



Contents lists available at ScienceDirect

Physica A

journal homepage: www.elsevier.com/locate/physa

Beyond the power law: Uncovering stylized facts in interbank networks



Benjamin Vandermarliere^{a,b,*}, Alexei Karas^{c,d}, Jan Ryckebusch^a,
Koen Schoors^b

^a Department of Physics and Astronomy, Ghent University, Belgium

^b Department of General Economics, Ghent University, Belgium

^c University College Roosevelt, The Netherlands

^d Utrecht University School of Economics, The Netherlands

HIGHLIGHTS

- Interbank measures are heavy-tailed.
- Evidence for power-law behavior in tail is not strong.
- Lognormal and stretched exponentials can fit the entire distributions of the interbank network measures.
- Robust for time-aggregation and distinctive phases.

ARTICLE INFO

Article history:

Received 24 September 2014

Received in revised form 4 January 2015

Available online 16 February 2015

Keywords:

Interbank networks

Power law

Lognormal

Stretched exponential

Distributions

Heavy tailed

ABSTRACT

We use daily data on bilateral interbank exposures and monthly bank balance sheets to study network characteristics of the Russian interbank market over August 1998–October 2004. Specifically, we examine the distributions of (un)directed (un)weighted degree, nodal attributes (bank assets, capital and capital-to-assets ratio) and edge weights (loan size and counterparty exposure). We search for the theoretical distribution that fits the data best and report the “best” fit parameters.

We observe that all studied distributions are heavy tailed. The fat tail typically contains 20% of the data and can be mostly described well by a truncated power law. Also the power law, stretched exponential and log-normal provide reasonably good fits to the tails of the data. In most cases, however, separating the bulk and tail parts of the data is hard, so we proceed to study the full range of the events. We find that the stretched exponential and the log-normal distributions fit the full range of the data best. These conclusions are robust to (1) whether we aggregate the data over a week, month, quarter or year; (2) whether we look at the “growth” versus “maturity” phases of interbank market development; and (3) with minor exceptions, whether we look at the “normal” versus “crisis” operation periods. In line with prior research, we find that the network topology changes greatly as the interbank market moves from a “normal” to a “crisis” operation period.

© 2015 Elsevier B.V. All rights reserved.

* Correspondence to: Proeftuinstraat 86, 9000 Ghent, Belgium.

E-mail addresses: Benjamin.Vandermarliere@UGent.be (B. Vandermarliere), a.karas@ucr.nl (A. Karas), Jan.Ryckebusch@UGent.be (J. Ryckebusch), Koen.Schoors@UGent.be (K. Schoors).

<http://dx.doi.org/10.1016/j.physa.2015.01.058>

0378-4371/© 2015 Elsevier B.V. All rights reserved.

1. Introduction

The frequency of an event follows a power law when that frequency varies as a power of some attribute of that event (e.g. its size). Power-law distributions have been claimed to occur in an extraordinarily diverse range of phenomena from the sizes of wars, earthquakes and computer files to the numbers of papers scientists write and citations those papers receive [1]. In economics and finance, power laws have been documented for income and wealth [2], the size of cities and firms, stock market returns, trading volume, international trade, and executive pay [3].

Most relevant to this paper, the tail parts of interbank network characteristics, such as degree distribution, have been shown to follow a power law too (see Ref. [4] for Brazil, [5,6] for Austria, [7] for Japan, and [8] for the commercial banks in the US). This ubiquitous presence of power laws has resulted in an extensive search for universal dynamics that can explain their existence (see Refs. [9,10] for examples of such search in interbank networks).

Recently, however, Clauset et al. [11] (followed by Ref. [12]) call these findings into question. In particular, they criticize the commonly used methods for analyzing power-law data, such as least-squares fitting, which can produce inaccurate estimates of parameters for power-law distributions or provide no indication of whether the data obey a power law at all. Clauset et al. propose a statistical framework for discerning and quantifying power-law behavior in empirical data and apply that framework to twenty-four real-world datasets, each of which has been conjectured to follow a power law. For most datasets they find moderate to weak evidence in favor of power laws.

This debate about the potential of power laws to capture the underlying network dynamics is important for economic policy. For example, during the recent financial crisis, the interbank lending market was one of the most important channels of financial contagion. The malfunctioning of the interconnectivity of the interbank lending network, caused a liquidity drought with consequences reverberating throughout the entire economy. Since then, interbank markets research has proliferated. In those studies one wishes to uncover the network topology of interbank markets, to understand how they function, and how they could catalyze a systemic meltdown [13,14]. Current research on contagion in interbank markets often relies on a scale-free topology to simulate the interbank network [15,16]. This choice likely affects the outcome of conducted stress tests (as is explicitly confirmed by Ref. [16]) and, therefore, the policy implications stemming from them. Yet the evidence supporting this choice is not ironclad. Understanding the properties of the tail is crucial to understand shock propagation in dynamic networks. The authors of Refs. [17,18], among others, find that only a small fraction of possible network structures may spread relatively sizable contagion losses across the system, thus highlighting the non-linear nature of shock propagation effects and stressing that contagion is to a considerable extent a tail risk problem.

This paper contributes to the debate by providing a detailed analysis of the network characteristics of a real interbank network over an extended period of time. We use daily data on bilateral interbank exposures and monthly bank balance sheets to study network characteristics of the Russian interbank market over August 1998–October 2004. Among other things, the analysis allows one to determine the theoretical distributions of connectivity among banks via interbank loans, crucial to assess efficiency and stability of the Russian interbank market. We focus on measures that represent essential input for most of the interbank contagion simulations. Specifically, we examine the distributions of (un)directed (un)weighted degree, nodal attributes (bank assets, capital and capital-to-assets ratio) and edge weights (loan size and counterparty exposure). Using the methodology of Ref. [11] we set up a horse race between the different theoretical distributions to find one that fits the data best. We then study the time evolution of the best-fit parameters.

We observe that all studied distributions are heavy tailed. The fat tail typically contains 20% of the data and can be systematically described by a truncated power law. In most cases, however, separating the bulk and tail parts of the data is hard, so we proceed to study the full range of the events. We find that the stretched exponential and the log-normal distributions fit the full range of the data best. Our conclusions turn out to be robust to whether we aggregate the data over a week, month, quarter or year. Further, we find no qualitative difference between the “growth” and “maturity” phases of interbank market development and little difference between the “crisis” and “non-crisis” periods.

Section 2 describes our data, defines the network measures we study, and summarizes the conclusions from previous studies of those measures. Sections 3 and 4, respectively, describe and illustrate the methodology we use to find the theoretical distribution that fits the data best. Section 5 reports the results. Section 6 concludes.

2. Data and definitions

2.1. Data source

Mobile, a private financial information agency, provided us with two datasets for the period August 1998–October 2004.¹ The information in the datasets is a part of standard disclosure requirements and is supplied to the regulator on a monthly basis. The first dataset, described in Ref. [19], contains monthly bank balances for most Russian banks. From this dataset we take two variables: total assets and capital. The second dataset contains monthly reports “On Interbank Loans and Deposits” (official form’s code 0409501) and represents a register of all interbank loans issued in the Russian market. For each loan we know its size, interest rate, issuer, receiver, reporting date and maturity date. On average, about half of the Russian banks are

¹ For more information on the data provider see its website at www.mobile.ru.

Table 1

The levels of aggregation and the number of banks for earlier empirical studies of interbank markets (overnight as well as credit networks).

Country	Reference	Sample period	Aggregation level	# Banks
Mexican	[22]	2005–2010	Day	40
Italy	[23]	1999–2002	Day	±200
Italy	[24]	1999–2010	Day/quarter	±200
Japan	[7]	2009m3–2013m3	Month	125
Brazil	[4]	2007–2008	Month	2400
Austria	[5]	2000–2003	Month (only 10)	900
Germany	[20,21]	2002q1–2012q3	Quarter	2000+

active on the interbank market. Consequently, the analysis of interbank network measures includes fewer banks than the analysis of balance sheet indicators. With regard to the maturity of the loans, we discriminate between short-term and long-term loans. Short-term loans are defined as loans with a one-day or a one-week maturity. In this paper, we restrict ourselves to short-term loans which account for more than 80% of the transactions both in terms of the number and of the volume. The reasons for this restriction is because the data provide information about the repayment date and not about the issuance date of the loans. This makes it hard to infer the exact duration of the connection between two banks for the long-term loans.

2.2. Network measures

Before we start constructing an interbank network from our data, we stress that there are several different ways of how banks can interact. For example, one can have liability [7,20,21,4,5,22] and overnight lending networks [23,24], but one can also construct networks of the financial payment flows between banks, or in aggregated form between countries [25–29,22]. Banks are also connected by mutual cross holdings [30]. In this work we choose to consider all loans between banks with a maturity of less than a week. Hence, we construct an interbank network which combines the overnight market with the short-term liability market. This enables us to keep our finger on the pulse of the interbank market, by not including longer and hence stickier contracts.

We use the two aforementioned datasets to construct an interbank network with banks as nodes and mutual contracts representing edges. This procedure is performed for aggregated data covering various time intervals. In every situation, we construct three versions of the interbank network, which differ in the level of detail in quantifying the edges. In the undirected version, an undirected unweighted edge is established between banks if they exchange money on the interbank market during the considered time period. For the directed versions, we discriminate between the issuer and the receiver of the interbank loan. A directed edge, which points from issuer to receiver, is established if the issuer has lent money to the receiver in the considered period. In the multidirected version, every interbank loan is represented by a directed edge pointing from issuer to receiver. This implies that there can be several edges between pairs of nodes. The directed and undirected versions represent complementary views of the interbank network. The directed version captures the lending/borrowing activity by looking at the direction of the flows and therefore the contribution of each bank, whereas the undirected version merely captures the existence of interbank relations.

In this paper we focus on network characteristics that represent essential input for typical interbank contagion simulations [31,32,15,33,18,17,7,34]. We distinguish three types of network characteristics: nodal attributes, edge weights and various measures of a node's centrality.

We consider three nodal attributes: bank capital, total assets and leverage. Total assets proxy for bank size. Capital measures its financial buffer. Leverage is defined as capital divided by total assets. Upon evaluating nodal attributes, we exclude banks with negative assets (which we attribute to data errors) or negative capital (banks in effective default). In determining the distributions of edge weights and of node centrality measures we do not exclude banks with negative assets or negative capital. Edge weights are measured separately for the directed and the multidirected networks. For the former, the edge weight equals counterparty exposure, that is, the total amount of money the issuer has lent to the receiver in the considered time period. For the multidirected network, the edge weight equals the size of the issued loan.

We consider various measures of a node's centrality. For an undirected network, we define the degree of a node as the number of edges connected to that node. It measures the number of counterparties a bank has on the interbank market. For the directed network, an in-degree (out-degree) of a node is the number of incoming (outgoing) edges; it measures the number of interbank lenders (borrowers) a bank has. For the multidirected network, an in-degree (out-degree) is the number of received (issued) loans. Finally, for each bank we define the total in-exposure (out-exposure) as the total amount of money borrowed from (lent to) the market.

The interbank credit network characteristics we consider in this paper have already been studied for a number of countries. Before sketching the main findings of those studies, we would like to draw the attention to Table 1 which summarizes the sheer variety of considered aggregation periods and of the actual networks' size. Because of this variation, care needs to be taken when comparing our results with those of other studies.

The majority of studies find numerous heavy tailed variables in the interbank network. In addition, many authors propose a power law as the best-fit candidate to (at least the tail of) the empirical distribution. In particular, Goddard et al. [8] find that the asset size distribution of US commercial banks is well described by a truncated lognormal, while the tail part is well

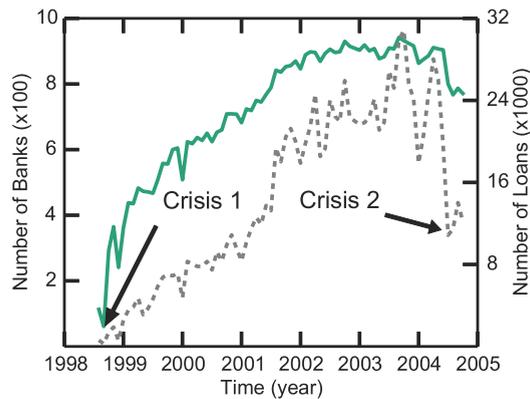


Fig. 1. (Color online) The time dependence of the number of active banks (solid line) and the number of interbank loans (dashed line) in the Russian interbank network between 1998 and 2005. Data are aggregated over a month. The arrows indicate the start of the two “crises”.

fitted with a power law. Cont et al. [4] obtain fair fits to the tails of the in-degree, out-degree, total degree, and exposure size distributions of the Brazilian interbank network with power laws. In the Austrian interbank market the loan size distribution is well described by a power law [5,6]. That same study also examines the degree distribution of the undirected network and the in- and out-degree distributions of the directed alternative. Each of these three distributions is seemingly well described by two power laws, one for the low-degree region and one for the tail part. For Japan, Kanno et al. [7] find a power-law distribution for the entire range of the counter-party exposure size distribution. Finally, Iori et al. [23] investigate the distributions of the in-degree, out-degree and exposure of the individual banks in the Italian market. The authors do not attempt to fit their data with some theoretical distribution. They do find, however, a structural difference in the network topology between the months of the global financial crisis of 2008 and normal operating months.

Few studies, however, consider alternate theoretical distributions as candidates to describe their data. The ones that do [24], in fact, cast doubt on the ability of power laws to provide the best empirical fit. In what follows, we too consider an extensive list of alternative theoretical distributions and identify those that describe both the tail and the entire distribution better than the power law. Along the way, we apply a criterion which allows us to discriminate between the tail and the bulk parts of the data.

2.3. Descriptive statistics

Different phases of the interbank market development (for example, growth versus maturity, or crisis versus non-crisis) may be guided by a different data generating mechanism. For example, Iori et al. [23] find a structural difference in the Italian interbank network topology before versus during the 2008 financial crisis. This finding suggests that the choice of the distribution that best fits the data may vary over time, and in particular may differ between crisis and non-crisis periods. In Section 5 we check whether such variation exists. In this section, we provide evidence that the Russian interbank market, indeed, went through some distinct development phases.

Fig. 1 shows the time series of the number of banks active on the interbank market and of the number of issued loans. For this figure we use data aggregated over a month. Over 1998–2002 the interbank network experiences growth: we observe a steady and comparable increase in the number of active banks and of issued loans. After 2002, the interbank network gradually matures: the number of active banks flattens out while the number of loans per bank grows. Note, however, the strong variation in the number of issued loans from the second half of 2003 onwards.

As is clearly indicated in Fig. 1, our sample period includes two crises: one in August 1998 and one in the summer of 2004. Both crises resulted in a partial meltdown of the Russian interbank market. They coincide with the edges of the sample period and are clearly marked by a reduction in the number of active banks and issued loans. The first crisis got triggered on August 17, 1998 when Russia abandoned its exchange rate regime, defaulted on its domestic public debt and declared a moratorium on all private foreign liabilities. The second crisis was ignited by an investigation of banks accused of money laundering and sponsorship of terrorism. This gave rise to a wave of distrust among banks and a consequent liquidity drought.

Fig. 2 illustrates the impact of August 1998 crisis on the Russian interbank network. Nodes are banks, and directed edges represent issued loans. The first panel shows the activity in the market in the two weeks leading up to the seizure, whereas the second panel covers the two weeks after the collapse. Evidently, we see a decrease in the number of active banks, from 87 to 65, and in the number of loans granted, from 507 to 96. When considering the structure in this unweighted multidirected network we can find a clear distinction between the “normal” and “crisis” periods. The technique of Ref. [35] can be used to uncover groups of nodes, called blocks, which fulfil similar roles in the network.² In the two weeks leading up to the crisis

² The technique of Ref. [35] makes use of a hierarchical version of stochastic block modeling and offers an alternative to the usual modularity optimization algorithms.

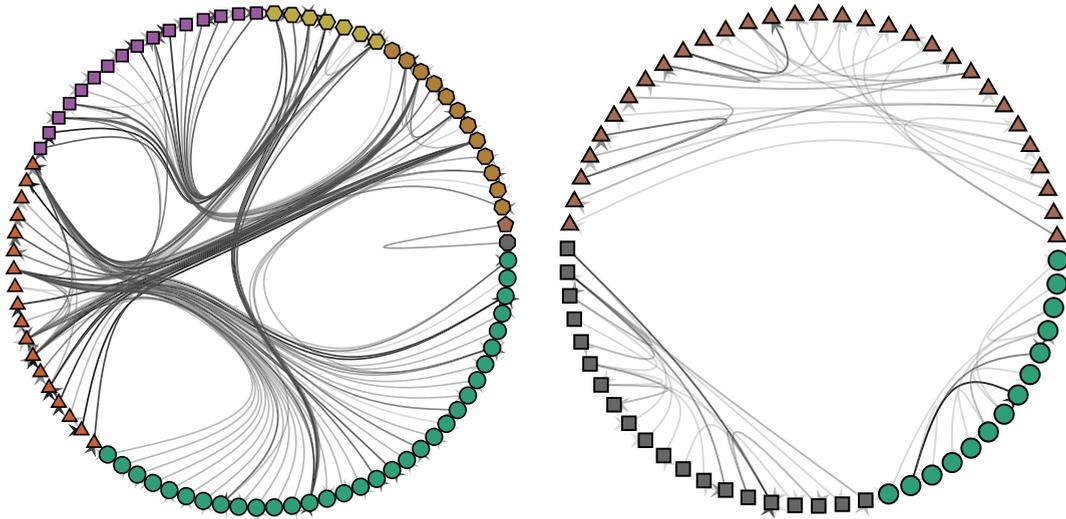


Fig. 2. The activity on the interbank market in the two weeks leading up to August 17, 1998 (left panel) and the two consecutive weeks (right panel). A bank is represented as a node and each link represents an issued loan. Nodes are colored per block as defined in Ref. [35].

we get an interbank market with 5 blocks (and two isolated banks). After the crisis the number of identified blocks goes down to 3. Notably, the identified blocks interact with each other before the crisis but not after. So the crisis disintegrates the network and reconfigures banks’ role in it.

We can also see evidence of different development phases by studying the time evolution of two basic network measures: the average local clustering coefficient, which is a potential indicator of systemic risk [36], and the average shortest path length. For an undirected network, the local clustering coefficient $0 \leq c_u \leq 1$ of a node u is the number of edges between the nodes of the neighborhood of u divided by the maximum amount that could possibly exist between them. It can be conveniently defined as in Ref. [37]

$$c_u = \frac{2T(u)}{\text{deg}(u) (\text{deg}(u) - 1)}, \tag{1}$$

where $T(u)$ is the number of triangles (a subgraph with three nodes and three edges) attached to u and $\text{deg}(u)$ is the degree of u . The average of this local clustering coefficient is

$$C = \frac{1}{n} \sum_{u \in G} c_u, \tag{2}$$

where n is the number of nodes in the network G [38]. The average shortest path length D is defined as

$$D = \sum_{u,v \in V} \frac{d(u,v)}{n(n-1)} \tag{3}$$

where V is the set of nodes in the network G , and $d(u,v)$ is the length of the shortest path from node u to v [39]. In a disconnected graph D can only be computed for the largest connected component.

Fig. 3 shows the time evolution of C and D computed per month for the largest connected component of the undirected interbank network. First, note that the local clustering coefficient averaged over all nodes and time periods ($C = 0.198$) is nearly identical to the one reported by Ref. [4] for the Brazilian interbank market ($C = 0.2$) but is much lower than for the German market (directed clustering $C = 0.80$) [20]. In contrast, the average shortest path length ($D \approx 3$) is notably higher compared to German ($D = 2.14$) [20], to Austrian ($D = 2$) [5] and Mexican ($D = 1.7$) [22] interbank networks. In the case of the Mexican and Austrian ones, this difference is likely driven by the fact that the Russian network is sizably larger.

The C and D tend to move in the opposite direction: three months after the first crisis hits C drops while D spikes; during the growing phase of the network (1999–2002) C grows while D falls; during the mature phase (2002–2004) both measures stabilize at $C \approx 0.22$ and $D \approx 3$; finally, during the 2004 crisis C drops while D spikes again. The average local clustering coefficient C , however, tends to have bigger fluctuations from period to period. In contrast, the time series of the average shortest path length is very smooth, and the only two obvious spikes occur around the two crises. Clearly, those crises disrupted the overall network structure: the liquidity drought, which is equivalent to the pruning of links, caused a significant increase in the average shortest path length. Even the decrease in the number of nodes, and hence shrinking of the network, could not offset this effect. The 1998 spike in D is particularly remarkable given we only consider the largest connected component, that is, about one third of the nodes (see Fig. 2).

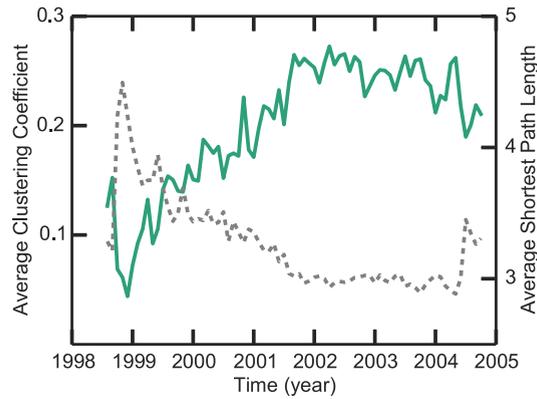


Fig. 3. (Color online) The time dependence of the average clustering coefficient (solid line) and the average shortest path length (dashed line) calculated for the most connected component of the undirected version of the Russian interbank network between 1998 and 2005. Data are aggregated over a month.

Table 2

Definition of the normalized distributions used in this work. The constant C is defined by the normalization condition $\int_{x_{min}}^{\infty} f(x)dx = 1$.

Distribution	$f(x)$
Power law (PL)	$Cx^{-\alpha}$
Truncated power law (TPL)	$Cx^{-\alpha} e^{-\lambda x}$
Exponential (Exp)	$Ce^{-\lambda x}$
Stretched exponential (SEXP)	$Cx^{\beta-1} e^{-(\lambda x)^{\beta}}$
Log-normal (LN)	$\frac{C}{x} \exp\left[-\frac{(\ln x - \mu)^2}{2\sigma^2}\right]$

3. Methodology

For each network measure we test whether the power law or an alternate fat-tailed distribution fits the data best. First, we fit distributions to the tail, then to the entire data range. Along the way, we apply a criterion which allows us to discriminate between the tail and the bulk parts of the data.

To study the tail we adopt the methodology of Ref. [11]. First, we fit a power law (PL) to the data to determine the starting point of the tail part (the so-called cut-off x_{min} of the scaling range). Then we fit each of the theoretical distributions to the tail using maximum likelihood (ML). Finally, we use a relative goodness-of-fit test to select the distribution that fits the data best. We consider the same selection of theoretical distributions as in Ref. [11]. Their functional forms are listed in Table 2 and involve one or two free parameters.

For completeness and to introduce the notation we briefly sketch the methodology of Ref. [11]. Consider a given ordered dataset $\{x_j, j = 1, \dots, N\}$. Every entry x_j is a potential x_{min} and for each of those we compute the ML estimate of the power-law exponent α

$$\hat{\alpha}(x_j = x_{min}) = 1 + (N - j + 1) \left[\sum_{i=j}^N \ln \frac{x_i}{x_{min}} \right]. \quad (4)$$

We then use the Kolmogorov–Smirnov test to select the optimum \hat{x}_{min} . It is defined as the cut-off which minimizes the quantity

$$Z = \max_{x \geq x_{min}} |S(x) - P(x)|. \quad (5)$$

Here, $S(x)$ is the cumulative distribution function (CDF) of the observed values for $x_j \geq x_{min}$, and $P(x) = \frac{C}{-\alpha+1} (x^{-\alpha+1} - x_{min}^{-\alpha+1})$ is the CDF of the power-law fit to the tail part of the data.

In the next step, for given x_{min} , we perform ML fits to the data with the other four candidate distributions of Table 2. As argued by Ref. [11], it is more useful to know which distribution is the best possible fit candidate, rather than the goodness of fit for each distribution individually. To compare the relative goodness of the different fits, we compute the likelihood ratios R for the pairs of probability density functions (PDFs) $p_1(x_i)$ and $p_2(x_i)$

$$R(p_1, p_2) = \frac{L_1}{L_2} = \prod_{i=j}^N \frac{p_1(x_i)}{p_2(x_i)}. \quad (6)$$

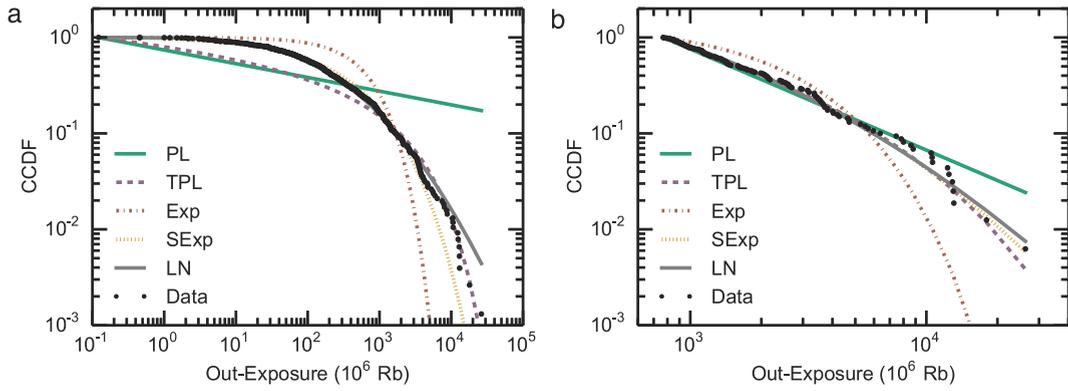


Fig. 4. (Color online) The complementary cumulative distribution function (CCDF) of the entire (a) and tail part (b) of the monthly-aggregated out-exposure distribution for January 2002. Also shown are the ML fits to the CCDFs for the five distributions of Table 2.

The corresponding normalized loglikelihood ratios $\mathcal{R}(p_1, p_2)$ read

$$\begin{aligned} \mathcal{R}(p_1, p_2) &= \frac{1}{\sigma_{12}\sqrt{N-j+1}} \sum_{i=j}^N \left[\ln \frac{p_1(x_i)}{p_2(x_i)} \right] \\ &= \frac{1}{\sigma_{12}\sqrt{N-j+1}} \sum_{i=j}^N \left[l_i^{(1)} - l_i^{(2)} \right], \end{aligned} \tag{7}$$

where $l_i^{(k)} = \ln p_k(x_i)$ and σ_{12} is defined as

$$\sigma_{12}^2 = \frac{1}{N-j+1} \sum_{i=j}^N \left[\left(l_i^{(1)} - l_i^{(2)} \right) - \left(\bar{l}^{(1)} - \bar{l}^{(2)} \right) \right]^2. \tag{8}$$

The $\mathcal{R}(p_1, p_2)$ is positive (negative) if the data is more likely in the p_1 (p_2) distribution.

In order to guarantee that the value of \mathcal{R} is not merely a product of fluctuations and that the true expectation value of \mathcal{R} is zero, we compute the probability that the measured normalized log likelihood ratio has a magnitude as large or larger than the observed value $|\mathcal{R}|$. This so-called p -value is defined as

$$p(\mathcal{R}) = \frac{1}{\sqrt{2\pi}} \left[\int_{-\infty}^{-|\mathcal{R}|} e^{-\frac{t^2}{2}} dt + \int_{|\mathcal{R}|}^{\infty} e^{-\frac{t^2}{2}} dt \right]. \tag{9}$$

The distribution p_k with the highest value of

$$g(p_k) = \sum_{l \neq k} \mathcal{R}(p_k, p_l), \tag{10}$$

is considered as the most suitable distribution among the different candidates $\{p_i\}$. Bootstrapping and the Kolmogorov–Smirnov test are alternate methods to compute the p values. Both methods, however, are subject to some pitfalls as outlined in Ref. [40]. By using the PL as a benchmark and letting it decide on the value of x_{min} , we are confident that our methodology produces the best possible PL fit to the data. If it turns out that an alternate distribution provides a significantly better fit to the tail part of the data than the PL, strong evidence emerges that this theoretical distribution better accounts for reality than a power law.

We stress that the above-sketched methodology of Ref. [11] cannot only be applied to the analysis of the tail part of the data but also to the full range of the data by setting $x_{min} = x_1$. When considering the full range of data it is natural to select a set of theoretical distributions that encompass a Gaussian-like regime (thermal or bulk part) supplemented with a fat tail (superthermal part). Examples of such distributions include the stretched exponential and the log-normal.

4. Illustration of methodology for out-exposure distribution

As a prototypical example of the methodology sketched in Section 3, we explore the out-exposure distribution for the monthly aggregated data. The data for the out-exposure distribution of January 2002 shown in Fig. 4 are exemplary for all considered 70 months (January 1999–October 2004). The considered values for the out exposure extend over more than five orders of magnitude and the distribution can be labeled as heavy-tailed.³

³ A distribution is defined as heavy-tailed if it is not exponentially bounded [40].

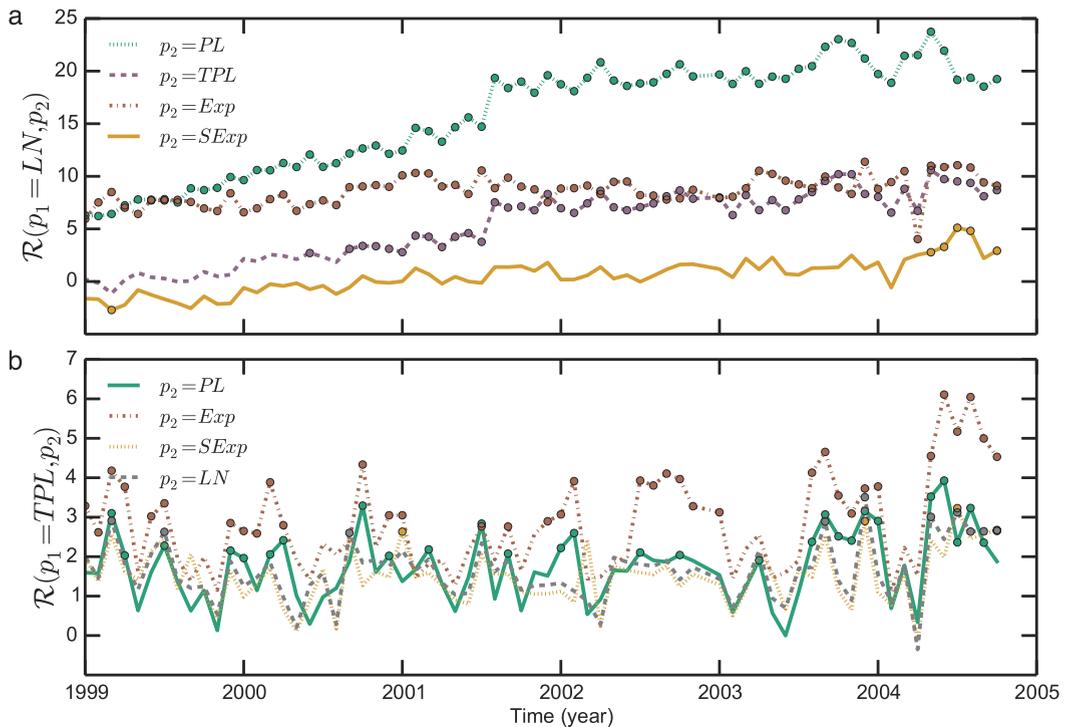


Fig. 5. (Color online) (a) The time series of the normalized loglikelihood ratios $\mathcal{R}(p_1 = \text{LN}, p_2)$ of Eq. (7) for the different fits to the total distribution of the monthly-aggregated out-exposure. (b) The time series of $\mathcal{R}(p_1 = \text{TPL}, p_2)$ for the fit to the tail of the monthly-aggregated out-exposure distribution.

Using the methodology outlined in Section 3 we identify the tail part of the data. From Fig. 4(b) it is clear that fits to the tail part of the considered data with the two-parameter log-normal (LN), stretched exponential (SExp) and truncated power law (TPL) distributions outperform those with the one-parameter exponential (Exp) and power law (PL) distributions. The parameter λ in the TPL accounts for the finite-size effects near the upper edge of the out-exposure distribution. For each of the 70 months in our sample we compute the ratios $\mathcal{R}(p_1, p_2)$ of Eq. (7) for all pair combinations out of the list of five distributions of Table 2. In Fig. 5 we display the time series of the $\mathcal{R}(\text{TPL}, p_2)$ for $p_2 = \text{PL}, \text{Exp}, \text{SExp}$, and LN. We observe that $\mathcal{R}(\text{TPL}, p_2)$ is mostly positive, which indicates that the TPL offers the best overall description of the tail of the out-exposure data. In order to test the significance of this observation, we evaluate the p values. If the normalized ratio for a pair of distributions in a given month is significant at a 1% level, the data points in Fig. 5 are dotted. The figure indicates that the two-parameter TPL is a significantly better fit to the tail of the monthly-aggregated out-exposure distribution than a power law and an exponential. The TPL, however, does not provide a significantly better fit than the LN or SExp for most of the months.

Fig. 6(b) shows the time evolution of the TPL parameters. Whereas α fluctuates around the same average value, λ falls over time. Hence, as time proceeds, the exponential cutoff to the power law shifts to larger values of the out-exposure.

To weigh the relative performance of the different distributions we compute their monthly $g(p_i)$ scores (see Eq. (10)). We dub p_i as the “best” overall fit candidate when its $g(p_i)$ score is highest for the largest fraction of the 70 months in our sample. As can be seen in Table 3, the truncated power law is the best-fit candidate in 94% of the considered months. Nevertheless, $\mathcal{R}(\text{TPL}, \text{SExp})$ is not significantly different from zero in 94% of the considered months as is the case with $\mathcal{R}(\text{TPL}, \text{PL})$ in 59% and with $\mathcal{R}(\text{TPL}, \text{LN})$ in 83% of the months. So we list the SExp as well as PL and LN as alternate best-fit candidates. In general, if a distribution is not significantly worse than the best-fit candidate in more than half of the months, it is mentioned in the last column of Table 3. We conclude that although the tail is described best by a TPL, the fit is not significantly better than the power-law, stretched-exponential and log-normal fits.

We now turn our attention to the distributions covering all out-exposure data points, from small to large. Again we consider the five PDFs of Table 2. From the ML fits, for example for January 2002 displayed in Fig. 4(a), we immediately notice that the exponential and power law do not fit the entire range of data well. This is a clear indication of the fact that the distribution of the out-exposure has both a thermal (Gaussian-like) and a superthermal (fat tail) part. Fig. 5 shows the normalized loglikelihood ratios $\mathcal{R}(p_1 = \text{LN}, p_2)$ over time. Using the same criteria as for the tail part, we conclude that the log-normal is the best fit candidate for 80% of the 70 months studied.

We also find that the stretched exponential is not significantly worse in 57% of the cases, which is reflected in the persistently small values of their likelihood ratios in the top panel of Fig. 5. During the 2004 crisis the stretched exponential has a significantly better fit than the log-normal. As the number of interbank loans is subject to a sudden drop in the summer

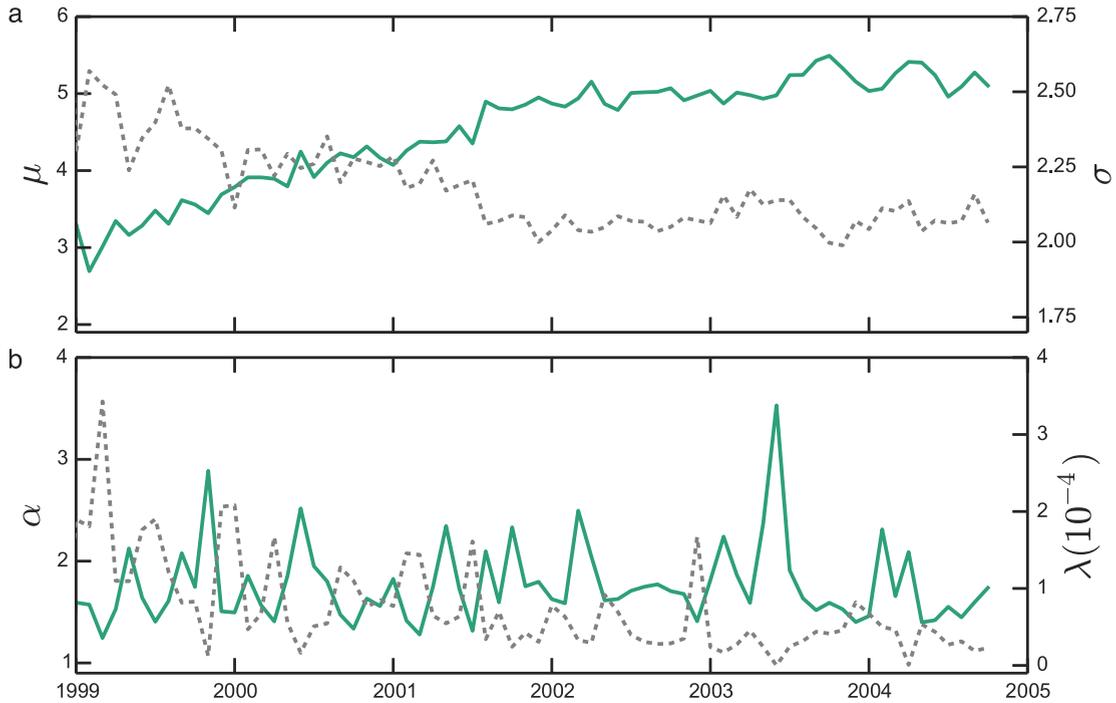


Fig. 6. (Color online) The time evolution of the parameters for the best fit to the total (a) and tail (b) part of the out-exposure distributions. The total data range is fitted with a log-normal distribution with parameters μ (solid line) and σ (dashed line). The tail part is fitted with a truncated power-law distribution with parameters α (solid line) and λ (dashed line).

Table 3

Summary of the results of the fits to the tails of the monthly-aggregated data of all the network features considered in this work. We list the average fraction of the events in the tail and the distribution which provides the best fit, together with the percentage of months for which it comes out “best” according to the criterion of Eq. (10). Columns 4 and 5 provide the average values of the parameters of the best-fit distributions. The last column lists the distributions which are not significantly (1% level) worse for more than half of the 70 months (January 1999–October 2004) considered, together with the percentage of months for which this is the situation. The measures with an extra label “(B)” refer to an analysis of a sample obtained after bootstrapping the real data.

Measure	Tail (%)	Best fit (%)	α	λ	Alternate fits (%)
Asset size	23 ± 9	TPL (77)	1.80 ± 0.06	$(4.47 \pm 9.56) \times 10^{-7}$	PL (100), SExp (100), LN (97)
Capital size	31 ± 18	TPL (58)	1.91 ± 0.14	$(1.25 \pm 1.71) \times 10^{-5}$	PL (100), SExp (100), LN (89)
Leverage	24 ± 9	PL (57)	3.67 ± 0.76	None	TPL (100), Exp (86), SExp (100), LN (100)
Leverage (B)	28 ± 10	TPL (57)	3.77 ± 0.10	$(1.81 \pm 2.77) \times 10^{-6}$	PL (100), Exp (93), SExp (100), LN (100)
Loan size	5 ± 9	TPL (75)	2.49 ± 0.47	$(3.22 \pm 3.54) \times 10^{-4}$	PL (91), SExp (87), LN (71)
Counterparty-exposure	12 ± 8	TPL (80)	2.17 ± 0.51	$(2.53 \pm 2.44) \times 10^{-4}$	PL (53), SExp (81), LN (71)
Counterparty-exposure (B)	16 ± 9	TPL (86)	2.02 ± 0.36	$(3.13 \pm 2.33) \times 10^{-4}$	PL (91), SExp (100), LN (100)
Undirected degree	21 ± 16	TPL (100)	1.80 ± 0.36	$(1.83 \pm 0.83) \times 10^{-2}$	PL (70), Exp (68), SExp (97), LN (97)
Directed in-degree	25 ± 20	TPL (94)	1.72 ± 0.43	$(2.65 \pm 2.74) \times 10^{-2}$	PL (71), Exp (68), SExp (91), LN (95)
Directed out-degree	13 ± 8	TPL (89)	2.07 ± 0.66	$(2.93 \pm 1.92) \times 10^{-2}$	PL (85), Exp (98), SExp (100), LN (100)
Directed out-degree (B)	28 ± 18	TPL (77)	1.94 ± 0.71	$(4.13 \pm 3.99) \times 10^{-2}$	PL (99), Exp (100), SExp (99), LN (99)
Multidirected in-degree	23 ± 24	TPL (97)	1.83 ± 0.48	$(5.05 \pm 3.93) \times 10^{-3}$	PL (74), Exp (77), SExp (99), LN (100)
Multidirected out-degree	23 ± 17	TPL (93)	1.74 ± 0.50	$(7.80 \pm 5.83) \times 10^{-3}$	PL (63), Exp (88), SExp (93), LN (91)
In-exposure	30 ± 17	TPL (97)	1.49 ± 0.43	$(7.68 \pm 5.11) \times 10^{-5}$	SExp (57)
Out-exposure	20 ± 11	TPL (94)	1.76 ± 0.39	$(7.49 \pm 6.23) \times 10^{-5}$	PL (59), SExp (94), LN (83)

of 2004 (Fig. 1), this change in the best-fit candidate illustrates how the network adapts to changing overall conditions. Fig. 6(a) shows the time evolution of the log-normal fit parameters. The parameter μ increases during the “growth” stage of the network whereas σ decreases, and both tend to flatten out in the “mature” phase of the network when few nodes or links are added (see Fig. 1).

Now, we investigate in how far the size of the time bin widths affects the conclusions with regard to the “best” theoretical distributions describing the data. To this end, we study the out-exposure and consider different time intervals to aggregate the data. Besides the networks build with all loans issued in one month, we can do this for all loans issued in a week, quarter or year. Fig. 7 shows the out-exposure distribution for the second week of January 2002, January 2002, the first quarter of 2002, and the entire year of 2002. We also include the ML fit with the stretched exponential and log-normal for each of the aggregates. If we calculate $g(p_k)$ for every distribution in Table 2, we find that the log-normal is the preferred candidate to fit each dataset, although, again, the stretched exponential is not significantly worse. For the tail part, the truncated

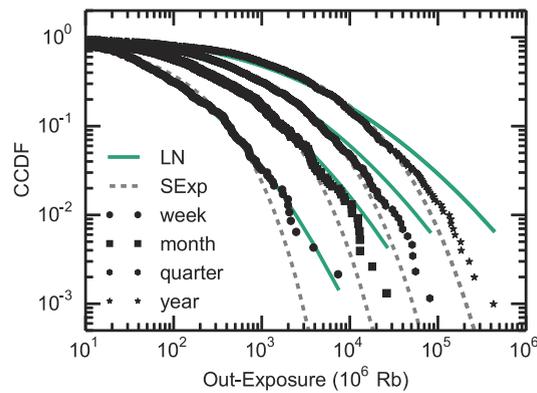


Fig. 7. (Color online) The CCDF of the out-exposure distribution from a network aggregated with data from one week (week 2 of January 2002), from one month (January 2002), from one quarter (first quarter of 2002), and one year (2002). We also show the respective stretched exponential and log-normal fits.

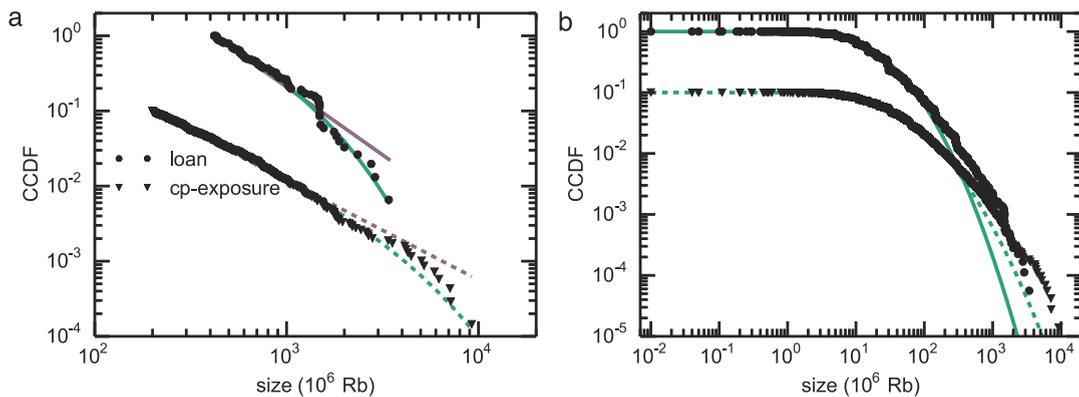


Fig. 8. The CCDF for the loan-size and the counterparty(cp)-exposure for January 2004. Panel (a) shows the tail part of the data together with a PL (gray) and TPL (green) fit. Panel (b) shows the data over the entire range together with LN fit. For illustrative reasons, the data for the cp-exposure have been scaled by a factor of 0.1. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

power law stays on top too. Because the data has a monthly periodicity, linked to the monthly compliance with regulatory requirements, aggregating the data monthly is the most natural thing to do.

To check whether these observations can be generalized, we repeat the illustrated procedure for weekly, quarterly, and yearly aggregates. Table 5 shows the “best” fit candidate for the tail and entire distribution over our data sample. For the tail part, the TPL is preferred over the others for each aggregate. It does so in 85% of the weeks and up to 100% of the years. For the entire distribution, the LN is always preferred.

5. Results

The methodology used for analyzing the out-exposures in the previous section, is used for each of the 12 complementary interbank network measures introduced in Section 2 and listed in Table 3. For a few combinations of network variables and time instances, we find that the x_{min} , which marks the lower boundary of the tail parts of the distributions and results from minimizing the quantity Z of Eq. (5), is close to the upper edge of the distribution. Under those circumstances the identified tail does not hold a sufficient amount of datapoints to perform a meaningful fit. This is particularly common during the 1998 interbank network collapse. For this reason, we only study the best-fit parameters from January 1999 onwards.

A prototypical example of a dataset for which the fit with the functions of Table 2 is not fully satisfactory is shown in Fig. 8. The figure includes the loan size and counterparty-exposure measures for January 2004. Whereas the TPL is a good match for the tail parts of the data, the “best” fit to the entire range of the data with an LN distribution clearly underestimates the probability of events in the tail. Fig. 9 displays for January 2004 the “tail” and “bulk + tail” parts of the 10 other interbank network measures considered. It is clear that the data can be reasonably well described by the adopted procedure and proposed theoretical distributions.

Table 3 reports our findings on the tail parts of the empirical distributions. The second column shows the percentage of data points assigned to the tail via the Kolmogorov–Smirnov criterion of Eq. (5). This number averages to about 20% across the considered network measures but varies a lot. In particular, the loan size, counterparty exposure, and directed out-degree have fewer than 20% of the data-points in the tail. For most of the 70 months included in the analysis, the truncated

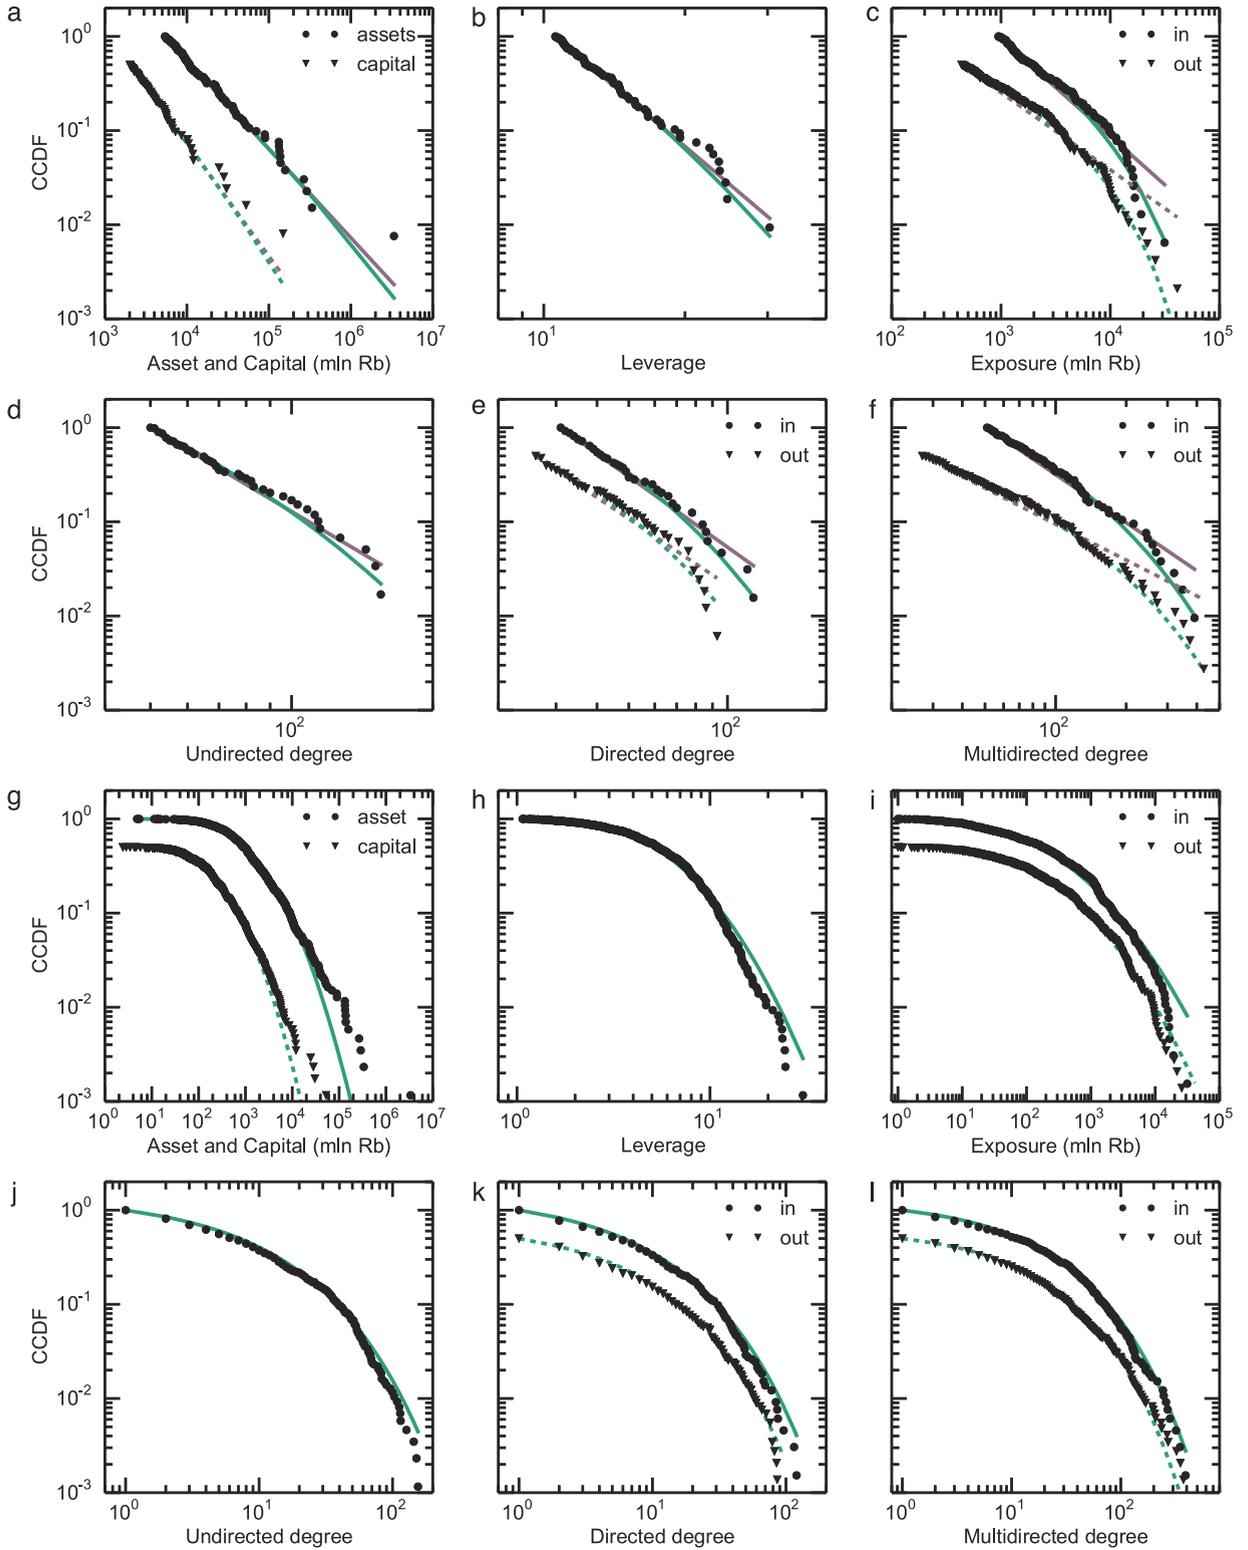


Fig. 9. The CCDF for 10 interbank network measures for the monthly aggregated data in January 2004. Panels (a)–(f) display the tails parts of the data together with the PL (gray) and TPL (green) fits. Panels (g)–(l) contain the data over the entire range together with the “best” fit from Table 4. For illustrative reasons, in those panels with two measures, the data of the second measure have been shifted vertically by multiplying them by a scaling factor of 0.5. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 4

Summary of the results of the fits to the monthly-aggregated data of all the network features considered in this work. We list the distribution which provides the best fit to the full range of the data (“bulk” + “tail”), together with the percentage of months for which it comes out “best” according to the criterion of Eq. (10). Columns 3 and 4 provide the average values of the parameters of the best-fit distribution. The last column lists the distributions which are not significantly (at the 1% level) worse for more than half of the 70 months (January 1999–October 2004) considered, together with the percentage of months for which this is the situation.

Measure	Best fit (%)	Parameter 1	Parameter 2	Alternative fit (%)
Asset size	LN (100)	$\mu = 5.45 \pm 0.63$	$\sigma = 1.91 \pm 0.02$	None
Capital size	LN (100)	$\mu = 4.64 \pm 0.53$	$\sigma = 1.51 \pm 0.03$	None
Leverage	LN (99)	$\mu = 1.57 \pm 0.05$	$\sigma = 0.74 \pm 0.01$	None
Loan size	LN (100)	$\mu = 2.54 \pm 0.34$	$\sigma = 1.27 \pm 0.08$	None
Counterparty-exposure	LN (99)	$\mu = 3.33 \pm 0.31$	$\sigma = 1.56 \pm 0.13$	None
Undirected degree	SExp (54)	$\lambda = 0.29 \pm 0.20$	$\beta = 0.55 \pm 0.06$	LN (94)
Directed in-degree	SExp (66)	$\lambda = 0.41 \pm 0.27$	$\beta = 0.46 \pm 0.07$	TPL (91), LN (71)
Directed out-degree	SExp (83)	$\lambda = 0.43 \pm 0.39$	$\beta = 0.60 \pm 0.05$	TPL (51), LN (97)
Multidirected in-degree	SExp (80)	$\lambda = 0.29 \pm 0.35$	$\beta = 0.43 \pm 0.07$	TPL (90)
Multidirected out-degree	SExp (73)	$\lambda = 0.19 \pm 0.26$	$\beta = 0.53 \pm 0.09$	LN (80)
In-exposure	LN (83)	$\mu = 4.26 \pm 0.83$	$\sigma = 2.37 \pm 0.09$	None
Out-exposure	LN (80)	$\mu = 4.55 \pm 0.65$	$\sigma = 2.13 \pm 0.10$	SExp (57)

power law outperforms, according to our methodology, the other four candidates for every measure. In Table 3 the time-averaged values for the α and λ parameters are also reported. The power α is consistently of the order of two. It emerges from our studies that a fat tail is a characteristic feature of the measurable quantities of an interbank system. We notice that λ is very volatile.

Upon evaluating the last column of Table 3, we find that although the truncated power law systematically provides the best fit, the other candidates are hardly significantly worse at the 1% level. This observation is particularly true for the unweighted degree measures which are discrete. In fact, for the degree measures all directed and undirected versions lead to similar results. We find that none of the five distribution candidates can be conclusively labeled as representing the best fit to the degree data. We note that the fraction of events located in the tail fluctuates a lot.

In order to evaluate the degree of sensitivity of our results to the amount of observations in the tail, we have made use of bootstrapping. For some prototypical measures we have drawn 1000 bootstrapped samples (random samples with replacement) of the real observations. The results of the analysis of the bootstrapped sample for the leverage, counterparty-exposure and directed out-degree are also included in Table 3. For the directed out-degree and the counterparty exposure, the fits to the real and bootstrapped data provide results which are compatible within the error bars. The leverage is the only studied network measure for which the best fit to the real data is provided by a PL. The PL, however, is the best fit for only 57% of the studied months. For all studied months, the TPL emerges as a good alternate fit to the leverage data. Accordingly, it is not all that surprising that a TPL emerges as the best fit to the bootstrapped leverage data. Overall, we may conclude that despite the fact that the number of observations in the tails is rather volatile, the fits to the tails are rather robust.

The fact that the fraction of events in the tail is very volatile hints that a more integrated approach whereby the bulk and the tail parts of the data are simultaneously fitted with a non-Gaussian distribution may lead to a more robust description of the network features.

Table 4 summarizes our fits to the entire distributions of all 12 network measures over all 70 months using monthly-aggregated data. The asset, capital, leverage, loan sizes, as well as the counterparty-exposure and in-exposure distributions are described significantly best by a log-normal. The various degree distributions prefer the stretched exponential, which is on average not significantly better than the log-normal and/or the truncated power law. Similar results are obtained for the directed and undirected versions of the degree distributions. For those network measures for which the stretched exponential represents the best fit, we observe that the parameter λ is very volatile. It is worth noting that the stretched exponential has been put forward as a natural fat tail distribution for many physical and economic phenomena which cannot be satisfactorily described by a power law [41]. In general, the parameters entering the fits to the tail parts are subject to larger fluctuations than those entering the distributions for the complete dataset.

Our results compare well with some of the existing studies of interbank networks, but differ from others. Specifically, our findings are in line with Goddard et al. [8] who study US banks: the log-normal distribution fits the bulk part of the asset-size data well, while the power law does the same for the tail part. Our findings are not too different from those of Cont et al. [4] for the Brazilian interbank network: while our TPL fits to the tails of various degree distributions (see Table 3) deliver values of α between 1.7 and 2.2, Cont et al. [4] report power-law exponents between 2.2 and 2.8 ($\alpha = 2.54$ for total degree, $\alpha = 2.46$ for in-degree, $\alpha = 2.83$ for out-degree, and $\alpha = 2.27$ for out-exposure size). Finally, in contrast to Ref. [5] studying the Austrian interbank market we do not find that the power law provides a good fit to the entire loan-size distribution ([5] report $\alpha = 1.87$). Their power-law exponents reported for the tails of total and in-degree distributions (resp. $\alpha = 2.0$ and $\alpha = 3.11$) are comparable to the TPL values in Table 3, while their out-degree $\alpha = 1.72$ differs substantially.

The most extensive study of interbank distributions was performed for the e-mid market [24]. In line with this work, the authors consider both the complete and tail parts of the distributions and report results for daily and quarterly time aggregates. For the quarterly aggregates of the in-, out and total degree distribution, a stretched exponential emerged as the best fit to the entire distribution. This confirms our findings. A log normal was reported as a best fit to the tail parts of the

Table 5

Robustness of the “best” fits for different aggregates. This table shows for each measure–aggregate pair the preferred fit candidate as well as the percentage of the bins it was considered so. This is done for fit of tail as well as total.

Measure	Tail				Bulk + tail			
	Week	Month	Quarter	Year	Week	Month	Quarter	Year
Loan size	TPL (70)	TPL (75)	TPL (78)	TPL (80)	LN (95)	LN (100)	LN (100)	LN (100)
Counterparty-exposure	TPL (82)	TPL (80)	TPL (83)	TPL (80)	LN (94)	LN (99)	LN (100)	LN (100)
Undirected degree	TPL (94)	TPL (100)	TPL (90)	TPL (100)	SExp (66)	SExp (54)	SExp (52)	SExp (80)
Directed in-degree	TPL (93)	TPL (94)	TPL (90)	TPL (100)	SExp (58)	SExp (66)	SExp (65)	SExp (60)
Directed out-degree	TPL (83)	TPL (89)	TPL (95)	TPL (100)	SExp (64)	SExp (83)	SExp (87)	SExp (100)
Multidirected in-degree	TPL (92)	TPL (97)	TPL (95)	TPL (100)	SExp (78)	SExp (80)	SExp (87)	SExp (100)
Multidirected out-degree	TPL (79)	TPL (93)	TPL (90)	TPL (100)	SExp (67)	SExp (73)	SExp (87)	SExp (100)
In-exposure	TPL (94)	TPL (97)	TPL (100)	TPL (100)	LN (74)	LN (83)	LN (74)	LN (60)
Out-exposure	TPL (85)	TPL (94)	TPL (100)	TPL (100)	LN (83)	LN (80)	LN (57)	LN (80)

data. We also find that the LN can reasonably account for the tail parts. It is slightly outperformed by the truncated power law (which was not included in the set of possible distributions in the work of Ref. [24]) and flags as not significantly worse in our table. For the quarterly aggregated number of transactions, a network feature comparable to the multidirected degree considered here, they also find similar results as for the regular degrees.

Upon scrutiny of the time evolution of the fitted distributions and their parameters, we do not find any systematic changes in the distributions which emerge as best fit to the data, between the “growth” phase (1999–2002) and the “mature” phase (2002–2004). In line with the expectations, however, the extracted parameters are subject to smaller variations as the network grows and the number of nodes and edges increases.

As already mentioned, our dataset covers two crises, the first in August 1998 and the second in the summer of 2004. We would like to know if the best-fit candidates for the crisis periods are different than the ones for the normal operation periods.

As a consequence of the moratorium in August 1998, the interbank market collapsed in September, and then again in December 1998 (see Fig. 1). Due to this collapse, the network is too small and there are too few transactions to gather proper statistics. Fits to the data in this period are inconclusive for most of the measures, and hence are not included in the overall discussion.

During the second crisis period, the network does stay large enough to have proper statistics. From Fig. 5 it is clear that the stretched exponential distribution is significantly better than the log-normal from mid 2004 onwards. Thus, for the out-exposure (as well as the in-exposure), there is a difference between pre- and post-crisis best-fit candidates. For the other measures, however, we find that the preferred distributions are identical in “crisis” and in “normal” operation periods.

To end this section, we test the qualitative robustness of the obtained results using different time windows. This is particularly important in view of the fact that a recent study has shown that the interbank network properties depend on the aggregation period [42]. We stress that on average the number of banks active in one week is 77% of the monthly active ones. Accordingly, some banks are active in the interbank credit market during one week of a particular month. For the quarters and years this is respectively 114% and 132%. This increase is mainly driven by new entrants to the market. Hence, banks tend to participate in the market at least once per month. To check if the preferred distribution is the same for different time aggregates, we repeat the monthly procedure for weekly, quarterly, and yearly aggregates. Table 5 shows the “best” fit for each measure–aggregate pair and the percentage of the bins it is considered so. For the tail parts, we find that the TPL can be considered as the best overall fit candidate. We notice that in general the TPL becomes more preferred as the bin width increases. For each measure, the favored fits to the entire distribution stay the same for different bin widths. We emphasize that the stability of the best-fit distribution with respect to variations in time does not imply that the parameters are not subject to time variations.

6. Conclusion

In this paper we use daily data on bilateral interbank exposures and monthly bank balance sheets to study network characteristics of the Russian interbank market over Aug 1998–Oct 2004. Specifically, we examine the distributions of (un)directed (un)weighted degree, nodal attributes (bank assets, capital and capital-to-assets ratio) and edge weights (loan size and counterparty exposure). Using the methodology of Ref. [11] we set up a horse race between the different theoretical distributions to find one that fits the data best.

In line with the existing literature, we observe that all studied distributions are heavy tailed with the tail typically containing 20% of the data. The tail is best described by a truncated power law, although the fit of other candidate distributions is only marginally worse. In most cases, separating the bulk and tail parts of the data turns out to be hard, and the proportion of observations assigned to the tail varies a lot. More stable fits to the data are obtained in an integrated approach that accounts for both the Gaussian and the non-Gaussian parts of the distributions. We find two distributions that fit the full range of the data best: the stretched exponential for measures related to unweighted degree and the log-normal for everything else. In case of the former, the log-normal performs only marginally worse. The power law distribution is rather ill suited to represent the full range of the studied characteristics of the Russian interbank market.

Our conclusions with regard to the best-fit distributions turn out to be robust to whether we aggregate the data over a week, month, quarter or year. Further, we find no qualitative difference between the “growth” and “maturity” phases of interbank market development and little difference between the “crisis” and “non-crisis” periods.

Our findings support the recent call [11, 12] for more rigorous statistical tests to detect power-law behavior in empirical data. While for most variables we find that the power law fits the tail of the distribution reasonably well, it is:

1. *almost never the best* candidate to describe the tail;
2. *typically not a good* candidate to describe the whole distribution.

Our findings echo those of Ref. [24] who also find that the power law provides an inferior fit, compared to alternative distributions, to their overnight money market data coming from the e-MID trading platform. We thus tend to conclude that the evidence on power laws is not (yet) strong enough to warrant their widespread use in policy simulations. From the study presented in this work we provide alternate distributions and corresponding parameters which can systematically and robustly capture the interbank network measures. We deem that those distributions represent a more realistic account of the interbank network structure than the widely used power laws. They could facilitate more realistic contagion modeling and provide more realistic estimates of interbank network measures.

Acknowledgments

This work was supported by the Research Foundation Flanders (FWO-Flanders), the Research Foundation of Ghent University (BOF) and the Bank of Finland (BOFIT).

References

- [1] Mark E.J. Newman, Power laws, Pareto distributions and Zipf's law, *Contemp. Phys.* 46 (5) (2005) 323–351.
- [2] Victor M. Yakovenko, J. Barkley Rosser Jr., Colloquium: statistical mechanics of money, wealth, and income, *Rev. Modern Phys.* 81 (4) (2009) 1703.
- [3] Xavier Gabaix, Power laws in economics and finance. Working paper 14299, NBER, 2008.
- [4] Rama Cont, Amal Moussa, Edson B. Santos, et al. Network structure and systemic risk in banking systems, 2013.
- [5] Michael Boss, Helmut Elsinger, Martin Summer, Stefan Thurner, Network topology of the interbank market, *Quant. Finance* 4 (6) (2004) 677–684. 4.
- [6] Daniel O. Cajueiro, Benjamin M. Tabak, The role of banks in the brazilian interbank market: does bank type matter? *Physica A* 387 (27) (2008) 6825–6836.
- [7] Masayasu Kanno, Assessing systemic risk based on interbank exposures in the japanese banking system, 2014. Available at SSRN 2482075.
- [8] John Goddard, Hong Liu, Donal McKillop, John O.S. Wilson, The size distribution of us banks and credit unions, *Int. J. Econ. Bus.* 21 (1) (2014) 139–156.
- [9] Giulia De Masi, Giulia Iori, Guido Caldarelli, Fitness model for the italian interbank money market, *Phys. Rev. E* 74 (6) (2006) 066112.
- [10] Giulia Iori, Roberto Reno, Giulia De Masi, Guido Caldarelli, Trading strategies in the italian interbank market, *Physica A* 376 (2007) 467–479.
- [11] Aaron Clauset, Cosma Rohilla Shalizi, Mark E.J. Newman, Power-law distributions in empirical data, *SIAM Rev.* 51 (4) (2009) 661–703.
- [12] Michael P.H. Stumpf, Mason A. Porter, Critical truths about power laws, *Science* 335 (6069) (2012) 665–666.
- [13] A.G. Haldane, Rethinking the financial network, Apr 2009. Speech delivered at the Financial Student Association, Amsterdam.
- [14] M. Buchanan, Meltdown modelling, *Nature* 460 (7256) (2009) 680–682.
- [15] Andreas Krause, Simone Giansante, Interbank lending and the spread of bank failures: a network model of systemic risk, *J. Econ. Behav. Organ.* 83 (3) (2012) 583–608.
- [16] Tarik Roukny, Hugues Bersini, Hugues Pirotte, Guido Caldarelli, Stefano Battiston, Default cascades in complex networks: topology and systemic risk, *Sci. Rep.* 3 (2013).
- [17] Grzegorz Halaj, Christoffer Kok, Assessing interbank contagion using simulated networks, *Comput. Manag. Sci.* 10 (2–3) (2013) 157–186.
- [18] Co-Pierre Georg, The effect of the interbank network structure on contagion and common shocks, *J. Bank. Finance* 37 (7) (2013) 2216–2228.
- [19] Alexei Karas, Koen Schoors, Heracles or sisyphus? Finding, cleaning and reconstructing a database of russian banks. Working paper 327, Ugent, 2005.
- [20] Tarik Roukny, Co-Pierre Georg, Stefano Battiston, A network analysis of the evolution of the german interbank market. Technical Report, Discussion Paper, Deutsche Bundesbank, 2014.
- [21] Ben Craig, Goetz Von Peter, Interbank tiering and money center banks, *J. Financ. Intermed.* 23 (3) (2014) 322–347.
- [22] Serafin Martinez-Jaramillo, Biliiana Alexandrova-Kabadjova, Bernardo Bravo-Benitez, Juan Pablo Solórzano-Margain, An empirical study of the mexican banking system's network and its implications for systemic risk, *J. Econom. Dynam. Control* 40 (2014) 242–265.
- [23] Giulia Iori, Giulia De Masi, Ovidiu Vasile Precup, Giampaolo Gabbi, Guido Caldarelli, A network analysis of the italian overnight money market, *J. Econom. Dynam. Control* 32 (1) (2008) 259–278.
- [24] Daniel Fricke, Thomas Lux, On the distribution of links in the interbank network: evidence from the e-mid overnight money market. Working paper 1819, IfW, 2013.
- [25] F. Kyriakopoulos, S. Thurner, C. Pühr, Stefan W. Schmitz, Network and eigenvalue analysis of financial transaction networks, *Eur. Phys. J. B* 71 (4) (2009) 523–531.
- [26] Kimmo Soramäki, Morten L. Bech, Jeffrey Arnold, Robert J. Glass, Walter E. Beyerle, The topology of interbank payment flows, *Physica A* 379 (1) (2007) 317–333.
- [27] Samantha Cook, Kimmo Soramäki, The global network of payment flows. SWIFT Institute Working Paper, 2014.
- [28] Kei Imakubo, Yutaka Soejima, The transaction network in Japan's interbank money markets, *Bank Jpn. Monet. Econ. Stud.* 28 (2010) 107–150.
- [29] Camelia Minoiu, Javier A. Reyes, A network analysis of global banking: 1978–2010, *J. Financ. Stab.* 9 (2) (2013) 168–184.
- [30] Xuqing Huang, Irena Vodenska, Shlomo Havlin, H. Eugene Stanley, Cascading failures in bi-partite graphs: model for systemic risk propagation, *Sci. Rep.* 3 (2013).
- [31] Prasanna Gai, Andrew Haldane, Sujit Kapadia, Complexity, concentration and contagion, *J. Monetary Econ.* 58 (5) (2011) 453–470.
- [32] Erlend Nier, Jing Yang, Tanju Yorulmazer, Amadeo Aleontos, Network models and financial stability, *J. Econom. Dynam. Control* 31 (6) (2007) 2033–2060.
- [33] Stefano Battiston, Domenico Delli Gatti, Mauro Gallegati, Bruce Greenwald, Joseph E. Stiglitz, Default cascades: when does risk diversification increase stability? *J. Financ. Stab.* 8 (3) (2012) 138–149.
- [34] Michael Gofman, Efficiency and stability of a financial architecture with too interconnected to fail institutions, 2012. Available at SSRN.
- [35] Tiago P. Peixoto, Hierarchical block structures and high-resolution model selection in large networks, *Phys. Rev. X* 4 (1) (2014) 011047.
- [36] Benjamin M. Tabak, Marcelo Takami, Jadson Rocha, Daniel O. Cajueiro, Sergio R.S. Souza, Directed clustering coefficient as a measure of systemic risk in complex banking networks, *Physica A* 394 (2014) 211–216.
- [37] Duncan J. Watts, Steven H. Strogatz, Collective dynamics of ‘small-world’ networks, *Nature* 393 (6684) (1998) 440–442.

- [38] Jari Saramäki, Mikko Kivelä, Jukka-Pekka Onnela, Kimmo Kaski, Janos Kertesz, Generalizations of the clustering coefficient to weighted complex networks, *Phys. Rev. E* 75 (2) (2007) 027105.
- [39] M.E.J. Newman, The structure and function of complex networks, *SIAM Rev.* 45 (2) (2003) 167–256.
- [40] Jeff Alstott, Ed. Bullmore, Dietmar Plenz, Powerlaw: a python package for analysis of heavy-tailed distributions, *PLoS One* 9 (1) (2014) e85777.
- [41] Jean Laherrere, Didier Sornette, Stretched exponential distributions in nature and economy: 'fat tails' with characteristic scales, *Eur. Phys. J. B* 2 (4) (1998) 525–539.
- [42] Karl Finger, Daniel Fricke, Thomas Lux, Network analysis of the e-mid overnight money market: the informational value of different aggregation levels for intrinsic dynamic processes, *Comput. Manag. Sci.* 10 (2–3) (2013) 187–211.