Tensorial Factor Analysis of International Bilateral Claims

Paolo Giudici[*]
University of Pavia, Pavia, Italy

Nicolò Pecora[†]
Catholic University, Piacenza, Italy

Alessandro Spelta[‡]
Catholic university, Milano, Italy

May 15, 2016

Abstract

The paper presents a new methodology aimed at detecting the modularity structure of an evolving weighted directed network, identifing communities and central nodes inside each of them, and tracking their common activity over time. The proposed methods are based on non-negative tensorial factorizations, and performance measures that evaluate their goodness of fit. The proposed methods are applied to the Consolidated Banking Statistics (CBS), provided by the Bank of International settlements (BIS). The empirical findings show that the introduction of a temporal dimension gives the opportunity to study the evolution of the communities over time. Data are well represented by three communities, dependent on the banking business models employed in different countries, as well as on their geographical proximity. The temporal activation of each community varies according to events involving the member nodes, showing a substantial decrease during crisis periods, such as the 2008 financial crisis and the European sovereign debt crisis.

1 Introduction

Many financial systems can be fruitfully represented as networks involving elementary structural entities, such as banks, hedge funds, etc. and specific relations between them. Studies that analyze the empirical characteristics of financial networks in different jurisdictions have systemically found the existence of a community structure (see Soramäki et al. 2007, Iori et al. 2007 and 2008, Cocco, Gomes and Martins 2009, Craig and Von Peter 2014, Fricke 2012 and Fricke and Lux 2013 among others). The community structure reveals how a network is internally organized, and indicates the presence of special relationships between nodes, that may not be easily accessible from direct empirical tests. In other words the community structure refers to the occurrence of groups of nodes that are more densely connected internally than with the rest of the network. This definition, suitable for undirected networks, has usually been (wrongly) applied to directed networks by assuming a member of a community having balanced out-links and in-links connections to other members. This symmetric assumption can be seriously violated in cases where a member may only play one main role, source (lender) or terminal (borrower), in the community. Asymmetric communities are common in directed networks where the direction implicitly express an asymmetric relationship among its nodes. A recent survey (Malliaros and

[*]paolo.giudici@unipv.it.
[†]nicolo.pecora@unicatt.it
[‡]alessandro.spelta@unicatt.it

Vazirgiannis 2013) provides a broader definition of community structure as *set of nodes that share common or similar feature together.* Financial networks, where links represents flow of funds and institutions have exposures on both sides of their balance sheets, are typical examples.

In addition, much of the focus of community detection algorithms has been on identifying disjoint communities. However, it is well known that nodes in a network are naturally characterized by multiple community memberships (Xie, Kelley and Szymanski 2013). Also in financial network, it is very common for an institutions to participate in more than one community, i.e., communities are often overlapped.

Beside community membership distributions, not all vertices are equal in a community and some vertices might be special in the sense that they are linked with almost all others. In literature, such a vertex is known as hub, leader, or center. In financial market, determining who are this central players is key to design policies that try to prevent and mitigate contagion. Not surprisingly, the research in network theory has dedicated a vast amount of effort to deal whit this question (Financial Stability Board 2013, Battiston et al. 2012).

Various measures of centrality had been proposed in network theory such as those based on counting the first neighbors of a node (degree centrality), or those based on the spectral properties of the graph (see Perra and Fortunato 2008). Spectral centrality measures include the eigenvector centrality (Bonacich 1972, Bonacich and Lloyd 2001), Katz's centrality (Katz, Blumler and Gurevitch 1973), PageRank (Brin and Page 1998), Hyperlink-Induced Topic Search (HITS) algorithm (Kleinberg 1999). These measures provide information on the position of each node relative to all the others.

In general, centrality measures rank vertices without paying attention to whether the network is characterized by a community structure and how this structure in turn affects the ranking. Only in Cao et al. 2013, the authors propose a novel model to identify overlapping communities and central nodes, in case of undirected static network.

In the context of financial markets, developing a measure that captures the institutions' systemic importance by revealing also the community structure, that proxies the most plausible areas of contagion of institutions' distress, may enhance a better understanding of the system functioning

Finally financial networks are highly dynamic. Although it is always possible to create static network representations by aggregating over the temporal evolution of the system, such temporally-aggregated representations may overlook essential features of the system or may confound structures that can be teased apart only by retaining the time-varying nature of the data. This should lead to models of financial networks that dynamically adapt to critical events, with a continuous feedback between topology and dynamics. Few studies pioneered approaches to community detection in temporal networks (Gauvin, Panisson and Cattuto 2014) but none of these works address at the same time the issue of identifying communities and central nodes inside each community in evolving directed weighted networks.

In this paper we employ a technique for multiway-data proposing a new methodology to detect the modularity structure of an evolving weighted directed network and to identify central nodes inside each community, tracking their common activity over time. This method is based on the fact that a temporal network is naturally represented as a time-ordered sequence of adjacency matrices, each one describing the state of the financial network at a given point in time. The adjacency matrices are thus combined in a single mathematical object: a three-way tensor. In the last ten years, interest in tensor analysis and their decomposition have expanded to different fields. Examples include signal processing (Chen, Petropulu and De Lathauwer 2002, Comon 2002), numerical linear algebra (De Lathauwer, De Moor and Vandewalle 2000 and 2001), computer vision (Vasilescu and Terzopoulos 2002a and 2002b) numerical analysis (Beylkin and Mohlenkamp 2002 and 2005), data mining (Acar et al. 2005, Acar, Amtepe and Yener 2006) graph analysis (Kolda and Bader 2006), and more, but applications to economics are still scarce (Bonacina et al. 2014).

The approach described in this paper is based on a particular tensor decomposition technique, the so-called CP decomposition (named after the two most popular and general variants, CANDECOMP developed in Carroll and Chang (1970) and PARAFAC developed by Harshman (1970)). It can be regarded as a generalization of the singular value decomposition (SVD) applied to tensors. In particular, we focused on non-negative tensor decomposition (Cichocki, Phan and Zdunek 2009, Morup 2011, Kolda and Bader 2009, Shashua and Hazan 2005). As already observed for non-negative matrix factorization (Lee and Seung 1999), it is a powerful tool for learning parts-based representation of a dataset, resulting in a more interpretable models (Nickel, Tresp and Kriegel 2011, Wang et al. 2011).

Our methodology can be seen as a multidimensional extension of the HITS algorithm (Kleinberg 1999). This algorithm provides two attributes for each node: an authority score and a hub score. Authority measures prestige: nodes who many other nodes point to are called authorities. If a node has a high number of nodes pointing to it, such a node has a high authority value and this quantifies its role as a source of information. On the contrary, a hub is an actor referring to many authorities and its score measures acquaintance. Hub and authority scores have different interpretation in term of the systemic importance associated to each financial institution. Institutions with high authority score are the main systemically important debtors (borrower) while those having high hub scores are the main systemically important creditors (lenders). In particular, as the HITS algorithm, our technique assigns to each node a centrality score proportional to the sum of the scores of its neighbors, and centrality results from a node having many neighbors, or from having some central neighbors, or both. Thus, two players will be ranked differently as hubs even if they lend the same amount of funds, depending on the behavior of their borrowers. The algorithm will rank higher those that lend to the most systemically important borrower. The same happens for the authority score with respect to the lender: two players that borrow the same amount of funds will be ranked differently depending from the lender they borrow from.

Kleinberg notes also the property of *reinforcement relationships* among nodes, which means that good hubs are the nodes that referred to good authorities and good authorities are nodes that are referred by good hubs. This reinforcement relationship suggests that nodes that make themselves good hubs and good authorities each other can be placed together in the same community.

When including also the temporal dimension, the CP decomposition provides a further score related to the temporal evolution of the activity level of the communities over time. The value of the activity pattern of a community in a time span is related to both hub and authority scores of all the players involved in the transactions occurred in that community during a certain period.

Once we have the hub and the authority vectors for each community we can retrieve the adjacency matrix describing the sub-network topology related to the community as the outer product of the two vectors. In order to assign nodes to the community to which they belong two schemes are suggested. The soft partition scheme is proposed by assigning to each node the percentage of its strength centrality that belong to that community. Such an edge decomposition can then be used also to assign nodes to communities according to a hard partition scheme, assigning each node to the community in which it has the highest impact in terms of strength.

We apply this technique to the consolidated banking statistics (CBS) compiled by the Bank of International settlements (BIS), which include countries bilateral claims. We consider a set of 30 countries each of them represented by the amounts of its foreign claims vis-a-vis the other countries, measured on quarterly basis, from the first quarter of 2005 (Q1-2005) to the last quarter of 2013 (Q4-2013).

We find that the emerging communities depend on the creditor's business model and on the geographical proximity between countries. The first community is related to the business banking model adopted in the United Kingdom and Japan, where banks pool funds at major offices and redistribute them around the banking group with the United States playing the role of the major hub. The time evolution of the statistics associated with this community shows a stable trend all over the sample.

The second community is specific to countries belonging to the European Union: France, Belgium, Germany and the Netherlands, that share cross-border activity with the two largest banking centers, the United States and the United Kingdom. Switzerland belongs to this community, but also to the first one. The time evolution of the statistics of this community show that its time activity increased substantially in the second part of the sample, when the European sovereign debt crisis erupted.

Finally the third community is mostly country specific, associated with the United States, having the largest credit risk against the United Kingdom, Japan and Germany. This community also encompasses most of the non-reporting countries in the BIS statistics. The time scores related to this community reveal that its activity was high in the first part of the sample, reflecting the subprime mortgage crisis of 2007 and the subsequent financial crisis of 2008-2009.

The paper is organized as follows: section 2 provides the data context and the proposed methodology, section 3 illustrates the main results, section 4 discusses the economic implications and the stylized facts that drive the obtained outcomes. Section 5 concludes.

2 Proposal

2.1 Bank of International Settlements Data

The Bank of International Settlements (BIS) [1] compiles statistics on international banking activity. The International Banking Statistics (IBS) comprises consolidated banking statistics (CBS), which measure worldwide consolidated claims of banks headquartered in reporting countries, including claims of their own foreign affiliates but excluding interoffice positions. These statistics build on measures used by banks in their internal risk management systems. The CBS include data on off-balance sheet exposures, such as risk transfers, guarantees and credit commitments. They also capture the maturity structure of some claims, but provide very limited currency information. The number of countries participating in the CBS increased from 15 in the early 1980s to 31 in 2012.

In this work we employ the consolidated banking statistics on ultimate risk basis (CBS/UR) that are based on the country where the ultimate risk or obligor resides, after taking into account risk transfers. The CBS provide information about banking systems' risk exposures, in particular country risk. The expression "CBS reporting area" refers collectively to the set of reporting countries. Since the statistics capture banks' worldwide consolidated positions, the CBS reporting area is not synonymous with the location of the banking offices participating in the data collection. That is, a reporting country should consolidate the positions of all banking entities owned or controlled by a parent institution located in the reporting country, and thus will include banking entities which are actually domiciled elsewhere.

Reporting institutions are financial entities whose business is to receive deposits, or close substitutes for deposits, and to grant credits or invest in securities on their own account. Thus, the community of reporting institutions should include not only commercial banks but also savings banks, credit unions or cooperative credit banks, and other financial credit institutions. Unfortunately a number of countries do not report their statistics on the asset side (out-flows).

[1] www.bis.org

In our available dataset there are only 15 fully reporting countries and more than 240 that do not report. In addition, for historical reasons among others the time series contain varying starting dates, as well as a number of missing values. To address the above data quality issues first we sum over time the exposures of all the countries, then we symmetrize the obtained matrix and finally we restrict the analysis to the 30 largest economies for which the borrowed/lent loans sum up to 2863750 billion dollars for the period from 1998 to 2013 (that is the lower value of the upper 50% quantile of the degree distribution computed on total exposure matrix). We consider a set of 30 countries each of them represented by the amounts of its foreign claims vis-a-vis the other countries (reporting and non-reporting)[2], measured on quarterly basis, from the fist quarter of 2005 (Q1-2005) to the last quarter of 2013 (Q4-2013). Starting from this dataset, we build a 3-way *ultimate risk tensor* $\mathcal{X} \in \mathbb{R}^{I \times I \times K}$ where the layer $I = 30$ represents countries and $K = 36$ denotes quarters respectively (see Figure 1), and accordingly the generic element $x_{i,j,k}$, represents the amount of funds borrowed by country i from country j at time k. Thus, the ultimate risk tensor is composed by 36 slices $\mathbf{X} \in \mathbb{R}^{30 \times 30}$ each of which represents the financial transactions between countries in each quarter.

2.2 Results

A tensor is a multidimensional array. More formally, a N-way or N-th order tensor is an element of the tensor product of N vector spaces, each of which has its own coordinate system. The order of a tensor is the number of dimensions, also known as ways or modes. Vectors (tensors of order one) are denoted by boldface lowercase letters, e.g., $\mathbf{x}$. Matrices (tensors of order two) are denoted by boldface capital letters, e.g., $\mathbf{X}$. Higher-order tensors (order three or higher) are denoted by boldface Euler script letters, e.g., $\mathcal{X}$. Scalars are denoted by lowercase letters, e.g., x. The i-th entry of a vector $\mathbf{x}$ is denoted by x_i; element (i, j) of a matrix $\mathbf{X}$ is denoted by x_{ij} , and element (i, j, k) of a third-order tensor $\mathcal{X}$ is denoted by x_{ijk}. The n-th element in a sequence is denoted by a superscript in parentheses, e.g., $\mathbf{X}^{(n)}$ denotes the n-th matrix in a sequence. Sub-arrays are formed when a subset of the indices is fixed. Fibers are the higher-order analogue of matrix rows and columns. A fiber is defined by fixing every index but one. A matrix column is a mode-1 fiber and a matrix row is a mode-2 fiber. Third-order tensors have column, row, and tube fibers, denoted by $x_{:jk}$, $x_{i:k}$, and $x_{ij:}$, respectively. Slices are two-dimensional sections of a tensor, defined by fixing all but two indices: $\mathbf{X}_{i::}$, $\mathbf{X}_{:j:}$, and $\mathbf{X}_{::k}$, denotes horizontal, lateral, and frontal slides of a third-order tensor $\mathcal{X}$, alternatively, the k-th frontal slice of a third-order tensor, $\mathbf{X}_{::k}$, may be denoted more compactly as $\mathbf{X}_k$.

The norm of a tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \cdots \times I_N}$ is the square root of the sum of the squares of all its elements,

$$\parallel \mathcal{X} \parallel = \sqrt{\sum_{i_1=1}^{I_1} \sum_{i_2=1}^{I_2} \dots \sum_{i_N=1}^{I_N} x^2_{i_1, i_2, \dots i_N}} \tag{1}$$

This is analogous to the matrix Frobenius norm.

The inner product of two same-sized tensors $\mathcal{X}, \mathcal{Y} \in \mathbb{R}^{I_1 \times I_2 \times \cdots \times I_N}$ is the sum of the products

[2] Note that the 21 reporting countries have both incoming and outcoming links whereas other 9 countries, the non-reporting ones, hold only incoming links.
Abbreviations for country names are: AU (Australia), AT (Austria), BE (Belgium), CA (Canada), CH (Switzerland), CN (China), DK (Denmark), ES (Spain), FI (Finland), FR (France), DE (Germany), GB (Great Britain), GR (Greece), HK (Hong Kong), IE (Ireland), IN (India), IT (Italy), JE (Jersey), JP (Japan), KY (Caiman Islands), LU (Luxembourg), NL (Netherlands), NO (Norway), NZ (New Zealand), PT (Portugal), RU (Russia), SE (Sweden), SG (Singapore), TR (Turkey), US (United States).

of their entries, i.e.,

$$\langle \mathcal{X}, \mathcal{Y} \rangle = \sum_{i_1=1}^{I_1} \sum_{i_2=1}^{I_2} \dots \sum_{i_N=1}^{I_N} x_{i_1,i_2,\dots i_N} y_{i_1,i_2,\dots i_N} \tag{2}$$

An N-way tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is rank one if it can be written as the outer product of N vectors, i.e.,

$$\mathcal{X} = \mathbf{a}^{(1)} \circ \mathbf{a}^{(2)} \circ \dots \circ \mathbf{a}^{(N)} \tag{3}$$

The symbol "$\circ$" represents the vector outer product. A tensor $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is diagonal if $x_{i_1 i_2 \cdots i_N} \neq 0$ only if $i_1 = i_2 = \dots = i_N$. Matricization, also known as unfolding or flattening, is the process of reordering the elements of an N-way array into a matrix. The mode-n matricization of a tensor X$\in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ is denoted by $\mathbf{X}_{(n)}$ and arranges the mode-n fibers to be the columns of the resulting matrix. Tensor element $(i_1, i_2, \dots, i_N)$ maps to matrix element (i_n, j), where

$$j = 1 + \sum_{k=1, k\neq n}^{N} (i_k - 1) J_k \quad \text{with} \quad J_k = \prod_{m=1, m\neq n}^{k-1} I_m \tag{4}$$

Several matrix products are important in the following sections, so we briefly define them here. The Kronecker product of matrices $\mathbf{X} \in R^{I \times J}$ and $\mathbf{Y} \in R^{K \times L}$ is denoted by $\mathbf{X} \otimes \mathbf{Y}$. The result is a matrix of size $(IK) \times (JL)$. The Khatri–Rao product is the "matching columnwise" Kronecker product. Given matrices $\mathbf{X} \in R^{I \times J}$ and $\mathbf{Y} \in R^{K \times L}$, their Khatri–Rao product is denoted by $\mathbf{X} \odot \mathbf{Y}$. The result is a matrix of size $(IJ) \times K$. The Hadamard product is the elementwise matrix product, and it is denoted by "$*$".

2.3 Tensor decompositions

The ultimate risk tensor $\mathcal{X}$, encompasses both the topological and temporal information of the evolving directed weighted network under study. Uncovering the community structures of such network together with the related activity patterns requires the identification and extraction of lower-dimensional factors.

To this end, we use tensor factorization techniques, i.e., we choose to represent the tensor as a suitable product of lower-dimensional factors. This can be achieved by means of the so-called CP decomposition. Solving this problem consists in finding the R rank-1 tensors that best approximate the tensor $\mathcal{X}$[3]. In this respect, the tensor factorization method is similar to community detection techniques where the number of communities is fixed a priori: the number of components we choose to approximate the tensor is the number of communities or activity patterns we extract.

Assuming the number of components is fixed, let $\mathcal{X} \in \mathbb{R}^{I \times I \times K}$ be the third-order ultimate risk tensor (we relax this assumption at the end of the section while showing a test to choose the proper number of components). The goal is to compute a CP decomposition with R components that best approximates $\mathcal{X}$, i.e., to find

$$\min_{\widehat{\mathcal{X}}} \left\| \mathcal{X} - \widehat{\mathcal{X}} \right\| \quad \text{with} \quad \widehat{\mathcal{X}} = [\boldsymbol{\sigma}; \mathbf{U}, \mathbf{V}, \mathbf{W}] = \sum_{r=1}^{R} \sigma_r \mathbf{u}_r \circ \mathbf{v}_r \circ \mathbf{w}_r \tag{5}$$

[3] The rank of a tensor $\mathcal{X}$, denoted $rank(\mathcal{X})$, is defined as the smallest number of rank-one tensors that generate $\mathcal{X}$ as their sum. In other words, this is the smallest number of components in an exact CP decomposition. An exact CP decomposition with $R = rank(\mathcal{X})$ components is called the rank decomposition. The definition of tensor rank is quite similar to the definition of matrix rank, but the properties of matrix and tensor ranks are different. Principally there are no straightforward algorithm to determine the rank of a specific tensor; in fact, the problem is NP-hard. In practice, the rank of a tensor is determined numerically by fitting various rank-R CP models.

where R is a positive integer and $\mathbf{V} \in \mathbb{R}^{I \times R}$, $\mathbf{U} \in \mathbb{R}^{I \times R}$, $\mathbf{W} \in \mathbb{R}^{K \times R}$ and $\boldsymbol{\sigma} = \|\mathbf{V}\| \, \|\mathbf{U}\| \, \|\mathbf{W}\|$.

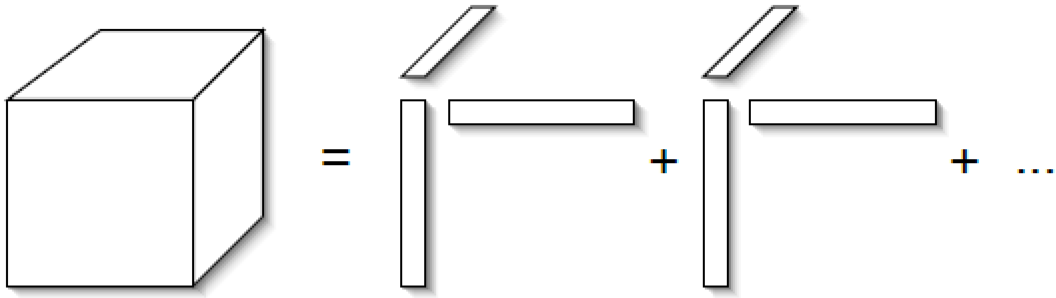

Figure 1: The CP model provides a 3-way decomposition that yields authority, hub, and activity patterns.

Since data are non negative, representing flows of funds between borrowers and lenders, we solve problem (5) applying a non-negative tensor decomposition. Thus, for $r = 1, .., R$, we obtain non negative hub, authority and activity patterns vectors $\mathbf{u}_r \geq 0$, $\mathbf{v}_r \geq 0$, and $\mathbf{w}_r \geq 0$. This is customarily used to achieve a purely additive representation of the tensor in terms of components, which greatly simplifies the interpretation of the resulting decomposition (Lee and Seung 1999).

The problem is therefore solved using a Projected Gradient Local Hierarchical Alternating Least Squares (HALS) Algorithm, (Cichocki, Phan and Zdunek 2009). The 3-dimensional problem is divided into 3-dimensional sub-problems by unfolding the tensor $\mathcal{X}$. This means reordering the elements of a tensor into three matrices $\mathbf{X}_{(1)}, \mathbf{X}_{(2)}$ and $\mathbf{X}_{(3)}$ as formula (4) suggests. The three resulting matrices have respectively a size of $I \times IK$, $I \times IK$ and $K \times II$. In this way problem (5) is equivalent to minimizing the difference between each of the modes and their respective approximation in terms of factors. Problem (5) is thus converted into two three-problems

$$\min_{\mathbf{U} \geq 0} \left\| \mathbf{X}_{(1)} - \mathbf{U}\mathbf{Y}_{VW}^T \right\|_F^2$$

where $\mathbf{Y}_{VW} = \mathbf{V} \odot \mathbf{W}$ and $\mathbf{X}_{(1)}$ is the $I \times IK$ unfolded matrix of $\mathcal{X}$ and

$$\min_{\mathbf{V} \geq 0} \left\| \mathbf{X}_{(2)} - \mathbf{V}\mathbf{Y}_{UW}^T \right\|_F^2$$

where $\mathbf{Y}_{UW} = \mathbf{U} \odot \mathbf{W}$ and $\mathbf{X}_{(2)}$ is the $I \times IK$ unfolded matrix of $\mathcal{X}$ and

$$\min_{\mathbf{W} \geq 0} \left\| \mathbf{X}_{(3)} - \mathbf{W}\mathbf{Y}_{UV}^T \right\|_F^2$$

where $\mathbf{Y}_{UV} = \mathbf{U} \odot \mathbf{V}$ and $\mathbf{X}_{(3)}$ is the $K \times II$ unfolded matrix of $\mathcal{X}$.

After computing the Karush-Kuhn-Tucker conditions, in order to estimate the stationary points, we simply compute the gradients of the local cost functions. The stationary points can be estimated via the following simple updates

$$\mathbf{U} \leftarrow \frac{1}{\mathbf{Y}_{VW}^T \mathbf{Y}_{VW}} \left[\mathbf{X}_{(1)} \mathbf{Y}_{VW}\right]_+ = \frac{1}{\mathbf{Y}_{VW}^T \mathbf{Y}_{VW}} \max\left\{\epsilon, \mathbf{X}_{(1)} \mathbf{Y}_{VW}\right\}$$

$$\mathbf{V} \leftarrow \frac{1}{\mathbf{Y}_{UW}^T \mathbf{Y}_{UW}} \left[\mathbf{X}_{(2)} \mathbf{Y}_{UW}\right]_+ = \frac{1}{\mathbf{Y}_{UW}^T \mathbf{Y}_{UW}} \max\left\{\epsilon, \mathbf{X}_{(2)} \mathbf{Y}_{UW}\right\}$$

$$\mathbf{W} \leftarrow \frac{1}{\mathbf{Y}_{UV}^T \mathbf{Y}_{UV}} \left[\mathbf{X}_{(3)} \mathbf{Y}_{UV}\right]_+ = \frac{1}{\mathbf{Y}_{UV}^T \mathbf{Y}_{UV}} \max\left\{\epsilon, \mathbf{X}_{(3)} \mathbf{Y}_{UV}\right\}$$

with $\epsilon > 0$. The number R of components is chosen on the basis of the desired level of detail: a low number of components only yields the strongest structures, whereas using a high number of components faces the risk of overfitting noise[4]. Increasing the number R of components allows us to represent more and more details of the temporal structure of the network. However, as the number of components increases, we go from underfitting to overfitting these structures, i.e., we face the usual trade-off between approximating complex structures and overfitting them, potentially capturing noise.

2.4 Community identification

It is straightforward to interpret $u_{ir}v_{ir}w_{kr}$ as the contribution, in terms of model fitting, of the r-th community to the edge $\mathcal{X}_{ijk}$. In other words, the interaction $\hat{\mathcal{X}}_{ijk} = \sum_{r=1}^{R} \hat{\mathcal{X}}_{ijk}^{r} = \sum_{r=1}^{R} u_{:r}v_{:r}w_{:r}$ between nodes i and j at time k is the result of the sum of their participation in the same communities (Psorakis et al. 2011, Mankad and Michailidis 2013). Therefore, $\hat{\mathcal{X}}$ is a summation of R tensors and each $\hat{\mathcal{X}}_{ijk}^{r}$ denotes the number of pairwise interactions in the context of community r at time k. Thus $\hat{\mathcal{X}}$ is an approximation of the original tensor $\mathcal{X}$. In order to assign nodes to the communities first we average over time the sub-tensor representing each community $\hat{\mathcal{X}}^r$ obtaining a matrix $\hat{\mathbf{X}}^r = \frac{1}{K}\sum_{k=1}^{K} \hat{\mathcal{X}}_{::k}^{r}$ then we calculate for each node its strength in each $\hat{\mathbf{X}}^r$ and we stack this measure in a matrix $\mathbf{D}^{I \times R}$ where each element d_{ir} represents the weighted degree centrality of node i in community r. Normalizing each row of $\mathbf{D}$ by $\boldsymbol{\delta} = \sum_{r=1}^{R} d_{:r}$ we obtain the soft partition solution of the so-called degree membership: $a_{ir} = d_{ir}/\delta_r$. Such an edge decomposition can then be used also to assign nodes to communities according to a hard partition scheme, assigning each node to the community in which it has the highest impact in terms of strength.

The first issue that arises in computing a CP decomposition is the choice of the number of rank-one components. Most procedures fit multiple CP decompositions with different number of components until one is “good” enough. Theoretically, if the data are noise-free and we have a procedure for calculating the CP decomposition with a given number of components, then we can do that computation for $R = 1, 2, 3, \ldots$ components and stop at the first value of R that gives a fit of 100%. However, in practice, there are many problems with this procedure. When the data are noisy (as is frequently the case), then fit alone cannot determine the rank in any case; instead, Bro and Kiers (2003) proposed the Core Consistency Diagnostic (CORCONDIA) to compare different numbers of components. The CORCONDIA test is proposed to determine whether, given a fixed number of factors, the model is better explained by a Tucker decomposition[5] (Tucker 1966) or by a CP decomposition. The CORCONDIA test should be ideally very close to one providing CP decomposition to better explain the data w.r.t. a Tucker model. In other words, if the CP

[4]In the extreme case of $R = 1$ this means to discover hubs, authorities and time scores, for the whole network, without assessing the presence of a community structure inside the network.

[5]The Tucker method decomposes a tensor into a set of matrices and one small core tensor. Initially described as a three-mode extension of factor analysis and principal component analysis, it may actually be generalized to higher mode analysis.

It may be regarded as a more flexible CP model. In CP the core tensor is restricted to be "diagonal".

model is valid then the core consistency is close to 100%. If the data cannot approximately be described by a trilinear model or too many components are used, the core consistency will be close to zero (or even negative). If the consistency is in-between (around 50%) the model is unstable, in which case imposing valid constraints may help stabilizing the model. In practice the core consistency levels off slowly for an increasing number of components and then sharply when the correct number of components is exceeded. The number of components corresponding to the last high consistency value should be chosen.

3 Application

3.1 Tensor decomposition

In general it is difficult to decide the best rank of a CP model, i.e. the column-dimension of the loading matrices. However, it is by far easier than for two-way problems because of the uniqueness properties of CP. While applying the CP algorithm to compute the scores for hubs, authorities and time activity patterns, an appropriate number of factors has to be found in order to guarantee a good fit and avoiding overfitting at the same time. In our case, we found quite clearly that the CORCONDIA test is consistently close to one (i.e. a CP model provides a good fit w.r.t. a Tucker model) for a number of factors $R = 3$ (see Figure 2). The decomposition obtained increasing the number of components stops being a proper description of the original tensor because it overfits or the components become redundant. As such, we focused on the number of factors by taking into account both a good level of fit and parsimony. Figure 2 (a) shows the goodness of fit obtained for $R = 3$ and Figure 2 (b) displays the core consistency diagnostic test for different number of factors.

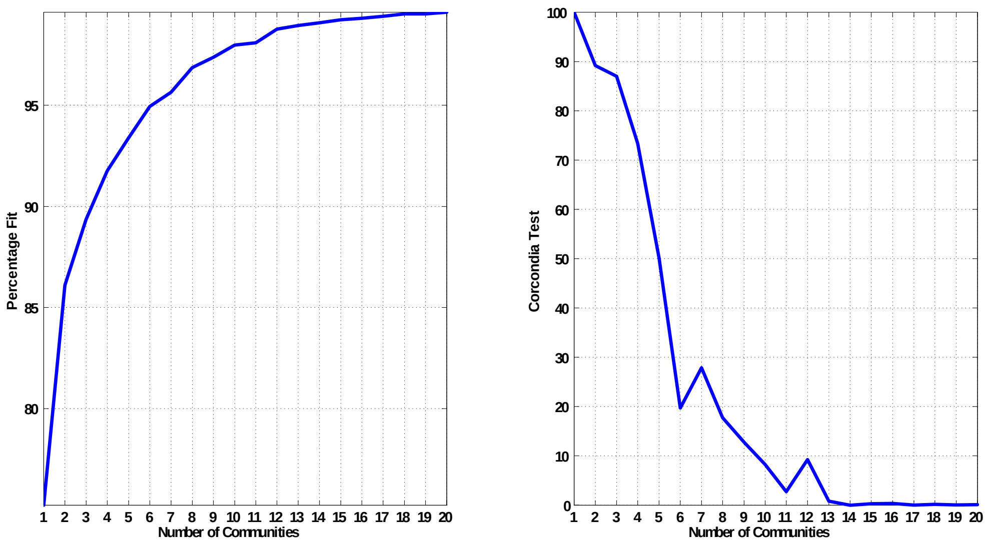

(a) (b)

Figure 2: Model fit for different number of factors (a). Core Consistency Diagnostic for different number of factors (b).

From these two pictures we decide that three are the factors that drive our data. In particular, a 3-factors model explains 88% of the data and the CORCONDIA test, being approximately around 90%, suggests that the CP model is adequate to describe the structure of the data we employ.

3.2 Community detection

Once we have identified the CP model as a suitable tool to analyze and describe our data, we can show the results of the adopted tensor decomposition technique. The factor matrices $\mathbf{U}, \mathbf{V}, \mathbf{W}$ all have R columns, each of them corresponding to the hub, authority and time activity pattern vector of one community. The matrix element u_{ir} indicates the lending systemic importance of country i into community r. Similarly v_{ir} describes the borrowing systemic importance of banks i into community r. The element w_{kr}, on the other hand, associates each component r to the time intervals k that it spans, and the matrix values for a given component indicate the activity level of that community as a function of time (index k), i.e., its temporal activity pattern. The outer product $u_{:r}v_{:r}w_{:r}$ approximate the sub-tensor representing the spatiotemporal connections inside the r-th community $\hat{\mathcal{X}}$ and the normalized strength $a_{ir} = d_{ir}/\delta_r$ computed over the time averaged matrix $\hat{\mathbf{X}}^r$, describing the sub-network topology related to each community, represent the degree of membership of country i in community r. We remark that individual nodes can be members of different components, with different weights. That is, the non-negative factorization of the temporal network tensor can naturally capture overlapping communities.

Figure 3 shows the degree of membership of each node with respect to the three communities. Figure 4 and 5 display the hub and the authority scores for each country. Finally Figure 6 reports the activity pattern of each community during time. In all figures, a particular color is associated to a specific community and to the nodes' characteristics in therms of systemic importance inside that community. The first module is identified with blue color, the second in green and the third in red.

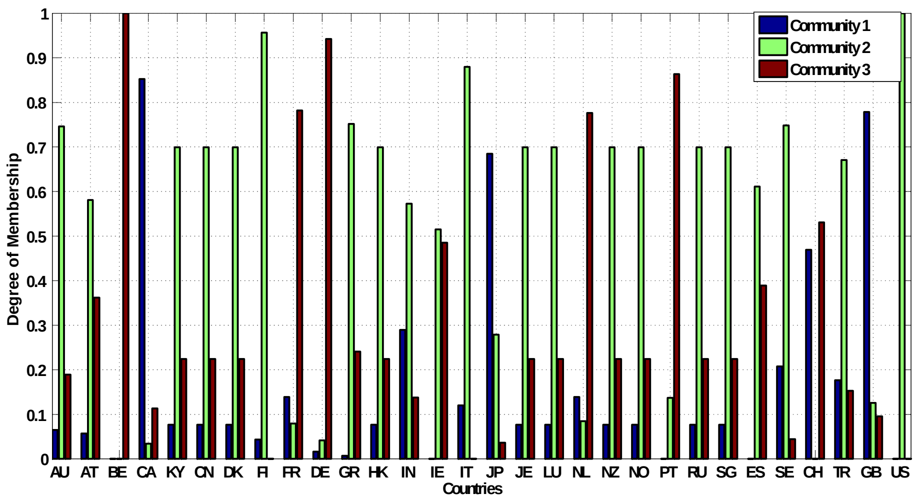

Figure 3: Degree membership of each country along the three communities

Figure 3 shows the fuzzy assignment (soft partition) of each country along the three communities. The crisp assignment (hard partition), in which the relationship between a node and a cluster is binary, could be directly inferred from Figure 3, as a special case, looking at the community for which each node has the maximum degree of membership.

With fuzzy assignment, each node is associated with communities in proportion to a membership weight between zero and one. From Figure 3 it becomes evident that the membership of most countries is unbalanced in the direction of one particular community; Belgium (BE) and The United States (US) are extreme examples belonging to only one community, but also Canada (CA), France (FR), Germany (DE), Portugal (PT) and the United Kingdom (GB) display strong membership associated to just one particular community. In other words, those countries have an unimodal membership distribution. On the other extreme, Ireland (IE), Switzerland (CH) and to a less extent Austria (AT) and Spain (ES) have a membership distribution that is closer to uniform, indicating membership to different communities at the same time.

Interpreting the clusters, we find that the found communities depend on the creditor's business model and on the geographical proximity between countries. The first community (denoted in blue) is related to the business banking model adopted in the United Kingdom, Japan (and also in Switzerland) where banks pool funds at major offices and redistribute them around the banking group. The second community encompasses most of the countries, with the United States, Finland and Italy having the strongest membership. Note that the community includes also most of the non-reporting countries in the BIS statistics. Finally, the third community is specific to countries belonging to the European Union. Most of them are core countries, such as France, Belgium, Germany, and the Netherlands. Note that Switzerland, with a banking model similar to the United Kingdom, also belongs to this community. Ireland and Spain, on the other hand, belong to both the second and the third community.

Figure 4 and 5 clarify the role of each country, as a borrower and as a lender inside each community. Since countries with high authority score are the main debtors and countries having high hub scores are the main creditors, hubness can be associated with country credit risk, being this measure related to number of outgoing links and representing claims on other countries. On the other hand, authoritativeness can be associated with funding risk, as countries with high authority scores are more prone to a sudden withdrawal of funds. Note that a country with a strong membership in one community can have an high hub/authority score inside another community. This feature means that the countries dry funds inside/outside its community redistributing them outside/inside the module.

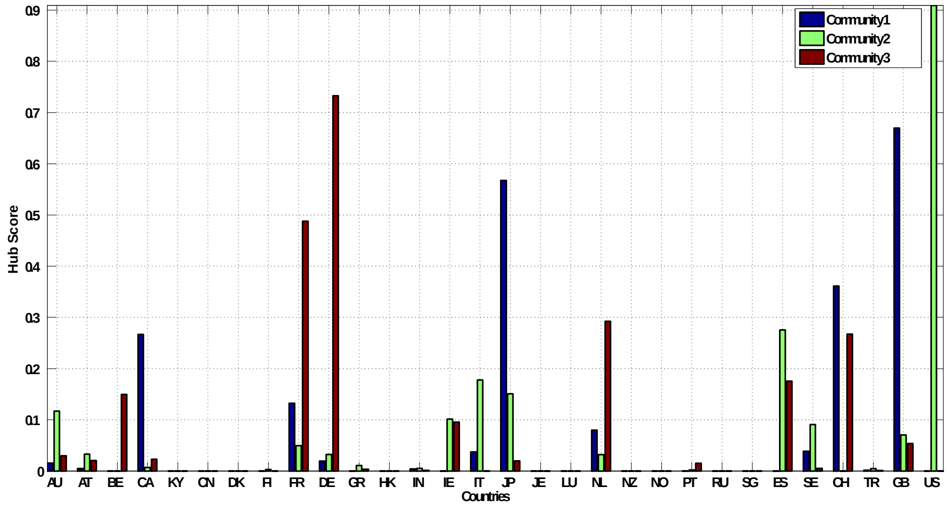

Figure 4: Hub score of each country inside each community

The United Kingdom (GB), Japan (JP), Switzerland (CH) and Canada (CA) stand out as the main hubs (see Figure 4) in the first community (blue), pointing to the United States (US) which is the main authority (see Figure 5). The high authority score associated with the United States reflects its large net debtor position, with the large current account deficits accumulated during the past decade and the role of their financial linkages in the transmission of the global financial crisis (Milesi-Ferretti, Strobbe and Tamirisa 2010). Overall this community points out the importance of funding risk for the United States, from the largest international hubs.

The second community (green) is more country specific. The highest hubness is associated with the United States (US), with Spain (ES) holding the second highest score, however relatively small. These countries point to the United Kingdom (GB), Japan (JP), Germany (DE) and France (FR). Overall, the existence of this community implies that credit risk shocks originating in the banking systems in the United States may have the potential to cause failures in the banking systems of most of countries under analysis, through the regional hubs GB,JP and DE/FR. In this perspective, risks arising from the previous countries may be more systematic than others. Indeed, systemic risks from shocks originated in the US, GB and European banking system appears to have increased in the recent crisis years (as Figure 7 shows).

The Third community (red) is specific to countries belonging to the European Monetary Union. Core countries such as Germany (DE), France (FR), the Netherlands (NL) hold the highest hub scores whereas the main authorities (in Figure 5) are the United States (US) together with the United Kingdom (GB), followed by peripheral countries such as Italy (IT) and Spain (ES). Overall, the third community reflect to a significant extent the cross-border activity of main Euro area banks conducted through their US and GB subsidiaries (Gyntelberg, McGuire and Von Peter 2009), or their peripheral ones. It shows the potential importance of credit risks shocks transmitted from the core countries to the periphery and to international authority nodes.

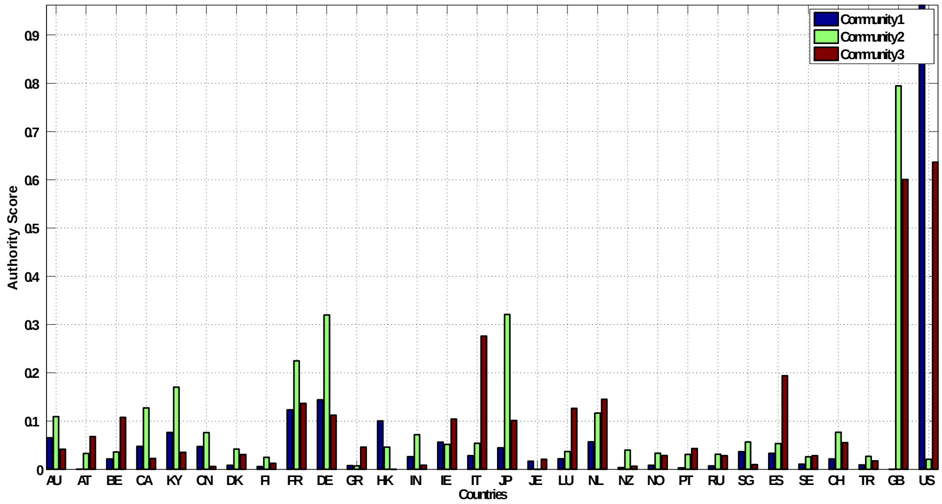

Figure 5: Authority score of each country inside each community

Since this technique is intrinsically temporal, the introduction of the third dimension gives us the opportunity to relate the temporal distributions of the links inside each communities with the different events that affected the behavior of the countries. A simple, yet interesting, economic interpretation, of the temporal activity patterns of the communities is that, an increase of this variable reflects the build-up of the systemic risk inside communities whereas its decrease is associated with the outbreak of financial crisis. For instance, the risk associated to banks located in the United states (green) was high in the first part of the sample, and the activity pattern fall of 2006 and 2008-09 reflects the subprime mortgage crisis and the subsequent global financial crisis. The risk associated to the third community (red) encompassing important EU countries, increases substantially in the second part of the sample abruptly decreasing when the European sovereign debt crisis erupted. Indeed, from the beginning of 2010, fears of a sovereign debt crisis developed among investors as a result of the rising private and government debt levels around the world together with a wave of downgrading of government debt in some European states. Finally the first community (blue) shows a more stable behavior. Moreover it has to be noticed that the activity level of all the communities displays a decreasing trend from the beginning of 2012 that outlines an unstable financial situation.

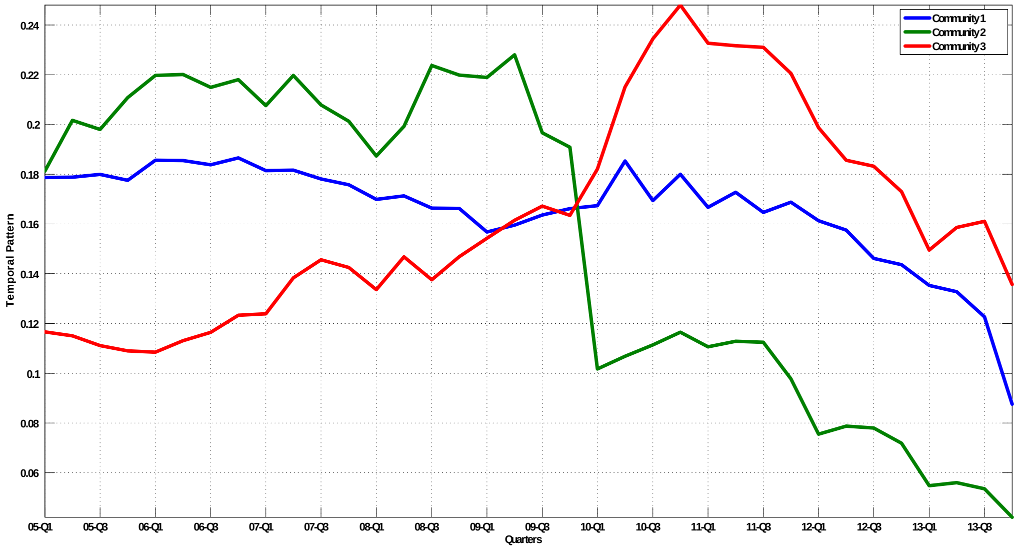

Figure 6: Activity patterns of the three communities along time

3.3 Model comparison

In this subsection we compare our algorithm with several well-known community detection methods. Since our method produces both soft and hard partition schemes, we compare the goodness of the communities obtained by both solutions against methods that produce crisp assignment (non-fuzzy) or fuzzy assignment. With crisp assignment, the relationship between a node and a cluster is binary. That is, a node i either belongs to cluster c or does not. With fuzzy assignment, each node is associated with communities in proportion to a belonging factor. Thus we compare our hard partition solution against methods that produce crisp assignment and the soft partition solution against methods that produce fuzzy assignment.

In particular we consider the modularity maximization method (Girvan and Newman 2002, Newman 2004) , the Louvain method (Blondel et al. 2008) and the K-means algorithm (MacQueen 1967) for crisp assignment; the C-means algorithm (Dunn 1967, Bezdek 2013) and the algorithms developed by Lancichinetti, Fortunato and Kertész (2009) and by Huang et al (2011) for fuzzy assignment.

Since the our algorithm is applied to networks for which the communities are not known in advance, we need a measure to quantify the goodness of the communities detected by each technique. In other word, we would like to know which of the divisions produced by the different algorithms are the best ones for the given network. To answer this question, we define two modularity measures that show the quality of a particular division of a network. For each time stamp, these two measures are the crisp and the fuzzy modularity for directed weighted network which are defined as:

$$Q^C = \frac{1}{m}\sum_{i,j}\left[\mathbf{W}_{ij} - \frac{\mathbf{s}_i^{in}\mathbf{s}_j^{out}}{m}\right]\delta\left(c_i, c_j\right)$$

$$Q^F = \frac{1}{m} \sum_c \sum_{i,j} \left[\mathbf{W}_{ij} - \frac{\mathbf{s}_i^{in} \mathbf{s}_j^{out}}{m} \right] \mathbf{a}_{ic} \mathbf{a}_{jc}$$

respectively. Where $\mathbf{s}^{in}$ and $\mathbf{s}^{out}$ are the in- and out-strength respectively, $m = \sum_i \mathbf{s}_i^{in} = \sum_j \mathbf{s}_j^{out}$. The difference between the two measures relies on the last term: $\delta(c_i, c_j)$ is the Kronecker delta symbol, and c_i (c_j) is the label of the community to which node i (j) is assigned; $\mathbf{a}_{ic}$ ($\mathbf{a}_{jc}$) is the degree of membership of node i (j) in the community c.

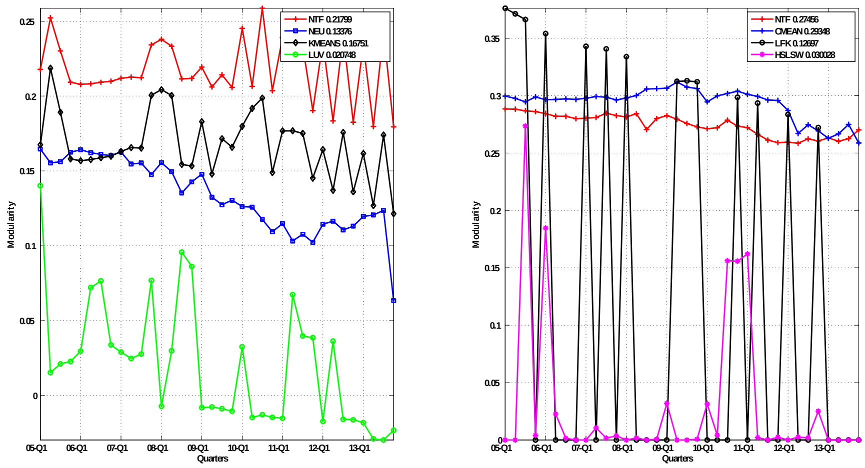

(a) (b)

Figure 7: The modularity measure of different algorithms for hard partition solutions (a) and for soft partition solutions (b). Near the names of each algorithm we report the average of the modularity values over the entire data sample. In the legends, NTF refers to our algorithm, Von Newman (NEW), K-means (KMEANS) and Louvain method (LUV) for hard partition solution. The soft partition solution of our method (NTF) is also compared with the C-means algorithm (CMEANS) with Lancichinetti, Fortunato, Kertész method (LKF) and with the Huang et al. method (HSLS)

Figure 7 reports the results for the different methods. We calculate the modularity metrics for each period in the data sample and in the legend, near the name of each algorithm, we report the average modularity value. Fig. 7 (a) encompasses the results about the hard partition solutions, while in the Figure 7 (b) we show the modularity for the soft partition solutions. In the first case, on average, our method outperforms the other algorithms. Figure 7 (b) on the other hand shows that, on average c-means outperform the method proposed in the paper, and in some period the other techniques provide a higher modularity.

4 Discussion

In this section we discuss the main implications of our findings, from an economic view-point and we show how our results are related to some stylized facts occurred in the last decade. As we stated in the introductory section, hubs and authorities have a clear interpretation in terms of risk associated to countries' banking system. We also found three different communities and the pictures above portray the basic framework for our results discussion.

In the first community, the high hubness score identified for the United Kingdom, Japan and Switzerland is specific for the business banking model adopted in such a countries. Here banks pool funds at major offices and redistributes them around the banking group, adn especially in the U.S. This behavior is specific to centralized funding models whereas decentralized banks let affiliates raise funds autonomously to finance assets in each location (McCauley, McGuire and Von Peter 2012). These countries, having foreign claims mainly denominated in US dollars, point to the United States which is the biggest authority (McGuire and Von Peter 2009). The time evolution of the statistics associated with this characteristic shows a stable trend all over the sample. This does not mean they have not been affected by the crisis but that their banking model is quite stable.

Indeed, Britain's 'Great Recession' started in the spring of 2008 and ended in the summer of 2009. The economy then raced out of recession faster than expected but faltered again late in 2010. But apart from this, the Bank of England shows no change in credit and interest rate exposures and UK banks maintained largely balanced net interbank US dollar position.

Japan, on the other hand, adopted a number of measures to overcome the crisis as well: a recent study (see Berkmen 2009) finds that the Bank of Japan's policy had a positive effect on economic activity, aided by the improvement in the banking sector and corporate deleveraging. In 2012, the Bank of Japan announced the introduction of the Stimulating Bank Lending Facility to provide unlimited long-term funds at 0.1% to financial institutions, beginning in June 2013.

Concerning the Swiss economy, the recent global financial crisis highlighted its vulnerability and that of the financial system of its two largest banks (UBS and Credit Suisse). The collapse in the market prices generated substantial losses, particularly for UBS that had substantial investments in US subprime mortgage-backed securities. By contrast, small and medium-sized Swiss banks weathered the crisis fairly well and were able to gain market share at the expense of the two large banks. However the Swiss economy recovered fairly quickly from the financial crisis, with a broad-based recovery driven by robust domestic demand and a rebound in exports. Given the lessons from the financial crisis, the Swiss authorities have concluded that it is crucial to strengthen the resilience of the systemically important banks in order to reduce the possibility of future bailouts and emergency assistance measures for them[6].

The second community is mostly related to country specific characteristics, and is quite relevant in terms of systemic risk. The largest hubs is the United States, having the largest credit risk against the United Kingdom, Japan and Germany, which are the countries with the highest authority score. The time-attribute scores related to those countries reveals that the risk associated to those banking system was high in the first part of the sample, reflecting the subprime mortgage crisis of 2006 and the subsequent financial crisis of 2008-2009. Between the second half of 2007 and the last quarter of 2008 many events occurred supporting the trend observed in the time attribute. By early November 2008, the S&P 500, was down 45 percent from its 2007 high. Housing prices had dropped 20% from their 2006 peak, with futures markets signaling a 30-35% potential drop. The 2008 crisis also led to the failure and bailout of a large number of banks and financial services firms in the US (Lehman Brothers, Washington Mutual Bank, Fanny Mae, Freddie Mac). On the contrary the second part of the sample influences less the measure, being the European countries the most contributors for systemic risk, as we stated above. Indeed starting from 2009, we observe a turnaround: the U.S. recession that began in December 2007 ended in June 2009, according to the U.S. National Bureau of Economic Research

[6] See the Financial Stability Board report at http://www.financialstabilityboard.org/publications

(NBER)[7] and the financial crisis appears to have ended about the same time. In April 2009 TIME Magazine declared "More Quickly Than It Began, The Banking Crisis Is Over".

The third community is mostly specific to continental countries belonging European Union, and it also bears a high potential of systemic risk transmission. It contains countries such as France, Germany and the Netherlands[8] are the main hubs, sharing cross-border activity with the two largest banking centers (Von Peter 2010): the United States and the United Kingdom, which turn out to be main authorities, followed by Italy and Spain. The temporal activity of this community decreased substantially in the second part of the sample when the European sovereign debt crisis erupted. Indeed, from the beginning of 2010, fears of a sovereign debt crisis developed among investors as a result of the rising private and government debt levels around the world along with a wave of downgrading of government debt in some European states. The elevated systemic risk of European banks during the sovereign debt crisis - reaching its peak at the end of 2010 - was largely due to increased default risk. Systemic risk quickly rose with the Greek bailout agreement in May 2010 as the European sovereign debt crisis disentangled (Black et al. 2013). The default probabilities of European banks fall in the second half of 2011 which may be attributable to additional liquidity injections from the European Central Bank (ECB), and this trend has been continuing up to the end of the sample.

Summing up, our work suggests that differences in banking systems and funding models could explain why some characteristics are associated to specific countries. Additionally, for some European banking systems, foreign claims are primarily denominated in the home country (or "domestic") currency, representing intra-euro area cross border positions (e.g. Belgian, Dutch, French and German banks). For others (e.g. Japanese, Swiss and UK banks), foreign claims are predominantly in foreign currencies, mainly US dollars. Moreover these characteristics map onto vulnerability to funding disruptions, measuring banks' cross-currency funding.

The results suggest a possible interpretation of the role played by the United Kingdom and the United States. These countries act as risk transfer, being authorities in one community and hubs in an other. Indeed capital flows saw sudden stops and a reversal in the wake of the recent crisis, in particular out of the United States.

In the first years of the sample, banks facilitated international capital flows out of Japan and the euro area. Banks routed these funds via offices in the United Kingdom, ultimately transferring them to borrowers in the United States. After the start of the crisis, the direction of many of the bilateral flows reversed, in part generated by capital movements back to the United Kingdom, and in part reflecting asset writedowns. Between the second quarter of 2007 and fourth quarter of 2008, the cumulative net flows from the United States to the United Kingdom totaled $482 billion. In contrast to the US-UK bilateral pair, the net flow of funds between Japan and the United States overall did not change direction. Throughout the crisis, banks in Japan have continued to channel money to US non-banks (see McGuire and Von Peter 2009).

5 Concluding remarks

In this paper we have introduces a new approach to simultaneously detect network community structures and systemically important nodes inside each community, tracking their common activity over time. We have shown how the introduction of the temporal dimension gives the opportunity to investigate the activity level of the communities as a function of time. Our approach is based on tensor decomposition, where a temporal network is naturally represented

[7] See the Business Cycle Dating Committee, National Bureau of Economic Research at http://www.nber.org/cycles/sept2010.pdf

[8] Switzerland has also a high hub score but its relative importance is lower than the one in the first community.

as a time-ordered sequence of adjacency matrices, each one describing the state of the financial network at a given point in time. We have applied the tensor decomposition technique to the consolidated banking statistics (CBS) provided by the Bank of International settlements (BIS), which include countries bilateral claims measured on quarterly basis, from the fist quarter of 2005 to the last quarter of 2013.

We found that a country can be systemically important because of its role as a borrower and also due to its role as a lender. Countries with high authority score face the risk of possible sudden funding withdrawal, while countries having high hub scores deal with a possible borrower default. BIS data is well represented by three communities. The first one is related to the business banking model adopted in the United Kingdom, Japan and Switzerland. These countries point to the United States and the time evolution of the statistics associated with this community indicate a stable model. The second community is mostly country specific, and especially associated with the United States, having the largest credit risk against the United Kingdom, Japan and Germany. Time scores reveal that the temporal pattern of this community is high in the first part of the sample, reflecting the subprime mortgage crisis of 2007 and the subsequent financial crisis of 2008-2009, indicating potential systemic risks arising from funding risk originating from the United States and, indirectly, from regional hubs such as UK, JP and DE.

The last community refers to the Euro area: France, Germany and the Netherlands share cross-border activity with the United States, the United Kingdom and peripheral countries such as Italy and Spain. The temporal activity associated to this community increased substantially when the European sovereign debt crisis erupted. This community indicates potential systemic risk arising from credit risk in the global authorities US and Uk and also in large European peripheral countries.

Overall, our results clearly shows that a promising way to analyze complex inter-relations is by tensorial factor analysis. The same technique could be applied to more disaggregated dataset to get a deeper understanding of the economic interactions among financial institutions.

References

[1] Soramäki, K., Bech, M. L., Arnold, J., Glass, R. J., & Beyeler, W. E. (2007). The topology of interbank payment flows. Physica A: Statistical Mechanics and its Applications, 379(1), 317-333.

[2] Iori, G., De Masi, G., Precup, O. V., Gabbi, G., & Caldarelli, G. (2008). A network analysis of the Italian overnight money market. Journal of Economic Dynamics and Control, 32(1), 259-278.

[3] Iori, G., Reno, R., De Masi, G., & Caldarelli, G. (2007). Trading strategies in the Italian interbank market. Physica A: Statistical Mechanics and its Applications, 376, 467-479.

[4] Cocco, J. F., Gomes, F. J., & Martins, N. C. (2009). Lending relationships in the interbank market. Journal of Financial Intermediation, 18(1), 24-48.

[5] Craig, B., & Von Peter, G. (2014). Interbank tiering and money center banks. Journal of Financial Intermediation.

[6] Fricke, D. (2012). Trading strategies in the overnight money market: Correlations and clustering on the e-MID trading platform. Physica A: Statistical Mechanics and its Applications, 391(24), 6528-6542.

[7] Fricke, D., & Lux, T. (2013). On the distribution of links in the interbank network: Evidence from the e-mid overnight money market (No. 1819). Kiel Working Paper.

[8] Malliaros, F. D., & Vazirgiannis, M. (2013). Clustering and community detection in directed networks: A survey. Physics Reports, 533(4), 95-142.

[9] Xie, J., Kelley, S., & Szymanski, B. K. (2013). Overlapping community detection in networks: The state-of-the-art and comparative study. Acm computing surveys (csur), 45(4), 43.

[10] Board, F. S. (2013). Global systemically important banks: updated assessment methodology and the higher loss absorbency requirement. Bank for International Settlements, Basel.

[11] Battiston, S., Puliga, M., Kaushik, R., Tasca, P., & Caldarelli, G. (2012). Debtrank: Too central to fail? financial networks, the fed and systemic risk. Scientific reports, 2.

[12] Perra, N., & Fortunato, S. (2008). Spectral centrality measures in complex networks. Physical Review E, 78(3), 036107.

[13] Bonacich, E. (1972). A theory of ethnic antagonism: The split labor market. American sociological review, 547-559.

[14] Bonacich, P., & Lloyd, P. (2001). Eigenvector-like measures of centrality for asymmetric relations. Social networks, 23(3), 191-201.

[15] Katz, E., Blumler, J. G., & Gurevitch, M. (1973). Uses and gratifications research. The Public Opinion Quarterly, 37(4), 509-523.

[16] Brin, S., & Page, L. (2012). Reprint of: The anatomy of a large-scale hypertextual web search engine. Computer networks, 56(18), 3825-3833.

[17] Kleinberg, J. M. (1999). Authoritative sources in a hyperlinked environment. Journal of the ACM (JACM), 46(5), 604-632.

[18] Cao, X., Wang, X., Jin, D., Cao, Y., & He, D. (2013). Identifying overlapping communities as well as hubs and outliers via nonnegative matrix factorization. Scientific reports, 3.

[19] Gauvin L, Panisson A, Cattuto C (2014) Detecting the Community Structure and Activity Patterns of Temporal Networks: A Non-Negative Tensor Factorization Approach. PLoS ONE 9(1): e86028. doi:10.1371/journal.pone.0086028

[20] Chen, B., Petropulu, A. P., & De Lathauwer, L. (2002). Blind identification of convolutive MIMO systems with 3 sources and 2 sensors. EURASIP Journal on Applied Signal Processing, 2002(1), 487-496.

[21] Comon, P. (2002). Tensor Decompositions, State of the Art and Applications. Mathematics in Signal Processing V, 1-24.

[22] De Lathauwer, L., De Moor, B., & Vandewalle, J. (2000). A multilinear singular value decomposition. SIAM journal on Matrix Analysis and Applications, 21(4), 1253-1278.

[23] De Lathauwer, L., De Moor, B., & Vandewalle, J. (2001). Independent component analysis and (simultaneous) third-order tensor diagonalization. Signal Processing, IEEE Transactions on, 49(10), 2262-2271.

[24] Vasilescu, M. A. O., & Terzopoulos, D. (2002). Multilinear analysis of image ensembles: Tensorfaces. In Computer VisionÑECCV 2002 (pp. 447-460). Springer Berlin Heidelberg.

[25] Vasilescu, M. A. O., & Terzopoulos, D. (2002, August). Multilinear image analysis for facial recognition. In Pattern Recognition, International Conference on (Vol. 2, pp. 20511-20511). IEEE Computer Society.

[26] Beylkin, G., & Mohlenkamp, M. J. (2002). Numerical operator calculus in higher dimensions. Proceedings of the National Academy of Sciences, 99(16), 10246-10251.

[27] Beylkin, G., & Mohlenkamp, M. J. (2005). Algorithms for numerical analysis in high dimensions. SIAM Journal on Scientific Computing, 26(6), 2133-2159.

[28] Acar, E., Camtepe, S. A., Krishnamoorthy, M. S., & Yener, B. (2005). Modeling and multiway analysis of chatroom tensors. In Intelligence and Security Informatics (pp. 256-268). Springer Berlin Heidelberg.

[29] Acar, E., ,Amtepe, S. A., & Yener, B. (2006). Collective sampling and analysis of high order tensors for chatroom communications. In Intelligence and security informatics (pp. 213-224). Springer Berlin Heidelberg.

[30] Kolda, T. G., & Bader, B. (2006, April). The TOPHITS model for higher-order web link analysis. In Workshop on link analysis, counterterrorism and security (Vol. 7, pp. 26-29).

[31] Bonacina, F., D'Errico, M., Moretto, E., Stefani, S., & Torriero, A. (2014). A Multiple Network Approach to Corporate Governance. arXiv preprint arXiv:1401.4387.

[32] J. D. Carroll and J. J. Chang, Analysis of individual differences in multidimensional scaling via an N-way generalization of "Eckart-Young" decomposition, Psychometrika, 35 (1970), pp. 283–319.

[33] R. A. Harshman, Foundations of the PARAFAC procedure: Models and conditions for an "explanatory" multi-modal factor analysis, UCLA Working Papers in Phonetics, 16 (1970), pp. 1–84. Available at http http://publish.uwo.ca/harshman/wppppfac0.pdf.

[34] Cichocki A, Phan AH, Zdunek R (2009) Nonnegative Matrix and Tensor Factorizations: Applications to Exploratory Multi-way Data Analysis and Blind Source Separation. Chichester: Wiley.

[35] 11. Mørup M (2011) Applications of tensor (multiway array) factorizations and decompositions in data mining. Wiley Interdisc Rew: Data Mining and Knowledge Discovery 1: 24–40.

[36] Kolda TG, Bader BW (2009) Tensor decompositions and applications. SIAM Rev 51: 455–500.

[37] Shashua A, Hazan T (2005) Non-negative tensor factorization with applications to statistics and computer vision. In: Proceedings of the 22nd international conference on Machine learning. ICML'05, 792–799.

[38] Lee, D. D., & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. Nature, 401(6755), 788-791.

[39] Nickel M, Tresp V, Kriegel HP (2011) A three-way model for collective learning on multi-relational data. In: Getoor L, Scheffer T, editors, Proceedings of the 28th International Conference on Machine Learning (ICML-11). New York, NY, USA: ACM, ICML '11, 809–816.

[40] Wang F, Li T, Wang X, Zhu S, Ding C (2011) Community discovery using nonnegative matrix factorization. Data Min Knowl Discov 22: 493–521.

[41] Bro, R., & Kiers, H. A. (2003). A new efficient method for determining the number of components in PARAFAC models. Journal of chemometrics, 17(5), 274-286.

[42] Tucker, L. R. (1966). Some mathematical notes on three-mode factor analysis. Psychometrika, 31(3), 279-311.

[43] Psorakis, I., Roberts, S., Ebden, M., & Sheldon, B. (2011). Overlapping community detection using bayesian non-negative matrix factorization. Physical Review E, 83(6), 066114.

[44] Mankad, S., & Michailidis, G. (2013). Structural and functional discovery in dynamic networks with non-negative matrix factorization. Physical Review E, 88(4), 042812.

[45] Milesi-Ferretti, G. M., Strobbe, F., & Tamirisa, N. T. (2010). Bilateral financial linkages and global imbalances: a view on the eve of the financial crisis. International Monetary Fund.

[46] Gyntelberg, J., McGuire, P., & von Peter, G. (2009). Highlights of international banking and financial market activity. BIS Quarterly Review, 19.

[47] Girvan, M., & Newman, M. E. (2002). Community structure in social and biological networks. Proceedings of the national academy of sciences, 99(12), 7821-7826.

[48] Newman, M. E. (2004). Analysis of weighted networks. Physical Review E, 70(5), 056131.

[49] Blondel, V. D., Guillaume, J. L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. Journal of statistical mechanics: theory and experiment, 2008(10), P10008.

[50] MacQueen, J. (1967, June). Some methods for classification and analysis of multivariate observations. In Proceedings of the fifth Berkeley symposium on mathematical statistics and probability (Vol. 1, No. 14, pp. 281-297).

[51] Dunn, J. C. (1973). A fuzzy relative of the ISODATA process and its use in detecting compact well-separated clusters

[52] Bezdek, J. C. (2013). Pattern recognition with fuzzy objective function algorithms. Springer Science & Business Media.

[53] Huang, J., Sun, H., Liu, Y., Song, Q., & Weninger, T. (2011). Towards online multiresolution community detection in large-scale networks. PloS one, 6(8), e23829.

[54] Lancichinetti, A., Fortunato, S., & Kertész, J. (2009). Detecting the overlapping and hierarchical community structure in complex networks. New Journal of Physics, 11(3), 033015.

[55] McCauley, R., McGuire, P., & Von Peter, G. (2012). After the global financial crisis: From international to multinational banking? Journal of Economics and Business, 64(1), 7-23.

[56] McGuire, P., & Von Peter, G. (2009). The US dollar shortage in global banking. BIS Quarterly Review, March.

[57] Berkmen, S. P. (2012). Bank of Japan's quantitative and credit easing: Are they now more effective? Washington, DC : Internat. Monetary Fund.

[58] Von Peter, G. (2010). International banking centres: a network perspective. Research on global financial stability: the use of BIS international financial statisti cs. CGFS Papers, (40), 71-82.

[59] Black, L., Correa, R., Huang, X., & Zhou, H. (2013). The systemic risk of European banks during the financial and sovereign debt crises. Federal Reserve Board.