Bargigli, Leonardo; Gallegati, Mauro

Working Paper

Finding communities in credit networks

Economics Discussion Papers, No. 2012-41

Provided in Cooperation with:
Kiel Institute for the World Economy – Leibniz Center for Research on Global Economic Challenges

Suggested Citation: Bargigli, Leonardo; Gallegati, Mauro (2012) : Finding communities in credit networks, Economics Discussion Papers, No. 2012-41, Kiel Institute for the World Economy (IfW), Kiel

This Version is available at:
https://hdl.handle.net/10419/62119

Economics
The Open-Access, Open-Assessment E-Journal

Discussion Paper
No. 2012-41 | August 30, 2012 | http://www.economics-ejournal.org/economics/discussionpapers/2012-41

Finding Communities in Credit Networks

Leonardo Bargigli and Mauro Gallegati
Università Politecnica delle Marche

Abstract In this paper the authors focus on credit connections as a potential source of systemic risk. In particular, they seek to answer the following question: how do we find densely connected subsets of nodes within a credit network? The question is relevant for policy, since these subsets are likely to channel any shock affecting the network. As it turns out, a reliable answer can be obtained with the aid of complex network theory. In particular, the authors show how it is possible to take advantage of the 'community detection' network literature. The proposed answer entails two subsequent steps. Firstly, the authors need to verify the hypothesis that the network under study truly has communities. Secondly, they need to devise a reliable algorithm to find those communities. In order to be sure that a given algorithm works, they need to test it over a sample of random benchmark networks with known communities. To overcome the limitation of existing benchmarks, the authors introduce a new model and test alternative algorithms, obtaining very good results with an adapted spectral decomposition method. To illustrate this method they provide a community description of the Japanese bank-firm credit network, getting evidence of a strengthening of communities over time and finding support for the well-known Japanese 'main bank' system. Thus, the authors find comfort both from simulations and from real data on the possibility to apply community detection methods to credit markets. They believe that this method can fruitfully complement the study of contagious defaults, since the likelihood of intracommunity default contagion is expected to be high.

Paper submitted to the special issue Coping with Systemic Risk

JEL C49, C63, D85, E51, G21

Keywords *C*redit networks; communities; contagion; systemic risk

Correspondence Leonardo Bargigli, Università Politecnica delle Marche, Department of Economics and Social Sciences, Ancona, Italy; email: leonardo.bargigli@gmail.com

The authors acknowledge the financial support from the European Community Seventh Framework Programme (FP7/2007-2013) under Socio-economic Sciences and Humanities, grant agreement no. 255987 (FOC-II). They are especially thankful to Hideaki Aoyama and Yoshi Fujiwara for sharing Japanese network data.

1 Introduction

Since the outbreak of the global crisis, policy makers have been haunted by the nightmare of a global financial meltdown breaking out of incontrollable feedbacks spreading across financial markets. In order to avoid this scenario, it became imperative to identify the potential sources of system-wide shocks, i.e. the sources of systemic risk.

Although many transmission channels for systemic risk have been listed by the literature (Brunnermeier, 2009), in this paper we will focus on those determined by the reciprocal claims between institutions over credit markets. The interconnectedness of credit institutions is a source of counterparty risk on interbank credit markets, which has been addressed recently by a number of theoretical models tackling the problem of contagious defaults (Gai and Kapadia, 2010; Amini et al., 2010, 2012; Battiston et al., 2012). These models, which go beyond previous simulation based works (Nier et al., 2007; Elsinger et al., 2006), rely on complex network theory, which has become a prominent tool in this field. In fact, network theory provides a rigorous representation of complex interactions, as well as the possibility to predict their static and dynamic behavior in many cases[1].

Regarding the case in point, it's natural to conceive credit markets as networks in which nodes represent agents and links represent credit claims and liabilities. Most works in this field focus specifically on the interbank market, since the latter is particularly relevant for financial stability and, at the same time, well suited for a representation with basic network theory. While earlier contributions (Allen and Gale, 2000) stressed the benefits of increasing diversification, suggesting that the more connections, the better for financial stability, more recently a growing number of works have challenged this view, showing that diversification is not always beneficial for stability. For instance, the recent model of Battiston et al. (2012) shows that, if market-related effects are considered along with credit-related effects by introducing a financial accelerator mechanism, then a potential trade-off between individual risk and systemic risk may exist for increasing connectivity of the network. Similar results are provided by Gai and Kapadia (2010), who show that financial systems exhibit a robust-yet-fragile tendency: while the probability of contagion may be low, once a default cascade is started its spread may be quite large. This effect is non monotonic in connectivity: for a given range of values, connectivity increases the chances that institutions surviving the effects of the initial default will be exposed to more than one defaulting counterpart after the first round of contagion, thus making them

[1]For an economics oriented introduction see Vega-Redondo (2007) or Jackson (2008).

more vulnerable to a second-round default.

The empirical support for the relevance of contagious defaults in the interbank market is mixed. This is not surprising at all, since empirical works in this field rely on a variety of simulation based approaches and diverse behavioral assumptions [2]. For instance, those works which examine the effects of idiosyncratic shocks affecting a single bank, come to the conclusion that the scope of contagion is limited (Elsinger et al., 2006; Upper and Worms, 2004; Mistrulli, 2011). By adopting a more realistic setting, e.g. taking into account correlated market shocks and short-term 100% losses for creditors, quite different results have been obtained (Cont et al., 2012). Notwithstanding this uncertainty, central banks are getting more and more interested in network analysis, supporting network-related research and dissemination, although most empirical work in this direction still looks merely descriptive (Castrén and Kavonius, 2009; ECB, 2010).

In this paper, instead of dealing directly with contagious defaults, we wish to answer the following related question: how do we find densely connected subsets of nodes within a given credit network? As it turns out, we can provide a rigorous answer by drawing from a growing body of literature devoted to community detection in complex networks (Fortunato, 2010). This answer is very important for any analysis of contagious defaults. In fact, the dynamics of any contagion process depends crucially on the topology of the network at the moment of the initial shock. This fact agrees with the following simple intuition: whenever a shock affects a node of a financial network, it will be transmitted to her neighbors with a probability that is proportional to the strength of their linkage to the shocked node. Thus, finding strongly connected subsets of nodes allows us to identify those regions of the network which are most likely to be hit when a specific subset of nodes is shocked initially. In this sense, community detection allows to outline the general "risk maps" invoked by policy makers.

The rest of the paper is organized as follows. As a first step, we shortly review the related literature on complex networks and community detection (sec. 2). Subsequently (sec. 3) we focus on a benchmark model which displays modular structure, i.e. it is endowed with significant communities as defined at the end of sec. 2. Then we test alternative community detection algorithms over the benchmark model examining both the case of strong, i.e. non overlapping, and weak, overlapping, communities (sec. 4). After having verified that a spectral decomposition algorithm provides the best results, we apply the same algorithm to real data in order to provide a community description of the Japanese bank-firm credit network (sec. 5). Finally, the

[2] For a survey see Upper (2011).

conclusive section summarizes the main results of the paper.

2 Related literature

In this paper we adopt a statistical equilibrium approach to the representation of credit markets [3]. Statistical means that the claim w_{ij} of agent i towards agent j is viewed as the realization of a random variable, defined over a discrete nonnegative domain. The market as a whole is nothing more than the collection of all these variables, that can be represented as a random matrix W with entries which are statistically independent but non necessarily equally distributed. Each realization of W represents a possible state or configuration of the market, and the collection of all these market states, together with a probability distribution over states $P(W)$, is called a *statistical ensemble*. Equilibrium means instead that, if the market is allowed to relax without external disturbances, it will converge to the stable probability distribution $P^*(W)$ which is obtained by solving the model itself[4]. As explained below, when the model is adequately constrained, it's possible to make the ensemble reflect some desirable economic property, like the expected balancing of supply and demand (see note 7).

Against this backdrop, the transition to complex networks theory appears very natural. In fact, the representation of markets outlined above is nothing different from the matrix representation of a directed network G, according to which the strength of the link $i \to j$ is given by w_{ij} [5]. In particular, while the interbank market is a simple (weighted and directed) network, credit relationships between banks and firms are represented by a bipartite (weighted and directed) network, i.e. a network whose nodes can be divided into two disjoint sets F and B such that every link connects an element of F with an element of B. The usual matrix representation of a bipartite network is given by a rectangular matrix $W \in \mathbb{N}_+^{n \times m}$, where $n = |F|$ and $m = |B|$. The analysis of these networks is addressed either with specific tools or by generalizing the standard tools in various ways. In this paper we

[3]For the general idea of using statistical equilibrium as a key tool in macroeconomics see Aoki and Yoshikawa (2006)

[4]There is a bijective relationship between a given model and the corresponding ensemble, since solving a given model means actually to find the particular P^* which is consistent with the constraints of that model. For this reason the terms "model" and "market ensemble", "network ensemble" or simply "ensemble" can be used as equivalents.

[5]For this reason, in this paper the terms "market" and "network" are used as equivalents even if, strictly speaking, a market is a directed weighted network. If G instead is a binary network, its links can take only binary values, and thus its matrix representation is given by the adjacency matrix A with binary entries.

follow mainly the latter route, because we wish to provide a general method which may be readily applied both to interbank networks and to bank-firms credit networks. Indeed, most of the tools we provide can be applied readily to any interaction occurring among economic agents or entities, provided that this interaction can be represented by a nonnegative matrix of arbitrary shape.

As mentioned above, we are interested in network ensembles displaying desirable properties. Park and Newman (2004) have proposed a general methodology for building ensembles of networks, with a fixed number n of nodes, satisfying linear and non linear constraints over the expected values of network observables. We may translate these constraints into equivalent ones defined over the strength matrix W:

$$E\left[\vec{F}(W)\right] = \vec{c} \tag{2.1}$$

where $\vec{c}$ is a vector of arbitrary values for the constrained observables. Park and Newman (2004) have showed, in analogy with equilibrium statistical physics, that P^* is a Boltzmann-Gibbs probability distribution over network configurations. Since the Boltzmann-Gibbs distribution belongs to the exponential family of probability distributions, networks belonging in the resulting ensemble are labeled as *exponential networks*. In particular, they provide a solution of this model when the constraints are represented by the strength distribution $\vec{w}$ of the network [6]. Given our previous market description, strengths represent in credit networks the expected lending and borrowing of the agents. By allowing for arbitrary strengths, this approach provides a decisive improvement with respect to current economic models, which usually adopt the unrealistic assumption of agents with identical strengths (Allen and Gale, 2000), introducing instead heterogeneity in the modeling of market interactions [7].

Exponential networks don't represent the only ensemble which may comply with the heterogeneity of strengths. In the already mentioned community detection literature, a different model is usually employed as statistical

[6]By strength of a node i in a weighted network we define the sum $w_i = \sum_{j \neq i} w_{ij}$. If W is asymmetric, i.e. G is directed, we need to distinguish between the out-strength and in-strength of the node i, and thus we have two distinct distributions $\vec{w}_{out}$ and $\vec{w}_{in}$. The degree of a node i, instead, is defined over the binary matrix A as the sum $d_i = \sum_{j \neq i} a_{ij}$. If A is asymmetric, again we need to distinguish between the out-degree and in-degree of the node i. In the case of weighted networks, it's useful to define $v = \sum_i w_i$, where v is said to be the *volume* of G.

[7]Incidentally, we observe that heterogeneity doesn't necessarily imply disequilibrium in this context. In fact, for any two out- and in-strength distributions we have that $\sum_i w_i^{out} = \sum_i w_i^{in} = v$, i.e. that the market is expected to clear.

benchmark. This model is related to the modularity function Q, originally introduced by Newman and Girvan (2004) in order to provide a stopping criterion for their community detection algorithm [8].

In order to introduce this model, it is convenient to adjust the definition of Q. We employ the normalized matrix $K^* = K - E[K]$ instead of the modularity matrix $B = W - E[W]$ as defined by Newman and Girvan, since the former displays some useful properties [9]. Thus, our definition of modularity reads

$$Q(W) \propto \sum_{ij} k^*_{ij} \delta(h_i, h_j) \tag{2.2}$$

where h_i, h_j denote community membership.

In general, we know that $W = \sum_{i=0}^{r-1} \sigma_i u_i v_i^T$ for $W \in \mathbb{R}^{n \times m}$, where the σ_i are the (decreasing) singular values (SV) of W, and the u_i and v_i are its left and right singular vectors. In particular, we know that $\sigma_0(K) = 1$ for all K and consequently that $u_0^T = (\sqrt{w_1^{out}/v}, \ldots \sqrt{w_n^{out}/v})$ and $v_0 = (\sqrt{w_1^{in}/v}, \ldots \sqrt{w_m^{in}/v})$ (Bargigli and Gallegati, 2011). It's easy to see that the matrix K^* is the normalized counterpart of B. In fact, we have that $B = W - E[W] = D_{out}^{\frac{1}{2}} \left[K - u_0 v_0^T\right] D_{in}^{\frac{1}{2}}$, since the elements of $E[W]$ are defined, in the directed case, as follows

$$E[w_{ij}] = \frac{w_i^{out} w_j^{in}}{v} \tag{2.3}$$

Squartini and Garlaschelli (2011) have showed that this expectation cannot be derived from exponential networks. For this purpose we need to introduce a different ensemble which can be labeled as *binomial networks*. In this ensemble the w_{ij} are binomially distributed with parameters v (see note 6) and p_{ij} (Bargigli and Gallegati, 2011). In the directed case, the parameters p_{ij} are obtained by solving the following maximum entropy problem:

$$\max_{\mathbf{p}} g(\mathbf{p}) = -\sum_{i=1}^{n} \sum_{j=1}^{m} p_{ij} \ln p_{ij} \tag{2.4}$$

subject to the following constraints:

[8] For further details see below, note 12.

[9] Normalization is obtained, e.g. for a symmetric weighted network, by introducing the normalized matrix $K = D^{-\frac{1}{2}} W D^{-\frac{1}{2}}$ where D is a diagonal matrix with elements $\{w_1, w_2, \ldots, w_n\}$. In case of directed networks, normalization becomes $K = D_{out}^{-\frac{1}{2}} W D_{in}^{-\frac{1}{2}}$ with obvious adaptation of the diagonal matrices.

$$\sum_{j=1}^{m} p_{ij} = r_i$$
$$\sum_{i=1}^{n} p_{ij} = c_j$$
$$p_{ij} > 0$$

where $r_i = w_i^{out}/v$, $c_j = w_j^{in}/v$.

Using the linear constraints of the problem, we can obtain the following explicit solution (Bargigli and Gallegati, 2011):

$$p_{ij} = \frac{w_i^{out} w_j^{in}}{v^2} \tag{2.5}$$

from which we get the required expectation (2.3).

The fundamental property of binomial networks is that their expectation $E[W]$ is a rank-1 matrix, since its rows and columns are all linearly dependent. Thanks to this property, the model defines a "community free" expectation of a (weighted and directed) network with given strength distribution. In fact, it's easy to see from eq. (2.2) that $Q\left(E[W]\right) \equiv 0$. On the other hand, it is also true that $Q\left(E[W]\right) \neq E[Q(W)] > 0$, since networks in the ensemble will be obviously different from $E[W]$. We may systematize these arguments with the following three definitions:

- $E[W]$ is the *community free network*
- the communities of all networks $G \in \mathcal{G}$, where $\mathcal{G}$ stands for the binomial ensemble, are *statistically non significant*, since they are the outcome of a random fluctuation around $E[W]$.
- the communities of networks $G \notin \mathcal{G}$ are *statistically significant*.

Thus, to be sure that we detect significant communities, we need to verify the null hypothesis $H_0 : G \in \mathcal{G}$. In the next section, we will show how it is possible to obtain this result.

Finally, we may wonder why we are so interested in an expectation of the form (2.3). From the very nature of the ME problem (2.4) we see that binomial networks provide the maximally diversified model which satisfy the constraint over strength distributions. From a systemic risk perspective, this is convenient since it allows us to focus on the specific network risk *conditioned to* a given distribution of debt and credit within the network. By

comparing a given real network with a sample of binomial networks with the same expected strength distributions, we are able to detect the risk in excess of the risk level observed on a maximally diversified network. The strength distributions represent an independent source of risk, which needs to be clearly separated from *network risk* itself. Although a binomial network contains no risky exposures by construction, it may be nonetheless very fragile because of an excessive level of lending or borrowing. Thus, we may define *network risk* as the default rate in excess of the default rate which may be observed on a sample of binomial networks [10].

3 The Modular Binomial Model

Artificial networks displaying a modular structure represent a very precious tool for testing community detection algorithms (Lancichinetti and Fortunato, 2009a). In these networks the assignment of nodes to communities is specified in advance by the researcher and the algorithm's task becomes to recover the known community subdivision of agents in the network. Then we can compare how good are alternative algorithms at this job with the help of some suitable measure, like normalized mutual information (see below).

We can induce a modular structure in the binomial network of the previous section by fixing the weights between agents belonging to different arbitrary defined communities. Instead of providing different weight values for each combination of communities, like in Bianconi (2009, 2008), we simplify by adding to the problem (2.4) the following constraint:

$$\sum_{ij}(1-\delta_{ij})p_{ij} = s \tag{3.1}$$

where $\delta_{ij} = 1$ if the two agents belong to the same community. The matrix Δ with elements δ_{ij} is a *co-membership matrix* which can be obtained by self-multiplication of the community matrix $M = (m_0|m_2|\dots|m_{c-1}) \in \{0,1\}^{n\times q}$, where q stands for the number of communities and we have $\forall i$ that $\sum_j m_{ij} = 1$, i.e. nodes belong to one community only. By solving for p_{ij} in the first order conditions we obtain

$$p_{ij} = \exp(\lambda_i + \mu_j + \theta(1-\delta_{ij})) = x_i y_j z^{1-\delta_{ij}} \tag{3.2}$$

[10] As it turns out, binomial networks are usually found to underestimate the extent of contagion (van Lelyveld and Liedorp, 2006; Mistrulli, 2011), which means that real networks are indeed an independent source of risk in the sense just explained. But it is also conceivable that real networks are less risky than binomial networks in some cases.

where λ_i, μ_j and θ are Lagrange multipliers. From this expression we derive a system of $n + m + 1$ coupled equations

$$\begin{aligned} x_i \sum_j y_j z^{1-\delta_{ij}} &= r_i && \forall i = 1 \dots n \\ y_j \sum_i x_i z^{1-\delta_{ij}} &= c_j && \forall j = 1 \dots m \\ \sum_{ij} (1-\delta_{ij}) x_i y_j z^{1-\delta_{ij}} &= s \end{aligned}$$

which, once solved numerically, delivers what we call the modular binomial (MB) model. In practice, it is convenient to specify s as follows

$$s = \mu \sum_{ij} (1-\delta_{ij}) \bar{p}_{ij} \tag{3.3}$$

with $\mu \in (0, 1]$ and $\bar{p}_{ij}$ specified as in eq. (2.5). In this way we assign to inter-modular connections a fraction μ of the ME probability obtained in the binomial model. Thus, it becomes possible to produce a continuum of networks ensembles with increasingly significant modular structure, ranging from a (significant) community-free network coincident with the binomial model ($\mu = 1$) to a quasi-disconnected union of (significant) communities ($\mu = \epsilon$ for a small ϵ). The results of this procedure for a sample of 100 randomly generated MB networks with $q = 10$, for each value of μ and different size distributions of communities, are depicted in Fig. 1. While agents are randomly assigned to communities, the strength sequence used for the simulations is taken from real data, namely from the Japanese credit network data analyzed below (see section 5). In particular, we employ the strength distribution of the 211 Japanese banks recorded during the year 2000.

The lines of Fig. 1 represent the average largest singular values of $K \in \mathbb{R}_+^{211\times 211}$ for these samples [11]. As underlined in many contributions (e.g. Chauhan et al. (2009) or Mitrovi and Tadi (2009)), the spectral, and SV related, properties of different network-related matrices are connected to the modular structure of the network itself. In particular, the spectral or singular gap, i.e. the largest difference between eigenvalues or SVs, clearly identifies the number of communities (Ng et al,, 2001). We can verify from Fig. (1) that in our samples the number of non unitary SV which are bounded away

[11] Regarding the statistics collected from MB simulations, we stress that we observed always a low variability within samples, resulting in very small standard errors. Thus we don't need to build larger samples in order to obtain robust results.

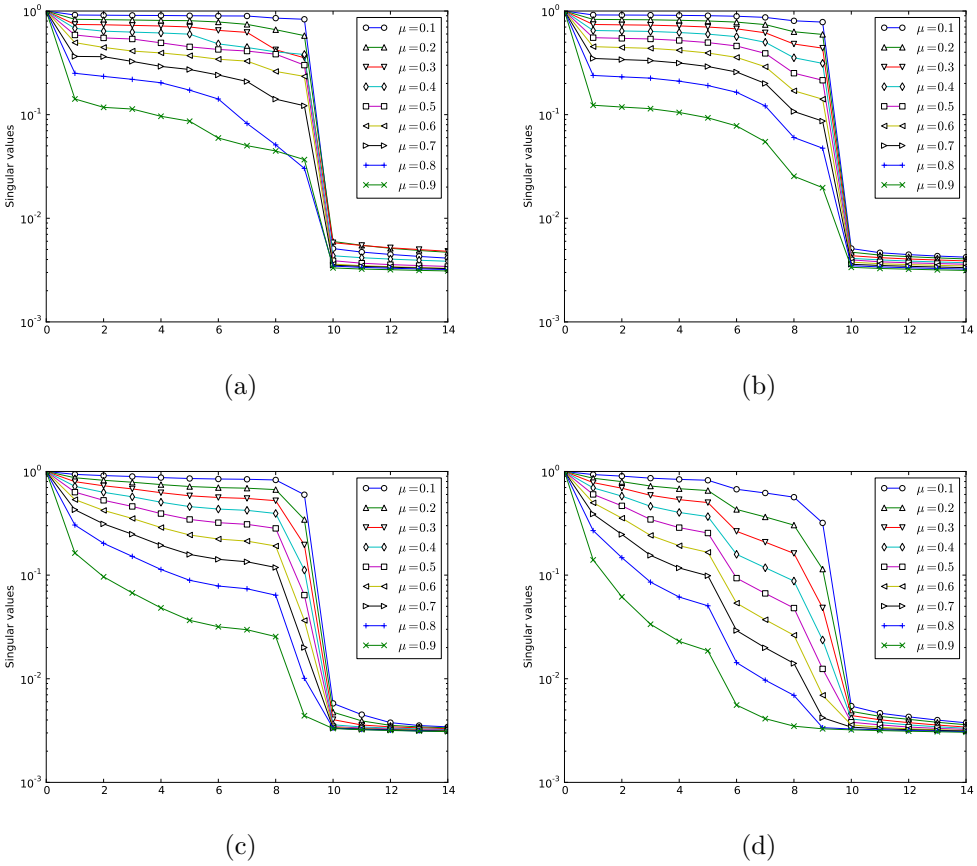

Figure 1: Average SVs of K for randomly generated MB networks with 10 communities of different size distributions: (a) uniform; (b) power law ($\alpha = 1.4$) with sizes ranging from 10 to 50 nodes; (c) power law ($\alpha = 1.4$) with sizes ranging from 5 to 100 nodes; (d) power law ($\alpha = 1.4$) with unconstrained size

Table 1: p-values for the null hypothesis that MB networks belong to the binomial ensemble

μ	max p-value	min p-value	μ	max p-value	min p-value
0.1	0.0010	0.0010	0.91	0.0959	0.0954
0.2	0.0012	0.0012	0.92	0.1211	0.1205
0.3	0.0014	0.0014	0.93	0.1580	0.1569
0.4	0.0020	0.0020	0.94	0.2142	0.2127
0.5	0.0030	0.0030	0.95	0.3066	0.3035
0.6	0.0047	0.0047	0.96	0.4729	0.4672
0.7	0.0085	0.0085	0.97	0.8156	0.8050
0.8	0.0193	0.0193	0.98	0.9999	0.9999
0.9	0.0777	0.0773	0.99	0.9999	0.9999

from zero is equal to $q - 1$, as expected, although q may become harder to detect when the size distribution of communities becomes more uneven (see subfigs (c) and (d)), since in this case we have very small communities which are more difficult to identify.

It is instructive to verify the null hypothesis $H_0 : G \in \mathcal{G}$ for MB networks (see end of sec. 2). Of course, we expect H_0 to be rejected if μ is not too close to unity. We can employ for this task the statistical test devised in Bargigli and Gallegati (2011):

$$P\left(\Sigma^2(K) \geqslant \lambda\right) \leqslant \frac{nm}{\lambda v} \tag{3.4}$$

for $K \in \mathbb{R}_+^{n \times m}$, where v is the volume of the network (see note 6) and $\Sigma^2(K) = \sum_{ij}(k_{ij} - E[k_{ij}])^2$. Following the definitions of sec. 2, if H_0 is rejected the network contains statistically significant communities. In order to compute the test of eq. (3.4), we need to know the singular values of K. In fact, from the Eckart-Young low rank approximation theorem we derive that $\Sigma^2(K) = \sum_{k=1}^{r-1} \sigma_k^2$, where the sigmas now stand for the SVs of K and r is the rank of K. In tab. (1) we reproduce the maximal and minimal p-value obtained on samples of 100 networks, for each value of μ, using eq. (3.4). The p-values are quite stable within samples, and increasing in μ as expected. In particular, the null hypothesis is clearly rejected for $\mu \leqslant 0.9$.

We can employ the same test in order to recover the number of communities in MB networks in a more rigorous way than by visual inspection. The basic idea is simple: if we order the singular modes of K by the decreasing magnitude of the associated SVs (excluding the largest one), and subtract from K iteratively those singular modes, we let K converge to its expectation under the binomial model. Then, if we compute the test (3.4) at each

step and record the resulting p-value, at some step we obtain that the null hypothesis is not rejected anymore, which means that all the significant communities have been subtracted away from the network. In practice, it turns out that q is best detected by observing the increments of the p-values, since the maximum of the increments is associated with the singular gap. As we see from fig. 2, this method works pretty well when the modular structure is clear enough (2a and 2b) and for $\mu \leqslant 0.5$.

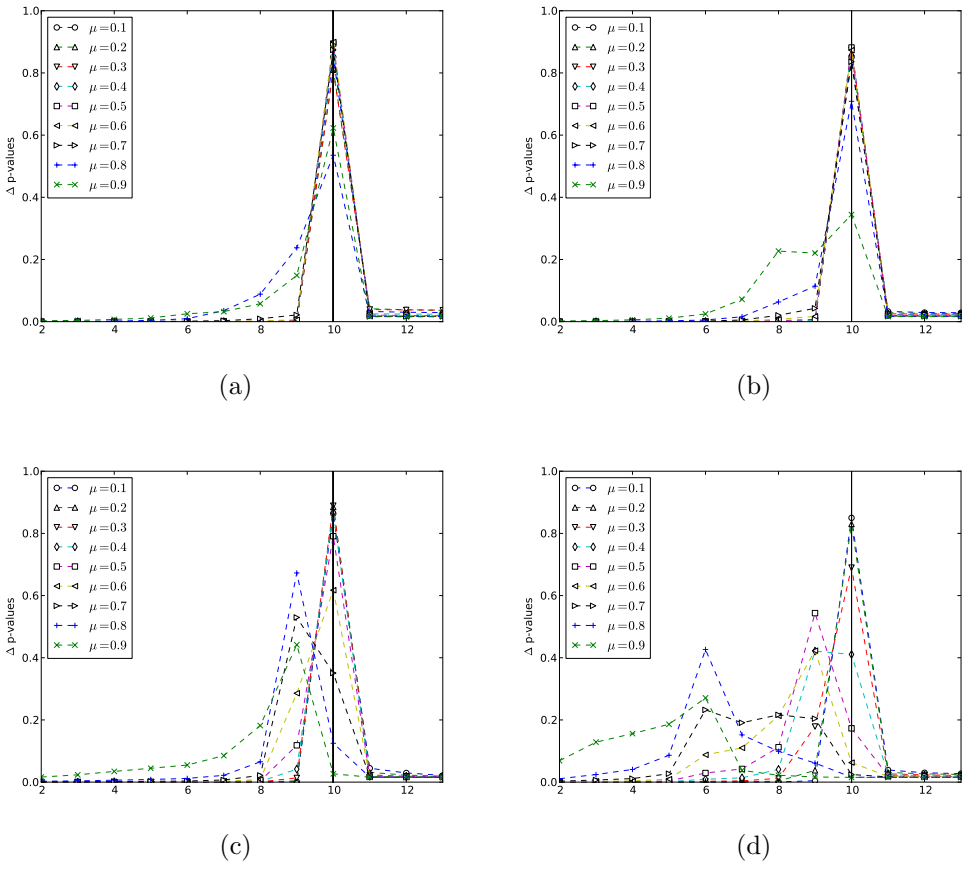

(a) (b)

(c) (d)

Figure 2: Average Δ p-values computed from (3.4) for MB networks with 10 communities of different size distributions: (a) uniform; (b) power law ($\alpha = 1.4$) with sizes ranging from 10 to 50 nodes; (c) power law ($\alpha = 1.4$) with sizes ranging from 5 to 100 nodes; (d) power law ($\alpha = 1.4$) with unconstrained size. The true number of communities is signaled by the black vertical line.

As it turns out from the analysis of real data (see section 5), the constraint of non-overlapping community membership of nodes, introduced above, may not describe adequately real credit networks. By relaxing this assumption,

we introduce *weak* (i.e. overlapping) communities instead of *strong* (i.e. non overlapping) communities. In practice, we obtain the latter by allowing $\sum_j m_{ij} > 1$ for some i in our ME problem. From fig. 3, we see that the behavior of simulated networks changes drastically: as the fraction β of nodes belonging to 2 different communities increases, the modular structure of the network becomes less clear.

From a geometrical point of view, the reason is intuitive. With strong communities the rows of M, representing nodes, may be divided in orthogonal subsets, corresponding to the communities, which give rise to clearly separated singular modes. With weak communities orthogonality is violated, and the stronger this violation the harder it is to recover clearly separated singular modes. In the limit case when all nodes belong to all communities, we simply recover binomial networks, which don't display significant singular modes. Thus, we may conclude that weak communities can cohabit with a clear modular structure only if the fraction of overlapping nodes is low. Not surprisingly, the more blurred modular structure of weak communities makes it harder to recover q with the help of eq. (3.4). From fig. 4 we see that q is correctly recovered only for low overlapping rates (4c and 4d).

4 Test of community detection algorithms

As underlined above, the availability of artificial modular benchmarks is of utmost importance for developing reliable community detection methods. In fact, given the complexity of the task, we are never sure of the result we obtain with a given procedure unless we can trace back this result to a previously known correct answer (Fortunato, 2010). From this perspective, MB networks are particularly effective because they deliver the same features of existing benchmarks (Lancichinetti and Fortunato, 2009a) by allowing for weighted, directed and even bipartite networks, as well as for both strong and weak communities (see sec. 3), while they don't require to make hypotheses on the strength distributions. Thus, the confidence in the tests performed on MB networks is bolstered by the fact that we build artificial modular networks which display important properties, such as strength distributions, that are equal on average to those of some real network under study.

In this section we don't aim to perform a systematic comparison of the main community detection procedures, since such comparison is already available in the literature (Lancichinetti and Fortunato, 2009b). Instead, we will focus on two different algorithms. The first one is Infomap (Rosvall and Bergstrom, 2008), which is widely held as the most reliable method since modularity-based methods have been challenged by the discovery of a

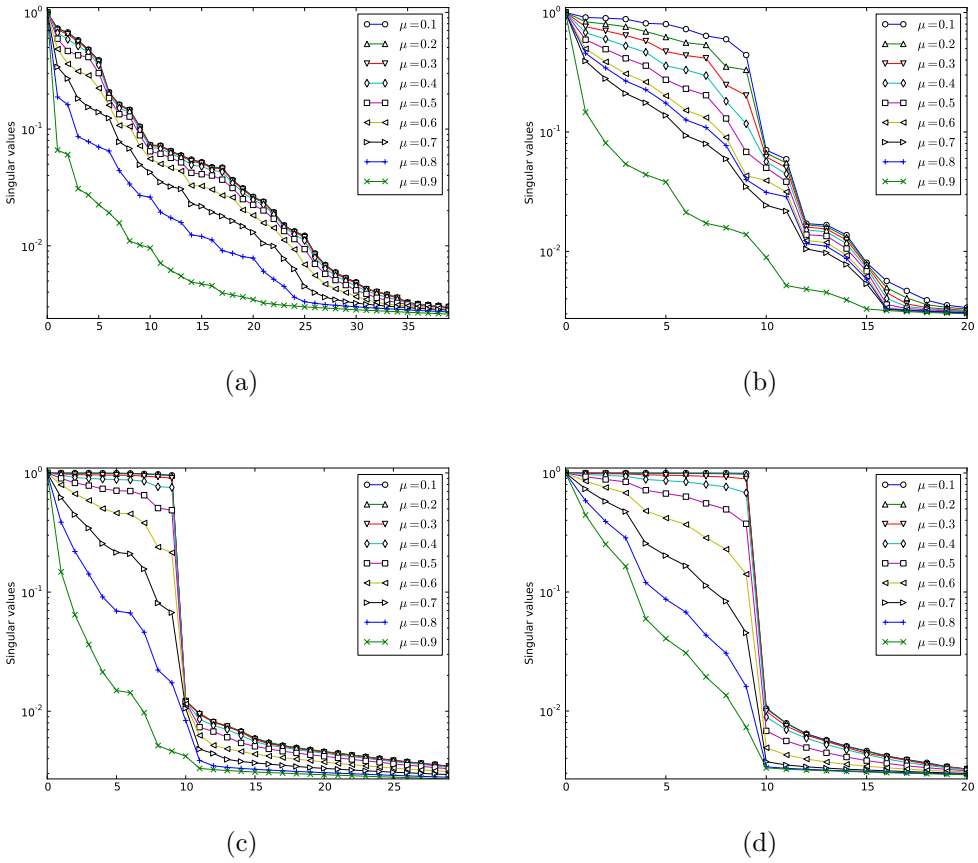

Figure 3: Average normalized singular values of MB networks with 10 bounded (max: 100; min: 5) power-law distributed communities for different overlapping rates: (a) $\beta = 1$; (b) $\beta = 0.1$; (c) $\beta = 0.05$; (d) $\beta = 0.01$

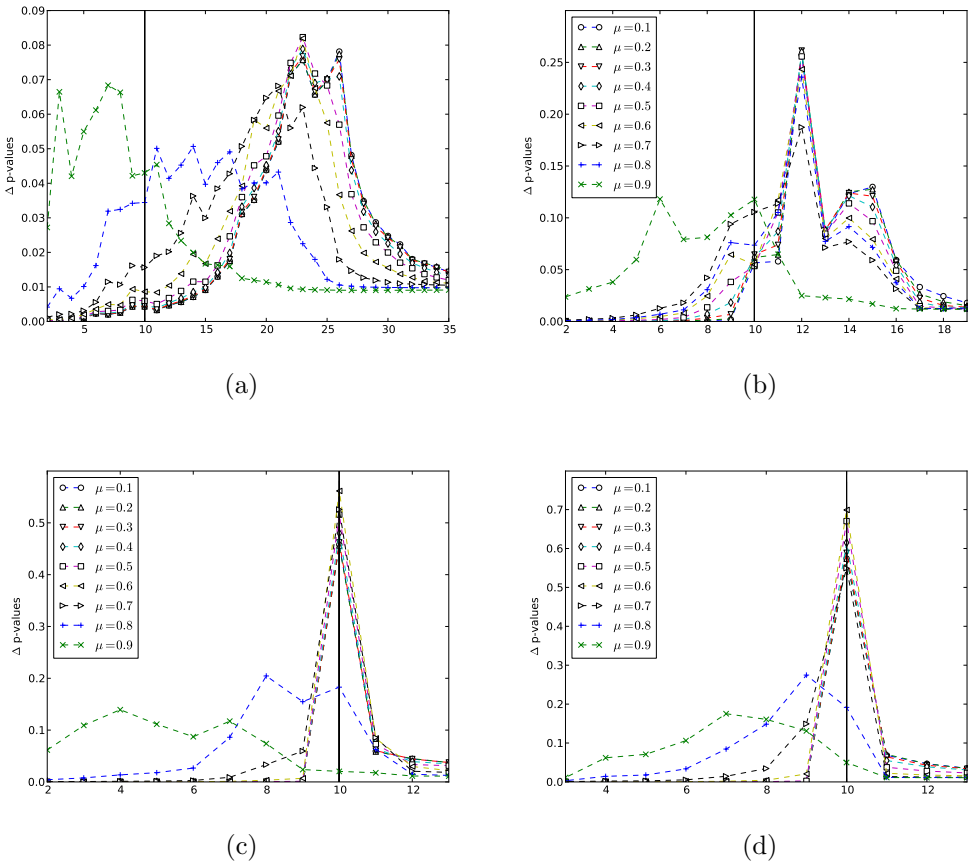

Figure 4: Average Δ p-values computed from (3.4) for MB networks with 10 bounded (max: 100; min: 5) power-law distributed communities for different overlapping rates: (a) $\beta = 1$; (b) $\beta = 0.1$; (c) $\beta = 0.05$; (d) $\beta = 0.01$

resolution bias [12]. The second one is an adaptation of the spectral method proposed by Donetti and Muñoz (2005).

The main motivation to perform this comparison is that spectral methods can be applied also to bipartite networks, while this is not possible for Infomap. In fact, this method employs the properties of Markov chains by defining a random walk over a directed network by means of the stochastic matrix S obtained by suitably normalizing W. But the convergence of the random walk over a directed network requires the latter to be strongly connected, a condition which, by construction, cannot be fulfilled by bipartite networks. On the other hand, the available community detection methods for bipartite networks fail to deliver very good results or otherwise are subject to severe computational limits (Sawardecker et al., 2009), so that improvements are strongly needed in this area.

We complement these methods with the approach of statistically validated networks (SVNet) proposed by Tumminello et al. (2011). This approach evaluates the probability to observe a given strength value $w_{ij} = x$, e.g. in a symmetric network, against the hypergeometric probability distribution:

$$w_{ij} \sim H(x|w_i, w_j, v) = \frac{\binom{w_i}{x}\binom{v-w_i}{w_j-x}}{\binom{v}{w_j}} \tag{4.1}$$

The intuition is clear: the probability is obtained by enumerating the possible extraction sequences (without replacement) of length w_j in which we obtain exactly x successes out of w_i favorable cases contained in the urn with v elements. The probability of the event $\{w_{ij} \geqslant x\}$ is

$$P\{w_{ij} \geqslant x\} = 1 - \sum_{w_{ij}=0}^{x-1} H(x|w_i, w_j, v) \tag{4.2}$$

The link between i and j is validated only if it is highly unlikely when evaluated using eq. (4.2), i.e. when the null hypothesis that w_{ij} is distributed according to $H(x|w_i, w_j, v)$ is rejected. The link validation procedure involves multiple hypothesis testing. Therefore the statistical threshold must be corrected for multiple comparisons by employing Bonferroni correction, i. e. by

[12] Q was introduced with the double purpose to evaluate the assignment of nodes to a given number q of communities and to expand or shrink q itself (Fortunato, 2010). The maximization of Q has become a popular method of community detection, until Fortunato and Barthélemy (2007) have showed that this method is plagued by a resolution bias, i.e. it is unable to detect correctly small communities. To overcome the resolution bias many solutions have been proposed, none of which delivers an equivalent solution to the original problem that modularity was meant to solve.

setting the rejection threshold for the null hypothesis at $\frac{p}{T}$ where T is the number of hypotheses to be tested ($T = n^2$ in the case of a network with n nodes) and p is the chosen confidence level. Once the procedure is completed, we obtain a binary, possibly directed, validated network which can be useful to solve the community detection problem in two ways: firstly, by looking at the components of the validated network [13], which may reflect, at some extent, the underlying community assignment; secondly, by letting the algorithms work more efficiently than on the original network. Next we are going to pursue both strategies, in order to verify the contribution of this approach over MB networks.

By adopting the hypergeometric distribution we refer our expectation to the so-called configuration model (Molloy and Reed, 1995; Newman et al., 2001), in which the strength distributions are fixed exactly and not on average. In order to adapt the SVNet approach to our needs we only need to replace the hypergeometric distribution with the binomial distribution:

$$w_{ij} \sim B(x|p_{ij}, v) = \binom{v}{x} p_{ij}^{x}(1 - p_{ij})^{v-x} \tag{4.3}$$

If we evaluate the strengths of a MB network using eq. (4.3), we expect only within-community connections to be validated. Thus, as claimed above, the components of the valid network may reflect communities.

Regarding the spectral algorithm, we introduce the following two modifications with respect to Donetti and Muñoz (2005): 1. the number q of communities is determined by means of eq. (3.4) and not by means of modularity optimization; 2. spectral decomposition is performed over K and not over the Laplacian L. Except for these differences, we perform the same steps: 1. the eigenvectors associated to the q largest eigenvalues are computed; 2. each node in the network is treated like a point in the $q-1$ - dimensional space with coordinates given by the node's projections onto the first $q-1$ nontrivial eigenvectors; 3. hierarchical clustering is applied over the nodes' coordinates in order to obtain q clusters. In particular, we observed the best performance using complete-linkage clustering over angular distances.

To sum up, the methods employed in the tests are listed as follows: "Infomap" denotes the algorithm introduced in Rosvall and Bergstrom (2008); "Natural partition" denotes the partition corresponding to the components of the SVnet; "Spectral" denotes our adaptation of the algorithm introduced in Donetti and Muñoz (2005); the suffix "on SV net" denotes the fact that

[13] The components of a network are defined as its maximal connected subsets of nodes.

the (Infomap or spectral) algorithm is applied to the SVnet. In order to verify the effectiveness of all these algorithms, we adopt the widely used measure of *normalized mutual information* between two random variables (Fortunato, 2010):

$$I_{norm}(X,Y) = \frac{I(X,Y)}{\sqrt{|H(X)H(Y)|}} \tag{4.4}$$

where H stands for the entropy of the random variable and $I(X,Y) = H(X) - H(X|Y) = H(X) + H(Y) - H(X,Y)$. In our case, I_{norm} is used to compare the partition obtained employing some community detection algorithm with the one originally employed to solve the MB problem or, as we can say more shortly, to compare the *recovered* partition P with the *true* partition P^*. In practice, for each node the assignment to a community is considered as a random variable taking values in $\{0, 1, \ldots, q-1\}$. Then, if the true and recovered partition are equal, we obtain that $I_{norm}(P, P^*) = 1$.

The results of the tests performed over MB networks are depicted in fig. 5. The lines represent the average values of $I_{norm}(P, P^*)$ over samples of 100 artificial networks. Again, we remark that within-sample variability is very low, so that we don't need to build larger samples. Our results provide a unequivocal indication: the best performance is obtained with the spectral algorithm applied to the original network. Further, we observe that SVnets help improve significantly the performance of Infomap for low μ, while they generally worsen the performance of the spectral algorithm. Also, the natural partition is showed to represent faithfully the underlying community structure, except for uneven community size distributions (subfig. 5d).

In order to extend our testing framework to the case of weak communities, we need to complement our partition methods with a procedure to detect the overlapping nodes between two or more different communities. For this purpose we adapt the method proposed by Lancichinetti et al. (2009), which evaluates the probability that a given node external to a given community has k links (or equivalently a link with weight k) with that community. In particular, the procedure runs as follows: 1. we start with a preliminary partition of the network into q communities; 2. for each community we compute, either on a binary version of the original network or on the SVnet, the number k of links running between that community and each external node; 3. we evaluate the probability of observing k links under the null model using the binomial distribution (4.3) with suitably adapted parameters; 4. we correct for multiple hypothesis testing by requiring that, in order to validate the inclusion of r nodes into the q communities, the overall probability of their links doesn't exceed a threshold p, i. e. we employ again Bonferroni

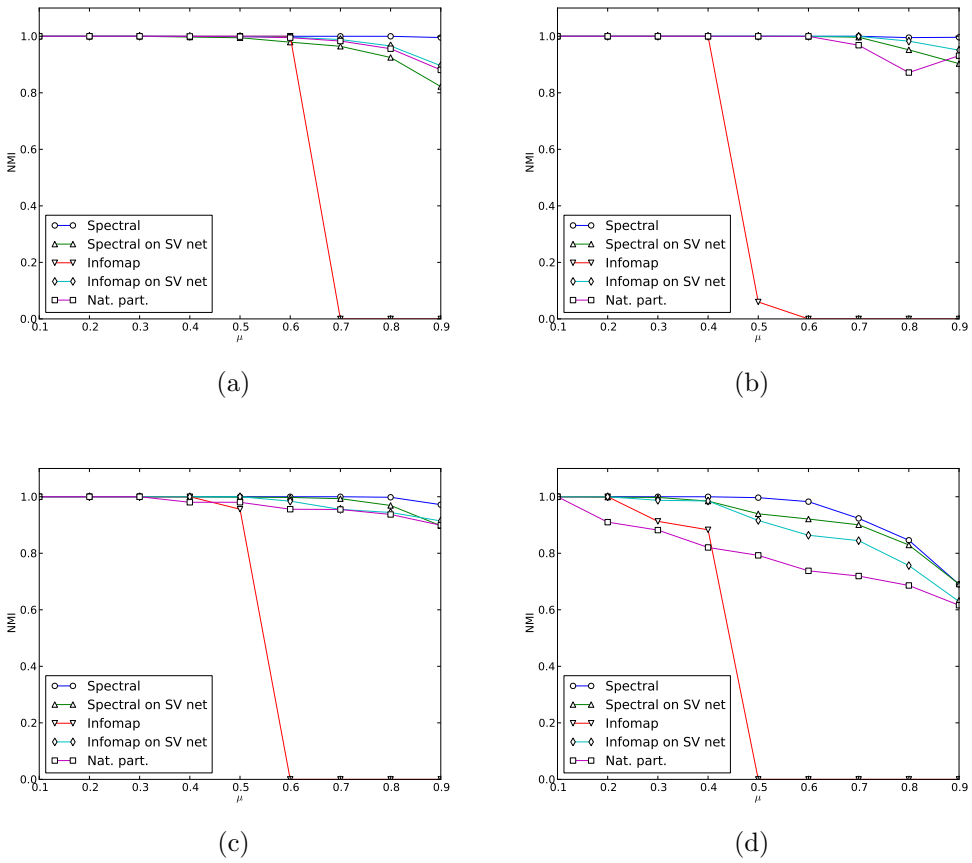

Figure 5: Tests over samples of 100 MB networks: (a) 422 nodes with 10 communities of uniform size distribution; (b) 211 nodes with 10 communities of uniform size distribution; (c) 211 nodes with 10 bounded (max: 50; min: 10) power-law distributed communities; (d) 211 nodes with 10 bounded (max: 100; min: 5) power-law distributed communities.

correction, lowering the acceptance probabilities as r increases. As a final step, we need also to adapt $I_{norm}(P, P^*)$ as suggested by Radicchi et al. (2011) in order to compare correctly the resulting assignment.

The results obtained are summarized in fig. 6. In this context the suffix "on SV net", when referred to the spectral algorithm, has a different meaning than before. Since we know that performing spectral partition over the SVnet has a negative effect, we proceed as follows: we apply the spectral algorithm to the original network, while the procedure for the recovery of overlapping nodes is performed either on the binarized version of the original network or on the SVnet. Instead, in the case of Infomap, both the partition procedure and the recovery of overlapping nodes are applied either on the former or on the latter. We observe that the best performance is obtained with spectral partition, and particularly when the recovery of overlapping nodes is performed over the SVnet.

5 Real credit networks

In this section we are going to apply the methods devised above to a dataset describing credit relationships between Japanese firms and banks. This dataset, which has been analyzed under different perspectives in previous works (De Masi et al., 2011; Fujiwara, 2009; Fujiwara and Aoyama, 2008; Bargigli and Gallegati, 2011), includes firms listed in the Japanese stock-exchange markets. Data are compiled from firms' financial statements, integrated by a survey of *Nikkei Media Marketing, Inc.* in Tokyo. They include the indication of the amount of borrowing obtained from each financial institution, subdivided in short-term debt and long-term debt. Financial institutions, which for sake of simplicity are referred to as "banks", consist of long-term credit banks, city banks, regional banks (primary and secondary), trust banks and insurance companies, all of which represent the universe of financial institutions in Japan [14].

We already know that bank-firm relationships are represented by a (weighted and directed) bipartite network, where the strengths are given by the amounts of credit/debt and the bipartite sets are represented by firms F and banks B respectively. In this context, a correct assignment requires that each community contains at least one element from each bipartite set. In fact, two firms (banks) may belong to the same community only if both are strongly connected to the same banks (firms). Thus, the smallest bipartite community is made of a single firm and a single bank. For this reason, we follow Barber

[14] For a more extensive description of this network the reader can refer to De Masi et al. (2011).

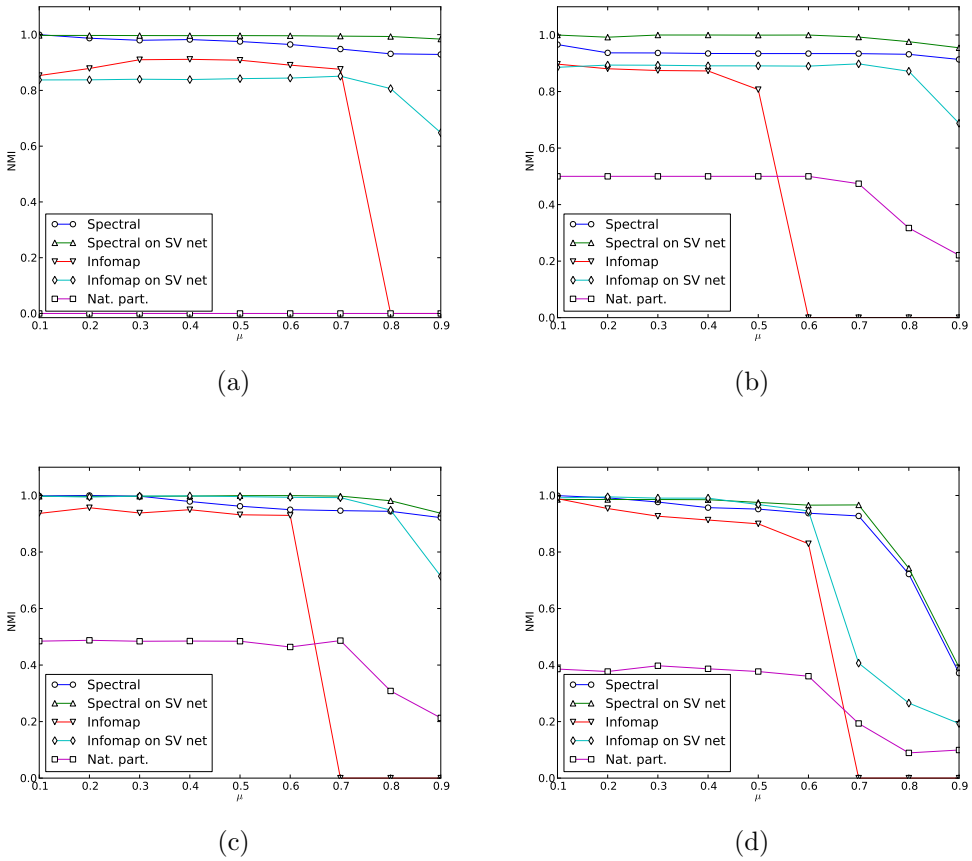

Figure 6: Tests over samples of 100 MB networks with overlapping rate $\beta = 0.05$: (a) 422 nodes with 10 communities of uniform size distribution; (b) 211 nodes with 10 communities of uniform size distribution; (c) 211 nodes with 10 bounded (max: 50; min: 10) power-law distributed communities; (d) 211 nodes with 10 bounded (max: 100; min: 5) power-law distributed communities.

Table 2: Descriptive statistics of the Dataset

Year	Firms	Banks	Links
2000	2,629	211	27,389
2001	2,714	204	26,597
2002	2,739	202	24,555
2003	2,700	192	22,585
2004	2,700	190	21,919
2005	2,674	182	21,811

(2007) in classifying both bipartite sets simultaneously. To this purpose, we can work directly with K obtained from W as defined in sec. 2, which has shape $|F| \times |B|$ and entries equal to the amounts of credit/debt between each bank and firm in G. In this case, we employ the singular vectors of K as inputs for the spectral decomposition method described in section 4. We resort to spectral decomposition since, as explained in the same section, Infomap cannot be applied to bipartite networks. Finally, we detect multi-community nodes by means of the probabilistic procedure described in the final part of the previous section.

The first step of the analysis is to determine the number of communities by means of the procedure based on eq. (3.4). In fig. 7 we draw for total credit data the squared SVs and the first differences of p-values as we did for artificial networks (we obtained equivalent graphs for partial data). We observe that the tail of the ranked SVs line is regular like in fig. 1. This fact suggests that Japanese real networks have a clear community structure, with a very limited overlap of communities. On the other hand, the p-values grow in a very regular way and display a very clear peak, as we observed in fig 2. Thus the number of communities can be safely detected from SV data according to our previous tests.

By comparing the complete results displayed in tab. 3 with tab. 2, we see that the number of communities tracks quite closely the number of banks in the dataset. This suggests that, while Japanese quoted firms may entertain relationships with a number of banks, at the same time they tend to cluster either around a single bank or, possibly, around a limited number of banks as their "special" partners. This result is consistent with previous analysis performed on the same dataset, as well as with the widely known Japanese "main bank system" (De Masi et al., 2011).

In tabs. 4, 5, 6 we report some summary statistics computed as a result

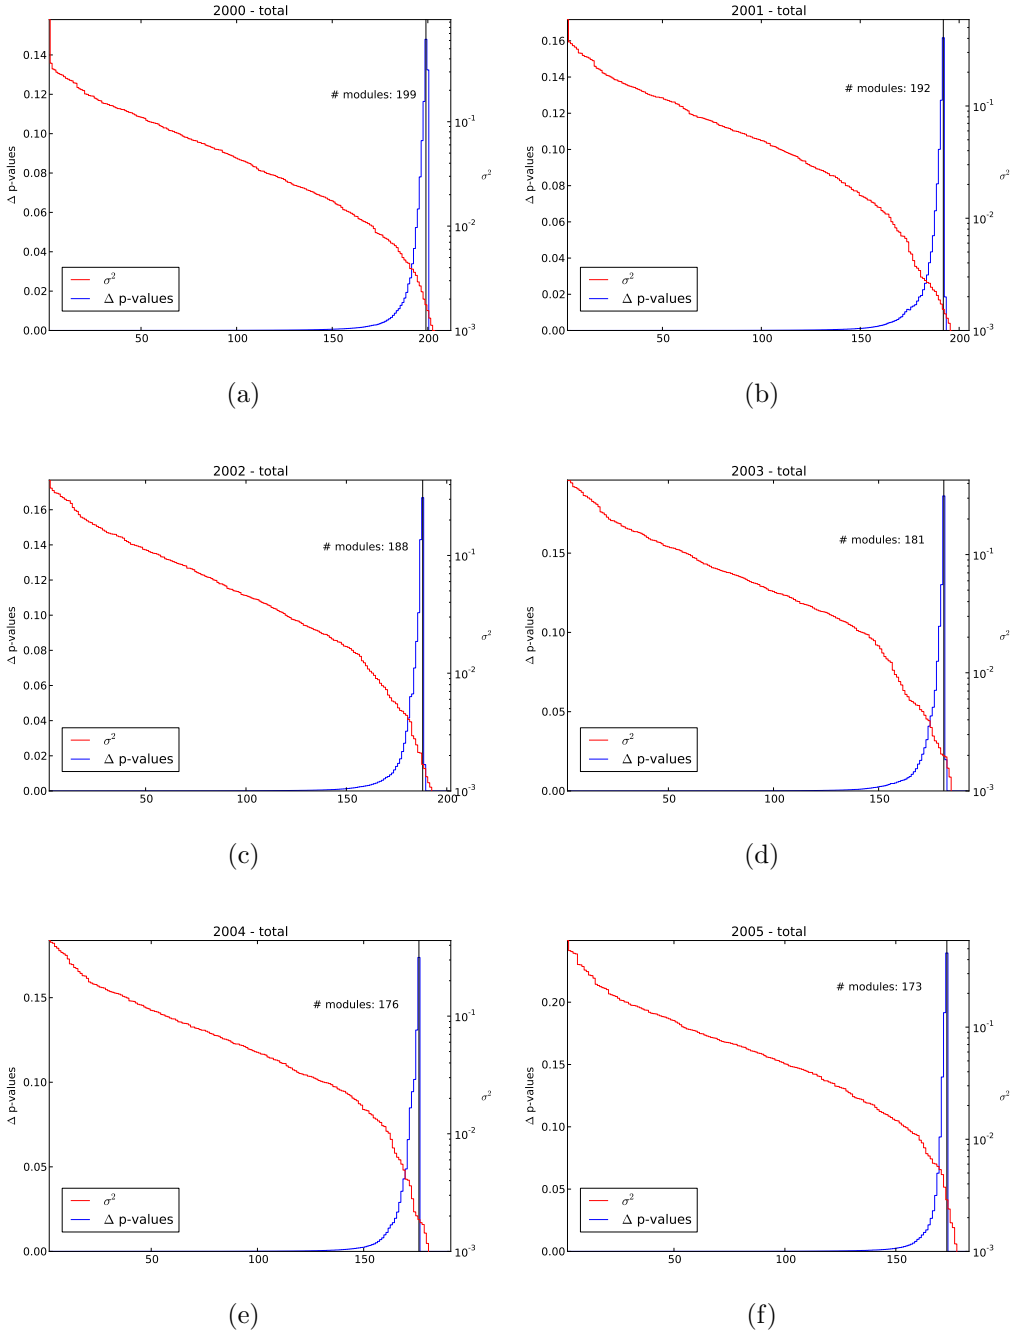

(a) (b) (c) (d) (e) (f)

Figure 7: Number of communities in the Japanese credit networks - total credit data.

Table 3: Number of communities

Year	Short term	Long term	Total
2000	169	192	199
2001	164	184	192
2002	158	182	188
2003	155	175	181
2004	149	174	176
2005	144	171	173

of our community detection procedure. From the second line of each table we see that our hypothesis of a limited community overlap is confirmed, as the fraction of multi-community nodes never exceeds significantly the 3% threshold. The large majority of nodes within these subsets are represented by banks, as expected. Further, we see that the total fraction is mostly higher than the partial ones, which would suggest that the overlapping patterns for different maturities tend to be complementary, i.e. multi-community nodes at one level don't act as such at a different level. It's noteworthy that the very limited degree of overlap coexists with high connectivity, since almost all of the nodes in each network are included in the largest component of the SVnet, as we observe from the first line of each table. Since a comprehensive planning of connections at the network level is excluded, here we find a clear sign of self-organization, as nodes behave spontaneously in such a way as to obtain a (quasi) optimal state in which connectivity is maximized while community overlap is minimized.

Links running between nodes in the same community carry only a minor portion of the total weight of the network. This is not surprising since "internal" links themselves represent a minority of total links. The latter is an interesting outcome: the participants of Japanese credit networks are likely to entertain a large number of relationship which may be defined as "weak" instead of focusing only on preferential relationships, and cutting all the remaining ties. This behavior is likely to be explained either by an implicit assurance against the eventuality of a disruption of preferential relationships or by complementary factors like, for instance, information gathering (De Masi et al., 2011). At any rate, we can detect a first clue of the effectiveness of our community assignment in the fact that the fraction of internal weight is always higher than the fraction of internal links. A more cogent measure is obtained by adapting Newman's modularity:

$$Q = \mathrm{Tr}(D) = \mathrm{Tr}(R^T B S) \tag{5.1}$$

where R and S are the community matrices obtained for the firms and banks bipartite sets respectively and B is the Newman's modularity matrix (see sec. 3). In order to obtain comparable values across different networks, we divide Q by the number of internal links. Since $\sum_{ij} d_{ij} = 0$, the sum of the off-diagonal cells of D is equal to $-Q$ by construction. Thus we divide the latter value by the (higher) number of external links in the last line of each table. The resulting values confirm the argument, advanced in Bargigli and Gallegati (2011), that in the Japanese credit network preferential connections with a small number of nodes are traded against the relatively modest downsizing of a large number of non preferential connections, instead of being traded against a reduction in the overall number of links.

Our community statistics are also informative of the temporal evolution of the networks. For instance, we see that the fraction of internal links and internal weight is increasing over time, especially for short-term data. The tendency towards a more pronounced community structure is consistent with the results of Bargigli and Gallegati (2011), where it was found that the distance between the observed networks and their expected configuration under the binomial model was growing over time. Curiously, this tendency is associated more to a decrease of external Q than to an increase of internal Q.

Table 4: Statistics: total data

Stats	2000	2001	2002	2003	2004	2005
Size of largest valid component	2,836	2,914	2,936	2,889	2,882	2,850
Fraction of multi-community nodes	0.0303	0.0260	0.0282	0.0239	0.0242	0.0280
Fraction of internal links	0.1359	0.1423	0.1584	0.1728	0.1736	0.1748
Fraction of internal weight	0.1681	0.1828	0.1905	0.2144	0.2229	0.2103
Mean Q per internal link	2,992.4	3,022.8	2,866.3	2,887.1	2,895.9	2,676.4
Mean Q per external link	-468.6	-502.3	-540.4	-602.6	-608.4	-566.5

As a final step, we plot the distributions of community size (fig. 8) and of the number of communities to which each nodes belongs (fig. 9). From the figures we see no sign of fat tails, since all distributions appear linear on a semilog scale. Regarding community size, we don't find in our data very large communities, although the tail of the distribution tends to become higher over time. This result is consistent with the previous remarks on the

Table 5: Statistics: long-term data

Stats	2000	2001	2002	2003	2004	2005
Size of largest valid component	2,372	2,457	2,495	2,467	2,474	2,449
Fraction of multi-community nodes	0.0231	0.0252	0.0244	0.0243	0.0221	0.0232
Fraction of internal links	0.1526	0.1769	0.1907	0.1885	0.2081	0.2020
Fraction of internal weight	0.1748	0.2113	0.2076	0.2237	0.2408	0.2131
Mean Q per internal link	2,497.2	2,514.2	2,399.4	2,641.0	2,491.9	2,320.7
Mean Q per external link	-450.0	-542.3	-568.8	-613.6	-655.0	-587.3

Table 6: Statistics: short-term data

Stats	2000	2001	2002	2003	2004	2005
Size of largest valid component	2,564	2,595	2,585	2,503	2,378	2,324
Fraction of multi-community nodes	0.0230	0.0227	0.0224	0.0231	0.0231	0.0236
Fraction of internal links	0.1640	0.1722	0.1920	0.2036	0.2145	0.2228
Fraction of internal weight	0.1914	0.2158	0.2491	0.2984	0.3010	0.2879
Mean Q per internal link	1,788.0	1,811.9	1,859.8	1,989.1	1,708.7	1,544.0
Mean Q per external link	-350.6	-377.0	-442.9	-509.5	-469.6	-442.6

high number of external links (which is reflected in the relatively small size of communities), as well as with the evolution towards a more pronounced community structure over time. It's also interesting to observe that very small communities are detected with our procedure (down to the theoretical minimum), so that, as expected, we discover no sign of a resolution bias in our results.

Finally, regarding fig. 9, we remark that some nodes are likely to be part not just of two, but of many communities. The role of these small subsets could turn out to be qualitatively and/or quantitatively relevant within the overall system, e.g. with respect to shock transmission. For this reason, it might be worth to carry out a more detailed analysis of this phenomenon.

6 Conclusions

It is worth to summarize the main findings of our analysis. As a first step, we have shortly presented, in sec. 2, a variety of techniques for building artificial network models satisfying sets of constraints on network observables. Then, in sec. 3 we have applied one of these techniques to obtain ensembles of random networks endowed with a community structure, which we called modular binomial (MB) networks. We showed that the SVD properties of MB networks reveal their community structure and, in particular, that it's possible to detect the number of communities contained in MB networks by looking at their singular values.

This result helps us to overcome the limitations of modularity as a device to detect the correct number of communities (note 12). Instead, in sec. 4 we resort to alternative algorithms, such as Infomap and spectral decomposition, complemented by the "Statistically validated networks" approach. Further, since we cannot exclude the possibility that communities in real networks overlap, we introduce the notion of *weak* communities by letting some nodes be part of more than a single community. In order to detect the overlapping nodes, we adopt the solution proposed by Radicchi et al. (2011) of adding single nodes to an initial partition with the help of statistical inference. Both in the case of strong and of weak communities, we obtain the best results by applying the spectral decomposition algorithm to the original (non validated) network.

In order to illustrate our method, in sec. 5 we have applied the spectral decomposition algorithm to the Japanese credit dataset. Firstly, we have obtained the number of communities by using singular values, and observed that Japanese firms tend to cluster around a single bank as their major partner, a result which is consistent with the Japanese "main bank" system (De

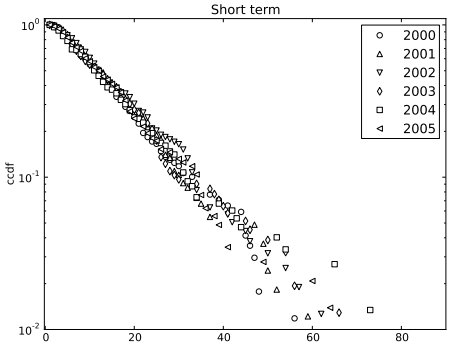

(a)

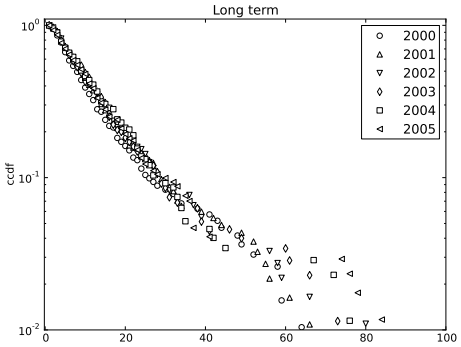

(b)

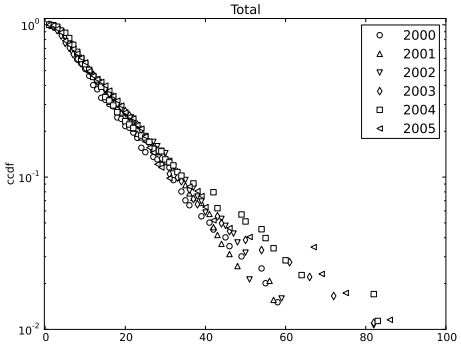

(c)

Figure 8: Community size distribution

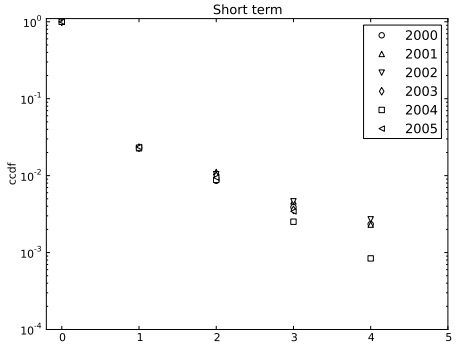

(a)

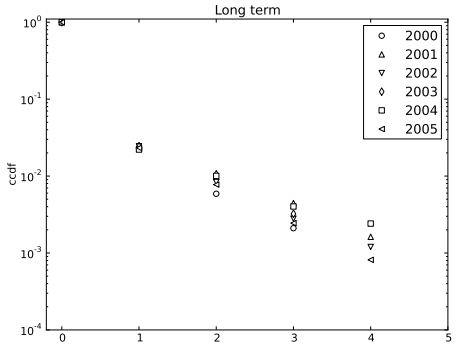

(b)

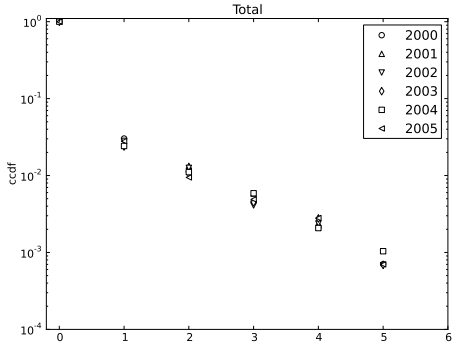

(c)

Figure 9: Distribution of nodes by the number of their communities

Masi et al., 2011). Secondly, we showed that the real networks under study are highly connected (since the largest valid component spans most of the nodes), while at the same time the fraction of overlapping nodes between communities is quite limited (and mostly represented by banks). Thirdly, we observed that the majority of links is external to communities. Taken together, these findings show that the connectivity of Japanese credit network is much higher of the one we could expect by taking into account only within-community (internal) connections. In other terms, Japanese firms and banks entertain a high number of links which we could define as "weak" since they connect those nodes to other nodes belonging to different communities. Lastly, we confirm the findings of Bargigli and Gallegati (2011) in that we find evidence of a strengthening of the community structure over time, with an increasing number of internal links and a growing size of communities themselves.

Thus, we find comfort both from simulations and from real data regarding the possibility to apply a suitably adapted community detection method to credit markets. As explained in section 2, the existence of statistically significant communities is assessed by comparing real markets with the expectation of the binomial network ensemble. The latter provides the maximally diversified model satisfying a constraint over strength distributions, i.e. over the expected lending or borrowing of the agents. Thanks to this property, binomial networks allow us to define, for a real credit market, *network risk* as the default rate in excess of the default rate which is expected in the binomial ensemble. The comparison between the binomial ensemble and real markets has already showed, in some cases, that the latter are indeed a source of network risk (van Lelyveld and Liedorp, 2006; Mistrulli, 2011). Starting from these premises, we believe that our method can fruitfully complement the study of contagious defaults. The next step in this direction is to estimate the weight of intracommunity default contagion, which is expected to be high, under different scenarios. We leave this task for future research.

References

Allen, F., Gale, D., 2000. Financial contagion. Journal of Political Economy, 108(1): 1–33.

Amini, H., Cont, R., Minca, A., 2010. Resilience to Contagion in Financial Networks. *SSRN eLibrary*.

Amini, H., Cont, R., Minca, A., 2012. Stress Testing the Resilience of Finan-

cial Networks. *International Journal of Theoretical and Applied Finance*, 15(1).

Aoki, M. , Yoshikawa, H., 2006. Reconstructing Macroeconomics, Cambridge University Press.

Barber, M.J., 2007. Modularity and community detection in bipartite networks. *Physical Review E*, 76:066102, December.

Bargigli, L., Gallegati, M., 2011. Random digraphs with given expected degree sequences: A model for economic networks. *Journal of Economic Behavior & Organization*, 78(3):396 – 411.

Battiston, S., Delli Gatti, D., Gallegati, M., Greenwald, B.C., Stiglitz, J.E., 2012. Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. *Journal of Economic Dynamics & Control*, 36(8):1121 – 1141.

Bianconi, G., 2008. The entropy of randomized network ensembles. *EPL (Europhysics Letters)*, 81(2):28005.

Bianconi, G., 2009. Entropy of network ensembles. *Physical Review E*, 79(3):036114, March.

Brunnermeier, M.K., 2009. Deciphering the liquidity and credit crunch 2007-2008. Journal of Economic Perspectives, 23(1): 77–100, Winter.

Castrén, O., Kavonius, I.K., 2009. Balance sheet interlinkages and macro-financial risk analysis in the euro area. Working Paper Series 1124, European Central Bank, December.

Chauhan, S., Girvan, M., Ott, E., 2009. Spectral properties of networks with community structure. *Physical Review E*, 80(5):056114, November.

Cont, R., Moussa, A., Santos, E.B., 2012. Network Structure and Systemic Risk in Banking Systems, in Fouque, J.P., & Langsam, J. (eds.),*Handbook of Systemic Risk*, Cambridge University Press.

De Masi, G., Fujiwara, Y., Gallegati, M., Greenwald, B.C., Stiglitz, J.E., 2011. An analysis of the Japanese credit network. *Evolutionary and Institutional Economics Review*, 7(2): 209-232.

Donetti, L., Muñoz, M. A., 2005. Improved spectral algorithm for the detection of network communities. *AIP Conference Proceedings*, 779:104–107, July 2005.

Elsinger, H., Lehar, A., Summer, M., 2006. Systemically important banks: an analysis for the European banking system. *International Economics and Economic Policy*, 3(1):73–89.

Fortunato, S., Barthélemy, M., 2007. Resolution limit in community detection. *Proceedings of the National Academy of Sciences*, 104(1):36–41.

European Central Bank. *Recent Advances in Modelling Systemic Risk Using Network Analysis*, Frankfurt am Main, January 2010. European Central Bank.

Fortunato, S., 2010. Community detection in graphs. *Physics Reports*, 486(3-5):75–174.

Fujiwara, Y., Aoyama, H., 2008. Large-scale structure of a nation-wide production network. The European Physical Journal B, 77(4): 565-580.

Fujiwara, Y., Aoyama, H., Ikeda, Y., Iyetomi, H., Wataru, S., 2009. Structure and temporal change of the credit network between banks and large firms in Japan. Economics e-journal, 3(7).

Gai, P., Kapadia, S., 2010 Contagion in financial networks. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science, 466(2120): 2401–2423.

Jackson, Matthew O. Social and Economic Networks. Princeton: Princeton University Press.

Lancichinetti, A., Fortunato., S., 2009a. Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities. *Physical Review E*, 80(1):016118, July.

Lancichinetti, A., Fortunato, S., 2009b. Community detection algorithms: a comparative analysis. *Physical Review E*, 80(5):056117, November.

Lancichinetti, A., Fortunato, S., KertÃ©sz, J, 2009. Detecting the overlapping and hierarchical community structure in complex networks. *New Journal of Physics*, 11(3):033015.

Mistrulli, P.E., 2011. Assessing financial contagion in the interbank market: Maximum entropy versus observed interbank lending patterns. *Journal of Banking & Finance*, 35(5):1114 – 1127.

Mitrović, M., Tadić, B., 2009. Spectral and dynamical properties in classes of sparse networks with mesoscopic inhomogeneities. Physical Review E, 80(2): 026123, August.

Molloy, M., Reed, B., 1995. A critical point for random graphs with a given degree sequence. *Random Structures and Algorithms*, 6:161–179.

Newman, M.E.J., Girvan, M., 2004. Finding and evaluating community structure in networks. Physical Review E, 69(2): 026113, February.

Newman, M. E. J., Strogatz, S. H., Watts, D. J., 2001. Random graphs with arbitrary degree distributions and their applications. *Physical Review E*, 026118.

Ng, A. Y., Jordan, M. I., Weiss, Y., 2001. On spectral clustering: Analysis and an algorithm. In *ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS*, pp. 849–856. MIT Press.

Nier, E., Yang, J., Yorulmazer, T., Alentorn, A., 2007. Network models and financial stability. *Journal of Economic Dynamics and Control*, 31(6):2033 – 2060.

Park, J., Newman, M. E. J., 2004. Statistical mechanics of networks. *Phys. Rev. E*, 70:066117, Dec 2004.

Radicchi, F.,Lancichinetti, A., Ramasco, J., Fortunato, S., 2011. Finding statistically significant communities in networks. *PLoS ONE*, 6(4):e18961, 04.

Rosvall, M., Bergstrom, C. T., 2008. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123.

Sawardecker, E. N., Amundsen, C. A., Sales-Pardo, M., Amaral, L. A.N., 2009. Comparison of methods for the detection of node group membership in bipartite networks. *European Physical Journal B*, 72(4):671–677, 2009.

Squartini,T. Garlaschelli, D., 2011. Analytical maximum-likelihood method to detect patterns in real networks. *New Journal of Physics*, 13(8):083001, 2011.

Tumminello, M., Miccichè, S., Lillo, F., Piilo, J., Mantegna, R. N.,Ben-Jacob, E., 2011. Statistically Validated Networks in Bipartite Complex Systems. *PLoS ONE*, 6, March.

Upper, C., 2011. Simulation methods to assess the danger of contagion in interbank markets. *Journal of Financial Stability*, 7(3):111 – 125.

Upper, C., Worms, A., 2004. Estimating bilateral exposures in the german interbank market: Is there a danger of contagion? *European Economic Review*, 48(4):827 – 849.

van Lelyveld, I., Liedorp, F., 2006. Interbank contagion in the dutch banking sector: A sensitivity analysis. *International Journal of Central Banking*, 2(2), May.

Vega-Redondo, F., 2007. Complex Social Networks. Cambridge: Cambridge University Press.