal{U}}$ is essentially the same as the convergence considered by by Lovász and Szegedy [19], but without any identification of non-isomorphic graphs of different orders.

Alternatively, we can consider τ_{inj} or τ_{ind} defined by

$$\tau_{\mathrm{inj}}(G) := (t_{\mathrm{inj}}(F, G))_{F \in \mathcal{U}} \in [0,1]^{\mathcal{U}},$$
$$\tau_{\mathrm{ind}}(G) := (t_{\mathrm{ind}}(F, G))_{F \in \mathcal{U}} \in [0,1]^{\mathcal{U}}.$$

It is easy to see that both τ_{inj} and τ_{ind} are injective mappings $\mathcal{U} \to [0,1]^{\mathcal{U}}$. (If $\tau_{\mathrm{inj}}(F, G_1) = \tau_{\mathrm{inj}}(F, G_2)$ for all F, we take $F = G_1$ and $F = G_2$ and conclude $G_1 = G_2$, using our special definition above when $v(F) > v(G)$.) Hence, we can again identify $\mathcal{U}$ with its image and consider its closure $\overline{\mathcal{U}}$ in $[0,1]^{\mathcal{U}}$. Moreover, using (2.4), (2.5), and (2.6), it is easily shown that if (G_n) is a sequence of unlabelled graphs, then

$$\tau^+(G_n) \text{ converges} \iff \tau_{\mathrm{ind}}(G_n) \text{ converges} \iff \tau_{\mathrm{inj}}(G_n) \text{ converges.}$$

Hence, the three compactifications $\overline{\tau^+(\mathcal{U})}$, $\overline{\tau_{\mathrm{inj}}(\mathcal{U})}$, $\overline{\tau_{\mathrm{ind}}(\mathcal{U})}$ are homeomorphic and we can use any of them for $\overline{\mathcal{U}}$. We let $\mathcal{U}_\infty := \overline{\mathcal{U}} \setminus \mathcal{U}$; this is the set of all limit objects of sequences (G_n) in $\mathcal{U}$ with $v(G_n) \to \infty$. (I.e., it is the set of all proper graph limits.)

We will in the sequel prefer to use $\overline{\mathcal{U}}$ rather than $\overline{\mathcal{U}^*}$, thus not identifying some graphs of different orders, nor identifying finite graphs with some limit objects in $\mathcal{U}_\infty$.

For every fixed graph F, the functions $t(F,\cdot)$, $t_{\mathrm{inj}}(F,\cdot)$ and $t_{\mathrm{ind}}(F,\cdot)$ have unique continuous extensions to $\overline{\mathcal{U}}$, for which we use the same notation. We similarly extend $v(\cdot)^{-1}$ continuously to $\overline{\mathcal{U}}$ by defining $v(G) = \infty$ and thus $v(G)^{-1} = 0$ for $G \in \mathcal{U}_\infty := \overline{\mathcal{U}} \setminus \mathcal{U}$. Then (2.4), (2.5) and (2.6) hold for all $G \in \overline{\mathcal{U}}$, where (2.4) means that

$$t_{\mathrm{inj}}(F,G) = t(F,G), \qquad G \in \mathcal{U}_\infty. \tag{2.10}$$

<div style="text-align:right">a4x</div>

Note that $\overline{\mathcal{U}}$ is a compact metric space. Different, equivalent, metrics are given by the embeddings τ^+, τ_{inj}, τ_{ind} into $[0,1]^{\mathcal{U}^+}$ and $[0,1]^{\mathcal{U}}$. Another equivalent metric is, by Remark 2.1 and the definition of τ^+, $\delta_\square(G_1, G_2) + |v(G_1)^{-1} - v(G_2)^{-1}|$.

We summarize the results above on convergence.

Theorem 2.1. *A sequence (G_n) of graphs converges in the sense of Lovász and Szegedy [19] if and only if it converges in the compact metric space $\overline{\mathcal{U}^*}$. Moreover, if $v(G_n) \to \infty$, the sequence (G_n) converges in this sense if and only if it converges in $\overline{\mathcal{U}}$.*

<div style="text-align:left">T1</div>

The projection $\pi : [0,1]^{\mathcal{U}^+} = [0,1]^{\mathcal{U}} \times [0,1] \to [0,1]^{\mathcal{U}}$ maps $\tau^+(G)$ to $\tau(G)$ for every graph G, so by continuity it maps $\overline{\mathcal{U}}$ into $\overline{\mathcal{U}^*}$. For graph $G \in \mathcal{U}$, $\pi(G) = \tau(G)$ is the object in $\overline{\mathcal{U}^*}$ corresponding to G considered above, and we will in the sequel denote this object by $\pi(G)$; recall that this projection $\mathcal{U} \to \overline{\mathcal{U}^*}$ is not injective. (We thus distinguish between a graph G and its "ghost" $\pi(G)$ in $\overline{\mathcal{U}^*}$. Recall that when graphs are considered as elements of $\overline{\mathcal{U}^*}$ as in [19] and [7], certain graphs are identified with each other; we avoid this.) On the other hand, an element G of $\overline{\mathcal{U}}$ is by definition determined by $\tau(G)$ and $v(G)^{-1}$, cf. (2.9), so the restriction $\pi : \mathcal{U}_n \to \overline{\mathcal{U}^*}$ is injective for each $n \leq \infty$. In particular, $\pi : \mathcal{U}_\infty \to \overline{\mathcal{U}^*}$ is injective. Moreover, this map is surjective because every element $G \in \overline{\mathcal{U}^*}$ is the limit of some sequence (G_n) of graphs in $\mathcal{U}$ with $v(G_n) \to \infty$; by Theorem 2.1, this sequence converges in $\overline{\mathcal{U}}$ to some element G', and then $\pi(G') = G$. Since $\mathcal{U}_\infty$ is compact, the restriction of π to $\mathcal{U}_\infty$ is thus a homeomorphism, and we have the following theorem, saying that we can identify the set $\mathcal{U}_\infty$ of proper graph limits with $\overline{\mathcal{U}^*}$.

<div style="text-align:left">or "shadow"?</div>

Theorem 2.2. *The projection π maps the set $\mathcal{U}_\infty := \overline{\mathcal{U}} \setminus \mathcal{U}$ of proper graph limits homeomorphically onto $\overline{\mathcal{U}^*}$.*

<div style="text-align:left">Tcuoo</div>

3. Convergence of random graphs

<div style="text-align:left">Sconv</div>

A *random unlabelled graph* is a random element of $\mathcal{U}$ (with any distribution; we do not imply any particular model). We consider convergence of a sequence (G_n) of random unlabelled graphs in the larger space $\mathcal{U}$; recall

that this is a compact metric space so we may use the general theory set forth in, for example, Billingsley [2].

We use the standard notations $\xrightarrow{\mathrm{d}}$, $\xrightarrow{\mathrm{p}}$, $\xrightarrow{\mathrm{a.s.}}$ for convergence in distribution, probability, and alsmost surely, respectively. We will only consider the case when $v(G_n) \to \infty$, at least in probability. (The reader may think of the case when G_n has n vertices, although that is not necessary in general.)

We begin with convergence in distribution.

T2 **Theorem 3.1.** *Let G_n, $n \geq 1$, be random unlabelled graphs and assume that $v(G_n) \xrightarrow{\mathrm{p}} \infty$. The following are equivalent, as $n \to \infty$.*

T2a   (i) *$G_n \xrightarrow{\mathrm{d}} \Gamma$ for some random $\Gamma \in \overline{\mathcal{U}}$.*
T2b   (ii) *For every finite family $F_1, \ldots, F_m$ of (non-random) graphs, the random variables $t(F_1, G_n), \ldots, t(F_m, G_n)$ converge jointly in distribution.*
T2c   (iii) *For every (non-random) $F \in \mathcal{U}$, the random variables $t(F, G_n)$ converge in distribution.*
T2d   (iv) *For every (non-random) $F \in \mathcal{U}$, the expectations $\mathbb{E}\, t(F, G_n)$ converge.*

If these properties hold, then the limits in (ii), (iii) *and* (iv) *are $\left(t(F_i, \Gamma)\right)_{i=1}^{m}$, $t(F, \Gamma)$ and $\mathbb{E}\, t(F, \Gamma)$, respectively. Furthermore, $\Gamma \in \mathcal{U}_\infty$ a.s.*

The same results hold if t is replaced by t_{inj} or t_{ind}.

Proof. (i) $\Longleftrightarrow$ (ii). Since $\overline{\mathcal{U}}$ is a closed subset of $[0,1]^{\mathcal{U}^+}$, convergence in distribution in $\overline{\mathcal{U}}$ is equivalent to convergence of $\tau^+(G_n) = \left((t(F, G_n))_{F \in \mathcal{U}},\, v(G_n)^{-1}\right)$ in $[0,1]^{\mathcal{U}^+}$, Since we assume $v(G_n)^{-1} \xrightarrow{\mathrm{p}} 0$, this is equivalent to convergence of $(t(F, G_n))_{F \in \mathcal{U}}$ in $[0,1]^{\mathcal{U}}$ [2, Theorem 4.4], which is equivalent to convergence in distribution of all finite families $(t(F_i, G_n))_{i=1}^{m}$.

(ii) $\Longrightarrow$ (iii). Trivial.

(iii) $\Longrightarrow$ (iv). Immediate, since t is bounded (by 1).

(iv) $\Longrightarrow$ (ii). Let $F_1, \ldots, F_m$ be fixed graphs and let $\ell_1, \ldots, \ell_m$ be positive integers. Let F be the disjoint union of ℓ_i copies of F_i, $i = 1, \ldots, m$. Then, for every $G \in \mathcal{U}$, from the definition of t,

$$t(F, G) = \prod_{i=1}^{m} t(F_i, G)^{\ell_i},$$

and hence

$$\mathbb{E} \prod_{i=1}^{m} t(F_i, G)^{\ell_i} = \mathbb{E}\, t(F, G). \qquad (3.1) \quad \boxed{\text{b3}}$$

Consequently, if (iv) holds, then every joint moment $\mathbb{E} \prod_{i=1}^{m} t(F_i, G)^{\ell_i}$ of $t(F_1, G_n), \ldots, t(F_m, G_n)$ converges. Since $t(F_i, G_n)$ are bounded (by 1), this implies joint convergence in distribution by the method of moments.

The identification of the limits is immediate. Since $v(G_n) \xrightarrow{\mathrm{p}} \infty$, (i) implies that $v(\Gamma) = \infty$ a.s., and thus $\Gamma \in \mathcal{U}_\infty$.

Finally, it follows from (2.4), (2.5) and (2.6) that we can replace t by t_{inj} or t_{ind} in (ii) and (iv), and the implications (ii) $\implies$ (iii) and (iii) $\implies$ (iv) are immediate for t_{inj} and t_{ind} too. $\qquad\square$

Specializing to the case of a non-random limit $G \in \mathcal{U}_\infty$, we obtain the corresponding result for convergence in probability.

C2

Corollary 3.2. *Let G_n, $n \geq 1$, be random unlabelled graphs such that $v(G_n) \xrightarrow{\mathrm{p}} \infty$, and let $G \in \mathcal{U}_\infty$. The following are equivalent, as $n \to \infty$.*

C2a   *(i) $G_n \xrightarrow{\mathrm{p}} G$.*
C2c   *(ii) $t(F, G_n) \xrightarrow{\mathrm{p}} t(F, G)$ for every (non-random) $F \in \mathcal{U}$.*
C2d   *(iii) $\mathbb{E}\, t(F, G_n) \to t(F, G)$ for every (non-random) $F \in \mathcal{U}$.*

The same result holds if t is replaced by t_{inj} or t_{ind}.

Note further that under the same assumptions, it follows directly from Theorem 2.1 that $G_n \xrightarrow{\mathrm{a.s.}} G$ if and only if $t(F, G_n) \xrightarrow{\mathrm{a.s.}} t(F, G)$ for every $F \in \mathcal{U}$.

We observe another corollary to Theorem 3.1 (and its proof).

C2aa

Corollary 3.3. *If Γ is a random element of $\mathcal{U}_\infty = \overline{\mathcal{U}} \setminus \mathcal{U} \cong \overline{\mathcal{U}^*}$, then, for every sequence $F_1, \ldots, F_m$ of graphs, possibly with repetitions,*

$$\mathbb{E}\prod_{i=1}^m t(F_i, \Gamma) = \mathbb{E}\, t\left(\oplus_{i=1}^m F_i, \Gamma\right), \qquad (3.2)$$

b3x

where $\oplus_{i=1}^m F_i$ denotes the disjoint union of $F_1, \ldots, F_m$. As a consequence, the distribution of Γ is uniquely determined by the numbers $\mathbb{E}\, t(F, \Gamma)$, $F \in \mathcal{U}$. Alternatively, the distribution of Γ is uniquely determined by the numbers $\mathbb{E}\, t_{\mathrm{ind}}(F, \Gamma)$, $F \in \mathcal{U}$.

Proof. Since $\mathcal{U}$ is dense in $\overline{\mathcal{U}} \supseteq \mathcal{U}_\infty$, there exists random unlabelled graphs G_n such that $G_n \xrightarrow{\mathrm{a.s.}} \Gamma$. In particular, $G_n \xrightarrow{\mathrm{d}} \Gamma$ and $v(G_n) \xrightarrow{\mathrm{p}} \infty$ (in fact, we may assume $v(G_n) = n$), so Theorem 3.1 and its proof apply, and (3.2) follows from (3.1) applied to G_n by letting $n \to \infty$.

For the final statement, note that (3.2) shows that the expectations $\mathbb{E}\, t(F, \Gamma)$, $F \in \mathcal{U}$, determine all moments $\mathbb{E}\prod_{i=1}^m t(F_i, \Gamma)$, and thus the joint distribution of $t(F, \Gamma)$, $F \in \mathcal{U}$, which is the same as the distribution of $\tau(\Gamma) = \left(t(F, \Gamma)\right)_{F \in \mathcal{U}} \in [0, 1]^{\mathcal{U}}$, and we have defined $\mathcal{U}_\infty$ such that we identify Γ and $\tau(\Gamma)$. Finally, the numbers $\mathbb{E}\, t_{\mathrm{ind}}(F, \Gamma)$, $F \in \mathcal{U}$, determine all $\mathbb{E}\, t(F, \Gamma)$ by (2.5), recalling that $t_{\mathrm{inj}}(F, \Gamma) = t(F, \Gamma)$ by (2.10). $\qquad\square$

Remark 3.1. The numbers $\mathbb{E}\, t(F, \Gamma)$ for a random $\Gamma \in \mathcal{U}_\infty$ thus play a role similar to the one played by moments for a random variable. (And the relation between $\mathbb{E}\, t(F, \Gamma)$ and $\mathbb{E}\, t_{\mathrm{ind}}(F, \Gamma)$ has some resemblance to the relation between moments and cumulants.)

4. Convergence to infinite graphs

Sinfinite

We will in this section consider also labelled *infinite* graphs with the vertex set $\mathbb{N} = \{1, 2, \dots\}$. Let $\mathcal{L}_\infty$ denote the set of all such graphs. These graphs are determined by their edge sets, so $\mathcal{L}_\infty$ can be identified with the power set $\mathcal{P}(E(K_\infty))$ of all subsets of the edge set $E(K_\infty)$ of the complete infinite graph K_∞, and thus with the infinite product set $\{0, 1\}^{E(K_\infty)}$. We give this space, and thus $\mathcal{L}_\infty$, the product topology. Hence, $\mathcal{L}_\infty$ is a compact metric space.

It is sometimes convenient to regard $\mathcal{L}_n$ for a finite n as a subset of $\mathcal{L}_\infty$: we can identify graphs in $\mathcal{L}_n$ and $\mathcal{L}_\infty$ with the same edge set. In other words, if $G \in \mathcal{L}_n$ is a graph with vertex set $[n]$, we add an infinite number of isolated vertices $n + 1, n + 2, \dots$ to obtain a graph in $\mathcal{L}_\infty$.

Conversely, if $H \in \mathcal{L}_\infty$ is an infinite graph, we let $H|_{[n]} \in \mathcal{L}_n$ be the induced subgraph of H with vertex set $[n]$.

If G is a (finite) graph, let $\widehat{G}$ be the random labelled graph obtained by a random labelling of the vertices of G by the numbers $1, \dots, v(G)$. (If G is labelled, we thus ignore the labels and randomly relabel.) Thus $\widehat{G}$ is a random finite graph with the same number of vertices as G, but as just said, we can (and will) also regard $\widehat{G}$ as a random graph in $\mathcal{L}_\infty$.

We use the same notation $\widehat{G}$ also for a random (finite) graph G given a random labelling.

TC1
Theorem 4.1. *Let (G_n) be a sequence of random graphs in $\mathcal{U}$ and assume that $v(G_n) \xrightarrow{\mathrm{p}} \infty$. Then the following are equivalent.*

(i) $G_n \xrightarrow{\mathrm{d}} \Gamma$ *in $\overline{\mathcal{U}}$ for some random $\Gamma \in \overline{\mathcal{U}}$.*

(ii) $\widehat{G_n} \xrightarrow{\mathrm{d}} H$ *in $\mathcal{L}_\infty$ for some random $H \in \mathcal{L}_\infty$.*

If these hold, then $\mathbb{P}(H|_{[k]} = F) = \mathbb{E}\, t_{\mathrm{ind}}(F, \Gamma)$ for every $F \in \mathcal{L}_k$. Furthermore, $\Gamma \in \mathcal{U}_\infty$ a.s.

Proof. Let G be a labelled graph and consider the graph $\widehat{G}|_{[k]}$, assuming $k \leq v(G)$. This random graph equals $G[k]' = G(v'_1, \dots, v'_k)$, where $v'_1, \dots, v'_k$ are k vertices sampled at random without replacement as in Section 2. Hence, by (2.3), for every $F \in \mathcal{L}_k$,

$$\mathbb{P}(\widehat{G}|_{[k]} = F) = t_{\mathrm{ind}}(F, G), \qquad \text{if } k \leq v(G).$$

Applied to the random graph G_n, this yields

$$\mathbb{E}\, t_{\mathrm{ind}}(F, G_n) \leq \mathbb{P}(\widehat{G_n}|_{[k]} = F) \leq \mathbb{E}\, t_{\mathrm{ind}}(F, G_n) + P\big(v(G_n) < k\big). \qquad (4.1) \quad \boxed{\text{c3}}$$

By assumption, $P\left(v(G_n) < k\right) \to 0$ as $n \to \infty$, and it follows from (4.1) and Theorem 3.1 that $G_n \xrightarrow{\mathrm{d}} \Gamma$ in $\overline{\mathcal{U}}$ if and only if

$$\mathbb{P}(\widehat{G_n}|_{[k]} = F) \to \mathbb{E}\, t_{\mathrm{ind}}(F, \Gamma) \qquad (4.2) \quad \boxed{\text{c3a}}$$

for every $k \geq 1$ and every $F \in \mathcal{L}_k$.

Since $\mathcal{L}_k$ is a finite set, (4.2) says that, for every k, $\widehat{G_n}|_{[k]} \xrightarrow{\mathrm{d}} H_k$ for some random graph $H_k \in \mathcal{L}_k$ with $\mathbb{P}(H_k = F) = \mathbb{E}\, t_{\mathrm{ind}}(F, \Gamma)$ for $F \in \mathcal{L}_k$. Since $\mathcal{L}_\infty$ has the product topology, this implies $\widehat{G_n} \xrightarrow{\mathrm{d}} H$ in $\mathcal{L}_\infty$ for some random $H \in \mathcal{L}_\infty$ with $H|_{[k]} \overset{\mathrm{d}}{=} H_k$.

Conversely, if $\widehat{G_n} \xrightarrow{\mathrm{d}} H$ in $\mathcal{L}_\infty$, then $\widehat{G_n}|_{[k]} \xrightarrow{\mathrm{d}} H|_{[k]}$ so the argument above shows that

$$\mathbb{E}\, t_{\mathrm{ind}}(F, G_n) = \mathbb{P}(\widehat{G_n}|_{[k]} = F) + o(1) \to \mathbb{P}(H|_{[k]} = F)$$

as $n \to \infty$, for every $F \in \mathcal{L}_k$, and Theorem 3.1 yields the existence of some random $\Gamma \in \mathcal{U}_\infty \subset \overline{\mathcal{U}}$ with $G_n \xrightarrow{\mathrm{d}} \Gamma$ and $\mathbb{E}\, t_{\mathrm{ind}}(F, \Gamma) = \mathbb{P}(H|_{[k]} = F)$. $\qquad \square$

5. Exchangeable random graphs

$\boxed{\text{Sexch}}$

finitary?

A permutation $\sigma : \mathbb{N} \to \mathbb{N}$ is *finite* if $\sigma(n) = n$ for all sufficiently large n.

Definition. A random infinite graph $H \in \mathcal{L}_\infty$ is *exchangeable* if its distribution is invariant under every finite permutation of the vertices.

$\boxed{\text{LE}}$ **Lemma 5.1.** *Let H be a random infinite graph in $\mathcal{L}_\infty$. Then the following are equivalent.*

    (i) *H is exchangeable.*
    (ii) *$H|_{[k]}$ has a distribution invariant under all permutations of $[k]$, for every $k \geq 1$.*
    (iii) *$\mathbb{P}(H|_{[k]} = F)$ depends only on the isomorphism type of F, and can thus be seen as a function of F as an unlabelled graph in $\mathcal{U}_k$, for every $k \geq 1$.*

Proof. (i) $\Longrightarrow$ (ii). Immediate.

(ii) $\Longrightarrow$ (i). If σ is a finite permutation of $\mathbb{N}$, then σ restricts to a permutation of $[k]$ for every large k, and it follows that if $H \circ \sigma$ is H with the vertices permuted by σ, then, for all large k $H \circ \sigma|_{[k]} = H|_{[k]} \circ \sigma \overset{\mathrm{d}}{=} H|_{[k]}$, which implies $H|_{[k]} \overset{\mathrm{d}}{=} H$.

(ii) $\Longleftrightarrow$ (iii). Trivial. $\qquad \square$

$\boxed{\text{TC1E}}$ **Theorem 5.2.** *The limit H is Theorem 4.1 is exchangeable.*

Proof. H satisfies Lemma 5.1(iii). $\qquad \square$

Moreover, Theorem 4.1 implies the following connection with random elements of $\mathcal{U}_\infty$.

$\boxed{\text{TE}}$ **Theorem 5.3.** *There is a one-to-one correspondence between distributions of random elements $\Gamma \in \mathcal{U}_\infty$ (or $\overline{\mathcal{U}^*}$) and distributions of exchangeable random infinite graphs $H \in \mathcal{L}_\infty$ given by*

$$\mathbb{E}\, t_{\mathrm{ind}}(F, \Gamma) = \mathbb{P}(H|_{[k]} = F) \tag{5.1}$$

$\boxed{\text{e2a}}$

for every $k \geq 1$ and every $F \in \mathcal{L}_k$, or, equivalently,

$$\mathbb{E}\, t(F, \Gamma) = \mathbb{P}(H \supset F) \qquad (5.2)$$ `e2b`

for every $F \in \mathcal{L}$. Furthermore, $H|_{[n]} \xrightarrow{\mathrm{d}} \Gamma$ in $\overline{\mathcal{U}}$ as $n \to \infty$.

Proof. Note first that (5.1) and (5.2) are equivalent by (2.5) and (2.6), since $t(F, \Gamma) = t_{\mathrm{inj}}(F, \Gamma)$ by (2.10), and $H \supset F$ if and only if $H|_{[k]} \supseteq F$ when $F \in \mathcal{L}_k$.

Suppose that Γ is a random element of $\mathcal{U}_\infty \subset \overline{\mathcal{U}}$. Since $\mathcal{U}$ is dense in $\overline{\mathcal{U}}$, there exist (as in the proof of Corollary 3.3) random unlabelled graphs G_n such that $G_n \xrightarrow{\mathrm{a.s.}} \Gamma$ in $\overline{\mathcal{U}}$ and thus $v(G_n) \xrightarrow{\mathrm{a.s.}} \infty$ and $G_n \xrightarrow{\mathrm{d}} \Gamma$. Hence, Theorems 4.1 and 5.2 show that $\widehat{G_n} \xrightarrow{\mathrm{d}} H$ for some random exchangeable infinite graph H satisfying (5.1). Furthermore, (5.1) determines the distribution of $H|_{[k]}$ for every k, and thus the distribution of k.

Conversely, if H is an exchangeable random infinite graph, let $G_n = H|_{[n]}$. By Lemma 5.1(ii), the distribution of each G_n is invariant under permutations of the vertices, so if $\widehat{G_n}$ is G_n with a random (re)labelling, we have $\widehat{G_n} \overset{\mathrm{d}}{=} G_n$. Since $G_n \xrightarrow{\mathrm{d}} H$ in $\mathcal{L}_\infty$ (because $\mathcal{L}_\infty$ has a product topology), we thus have $\widehat{G_n} \xrightarrow{\mathrm{d}} H$ in $\mathcal{L}_\infty$, so Theorem 4.1 applies and shows the existence of a random $\Gamma \in \mathcal{U}_\infty$ such that $G_n \xrightarrow{\mathrm{d}} \Gamma$ and (5.1) holds. Finally (5.1) determines the distribution of Γ by Corollary 3.3. $\qquad\square$

`Ras` **Remark 5.1.** Moreover, $H|_{[n]}$ converges a.s. to some random variable $\Gamma \in \mathcal{U}_\infty$, because $t_{\mathrm{ind}}(F, H|_{[n]})$, $n \geq v(F)$, is a reverse martingale for every $F \in \Gamma$. Alternatively, this follows by concentration estimates from the representation in Section 6, see Lovász and Szegedy [19, Theorem 2.5].

I hope you agree it is a reverse martingale. Do we need more details?

`CE` **Corollary 5.4.** *There is a one-to-one correspondence between elements Γ of $\mathcal{U}_\infty \cong \overline{\mathcal{U}}^*$ and extreme points of the set of distributions of exchangeable random infinite graphs $H \in \mathcal{L}_\infty$. This correspondence is given by*

$$t(F, \Gamma) = \mathbb{P}(H \supset F) \qquad (5.3)$$ `ce`

for every $F \in \mathcal{L}$. Furthermore, $H|_{[n]} \xrightarrow{\mathrm{a.s.}} \Gamma$ in $\overline{\mathcal{U}}$ as $n \to \infty$.

Proof. The extreme points of the set of distributions on $\mathcal{U}_\infty$ are the point masses, which are in one-to-one correspondence with the elements of $\mathcal{U}_\infty$. $\quad\square$

We can characterize these extreme point distributions of exchangeable random infinite graphs as follows.

`TE2` **Theorem 5.5.** *Let H be an exchangeable random infinite graph. Then the following are equivalent.*

`te2a`   (i) *The distribution of H is an extreme point in the set of exchangeable distributions in $\mathcal{L}_\infty$.*

te2b'   (ii) *If F_1 and F_2 are two (finite) graphs with disjoint vertex sets $V(F_1)$, $V(F_2) \subset \mathbb{N}$, then*

$$\mathbb{P}(H \supset F_1 \cup F_2) = \mathbb{P}(H \supset F_1)\,\mathbb{P}(H \supset F_2).$$

te2b   (iii) *The restrictions $H|_{[k]}$ and $H|_{[k+1,\infty)}$ are independent for every k.*

te2c   (iv) *Let $\mathcal{F}_n$ be the σ-field generated by $H|_{[n,\infty)}$. Then the tail σ-field $\bigcap_{n=1}^{\infty} \mathcal{F}_n$ is trivial, i.e., contains only events with probability 0 or 1.*

Proof. (i) $\implies$ (ii). By Corollary 5.4, H corresponds to some (non-random) $\Gamma \in \mathcal{U}_\infty$ such that

$$\mathbb{P}(H \supset F) = t(F, \Gamma) \qquad\qquad (5.4) \quad \boxed{\texttt{e6}}$$

for every $F \in \mathcal{L}$. We have defined $\mathcal{L}$ such that a graph $F \in \mathcal{L}$ is labelled by $1, \ldots, v(F)$, but both sides of (5.4) are invariant under relabelling of F by arbitrary positive integers; the left hand side because H is exchangeable and the right hand side because $t(F, \Gamma)$ only depends on F as an unlabelled graph. Hence (5.4) holds for every finite graph F with $V(F) \subset \mathbb{N}$.

Furthermore, since Γ is non-random, Corollary 3.3 yields $t(F_1 \cup F_2, \Gamma) = t(F_1, \Gamma)t(F_2, \Gamma)$. Hence,

$$\mathbb{P}(H \supset F_1 \cup F_2) = t(F_1 \cup F_2, \Gamma) = t(F_1, \Gamma)t(F_2, \Gamma) = \mathbb{P}(H \supset F_1)\,\mathbb{P}(H \supset F_2).$$

(ii) $\implies$ (iii). By inclusion-exclusion, as for (2.3), (ii) implies that if $1 \le k < l < \infty$, then for any graphs F_1 and F_2 with $V(F_1) = \{1, \ldots, k\}$ and $V(F_2) = \{k+1, \ldots, k+l\}$, the events $H|_{[k]} = F_1$ and $H|_{\{k+1, \ldots, l\}} = F_2$ are independent. Hence $H|_{[k]}$ and $H|_{\{k, \ldots, l\}}$ are independent for every $l > k$, and the result follows.

(iii) $\implies$ (iv). Suppose A is an event in the tail σ-field $\bigcap_{n=1}^{\infty} \mathcal{F}_n$. Let $\mathcal{F}_n^*$ be the σ-field generated by $H|_{[n]}$. By (iii), A is independent of $\mathcal{F}_n^*$ for every n, and thus of the σ-field $\mathcal{F}$ generated by $\bigcup \mathcal{F}_n^*$, which equals the σ-field $\mathcal{F}_1$ generated by H. However, $A \in \mathcal{F}_1$, so A is independent of itself and thus $\mathbb{P}(A) = 0$ or 1.

(iv) $\implies$ (i). Let $F \in \mathcal{L}_k$ for some k and let F_n be F with all vertices shifted by n. Consider the two indicators $I = \mathbf{1}[H \supseteq F]$ and $I_n = \mathbf{1}[H \supseteq F_n]$. Since I_n is $\mathcal{F}_n$-measurable,

$$\mathbb{P}(H \supset F \cup F_n) = \mathbb{E}(II_n) = \mathbb{E}\big(\mathbb{E}(I \mid \mathcal{F}_n)I_n\big). \qquad (5.5) \quad \boxed{\texttt{e70}}$$

Moreover, $\mathbb{E}(I \mid \mathcal{F}_n)$, $n = 1, 2, \ldots$, is a reverse martingale, and thus

$$\mathbb{E}(I \mid \mathcal{F}_n) \to \mathbb{E}\Big(I \mid \bigcup_{n=1}^{\infty} \mathcal{F}_n\Big) = \mathbb{E}\,I$$

a.s. Hence, $\big(\mathbb{E}(I \mid \mathcal{F}_n) - \mathbb{E}\,I\big)I_n \to 0$ a.s., and by dominated convergence

$$\mathbb{E}\Big(\big(\mathbb{E}(I \mid \mathcal{F}_n) - \mathbb{E}\,I\big)I_n\Big) \to 0.$$

Consequently, (5.5) yields

$$\mathbb{P}(H \supset F \cup F_n) = \mathbb{E}\,I\,\mathbb{E}\,I_n + o(1) = \mathbb{P}(H \supset F)\,\mathbb{P}(H \supset F_n) + o(1).$$

Moreover, since H is exchangeable, $\mathbb{P}(H \supset F \cup F_n)$ (for $n \geq v(F)$) and $\mathbb{P}(H \supset F_n)$ do not depend on n, and we obtain as $n \to \infty$

$$\mathbb{P}(H \supset F \cup F_k) = \mathbb{P}(H \supset F)^2. \qquad (5.6) \quad \boxed{\texttt{e7}}$$

Let Γ be a random element of $\mathcal{U}_\infty$ corresponding to H as in Theorem 5.3. By (5.2) and (3.2), (5.6) can be written

$$\mathbb{E}\, t(F,\Gamma)^2 = \big(\mathbb{E}\, t(F,\Gamma)\big)^2.$$

Hence the random variable $t(F,\Gamma)$ has variance 0 so it is a.s. constant. Since this holds for every $F \in \mathcal{L}$, it follows that Γ is a.s. constant, i.e., we can take Γ non-random, and (i) follows by Corollary 5.4. $\qquad\square$

$\boxed{\texttt{Srep}}$

6. Representations of graph limits and exchangeable graphs

As said in the introduction, the exchangeable infinite random graphs were characterized by Aldous [1], see also Kallenberg [12], and the graph limits in $\mathcal{U}_\infty \cong \overline{\mathcal{U}^*}$ were characterized in a very similar way by Lovász and Szegedy [19]. We can now make the connection between these two characterizations explicit.

Let $\mathcal{W}$ be the set of all measurable functions $W : [0,1]^2 \to [0,1]$ and let $\mathcal{W}_{\mathsf{s}}$ be the subset of symmetric functions. For every $W \in \mathcal{W}_{\mathsf{s}}$, we define an infinite random graph $G(\infty, W) \in \mathcal{L}_\infty$ as follows: we first choose a sequence $X_1, X_2, \ldots$ of i.i.d. random variables uniformly distributed on $[0,1]$, and then, given this sequence, for each pair (i,j) with $i < j$ we draw an edge ij with probability $W(X_i, X_j)$, independently for all pairs (i,j) with $i < j$ (conditionally given $\{X_i\}$). Further, let $G(n, W)$ be the restriction $G(\infty, W)|_{[n]}$, which is obtained by the same construction with a finite sequence $X_1, \ldots, X_n$.

It is evident that $G(\infty, W)$ is an exchangeable infinite random graph, and the result by Aldous [1] is that every exchangeable infinite random graph is obtained as a mixture of such $G(\infty, W)$; in other words as $G(\infty, W)$ with a random W.

Considering again a deterministic $W \in \mathcal{W}_{\mathsf{s}}$, it is evident that Theorem 5.5(ii) holds, and thus Theorem 5.5 and Corollary 5.4 show that $G(\infty, W)$ corresponds to an element $\Gamma_W \in \mathcal{U}_\infty$. Moreover, by Theorem 5.3 and Remark 5.1, $G(n, W) \to \Gamma_W$ a.s. as $n \to \infty$, and (5.3) shows that if $F \in \mathcal{L}_k$, then

$$t(F, \Gamma_W) = \mathbb{P}\big(F \subseteq G(k, W)\big) = \int_{[0,1]^k} \prod_{ij \in \mathbb{E}(F)} W(x_i, x_j)\, \mathrm{d}x_1 \ldots \mathrm{d}x_k. \quad (6.1) \quad \boxed{\texttt{tuw}}$$

The main result of Lovász and Szegedy [19] is that every element of $\mathcal{U}_\infty \cong \overline{\mathcal{U}^*}$ can be obtained as Γ_W satisfying (6.1) for some $W \in \mathcal{W}_{\mathsf{s}}$.

It is now clear that the representation theorems by Aldous [1] and Lovász and Szegedy [19] are connected by Theorem 5.3 and Corollary 5.4 above, and that one characterization easily follows from the other.

Remark 6.1. The representation by W is far from unique. Clearly, any measure preserving bijection σ of $[0,1]$ onto itself transforms W into a function W^σ with the same $G(\infty, W)$ and Γ_W. For a precise result, see Aldous [1]??

Borgs, Chayes, Lovász, Sós and Vesztergombi [7] call an element $W \in \mathcal{W}_{\mathsf{s}}$ a *graphon*. They further define a pseudometric (called the *cut-distance*) on $\mathcal{W}_{\mathsf{s}}$ and show that if we consider the quotient space $\widehat{\mathcal{W}_{\mathsf{s}}}$ obtained by identifying elements with cut-distance 0, we obtain a compact metric space, and the mapping $W \mapsto \Gamma_W$ yields a bijection $\widehat{\mathcal{W}_{\mathsf{s}}} \to \overline{\mathcal{U}^*} \cong \mathcal{U}_\infty$, which further more is a homeomorphism.

Remark 6.2. As remarked in Lovász and Szegedy [19], we can more generally consider a symmetric measurable function $W : \mathcal{S}^2 \to [0,1]$ for any probability space $(\mathcal{S}, \mu)$, and define $G(\infty, W)$ as above with X_i i.i.d. random variables in $\mathcal{S}$ with distribution μ. This does not give any new limit objects $G(\infty, W)$ or Γ_W, since we just said that every limit object is obtained from some $W \in \mathcal{W}_{\mathsf{s}}$, but they can sometimes give useful representations.

An interesting case is when W is the adjacency matrix of a (finite) graph G, with $\mathcal{S} = V(G)$ and μ the uniform measure on $\mathcal{S}$; we thus let X_i be i.i.d. random vertices of G and $G(n, W)$ equals the graph $G[n]$ defined in Section 2. It follows from (6.1) and (2.1) that $t(F, \Gamma_W) = t(F, G)$ for every $F \in \mathcal{U}$, and thus $\Gamma_W = G$ as elements of $\overline{\mathcal{U}^*}$. In other words, $\Gamma_W \in \mathcal{U}_\infty = \pi(G)$, the "ghost" of G in $\mathcal{U}_\infty \cong \overline{\mathcal{U}^*}$.

Remark 6.3. For the asymptotic behavior of $G(n, W)$ in another, sparse, case, with W depending on n, see [3].

7. Bipartite graphs

The definitions and results above have analogues for bipartite graphs, which we give in this section, leaving some details to the reader. The proofs are straightforward analogues of the ones given above and are omitted.

A *bipartite graph* will be a graph with an explicit bipartition; in other words, a bipartite graph G consists of two vertex sets $V_1(G)$ and $V_2(G)$ and an edge set $E(G) \subseteq V_1(G) \times V_2(G)$; we let $v_1(G) := |V_1(G)|$ and $v_2(G) := |V_2(G)|$ be the numbers of vertices in the two sets. Again we consider both the labelled and unlabelled cases; in the labelled case we assume the labels of the vertices in $V_j(G)$ are $1, \ldots, v_j(G)$ for $j = 1, 2$. Let $\mathcal{B}_{n_1 n_2}^L$ be the set of the $2^{n_1 n_2}$ labelled bipartite graphs with vertex sets $[n_1]$ and $[n_2]$, and let $\mathcal{B}_{n_1 n_2}$ be the quotient set $\mathcal{B}_{n_1 n_2}^L / \cong$ of unlabelled bipartite graphs with n_1 and n_2 vertices in the two parts; further, let $\mathcal{B}^L := \bigcup_{n_1, n_2 \geq 1} \mathcal{B}_{n_1 n_2}^L$ and $\mathcal{B} := \bigcup_{n_1, n_2 \geq 1} \mathcal{B}_{n_1 n_2}$.

We let $G[k_1, k_2]$ be the random graph in $\mathcal{B}_{k_1 k_2}^L$ obtained by sampling k_j vertices from $V_j(G)$ ($j = 1, 2$), uniformly with replacement, and let, provided $k_j \leq v_j(G)$, $G[k_1, k_2]'$ be the corresponding random graph obtained by sampling without replacement. We then define $t(F, G)$, $t_{\mathrm{inj}}(F, G)$ and

$t_{\mathrm{ind}}(F, G)$ for (unlabelled) bipartite graphs F and G in analogy with (2.1)–(2.3). Then (2.4)–(2.6) still hold, *mutatis mutandis*; for example,

$$\left| t(F, G) - t_{\mathrm{inj}}(F, G) \right| \leq \frac{v_1(F)^2}{2v_1(G)} + \frac{v_2(F)^2}{2v_2(G)}. \tag{7.1}$$ `a4xx`

In analogy with (2.7), we now define $\tau : \mathcal{B} \to [0, 1]^{\mathcal{B}}$ by

$$\tau(G) := (t(F, G))_{F \in \mathcal{B}} \in [0, 1]^{\mathcal{B}}. \tag{7.2}$$ `taub`

We define $\mathcal{B}^* := \tau(\mathcal{B}) \subseteq [0, 1]^{\mathcal{B}}$ to be the image of $\mathcal{B}$ under this mapping τ, and let $\overline{\mathcal{B}^*}$ be the closure of $\mathcal{B}^*$ in $[0, 1]^{\mathcal{B}}$; this is a compact metric space.

Again, τ is not injective; we may consider a graph G as an element of $\mathcal{B}^*$ by identifying G and $\tau(G)$, but this implies identification of some graphs of different orders and we prefer to avoid it. We let $\mathcal{B}^+$ be the union of $\mathcal{B}$ and some two-point set $\{*_1, *_2\}$ and consider the mapping $\tau^+ : \mathcal{B} \to [0, 1]^{\mathcal{B}^+} = [0, 1]^{\mathcal{B}} \times [0, 1] \times [0, 1]$ defined by

$$\tau^+(G) = \left(\tau(G),\, v_1(G)^{-1},\, v_2(G)^{-1} \right). \tag{7.3}$$ `tab+`

Then τ^+ is injective and we can identify $\mathcal{B}$ with its image $\tau^+(\mathcal{B}) \subseteq [0, 1]^{\mathcal{B}^+}$ and define $\overline{\mathcal{B}} \subseteq [0, 1]^{\mathcal{B}^+}$ as its closure; this is a compact metric space.

The functions $t(F, \cdot)$, $t_{\mathrm{inj}}(F, \cdot)$, $t_{\mathrm{ind}}(F, \cdot)$ and $v_j(\cdot)^{-1}$, for $F \in \mathcal{B}$ and $j = 1, 2$, have unique continuous extensions to $\overline{\mathcal{B}}$.

We let $\mathcal{B}_{\infty\infty} := \{G \in \overline{\mathcal{B}} : v_1(G) = v_2(G) = \infty\}$; this is the set of all limit objects of sequences (G_n) in $\mathcal{B}$ with $v_1(G_n), v_2(G_n) \to \infty$. By (7.1), $t_{\mathrm{inj}}(F, G) = t(F, G)$ for every $G \in \mathcal{B}_{\infty\infty}$ and every $F \in \mathcal{B}$. The projection $\pi : \overline{\mathcal{B}} \to \overline{\mathcal{B}^*}$ restricts to a homeomorphism $\mathcal{B}_{\infty\infty} \cong \overline{\mathcal{B}^*}$.

Remark 7.1. Note that in the bipartite case there are other limit objects too in $\overline{\mathcal{B}}$; in fact, $\overline{\mathcal{B}}$ can be partitioned into $\mathcal{B}$, $\mathcal{B}_{\infty\infty}$, and the sets $\mathcal{B}_{n\infty}$, $\mathcal{B}_{\infty n}$, for $n = 1, 2, \ldots$, where, for example, $\mathcal{B}_{n_1\infty}$ is the set of limits of sequences (G_n) of bipartite graphs such that $v_2(G_n) \to \infty$ but $v_1(G_n) = n_1$ is constant. We will not consider such degenerate limits further here, but we remark that in the simplest case $n_1 = 1$, a bipartite graph in $\mathcal{B}_{1n_2}^L$ can be identified with a subset of $[n_2]$, and an unlabelled graph in $\mathcal{B}_{1n_2}$ thus with a number in $m \in \{0, \ldots, n_2\}$, the number of edges in the graph, and it is easily seen that a sequence of such unlabelled graphs with $n_2 \to \infty$ converges in $\overline{\mathcal{B}}$ if and only if the proportion m/n_2 converges; hence we can identify $\mathcal{B}_{1\infty}$ with the interval [0,1].

We have the following basic result, cf. Theorem 2.1.

`T1B` **Theorem 7.1.** *Let (G_n) be a sequence of bipartite graphs with $v_1(G_n)$, $v_2(G_n) \to \infty$. Then the following are equivalent.*

`T1Bt`   (i) *$t(F, G_n)$ converges for every $F \in \mathcal{B}$.*

`T1Btinj`   (ii) *$t_{\mathrm{inj}}(F, G_n)$ converges for every $F \in \mathcal{B}$.*

`T1Btind`   (iii) *$t_{\mathrm{ind}}(F, G_n)$ converges for every $F \in \mathcal{B}$.*

    (iv) *G_n converges in $\overline{\mathcal{B}}$.*

In this case, the limit G of G_n belongs to $\mathcal{B}_{\infty\infty}$ and the limits in (i), (iii) and (iii) are $t(F, G)$, $t_{\mathrm{inj}}(F, G)$ and $t_{\mathrm{ind}}(F, G)$.

For convergence of random unlabelled bipartite graphs, the results in Section 3 hold with trivial changes.

Theorem 7.2. *Let G_n, $n \geq 1$, be random unlabelled bipartite graphs and assume that $v_1(G_n), v_2(G_n) \xrightarrow{\mathrm{p}} \infty$. The following are equivalent, as $n \to \infty$.*

    (i) *$G_n \xrightarrow{\mathrm{d}} \Gamma$ for some random $\Gamma \in \overline{\mathcal{B}}$.*

    (ii) *For every finite family $F_1, \ldots, F_m$ of (non-random) bipartite graphs, the random variables $t(F_1, G_n), \ldots, t(F_m, G_n)$ converge jointly in distribution.*

    (iii) *For every (non-random) $F \in \mathcal{B}$, the random variables $t(F, G_n)$ converge in distribution.*

    (iv) *For every (non-random) $F \in \mathcal{B}$, the expectations $\mathbb{E}\, t(F, G_n)$ converge.*

If these properties hold, then the limits in (ii), (iii) and (iv) are $\left(t(F_i, \Gamma)\right)_{i=1}^{m}$, $t(F, \Gamma)$ and $\mathbb{E}\, t(F, \Gamma)$, respectively. Furthermore, $\Gamma \in \mathcal{B}_{\infty\infty}$ a.s.

    The same results hold if t is replaced by t_{inj} or t_{ind}.

Corollary 7.3. *Let G_n, $n \geq 1$, be random unlabelled bipartite graphs such that $v_1(G_n), v_2(G_n) \xrightarrow{\mathrm{p}} \infty$, and let $G \in \mathcal{B}_{\infty\infty}$. The following are equivalent, as $n \to \infty$.*

    (i) *$G_n \xrightarrow{\mathrm{p}} G$.*

    (ii) *$t(F, G_n) \xrightarrow{\mathrm{p}} t(F, G)$ for every (non-random) $F \in \mathcal{B}$.*

    (iii) *$\mathbb{E}\, t(F, G_n) \to t(F, G)$ for every (non-random) $F \in \mathcal{B}$.*

    The same result holds if t is replaced by t_{inj} or t_{ind}.

As above, the distribution of Γ is uniquely determined by the numbers $\mathbb{E}\, t(F, \Gamma)$, $F \in \mathcal{B}$.

Let $\mathcal{B}_{\infty\infty}^{L}$ denote the set of all labelled infinite bipartite graphs with the vertex sets $V_1(G) = V_2(G) = \mathbb{N}$. $\mathcal{B}_{\infty\infty}^{L}$ is a compact metric space with the natural product topology.

If G is a bipartite graph, let $\widehat{G}$ be the random labelled bipartite graph obtained by random labellings of the vertices in $V_j(G)$ by the numbers $1, \ldots, v_j(G)$, for $j = 1, 2$. This is a random finite bipartite graph, but we can also regard it as a random element of $\mathcal{B}_{\infty\infty}^{L}$ by adding isolated vertices.

Definition. A random infinite bipartite graph $H \in \mathcal{B}_{\infty\infty}^{L}$ is *exchangeable* if its distribution is invariant under every pair of finite permutations of $V_1(H)$ and $V_2(H)$.

Theorem 7.4. *Let (G_n) be a sequence of random graphs in $\mathcal{B}$ and assume that $v_1(G_n), v_2(G_n) \xrightarrow{\mathrm{p}} \infty$. Then the following are equivalent.*

    (i) *$G_n \xrightarrow{\mathrm{d}} \Gamma$ in $\overline{\mathcal{B}}$ for some random $\Gamma \in \overline{\mathcal{B}}$.*

    (ii) *$\widehat{G_n} \xrightarrow{\mathrm{d}} H$ in $\mathcal{B}_{\infty\infty}^{L}$ for some random $H \in \mathcal{B}_{\infty\infty}^{L}$.*

If these hold, then $\mathbb{P}(H|_{[k_1] \times [k_2]} = F) = \mathbb{E}\, t_{\mathrm{ind}}(F, \Gamma)$ *for every* $F \in \mathcal{B}_{k_1 k_2}^L$. *Furthermore,* $\Gamma \in \mathcal{B}_{\infty\infty}$ *a.s., and* H *is exchangeable.*

Theorem 7.5. *There is a one-to-one correspondence between distributions of random elements* $\Gamma \in \mathcal{B}_{\infty\infty}$ *(or* $\overline{\mathcal{B}^*}$) *and distributions of exchangeable random infinite graphs* $H \in \mathcal{B}_{\infty\infty}^L$ *given by*

$$\mathbb{E}\, t_{\mathrm{ind}}(F, \Gamma) = \mathbb{P}(H|_{[k_1] \times [k_2]} = F) \tag{7.4}$$

for every $k_1, k_2 \geq 1$ *and every* $F \in \mathcal{B}_{k_1 k_2}^L$, *or, equivalently,*

$$\mathbb{E}\, t(F, \Gamma) = \mathbb{P}(H \supset F) \tag{7.5}$$

for every $F \in \mathcal{B}^L$. *Furthermore,* $H|_{[n_1] \times [n_2]} \xrightarrow{\mathrm{d}} \Gamma$ *in* $\overline{\mathcal{B}}$ *as* $n_1, n_2 \to \infty$.

Corollary 7.6. *There is a one-to-one correspondence between elements* Γ *of* $\mathcal{B}_{\infty\infty} \cong \overline{\mathcal{B}^*}$ *and extreme points of the set of distributions of exchangeable random infinite graphs* $H \in \mathcal{B}_{\infty\infty}^L$. *This correspondence is given by*

$$t(F, \Gamma) = \mathbb{P}(H \supset F) \tag{7.6}$$

for every $F \in \mathcal{B}^L$. *Furthermore,* $H|_{[n_1] \times [n_2]} \xrightarrow{\mathrm{p}} \Gamma$ *in* $\overline{\mathcal{B}}$ *as* $n_1, n_2 \to \infty$.

Remark 7.2. We have not checked whether $H|_{[n_1] \times [n_2]} \xrightarrow{\mathrm{a.s.}} \Gamma$ in $\overline{\mathcal{B}}$ as $n_1, n_2 \to \infty$. This holds at least for a subsequence $(n_1(m), n_2(m))$ with both $n_1(m)$ and $n_2(m)$ non-decreasing because then $t_{\mathrm{inj}}(F, H|_{[n_1] \times [n_2]})$ is a reverse martingale.

<div style="text-align:right">Do we care about this?</div>

Theorem 7.7. *Let* H *be an exchangeable random infinite bipartite graph. Then the following are equivalent.*

(i) *The distribution of* H *is an extreme point in the set of exchangeable distributions in* $\mathcal{B}_{\infty\infty}^L$.

(ii) *If* F_1 *and* F_2 *are two (finite) bipartite graphs with the vertex sets* $V_j(F_1)$ *and* $V_j(F_2)$ *disjoint subsets of* $\mathbb{N}$ *for* $j = 1, 2$, *then*

$$\mathbb{P}(H \supset F_1 \cup F_2) = \mathbb{P}(H \supset F_1)\,\mathbb{P}(H \supset F_2).$$

The construction in Section 6 takes the following form; note that there is no need to assume symmetry of W. For every $W \in \mathcal{W}$, we define an infinite random bipartite graph $G(\infty, \infty, W) \in \mathcal{B}_{\infty\infty}^L$ as follows: we first choose two sequence $X_1, X_2, \ldots$ and $Y_1, Y_2, \ldots$ of i.i.d. random variables uniformly distributed on $[0, 1]$, and then, given these sequences, for each pair $(i, j) \in \mathbb{N} \times \mathbb{N}$ we draw an edge ij with probability $W(X_i, Y_j)$, independently for all pairs (i, j). Further, let $G(n_1, n_2, W)$ be the restriction $G(\infty, \infty, W)|_{[n_1] \times [n_2]}$, which is obtained by the same construction with finite sequences $X_1, \ldots, X_{n_1}$ and $Y_1, \ldots, Y_{n_2}$.

It is evident that $G(\infty, \infty, W)$ is an exchangeable infinite random bipartite graph. Furthermore, it satisfies Theorem 7.7(ii). Theorem 7.5 and

<div style="text-align:right">I'm too lazy to state the analogues of Theorem 5.5(iii) and (iv) correctly</div>

Corollary 7.6 yield a corresponding element $\Gamma''_W \in \mathcal{B}_{\infty\infty} \cong \overline{\mathcal{B}^*}$ such that $G(n_1, n_2, W) \xrightarrow{\mathrm{p}} \Gamma''_W$ as $n_1, n_2 \to \infty$ and, for every $F \in \mathcal{B}^L_{k_1 k_2}$,

$$t(F, \Gamma''_W) = \int_{[0,1]^{k_1+k_2}} \prod_{ij \in \mathbb{E}(F)} W(x_i, y_j) \, \mathrm{d}x_1 \ldots \mathrm{d}x_{k_1} \, \mathrm{d}y_1 \ldots \mathrm{d}y_{k_2}. \qquad (7.7) \quad \boxed{\texttt{tuww}}$$

The result by Aldous [1] in the non-symmetric case is that every exchangeable infinite random bipartite graph is obtained as a mixture of such $G(\infty, \infty, W)$; in other words as $G(\infty, \infty, W)$ with a random W.

By Theorem 7.5 and Corollary 7.6 above, this implies (and is implied by) the fact that every element of $\overline{\mathcal{B}}$ equals Γ''_W for some (non-unique) $W \in \mathcal{W}$; the bipartite version of the characterization by Lovász and Szegedy [19].

References

$\boxed{\texttt{Aldous}}$ [1] D. Aldous, Representations for partially exchangeable arrays of random variables. *J. Multivar. Anal.* **11**, 581–598, 1981.

$\boxed{\texttt{Bill}}$ [2] P. Billingsley, *Convergence of Probability Measures.* Wiley, New York, 1968.

$\boxed{\texttt{SJ178}}$ [3] B. Bollobás, S. Janson and O. Riordan, The phase transition in inhomogeneous random graphs. *Random Struct. Alg.*, to appear, 2007.

$\boxed{\texttt{BR}}$ [4] B. Bollobás and O. Riordan, 2007.

$\boxed{\texttt{BCL3}}$ [5] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, B. Szegedy and K. Vesztergombi, Graph limits and parameter testing. *STOC*, 2006.

$\boxed{\texttt{BCL2}}$ [6] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós and K. Vesztergombi, Counting graph homomorphisms. In *Topics in Discrete Mathematics* (ed. M. Klazar, J. Kratochvil, M. Loebl, J. Matousek, R. Thomas, P. Valtr), 315–371. Springer, New York, 2006.

$\boxed{\texttt{BCL1}}$ [7] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós and K. Vesztergombi, Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing. II: Multiway cuts and statistical physics. Preprint, January 2007. `http://arxiv.org/math.CO/0702004`

$\boxed{\texttt{DF1981}}$ [8] P. Diaconis and D. Freedman, On the statistics of vision: The Julesz conjecture. *J. Math. Psychol.* **24**, 112–138, 1981.

$\boxed{\texttt{DF1984}}$ [9] P. Diaconis and D. Freedman, Partial exchangeability and sufficiency. In *Statistics: Applications and New Directions* (eds. J. K. Jhosh and J. Roy), 205–236. Indian Statistical Institute, Calcutta, 1984.

$\boxed{\texttt{Freed}}$ [10] M. Freedman, L. Lovász and A. Schrijver, Reflection positivity, rank connectivity, and homomorphism of graphs. *J. Amer. Math. Soc.* **20**, 37–51, 2007.

$\boxed{\texttt{Hoover}}$ [11] D. Hoover, *Relations on Probability Spaces and Arrays of Random Variables.* Preprint, Institute for Advanced Study, Princeton, NJ, 1979.

$\boxed{\texttt{Kallenberg:exch}}$ [12] O. Kallenberg, *Probabilistic Symmetries and Invariance Principles.* Springer, New York, 2005.

$\boxed{\texttt{LL2006}}$ [13] L. Lovász, The rank of connection matrices and the dimension of graph algebras. *Eur. J. Comb.* **27**, 962–970, 2006.

LL2007 [14] L. Lovász, Connection matrices. In *Combinatorics, Complexity, and Chance: A Tribute to Dominic Welsh* (ed. G. Grimmet and C. McDiarmid), 179–190. Oxford University Press, Oxford, 2007.

LSos [15] L. Lovász and V. T. Sós, Generalized quasirandom graphs. *J. Comb. Theory B.*

LSzcont [16] L. Lovász and B. Szegedy, Contractors and connectors in graph algebras. *J. Comb. Theory B.*

LSzszem [17] L. Lovász and B. Szegedy, Szemerédi's Lemma for the analyst. *J. Geom. Func. Anal.* **17**, 252–270, 2007.

LSztest [18] L. Lovász and B. Szegedy, Testing properties of graphs and functions.

LSz [19] L. Lovász and B. Szegedy, Limits of dense graph sequences. *J. Comb. Theory B* **96**, 933–957, 2006. `http://arxiv.org/math.CO/0408173`

STANFORD AND NICE ???

DEPARTMENT OF MATHEMATICS, UPPSALA UNIVERSITY, PO BOX 480, SE-751 06 UPPSALA, SWEDEN
E-mail address: `svante.janson@math.uu.se`
URL: `http://www.math.uu.se/∼svante/`