Supporting Document for “A Bayesian methodology for systemic risk assessment in financial networks”

Axel Gandy *
Imperial College London

Luitgard A. M. Veraart †
London School of Economics and Political Science

May 3, 2016

1 Simulations checking the algorithm

In this section we present simulations that confirm that the MCMC sampler is sampling from the correct conditional distribution. For this, we compare the unconditional distribution of (L, θ) (which we can easily simulate in our models) to the following sampling method that involves the MCMC method:

1. Generate a sample $(\tilde{L}, \tilde{\theta})$ from the unconditional model (i.e., either model (5)) or model (8)).
2. Compute the observations $l = r(\tilde{L})$ and $a = c(\tilde{L})$.
3. Generate one sample (L, θ) conditional on $r(L) = l$, $c(L) = a$ using the MCMC sampler.

We will call samples (L, θ) from the above method “MCMC samples”. If the MCMC sampler is sampling from the correct conditional distribution then the distribution of the MCMC samples must be the same as the unconditional distribution.

1.1 Basic model

Consider the basic model (5) with $p_{ij} = 0.3\mathbb{I}(i \neq j)$ and $\lambda_{ij} = 1/10$ for a network with $n = 11$ banks. We consider the minimal observation setting, i.e. only row and column sums are observed. We generate 1000 samples from the unconditional distribution as well as 1000 MCMC samples (each of which involves running the MCMC chain). To ensure that the sample in step 3 is close to the target distribution we perform a large number of steps of the MCMC chain (1210000 individual updates - 10000 times the number of elements in L; here 11^2) before taking the sample.

Figure 1(a) shows the distribution of the mean out-degree of the network over the 1000 simulations, i.e. the empirical distribution of $\frac{1}{n}\sum_{i=1}^{n} r_i(\mathcal{A})$, where $\mathcal{A}$ is the adjacency matrix corresponding to the sample liabilities matrix L.

Figure 1(b) shows the marginal distribution of the out-degrees, i.e. the distribution of $r_i(\mathcal{A})$. Finally, Figure 1(c) shows the marginal distribution of the individual matrix entries, i.e. the distribution of L_{ij}, where i, j are independently uniformly chosen in $\mathcal{N}$.

There is good agreement between the distributions, which supports that the MCMC sampler is indeed sampling from the correct distribution.

*Imperial College London, Department of Mathematics, South Kensington Campus, London, SW7 2AZ, UK, a.gandy@imperial.ac.uk

†London School of Economics and Political Science, Department of Mathematics, Houghton Street, London WC2A 2AE, UK, l.veraart@lse.ac.uk

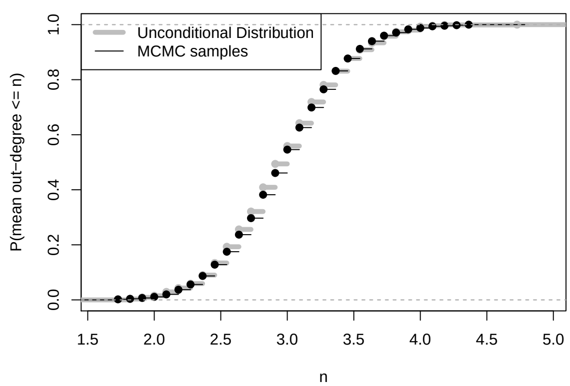

(a) CDF of mean out-degree

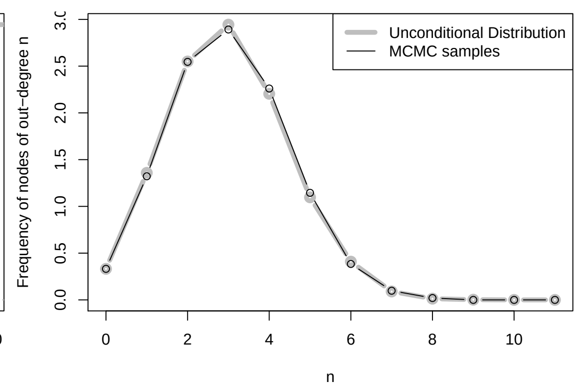

(b) Average frequency of out-degrees

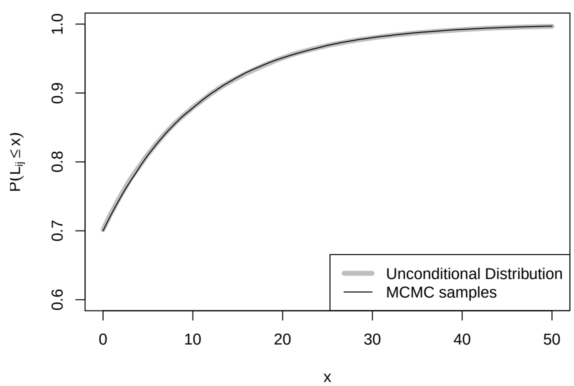

(c) CDF of marginal distribution of the entries of the liabilities matrix

Figure 1: Verification of the MCMC sampler in the basic model: Unconditional distribution and MCMC samples.

1.2 Fitness model

Consider the fitness model (8) with parameter choices $\alpha = -2.5$, $\beta = 0.2$, $\gamma = 1$ and priors $\zeta \sim U(0.5, 2)$, $\eta \sim \text{Exp}(1000)$ for a network with $n = 11$ banks. Again, we consider the minimal observation setting, i.e. only row and column sums are observed.

We generate 1000 samples from the unconditional distribution as well as 1000 MCMC samples (each of which involves running the MCMC chain). To ensure that the sample in step 3 is close to the target distribution we again perform a large number of steps of the MCMC chain (121000 individual updates - 1000 times the number of elements in L; here 11^2) before taking the sample.

As for the basic model, Figures 2(a), 2(b), and 2(c) show the distribution of the mean out-degree, the marginal distribution of the out-degrees and the marginal distribution of the entries of L, respectively. In addition to that, Figure 2(d) shows the marginal distribution of the fitness X_i, Figure 2(e) shows the distribution of the shape parameter ζ and Figure 2(f) shows the distribution of the scale parameter η.

Again there is good agreement between the samples from the unconditional distribution and the MCMC samples, supporting the validity of the implemented algorithm.

2 Convergence diagnostic

Consider the basic setup of Section 5.3.3, i.e. the fitness model (8) with parameter choices $\alpha = -2.5$, $\beta = 0.2$, $\gamma = 1$ and priors $\zeta \sim U(0.5, 2)$, $\eta \sim \text{Exp}(1000)$. In this section we describe some basic diagnostic for this situation. All of these are standard diagnostics for MCMC models.

Figure 3 shows traceplots of the first 200 recorded samples (after thinning) of some of the variables involved in the sampling. We deliberately pick variables involving banks with different levels of interbank assets/liabilities. The trace plots seem to indicate a reasonable mixing of the chain.

Figure 4 shows, for the same components, autocorrelation plots based on the full 10000 samples. Most autocorrelation is relatively low, for η it is moderate. These plots support that the chain, with the chosen amount of thinning, is mixing well.

Other convergence diagnostics from the CODA package (Plummer et al., 2006) (e.g. Geweke's convergence diagnostic) did not point towards problems with the convergence of the thinned chain.

3 MCMC updates in the hierarchical model

We iterate between updating $L|\theta$ and $\theta|L$. For the updates of $L|\theta$ we use the cycle updates of the basic model. As every cycle update of L only updates a small part of L, we perform a large number of updates for L for every update of θ. We have chosen to perform n^2 randomly chosen cycle updates of L (where n is the number of rows/columns of L) for every update of θ.

In the fitness model, $\theta = (x_1, \dots, x_n, \eta, \zeta)$ consists of two parts, the n-dimensional vector x containing the fitness and the shape and scale parameter for the weights (η, ζ). We update the fitness x and parameters (η, ζ) sequentially.

To update the fitness x we use a multiplicative Metropolis Hastings step, where the new proposed fitness $x^* = (x_1^*, \dots, x_n)$ is given by

$$x_i^* = x_i \exp(\epsilon_i),$$

where $\epsilon_i \sim N(0, \sigma^2)$ independently, and the variance σ^2 is tuned to give good acceptance rates. The new fitness is then accepted or rejected with the usual acceptance probabilities of the Metropolis Hastings algorithm.

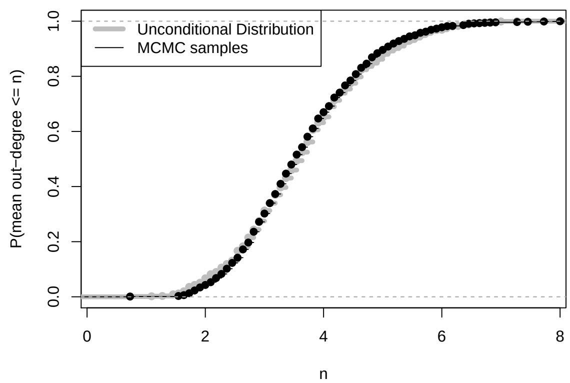

(a) CDF of mean out-degree

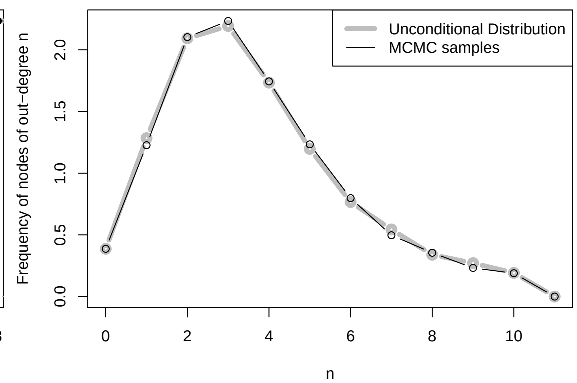

(b) Average frequency of out-degrees

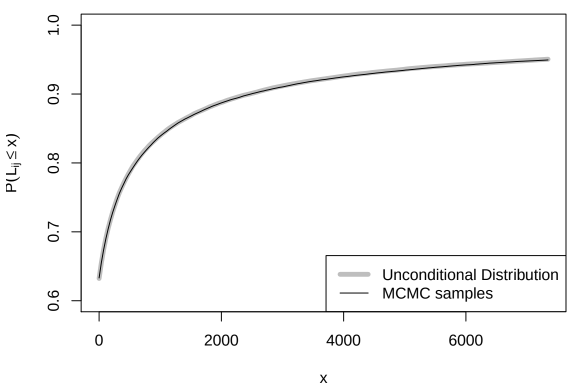

(c) CDF of marginal distribution of the entries of the liabilities matrix

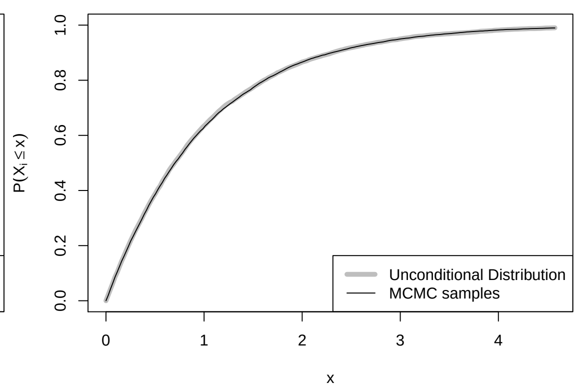

(d) CDF of marginal distribution of the fitness

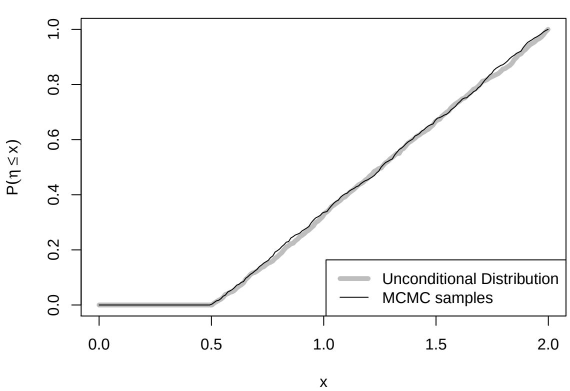

(e) CDF of shape parameter for λ

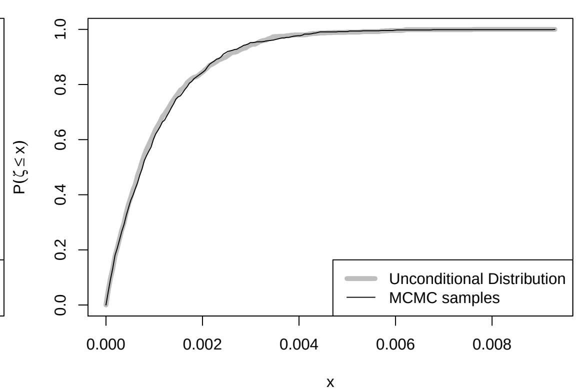

(f) CDF of scale parameter for λ

Figure 2: Verification of the MCMC sampler in a fitness model: Unconditional distribution and MCMC samples.

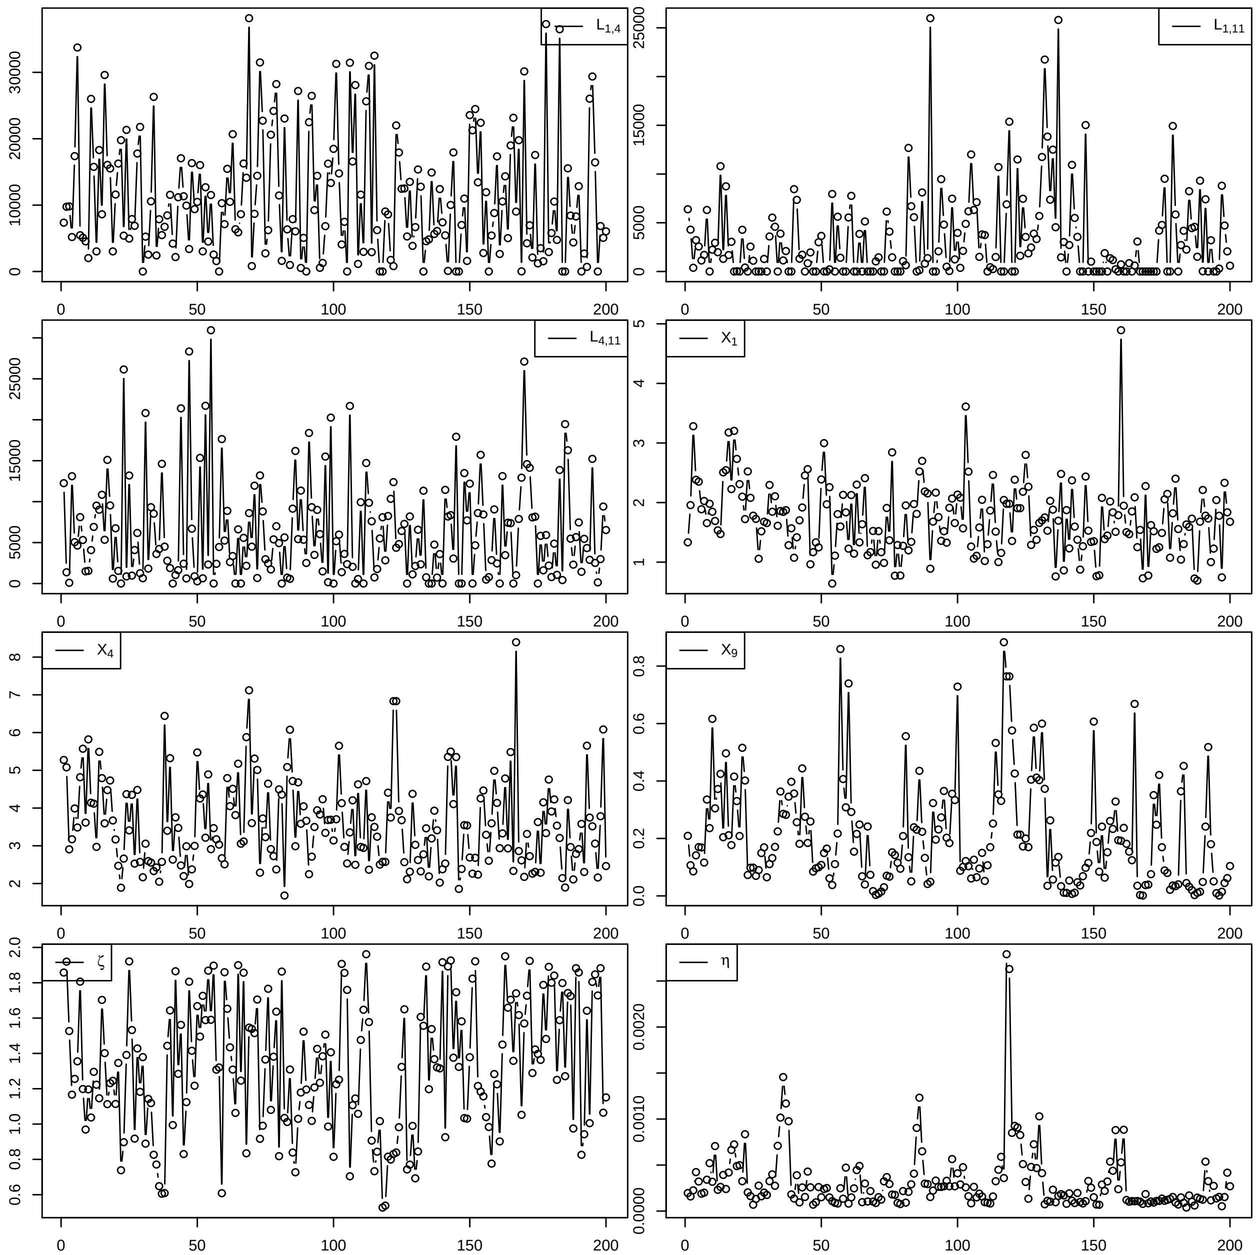

Figure 3: Trace plot for selected components of the model in the setup for the fitness plots with $\alpha = -2.5$, $\beta = 0.2$, $\gamma = 1$.

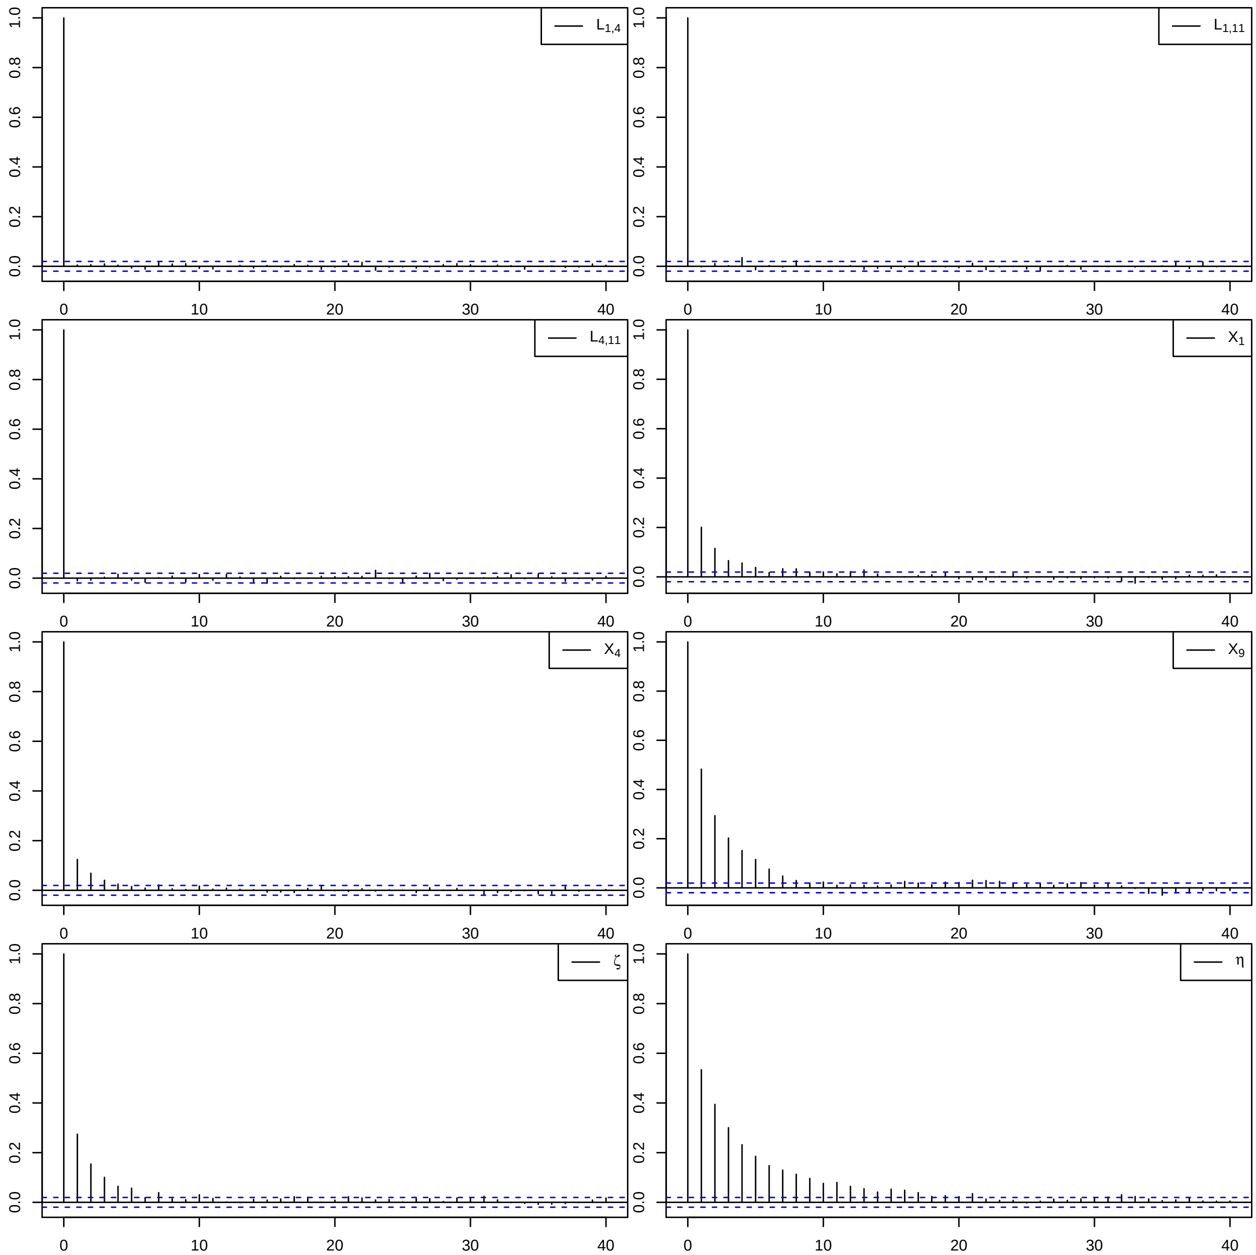

Figure 4: Estimated autocorrelation functions for selected components of the model in the setup for the fitness plots with $\alpha = -2.5$, $\beta = 0.2$, $\gamma = 1$.

To update the scale and shape parameter (η, ζ) we again use a Metropolis Hasting step. We propose a new (η^*, ζ^*) via

$$\begin{aligned}\eta^* =& \eta \exp(\epsilon_\eta), \quad \epsilon_\eta \sim N(0, \sigma_\eta^2),\\ \zeta^* =& \zeta + \epsilon_\zeta, \quad \epsilon_\zeta \sim N(0, \sigma_\zeta^2).\end{aligned}$$

Again, σ_η^2 and σ_ζ^2 are tuned to give desirable acceptance probabilities. We accept or reject the new (η^*, ζ^*) with the acceptance probabilities derived through standard calculations for the MCMC step.

4 Generating power laws using a fitness model

In the following we show how a power law degree distribution can be obtained with the fitness model introduced in Subsection 3.3.2. The parts of the model concerning the adjacency matrix were developed by Servedio et al. (2004) (for undirected, non-weighted networks) and they derived specific formulae for the case $\alpha = -2$. We use their methodology and adapt it to the directed/weighted network case in the following.

4.1 Power law for the out- and in-degree distributions

Assuming that a bank cannot have liabilities to itself but only to the $n-1$ remaining banks in the system, the expected out-degree $d^{\text{out}}(x)$ of a node with fitness x is given by

$$d^{\text{out}}(x) = (n-1) \int_0^\infty f(x,z)\rho(z)dz. \tag{1}$$

We focus on the expected out-degree, since the calculations for the in-degree and the degree are analogous. Since the fitness X has pdf ρ it follows from the transformation theorem of probability densities that the random variable $Y := d^{\text{out}}(X)$ has pdf

$$f_Y(y) = \left|\frac{\partial (d^{\text{out}})^{-1}(y)}{\partial y}\right| \rho((d^{\text{out}}(y))^{-1}).$$

Using the result on differentiation of the inverse function we can simplify this expression to

$$f_Y(d^{\text{out}}(x)) = \left|\frac{1}{d^{\text{out}\prime}(x)}\right| \rho(x) = \frac{\rho(x)}{d^{\text{out}\prime}(x)}.$$

If we now want this pdf to be a power law with parameter α on the finite range $[d_0, d_\infty]$, where $d_z = \lim_{x\to z} d^{\text{out}}(x)$, we require that

$$\begin{aligned} f_Y(d^{\text{out}}(x)) = \frac{\rho(x)}{d^{\text{out}\prime}(x)} &\stackrel{!}{=} c d^{\text{out}}(x)^\alpha \\ \Leftrightarrow \rho(x) &= c d^{\text{out}\prime}(x) d^{\text{out}}(x)^\alpha, \end{aligned} \tag{2}$$

and c is the constant such that $\int_{d_0}^{d_\infty} c y^\alpha dy = 1$. In particular,

$$c = \begin{cases} \left(\log\left(\frac{d_\infty}{d_0}\right)\right)^{-1}, & \text{if } \alpha = -1, \\ \frac{\alpha+1}{d_\infty^{\alpha+1} - d_0^{\alpha+1}}, & \text{if } \alpha \neq -1. \end{cases}$$

From (1) we see that $d_0 = \beta(n-1)$ and $d_\infty = \gamma(n-1)$ for some $0 < \beta < \gamma \leq 1$.

Integrating both sides in (2) from 0 to x gives

$$R(x) := \int_0^x \rho(y)dy = c\int_0^x c{d^{\text{out}}}'(y)d^{\text{out}}(y)^{\alpha}dy = \begin{cases} c(\log(d^{\text{out}}(x)) - \log(d^{\text{out}}(0))), & \text{if } \alpha = -1, \\ \frac{1}{\alpha+1}(d^{\text{out}}(x) - d^{\text{out}}(0)), & \text{if } \alpha \neq -1. \end{cases}$$

Solving for $d^{\text{out}}(x)$ yields

$$d^{\text{out}}(x) = \begin{cases} d^{\text{out}}(0)\exp(R(x)/c), & \text{if } \alpha = -1, \\ (d^{\text{out}}(0)^{\alpha+1} + \frac{\alpha+1}{c}R(x))^{\frac{1}{\alpha+1}}, & \text{if } \alpha \neq -1. \end{cases} \tag{3}$$

In the following we derive the particular form for the link function f assuming that $f(x,y) = \tilde{f}(x+y)$ for a function $\tilde{f}$ that we determine in the following. We assume that the fitness has exponential distribution with parameter 1. Therefore,

$$d^{\text{out}}(x) = (n-1)\int_0^\infty f(x,y)\rho(y)dy = \int_0^\infty \tilde{f}(x+y)e^{-y}dy.$$

Setting $z := x+y$ gives

$$\int_x^\infty \tilde{f}(z)e^{-(z-x)}dz = \frac{d^{\text{out}}(x)}{n-1} \Leftrightarrow \int_x^\infty \tilde{f}(z)e^{-z}dz = e^{-x}\frac{d^{\text{out}}(x)}{n-1}.$$

Differentiating wrt x yields

$$\tilde{f}(x) = \frac{d^{\text{out}}(x) - {d^{\text{out}}}'(x)}{n-1}.$$

in particular,

$$f(x,y) = \tilde{f}(x+y) = \frac{d^{\text{out}}(x+y) - {d^{\text{out}}}'(x+y)}{n-1}, \tag{4}$$

where d^{out} is given by (3). Using the expression in (3) together with its derivative and formula (4) yields formula (9).

Since we consider the link function $f(x,y) = \tilde{f}(x+y)$, which only depends on the two fitnesses via their sum, the corresponding results for the in-degree distribution are exactly the same as for the out-degree distribution. Hence, the in-degree distribution also follows a power law.

4.2 Power law from exponential mixture model

In our model for the parameter matrix λ we assume that λ_{ij} is the sum of two independent Gamma distributed random variables with shape parameter $\zeta > 0$ and scale parameter $\eta > 0$. Hence (conditional on ζ, η), this sum follows again a Gamma distribution with shape parameter 2ζ and scale parameter η. Harris (1968) has shown that exponential mixture models in which the rate parameter of the exponential distribution follows a Gamma distribution have a Pareto II (Lomax) distribution. This is the statement of the following lemma.

Lemma 4.1. *Let λ be a random variable with Gamma distribution with shape parameter $\zeta > 0$ and scale parameter $\eta > 0$. Let X be a random variable such that $X|\lambda$ has exponential distribution with parameter λ. Then, X follows a Pareto II (Lomax) distribution with scale parameter $1/\eta$ and shape parameter ζ.*

Proof. For $y > 0$ the pdf of λ is given by $f_\lambda(\lambda) = \frac{1}{\Gamma(\zeta)\eta^\zeta}\lambda^{\zeta-1}\exp(-\frac{\lambda}{\eta})$. Let $x > 0$, then

$$\begin{aligned}\mathbb{P}(X > x) &= \int_0^\infty \mathbb{P}(X > x|\lambda)f_\lambda(\lambda)d\lambda = \int_0^\infty e^{-\lambda x}f_\lambda(\lambda)d\lambda = \frac{1}{\Gamma(\zeta)\eta^\zeta}\int_0^\infty \lambda^{\zeta-1}\exp(-\lambda(x+\frac{1}{\eta}))d\lambda \\ &= \frac{(x+\frac{1}{\eta})^{-\zeta+1}}{\Gamma(\zeta)\eta^\zeta}\int_0^\infty ((x+\frac{1}{\eta})\lambda)^{\zeta-1}\exp(-\lambda(x+\frac{1}{\eta}))d\lambda = (1+\eta x)^{-\zeta}.\end{aligned}$$

□

5 Data used in simulation study in Subsection 5.3

Table 1: Balance sheet data (in million Euros) from banks in the EBA 2011 stress test used in Section 5.3:

Bank code	Bank	$a^{(e)} + a$	a	w
DE017	DEUTSCHE BANK AG	1,905,630	47,102	30,361
DE018	COMMERZBANK AG	771,201	49,871	26,728
DE019	LANDESBANK BADEN-WURTTEMBERG	374,413	91,201	9,838
DE020	DZ BANK AG	323,578	100,099	7,299
DE021	BAYERISCHE LANDESBANK	316,354	66,535	11,501
DE022	NORDDEUTSCHE LANDESBANK -GZ-	228,586	54,921	3,974
DE023	HYPO REAL ESTATE HOLDING AG	328,119	7,956	5,539
DE024	WESTLB AG, DUSSELDORF	191,523	24,007	4,218
DE025	HSH NORDBANK AG, HAMBURG	150,930	4,645	4,434
DE027	LANDESBANK BERLIN AG	133,861	27,707	5,162
DE028	DEKABANK DEUTSCHE GIROZENTRALE	130,304	30,937	3,359

6 Empirical evidence of the missing data problem

Typical data sets available on interbank exposures fall essentially into three categories: “data from large exposures, payment systems and credit registers”, (Langfield et al., 2014, p. 302). (Langfield et al., 2014, Appendix A) contains an overview of these data sets and the literature that analyses them.

The 2007/2008 financial crisis lead to several new initiatives to collect more financial data both on national and international level. As part of these initiatives collection of data on financial networks has also been considered and partially improved. In the following we provide some brief examples which show that missing data are still a major concern to regulators and policy makers.

As one example where new data have been collected in a national initiative, we look at the UK and the recent description of data available to Bank of England, see Langfield et al. (2014). The new data set that is available to Bank of England contains far more information than most of the existing data sets. “UK banks report their exposures to other banks and broker dealers by financial instruments, including lending (unsecured, secured and undrawn); holdings of equity and fixed-income securities issued by banks; credit default swaps bought and sold; securities lending and borrowing (gross and net of collateral); and derivatives exposures (with breakdown by asset class). Moreover, banks report exposures with breakdown by maturity of the instruments. Banks’ internal risk management limits with respect to counterparties and instruments are also supplied. Each bank reports exposures by instruments to their top-20 bank and broker-dealer counterparties”, (Langfield et al., 2014, Section 2).

The new data set still has two major limitations: There are jurisdictional data constraints: “Banks in the UK which are subsidiaries of a foreign parent - comprising 43% of all UK banks - only report their UK subsidiaries’ interbank exposures, not those of the foreign group. Nevertheless, these UK subsidiaries account for a sizable share (41%) of their groups’ global assets. In addition, we do not observe interbank positions held by banks with no regulated subsidiary in the UK”, (Langfield et al., 2014, Appendix A).
In addition there are limitations on the number of reported counterparties: “UK banks report exposures to their top 20 bank and broker-dealer counterparties. Exposures to counterparties beyond the top 20 are not observed”, (Langfield et al., 2014, Appendix A).

The jurisdictional constraints will likely remain a problem in the future. Some exceptions will arise due to the G-20 Data Gaps Initiative which contains recommendations by the International Monetary Fund (IMF) and the Financial Stability Board (FSB) to enhance data collection on an international level. A particular focus is on global systemically important banks (G-SIBs): "Since 2009, a new conceptual framework has been built to assess global network connections of G-SIBs and their linkages with financial systems. [...] Recommendation I.8 of DGI-1 is considered completed with the collection of consistent institution to institution (I-I) data through a template that identifies bilateral credit exposures and funding liabilities of G-SIBs, combined with the launching of the BIS International Data Hub to host the database (Phase 1 in March 2013 and Phase 2 in June 2015). Recommendation I.9 of DGI-1 is close to completion pending the FSB Plenary approval of the Phase 3 template which focuses on the granular institution to aggregate (I-A) exposure on funding data", (Financial Stability Board & International Monetary Fund, 2015, p. 25). The second phase of the G-20 Data Gaps Initiative contains a recommendation (Financial Stability Board & International Monetary Fund, 2015, Recommendation II.4, p. 26) to ensure the regular collection and sharing of data on G-SIBs and to possibly extend the data collection to global systemically important non-bank financial institutions (e.g. insurance companies).

Hence we see that despite significant changes and improvements on data collection missing data are still a concern to regulators as of 2015.

References

Financial Stability Board & International Monetary Fund (2015). The financial crisis and information gaps - sixth progress report on the implementation of the G-20 data gaps initiative. Available at http://www.imf.org/external/np/g20/pdf/2015/6thprogressrep.pdf.

Harris, C. M. (1968). The Pareto distribution as a queue service discipline. *Operations Research* **16**, 307–313.

Langfield, S., Liu, Z. & Ota, T. (2014). Mapping the UK interbank system. *Journal of Banking & Finance* **45**, 288–303.

Plummer, M., Best, N., Cowles, K. & Vines, K. (2006). CODA: Convergence diagnosis and output analysis for MCMC. *R news* **6**, 7–11.

Servedio, V. D. P., Caldarelli, G. & Buttà, P. (2004). Vertex intrinsic fitness: How to produce arbitrary scale-free networks. *Phys. Rev. E* **70**, 056126.