

ARE EXTREME VALUE ESTIMATION METHODS USEFUL FOR NETWORK DATA?

PHYLLIS WAN, TIANDONG WANG, RICHARD A. DAVIS, AND SIDNEY I. RESNICK

ABSTRACT. Preferential attachment is an appealing edge generating mechanism for modeling social networks. It provides both an intuitive description of network growth and an explanation for the observed power laws in degree distributions. However, there are often limitations in fitting parametric network models to data due to the complex nature of real-world networks. In this paper, we consider a semi-parametric estimation approach by looking at only the nodes with large in- or out-degrees of the network. This method examines the tail behavior of both the marginal and joint degree distributions and is based on extreme value theory. We compare it with the existing parametric approaches and demonstrate how it can provide more robust estimates of parameters associated with the network when the data are corrupted or when the model is misspecified.

1. INTRODUCTION

Empirical studies [18] suggest that the distribution of in- and out-degrees of the nodes of many social networks have Pareto-like tails. The indices of these distributions control the likelihood of nodes with large degrees appearing in the data. Some social network models, such as preferential attachment, theoretically exhibit these heavy-tailed characteristics. This paper estimates heavy tail parameters using semi-parametric extreme value (EV) methods and compares such EV estimates with model-based likelihood methods. The EV estimates only rely on the upper tail of the degree distributions so one might expect these estimates to be robust against model error or data corruption.

Preferential attachment (PA) describes the growth of a network where edges and nodes are added over time based on probabilistic rules that assume existing nodes with large degrees attract more edges. This property is attractive for modeling social networks due to intuitive appeal and ability to produce power-law networks with degrees matched to data [3, 9, 16, 17, 23]. Elementary descriptions of the preferential attachment model can be found in [10] while more mathematical treatments are available in [1, 9, 23]. Also see [15] for a statistical survey of methods for network data and [11] for inference for an undirected model.

The linear preferential attachment model has received most attention. Marginal degree power laws were established in [3, 16, 17], while joint power-law behavior, also known as joint regular variation, was proved in [21, 22, 26] for the directed linear PA model. Given observed network data, [25] proposed parametric inference procedures for the model in two data scenarios. For the case where the history of network growth is available, the MLE estimators

Key words and phrases. power laws, multivariate heavy-tailed statistics, preferential attachment, regular variation, estimation.

Research of the four authors was partially supported by Army MURI grant W911NF-12-1-0385 to Cornell University.

of model parameters were derived and shown to be strongly consistent, asymptotically normal and efficient. For the case where only a snapshot of the network is available at a single time point, the estimators based on moment methods as well as an approximation to the likelihood were shown to be strongly consistent. The loss of efficiency relative to full MLE was surprisingly mild.

The drawback of these two methods is that they are model-based and sensitive to model error. To overcome this lack of robustness, this paper describes an EV inference method applied to a single snapshot of a network and where possible, compares the EV method to model-based MLE methods. The EV method is based on estimates of in- and out-degree tail indices, ι_{in} and ι_{out} , using a combination of the Hill estimator [13, 20] coupled with a minimum distance threshold selection method [5]. We also describe estimation of model parameters using the joint tail distribution of in- and out-degrees relying on the asymptotic angular measure [20, page 173] density obtained after standardizing [20, page 203] the data.

If the data are generated by the linear PA model, the EV estimators can be applied to estimate the parameters of the model and compared with MLE estimates and not surprisingly, the EV estimates exhibit larger variance. However, if there is model error or data corruption, the EV estimates more than hold their own and we illustrate the comparison in two ways:

- The data is corrupted; linear PA data have edges randomly deleted or added. The EV approach reliably recovers the original preferential attachment parameters while parametric methods degrade considerably.
- The data comes from a misspecified model, namely a directed edge superstar model [2] but is analyzed as if it comes from the linear PA model. The EV method gives good estimates for superstar model tail indices and outperforms MLE based on a misspecified linear PA model if the probability of attaching to the superstar is significant.

The rest of the paper is structured as follows. Section 2 formulates the power-law phenomena in network degree distributions along with joint dependency in the in- and out-degrees. We describe two network models which exhibit such heavy tail properties, the linear PA and the superstar linear PA models. The EV inference method for networks is described in Section 3 where we discuss its use for estimating the parameters of the linear PA model. Section 4 gives EV estimation results for simulated data from the linear PA model. Since the generating model is correctly specified, we use the previous parametric methods as benchmarks for comparison in Section 4.1. Section 4.2 analyzes network data generated from the linear PA model but corrupted by random edge addition or deletion. Pretending ignorance of the perturbation, we compare the performance of the extreme value method with the MLE and snapshot methods to recover the original model. In Section 4.3, we use our EV inference approach on data from the directed superstar model and attempt to to recover the tail properties of the degree distributions. A concluding Section 5 summarizes the discussion and reasons why EV methods have their place. Appendices give proofs and a fuller discussion of MLE and the snapshot method for linear PA models abstracted from [25].

2. NETWORKS AND HEAVY-TAILED DEGREE DISTRIBUTIONS

2.1. General discussion. We begin with a general discussion of power laws and networks. Let $G(n) = (V(n), E(n))$ denote a directed network, where $V(n)$ is the set of nodes, $E(n)$ is the set of edges, and n is the number of edges. Let $N(n)$ denote the number of nodes in $G(n)$ and $N_n(i, j)$ be the number of nodes with in-degree i and out-degree j . The marginal counts of nodes with in-degree i and out-degree j are given by

$$N_i^{\text{in}}(n) := \sum_{j=0}^{\infty} N_n(i, j) \quad \text{and} \quad N_j^{\text{out}}(n) := \sum_{i=0}^{\infty} N_n(i, j),$$

respectively. For many network data sets, log-log plots of the in- and out-degree distributions, i.e., plots of $\log i$ vs. $\log N_i^{\text{in}}(n)$ and $\log j$ vs. $\log N_j^{\text{out}}(n)$, appear to be linear and generative models of network growth seek to reflect this. Consider models such that the empirical degree frequency converges almost surely,

$$(2.1) \quad N_n(i, j)/N(n) \rightarrow p_{ij}, \quad (n \rightarrow \infty)$$

where p_{ij} is a bivariate probability mass function (pmf). The network exhibits power-law behavior if

$$(2.2) \quad p_i^{\text{in}} := \sum_{j=0}^{\infty} p_{ij} \sim C_{\text{in}} i^{-(1+\iota_{\text{in}})} \text{ as } i \rightarrow \infty,$$

$$(2.3) \quad p_j^{\text{out}} := \sum_{i=0}^{\infty} p_{ij} \sim C_{\text{out}} j^{-(1+\iota_{\text{out}})} \text{ as } j \rightarrow \infty,$$

for some positive constants $C_{\text{in}}, C_{\text{out}}$. Let (I, O) be a fictitious random vector with joint pmf p_{ij} , then

$$\begin{aligned} \mathbf{P}(I \geq i) &\sim C_{\text{in}}(1 + \iota_{\text{in}})^{-1} \cdot i^{-\iota_{\text{in}}} \text{ as } i \rightarrow \infty, \\ \mathbf{P}(O \geq j) &\sim C_{\text{out}}(1 + \iota_{\text{out}})^{-1} \cdot j^{-\iota_{\text{out}}} \text{ as } j \rightarrow \infty. \end{aligned}$$

In the linear PA model, the joint distribution of (I, O) satisfies non-standard regular variation. Let $\mathbb{M}(\mathbb{R}_+^2 \setminus \{\mathbf{0}\})$ be the set of Borel measures on $\mathbb{R}_+^2 \setminus \{\mathbf{0}\}$ that are finite on sets bounded away from the origin. Then (I, O) is *non-standard regularly varying* on $\mathbb{R}_+^2 \setminus \{\mathbf{0}\}$ means that as $t \rightarrow \infty$,

$$(2.4) \quad t\mathbf{P} \left[\left(\left(\frac{I}{t^{1/\iota_{\text{in}}}}, \frac{O}{t^{1/\iota_{\text{out}}}} \right) \in \cdot \right) \right] \rightarrow \nu(\cdot), \quad \text{in } \mathbb{M}(\mathbb{R}_+^2 \setminus \{\mathbf{0}\}),$$

where $\nu(\cdot) \in \mathbb{M}(\mathbb{R}_+^2 \setminus \{\mathbf{0}\})$ is called the limit or tail measure [7, 14, 19]. Using the power transformation $I \mapsto I^a$ with $a = \iota_{\text{in}}/\iota_{\text{out}}$, the vector (I^a, O) becomes standard regularly varying, i.e.,


$$(2.5) \quad t\mathbf{P} \left[\left(\left(\frac{I^a}{t^{1/\iota_{\text{out}}}}, \frac{O}{t^{1/\iota_{\text{out}}}} \right) \in \cdot \right) \right] \rightarrow \tilde{\nu}(\cdot), \quad \text{in } \mathbb{M}(\mathbb{R}_+^2 \setminus \{\mathbf{0}\}),$$

where $\tilde{\nu} = \nu \circ T^{-1}$ with $T(x, y) = (x^a, y)$. With this standardization, the transformed measure $\tilde{\nu}$ is directly estimable from data [20].

In the following we describe two classes of preferential attachment models that generate networks with power-law degree distributions.

2.2. The linear preferential attachment (linear PA) model. The directed linear PA model [3, 17] constructs a growing sequence of directed random graphs $G(n)$'s whose dynamics depend on five nonnegative parameters $\alpha, \beta, \gamma, \delta_{\text{in}}$ and δ_{out} , where $\alpha + \beta + \gamma = 1$ and $\delta_{\text{in}}, \delta_{\text{out}} > 0$. To avoid degenerate situations, assume that each of the numbers α, β, γ is strictly smaller than 1.

We start with an arbitrary initial finite directed graph $G(n_0)$ with at least one node and n_0 edges. Given an existing graph $G(n-1)$, a new graph $G(n)$ is obtained by adding a single edge to $G(n-1)$ so that the graph $G(n)$ contains n edges for all $n \geq n_0$. Let $I_n(v)$ and $O_n(v)$ denote the in- and out-degree of $v \in V(n)$ in $G(n)$, that is, the number of edges pointing into and out of v , respectively. We allow three scenarios of edge creation, which are activated by flipping a 3-sided coin with probabilities α, β and γ . More formally, let $\{J_n, n > n_0\}$ be an iid sequence of trinomial random variables with cells labelled 1, 2, 3 and cell probabilities α, β, γ . Then the graph $G(n)$ is obtained from $G(n-1)$ as follows.



- If $J_n = 1$ (with probability α), append to $G(n-1)$ a new node $v \in V(n) \setminus V(n-1)$ and an edge (v, w) leading from v to an existing node $w \in V(n-1)$. Choose the existing node $w \in V(n-1)$ with probability depending on its in-degree in $G(n-1)$:

$$(2.6) \quad \mathbf{P}[\text{choose } w \in V(n-1)] = \frac{I_{n-1}(w) + \delta_{\text{in}}}{n-1 + \delta_{\text{in}}N(n-1)}.$$

- If $J_n = 2$ (with probability β), add a directed edge (v, w) to $E(n-1)$ with $v \in V(n-1) = V(n)$ and $w \in V(n-1) = V(n)$ and the existing nodes v, w are chosen independently from the nodes of $G(n-1)$ with probabilities

$$(2.7) \quad \mathbf{P}[\text{choose } (v, w)] = \left(\frac{O_{n-1}(v) + \delta_{\text{out}}}{n-1 + \delta_{\text{out}}N(n-1)} \right) \left(\frac{I_{n-1}(w) + \delta_{\text{in}}}{n-1 + \delta_{\text{in}}N(n-1)} \right).$$

- If $J_n = 3$ (with probability γ), append to $G(n-1)$ a new node $v \in V(n) \setminus V(n-1)$ and an edge (w, v) leading from the existing node $w \in V(n-1)$ to the new node v . Choose the existing node $w \in V(n-1)$ with probability

$$(2.8) \quad \mathbf{P}[\text{choose } w \in V(n-1)] = \frac{O_{n-1}(w) + \delta_{\text{out}}}{n-1 + \delta_{\text{out}}N(n-1)}.$$

For convenience we call these scenarios the α -, β - and γ -schemes. Note that this construction allows for the possibility of multiple edges between two nodes and self loops. This linear preferential attachment model can be simulated efficiently using the method described in [25, Algorithm 1] and linked to <http://www.orie.cornell.edu/orie/research/groups/multheavytail/software.cfm>.

It is shown in [21, 22, 26] that the empirical degree distribution

$$\frac{N_n(i, j)}{N(n)} \xrightarrow{\text{a.s.}} p_{ij},$$

and the marginals satisfy (2.2) and (2.3), where the tail indices are

$$(2.9) \quad \iota_{\text{in}} =: \frac{1 + \delta_{\text{in}}(\alpha + \gamma)}{\alpha + \beta}, \quad \text{and} \quad \iota_{\text{out}} =: \frac{1 + \delta_{\text{out}}(\alpha + \gamma)}{\beta + \gamma}.$$

Furthermore, the joint regular variation condition (2.5) is satisfied by the limit degree distribution and the limit measure [22] or its density [26] can be explicitly derived. We shall use this property for parameter estimation in Section 3.

2.3. The superstar linear PA model. The key feature of the superstar linear PA model that distinguishes it from the standard linear PA model is the existence of a superstar node, to which a large proportion of nodes attach. A new parameter p represents the attachment probability. The α -, β - and γ -schemes of the linear PA model are still in action. However, for the α - and β -schemes, an outgoing edge will attach to the superstar node with probability p , while with probability $1 - p$ it will attach to a non-superstar node according to the original linear PA rules.

For simplicity, the network is initialized with two nodes $V(1) = \{0, 1\}$ where node 0 is the superstar node. We assume at the first step, there is an edge pointing from $1 \rightarrow 0$ so $E_1 = \{(1, 0)\}$. Again each graph $G(n)$ contains n edges for all $n \geq 1$. Let

$$V^0(n) := V(n) \setminus \{0\}, \quad \text{and} \quad E^0(n) := E(n) \setminus \{(u, 0) : u \in V^0(n)\},$$

so that $E^0(n)$ is the set of edges in $G(n)$ that do not point to the superstar. Let $|V^0(n)|$ and $|E^0(n)|$ denote the number of nodes and edges in the non-superstar subgraph of $G(n)$, respectively.

The model is specified through the parameter set $(p, \alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}})$. Let $\{B_n : n \geq 1\}$ be another iid sequence of Bernoulli random variables where

$$\mathbf{P}(B_n = 1) = p = 1 - \mathbf{P}(B_n = 0).$$

The Markovian graph evolution from $G(n-1)$ to $G(n)$ is modified from the linear PA model as follows.

- If $J_n = 1$ (with probability α), append to $G(n-1)$ a new node $v \in V(n) \setminus V(n-1)$ and an edge (v, w) leading from v to an existing node w .
 - If $B_n = 1$ (with probability p), $w = 0$, the superstar node;
 - If $B_n = 0$ (with probability $1 - p$), $w \in V^0(n-1)$ is chosen according to the linear PA rule (2.6) applied to $(V^0(n-1), E^0(n-1))$.
- If $J_n = 2$ (with probability β), add a directed edge (v, w) to $E(n-1)$ where

- If $B_n = 1$ (with probability p), $v = 0$ and $w \in V^0(n-1) = V^0(n)$ is chosen with probability (2.6) applied to $(V^0(n-1), E^0(n-1))$;
- If $B_n = 0$ (with probability $1-p$), $v, w \in V^0(n-1) = V^0(n)$ are chosen with probability (2.7) applied to $(V^0(n-1), E^0(n-1))$.
- If $J_n = 3$ (with probability γ), append to $G(n-1)$ a new node $w \in V^0(n) \setminus V^0(n-1)$ and an edge (v, w) leading from the existing node $v \in V^0(n-1)$ to w , where $v \in V^0(n-1)$ is chosen with probability (2.8) applied to $(V^0(n-1), E^0(n-1))$.

If we use $N_i^{\text{in}}(n)$ and $N_j^{\text{out}}(n)$ to denote the number of *non*-superstar nodes that have in-degree i and out-degree j , respectively, then Theorem 2.1 shows that $(N_i^{\text{in}}(n)/n, N_j^{\text{out}}(n)/n) \rightarrow (q_i^{\text{in}}, q_j^{\text{out}})$ almost surely where the limits are deterministic constants that decay like power laws.

Theorem 2.1. *Let $(N_i^{\text{in}}(n), N_j^{\text{out}}(n))$ be the in- and out-degree counts of the non-superstar nodes of the superstar model. There exists constants q_i^{in} and q_j^{out} such that as $n \rightarrow \infty$,*

$$\frac{N_i^{\text{in}}(n)}{n} \xrightarrow{\text{a.s.}} q_i^{\text{in}}, \quad \frac{N_j^{\text{out}}(n)}{n} \xrightarrow{\text{a.s.}} q_j^{\text{out}}.$$

Moreover,

(i) As $i \rightarrow \infty$,

$$(2.10) \quad q_i^{\text{in}} \sim C'_{\text{in}} i^{-(1+\iota_{\text{in}})},$$

where C'_{in} is a positive constant and

$$(2.11) \quad \iota_{\text{in}} := \frac{1 - (\alpha + \beta)p + \delta_{\text{in}}(\alpha + \gamma)}{(\alpha + \beta)(1 - p)}.$$

(ii) As $j \rightarrow \infty$,

$$(2.12) \quad q_j^{\text{out}} \sim C'_{\text{out}} j^{-(1+\iota_{\text{out}})},$$

where C'_{out} is a positive constant and

$$(2.13) \quad \iota_{\text{out}} := \frac{1 + \delta_{\text{out}}(\alpha + \gamma)}{\beta + \gamma}.$$

The proof of Theorem 2.1 is provided in Appendix B.

3. ESTIMATION USING EXTREME VALUE THEORY

In this section, we consider network parameter estimation using extreme value theory. Given a graph $G(n)$ at a fixed timestamp, the data available for estimates are the in- and out-degrees for each node denoted by $(I_n(v), O_n(v))$, $v = 1, \dots, N(n)$. Let $F_n(\cdot)$ be the empirical distribution of this data on $\mathbb{N} \times \mathbb{N}$. Then from (2.1), almost surely F_n converges weakly to a limit distribution F on $\mathbb{N} \times \mathbb{N}$ which is the measure corresponding to the mass function $\{p_{ij}\}$. Let $\epsilon_{(i,j)}(\cdot)$ be the Dirac measure concentrating on (i, j) and we have from (2.1),

$$(3.1) \quad F_n(\cdot) = \frac{1}{N(n)} \sum_{v=1}^{N(n)} \epsilon_{(I_n(v), O_n(v))}(\cdot) = \sum_{i,j} \frac{N_n(i, j)}{N(n)} \epsilon_{(i,j)}(\cdot) \xrightarrow{w} \sum_{i,j} p_{ij} \epsilon_{(i,j)}(\cdot) =: F(\cdot).$$

3.1. Estimating tail indices; Hill estimation. We review tail index estimation of ι_{in} (ι_{out} is similar) using the Hill estimator [13, 20] applied to in-degree data $I_n(v)$, $v = 1, \dots, N(n)$. From (2.2), the marginal of F , called F_{in} is regularly varying with index $-\iota_{\text{in}}$. From Karamata's theorem ι_{in}^{-1} can be expressed as a function of F_{in} [8, page 69],

$$(3.2) \quad \iota_{\text{in}}^{-1} = \lim_{t \rightarrow \infty} \frac{\int_t^\infty (\log(u) - \log(t)) F_{\text{in}}(du)}{1 - F_{\text{in}}(t)}.$$

The Hill estimator of ι_{in}^{-1} replaces $F_{\text{in}}(\cdot)$ with the marginal of the empirical distribution in (3.1) of in-degrees, called $F_{\text{in},n}$, and t with $I_{(k_n+1)}$ in (3.2). Let $I_{(1)} \geq \dots \geq I_{(N(n))}$ be the decreasing order statistics of $I_n(v)$, $v = 1, \dots, N(n)$. The resulting estimator is

$$\begin{aligned} \hat{\iota}_{\text{in}}^{-1}(k_n) &= \frac{\int_{I_{(k_n+1)}}^\infty (\log(u) - \log(I_{(k_n+1)})) F_{\text{in},n}(du)}{k_n/N(n)} \\ &= \frac{1}{k_n} \sum_{j=1}^{k_n} (\log(I_{(j)}) - \log(I_{(k_n+1)})). \end{aligned}$$

With iid data, if we assume $k_n \rightarrow \infty$ and $k_n/N(n) \rightarrow 0$, then the Hill estimator is consistent. Of course, our network data is not iid but Hill estimation still works in practice. Consistency for an undirected graph is proven in [27] but for directed graphs, this is an unresolved issue.

To select k_n in practice, [5] proposed computing the Kolmogorov-Smirnov (KS) distance between the empirical distribution of the upper k observations and the power-law distribution with index $\hat{\iota}_{\text{in}}(k)$:

$$D_k := \sup_{y \geq 1} \frac{1}{k} \sum_{j=1}^k \mathbf{1}_{\{I_{(j)}/I_{(k+1)} > y\}} - y^{-\hat{\iota}_{\text{in}}(k)}, \quad 1 \leq k \leq n-1.$$

Then the optimal k^* is the one that minimizes the KS distance

$$k^* := \operatorname{argmin}_{1 \leq k \leq n} D_k,$$

and the tail index is estimated by $\hat{\iota}_{\text{in}}(k^*)$. This estimator performs well if the thresholded portion comes from a Pareto tail and also seems effective in a variety of non-iid scenarios. It is widely used by data repositories of large network datasets such as KONECT (<http://konect.uni-koblenz.de/>) [18] and is realized in the R-package *powerLaw* [12].

We refer to the above procedure as the *minimum distance method* in estimating ι_{in} , ι_{out} for network data. There are two issues when applying this method. First, the data is node-based and not collected from independent repeated sampling. Secondly, degree counts are discrete and do not exactly comply with the Pareto assumption made in the minimum distance method. Our analysis shows that even if we ignore these two issues, the tail estimates are still reasonably good.

3.2. Estimating dependency between in- and out-degrees. If the limiting random vector $(I, O) \sim F$ corresponding to p_{ij} in (2.1) is jointly regularly varying and satisfies (2.5), we may apply a polar coordinate transformation, for example, with the L_2 -norm,

$$(I^a, O) \mapsto (\sqrt{I^{2a} + O^2}, \arctan(O/I^a)) := (R, T),$$

where $a = \iota_{\text{in}}/\iota_{\text{out}}$. Then, with respect to F in (3.1), the conditional distribution of T given $R > r$ converges weakly (see, for example, [20, p. 173]),

$$F[T \in \cdot | R > r] \rightarrow S(\cdot), \quad r \rightarrow \infty,$$

where S is the *angular measure* and describes the asymptotic dependence of the standardized pair (I^a, O) . Since for large r , $F[T \in \cdot | R > r] \approx S(\cdot)$ and for large n , $F_n \approx F$, it is plausible that for r and n large $F_n[T \in \cdot | R > r] \approx S(\cdot)$. Skeptics may check [20, p. 307] for a more precise argument and recall F_n is the empirical measure defined in (3.1).

Based on observed degrees $\{(I_n(v), O_n(v)); v = 1, \dots, N(n)\}$, how does this work in practice? First a is replaced by $\hat{a} = \hat{\iota}_{\text{in}}/\hat{\iota}_{\text{out}}$ estimated from Section 3.1. Then the distribution S is estimated via the empirical distribution of the sample angles $T_n(v) := \arctan(O_n(v)/I_n(v)^{\hat{a}})$ for which $R_n(v) := \sqrt{I_n(v)^{2\hat{a}} + O_n(v)^2} > r$ exceeds some large threshold r . This is the POT (Peaks Over Threshold) methodology commonly employed in extreme value theory [6].

In the cases where the network model is known, S may be specified in closed form. For the linear PA model, S has a density that is an explicit function of the linear PA parameters [22]. After estimating ι_{in} and ι_{out} by the minimum distance method, the remaining parameters can then be estimated by an approximate likelihood method that we now explain.

3.3. EV estimation for the linear PA model. From (2.9),

$$\delta_{\text{in}} = \frac{\iota_{\text{in}}(\alpha + \beta) - 1}{\alpha + \gamma}, \quad \delta_{\text{out}} = \frac{\iota_{\text{out}}(\beta + \gamma) - 1}{\alpha + \gamma},$$

so that the linear PA model may be parameterized by $\boldsymbol{\theta} = (\alpha, \beta, \gamma, \iota_{\text{in}}, \iota_{\text{out}})$. To construct the EV estimates, begin by computing the minimum distance estimates $\hat{\iota}_{\text{in}}^{EV}, \hat{\iota}_{\text{out}}^{EV}$ of the in- and out-degree indices. The parameter β , which represents the proportion of edges connected between existing nodes, is estimated by $\hat{\beta}^{EV} = 1 - N(n)/n$.

From (2.5), $\arctan(O/I^a)$ given $I^{2a} + O^2 > r^2$ converges weakly as $r \rightarrow \infty$ to the distribution of a random variable Θ [22, Section 4.1.2], whose pdf is given by $(0 \leq x \leq \pi/2)$

$$\begin{aligned} f_{\Theta}(x; \alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}}) &\propto \frac{\gamma}{\delta_{\text{in}}} (\cos x)^{\frac{\delta_{\text{in}}+1}{a}-1} (\sin x)^{\delta_{\text{out}}-1} \int_0^{\infty} t^{\iota_{\text{in}}+\delta_{\text{in}}+a\delta_{\text{out}}} e^{-t(\cos x)^{1/a}-t^a \sin x} dt \\ &+ \frac{\alpha}{\delta_{\text{out}}} (\cos x)^{\frac{\delta_{\text{in}}}{a}-1} (\sin x)^{\delta_{\text{out}}} \int_0^{\infty} t^{a-1+\iota_{\text{in}}+\delta_{\text{in}}+a\delta_{\text{out}}} e^{-t(\cos x)^{1/a}-t^a \sin x} dt. \end{aligned} \quad (3.3)$$

By replacing $\beta, \iota_{\text{in}}, \iota_{\text{out}}$ with their estimated values $\hat{\beta}^{EV}, \hat{\iota}_{\text{in}}^{EV}$, and $\hat{\iota}_{\text{out}}^{EV}$ and setting $\gamma = 1 - \alpha - \hat{\beta}^{EV}$, the density (3.3) can be viewed as a profile likelihood function (based on a single observation x) of the unknown parameter α , which we denote by

$$l(\alpha; x) = f_{\Theta}(x; \alpha, \hat{\beta}^{EV}, 1 - \alpha - \hat{\beta}^{EV}, \hat{\delta}_{\text{in}}^{EV}, \hat{\delta}_{\text{out}}^{EV}).$$

Given the degrees $((I_n(v), O_n(v)), v \in V(n))$, $\hat{\alpha}^{EV}$ can be computed by maximizing the profile likelihood based on the observations $(I_n(v), O_n(v))$ for which $R_n(v) > r$ for a large threshold r . That is,

$$\hat{\alpha}^{EV} := \arg\max_{0 \leq \alpha \leq 1} \sum_{v=1}^{N(n)} \log l \left(\alpha; \arctan \left(\frac{O_n(v)}{(I_n(v))^{\hat{a}}} \right) \right) \mathbf{1}_{\{R_n(v) > r\}}, \quad (3.4)$$

where r is typically chosen as the $(n_{\text{tail}} + 1)$ -th largest $R_n(v)$'s for a suitable n_{tail} . This estimation procedure is sometimes referred to as the ‘‘independence estimating equations’’ (IEEs) method [4, 24], in which the dependence between observations is ignored. This technique is often used when the joint distribution of the data is unknown or intractable. Finally, using the constraint, $\alpha + \beta + \gamma = 1$, we estimate γ by $\hat{\gamma}^{EV} = 1 - \hat{\alpha}^{EV} - \hat{\beta}^{EV}$.

4. ESTIMATION RESULTS

In this section, we demonstrate the estimation of the linear PA and related models through the EV method described in Section 3.3. In Section 4.1, data are simulated from the standard linear PA model and used to estimate the true parameters of the underlying model. Section 4.2 considers data generated from the linear PA model but corrupted by random addition or deletion of edges. Our goal is to estimate the parameters of the original linear PA model. In Section 4.3, we simulate data from the superstar linear PA model and attempt to use the standard linear PA estimation to recover the degree distributions.

Throughout the section, the EV method is compared with two parametric estimation approaches for the linear PA model, namely the MLE and snapshot (SN) methods, proposed in [25]. For a given network, when the network history is available, that is, each edge is marked with the timestamp of its creation, MLE estimates are directly computable. In the case where only a snapshot of the network is given at a single point in time (i.e., the timestamp information for the creation of the edges is unavailable), we have an estimation procedure combining elements of method of moments with an approximation to the likelihood. A brief summary of the MLE and SN methods is in Appendix A and desirable properties of these estimators are in [25].

Note that a main difference between the MLE, SN and EV methods lies in the amount of data utilized. The MLE approach requires the entire growth history of the network while the SN method uses only a single snapshot of the network. The EV method, on the other hand, requires only a subset of a snapshot of the network; only those degree counts of nodes with large in- or out-degrees. When the underlying model is true, MLE is certainly the most efficient, but also hinges on having a complete data set. As we shall see, in the case where the model is misspecified, the EV method provides an attractive and reliable alternative.

4.1. Estimation for the linear PA model.

4.1.1. *Comparison of EV with MLE and SN.* Figure 4.1 presents biases for estimates of $(\alpha, \iota_{\text{in}}, \iota_{\text{out}})$ using EV, MLE, and SN methods on data simulated from the linear PA model.

We held $(\beta, \delta_{\text{in}}, \delta_{\text{out}}) = (0.4, 1, 1)$ constant and varied $\alpha = 0.1, 0.2, 0.3, 0.4$ so that the true values of $\gamma, \iota_{\text{in}}, \iota_{\text{out}}$ were also varying. For each set of parameter values $(\alpha, \iota_{\text{in}}, \iota_{\text{out}})$, 200 independent replications of a linear PA network with $n = 10^5$ edges were simulated and the true values of $(\iota_{\text{in}}, \iota_{\text{out}})$ were computed by (2.9). We estimated $(\iota_{\text{in}}, \iota_{\text{out}})$ by the minimum distance method $(\hat{\iota}_{\text{in}}^{EV}, \hat{\iota}_{\text{out}}^{EV})$, MLE and the one-snapshot methods applied to the parametric model (cf. Section A), denoted by $(\hat{\iota}_{\text{in}}^{MLE}, \hat{\iota}_{\text{out}}^{MLE})$ and $(\hat{\iota}_{\text{in}}^{SN}, \hat{\iota}_{\text{out}}^{SN})$, respectively. With $(\hat{\iota}_{\text{in}}^{EV}, \hat{\iota}_{\text{out}}^{EV})$, $\hat{\alpha}^{EV}$ is calculated by (3.4) using $n_{\text{tail}} = 200$.

As seen here, for simulated data from a known model, MLE outperforms other estimation procedures. The EV procedure tends to have much larger variance than both MLE and SN

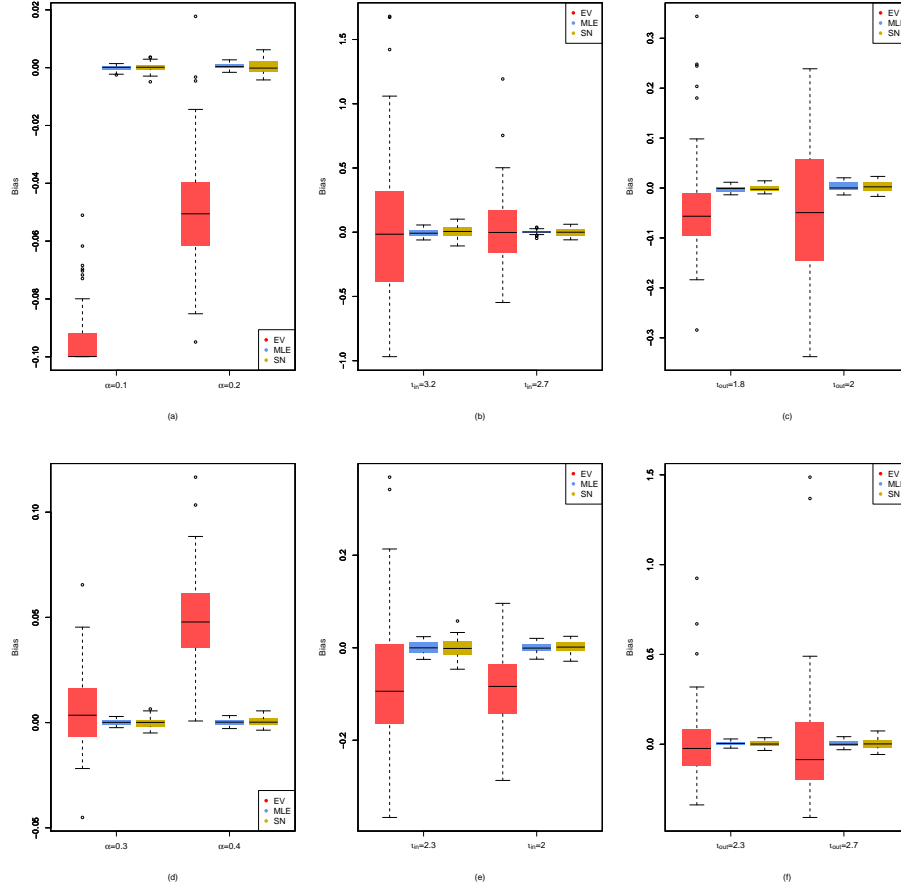


FIGURE 4.1. Boxplots of biases for estimates of $(\alpha, \ell_{in}, \ell_{out})$ using EV, MLE and SN methods. Panels (a)–(c) correspond to the case where $\alpha = 0.1, 0.2$ and (d)–(f) are for $\alpha = 0.3, 0.4$, holding $(\beta, \delta_{in}, \delta_{out}) = (0.4, 1, 1)$ constant.

with slightly more bias. This is not surprising as the performance of the EV estimators is dependent on the quality of the following approximations:

- (1) The number of edges in the network, n , should be sufficiently large to ensure a close approximation of $N_n(i, j)/N(n)$ to the limit joint pmf p_{ij} .
- (2) The choice of thresholds must guarantee the quality of the EV estimates for the indices and the limiting angular distribution. The thresholding means estimates are based on only a small fraction of the data and hence have large uncertainty.
- (3) The parameter a used to transform the in- and out-degrees to standard regular variation is estimated and thus subject to estimation error which propagates throughout the remaining estimation procedures.

4.1.2. Sensitivity analysis. We explore how sensitive EV estimates are to choice of r , the threshold for the approximation to the limiting angular density in (3.4). Equivalently, we consider varying n_{tail} , the number of tail observations included in the estimation.




FIGURE 4.2. (a) Boxplots of biases of $\hat{\alpha}$ and $\hat{\alpha}^*$ for different n_{tail} and n_{tail}^* over 50 replications, where $(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}}) = (0.3, 0.4, 0.3, 1, 1)$. (b) Linearly interpolated trajectories of biases of $\hat{\alpha}$ and $\hat{\alpha}^*$ from 10 randomly picked realizations.

For the sensitivity analysis, 50 linear PA networks with 10^5 edges and parameter set

$$(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}}) = (0.3, 0.4, 0.3, 1, 1),$$

or equivalently,

$$(\alpha, \beta, \gamma, \iota_{\text{in}}, \iota_{\text{out}}) = (0.3, 0.4, 0.3, 2.29, 2.29)$$

are generated. We use $n_{\text{tail}} = 50, 100, 200, 300, 500, 1000, 1500$ to calculate the EV estimates for α . The performances of $\hat{\alpha}^{EV}$ across different values of n_{tail} are demonstrated by the blue boxplots in Figure 4.2(a). We see that the biases of $\hat{\alpha}$ remain small until n_{tail} increases to 300, and for larger values of n_{tail} , $\hat{\alpha}$ considerably underestimates α .

We note that the angular components $R_n(v)$, $1 \leq v \leq N(n)$ are also power-lawed. As an attempt to select the optimal value of n_{tail} , we apply the minimum distance method to the $R_n(v)$'s and use the selected threshold, n_{tail}^* , as the truncation threshold. The boxplot of n_{tail}^* for the 50 simulated networks are represented by the horizontal boxplot in Figure 4.2(a). The EV estimator with respect to this threshold for each simulation, denoted by $\hat{\alpha}^*$, is shown by the red boxplot and plotted at $n_{\text{tail}} = 875$, the mean of n_{tail}^* . Overall, n_{tail}^* varies between 300 and 1500 and results in an underestimated $\hat{\alpha}^*$.

In Figure 4.2(b), we randomly choose 10 realizations (among the 50 replications) and plot the linearly interpolated trajectories of $\hat{\alpha}$, based on different values of n_{tail} . Black points are the estimation results using fixed thresholds $n_{\text{tail}} = 50, 100, 200, 300, 500, 1000, 1500$ and red ones are determined by $(\hat{\alpha}^*, n_{\text{tail}}^*)$ using the minimum distance method. Black and red points denoted by the same symbol belong to the same realization. Comparison among estimation results for different values of n_{tail} reveals that choosing a fixed threshold $n_{\text{tail}} \leq 300$

outperforms selecting a n_{tail}^* using the minimum distance method, as it produces estimates with smaller biases and variances.

4.2. Data corrupted by random edge addition/deletion. PA models are designed to describe human interaction in social networks but what if data collected from a network is corrupted or usual behavior is changed? Corruption could be due to collection error and atypical behavior could result from users hiding their network presence or trolls acting as provocateurs. In such circumstances, the task is to unmask data corruption or atypical behavior and recover the parameters associated with the original preferential attachment rules.

In the following, we consider network data that are generated from the linear PA model but corrupted by random addition or deletion of edges. For such corrupted data, we attempt to recover the original model and compare the performances of MLE, SN, and EV methods.

4.2.1. Randomly adding edges. We consider a network generating algorithm with linear PA rules but also a possibility of adding random edges. Let $G(n) = (V(n), E(n))$ denote the graph at time n . We assume that the edge set $E(n)$ can be decomposed into two disjoint subsets: $E(n) = E^{PA}(n) \cup E^{RA}(n)$, where $E^{PA}(n)$ is the set of edges resulting from PA rules, and $E^{RA}(n)$ is the set of those resulting from random attachments. This can be viewed as an interpolation of the PA network and the Erdős-Rényi random graph.

More specifically, consider the following network growth. Given $G(n-1)$, $G(n)$ is formed by creating a new edge where:

- (1) With probability p_a , two nodes are chosen randomly (allowing repetition) from $V(n-1)$ and an edge is created connecting them. The possibility of a self loop is allowed.
- (2) With probability $1-p_a$, a new edge is created according to the preferential attachment scheme $(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}})$ on $G^{PA}(n-1) := (V(n-1), E^{PA}(n-1))$.

The question of interest is, if we are unaware of the perturbation effect and pretend the data from this model are coming from the linear PA model, can we recover the PA parameters? To investigate, we generate networks of $n = 10^5$ edges with parameter values

$$(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}}) = (0.3, 0.4, 0.3, 1, 1), \quad p_a \in \{0.025, 0.05, 0.075, 0.1, 0.125, 0.15\}.$$

For each network, the original PA model is fitted using the MLE, SN and EV methods, respectively. The angular MLE in (3.4) in the extreme value estimation is performed based on $n_{\text{tail}} = 500$ tail observations. In order to compare these estimators, we repeat the experiment 200 times for each value of p_a and obtain 200 sets of estimated parameters for each method. Figure 4.3 summarizes the estimated values for $(\delta_{\text{in}}, \delta_{\text{out}}, \alpha, \gamma, \iota_{\text{in}}, \iota_{\text{out}})$ for different values of p_a . The mean estimates are marked by crosses and the 2.5% and 97.5% empirical quantiles are marked by the bars. The true value of parameters are shown as the horizontal lines.

While all parameters deviate from the true value as p_a increases and the network becomes more “noisy”, the EV estimates for $(\delta_{\text{in}}, \delta_{\text{out}})$ exhibit smaller bias than the MLE and SN methods (Figure 4.3 (a) and (b)). All three methods give underestimated probabilities (α, γ) (Figure 4.3 (c) and (d)). This is because the perturbation step (1) creates more edges between existing nodes and consequently inflates the estimated value of β .

Also note that the mean EV estimates of $(\iota_{\text{in}}, \iota_{\text{out}})$ stay close to the theoretical values for all choices of p_a ; see Figure 4.3 (e) and (f). The MLE and SN estimates of $(\iota_{\text{in}}, \iota_{\text{out}})$, which

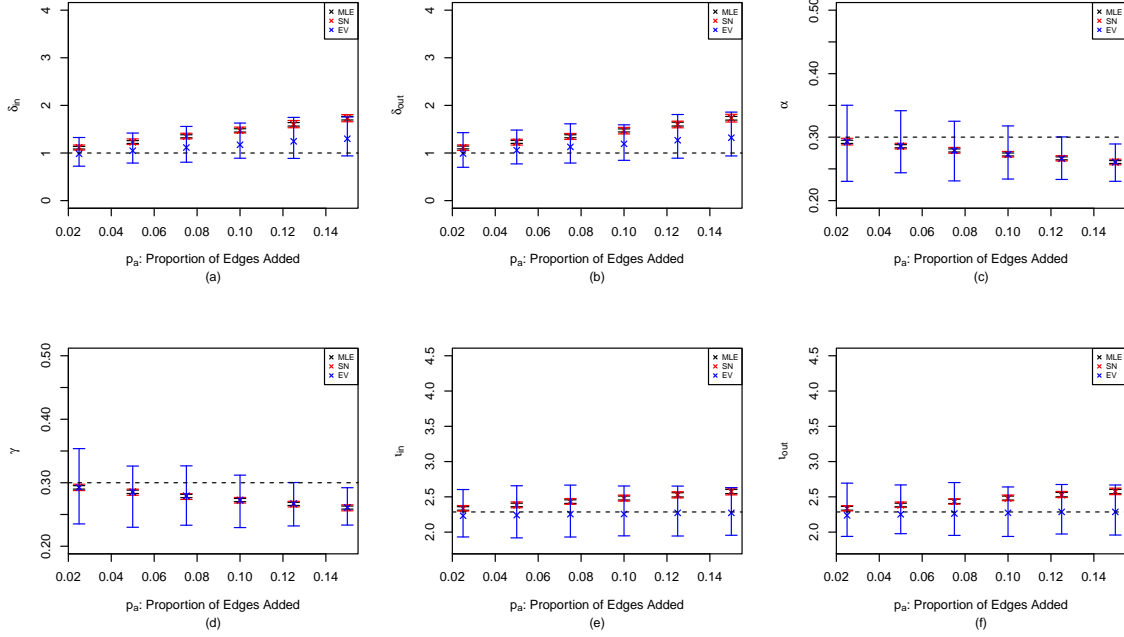


FIGURE 4.3. Mean estimates and 2.5% and 97.5% empirical quantiles of (a) δ_{in} ; (b) δ_{out} ; (c) α ; (d) γ ; (e) ι_{in} ; (f) ι_{out} , using MLE (black), SN (red) and EV (blue) methods over 200 replications, where $(\alpha, \beta, \gamma, \delta_{in}, \delta_{out}) = (0.3, 0.4, 0.3, 1, 1)$ and $p_a = 0.025, 0.05, 0.075, 0.1, 0.125, 0.15$. For the EV method, 500 tail observations were used to obtain $\hat{\alpha}^{EV}$.

are computed from the corresponding estimates for $(\alpha, \beta, \gamma, \delta_{in}, \delta_{out})$, show strong bias as p_a increases. In this case, the EV method is robust for estimating the PA parameters and recovering the tail indices from the original model.

4.2.2. Randomly deleting edges. We now consider the scenario where a network is generated from the linear PA model, but a random proportion p_d of edges are deleted at the final time. We do this by generating $G(n)$ and then deleting $[np_d]$ edges by sampling without replacement. For the simulation, we generated networks with parameter values

$$(\alpha, \beta, \gamma, \delta_{in}, \delta_{out}) = (0.3, 0.4, 0.3, 1, 1), \quad p_d \in \{0.025, 0.05, 0.075, 0.1, 0.125, 0.15\}.$$

Again, for each value of p_d , the experiment is repeated 200 times and the resulting parameter plots are shown in Figure 4.4 using the same format as for Figure 4.3. For the EV method, 100 tail observations were used to compute an $\hat{\alpha}^{EV}$.

Surprisingly, for all six parameters considered, MLE estimates stay almost unchanged for different values of p_d while SN and EV estimates underestimate $(\delta_{in}, \delta_{out})$ and overestimate (α, γ) , with increasing magnitudes of biases as p_d increases. For tail estimates, the minimum distance method still gives reasonable results (though with larger variances), whereas the SN method keeps underestimating ι_{in} and ι_{out} .

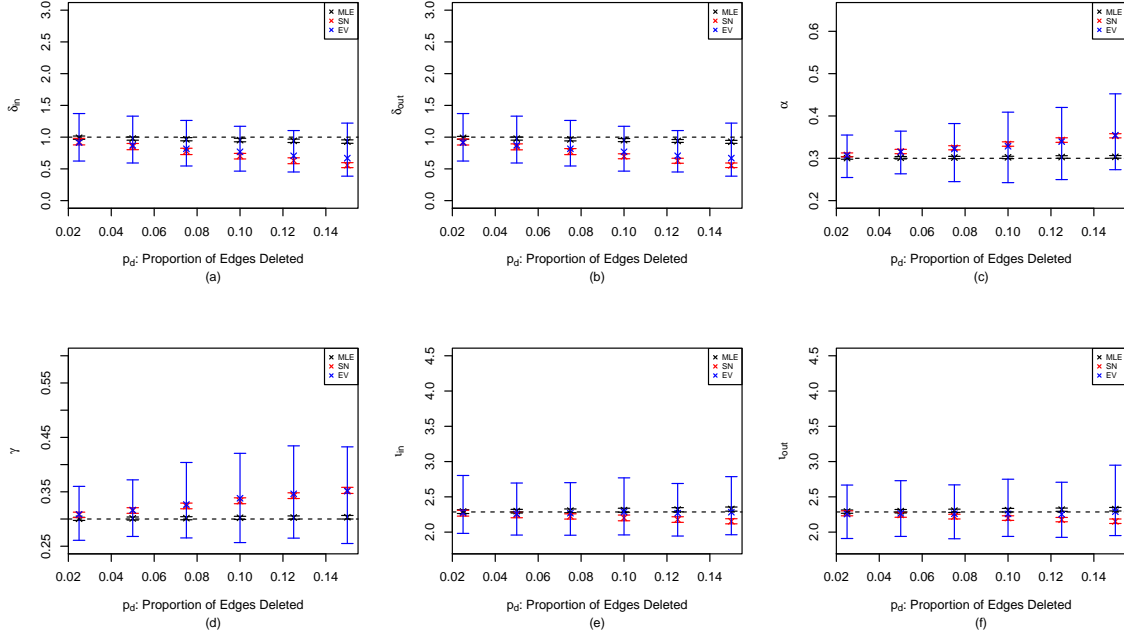


FIGURE 4.4. Mean estimates and 2.5% and 97.5% empirical quantiles of (a) δ_{in} ; (b) δ_{out} ; (c) α ; (d) γ ; (e) ι_{in} ; (f) ι_{out} , using MLE (black), SN (red) and EV (blue) methods over 50 replications, where $(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}}) = (0.3, 0.4, 0.3, 1, 1)$ and $p_d = 0.025, 0.05, 0.075, 0.1, 0.125, 0.15$. For the EV method, 100 tail observations were used to compute $\hat{\alpha}^{\text{EV}}$.

The performance of MLE in this case is surprisingly competitive. This is intriguing and in ongoing work, we will think about why this is the case.

4.3. Superstar model. In this section, we consider network data generated from the superstar model. We compare the accuracy of tail index estimates under parametric methods applied to the linear PA model with extreme value estimates applied directly to data.

Networks are simulated from the superstar model with the following parameter values:

$$(\alpha, \beta, \delta_{\text{in}}, \delta_{\text{out}}, n) = (0.3, 0.4, 0.3, 1, 1, 10^6), \quad p \in \{0.1, 0.15, 0.2, 0.25, 0.3\}.$$

The MLE estimates of the tail indices based on (2.9), $(\hat{\iota}_{\text{in}}^{\text{MLE}}, \hat{\iota}_{\text{out}}^{\text{MLE}})$, are compared to the EV estimates calculated directly from the node degree data, $(\hat{\iota}_{\text{in}}^{\text{EV}}, \hat{\iota}_{\text{out}}^{\text{EV}})$. According to Theorem 2.1, the theoretical marginal tail indices for $I_n(v)$ and $O_n(v)$, $1 \leq v \leq N(n)$, based on a superstar PA model are given by (2.11), (2.13). This experiment is repeated 50 times and Table 1 records the mean estimates for $(\iota_{\text{in}}, \iota_{\text{out}})$ over these 50 replications.

As p increases and the influence of the superstar node becomes more profound, the MLE method does not give an accurate estimate of tail indices, while the EV method stays more robust. However, when p becomes too large, the in-degrees of non-superstar nodes will be greatly restricted, which increases the finite sample bias in the EV estimates.

p	$(\iota_{\text{in}}, \iota_{\text{out}})$	$(\hat{\iota}_{\text{in}}^{MLE}, \hat{\iota}_{\text{out}}^{MLE})$	$(\hat{\iota}_{\text{in}}^{EV}, \hat{\iota}_{\text{out}}^{EV})$
0.1	(2.43, 2.29)	(2.11, 2.31)	(2.24, 2.25)
0.15	(2.51, 2.29)	(2.03, 2.33)	(2.28, 2.20)
0.2	(2.61, 2.29)	(1.97, 2.34)	(2.35, 2.18)
0.25	(2.71, 2.29)	(1.91, 2.36)	(2.43, 2.18)
0.3	(2.84, 2.29)	(1.86, 2.38)	(2.51, 2.15)

TABLE 1. Mean estimates for $(\iota_{\text{in}}, \iota_{\text{out}})$ using both MLE and minimum distance methods, with $(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}}, n) = (0.3, 0.4, 0.3, 1, 1, 10^6)$.

Note that the theoretical indices $(\iota_{\text{in}}, \iota_{\text{out}})$ in Table 1 are for the in- and out-degrees of the non-superstar nodes. In the EV methods, the inclusion of the superstar node can severely bias the estimation of ι_{in} . Let k_n be some intermediate sequence such that $k_n \rightarrow \infty$ and $k_n/n \rightarrow 0$ as $n \rightarrow \infty$ and use $I_{(1)} \geq \dots \geq I_{(k_n+1)}$ to denote the upper $k_n + 1$ order statistics of $\{I_n(v) : 0 \leq v \leq N(n)\}$. Then the corresponding Hill estimator is

$$\begin{aligned}
 1/\hat{\iota}_{\text{in}}^{EV}(k_n) &:= \frac{1}{k_n} \sum_{i=1}^{k_n} \log \frac{I_{(i)}}{I_{(k_n+1)}} \\
 (4.1) \qquad &= \frac{1}{k_n} \log I_{(1)} - \frac{1}{k_n} \log I_{(k_n+1)} + \frac{1}{k_n} \sum_{i=2}^{k_n} \log \frac{I_{(i)}}{I_{(k_n+1)}}.
 \end{aligned}$$

From the construction of the superstar model, we know that the superstar node likely has the largest in-degree, which is approximately equal to np for large n . Hence, the first term in (4.1) goes to 0, as long as

$$k_n / \log n \rightarrow \infty, \quad \text{as } n \rightarrow \infty,$$

and the third term in (4.1) is the Hill estimator computed from the in-degrees of non-superstar nodes. In [27], the consistency of the Hill estimator has been proved for a simple undirected linear PA model, but consistency for $\hat{\iota}_{\text{in}}^{EV}(k_n)$ is not proven for either of the two models we consider here. However, with the belief on the consistency of $\hat{\iota}_{\text{in}}^{EV}(k_n)$, (4.1) suggests that choosing a larger k_n will reduce the bias when estimating ι_{in} in the superstar model.

To illustrate this point numerically, we choose $k_n = 200, 500, 1000, 1500, 2000$ for a superstar network with 10^6 edges and probability of attaching to the superstar node $p = 0.1, 0.15, 0.2, 0.25, 0.3$. For each value of p , we again simulate 50 independent replications of the superstar PA model with parameters $(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}}, n) = (0.3, 0.4, 0.3, 1, 1, 10^6)$. Then for each replication generated, Hill estimates of the in- and out-degree tail indices are calculated under different choices of k_n . The mean values of the 50 pairs of estimates are recorded in Table 2, where the first entry is the in-degree tail estimate and the second is for out-degree.

From the in-degree estimates in Table 2, we observe that for most values of p increasing k_n to 500 improves the estimation results, but further increase in k_n has adverse effects. One reason is that large k_n means smaller in-degrees are taken into the calculation of the Hill estimator; these smaller in-degrees might not be large enough to be considered as following the power law in (2.10). This also explains the increasing biases for the out-degree estimates,

	Number of Upper Order Statistics k_n				
	200	500	1000	1500	2000
$p = 0.1$	(2.16, 2.22)	(2.26, 2.19)	(2.27, 2.16)	(2.28, 2.14)	(2.27, 2.15)
$p = 0.15$	(2.25, 2.18)	(2.32, 2.17)	(2.29, 2.14)	(2.31, 2.15)	(2.28, 2.14)
$p = 0.2$	(2.32, 2.17)	(2.39, 2.16)	(2.37, 2.15)	(2.39, 2.11)	(2.33, 2.13)
$p = 0.25$	(2.36, 2.18)	(2.47, 2.16)	(2.43, 2.12)	(2.49, 2.11)	(2.52, 2.12)
$p = 0.3$	(2.41, 2.17)	(2.58, 2.13)	(2.56, 2.11)	(2.47, 2.11)	(2.51, 2.12)

TABLE 2. Mean values of EV estimates of tail indices $(\iota_{\text{in}}, \iota_{\text{out}})$ over 50 replications, with $(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}}, n) = (0.3, 0.4, 0.3, 1, 1, 10^6)$. The true values are given in Table 1.

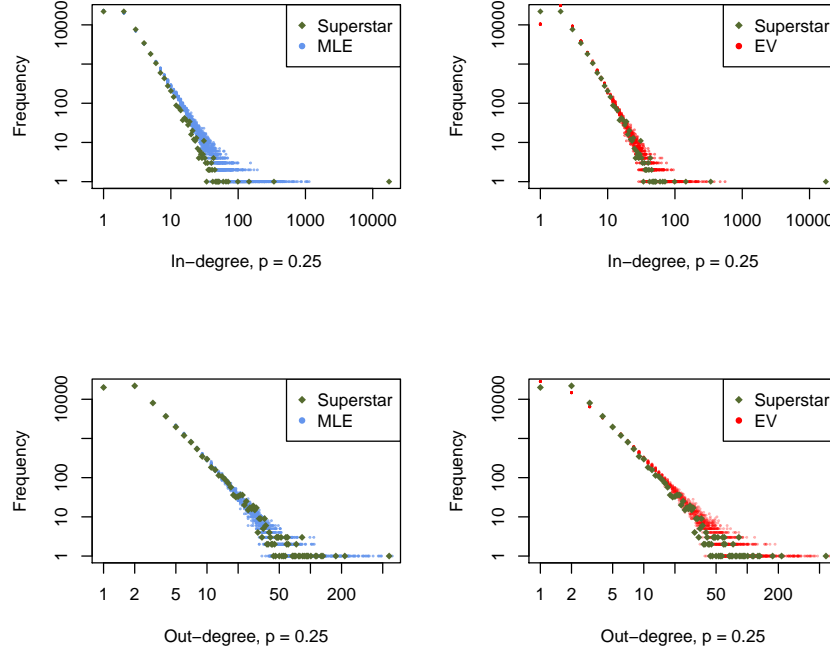


FIGURE 4.5. Empirical in- and out-degree distributions, with $(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}}, n, p) = (0.3, 0.4, 0.3, 1, 1, 10^5, 0.25)$.

where the superstar node does not have any impact. Comparing the results in Table 2 to those EV estimates in Table 1, we see that the minimum distance method seeks a good balance between eliminating the effect of the superstar nodes and choosing a reasonably large threshold.

The next question is how the model misspecification affects the empirical distributions of in- and out-degrees. To evaluate this, we generated a superstar PA model with parameters

$$(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}}, n, p) = (0.3, 0.4, 0.3, 1, 1, 10^5, 0.25).$$

We estimated parameters by both MLE and EV methods from simulated superstar data, pretending that the data was generated from an ordinary PA rule. For the EV approach, 200 tail observations were used while computing $\hat{\alpha}^{EV}$. Denote the MLE and EV estimates by

$$\begin{aligned}\hat{\boldsymbol{\theta}}_n^{MLE} &:= (\hat{\alpha}^{MLE}, \hat{\beta}^{MLE}, \hat{\gamma}^{MLE}, \hat{\delta}_{\text{in}}^{MLE}, \hat{\delta}_{\text{out}}^{MLE}), \\ \hat{\boldsymbol{\theta}}_n^{EV} &:= (\hat{\alpha}^{EV}, \hat{\beta}^{EV}, \hat{\gamma}^{EV}, \hat{\delta}_{\text{in}}^{EV}, \hat{\delta}_{\text{out}}^{EV}).\end{aligned}$$

We then simulated 20 independent replications of a linear PA model with parameters $\hat{\boldsymbol{\theta}}_n^{MLE}$ and 20 with parameters $\hat{\boldsymbol{\theta}}_n^{EV}$. For each set of replicates we computed the empirical frequency distributions. Comparisons of degree distributions are provided in Figure 4.5.

In all 4 panels, the green dots represent the empirical degree frequencies for the simulated superstar data, top for in-degree and bottom for out-degree. Blue in the two left panels represents overlaid frequency distributions for the 20 simulated data sets from the linear PA replicates using $\hat{\boldsymbol{\theta}}_n^{MLE}$. Red in the right two panels does the same thing for 20 replicates of the linear PA model using parameter $\hat{\boldsymbol{\theta}}_n^{EV}$.

The EV method seems to give better fit for in-degrees. Based on out-degrees, it is difficult to visually discern an advantage for either approach. While not obvious in the plots, we again expect the estimated degrees from the EV method to have higher variance than those from MLE, as much less data were used for the model fitting.

5. CONCLUSION

In this paper, we proposed a semi-parametric extreme value (EV) estimation method for network models. We compared the performance of this method to the two parametric approaches (MLE and snapshot methods) given in [25] under three scenarios: (1) data generated from a linear preferential attachment (linear PA) model; (2) data generated from a linear PA model with corruption; (3) data generated from a superstar linear PA model.

To summarize our findings and experience, EV estimation methods play important roles while applied to social network data. The method provides a robust procedure for estimating parameters of the network related to heavy-tailedness of the marginal and joint distributions of the in- and out-degrees. Also EV methods play a confirmatory role to other estimation procedures that are likelihood based, such as MLE or the snapshot (SN) method, which require that the model is correctly specified. If, for example, MLE or SN produces estimates of tail indices different from those given by the EV procedure, then this might suggest a lack of fit of the underlying model.

In practice, data are not as *clean* as those produced in simulations and one expects deviations from a base model such as the linear PA. As seen in this paper, these deviations can lead to sharply biased MLE and SN estimates especially when compared to EV estimates. As in classical EV estimation in the iid setting, the choice of threshold upon which to base the estimation remains a thorny issue in the network context. The minimum distance method based on [5] for estimating marginal tail indices works well for the examples considered here, but worse for multivariate data where it is employed to set thresholds based on radius vectors.

APPENDIX A. PARAMETER ESTIMATION FOR LINEAR PA MODEL

Parameter estimation for the linear PA model was studied in [25]. If the complete history of the network evolution is available (i.e., timestamps of edge creation are known), then MLE estimates exist and are computable. On the other hand, if only a snapshot of the network is given at a single point in time (i.e., timestamp information for the creation of the edges is unavailable), an approximate MLE was proposed. This procedure combined elements of method of moments with an approximation to the likelihood. In the following we provide a brief summary of these two estimation methods. Asymptotic properties of these estimators can be found in [25].

A.1. MLE. Given the full evolution of the network $G(n)$, assuming the graph began with n_0 initial edges, the MLE estimator of $\boldsymbol{\theta} = (\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}})$,

$$\hat{\boldsymbol{\theta}}_n^{MLE} := (\hat{\alpha}^{MLE}, \hat{\beta}^{MLE}, \hat{\gamma}^{MLE}, \hat{\delta}_{\text{in}}^{MLE}, \hat{\delta}_{\text{out}}^{MLE}),$$

is obtained by setting

$$\begin{aligned}\hat{\alpha}^{MLE} &= \frac{1}{n - n_0} \sum_{t=n_0+1}^n \mathbf{1}_{\{J_t=1\}}, \\ \hat{\beta}^{MLE} &= \frac{1}{n - n_0} \sum_{t=n_0+1}^n \mathbf{1}_{\{J_t=2\}}, \\ \hat{\gamma}^{MLE} &= 1 - \hat{\alpha}^{MLE} - \hat{\beta}^{MLE},\end{aligned}$$

and solving for $(\hat{\delta}_{\text{in}}^{MLE}, \hat{\delta}_{\text{out}}^{MLE})$ from

$$\begin{aligned}\sum_{i=0}^{\infty} \frac{N_{>i}^{\text{in}}(n) - N_{>i}^{\text{in}}(n_0)}{i + \hat{\delta}_{\text{in}}^{MLE}} &= \frac{n - n_0}{\hat{\delta}_{\text{in}}^{MLE}} \hat{\gamma}^{MLE} + \sum_{t=n_0+1}^n \left(\frac{N(t-1)}{t-1 + \hat{\delta}_{\text{in}}^{MLE} N(t-1)} \mathbf{1}_{\{J_t \in \{1,2\}\}} \right), \\ \sum_{j=0}^{\infty} \frac{N_{>j}^{\text{out}}(n) - N_{>j}^{\text{out}}(n_0)}{j + \hat{\delta}_{\text{out}}^{MLE}} &= \frac{n - n_0}{\hat{\delta}_{\text{out}}^{MLE}} \hat{\alpha}^{MLE} + \sum_{t=n_0+1}^n \frac{N(t-1)}{t-1 + \hat{\delta}_{\text{out}}^{MLE} N(t-1)} \mathbf{1}_{\{J_t \in \{2,3\}\}},\end{aligned}$$

where

$$N_{>i}^{\text{in}}(n) := \sum_{i'>i} N_{i'}^{\text{in}}(n), \quad N_{>j}^{\text{out}}(n) := \sum_{j'>j} N_{j'}^{\text{out}}(n).$$

By [25, Theorem 3.3], $\hat{\boldsymbol{\theta}}_n^{MLE}$ is strongly consistent, asymptotically normal and efficient.

A.2. Snapshot. The estimation method for $\boldsymbol{\theta}$ from the snapshot $G(n)$ is summarized in the following 7-step procedure:

1. Estimate β by $\hat{\beta}^{SN} = 1 - N(n)/n$.
2. Obtain $\hat{\delta}_{\text{in}}^0$ by solving

$$\sum_{i=1}^{\infty} \frac{N_{>i}^{\text{in}}(n)}{n} \frac{i}{i + \hat{\delta}_{\text{in}}^0} (1 + \hat{\delta}_{\text{in}}^0 (1 - \hat{\beta}^{SN})) = \frac{\frac{N_0^{\text{in}}(n)}{n} + \hat{\beta}^{SN}}{1 - \frac{N_0^{\text{in}}(n)}{n} \frac{\hat{\delta}_{\text{in}}^0}{1 + (1 - \hat{\beta}^{SN}) \hat{\delta}_{\text{in}}^0}},$$

where $N_0^{\text{in}}(n)$ denotes the number of nodes with in-degree 0 in $G(n)$.

3. Estimate α by

$$\hat{\alpha}^0 = \frac{\frac{N_0^{\text{in}}(n)}{n} + \hat{\beta}^{SN}}{1 - \frac{N_0^{\text{in}}(n)}{n} \frac{\hat{\delta}_{\text{in}}^0}{1 + (1 - \hat{\beta}^{SN})\hat{\delta}_{\text{in}}^0}} - \hat{\beta}^{SN}.$$

4. Obtain $\hat{\delta}_{\text{out}}^0$ by solving

$$\sum_{j=1}^{\infty} \frac{N_{>j}^{\text{out}}(n)}{n} \frac{j}{j + \hat{\delta}_{\text{out}}^0} (1 + \hat{\delta}_{\text{out}}^0 (1 - \hat{\beta}^{SN})) = \frac{\frac{N_0^{\text{out}}(n)}{n} + \hat{\beta}^{SN}}{1 - \frac{N_0^{\text{out}}(n)}{n} \frac{\hat{\delta}_{\text{out}}^0}{1 + (1 - \hat{\beta}^{SN})\hat{\delta}_{\text{out}}^0}},$$

where $N_0^{\text{out}}(n)$ denotes the number of nodes with out-degree 0 in $G(n)$.

5. Estimate γ by

$$\hat{\gamma}^0 = \frac{\frac{N_0^{\text{out}}(n)}{n} + \hat{\beta}^{SN}}{1 - \frac{N_0^{\text{out}}(n)}{n} \frac{\hat{\delta}_{\text{out}}^0}{1 + (1 - \hat{\beta}^{SN})\hat{\delta}_{\text{out}}^0}} - \hat{\beta}^{SN}.$$

6. Re-normalize the probabilities

$$(\hat{\alpha}^{SN}, \hat{\beta}^{SN}, \hat{\gamma}^{SN}) \leftarrow \left(\frac{\hat{\alpha}^0(1 - \hat{\beta}^{SN})}{\hat{\alpha}^0 + \hat{\gamma}^0}, \hat{\beta}^{SN}, \frac{\hat{\gamma}^0(1 - \hat{\beta}^{SN})}{\hat{\alpha}^0 + \hat{\gamma}^0} \right) \left($$

7. Solve for $\hat{\delta}_{\text{in}}^{SN}$ from

$$\sum_{i=0}^{\infty} \frac{N_{>i}^{\text{in}}(n)/n}{i + \hat{\delta}_{\text{in}}^{SN}} - \frac{1 - \hat{\alpha}^{SN} - \hat{\beta}^{SN}}{\hat{\delta}_{\text{in}}^{SN}} - \frac{(\hat{\alpha}^{SN} + \hat{\beta}^{SN})(1 - \hat{\beta}^{SN})}{1 + (1 - \hat{\beta}^{SN})\hat{\delta}_{\text{in}}^{SN}} = 0.$$

Similarly, solve for $\hat{\delta}_{\text{out}}^{SN}$ from

$$\sum_{j=0}^{\infty} \frac{N_{>j}^{\text{out}}(n)/n}{j + \hat{\delta}_{\text{out}}^{SN}} - \frac{1 - \hat{\gamma}^{SN} - \hat{\beta}^{SN}}{\hat{\delta}_{\text{out}}^{SN}} - \frac{(\hat{\gamma}^{SN} + \hat{\beta}^{SN})(1 - \hat{\beta}^{SN})}{1 + (1 - \hat{\beta}^{SN})\hat{\delta}_{\text{out}}^{SN}} = 0.$$

Note that Step 6 ensures that

$$\hat{\alpha}^{SN} + \hat{\beta}^{SN} + \hat{\gamma}^{SN} = 1.$$

It is shown in [25, Theorem 4.1] that $\hat{\boldsymbol{\theta}}_n^{SN} := (\hat{\alpha}^{SN}, \hat{\beta}^{SN}, \hat{\gamma}^{SN}, \hat{\delta}_{\text{in}}^{SN}, \hat{\delta}_{\text{out}}^{SN}) \xrightarrow{\text{a.s.}} \boldsymbol{\theta}$. Its asymptotic normality and efficiency are analyzed through simulation studies in the same paper.

APPENDIX B. PROOF OF THEOREM 2.1

Proof. We first prove the out-degree part of Theorem 2.1. Note that

$$(B.1) \quad \mathbf{E} (N_j^{\text{out}}(n+1) | G(n)) = N_j^{\text{out}}(n) + \gamma \mathbf{1}_{\{j=0\}} + \alpha \mathbf{1}_{\{j=1\}} + (\beta + \gamma) \left(N_{j-1}^{\text{out}}(n) \frac{j-1 + \delta_{\text{out}}}{n + \delta_{\text{out}} |V^0(n)|} - N_j^{\text{out}}(n) \frac{j + \delta_{\text{out}}}{n + \delta_{\text{out}} |V^0(n)|} \right) \left($$

Meanwhile, by the definition of $V^0(n)$, we have

$$(B.2) \quad |V^0(n)| + 1 = N(n) \sim \text{Binomial}(n, 1 - \beta).$$

Applying the arguments in the proof of Theorem 3.1 of [3], it follows that the out-degree distribution of a linear superstar model coincides with that of a standard linear preferential attachment network with parameters $(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}})$. Moreover,

$$\frac{N_j^{\text{out}}(n)}{n} \xrightarrow{\text{a.s.}} q_j^{\text{out}}, \quad j > 0, \quad n \rightarrow \infty,$$

where $\{q_j^{\text{out}}\} := \{p_j^{\text{out}}\}$ is the limiting out-degree distribution of $\text{PA}(\alpha, \beta, \gamma, \delta_{\text{in}}, \delta_{\text{out}})$. In particular,

$$q_j^{\text{out}} \sim C'_{\text{out}} j^{-(1+\iota_{\text{out}})} \quad \text{as } j \rightarrow \infty,$$

for C'_{out} positive and

$$\iota_{\text{out}}^{-1} = \frac{\beta + \gamma}{1 + \delta_{\text{out}}(\alpha + \gamma)}.$$

Next we consider the in-degree counts of non-superstar nodes. Observe also from the construction of the superstar model that

$$(B.3) \quad |E^0(n)| \sim \text{Binomial}(n, 1 - (\alpha + \beta)p).$$

Applying the Chernoff bound to both (B.2) and (B.3) gives

$$\begin{aligned} V^0(n) &= (1 - \beta)n + O(n^{1/2} \log n), \\ E^0(n) &= (1 - (\alpha + \beta)p)n + O(n^{1/2} \log n). \end{aligned}$$

Taking expectation on both sides of (B.1) then gives

$$(B.4) \quad \begin{aligned} &\mathbf{E} \left((\alpha + \beta)(1 - p) N_i^{\text{in}}(n) \frac{i + \delta_{\text{in}}}{|E^0(n)| + \delta_{\text{in}} |V^0(n)|} \right) \left(\right. \\ &= (\alpha + \beta)(1 - p) \frac{i + \delta_{\text{in}}}{n(1 - (\alpha + \beta)p) + \delta_{\text{in}} \cdot n(1 - \beta)} \mathbf{E}(N_i^{\text{in}}(n)) + O(n^{-1/2} \log n). \end{aligned}$$

By the rule of the superstar model, given $G(n)$, $N_i^{\text{in}}(n)$ will increase by 1 if either scenario (1b) or (2b) happens and a node with $I_n^{(n)}(v) = i - 1$ is chosen as the ending point of the edge. Also, it will decrease by 1 if either scenario (1b) or (2b) happens, but a node with $I_n^{(n)}(v) = i$ is chosen as the ending point of the edge. Moreover, with probability α a new node with in-degree 0 will be added to the graph, and with probability γ a new node with in-degree 1 is created in the next step. Hence, $\{N_i^{\text{in}}(n)\}_{n \geq 1}$ satisfies the following:

$$\begin{aligned} \mathbf{E} (N_i^{\text{in}}(n+1) | G(n)) &\left(\begin{aligned} &N_i^{\text{in}}(n) + \alpha \mathbf{1}_{\{i=0\}} + \gamma \mathbf{1}_{\{i=1\}} \\ &+ (\alpha + \beta)(1 - p) N_{i-1}^{\text{in}}(n) \frac{i - 1 + \delta_{\text{in}}}{|E^0(n)| + \delta_{\text{in}} |V^0(n)|} \\ &- (\alpha + \beta)(1 - p) N_i^{\text{in}}(n) \frac{i + \delta_{\text{in}}}{|E^0(n)| + \delta_{\text{in}} |V^0(n)|} \end{aligned} \right. \end{aligned}$$

Now let $q_{-1}^{\text{in}} = 0$, and define $\{q_i^{\text{in}}\}_{i \geq 0}$ by

$$(B.5) \quad q_i^{\text{in}} = \frac{(\alpha + \beta)(1 - p)}{1 - (\alpha + \beta)p + \delta_{\text{in}}(\alpha + \gamma)} \left((i - 1 + \delta_{\text{in}}) q_{i-1}^{\text{in}} - (i + \delta_{\text{in}}) q_i^{\text{in}} \right) \left(\begin{aligned} &\alpha \mathbf{1}_{\{i=0\}} + \gamma \mathbf{1}_{\{i=1\}} \end{aligned} \right).$$

According to the approximation in (B.4), we use the same proof technique as in [3, Theorem 3.1] to obtain

$$\frac{N_i^{\text{in}}(n)}{n} \xrightarrow{\text{a.s.}} q_i^{\text{in}}, \quad \text{as } n \rightarrow \infty.$$

Also, solving the recursion in (B.5) yields

$$(B.6) \quad q_0^{\text{in}} = \frac{\alpha}{1 + \iota_{\text{in}}^{-1} \delta_{\text{in}}},$$

$$(B.7) \quad \begin{aligned} q_1^{\text{in}} &= (1 + \delta_{\text{in}} + \iota_{\text{in}})^{-1} \left(\frac{\alpha \delta_{\text{in}}}{1 + \iota_{\text{in}}^{-1} \delta_{\text{in}}} + \frac{\gamma}{\iota_{\text{in}}^{-1}} \right) \left(\right. \\ q_i^{\text{in}} &= \frac{\Gamma(i + \delta_{\text{in}})}{\Gamma(i + \delta_{\text{in}} + \iota_{\text{in}} + 1)} \frac{\Gamma(2 + \delta_{\text{in}} + \iota_{\text{in}})}{\Gamma(1 + \delta_{\text{in}})} q_1^{\text{in}}, \quad i \geq 2, \end{aligned}$$

where

$$\iota_{\text{in}}^{-1} := \frac{(\alpha + \beta)(1 - p)}{1 - (\alpha + \beta)p + \delta_{\text{in}}(\alpha + \gamma)}.$$

Therefore, applying Stirling's approximation to (B.6)–(B.7) gives

$$q_i^{\text{in}} \sim C'_{\text{in}} i^{-(1+\iota_{\text{in}})}, \quad \text{as } n \rightarrow \infty,$$

for some positive constant C'_{in} . This completes the proof. \square

REFERENCES

- [1] S. Bhamidi. Universal techniques to analyze preferential attachment trees: Global and local analysis. *available: <http://www.unc.edu/~bhamidi/preferent.pdf>*, 2007. Preprint.
- [2] S. Bhamidi, J. M. Steele, and T. Zaman. Twitter event networks and the superstar model. *Ann. Appl. Probab.*, 25(5):2462–2502, 10 2015.
- [3] B. Bollobás, C. Borgs, J. Chayes, and O. Riordan. Directed scale-free graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, 2003)*, pages 132–139, New York, 2003. ACM.
- [4] R.E. Chandler and S. Bate. Inference for clustered data using the independence log-likelihood. *Biometrika*, 95:167–183, 2007.
- [5] A. Clauset, C.R. Shalizi, and M.E.J. Newman. Power-law distributions in empirical data. *SIAM Rev.*, 51(4):661–703, 2009.
- [6] S.G. Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer Series in Statistics. London: Springer. xiv, 210 p. , 2001.
- [7] B. Das, A. Mitra, and S. Resnick. Living on the multi-dimensional edge: Seeking hidden risks using regular variation. *Advances in Applied Probability*, 45(1):139–163, 2013.
- [8] L. de Haan and A. Ferreira. *Extreme Value Theory: An Introduction*. Springer-Verlag, New York, 2006.
- [9] R.T. Durrett. *Random Graph Dynamics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 2010.
- [10] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets*. Cambridge University Press, Cambridge, 2010.

- [11] F. Gao and A. van der Vaart. On the asymptotic normality of estimating the affine preferential attachment network models with random initial degrees. *Stochastic Process. Appl.*, 127(11):3754–3775, 2017.
- [12] C.S. Gillespie. Fitting heavy tailed distributions: The powerLaw package. *Journal of Statistical Software*, 64(2):1–16, 2015.
- [13] B.M. Hill. A simple general approach to inference about the tail of a distribution. *Ann. Statist.*, 3:1163–1174, 1975.
- [14] H. Hult and F. Lindskog. Regular variation for measures on metric spaces. *Publ. Inst. Math. (Beograd) (N.S.)*, 80(94):121–140, 2006.
- [15] E.D. Kolaczyk and G. Csárdi. *Statistical Analysis of Network Data with R. Use R!* Springer, New York, 2014.
- [16] P. Krapivsky, G. Rodgers, and S. Redner. Degree distributions of growing networks. *Phys. Rev. Lett*, 86, 2001.
- [17] P.L. Krapivsky and S. Redner. Organization of growing random networks. *Physical Review E*, 63(6):066123:1–14, 2001.
- [18] J. Kunegis. Konect: the Koblenz network collection. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 1343–1350. ACM, 2013.
- [19] F. Lindskog, S.I. Resnick, and J. Roy. Regularly varying measures on metric spaces: Hidden regular variation and hidden jumps. *Probab. Surv.*, 11:270–314, 2014.
- [20] S.I. Resnick. *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*. Springer Series in Operations Research and Financial Engineering. Springer-Verlag, New York, 2007. ISBN: 0-387-24272-4.
- [21] S.I. Resnick and G. Samorodnitsky. Tauberian theory for multivariate regularly varying distributions with application to preferential attachment networks. *Extremes*, 18(3):349–367, 2015.
- [22] G. Samorodnitsky, S. Resnick, D. Towsley, R. Davis, A. Willis, and P. Wan. Nonstandard regular variation of in-degree and out-degree in the preferential attachment model. *Journal of Applied Probability*, 53(1):146–161, March 2016.
- [23] R. van der Hofstad. *Random Graphs and Complex Networks. Vol. 1*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 2017.
- [24] C. Varin, N. Reid, and D. Firth. An overview of composite likelihood methods. *Statist. Sinica*, 21:5–42, 2011.
- [25] P. Wan, T. Wang, R. A. Davis, and S. I. Resnick. Fitting the linear preferential attachment model. *Electron. J. Statist.*, 11(2):3738–3780, 2017.
- [26] T. Wang and S.I. Resnick. Multivariate regular variation of discrete mass functions with applications to preferential attachment networks. *Methodology and Computing in Applied Probability*, pages 1–14, 2016.
- [27] T. Wang and S.I. Resnick. Consistency of Hill estimators in a linear preferential attachment model. *ArXiv e-prints*, 2017. Submitted.

PHYLLIS WAN, DEPARTMENT OF STATISTICS, 1255 AMSTERDAM AVENUE, MC 4690, COLUMBIA UNIVERSITY, NEW YORK, NY 10027

E-mail address: phyllis@stat.columbia.edu

TIANDONG WANG, SCHOOL OF OPERATIONS RESEARCH AND INFORMATION ENGINEERING, CORNELL UNIVERSITY, ITHACA, NY 14853

E-mail address: tw398@cornell.edu

PROF. RICHARD A. DAVIS, DEPARTMENT OF STATISTICS, 1255 AMSTERDAM AVENUE, MC 4690, ROOM 1004 SSW, COLUMBIA UNIVERSITY, NEW YORK, NY 10027

E-mail address: rdavis@stat.columbia.edu

PROF. SIDNEY I. RESNICK, SCHOOL OF OPERATIONS RESEARCH AND INFORMATION ENGINEERING, CORNELL UNIVERSITY, ITHACA, NY 14853

E-mail address: sir1@cornell.edu