IMT SCHOOL FOR ADVANCED STUDIES LUCCA

Three essays on the applications of multiplex networks in economics

PhD in Institutions, Markets and Technologies

Curriculum in Economics, Management and Data Science

XXXI Cycle

By

Giovanni Bonaccorsi

2020

The dissertation of Giovanni Bonaccorsi is approved.

Program Coordinator:
Prof. Dr. Pietro Pietrini, IMT School for Advanced Studies Lucca

Supervisors:
Prof. Dr. Massimo Riccaboni, IMT School for Advanced Studies Lucca
Prof. Dr. Giorgio Fagiolo, Sant'Anna School of Advanced Studies

The dissertation of Giovanni Bonaccorsi has been reviewed by:

Prof. Dr. Stefano Schiavo, University of Trento
Prof. Dr. Lucia Tajoli, Polytechnic University of Milan
Dr. Zhen Zhu, University of Greenwich

IMT School for Advanced Studies, Lucca
2020

A nonna Osvalda,
la mia prima prof

Contents

List of Figures

List of Tables

Acknowledgements

This thesis is based on the work of the author and coauthors as follows:

- Chapter 1 is the reproduction of the paper *Country centrality in the international exchange multiplex* published in the journal Applied Network Science, co-authored with Massimo Riccaboni (IMT School for Advanced Studies Lucca), Giorgio Fagiolo (Sant'Anna School of Advanced Studies) and Gianluca Santoni (CEPII)
- Chapter 2 is a joint work with Massimo Riccaboni (IMT School for Advanced Studies Lucca) and Giorgio Fagiolo (Sant'Anna School of Advanced Studies).
- Chapter 3 is a joint work with Massimo Riccaboni (IMT School for Advanced Studies Lucca).

This dissertation would not have been possible without the help of many persons. I would like to thank my supervisors, Massimo Riccaboni and Giorgio Fagiolo, for the helpful advices they have given me during the the compilation of this work. Together with them I would also to thank Gianluca Santoni who has helped me at the start of this project.

I'm thankful also to the INET community in Oxford which has welcomed me for a (too) short visiting period and has provided a many helpful feedbacks on my work and a lot of inspiration for the future. Among them I would like to thank especially Marco Pangallo for all the help he has given me.

Finally this work would not have been possible without the infinite patience of my parents, the company of my brothers and sister, and the continuous support of Chiara, who has saved me countless times from going crazy. To them goes all my love.

Vita

September 18, 1990	Born, Pisa, Italy

Education

2012	Bachelor's Degree in Economics Final mark: 110/110 cum laude University of Pisa, Pisa, Italy
2015	Master's Degree in Economics Final mark: 110/110 cum laude University of Pisa and Scuola Superiore Sant'Anna, Italy

Experience

2018	Research collaborator IMT School of Advanced Studies Lucca, Italy

Publications

1. Bonaccorsi, Giovanni, Massimo Riccaboni, Giorgio Fagiolo, and Gianluca Santoni. "Country Centrality in the International Multiplex Network." Applied Network Science 4, no. 1 (December 2019): 126. `https://doi.org/10.1007/s41109-019-0207-3`.

Presentations

1. G. Bonaccorsi, "Finding the most central countries in the international exchange multiplex using multirank," at *Complex Network 2018, University of Cambridge*, Cambridge, United Kingdom, 2018.
2. G. Bonaccorsi, "The effect of the network structure of the economy on the performance of countries during the 2008-2009 crisis," at *YETI Meeting 2019, University of Siena*, Siena, Italy, 2019.
3. G. Bonaccorsi, "The effect of the network structure of the economy on the performance of countries during the 2008-2009 crisis," at *ARS Workshop 2019, University of Salerno*, Vietri Sul Mare (SA), Italy, 2019.

Abstract

In the last years network theory has seen several theoretical advancements and an increasing number of interesting applications in various fields of knowledge such as in social, biological, human and economic networks.

The use of network results in economics has led to fruitful developments in the theory of trade, of the economic effect of migration and of financial distress contagion. Moreover, in agent based modelling, a network structure is often employed as foundation for the behaviour of agents. Hence it has been demonstrated that the applications of network findings to different economic models can lead to new discoveries, showing that economic phenomena may obtain interesting explanations when network diffusion processes are taken into consideration.

However, the main economic applications of network theory are often limited to single layer network results, where the networks employed represent one single type of relationship among the nodes and if more layers are analysed, they are considered independent. On the contrary an increasing number of publications by leading network scholars is focused on studying multilayer networks, where the same nodes have different types of links between them and their respective interdependence is recognized and studied. As a consequence, many of the single layer network concepts have been generalized to multilayer networks, improving previous analysis by adding the possibility to study different types of relations in an organic manner.

In economics, in particular, network data regarding multiple relations among world countries has been employed over

time but only recently the focus has shifted towards a more systematic approach. The first contribution of the present work is the harmonization of the majority of these sources in a consolidated dataset, the first merging together information from different fields: from flows of goods, to flows of financial contracts, to flows of people, to flows of citations. The final dataset spans over 40 years and 211 countries and reaches, in the more rich cross sections, 19 layers of data (ignoring duplicated and redundant sources). Since nodes are common across all layers the particular type of multilayer network we are using is a multiplex.

In our first study on this new dataset we have measured the centrality of countries over time. We have identified two cross sections of layers, the years 2003 and 2010 (before and after the Great Financial Crisis), where the majority of the sources was present. Then we have harmonized the data filtering out excessive differences among the layers. Finally, we have applied on the dataset two recent multilayer algorithms which have generalized two of the most common centrality measures. The first is the MultiRank, the multilayer generalization of the PageRank algorithm, the second is the MD-HITS (MultiDimensional HITS) which generalizes the hubs and authority algorithm. Both the algorithms have been used to rank the importance of page results on the web and they highlight different features of the nodes: the first one refers to the property of webpages of being linked (cited) from other important ones, the second one instead is related to the status of a page as an important source of information (an authority) or as an important hub redirecting to authority pages. The interesting feature of both the multilayer generalizations of the algorithms is that they produce automatically two types of rankings: one for the nodes of the multiplex and one for the layers. This allows us to also identify which are the sources of importance of a certain country in the whole multiplex in an unsupervised manner. To obtain a measure of the relevance of these new methods we have compared the ranking of nodes obtained using the multiplex centrality measures with the ranking of countries by per capita GDP. We have found similarity in the rankings but not perfect correlation, signalling that our new dataset may contain some additional source of information to be exploited in explaining country development.

After measuring country centrality, the second research question we have addressed with the aid of our new data sources regards the Great Financial Crisis. The collapse of the world financial markets in 2009, symbolically kick-started by the default of Lehman Brothers, made clear that economic theories were missing something, otherwise a crisis so deep and pervasive would have been avoided. One of the streams of research originated by this event is tightly related with network theory and it is the study of the propagation of contagious phenomena over networks, in particular financial distresses. However, contagion models are mostly theoretical and the empirical evidence on financial contagion is still scarce. Moreover, the econometric studies on financial crisis have yet to find a consistent and persistent explanation of why some countries are more affected than others during these events. Finally multilayer studies in this field are still rare.

In this work we have used as starting point a consolidated set of evidences obtained in Feldkircher (2014). From a set of 95 economic and financial measures regarding world countries, they found only one which was significantly present in every model when trying to explain why countries have had different performances after the GFC: the growth of credit supply from domestic banks. Starting from this element we have integrated their analysis with a set of network variables obtained from each of our layers regarding topological features of the networks such as centrality (both at single and multilayer level), clustering and community structure. We have used their same methodology, Bayesian Model Averaging, to solve the issue of model selection and avoid bias in selecting the explanatory variables. With our final results we have improved on Feldkircher 2014 by finding a new variable which is consistently present in the majority of the analysed models: the kcore centrality of the investment layer. This result is important both because it confirms the relevance of network variables as explanatory candidates for economic models and because it introduces a new explanation for the different performance of countries after the crisis.

Our last research question regards network embeddings and their use to predict missing links. In our dataset we have missing information due to unreported or censored data, to reconstruct it we have used the information available from the known part of the network to obtain predictions on the existence of the unobserved part. This is achieved employing the method of network embeddings both at single and multiple layer level: an embedding of a network is a mapping of the rich structure of the graph at node level to a lower-dimensional latent space where projections of nodes are optimized to be closer when they map to closer relations at graph level. By doing so networks can be used as features for machine learning tasks in a very flexible way. Among the network embeddings literature we have seen a recent development of several multilayer methods among which we have found the scalable multilayer network embeddings method (MNE, D. Zhang et al. (2018)) to be our best option for making predictions. By pairing the MNE binary prediction to the method of weighted stochastic block models (Peixoto, 2018a) to assign a weight to links we have predicted missing links in all the multiplex layers. Our results show that a certain level of reconstruction can be achieved, even though with wide variability by layer.

To conclude, by answering these three research question we have shown how network measures can be of great help to improve the analysis of economic issues and in particular how the integration of different data sources mapping relation among countries can alter vividly the picture of the world that we have.

Chapter 1

Country centrality in the international multiplex network

The content of this chapter was published in the journal Applied Network Science as Bonaccorsi et al. (2019)

1.1 Introduction

Over the last few years we have witnessed an increase in the use of network models in economics, finance and business studies: details coming from the topology of specific networks, either on micro level of social interactions among individuals and collaborations between firms, or at a macro level of international economics, help scholars identify new phenomena and clarify how they diffuse and distribute (Easley and J. Kleinberg, 2010).

Usually, though, one network is examined at a time (i.e. trade) or the relationship between two networks is investigated (i.e. migration and trade). Network features are frequently used as explanatory or descriptive variables in gravity models of international activities. For instance, we have several studies on the effect of the centrality of countries in the trade and migration networks on the magnitude of countries' exports (Fagiolo and Mastrorillo, 2014, Sgrignoli et al., 2015, Metulini et al., 2018). Occasionally, other networks have been considered in the literature, such as the network of foreign direct investments, human mobility, information, knowledge and financial flows.

On the one hand, country centrality in networks has been used in the study of international relations as a proxy of power (Hafner-Burton, Kahler, and Montgomery, 2009). On the other hand, new developments in multilayer network analysis suggest that international studies may greatly benefit from jointly analyzing multiple international relations in a unified framework (De Domenico, Solé-Ribalta, et al., 2015).

In this Chapter we study the international network of countries by looking at several dimensions of development from a multilayer network perspective. We assemble a multiplex network of 112 countries as nodes and their connections across 19 heterogeneous layers, observed at two moments in time before and after the Great Recession (2003 and 2010). We compute the centrality of countries in the multilayer network by using the multilayer generalization of the PageRank algorithm (Iacovacci and Bianconi, 2016) – MultiRank from now on – and of the hubs and authority score (Arrigo and Tudisco, 2018) – MD-HITS from now on.

The ranking we obtain based on MultiRank centrality is consistent with the common north-south divide, that is to say the division of countries among developed and underdeveloped, and depicts a positive correlation with relevant economic variables such as GDP per capita purchasing power parity (PPP) (pcGDP from now on). This configuration of MultiRank gives more importance to those layers which are commonly considered as relevant for the performance of countries, predominantly trade and financial networks. Moreover it is very similar to the ranking selected by the MD-HITS algorithm, with the latter actually having an even better fit with relevant macroeconomic variables.

Our work sheds new light on the relationship between position and centrality of countries in the international multiplex network and their development by analyzing, for the first time, a vast set of international relations: trade, finance, transportation, human mobility and migrations, information and knowledge flows, international alliances. Increasing data availability coupled with the recent advancements in the analysis of multiple networks open up new possibilities for the analysis of complex global phenomena and new dimensions can be added to our framework to cover an even broader set of international linkages.

The rest of the Chapter is structured as follows: section 1.2 describes how our work fits in the extant literature, section 1.3 contains a detailed description of the data we have collected, section 1.4 explains the methodology we apply and finally section 1.5 illustrates our results. Finally the last section discusses our contribution to the literature and future work.

1.2 Related literature

Following previous applications of network theory in macroeconomics, we define a single layer economic network $G(N, E)$ as a set of nodes N representing countries, coupled with a set of edges E describing the relations between them. In so doing we follow a long tradition of studies in trade (Chaney, 2014), migration (Fagiolo and Mastrorillo, 2014), international aid (Kali, Horowitz, and Song, 2017), banking (Minoiu and J. A. Reyes, 2013, S. Battiston et al., 2012; Bardoscia et al., 2015), mergers and acquisitions (Campi et al., 2016). Usually, the edges between nodes are weighted and directed, representing the amount of monetary exchanges from a source country to a target country.

However, there are some international networks in which directed edges represent relationships other than monetary exchanges, such as the networks of migration and human mobility (Sgrignoli et al., 2015; Fagiolo and Santoni, 2016; Fagiolo and Santoni, 2015; Fagiolo and Mastrorillo, 2014; Riccaboni, Rossi, and Schiavo, 2013).

Sometimes edges may also be undirected (symmetric) and unweighted. A stream of literature has investigated international infrastructures including, among others, the international airport transportation network (Colizza et al., 2006), the trade shipping network (Kaluza et al., 2010) and the overseas telephone and fiber cables linkage (Rossello, 2015) . Intangible networks of knowledge flows have also been studied by looking at patent citations (Hall, Jaffe, and Trajtenberg, 2001) and international collaboration among scholars (Pan, Kaski, and Fortunato, 2012). Finally, international relationships are shaped by the networks of diplomatic relationships, such as alliances, trade agreements and intergovernmental organization memberships, as reported in the "Correlates of war" project (Bayer, 2006; Pevehouse, Nordstrom, and Warnke, 2003; Gibler and CQ Press, 2009; Sousa and Lochard, 2011)

Building upon previous studies, which have analyzed the different dimensions of international connectivity, and thanks to multiple available data sources, in this paper we reconstruct the first comprehensive multilayer network among countries: instead of analyzing each of the layers in isolation, focusing on standard network statistics and comparing them, we analyze the structure of international networks considered as a whole. Following Kivela et al., 2014 we define our multilayer network as a multiplex, i.e. a network where no link connects directly the layers but all the layers share the same set of nodes, e.g. countries. Therefore, a multiplex is a colored network where edges of different colors represent different types of international relationships (i.e. trade, migration, investments, knowledge flows etc) between countries.

In recent years multilayer networks have become the subject of several scholarly works aimed at generalizing network statistics and algorithms from single to multiple layer analysis (Lee et al., 2012; De Domenico, Nicosia, et al., 2015; Kivela et al., 2014; F. Battiston, Nicosia, and Latora, 2014; Boccaletti et al., 2014; Bródka Piotr et al., 2018; De Bacco et al., 2017; Aleta and Moreno, 2019).

Among the numerous methods which have been proposed to measure the centrality of nodes in a multilayer network we have selected two measures which can be computed for weighted and directed multilayer networks and do not require any pre-defined ranking of the importance of network layers (i.e. different types of international relationships), but provide it as an output.

The first is the generalization to multiplex network of the PageRank centrality of nodes (Brin and Page, 1998). According to it, nodes are ranked higher if they receive links from other important nodes: hence it focuses on the attractiveness of nodes only. To apply it to our multilayer network we closely followed a set of recently published papers (Rahmede et al., 2018; Halu et al., 2013; Iacovacci, Rahmede, et al., 2016; Iacovacci and Bianconi, 2016)[1] which have developed a way of defining the attractiveness of nodes by taking into account also the relevance of each layer in the whole multiplex. In the multiplex version of the Pagerank nodes are rewarded for their capacity of receiving links from other central nodes in more relevant layers.

Our second choice is the hubs and authority algorithm, also known as HITS, Hyperlink-Induced Topic Search, (J. M. Kleinberg, 1999). Referring to the problem of getting access to information, it defines as authorities those nodes which hold essential information to the user whereas hubs are those nodes which redirect the user to the authorities.[2] In our case some countries act as brokers among important partners whereas other countries acquire importance for being connected to many hubs. To apply this algorithm to our data we follow the work by Arrigo and Tudisco (2018) where not only nodes are ranked but also layers receive a score, ranking their capacity of linking important layers or being linked by hubs layers. [3]

One challenge we face is to analyze a multiplex of heterogeneous layers to characterize the relevance of countries over a set of different features (flows of goods, services, investments, knowledge, human mobility and so on).[4]. So far, only multiplexes of homogeneous layers have been considered, where each layer represents a different typology of the same relation (different trade commodities in the trade network, different airlines in the airport network and so on). To deal with this issue we have filtered our data both by node, i.e. by deciding which nodes were con-

[1]Specifically we apply the last version of the algorithm where both layer relevance and node centrality are calculated at each iteration and for multiplexes of many layers.

[2]The classification reflects how information is searched and found: it is difficult to find the right answer at first try, hence there are nodes which collect information and redirect the user to the right destination gaining in centrality for this reason. Hence to have a high authority score a node must have links from nodes with high hub score, whereas higher hub score is awarded to nodes linking to high authorities nodes.

[3]More details on this are provided in section 1.4.

[4]By doing so our work is related to social network studies of international relationships such as Smith and White, 1992.

stantly measured over time, and by edge, i.e. by selecting only relevant edges and discarding less probable ones.

The position and role of countries in the multiple network of international relations is then compared to pcGDP as a measure to summarize the relative economic performance of countries. By doing so our work contributes also to the ongoing debate on the appropriateness of pcGDP as a measure of the well being of populations (Sen, 1985; Dasgupta and Mäler, 2000; Stevenson and Wolfers, 2013; Lange, Wodon, and Carey, 2018). While on one side it is obvious to see that pcGDP is an excessively synthetic measure which discards many details and should never used as the only indicator of development, on the other side there have been many different proposals for substitutes of pcGDP, all revolving around what one should consider as the right measure of growth for countries (Costanza et al., 2009), but without a clear consensus on which should be used. As consequence the trade-off between accuracy and simplicity for a world development index is yet to be solved: while most scholars agree on the shortcomings of pcGDP it still represents the best candidate to summarize efficiently the position of a country in the world landscape.

Instead of trying to solve this trade-off we recognize the informative content of pcGDP and use it as reference for the capacity of our new measures to capture the overall state of development of world countries. We will consider good centrality measures those which can summarize in a single index the information content of all the layers in our dataset with the same efficiency of pcGDP and with a sufficient correlation with it. Being the content of the layers heterogeneous and not strictly related to pcGDP (even more given the network nature of our data) we consider a positive sign that multilayer centrality measures can reconstruct a ranking of countries similar to a widely used development index.

However we are not concerning ourselves with looking for perfect correlation with pcGDP, first because of its well-known shortcomings and secondly because a complete matching will also deprive our measures of any new information content. Hence we will consider the differences in the rankings of countries between pcGDP and multilayer centralities as the proof that we are measuring dimensions of development which are not included in pcGDP and hence may be worth exploring.

We are aware that there is not a clear threshold which is able to quantify how much our measures should be similar or different with respect to pcGDP and hence we cannot fully validate our results. To do so we should require further evidences to which test both rankings against and this would also imply to choose a specific framework to evaluate our new well-being measures. We think that this task is outside the scope of this paper and we regard our contribute not as definitive conclusion of the GDP debate but as a first step toward a better understanding of which network measures should be included (and how) in order to fully account for the development of a certain country.

1.3 Data

As acknowledged in the previous section, this work takes advantage of previous applied network studies mapping international relations among countries. We have collected and integrated multiple data sources in a unique dataset spanning several years and we have selected among the different layers the ones we thought were the most significant to describe the role of a country in the global system. Figure A.7 in the Appendix provides a full overview of all the data we collected for years 1990-2015, also with examples of layers which have not been included in our multiplex. In this section instead we first explain the criteria we have followed to select layers, then we describe their main structural properties and finally we describe the method we have used to normalize them. All the details about the sources we have consulted for the selected layers are reported in Table A.7 with references to the literature and URL links to online sources when retrievable. [5]

[5] A version of the dataset is available also at: `https://github.com/gibbbone/international_multiplex_network`

1.3.1 A selection of international networks

Our main objective in building the international multiplex has been to identify two cross sections with a sufficient number of nodes and layers to compare the structure of the multiplex before and after the Great Recession.

From the larger dataset showed in Figure A.7 in the Appendix we have first selected a set of layers with non overlapping characteristics, avoiding duplicates and partial sources.

The selection of the two years in the dataset to be compared is based on two criteria: first network measures often evolve slowly, thus a sufficiently long period is needed to capture some changes; second depending on the choice of the years some layers might not be present in the final selection, due to lack of data. To partially mitigate this second issue we have relaxed the constraint on some significant layers including them in the final cross sections even when they are missing in a given year, but data are available for an interval around it.

By choosing years 2003 and 2010 as reference point we have obtained the results shown in the last two columns of Table 1.1, t_0 and t_1, where we can see that the two cross sections are scattered: migration stock, FDI and aid layers are not matching the 2003 reference year [6]. In the rest of the Chapter we will occasionally refer to the cross section in 2003 as first cross section and to the cross section in 2010 as second cross section.

Our main goal in the construction of the international multiplex was to include as many countries as possible. By looking at the column *nodes* of Table 1.1 one can see that the number of reported nodes for some layers differs across time. To obtain a stable and large set of countries we proceeded in two steps. First, we have identified non missing observations: if the node is the target or the source of at least one non-zero edge in any of the years available for a given layer it belongs to the dataset. Second we have calculated the common subset of nodes for all layers. Some layers have been discarded from the analysis since they do not have enough observations. A more detailed discussion of our two steps

[6] Migration flows data covers five years intervals by construction.

strategy is provided in the Appendix, section "Data selection".

As a result we selected a multiplex of 19 layers with the same set of 112 nodes in each of them, sampled at two time periods around 2003 and 2010. Table 1.1 shows the features of our final selection and in Section 1.3.2 we describe the content of each layer. By looking at the distribution of weights and at the density of each layer we can see that, even though balancing our two cross sections on the same set of nodes and layers ensures that we can make meaningful comparisons before and after the Crisis, some further steps are required to normalize the content of the dataset, across multiple and heterogeneous dimensions. In subsection 1.3.3 we explain how we have filtered the edges to make layers comparable. For additional details on how we processed the data before filtering we refer to section "Data preprocessing" in the Appendix.

1.3.2 Summary statistics

Table 1.1: Variable summary: descriptive statistics of the raw data

layer code	nodes	edges	weights [a]				U[b]	S[c]	t_0	t_1
	t_0 t_1	t_0 t_1	min	max	avg.	std.				
fta_wto	211 / 211	2800 / 5004	1.00E+00	1.00E+00	1.00E+00	0.00E+00	✓	✓	2003	2010
expv	211 / 211	24276 / 25862	1.00E+00	3.49E+08	4.24E+05	4.24E+06			2003	2010
serv_exp	207 / 207	4903 / 7126	1.00E+00	5.72E+10	3.33E+08	1.92E+09			2003	2010
arms	153 / 153	274 / 396	1.00E+00	2.33E+03	6.60E+01	1.98E+02			2003	2010
invest	209 / 208	2892 / 3997	1.00E+00	1.14E+06	6.19E+03	3.72E+04			2003	2010
FDI	210 / 210	452 / 720	0.00E+00	1.25E+05	3.19E+03	1.11E+04			2005	2010
FDI_Greenfield	158 / 158	1767 / 2233	7.00E-02	2.19E+04	3.95E+02	1.25E+03			2003	2010
value	177 / 177	672 / 906	3.00E-03	3.47E+04	5.01E+02	1.74E+03			2003	2010
BIS_flow_claims	183 / 183	1993 / 2720	4.00E-03	8.97E+05	3.90E+03	2.81E+04			2003	2010
aid	170 / 136	2726 / 2887	1.00E-02	4.78E+03	1.98E+01	1.15E+02			2006	2010
migration_flow	179 / 179	11544 / 11893	1.00E+00	2.68E+06	3.47E+03	3.24E+04			2000	2005
migration_stock	205 / 202	10787 / 10991	1.00E+00	1.16E+07	1.81E+04	1.53E+05			2005	2010
out_tour	189 / 191	9842 / 10789	1.00E+00	7.93E+07	7.38E+04	8.80E+05			2003	2010
mobility	191 / 189	6014 / 7215	1.00E+00	1.26E+05	3.98E+02	2.82E+03			2003	2010
citation	196 / 195	24054 / 24054	0.00E+00	1.16E+07	2.52E+03	8.01E+04			2003	2010
collaboration	191 / 189	8854 / 8854	0.00E+00	2.04E+06	8.97E+02	2.46E+04			2003	2010
pat_cit_inv	164 / 164	2668 / 2244	1.00E+00	1.68E+06	9.90E+02	2.71E+04			2003	2010
totIC	187 / 187	627 / 833	0.00E+00	2.91E+01	7.87E-01	2.42E+00		✓	2003	2010
cow_alliances	124 / 112	2455 / 2505	1.00E+00	5.00E+00	1.24E+00	5.26E-01			2003	2010

Legend: ✓ = `True`, empty space = `False`, - = `NaN`.
[a] Weights are calculated on values greater than 0. Reported zeros are positive values lower than 0.001.
[b] If True the layer is unweighted.
[c] If True the layer is symmetric.

Table 1.1 reports the list of selected networks in our multiplex, with the short name of each layer used in our dataset in the first column and some summary statistics for each layer of the multiplex: number of nodes and edges in each of the cross sections analyzed, symmetry and weight checks on the edges and actual years used to construct the database.

Layers can be classified in six categories according to the type of international relationship they represent:

- trade: trade agreements among countries (`fta_wto`), commodity and services exchanges (`expv, serv_exp`) and arms transfers (`arms`)
- investment: foreign direct investments (`FDI, FDI_Greenfield`), total portfolio investments (TPI) (`invest`), value of mergers and acquisitions (M&A) (`value`), international aid (`aid`) and international bank loans (`BIS_flow_claims`)
- human mobility: movement of individuals between states as measured by permanent migration in flows and stocks (`migration_flow`, `migration_stock`) and temporary mobility, i.e. tourism (`out_tour`) and students abroad (`mobility`)
- knowledge flows through patent citations (`pat_cit_inv`), citations of scientific papers (`citation`) and paper coauthorships (`collaborati`
- common infrastructures between countries as measured by capacity of internet cable routing (`totIC`)
- diplomatic relationships (`cow_alliances`).

Table A.7 shows that trade layers are constituted mostly by monetary flows, hence they are represented by directed weighted graphs where edges are flows of money from source country to target country. The network of trade agreements (`fta_wto`) is one exception: it is a symmetric and unweighted layer with edges equal to 1 when two country have signed a free trade agreement.

Another exception regards arms exchange (`arms`): it is an interesting network to political scientists, hence a particular effort has been devoted in tracing and accounting arms trade. They come from separate sources collected in the the SIPRI database.

Financial layers too have edges representing monetary flows, such as FDI, TPI and M&A. It is worth noticing that FDI flows are reported twice, since we integrate OECD data with a separate source (`FDI_Greenfield`) collected in Kirkegaard (2013) focusing on specific type of firms which are difficult to measure in the official FDI statistic. Given the extent of this second source and the relevance of its content (which differs significantly from OECD) we added it as a separate source. As for BIS flows, which represent data collected from the Bank of International Settlements on loans among international banks, we selected only the claims layer, which is also automatically adjusted for changes in exchange rates. Finally, flows of international aid constitute another exception. These specific flows of capital are named Official Development Assistance (ODA) Disbursments and are measured by the OECD. While in the other layers edges represented monetary payments from source countries to target ones, here the relation represents voluntary donation from source countries to foster the development of their counterparts. Hence ODA flows represent a reversed measure of the development of a countries: central nodes (in terms of incoming edges) are those which are more underdeveloped. For this reason in our final multiplex the aid layer has been reversed (more details on this in the section "Data preprocessing" in the Appendix).

The third category of layers concerns the movement of people. Layers of this group can be divided in temporary migration (tourism and residencies lasting less than one year) and permanent stock of migrants. Both measures are collected by the United Nations statistical division. A third measure, flows of migrants among countries, has been produced for every five years from 1990 to 2010 by G. J. Abel and Sander (2014) through estimates based on the UN stock data [7].

Knowledge flows are represented by citations of patents and papers by countries. Patent citations come from the NBER database and refer

[7] Another set of estimates, this time in ten years interval from 1960 to 2000, was also provided in a previous work (G. Abel, 2013)

to the country of patent inventors. Paper citations instead have been collected by Pan, Kaski, and Fortunato (2012) and represent an average over a five year period. These layers have a common unit of measurement: each edge between two nodes represents the number of papers/patents of the destination country cited by the source one.

The last two layers are cable route capacity and diplomatic alliances. The first is a measure of the common telecommunication infrastructure linking two countries: each edge represents how many terabyte of data can travel between two countries via cable connection. Hence by construction this layer is weighted and symmetric. The alliances layer is symmetric too but it is unweighted: each edge represents an international diplomatic alliance between two countries. This is one of the measures of diplomatic interaction from the Correlates of War project. [8]

1.3.3 Filtering

The layers of the multiplex are very different and there is a need to standardize them before proceeding in any further analysis. In fact, the distribution of weights of layers in Table 1.1 (more details in Figure A.8 and A.9 in the Appendix) shows that they have very different mean value, range and skewness. Moreover, in many of our sources there is a bias in the reporting countries: due to their economic (and sometimes territorial) size, developed countries take part in exchanges with other countries more frequently.

To solve these issues we apply an hypergeometric filter on the data, as in Riccaboni, Rossi, and Schiavo (2013) . The reasoning goes like this: since edge weights are affected by the size of nodes, a filter must be applied to the edges in order to distinguish the significant ones. Following Riccaboni, Rossi, and Schiavo (2013) and Armenter and Koren (2014) we take as as starting point a null model where we assume that edges are randomly assigned to all nodes with probability proportional to their connectivity. The resulting probability distribution is an hypergeometric one. Next we fix a significance threshold and we test if the weights

[8] The others are: exchange of diplomatic representatives, membership in international organizations, armed conflicts and militarized disputes

of the actual edges in our layers passes the threshold. If not so, we discard them. The outcome of this procedure is a filtered layer where insignificant edges are removed. This approach is used to standardize the layers using the probabilities created with the filter, solving at the same time both the size bias and the heterogeneity of layers. In a nutshell, we replace heterogeneous links with the probability of them being more intense than expected under a random null model.

The final effect of the hypergeometric filter is to reduce the density variance by smoothing out denser layers and leaving almost untouched sparse layers. This follows our assumption that more connected nodes (in denser layers) have a greater share of irrelevant edges, while less connected ones have few relevant links which pass the filter. Figures A.10 and A.11 in the Appendix demonstrate this by showing the amount of observations eliminated by the filter in every layer and how this quantity is positively correlated with the layer density before filtering.

Our choice of the filter was made to reduce the variation in density across layers to make them comparable. However, we have not controlled for variation in other possible features, such as the clustering or the weak-tie structure, which might have been altered as a consequence of our procedure.

Other filters could have been chosen to preserve different characteristics of the networks. Some examples are the disparity filter of Serrano, Boguñá, and Vespignani (2009) , the GloSS filter of Radicchi, Ramasco, and Fortunato (2011) or the Maximum Entropy approach proposed in Gemmetto, Cardillo, and Garlaschelli (2017). While some of these methods are certainly better to avoid unwanted alterations in the structure of the filtered networks, as demonstrated in Gemmetto, Cardillo, and Garlaschelli (2017), we see two possible difficulties in our case. First, the disparity filter relies on a different theoretical assumption about the null model used to validate the network, a uniform distribution of links, whereas in our case we have assumed an hypergeometric one. Second, in the other two cases, the full information on the network structure is used to construct the null model. Even though this is in principle a desired property of the filter, we cannot rule out some reporting error in our data which could make the above mentioned methods less reliable. In our approach, by using only information on a given country, we avoid spreading the reporting error to the full specification of the null model.

In the next subsections we show some general network statistics calculated on the filtered multiplex using a very conservative filtering threshold.[9]

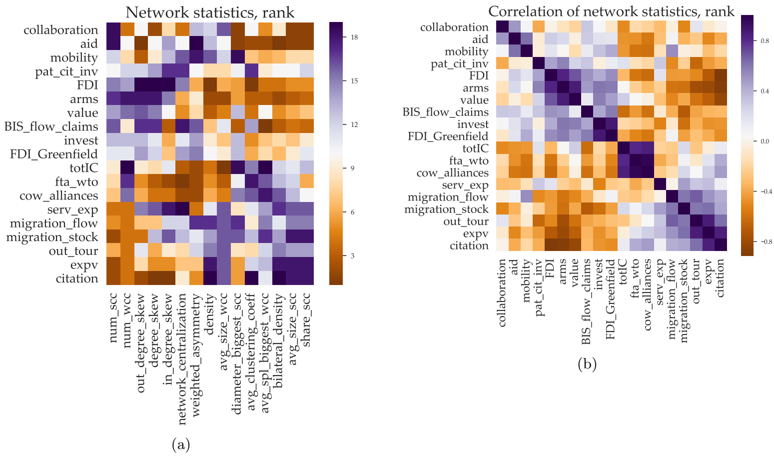

Figure 1.1: Network statistics by layer of the cross section in 2003: ranking of layers on panel (a), correlation among rankings on panel (b)

1.3.4 A comparative analysis of the international networks

Given that our interest is in the aggregated multiplex and not in the single layers composing it, we have simplified the presentation of the single layer analysis by computing network statistics on the unweighted version of our graphs, without taking into consideration those assortativity and clustering measures which rely on the directionality of edges. We also do not perform a full fledged analysis of similarity of layers, as in Bródka Piotr et al. (2018).

[9] Cfr. Figures A.12 and A.12 in the Appendix to see the effect of different thresholds.

Table 1.2 reports the network statistics for the cross section in 2003 of our multiplex. A similar table is provided in the Appendix for the cross section of year 2010 (Table A.8). To help us in visualizing some properties of layers we show in Figure 1.1 the rank of each layer with respect to each statistic (left hand side plot) and the correlation matrix measured on these rankings (right hand side plot). A lighter (darker) color in the left panel means that the selected layer (y-axis) ranks lower (higher) with respect to the selected statistics (x-axis). A lighter (darker) color in the right panel means that the two selected layers have lower (higher) correlation with respect to their position in the statistics ranking (i.e. if correlation is high the two layers rank similar in the same network statistic).

While we see some patterns of similarity, no clear cluster structure emerges: only few layers have a distinct profile (the symmetric layers) while all the others cannot be grouped in a clear manner.

This is evident in plot (a) of Figure 1.1: layers rank similarly only regarding some particular statistics, but look very different in all the others. For this reason we do not perform a proper cluster analysis, which will only increase the complexity of the exposition, and we inspect only the macro structure of the layers.

We identify two macro groups of layers with somewhat similar ranking in the statistics. The first one comprehends trade and migration data with the paper citations layer and the three symmetric layers. These last three also form another more compact ensemble in the picture. The other big group is mainly composed by financial layers together with mobility, papers collaboration and patent citation data. While the division in two groups is not clear-cut, it approximately corresponds to the division of layers by typology we introduced at the beginning with layer of different types behaving differently. This suggests that integrating these different data sources may be beneficial.

Moving to the network statistics graph (left hand side of Figure 1.1) the picture is more detailed and we can observe differences in the previously identified groups. The first one is the more defined: migration and trade data, together with the citation layer, are usually denser, both in the simple and in the bilateral sense, they have bigger average size and diameter of their bigger strongly connected components, longer average short path length and larger size of their weakly connected components. On the other hand, they have fewer components, their degree (including outdegree and indegree) is usually more skewed to the right and they have lower values of network centralization and weighted asymmetry. These characteristics are also shared by the next three layers (`cow_alliances`, `fta_wto`, `totIC`) which are also less dense.

Other layers exhibit an opposite behavior. There is a set of layers (`FDI`, `arms`, `value`) which is highly anti-correlated with the previous ones: they all rank higher in the left hand side statistics, i.e they have greater number of weakly and strongly connected components and degrees skewed to the left. Moreover they show lower rankings in the statistics where previous layers showed higher rankings: they are less dense, have lower average clustering coefficient and their biggest weakly connected component is smaller and has lower average short path length. Moreover their largest strongly connected components have smaller diameter and size.

Finally, the remaining layers have less sparkly defined characteristics: while, on average, they have opposite behavior with respect to the initial ones the differences are not so clear. In particular, like the first layers, they have higher density and higher average size of their biggest weakly connected component. However, for the remaining statistics, they differ from the first set of layers.

Table 1.2: Network statistics - cross section in 2003

	degree skew.	indegree skew.	outdegree skew.	density	bilateral density	is strongly connected	num. scc	avg. size scc	share scc	is weakly connected	num. wcc	avg. size wcc	diameter	diameter biggest scc	avg. spl	avg. spl biggest wcc	avg. clustering coeff.	network centralization	weighted asymmetry
fta_wto	0.390	0.390	0.390	0.034	0.034		48.000	2.333	0.196		48.000	2.333	-	6.000	-	2.810	0.526	0.132	0.000
expv	0.600	0.777	1.128	0.209	0.089	✓	1.000	112.000	1.000	✓	1.000	112.000	4.000	4.000	1.973	1.973	0.393	0.382	0.558
serv_exp	3.501	3.711	3.275	0.077	0.047		2.000	56.000	0.991	✓	1.000	112.000	-	5.000	2.174	2.174	0.387	1.119	0.358
arms	3.967	1.782	5.083	0.013	0.000		107.000	1.047	0.054		49.000	2.286	-	4.000	-	0.281	0.078	0.341	0.591
invest	1.684	0.870	2.263	0.058	0.007		74.000	1.514	0.348		19.000	5.895	-	7.000	-	1.069	0.269	0.524	0.775
FDI	6.883	7.107	5.182	0.012	0.002		96.000	1.167	0.134		23.000	4.870	-	4.000	-	0.687	0.000	0.663	0.486
FDI_Greenfield	1.789	0.935	2.826	0.060	0.007		63.000	1.778	0.420		13.000	8.615	-	7.000	-	1.345	0.285	0.612	0.786
value	2.825	1.125	3.850	0.018	0.003		78.000	1.436	0.312		37.000	3.027	-	7.000	-	1.531	0.125	0.294	0.578
BIS_flow_claims	3.951	0.418	4.084	0.045	0.001		105.000	1.067	0.054		3.000	37.333	-	3.000	-	0.152	0.415	0.826	0.826
aid	1.592	1.689	-0.593	0.120	0.000		112.000	1.000	0.009		4.000	28.000	-	0.000	-	0.163	0.031	0.518	0.947
migration_flow	0.970	1.431	0.602	0.194	0.021		17.000	6.588	0.830	✓	1.000	112.000	-	10.000	2.430	2.430	0.339	0.559	0.867
migration_stock	1.795	2.141	0.186	0.144	0.036	✓	1.000	112.000	1.000	✓	1.000	112.000	9.000	9.000	3.024	3.024	0.403	0.598	0.717
out_tour	0.909	1.046	1.290	0.152	0.037		22.000	5.091	0.812	✓	1.000	112.000	-	5.000	1.906	1.906	0.338	0.397	0.729
mobility	1.498	1.679	0.102	0.159	0.007		73.000	1.534	0.339		2.000	56.000	-	12.000	-	0.864	0.309	0.685	0.924
citation	-0.182	-0.019	-0.083	0.383	0.237	✓	1.000	112.000	1.000	✓	1.000	112.000	4.000	4.000	1.677	1.677	0.543	0.449	0.362
collaboration	0.300	0.946	1.213	0.110	0.009		112.000	1.000	0.009	✓	1.000	112.000	-	0.000	0.728	0.728	0.340	0.234	0.751
pat_cit_inv	2.247	2.469	1.754	0.051	0.010		59.000	1.898	0.482		17.000	6.588	-	6.000	-	1.469	0.500	0.693	0.616
totIC	1.189	1.189	1.189	0.025	0.025		50.000	2.240	0.527		50.000	2.240	-	12.000	-	5.169	0.441	0.188	0.000
cow_alliances	0.475	0.475	0.475	0.042	0.042		37.000	3.027	0.384		37.000	3.027	-	8.000	-	3.564	0.465	0.153	0.000

Legend: ✓ = `True`, empty space = `False`, - = `NaN`.
Abbreviations: *scc*: strongly connected components, *wcc*: weakly connected components, *spl*: shortest path length.

1.4 Multiplex centrality

We define a multiplex M as a set L of layers, where each layer l is a network defined as two coupled set of nodes N^l and edges E^l connecting nodes. In our case nodes are countries and edges represent different types of relationships among them. For instance countries like USA, $Russia$ and $India$ are simultaneously connected by trade, migration and knowledge flows (see Figure 1.2). Therefore, two nodes are connected in the multiplex if they have a link in at least one layer.

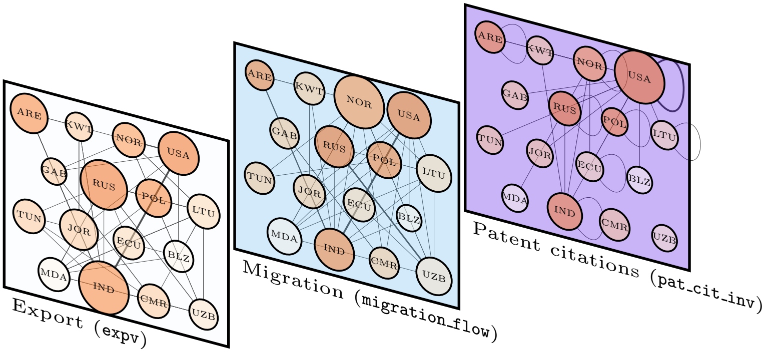

Figure 1.2: Subset of the international multiplex. Four nodes have been selected from each of the pcGDP quartiles. Node size is proportional to node degree in the whole layer, node color intensity is proportional to MultiRank centrality in the whole multiplex, edge width is proportional to edge weights

To exemplify this reasoning in Figure 1.2 we can see a selection of nodes and links from the international multiplex: 3 layers have been represented (export, migration and patent citations). For each layer, four nodes have been chosen for each of the 4 quartiles of the GDP per capita distribution (PPP) and placed in ascending order on the vertical axis. The size of the nodes is proportional to their degree in each layer while the color intensity of the nodes corresponds to the centrality of each of them in the whole multiplex calculated using the MultiRank centrality explained in Section 1.4.1. We can see that Uzbekistan (UZB) and Camerun (CMR) are isolated nodes in the patent citations layer, but not in the migration and export ones. Moreover we can see that, between the three depicted layers, self-loops are present only in the patent citation one.[10] In our multiplex there is no link across layers, but the same countries are present in all layers (Kivela et al., 2014).

The centrality of a node in a multiplex network is no more a local measure related to its role in each single layer, but becomes a global measure affected by all relationships in which the node takes part. To measure country centrality we have selected two different algorithms which generalize single layer centrality measures to multilayer networks and have the nice property to provide also a ranking of layers as output. This allows us to avoid any a priori assumption of the relevance of layers.[11]

The two algorithms are the multiplex PageRank (MultiRank) (Rahmede et al., 2018) and the multiplex HITS(MD-HITS) (Arrigo and Tudisco, 2018) also known as "hubs and authorities" algorithm. In the next subsections we provide a brief description of the two methods, while we refer to the original works for any further details.

1.4.1 MultiRank

The (single layer) PageRank of a node is a measure of node centrality that accounts for the importance of the node by looking at its centrality and the centrality of all the in-neighbours pointing to it.

Formally, given an unweighted network of N nodes, the PageRank of node i can be defined as (cfr. Halu et al., 2013):

$$x_i = d \sum_j \frac{A_{ij}}{g_j} \cdot x_j + \frac{1-d}{N} \tag{1.1}$$

Where d is the so called damping factor, A_{ij} is an element of the unweighted adjacency matrix of our network A and g_j is the out-degree

[10] More generally, self-loops have not been removed from our dataset, when reported by the original sources

[11] Obviously, there is still our selection of the layers and nodes in the dataset but, as we have shown, this is related to data availability and comparability.

of node j (i.e. the number of its out-neighbors) when its greater than 0, 1 otherwise. Hence a random walker will choose among the out-neighbours of node j with probability d and with probability $1 - d$ will switch to one of all the other nodes in the network. Starting with a uniform distribution and running iteratively the algorithm we should obtain a stable distribution of the PageRank for all nodes.

When we move from single to multiple layer networks, nodes will have multiple dimensions over which nodes share links, hence the PageRank centrality of nodes is affected both by the single-layer centrality and by the centrality of a layer in the multiplex. Therefore, a rank of the layers of the multiplex is needed to compute the multiplex version of the PageRank.

Following Rahmede et al. (2018) we summarize the multiplex from two perspectives: first as a colored network with links of different colors having different influences, second as a bipartite network of nodes and layers. Then we use both dimensions to obtain a generalized multiplex PageRank. This multilayer version of the PageRank algorithm can be defined for directed and undirected networks as well as weighted and binary networks, hence the adjacency matrix $\mathbf{A}^{\alpha}$ will refer without distinction to any of these types of layers.

From the colored network perspective we obtain matrix $\mathbf{G}$ as the sum of adjacency matrices across M layers weighted by their respective influence z^{α}. So, given layers $\alpha = 1, 2, \ldots M$, the elements of $\mathbf{G}$ are given by:

$$G_{ij} = \sum_{\alpha=1}^{M} A_{ij}^{\alpha} z^{\alpha} \tag{1.2}$$

From the bipartite view of a multiplex we obtain the $M \times N$ incidence matrices $\mathbf{B}^{in}$ and $\mathbf{B}^{out}$ representing the normalized in-strength and out-strength (respectively in-degree and out-degree for binary networks) of each node i in each layer α:

$$B_{\alpha i}^{in} = \frac{\sum_j A_{ji}^{\alpha}}{\sum_{i=1}^{N} \sum_{j=1}^{N} A_{ij}^{\alpha}} = \frac{\sum_j A_{ji}^{\alpha}}{W^{\alpha}} \qquad B_{\alpha i}^{out} = \frac{\sum_j A_{ij}^{\alpha}}{\sum_{i=1}^{N} \sum_{j=1}^{N} A_{ij}^{\alpha}} = \frac{\sum_j A_{ij}^{\alpha}}{W^{\alpha}} \tag{1.3}$$

Here $W^{\alpha} = \sum_{i=1}^{N} \sum_{j=1}^{N} A_{ij}^{\alpha}$ represents the sum of all edges weights in layer α for directed weighted networks or twice this number if the layer is undirected. For unweighted networks it becomes the total number of links (directed case) or its double (undirected case). Similarly, for undirected multiplex networks $B_{\alpha i}^{in}$ and $B_{\alpha i}^{out}$ are identical since the $\mathbf{A}^{\alpha}$ matrix is symmetric.

Our specification of the multiplex PageRank works on the assumption that, given a certain node in the multiplex, its centrality will be affected both by the centrality of nodes pointing to it and by the influence of the layers to which these in-neighbors of the node belong. A random walker moves from node j to a neighbour of i along all layers with a probability $\tilde{d}$ and proportionally to G_{ij}. Otherwise with probability $1 - \tilde{d}$ it jumps randomly on another node of $\mathbf{G}$. The stable distribution we get at the end of this process is the multiplex PageRank centrality of the nodes.

Similarly to the single layer PageRank equation we get the multilayer PageRank equation:

$$X_i = \tilde{d} \sum_{j=1}^{N} \frac{G_{ji} X_j}{k_j} + \beta v_i \tag{1.4}$$

Where $\tilde{d}$ is the damping factor and:

$$k_j = \max\left(1, \sum_{i=1}^{N} G_{ji}\right) \tag{1.5}$$

$$v_i = \theta\left[\sum_{j=1}^{N} (G_{ij} + G_{ji})\right] \tag{1.6}$$

$$\beta = \frac{\sum_{j=1}^{N}\left[1 - \tilde{d}\theta\left(\sum_{i=1}^{N} G_{ji}\right)\right]}{\sum_{i=1}^{N} v_i} X_j \tag{1.7}$$

$$\tag{1.8}$$

Here $\theta(\cdot)$ is the Heaviside step function. Equation 1.4 is a function of the set of layers' influences. To avoid calculating all these values, we couple the above equation with another one aimed at determining the influence of layers. This time we define layers' influence (relevance) as a function of the centrality of their nodes. Hence:

$$z^{\alpha} = \frac{W^{\alpha}}{\mathcal{N}} \sum_{i=1}^{N} B_{\alpha i}^{in} X_i \tag{1.9}$$

Here $\mathcal{N}$ represents a normalization constant. A more flexible specification of this equation can be obtained by introducing some tuning parameters:

$$z^{\alpha} = \frac{(W^{\alpha})^{a}}{\mathcal{N}} \left[\sum_{i=1}^{N} B_{\alpha i}^{in} (X_i)^{s\gamma}\right]^{s} \tag{1.10}$$

where a regulates the effect of total weight of layers W^{α} on the influence: with $a = 1$ layers with larger W^{α} become more influential, while with $a = 0$ the layer influence is rescaled with respect to W^{α}.

The s parameters instead indicates if more influential layers are those with fewer ($s = -1$) or more central nodes ($s = 1$).

Finally, once s is settled, the parameter γ allows us to suppress or enhance the contribution of low centrality nodes: with $s = 1$ values of $\gamma > 1$ ($\gamma < 1$) suppress (enhance) their contribution. Conversely, when $s = -1$ values of $\gamma > 1$ ($\gamma < 1$) enhance (suppress) the contribution of less central nodes.

Solving simultaneously the two coupled equations 1.4 and 1.10, given a set of parameters a, s, γ, allows us to assign a centrality X_i to each node and an influence z^α to each layer α of the multiplex.

Clearly, different choices of parameters will return different rankings.Nevertheless, as we will show in the next section, the configurations can be refined once we require them to be stable enough with respect to the choice of γ.

1.4.2 MD-HITS

The (single layer) hubs and authority scores of a set of nodes are two measures of centrality which depend one on the other recursively. For a node to have an high authority score it is necessary to have many high hub nodes pointing to it. Similarly, for a node to have an high hub score it must have a lot of high authority score nodes to point to.

Given the adjacency matrix A of an (unweighted) graph, the hub and authority scores are defined as:

$$h_i = \sum_j A_{ij} a_j \quad \text{and} \quad a_i = \sum_j A_{ji} h_j \Rightarrow \mathbf{h} = A\mathbf{a} \quad \text{and} \quad \mathbf{a} = A^T \mathbf{h} \tag{1.11}$$

Usually, the algorithm is calculated by setting all scores to one and iterating for a sufficient number of times, which requires after each step to normalize to one the sum of all the scores.

The generalization of the algorithm for multilayer networks provided by Arrigo and Tudisco (2018) includes layers and time stamps as dimensions to be used to compute the centrality. This results in five scores: two for nodes (hub and centrality scores), two for layers (broadcasting and receiving scores) and one for the time dimension. In our case the last score would not be used since we have only two cross sections.

Similarly to the MultiRank, a node receives an high hub score if it belongs to an high broadcasting layer and has many links towards high authority nodes in layers with high receiving capabilities. Conversely, high authority would be awarded to nodes in high receiving layers with high hub nodes from high broadcasting layers pointing to them.

Finally the definition of the broadcasting and receiving scores of the layers follows when the focus is toward layers instead of nodes.

As for the MultiRank, also the MD-HITS algorithm can be defined on both directed and undirected network as well as weighted and binary networks. Hence, we refer to the adjacency matrix $\mathbf{A}$ without making distinction among them.

For each layer α the multilayer generalization of the HITS algorithm reads as:

$$h_i^\alpha = b^\alpha \sum_\alpha \sum_j r^\alpha A_{ij}^\alpha a_j^\alpha \tag{1.12}$$

$$a_i^\alpha = r^\alpha \sum_\alpha \sum_j b^\alpha A_{ji}^\alpha h_j^\alpha \tag{1.13}$$

$$b^\alpha = \sum_\alpha \sum_j r^\alpha h_j^\alpha \tag{1.14}$$

$$r^\alpha = \sum_\alpha \sum_j b^\alpha a_j^\alpha \tag{1.15}$$

where b and r are the broadcasting and receiving scores of each layer. It is straightforward to notice that the overall hub centrality of a node in the multiplex is $h_i = \sum_\alpha b^\alpha a_i^\alpha$ while its overall authority score is $a_i = \sum_\alpha r^\alpha h_i^\alpha$. Since we are dealing with a multiplex, our formulation of the layer scores does not allow for some layer to not be connected to others: in other words the inter-layer network in our case is fully connected.

As we can see, there are no parameters to choose in the MD-HITS, hence we will always have one and only one ranking of nodes and layers (according to the score we are interested in). To obtain the centralities of nodes and layers we need to solve the recursive equations defined above, which requires us to solve the eigenvector problem on the whole tensor which represents the multiplex. This can be done by finding an adequate multi-homogeneous map on the adjacency tensor and its unique Perron eigenvector (cfr. Definition 1 in Arrigo and Tudisco (2018)). In Arrigo and Tudisco (2018) the authors then demonstrate the existence, uniqueness and maximality of the MD-HITS measure and provide a converging and fast parallel algorithm to compute it. MD-HITS is an ideal candidate algorithm to compute centrality in a multiplex, since it exists and is unique regardless of connectivity of the graph, while other centrality measures based on eigenvectors require the graph to be strongly connected, a property which is rarely satisfied.

In the next section we analyze the centrality of countries in the international multiplex.

1.5 Results

In this Section we summarize our main results concerning country ranking obtained computing MultiRank and MD-HITS centrality indicators. In particular, we shall focus on two setups regarding the free parameters of the MultiRank algorithm, i.e. $(\alpha, s, \gamma) = (1, 1, 1)$ and $(\alpha, s, \gamma) = (0, 1, 1)$. In general, results are fairly robust to alternative specifications, see section "Choice of the MultiRank parameters" in the Appendix for more details. Notice that when $(\alpha, s, \gamma) = (1, 1, 1)$ we measure the layer importance z without rescaling layers by their weights ($\alpha = 1$) and without giving more importance to layers with less central nodes ($s = 1$). Furthemore, the contribution of low centrality nodes is neither enhanced nor suppressed ($\gamma = 1$).

1.5.1 Preliminary results

Our dataset contains 19 different layers representing heterogeneous relations among countries. Our objective is to show how different information sources can be integrated in order to obtain synthetic measures of importance of countries and which are the possible differences among these measures. Hence we have focused on using the full set of layers at our disposal, discarding unsuitable or repeated layers before constructing the multiplex (more details on this in section A.1 in the Appendix).

However another possible question would be to check how much of the final results is correlated with the single layer composing our network, i.e. to investigate the redundancy of some layers in determining the multilayer centrality of the nodes. The intuition would be that if layers are heavily correlated among themselves the analysis could be carried on by removing some of them without loosing explanatory power.

Given that this objective is only partially consistent with our focus on analyzing the multiplex as a whole we show here some preliminary results on the correlation among single layer centrality measures of our layers and their multiplex equivalent. It is worth noticing that, however, a clear relation among the multiplex centrality rankings and the single layer ones cannot be found directly given the recursive nature of the multilayer algorithms.

In Figure 1.3 we show how much the single layers are correlated among themselves in their centrality rankings and also with respect to the multilayer version of centrality. We can see that the MultiRank algorithm is positively correlated with almost every layer except for two, even though rarely with correlation over 70%. Moreover the layers positively correlated with MultiRank are also correlated among themselves. Conversely for the MD-HITS algorithm the set of layers negatively correlated with the multilayer centrality measures contains more layers. Interestingly the MultiHub and MultiAuth are more correlated with different set of layers hence meaning that removing a set of them or another would affect the centrality measures differently.

In Figure 1.4 we have calculated for each layer the centrality scores of its nodes according to the single layer version of our centrality measures. Then for each measure we have decomposed the set of layer scores in their principal components and plotted the number of components necessary to explain at least 90% of the variance in the data (the red step line). We can see that there is some redundancy in the dataset and that the number of required components is almost half the one we are analyzing in our work for all the centrality measures. Moreover we have performed a principal component regression of the multilayer centrality measures. In principal component regression the dependent variable, the multilayer centrality in our case, is not regressed directly on the explanatory variables, the single layer centrality scores, but on their principal components. Since we want to test the best number of components in order to explain the multilayer measures, we have added an increasing number of components to the regression registering the Mean Squared Error of the fitting. [12] The blue line in Figure 1.4 shows how the Mean Squared Error of the regression behaves. We can see that it roughly agrees with the explained variance result and reaches a minimum for a number of component near half the ones we are using (even though with more variability).

All in all, we find that it exists some level of redundancy among the layers in our dataset and that a more parsimonious choice of layers could deliver almost similar results than the one we have made. However how to choose those layers is left to be decided, knowing that some choices could penalize some of the algorithms more than the others. [13] Given this uncertainty we think our full-fledged approach would at least avoid this type of bias at the cost of more redundancy in the data.

1.5.2 Layer rankings

The ranking of the layers resulting from the first combination of parameters for the MultiRank is shown in Figure 1.5 at the left of the two panels. We can see that this choice of parameters rewards the trade layers (value of exports and service exports), some of the financial layers (BIS loans, FDI flows of both types and the value of M&As) and some of the mobility layers (temporary migrations and migration stocks). Even though some choices of the algorithm are not obvious (especially among layers of similar kind) one can find a final ranking which fulfills some of the

[12] Given the multilayer centrality Y^k for type k of centrality (pagerank, hub and authority) and the scores of its single layer version, y^k_α, for each layer α, we obtain the principal components of the single layer scores, $f_{PCA}(\mathbf{y}^k) = \mathbf{X}^k_{PCA}$, and then perform OLS as follows: $Y^k = \mathbf{X}^k_{PCA}\beta + \epsilon$

[13] Indeed while the number of sufficient principal components is similar across all centrality measures, this does not give us any hint about which of the layers we should remove.

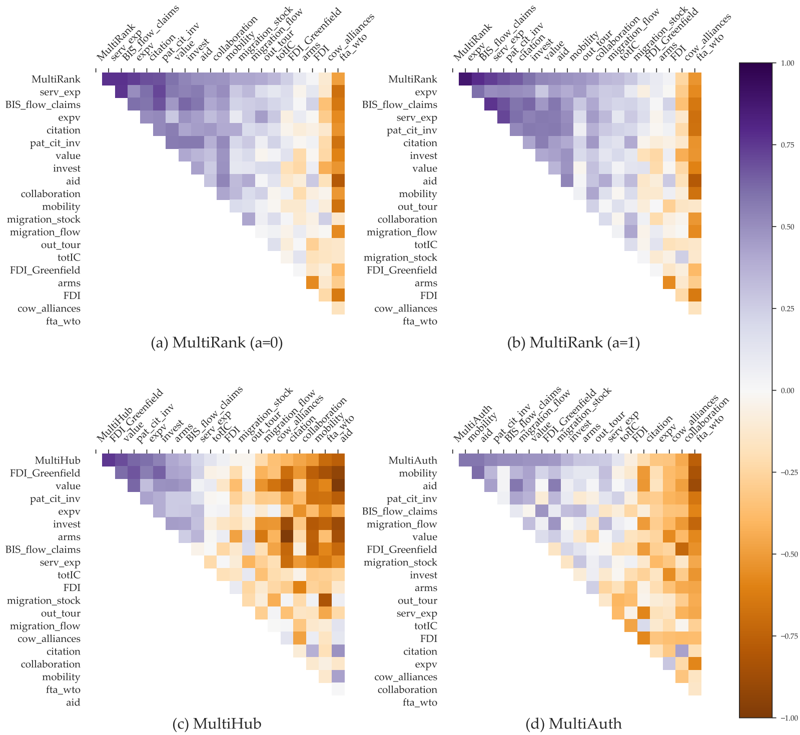

Figure 1.3: Spearman correlation of the node rankings obtained by calculating centrality measures on both the single layers and the multiplex (hence using the multiplex version of the algorithms). For the MultiRank rankings the other parameters are in both cases $s = 1$ and $\gamma = 1$.All results refer to the cross section in 2003, results for the cross section in 2010 available in the Appendix.

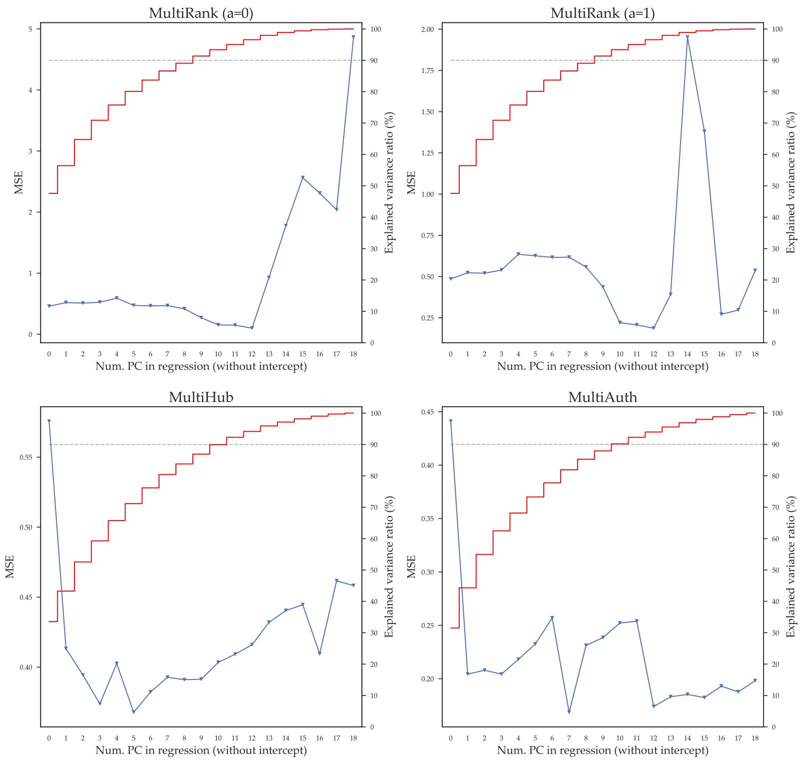

Figure 1.4: In red: number of principal components sufficient to explain 90% of the variance of the measures of centrality calculated on the single layers (scale reported on the right y-axis). In blue: mean square error obtained by regressing the multilayer measures of centrality against the principal components of the single layer centralities, added one by one (scale reported on the left y-axis). All results refer to the cross section in 2003, results for the cross section 2010 are available in the Appendix

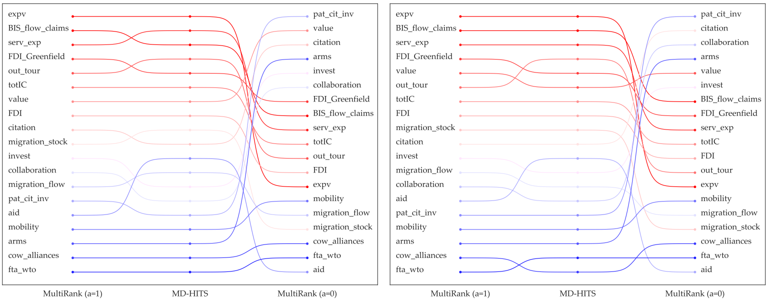

(a) (b)

Figure 1.5: Layer rankings in different configurations: cross section 2003 (a), cross section 2010 (b). For the MultiRank rankings the other parameters are in both cases $s = 1$ and $\gamma = 1$.

usual expectations on which aspects of a country are more relevant when looking at a global level: exports, foreign direct investment, international M&A, tourism and stock of migration.

The MD-HITS rankings are reported at the center of Figure 1.5. The colouring of the lines helps us see that this new ranking is very similar to the MultiRank one with $a = 1$, with some layers switching only few positions (the only exception being international aid).

Finally by looking at the two panels in Figure 1.5 we can see that over time the two previous observations hold: the ranking of layers does not change and the two algorithms follow each other.

The second configuration of the MultiRank is reported on the right of Figure 1.5. In this setting we leave $s = 1$ and $\gamma = 1$ but now $a = 0$, hence all layers are normalized by their total weight. This corresponds to define the layer importance as $z^\alpha = \frac{1}{N}\sum_{i=1}^{N} B_{\alpha i}^{in}(X_i)$.

This new configuration reduces the effect of different topological properties among layers and makes them more comparable. The result is that we have a new ranking with stark differences from the other two: now trade and migration layers are less important, loosing positions in favor of other layers such as arms trade and patent citations.

1.5.3 Node rankings

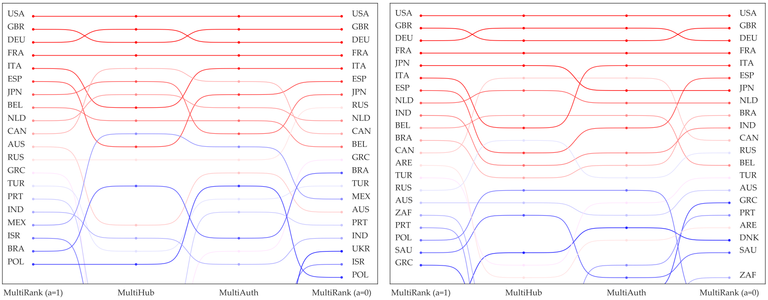

(a) (b)

Figure 1.6: Node rankings in different configurations: cross section in 2003 (a), cross section in 2010 (b). For the MultiRank rankings the other parameters are in both cases $s = 1$ and $\gamma = 1$.

We now move to the comparison of node rankings. Our final goal is to understand how different centrality indicators perform in ranking countries in the multiplex. In Figure 1.6 we report rankings for the top 20 countries obtained employing the two parameter setups used for the MultiRank, as well as the multiplex hub and authority scores (MultiHub and MultiAuth).

Results show that, as opposed to the layer rankings, on average the differences in rankings are less pronounced (the blue lines rarely cross the red ones). In fact for the top ten countries we see an almost stable distribution, with countries switching few positions by changing algorithms with the exception of RUS and CAN (this last being severely penalized by the MultiRank and rewarded by MD-HITS) and all algorithms agreeing on the same nodes. In the bottom 10 positions instead we see more differences, especially between the MultiRank and MD-HITS algorithms. In particular the ranking of the multiplex authority score rewards the countries at the bottom at the distribution and penalizes those at the middle. Finally results for the two cross sections show that some countries experience strong catching up or falling behind behaviour. For instance Poland (POL) in the 2010 cross section does belong to the top twenty countries by MultiRank when $a = 1$ but with $a = 0$ it loses several positions. On the contrary the United Arab Emirates (ARE) were not in the top twenty countries for any of the algorithms in 2003, while they are in the 2010.

We now ask whether country centrality rankings correlate with country income, as measured by pcGDP. Note that, should one find a perfect correlation, our centrality indicators would not provide any additional insight.

	MultiHub	MultiAuth	MultiRank ($a = 0$)	MultiRank ($a = 1$)	PageRank	Hub	Authority
pcGDP 2003	0.7623	0.7607	0.6372	0.6899	0.6745	0.7017	0.7594
pcGDP 2010	0.7779	0.748	0.6414	0.6264	0.6873	0.7251	0.7373
Difference 2003-2010	0.31	0.137	-0.1715	0.1228	-0.0772	0.3568	-0.1638

Table 1.3: First two lines: Spearman correlation coefficient between the rankings from pcGDP of the reference year and those obtained by different algorithms. Last line: correlation coefficient between evolution in rankings over time. First 4 columns use multilayer algorithms, last 3 use single layer algorithms on the aggregate network.

Table 1.3 shows the Spearman rank correlation coefficient between the nodes ranking resulting from pcGDP and the one resulting from different centrality algorithms. For robustness purposes we have added to our four multilayer measures their single layer version (PageRank, Hubs and Authority) calculated on the aggregate multiplex obtained by summing over all the layers the corresponding entries of their adjacency matrices. The first two rows report the rank correlations for our two cross sections, while the third contains the (simple) correlation between the evolution in ranking of pcGDP over time and of the centrality measures. The MultiHub centrality is the one with the highest correlation in ranking with pcGDP consistently between cross sections and over time. It is followed by MultiAuth centrality and the two aggregate multiplex Hub and Authority scores. Both MultiRank indicators perform relatively worse together with the PageRank centrality calculated and the aggregate network. These results show that some of the multilayer centrality algorithms rank countries in a way consistent with the ranking of pcGDP, but with some relevant differences. Specifically for the MD-HITS algorithm they are on average better than their single layer versions calculated on the aggregated network.

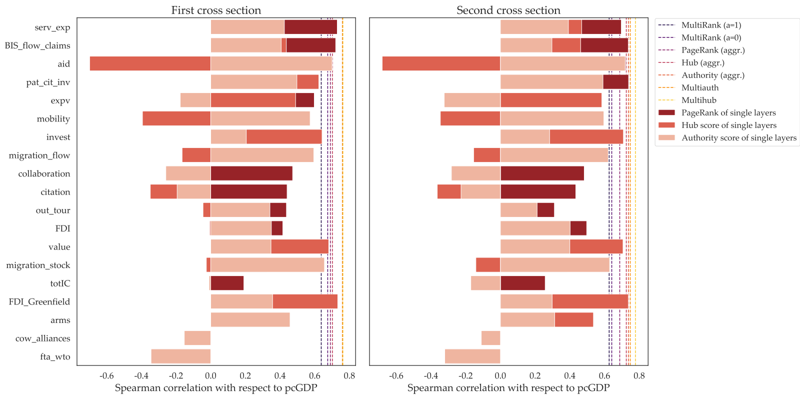

Figure 1.7: Performance of single layers centrality measures (stacked barplots) against multiplex ones (dashed lines) evaluated by correlation of their ranking of nodes with respect to pcGDP.

Another interesting result is that not only our multilayer measures perform better than their aggregated multiplex analogues, but they also have better performance than single layer centrality measures calculated on each layer. This can be seen in Figure 1.7 where we plot the the first two rows of Table 1.3 as dashed lines and the single layer centrality measures as barplots. We can see that only some layer centralities have a comparable performance with our measures, but none of them is better than our best ones, i.e. MultiHub for both cross sections and MultiAuth for the first only. In the Appendix we are reporting also the geographic distribution of MultiHub, for the cross section of year 2003 and for the evolution of rankings over time (Figure A.16).

Finally to further inspect the relation between the node rankings of our algorithms and pcGDP over time, in Figure 1.8 we have plotted the change in rank of each country from one cross section to the other for the MultiRank score (x-axis), the MultiAuth score (y-axis), the MultiHub score (coloring of dots) and pcGDP (size of dots). One can observe that

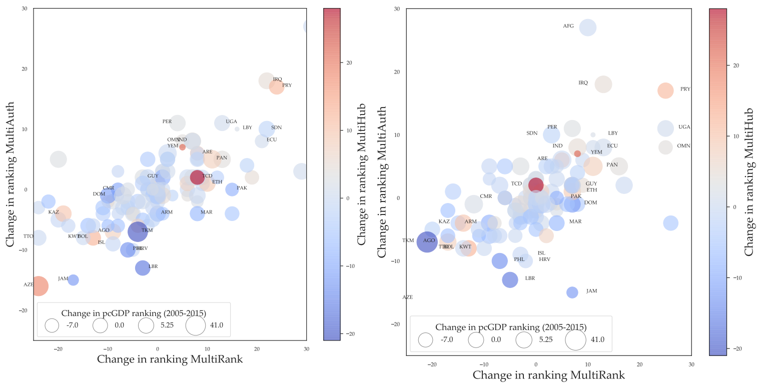

(a) (b)

Figure 1.8: Evolution over time of node centrality rankings from different algorithms for $a = 1$ (left panel) and $a = 0$ (right panel). Labels are provided for the top and bottom 10 nodes in the evolution of the MD-HITS scores.

while for the pcGDP and MultiHub score rankings we cannot see a clear pattern, for MultiRank and MultiAuth score there is a linear correlation among their rank evolution over time. Moreover this is not affected by the choice of the parameter a.

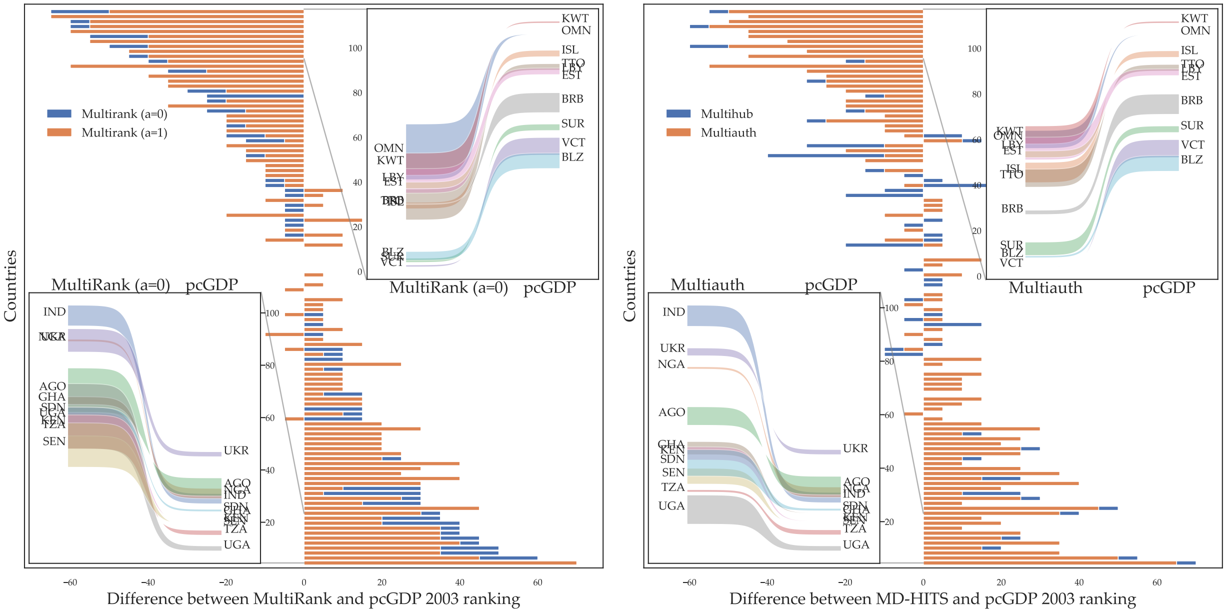

(a) (b)

Figure 1.9: Difference between ranking of countries by MultiRank (a) and MD-HITS (b) and by pcGDP in year 2003. The alternative measure for each algorithm is shown in darker color. In the insets only the top and bottom ten countries by rank difference are plotted, with ranking by multiplex centrality on the left and ranking by pcGDP on the right. The width of the line represents the evolution over time of the measures

1.5.4 Difference between pcGDP and multiplex rankings

As mentioned before, similarity with pcGDP is one possible criterion to analyze multiplex centrality. On top of this, one can get additional insights by exploring how centrality rankings deviate from those obtained using pcGDP. Figure 1.9 reports the difference in rank with respect to pcGDP ranking in 2003 for the two configurations of MultiRank (Fig 1.9 (a)) and the two MD-HITS scores (Fig 1.9(b)).[14] Observations at the top of the y-axis have greater negative difference between their ranking and the pcGDP one, hence they are being rewarded by pcGDP; while elements at the bottom of the graph have greater positive difference, hence their rank is better according to MultiRank while pcGDP ranking penalizes them.

We can see that there are wide differences between the ranking assigned by pcGDP and those assigned by the algorithms: in both cases there are nodes gaining or loosing more than 40 positions in the ranking with respect to pcGDP. Moreover, since the two graphs share the same ordering of rows (given by the Multirank differences), we can see that the rank divergences do not correspond between the MultiRank and the MD-HITS.

To further inspect these differences the behavior of the top ten and bottom ten nodes of the distributions has been analyzed in the inset of the graphs. In the insets we position on the left y-axis the ranking by our algorithm and on the right y-axis the pcGDP ranking. The "bump" from left to right represents the divergence of rankings with respect of pcGDP, while the thickness of the line represents how much the rankings have changed over time.

Two observations are in order. First, by looking at the insets in the MultiRank graph, the nodes with greater divergence have very different starting MultiRank ranking: we can see that countries being penalized by pcGDP (bottom inset) are coming both from the bottom part of the distribution (Benin, BEN) and from the top part (India, IND). And the same holds in the top inset: see Mongolia (MNG) and Kuwait (KWT).

The second observation is that the same countries evaluated by the

[14]Notice that values have been binned in five element intervals.

MD-HITS algorithm have less extreme behaviour. For instance by looking at the bottom inset of the MD-HITS graph we can see that nodes that were ranked higher in MultiRank with respect to pcGDP, in MD-HITS do not show a great divergence.

Finally in each graph we can also observe with a different color the other ranking obtained by the algorithms: MultiRank with a=1 and MultiHub. There are some differences with respect to the principal rankings but less pronounced than the ones between algorithms.

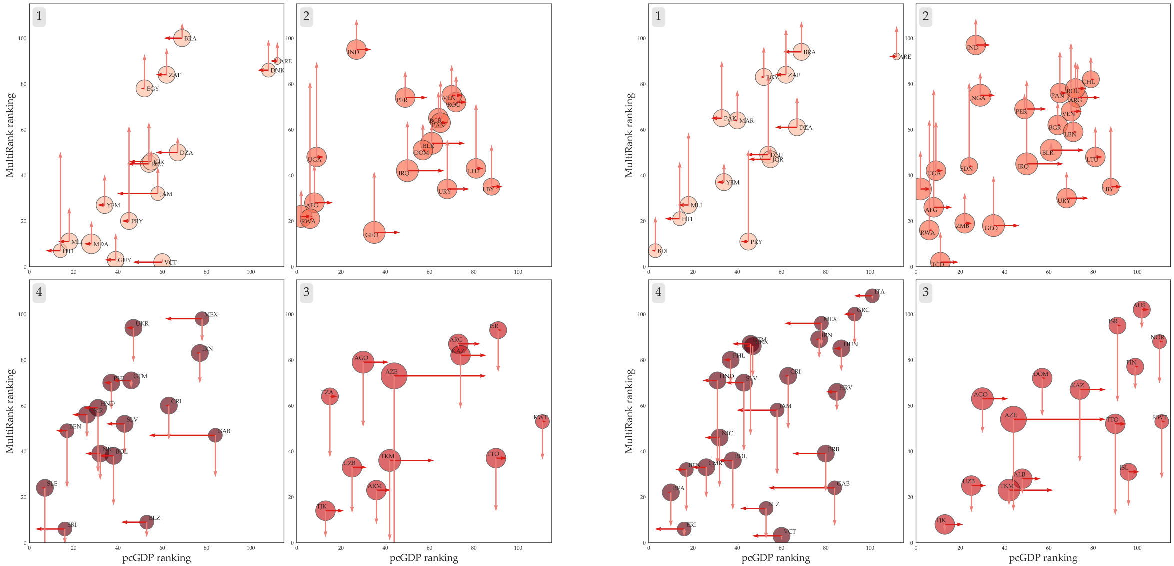

(a) (b)

Figure 1.10: Evolution of ranking over time of pcGDP (x-axis) and MultiRank (y-axis). MultiRank calculated with $a = 1$ in (a) and with $a = 0$ in (b). Countries have been split by the direction of evolution of MultiRank and pcGDP; starting from the top left corner clockwise: 1) increase/decrease, 2) increase/increase, 3) decrease/increase, 4) decrease/decrease.

1.5.5 Difference between pcGDP and multiplex rankings over time

In the last subsection, we have compared rankings obtained using our centrality indicators with the country ranking based on pcGDP. We now explore how country centrality and pcGDP rankings change across the two cross sections. For the sake of simplicity, we focus on countries gaining or losing at least five positions over time in either centrality or pcGDP rankings. We then create four country groups according to whether countries experience a rank improving (worsening) change in the centrality or pcGDP rank.

In Figure 1.10 each group of countries has been plotted in a separate graph. The position of the points represents their initial rankings by pcGDP and MultiRank of 2003 while the arrows show the movement in ranking for pcGDP(x-axis) and MultiRank (y-axis) over time. In subplots 2 and 4 their movements in pcGDP and MultiRank are correlated (in direction, not in magnitude) while in the other quadrants they are anticorrelated: while for pcGDP a certain country is expected to gain (lose) positions for MultiRank is the opposite. As a reference to pcGDP growth the size of points is proportional to the compounded growth rate of the pcGDP of the country. Finally Figures 1.10 (a) and 1.10 (b) show the differences when we change from $a = 1$ to $a = 0$ in the MultiRank.

The top right (2) and bottom left (4) plots represent those countries experiencing a similar change in centrality and pcGDP. We can see that this is not related to their initial rankings: this happens for Venezuela (VEN), Iraq (IRQ) and Georgia (GEO) in the top graph as well as Mexico (MEX), El Salvador (SLV) and Eritrea (ERI) in the bottom graph which are in different parts of the distribution. At the same time, we notice that in the top right and bottom left panels nodes tend to be distributed along the bisector (hence their initial rankings among pcGDP and MultiRank coincided as well), but with some notable exceptions such as India (IND) and Gabon (GAB).

The top left (1) and bottom right (3) plots instead show countries whose evolution over time did not coincide (in direction) by MultiRank and pcGDP. Top left countries have lost positions by pcGDP but gained by MultiRank, while the contrary happens to bottom right countries. We can see that affected countries have initial position either at the top of the graph (Brasil (BRA) in plot (1) and Israel (ISR) in plot (3)) or at the bottom part of the distribution (Haiti (HTI) in plot (1) and Tajikistan (TJK) in plot (3)) with less countries in the middle of the distibution.

But the most interesting findings come from the difference between plot 1.10 (a) with MultiRank using $a = 0$ and plot 1.10 (b) where MultiRank is calculated using $a = 1$.

These differences show that the choice of the MultiRank parameter a and the resulting ranking of layers has an effect on the final ranking of nodes and on their expected evolution. So if we assume that layers have to be rewarded for their total weight ($a = 1$), we get a ranking of layers which rewards trade layers the most (cfr. Figure 1.5) and this in turn affects the ranking of countries and affects which of them seem to gain positions over their pcGDP rankings. If instead we prefer to normalize layers by their weight ($a = 0$) we get a different ranking of layers.

While the relation between the choice of parameters, layer and node rankings is clear for MultiRank, we do not have similar insights for the MD-HITS algorithm: we have a unique ranking of layers and two ranking of nodes which take into consideration two different aspects of country centrality.

To have a clear overview on how differently the algorithms we used classify nodes with respect to their rank evolution in Figure 1.11 the four quadrants of plot Figure 1.10 have been assigned four colour codes as indicated by the colorbar on the right and for each of our algorithms nodes have been sorted in their respective categories.[15] We can see that there are few nodes where all algorithms agree on the classification, a bigger set where most of them agree and finally for more than half nodes the algorithms provide wildly different classifications.

So if we want to use a very conservative criterion and think that the

[15]Hence to see which countries are "appearing" in the plots, as we discussed in the previous paragraphs, it is sufficient to check the columns where one country is present for one algorithm and not for another (hence the third and the fourth column in our case)

agreement of the all four algorithms represent a good way to identify some stylized facts, we get the following picture: while by their pcGDP ranking Yemen (YEM) United Arab Emirates (ARE), Ecuador (ECU) and Paraguay (PRI) have not been growing, by their multiplex ranking they have been becoming more central. Another way to see this is that all four algorithm would have placed these countries in panel 1 (top left) of Figure 1.10. On the opposite all algorithms agree to place only Turkmenistan (TKM) in panel 3, which identifies countries which have become less central than what their pcGDP growth would have predicted. Furthermore in panel 2 (both the rankings resulting by pcGDP and by multiplex centrality have positive growth) the algorithms would place Uruguay (URY), Iraq (IRQ), Georgia(GEO), India (IND), Perù (PER) and Panama (PAN). Finally in panel 4 (negative growth of rankings) they would place Gabon (GAB), Belize (BLZ), Philippines (PHL), Costa Rica (CRI) and El Salvador (SLV).

1.5.6 Economic interpretation of the results

As we already mentioned in section 1.2 this work tries to contribute to the debate on country development indicators by bridging gap between economic-based indicators (pcGDP in primis) and network based indicators. Given the novelty of the approach and the ongoing debate regarding the appropriateness of pcGDP as a measure of well-being it is difficult to obtain an economic interpretation of the results without introducing further evidences or assumptions, i.e without anchoring our analysis to specific case studies or specific frameworks of interpretation of the results.

Moreover, unlike other works on multilayer networks (see De Domenico, Solé-Ribalta, et al., 2015 for instance), our work utilizes heterogeneous layers where edges represent very different types of relation (financial flows vs. human migrations vs. paper citations and so on), and hence it is difficult to give a simple interpretation of the centrality indicators without sacrificing the complexity of the information that has been analyzed: for example the centrality of nodes is strictly related to the influence of

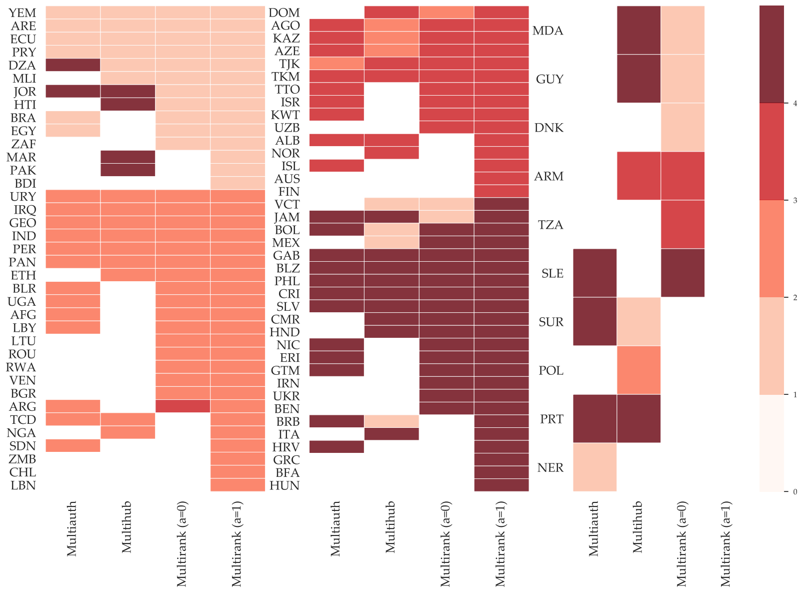

Figure 1.11: Expected direction of movement of nodes with respect to pcGDP for all the algorithms: color corresponds to the 4 quadrants in Figure 1.10, for each of the algorithm each country is classified by the joint direction of evolution of pcGDP and Multirank. Order of rows given by the Multirank $(a = 1)$ algorithm

layers which in turns may depend, in the MultiRank case, on the choice of the parameters of the algorithm equation.

Given the previous points we think that the appropriate interpretation of our results may be to consider the centrality algorithms as particular functions of aggregation of factors: the more central countries are those which are capable to obtain relevant factors more easily and the relevance of factors is defined by the importance of layers which, we showed, is a byproduct of the procedure. Hence the process we are looking at is that of a hierarchical random search of factors: if we randomly look for factors we will more probably end up in one of the more central nodes of the more relevant layers, on the reverse more peripheral nodes have less probability to be visited randomly if the search starts from other parts of the multiplex.

While from the economic point of view a random search of factor is not a likely process since agents usually know what they are looking for, what is more likely is instead the random emergence of opportunities to exploit factors in a more productive manner: countries were relevant factors are not abundant may not enjoy the same economic opportunities of others and moreover they may see their important factors move towards central nodes were demand is stronger and there is more probability of profitable employment.

Finally we highlight the fact that, even though some of our networks have edges weighted by their economic value, our measures are affected by the network position of countries and not only by their economic relevance, hence we are measuring the diffusion of factors more than their efficient (or profitable) use. The differences in the rankings of pcGDP and of the multilayer centralities may then be explained by their different focus: on one hand pcGDP may highlight country specific features reflected by national accounting but also distorted by the focus on measurable economic flows, on the other multilayer centrality may highlight the reciprocal position of countries in a wide set of different networks of exchanges ignoring economic prevalence but also all the non-network features of development of a country.

In this sense the analysis of Section 1.5.5 (and specifically in Figure

1.10) highlights countries with different path of developments: quadrants where the expected evolution of countries differ (top left and bottom right in the figure) contain countries which have gained (lost) position as attractors of relevant factors while loosing (gaining) positions from the static economic profile. As we already said in section 1.2, we don't have enough evidence to validate these hypothesis but we think that by starting from the pattern of development of these countries the effectiveness of the centrality algorithms will be confirmed.

1.6 Conclusions

In this work we have collected a new dataset of international network measures, measuring the connections between countries over different financial and non-financial dimensions and using multilayer network analysis for the first time.

We have computed different centrality measures and we have shown how different aspects of the data can be highlighted with little change in the parameters of the algorithms and how the resulting change in ranking can be interpreted. Next, we have compared the ranking of countries obtained by network measures and a common measure of macroeconomic performance, such as pcGDP. We have found a satisfying correlation for some of the algorithms, the MultiHub score especially, both at the static cross section level and for the change in pcGDP.

To show the differences between network measures and pcGDP, we have analyzed how each country development trajectory compares to changes in network centrality. We have found only few countries for which all measures agree, while for most of countries different possible trajectories were drawn, not always consistent to those measured by pcGDP.

This may reflect the fact that the development of countries is not characterized by a single path and multiple factors concur to define their growth trajectories, hence new measures are required to capture those differences. However we are conscious that any index that would try to summarize all these aspects over one dimension is by definition an approximation and entails a loss of information.

This work is a first step in order to provide meaningful multiplex measures for the centrality of countries, hence we have left many issues open for future work. We have not addressed any issue of community detection of nodes and layer similarity. Moreover we have not addressed any of the issues of causality over the network nor network reconstruction which we think are too broad topics to be discussed here. Further work will also be required to integrate other sources in the dataset of relevant international relationships.

Chapter 2

The effect of the network structure of the economy on the performance of countries during the 2008-2009 crisis

2.1 Introduction

The Great Financial Crisis (GFC from now on) is regarded as the most serious global financial crisis in history since the Black Tuesday of 1929. Starting from the collapse of the US market of housing mortgages, weakened by the use of subprime mortgages, it then transmitted to the banking sector peaking with the failure of the Lehman Brothers investment bank. From there it rapidly spread all over the world.

Officially the NBER established the duration of the GFC for the U.S as 18 months from December 2007 to June 2009. In this period we can observe one of the most impressive feature of the GFC: starting from an initial moment when few countries were affected, it rapidly involved the majority of them with a share of more than 80% of the total at its peak, then deflating to a share of around 20% in the immediate subsequent period. In Figure 2.1 we can observe the evolution of the GFC and its aftermath in the period from 2007Q1 to 2015Q1 as measured as months of recession by country [1].

Given the international spread of this phenomenon it would be meaningful to extend the observations regarding the GFC in two ways.

First, by considering the whole sample of affected countries, we could move the end date of the GFC as the first period where the majority of the sample exited from recession. This can be identified as the end of the first quarter of 2010 (the dashed line) when only 5 countries out of 70 (the 7%) were still in recession.

Secondly, we should account for the different effect of the GFC on countries. In fact if we look at the long period and extend our time range up to 2015Q1 we can see that a relevant share of countries (up to a 25% of the sample) did not recover fully from the crisis and fell in a second period of recessions from the end of 2011 up to half of 2013. For one of them, Greece, the recession period lasted uninterrupted from 2009 to 2014. In Figure 2.1 this phenomenon is highlighted by the position of countries in the plot and by their colour: more affected countries are plotted on top with darker colours, less affected ones on the bottom with lighter tones.

This is reflected also by the different economic performances of these countries as measured as the losses they bore during the GFC period. In Figure 2.3 (top panel) we have plotted the cumulative loss of each country recession during the GFC (accounting for differences in their timing): we can see a (rough) correspondence between the colouring we obtained in Figure 2.1 and the amount of losses for each countries.

The reason of these differences has been the object of several studies in the macroeconomic literature in the past years aimed at explaining why some countries have been more resilient during the crisis and have been less affected by the GFC. Moreover given that the recession started as a local banking crisis and then rapidly escalated at national and in-

[1] Recessions in Figure 2.1 have been measured as negative differences quarter-on-quarter in adjusted GDP. In the rest of the paper we will use a separate definition of recession.

ternational level we have observed the development of a new field of studies focused on contagion over financial networks. We will briefly review both these strains of literature in Section 2.2.

However there is still some space of improvement in explaining country resilience given by the interplay of these two fields of studies. On one hand network measures are still uncommon in macroeconomic studies on country resilience with only few examples of works integrating them with other economic and financial indicators. On the other hand the literature on banking contagion is mostly focused on modeling national banking networks and not international contagion. While some models could be extended to international financial flows, it seems that more simple mechanisms may help us explain the origin of the international spillover of the GFC and its different effect on world countries.

In what follows we will look for explanatory variables for different measures of economic losses during the GFC checking if network measures can help explain the different resilience of certain countries. We will show that by adding several of them to the econometric specification and testing them in an agnostic way we can improve on previous analyses. In particular we have found that kcore centrality of the network of investments among countries has an high explanatory value for the performance of countries after GFC both in the long and in the short run, something previously not discovered. We will then try to see which kind of theoretical model is useful in explain some of our findings.

The rest of the Chapter is organized as follows: in Section 2.2 we briefly review other works related to our research question, in Section 2.3 we describe the data we are using, in Section 2.4 we explain our methodology and finally in Section 2.5 we show our results, 2.6 is left for concluding remarks.

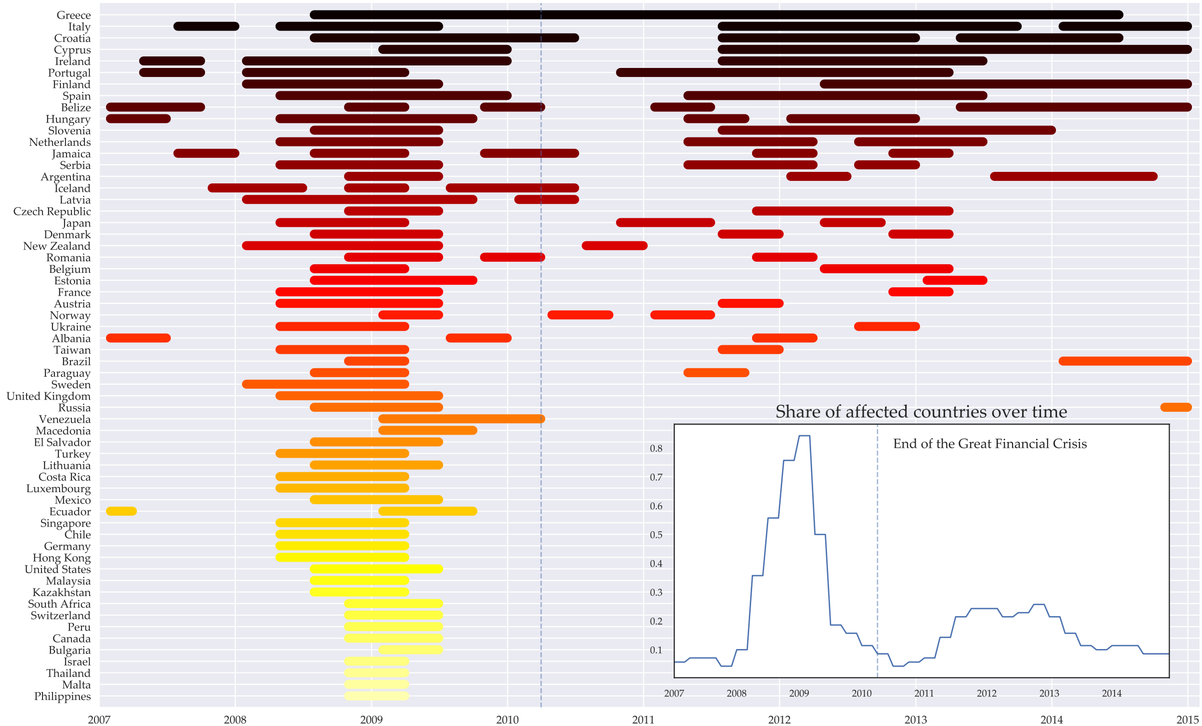

Figure 2.1: Months of recessions by country identified as quarter-on-quarter changes in real GDP (seasonally adjusted), in the years from 2007 to 2015. The share of affected countries with respect to the total is shown in the inset. Countries affected for longer periods are shown on top with darker colours while countries affected for shorter periods are shown at the bottom with lighter colours. Note: the inset is not covering any data point.

2.2 Related literature

Our work is related to two particular field of literature at the intersection of economics and network theory: models of financial contagion in financial networks and empirical analyses on banking crisis.

2.2.1 Contagion on financial networks

During the GFC and in its aftermath what became evident was that the interconnectedness of global financial institutions while providing a channel for economic development also provided routes for financial contagion which could rapidly spread across countries. As a consequence network scientists and economists have focused on the study of the properties of financial networks and of their sensitivity to contagion (Glasserman and Young, 2016 for a review).

Models of financial contagion however provide only a partial explanation of the phenomena witnessed during GFC. Firstly the focus of contagion literature is mainly theoretical (Chinazzi and Fagiolo, 2013), either because data from financial institutions are not publicly or entirely available (Anand et al., 2018), or because contagion phenomena are rare to witness (Iyer and Peydró, 2011). Two exceptions in this case are Caballero (2015) and Minoiu, Kang, et al. (2015) which use data from banking crises in the years 1980-2010 to look for network properties capable of explaining the emergence of recessions. Using data on international syndicated loans Caballero (2015) finds (among other network indicators) that the betweenness in the international network of the average bank in a country is a good explanatory variable for the number of banking crisis the country has experienced. Minoiu, Kang, et al. (2015) using instead data from BIS have shown that the degree, clustering and connectedness of a country in the network of BIS loans help explain the probability of a crisis in that country.

A second reason why contagion models are still unsatisfactory in explaining the GFC is the fact that they rarely focus on the transmission of financial shocks to the real economic system. This calls for further investigation since, as we will show in the next subsection, the debate is still open on whether banking crises necessarily lead to recessions and if this phenomenon is necessarily contagious.

One possible solution may be to adapt one of the existing models to an international scenario. For instance in Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) the collapse of a financial network is represented in a simple yet effective manner.

Financial actors have a three periods horizon: in the first period they allocate resources to risk-free assets (cash), risky investments and lending to others. If they find themselves liquidity constrained, they borrow resources from other actors for their investments. In the second period they obtain a return from their investments and have to repay their debts (or collect their credits), eventually deciding to liquidate (at a cost) the asset they will receive in the next period. Finally if they reach the third period without defaulting they receive an asset as return from their investment.

The network structure emerges from the borrowing and lending relations in period 1 and the contagion mechanism emerges as the consequence of three features of the model. First actors have senior external creditors, for instance employees or the government through taxes, which have precedence over junior creditors involved in the credit relation. Second a shock may hit randomly each of the actors with low or high severity. Finally the network of lending may be more or less sparse.

Here is how the contagion works. Since actors have senior creditors when a shock hits their current resources they automatically are unable to repay their junior creditors and are faced with a choice: liquidating their future assets or defaulting. If they default their repayment is 0, otherwise it is a share of their initial borrowed sum. Hence if the shock is high enough they will default automatically and the shock will be transmitted to their creditors. In this case the network structure play an important role: if the network is maximally sparse (a ring) or maximally dense (a complete graph) then the shock propagation will be maximum. Any intermediate configuration obtained by isolating a set of the nodes from the others will instead attenuate the propagation and reduce the global losses.

For small shocks instead the reverse situation appears: maximally dense graphs are the more resilient to shocks, while rings are the more fragile. Hence any intermediate configuration moving from the ring toward the complete graph will improve the stability of the network.

These two different behaviours of the network, according to its density and to the shock its actors receive, highlight the robust-yet-fragile feature of highly dense networks: when shocks are small a very dense network may distribute the consequences of defaults among many different actors hence reducing the effect of the shocks, on the contrary when shocks are large the same dense structure may reverberate the heavy losses and transmit them to the full network.

The model of Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) does not impose that all banks in the network must belong to the same country, hence it could be easily employed to model international borrowing relations. If we abstract from the differences among regulations of banks of different nationality and we assume that shocks are propagating across national borders without any modification, we can see the GFC as the catastrophic reaction of a dense network of international borrowing to a regime of large shocks. However what we are interested to find is how the world network structure of the economy has affected the country performance after the shocks and on this issue the work of Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) does not give any particular insights. While the model is effective in pointing out that network variables must have a role in international resilience analysis, it does not give us any hint on which of the possible network variables we should use.

An intuition of this can be obtained by looking at Figure 2.2. If we consider countries in recession as infected units of an epidemic model we can plot the share of units at risk of recession (susceptible), infected and recovered in each moment of the GFC and obtain a graph resembling a very basic SIR epidemic model: a rapid outburst of the epidemic followed by an equally fast decay in the number of infected units.

As our results will show, simpler epidemiological models (Moreno, Pastor-Satorras, and Vespignani, 2002; Pastor-Satorras and Vespignani, 2002; Pastor-Satorras, Castellano, et al., 2015) similar to the one shown

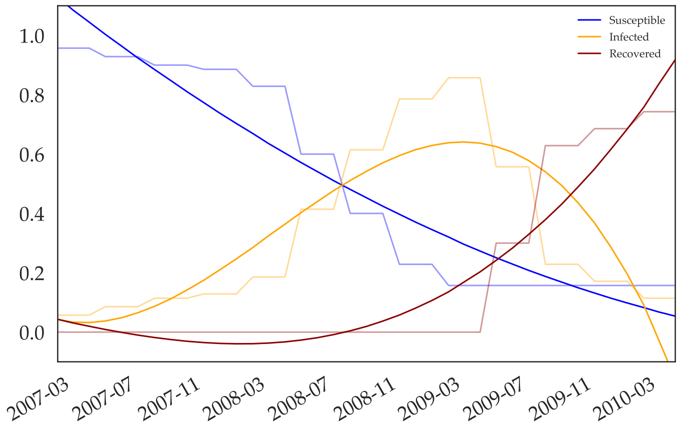

Figure 2.2: SIR model of the Great Financial Crisis. Share of susceptible countries in blue, share of infected countries (i.e. in recession) in yellow, share of recovered countries in red. Stepwise plots represent actual data. Smooth lines represent spline interpolations with $k = 3$.

in Figure 2.2 can provide the theoretical insight on where to look for explanatory network measures regarding resilience. In fact it has been demonstrated that in epidemiological models one of the most effective measure to recognize strong spreaders in the network is k-core centrality (Kitsak et al., 2010; Lin et al., 2014) which is exactly the kind of measure we will find relevant in our analyses.

Hence while a banking contagion model such as the one of Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) can be useful to explain a first phase of the GFC when the cause of the crisis has been developed, for its second phase of contagion across countries a simpler epidemiological model may be the best choice.[2] An integration of the two phases in single model might be an interesting development in the financial contagion field.

[2]In a certain sense an epidemiological model, by modeling each of the recessions as the same type of epidemics is making an even stricter assumption of similarity among countries with respect of the banking models

2.2.2 Early warnings and effect of banking crises on output

A long standing series of studies on banking crises in the last 100 years (Reinhart and Rogoff, 2009; Reinhart and Rogoff, 2014) has focused on demonstrating similarities among the effects of those episodes on country performances. The main takeaway from these works has always been that "This time is *not* different": crises have common features and identifiable policy solutions.

However following the recent evidences it has been shown that recessions after banking crises are not inevitable (Devereux and Dwyer, 2016) and that not all the recessions resulting from a banking crisis have to be severe (C. D. Romer and D. H. Romer, 2017). This confirms the early findings in the literature that financial crises are difficult to study since they are heterogeneous and rare events and that the GFC represented an exceptionally severe and diffuse episode (Cecchetti, Kohler, and Upper, 2009). For instance, an interesting approach in this sense is given by the study of the increased probability of extreme events as measured by the changes in the tail probability distribution of the stock exchange returns. As shown in Bee, Riccaboni, and Trapin (2017) for US, the detection of tail change points can be used as an early warning signal, since they anticipates by months the start of a crisis.

Given the previous observations a consistent theory of the resilience of countries after a crisis has not emerged yet and for the GFC different competing hypothesis have tried to explain the variable performance of countries (Frankel and Saravelos, 2012 for a review). On one hand all these studies agree that credit growth before the crisis is the leading explanatory variable, on the other they wildly disagree on further possible explanatory factors.[3]

Furthermore, among the different proposed explanations only few works have employed network measures aimed at addressing country interconnection, even though they are commonly used for financial contagion models. The few exceptions are Kali and J. Reyes (2010) for trade

[3]Moreover different methodology for measuring crisis have been proposed and this obviously affects the final results and makes it difficult to make comparisons.

data, Minoiu and J. A. Reyes (2013) for data regarding the interbank loans and Chinazzi, Fagiolo, et al. (2013) for data on cross-border portfolio investment holdings. In all these works adding measures of country connectivity to the econometric analysis has shown to be effective in explaining the performance of countries after a crisis.

Since no valid theory has emerged on country resilience after the GFC, any empirical effort to explain it is potentially exposed to model uncertainty bias, i.e. to the possibility of testing a theory on a suboptimal set of models while ignoring other equivalently good explanations in the models space. To address this issue in this work we follow the approach showed in Feldkircher (2014) where a Bayesian Model Averaging (BMA from now on) methodology has been used to solve the issue of model uncertainty in the study of country resilience after the crisis. [4]

In Feldkircher (2014) country resilience is measured with 4 different loss indicators capturing both short run and long run deviations in the development of a country. As explanatory variables a set of 97 financial and economic indicators has been collected, drawing from previous studies on the different reaction of countries after financial crises [5]. Each of the dependent variables has been regressed against all the 97 explanatory variables using BMA as variable selection mechanism. The results of this process highlight one principal element affecting all countries regardless of the measure of the losses and several country specific and measure specific effects. The main variable affecting negatively country resilience is the growth of credit in the years before the crisis: countries with higher credit growth before crisis bore greater losses during the GFC. This phenomenon is ubiquitous in the whole sample. On the contrary there are two countries which have specific features: in Ukraine losses are significantly higher than the rest of the sample while in Belarus they are lower. While the severe crisis of Ukraine has been historically recognized, the exception of Belarus can be explained by its strict relation with Russia, through trade arrangements, which has sheltered it

[4]BMA has been employed in growth economics for a long time and only recently has started to be used also in different contexts, for a review on its applications in economics we refer to Moral-Benito (2015)

[5]More details in Section 2.3

from more severe losses. The use of several different measures of crisis resilience helps Feldkircher (2014) recover other specific effects, in particular the positive effect of food exports which signal country less embedded in financial markets, the negative effect of free trade agreements which signal a probable channel of contagion and another channel of contagion through the exposure to external funding [6]. However these effect are specific to the type of loss measure used, hence the main takeaway of Feldkircher (2014) is that only one variable among the 97 tested can consistently explain cross country resilience, our hypothesis is that instead we can discover further effects when network variables are used. Hence we will first replicate Feldkircher (2014) results and then expand on them with newer variables.

To summarize in this Chapter we make three contributions to the debate on cross-country performance after the crisis. First, by using the same methodology and baseline models of Feldkircher (2014) we expand their analysis by testing and demonstrating the importance of country network variables, effectively being the first work to use them in the BMA framework. Secondly we employ a newly collected dataset on 18 network layers among international countries, expanding over several dimensions the previous analyses on country connectedness. Third we do not only use single layer centrality measures but employ them also at a multilayer level. While the use of multilayer measures is common place in the contagion literature we have not seen any work using them in the context of cross country resilience analysis.

2.3 Data

2.3.1 Variables from Feldkircher, 2014

In order to replicate the analysis in Feldkircher (2014) we are using the same dataset of variables they employ. Here we will briefly describe them and refer for further details to the original paper. We have four types of dependent variables capturing the effect of the crisis on output

[6] We refer to the original work for a detailed analysis.

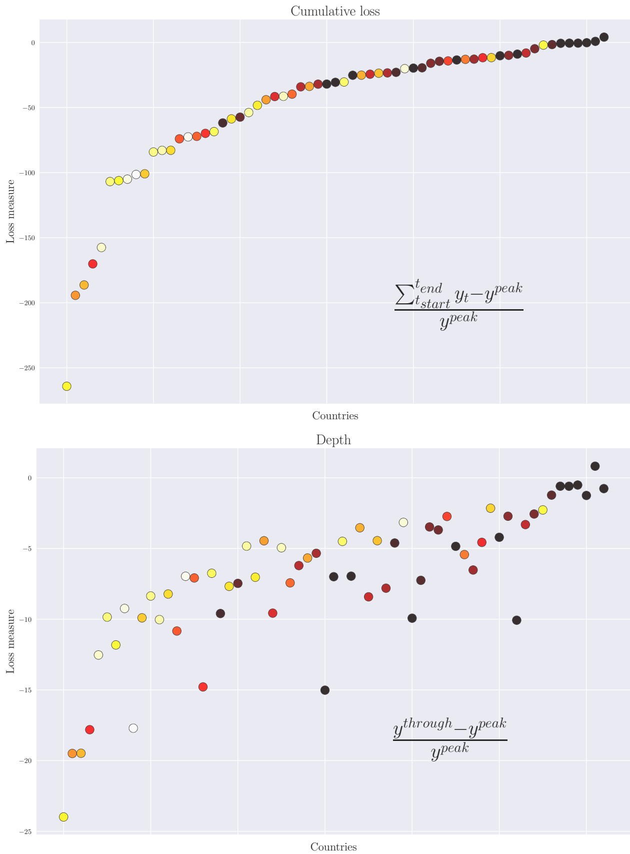

Figure 2.3: Short-run measures of economic loss. Cumulative loss on top panel, depth in the bottom one. Coloring reflects the ordering of countries obtained in Figure 2.1

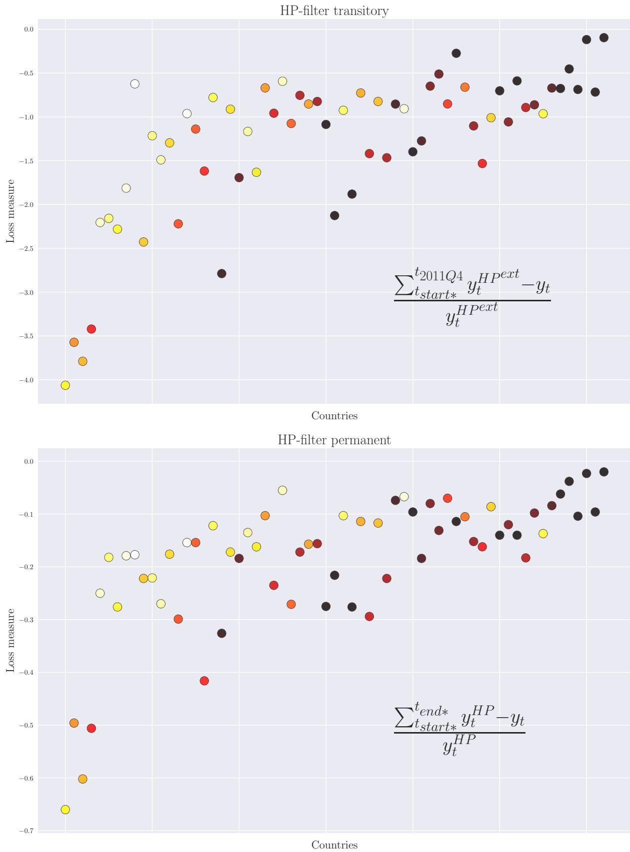

Figure 2.4: Long-run measures of economic loss. Transitory divergence from (Hodrick-Prescott) trend in the top panel, permanent divergence in the bottom one. Coloring reflects the ordering of countries obtained in Figure 2.1

both in the short and in the long run and a set of 97 possible explanatory variables.

Measures of the severity of the crisis

The static approach to define the period of crisis for a country consist in using a fixed period of time for all countries in the dataset and compare their levels of GDP or GDP growth to the ones before the crisis. In Feldkircher (2014) instead a more dynamic approach has been used: the period of crisis for each country starts with the second quarter of consecutive negative growth and ends in the quarter in which the economy reach again the maximum of its output as it was recorded in the four quarters before the start of the recession [7]. The dynamic approach allows us to analyze the effect of the crisis without being affected by the different timing of recessions across the world.

Let's define as t_{start} and t_{end} the country specific quarters delimiting the crisis period. Then we can define the following two short run measures of the impact of the crisis, the cumulative loss of output and the depth of the crisis:

$$\text{cum.loss} = \frac{\sum_{t_{start}}^{t_{end}} y_t - y^{peak}}{y^{peak}} * 100 \tag{2.1}$$

$$\text{Depth} = \frac{y^{through} - y^{peak}}{y^{peak}} * 100 \tag{2.2}$$

Here y_t is real quarterly GDP and peak and through follow the definitions given before. Next two long run measures of the impact of the crisis can be defined by looking at the deviation after the crisis from the long run trend of output obtained via the Hodrick-Prescott filter. To do so we need to redefine the period of crisis as the range from the first quarter the real GDP schedule detached from the trend trajectory (t_{start*}) and the first quarter it surpassed the trend again (t_{end*}). Moreover we assume two types of effect: a transitory effect ($HP.trans$) and a permanent

[7] Similarly the trough of the output is defined as the minimum value of output in the four quarter before the crisis

effect ($HP.per$). In the first case we assume that the output long run trend is not affected by the crisis and hence we use as a reference point the extrapolation of the output trend from the start of the crisis $y_t^{HP^{ext}}$. In the second case the long run trend is affected and hence we use as reference the trend obtained via the Hodrick-Prescott filter on the whole period y_t^{HP}:

$$\text{HP.trans} = \frac{\sum_{t_{start*}}^{t_{2011Q4}} y_t^{HP^{ext}} - y_t}{y_t^{HP^{ext}}} * 100 \tag{2.3}$$

$$\text{HP.per} = \frac{\sum_{t_{start*}}^{t_{end*}} y_t^{HP} - y_t}{y_t^{HP}} * 100 \tag{2.4}$$

Notice that in $HP.trans$ the end period is the same for all countries because none of them reached their pre-crisis trend in output. The four crisis measures are plotted in Figures 2.3 and 2.3: we can see that they have a general similar schedule but not a perfect correlation.

Explanatory variables for country performance

The set of explanatory variables contains 97 different measures calculated annually before the crisis period. The dataset aims to contain a broader set of causes in order to capture all the possible explanations provided in the literature and reduce the uncertainty. The variables belong to different categories of risk: macroeconomic, external, fiscal, financial and risk of contagion and spillover. We refer to the Appendix of Feldkircher (2014) for a detailed description.

Table 2.1: Variables summary

layer code	nodes	edges	weights [a]				U[b]	S[c]	t
			min	max	avg.	std.			
fta_wto	211	3742	1.0	1.0	1.0	0.0	✓	✓	2006
expv	211	25450	1.0	295528891.206	444044.45	4421265.2			2006
serv_exp	207	6826	100.0	51692128612.0	305982312.319	1823159056.15			2006
arms	153	331	1.0	2472.0	72.408	203.262			2006
invest	208	3716	1.0	1075579.0	6620.332	39333.797			2006
FDI	210	493	0.01	116745.324	3710.756	12378.164			2006
FDI_Greenfield	158	1873	0.25	25024.721	504.81	1615.989			2006
value	177	937	0.001	56561.026	894.909	3291.988			2006
BIS_flow_claims	183	2483	0.001	2829050.054	7880.348	82954.496			2006
aid	136	2726	0.01	4781.82	21.033	137.958			2006
migration_flow	179	11893	1.0	1957397.0	3488.237	30348.68			2005
migration_stock	206	10787	1.0	10309054.0	16902.973	145370.344			2005
out_tour	189	10586	1.0	73909666.0	76076.011	914418.922			2006
mobility	191	5932	1.0	93672.0	398.405	2734.95			2006
citation	196	24054	0.0	11594656.0	2519.782	80055.601			2006
collaboration	191	8854	0.0	2041948.75	896.614	24557.128			2006
pat_cit_inv	164	2891	1.0	1626451.0	1148.407	30879.491			2006
totIC	187	697	0.0	22.747	0.691	2.323		✓	2006
cow_alliances	112	2503	1.0	5.0	1.242	0.527			2006

Legend: ✓ = `True`, empty space = `False`, - = `NaN`.
[a] Weights are calculated on values greater than 0. Reported zeros are positive values lower than 0.001.
[b] If True the layer is unweighted.
[c] If True the layer is symmetric.

2.3.2 Network variables

We extend the analysis of Feldkircher (2014) by using a set of network variables selected from the international exchange multiplex (IEM from now on). As we already shown in Chapter 1 the IEM contains more than 20 network layers where nodes represent countries and edges represent different type of relations among them. For each layer we have observations in a different set of years, with maximum coverage achieved in years 2000-2010. Countries are not uniformly present in each layers but a minimum set of 112 nodes can be recovered in the most dense snapshots of data.

Layers in the IEM can be classified in six categories according to the type of international relationship they represent.

- trade: trade agreements among countries (`fta_wto`), commodity and services exchanges (`expv`, `serv_exp`) and arms transfers (`arms`)

- investment: foreign direct investments (`FDI`, `FDI_Greenfield`), total portfolio investments (TPI) (`invest`), value of mergers and acquisitions (M&A) (`value`), international aid (`aid`) and international bank loans (`BIS_flow_claims`)

- human mobility: movement of individuals between states as measured by permanent migration in flows and stocks (`migration_flow`, `migration_stock`) and temporary mobility, i.e tourism (`out_tour`) and students abroad (`mobility`)

- knowledge flows through patent citations (`pat_cit_inv`), citations of scientific papers (`citation`) and paper coauthorships (`collaboration`)

- common infrastructures between countries as measured by capacity of internet cable routing (`totIC`)

- diplomatic relationships (`cow_alliances`).

For the sources from which we have collected the data and a detailed description of the data and the procedure we have followed to homogenize its content we refer to Section 1.3 in Chapter 1

For this study we have selected a single year of the dataset, 2006. Moreover, in order to achieve a sufficient number of nodes in each layer, we have selected only 18 layers among the ones available in the IEM. Table 2.1 reports the list of selected networks included in our cross section of layers, with the short name used in our dataset in the first column and some summary statistics for each layer of the multiplex: number of nodes and edges, symmetry and weight checks on the edges and actual years used to construct the database. For migration data we have used observations from the closest year available, 2005 in this case.

Using the data described in Table 2.1 we have expanded on the analysis made in Feldkircher (2014) adding to the BMA models several network measures at node level drawing from different strand of literature. Starting from the previous analyses on the effect of network structure (Kali and J. Reyes, 2010; Minoiu and J. A. Reyes, 2013; Chinazzi, Fagiolo, et al., 2013) we have included for each country measures of connectedness (indegree, outdegree, instrength, outstrength, betweennes) and clustering (weighted and binary clustering coefficient). Moreover we have included different type of node centrality at single layer level (PageRank, HITS, K-core) and at multilayer level (MultiRank and MD-HITS). These last two measures represent generalizations of the well known PageRank and HITS algorithms which take into account the multilayer structure of our data and obtain centrality measures considering also the importance of different layers. While for node and single layer centrality measures the definitions are the usual ones available in most network science textbook (M. Newman, 2010), for the multilayer ones we refer to Rahmede et al. (2018) for the MultiRank and to Arrigo and Tudisco (2018) for the MD-HITS.

2.4 Methods

The econometric models we are going to estimate read as follows:

$$y_{i,t} = \beta_i x_{i,t-1} + \gamma_i z_{i,t-1} + \epsilon_{i,t} \tag{2.5}$$

Where y_i represents one of the four crisis measures, x_i is a subset of the 97 macroeconomic regressors variables chosen by Feldkircher (2014) and z_i is a subset of our network variables. $i = 1 \dots N$ are the countries of our sample with $N = 55$ obtained as the intersection of our network measures and the original macroeconomic ones. The timing of measurement is not the same for each measures, since crisis measures have been measured on a moving window of time for each country and not all macroeconomic and network measures were available exactly in the same year: in what follows independent variables will be measured in 2006, if available, otherwise the earliest measurement in a range of 3 years will be used. Finally each of the network measures we add is calculated in every layer of the dataset, hence we are adding 18 new variables to the estimation each time.[8]

[8]For multilayer measures we have included them in the relative single layer set, hence MultiRank with single layer PageRank measures and MD-HITS with HITS ones. Hence in those cases the variable added to the model are 20

2.4.1 Preliminary tests

If we assume that the true model in our analysis has a low number of explanatory variables we can frame our problem as an high dimensional approximately sparse linear regression model, i.e. where there is an high number of regressors p, a low number of observations n and the number of significant variables selected must be relatively low. The first way to address this problem is to use LASSO regressions (Tibshirani, 1996), i.e. to shrink to zero the coefficients of all the variables which do not contribute to the fitting of the model. Moreover a more robust method has been proposed in Belloni and Chernozhukov (2013), the RLASSO ("robust" lasso), where λ, the shrinking threshold of the LASSO, is automatically chosen by the algorithm and a 2 stage least square estimation is performed on the LASSO selected variables.

Hence as preliminary steps in the analysis we have employed two different types of LASSO regressions to see which regressors were selected and after that we have moved to the proper BMA model. This has the advantage that LASSO regressions are less computationally intensive and allowed us to have some preliminary insights on which variables were more relevant to maximize the fit of the models. We will show the LASSO and RLASSO results together with BMA ones in the results section to compare which variables have been selected by the different methods.

2.4.2 Bayesian Model Averaging

Since no specific theory on country resilience has emerged, any combination of our variables may represent a good candidate to explain the different reaction of countries after the crisis. With at least 115 variables to test at each step (97 macroeconomic variables and 18 network measures) the space of models to explore is enormous (2^{118} possible combinations of parameters). One possible approach may be to select as baseline a small subset of significant variables, already tested in the literature, and then enhance the initial model with some variations chosen ad-hoc. This "kitchen-sink" approach though has revealed to be fallacious since it is not robust to the addition/exclusion of variables from the chosen model. A more agnostic approach instead is to find a method to explore efficiently the whole space of models and then base inference on the weighted average of all the results. Bayesian Model Averaging is one of the possible solutions to this problem.

Inference on regression coefficients in the BMA approach reads as follows:

$$p\left(\beta_i|y\right) = \sum_{j=1}^{2^K} p\left(\beta_i|M_j, y\right) p\left(M_j|y\right) \tag{2.6}$$

Where M_j represents a single model among the 2^K models available in the model space $\mathcal{M}$ (with $K = 116$ in our case). The first term of Equation 2.6 is the posterior distribution of β_i under the hypothesis that the selected model M_j is the true one. The second term instead, $p\left(M_j|y\right)$ is the posterior model probability (PMP), the probability that the model is the true one. Hence each estimated coefficient is obtained as the average of single model coefficients weighted by their PMP.

To obtain the posterior model probabilities we must use the the Bayes update formula with the prior probability of each model, $\bar{p}$:

$$PMP_j = p\left(M_j|y\right) = \frac{p\left(y|M_j\right)\bar{p}(M_j)}{p(y)} = \frac{p\left(y|M_j\right)\bar{p}(M_j)}{\sum_{l=1}^{2^K} p\left(y|M_l\right)\bar{p}(M_l)} \tag{2.7}$$

At the first term of the numerator we find the likelihood of the observations under the selected model:

$$p\left(y|M_j\right) = \int p\left(y|M_j, \phi_j\right) p\left(\phi_j|M_j\right) \, d\phi_j \quad \text{with} \quad \phi_j = \left(\beta_j, \epsilon_j\right) \tag{2.8}$$

Hence we find the PMP of each models as the ratio of likelihood under different model specification: models with higher PMP are those with better fit to the data. Next, to obtain a measure of the importance of each variable, we sum the PMP of each model which includes the selected variable: highly significant variables are then defined as those which are included in the best fitting models. The posterior inclusion probability

(PIP) of a determinate variable z is then defined as follows:

$$PIP_z \equiv \sum_{j:z\in M_j}^{2^K} p\left(M_j|y\right) = \sum_{j:z\in M_j}^{2^K} PMP_j \tag{2.9}$$

Finally we must specify a prior distribution for the model parameters and for the distribution of models. In the first case we follow the most common option: model parameters β_j are distributed as normal around zero means. This is equivalent to assume that a priori the regressor has no effect on the dependent variable. Formally we have:

$$\beta_j \Big| \sigma^2, M_j, g \sim \mathcal{N}\left(0, \sigma^2 g \left(X_j' X_j\right)^{-1}\right) \tag{2.10}$$

Where g is the Zellner's hyperparameter (Zellner, 1986) for rescaling the number of variables. Similarly to the a priori distribution of model parameters also for the distribution of models we make an agnostic choice and assume that all models are equally likely.

Being our work an extension of the one Feldkircher (2014) we will not explain all the details of the technical implementation of the BMA estimation. For further technical details we refer to the original work.

2.5 Results

2.5.1 Description of the results

We run four different tests, one for each of the different crisis measures. In each test 10 millions models are extracted from the distribution via a Monte Carlo Markov Chain procedure, half of them are used to initialize the analysis (the burn-in phase) while only the other half are actually analyzed. This is required since the Monte Carlo procedure may visit parts of the model distribution which are not significant and hence collecting results would not be meaningful.

For each test we have run three preliminary LASSO regressions: an ordinary LASSO regression (the L column), a RLASSO regression without constraints (the RL column) and finally in the third one we have constrained the RLASSO analysis to include our netwok variables (the ERL column). For the first one we are reporting the λ threshold at which our network variables are included in the final model: b represents the best threshold, the one automatically chosen by the estimation, l represents a low threshold, which corresponds to .75% of the best one, hence a very permissive bound. Ideally we would like our variables to be selected with the best threshold. Similarly for the RLASSO regression we observe both the automatic result of the estimation and the one where we constrain the regression to include our variables and check if they are statistically significant. In both cases we report the level of statistical significance.

While we do not necessarily expect the LASSO and RLASSO procedures to agree on the inclusion of certain variables, we would like to see the same variables selected by both the BMA and the more robust RLASSO procedures, especially the unconstrained one. However since the two different methodologies rely on different principles to evaluate the importance of variables it is interesting to observe the differences between the two selection criteria.

For the BMA analysis we report three measures: the first is the posterior inclusion probability of our variables (the PIP column), the second is the posterior mean coefficient ($PostMean$) and the third is a measure of "effective estimation" or posterior significance ($PostSign.$). Following Masanjala and Papageorgiou (2008) variables which have a ratio of posterior mean with respect to the posterior standard deviation greater than 1.3 are dubbed as "efficiently estimated" and this is reported in this last column. Variables with high PIP (over .75) and which are efficiently estimated are good candidates as explanatory variables for the different performance of countries after the crisis.

For each of the variables included in the analysis the BMA reports a posterior inclusion probability. However most of the regressors are often discarded resulting in very low PIP. Hence while we report only the top 15 variables by PIP for each tests, these are always sufficient to show the most important variables in the models. As a matter of fact, finding more than one variable with a PIP over .9 would be very rare. We highlight network variables in the BMA results with a gray shade.

Finally for each experiment we report a comparison between an OLS regression for the top 5 selected models with network variables and the best model selected without network variables. Since the sample size is the same for both procedure we can directly compare the models, hence we can make observation about the best fitting of each model to the data.

2.5.2 Summary of the results

In the next subsection we will report only the main results related to the cumulative loss dependent variable (Tables B.21 and 2.3). The other results are reported in the Appendix B and are organized as follows:

- **Original BMA results on new sample**: as starting point we replicated the exact Feldkircher (2014) analysis on our new sample of countries. Results include models with and without interaction terms, which we label M1 and M3.[9]

- **BMA results with full set of network variables**: as second step we ran BMA with the full sample of network variables as first model selection stage. Results include only the model without interaction terms (M1).

- **BMA results with only kcore centralities**: after the first selection stage we identified a set of network measures which consistently appeared in the top positions by PIP, the k-core centralities of layers, so we repeated the analysis only with them including models with and without interaction terms (M1 and M3).

- **Regression results with only kcore centralities**: finally to compare the fitting of the new models with respect to the old ones selected in Feldkircher (2014) we have run an OLS regression using the first 5 models selected with the new procedures against the best one selected by the original analysis.

[9]Considered interactions are among 15 relevant economic variables and the growth of private credit (as reported in Feldkircher (2014)). For network variables we consider the interaction of all k-core layer centralities and the growth of private credit.

2.5.3 Comments on the findings

Table 2.2: Dependent variable: cum.loss. Network measure: kcore

	PIP	Post Mean	Post Sign.	RL	ERL	L
chg.dom.cred.bank.0006	0.996589	-0.193375	True	***	***	
invest_kcore_nr	0.913447	-4.081269	True	***	***	
net.pf.equ.infl.0006	0.092341	-0.081941				
collaboration_kcore_nr	0.079319	0.506029			**	b
dummy_blr	0.078006	7.602253		***	***	b
migration_stock_kcore_nr	0.069573	0.252032			**	b
gen.govDebt.06	0.042261	-0.012355		***	***	
dummy_ukr	0.042036	-2.442006		***	***	
twin.fis	0.040403	0.001346				b
net.pf.debt.infl.0006	0.039193	-0.022468				
mobility_kcore_nr	0.038861	0.159460			*	
chg.stocks.gdp.0006	0.035835	0.000316				b
outputGap_0006Exo	0.035443	-3.538792				
trade.freedom.06	0.033166	-0.026626		*	*	
arms_kcore_nr	0.030527	0.138219			.	b

Reproduction of the original results

First from Section B.1 we can see that even with a smaller sample we are able to reproduce the results of Feldkircher (2014). Hence when network variables are not added to the regressions the conclusions are the same: in almost all the tests the variable with a level of PIP great enough to be considered is always the variation in the supply of domestic credit by banks in the years between 2000 and 2006 (*chg.dom.cred.bank.0006*) which has a negative sign. This is consistent over the majority of tests we ran where *chg.dom.cred.bank.0006* is always in the top 5 variables by PIP, often the first by several orders of magnitude with respect to the others and almost always efficiently estimated. Also the RLASSO regressions correctly identify the change in domestic credit as a relevant variable.

Moreover, similarly to Feldkircher (2014), we also find that the dummy variables for Belarus (*dummy_blr*) is most of the time the second most important variable by PIP with a positive sign, even though almost never efficiently estimated and with a level of PIP comparable to credit growth.

Table 2.3: Regression results. Comparison of top 5 models with the best original model.

	Dependent variable:					
	cum_loss					
	1	2	3	4	5	ref.
invest_kcore_nr	-4.503^{***}	-5.182^{***}	-4.691^{***}	-4.223^{***}		
	(0.923)	(0.932)	(0.891)	(0.896)		
migration_stock_kcore_nr		3.869^{**}				
		(1.652)				
collaboration_kcore_nr			6.464^{**}			
			(2.808)			
net.pf.equ.infl.0006				-0.855^{**}		
				(0.373)		
chg.dom.cred.bank.0006	-0.201^{***}	-0.185^{***}	-0.213^{***}	-0.200^{***}	-0.213^{***}	-0.245^{***}
	(0.021)	(0.022)	(0.021)	(0.020)	(0.025)	(0.022)
chg.stocks.gdp.0006						0.013^{***}
						(0.005)
dummy_blr						122.978^{***}
						(32.954)
Constant	63.521^{***}	4.390	37.736^{*}	61.313^{***}	-10.658	-13.881^{**}
	(16.281)	(29.698)	(19.244)	(15.683)	(6.970)	(6.215)
Observations	55	55	55	55	55	55
R^2	0.708	0.736	0.735	0.735	0.574	0.705
Adjusted R^2	0.697	0.721	0.720	0.720	0.566	0.688
Residual Std. Error	31.020	29.764	29.813	29.823	37.099	31.471
F Statistic	63.017^{***}	47.459^{***}	47.250^{***}	47.205^{***}	71.470^{***}	40.656^{***}
BF	0.57865	1	0.92427	0.90871	0.00035	1
Logmarg	26.21542	26.76248	26.68373	26.66675	18.80165	24.06065
Post. Prob	0.28666	0.01304	0.01208	0.01161	0.00987	0.01373
Dim	3	4	4	4	2	4
Note:						*p<0.1; **p<0.05; ***p<0.01

The only difference can be found for the BMA test run with depth as dependent variable: here we obtain different results with respect to the original work, however the PIP of the variables and their coefficients are in line with our previous observations.

When instead we run BMA with interaction terms we are unable to reproduce the level of PIP obtained by Feldkircher (2014), even though we are following similar procedures.[10] However the ranking of variables by PIP and the coefficients are consistent with the original results and do not alter what we have seen so far.

Network variables as regressors

Our most important contribution regards the effect of some of our network measures on the BMA estimation. For some of them in fact we find high levels of PIP and efficient estimates which suggest that these variables may have a good explanatory power. When compared to original results in terms of fitting of the data, i.e. when used as regressors in Table 2.3 and in Section B.4, we find that they have better or comparable performance.

The first step in this process has been to perform the BMA analysis with all the network variables available to select the best models among all the possible ones. These results are shown in Section B.2. We can see that for all the four dependent variables the kcore centralities appear among the top five network regressors by PIP, while only strength and betweenness among the other types of variables are sometimes present but not with the same consistency. For this reason in the second step of model selection we have restricted the analysis to kcore variables only.

Following the literature on contagion models, the relevance of kcore centrality as an indicator to identify influential spreaders in the network was expected. Hence it is a confirmation that, while using an agnostic procedure as Bayesian model averaging, we found it as the main type of explanatory variable.

[10]This is confirmed even when we ran the test on the original data hence ruling out an effect deriving from different sample sizes

The results of the second step BMA procedure using kcore centralities only are presented in Table B.21 for the cumulative loss dependent variable and in Section B.3 for all the others variables. The outcomes are interesting for several reasons.

First, we finally find a new set of variables which have posterior inclusion probability greater than .5, something which was rare before, as we already observed in the subsection on replicated results. In Table B.21 for instance, the kcore centrality of the investment layer has PIP of 0.91. This confirms that network variables are useful in general and were previously ignored.

Secondly we can see that these new variables do not substitute the previously found most important explanatory variable, growth of private credit, but add to it. In fact if we observe the first four top models selected by the BMA procedure in Table B.21 we can see that all of them include *chg.dom.cred.bank.0006* together with the kcore centrality of the investment layer. On the contrary in the case of the Belarus dummy variable we have a substitution effect: the new network variables capture the statistical significance of the Belarus dummy which is retained less often by the BMA procedure. This is preferable since a dummy variable is a single data point regarding a single observation, hence with little content regarding the rest the sample and highly dependent on the choice of countries in the analysis. Substituting it with an index which contains information on all the dataset may be a better choice if the new model has comparable explanatory power, as we will show next.

A third reason why our results are interesting regards the type of layer which are selected more often and their posterior means. In fact we can see these layers are the investment layer with a negative coefficient (selected for all types of dependent variables: see Section B.4) and the stock of migration with a positive one (selected for all dependent variables except HP_trans). Hence we find that higher kcore centrality in the investment layer amplifies losses together with the growth of private credit, while higher kcore centrality in the migration stock layer attenuates losses.

The inclusion of the investment layer and its sign was expected and is explainable in a straightforward way: total portfolio investments represent the international counterparts of local financial networks, hence the more probable main channel of contagion among countries. Indeed already in Chinazzi, Fagiolo, et al. (2013) it was chosen as a candidate layer for explaining losses after the crisis, even though without choosing kcore as one of the target variables.

On the other hand the migration stock layer is less straightforward to interpret, especially given its positive coefficient. Since we are using the undirected version of kcore centrality in the migration layer, we are measuring the number of countries with which each node is involved in bilateral migration and which have a degree equal to the one of the node. Our results show that the higher is the kcore of a country the less it has been affected by the GFC, hence greater fluxes of people are beneficial to the resilience of a country. The explanation of this phenomenon may be related to the role of remittances. In fact a consolidated strain of literature (Aggarwal, Demirgüç-Kunt, and Pería, 2011; Ziesemer, 2012; Fromentin, 2017) has shown that financial flows of emigrants from host countries to home countries have a positive effect on the growth and development of home countries. Hence the positive coefficient we have found confirms that countries with a consistent stock of migrants abroad have enjoyed from an external help from remittances during the GFC.

Similarly another layer with positive effect but ambiguous interpretation is the layer of exchange of weapons (*arms*) which seems to be relevant in explaining transitory deviation from the trend (cfr. Table B.24). Here we are capturing the effect of weapons expenditures which has been shown repeatedly to have a negative effect on growth (Yakovlev, 2007; Alptekin and Levine, 2012; Dunne and Tian, 2015). Countries with high defense expenditures were already growing less than others and hence their transitory deviation from the trend has been smaller than the rest of the sample, which is captured by the positive effect of the arms variables.

The last reason why our results are interesting can be observed in the regression tables in Section B.4 and in Table 2.3: when compared with the original models obtained without network variables the new models which include kcore centrality have always a greater or comparable fit with respect to the data, hence signaling that they have explanatory power and are suitable candidate for alternative theories on the resilience of countries after the GFC. Indeed in Table 2.3 all the first four selected models have an adjusted R-squared statistic greater than the best reference.

Finally, one last note: all the other network measures in our dataset do not show particularly high PIPs, including the multilayer ones. This is interesting because most of the recent literature in network science seems to emphasize the better performance of multilayer measures over single layer ones (De Domenico, Solé-Ribalta, et al., 2015), while in this case we do not observe any particular advantage in using them.

We have represented our results in Figures B.19,B.20 and B.21 in the Appendix where we have plotted the investment layer with nodes as countries and edge as investment flows. We can observe that the correlation between the cumulative loss variable (the size of the nodes) and the dependent variables (the colour of the nodes) is positive (bigger nodes have darker colours) in the case of the change in domestic credit provided by bank (Figure B.19) and in the case of the kcore of the investment layer (Figure B.20), while it is negative (bigger nodes have lighter colours) for the kcore of the stock of migration (Figure B.21). Moreover we can see how the degree of the nodes (the position of node on the circle layout) is not enough to capture the contagion effect: many nodes with high degree have small size, indicating no relation with the cumulative loss variable.

Network variables as regressors with interaction terms

As robustness test we have also included interaction terms in the BMA analysis both for economic variables and kcore centralities. We have used the strong heredity prior (Chipman, 1996; Moser and Hofmarcher, 2014) which selects an interaction term only if both its interacting terms have already been selected. This restricts the space of models but avoid the situation where the three terms appear independently, making impossible to test the actual effect of the interaction as separated from the effect of the other terms. Hence while adding common variables to the model does not restrict the selection of particular models (as when we added network variables to economic ones), in this case instead some of the previous candidate models are visited with lower probability than before since the space of models is different and favours model with interaction terms, when they have sufficient PIP.

This is the case only for the models where we used cumulative loss and permanent deviation from the trend as dependent variables. In these cases one interaction term has been selected as relevant but only in the *cum_loss* model we have a better fit with the data. Indeed in this case the model where domestic credit growth and claims on GDP interact is the one with better fit with the data, even better than models with network variables, however only the interaction term is statistically significant and not the two single terms, so that paradoxically the best model is the one where domestic credit growth is not statistically significant anymore but only if interacted.

For all the other cases interaction terms are not relevant and models with no interaction have a better explanatory power.

2.6 Conclusions

In this Chapter we contributed to the effort in explaining the cross-country differences in performances after the GFC by adding network measures as explanatory variables. We employed a Bayesian Model Averaging approach to account for different model hypotheses and extended the work carried on in Feldkircher (2014). We have found that some of our network measures have good explanatory power with respect to particular indexes of after-crisis performance. This suggests that these variables should be added in future studies on country resilience after financial crises.

Specifically following the literature on epidemic contagion we have found that kcore measures of centrality are relevant in explaining the resilience of countries after the GFC with the investment layer adding to the negative effect of the growth of domestic credit in a country and the stock of migration layer reducing the effect of both those variables. While the effect of investment was expected and it is easy to explain, the effect of the migration layer is surprising and fosters further investigation.

With respect of models of banking contagion these findings reflect the intuition of more simple models of epidemics which recommend to focus on kcore centrality measures to identify potential spreaders. If we use as reference Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) we can think that, once shocks reach a critical threshold at a local level and transmit in the whole network, then they can spillover internationally with little friction; hence it is realistic to model each country as a similar unit susceptible of contagion. As a consequence while the global network statistics on which Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) focused was the density of the network, in our case we need to look at the distance between low central nodes and high central nodes, i.e. their degrees of separation: indeed if contagion is immediate for a node to be infected it is not necessary to be central but it is sufficient to be connected to a very influential node.

In fact, if we look at the investment layer we find a very low density but high "small-wordness" of the layer: the small-world omega and sigma (Hurlin, 2006) are respectively 0.06 and 0.96 indicating that the investment layer has a small-world property with respect to a random graph or a random lattice. Moreover if we use the empirical estimate of the degrees of separation among the nodes of the investment network, $\frac{\log N}{\log <k>}$ where N is the number of nodes and k is their degree, we obtain that each node is 1.23 steps away from the others: from the point of view of the investment layer, the world is really a small (global) village.

Possible extensions of this work may be directed to expand the sample size of the collected variables either by taking into account other crisis events or by increasing the number of countries in the used network; moreover given the framework we have worked with other network variables may be tested, including more refined index of contagion spreading. Finally, in case a sample covering longer time is obtained, the recent advances in BMA panel estimation should be used to address possible biases, other than model selection uncertainty, which we have ignored in this work.

Chapter 3

Multiplex network reconstruction using network embeddings on the international exchange multiplex

3.1 Introduction

One of the recurrent issues in network analysis is the missing data problem, i.e. the fact that portions of the network may not be observable. The most common causes of this problem are two: first, being network data the result of a measurement or sampling procedure, there is the possibility of inaccuracy in the measurements or bias in the sampling procedure (Peixoto, 2018c); second privacy issues may force the researcher to use censored or only partially available data (Anand et al., 2018).

More precisely, when we observe a missing link between two nodes (a 0 in the adjacency matrix), we can formulate two possible hypotheses: either we are observing the relation between the two nodes and we are measuring its absence (there is no edge between the nodes or there is an edge with weight 0) or we are not observing the interaction and hence the measurement is missing, there could be a possible link between the two nodes but we cannot know if it exists or not (Marchette and Priebe, 2008; Guimera and Sales-Pardo, 2009).

Moreover when censored data is used a second type of issue may emerge related to the missing nodes problem: our network may contain more nodes than the ones we have sampled, hence we have only partial information and we need to reconstruct the rest. Now we have missing rows and columns in our adjacency matrix but we cannot know if they are empty because they do not actually exist or because they have not been sampled (Kim and Leskovec, 2011).

Different approaches have been proposed to solve these issues. For link prediction a common distinction is between methods which are based on measures of similarity of nodes and maximum likelihood methods (Lü and Zhou, 2011).

The first methods rely on the structural properties of the nodes inside the network to obtain a score of their respective similarity, then use the score as a proxy of the probability that a link exists among the nodes. According to the scope to which each measure refer to they can be classified as local (referring to the neighbourhood of the node), global (referring to the whole graph) or quasi-local. Three of the most common local similarity measures are the common neighbor score (M. E. J. Newman, 2001) the Jaccard coefficient and the Adamic-Adar score (Adamic and Adar, 2003).

The maximum likelihood based methods instead obtain the probability of a link between two nodes as the solution of a maximum likelihood problem where some information on the network is used as constraint and the most likely graph is the solution of the constrained maximization. These methods have roots in generative network models such as the classical and weighted configuration models (Park and M. E. J. Newman, 2004; Serrano and Boguñá, 2005) and the binary and weighted stochastic block models (Holland, Laskey, and Leinhardt, 1983; Aicher, Jacobs, and Clauset, 2015), and have demonstrated to be very effective for link prediction (Squartini and Garlaschelli, 2011; Clauset, Moore, and M. E. J. Newman, 2008). However they are computational intensive (hence less scalable) and require additional information to act as constraint (usually the degree or strength sequences).

Finally maximum likelihood methods are also useful for network reconstruction (Squartini, Mastrandrea, and Garlaschelli, 2015; Mastrandrea et al., 2014; Peixoto, 2018c) especially for economics and financial networks (Almog, Squartini, and Garlaschelli, 2015; Squartini, Caldarelli, et al., 2018; Anand et al., 2018).

A completely different approach is the one of network embeddings. In this case the network data is projected on a lower dimensional space (the embedding) which is optimized in order to obtain the best fit on a proximity measure of the nodes (Hamilton, Ying, and Leskovec, 2018). The outcome is a representation of the nodes in the graph which has less complexity but preserves the main information on the structural properties of the nodes in the network. This allows to use the embeddings for link prediction (via the Hadamard product of the embeddings, usually) but also for a larger set of tasks such as node classification and clustering.

Another advantage of network embeddings is that we already have several multilayer methods developed for link prediction on multidimensional graphs and hence we are able to run an horse race to determine which is the best for our use case. Moreover multilayer methods for link prediction have a natural advantage: while with single layer methods only the information of the single network is available, by using different sources of information (the different layers of the network) multilayer methods are able to measure the relation between two nodes on multiple dimensions allowing the prediction to be refined. For instance if we have no information about patent citations between two countries, in single layer methods we must rely only on similar patterns of citations among the nodes on the layer, whereas with multilayer methods we can infer a link between the two countries by their trade relation or by the patterns of their paper citation behaviour, exploiting up to other 17 types of relations (when available). By doing so this approach is similar to other "enhanced" method of prediction where metadata about the nodes is used to predict their behaviour, however in our case we are using information about edges which is different for each possible nodes dyad and hence has a greater level of detail.

In this Chapter we will use the methodology of network embeddings to predict the existence of missing links in the 2003 cross section of our data. We will employ both single layer and multilayer network embeddings, first testing their performance on our data and then using the best one to predict missing links. Then we will assign a weight to each predicted link with a weighted stochastic block model and measure multilayer centralities comparing the reconstructed multilayer network with the constrained one used in Chapter 1.

In our case the problem of missing nodes is irrelevant: since our reference units are countries, we do know that they exist and it is just a matter of their interactions being correctly measured in the layers of our multiplex.

As we have shown in Chapter 1 our dataset does not account for all world countries and, more importantly, relation among countries are measured at different scale in different layers. As a consequence to obtain a subset of common nodes among all layers we were forced to employ a reduced country set of 112 nodes instead of the bigger possible set of 213 countries.

In building the country subset, to account for the difference between measured missing data and unmeasured and (probably) missing data we have exploited the time dimension of our dataset and checked if edges linking to and from a certain country were measured before and after the cross section we were dealing with (in an interval of maximum 20 years if the data was available). Hence if in this interval any measurement was registered, a zero in the adjacency matrix represented a real (observed) absence of relation, on the contrary if no measurement was found that zero counted as missing data and hence the node was discarded from the common set.

Since in Chapter 1 our main goal was to ensure compatibility among layers we have used a restricted dataset, however now we would like to attempt a partial reconstruction of the missing links in order to use the full extent of our layers to calculate centrality measures. The main contribution of this Chapter is the prediction of the missing links in our multiplex with a state of the art multilayer embedding algorithm and the use of these links (after weight assignment) to check the effect of missing data on our previous analyses.

The rest of the Chapter is organized as follows: in Section 3.2 we briefly review the related literature on link prediction with network embeddings, in Section 3.3 we describe the methodology we have used and finally in Section 3.4 we show our findings. Section 3.5 is left for concluding remarks.

3.2 Related works

3.2.1 Network embeddings

In what follows we will use the categorization outlined in Hamilton, Ying, and Leskovec, 2018 (and their notation) to introduce the basic concepts of network embeddings. Other similar overviews can be found in Cai, Zheng, and Chang (2017), Goyal and Ferrara (2018) and D. Zhang et al. (2018). In the next subsection we will briefly describe the methods we will test.

Network embeddings are methods developed in the field of representation learning, which focus on summarizing the topological features of the nodes of a network in a smaller vector space in order to provide feature inputs for machine learning tasks. This solves the problem of including network information as input for machine learning algorithms without relying on summary statistics but instead with a data-driven approach: node embeddings are learned from the specific structure of the graph and are able to capture more nuanced characteristics.

Obtaining the embedding of a node corresponds to project the node in a latent space with less dimensions where the more geometrically close are the projections the more they should reflect a stronger relation between the nodes in the original graph. Since network relations may have different degree of complexity, different types of proximity have been defined and algorithms may be ranked according to their ability to preserve them (Tang et al., 2015). The main types of proximity are of the first, second and high order. First-order proximity describes the relation among vertices via their the direct neighborhood: if two edges are connected then their proximity of first order is the weight of their edge, otherwise is 0. Second-order proximity expands this relation to two-step paths: it is the number of common neighbors shared by the two vertices. Finally higher-order proximity captures the similarity of indirectly connected nodes embedded in a similar set of of neighbors and it can be measured by the number of path of length greater than 2 connecting the nodes (Cao, Lu, and Xu, 2015).

To project nodes in the embedding space the first step is to use an encoder function, f_{ENC}, which maps $\mathcal{V}$, the set of nodes of our graph $\mathcal{G}$, into a lower-dimensional space $\mathbb{R}^d$ composed of vectors $\mathbf{z}$. The second step is to devise a decoder function which maps the obtained encoding to specific graph features at higher level, hence allowing us to recover the richness of the graph topology from a lower dimensional vectors.

The most common decoder function is the pairwise decoder which maps two node embeddings to $s_{\mathcal{G}}$, a measure of proximity of the nodes in the graph: $f_{DEC} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^+$.

By applying the decoder to the encoded nodes vectors we obtain an approximation of the proximity measure of the two original nodes:

$$f_{DEC}(\mathbf{z}_i, \mathbf{z}_j) = f_{DEC}\left(f_{ENC}(v_i), f_{ENC}(v_j)\right) \approx s_{\mathcal{G}}(v_i, v_j) \tag{3.1}$$

The aim of the network embedding methodology is to minimize the error in approximating the $s_{\mathcal{G}}$ measure, hence the objective of most approaches is to minimize the empirical loss $\mathcal{L}$ given a random selection of training nodes pairs $\mathcal{D}$ by choosing the appropriate set of parameters for the encoding function $\Theta_{f_{ENC}}$:

$$\min_{\Theta_{f_{ENC}}} \mathcal{L} = \sum_{(v_i, v_j) \in \mathcal{D}} \ell(f_{DEC}(\mathbf{z}_i, \mathbf{z}_j), s_{\mathcal{G}}(v_i, v_j)) \tag{3.2}$$

where ℓ is a a suitable loss function: $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$.

The minimization problem is then usually solved by optimizing the parameters via stochastic gradient descent.

Each of the 4 element composing the embedding methodology (encoder function, decoder function, proximity measure, loss function) can characterize different approaches. In this work we will focus only on a subset of these approaches, the random walk based, which share these common features: the use of direct encoding of nodes, inner product decoder, random walk based proximity functions and entropy loss function.

3.2.2 Random walk based embeddings

Random walk based methods use direct encoding of nodes. The direct encoding function is the most common encoder in node embedding methods, it consists in just an embedding lookup: $\mathbf{z}_i$, the encoding of node v_i, is simply the corresponding column in the embedding matrix $\mathbf{Z} \in \mathbb{R}^{d \times \mathcal{V}}$:

$$f_{ENC}(v_i) = \mathbf{Z}\mathbf{v}_i \tag{3.3}$$

Then the encoding parameters to be optimized, $\Theta_{f_{ENC}}$, correspond to the embedding matrix itself.

Next, to decode the embedded nodes the random walk based methods start by sampling numerous random-walks at fixed length $T \in \{2, \ldots, 10\}$ from each node v_i obtaining the probability $p_{\mathcal{G},T}(v_j|v_i)$ of visiting node j on a random path of length T starting from i. Then the embeddings are optimized in a manner that the dot product among them is proportional to $p_{\mathcal{G},T}(v_j|v_i)$.

Hence the decoder for random walks methods reads as:

$$f_{DEC}(\mathbf{z}_i, \mathbf{z}_j) \triangleq \frac{e^{\mathbf{z}_i{}^\top \mathbf{z}_j}}{\sum_{v_k \in \mathcal{V}} e^{\mathbf{z}_i{}^\top \mathbf{z}_k}} \approx p_{\mathcal{G},T}(v_j|v_i) \tag{3.4}$$

To conclude the last feature of these methods is type of loss function they use in the minimization problem, the cross-entropy loss:

$$\min_{\Theta_{f_{ENC}}} \mathcal{L} = \sum_{(v_i, v_j) \in \mathcal{D}} -\log(f_{DEC}(\mathbf{z}_i, \mathbf{z}_j)) \tag{3.5}$$

Here $\mathcal{D}$, the training set for each node v_i, is extracted from the distribution of random walks starting from v_i.

The novelty of the random walk approach is that the graph proximity measure is stochastic and asymmetric, allowing a greater flexibility which has led to improvements in the performance of these embeddings with respect to other more deterministic proximity measures.In particular random walk based methods are known to preserve proximities up to the higher order exactly because they optimize their encodings by extracting paths of different length originating from the same node.

These achievements stem from the intuition that the word embedding methods which proved to be effective in natural language processing and in particular word2vec (Mikolov, Yih, and Zweig, 2013; Mikolov, Sutskever, et al., 2013), could be applied to network data. As sentences are considered similar when they share similar context words, so path sequences are corresponding if they link together nodes with the same neighbourood. The first application of this intuition has been applying the skip-gram algorithm of Word2Vec on networks in Perozzi, Al-Rfou, and Skiena (2014) followed by Grover and Leskovec (2016) where a biased variation of the algorithm has been proposed (more details in Section 3.2.3).

Following the growing interest in multilayer networks several multilayer embedding methods have been developed, including random walk based ones. These methods can be distinguished by the type of embeddings they try to obtain: one common node embedding which summarizes the collapsed information from all the layers or a different node embeddings for each layer which however are optimized including information from all the others. We have used one of the first type (PMNE, (Liu et al., 2017)) and two of the second type in this work (Ohmnet, (Zitnik and Leskovec, 2017), MNE, (H. Zhang et al., 2018)) obtaining results which are comparable or better than single layer embeddings. In the next subsection we will add some details to describe them before showing the type of experiments we have run.

3.2.3 An overview of our embedding methods

Single layer network embeddings

- **DeepWalk** (Perozzi, Al-Rfou, and Skiena, 2014). The first random walk based method. It was characterize by the use of unbiased walks over the graph.

- **Node2Vec** (Grover and Leskovec, 2016). With respect of DeepWalk, Node2Vec has a major and a minor difference.

  The major difference is the introduction of two hyperparameters regulating the bias of the random walk: q tuning the likelihood of the random walk visiting the 1-step neighborhood of the node and p controlling the probability of the path to revisit a previous node. These two parameters allow the algorithm to simulate walks similar to either bread-first or depth-first searches hence reconstructing features of nodes which are more related the role of the node in its immediate neighborhood (BFS) or to the reachability of the node from longer path distances (DFS).

  The second minor difference is the introduction of negative sampling to optimize the approximation of Equation 3.5, which is another improvement borrowed from Word2Vec.

Multilayer network embeddings

- **PMNE** (Liu et al., 2017). The Principled Multilayer Network Embedding (PMNE) is a multilayer embedding method which produces a single embedding from the whole multilayer. Three different types of aggregate embeddings are available: network aggregation, where Node2Vec is calculated on the collapsed network; embedding aggregation, where Node2Vec is run on each layer and then the resulting embedding are joined together; co-analysis aggregation. This last method introduces a new hyperparameter r which manages the "hopping" of random walks across layers: with $r \to 0$ the algorithm will favour inter-layer hops while with $r \to 1$ the random walk will remain on the same layer more often. After the random walks are collected Node2Vec is run as in the previous methods.

- **Ohmnet** (Zitnik and Leskovec, 2017) The Ohmnet method has been developed for predicting the function of proteins in a multilayer network of cellular tissues. Since cellular tissues have a natural hierarchy Ohmnet employs this external feature to impose a further constraint on the embedding: while on one hand it learns feature at each layer level via Node2Vec, at global level proteins which belong to layers closer in the hierarchy are assumed to behave similarly and hence have similar embeddings. This is concretely implemented by adding a regularization on the loss function which pushes children nodes towards the embeddings representation of their parents, recursively.

- **MNE** (H. Zhang et al., 2018) The scalable multilayer network embedding (MNE) is a multilayer embedding method which uses a methodology similar to the one of other random walk methods to produce for each node $n \in \mathcal{N}$ in layer $i \in \mathcal{M}$ the following specific embedding:
$$\mathbf{v}_n^i = \mathbf{b}_n + w^i \cdot \mathbf{X}^{i\top} \mathbf{u}_n^i \tag{3.6}$$
Here $\mathbf{b}_n$ is a common embedding for the node which is shared among all the layers of the network as a channel of influence among them. On the other hand the embedding $\mathbf{u}_n^i$ represents the embedding of the node from the specific features of the single layer network. Then the transformation matrix $\mathbf{X}^i$ provides the coordination between the common and specific embedding via an influence parameter w^i for the single layer i.

Other methods

As baselines and references for other types of methodology we have also included the following non random walk based prediction methods:

- **LINE** (Tang et al., 2015) The Large-scale Information Network Embeddings (LINE) method obtains embeddings by using two objective functions which optimize proximity of the first and second order on a direct encoding of nodes. Then they use a Kullback-Liebler divergence as loss function between the probability distribution obtained by the embeddings and the one obtained by the adjacency matrix.

- **Common Neighbor** (M. E. J. Newman, 2001) The Common Neighbor (CN) is a network structure based measure which use as proxy of the probability of a link between two nodes the number of common neighbours between the two.

- **Jaccard Coefficient** The Jaccard Coefficient (JC) is a rescaled version of CN where common neighbours are rescaled by the total number of neighbours of the node.

- **Adamic-Adar** (Adamic and Adar, 2003)The Adamic-Adar (AA) is a network structure measure which similarly to CN and JC focus on common neighbors between nodes, however, by giving more importance to the nodes with fewer neighbors, it achieves greater accuracy than the others.

3.3 Methods

3.3.1 Prediction test

As first step in our analysis we have ran an horse race between all the link prediction methods on the cross section of 2003 of our international layer multiplex. We have used the same dataset of Chapter 1, with a set of 112 common nodes. Each algorithm has been trained on a portion of the dataset and evaluated on a prediction task on the remaining unused share. We have used a random 4-to-1 train/test split and a five-fold cross validation procedure.

The PMNE and network-structure methods have been trained on the full set of training edges for all the layers and the two LINE embeddings (first order and second order) have been concatenated. [1].

The size of the embeddings has been fixed to be 200 for all methods, with 10 step window for the random walks and negative sampling equal to 5. For Node2Vec the hyperparameters have been fixed at $p = 2$ and $q = 0.5$.

Particular attention has been devoted to the Ohmnet method since a measure of hierarchy among the layers is required in order to obtain the embeddings. Given that both multilayer centrality algorithms we have employed in Chapter 1 produce a measure of the influence of each layer in the multiplex we have used the unnormalized layer scores from the two algorithms to produce a tree hierarchy based on the distance among the scores. The three hierarchies (two for the two possible configurations of the MultiRank and one for MD-HITS) are shown in Figure 3.1.

3.3.2 Link prediction

We have used the best model selected by the horse race, MNE, to obtain predictions for missing links in our full dataset. The procedure we have followed is highlighted in Figure 3.2.

As first step we have selected five "full" layers, i.e. layers which contain the maximum common set of nodes. Their final number of nodes is 208[2] and they are: free trade agreement (`fta_wto`), the value of export (`expv`), the value of exported services (`serv_exp`), foreign direct investment (`FDI`) and investment (`invest`).

The remaining layers have the been ordered increasingly by the number of missing nodes from their country set with respect to the common country set among the five "full" layers, so that the first layer is the one which has fewer nodes missing with respect to the five reference layers and the last one is the one with more nodes missing. [3]

[1]Notice that by doing so we are replicating the same analysis of H. Zhang et al. (2018) using a new dataset.

[2]Selecting the full sample of country nodes, with dimension equal to 213 as we said in Section 3.1, would have made necessary to predict edges on the initial layers too, making the prediction procedure less robust

[3]Note: here we are using the periphrasis "missing nodes" to say that some edges were not measured in a certain layer. Since the set of nodes of a layer has been defined by the

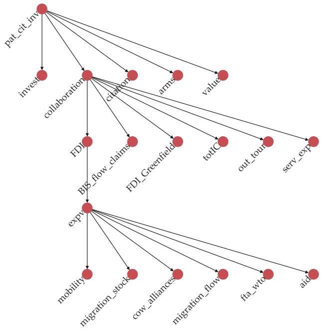

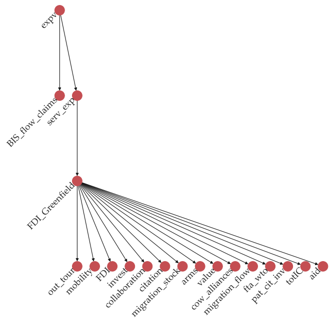

(a) Multirank influence of layers with $a = 0$ and $s = 1$

(b) Multirank influence of layers with $a = 1$ and $s = 1$

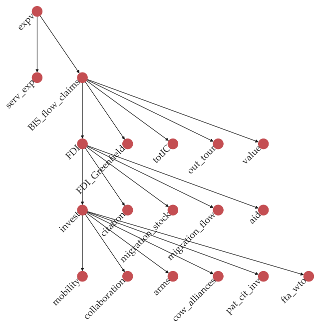

(c) MD-HITS layer centrality.

Figure 3.1: Hierarchies of layers used with Ohmnet to obtain embeddings.

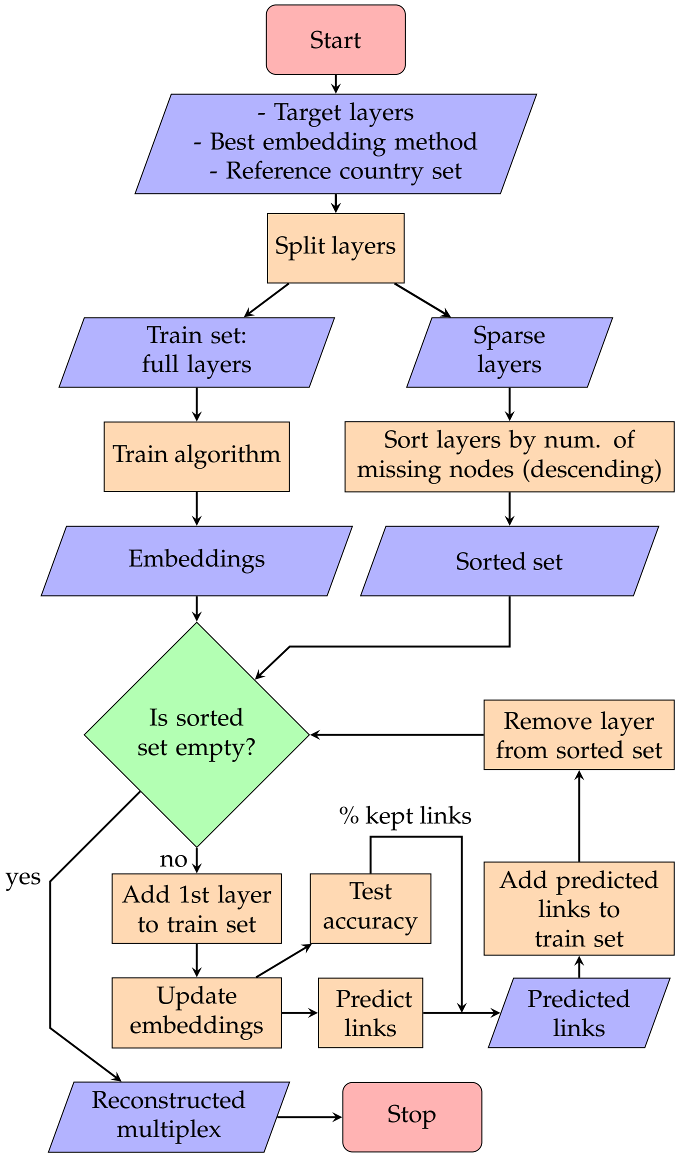

Figure 3.2: Algorithm for link prediction

We have used the five full layers as starting training set for MNE using all their nodes. Then starting from the most similar layer (by number of nodes) we have updated the embeddings by training on the new layer data with less available nodes. After the training we have predicted the missing links from the newly added layer and finally added the newly predicted links to training set of the following step. Hence starting from the second step (since at first there are no predicted links) the training set consists of the union of the previously predicted edges and the new layer edges and the used nodes are always the maximum available from start.

At each step we have run a prediction test on the layer data and used the best prediction from the algorithm to fix a threshold in order to accept the predicted links.

This procedures ensures that at each step the algorithm has been trained on the more complete data available from the previous iterations, resulting in more reliable predictions. In fact even if the layer data for the current step is very noisy or the predictions are not good this is mitigated by the compound effect of the previous "good" steps.

3.3.3 Weight assignment

The newly predicted links are only binary since the task we have tested the algorithms on is only a binary classifier. To obtain a weight for the new links we have employed a different strategy.

Since, as we will show in Section 3.4, the number of predicted links is always a small share of the original size of the layer, we can assume the network macro structure would not be altered by their addition.

Hence if we could obtain a reliable estimate of the latent block structure of the network before adding the links, we could assign a weight on them by using the inferred structure and the previously existing nodes to which the newly added links are pointing to.

In other words the predicted links will have a weight corresponding to the average weight of the block of nodes to which they have been

edges measured in that layer (in an interval of 20 years), saying that a node is missing means to say that edges have not been measured, hence actually the links are missing and as a consequence the nodes have been dropped from the dataset.

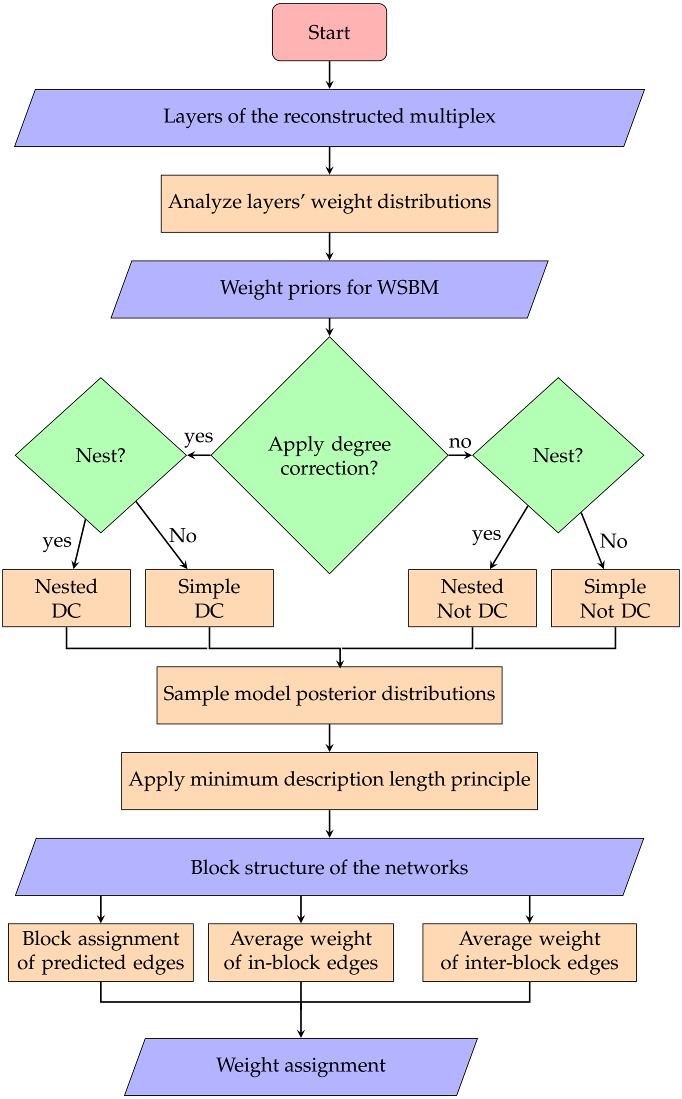

Figure 3.3: Algorithm for weight assignment

layer	distribution	nested	degree corrected
value	real-normal	✓	
totIC	real-exponential	✓	
BIS_flow_claims	real-normal	✓	✓
aid	real-exponential		✓
FDI_Greenfield	real-normal		
cow_alliances	discrete-binomial		
pat_cit_inv	real-exponential		
mobility	real-exponential	✓	
out_tour	real-normal		
migration_stock	real-normal		
citation	real-exponential		
migration_flow	real-exponential		
collaboration	real-normal	✓	
arms	real-exponential		

Table 3.1: Characteristic of WSBM models obtained via model selection for each of the layers

attached, as if they always belonged to it. More precisely we need also to allow for the possibility of the new links to be casted as inter-block links (links between two different blocks) hence their weight would be the weighted average of the weights inside the block and of all the weights pointing to it from other blocks, averaged by their respective frequencies.

Since weight is a fundamental feature of the structure of the network we want to estimate we are going to use weighted stochastic block models (WSBM) (Peixoto, 2014; Peixoto, 2018b). These methods try to obtain a reliable estimate of the block structure of a network by maximizing the maximum likelihood of the observed data via the choice of different partitions of the network under the hypothesis of a prior distribution of the edges (a generative model). Moreover by including a prior on the distribution of weights they can also model weighted networks (Aicher, Jacobs, and Clauset, 2015). Finally since all these method are Bayesian a model averaging approach may be employed to select the best partition among all the one obtainable via sampling on the distribution of the posterior probability (in a similar way of what we have shown in Chapter 2)

(Peixoto, 2018a).

Due to their ability to generalize the structure of a network, SBMs can be employed also for link and weight prediction with significant results (Clauset, Moore, and M. E. J. Newman, 2008; Aicher, Jacobs, and Clauset, 2015; Peixoto, 2018c), however since the focus of this Chapter is on network embeddings we will employ them only as an assignment mechanism for the predicted edges weight.

The procedure to assign weights that we have used is depicted in Figure 3.3 and works as follows: we fit a WSBM on each of the layer of the dataset with a weight prior chosen appropriately for each layer among exponential and lognormal distributions for continuous data and Bernoulli distribution for discrete data. After this we have ran 4 different WSBMs to test the best fit of simple and nested block models, with or without degree correction. We evaluate the model performances according to the minimum description length principle, according to which the best model is the one that uses the less information to describe the available data. To select one specific class of models we approximated the posterior entropy by the Bethe measure (Mézard and Montanari, 2009) and calculated the model evidence after 200000 sampling from the posterior distribution.[4] The results of the model choice procedure are reported in Table 3.1.

Finally after the best model has been chosen we have used the resulting block structure to obtain the average edge weight inside each block and between them and used these averages to assign weights to the predicted links.

3.4 Results

3.4.1 Comparison of the performance of the different embeddings methods

The results of our first step are shown in Figure 3.4 and 3.5.

[4]Comparing the results with others obtained with an alternative measure of entropy, the mean field approximation, we have found both measures agreeing on the majority of the model class selections except for a pair of cases.

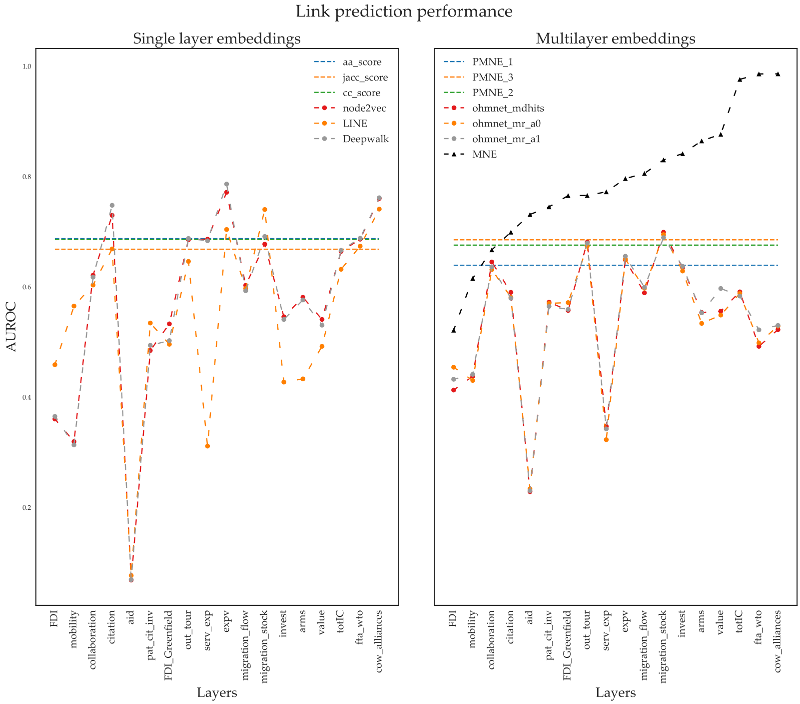

Figure 3.4: Area under the ROC curve for 5-fold link prediction experiment for all the embedding methods tested. Showed results are the averages over all the iterations. Embeddings have 200 features and have been trained for 100 epochs for each method.

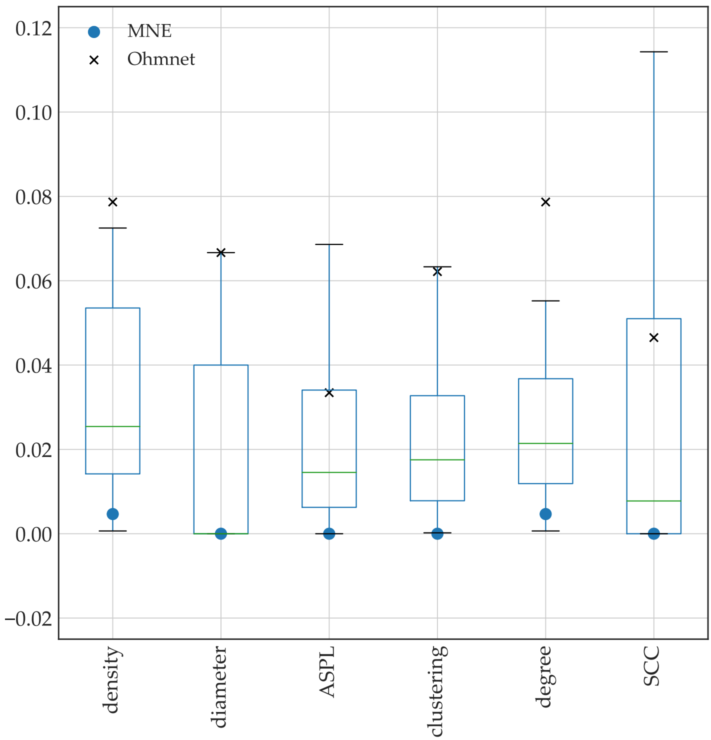

Figure 3.5: Boxplot of the average relative difference between the network structural measures calculated on the predicted graph and measured on the actual ones. Blue dots and black crosses represent the median values of the MNE and Ohmnet differences. Averages have been calculated over the five fold iterations for each of the algorithms.

In the first one we are showing the area under the ROC curve obtained in the prediction task described in Section 3.3.1 for each layer and for each algorithm. The PMNE and the network structure methods are shown as constant since they train on the whole set of edges of the multiplex. It is immediate to observe two things: first for all the methods there are layers which are more difficult to embed, hence where all the embedding methods fail to predict accurately links. These are in particular: `aid`, `FDI`, `mobility` and export of services. The second observation is the stark difference in the performance of MNE with respect to all the other methods: in Table 3.2 we have reported also the averages and standard deviation of each of the methods across all layers and it is clear to see how better is MNE. The main explanation for this is the way in which the MNE embedding method integrates information about all the layers, i.e. the common embeddings $\mathbf{b}_n$ of Equation 3.6. This allows the method to generalize information about all the layers and use it to characterize the nodes on a variety of levels, effectively exploiting all the dimensions of our data.

Our second test is shown in Figure 3.5. Here we have used the embeddings to reconstruct the full network and we have calculated the average difference between several structural network measures in the predicted and in the actual networks. In the plot we are showing the averages of all the methods for all the layers across all the 5 fold validations. To allow better comparisons we have only plotted the median value for MNE and we have also excluded the Ohmnet results from the average of the other methods since they were clearly outliers as it is shown from the position of their median value (the black crosses in the plot).

Again we can see the same patterns commented above. The MNE performance is among the best but there is a wide variability in the performance of the other methods: while their median error is usually close to the one of MNE we can see that the minimum average error (the lower "whisker") is often better than MNE, while their maximum error is several times worse than MNE. To better observe and explain this variability, for each method we have reported the mean and standard deviation of the errors in Table 3.3.

We can see that the overall performance visualized in the boxplots of Figure 3.5 is composed of two trends: the first one is the worse performance of the single layer methods, which have all higher mean errors and higher variability; the second trend instead is the better performance of the aggregated methods (structural scores and PMNE embedding) which are calculated on the full network and in one case, PMNE, make less errors than MNE. However, as in the prediction test, these methods benefit from making a single estimate for all the multiplex, hence avoiding to make big errors on the most difficult layers and scoring on average better. The downside of this is that we do not have a real prediction for each single layer but an aggregate prediction for all the multiplex which is way less useful once we want to find missing links in separate layers.

	mean	std
node2vec	0.573	0.0143
LINE	0.544	0.0121
Deepwalk	0.573	0.0131
aa_score	0.687	0.00595
jacc_score	0.668	0.00357
cc_score	0.686	0.00595

(a) Single layer methods

	mean	std
PMNE_1	0.639	0.0097
PMNE_3	0.685	0.00764
PMNE_2	0.676	0.00444
ohmnet_mdhits	0.542	0.12
ohmnet_mr_a0	0.541	0.1186
ohmnet_mr_a1	0.546	0.1174
MNE	0.791	0.0238

(b) Multilayer methods

Table 3.2: Mean and standard deviation of the area under the ROC curve scores for each method, over all layers and iterations.

3.4.2 Missing links prediction using MNE

Given the good performance of the MNE method we have chosen it to obtain a prediction of the missing links on our multiplex. Using the methodology explained in Section 3.3.2 we have obtained the results shown in Table 3.5. Here we show how many links have been predicted in each layer by the algorithm and how many of them we have accepted: since the accuracy of the algorithm is highly unstable with respect to the layer in which it is applied we are preserving only a share of the pre-

	density	diameter	ASPL	clustering	avg. degree	SCC
Deepwalk	0.0506 (0.0559)	0.037 (0.0671)	0.0312 (0.0451)	0.0408 (0.064)	0.0438 (0.0583)	0.038 (0.0545)
node2vec	0.0488 (0.0547)	0.0473 (0.0541)	0.0397 (0.049)	0.043 (0.0654)	0.0423 (0.0575)	0.0377 (0.0558)
LINE	0.0439 (0.0536)	0.0207 (0.0374)	0.0229 (0.0255)	0.0277 (0.0319)	0.0426 (0.0543)	0.0462 (0.0809)
aa_score	0.0551 (0.0)	0.0 (0.0)	0.0287 (0.0)	0.0243 (0.0)	0.0551 (0.0)	0.0 (0.0)
cc_score	0.0552 (0.0)	0.0 (0.0)	0.0287 (0.0)	0.0244 (0.0)	0.0552 (0.0)	0.0 (0.0)
jacc_score	0.0519 (0.0)	0.0 (0.0)	0.027 (0.0)	0.024 (0.0)	0.0519 (0.0)	0.0 (0.0)
PMNE_1	0.0075 (0.0)	0.0 (0.0)	0.0037 (0.0)	0.0034 (0.0)	0.0075 (0.0)	0.0 (0.0)
PMNE_2	0.0007 (0.0)	0.0 (0.0)	0.0001 (0.0)	0.0002 (0.0)	0.0007 (0.0)	0.0 (0.0)
PMNE_3	0.0116 (0.0)	0.0 (0.0)	0.0059 (0.0)	0.0077 (0.0)	0.0116 (0.0)	0.0 (0.0)
ohmnet_mr_a0	0.1929 (0.2921)	0.076 (0.0647)	0.0744 (0.0942)	0.1746 (0.2874)	0.1753 (0.2533)	7.01 (19.67)
ohmnet_mr_a1	0.1866 (0.2892)	0.0719 (0.0814)	0.0649 (0.072)	0.1681 (0.2862)	0.1715 (0.2514)	6.808 (19.69)
ohmnet_mdhits	0.1973 (0.3002)	0.0936 (0.084)	0.0676 (0.0596)	0.1813 (0.2934)	0.1812 (0.2576)	7.263 (19.89)
MNE	0.0294 (0.0477)	0.0 (0.0)	0.0107 (0.0186)	0.0127 (0.0233)	0.0106 (0.0137)	0.0284 (0.0425)

Table 3.3: Mean and standard deviation (in parenthesis) of the average error between structural measures calculated on the actual graph and on the reconstructed ones using all embedding methods.

dicted edges proportional to that accuracy. Finally in the last column we are showing which percentage of the original layer the kept edges represent.

As comparison we have carried on the same task also with another of the embeddings methods, the Jaccard score, which we show in Table 3.4. We can see that the Jaccard score is increasingly generous with its predictions with links predicted in all the layers and a very high share of the original layers predicted in certain cases. On the contrary the MNE algorithm is very conservative: it does not predict any link in the collaboration and in the infrastructure networks and for many others predict very few edges. Moreover for the layers which have shown to be difficult to predict, aid for instance, the MNE seems to predict more edges than in the other layers, a sign maybe of a difficulty in embedding their information. Accordingly also the accuracy of the algorithm is lower on these networks, hence we have accepted a smaller share of predicted links. As we have already observed, the share of newly predicted edges is very small in each layer, allowing us to use the WSBM methodology under the assumption that the new information will not alter the original block structure of the network.

	predicted	kept	share kept
migration_stock	16	13	0.13 %
citation	352	246	1.04 %
collaboration	411	273	3.14 %
mobility	586	348	6.31 %
out_tour	203	157	1.62 %
totIC	32	31	4.93 %
BIS_flow_claims	71	60	5.50 %
migration_flow	819	658	6.53 %
value	273	231	36.32 %
aid	5992	4422	179.17 %
pat_cit_inv	1985	1486	58.41 %
FDI_Greenfield	707	537	32.80 %
arms	2310	2017	766.92 %
cow_alliances	9634	9502	394.76 %

Table 3.4: Missing link prediction using Jaccard score method. In percentages it is shown which proportion of the old network is covered by the new links in the layer.

	predicted	kept	share_new
migration_stock	170	125	1.30%
citation	558	430	1.82%
collaboration	0	0	0.00%
mobility	228	123	2.23%
out_tour	362	261	2.70%
totIC	0	0	0.00%
BIS_flow_claims	16	12	1.10%
migration_flow	1344	936	9.29%
value	2	1	0.16%
aid	9280	1714	69.45%
pat_cit_inv	220	179	7.04%
FDI_Greenfield	110	84	5.13%
arms	184	137	52.09%
cow_alliances	82	76	3.16%

Table 3.5: Missing link prediction using MNE multilayer method. In percentages it is shown which proportion of the old network is covered by the new links in the layer.

3.4.3 Multirank centrality on the full multiplex

Having obtained the binary and weighted information on the new links via the MNE and the WSBM methods we have added the predicted links to our multiplex and repeated the multirank analysis of Chapter 1. However the great advantage of having an estimate of the missing links in our network is that we can safely assume (if the predictions are reliable) that all the 0s in the adjacency matrices are really observed so that we do not need to restrict ourselves to a common subset of all the nodes. As a consequence the new analysis has been run on the full multiplex of 208 nodes and to further inspect the effect of the new sample size without other possible confounding factors, we have also avoided filtering the graphs. [5]

Not surprisingly the results are not affected neither by the addition of the predicted missing links, which was predictable given their small size, nor by the inclusion of the previously missing nodes.

As we can see from Figure 3.6 the correlation between the scores of the Multirank calculated on the reduced ($\mathcal{N} = 112$) and on the extended node set ($\mathcal{N} = 208$) are really similar both at node and at layer level and with a parameter equal to 0 and 1.

Similarly the set of the top 20 countries by multirank centrality does not see huge shifts, with the notable exception that with the extended dataset finally China is included in the sample (and with a very high centrality).

Apart from that, it seems that our initial choice of nodes was sufficient for capturing the essential relationship among the countries of our international multiplex. With the prediction of the missing links we now have a more stable foundation for the use of a reduced set of countries even in absence of a portion of the data.

[5] A version of the reconstructed dataset is available at: `https://github.com/gibbbone/international_multiplex_network`

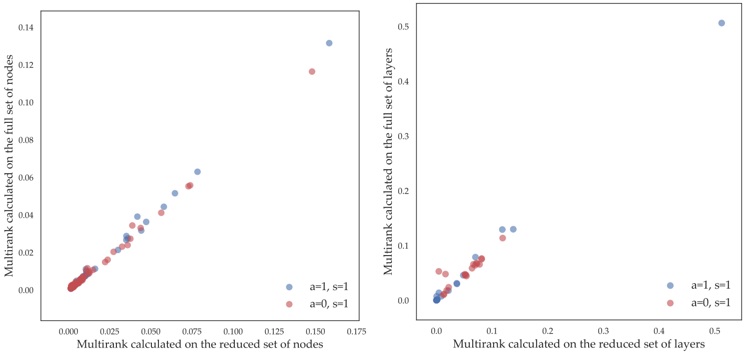

(a) Multirank centrality of nodes **(b)** Multirank influence of layers

Figure 3.6: Comparison of the multirank with full and restricted set of nodes.

full (a=1)	full (a=0)	reduced (a=1)	reduced (a=0)
USA	USA	USA	USA
DEU	GBR	DEU	GBR
GBR	DEU	GBR	DEU
FRA	FRA	FRA	FRA
JPN	CHN	ITA	ITA
CHN	JPN	CAN	JPN
ITA	ITA	JPN	ESP
CAN	ESP	BEL	CAN
ESP	CAN	ESP	NLD
BEL	NLD	NLD	BEL
NLD	BEL	MEX	RUS
MEX	HKG	RUS	MEX
KOR	RUS	TUR	GRC
HKG	KOR	IND	TUR
SGP	MEX	POL	BRA
RUS	IRL	AUS	PRT
AUS	SGP	PRT	AUS
IRL	AUS	GRC	IND
IND	GRC	BRA	POL
SWE	IND	DNK	DNK

Table 3.6: Top 20 countries by multirank calculated on the full and reduced dataset with parameter a equal to 0 and 1.

3.5 Conclusions

In this work we have employed different random-walk based network embeddings to obtain a reliable estimate of the missing links in our dataset. We have tested their performance on our data and elected as candidate for the prediction the Scalable Multiplex Network Embedding of H. Zhang et al. (2018). We have then employed weighted stochastic block-models to assign a weight to our predicted links and reanalyzed the Multirank centrality with an extended and more complete set of nodes. We have found that very little changes after these steps in the final results of the Multirank. However this allows us to have a measure of the (probable) missing information in our data and how this could impact our analyses: now we have more evidence that a restricted set of highly influential nodes can allow a sufficient analysis without the need of the full sample of the observations.

Appendix A

Appendix to Chapter 1

A.1 Data selection

Reporting agencies collect data on a fixed set of countries in different years. Hence an empty report would not imply the absence of the country from the dataset but for a given year only. To translate in network theory terms: there may be some nodes in a network whose edge weights with respect to all other nodes are 0 at a certain time t_i, but are still present in the network with positive edge weights at other t_j with $j \neq i$. These nodes belong to the network and we call them isolated nodes at time t_i. On the other hand, missing nodes don't have link to other nodes at any t.

To tell if the absence of a certain node from a specific year means its absence in the overall layer we needed to define when a node is part of our dataset. To do that we have made the following assumption: if a node is present as origin or destination at least once in the whole reported set of years for a certain layer, that node has to be considered as existing in the dataset.[1] To calculate the set of existing nodes in each layer we have used the whole span of its observations in our dataset, hence not only years 2003 and 2010.

Finally, to get a constant set of nodes across all the multiplex, the ex-

[1] A more strict assumption would be to use only nodes existing at least once as sources.

isting nodes of each layer must match across all the others. The result of this process is a set of 112 nodes common to all our layers. By referring to the *nodes* column in Table 1.1 we can get a good approximation of this process: we define the set of nodes for each layer as the union of the existing nodes in the years 2003 and 2010. Hence their size corresponds to the maximum number between t_0 and t_1 in Table 1.1 and ranges from 211 in the free trade agreement layer to 124 in the alliances layer. This leaves us with 19 different sets of nodes, one for each layer, partially overlapping. To obtain a common set of countries for the whole multiplex we calculate the subset of all the individual layers' node sets. Hence our final set is a group of 119 countries and the difference between 124 and 119 is given by five countries belonging only to the alliances layer and not to the rest of the multiplex, which we have removed.

It must be noted that this process leads automatically to select the smallest set of common countries across all layers. To avoid reducing this number too much we had to select a minimum threshold of countries to include in the multiplex with the consequence of eliminating some layers which would have had observations in the required years but with a reduced number of nodes (Figure A.7 in the Appendix shows some of the layers we have eliminated with this last criterion).

A.2 Data preprocessing

Our final result is a multiplex of 19 layers with 112 nodes in each of them, observed at two "rugged" time snapshots, i.e. with not exactly the same starting and ending year for all layers. Balancing our two cross sections on the same selection of nodes and layers ensures that we can make meaningful comparisons over time, but some further steps are required before moving on to the actual analysis.

First we have preprocessed our data to remove inconsistent observations. While in principle network weights do not have strict requirements to satisfy, in our case we needed to impose some restrictions: since our network layers are mostly directed, negative weights have no meaningful interpretation. They should instead be rewritten as positive edges in the reverse direction: if country A receives a negative flow from country B, it actually means that country B is the source of a positive flow towards A. Hence we have reversed negative flows when possible, i.e when the whole layer showed a consistent pattern of reversed flows and there were no duplicates after reversion. If instead most negative weights seemed randomly distributed or reverse flows were already recorded in the dataset we have dismissed the negative values as computational errors. [2]

Another type of requirements we imposed is to have edge weights not smaller than one, in order to avoid problems with eventual filtering operations. As one can see from the minimum weight column in Table 1.1 several layers have fractional values. Most of the time this was due to the unit of scale employed, hence it has been sufficient to rescale them with the right multiplier. This is the case for `FDI`, `FDI_Greenfield`, `BIS_flow_claims`, `aid`, `value`, `serv_exp` and `totIC` layers.

The other cases of fractional weights are the results of citations of multi-authored works which are attributed to certain countries only with the fraction related to a specific author. To deal with this we have added by default a single citation to all the fractional citations. [3].

The last restriction we have imposed on our data regards the way we intend to interpret our results. Since the basic stylized fact we want to recover from the MultiRank and MD-HITS algorithms is the division of countries in two groups, developed and underdeveloped countries (a north-south view of the word), we need our layers to be oriented in the right way: the concentration of flows toward some countries must reflect their centrality as developed nodes in the network. To achieve this we have inverted the international aid layer, as we said before, and the migration stock one which, by construction, tells us how many persons from the source country migrated to the destination one. On the contrary in a more north-centric view of the world we will need to know how many migrants the source country hosted from the destination one, since

[2]This issue affected a small number of layers with marginal effects. Only in one case, the BIS layer, the share of negative values was substantial (around 15%)

[3]Another option would have been to rescale all the citations by the minimum, but this would have made the whole distribution of weights explode.

migration flows usually points from south countries to north ones.

A.3 Choice of the MultiRank parameters

Results of the MultiRank algorithm with respect to changes in each of the three parameters a, s, γ (for nodes only the top 18 elements are shown) are reported in Figure A.13 (a) and A.13 (b). We can see that choosing a different combination of parameters alters the final ranking of nodes and layers and that some combinations experience more abrupt changes than others. While for the MD-HITS we don't need to provide any parameters in order to obtain the final rankings, for the MultiRank we need to choose how to specify them. In what follows we will show how we choose our two configurations.

The process which led us to the choice of the final configurations of the MultiRank starts by making sensible assumptions on how the ranking of layers should be performed: as we already explained in the main text the safest configuration is (1,1,1) where the content of the layers is considered as it is, without making any adjustments. The next best candidate is the configuration where $a = 0$, hence where we try to adjust for differences in the total weight of layers, which we have shown to be non negligible even after filtering, still without making any other assumptions on the importance of central nodes inside each layer, which will require us to introduce further assumptions on the data.

In Figure 1.5 one can see the two rankings of layers which are the results of the previous assumptions. Now the second step in the evaluation of the two configurations is to move to the node rankings deriving from them and compare their stability with respect to the parameter γ which from Figure A.13 seems to create disruptions in the rankings even when keeping fixed the other two parameters. We would like to have two configurations of the MultiRank which are not heavily affected by the choice of γ in at least the majority of the range of its values.

Finally the third criterion to evaluate our choice of parameters is by comparing the similarity of the node rankings with respect to the one stemming from pcGDP. We would like our choice of parameters to have a good correlation (in rank) with pcGDP which will mean that our centrality measures are capturing a good signal from the underlying data. However we do not aim to completely reconstruct the ranking of nodes by pcGDP, otherwise our measures will not have an informative content.

Our test for these last two criteria is shown in Figure A.14. Here we have plotted the Spearman rank correlation coefficient (ρ) between the nodes ranking resulting from pcGDP and the one resulting from different centrality algorithms while we let the parameter γ take different values (x-axis). The coloured lines are 4 different configuration of the MultiRank while the dashed lines are the MD-HITS hubs and authority scores (here called *MultiHub* and *MultiAuth*) which by definition are not affected by γ. For robustness we have added the single layer version of the three previous measures (PageRank, Hubs and Authority) calculated on the aggregate multiplex obtained by summing over all the layers the corresponding entries of their adjacency matrices, which are constant too.

The first thing to notice is that the ranking deriving from hub scores for the MD-HITS algorithm is the one which follows more closely the pcGDP ranking with a Spearman ρ between 0.76 and 0.78. In the cross section in 2003 also the MD-HITS authority scores and the authority scores calculated on the aggregated network have a good fit too but the same strong similarity does not hold in the cross section in 2010.

Moreover among the ranking obtained from the MultiRank we see different behaviours according to different choice of parameters in the two cross sections: only for some particular combinations of them in the 2003 cross section the MultiRank outperforms the ranking resulting from PageRank calculated on the aggregated multiplex, while in the 2010 cross section for the majority of the values of γ the ranking resulting from PageRank is better than all possible combinations of the MultiRank making it the worse choice.

We observe in fact that in the 2003 cross section the combinations with $a = 1$ are rewarded, while for $a = 0$ the Spearman correlation with pcGDP ranking is lower. On the contrary in the second cross section combination with $s = -1$ are rewarded irrespective of the magnitude of a.

Even though in these last cases we obtain a better fit with respect to pcGDP, still choosing $s = -1$, which corresponds to assume that low centrality nodes are more important is difficult to justify without further evidence and does not account for a sufficient improvement in fitting in cross section 2010. Moreover the combination $a = 1$, $s = -1$ which would be ideal in both cross section is the more affected by the choice of γ in cross section 2010, hence making it too volatile. Finally for $\gamma = 1$ has the same fitting to pcGDP ranking as our second choice $a = 0$, $s = 1$.

Hence we have two configurations of the MultiRank which are worth analyzing since they stem from sensible assumptions on the data and are sufficiently stable with respect to γ. Once we fix the parameter γ to 1, one of them, ($a = 1$,$s = 1$), has the best fitting with respect to pcGDP in cross section 2003, while the other ($a = 0$,$s = 1$) represents a good second best option in cross section 2010. Similar trade-offs could have arisen by choosing one of the other combinations, which however have stronger implications that need further empirical justifications.

A.4 Figures

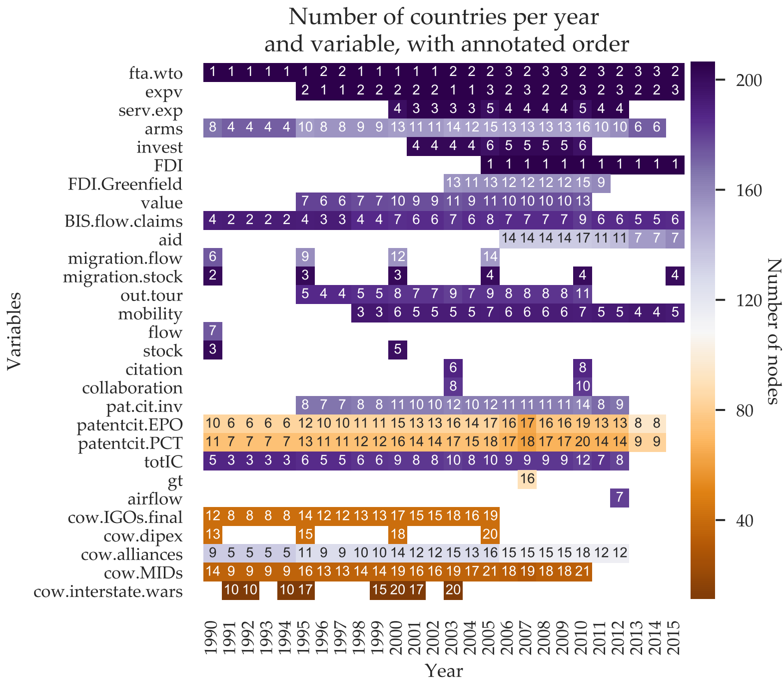

Figure A.7: Overview of the full dataset.

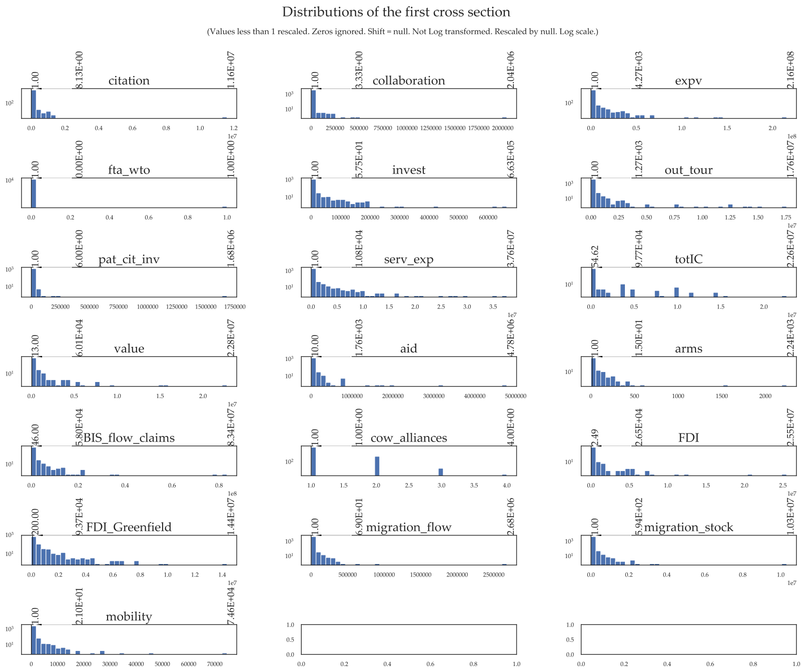

Figure A.8: Positive weights distribution in each layer of the cross section in 2003. Values rescaled to 1 if lower than 1. Logarithmic scale on the y axis.

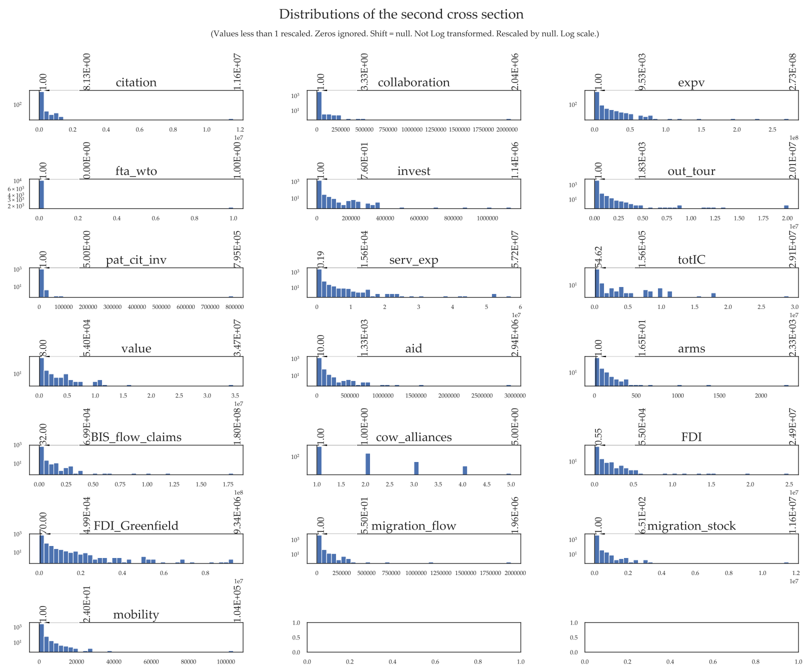

Figure A.9: Positive weights distribution in each layer of the cross section in 2010. Values rescaled to 1 if lower than 1. Logarithmic scale on the y axis.

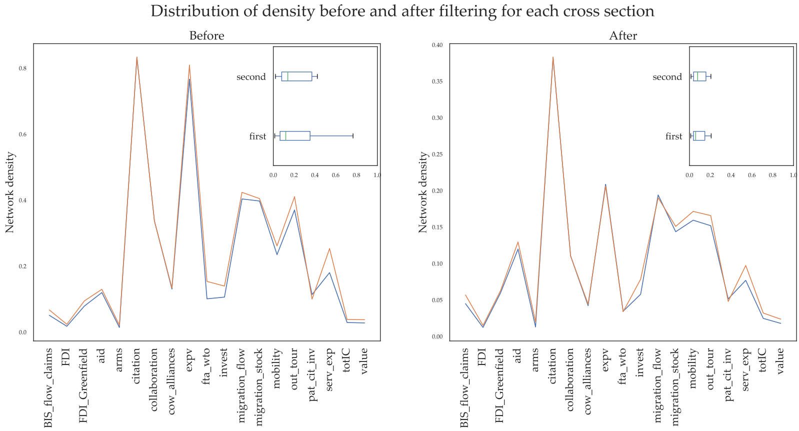

Figure A.10: How density of layers has changed: before the filter (left panel) and after (right panel). Cross section in 2003 in blue, cross section in 2010 in red. On the inset panels boxplots of the density values of the two cross sections.

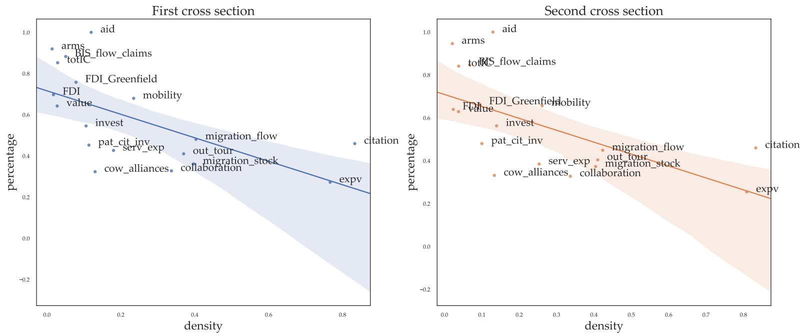

Figure A.11: Relation between percentage of values preserved by filter and density before filtering.

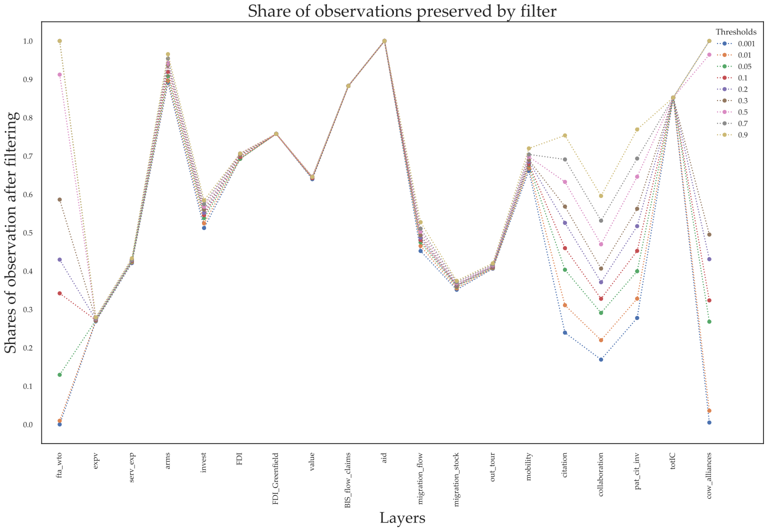

(a)

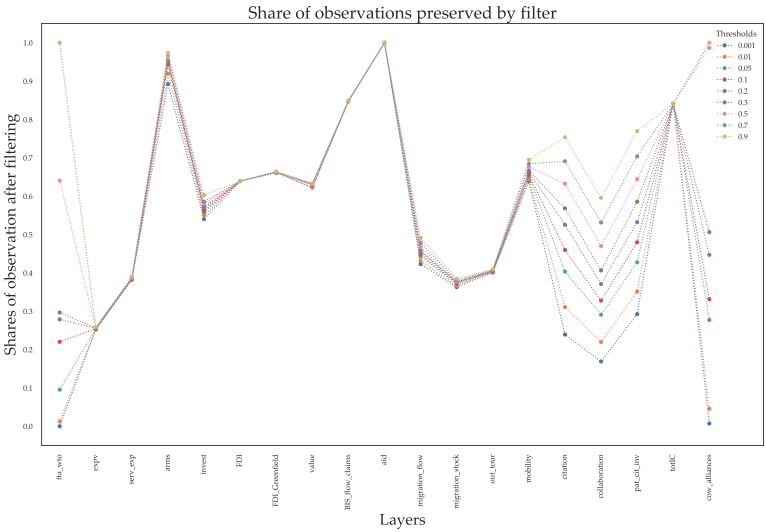

(b)

Figure A.12: Sensitivity of the preserved values after filtering with respect to the change of the threshold of the filter. Cross section in 2003 (a) and in 2010 (b).

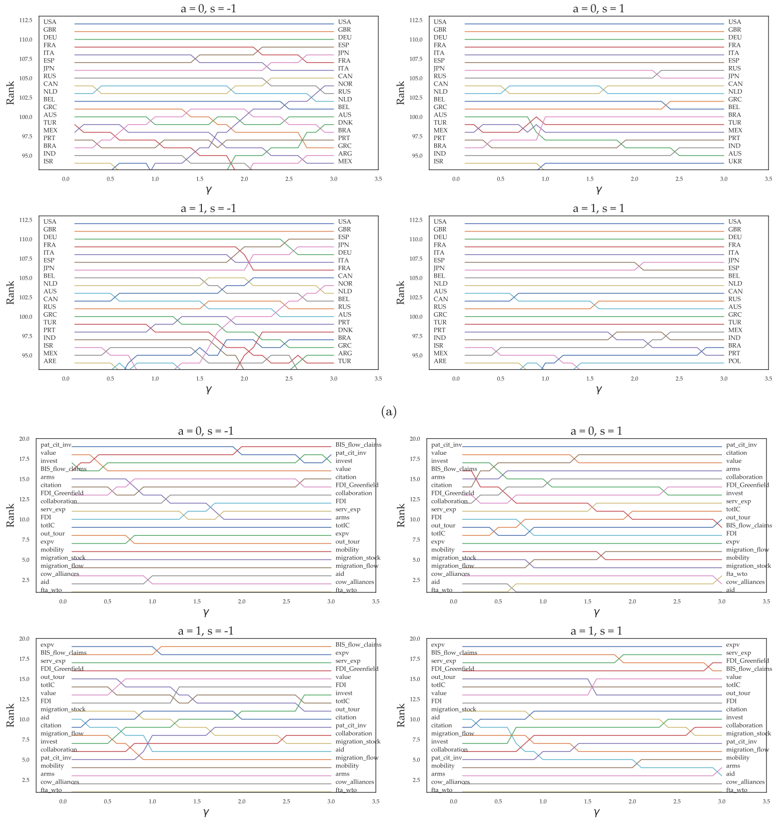

Figure A.13: Evolution of MultiRank for the 18 top countries (a) and for all layers (b) in the dataset with respect to different parameters choice.

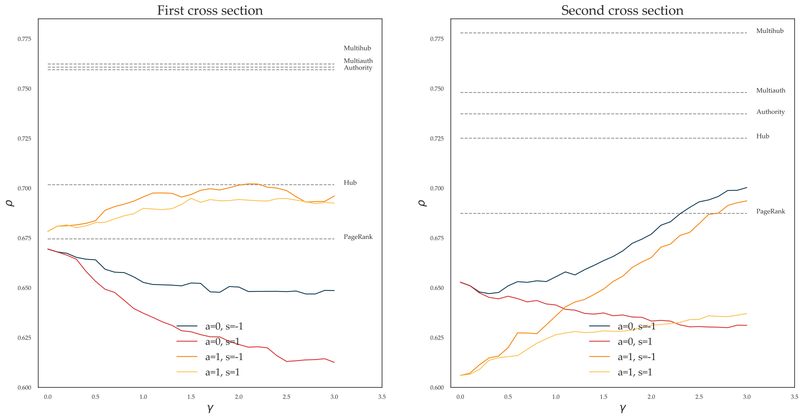

Figure A.14: Spearman correlation of rankings of different measures of node centrality with respect to ranking by pcGDP.

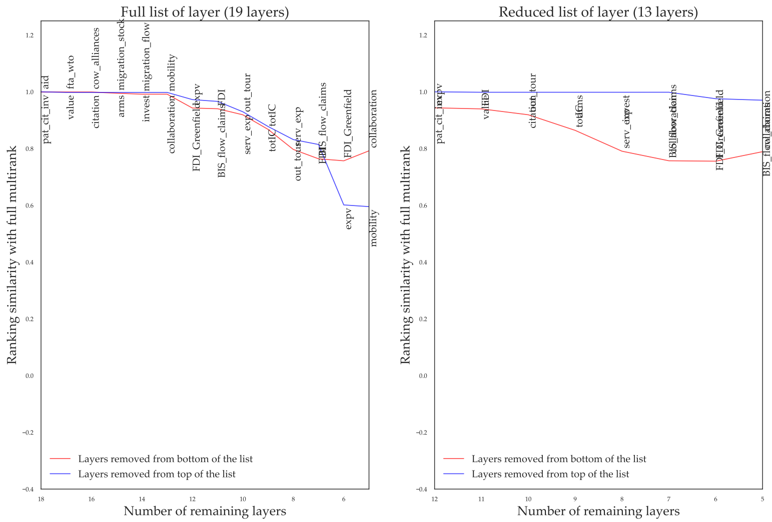

(a)

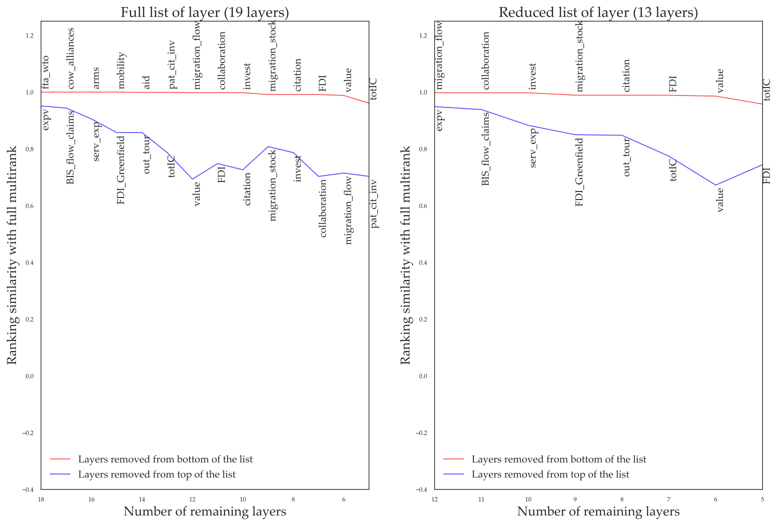

(b)

Figure A.15: Change in ranking with respect to MultiRank with parameters $s = 1, \gamma = 1$ and $a = 0$ (a) or $a = 1$ (b)

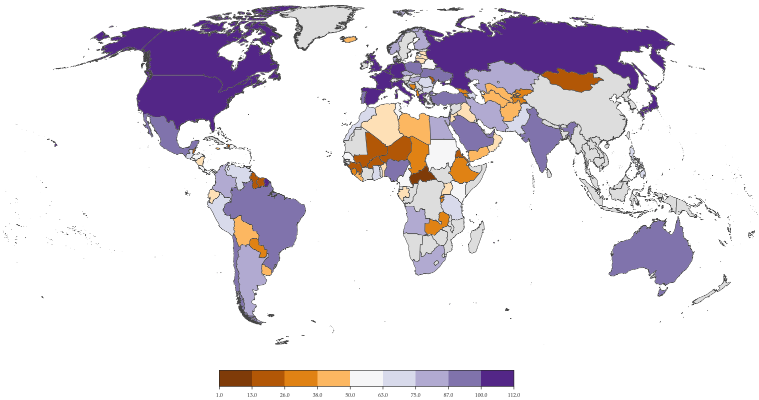

(a)

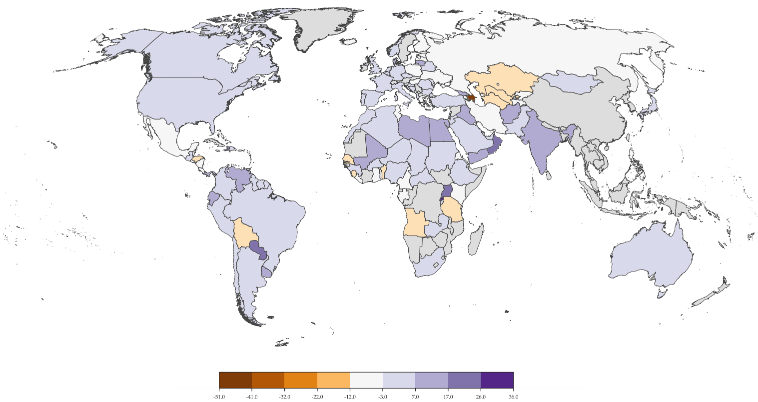

(b)

Figure A.16: Geographic distribution of the multiplex hub score (Multi-Hub). Cross section in 2003 (a) and evolution of the ranking from the first cross section to the cross section in 2010 (b).

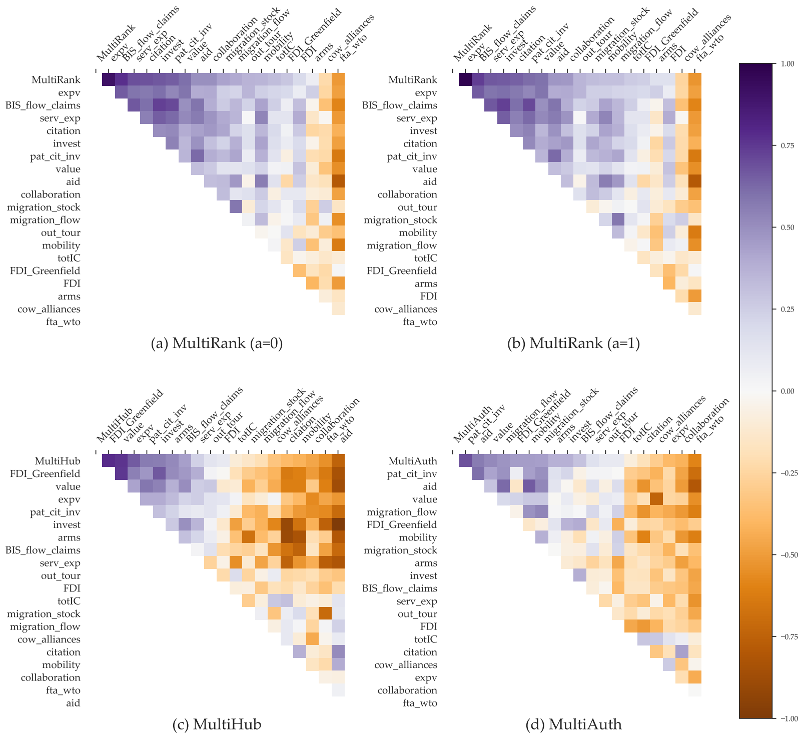

Figure A.17: Spearman correlation of the node rankings obtained by calculating centrality measures on both the single layers and the multiplex (hence using the multiplex version of the algorithms). For the MultiRank rankings the other parameters are in both cases $s = 1$ and $\gamma = 1$. All results refer to the cross section in 2010

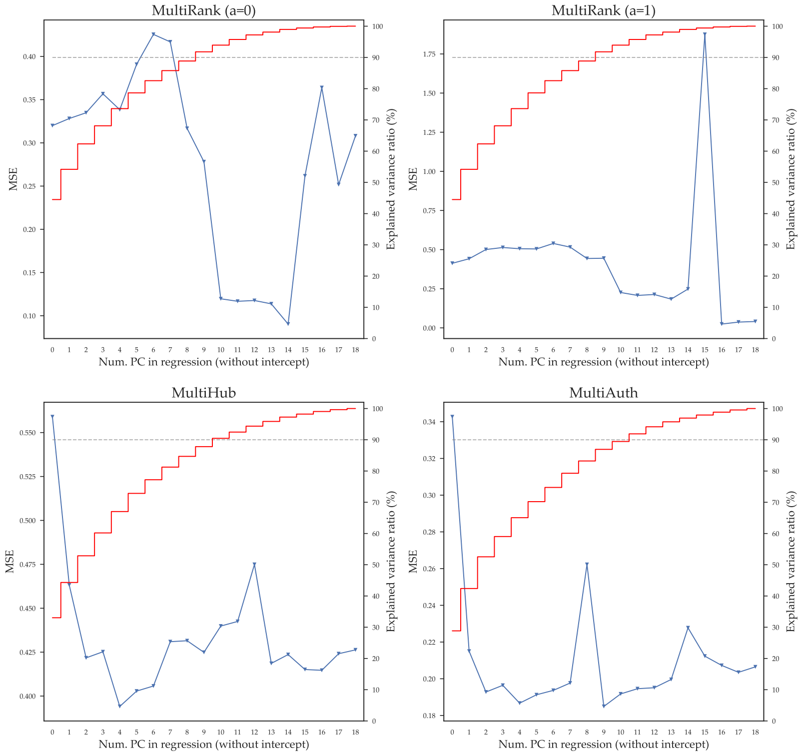

Figure A.18: In red: number of principal components sufficient to explain 90% of the variance of the measures of centrality calculated on the single layers (scale reported on the right y-axis). In blue: mean square error obtained by regressing the multilayer measures of centrality against the principal components of the single layer centralities, added one after one (scale reported on the left y-axis). All results refer to the cross section in 2010.

A.5 Tables

Table A.7: Sources and references of the variables

layer	description	units	data_source	reference
fta_wto	Trade agreements	1=RTA	WTO	**de_sousa_does_2011**
expv	Trade of commodities	Export value, thous. of US$	BACI	Gaulier and Zignago, 2010
serv_exp	Trade of services	Export value, thous. of US$	COMTRADE	-
arms	Transfers of arms	SIPRI trend-indicator value in Mil.	SIPRI	SIPRI Arms Transfers Database (website), 2019
invest	TPI: total portfolio investment	Mil. of US$	IMF	Coordinated Portfolio Investment Survey(IMF website), 2019
FDI	Total FDI financial flow	Mil. of US$	OECD	FDI financial flows - By partner country (OECD online API), 2019
FDI_Greenfield	Total FDI financial flow	Mil. of US$	-	Kirkegaard, 2013
value	Value of merge and aquisition operations	Monetary value, US$	Thomson Reuters	Worldwide Mergers, Acquisitions, and Alliances Databases SDC Platinum
BIS_flow_claims	Flow of financial claims across international banks	Mil. of US$	BIS	Locational banking statistics (BIS website), 2019
aid	Official Development Assistance Disbursements	Mil. of US$	OECD	OECD ODA disbursment (web API), n.d.
migration_flow	Flow of migrants (at 5 year intervals)	N°of people	-	G. J. Abel and Sander, 2014
migration_stock	Stock of migrants (at 5 year intervals)	N°of people	UN	United Nations, 2013, International migrant stock (direct download), 2015
out_tour	Temporary migration bilateral flows	N°of people	UN	Tourism Statistics (UN Tourism office website), 2019
mobility	Flow of mobile students at the tertiary level	N°of people	UNESCO	-
citation	Paper citations among scholars	N°of citations	-	Pan, Kaski, and Fortunato, 2012
collaboration	Coauthorship among scholars	N°of coauthorships	-	Pan, Kaski, and Fortunato, 2012
pat_cit_inv	Patent citations (by inventor country)	N°of citations	NBER	Hall, Jaffe, and Trajtenberg, 2001, NBER patent data (download), 2019
totIC	Cumulative sum of initial capacity of cable routing	Cable capacity	TeleGeography	Rossello, 2015, TeleGeography Submarine Cable Map, 2019
cow_alliances	Membership in international alliances	1= alliance	COW	Correlates of war datasets (website), 2019

* When missing the source comes only from a paper.

Table A.8: Cross section in 2010 - Network statistics, part 1

	degree skew.	indegree skew.	outdegree skew.	density	bilateral density	is strongly connected	num. scc	avg. size scc	share scc	is weakly connected	num. wcc	avg. size wcc	diameter	diameter biggest scc	avg. spl	avg. spl biggest wcc	avg. clustering coeff.	network centralization	weighted asymmetry
fta_wto	0.506	0.506	0.506	0.034	0.034		44.000	2.545	0.348		44.000	2.545	-	8.000	-	3.537	0.428	0.151	0.000
expv	0.673	0.666	0.660	0.206	0.087	✓	1.000	112.000	1.000	✓	1.000	112.000	4.000	4.000	1.977	1.977	0.389	0.415	0.562
serv_exp	2.668	2.784	2.444	0.097	0.052		3.000	37.333	0.982	✓	1.000	112.000	-	5.000	2.109	2.109	0.391	0.949	0.434
arms	3.265	1.025	3.825	0.020	0.001		107.000	1.047	0.054		37.000	3.027	-	2.000	-	0.196	0.126	0.400	0.671
invest	1.639	0.679	2.126	0.078	0.009		74.000	1.514	0.348		9.000	12.444	-	6.000	-	0.806	0.306	0.638	0.808
FDI	5.934	7.214	4.076	0.015	0.004		78.000	1.436	0.312		25.000	4.480	-	10.000	-	1.634	0.000	0.621	0.465
FDI_Greenfield	1.847	0.869	2.751	0.063	0.006		46.000	2.435	0.598		8.000	14.000	-	10.000	-	2.389	0.270	0.542	0.805
value	1.861	1.387	2.695	0.024	0.003		74.000	1.514	0.339		25.000	4.480	-	8.000	-	1.393	0.114	0.264	0.636
BIS_flow_claims	3.371	0.379	3.448	0.057	0.002		97.000	1.155	0.143		2.000	56.000	-	5.000	-	0.284	0.400	0.820	0.846
aid	1.503	1.565	-0.648	0.130	0.000		112.000	1.000	0.009		4.000	28.000	-	0.000	-	0.179	0.035	0.599	0.952
migration_flow	1.088	1.657	1.288	0.190	0.019		12.000	9.333	0.902	✓	1.000	112.000	-	10.000	2.862	2.862	0.359	0.585	0.877
migration_stock	1.859	2.105	-0.024	0.151	0.037		2.000	56.000	0.991	✓	1.000	112.000	-	9.000	2.933	2.933	0.397	0.592	0.728
out_tour	0.739	1.085	1.165	0.166	0.039		20.000	5.600	0.830	✓	1.000	112.000	-	6.000	1.945	1.945	0.345	0.396	0.737
mobility	1.348	1.486	0.261	0.171	0.007		72.000	1.556	0.348	✓	1.000	112.000	-	10.000	0.854	0.854	0.283	0.669	0.928
citation	-0.182	-0.019	-0.083	0.383	0.237	✓	1.000	112.000	1.000	✓	1.000	112.000	4.000	4.000	1.677	1.677	0.543	0.449	0.362
collaboration	0.300	0.946	1.213	0.110	0.009		112.000	1.000	0.009	✓	1.000	112.000	-	0.000	0.728	0.728	0.340	0.234	0.751
pat_cit_inv	2.308	2.529	2.556	0.048	0.009		64.000	1.750	0.429		17.000	6.588	-	6.000	-	1.444	0.505	0.673	0.613
totIC	1.214	1.214	1.214	0.032	0.032		40.000	2.800	0.598		40.000	2.800	-	9.000	-	3.872	0.547	0.247	0.000
cow_alliances	0.243	0.243	0.243	0.044	0.044		32.000	3.500	0.438		32.000	3.500	-	11.000	-	4.060	0.468	0.130	0.000

Legend: ✓ = `True`, empty space = `False`, - = `NaN`.
Abbreviations: *scc*: strongly connected components, *wcc*: weakly connected components, *spl*: shortest path length.

Appendix B

Appendix to Chapter 2

B.1 Original BMA results on new sample

B.1.1 Without interaction terms (M1)

Table B.9: Dependent variable: cum.loss. Network measure: original

	PIP	Post Mean	Post Sign.	RL	L
chg.dom.cred.bank.0006	0.990196	-0.209281	True	***	
dummy_blr	0.652980	82.809459		***	b
chg.stocks.gdp.0006	0.226506	0.002769			b
net.pf.debt.infl.0006	0.156034	-0.132258			
net.pf.equ.infl.0006	0.125816	-0.116705			
gen.govDebt.06	0.118906	-0.044018		***	
gr.savings.gdp.06	0.110369	-0.221302			
twin.fis	0.101405	0.005114			l
ext.debt.gdp.06	0.092907	-0.010232			
food.exports.0006	0.089996	0.074366			b
chg.rgdpcap.0006	0.081463	-0.095454		.	
int.res.extDebt.06	0.077147	0.027441			
adv.claims.gdp.06	0.074040	-0.017504			
trade.freedom.06	0.073205	-0.069921			
dummy_ukr	0.069632	-4.295198		***	

Table B.10: Dependent variable: Depth. Network measure: original

	PIP	Post Mean	Post Sign.	RL	L
chg.dom.cred.bank.0006	0.850753	-0.014609	True	***	
chg.stocks.gdp.0006	0.310245	0.000476			b
chg.rgdpcap.0006	0.174569	-0.030507			
outputGap_0006Exo	0.093592	-1.368720			
gen.govDebt.06	0.088946	-0.003213		**	
trade.freedom.06	0.083414	-0.008992			
food.exports.0006	0.076132	0.006881			l
twin.fis	0.071166	0.000300			
m.growth.06	0.067807	0.048371			l
dummy_blr	0.062649	0.492534		***	
outputGap_06Exo	0.058185	2.977258			
pop.gr.0006	0.054681	0.011210			b
ext.debt.gdp.06	0.050717	-0.000415			
net.pf.debt.infl.0006	0.048776	-0.003319			
size.of.gov.06	0.047985	-0.031968			

Table B.11: Dependent variable: HP.trans. Network measure: original

	PIP	Post Mean	Post Sign.	RL	L
chg.dom.cred.bank.0006	0.915299	-0.000411	True	***	
dummy_blr	0.860033	0.355244	True	***	l
rgdpgr.06	0.293219	-0.006580			
food.exports.0006	0.291387	0.000807			
gr.savings.gdp.06	0.286037	-0.001591			
int.res.extDebt.06	0.282461	0.000352			
unempl.06	0.207220	0.001478			
trade.freedom.06	0.182201	-0.000543			
sound.money.06	0.177432	0.006343			
m.growth.06	0.173514	0.003314			l
twin.fis	0.168617	0.000020			
chg.rgdpcap.0006	0.168076	-0.000516			
infl.0006	0.164447	-0.000690			
outputGap_0006Exo	0.160363	-0.053236			
depRate.06	0.143529	-0.001094			

Table B.12: Dependent variable: HP.per. Network measure: original

	PIP	Post Mean	Post Sign.	RL	L
trade.freedom.06	0.792734	-0.024236	True	***	
chg.rgdpcap.0006	0.551611	-0.018957			
chg.dom.cred.bank.0006	0.451667	-0.001147		**	
dummy_blr	0.237513	0.357576		***	b
rgdpgr.06	0.169411	-0.020510			
rgdpgr.0006	0.148592	-0.032245			
gr.savings.gdp.06	0.136961	-0.004457			
unempl.06	0.096257	0.003713			b
infl.06	0.084605	-0.006007			
emp.0006	0.072777	0.003232			b
gov.spend.06	0.065343	-0.000487			
ext.debt.gdp.06	0.055509	-0.000067			
cred.mark.reg.06	0.049728	-0.010154			
for.bank.comp.06	0.049685	-0.004683			
rgdpcap.06	0.048935	0.011745			

B.1.2 With interaction terms (M3)

Table B.13: Dependent variable: cum.loss. Network measure: original

	PIP	Post Mean	Post Sign.	RL	L
chg.dom.cred.bank.0006	0.997954	-0.169951	True	***	
dummy_blr	0.426877	52.156765		***	b
dummy_ukr	0.236893	-29.151510		***	
food.exports.0006	0.225583	0.319547			b
adv.claims.gdp.06	0.224833	0.018659			
chg.dom.cred.bank.0006 : adv.claims.gdp.06	0.189386	-0.000405			
chg.stocks.gdp.0006	0.137779	0.001496			b
net.pf.debt.infl.0006	0.085029	-0.067070			
net.pf.equ.infl.0006	0.076332	-0.071024			
gen.govDebt.06	0.067353	-0.022041		***	
infl.targeter	0.049542	0.686335			b
twin.fis	0.047408	0.001805			b
trade.freedom.06	0.047403	-0.041225			
ext.debt.gdp.06	0.047093	-0.004142			
gr.savings.gdp.06	0.043365	-0.067669			

Table B.14: Dependent variable: Depth. Network measure: original

	PIP	Post Mean	Post Sign.	RL	L
chg.dom.cred.bank.0006	0.923423	-0.016297	True	***	
chg.stocks.gdp.0006	0.203675	0.000304			b
chg.rgdpcap.0006	0.095781	-0.016611			
outputGap_0006Exo	0.052816	-0.743767			
trade.freedom.06	0.041981	-0.004620			
food.exports.0006	0.039840	0.003497			l
gen.govDebt.06	0.039292	-0.001303		**	
twin.fis	0.036980	0.000138			
m.growth.06	0.029470	0.018370			l
Mark.cap.06	0.026305	0.000395			
freedom.from.corr.06	0.024380	0.001114			
dummy_blr	0.023373	0.150973		***	
size.of.gov.06	0.022795	-0.015504			
pop.gr.0006	0.020394	0.003562			b
net.pf.debt.infl.0006	0.019923	-0.001058			

Table B.15: Dependent variable: HP.trans. Network measure: original

	PIP	Post Mean	Post Sign.	RL	L
chg.dom.cred.bank.0006	0.968449	-0.000332		***	
dummy_blr	0.694773	0.281924		***	b
food.exports.0006	0.212552	0.000641			b
infl.0006	0.198211	-0.001409			
gr.savings.gdp.06	0.184799	-0.000908			
rgdpgr.0006	0.148429	-0.001626			
rgdpgr.06	0.130429	-0.002179			
int.res.extDebt.06	0.113679	0.000108			b
m.growth.06	0.111229	0.002127			b
trade.freedom.06	0.107619	-0.000298			
outputGap_0006Exo	0.101551	-0.034471			
chg.rgdpcap.0006	0.093403	-0.000281			
depRate.06	0.090802	-0.000722			
sound.money.06	0.090735	0.002771			l
unempl.06	0.090162	0.000539			b

Table B.16: Dependent variable: HP.per. Network measure: original

	PIP	Post Mean	Post Sign.	RL	L
trade.freedom.06	0.702028	-0.021293	True	***	
chg.rgdpcap.0006	0.549763	-0.019617			
chg.dom.cred.bank.0006	0.445992	-0.001207		**	
dummy_blr	0.138837	0.197353		***	b
rgdpgr.0006	0.127793	-0.029715			
rgdpgr.06	0.100439	-0.011456			
gr.savings.gdp.06	0.069836	-0.001986			
infl.06	0.048814	-0.003304			
unempl.06	0.045674	0.001572			b
gov.spend.06	0.043677	-0.000334			
emp.0006	0.034943	0.001424			b
ext.debt.gdp.06	0.029005	-0.000031			
size.of.gov.06	0.028833	-0.002998			
dummy_ukr	0.027156	-0.025210		***	
net.pf.debt.infl.0006	0.027067	-0.000223			

B.2 BMA results with full set of network variables

B.2.1 Without interaction terms (M1)

Table B.17: Dependent variable: cum.loss. Network measure: centr

	PIP	Post Mean	Post Sign.	RL	ERL	L
chg.dom.cred.bank.0006	0.996004	-0.193969	True	***	***	
invest_kcore_nr	0.882947	-3.957268	True	*	*	
dummy_blr	0.055094	5.704390		**	**	b
collaboration_kcore_nr	0.045075	0.302184			*	b
net.pf.equ.infl.0006	0.044412	-0.037697				
migration_stock_kcore_nr	0.035613	0.129549			*	b
migration_flow_betweenness_nr	0.032843	-5.531874				
value_pagerank_nr	0.032136	15.943713				l
dummy_ukr	0.023023	-1.366911		***	***	
totIC_instrength_nr	0.022768	-1.085531			**	
mobility_kcore_nr	0.021153	0.091134				
mobility_authority_nr	0.020691	12.464149				
chg.stocks.gdp.0006	0.019599	0.000180				l
twin.fis	0.019445	0.000615				b
gen.govDebt.06	0.019028	-0.005205		***	***	

Table B.18: Dependent variable: Depth. Network measure: centr

	PIP	Post Mean	Post Sign.	RL	ERL	L
chg.dom.cred.bank.0006	0.914681	-0.014326	True	***	***	
migration_stock_kcore_nr	0.539577	0.381900			***	b
invest_kcore_nr	0.538689	-0.248409				
pat_cit_inv_kcore_nr	0.476569	-0.266488				
FDI_Greenfield_instrength_nr	0.476027	4.551346			.	
totIC_instrength_nr	0.463028	-4.848323			**	
rgdpcap.06	0.277205	0.864655				
chg.stocks.gdp.0006	0.218597	0.000291				b
arms_pagerank_nr	0.138018	13.480885			.	
citation_betweenness_nr	0.127088	-1.185583				
outputGap_0006Exo	0.124427	-1.756882				
expv_hubs_nr	0.106457	-11.590491				
floater	0.105607	0.196811				
size.of.gov.06	0.098223	-0.145490				
rgdpgr.0006	0.096203	-0.120991				

Table B.19: Dependent variable: HP.trans. Network measure: centr

	PIP	Post Mean	Post Sign.	RL	ERL	L
chg.dom.cred.bank.0006	0.864131	-0.000376	True	***	***	
dummy_blr	0.622907	0.242757		***	***	h
migration_stock_kcore_nr	0.375009	0.005562			**	h
FDI_Greenfield_kcore_nr	0.202929	-0.003388			.	
rgdpgr.0006	0.175083	-0.006417				
aid_betweenness_nr	0.163167	-0.552310			***	
invest_kcore_nr	0.148889	-0.001506				
value_kcore_nr	0.138075	0.002586				
rgdpgr.06	0.132924	-0.002275				
trade.freedom.06	0.126042	-0.000451				
gr.savings.gdp.06	0.116358	-0.000611				
FDI_Greenfield_hubs_nr	0.105495	-0.283014			***	
int.res.extDebt.06	0.101886	0.000118				
depRate.06	0.101780	-0.000965				
pat_cit_inv_kcore_nr	0.101513	-0.000955				

Table B.20: Dependent variable: HP.per. Network measure: centr

	PIP	Post Mean	Post Sign.	RL	ERL	L
trade.freedom.06	0.907635	-0.032518	True			
rgdpgr.0006	0.737155	-0.238038	True			
arms_kcore_nr	0.515491	0.084696			.	b
invest_kcore_nr	0.426259	-0.022991			*	
migration_stock_outstrength_nr	0.324349	-0.409899			.	
chg.dom.cred.bank.0006	0.271497	-0.000574		*	*	
migration_flow_betweenness_nr	0.228866	-0.641719		.	.	
arms_authority_nr	0.184124	2.673576				b
FDI_Greenfield_hubs_nr	0.173072	-2.967023			***	
FDI_Greenfield_kcore_nr	0.157677	-0.012578				
chg.rgdpcap.0006	0.142150	-0.004653				
pat_cit_inv_kcore_nr	0.103951	-0.005218				
dummy_blr	0.097088	0.118411		***	***	b
migration_stock_kcore_nr	0.074789	0.004357			.	b
gov.spend.06	0.071056	-0.000940				

B.3 BMA results with only kcore centralities

B.3.1 Without interaction terms (M1)

Table B.21: Dependent variable: cum.loss. Network measure: kcore

	PIP	Post Mean	Post Sign.	RL	ERL	L
chg.dom.cred.bank.0006	0.996589	-0.193375	True	***	***	
invest_kcore_nr	0.913447	-4.081269	True	***	***	
net.pf.equ.infl.0006	0.092341	-0.081941				
collaboration_kcore_nr	0.079319	0.506029			**	b
dummy_blr	0.078006	7.602253		***	***	b
migration_stock_kcore_nr	0.069573	0.252032			**	b
gen.govDebt.06	0.042261	-0.012355		***	***	
dummy_ukr	0.042036	-2.442006		***	***	
twin.fis	0.040403	0.001346				b
net.pf.debt.infl.0006	0.039193	-0.022468				
mobility_kcore_nr	0.038861	0.159460			*	
chg.stocks.gdp.0006	0.035835	0.000316				b
outputGap_0006Exo	0.035443	-3.538792				
trade.freedom.06	0.033166	-0.026626		*	*	
arms_kcore_nr	0.030527	0.138219			.	b

Table B.22: Dependent variable: Depth. Network measure: kcore

	PIP	Post Mean	Post Sign.	RL	ERL	L
chg.dom.cred.bank.0006	0.792740	-0.011859	True	***	***	
invest_kcore_nr	0.602631	-0.250068			***	
migration_stock_kcore_nr	0.558166	0.361456			*	h
chg.stocks.gdp.0006	0.214344	0.000286				b
outputGap_0006Exo	0.187540	-3.001745				
rgdpgr.0006	0.160131	-0.197833				
chg.rgdpcap.0006	0.158581	-0.023116				
emp.0006	0.140615	-0.041960				
trade.freedom.06	0.134988	-0.015629				
expv_kcore_nr	0.129593	-0.060104			*	
gen.govDebt.06	0.110795	-0.003966		*	*	
pat_cit_inv_kcore_nr	0.104376	-0.030735				
FDI_Greenfield_kcore_nr	0.090430	-0.037776			**	
collaboration_kcore_nr	0.089778	0.048960				
mobility_kcore_nr	0.089603	0.044919				

Table B.23: Dependent variable: HP.trans. Network measure: kcore

	PIP	Post Mean	Post Sign.	RL	ERL	L
chg.dom.cred.bank.0006	0.883112	-0.000315	True	***	***	
migration_stock_kcore_nr	0.650883	0.010118				h
dummy_blr	0.649971	0.256761		***	***	h
trade.freedom.06	0.408269	-0.001623				
FDI_Greenfield_kcore_nr	0.405665	-0.007254				
value_kcore_nr	0.382445	0.007303				b
invest_kcore_nr	0.370206	-0.003828			**	
bus.reg.06	0.361473	-0.016327				
rgdpgr.06	0.325581	-0.006727				
rgdpcap.06	0.313460	0.026449				l
int.res.extDebt.06	0.243184	0.000287				b
depRate.06	0.202520	-0.001729				
rgdpgr.0006	0.197799	-0.005974				
migration_flow_kcore_nr	0.190862	0.002867				l
pat_cit_inv_kcore_nr	0.181733	-0.001480			*	

Table B.24: Dependent variable: HP.per. Network measure: kcore

	PIP	Post Mean	Post Sign.	RL	ERL	L
trade.freedom.06	0.934515	-0.030406	True	***	***	
arms_kcore_nr	0.750060	0.118763	True		***	b
rgdpgr.0006	0.739034	-0.216284	True			
invest_kcore_nr	0.495432	-0.027608			**	
FDI_Greenfield_kcore_nr	0.373809	-0.035036				
chg.dom.cred.bank.0006	0.326871	-0.000582		**	**	
migration_stock_kcore_nr	0.253954	0.016486				b
trade.to.eu15.totExp.0006	0.163201	0.000533				b
dummy_blr	0.142938	0.172905		***	***	b
chg.rgdpcap.0006	0.134113	-0.003819				
rgdpgr.06	0.129033	-0.016883				
for.bank.comp.06	0.115329	-0.012644				
infl.06	0.112519	-0.009257				
mon.freedom.06	0.103421	-0.003008				
pop.06	0.103283	-0.013911				

B.3.2 With interaction terms (M3)

Table B.25: Dependent variable: cum.loss. Network measure: kcore

	PIP	Post Mean	Post Sign.	RL	ERL	L
chg.dom.cred.bank.0006	0.999049	-0.182986	True			
invest_kcore_nr	0.925301	-4.019574	True	*	*	
net.pf.equ.infl.0006	0.064062	-0.057175				
collaboration_kcore_nr	0.054527	0.345634				b
migration_stock_kcore_nr	0.053847	0.178639			*	b
dummy_blr	0.047688	4.392633		**	**	
invest_kcore_nr : chg.dom.cred.bank.0006	0.045940	-0.000425				
dummy_ukr	0.035800	-2.706387				
mobility_kcore_nr	0.027007	0.110260		*	*	
chg.stocks.gdp.0006	0.025753	0.000126				
net.pf.debt.infl.0006	0.024805	-0.014114				
twin.fis	0.024493	0.000769				b
gen.govDebt.06	0.024191	-0.006489				
adv.claims.gdp.06	0.023263	0.000609				
outputGap_0006Exo	0.022745	-2.288245				

Table B.26: Dependent variable: Depth. Network measure: kcore

	PIP	Post Mean	Post Sign.	RL	ERL	L
chg.dom.cred.bank.0006	0.908059	-0.014878	True			
invest_kcore_nr	0.379052	-0.146408			**	
migration_stock_kcore_nr	0.347743	0.218676			**	b
chg.stocks.gdp.0006	0.127792	0.000175				b
outputGap_0006Exo	0.101993	-1.631124				
chg.rgdpcap.0006	0.081640	-0.012639				
trade.freedom.06	0.055667	-0.006245				
rgdpgr.0006	0.054217	-0.059433				
emp.0006	0.048670	-0.013170				
expv_kcore_nr	0.046464	-0.016348			**	
twin.fis	0.034418	0.000115				
gen.govDebt.06	0.033738	-0.001038				
pat_cit_inv_kcore_nr	0.033444	-0.008249			*	
collaboration_kcore_nr	0.029339	0.015106				l
int.res.extDebt.06	0.028631	0.000875				

Table B.27: Dependent variable: HP.trans. Network measure: kcore

	PIP	Post Mean	Post Sign.	RL	ERL	L
chg.dom.cred.bank.0006	0.951186	-0.000299				
migration_stock_kcore_nr	0.557428	0.006484		*	*	b
dummy_blr	0.452533	0.174701		*	*	b
invest_kcore_nr	0.345099	-0.003377				
bus.reg.06	0.292619	-0.011964				
pat_cit_inv_kcore_nr	0.280455	-0.002895			**	
FDI_Greenfield_kcore_nr	0.229537	-0.003380			.	
migration_stock_kcore_nr : chg.dom.cred.bank.0006	0.220308	0.000013				
trade.freedom.06	0.208082	-0.000672				
rgdpcap.06	0.202407	0.015173				b
rgdpgr.06	0.175851	-0.002702				
rgdpgr.0006	0.168357	-0.003971				
value_kcore_nr	0.148250	0.002448				
FDI_kcore_nr	0.139905	0.000136				
infl.0006	0.134302	-0.000807				

Table B.28: Dependent variable: HP.per. Network measure: kcore

	PIP	Post Mean	Post Sign.	RL	ERL	L
trade.freedom.06	0.850759	-0.027122	True			
rgdpgr.0006	0.636798	-0.187927				
arms_kcore_nr	0.568257	0.088220			.	h
invest_kcore_nr	0.405075	-0.022767			*	
chg.dom.cred.bank.0006	0.395451	-0.000234				
FDI_Greenfield_kcore_nr	0.247951	-0.021898				
chg.rgdpcap.0006	0.186478	-0.006272				
migration_stock_kcore_nr	0.150872	0.009509				h
trade.to.eu15.totExp.0006	0.101056	0.000312				h
dummy_blr	0.096502	0.117861		.	.	b
pat_cit_inv_kcore_nr	0.081135	-0.004232			*	
rgdpgr.06	0.076507	-0.009150				
for.bank.comp.06	0.067749	-0.007532				
pop.06	0.061641	-0.007813				
expv_kcore_nr	0.054860	0.001751			.	

B.4 Regression results with only kcore centralities

B.4.1 Without interaction terms (M1)

Table B.29: Regression results. Comparison of top 5 models with the best original model.

	Dependent variable:					
	Depth					
	1	2	3	4	5	ref.
invest_kcore_nr		-0.392^{***}			-0.446^{***}	
		(0.097)			(0.093)	
migration_stock_kcore_nr		0.716^{***}		0.499^{***}	0.607^{***}	
		(0.172)		(0.186)	(0.166)	
chg.dom.cred.bank.0006	-0.019^{***}	-0.015^{***}	-0.020^{***}	-0.017^{***}	-0.015^{***}	-0.020^{***}
	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)
chg.stocks.gdp.0006			0.001^{***}			0.001^{***}
			(0.001)			(0.001)
outputGap_0006Exo					-15.972^{***}	
					(5.589)	
Constant	-3.524^{***}	-10.069^{***}	-4.098^{***}	-12.592^{***}	-4.767	-4.098^{***}
	(0.701)	(3.100)	(0.690)	(3.453)	(3.445)	(0.690)
Observations	55	55	55	55	55	55
R^2	0.506	0.671	0.572	0.566	0.717	0.572
Adjusted R^2	0.497	0.652	0.555	0.549	0.695	0.555
Residual Std. Error	3.733	3.107	3.510	3.533	2.909	3.510
F Statistic	54.378^{***}	34.689^{***}	34.722^{***}	33.931^{***}	31.715^{***}	34.722^{***}
BF	0.00027	0.15348	0.00147	0.00107	1	1
Logmarg	15.08684	21.42833	16.78135	16.46169	23.30254	16.78135
Post. Prob	0.07166	0.01886	0.00681	0.00534	0.00449	0.02899
Dim	2	4	3	3	5	3
Note:						$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Table B.30: Regression results. Comparison of top 5 models with the best original model.

	Dependent variable:					
	HP_per					
	1	2	3	4	5	ref.
value_kcore_nr			0.018***			
			(0.005)			
FDI_Greenfield_kcore_nr			−0.022***			
			(0.004)			
invest_kcore_nr				−0.013***		
				(0.002)		
pat_cit_inv_kcore_nr				−0.010***		
				(0.002)		
migration_stock_kcore_nr			0.017***	0.017***		
			(0.003)	(0.003)		
gr.savings.gdp.06						−0.005***
						(0.002)
chg.dom.cred.bank.0006	−0.0005***	−0.001***	−0.0004***	−0.0002***	−0.001***	−0.001***
	(0.0001)	(0.0001)	(0.00005)	(0.00005)	(0.0001)	(0.0001)
depRate.06			−0.010***			
			(0.002)			
infl.0006					−0.007***	−0.008***
					(0.002)	(0.002)
dummy_blr		0.302***	0.427***		0.620***	0.707***
		(0.089)	(0.065)		(0.133)	(0.126)
rgdpgr.06				−0.020***		
				(0.004)		
trade.freedom.06				−0.004***		
				(0.001)		
bus.reg.06				−0.040***		
				(0.008)		
rgdpcap.06				0.088***		
				(0.015)		
Constant	−0.092***	−0.087***	−0.224***	−0.303**	−0.055***	0.059
	(0.018)	(0.016)	(0.062)	(0.125)	(0.018)	(0.041)
Observations	55	55	55	55	55	55
R^2	0.522	0.608	0.829	0.883	0.670	0.721
Adjusted R^2	0.513	0.593	0.807	0.862	0.650	0.699
Residual Std. Error	0.093	0.085	0.059	0.050	0.079	0.073
F Statistic	57.980***	40.407***	38.686***	43.310***	34.442***	32.321***
BF	0	0	0.00215	0.26733	0	1
Logmarg	15.91348	18.97027	31.39138	36.21654	21.31348	23.6231
Post. Prob	0.00675	0.00257	0.00137	0.00073	0.00073	0.00174
Dim	2	3	7	9	4	5
Note:						*p<0.1; **p<0.05; ***p<0.01

Table B.31: Regression results. Comparison of top 5 models with the best original model.

	Dependent variable:					
	HP_trans					
	1	2	3	4	5	ref.
invest_kcore_nr			-0.060^{***}		-0.081^{***}	
			(0.013)		(0.013)	
FDI_Greenfield_kcore_nr				-0.097^{***}		
				(0.021)		
arms_kcore_nr			0.123^{***}	0.212^{***}	0.138^{***}	
			(0.033)	(0.039)	(0.030)	
rgdpgr.0006			-0.365^{***}	-0.317^{***}	-0.363^{***}	
			(0.031)	(0.032)	(0.029)	
trade.freedom.06	-0.031^{***}		-0.037^{***}	-0.035^{***}	-0.038^{***}	-0.030^{***}
	(0.006)		(0.006)	(0.006)	(0.005)	(0.006)
chg.rgdpcap.0006	-0.039^{***}					-0.042^{***}
	(0.004)					(0.004)
chg.dom.cred.bank.0006		-0.003^{***}				
		(0.0004)				
trade.to.eu15.totExp.0006					0.003^{***}	
					(0.001)	
dummy_blr						1.225^{**}
						(0.502)
Constant	5.917^{***}	-0.614^{***}	3.591^{***}	3.030^{***}	3.601^{***}	6.122^{***}
	(0.721)	(0.105)	(0.478)	(0.465)	(0.436)	(0.694)
Observations	55	55	55	55	55	55
R^2	0.697	0.606	0.800	0.799	0.837	0.729
Adjusted R^2	0.685	0.599	0.784	0.783	0.820	0.713
Residual Std. Error	0.497	0.560	0.411	0.412	0.375	0.475
F Statistic	59.737^{***}	81.650^{***}	50.135^{***}	49.816^{***}	50.286^{***}	45.616^{***}
BF	1e-05	0	0.00467	0.00413	0.07992	1
Logmarg	25.28173	20.7911	31.56105	31.43887	34.40068	26.05862
Post. Prob	0.02075	0.0141	0.00957	0.0086	0.0075	0.00808
Dim	3	2	5	5	6	4
Note:						*p<0.1; **p<0.05; ***p<0.01

B.4.2 With interaction terms (M3)

Table B.32: Regression results. Comparison of top 5 models with the best original model.

	Dependent variable:					
	cum_loss					
	1	2	3	4	5	ref.
invest_kcore_nr	-4.503^{***}		-5.182^{***}	-2.384^{*}	-4.691^{***}	
	(0.923)		(0.932)	(1.304)	(0.891)	
migration_stock_kcore_nr			3.869^{**}			
			(1.652)			
collaboration_kcore_nr					6.464^{**}	
					(2.808)	
chg.dom.cred.bank.0006	-0.201^{***}	-0.213^{***}	-0.185^{***}	-0.037	-0.213^{***}	-0.019
	(0.021)	(0.025)	(0.022)	(0.077)	(0.021)	(0.035)
invest_kcore_nr:chg.dom.cred.bank.0006				-0.009^{**}		
				(0.004)		
adv.claims.gdp.06						0.142
						(0.111)
food.exports.0006						1.676^{***}
						(0.349)
dummy_ukr						-133.524^{***}
						(28.570)
chg.dom.cred.bank.0006:adv.claims.gdp.06						-0.002^{***}
						(0.0004)
Constant	63.521^{***}	-10.658	4.390	27.311	37.736^{*}	-42.443^{***}
	(16.281)	(6.970)	(29.698)	(22.615)	(19.244)	(8.623)
Observations	55	55	55	55	55	55
R^2	0.708	0.574	0.736	0.734	0.735	0.811
Adjusted R^2	0.697	0.566	0.721	0.718	0.720	0.792
Residual Std. Error	31.020	37.099	29.764	29.906	29.813	25.707
F Statistic	63.017^{***}	71.470^{***}	47.459^{***}	46.849^{***}	47.250^{***}	42.048^{***}
BF	0.57865	0.00035	1	0.79405	0.92427	1
Logmarg	26.21542	18.80165	26.76248	26.53187	26.68373	30.95571
Post. Prob	0.38085	0.01484	0.01339	0.01242	0.01237	0.02435
Dim	3	2	4	4	4	6

Note: *p<0.1; **p<0.05; ***p<0.01

Table B.33: Regression results. Comparison of top 5 models with the best original model.

	Dependent variable:					
	Depth					
	1	2	3	4	5	ref.
invest_kcore_nr		-0.392^{***}			-0.267^{**}	
		(0.097)			(0.106)	
migration_stock_kcore_nr		0.716^{***}		0.499^{***}		
		(0.172)		(0.186)		
chg.dom.cred.bank.0006	-0.019^{***}	-0.015^{***}	-0.020^{***}	-0.017^{***}	-0.018^{***}	-0.020^{***}
	(0.003)	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)
chg.stocks.gdp.0006			0.001^{***}			0.001^{***}
			(0.001)			(0.001)
Constant	-3.524^{***}	-10.069^{***}	-4.098^{***}	-12.592^{***}	0.868	-4.098^{***}
	(0.701)	(3.100)	(0.690)	(3.453)	(1.867)	(0.690)
Observations	55	55	55	55	55	55
R^2	0.506	0.671	0.572	0.566	0.560	0.572
Adjusted R^2	0.497	0.652	0.555	0.549	0.543	0.555
Residual Std. Error	3.733	3.107	3.510	3.533	3.558	3.510
F Statistic	54.378^{***}	34.689^{***}	34.722^{***}	33.931^{***}	33.104^{***}	34.722^{***}
BF	0.00027	0.15348	0.00147	0.00107	0.00076	1
Logmarg	15.08684	21.42833	16.78135	16.46169	16.12383	16.78135
Post. Prob	0.2559	0.03616	0.01769	0.01304	0.00891	0.0362
Dim	2	4	3	3	3	3
Note:					$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$	

Table B.34: Regression results. Comparison of top 5 models with the best original model.

	Dependent variable:					
	HP_per					
	1	2	3	4	5	ref.
FDI_kcore_nr			0.020 (0.016)			
migration_stock_kcore_nr				0.013^{***} (0.004)		
gr.savings.gdp.06						-0.005^{***} (0.002)
chg.dom.cred.bank.0006	-0.0005^{***} (0.0001)	-0.001^{***} (0.0001)	0.001^{**} (0.0002)	-0.0005^{***} (0.0001)	-0.001^{***} (0.0001)	-0.001^{***} (0.0001)
infl.0006					-0.007^{***} (0.002)	-0.008^{***} (0.002)
dummy_blr		0.302^{***} (0.089)		0.344^{***} (0.084)	0.620^{***} (0.133)	0.707^{***} (0.126)
FDI_kcore_nr:chg.dom.cred.bank.0006			-0.0003^{***} (0.0001)			
Constant	-0.092^{***} (0.018)	-0.087^{***} (0.016)	-0.167^{***} (0.041)	-0.322^{***} (0.078)	-0.055^{***} (0.018)	0.059 (0.041)
Observations	55	55	55	55	55	55
R^2	0.522	0.608	0.681	0.669	0.670	0.721
Adjusted R^2	0.513	0.593	0.662	0.650	0.650	0.699
Residual Std. Error	0.093	0.085	0.078	0.079	0.079	0.073
F Statistic	57.980^{***}	40.407^{***}	36.316^{***}	34.425^{***}	34.442^{***}	32.321^{***}
BF	0	0	0	0	0	1
Logmarg	15.91348	18.97027	22.17321	21.3053	21.31348	23.6231
Post. Prob	0.04339	0.01337	0.00992	0.00277	0.00263	0.00325
Dim	2	3	4	4	4	5
Note:						*p<0.1; **p<0.05; ***p<0.01

Table B.35: Regression results. Comparison of top 5 models with the best original model.

	Dependent variable:					
	HP_trans					
	1	2	3	4	5	ref.
chg.dom.cred.bank.0006	-0.003^{***}					
	(0.0004)					
invest_kcore_nr			-0.060^{***}		-0.057^{***}	
			(0.013)		(0.015)	
FDI_Greenfield_kcore_nr				-0.097^{***}		
				(0.021)		
arms_kcore_nr			0.123^{***}	0.212^{***}		
			(0.033)	(0.039)		
rgdpgr.0006			-0.365^{***}	-0.317^{***}	-0.390^{***}	
			(0.031)	(0.032)	(0.034)	
trade.freedom.06		-0.031^{***}	-0.037^{***}	-0.035^{***}	-0.033^{***}	-0.030^{***}
		(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
chg.rgdpcap.0006		-0.039^{***}				-0.042^{***}
		(0.004)				(0.004)
dummy_blr						1.225^{**}
						(0.502)
Constant	-0.614^{***}	5.917^{***}	3.591^{***}	3.030^{***}	3.676^{***}	6.122^{***}
	(0.105)	(0.721)	(0.478)	(0.465)	(0.535)	(0.694)
Observations	55	55	55	55	55	55
R^2	0.606	0.697	0.800	0.799	0.744	0.729
Adjusted R^2	0.599	0.685	0.784	0.783	0.729	0.713
Residual Std. Error	0.560	0.497	0.411	0.412	0.461	0.475
F Statistic	81.650^{***}	59.737^{***}	50.135^{***}	49.816^{***}	49.466^{***}	45.616^{***}
BF	0	1e-05	0.00467	0.00413	8e-05	1
Logmarg	20.7911	25.28173	31.56105	31.43887	27.50779	26.05862
Post. Prob	0.06118	0.05912	0.0183	0.01718	0.01234	0.00991
Dim	2	3	5	5	4	4
Note:						*p<0.1; **p<0.05; ***p<0.01

B.5 Network representation of the results

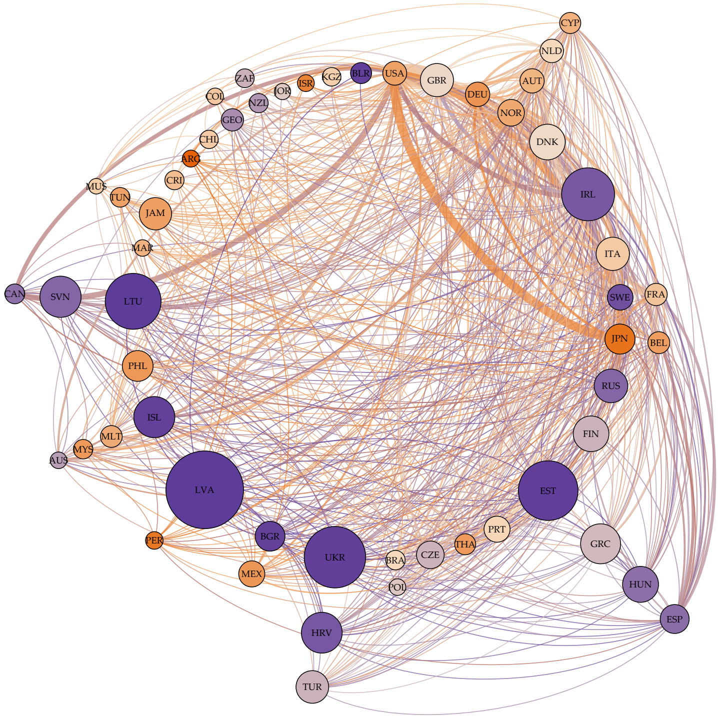

Figure B.19: Network representation of the investment layer with colours based on the intensity of the change in domestic credit variable. Nodes are ordered on the circle by their degree with size representing the cumulative loss variable and edge width representing the flow of investments among them.

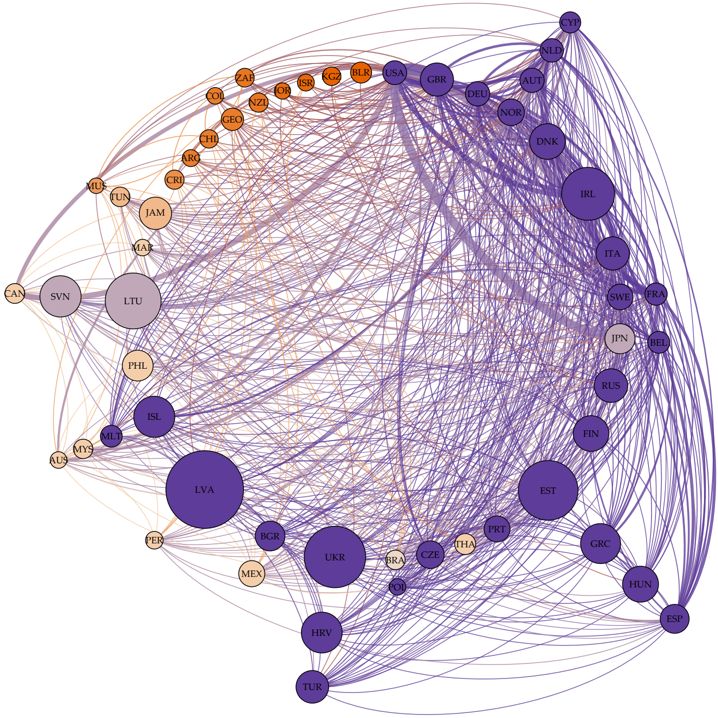

Figure B.20: Network representation of the investment layer with colours based on the kcore of nodes on the investment layer. Nodes are ordered on the circle by their degree with size representing the cumulative loss variable and edge width representing the flow of investments among them.

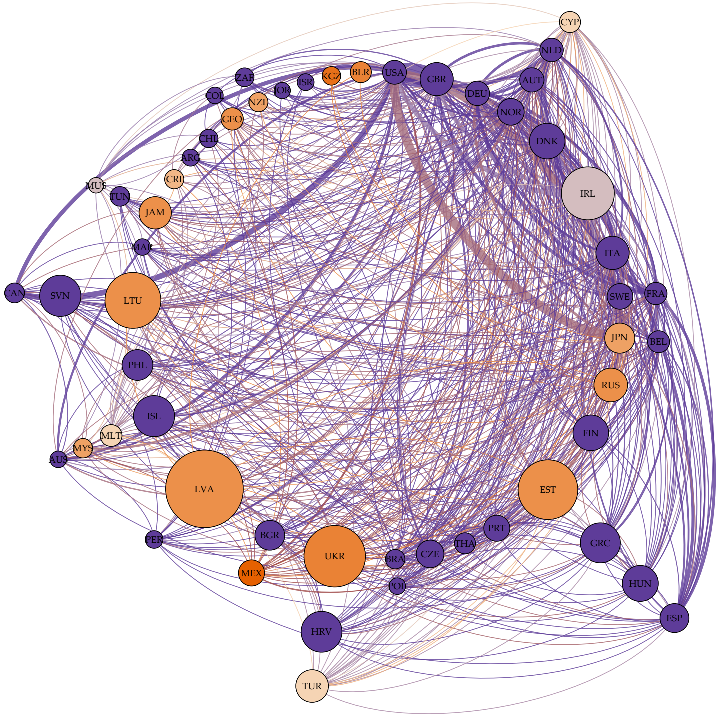

Figure B.21: Network representation of the investment layer with colours based on the kcore of nodes on the stock of migration layer. Nodes are ordered on the circle by their degree with size representing the cumulative loss variable and edge width representing the flow of investments among them.

Bibliography

Abel, Guy (2013).
"Estimating global migration flow tables using place of birth data"
In: *Demographic Research* 28, pp. 505–546.

Abel, Guy J. and Nikola Sander (Mar. 2014).
"Quantifying Global International Migration Flows". en
In: *Science* 343.6178, pp. 1520–1522. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1248676. URL: http://science.sciencemag.org/content/343/6178/1520 (visited on 03/29/2019).

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (Feb. 2015).
"Systemic Risk and Stability in Financial Networks". en
In: *American Economic Review* 105.2, pp. 564–608. ISSN: 0002-8282. DOI: 10.1257/aer.20130456. URL: http://pubs.aeaweb.org/doi/10.1257/aer.20130456 (visited on 04/02/2019).

Adamic, Lada A and Eytan Adar (July 2003).
"Friends and neighbors on the Web". en
In: *Social Networks* 25.3, pp. 211–230. ISSN: 0378-8733. DOI: 10.1016/S0378-8733(03)00009-1. URL: http://www.sciencedirect.com/science/article/pii/S0378873303000091 (visited on 11/19/2019).

Aggarwal, Reena, Asli Demirgüç-Kunt, and Maria Soledad Martínez Pería (Nov. 2011).
"Do remittances promote financial development?" en
In: *Journal of Development Economics* 96.2, pp. 255–264. ISSN: 0304-3878. DOI: 10.1016/j.jdeveco.2010.10.005. URL: http://www.sciencedirect.com/science/article/pii/S0304387810001161 (visited on 02/02/2020).

Aicher, Christopher, Abigail Z. Jacobs, and Aaron Clauset (June 2015).

"Learning Latent Block Structure in Weighted Networks". en
In: *Journal of Complex Networks* 3.2, pp. 221–248. ISSN: 2051-1310, 2051-1329. DOI: 10.1093/comnet/cnu026. URL: http://arxiv.org/abs/1404.0431 (visited on 11/13/2019).

Aleta, Alberto and Yamir Moreno (2019).
"Multilayer Networks in a Nutshell"
In: *Annual Review of Condensed Matter Physics* 10.1, pp. 45–62. DOI: 10.1146/annurev-conmatphys-031218-013259. URL: https://doi.org/10.1146/annurev-conmatphys-031218-013259 (visited on 03/22/2019).

Almog, Assaf, Tiziano Squartini, and Diego Garlaschelli (Jan. 2015).
"A GDP-driven model for the binary and weighted structure of the International Trade Network". en
In: *New Journal of Physics* 17.1, p. 013009. ISSN: 1367-2630. DOI: 10.1088/1367-2630/17/1/013009. URL: http://stacks.iop.org/1367-2630/17/i=1/a=013009?key=crossref.1e89a8969a471b399a49503bac748447 (visited on 11/12/2019).

Alptekin, Aynur and Paul Levine (Dec. 2012).
"Military expenditure and economic growth: A meta-analysis". en
In: *European Journal of Political Economy* 28.4, pp. 636–650. ISSN: 0176-2680. DOI: 10.1016/j.ejpoleco.2012.07.002. URL: http://www.sciencedirect.com/science/article/pii/S0176268012000432 (visited on 02/01/2020).

Anand, Kartik et al. (Apr. 2018).
"The missing links: A global study on uncovering financial network structures from partial data". en
In: *Journal of Financial Stability* 35, pp. 107–119. ISSN: 15723089. DOI: 10.1016/j.jfs.2017.05.012. URL: https://linkinghub.elsevier.com/retrieve/pii/S1572308917303649 (visited on 02/18/2019).

Armenter, Roc and Miklós Koren (July 2014).
"A Balls-and-Bins Model of Trade". en
In: *American Economic Review* 104.7, pp. 2127–2151. ISSN: 0002-8282. DOI: 10.1257/aer.104.7.2127. URL: http://pubs.aeaweb.org/doi/10.1257/aer.104.7.2127 (visited on 04/10/2018).

Arrigo, Francesca and Francesco Tudisco (Sept. 2018).
"Multi-Dimensional, Multilayer, Nonlinear and Dynamic HITS". en
In: *arXiv:1809.08004 [physics]*. URL: http://arxiv.org/abs/1809.08004 (visited on 02/18/2019).

Bardoscia, Marco et al. (June 2015).

"DebtRank: A Microscopic Foundation for Shock Propagation". en
In: *PLOS ONE* 10.6. Ed. by Matjaz Perc, e0130406. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0130406. URL: https://dx.plos.org/10.1371/journal.pone.0130406 (visited on 02/18/2019).

Battiston, Federico, Vincenzo Nicosia, and Vito Latora (Mar. 2014).
"Structural measures for multiplex networks". en
In: *Physical Review E* 89.3. ISSN: 1539-3755, 1550-2376. DOI: 10.1103/PhysRevE.89.032804. URL: https://link.aps.org/doi/10.1103/PhysRevE.89.032804 (visited on 04/10/2018).

Battiston, Stefano et al. (Dec. 2012).
"DebtRank: Too Central to Fail? Financial Networks, the FED and Systemic Risk". en
In: *Scientific Reports* 2.1. ISSN: 2045-2322. DOI: 10.1038/srep00541. URL: http://www.nature.com/articles/srep00541 (visited on 02/18/2019).

Bayer, Reşat (2006).
Diplomatic Exchange Data Set. v2006.1. Correlates of War Project. URL: http://correlatesofwar.org/.

Bee, Marco, Massimo Riccaboni, and Luca Trapin (Aug. 2017).
"An extreme value analysis of the last century crises across industries in the U.S. economy". en
In: *Journal of Economic Dynamics and Control*. International Conference "Large-scale Crises: 1929 vs. 2008" 81, pp. 65–78. ISSN: 0165-1889. DOI: 10.1016/j.jedc.2017.01.012. URL: http://www.sciencedirect.com/science/article/pii/S0165188917300209 (visited on 11/27/2019).

Belloni, Alexandre and Victor Chernozhukov (May 2013).
"Least squares after model selection in high-dimensional sparse models". EN
In: *Bernoulli* 19.2, pp. 521–547. ISSN: 1350-7265. DOI: 10.3150/11-BEJ410. URL: https://projecteuclid.org/euclid.bj/1363192037 (visited on 04/10/2019).

Boccaletti, S. et al. (Nov. 2014).
"The structure and dynamics of multilayer networks". en
In: *Physics Reports* 544.1, pp. 1–122. ISSN: 03701573. DOI: 10.1016/j.physrep.2014.07.001. URL: http://arxiv.org/abs/1407.0742 (visited on 05/27/2019).

Bonaccorsi, Giovanni et al. (Dec. 2019).

"Country centrality in the international multiplex network". en
In: *Applied Network Science* 4.1, p. 126. ISSN: 2364-8228. DOI: 10.1007/s41109-019-0207-3. URL: https://appliednetsci.springeropen.com/articles/10.1007/s41109-019-0207-3 (visited on 01/17/2020).

Brin, Sergey and Lawrence Page (Apr. 1998).
"The anatomy of a large-scale hypertextual Web search engine"
In: *Computer Networks and ISDN Systems*. Proceedings of the Seventh International World Wide Web Conference 30.1, pp. 107–117. ISSN: 0169-7552. DOI: 10.1016/S0169-7552(98)00110-X. URL: http://www.sciencedirect.com/science/article/pii/S016975529800110X (visited on 04/13/2018).

Bródka Piotr et al. (Apr. 2018).
"Quantifying layer similarity in multiplex networks: a systematic study"
In: *Royal Society Open Science* 5.8, p. 171747. DOI: 10.1098/rsos.171747. URL: https://royalsocietypublishing.org/doi/full/10.1098/rsos.171747 (visited on 02/25/2019).

Caballero, Julian (Jan. 2015).
"Banking crises and financial integration: Insights from networks science". en
In: *Journal of International Financial Markets, Institutions and Money* 34, pp. 127–146. ISSN: 10424431. DOI: 10.1016/j.intfin.2014.11.005. URL: https://linkinghub.elsevier.com/retrieve/pii/S1042443114001309 (visited on 04/02/2019).

Cai, Hongyun, Vincent W. Zheng, and Kevin Chen-Chuan Chang (Sept. 2017).
"A Comprehensive Survey of Graph Embedding: Problems, Techniques and Applications". en
In: *arXiv:1709.07604 [cs]*. URL: http://arxiv.org/abs/1709.07604 (visited on 10/15/2019).

Campi, Mercedes et al. (2016).
"Do Intellectual Property Rights Influence Cross-Border Mergers and Acquisitions?" en
In: *SSRN Electronic Journal*. ISSN: 1556-5068. DOI: 10.2139/ssrn.2805067. URL: https://www.ssrn.com/abstract=2805067 (visited on 04/10/2018).

Cao, Shaosheng, Wei Lu, and Qiongkai Xu (2015).
"GraRep: Learning Graph Representations with Global Structural Information"
In: *Proceedings of the 24th ACM International on Conference on Infor-*

mation and Knowledge Management. CIKM '15. New York, NY, USA: ACM, pp. 891–900. ISBN: 978-1-4503-3794-6. DOI: 10.1145/2806416.2806512. URL: http://doi.acm.org/10.1145/2806416.2806512 (visited on 11/20/2019).

Cecchetti, Stephen, Marion Kohler, and Christian Upper (Sept. 2009). *Financial Crises and Economic Activity*. en. Tech. rep. w15379. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w15379. URL: http://www.nber.org/papers/w15379.pdf (visited on 02/28/2019).

Chaney, Thomas (Nov. 2014). "The Network Structure of International Trade". en In: *American Economic Review* 104.11, pp. 3600–3634. ISSN: 0002-8282. DOI: 10.1257/aer.104.11.3600. URL: http://pubs.aeaweb.org/doi/10.1257/aer.104.11.3600 (visited on 04/10/2018).

Chinazzi, Matteo and Giorgio Fagiolo (2013). "Systemic Risk, Contagion, and Financial Networks: A Survey". en In: *SSRN Electronic Journal*. ISSN: 1556-5068. DOI: 10.2139/ssrn.2243504. URL: http://www.ssrn.com/abstract=2243504 (visited on 04/02/2019).

Chinazzi, Matteo, Giorgio Fagiolo, et al. (Aug. 2013). "Post-mortem examination of the international financial network". en In: *Journal of Economic Dynamics and Control* 37.8, pp. 1692–1713. DOI: 10.1016/j.jedc.2013.01.010. URL: https://linkinghub.elsevier.com/retrieve/pii/S0165188913000183 (visited on 04/02/2019).

Chipman, Hugh (Mar. 1996). "Bayesian variable selection with related predictors". en In: *Canadian Journal of Statistics* 24.1, pp. 17–36. ISSN: 03195724, 1708945X. DOI: 10.2307/3315687. URL: http://doi.wiley.com/10.2307/3315687 (visited on 01/14/2020).

Clauset, Aaron, Cristopher Moore, and M. E. J. Newman (May 2008). "Hierarchical structure and the prediction of missing links in networks" In: *Nature* 453.7191, pp. 98–101. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature06830. URL: http://arxiv.org/abs/0811.0484 (visited on 11/21/2019).

Colizza, Vittoria et al. (2006). "The role of the airline transportation network in the prediction and predictability of global epidemics"

In: *Proceedings of the National Academy of Sciences of the United States of America* 103.7, pp. 2015–2020.
Coordinated Portfolio Investment Survey(IMF website) (2019).
Coordinated Portfolio Investment Survey. URL: http://data.imf.org/regular.aspx?key=60587804 (visited on 10/31/2019).
Correlates of war datasets (website) (2019).
Correlates of war datasets. URL: http://www.correlatesofwar.org/data-sets (visited on 10/31/2019).
Costanza, Robert et al. (Jan. 2009).
"Beyond GDP: The Need for New Measures of Progress". en
In: p. 46.
Dasgupta, Partha and Karl-Göran Mäler (Feb. 2000).
"Net national product, wealth, and social well-being". en
In: *Environment and Development Economics* 5.1, pp. 69–93. ISSN: 1355770X. DOI: 10.1017/S1355770X00000061. URL: http://www.journals.cambridge.org/abstract_S1355770X00000061 (visited on 03/16/2019).
De Bacco, Caterina et al. (Apr. 2017).
"Community detection, link prediction, and layer interdependence in multilayer networks". en
In: *Physical Review E* 95.4, p. 042317. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.95.042317. URL: http://arxiv.org/abs/1701.01369 (visited on 10/15/2019).
De Domenico, Manlio, Vincenzo Nicosia, et al. (Dec. 2015).
"Structural reducibility of multilayer networks". en
In: *Nature Communications* 6.1. ISSN: 2041-1723. DOI: 10.1038/ncomms7864. URL: http://www.nature.com/articles/ncomms7864 (visited on 04/10/2018).
De Domenico, Manlio, Albert Solé-Ribalta, et al. (Dec. 2015).
"Ranking in interconnected multilayer networks reveals versatile nodes". en
In: *Nature Communications* 6.1. ISSN: 2041-1723. DOI: 10.1038/ncomms7868. URL: http://www.nature.com/articles/ncomms7868 (visited on 02/18/2019).
Devereux, John and Gerald P. Dwyer (Dec. 2016).
"What determines output losses after banking crises?" en
In: *Journal of International Money and Finance* 69, pp. 69–94. ISSN: 02615606. DOI: 10.1016/j.jimonfin.2016.10.004. URL: https://linkinghub.elsevier.com/retrieve/pii/S0261560616301206 (visited on 04/02/2019).

Dunne, J. Paul and Nan Tian (Jan. 2015).
"Military Expenditure, Economic Growth and Heterogeneity"
In: *Defence and Peace Economics* 26.1, pp. 15–31. ISSN: 1024-2694. DOI: 10.1080/10242694.2013.848575. URL: https://doi.org/10.1080/10242694.2013.848575 (visited on 02/01/2020).

Easley, David and Jon Kleinberg (July 2010).
Networks, Crowds, and Markets: Reasoning about a Highly Connected World. English. 1 edition. New York: Cambridge University Press. ISBN: 978-0-521-19533-1.

Fagiolo, Giorgio and Marina Mastrorillo (May 2014).
"Does Human Migration Affect International Trade? A Complex-Network Perspective". en
In: *PLoS ONE* 9.5. Ed. by Luís A. Nunes Amaral, e97331. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0097331. URL: http://dx.plos.org/10.1371/journal.pone.0097331 (visited on 04/10/2018).

Fagiolo, Giorgio and Gianluca Santoni (Sept. 2015).
"Human-mobility networks, country income, and labor productivity". en
In: *Network Science* 3.03, pp. 377–407. ISSN: 2050-1242, 2050-1250. DOI: 10.1017/nws.2015.25. URL: http://www.journals.cambridge.org/abstract_S2050124215000259 (visited on 04/10/2018).

— (Feb. 2016).
"Revisiting the role of migrant social networks as determinants of international migration flows". en
In: *Applied Economics Letters* 23.3, pp. 188–193. ISSN: 1350-4851, 1466-4291. DOI: 10.1080/13504851.2015.1064072. URL: http://www.tandfonline.com/doi/full/10.1080/13504851.2015.1064072 (visited on 04/10/2018).

FDI financial flows - By partner country (OECD online API) (2019).
FDI financial flows. URL: http://stats.oecd.org/Index.aspx?QueryId=64194 (visited on 10/31/2019).

Feldkircher, Martin (May 2014).
"The determinants of vulnerability to the global financial crisis 2008 to 2009: Credit growth and other sources of risk". en
In: *Journal of International Money and Finance* 43, pp. 19–49. ISSN: 02615606. DOI: 10.1016/j.jimonfin.2013.12.003. URL: https://linkinghub.elsevier.com/retrieve/pii/S0261560613001770 (visited on 02/17/2019).

Frankel, Jeffrey and George Saravelos (July 2012).

"Can leading indicators assess country vulnerability? Evidence from the 2008–09 global financial crisis". en
In: *Journal of International Economics* 87.2, pp. 216–231. ISSN: 00221996. DOI: 10.1016/j.jinteco.2011.12.009. URL: https://linkinghub.elsevier.com/retrieve/pii/S0022199611001735 (visited on 04/09/2019).
Fromentin, Vincent (Nov. 2017).
"The long-run and short-run impacts of remittances on financial development in developing countries". en
In: *The Quarterly Review of Economics and Finance* 66, pp. 192–201. ISSN: 1062-9769. DOI: 10.1016/j.qref.2017.02.006. URL: http://www.sciencedirect.com/science/article/pii/S1062976917300807 (visited on 02/01/2020).
Gaulier, Guillaume and Soledad Zignago (2010).
"BACI: International Trade Database at the Product-Level (the 1994-2007 Version)". en
In: *SSRN Electronic Journal*. ISSN: 1556-5068. DOI: 10.2139/ssrn.1994500. URL: http://www.ssrn.com/abstract=1994500 (visited on 10/30/2019).
Gemmetto, Valerio, Alessio Cardillo, and Diego Garlaschelli (June 2017).
"Irreducible network backbones: unbiased graph filtering via maximum entropy". en
In: *arXiv:1706.00230 [physics]*. URL: http://arxiv.org/abs/1706.00230 (visited on 07/06/2019).
Gibler, Douglas M and CQ Press (2009).
International military alliances, 1648-2008. English. Washington, D.C.: CQ Press. ISBN: 978-1-60426-684-9. URL: http://site.ebrary.com/id/10314601 (visited on 10/30/2019).
Glasserman, Paul and H. Peyton Young (Sept. 2016).
"Contagion in Financial Networks". en
In: *Journal of Economic Literature* 54.3, pp. 779–831. ISSN: 0022-0515. DOI: 10.1257/jel.20151228. URL: http://pubs.aeaweb.org/doi/10.1257/jel.20151228 (visited on 02/28/2019).
Goyal, Palash and Emilio Ferrara (July 2018).
"Graph Embedding Techniques, Applications, and Performance: A Survey". en
In: *Knowledge-Based Systems* 151, pp. 78–94. ISSN: 09507051. DOI: 10.1016/j.knosys.2018.03.022. URL: http://arxiv.org/abs/1705.02801 (visited on 07/09/2019).
Grover, Aditya and Jure Leskovec (2016).

"Node2Vec: Scalable Feature Learning for Networks"
In: *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '16. New York, NY, USA: ACM, pp. 855–864. ISBN: 978-1-4503-4232-2. DOI: 10.1145/2939672.2939754. URL: http://doi.acm.org/10.1145/2939672.2939754 (visited on 11/19/2019).

Guimera, R. and M. Sales-Pardo (Dec. 2009).
"Missing and spurious interactions and the reconstruction of complex networks"
In: *Proceedings of the National Academy of Sciences* 106.52, pp. 22073–22078. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.0908366106. URL: http://arxiv.org/abs/1004.4791 (visited on 11/21/2019).

Hafner-Burton, Emilie M., Miles Kahler, and Alexander H. Montgomery (July 2009).
"Network Analysis for International Relations". en
In: *International Organization* 63.3, pp. 559–592. ISSN: 1531-5088, 0020-8183. DOI: 10.1017/S0020818309090195. URL: https://www.cambridge.org/core/journals/international-organization/article/network-analysis-for-international-relations/DE2910979C1B5C44C4CC13F336C5DE97 (visited on 03/25/2019).

Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg (Dec. 2001).
The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools. en. SSRN Scholarly Paper ID 296162. Rochester, NY: Social Science Research Network. URL: https://papers.ssrn.com/abstract=296162 (visited on 10/30/2019).

Halu, Arda et al. (Oct. 2013).
"Multiplex PageRank". en
In: *PLoS ONE* 8.10. Ed. by Yamir Moreno, e78293. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0078293. URL: http://dx.plos.org/10.1371/journal.pone.0078293 (visited on 04/10/2018).

Hamilton, William L., Rex Ying, and Jure Leskovec (Apr. 2018).
"Representation Learning on Graphs: Methods and Applications"
In: *arXiv:1709.05584 [cs]*. URL: http://arxiv.org/abs/1709.05584 (visited on 12/02/2019).

Holland, Paul W., Kathryn Blackmond Laskey, and Samuel Leinhardt (June 1983).
"Stochastic blockmodels: First steps". en
In: *Social Networks* 5.2, pp. 109–137. ISSN: 03788733. DOI: 10.1016/0378-8733(83)90021-7. URL: https://linkinghub.el

sevier.com/retrieve/pii/0378873383900217 (visited on 12/01/2019).
Hurlin, Christophe (Jan. 2006).
Network Effects Of The Productivity Of Infrastructure In Developing Countries. en. Policy Research Working Papers. The World Bank. DOI: 10.1596/1813-9450-3808. URL: http://elibrary.worldbank.org/doi/book/10.1596/1813-9450-3808 (visited on 11/04/2019).
Iacovacci, Jacopo and Ginestra Bianconi (June 2016).
"Extracting Information from Multiplex Networks". en
In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 26.6, p. 065306. ISSN: 1054-1500, 1089-7682. DOI: 10.1063/1.4953161. URL: http://arxiv.org/abs/1602.08751 (visited on 03/07/2019).
Iacovacci, Jacopo, Christoph Rahmede, et al. (Oct. 2016).
"Functional Multiplex PageRank". en
In: *EPL (Europhysics Letters)* 116.2, p. 28004. ISSN: 0295-5075, 1286-4854. DOI: 10.1209/0295-5075/116/28004. URL: http://stacks.iop.org/0295-5075/116/i=2/a=28004?key=crossref.4c0ed62a515e4c93404e6e7de9476d6d (visited on 04/10/2018).
International migrant stock (direct download) (2015).
International migrant stock (direct download). URL: http://www.un.org/en/development/desa/population/migration/data/estimates2/data/UN_MigrantStockByOriginAndDestination_2015.xlsx (visited on 10/31/2019).
Iyer, Rajkamal and José-Luis Peydró (Apr. 2011).
"Interbank Contagion at Work: Evidence from a Natural Experiment". en
In: *Review of Financial Studies* 24.4, pp. 1337–1377. ISSN: 0893-9454, 1465-7368. DOI: 10.1093/rfs/hhp105. URL: https://academic.oup.com/rfs/article-lookup/doi/10.1093/rfs/hhp105 (visited on 04/02/2019).
Kali, Raja, Andrew Horowitz, and Hongwei Song (2017).
Rethinking the Aid-Growth Relationship: A Network Approach. Tech. rep.
Kali, Raja and Javier Reyes (Oct. 2010).
"Financial contagion on the international trade network". en
In: *Economic Inquiry* 48.4, pp. 1072–1101. ISSN: 00952583. DOI: 10.1111/j.1465-7295.2009.00249.x. URL: http://doi.wiley.com/10.1111/j.1465-7295.2009.00249.x (visited on 04/10/2018).
Kaluza, Pablo et al. (2010).

"The complex network of global cargo ship movements"
In: *Journal of the Royal Society Interface* 7.48, pp. 1093–1103.
Kim, Myunghwan and Jure Leskovec (Apr. 2011).
"The Network Completion Problem: Inferring Missing Nodes and Edges in Networks". en
In: *Proceedings of the 2011 SIAM International Conference on Data Mining*. Society for Industrial and Applied Mathematics, pp. 47–58. ISBN: 978-1-61197-281-8. DOI: 10.1137/1.9781611972818.5. URL: https://epubs.siam.org/doi/10.1137/1.9781611972818.5 (visited on 05/16/2019).
Kirkegaard, J.F. (2013).
"New Avenues for Empirical Analysis of Cross-Border Investments: An Application for the ASEAN Members and Middle and Low Income Country Outward Investments." PhD Thesis. Baltimore: Johns Hopkins University.
Kitsak, Maksim et al. (Nov. 2010).
"Identification of influential spreaders in complex networks". en
In: *Nature Physics* 6.11, pp. 888–893. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/nphys1746. URL: http://arxiv.org/abs/1001.5285 (visited on 02/15/2019).
Kivela, M. et al. (Sept. 2014).
"Multilayer networks". en
In: *Journal of Complex Networks* 2.3, pp. 203–271. ISSN: 2051-1310, 2051-1329. DOI: 10.1093/comnet/cnu016. URL: https://academic.oup.com/comnet/article-lookup/doi/10.1093/comnet/cnu016 (visited on 04/10/2018).
Kleinberg, Jon M. (Dec. 1999).
"Hubs, authorities, and communities". en
In: *ACM Computing Surveys* 31.4es, 5–es. ISSN: 03600300. DOI: 10.1145/345966.345982. URL: http://portal.acm.org/citation.cfm?doid=345966.345982 (visited on 02/24/2019).
Lange, Glenn-Marie, Quentin Wodon, and Kevin Carey, eds. (Jan. 2018).
The Changing Wealth of Nations 2018: Building a Sustainable Future. en. The World Bank. ISBN: 978-1-4648-1046-6. DOI: 10.1596/978-1-4648-1046-6. URL: http://elibrary.worldbank.org/doi/book/10.1596/978-1-4648-1046-6 (visited on 03/16/2019).
Lee, Kyu-Min et al. (Mar. 2012).
"Correlated multiplexity and connectivity of multiplex random networks". en
In: *New Journal of Physics* 14.3, p. 033027. ISSN: 1367-2630. DOI: 10.

1088/1367-2630/14/3/033027. URL: http://stacks.iop.org/1367-2630/14/i=3/a=033027?key=crossref.7734afd83ce1f03a4b0fdcb967388e36 (visited on 04/10/2018).

Lin, Jian-Hong et al. (Oct. 2014).
"Identifying the node spreading influence with largest k -core values". en
In: *Physics Letters A* 378.45, pp. 3279–3284. ISSN: 03759601. DOI: 10.1016/j.physleta.2014.09.054. URL: https://linkinghub.elsevier.com/retrieve/pii/S0375960114009839 (visited on 11/04/2019).

Liu, Weiyi et al. (Sept. 2017).
"Principled Multilayer Network Embedding". en
In: *arXiv:1709.03551 [physics]*. URL: http://arxiv.org/abs/1709.03551 (visited on 07/16/2019).

Locational banking statistics (BIS website) (2019).
Locational banking statistics. URL: http://www.bis.org/statistics/bankstats.htm?m=6%7C31%7C69 (visited on 10/31/2019).

Lü, Linyuan and Tao Zhou (Mar. 2011).
"Link prediction in complex networks: A survey". en
In: *Physica A: Statistical Mechanics and its Applications* 390.6, pp. 1150–1170. ISSN: 03784371. DOI: 10.1016/j.physa.2010.11.027. URL: https://linkinghub.elsevier.com/retrieve/pii/S037843711000991X (visited on 12/01/2019).

Marchette, David J. and Carey E. Priebe (Jan. 2008).
"Predicting unobserved links in incompletely observed networks". en
In: *Computational Statistics & Data Analysis* 52.3, pp. 1373–1386. ISSN: 0167-9473. DOI: 10.1016/j.csda.2007.03.016. URL: http://www.sciencedirect.com/science/article/pii/S0167947307001193 (visited on 11/21/2019).

Masanjala, Winford H. and Chris Papageorgiou (2008).
"Rough and lonely road to prosperity: a reexamination of the sources of growth in Africa using Bayesian model averaging". en
In: *Journal of Applied Econometrics* 23.5, pp. 671–682. ISSN: 1099-1255. DOI: 10.1002/jae.1020. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.1020 (visited on 04/11/2019).

Mastrandrea, Rossana et al. (Dec. 2014).
"Reconstructing the world trade multiplex: The role of intensive and extensive biases". en
In: *Physical Review E* 90.6. ISSN: 1539-3755, 1550-2376. DOI: 10.1103/

PhysRevE.90.062804. URL: https://link.aps.org/doi/10.1103/PhysRevE.90.062804 (visited on 02/17/2019).

Metulini, Rodolfo et al. (Dec. 2018). "The network of migrants and international trade". en In: *Economia Politica* 35.3, pp. 763–787. ISSN: 1973-820X. DOI: 10.1007/s40888-018-0106-6. URL: https://doi.org/10.1007/s40888-018-0106-6 (visited on 04/08/2019).

Mikolov, Tomas, Ilya Sutskever, et al. (2013). "Distributed Representations of Words and Phrases and Their Compositionality" In: *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*. NIPS'13. USA: Curran Associates Inc., pp. 3111–3119. URL: http://dl.acm.org/citation.cfm?id=2999792.2999959 (visited on 11/20/2019).

Mikolov, Tomas, Wen-tau Yih, and Geoffrey Zweig (June 2013). "Linguistic Regularities in Continuous Space Word Representations" In: *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Atlanta, Georgia: Association for Computational Linguistics, pp. 746–751. URL: https://www.aclweb.org/anthology/N13-1090 (visited on 11/20/2019).

Minoiu, Camelia, Chanhyun Kang, et al. (Apr. 2015). "Does financial connectedness predict crises?" en In: *Quantitative Finance* 15.4, pp. 607–624. ISSN: 1469-7688, 1469-7696. DOI: 10.1080/14697688.2014.968358. URL: http://www.tandfonline.com/doi/abs/10.1080/14697688.2014.968358 (visited on 04/04/2019).

Minoiu, Camelia and Javier A. Reyes (June 2013). "A network analysis of global banking: 1978–2010". en In: *Journal of Financial Stability* 9.2, pp. 168–184. ISSN: 15723089. DOI: 10.1016/j.jfs.2013.03.001. URL: http://linkinghub.elsevier.com/retrieve/pii/S1572308913000193 (visited on 04/10/2018).

Moral-Benito, Enrique (Feb. 2015). "MODEL AVERAGING IN ECONOMICS: AN OVERVIEW: MODEL AVERAGING IN ECONOMICS". en In: *Journal of Economic Surveys* 29.1, pp. 46–75. ISSN: 09500804. DOI: 10.1111/joes.12044. URL: http://doi.wiley.com/10.1111/joes.12044 (visited on 04/03/2019).

Moreno, Yamir, Romualdo Pastor-Satorras, and Alessandro Vespignani (Apr. 2002).
"Epidemic outbreaks in complex heterogeneous networks". en
In: *The European Physical Journal B* 26.4, pp. 521–529. ISSN: 1434-6028. DOI: 10.1140/epjb/e20020122. URL: http://arxiv.org/abs/cond-mat/0107267 (visited on 11/04/2019).
Moser, Mathias and Paul Hofmarcher (Mar. 2014).
"MODEL PRIORS REVISITED: INTERACTION TERMS IN BMA GROWTH APPLICATIONS: INTERACTION TERMS AND BMA MODEL PRIORS". en
In: *Journal of Applied Econometrics* 29.2, pp. 344–347. ISSN: 08837252. DOI: 10.1002/jae.2365. URL: http://doi.wiley.com/10.1002/jae.2365 (visited on 01/14/2020).
NBER patent data (download) (2019).
NBER patent data. URL: https://sites.google.com/site/patentdataproject/Home/downloads (visited on 10/31/2019).
Newman, M. E. J. (Apr. 2001).
"Clustering and preferential attachment in growing networks". en
In: DOI: 10.1103/PhysRevE.64.025102. URL: https://arxiv.org/abs/cond-mat/0104209v1 (visited on 11/21/2019).
Newman, Mark (Mar. 2010).
Networks: An Introduction. en. OUP Oxford. ISBN: 978-0-19-920665-0.
OECD ODA disbursment (web API) (n.d.).
OECD ODA disbursment. URL: http://stats.oecd.org/Index.aspx?DataSetCode=Table2A.
Pan, Raj Kumar, Kimmo Kaski, and Santo Fortunato (Nov. 2012).
"World citation and collaboration networks: uncovering the role of geography in science". en
In: *Scientific Reports* 2, p. 902. DOI: 10.1038/srep00902. URL: https://www.nature.com/articles/srep00902 (visited on 03/29/2019).
Park, Juyong and M. E. J. Newman (Dec. 2004).
"Statistical mechanics of networks"
In: *Physical Review E* 70.6, p. 066117. DOI: 10.1103/PhysRevE.70.066117. URL: https://link.aps.org/doi/10.1103/PhysRevE.70.066117 (visited on 12/01/2019).
Pastor-Satorras, Romualdo, Claudio Castellano, et al. (Aug. 2015).
"Epidemic processes in complex networks". en
In: *Reviews of Modern Physics* 87.3, pp. 925–979. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.87.925. URL: https://lin

k.aps.org/doi/10.1103/RevModPhys.87.925 (visited on 11/04/2019).
Pastor-Satorras, Romualdo and Alessandro Vespignani (Feb. 2002). "Immunization of complex networks" In: *Physical Review E* 65.3, p. 036104. DOI: 10.1103/PhysRevE.65.036104. URL: https://link.aps.org/doi/10.1103/PhysRevE.65.036104 (visited on 11/07/2019).
Peixoto, Tiago P. (Jan. 2014). "Efficient Monte Carlo and greedy heuristic for the inference of stochastic block models". en In: *Physical Review E* 89.1, p. 012804. ISSN: 1539-3755, 1550-2376. DOI: 10.1103/PhysRevE.89.012804. URL: http://arxiv.org/abs/1310.4378 (visited on 11/17/2019).
— (Nov. 2018a). "Bayesian stochastic blockmodeling". en In: *arXiv:1705.10225 [cond-mat, physics:physics, stat]*. URL: http://arxiv.org/abs/1705.10225 (visited on 11/17/2019).
— (Jan. 2018b). "Nonparametric weighted stochastic block models". en In: *Physical Review E* 97.1, p. 012306. ISSN: 2470-0045, 2470-0053. DOI: 10.1103/PhysRevE.97.012306. URL: http://arxiv.org/abs/1708.01432 (visited on 11/17/2019).
— (Oct. 2018c). "Reconstructing networks with unknown and heterogeneous errors". en In: *Physical Review X* 8.4, p. 041011. ISSN: 2160-3308. DOI: 10.1103/PhysRevX.8.041011. URL: http://arxiv.org/abs/1806.07956 (visited on 11/14/2019).
Perozzi, Bryan, Rami Al-Rfou, and Steven Skiena (2014). "DeepWalk: Online Learning of Social Representations" In: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '14*, pp. 701–710. DOI: 10.1145/2623330.2623732. URL: http://arxiv.org/abs/1403.6652 (visited on 11/19/2019).
Pevehouse, Jon, Timothy Nordstrom, and Kevin Warnke (2003). *Intergovernmental Organizations, 1815-2000: A New Correlates of War Data Set*. 2.1. Correlates of War Project. URL: http://correlatesofwar.org.
Radicchi, Filippo, José J. Ramasco, and Santo Fortunato (Apr. 2011).

"Information filtering in complex weighted networks". eng
In: *Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics* 83.4 Pt 2, p. 046101. ISSN: 1550-2376. DOI: 10.1103/PhysRevE.83.046101.
Rahmede, Christoph et al. (Oct. 2018).
"Centralities of nodes and influences of layers in large multiplex networks". en
In: *Journal of Complex Networks* 6.5. Ed. by Jesus Gomez-Gardenes, pp. 733–752. ISSN: 2051-1310, 2051-1329. DOI: 10.1093/comnet/cnx050. URL: https://academic.oup.com/comnet/article/6/5/733/4561631 (visited on 03/07/2019).
Reinhart, Carmen M. and Kenneth S. Rogoff (Sept. 2009).
This Time Is Different: Eight Centuries of Financial Folly. en. Princeton University Press. ISBN: 978-1-4008-3172-2.
— (May 2014).
"Recovery from Financial Crises: Evidence from 100 Episodes". en
In: *American Economic Review* 104.5, pp. 50–55. ISSN: 0002-8282. DOI: 10.1257/aer.104.5.50. URL: https://www.aeaweb.org/articles?id=10.1257/aer.104.5.50 (visited on 04/09/2019).
Riccaboni, Massimo, Alessandro Rossi, and Stefano Schiavo (Apr. 2013).
"Global networks of trade and bits". en
In: *Journal of Economic Interaction and Coordination* 8.1, pp. 33–56. ISSN: 1860-711X, 1860-7128. DOI: 10.1007/s11403-012-0101-x. URL: http://link.springer.com/10.1007/s11403-012-0101-x (visited on 04/10/2018).
Romer, Christina D. and David H. Romer (Oct. 2017).
"New Evidence on the Aftermath of Financial Crises in Advanced Countries". en
In: *American Economic Review* 107.10, pp. 3072–3118. ISSN: 0002-8282. DOI: 10.1257/aer.20150320. URL: http://pubs.aeaweb.org/doi/10.1257/aer.20150320 (visited on 04/09/2019).
Rossello, Giulia (July 2015).
Internet Infrastructure: a Multi-layer Network Analysis. it. URL: http://etd.adm.unipi.it/ (visited on 03/29/2019).
Sen, Amartya (1985).
Commodities and Capabilities. Amsterdam: North-Holland.
Serrano, M. Ángeles and Marián Boguñá (June 2005).
"Weighted Configuration Model"
In: *AIP Conference Proceedings* 776.1, pp. 101–107. ISSN: 0094-243X. DOI:

10.1063/1.1985381. URL: https://aip.scitation.org/doi/abs/10.1063/1.1985381 (visited on 12/01/2019).

Serrano, M. Ángeles, Marián Boguñá, and Alessandro Vespignani (Apr. 2009).
"Extracting the multiscale backbone of complex weighted networks"
In: *Proceedings of the National Academy of Sciences of the United States of America* 106.16, pp. 6483–6488. ISSN: 0027-8424. DOI: 10.1073/pnas.0808904106. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2672499/ (visited on 07/06/2019).

Sgrignoli, Paolo et al. (Jan. 2015).
"The relation between global migration and trade networks". en
In: *Physica A: Statistical Mechanics and its Applications* 417, pp. 245–260. ISSN: 03784371. DOI: 10.1016/j.physa.2014.09.037. URL: http://linkinghub.elsevier.com/retrieve/pii/S037843711400795X (visited on 04/10/2018).

SIPRI Arms Transfers Database (website) (2019).
SIPRI Arms Transfers Database. URL: https://www.sipri.org/databases/armstransfers (visited on 10/31/2019).

Smith, David A. and Douglas R. White (June 1992).
"Structure and Dynamics of the Global Economy: Network Analysis of International Trade 1965-1980". en
In: *Social Forces* 70.4, p. 857. ISSN: 00377732. DOI: 10.2307/2580193. URL: http://www.jstor.org/stable/2580193?origin=crossref (visited on 04/10/2018).

Sousa, José de and Julie Lochard (2011).
"Does the Single Currency Affect Foreign Direct Investment?*". en
In: *The Scandinavian Journal of Economics* 113.3, pp. 553–578. ISSN: 1467-9442. DOI: 10.1111/j.1467-9442.2011.01656.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9442.2011.01656.x (visited on 03/29/2019).

Squartini, Tiziano, Guido Caldarelli, et al. (Oct. 2018).
"Reconstruction methods for networks: the case of economic and financial systems". en
In: *Physics Reports* 757, pp. 1–47. ISSN: 03701573. DOI: 10.1016/j.physrep.2018.06.008. URL: http://arxiv.org/abs/1806.06941 (visited on 02/17/2019).

Squartini, Tiziano and Diego Garlaschelli (Aug. 2011).
"Analytical maximum-likelihood method to detect patterns in real networks"
In: *New Journal of Physics* 13.8, p. 083001. ISSN: 1367-2630. DOI: 10.

1088/1367-2630/13/8/083001. URL: http://arxiv.org/abs/1103.0701 (visited on 12/01/2019).

Squartini, Tiziano, Rossana Mastrandrea, and Diego Garlaschelli (Feb. 2015). "Unbiased sampling of network ensembles". en In: *New Journal of Physics* 17.2, p. 023052. ISSN: 1367-2630. DOI: 10.1088/1367-2630/17/2/023052. URL: http://stacks.iop.org/1367-2630/17/i=2/a=023052?key=crossref.707548a80edb2e237cd620c454060e24 (visited on 02/17/2019).

Stevenson, Betsey and Justin Wolfers (May 2013). "Subjective Well-Being and Income: Is There Any Evidence of Satiation?" en In: *American Economic Review* 103.3, pp. 598–604. ISSN: 0002-8282. DOI: 10.1257/aer.103.3.598. URL: http://pubs.aeaweb.org/doi/10.1257/aer.103.3.598 (visited on 03/16/2019).

Tang, Jian et al. (2015). "LINE: Large-scale Information Network Embedding" In: *Proceedings of the 24th International Conference on World Wide Web*. WWW '15. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, pp. 1067–1077. ISBN: 978-1-4503-3469-3. DOI: 10.1145/2736277.2741093. URL: https://doi.org/10.1145/2736277.2741093 (visited on 11/20/2019).

TeleGeography Submarine Cable Map (2019). *TeleGeography Submarine Cable Map*. URL: https://www.submarinecablemap.com/ (visited on 10/31/2019).

Tibshirani, Robert (1996). "Regression Shrinkage and Selection via the Lasso" In: *Journal of the Royal Statistical Society. Series B (Methodological)* 58.1, pp. 267–288. ISSN: 0035-9246. URL: https://www.jstor.org/stable/2346178 (visited on 04/10/2019).

Tourism Statistics (UN Tourism office website) (2019). *Tourism Statistics*. URL: http://www.e-unwto.org/loi/unwtotfb (visited on 10/31/2019).

United Nations, Population Division (2013). *Trends in International Migrant Stock: Migrants by Destination and Origin, The 2013 Revision*.

Yakovlev, Pavel (Aug. 2007). "Arms Trade, Military Spending, and Economic Growth" In: *Defence and Peace Economics* 18.4, pp. 317–338. ISSN: 1024-2694. DOI:

10.1080/10242690601099679. URL: https://doi.org/10.1080/10242690601099679 (visited on 02/01/2020).

Zellner, Arnold (June 1986).
"Bayesian Estimation and Prediction Using Asymmetric Loss Functions"
In: *Journal of the American Statistical Association* 81.394, pp. 446–451. ISSN: 0162-1459. DOI: 10.1080/01621459.1986.10478289. URL: https://www.tandfonline.com/doi/abs/10.1080/01621459.1986.10478289 (visited on 04/10/2019).

Zhang, Daokun et al. (2018).
"Network Representation Learning: A Survey". en
In: *IEEE Transactions on Big Data*, pp. 1–1. ISSN: 2332-7790. DOI: 10.1109/TBDATA.2018.2850013. URL: https://ieeexplore.ieee.org/document/8395024/ (visited on 07/16/2019).

Zhang, Hongming et al. (July 2018).
"Scalable Multiplex Network Embedding". en
In: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*. Stockholm, Sweden: International Joint Conferences on Artificial Intelligence Organization, pp. 3082–3088. ISBN: 978-0-9992411-2-7. DOI: 10.24963/ijcai.2018/428. URL: https://www.ijcai.org/proceedings/2018/428 (visited on 07/09/2019).

Ziesemer, Thomas H. W. (Mar. 2012).
"Worker remittances, migration, accumulation and growth in poor developing countries: Survey and analysis of direct and indirect effects". en
In: *Economic Modelling* 29.2, pp. 103–118. ISSN: 0264-9993. DOI: 10.1016/j.econmod.2011.08.013. URL: http://www.sciencedirect.com/science/article/pii/S0264999311002057 (visited on 02/01/2020).

Zitnik, Marinka and Jure Leskovec (July 2017).
"Predicting multicellular function through multi-layer tissue networks". en
In: *Bioinformatics* 33.14, pp. i190–i198. ISSN: 1367-4803, 1460-2059. DOI: 10.1093/bioinformatics/btx252. URL: https://academic.oup.com/bioinformatics/article/33/14/i190/3953967 (visited on 07/09/2019).