

DENSE GRAPH LIMITS UNDER RESPONDENT-DRIVEN SAMPLING

BY SIVA ATHREYA^{1,2} AND ADRIAN RÖLLIN²

Indian Statistical Institute and National University of Singapore

We consider certain respondent-driven sampling procedures on dense graphs. We show that if the sequence of the vertex-sets is ergodic then the limiting graph can be expressed in terms of the original dense graph via a transformation related to the invariant measure of the ergodic sequence. For specific sampling procedures, we describe the transformation explicitly.

1. Introduction. Respondent-driven sampling (RDS) of social networks has received a lot of attention since [9] and [10], and many studies have implemented the procedure in order to obtain estimates about properties of so-called “hidden” or “hard-to-reach” populations. The basic idea is to start with a convenience sample of participants, to ask the participants for *referrals* among their peers and then to iterate this process. It is intuitively clear that one cannot hope to obtain an unbiased sample in this manner as individuals with higher connectivity are more likely to appear in the sample than individuals with lower connectivity. In order to avoid this bias, one of the key assumptions of [9] is that each individual in the network has the same degree. Subsequent refinements of the procedure have been proposed to overcome such restrictions; see [22].

Respondent-driven sampling has also received quite some criticism. Besides inadequate control of biases for finite samples, another major issue can be the underestimation of sample variance; see, for example, [7] and [8].

The main purpose of this article is to take a first (and very preliminary) step in establishing a rigorous theory of RDS on dense graphs in order to understand the graphs produced under various sampling procedures. Our main contribution is that the limit of a dense graph sequence obtained through a specific respondent-driven sampling procedure, where the sequence of the vertex-sets is ergodic, can be expressed in terms of the original graph limit and a transformation related to the invariant measure of the ergodic sequence. The transformation, in essence, confirms the bias toward nodes with larger degrees.

In practice, researchers typically are interested in estimating certain quantities at population or subpopulation level, such as prevalence of STIs, sexual contact

Received April 2014; revised September 2015.

¹Supported in part by CPDA-grant from the Indian Statistical Institute, Bangalore.

²Supported in part by NUS Research Grant R-155-000-124-112.

MSC2010 subject classifications. Primary 05C80, 60J20; secondary 37A30, 9482.

Key words and phrases. Random graph, dense graph limits, ergodic, respondent driven sampling.

frequencies, condom use, etc. Hence, for each node in the network, additional data is collected, and the main question of RDS becomes how to obtain representative estimates of those quantities from the RDS sample. In this article, we will only be interested in the network itself and the question how specific RDS procedures bias the network. However, if, for example, a quantity of interest (such as STI prevalence) correlates with the degree that a node has in that network, then it is obviously important to understand the bias in the network itself in order to understand the resulting bias of that quantity of interest.

It is also important to note that the sampling procedure analysed in this article is not representative for what is mostly being done in practice. In particular, we assume that after the referral chain has been sampled (or rather “revealed”), all yet unknown connections between the subjects in the sample are also revealed. In other words, if Subject A refers to Subject B and Subject B refers to Subject C, we assume that, in a second step, the relationship between Subjects A and C be revealed, also. In practice that last relationship typically remains unknown, unless either A refers to C or C refers to A.

Our proof is based on subgraph counts convergence and ergodicity of the sampling procedure. Subgraph counts can be written as incomplete U -statistics or generalised U -statistics, but there does not seem to exist a well-established general theory that would cover ergodic sequences in the generality needed in this article. However, noticing that, in our model, the conditional expectation of a subgraph count, conditioned on the vertex set, is a *complete* U -statistic, we can resort to the well-established theory of U -statistics, in particular for ergodic sequences. We modify the arguments of [1] in order to deal with nonstationary sequences, which seems a more realistic assumption in the context of RDS.

The rest of the paper is organised in the following manner. We conclude this section with a focussed review of the dense graph literature. We state the model and main results in Section 2 and prove them in Section 3. We finally discuss some applications in Section 4.

1.1. *A brief introduction to dense graphs.* Dense graph theory has been introduced by [17]. Diaconis and Janson [6] made connections with earlier work of [2]. Let us briefly summarise those parts of dense graph theory which are needed in this paper; see the monograph [16] and [4, 5] for an in-depth discussion, or [3] for another introduction with extensions to sparse graphs.

Let $(G_n)_{n \geq 1}$ be a sequence of graphs, where for simplicity we assume that the vertex set of graph G_n is $\{1, \dots, n\}$. Assume further that the number of edges $E(G_n)$ in G_n is of order n^2 ; that is, $\liminf_{n \rightarrow \infty} (E(G_n)n^{-2}) > 0$. We call $(G_n)_{n \geq 1}$ a *dense graph sequence*.

Subgraph distance. Let K_n be the complete graph on $\{1, \dots, n\}$. For any (small) graph F on k vertices, let $X_F(G)$ be the number of copies of F in a (large) graph G on n vertices. Define the normalised subgraph count

$$t(F, G) = \frac{|\text{inj}(F, G)|}{(n)_k} = \frac{X_F(G)}{X_F(K_n)},$$

where $\text{inj}(F, G)$ denotes the set of injective graph homomorphisms of F into G , that is, the functions that map the vertices of F into the vertices of G injectively such that connected vertices in F remain connected in G . Here, as usual, $(n)_k = n(n - 1) \cdots (n - k + 1)$. When $|F| > |G|$, we define $t(F, G) = 0$.

Note that $0 \leq t(F, G) \leq 1$. Let \mathcal{F} denote the class of isomorphism classes on finite graphs and let $(F_i)_{i \geq 1}$ be a particular enumeration of \mathcal{F} , where each F_i is the representative of an isomorphism class. We can define a distance function between graphs by

$$d_{\text{sub}}(G, G') = \sum_{i \geq 1} 2^{-i} |t(F_i, G) - t(F_i, G')|.$$

A key feature of d_{sub} is that there is a natural completion of $(\mathcal{F}, d_{\text{sub}})$ by standard kernels. We call any function $\kappa : [0, 1]^2 \rightarrow [0, 1]$ that is measurable and symmetric a *standard kernel*. For F a graph on k vertices, we can extend the definition of $t(F, G)$ to kernels by means of

$$t(F, \kappa) = \int_{[0,1]^k} \prod_{\{i,j\} \in E(F)} \kappa(x_i, x_j) dx_1 \cdots dx_k,$$

where $E(F)$ is the set of edges in F . One of the key results in dense graph theory is the following theorem.

THEOREM 1.1. *Let $(G_n)_{n \geq 1}$ be a dense graph sequence which is Cauchy with respect to d_{sub} . Then there exists a standard kernel κ such that*

$$(1.1) \quad d_{\text{sub}}(G_n, \kappa) \rightarrow 0$$

as $n \rightarrow \infty$.

For a proof of the above (see [17]), which uses Szemerédi partitions and the Martingale convergence theorem, or [6], who show that it can be proved using results from [12] and [2]. Note that κ above is in general not unique, but this will not be of importance for what follows; we refer to [4, 5] for a discussion of this and related questions. We refer to [16], Chapter 11, for a detailed discussion of convergence of dense graph sequences.

2. Model and main results. A convenient way of “creating” finite (random) graphs on n vertices from a standard kernel κ is the following model, which we will denote by $G(n, \kappa)$. Firstly, let U_1, \dots, U_n be i.i.d. with uniform distribution on $[0, 1]$. Second, for each two vertices i and j , connect them with probability $\kappa(U_i, U_j)$, independently of all the other edges. It is not difficult to prove that

$$(2.1) \quad d_{\text{sub}}(G(n, \kappa), \kappa) \rightarrow 0$$

almost surely as $n \rightarrow \infty$. This is, in some sense, the basic law of large numbers in dense graph theory. In this article, instead of sampling the labels i.i.d. and uniformly from $[0, 1]$, we will allow the labels to be sampled in a more general way.

The random graph $G(x, \kappa)$. Let $x = (x_1, \dots, x_n) \in [0, 1]^n$ be fixed. Define the random graph $G(x, \kappa)$ by connecting vertices i and j with probability $\kappa(x_i, x_j)$ independently of all other vertices. Clearly, $G((U_1, \dots, U_n), \kappa)$ is equivalent to $G(n, \kappa)$. We will show a version of (2.1) for $G(X, \kappa)$, where—in essence—the labels X are allowed to come from a general ergodic sequence.

To this end, let κ be a standard kernel and let $g : [0, 1] \rightarrow [0, 1]$ be a Lebesgue-measurable function. Define the g -transformed kernel

$$\kappa_g(x, y) = \kappa(g(x), g(y)).$$

THEOREM 2.1. *Let $X^{(n)} = (X_{n,1}, \dots, X_{n,i}, \dots, X_{n,n}), n \geq 1$, be a triangular array of random variables taking values in $[0, 1]$. Assume that there is a probability measure π on $[0, 1]$ such that the following two conditions hold:*

(i) *for all bounded and measurable functions f , we have*

$$(2.2) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n f(X_{n,i}) = \int_0^1 f(x) d\pi(x)$$

almost surely;

(ii) $\kappa(\cdot, \cdot)$ *is continuous $(\pi \times \pi)$ -almost everywhere.*

Then

$$(2.3) \quad d_{\text{sub}}(G(X^{(n)}, \kappa), \kappa_{\tau^{-1}}) \rightarrow 0$$

almost surely, where $\tau(x) = \pi([0, x])$ is the distribution function of π and

$$\tau^{-1}(v) = \inf\{u \in [0, 1] : \tau(u) \geq v\}$$

its generalised inverse.

2.1. Respondent driven sampling. The way we think about respondent-driven sampling in this article is by means of the following two-step procedure. First, sample a set of subjects $X^{(n)} = (X_{n,1}, \dots, X_{n,n})$, where new subjects are added by referrals; each subject i obtains a unique label $X_{n,i} \in [0, 1]$ (note that the $X_{n,i}$ are just the labels of the nodes, not any additional observation related to that node). If i referred to j , or j referred to i , an edge between the two nodes is added; denote the resulting graph by H_n . Second, the remaining relationships are then revealed by connecting i and j with probability $\kappa(X_{n,i}, X_{n,j})$, unless they are already connected in H_n . We define the model precisely below.

The random graph $G(x, H_n, \kappa)$. Let $x = (x_1, \dots, x_n) \in [0, 1]^n$ be fixed and let H_n be a given graph on the vertices $\{1, \dots, n\}$. Define the random graph $G(x, H_n, \kappa)$ on the same set of vertices as follows:

- if there is an edge between i and j in H_n , then connect vertices i and j in $G(x, H_n, \kappa)$;

- if there is no edge between i and j in H_n , then connect i and j in $G(x, H_n, \kappa)$ with probability $\kappa(x_i, x_j)$ independently of all other vertices.

COROLLARY 2.2. *Let $X^{(n)} = (X_{n,i})_{1 \leq i \leq n}$, κ and τ be as in Theorem 2.1, satisfying conditions (i) and (ii). If the number of edges in H_n is $o(n^2)$, then*

$$(2.4) \quad d_{\text{sub}}(G(X^{(n)}, \kappa, H_n), \kappa_{\tau-1}) \rightarrow 0$$

almost surely.

The above corollary is an easy consequence of Theorem 2.1 and the counting lemma [16], Lemma 10.22. For completeness sake, we present a proof in the next section.

2.2. *Remarks.* Before concluding this section, we discuss some interesting aspects around Theorem 2.1.

Reference measure space. Using $[0, 1]$ and the Lebesgue measure as reference is only a matter of convenience and in line with the prevailing literature. However, in order to shed some light on the main result, let us state Theorem 2.1 in greater generality; we refer to [16], Chapter 13, for a more in-depth discussion of this setting.

Let $(\mathcal{X}, \mathcal{A}, \mu)$ be a probability space, and let $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$ be a symmetric and $(\mathcal{A} \times \mathcal{A})$ -measurable function. For any graph F on k vertices, where $k \geq 1$, we can easily generalise the definition of the subgraph density to

$$t_\mu(F, \kappa) = \int_{\mathcal{X}^k} \prod_{\{i,j\} \in E(F)} \kappa(x_i, x_j) d\mu(x_1) \cdots d\mu(x_k).$$

Moreover, for U_1, \dots, U_n being i.i.d. random elements taking values in \mathcal{X} with common distribution μ , the random graph model $G((U_1, \dots, U_n), \kappa)$ can be defined in a straightforward manner, and one can prove that

$$d_{\mu, \text{sub}}(G((U_1, \dots, U_n), \kappa), \kappa) \rightarrow 0, \quad n \rightarrow \infty.$$

It is important to emphasise that t_μ and, as a result, the metric $d_{\mu, \text{sub}}$ depend on the reference measure μ .

Now, assume $(X_{n,i})_{1 \leq i \leq n}$ is a triangular array of \mathcal{X} -valued random elements such that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n f(X_{n,i}) = \int_{\mathcal{X}} f(x) d\pi(x)$$

for some probability measure π on $(\mathcal{X}, \mathcal{A})$. Now, assume \mathcal{X} is a Polish space. If there is a function $g : \mathcal{X} \rightarrow \mathcal{X}$ such that

$$g(U_1) \sim \pi,$$

and if κ is continuous ($\pi \times \pi$)-almost everywhere, then

$$(2.5) \quad d_{\mu, \text{sub}}(G(X^{(n)}, \kappa), \kappa_g) \rightarrow 0, \quad n \rightarrow \infty.$$

In the case where $(\Omega, \mathcal{A}, \mu)$ is the interval $[0, 1]$ and μ the uniform distribution, g can be identified as the generalised inverse of the distribution function of π , but note that, for general spaces \mathcal{X} , it is difficult to find such g explicitly.

It is illuminating to consider the following alternative way to state (2.5). From the proof of Theorem 2.1 [cf. (3.6)], it becomes clear that, by changing the reference measure from μ to π , (2.5) can also be written as

$$(2.6) \quad d_{\pi, \text{sub}}(G(X^{(n)}, \kappa), \kappa) \rightarrow 0, \quad n \rightarrow \infty.$$

Although (2.5) and (2.6) are equivalent, the former statement is more important in the context of RDS, since we are interested in describing the distortion of the network through biased sampling.

Necessity of condition (ii). Condition (ii) in Theorem 2.1 can be replaced by other conditions, but that it cannot be dispensed with entirely can be seen from [1], Example 4.1. We state the example below with notation as applicable to our case.

Consider the interval $(0, 1)$ and define the mapping $\phi : (0, 1) \rightarrow (0, 1)$ as

$$\phi(x) = 2x \pmod{1}.$$

Let $X_1 \in (0, 1)$ fixed, and let $X_n = \phi^n(X_1)$. It follows from standard ergodic theory that X_1, X_2, \dots is ergodic with the Lebesgue measure as its invariant measure, that is,

$$\frac{1}{n} \sum_{i=1}^n f(X_i) \rightarrow \int_0^1 f(x) dx.$$

Define the set

$$L = \{(x_1, x_2) \in (0, 1)^2 : x_1 \in (0, 1) \text{ and } x_2 = \phi^n(x_1) \text{ for some } n \geq 1\}.$$

So, (2.2) is satisfied with π being the Lebesgue measure; hence $\tau(x) = x$ and $\tau^{-1}(x) = x$. Define the standard kernel $\kappa(x, y) = \mathbb{I}[(x, y) \in L \text{ or } (y, x) \in L]$. Now, on the one hand we have

$$\frac{\sum_{i=1}^n \kappa(X_i, X_j)}{n(n-1)} = 1 \quad \text{for all } n \geq 1.$$

On the other hand,

$$\int \kappa(x, y) dx dy = 0$$

since L is the countable union of null sets with respect to the two-dimensional Lebesgue measure. Thus,

$$\lim_{n \rightarrow \infty} t(F, G((X_1, \dots, X_n), \kappa)) = 1 \neq 0 = t(F, \kappa).$$

Since L is dense in $[0, 1]$, the standard kernel κ is nowhere continuous and does therefore not satisfy condition (ii).

3. Proof of Theorem 2.1. We will need a law of large numbers of a particular U -statistic for the proof of Theorem 2.1. This essentially allows us to go from a simple ergodic theorem to a higher order ergodic theorem. Toward that we define

$$(3.1) \quad \mu_F(x) = \frac{1}{\binom{n}{k}} \sum_{(i_1, \dots, i_k)} T_F(x_{i_1}, \dots, x_{i_k})$$

with

$$(3.2) \quad T_F(z_1, \dots, z_k) = \prod_{\{i,j\} \in F} \kappa(z_i, z_j),$$

for $x = (x_1, \dots, x_n) \in [0, 1]^n$ and $z_i \in [0, 1], 1 \leq i \leq k$. Here, the summation $\sum_{(i_1, \dots, i_k)}$ ranges over all vectors (i_1, \dots, i_k) with mutually different coordinates.

The following result was proved by [1] for ergodic stationary sequences, but we note that the key assumption is (2.2), so that their proof generalises to nonstationary triangular arrays, which is more appropriate for the applications we have in mind.

LEMMA 3.1. *Let $X^{(n)} = (X_{n,i})_{1 \leq i \leq n}$ be a triangular array of random variables taking values in $[0, 1]$ and satisfying (2.2) almost surely. Then, for any fixed graph F of size k ,*

$$\lim_{n \rightarrow \infty} \mu_F(X^{(n)}) \rightarrow \mathbb{E}T_F(V_1, \dots, V_k)$$

almost surely, where V_1, \dots, V_k are i.i.d. random variables with distribution π .

PROOF. Our proof is a close imitation of the proof of [1], Theorem U. Denote by \mathcal{C}_k the set of all functions from $[0, 1]^k$ to $[0, 1]$, continuous $\pi^{(k)}$ -almost everywhere, where $\pi^{(k)} = \otimes_{i=1}^k \pi$. For $x = (x_1, \dots, x_n) \in [0, 1]^n$ and $h \in \mathcal{C}_k$, define

$$v_h(x) = \frac{1}{n^k} \sum_{1 \leq i_1, \dots, i_k \leq n} h(x_{i_1}, \dots, x_{i_k}).$$

Let

$$\mathcal{P}_k = \left\{ h \in \mathcal{C}_k : \exists h_1, \dots, h_k \in \mathcal{C}_1 \text{ such that } h(x) = \prod_{i=1}^k h_i(x_i) \right\}.$$

For any $h \in \mathcal{P}_k$, we have

$$v_h(X^{(n)}) = \prod_{i=1}^k \left(\frac{1}{n} \sum_{j=1}^n h_i(X_{n,j}) \right).$$

So, by (2.2),

$$v_h(X^{(n)}) \rightarrow \prod_{i=1}^k \int_{[0,1]} h_i(x_i) d\pi(x_i) = \mathbb{E}h(V^{(k)})$$

almost surely as $n \rightarrow \infty$, where we set $V^{(k)} = (V_1, \dots, V_k)$ to shorten formulas. It is easily seen that the above holds whenever $h \in \text{span}(\mathcal{P}_k)$.

Fix $\varepsilon > 0$ and let $h \in \mathcal{C}_k$. As h is continuous $\pi^{(k)}$ -almost everywhere and $\pi^{(k)}$ -integrable, there exist $s_1, s_2 \in \text{span}(\mathcal{P}_k)$ such that

$$(a) \quad |h - s_1| \leq s_2 \quad \pi^{(k)}\text{-almost everywhere,} \quad (b) \quad \int s_2 d\pi^{(k)} \leq \varepsilon.$$

As $s_1, s_2 \in \text{span}(\mathcal{P}_k)$, there exists (random) N such that, for $n \geq N$,

$$|v_{s_1}(X^{(n)}) - \mathbb{E}s_1(V^{(k)})| \leq \varepsilon$$

and

$$v_{s_2}(X^{(n)}) \leq \mathbb{E}s_2(V^{(k)}) + \varepsilon \leq 2\varepsilon.$$

Hence, for $n \geq N$,

$$\begin{aligned} & |v_h(X^{(n)}) - \mathbb{E}h(V^{(k)})| \\ & \leq |v_h(X^{(n)}) - v_{s_1}(X^{(n)})| + |v_{s_1}(X^{(n)}) - \mathbb{E}s_1(V^{(k)})| \\ & \quad + |\mathbb{E}s_1(V^{(k)}) - \mathbb{E}h(V^{(k)})| \\ & \leq v_{|h-s_1|}(X^{(n)}) + |v_{s_1}(X^{(n)}) - \mathbb{E}s_1(V^{(k)})| + \mathbb{E}s_2(V^{(k)}) \\ & \leq v_{s_2}(X^{(n)}) + |v_{s_1}(X^{(n)}) - \mathbb{E}s_1(V^{(k)})| + \mathbb{E}s_2(V^{(k)}) \\ & \leq 4\varepsilon. \end{aligned}$$

Thus, for all $h \in \mathcal{C}_k$,

$$(3.3) \quad \lim_{n \rightarrow \infty} v_h(X^{(n)}) \rightarrow \mathbb{E}h(V^{(k)})$$

almost surely. Let T_F be as in (3.2). As T_F is bounded by 1, we see that there exists $c_1 > 0$ such that

$$\begin{aligned} & |\mu_F(X^{(n)}) - v_{T_F}(X^{(n)})| \\ & = \left| \frac{1}{\binom{n}{k}} \sum_{(i_1, \dots, i_k)} T_F(X_{i_1}, \dots, X_{i_k}) - \frac{1}{n^k} \sum_{1 \leq i_1, \dots, i_k \leq n} T_F(X_{i_1}, \dots, X_{i_k}) \right| \\ & \leq \frac{c_1}{n}. \end{aligned}$$

As $T_F \in \mathcal{C}_k$, the result follows. \square

PROOF OF THEOREM 2.1. It is enough to show that, for every graph F ,

$$(3.4) \quad \lim_{n \rightarrow \infty} |t(F, G(X^{(n)}, \kappa)) - t(F, \kappa_{\tau-1})| = 0$$

almost surely. Using the triangle inequality,

$$(3.5) \quad \begin{aligned} &|t(F, G(X^{(n)}, \kappa)) - t(F, \kappa_{\tau^{-1}})| \\ &\leq |t(F, G(X^{(n)}, \kappa)) - \mu_F(X^{(n)})| + |\mu_F(X^{(n)}) - t(F, \kappa_{\tau^{-1}})|. \end{aligned}$$

By definition of τ^{-1} , V_i has the same distribution as $\tau^{-1}(U_i)$, so that

$$(3.6) \quad \mathbb{E}T_F(V_1, \dots, V_k) = \mathbb{E}T_F(\tau^{-1}(U_1), \dots, \tau^{-1}(U_k)) = t(F, \kappa_{\tau^{-1}})$$

is immediate. Hence, Lemma 3.1 implies that the second term in (3.5) approaches 0 as $n \rightarrow \infty$.

We will use the main result of [18] to show that the first term in (3.5) also vanishes. If f is a function in N arguments such that changing the i th coordinate will change the value of f by at most c_i and if $Y = (Y_1, \dots, Y_N)$ are independent random variables, then

$$(3.7) \quad \mathbb{P}[|f(Y) - \mathbb{E}f(Y)| \geq \varepsilon] \leq 2 \exp\left(-\frac{2\varepsilon^2}{\sum_{i=1}^N c_i^2}\right).$$

Note now that, if G is a graph with n vertices, then $t(F, G)$ changes by at most $\frac{k(k-1)}{n(n-1)}$ if one edge is changed. Applying McDiarmid’s concentration inequality to $t(F, G(x, \kappa))$ [with f being a function of the $N = \binom{n}{2}$ random edges], we therefore have that, for every fixed $x \in [0, 1]^n$,

$$(3.8) \quad \mathbb{P}[|t(F, G(x, \kappa)) - \mu_F(x)| > \varepsilon] \leq 2 \exp\left(-\frac{2\varepsilon^2}{\binom{n}{2}(k(k-1)/(n(n-1)))^2}\right).$$

Using Borel–Cantelli, we can conclude that

$$(3.9) \quad |t(F, G(X^{(n)}, \kappa)) - \mu_F(X^{(n)})| \rightarrow 0$$

almost surely as $n \rightarrow \infty$. This proves the claim. \square

PROOF OF COROLLARY 2.2. We can essentially imitate the proof of Theorem 2.1 to obtain this result; the one difference being that we have to control

$$|t(F, G(X^{(n)}, H_n, \kappa)) - \mu_F(X^{(n)})|.$$

We need to be bit careful at (3.8) because of the dependencies introduced by H_n . Suppose $E(H_n) = m_n \equiv m$. Applying McDiarmid’s concentration inequality to $t(F, G(x, H_n, \kappa))$ [with f being a function of the $N = \binom{n}{2} - m$ random edges], we therefore have, with $G = G(x, H_n, \kappa)$,

$$\begin{aligned} &\mathbb{P}[|t(F, G) - \mu_F(x)| > \varepsilon] \\ &\leq \mathbb{P}[|t(F, G) - \mathbb{E}t(F, G)| > \varepsilon - |\mu_F(x) - \mathbb{E}t(F, G)|] \\ &\leq 2 \exp\left(-\frac{2(\varepsilon - m(k(k-1)/(n(n-1))))^2}{(\binom{n}{2} - m)(k(k-1)/(n(n-1)))^2}\right). \end{aligned}$$

As $m = o(n^2)$, it follows from Borel–Cantelli that

$$|t(F, G(X^{(n)}, H_n, \kappa)) - \mu_F(X^{(n)})| \rightarrow 0. \quad \square$$

4. Applications. In this section, we will discuss two different sampling schemes, namely a Markov chain model, where each respondent gives exactly one referral, and a Poisson branching process model, where each respondent gives a Poisson number of referrals (and thus, allowing for no referrals). For both procedures, we essentially need to establish (2.2). Once this is done, Theorem 2.1 automatically yields the corresponding convergence provided κ is continuous. We compare the two procedures for a concrete parametrised standard kernel under different parameter values.

In order to avoid that the standard kernel decomposes into two or more disconnected parts, it is natural to impose an irreducibility condition. We follow [15]. Denote by $\text{Vol}(A)$ the Lebesgue measure of $A \subset [0, 1]$ and let $A^c = [0, 1] \setminus A$. We say that a standard kernel is *connected*, if $0 < \text{Vol}(A) < 1$ implies

$$(4.1) \quad \int_A \int_{A^c} \kappa(x, y) dx dy > 0.$$

Loosely speaking, this condition guarantees that there can be links from any set A into its complement, so that no area can remain disconnected from the rest of the graph [at least as $n \rightarrow \infty$; for a finite realisation of $G(n, \kappa)$, it may of course happen that the graph consists of disconnected components]. Note that (4.1) implies in particular that

$$\int_A \int_0^1 \kappa(x, y) dy > 0$$

for all A with $\text{Vol}(A) > 0$. This only guarantees that almost all x have positive degree. In order to avoid technicalities, we shall assume that *all* x have positive degree, that is,

$$(4.2) \quad \int_0^1 \kappa(x, y) dy > 0 \quad \text{for all } x \in [0, 1].$$

If (4.2) is satisfied, we say that a standard kernel is *positive*.

4.1. *One-referral Markov chain sampling.* The first model is a procedure where each respondent is asked (or rather “forced”) to give exactly one referral, resulting in one single chain of referrals. We assume that these referrals happen in a Markovian way, and a respondent of type x chooses the referral proportional to $\kappa(x, y) dy$. More rigorously, define the Markov kernel

$$(4.3) \quad K_\kappa(x, dy) := \frac{\kappa(x, y) dy}{\int_0^1 \kappa(x, v) dv}.$$

Under (4.2), the kernel is well defined.

PROPOSITION 4.1. *Let κ be a positive and connected standard kernel, and let $X = (X_1, X_2, \dots)$ be a Markov chain with Markov kernel K_κ . Then X has a unique invariant probability measure π given by*

$$(4.4) \quad \frac{\pi(dx)}{dx} = \frac{\int_0^1 \kappa(x, v) dv}{\int_0^1 \int_0^1 \kappa(u, v) du dv}.$$

Furthermore, for every measurable and bounded function f and for almost every $x \in [0, 1]$ we have

$$(4.5) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n f(X_i) = \int_0^1 f(x)\pi(dx)$$

$\mathbb{P}[\cdot | X_1 = x]$ -almost surely.

PROOF. Let us first prove that K_κ does not have any invariant measures with atoms. Assume that ρ is an invariant measure. Write $\rho = \rho^* + \rho'$, where ρ^* is the atomic and ρ' is the nonatomic parts, and assume that ρ^* is not the zero measure. Let A^* be the support of ρ^* ; note that A^* is countable and that $\rho(A^*) > 0$. However, $K(x, A^*) = 0$ for all $x \in [0, 1]$ due to (4.3) and, therefore,

$$\rho(A^*) = \int K_\kappa(x, A^*) d\rho(x) = 0,$$

which is a contradiction.

We now use Yosida’s ergodic decomposition to prove that π is the only invariant probability measure with respect to K_κ and that (4.5) holds; see [23] and [11].

Recall that an *invariant set* is a set A such that $K_\kappa(x, A) = 1$ for all $x \in A$, that is,

$$(4.6) \quad \int_A \kappa(x, y) dy = \int_0^1 \kappa(x, y) dy \quad \text{for all } x \in A.$$

Hence, we must have $\int_{A^c} \kappa(x, y) dy = 0$ for all $x \in A$. This implies that $\int_A \int_{A^c} \kappa(x, y) dy dx = 0$, which by symmetry of κ and (4.1), implies that $\text{Vol}(A) = 0$ or $\text{Vol}(A) = 1$. The case $\text{Vol}(A) = 0$ can be excluded since the right-hand side of (4.6) is positive by (4.2). By the definition of π , it follows that, for every such invariant set A , we have $\pi(A) = 1$. Therefore, π is an ergodic measure in the Yosida sense. Now, Lemma 4.2 of [11] implies that π is unique on the invariant sets up to π -null sets, but since π cannot have any atoms, π is unique on $[0, 1]$; (4.5) now follows from Theorem 6.1(b) of [11]. \square

It is worthwhile mentioning that (4.5) holds even if the Markov chain exhibits certain periodic behaviour. For example, if κ is such that the resulting graph is bipartite, the resulting Markov chain does not converge to its stationary distribution, but it is still ergodic.

4.2. *A Poisson branching process model.* Let us consider a continuous-time, multi-type Galton–Watson branching process with type space $[0, 1]$ as follows. A particle of type $x \in [0, 1]$ is assumed to have a standard exponential lifetime and during that time it will give birth to new particles of type y at rate $\lambda\kappa(x, y) dy$ for some $\lambda > 0$ independently of all else. Let T_t be the random point measure on $[0, 1]$ given by all particles ever born up to and including time t . We denote by δ_x the point unit measure at $x \in [0, 1]$ and we write $\mathbb{P}_x[\cdot] = \mathbb{P}[\cdot | T_0 = \delta_x]$. Note that $T_t[0, 1]$ the total number of points in T_t .

In order to push all arguments through as easily as possible, we will not only assume that the standard kernel positive and connected, but make the (most likely unnecessarily) strong assumption that

$$(4.7) \quad \inf_{0 \leq x, y \leq 1} \kappa(x, y) > 0.$$

It is clear that (4.7) implies that κ is both, connected and positive.

PROPOSITION 4.2. *Let $\lambda > 0$, let κ be a standard kernel satisfying (4.7), and let T_t be the resulting branching process. Then there exists α^* and a unique probability measure π on $[0, 1]$ satisfying*

$$(4.8) \quad \frac{\pi(dx)}{dx} = \frac{\lambda}{1 + \alpha^*} \int_0^1 \kappa(x, u)\pi(du).$$

Furthermore, if $\alpha^* > 0$, then $\mathbb{P}_x[|T_t| \rightarrow \infty] > 0$, and, for any measurable and bounded function f and for almost all $x \in [0, 1]$,

$$(4.9) \quad \lim_{t \rightarrow \infty} \frac{1}{|T_t|} \int_0^1 f(y)T_t(dy) = \int_0^1 f(y)\pi(dy)$$

$\mathbb{P}_x[\cdot | |T_t| \rightarrow \infty]$ -almost surely.

PROOF. We follow the setup of [14]; see also [13]. Define the reproduction kernel

$$\mu(x, dy \times dt) = e^{-t}\lambda\kappa(x, y) dy dt,$$

which, loosely speaking, is the expected number of offspring of type y that a particle of type x , born at time 0, produces at time t (the prefactor e^{-t} is simply the probability that the x -particle survives until time t). Furthermore, define the transition kernel

$$\hat{\mu}_\alpha(x, dy) = \int_0^\infty e^{-\alpha t} \mu(x, dy \times dt) = \frac{\lambda}{1 + \alpha} \kappa(x, y) dy.$$

It is not difficult to see that (4.7) implies that the kernel $\hat{\mu}_0$ is irreducible with respect to the Lebesgue measure on $[0, 1]$ (cf. [21], Example 2.1(b), page 11). Hence, there is a number α^* such that the kernel μ_α has convergence radius 1

(cf. [20], Proposition 2.1). The parameter α^* is commonly called the *Malthusian parameter*. Moreover, (4.7) also implies that μ_α is recurrent (cf. [19], Lemma 2.3). Hence, there is a σ -finite measure π and a strictly positive function h defined on $[0, 1]$ (cf. [14], page 42 and [21], Theorem 5.1) such that

$$(4.10) \quad \int_0^1 \hat{\mu}_{\alpha^*}(x, dy)\pi(dx) = \pi(dy)$$

and

$$\int_0^1 h(y)\hat{\mu}_{\alpha^*}(x, dy) = h(x).$$

Note that (4.10) is just (4.8). It is also straightforward to show that μ is *positive recurrent* (cf. [14], page 43). Since

$$h(x) = \int_0^1 h(y) \frac{\lambda}{1 + \alpha^*} \kappa(x, y) dy \geq \frac{\lambda}{1 + \alpha^*} \inf_{x,y} \kappa(x, y) \int_0^1 h(y) dy,$$

it is clear that $\inf h(x) > 0$. This implies that π is finite and can be normed to a probability measure (cf. [14], page 43) and h can be chosen so that $\int h(x)\pi(dx) = 1$. Finally, it is clear that μ is nonlattice and that there is $\varepsilon > 0$ such that

$$\sup_x \mu(x, [0, 1] \times [0, \varepsilon]) < 1.$$

These conditions are summarised as μ being *nonlattice and strictly Malthusian*.

Note that, since $\kappa \leq 1$, $|T_t|$ can be dominated by a unitype branching process where each particle has standard exponential lifetime and produces offspring at rate λ . Therefore, for fixed t , $|T_t|$ is uniformly integrable in the type of the starting particle. Moreover, the usual “ $x \log x$ ” condition follows easily from the fact that the dominating branching process has finite variance. Applying [14], Theorem 2, it follows that, for almost all x and for $A \subset [0, 1]$,

$$(4.11) \quad e^{-\alpha^*t} T_t(A) \rightarrow \frac{\pi(A)}{\alpha^* \beta} W$$

\mathbb{P}_x -almost surely for some nonnegative random variable W that satisfies $\mathbb{E}_x W = h(x)$, and for some β (which is explicit, but not of interest here). Note that clearly $\{W > 0\} \subset \{|T_t| \rightarrow \infty\}$, but it is not immediate that the two sets are equal. In order to make statements about (4.11) with $e^{-\alpha^*t}$ replaced by $1/|T_t|$, we need that

$$(4.12) \quad \inf_x \mathbb{P}_x[W > 0] > 0,$$

which guarantees that $\{W > 0\} = \{|T_t| \rightarrow \infty\}$ by [14], Lemma 1.

In order to prove (4.12), note that there must be a set A with $\text{Vol}(A) > 0$ such that $p_A := \inf_{x \in A} \mathbb{P}_x[W > 0] > 0$, for otherwise we would have $\mathbb{P}_x[W > 0] = 0$ for almost all x which is in contradiction to $\mathbb{E}_x W = h(x) > 0$ for almost all x . Let $M = \inf_{x,y} \kappa(x, y)$, which by (4.7) is positive, and

$$E_A = \{1\text{st particle has exactly one child of some type } y \in A\}.$$

Now,

$$\begin{aligned} \mathbb{P}_x[W > 0] &\geq \mathbb{P}_x[W > 0, E_A] \\ &= \mathbb{P}_x[W > 0|E_A]\mathbb{P}_x[E_A] \\ &\geq p_A \int_0^\infty e^{-t} \lambda M \text{Vol}(A) t e^{-\lambda M \text{Vol}(A)t} dt, \end{aligned}$$

which is a positive lower bound independent of x . Hence, (4.12) follows. From [14], Corollary 4, we have for almost all x that

$$\frac{T_t(A)}{|T_t|} \rightarrow \pi(A)$$

$\mathbb{P}_x[\cdot || T_t| \rightarrow \infty]$ -almost surely for any measurable $A \subset [0, 1]$. Since π is finite, it is easy to extend this to (4.9) for bounded f . \square

4.3. *A concrete standard kernel.* In this section, we consider a particular standard kernel $\kappa : [0, 1]^2 \rightarrow [0, 1]$ given as

$$(4.13) \quad \kappa(x, y) = \begin{cases} \alpha, & \text{if } 0 \leq x \leq \gamma \text{ and } 0 \leq y \leq \gamma, \\ \beta, & \text{if } \gamma < x \leq 1 \text{ and } \gamma < y \leq 1, \\ \delta, & \text{otherwise,} \end{cases}$$

where $0 < \alpha, \gamma, \delta, \beta < 1$.

One could think of κ as a graph between two groups of vertices. The internal connections between a primary group A (say) are specified by α and a secondary group B (specified) by β . The inter-connections between the groups of vertices are specified by δ . If we sample the vertices ergodically with invariant measure π , then Theorem 2.1 specifies that our limit graph will be governed by

$$\kappa_{\tau^{-1}}(x, y) = \begin{cases} \alpha, & \text{if } 0 \leq x \leq \tau(\gamma) \text{ and } 0 \leq y \leq \tau(\gamma), \\ \beta, & \text{if } \tau(\gamma) < x \leq 1 \text{ and } \tau(\gamma) < y \leq 1, \\ \delta, & \text{otherwise,} \end{cases}$$

where $\tau(x) = \pi([0, x])$; see Figure 1 for a graphical representation of the distortion in κ .

We shall now compare $\kappa_{\tau^{-1}}$ in the sampling procedures discussed in Sections 4.1 and 4.2. In the procedure discussed in Section 4.1, we have π given by (4.4) and a routine calculation gives us that the value of the distortion, denoted by τ_M , at γ is given by

$$(4.14) \quad \tau_M(\gamma) = \frac{(\alpha\gamma - \delta\gamma + \delta)\gamma}{\alpha\gamma^2 - 2\gamma^2\delta + \beta\gamma^2 + 2\delta\gamma - 2\beta\gamma + \beta}.$$

In the procedure discussed in Section 2, we need to find $\pi(dx)/dx = \nu(x)$, which satisfies

$$\nu(x) = \frac{\lambda}{1 + \alpha^*} \int_0^1 \kappa(x, y)\nu(y) dy, \quad \int_0^1 \nu(x) dx = 1.$$

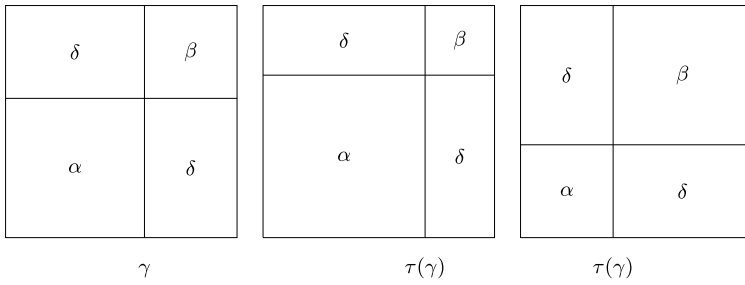


FIG. 1. Above is a pictorial representation of κ (left), $\kappa_{\tau^{-1}}$ when $\tau(\gamma) > \gamma$ (middle), and $\kappa_{\tau^{-1}}$ when $\tau(\gamma) < \gamma$ (right).

In the case of (4.13), this is equivalent to finding the largest eigenvalue and corresponding eigenvector of a (2×2) -matrix. A standard calculation then shows that the value of the distortion, denoted by τ_P , at γ is given by

$$(4.15) \quad \tau_P(\gamma) = \frac{(\alpha + \beta)\gamma - \beta + s}{2\delta + \gamma(\alpha - 2\delta + \beta) - \beta + s},$$

where

$$s = \sqrt{\gamma^2((\alpha + \beta)^2 - 4\delta^2) + 2\gamma(2\delta^2 - \alpha\beta - \beta^2) + \beta^2}.$$

In general, the formulae (4.14) and (4.15) do not compare in an obvious manner with themselves or with the unbiased sampling $[\tau(\gamma) = \gamma]$. In the one-referral Markov chain sampling model, a new vertex is chosen proportional to the values of κ , with the proportionality constant being the volume measure under κ . In contrast, in the Poisson branching process model, due to the branching effect, the offspring of a vertex will be from the regions governed by the sectional area of κ at the vertex. Thus, it is natural to expect differences in bias between the two procedures. We illustrate this via three examples of α, β, δ to illustrate the differences in distortion between the two sampling procedures.

The first example we consider is when $\alpha = 1/5, \delta = 1/200, \beta = 1/5$. In Figure 2, we plot $\tau(\gamma)$ as a function of γ . One can quickly observe that for $\gamma = 0.5$ there is no distortion in either sampling scheme as expected with $\tau_M(0.5) = \tau_P(0.5) = 0.5$. One observes that when $\gamma < 0.5$ then $\tau_P(\gamma) < \tau_M(\gamma) < \gamma$ and when $\gamma > 0.5$ then $\gamma < \tau_M(\gamma) < \tau_P(\gamma)$. This indicates that the Poisson branching process model will result in a larger bias toward the larger group (the secondary group B when $\gamma < 0.5$ and the primary group A when $\gamma > 0.5$). This is expected as both the primary group A and secondary group B are similarly well connected internally, but a small δ implies that they are poorly interconnected.

In the second example (see Figure 3), we consider $\alpha = 1/5, \delta = 1/5$ and $\beta = 1/200$. Group A has a fair number of connections within itself, and there are fair number of connections between the Groups A and B, but with a small β , Group

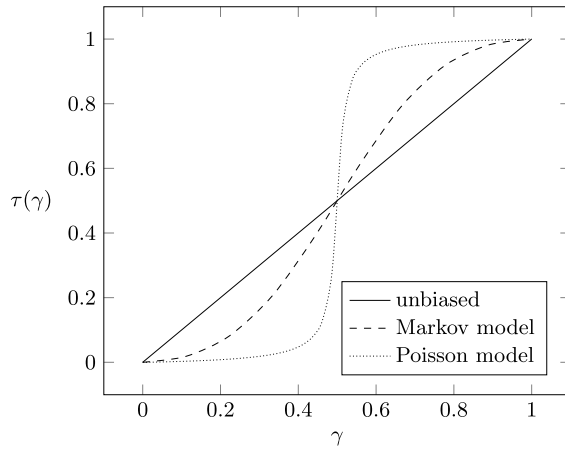


FIG. 2. Distortion plot for two large groups that are internally well connected, but where there are not many connections between the groups ($\alpha = 1/5, \delta = 1/200, \beta = 1/5$).

B has a smaller number of connections within itself. Note that for all $0 < \gamma < 1$, $\gamma < \tau_P(\gamma) < \tau_M(\gamma)$, indicating a larger bias in the one-referral Markov chain sampling procedure.

Finally, we consider the case $\alpha = 1/5, \delta = 1/200, \beta = 1/200$ (see Figure 4). With interconnection probability and within Group B connection probabilities being small this time there is a strong bias toward selecting vertices from the primary Group A. In plot shown in Figure 4, we can see that the bias is more pronounced this time in the Poisson branching process sampling procedure.

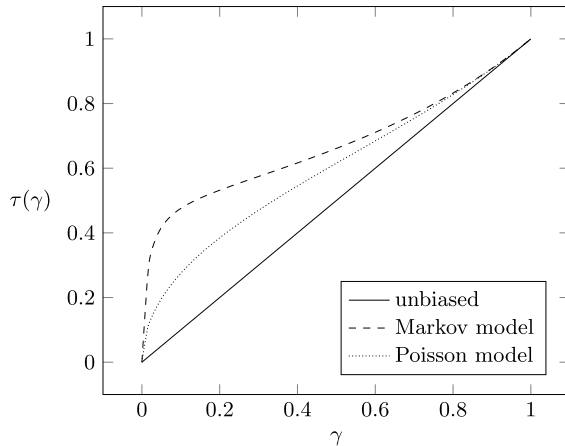


FIG. 3. Distortion plot for a graph where the primary group is well connected to itself and to a secondary group, but where the secondary group is not well connected within itself ($\alpha = 1/5, \delta = 1/5, \beta = 1/200$).

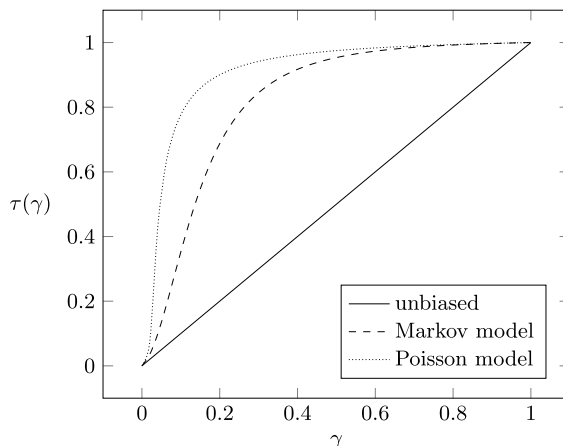


FIG. 4. Distortion plot for a graph where the primary group is well connected to itself, and where the secondary group is neither well connected to the primary group nor within itself ($\alpha = 1/5$, $\delta = 1/200$, $\beta = 1/200$).

In conclusion, depending on the size, connectedness of the groups and inter-connections between them, the sampling scheme has to be chosen appropriately to control the bias.

Acknowledgements. We thank Andrew Barbour and Don Dawson for helpful discussions, and we are grateful to the anonymous referee for making suggestions that resulted in improving the paper and for pointing out an important error in the proof of the main theorem; in particular, this resulted in the introduction of assumption (ii) in Theorem 2.1.

Parts of this work was done when S. Athreya visited the Institute of Mathematical Sciences, NUS. The authors would like to express their gratitude for the kind hospitality.

REFERENCES

- [1] AARONSON, J., BURTON, R., DEHLING, H., GILAT, D., HILL, T. and WEISS, B. (1996). Strong laws for L - and U -statistics. *Trans. Amer. Math. Soc.* **348** 2845–2866. [MR1363941](#)
- [2] ALDOUS, D. J. (1981). Representations for partially exchangeable arrays of random variables. *J. Multivariate Anal.* **11** 581–598. [MR0637937](#)
- [3] BOLLOBÁS, B. and RIORDAN, O. (2009). Metrics for sparse graphs. In *Surveys in Combinatorics 2009. London Mathematical Society Lecture Note Series* **365** 211–287. Cambridge Univ. Press, Cambridge. [MR2588543](#)
- [4] BORGS, C., CHAYES, J. T., LOVÁSZ, L., SÓS, V. T. and VESZTERGOMBI, K. (2008). Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing. *Adv. Math.* **219** 1801–1851. [MR2455626](#)
- [5] BORGS, C., CHAYES, J. T., LOVÁSZ, L., SÓS, V. T. and VESZTERGOMBI, K. (2012). Convergent sequences of dense graphs II. Multiway cuts and statistical physics. *Ann. of Math.* (2) **176** 151–219. [MR2925382](#)

- [6] DIACONIS, P. and JANSON, S. (2008). Graph limits and exchangeable random graphs. *Rend. Mat. Appl. (7)* **28** 33–61. [MR2463439](#)
- [7] GILE, K. J. and HANDCOCK, M. S. (2010). Respondent-driven sampling: An assessment of current methodology. *Sociological Methodology* **40** 285–327.
- [8] GOEL, S. and SALGANIK, M. J. (2010). Assessing respondent-driven sampling. *Proc. Natl. Acad. Sci. USA* **107** 6743–6747.
- [9] HECKATHORN, D. D. (1997). Respondent-driven sampling: A new approach to the study of hidden populations. *Social Problems* **44** 174–199.
- [10] HECKATHORN, D. D. (2002). Respondent-driven sampling II: Deriving valid population estimates from chain-referral samples of hidden populations. *Social Problems* **49** 11–34.
- [11] HERNÁNDEZ-LERMA, O. and LASSERRE, J. B. (1998). Ergodic theorems and ergodic decomposition for Markov chains. *Acta Appl. Math.* **54** 99–119. [MR1656784](#)
- [12] HOOVER, D. N. (1979). Relations on probability spaces and arrays of random variables. Preprint, Institute for Advanced Studies. Princeton, NJ.
- [13] JAGERS, P. (1989). General branching processes as Markov fields. *Stochastic Process. Appl.* **32** 183–212. [MR1014449](#)
- [14] JAGERS, P. and NERMAN, O. (1996). The asymptotic composition of supercritical multi-type branching populations. In *Séminaire de Probabilités, XXX. Lecture Notes in Math.* **1626** 40–54. Springer, Berlin. [MR1459475](#)
- [15] JANSON, S. (2008). Connectedness in graph limits. Technical Report 2008:8, Uppsala.
- [16] LOVÁSZ, L. (2012). *Large Networks and Graph Limits. American Mathematical Society Colloquium Publications* **60**. Amer. Math. Soc., Providence, RI. [MR3012035](#)
- [17] LOVÁSZ, L. and SZEGEDY, B. (2006). Limits of dense graph sequences. *J. Combin. Theory Ser. B* **96** 933–957. [MR2274085](#)
- [18] MCDIARMID, C. (1989). On the method of bounded differences. In *Surveys in Combinatorics, 1989 (Norwich, 1989). London Mathematical Society Lecture Note Series* **141** 148–188. Cambridge Univ. Press, Cambridge. [MR1036755](#)
- [19] NEY, P. and NUMMELIN, E. (1987). Markov additive processes. I. Eigenvalue properties and limit theorems. *Ann. Probab.* **15** 561–592. [MR0885131](#)
- [20] NIEMI, S. and NUMMELIN, E. (1986). On nonsingular renewal kernels with an application to a semigroup of transition kernels. *Stochastic Process. Appl.* **22** 177–202. [MR0860932](#)
- [21] NUMMELIN, E. (1984). *General Irreducible Markov Chains and Nonnegative Operators. Cambridge Tracts in Mathematics* **83**. Cambridge Univ. Press, Cambridge. [MR0776608](#)
- [22] VOLZ, E. and HECKATHORN, D. D. (2008). Probability based estimation theory for respondent driven sampling. *Journal of Official Statistics* **24** 79–97.
- [23] YOSIDA, K. (1980). *Functional Analysis*, 6th ed. *Grundlehren der Mathematischen Wissenschaften* **123**. Springer, Berlin. [MR0617913](#)

STATMATH UNIT
INDIAN STATISTICAL INSTITUTE
8TH MILE MYSORE ROAD
BANGALORE, 560059
INDIA
E-MAIL: athreya@isibang.ac.in

DEPARTMENT OF STATISTICS AND APPLIED PROBABILITY
NATIONAL UNIVERSITY OF SINGAPORE
6 SCIENCE DRIVE 2
SINGAPORE 117546
E-MAIL: adrian.roellin@nus.edu.sg