

**On the Impact of Network Topology on Systemic Risk:
Network Reconstruction, Stress Testing and Control**

Daniel Grigat

*Department of Computer Science, Faculty of Engineering Sciences
University College London*

First Supervisor: Dr. Fabio Caccioli
Second Supervisor: Dr. Giacomo Livan

*Dissertation submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy in Computer Science*



August 2017

"I, Daniel Grigat, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis."

Daniel Grigat, *August 2017*

Für Mama, who always believed in me. Sine qua non.

Acknowledgements

I am indebted to my first supervisor Dr. Fabio Caccioli. Fabio has been instrumental in every aspect of my PhD studies. His personal and professional support, advice and expertise are invaluable to me. Despite all the mathematical skills that I have acquired in my studies, I cannot measure the value of his support. Fabio has instilled in me a deeply inquiring mindset, that will hopefully continue to improve the quality of my life to come.

I thank Dr. Giacomo Livan. Giacomo is my second supervisor who supported me across my projects and has given me helpful advice throughout my studies.

I would like to thank the Financial Computing and Analytics Group. I have been a member of the group throughout my studies and benefited from the advice of its members and the frequent seminar series. I would like to particularly thank Prof. Tomaso Aste, who was my examiner in my first year viva and transfer viva. Tomaso has given me very insightful and constructive feedback that has greatly improved the quality of my thesis.

I am particularly grateful to Prof. Philip Treleaven who is the director of the Centre for Doctoral Training (CDT) in Financial Computing and Analytics. Philip admitted me onto the PhD programme and to a 4-year Engineering and Physical Sciences Research Council PhD scholarship, which supported me financially throughout my studies. I also thank the PhD candidates of the CDT and Yonita Carter, who have given me guidance and advice during my time in the centre.

Lastly, I am indebted to the administrative staff of the computer science department, in particular Sarah Turnbull, Abu Mohamed, JJ Giwa-Majekodunmi and Patricia Fenoy, who have made my three years in the department much more effective and enjoyable. I am especially thankful to the Technical Support Group who helped me greatly in the use of the high performance cluster and other computing tools. Without the high performance cluster, I could not have undertaken many of the computations that I performed in my research.

Abstract

In this thesis I undertake four projects that investigate the influence of network topology on financial contagion. In the first project I study the ability of three network models to reconstruct the topology and contagion dynamics of a real interbank network. This allows me to identify the main topological properties of networks that affect financial contagion. In the third and fourth projects I develop a framework to reverse stress test interbank networks. Based on this framework I propose a strategy to reduce systemic losses and I test the strategy in interbank networks with different topologies. I develop the methodological background to the reverse stress testing framework in the second project. In this project I study network control theory, to investigate the influence of network topology on the cost of managing the dynamics of networks.

Project 1: I focus on the problem of the required amount of information to reconstruct a real interbank network. I obtain data of the Austrian interbank network and compare the ability of three null models to reconstruct its topology. I test Erdős-Rényi, the directed configuration and the fitness null models, each of which requires a different amount of prior information about the real network. Since the same counterparty risk algorithms are employed to stress test each network, an understanding of the topological determinants of financial contagion is obtained by analysing the differences in the topologies of the real and reconstructed networks. I find that the null models capture the DebtRank algorithm well, which does not strongly depend on the local topology of the network, but rather on global topological properties. In contrast, the null models cannot reproduce the exact values of contagion observed due to the Furfine algorithm, that strongly depends on specific constellations of links in the local topology. The difference in the observed financial contagion due to the Furfine algorithm between the reconstructed and the real networks, is due to a higher concentration of risk in the real network. I explain the higher concentration of risk in the real network via the concept of vulnerable clusters of banks. Lastly, I find that the fitness model provided with information only on the interbank assets and liabilities of each bank outperforms the other null

models. This result suggests that the heterogeneity of the distribution of the inter-bank relationships is sufficient in determining financial contagion at the aggregate level.

Project 2: I perform numerical simulations to study the cost of controlling the dynamics of networks. In particular, I investigate Erdős-Rényi and scale-free networks, to identify the topological determinants of control energy. Control energy represents the cost of driving each node in a network to a specific target state. In interbank networks control energy translates into capital injections (losses) required to reduce (increase) systemic risk. I find that, when nodes have a common target state that they push one another to their common goal and thereby reduce the control energy requirement. This reduction of the control energy is also evident when the network connectivity increases, in this case more densely connected nodes help one another in reaching their common target state. This effect is more pronounced in scale-free networks, due to the presence of hub nodes that reach throughout the network. To provide an intuition of the effect of network connectivity on control energy, I provide an analytical approach based on mean-field theory to compute the control energy in Erdős-Rényi networks.

Project 3: I introduce a reverse stress testing framework that can be used to identify the potential triggers of systemic events and that removes the arbitrariness in the selection of shock scenarios in stress testing. I reverse engineer financial contagion to find the distribution of smallest exogenous shocks that, should it occur, would lead to a given final systemic loss. I consider the case of distress propagation in an interbank market based on the 44 largest European banks, which I construct using data collected from Bloomberg. By looking at the distribution across banks of the size of smallest exogenous shocks, I rank banks in terms of their systemic importance, and show the effectiveness of a policy with capital requirements based on this ranking. I also study the properties of the exogenous shocks as a function of the largest eigenvalue λ_{\max} of the matrix of interbank leverages, which determines the endogenous amplification of shocks. I find that the size of exogenous shocks reduces and that the distribution across banks becomes more localized as λ_{\max} increases.

Project 4: I study the impact of network topology on the reverse stress testing framework. I apply the framework to networks of various densities representing the largest European banks and to the Austrian interbank network. I then perform an analysis similar to that in the first project, in which I reconstruct the Austrian network using three network models to test their ability to estimate the exogenous shock applied via the reverse stress testing framework. Through these investigations

I find that the behaviour of the exogenous shock is robust across different topologies. Moreover, I find that, when the system is unstable, the exogenous shock decreases as a function of the network density. The null model reconstruction exercise surprisingly reveals that the fitness model performs worse than the other two network models. I link the performance of the fitness model to its overestimation of isolated nodes, which significantly increases the required exogenous shock.

Main contribution of each project:

Project 1: Null models capture general trend in financial contagion, but cannot accurately recreate local topology.

Project 2: Cost of controlling network dynamics decreases as a function of the network connectivity.

Project 3: Introduction of framework to reverse stress test interbank networks, which highlights sources of vulnerability in the network.

Project 4: Exogenous shock computed via reverse stress testing framework depends on the network connectivity.

Table of Contents

Acknowledgements	5
Abstract	7
List of Tables	23
List of Figures	23
1 Introduction	25
1.1 Motivation and Research Objective	25
1.2 Thesis Structure	26
1.3 Research Methodology	26
1.3.1 First project: Systemic Risk - Disentangling the influence of interbank network topology and dynamics	27
1.3.2 Second project: Topological determinants of control energy	28
1.3.3 Third project: Reverse stress testing interbank networks	29
1.3.4 Fourth project: Reverse stress testing interbank networks - The influence of topology	29
1.4 Contributions of this Thesis	30
1.4.1 Systemic Risk: Disentangling the influence of interbank net- work topology and dynamics	31
1.4.2 Topological determinants of control energy	31
1.4.3 Reverse stress testing interbank networks	31
1.4.4 Reverse stress testing interbank networks - The influence of topology	32
1.5 Publications and Talks	32
2 Background and Literature Review	35
2.1 Systemic Risk	35
2.1.1 General definitions and a financial perspective	35
2.1.2 Network theory underlying systemic risk	37
2.2 Null Models of Networks	39
2.2.1 General insights gained from null models	40
2.2.2 Maximum entropy approach	41
2.2.3 Fitness model	42

3	Systemic Risk: Topology and Dynamics	45
3.1	Introduction	45
3.1.1	Motivation	45
3.1.2	Research objective and approach	46
3.2	Related Work	47
3.2.1	Financial contagion	47
3.2.2	Interbank network topology and financial contagion	49
3.2.3	Global Cascade model	52
3.3	Methods	53
3.3.1	Counterparty default risk - Furfine model	53
3.3.2	Interbank stress propagation - DebtRank	54
3.3.3	Network topology characteristics	55
3.3.4	Null model networks	58
3.4	Analysis of the Austrian Interbank Network	61
3.4.1	Description of interbank data	61
3.4.2	Basic overview of network topology	62
3.4.3	Counterparty risk	63
3.5	Analysis of Reconstructed Networks	65
3.5.1	Network topological properties	65
3.5.2	Analysis of counterparty risk in the reconstructed networks	68
3.6	Vulnerable Cluster	75
3.6.1	Intuition	75
3.6.2	Vulnerable clusters in the real and reconstructed networks	76
3.7	Discussion and Future Developments	77
4	Control Theory on Networks	81
4.1	Introduction	81
4.1.1	Motivation	81
4.1.2	Research objective and approach	82
4.2	Related Work	83
4.2.1	Control theory on networks	83
4.2.2	Defining the controllability of a dynamical system	84
4.2.3	Control energy	84
4.2.4	The impact of network topology on control energy	86
4.3	Methods	87
4.3.1	Control system set-up	87
4.3.2	Static method for scale-free network generation	88
4.3.3	System stability	88
4.3.4	Control energy	89
4.4	Topological Properties of Network Control Energy	90
4.4.1	Simulation set-up	90
4.4.2	Increasing network connectivity reduces control energy	91
4.4.3	Analytical solution to optimal control	92
4.4.4	Topological determinants of control energy	96
4.4.5	Absence of coordination advantage when nodal states deviate	97
4.5	Discussion and Future Developments	100

5	Reverse Stress Tests	103
5.1	Introduction	103
5.1.1	Motivation and research objective	103
5.1.2	Approach	103
5.1.3	Contributions	104
5.2	Related Work	104
5.2.1	Measuring the contribution of banks to systemic risk	105
5.3	Methods	107
5.3.1	Reverse stress testing problem set-up	107
5.3.2	Link to network control theory	110
5.4	Analysis of Homogeneous System	110
5.5	Case Study Based on Largest European Banks	111
5.5.1	Description of data	112
5.5.2	Aggregate properties of optimal shocks	113
5.5.3	Concentration of risk	115
5.5.4	A simple policy experiment	118
5.5.5	Robustness of results	122
5.6	Analytical Approach to the Optimisation Problem	122
5.6.1	Comparison of analytical and numerical results	124
5.7	Discussion and Future Developments	125
5.7.1	Future developments	126
6	Reverse Stress Tests - Topology	129
6.1	Introduction	129
6.1.1	Motivation and approach	129
6.2	Density Comparison across STOXX Networks	130
6.2.1	Exogenous shock across densities	130
6.2.2	Equity capital policy	132
6.3	Comparison to Austrian Network	134
6.3.1	Reverse stress testing the Austrian network	134
6.3.2	Policy experiment	136
6.3.3	Robustness tests	139
6.4	Reconstruction of Austrian Network	139
6.4.1	Estimate of exogenous shock	139
6.4.2	Concentration of nodal shocks	141
6.4.3	Role of network topology in reverse stress testing	143
6.4.4	Robustness tests	144
6.5	Discussion and Future Developments	145
6.5.1	Future developments	146
7	Conclusion and Discussion	149
7.1	Conclusions of Individual Projects	149
7.1.1	First project: Systemic risk - Disentangling the influence of interbank network topology and dynamics	149
7.1.2	Second project: Topological determinants of control energy	150
7.1.3	Third project: Reverse stress testing interbank networks	151

7.1.4	Fourth project: Reverse stress testing interbank networks - The influence of topology	151
7.2	Conclusion and Discussion across Projects	152
7.3	Future Developments	154
Appendix		156
A.1	Properties of the Topology and Financial Contagion in the Austrian Network and its Reconstructed Networks in Different Quarters	157
A.1.1	Topological measures	157
A.1.2	Trajectory profiles	158
A.2	Identification of Vulnerable Clusters	159
A.2.1	Vulnerable cluster identification algorithm	159
A.2.2	Simulations to confirm the efficacy of the VC identification algorithm	160
A.2.3	Creation of synthetic bank balance sheets	161
A.3	Reverse Stress Testing the Austrian Network	162
A.4	Reverse Stress Testing Reconstructed Networks	164
List of References		167

List of Tables

3.1	Summary statistics of topological properties of the real and reconstructed networks. For the reconstructed networks the metrics are averaged over 10,000 simulations and rounded to the nearest significant integer that can be described with confidence. The figures in the brackets represent the 10% and 90% quantiles, respectively. Results presented refer to the first quarter of 2006.	66
3.2	Abundance and average maximum size of vulnerable clusters in real and reconstructed networks. There is no single giant vulnerable cluster in the reconstructed networks, but a much larger abundance of individual vulnerable clusters. This shows that the local concentration of risk is much larger in the real network. The values for the reconstructed networks are averaged over 10,000 simulations. Results shown for the first quarter of 2006.	77
6.1	Conversion table for λ_{max} to β in the real network.	140
6.2	Root mean squared distance of nodal energies with respect to the real network. The configuration model always performs best and the fitness model has the largest error when β is large.	143
A.1	Summary statistics of topological properties of the real and reconstructed networks. For the reconstructed networks, the different metrics are averaged over 10,000 simulations and rounded to the nearest significant integer that can be described with confidence. The figures in the brackets represent the 10% and 90% quantiles, respectively. Results presented refer to the first quarter of 2007.	157

A.2	Summary statistics of topological properties of the real and reconstructed networks. For the reconstructed networks, the different metrics are averaged over 10,000 simulations and rounded to the nearest significant integer that can be described with confidence. The figures in the brackets represent the 10% and 90% quantiles, respectively. Results presented refer to the first quarter of 2008.	158
-----	--	-----

List of Figures

3.1	Illustration of the triadic motifs studied in this chapter.	57
3.2	Rewiring steps in the directed configuration (rewired) model.	59
3.3	Distribution of equity capital, interbank assets and liabilities of each bank in the first quarter of 2006.	62
3.4	Degree distribution of the real network (black), Erdős-Rényi (blue) and fitness model (red) networks. The distribution corresponding to the configuration model is not shown because of its perfect overlap with the real network. Whilst the Erdős-Rényi network has a substantially different degree distribution to the real network, the fitness model closely matches the real network's heterogeneity. The degree distribution therefore appears to be induced by the heterogeneity of balance sheet sizes, which are the inputs of the fitness model.	63
3.5	Edge weight distribution of the real network (black), random Erdős-Rényi (ER) network (blue), configuration model (pink) and the fitness model (red). Due to the allocation of edge weights using the RAS algorithm for all null models, the edge weight distribution of the three null models is much more similar than the degree distribution shown in fig. 3.4. Nonetheless, it can clearly be seen that the ER networks consist of a much larger abundance of small edge weights and a longer tail that reflects larger edge weights.	67

- 3.6 Main contagion measures across 12 quarters from 2006 until 2008. Black line shows the measure of the real network. The shaded areas show the 10% and 90% quantiles of the null models respectively and the dotted lines their averages. Blue refers to Erdős-Rényi, pink to configuration and red to fitness model networks. The measured values of contagion in the reconstructed networks generally do not fit those observed in the real network, except for the DebtRank amplification shown in panel (d). The null model estimates are more robust across time than the values observed in the real network, due to their dependence on the relatively more stable total interbank assets and liabilities via the RAS algorithm. 70
- 3.7 Trajectories of financial contagion due to the Furfine algorithm. Contagion is more abrupt and concentrated in the real relative to the null model networks. Black line shows the measure of the real network. Blue refers to Erdős-Rényi, pink to rewired and red to fitness model networks. L.H.S. figure based on network of first quarter of 2006 only, which is representative of all other quarters. Trajectories are cut off at iteration 10, which in some simulations do not decay until iteration step 20. R.H.S. the shaded areas show the 10% and 90% quantiles of the null models respectively and the dotted lines their averages. . . . 72
- 3.8 Equity capital lost due to counterparty default risk as measured by the Furfine algorithm across 12 quarters from 2006 until 2008. Black line shows the measure of the real network. The shaded areas show the 10% and 90% quantiles of the null models respectively and the dotted lines their averages. Blue refers to Erdős-Rényi, pink to configuration and red to fitness model networks. Similarly to the contagion extent shown in fig. 3.6b, the null models underestimate the equity lost, thereby indicating a lower concentration of risk in their networks. . . . 74
- 3.9 A schematic illustration of the (extended) giant vulnerable cluster (GVC). Unfilled nodes are in the GVC, filled nodes are robust banks out-with the GVC, but part of the EGVC. Each robust and vulnerable bank is capable of triggering the failure of the entire GVC. Only "vulnerable" exposures are shown, which could lead to the default of a lender if the borrower is unable to repay. 76

4.1 Average total network control energy E across simulations as a function of the average degree $\langle k \rangle$ for ER and SF networks. Inset: for $n = 100$, in which the difference between ER and SF networks is less strong. 91

4.2 L.H.S.: Average across 5,000 ER network simulations (straight line) together with analytical solution (dashed line) from eq. (4.28) for $t_f = 3$ (blue and red respectively) and $t_f = 50$ (black and green respectively). The analytical and numerical solutions match perfectly. R.H.S.: Analytical solution as a function of t_f for $\langle k \rangle = 100$ 95

4.3 Distribution of total energy E for ER and SF networks, here shown for $\langle k \rangle = 25$, which is representative of all average degrees. 95

4.4 The figures represent the expected fractional nodal energy $\frac{E_i}{E}$ requirement for a given degree k_i averaged across all simulations for ER and SF networks, where the colours represent the different average degrees $\langle k \rangle$. The fraction is almost identical for nodes across all nodal degrees in ER networks, but highly heterogeneous for SF networks, with hub nodes requiring the largest fraction of the energy. 96

4.5 The inverse participation ratio (IPR) for ER (blue) and SF (red) networks. The concentration of nodal energies is much higher in SF relative to ER networks. 97

4.6 Total network control energy E and inverse participation ratio (IPR) as a function of the average network degree $\langle k \rangle$ for various target vector entry distributions. Blue lines reflect results for ER networks and red lines for SF networks. Straight lines reflect the case when each node has an identical target state, dashed lines when the target states are normally distributed around one and dashed-dot lines when target states are distributed in a uniform sphere around zero. When nodal target states deviate as shown by the normal and uniform sphere distributions, the energy requirement increases (L.H.S.) and the nodal energies are much more concentrated, in particular for low $\langle k \rangle$ (R.H.S.). 98

- 4.7 The figures represent the expected fractional nodal energy requirement $\frac{E_i}{E}$ for a given degree k_i averaged across all simulations for SF networks for normally distributed target vector entries and the uniform sphere. As opposed to the case of identical nodal target states, shown in fig. 4.4, the distribution of the energy requirements across nodes is less dependent upon k_i . In fact, in the case of the uniform sphere shown in panel (b) nodal energies are much more evenly split across nodal degrees. The colours represent the different average degrees $\langle k \rangle$. I do not show the respective figures for ER networks, as their nodal energies change very little as a function of k_i with respect to the different distributions, see fig. 4.4 as a reference. 99
- 4.8 Fractional nodal energies $\frac{E_i}{E}$ for each nodal target state $x_i(t_f)$ averaged across all simulations. The figures represent the expected fractional nodal energy requirement for a given nodal target state $x_i(t_f)$ for normally distributed target vector entries. The colours represent the different average degrees $\langle k \rangle$. Figures shown refer to SF networks, the figures for ER networks are very similar. 100
- 5.1 Distribution of equity capital, interbank assets and liabilities of each bank in 2015. 113
- 5.2 Size of shock K in STOXX network. 114
- 5.3 The loss fraction measures the total amount of equity externally shocked as a fraction of the total equity in system. Shown as a function of the target losses $\ell_i = \ell$ and λ_{\max} . The loss fraction always remains below one, indicating that the endogenous network dynamics propagate the losses. 115
- 5.4 Distribution of the standardized shock of each node (z-score) for different values of λ_{\max} . Shown for $T = 20$ and for $\ell_i = 0.1$. s is the standard deviation and μ the mean of the respective distribution of nodal shocks K_i for each λ_{\max} . As λ_{\max} increases larger relative deviations with respect to the mean amplitude are observed in the right tail of the distribution. 116
- 5.5 Inverse participation ratio (IPR) in the STOXX network. The nodal shocks become more concentrated as λ_{\max} increases. 117
- 5.6 Histogram of interbank leverages of each bank in the STOXX network under investigation. It can be seen that lenders have heterogeneous strategies for their interbank leverage, defined as the ratio of the total amount lent (assets) over a bank's equity capital. 118

5.7 Incremental addition of $u_i(t)$ to banks in a decreasing order of the size of their individual shocks K_i for various λ_{\max} . The circles represent the average across 500 simulations in which banks were randomly selected (same colours as for the three lines corresponding to the three different λ_{\max} shown in the legend). For $\lambda_{\max} = 1.5$ the shock on the first five banks already accounts for 70% of final losses. Results shown for $T = 20$ and $\ell_i = 0.1$. As noted before, the behaviour is very similar for larger values of T 119

5.8 Results of policy experiments. Losses R are recomputed after the equity was increased and expressed as a fraction of the original losses R_0 on the y-axis. These results are shown as a function of λ_{\max} along the x-axis. Results shown for $T = 20$ and $\ell_i = 0.1$ 121

5.9 Scatter plots comparing nodal shocks K_i for each bank across three years 2014, 2015 and 2016. It can clearly be seen that the nodal shocks are very similar across the years. Results shown for $T = 20$ and $\ell_i = 0.1$ 122

5.10 Results from analytical solution (5.26) for exogenous shocks K 125

6.1 Average exogenous shock $\sqrt{K/n}$ and inverse participation ratio (IPR) as function of network density for various β . The blue lines refer to $\lambda_{\max} \approx 0.5$, pink to $\lambda_{\max} \approx 1.0$ and red to $\lambda_{\max} \approx 1.5$. Shown for target loss $\ell_i = 0.1$ and time horizon $T = 20$. Dashed lines show the 10% and 90% quantiles, respectively. As it can be by the tight dashes lines, the results are stable across all simulations. 131

6.2 Average exogenous shock $\sqrt{K/n}$ and inverse participation ratio (IPR) for each of the 100 generated STOXX networks with density 10% and $\beta = 1.0$ ($\lambda_{\max} = 1.5$). Isolated nodes ($k^{in} = 0$) are a core driver of the average exogenous shock and the IPR. 132

6.3 Comparison of policies across different densities of the STOXX network. Results shown for target loss $\ell_i = 0.1$ and time horizon $T = 20$. Dashed lines show the 10% and 90% quantiles, respectively. As it can be by the tight dashes lines, the results are stable across all simulations. 133

6.4 Average exogenous shock $\sqrt{K/n}$ and inverse participation ratio (IPR) for the Austrian network. Results shown for $T = 20$ 135

6.5	Distribution of the standardized shock of each node (z-score) for different values of λ_{\max} in the Austrian network. Results shown for $T = 20$ and for $\ell_i = 0.1$. s is the standard deviation and μ the mean of the respective distribution of nodal shocks K_i for each λ_{\max} . As λ_{\max} increases larger relative deviations with respect to the mean amplitude are observed in the right tail of the distribution.	136
6.6	Incremental addition of $u_i(t)$ to banks in a decreasing order of the size of their individual shocks K_i for various λ_{\max} . Results shown for $T = 20$ and $\ell_i = 0.1$	137
6.7	Equity capital policies applied to the Austrian network. The K_i -based policy in the subnetwork (R.H.S.) is much more effective than in the full network (L.H.S.).	137
6.8	K_i -based policy for the subnetwork with $n = 225$ in comparison to the 5-Core with $n = 161$ and the 6-Core with $n = 147$. As the connectivity increases, indicated by the k-core, the policy becomes more effective.	138
6.9	Average exogenous shock $\sqrt{K/n}$ for the three null models and the real network. The error bars indicate the 5th and 95th percentiles, respectively. Maximum entropy network plotted as a benchmark. The configuration model performs best out of the three null models, the ER networks perform worse than the fitness networks when $\beta < 0.24$, with this relation being reversed when $\beta > 0.24$. Results shown for $T = 20$ and $\ell_i = 0.1$	141
6.10	Inverse participation ratio (IPR) as a function of β for the three null models and the real network. The error bars indicate the 5th and 95th percentiles, respectively. Maximum entropy network plotted as a benchmark. Results shown for $T = 20$ and $\ell_i = 0.1$	142
6.11	Nodal shocks K_i in log-scale for the three null model networks shown vis-à-vis the expectation from the real network. The fitness network has the largest outliers.	143
6.12	Average exogenous shock $\sqrt{K/n}$ as a function of β for the fitness networks and real network, in which the isolated nodes in the fitness networks have been removed from both networks. The match between the two networks is almost perfect, suggesting that the poor performance of the fitness networks is due to their overestimation of the abundance of isolated nodes.	144

A.1	Trajectories of financial contagion, measured as the total number of banks failed in each iteration of the dynamic scaled by the number of contagion events in the first iteration. Black line shows the measure of the real network. Blue refers to Erdős-Rényi, pink to configuration and red to fitness model networks. X-axis cut-off at 10 iterations.	159
A.2	Reproduction of experiments undertaken by Watts and Hurd, to prove the efficacy of the VC identification algorithm [135, 74]. Figures represent the averages over 1,000 simulations at each average degree $\langle k \rangle$ on the x-axis. The fractional EGVC size approximates the contagion probability almost perfectly. L.H.S.: In networks with $10 < \langle k \rangle < 20$ the fractional EGVC size does not match the contagion probability, as they are typically defined as $> 5\%$ of nodes failing.	161
A.3	Average exogenous shock $\sqrt{K/n}$ and inverse participation ratio (IPR) for the Austrian network based on data for the first quarter of 2007. Results shown for $T = 20$	163
A.4	Equity capital policies applied to the Austrian network based on the first quarter of 2007. The K_i -based policy is much more effective in the subnetwork (R.H.S.) than in the full network (L.H.S.).	164
A.5	Average exogenous shock $\sqrt{K/n}$ for the three null models and the real network as a function of β . The error bars indicate the 5th and 95th percentiles, respectively. Results shown for $T = 20$ and $\ell_i = 0.1$ based on data for the first quarter of 2007.	165
A.6	Inverse participation ratio (IPR) as a function of β for the three null models and the real network. The error bars indicate the 5th and 95th percentiles, respectively. Results shown for $T = 20$ and $\ell_i = 0.1$ based on data of the first quarter of 2007.	166
A.7	Nodal shocks K_i in log-scale for the three null model networks shown vis-à-vis the expectation from the real network. The fitness networks have the largest outliers.	166

Chapter 1

Introduction

1.1 Motivation and Research Objective

Financial stability is one of the main concerns for the regulators of financial markets. Only when financial markets are stable, do companies feel confident to invest and thereby keep an economy prosperous. Understanding the dynamics of financial stability is thus not only in the interest of regulators, but also policy-makers, businesses, banks and ultimately the general public. This research area is especially timely given the financial crisis of 2008 and the resulting economic uncertainty that still affects financial markets today.

In this thesis I consider a complex systems approach to study systemic risk. I focus in particular on the influence of the topology of interbank networks on financial contagion. Systemic risk describes the possibility of the breakdown of a financial system and emerges from the endogenous network dynamics resulting from the borrowing and lending interactions (edges) between banks (nodes). Crucially, when modelling the dynamics of systemic risk, the topology of interbank networks –that is the bilateral connections between banks– is very influential in the way financial contagion propagates throughout a network. An understanding of the interplay of network topology and dynamics is therefore essential in studies of systemic risk.

Typically analyses of financial contagion employ a forward-looking approach, by studying the impact of an initial exogenous shock on the stability of an interbank network. Within such approaches a shock is propagated throughout the network and the resulting losses of each bank are summed across all iteration steps of propagation. Using this approach it cannot be known whether a shock employed in a marginally different manner, for instance in terms of its size or the banks that are targeted, has a significantly different effect upon the stability of the system. With these limitations in mind, I introduce a framework to reverse stress test interbank networks. Using this framework I consider the opposite problem to forward-looking stress tests, by

computing the minimal exogenous shocks that need to affect each bank to cause a specific loss to the system.

The reverse stress testing framework is similar to an optimal control problem on a network. To develop the methodological background to framework, I use control theory, a branch in the engineering sciences, to study the cost of managing the dynamics of a network. The goal in control theory is to move a dynamical system from a given state to a target state. Since control theory has only recently been applied to networks, essential parts of the theory are not yet understood within networks. The computation of control energy, which measures the cost of control, is one of the central parts in any control problem. The properties of control energy on networks have not yet fully been understood [86, 128]. Since the reverse stress test framework requires the ability to manage, i.e. to compute the cost of controlling, systemic risk in an interbank network, this research is timely not only from a perspective of systemic risk, but also for control theory.

1.2 Thesis Structure

This section provides a brief overview of the structure of this thesis. In this thesis I undertake four projects. Each project is concerned with a different aspect of the impact of network topology upon financial contagion. In the subsequent second chapter I provide a background literature review to systemic risk and network reconstruction. The third chapter describes the first project, in which I study the required amount of information in order to reconstruct the topology and dynamics of interbank networks. In the fourth chapter I describe the second project, in which I investigate the influence of the topology of a network on the cost of controlling its dynamics. In the subsequent third project I use the insights from the first and second projects to develop a framework to reverse stress test interbank networks. I describe these results in the fifth chapter. In the sixth chapter I present the results of the fourth and final project, in which I study the robustness and properties of the reverse stress testing framework in interbank networks with different topologies. I summarise and discuss the results across all projects in the final seventh chapter and provide suggestions for future research.

1.3 Research Methodology

In this section I introduce the individual projects and thereby provide a more detailed overview of the research methodology that I use in this thesis.

1.3.1 First project: Systemic Risk - Disentangling the influence of interbank network topology and dynamics

In the first project I consider the problem of the amount of information required to estimate systemic risk due to counterparty risk. I obtain data of the Austrian interbank network to compare its topology and financial contagion properties to that of reconstructed networks, which are built from limited information. Through a comparison of the properties of the real network's dynamics to that of the reconstructed networks, a much deeper understanding of the topological determinants of financial contagion can be gained, than by considering a real network in isolation. Since the same counterparty risk algorithm is employed on the reconstructed and real networks, I can identify the topological factors that influence financial contagion by comparing the topology of the networks.

I consider the following models, which I chose due to the different amounts of prior information they require of a real network: i) Erdős-Rényi random networks in which the only information from the real network is the average degree across nodes, ii) the fitness model in which information on total interbank assets and liabilities of each bank is used to reconstruct the network, iii) the directed configuration model that preserves the entire degree sequence of the original network by rewiring each edge. The edge weights of all three reconstructed networks are fitted with the RAS algorithm [13, 23].

After I reconstruct the real network, I compare the results of stress tests performed on the real interbank network with those performed on the reconstructed networks. I test in particular the counterparty default risk algorithm due to Furfine [58] and the DebtRank algorithm [17]. For the Furfine algorithm my analysis shows that the fitness model can only measure the general trend in financial contagion, but cannot capture the exact values of contagion observed in the real network. This suggests that knowledge of the network density as well as of total interbank assets and liabilities of each bank is sufficient to assess the general trends in the stability of a financial network.

However, a comparison between the contagion trajectories of the real and reconstructed networks shows that there is a higher concentration of risk in the real network. This higher concentration of risk is exemplified in the real network by a vulnerable cluster that is much larger and triadic motifs that contain more risky banks compared to the reconstructed networks. The vulnerable cluster is a set of banks that default when only one of their counterparties fails to repay its liabilities. The inability of the null models to accurately measure financial contagion due to the Furfine algorithm is largely due to their inability to reconstruct the local topology

of the real network. In contrast, all three null models reconstruct financial contagion measured by the DebtRank algorithm well. The reason for this is that the DebtRank algorithm does not strongly depend on the local topology, but rather on global topological properties such as the network density. Moreover, the DebtRank algorithm is strongly influenced by interbank leverages, which the null models indirectly obtain via the RAS algorithm that allocates edge weights based on the total interbank assets and liabilities of each bank.

1.3.2 Second project: Topological determinants of control energy

Because the properties of control energy, the cost of controlling a network, are not yet understood, I undertake a theoretical study of the impact of the topology of networks upon control energy in the second project. The application of control theory to networks is a relatively recent phenomenon, which was introduced through a study by Liu et al in 2011 [86]. Since then the literature has focused upon structural controllability, which describes mathematical criteria that define when a network is controllable. A central aim of this effort is to find the minimum number of nodes that require to be externally controlled to ensure structural controllability.

Yan et al in 2012 and 2015 produced the first studies that systematically investigate the control energy in Erdős-Rényi (ER) and scale-free (SF) networks [137, 138]. These studies established the minimum and maximum bounds of control energy for both types of networks. These bounds are very wide and therefore do not provide a sufficiently detailed understanding of the influence of network topology and dynamics on control energy.

In this project I construct ER and SF networks with varying levels of connectivity and compute the control energy for a simple control task in these networks. I find that control energy decreases with an increasing network connectivity, because nodes that have a common target state push one another to their common goal and thus reduce the required exogenous control input. Importantly, this effect is more pronounced in SF networks due to hub nodes that can coordinate the state of most nodes in a network. I provide an analytical derivation based on a mean-field approximation, to prove the diminishing effect of network connectivity on control energy.

1.3.3 Third project: Reverse stress testing interbank networks

In the third project I use the DebtRank algorithm to propagate stress in an interbank network and develop a framework inspired by network control theory to reverse engineer these dynamics. Using this framework I can identify the optimal (minimal) exogenous shocks that would lead to a given cumulative loss of equity for each bank in the network. I test this framework in a network consisting of the 44 largest European banks that are members of the STOXX Europe 600 Banks index. Using the RAS algorithm I construct a full network of interbank assets and liabilities between all banks. I chose to apply the reverse stress testing framework to a full network in order to focus on the dynamics of the framework free of the influence of network topology.

I find that as the interbank leverage increases, that the size of the exogenous shock decreases. I measure interbank leverage via the largest eigenvalue of the matrix of interbank leverages, which describes the relative impact of the default of a borrower on its lender's equity capital. The reason for the decrease of the exogenous shock is that as the dynamics become faster, the interactions between nodes increase the systemic losses, such that a smaller exogenous shock is required to reach the targeted losses. An additional effect of faster network dynamics is that the shocks become increasingly concentrated in fewer banks.

The increased concentration of nodal shocks, i.e. the increased concentration of risk in a subset of banks, allows me to create a ranking of banks in terms of their systemic importance. Banks with the largest exogenous shocks are those that have the largest impact on the stability of the system. Based on this ranking of banks I design a simple policy to reduce the observed system losses. The policy allocates additional equity capital to banks based on their relative systemic importance. I show that this policy can significantly reduce the observed losses of the system by up to 95%.

1.3.4 Fourth project: Reverse stress testing interbank networks - The influence of topology

In this final project I analyse the reverse stress testing framework, that I developed in the previous project, in different network topologies. I apply the framework to the STOXX network with a range of different densities that I create using the RAS algorithm, and the Austrian network together with its reconstructed networks using the three null models that I study in the first project. This approach allows me to test the the influence of network topology on the properties of the exogenous shock

that is computed via the framework.

I find that the exogenous shock is independent of the density when the STOXX network is stable. In contrast, when the network is unstable, the exogenous shock decreases as a function of the density. I further find that the equity capital policy used to reduce the observed system losses is effective only when the density of the network is larger than 50%. The reason for this is that the additional equity capital injected by the policy can only be effectively spread throughout the network, when the connectivity of the network is sufficiently high.

The properties of the total exogenous shock are similar in the Austrian relative to the STOXX network. However, due to the higher heterogeneity of the nodes in the Austrian network, the individual nodal shocks are more concentrated in the Austrian than in the STOXX network. The reason for this is that the largest nodes in a network absorb the largest fraction of the exogenous shocks.

With respect to the reconstructed networks via the three null models, I surprisingly find in contrast to the first project, that the fitness networks have the largest measuring error. The configuration model performs best out of the three null models in terms of estimating the size of the total exogenous shock of the Austrian network. The Erdős-Rényi networks perform better than the fitness networks when the underlying network is unstable, and conversely, the fitness networks provide better estimates than the Erdős-Rényi networks when the system is stable. The poor performance of the fitness model is due to its overestimation of the abundance of isolated nodes. Isolated nodes require the largest fraction of the exogenous shock in a network, as they cannot be affected by the endogenous network dynamics. The total exogenous shock in the fitness networks is therefore much larger than in the Austrian network.

1.4 Contributions of this Thesis

This thesis contributes to complex systems theory by providing insights into the influence of the topology of a complex system on its dynamics. I use this connection between network topology and dynamics to develop a framework to reverse stress test interbank networks. This framework is a specific example of a new strand of research that develops methodologies to manage complex system dynamics with targeted interventions via their topology. In the remainder of this section I outline the contributions that I make in each of the individual research projects.

1.4.1 Systemic Risk: Disentangling the influence of interbank network topology and dynamics

1. Most previous research investigated reconstructed networks without comparing their ability to estimate financial contagion directly with other reconstructed networks. I compare and contrast the results of an analysis of two financial contagion algorithms, the Furfine and DebtRank algorithms, in reconstructed networks with those of a real interbank network.
2. In contrast to the previous literature that focused on the good ability of fitness model in reconstructing the higher-order properties of the topology of interbank networks, I find that the fitness model only captures the general trend of financial contagion and cannot accurately recreate the exact values of financial contagion observed in the real network due to the Furfine algorithm. In contrast, all three null models reconstruct counterparty risk measured with the DebtRank algorithm well.
3. In all three reconstructed networks the trajectory profiles over time due to the Furfine algorithm are quite different to the real network. I show that this difference is due to a higher concentration of systemic risk in the real network that results from a higher clustering of risky banks. I quantify the higher clustering of risk by identifying vulnerable clusters, i.e. sets of banks that default upon the failure of only one of their counterparties.

1.4.2 Topological determinants of control energy

1. Total network control energy, which measures the cost of control, decreases with the average degree of the network.
2. The decrease in energy is more pronounced in scale-free networks due the presence of large hub nodes, which can coordinate the states of a large fraction of nodes throughout the network.
3. I provide an analytical derivation of the control energy for Erdős-Rényi networks based on a mean-field approximation.

1.4.3 Reverse stress testing interbank networks

1. I introduce a framework to reverse stress test interbank networks. This framework allows the identification of the minimal exogenous shocks that lead to specific total losses in an interbank network. This framework provides a more

rigorous method with which to identify potential causes of financial instability than previous approaches, which a priori presume the causes of instability.

2. As the largest eigenvalue of the matrix of interbank leverages increases the optimal shocks become smaller and more concentrated in a smaller set of banks.
3. Based on the distribution of optimal shock sizes, I rank banks in terms of their systemic importance. This ranking can be used to make the system more robust through the implementation of a targeted equity capital policy.

1.4.4 Reverse stress testing interbank networks - The influence of topology

1. I find that the exogenous shock is independent of the network density when the network is stable, however when it is unstable then networks with higher densities require lower shocks.
2. The equity capital policy employed to reduce the observed system losses, is significantly less effective when the network connectivity is low.
3. The fitness model performs worst in estimating the size of the exogenous shock out of the three network null models studied. The reason for the poor performance of the fitness model is due to its underestimation of the connectivity of small nodes and its overestimation of the connectivity of large nodes.

1.5 Publications and Talks

In this section I briefly mention the publications and talks that resulted from my individual projects.

- Project 1 on systemic risk in the Austrian interbank network is currently in submission with the *Journal of Economic Interaction and Coordination*.
- Project 2 on control energy in networks is in preparation for submission. I have given invited talks of this project at:
 1. *6th NII-Kobe Collaborative Workshop* at the *University of Kobe*, Japan, 4th March 2016.
 2. *Inoue Laboratory for Inference and Learning at the National Institute of Informatics and the University of Tokyo*, Japan, 30th March 2016.

- Project 3 on reverse stress testing is currently in the second round of reviews with Scientific Reports. A preprint is available on arXiv: 1702.08744 with the title "Reverse stress testing interbank networks". I have given a talk on this project at the *13th Econophysics Colloquium at the University of Warsaw*, Poland on the 6th of July 2017.
- Project 4 on the impact of network topology on reverse stress testing is in preparation for submission to the *Journal of Network Theory in Finance*.

Chapter 2

Background and Literature Review

2.1 Systemic Risk

In this section I discuss the literature on systemic risk that is related to my thesis. After generally introducing systemic risk, I focus on applications of network theory to systemic risk. I put a particular emphasis on the literature that studies the topology of interbank networks.

2.1.1 General definitions and a financial perspective

Generally speaking systemic risk is the risk associated with a catastrophic breakdown of a financial system. A broad characterisation is to describe systemic risk as the risk not specific to an agent in a financial market, but rather as the risk that naturally emerges in a system due to the interaction of all agents, these interactions are referred to as the endogenous dynamics of a system [23, 36, 97].

Traditionally any risk associated with financial markets, including the relatively recent rise of the systemic risk concept, has been investigated from a purely financial perspective. It is therefore useful to briefly discuss the main insights from this financial research, to link the mechanisms that these studies identified to those discovered through the methods of network theory, which I discuss in the following section 2.1.2. In particular, I focus on the leverage cycle, which is a factor that is influenced by many driving mechanisms of systemic risk and that significantly changes the systemic risk profile of a market system.

The rise and fall of an economy is often explained with the leverage cycle, which was originally described by Minsky [62]. Leverage moves pro-cyclically: In prosperous times asset prices are high as easy borrowing conditions allow high leverages; however, during depressed times asset prices are low because of constrained borrowing conditions [62]. In times of financial distress margins increase and borrowing

becomes expensive. In response investments turn unprofitable, which results in more asset sales and further decreases in leverage. Many factors that drive financial markets therefore influence one another [34]. Especially in downturns the normally heterogeneous perspectives of investors align to a narrower pessimistic outlook, which is signified by an increasing correlation of asset prices [45]. Brunnermeier et al (2009) find a non-linear impact of the interaction of the ease of borrowing with asset prices. This interaction results in a negative feedback loop, such that asset prices often drop below their fundamental value in a market correction after prosperous times [34]. Distressed periods are also characterised by higher volatilities, due to the rising uncertainty relative to the optimism that characterises prosperous periods [62].

Leverage allows some investors to exert a much larger influence on the price of financial assets relative to other less leveraged market participants. From a perspective of systemic risk this means that the stability of a market depends on a few major connected agents. The failure of only one of these agents can cause significant disruptions to an economy [62]. In a very extensive study Jorda et al (2017) studied the history of the leverage cycle across 17 advanced economies over 150 years [79]. The authors find that larger leverages, defined as the ratio of total credit to gross domestic product (GDP), lead to slower recoveries from crises. However, larger leverages also have the effect of smoothing out smaller crises, in that they allow an economy to absorb small shocks by borrowing additional capital.

Traditional economic supply-demand models that define equilibrium interest rates and that are based on the assumption of a law of one price, provide very similar value estimates for different liabilities. Geanakoplos (2010) shows that these models fail to adequately forecast the dynamics of the leverage cycle, precisely because of their assumption of homogeneous investor perspectives [62]. Geanakoplos suggests different types of models, that forecast equilibrium leverage and margins based on individual price estimates by analysing each loan separately. The core rationale for this suggestion is that the same loan can have different prices for different investors, due to investors' varying ability to leverage returns. See Aymanns et al in 2014 and 2016 for two recent studies that define an optimal leverage policy [11, 12]. For a comprehensive overview of methods used in traditional finance to analyse systemic risk see Bisias et al (2012) [30].

Two insights from the preceding review of the financial perspective on systemic risk are of importance to the network theoretic approach to systemic risk. Firstly, heterogeneity influences the manner in which markets price assets; and secondly, several interacting mechanisms are at play in driving systemic risk, that can be roughly summarised with the leverage of financial markets. The focus of my thesis is to study the impact of the structure of interactions between market participants

on systemic risk. Networks allow these heterogeneous interactions to be integrated with the leverage dynamics of financial markets. In the following section, together with section 3.2, I discuss how these features of systemic risk are studied in the network theory literature.

2.1.2 Network theory underlying systemic risk

Network theoretic investigations of systemic risk have identified multiple channels of financial contagion in interbank networks, which describe the lending and borrowing activities between banks. These include amongst others, counterparty default risk, overlapping portfolio risk and liquidity risks. The topology of interbank networks is crucial in defining these channels of financial contagion. In this section I discuss the intricacies of interbank network topologies in preparation for a discussion of the channels of financial contagion in section 3.2.

Tasca et al (2016) provide a link between the procyclical behaviour of leverage and interbank networks [129]. They find that networks can amplify the leverage cycle, which is driven largely by banks' equity capital requirements and the liquidity of assets. For instance, low levels of equity capital and low liquidity result in much higher levels of systemic risk, because banks do not have sufficient capital buffers to absorb losses from their counterparties and cannot sell their assets in time to increase their equity capital [129].

The large-scale organisation of interbank networks

Many studies of national interbank markets have been undertaken, see Upper (2011) [132] and Glasserman et al (2016) [66] for two extensive reviews. For a major representative study on the topology of interbank networks based on the Mexican interbank network, with access to rare daily data of interbank exposures, see the study of Martinez-Jaramillo et al in 2014 [89]. Most of these studies have identified topological features that occur across interbank networks, which therefore appear to be universal to interbank networks. My focus in this section is to describe these common topological features of interbank networks.

A major advantage of network theory is that it enables scholars to model the emergence of the heterogeneity of financial data, such as losses and leverage [62, 122]. Networks with power-law (heterogeneous) degree distributions are called scale-free (SF) networks. SF networks tend to consist of many poorly connected nodes and a few extremely well connected nodes [75, 119, 122]. Interbank networks are typically disassortative, such that large banks are typically the borrowers and small banks the lenders [44, 75, 132].

Moreover, interbank networks tend to have a core-periphery structure. The core is a set of banks that coordinates the majority of the lending activity in an interbank network. An ideal core-periphery network is characterized by a set of core banks, that are fully interconnected, and a set of peripheral banks that are connected to the core but have no connections to one another. Core-periphery structures have been identified in many interbank markets including the German [44], Dutch [83], Italian [57] and Austrian in this thesis. The core-periphery structure has been shown to be fully determined by the degree sequence of a network [85], and thus not to be representative of the underlying financial motivations that drive interbank lending and borrowing. Moreover, the ability of algorithms to detect core-periphery structures depends on the a priori assumptions underlying these algorithms, especially the level of time aggregation that is used when constructing interbank networks. Barucca and Lillo (2016) show that with most time aggregations the structure of the e-MID market is bipartite rather than resembling a core-periphery structure [21]. The e-MID market is the Italian electronic market for interbank deposits, for a thorough introduction to its mechanism see Iori et al (2015) [77]. See also section 3.3.3 where I outline a methodology to detect core-periphery structures.

Nonetheless, from a financial perspective the core-periphery structure is sensible, as large core banks can offer services to smaller banks (e.g. clearing, lending of large amounts at lower rates, etc.) and generally coordinate the activity of the network as money centre banks. Craig and von Peter (2014) for instance show, that over 90% of transactions of periphery banks involve core banks, yet out of over 2,000 banks in the German interbank network, only 45 are in the core. This demonstrates the extraordinary significance of core banks in interbank activity [44].

Robust-yet-fragile structure of financial networks

Financial networks are characterised by relatively few but major disruptions. In complex systems theory this behaviour is described as robust-yet-fragile [2]. Scale-free networks tend to be robust with regards to random attacks, but are fragile with respect to targeted attacks. When 5% of the links in a scale-free network are randomly removed, the average shortest path length – an indicator of network connectivity, does not significantly reduce [2]. However, when the most connected, i.e. the nodes of highest degree, are disconnected from the giant connected component, then the network connectivity quickly breaks down. The giant connected component is the largest cluster of nodes in a network and strongly contributes the connectivity of a network. In contrast, due to the homogeneous degree distribution of ER networks, the fragmentation of the giant connected component in ER networks occurs only when a much larger fraction of links are removed relative to SF

networks. Lenzu and Tedeschi (2012) explore this behaviour with an agent based model employed on different topologies ranging from ER over SF to star networks [84]. The authors find that the reason for the apparent robustness of SF networks is due to their high abundance of poorly connected nodes. These nodes form isolated clusters that are insufficiently diversified with respect to counterparty default risk, which describes the risk of the lender, that one of its borrowers defaults on its liabilities. A particular reason for this higher risk is that banks in isolated clusters are unable to raise sufficient liquidity to continue their operations [84].

The core-periphery structure contributes to the robust-yet-fragile nature of interbank networks [31]. Because most banks in an interbank network are small, it is likely that a defaulting bank is a small bank in the periphery. The network appears robust with respect to the default of a small bank, since the failure of a small bank has a limited impact on other banks. In contrast, the network is not strong enough to cope with the failure of a bank from the core, which would cause many other banks to default [2, 39, 108]. This description mirrors the effect of leverage on the stability of financial markets, that I described in section 2.1.1: In a highly leveraged financial market a few investors are the major drivers of asset prices, and are therefore the largest sources of potential financial instability.

2.2 Null Models of Networks

In this section I discuss statistical procedures to estimate and generate networks. Statistical estimation of networks has two main purposes. Firstly to estimate unknown data; and secondly to gain insights into the factors and mechanisms that define network topology and dynamics.

The estimation of unknown data is particularly important for studies of systemic risk for three main reasons [97]. Firstly, banks are not required to publicly report all their bilateral connections (loans) to their counterparties. Secondly, interbank networks can evolve extremely quickly. The pace of interbank networks is amongst other reasons due to the maturities of loans, such as in overnight lending markets; or at times because the network is pushed into a different regime as a result of certain endogenous or exogenous events. Thirdly, most transactions are over-the-counter, which are not reported centrally and are therefore difficult to study comprehensively.

Methods that reconstruct networks can address these problems, as these methods are based on public data that is easily accessible and up-to-date. In this way the bilateral connections can be estimated and changes in interbank markets due to endogenous and exogenous factors can be captured in time.

In this section I focus on null models, which generate ensembles of networks in

which certain properties, such as the in- and out-degree of each node are preserved. Statistical estimation is used primarily to estimate networks when no information about bilateral connections exists. After reconstructing a network, analyses are run on this network to compare its topology and dynamics to the underlying real network. In the case of systemic risk studies, financial contagion mechanisms are often simulated on these networks, such as those discussed in section 3.2, to compare the ability of different null models to estimate the risk in a real interbank network.

2.2.1 General insights gained from null models

The null models studied in this thesis are static, i.e. time-invariant, in which individual nodes do not adjust their decisions to a new state of the network, for instance by selling assets to raise equity in a financial crisis. Instead, banks simply absorb the effect of the network dynamics, for instance losses resulting from the default of a counterparty. For a more thorough contrast between static and dynamic models see Montagna and Lux (2013) [97].

Erdős-Rényi (ER) random networks are the most simple type of null model, as they require information only of the average degree of the real network. In the previous section I already described through the robust-yet-fragile concept introduced by Albert et al, that homogeneously distributed ER networks are not as fragile as real interbank networks [2]. In fact, Montagna and Lux (2013) show that ER networks always underestimate the severity of financial contagion [97]. However, due to the homogeneous distribution of degrees and balance sheet sizes, these smaller contagion events happen more frequently in random networks [117]. Further studies of random networks have shown, that in fact contagion can be large when nodes are homogeneous [76]. ER networks therefore require several additional tuning parameters to perform well in recreating the topological features and systemic risk observed in real networks [105]. See Nier et al (2007) [105] and Sachs et al (2011) [117] for two of the first null model studies on random networks.

Analyses of simulations run on null models confirm that disassortative mixing and heterogeneous degree as well as balance sheet size distributions are defining features of interbank networks. These characteristics of networks are therefore a necessary feature for null models to recreate realistic topologies and dynamics [23, 36, 97, 117]. In fact, I discussed the influence of heterogeneous distributions of underlying data on leverage and volatility in financial markets in section 2.1.1 and section 2.1.2.

Anand et al (2017) study 7 different reconstruction methods in a very extensive study based on datasets from 13 jurisdictions with 25 different interbank markets,

that consist of each bilateral exposure between banks [9]. Whilst the study provides no new insights into the capabilities of null models, it provides the most rigorous study to date, that compares an extensive array of null models across a very large number of financial markets. Whilst I refer the reader to the paper for specific results, I note that Anand et al test two different categories of null models. The first category is an iterative null model, which builds upon an initial guess of the network structure and successively rescales the edge weights until the marginals of the matrix, i.e. the total interbank assets and liabilities of each bank, are allocated properly. An example of this approach is the RAS algorithm, which I describe in the methods section 3.3.4. The second category is a Monte Carlo type sampling approach, for instance based on copulas or the fitness model, that I describe below in section 2.2.3. This type of statistical sampling null model allocates edge weights based on a pre-defined probability distribution of bilateral exposures between banks.

2.2.2 Maximum entropy approach

The maximum entropy (ME) approach was the most common information theoretic approach used to estimate bilateral connections [132]. In brief, the maximum entropy method maximises the entropy of matrix L_{ij} that is the directed and weighted adjacency matrix, which represents the liability of bank i towards bank j . Typically the sums of each bank's interbank assets and liabilities are known, $a_i = \sum_j L_{ji}$ and $l_i = \sum_j L_{ij}$, respectively [95, 133, 132]. The ME approach iteratively places edge weights in the estimated matrix L in keeping with the constraints provided by a and l . Formally, the ME approach minimises the Kullback-Leibler divergence between the real adjacency matrix and its estimate L [132]. This minimisation is typically implemented using the RAS algorithm, which I outline in section 3.3.4.

The ME approach creates networks in which the bilateral exposures are spread as evenly as possible amongst all nodes. When each bank i in a network lends $a_i > 0$ and borrows $l_i > 0$, then the ME approach creates complete networks, in which each node is connected to every other node in the network. Such homogeneous networks, which disregard the diversity of bilateral connections, are not applicable to most interbank networks. As I noted in the previous section, network topology is a crucial determinant of network dynamics. ME can therefore not recreate properties of networks such as the borrowing and lending dependence of banks, clustering and assortativity [7]. Moreover, given the even spread of interbank positions across an ME generated network, these networks underestimate financial contagion relative to real networks. Due to the homogeneous distribution of interbank liabilities, the default of an individual bank does not have a strong effect on the system and can therefore

not create the necessary trajectories to cause events in which a large number of banks default, which are a common feature of interbank networks, see discussions in [91, 95, 97, 117, 133, 132].

In response to the inapplicability of the maximum entropy approach Anand et al (2015) introduced a method called Minimum Density (MD) that creates sparse networks [7]. MD assumes that links are costly to establish and maintain in real networks, based on the fact that most real interbank networks have a very small density around 1%. Given the additional cost factor attached to each link, the total interbank volume is allocated using a minimum number of links, whilst still properly allocating the total interbank assets and liabilities of each bank. In addition, a prior is assumed that contributes to the disassortativity of the estimated network, such that small banks link to large banks and vice versa.

MD creates networks with remarkably similar topological characteristics to real networks. Due to the objective function being tuned to the density of the underlying network, its networks have similar degree distributions to the real network that it aims to recreate. Moreover, MD produces disassortative networks, but struggles to estimate the local topology, for instance the clustering coefficient and the abundance of triadic motifs [7].

It is not the explicit goal of the MD method to construct interbank networks with similar systemic risk to a real network, but rather to provide an upper bound on the systemic risk in a given network. The reason for this is that poorly connected networks typically have a higher fragility, see also section 3.2.2. MD based networks have a very low connectivity and therefore fall into the category of fragile networks. As I noted above, ME generated networks underestimate systemic risk, together with the estimates of the MD networks, both methods may therefore provide reasonable lower and upper bounds of the likely systemic risk in a real interbank network. Whilst Anand et al show that this is the case for different levels of systemic risk within the same network, the authors do not show whether this holds for networks of different densities [7].

2.2.3 Fitness model

The fitness model goes further into the statistical mechanics of network estimation than the null models that I have discussed so far. The fitness model is a bootstrapping method that aims to fit a non-topological fitness metric, for instance the interbank activity of each bank, to a topological metric, for instance the degree distribution. A core assumption is that a monotonous relationship exists between the fitness and topological metric. A major advantage of this approach is that the data

required for the fitness values may be more readily available than the desired topological information. In contrast to the preferential attachment model, the fitness model fits all nodes at once rather than growing the network via a specified rule of attachment of nodes [16, 46].

In applications to systemic risk, the fitness model generates disassortative, scale-free and sparse networks, as it attaches a higher linking probability to larger banks, due to their higher interbank activity, i.e. their higher fitness value [38, 46, 61]. See section 3.3.4 where I outline the methodology of the fitness model in detail.

Cimini et al (2015) empirically test the fitness model on the World Trade Network and the Italian e-MID Italian interbank network [41]. It is a unique study in recreating and testing the ability of the fitness model to reconstruct not only reciprocity, DebtRank and shortest path lengths but also percolation thresholds and the size of the giant connected component. Cimini et al point out that errors in the networks created by the fitness model are due to two factors. Firstly, the assumption that the fitness values are proportional to the statistically fitted interbank activity and secondly the partial knowledge of in- and out-degrees assumed in scaling the model. For the second point it should be noted, that fitness model studies often assume knowledge of only a subset of nodes in a network.

Cimini et al find that the topology of the underlying real networks is recreated surprisingly well. The fitness model provided with a subset of only 10% of all nodes, achieves a small measuring error for the DebtRank in the reconstructed networks [41]. However, the measurement error for link reciprocity is high, even when a larger subset of nodes is known. This is explained by the fact that the fitness model can only capture the global topological properties of the real network, whereas link reciprocity is constrained by local rather than global dependencies. Similar results are found by Squartini et al (2016) who find that the fitness model can extrapolate an entire network from a sub-sample as small as 2% of the nodes of a network with an accuracy of 80% for the Italian e-MID market [123]. Other applications of the fitness model to systemic risk include the study of Montagna and Lux (2013), who use the fitness model approach to successfully recreate the core-periphery structure of interbank networks [97]. Musmeci et al (2012) in one of the earlier studies that use the fitness model on the World Trade Network, have similarly good results, with a maximum root mean squared error of 10% between the reconstructed and real network [98].

Chapter 3

Systemic Risk: Disentangling the Influence of Interbank Network Topology and Dynamics

3.1 Introduction

3.1.1 Motivation

It is difficult to obtain complete information on the exposures between financial counterparties, as I discuss in the literature review in section 2.2. Analyses of interbank networks are therefore often conducted on networks that are reconstructed from partial information. This partial information frequently consists of the total interbank assets and liabilities of each bank, which are publicly available. This information as the input for network reconstruction models allows users of interbank networks to conduct research on networks that are based on data that is easily obtainable and up-to-date. It is therefore important for practitioners, policy-makers and researchers to understand the ability of null models to reconstruct networks. Of further importance is an understanding of the ability of null models to estimate the financial contagion properties of real interbank networks, in order to classify one null model as superior to another. In this project I assess the ability of null models to reconstruct network topologies and to estimate financial contagion.

Researchers are rarely in a position to compare different null models based on data of both the complete balance sheets of banks and the networks that describe banks' bilateral exposures. Indeed, the empirical work on financial contagion can be broadly divided into two strands. Firstly, analyses of real interbank networks, see for instance [26, 36, 44, 54, 58]; and secondly, the introduction of models to recreate real network topologies from partial information, see for instance [7, 9, 41, 48, 98, 105].

An exception to this division of the literature is for instance the study of Mistrulli (2011), who analyses the ability of the maximum entropy approach to reconstruct financial contagion [95].

In this project I combine these two strands by comparing the results of an analysis of financial contagion in reconstructed networks with those of a real interbank network. I can therefore compare the ability of null models to reconstruct a real network and assess how well they can estimate the financial contagion of a real network. This approach allows me to link the contribution of the network topology –via the difference in the topology of the reconstructed networks– to financial contagion. My study is more extensive than that of Mistrulli, who analysed only the maximum entropy approach, whereas I analyse three different null models at the same time. Moreover, as I discuss in section 2.2.2, the maximum entropy approach has been shown to generate networks that exhibit topological properties and financial contagion that are very different to those observed in real interbank networks.

3.1.2 Research objective and approach

I obtained data of the Austrian interbank network, that has previously been studied by Caccioli et al (2015) [36] and for an earlier period by Boss et al (2003) [32]. I reconstruct the Austrian network using three null models, which are Erdős-Rényi random networks, the fitness model and the directed configuration (rewired) model. In all cases, after generating a binary adjacency matrix using one of the three models, I employ the RAS algorithm to assign edge weights [13, 23]. The RAS algorithm ensures that the total interbank assets and liabilities of each bank in the reconstructed network match those of the real system.

In this project I compare the topology and dynamics of the reconstructed networks to those of the Austrian network. Specifically, I compare the degree distribution, assortativity and clustering as well as triadic motifs of the real and reconstructed networks. I then perform stress tests on the real and reconstructed networks to assess the ability of the null models to produce networks with systemic risk characteristics that are similar to the real network. I study in particular the absorbing state dynamics and the trajectory profiles of contagion over time via the Furfine counterparty default risk model [58] and the DebtRank algorithm [17].

I then link the insights that I gained on the contagion dynamics to the networks' topological properties. One specific approach that I undertake is to identify a vulnerable cluster of banks as defined by Watts (2002) [135]. The vulnerable cluster is a set of vulnerable banks that default when just one of their borrowers is unable to repay its liabilities. This experiment represents the first study in which vulnerable

clusters are identified in a real interbank network.

3.2 Related Work

In this section I discuss the directly related previous work to the investigations that I undertake in this project. I focus on the part of the literature that investigates the influence of network topology on financial contagion in interbank networks. For a general literature review on systemic risk see section 2.1 and for null models of networks see section 2.2.

3.2.1 Financial contagion

Network theory can offer a deep understanding of financial contagion by disentangling the roles played by the topology of an interbank network and its dynamics. Due to its focus on endogenous factors that can significantly amplify shocks to financial systems, network theory is increasingly being used to gain an understanding of financial contagion [23].

Systemic risk can spread through multiple channels of contagion. The most frequently studied mechanism is counterparty risk, which describes the risk that borrowers are unable to repay their liabilities. In the Furfine algorithm losses propagate to lenders only after the default of a borrower [58]. The algorithm measures counterparty risk in terms the number of defaults of lenders [58], I outline the algorithm in section 3.3.1. If a lender does not have sufficient equity capital to cover the default of its borrowers, the lender itself may default, which can lead to domino effects in the entire network.

A different way of looking at counterparty risk is provided by the DebtRank algorithm. It measures the incremental build-up of stress in terms of the cumulative losses of banks' equity capital, as a result of an exogenous shock to lenders' interbank assets [26]. DebtRank in contrast to the Furfine algorithm therefore measures counterparty risk before the default of a borrower. The DebtRank algorithm was introduced in 2012 by Battiston et al [26]. Bardoscia et al (2015) developed a dynamic DebtRank version, which allows a bank to receive and propagate stress at each iteration of contagion [17]. The authors find that a network of 183 publicly-listed European banks can amplify an initial shock on its equity by a factor of up to 6. I outline the DebtRank algorithm in section 3.3.2.

Other channels of contagion exist that consider systemic risk introduced through mechanisms beyond lending networks. Fire-sales refer to system-wide losses induced through the selling of assets that are held by other banks, which are referred to as

overlapping portfolios [23, 37, 73]. In a study of the US interbank market Duarte et al (2013) were able to show, based on data from 2008 until 2014, that an exogenous shock of 1% to the common assets held by banks, leads to average losses of 8% of the entire equity capital of the system, that are solely due to fire-sales. When combining all channels of financial contagion the authors found that this shock led to a loss of 35% of the equity capital of system [50]. In a similar study, Huang et al (2013) study overlapping portfolio risk in the USA and describe how shocks to commonly held assets exacerbate systemic risk [73].

A further channel of contagion is roll-over risk, which occurs when a bank loses the trust of interbank markets, such that it is unable to refinance (roll-over) its liabilities [8, 36]. Roll-over risk is very similar to liquidity risk, which describes the inability of banks to access the interbank market to obtain funding [3, 40]. When banks have better access to additional liquidity, systemic risk is drastically reduced [3]. Amini et al (2016) provide analytical proofs for different channels of financial contagion, based on the number of highly leveraged exposures that constitute the largest sources of risk [6].

There are moreover many additional indirect channels of contagion that amplify systemic risk [42]. Such channels are for instance the reputational risk that banks are exposed to, for example via media articles describing their standing in interbank markets. These indirect channels can increase counterparty, liquidity, fire-sale and roll-over risks [42]. Uniquely, Glasserman et al (2014) point to the importance of bankruptcy costs and delays in the repayment of liabilities as significant additive factors to counterparty risk [65]. These factors are not commonly considered in the abstract studies of financial contagion. Glasserman et al hypothesise that the negative feedback loop created by fair value (mark-to-market) accounting is larger than all other counterparty risk effects. Mark-to-market accounting reflects the negative expectations of market participants in a crisis and can create negative feedback loops in asset prices.

Recent investigations have begun to study the interactions of these channels of financial contagion [23, 36, 81]. For instance, during the financial crisis banks reduced their connections to other banks, even though these banks did not constitute a significant counterparty risk. Banks decided to disconnect because of the additional risk imposed from overlapping portfolios. Since banks tend to hold the same assets, the default of one bank could lead to fire-sales, which would in turn reduce the equity base of all banks holding these assets [15, 36, 37, 70]. Caccioli et al (2015) show that overlapping portfolio risk can spread throughout the entire network, whereas counterparty risk affects the network only locally via neighbouring nodes. Importantly, when both mechanisms are combined a much larger frequency

of global contagion is observed [36].

Multilayer networks

In recent years studies have emerged that investigate interbank networks as multilayer networks. Each layer of these networks captures a different activity of financial institutions. These different activities are for instance interbank loans of different maturities such as overnight and three months maturities; overlapping portfolios, such as investments in companies listed on stock exchanges; and financial derivatives, for instance to hedge currency and credit risks. The layers are interconnected through the individual institutions that hold assets in the markets that are captured by each layer. The major advantage of multilayer networks is that they allow different types of financial contagion, that are unique to each asset class, to be studied together. The most important insight from these studies is that there are strong non-linear effects between the layers, due to which the total systemic risk of multilayer networks is larger than the sum of the systemic risk of the individual layers [20, 96, 110, 118]. These non-linear effects show that there are strong interactions across the different asset classes that increase systemic risk.

3.2.2 Interbank network topology and financial contagion

In this section I review the literature that addresses the influence of interbank network topology on financial contagion. I focus on the roles played by the local topology, such as triadic motifs; and global topological properties, such as diversification.

Impact of local topology on systemic risk

Squartini et al have shown in 2013 that the nature of interactions between Dutch banks significantly changed in the period leading up to the financial crisis in 2008 [125]. Before the crisis the Dutch interbank network was close to an equilibrium network structure, as banks arranged themselves in a way that maximises the lending and borrowing activities in their network [125]. As the financial crisis began to emerge, the Dutch interbank network moved away from this equilibrium topology. Financial markets became less stable as stable banks stopped interacting with riskier banks, which as a result did not have sufficient connections to maintain their liquidity [125].

The local topology of an interbank network describes the specific constellations of bilateral exposures between a small subset of banks. Triads are typically used to deconstruct the local topology of networks. Triads are clusters that involve three nodes, which can create 13 different motifs of connections, when taking into account

the in-, out- and reciprocated links [94]. I thoroughly define triads and outline an identification methodology in section 3.3.3. Research has shown that stronger reciprocity in triads significantly reduces financial risk, due to a more homogeneous spread of liabilities across the network [124, 125]. Reciprocity in networks describes two nodes that have both in- and outgoing links to one another. In contrast, in an unreciprocated triadic loop the lender is unaware of its borrowers' interbank positions and is indeed unaware that this risk loops back to itself. The lender therefore underestimates the risk resulting from a loan extended to a borrower, which results in the lender setting up insufficient equity capital buffers.

Squartini et al found that the abundance of unreciprocated triads was abnormally high in the period from 2000 until 2004 [125]. However, as the financial crisis began to emerge in 2008 this abundance was abnormally low, because banks reduced their exposures to risky counterparties. In addition to banks being prone to underestimate their counterparty exposures in the pre-crisis phase due to the high abundance of unreciprocated triads, the number of reciprocated dyads in the Dutch market significantly decreased in the pre-crisis period, leading to a further underestimation of risk. Unreciprocated dyads, as Squartini et al note, are the most risky topological structures from a perspective of financial stability, because they can become part of unreciprocated triads. This change is analogous to the explanations that I have given in section 2.1.1 of the leverage cycle: When margins increase, investments become more risky and less profitable, in response banks reduce their exposures to the most risky and vulnerable banks first and thereby cut the periphery off from the core. See also section 2.1.2 where I discuss the core-periphery structure of interbank networks. Similar insights to those of Squartini et al in the Dutch interbank market were also shown in a multilayer network of the Italian interbank market [20].

Bardoscia et al (2016) study the impact of the connectivity of an interbank network upon its stability [18]. With an increasing network density the number of loops in a network increases, which in turn increases the instability of the network as measured by the DebtRank algorithm. The instability of the network is measured by λ_{max} , the largest eigenvalue of the matrix of interbank leverages, see section 3.3.2 where I outline the methodology. Each entry in the matrix of interbank leverages specifies the leverage of each exposure. Interbank leverage is typically expressed as the amount of money lent by a bank relative to its equity capital [18, 24]. Bardoscia et al show that when $\lambda_{max} > 1$ the system is unstable, such that losses infinitely compound.

Impact of network diversification on systemic risk

An increasing number of studies investigate the amplification of risk that results from deeper integration of interbank markets [127]. Many studies investigate the impact of the connectivity of interbank networks on systemic risk, but may not explicitly address the issue of network diversification. I here synthesise the results across studies that directly or indirectly touch upon network diversification. This review of the literature provides an important background to the investigations on network connectivity that I undertake in this and the following projects.

Vitali et al (2016) study the influence of the connectivity between different world regions across interbank markets with an agent based model [134]. The authors find that a system is more resilient through diversification when the pace of financial activities is constant, when this pace is accelerating however, diversification leads to increased levels of systemic risk. The reason for this is that, as described in section 2.1.1, increases in financial activity may result from larger leverages that can become sources of risk. In the model of Vitali et al the size of banks changes over time, whereas the network structure is constant. Vitali et al thus from a different perspective confirm the results of Stiglitz (2010) [127] and Battiston et al (2012) [25], who instead study changing network structures and also find increased risks due to deeper market integration.

Nier et al (2007) assume a simple set-up of purely random graphs with equal asset sizes for all banks in the network. They find for low levels of network connectivity, that an increase in connectivity reduces the stability of the network, as it amplifies the transmission of shocks. For higher levels of connectivity, increases in connectivity stabilise networks, since shocks can be spread out to many other nodes [105]. Similarly, in a seminal study Allen and Gale (2000) show that a monotonic relationship exists between the average degree of a network and its financial stability. Full connectivity between banks reduces risk because of banks are fully diversified. As the average degree decreases, the diversification reduces, which results in larger levels of systemic risk [4].

A similar insight is provided by Tedeschi et al (2012) [130] and Grilli et al (2014) [68], who by studying agent based models show that larger heterogeneities in interbank networks lead to larger avalanches of bank failures. The reason for the increased avalanches is that the exposure of lenders to the network is highly varied, such that the failure of a large bank can lead to large avalanches of failures. These results were confirmed in a follow-up study by the authors in 2017 [28].

Battiston et al (2012) combine fire-sales and counterparty risk to discover that increasing diversification in networks initially makes networks more robust, but show that once systemic risk has increased beyond a certain level, diversification decreases

the stability of a network [25]. This behaviour occurs because after some banks have defaulted, the probability increased that a defaulted bank holds assets of the portfolio of another bank. Interestingly, Roukny et al (2013) show that network topology only matters to the stability of a network when markets are illiquid. When markets dry up, fire-sales result that are amplified in a scale-free network due to their robust-yet-fragile nature [114]. Caccioli et al (2012) show that contagion dynamics are not only defined by heterogeneous distributions of degree and balance sheet sizes, but also through the mixing patterns (degree correlations) of nodes [35].

In summary, this review of studies on the impact of market integration on systemic risk shows that it is very difficult to judge whether diversification in interbank networks increases or decreases financial stability. All studies are however in agreement that high leverages in poorly connected networks increase the probability of very large avalanches of bank failures.

3.2.3 Global Cascade model

Epidemic and cascading models provide additional methods to study network dynamics [103, 107, 135]. Watts' Global Cascade model is a threshold model that describes the conditions for and characteristics of contagion in complex systems [135]. Watts assumes that network dynamics depend on the actions of the nodes in a network, which can be in either of two states: active or inactive, such as being in default.

A node can only spread its state in the network if at least one of its neighbours is vulnerable with a threshold larger than $1/k$, where k is the degree of the neighbour. A cascade is defined as an event of at least one node defaulting as a result of the default of one vulnerable node. Typically a value of 5% of nodes defaulting is chosen in numerical analyses, I also choose this value in my analysis in appendix A.2.2. When a cluster of vulnerable banks exists that is larger than the average degree $\langle k \rangle$ of the network, global cascades will occur [135]. If the clusters are smaller than $\langle k \rangle$, then the network is too fragmented for a global cascade to occur. An important insight provided by Watts is that the frequency of global cascades is approximated by the size of largest cluster of vulnerable and immediately adjacent nodes.

Adaptation to interbank networks

Hurd (2016) [74, p. 117ff] as well as Gai and Kapadia (2010) [59] adapt the Global Cascade model to directed interbank networks. The size and probability of global cascades in Hurd's adaptation are very similar to Watts' experimental results [74, p. 132]. For a similar study see Amini et al (2012 and 2016), who study the stylised

facts of interbank networks using cascading models and statistically quantify the notion of contagious links based on the exposures of banks [5, 6].

The applicability of Watts' (and Hurd's) analytical model to real-world networks is limited because the model assumes that loops are absent in the generating function of networks, which is also referred to as the locally tree-like independence property [93]. Interbank networks, as described in the previous sections, are characterised by more frequent local clustering than expected from a random null model. In fact, even in ER networks loops exist, even though they may be very long when networks become large in the thermodynamic limit. It can therefore be questioned in how far Watts' and by extension Hurd's model can approximate the size of vulnerable clusters in interbank networks whose size does not reach the thermodynamic limit. I discuss this question in section 3.6.

3.3 Methods

3.3.1 Counterparty default risk - Furfine model

The counterparty default risk algorithm is based on the work by Furfine [58]. The dynamics are based on an interbank network that is represented by a weighted directed adjacency matrix L_{ij} , in which each entry represents the borrowing amount of bank i (the borrower) from bank j (the lender).

Each bank is associated with a balance sheet that consists of assets and liabilities including interbank assets and liabilities, as well as equity capital. A lender j is assumed to have defaulted when its borrower i defaults on its liabilities towards bank j and this liability is larger than the equity of bank j , formally: $L_{ij} > E_j$. Where E_j is the equity of the lender j , which is the difference between total assets and total liabilities. A lender may also default if the liability of a defaulted individual borrower is smaller than the lender's equity. This may occur when multiple borrowers default such that the resulting cumulative losses exceed the equity capital of the lender.

For each bank i a state variable indicates whether it is active $\theta_i = 0$ or in default $\theta_i = 1$. The following steps are used to measure the cascade resulting from the failure of a bank, which is referred to as the seed:

1. A seed node q is selected between 1 and n , the number of nodes in the network described by the liability matrix L . Set $\theta_q = 1$ and $\theta_i = 0$ for all other nodes.
2. Banks update their balance sheets. For each lender j set $\theta_j = 1$ if the following holds

$$E_j < \sum_i^n L_{ij}\theta_i, \quad (3.1)$$

presuming zero recovery from a failed bank.

3. If new lenders failed in step 2, return to step 2; otherwise exit.

I define a contagious event as an event in which the above algorithm results in at least one failure beyond the seed. I use the following three measures to describe counterparty default risk:

- Contagion probability: The probability of observing a contagion event, which indicates the robustness of the network. It is computed as the probability of observing a contagion event in the network across all seeds.
- Conditional extent of contagion: The average number of banks that failed when a contagion event occurs, averaged across all seeds.
- Maximum extent of contagion: The largest observed contagion event across all seeds.

3.3.2 Interbank stress propagation - DebtRank

The DebtRank algorithm measures the total relative fraction of equity lost following an initial shock on the equity of some or all the banks in an interbank network L . Shocks propagate via the leverage matrix $\Lambda_{ij} = \frac{L_{ji}}{E_i}$, such that the interbank assets of lender i towards its borrowers are expressed as a fraction of the lender's equity. This shift of indices in Λ_{ij} with respect to L_{ji} ensures that shocks propagate from the borrower to the lender.

The state of a lender i across time $h_i(t)$ measures the relative fraction of equity lost. At the first time step $t = 1$ a shock $0 < S_i < 1$ is applied to all or some of the lenders' (seeds') equity $h_i(1) = S_i$. The state of all affected banks is updated according to the following rule:

$$\Delta h(t+1) = \Lambda \Delta h(t), \quad (3.2)$$

where $\Delta h(t) = h(t+1) - h(t)$, which measures the losses that occurred between two time steps. The algorithm converges at time step T when the lenders experience no additional stress. The total DebtRank R_i of each lender i is computed as the sum of the relative fraction of equity lost excluding the initial shock S_i at $t = 1$:

$$R_i = \sum_{t=2}^T \Delta h_i(t). \quad (3.3)$$

The stability of the leverage matrix can be expressed with its largest eigenvalue λ_{\max} . When $\lambda_{\max} < 1$ the system is stable such that it will eventually settle to $h(t) = 0$, when the system is not stable $\lambda_{\max} > 1$, then the system will not settle such that losses will infinitely compound [17].

3.3.3 Network topology characteristics

This section introduces the most important metrics that I will use to describe the topologies of the networks that I study.

Degree The degree k_i of node i is the number of other nodes it is connected to in the network described by the binary adjacency matrix A with n number of nodes. For directed matrices a distinction is made between the in- k_i^{in} and out-degree k_i^{out} of node i .

$$k_i^{in} = \sum_{j=1}^n A_{ij}, \quad k_i^{out} = \sum_{j=1}^n A_{ji}. \quad (3.4)$$

The average in- $\langle k^{in} \rangle$ and out- $\langle k^{out} \rangle$ degree of a network are always equal, because the number of edges e in a network are equal to the sum of all in- (or out-) degrees, see Newman (2010) [103, p. 135]:

$$e = \sum_{j=1}^n k_j^{in} = \sum_{j=1}^n k_j^{out}, \quad (3.5)$$

$$\langle k^{in} \rangle = \frac{1}{n} \sum_{j=1}^n k_j^{in} = \frac{1}{n} \sum_{j=1}^n k_j^{out} = \langle k^{out} \rangle. \quad (3.6)$$

The total degree of a node is the sum of its in- and out-degree, and the average degree of the network $\langle k \rangle$ is the average over all nodes' total degree. $\langle k \rangle$ can also be computed as

$$\langle k \rangle = \frac{2e}{n}. \quad (3.7)$$

Density The density is a measure of what fraction of links between nodes that could be present, are actually present in the network:

$$d = \frac{e}{n(n-1)}, \quad (3.8)$$

where $n - 1$ is used as nodes are not allowed to have self-loops.

Clustering The clustering coefficient is a ratio that indicates how many connections between neighbouring nodes exist [103]. The clustering coefficient for node i

is computed as follows:

$$C_i = \frac{\text{number of pairs of neighbours of } i \text{ that are connected}}{\text{number of pairs of neighbours of } i}. \quad (3.9)$$

Mixing The mixing pattern of nodes in the binary adjacency matrix A describes the manner in which nodes link to one another. Assortativity refers to the tendency of similar nodes to link together, for instance high-degree nodes connecting to other high-degree nodes, disassortativity describes the opposite behaviour [101, 100].

The mixing pattern score is typically expressed in a range from -1 (perfectly disassortative) over 0 (neutral mixing) to 1 (perfectly assortative) and is computed as follows:

$$r = \frac{\sum_{ij}(A_{ij} - k_i k_j / 2n)x_i x_j}{\sum_{ij}(k_i \delta_{ij} - k_i k_j / 2n)x_i x_j}, \quad (3.10)$$

where x_i and x_j are the measures of interest of nodes i and j , in this case their degrees k_i and k_j , and δ is the Kronecker delta. The mixing pattern score is the Pearson correlation coefficient adapted to networks [103, p. 229].

Core-periphery structure

I use the core-periphery detection algorithm introduced by Lip [85]. In a core-periphery structure banks are divided into sets of core and periphery banks. An idealised core-periphery structure is a network in which all core banks are connected with one another; and the periphery banks have only links to the core and no links among themselves. In a network of size n the Lip algorithm identifies c the number of core banks, that minimizes the following error function:

$$e(c) = \frac{c(c-1)}{2} - \frac{\sum_{i=n-c+1}^n k_i}{2} + \frac{\sum_{i=1}^{n-c} k_i}{2}, \quad (3.11)$$

where nodes are sorted in ascending order of their total degree k_i . The first term in the above equation represents the maximum number of links in the core for a given number of c core banks, the second term reflects the number of actual links in the core, and the third term the number of links between peripheral nodes.

Lip's definition of core-periphery structures does not exactly correspond to the definition given by Craig and von Peter (2014), who describe another commonly used algorithm [44]. The authors add the requirement that each core bank must act as an intermediary that provides interbank lending between peripheral banks that do not interact among themselves. The final results of both core-periphery algorithms have been shown by Fricke et al (2014) to be very similar in practice in the Italian e-MID interbank market [57]. For a more detailed discussion of the

literature on core-periphery structures see my literature review in section 2.1.2.

Triadic Motifs

A triad is a constellation of three nodes [103, p. 199]. When taking into account the in-, out- and reciprocated links among three nodes, 13 different motifs are possible, as described by Milo et al in 2002 [94]. See fig. 3.1 for an overview of the motifs that I study in this project and section 3.2.2 for a literature review.

I measure the number of times t each of the following triadic motifs appears in the binary adjacency matrix A , which I chose for their different impact on financial stability that is reported by Squartini et al (2013) [125]:

- *Motif 1* Safe motif, in which one node has outgoing links to two others. These two vertices are not connected to each other:

$$t_1 = \sum_{i \neq j \neq k}^n (1 - A_{ij})A_{ji}A_{jk}(1 - A_{kj})(1 - A_{ik})(1 - A_{ki}); \quad (3.12)$$

- *Motif 5* Neutral risk motif, which is identical to triad motif 1 but for a connection between the two other vertices:

$$t_5 = \sum_{i \neq j \neq k}^n (1 - A_{ij})A_{ji}A_{jk}(1 - A_{kj})A_{ik}(1 - A_{ki}); \quad (3.13)$$

- *Motif 9* Highest risk motif, in which each of the three nodes is connected to one another node with a link in the same direction, forming a directed loop:

$$t_9 = \sum_{i \neq j \neq k}^n (1 - A_{ij})A_{ji}(1 - A_{jk})A_{kj}A_{ik}(1 - A_{ki}); \quad (3.14)$$

- *Motif 13* Less risky motif, as all three vertices are connected with reciprocated links:

$$t_{13} = \sum_{i \neq j \neq k}^n A_{ij}A_{ji}A_{jk}A_{kj}A_{ik}A_{ki}. \quad (3.15)$$

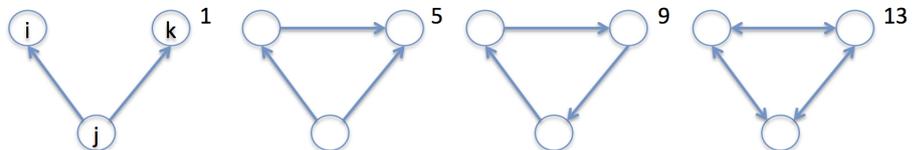


Figure 3.1: Illustration of the triadic motifs studied in this chapter.

Since each node can form a unique triad with two different nodes, the number of times t a motif occurs may be higher than the number of nodes n in the network. For instance motif 5 could occur six times despite it consisting of the same nodes, because the links point in the opposite direction in the other occurrences. In my analysis I am interested in assessing the financial stability of individual nodes that are members of a motif. I therefore identify and analyse the unique members of the individual motifs.

3.3.4 Null model networks

In this section I describe the three null models that I use to reconstruct the real network.

Erdős-Rényi networks

The simplest kind of random graph was introduced by Erdős and Rényi (ER) in 1959 [51]. For each possible pair of nodes a link is placed between two nodes with probability

$$p = c/(n - 1), \quad (3.16)$$

where c is the desired average degree of the graph. Equation 3.16 is easily proven since the average degree of a graph is $\langle k \rangle = 2e/n$ as shown in (3.7), the expected average degree of a random graph is

$$c = (n - 1)p. \quad (3.17)$$

The graph $G(n, p)$ occurs with probability

$$P(G) = p^e(1 - p)^{\binom{n}{2} - e}, \quad (3.18)$$

which is the binomial distribution [103, p. 400].

Random rewiring of directed networks

The directed configuration model (DCM) randomly reassigns (rewires) the directed edges between two nodes to another set of two nodes [48, 90]. This ensures that the degree distribution, including the in- and out-degrees of each node, is maintained, as it can be seen in fig. 3.2.

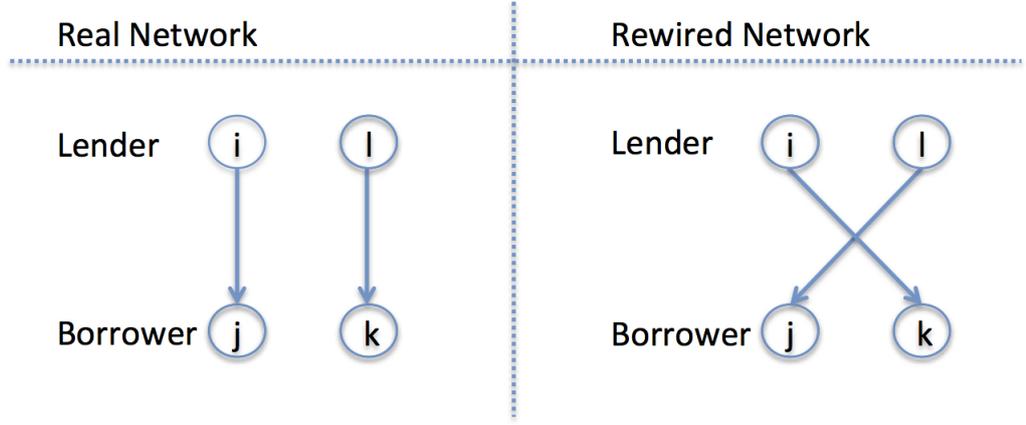


Figure 3.2: Rewiring steps in the directed configuration (rewired) model.

Fitness model networks

I use the fitness model as described by Musmeci et al (2012) [98] following the earlier introduction of the method by Caldarelli et al (2002) [38] and Garlaschelli (2004) [61]. Each node is assigned two fitness values, which determine the likelihood of this node to connect with other nodes. The essence of the fitness model is that the fitness values are assumed to be proportional to the topological metrics that are to be estimated. In my case the fitness values are the lending y_i^a and borrowing y_i^l propensities of bank i , that are assumed to be proportional to the in- and out-degree of bank i . Beyond these constraints the fitness model ensures that the generated ensemble of networks is maximally random.

The fitness model approach requires its users to have an understanding of the theoretical reason underlying the assumption that the fitness metrics are proportional to the topological metrics. The understanding that the degree of a bank scales with its interbank activity is well understood in the literature, and has in fact been proven in many studies [23, 38, 44, 61, 98]. Large banks in terms of their interbank activity are those banks that are the most interconnected in the interbank network. These banks therefore tend to have large in- and out-degrees, which correspond to their large interbank borrowing and lending. This applies in reverse to small banks, which borrow and lend less relative to large banks, and have correspondingly smaller in- and out-degrees in the interbank network.

The fitness values for bank i are calculated as follows:

$$y_i^l = \frac{l_i}{\sum_{j=1}^n l_j} \quad (3.19)$$

and

$$y_i^a = \frac{a_i}{\sum_{j=1}^n a_j}, \quad (3.20)$$

where l and a are vectors of length n representing the total interbank liabilities and assets of each bank i , respectively.

Next, the probability that node i borrows from node j is given by:

$$p_{ij} = \frac{zy_i^a y_j^l}{1 + zy_i^a y_j^l}, \quad (3.21)$$

where z is the only free parameter of the model. z is estimated by minimising the following function:

$$\min \left[\frac{1}{2} \sum_{i=1}^n \sum_{j \neq i}^n (p_{ij} - k_{ij})^2 \right], \quad (3.22)$$

since half over the sum of p_{ij} must equal e the number of edges by the definition given in (3.7).

RAS algorithm

After generating the binary directed adjacency matrix A via one of the three null models described above, I ensure that its diagonal is zero $A_{ii} = 0$ to avoid self-loops. Following this, I assign edge weights to A with the RAS algorithm.

The deterministic RAS or iterative proportional fitting (IPF) algorithm constructs the weighted directed adjacency matrix β of size n such that each entry β_{ij} corresponds to the relative weight of each exposure to the total interbank volume [13, 14, 23]. I firstly set $\beta = A$. I assign the relative edge weights in two steps:

1. Even Step: Update β with relative borrowing propensity of bank i adjusted by y_i^l , its relative liabilities:

$$\beta_{ij} = \frac{\beta_{ij}}{\sum_{k=1}^n \beta_{ik}} y_i^l. \quad (3.23)$$

2. Odd Step: Update β with the relative lending propensity of bank i adjusted by y_i^a , its relative assets:

$$\beta_{ji} = \frac{\beta_{ji}}{\sum_{k=1}^n \beta_{ki}} y_i^a. \quad (3.24)$$

This algorithm is run until convergence, which is the point at which the L_1 -error is equal to zero [112]. The L_1 -error is computed as follows:

$$e_{ij} = |\beta_{ij} - y_i^l| + |\beta_{ji} - y_i^a|. \quad (3.25)$$

The mathematical literature on the RAS algorithm assumes complete matrices, in which all elements are non-zero. It follows from this assumption that the matrix

has non-zero row and column marginals, because as $t \rightarrow \infty$ an element β_{ij} would vanish in a row or column with a marginal value of zero [63]. t is the iteration of the RAS algorithm. This assumption is unlikely to be met in sparse networks, such as interbank networks. I address the problem of zero elements by ensuring that no matrix element of a zero row or column marginal is of a non-zero value. To achieve this I remove all edges in A that are in rows or columns with zero marginals. I then address the problem of vanishing row and column marginals by running the algorithm only on rows and columns whose marginals are non-zero.¹

Moreover, with real world data it is very rare to achieve perfect convergence of the error function (3.25) to zero. I therefore break the L_1 -error function into two parts, that represent the odd and even step. Following this split, the threshold for the update from the odd step to be accepted depends on the preceding step, i.e. the even step and vice versa [63]. For the odd step the upper-bound is defined as follows:

$$e_j \leq \sum_i y_i^l \log \max_{j:\beta_{ij}>0} \frac{\beta_{ij}(t+3)}{\beta_{ij}(t+1)}, \quad (3.26)$$

i.e. the upper bound is the largest change of any value in β_{ij} after one iteration of both steps. In all 10,000 simulations that I undertook, the algorithm converged after 2 full iterations with an error smaller than 10^{-17} in each step. As a reference the average edge weight in β is 7×10^{-5} and the minimum edge weight is 8×10^{-8} .

Finally, I multiply the estimated matrix β by the total interbank volume $\sum_i a_i$ to assign the actual edge weights.

3.4 Analysis of the Austrian Interbank Network

In this section I provide an overview of the Austrian interbank network. I describe its basic topology and financial contagion and conclude the section by discussing the connection between its local topology and financial contagion.

3.4.1 Description of interbank data

The exposures in the dataset that I analyse are encoded into a matrix of interbank liabilities L , whose element L_{ij} represents the liabilities of bank i towards bank j .² The dataset in the first quarter of 2006 consists of $n = 846$ banks (nodes) with

¹Because of this I increase p for Erdős-Rényi graphs by the relative proportion of banks that have no in- or out-going links, as this is the rough proportion of edges that will be set to 0 by the RAS algorithm (this affects around 300 links in the Austrian network in the first quarter of 2006, which has $e = 14,519$ edges in total), specifically: $p = \frac{e}{(n*(n-1))-(n*s)}$ where s is the number of banks with $k^{out} = 0$ or $k^{in} = 0$.

²These exposures exclude short-term lending of less than one month.

$e = 14,519$ interbank loans (edges). In addition to the information on banks' mutual exposures, I have data of the total assets and liabilities as well as the equity capital of each bank. Figure 3.3 below shows the distribution of the equity capital, interbank assets and liabilities for each bank in the first quarter of 2006. The distributions look very similar for the other quarters. The average equity capital is \approx EUR81m and the average interbank assets and liabilities are \approx EUR169m. The average of interbank assets and liabilities are identical, because they consist entirely of bilateral lending relationships within the network.

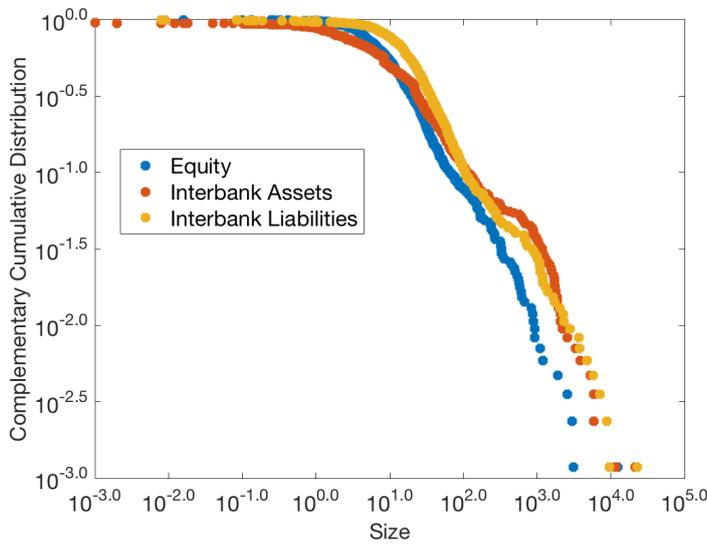


Figure 3.3: Distribution of equity capital, interbank assets and liabilities of each bank in the first quarter of 2006.

When I discuss the topology of the Austrian interbank network I refer to my analysis of the first quarter of 2006 only, whereas for my analysis of financial contagion I refer to all 12 quarters from 2006 until 2008. The topology of the Austrian network is relatively stable across the quarters and the topology of the network based on data of the first quarter of 2006 is representative of all other quarters under investigation. I show that the results are very similar across the quarters in appendix A.1.1, where I present the topological results for the first quarter of 2007 and 2008.

3.4.2 Basic overview of network topology

The network is relatively sparse with a density of 2% and an average degree $\langle k \rangle$ of 34.3. The average in- and out-degree of the network is about 17, and the distributions of in- and out-degrees of banks are heavy-tailed, as it can be seen in the black line in fig. 3.4. This shows that there are many banks with relatively low degree and

a few hubs with a very large degree. The mixing pattern score of -0.59 indicates that the network is disassortative, as small banks tend to connect to large banks, which is a typical behaviour observed in interbank networks [132, 44, 75]. The clustering coefficient, averaged over all nodes, at 0.37 is much larger than expected in a random graph with the same average degree, in the ER networks for instance the average clustering coefficient is 0.006 . 72 out of the 846 banks in the network are in the core. This indicates that less than 10% of banks, the core banks, account for the majority of the interbank activity in the Austrian network. A summary of the topological analysis can be found in table 3.1.

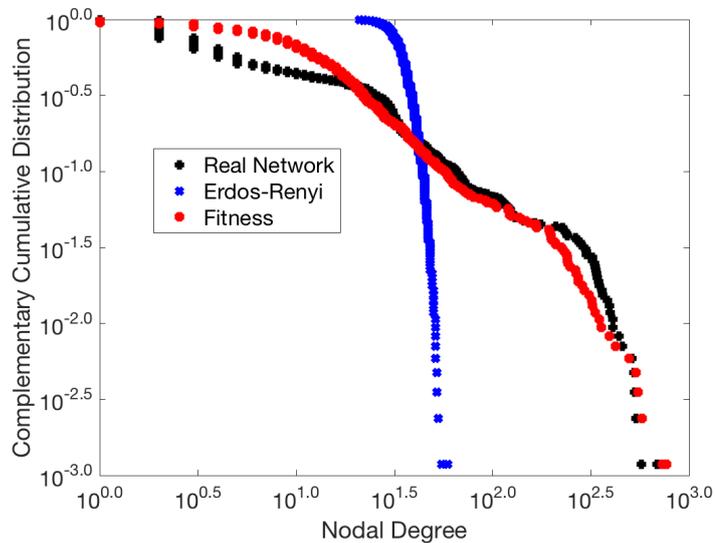


Figure 3.4: Degree distribution of the real network (black), Erdős-Rényi (blue) and fitness model (red) networks. The distribution corresponding to the configuration model is not shown because of its perfect overlap with the real network. Whilst the Erdős-Rényi network has a substantially different degree distribution to the real network, the fitness model closely matches the real network’s heterogeneity. The degree distribution therefore appears to be induced by the heterogeneity of balance sheet sizes, which are the inputs of the fitness model.

3.4.3 Counterparty risk

Counterparty default risk is the risk of a bank i , the borrower, defaulting on its liabilities towards bank j , the lender. Following the computations outlined in section 3.3.1, I find across the 12 quarters in 2006 to 2008 and across all nodes in the Austrian network, that the probability of contagion ranges from 5% to 14%, the average contagion extent from 19 to 91 and the maximum size of a cascade observed ranges from 157 to 343, see fig. 3.6 in the next section for a graphical illustration of these measures. It is interesting to note that at the time of the onset of the financial crisis, towards the end of 2007 and beginning of 2008, the contagion extent increases. This confirms that systemic risk increased during the financial crisis.

Counterparty risk can also be measured with the DebtRank algorithm. I shocked each bank by imposing a 1% loss on its equity and measured the total resulting DebtRank. Following this procedure, I found that the initial value of this shock is on average across all banks amplified by a factor between 1.4 and 8.6 by the interactions of the banks across the 12 quarters. See fig. 3.6d in the next section for a depiction of the DebtRank amplification. As with the contagion extent, the DebtRank increases at the onset of the financial crisis. The amplification that I observe is at times higher than the amplification across European listed banks of up to 6 times measured by Bardoscia et al (2015) [17].

In agreement with my literature review in section 2.1.2, the Austrian network exhibits a robust-yet-fragile topology. The network is robust against failures of a random bank as indicated by the relatively small contagion probability. It is however fragile should a systematically important bank default. This fragility is indicated by the relatively large maximum contagion extent, which shows that in the largest measured event of contagion 40% of the network defaulted. Speaking from a financial perspective the apparent robustness of banks is to be expected, as lenders anticipate a small fraction of their counterparties to default. Lenders in all developed countries are required by law to hold a minimum equity capital buffer to absorb potential losses, for further information see the Basel III rules [22].

The contribution of the local topology to systemic risk

To better understand the counterparty default risk in the Austrian network, I analyse its local topology through the contagiousness of banks that are members of specific triadic motifs. See fig. 3.1 for a depiction of the motifs that I study in this chapter.

To understand the contagiousness of the various triadic motifs I identified the unique individual banks that are its members. The number of unique banks is computed such that each bank is counted only once. I report the unique number of banks that are members of each motif in table 3.1. I then calculated the average extent of contagion of these subsets of banks for each motif in the Austrian network based on the first quarter of 2006. In the literature review in section 3.2.2, I described that motif 9 is the riskiest motif from a financial point of view. This is confirmed in the Austrian interbank network, as the 105 unique banks that are part of this motif have an average extent of contagion of 14.2, which is the largest contagion extent of the four motifs under investigation. The members of motif 5 have an average contagion extent of 6.9 and the banks in motif 13 have an average extent of 10.4. In comparison the "safe" triad type 1, is formed by 416 unique banks, which have the smallest average extent of contagion at 6.7.

3.5 Analysis of Reconstructed Networks

In this section I outline the performance of the Erdős-Rényi (ER), the directed configuration (rewired) and fitness null models as defined in section 3.3.4, in reconstructing the Austrian network. Using each null model I created 10,000 network reconstructions for each quarter from 2006 until 2008, on which I employed the RAS algorithm to assign edge weights. I then performed the same analysis presented for the real network in the previous section 3.4 on these reconstructed networks.

3.5.1 Network topological properties

The central question in comparing reconstructed and real networks is to what extent reconstructed networks are capable of recreating the topology and dynamics of real networks. The set-up of the three null models guarantees that the reconstructed networks have the same average degree as the real network. Significant differences emerge when considering other topological properties. I report a summary of the topological properties for each network in table 3.1. From the table it can be seen that Erdős-Rényi networks are very different from the real network, while the other two models can reproduce some properties of the real network. The strong difference of the Erdős-Rényi random networks to the real network is not surprising, since the degree distributions of the two networks have different properties, as it can be seen in fig. 3.4 in the previous section.

Both the fitness and configuration models display assortativity, clustering, and core-periphery patterns similar to the real network. The fact that these properties are preserved when the links of the real network are randomly rewired suggests that these properties are induced by the degree distribution of the network. In other words, given the sequence of in- and out-degrees observed in the real network, it is not possible to match links that would produce a network with assortativity, clustering, and core-periphery structure much different from the real network. This means that these values may not be associated with the decision-making process of banks underlying the formation of the network and implies a very constrained degree sequence in the real network. This result is surprising because the rewiring process destroys degree correlations. On the other hand, for the triadic motifs I observe strong deviations between the real and the reconstructed networks, as it can be seen in table 3.1 for the first quarter of 2006. In appendix A.1.1 I present the topological results for the first quarter of 2007 and 2008, to show that the results are very similar across the quarters.

Metric		Rewired	Erdős-Rényi	Fitness Model	Real Network
Basic	Nodes (n)	846	846	846	846
	Edges (e)	14, 519 (14, 519 14, 519)	14, 468 (14, 316 14, 549)	14, 519 (14, 397 14, 642)	14, 519
	Density (d)	0.0203 (0.0203 0.0203)	0.0202 (0.0200 0.0205)	0.0203 (0.0201 0.0205)	0.0203
	Degree ($\langle k \rangle$)	34.3239 (34.3239 34.3239)	34.2041 (33.8440 34.5674)	34.3240 (34.0355 34.6147)	34.3239
	Assortativity	-0.58 (-0.58 -0.58)	0.0 (0.00 0.00)	-0.40 (-0.40 -0.40)	-0.5946
	Global Clustering	0.39 (0.38 0.39)	0.0058 (0.0056 0.0059)	0.30 (0.29 0.31)	0.3672
	Core Size	72 (72 72)	45 (44 45)	68 (67 69)	72
	Core-Periphery Error	0.0495 (0.0495 0.0495)	0.9180 (0.9086 0.9274)	0.0741 (0.0701 0.0782)	0.0495
	Cosine Similarity	0.713 (0.705 0.722)	0.0454 (0.0073 0.0711)	0.6613 (0.6564 0.6665)	NA
	RMSD	0.072 (0.072 0.073)	0.1597 (0.1379 0.1859)	0.0787 (0.0783 0.0790)	NA
Triads	# of nodes in Motif 1	764 (753 774)	845 (845 845)	771 (764 779)	416
	# of nodes in Motif 5	499 (487 510)	844 (843 845)	766 (758 774)	406
	# of nodes in Motif 9	138 (127 149)	796 (792 800)	500 (485 515)	105
	# of nodes in Motif 13	174 (170 178)	0 (0 0)	287 (276 298)	187

Table 3.1: Summary statistics of topological properties of the real and reconstructed networks. For the reconstructed networks the metrics are averaged over 10,000 simulations and rounded to the nearest significant integer that can be described with confidence. The figures in the brackets represent the 10% and 90% quantiles, respectively. Results presented refer to the first quarter of 2006.

The similar values of assortativity, clustering and core-periphery patterns in the real and fitness model networks suggest, that the distribution of interbank lending volumes has a significant impact on the interbank network structure. In fact, in the fitness model the probability of two banks being connected solely depends on their total interbank assets and liabilities. This can clearly be observed in fig. 3.4, where I plot the degree distribution of the real and the fitness model networks.

In fig. 3.5 I plot the distributions of edge weights for the real and reconstructed networks. As for the degree distribution, the strongest deviation is observed between the real and Erdős-Rényi networks, while the distributions observed for the fitness and rewired networks are closer to the real network. Erdős-Rényi networks in particular have a much larger number of small edges and a heavier tail relative to the real network. This is due to the fact that in ER networks nodes have more similar degrees than in the real network. Banks that are hubs in the real network now have a much smaller degree in ER networks, which forces them to allocate their interbank assets and liabilities across fewer counterparties. This leads to the over-expression of large weights with respect to the real network. In contrast, as it can be seen in fig. 3.5 the configuration and fitness model networks have a smaller tail than the real network but a higher abundance of moderately sized edges. The RAS algorithm is used to assign edge weights to all three null model networks. Since the RAS algorithm attempts to maximise the entropy of a network, it creates an edge weight distribution that is less heterogeneous than that of the real network [7]. Despite the fact that the RAS algorithm is provided with the heterogeneous fitness values as marginal constraints.

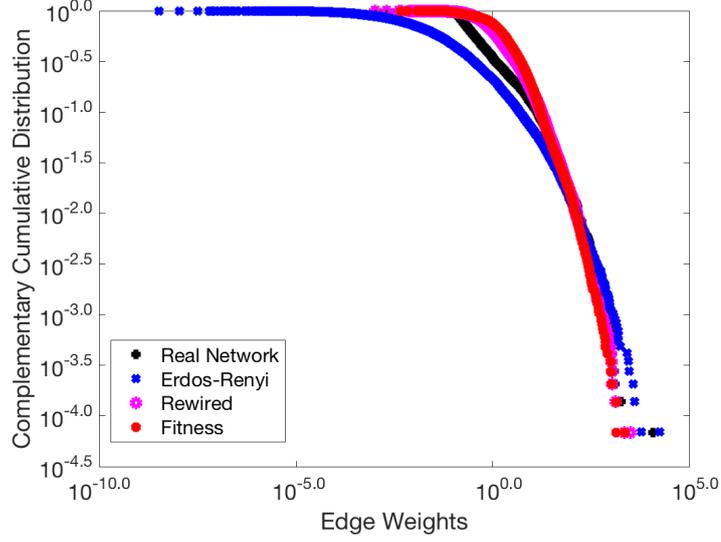


Figure 3.5: Edge weight distribution of the real network (black), random Erdős-Rényi (ER) network (blue), configuration model (pink) and the fitness model (red). Due to the allocation of edge weights using the RAS algorithm for all null models, the edge weight distribution of the three null models is much more similar than the degree distribution shown in fig. 3.4. Nonetheless, it can clearly be seen that the ER networks consist of a much larger abundance of small edge weights and a longer tail that reflects larger edge weights.

To achieve an understanding of the estimation error of each null model I compute the cosine similarity and the root mean squared deviation (RMSD). The cosine similarity C is based on the Euclidean dot product and can be interpreted as the weighted probability of observing an edge in a reconstructed network in the same position as in the real network. Denoting by L_{ij} the matrix of interbank liabilities of the real network and by $L_{ij}^{(0)}$ that of a reconstructed network, C is computed as

$$C = \frac{\sum_{ij} L_{ij} L_{ij}^{(0)}}{\sqrt{\sum_{ij} L_{ij}^2} \sqrt{\sum_{ij} L_{ij}^{(0)2}}}. \quad (3.27)$$

The cosine similarity is bound between 0 and 1, where $C = 1$ indicates that the two matrices are identical.

The RMSD, which measures the relative fraction of misallocated interbank volume, is defined as

$$RMSD = \frac{\sqrt{\sum_{ij} (L_{ij} - L_{ij}^{(0)})^2}}{\sum_{ij} L_{ij}}. \quad (3.28)$$

The cosine similarity between ER networks and the real network is 0.05, for the rewired it is 0.71 and for the fitness networks 0.66 on average. The RMSD for ER networks relative to the real network is 0.16, 0.07 for the rewired and 0.08 for the

fitness networks, as I also report in table 3.1. The very low similarity of the ER networks to the real network comes as no surprise, following the description of the performance of the ER networks in this section. However, since the fitness model requires much less prior information than the configuration model, the relatively equal performance of the configuration and fitness model in producing networks with high similarity and relatively little deviation from the real network is surprising.

3.5.2 Analysis of counterparty risk in the reconstructed networks

I now describe the financial contagion properties of the reconstructed networks. For each model I reconstruct 10,000 realizations and carry out the same stress test exercise on these networks, that I discussed in section 3.4 for the real network. In fig. 3.6 I present the results of these investigations across the 12 quarters from 2006 until 2008 for each null model. The shaded area in fig. 3.6 represents the respective 10% and 90% quantiles as boundaries and the respective average for each null model is shown as a dotted line. The black lines in these plots are drawn in correspondence to the values observed in the real network.

Overall it can be seen in fig. 3.6 that the null model estimates change in step with the values observed in the real network for contagion probability, contagion extent and maximum extent. However, all three metrics in the null models rarely correspond to the exact values observed in the real network across time. In particular, the null models consistently overestimate the contagion probability and underestimate the contagion extent. In contrast, all reconstructed networks offer a good estimate of the the DebtRank amplification observed in the real network. As I will discuss in the following sections, the reason for the good performance of the null models with respect to the DebtRank algorithm, is due to its dependence on interbank leverage and global topological properties, such as the density, which the null models reconstruct well. On the other hand, the Furfine algorithm depends strongly on the local topology, which the null models cannot reconstruct well.

Figure 3.6 shows that the fitness model provides an overall better estimation of the network stability relative to the other null models. In particular, for the contagion probability and extent the fitness model offers the best estimate of the real network. The real network is in fact an instance of the ensemble of fitness model networks in a small number of quarters, such as the first quarter of 2006 and the second quarter of 2008. However, in other quarters the measuring error of the fitness model is high at levels up to 100%. This is surprising because the previous literature has shown the fitness model to be very accurate [41, 97, 123]. A

possible reason for this is that previous studies have not studied the ability of null models to estimate counterparty default risk as defined by the Furfine algorithm. For the DebtRank amplification shown in fig. 3.6d, I find that the fitness model networks provide very good predictions of the values observed in the real network, and that the fitness model networks have the tightest prediction boundaries of the three null models. The good performance with respect to the DebtRank algorithm is in agreement with the previous literature [41].

Erdős-Rényi networks give quite different results with respect to the real network and also have relatively much wider confidence intervals than the configuration and fitness model networks. ER networks strongly overestimate the probability of contagion and they persistently underestimate the conditional and maximum extent of contagion to a larger degree than the other null models. This is consistent with the fact that the real network has a heavy-tailed degree distribution, which is on average more robust with respect to random failures, but more fragile when the seed of contagion is a large degree node, as I explain in section 2.1.2. ER networks on the other hand are characterised by binomial degree distributions that do not have a heavy tail. Whilst the estimates of the ER networks are much better for the DebtRank algorithm, its prediction boundaries are wide relative to the configuration and fitness model. The fact that the ER networks can reproduce the DebtRank amplification well, despite their strong topological differences to the real network, shows that the DebtRank algorithm is driven strongly by global topological properties such as the density and interbank leverage. Information on interbank leverage in particular, is provided to the ER networks via the RAS algorithm.

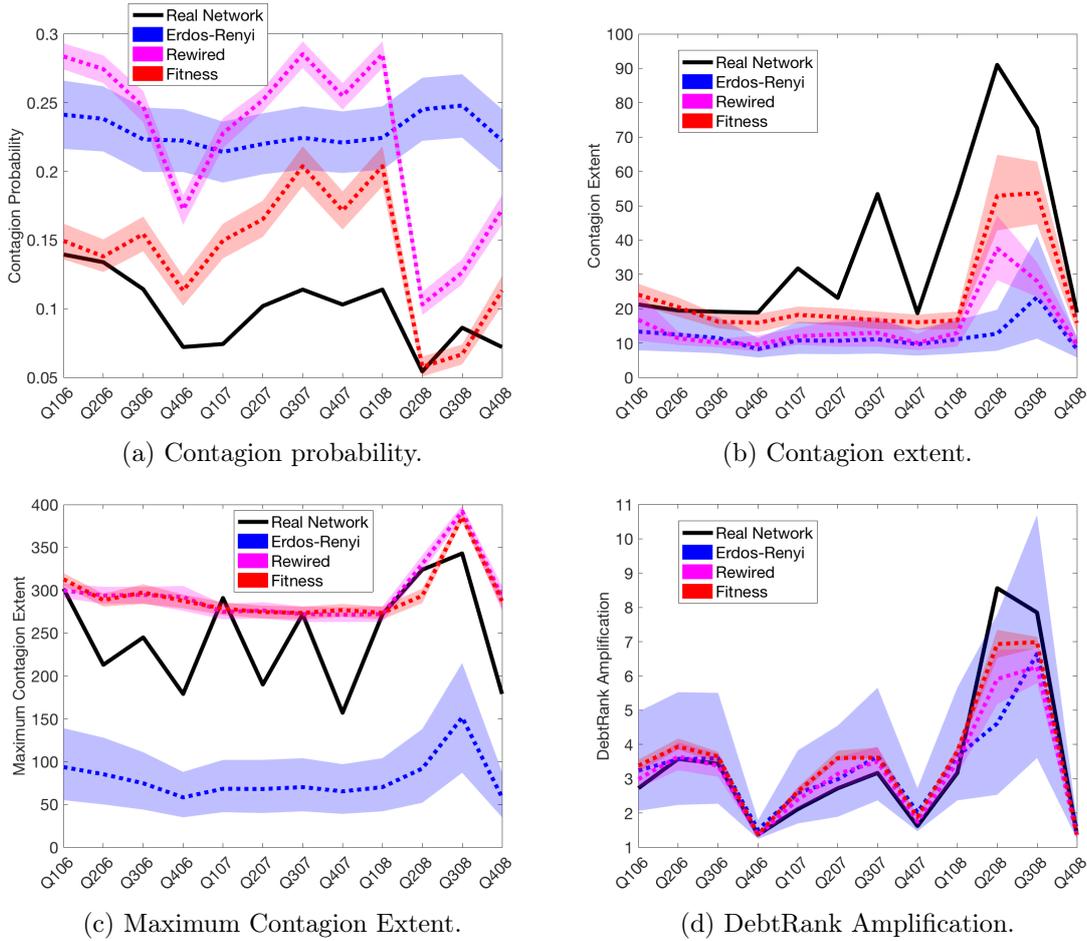


Figure 3.6: Main contagion measures across 12 quarters from 2006 until 2008. Black line shows the measure of the real network. The shaded areas show the 10% and 90% quantiles of the null models respectively and the dotted lines their averages. Blue refers to Erdős-Rényi, pink to configuration and red to fitness model networks. The measured values of contagion in the reconstructed networks generally do not fit those observed in the real network, except for the DebtRank amplification shown in panel (d). The null model estimates are more robust across time than the values observed in the real network, due to their dependence on the relatively more stable total interbank assets and liabilities via the RAS algorithm.

If it was expected that Erdős-Rényi networks cannot reproduce the contagion patterns of the real system, it is surprising that the rewired networks perform worse than those generated via the fitness model. Whilst the contagion probability as estimated by the configuration model does change in lock-step with that of the real network, its estimate is much larger. Equally, for the contagion extent its estimate is consistently smaller than that of the real network. Only for the DebtRank algorithm does the configuration model produce a good estimate. This is surprising because the configuration model is given the entire degree sequence of the real network.

The reason for the poor performance of the configuration model is the following. When links are randomly rewired it may happen that a link that goes from a large lender to a large borrower, is rewired in such a way that the large lender is now

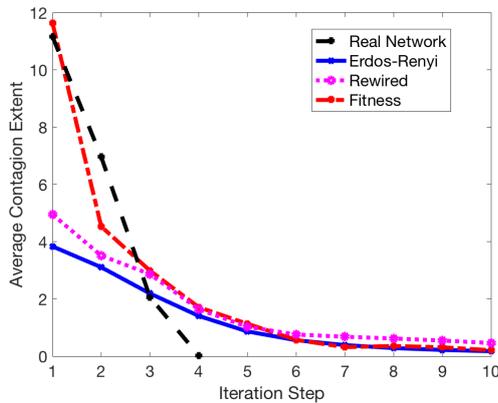
lending to a small borrower. This causes the new link to have a much smaller weight than the original link, and therefore both the large lender and the large borrower will have to reallocate their investments across their counterparties, which results in financially unsound lending allocations. This happens to a lesser extent in the fitness model, in which large lenders are connected preferentially to large borrowers. This is also reflected in fig. 3.8a and fig. 3.8b, which show that the average equity loss is much lower in the rewired compared to the fitness model.

It can also be observed in fig. 3.6 that the estimates of the null models are more similar across time, i.e. more constant, relative to the real network with respect to the results of the Furfine algorithm. Contagion probability and extent critically depend on specific arrangements in the local network topology. The null model networks, as discussed in the previous sections, perform well in reconstructing the higher-order properties of the real network, yet cannot accurately capture the local topology, which is a strong facilitator of financial contagion. In addition, the edge weights of all three null models are fitted with the RAS algorithm, which is based on the total interbank assets and liabilities of each bank. The total interbank assets and liabilities are relatively much more stable across time than individual interbank exposures. This contributes to the relative stability of the null model contagion estimates compared to the values observed in the real network. In contrast for the DebtRank algorithm shown in fig. 3.6d, this difference in the stability of the contagion probability and extent between the reconstructed and real networks does not occur. The reason for this is that unlike the Furfine algorithm, the DebtRank algorithm does not strongly depend on the specific arrangements of bilateral exposures between banks.

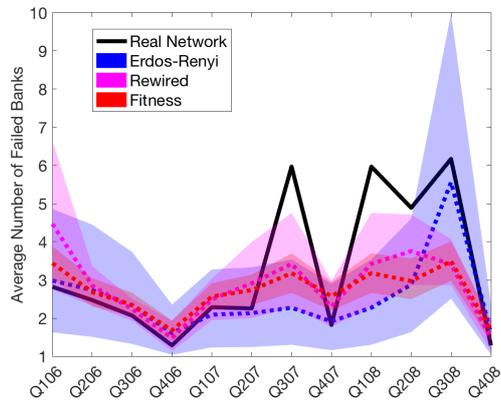
Trajectories of contagion across time

In addition to the final state of the contagion dynamics, I also compare the trajectories observed in the different null model networks due to the Furfine algorithm. Figure 3.7a shows for the first quarter of 2006, which is representative of the results across all quarters, the average number of banks that failed at each iteration step, scaled by the total number of contagion events, i.e. the number of banks (seeds) that cause a contagion event. This analysis shows that whilst in the real network contagion unfolds over 4 rounds, in the reconstructed networks there are on average about 7–8 rounds of contagion that involve a relatively smaller number of banks. The profiles of the trajectories are also different, contagion in the real network abruptly stops, as it remains relatively high at all iterations, while in the recon-

structured networks the contagion rounds decay exponentially.³ This wider spread of less impactful contagion events in the reconstructed networks may be due to the RAS algorithm, which aims at creating homogeneous edge weight distributions despite the heterogeneously distributed marginals, and is common to the three models used. Indeed, given the similarity of the trajectory profiles across the null models, financial contagion may mostly be influenced by the distribution of interbank assets and liabilities. In appendix A.1.2 I show this analysis also for the first quarter of 2007 and 2008 to prove that this behaviour is similar across the quarters.



(a) Total number of banks failed in each iteration of the dynamic scaled by the number of contagion events.



(b) Total number of banks failed across all iterations of the dynamic scaled by the total number of banks.

Figure 3.7: Trajectories of financial contagion due to the Furfine algorithm. Contagion is more abrupt and concentrated in the real relative to the null model networks. Black line shows the measure of the real network. Blue refers to Erdős-Rényi, pink to rewired and red to fitness model networks. L.H.S. figure based on network of first quarter of 2006 only, which is representative of all other quarters. Trajectories are cut off at iteration 10, which in some simulations do not decay until iteration step 20. R.H.S. the shaded areas show the 10% and 90% quantiles of the null models respectively and the dotted lines their averages.

Figure 3.7b shows the total number of failed banks across all iteration steps scaled by the total number of banks n in the network. Interestingly, whilst the null models can capture the average number of failed banks relatively well, from the third quarter of 2007 the values observed in the real network are suddenly much higher than those in the reconstructed networks. This corresponds to the onset of the financial crises, in which many more banks failed. Note that the counterparty default algorithm provides a stylised picture of the failure of banks, its results do not necessarily correspond to actual failures during the financial crisis. The inability of

³The trajectories are computed as follows. For each round of contagion I count the additional banks that defaulted for each seed, I then sum the total defaults for each round and divide this sum by the number of contagion events in the first round. Each sum is divided by the number of contagion events in the first round, rather than by the number of contagion events in each round, to achieve a value that is comparable across rounds.

the null models to reflect this change in the real network, corresponds to my previous explanation of the slow changing nature of the total interbank assets and liabilities on which the edge weights of the null models are based via the RAS algorithm. For this reason the local topology of the reconstructed network cannot quickly pick up on the regime change in the real network, such that the number of failures in the reconstructed networks is relatively stable compared to the real network. For this reason only the configuration model with access to the entire degree distribution, pink in fig. 3.7b, can pick up on the sudden regime change in the real network.

Concentration of risk

The shorter trajectory profiles of the real network, in which more banks fail in each iteration, point to a higher concentration of risk in the real relative to the reconstructed networks. The length of the trajectory profiles as well as the contagion extent and probability of the real and reconstructed networks provide only an indirect view of the concentration of risk. A more direct measure of the concentration of risk is the amount of equity capital lost in each of the networks due to counterparty default risk as measured by the Furfine algorithm. It can be seen in fig. 3.8 that in the reconstructed networks the average and maximum equity lost per node is lower compared to the real network. Indeed, it can be seen that similarly to the contagion extent shown in fig. 3.6b, that the null models underestimate the average equity failed and generally overestimate –except for ER networks– the maximum equity failed, as with the maximum contagion extent, shown in fig. 3.6c. It is also noticeable that the ER networks have the widest confidence intervals and that, again, the fitness model appears to provide the most realistic estimates of the real network.

From this analysis it can be concluded that the trajectories of the reconstructed networks are less impactful than in the real network, not only in terms of the number of banks that fail, but also in terms of the amount of equity capital that is lost.

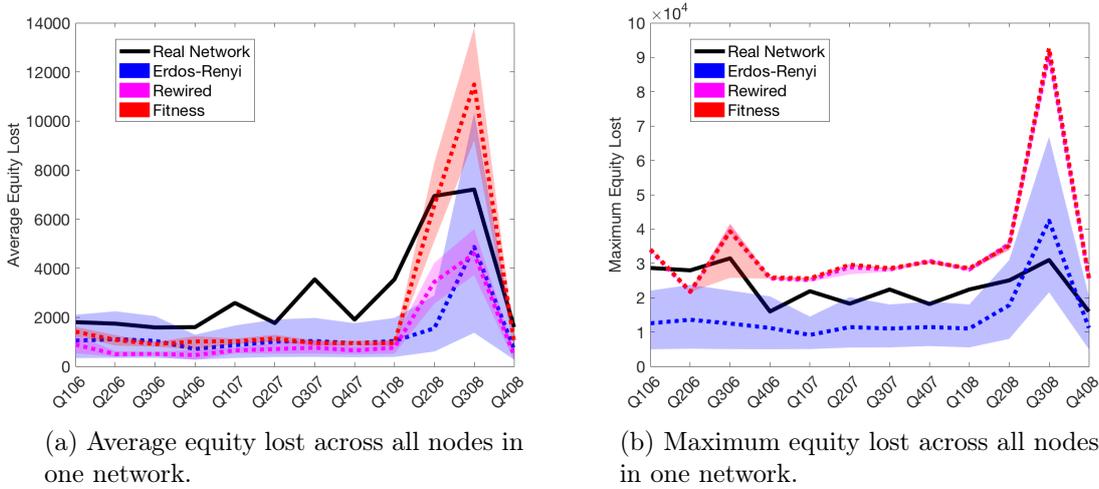


Figure 3.8: Equity capital lost due to counterparty default risk as measured by the Furfine algorithm across 12 quarters from 2006 until 2008. Black line shows the measure of the real network. The shaded areas show the 10% and 90% quantiles of the null models respectively and the dotted lines their averages. Blue refers to Erdős-Rényi, pink to configuration and red to fitness model networks. Similarly to the contagion extent shown in fig. 3.6b, the null models underestimate the equity lost, thereby indicating a lower concentration of risk in their networks.

Contribution of local network topology to systemic risk

Following my analysis of the influence of triadic motifs on financial contagion in the Austrian network in the previous section 3.4.3, I continue this analysis in this section for the reconstructed networks. In particular, I test the ability of the null models to reproduce the number of unique banks that are members of each triadic motif. As noted previously, for the topological analyses my results are based on the network representing the first quarter of 2006 as a representative case. See fig. 3.1 for an overview of the motifs that I analyse and table 3.1 for a summary of these results.

Overall, I find that the number of banks in each motif cannot be reproduced by the null models. This suggests that these motifs are related to the decision-making process of banks and are genuine features of the network formation process [121, 10].

Given that none of the null models include constraints on the triadic motifs, such as a requirement for reciprocity, it is unsurprising that all null models overestimate motif 1, the simplest kind of triad to generate by chance. In contrast, the null models overestimate to a much smaller extent the abundance of members in the most connected triad motif 13, in which each connection is reciprocated. Since the configuration model is provided with the entire degree distribution of the real network, its estimate of 174 unique banks as members of triad 13 comes close to the real network with a count of 187. Given the homogeneous degree distribution of the ER networks, due to which very few nodes exist that have a very low degree,

the ER networks significantly overestimate motifs 1, 5 and 9. For the same reason ER networks do not create nodes of sufficiently high connectivity to form the fully reciprocated motif 13.

Whilst the null models cannot recreate the exact values of contagion of the individual motif members in the real network, the relative ranking of riskiness of the motifs is consistent in the null models with the findings of the real network, that I presented in section 3.4.3. For instance, the average extent of contagion across the banks that form the risky motif 9 in the fitness networks is 6.5. This value is larger than the average extent of contagion of 4.8 across the safer motif 1 nodes in the fitness network, and thus confirms the relative ranking of the contagiousness of the motifs. However, the average contagion of 6.5 observed in the fitness networks for motif 9 is much lower than the average of 14.2 that I observe in the same motif in the real network.⁴

3.6 Vulnerable Cluster

I now develop a deeper intuition of the propagation of financial contagion by introducing the concept of a cluster of vulnerable banks. So far only a theoretical case has been made for the identification, analysis and interpretation of vulnerable clusters, as I described in the literature review in section 3.2.3. I first provide an intuition for vulnerable clusters. I then show how the concept of vulnerable clusters can be employed to understand financial contagion in real networks. In particular, I explain the higher concentration of risk in the real interbank network relative to the reconstructed networks, that I identified in section 3.5.2, via a larger single giant vulnerable cluster in the real network. In appendix A.2 I describe and test the method that I developed to identify the vulnerable clusters.

3.6.1 Intuition

The definition of a vulnerable cluster (VC) of banks is an extension of the definition of the Furfine counterparty default risk algorithm that I outline in section 3.3.1. I define a VC as follows:

1. A bank is vulnerable if it can be failed by the default of only one of its counterparties;

⁴As a reference, the 10% and 90% quantiles support this conclusion: The respective bounds of the average contagion extent for the fitness networks are 4.3 and 5.2 for motif 1 and 5.9 and 7.2 for motif 9.

2. A vulnerable bank is connected to at least one other vulnerable bank, to form a cluster;
3. It follows that the failure of one bank in this cluster defaults all members of its cluster.

Typically there are multiple VCs in a network, here I am interested in the largest of these, which also bears the largest risk and that I refer to as the Giant Vulnerable Cluster (GVC). In the following I will refer to robust banks as those banks that are not vulnerable. Robust banks that are adjacent to the GVC and can fail a member of the VC, constitute the extended giant vulnerable cluster (EGVC).

The sketch in fig. 3.9 provides an intuition of these definitions. A GVC of size 7 implies that a seed can set in motion a cascade of 7 bankruptcies. In the sketch the robust bank 10 for instance would fail bank 2, which in turn can fail banks 1;3; 6; and 7 these in turn would fail banks 4 and 5.

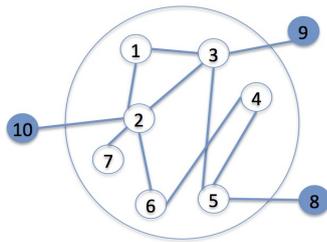


Figure 3.9: A schematic illustration of the (extended) giant vulnerable cluster (GVC). Unfilled nodes are in the GVC, filled nodes are robust banks out-with the GVC, but part of the EGVC. Each robust and vulnerable bank is capable of triggering the failure of the entire GVC. Only "vulnerable" exposures are shown, which could lead to the default of a lender if the borrower is unable to repay.

3.6.2 Vulnerable clusters in the real and reconstructed networks

In this section I describe the results of the VC identification algorithm that I employed on the real and reconstructed networks. I describe the algorithm in appendix A.2.

The identification of the GVC among the vulnerable clusters in the reconstructed networks is not as straightforward as in the case of the Austrian network. In the Austrian network, as I describe in appendix A.2, all but 3 vulnerable nodes are in one giant vulnerable cluster of size 47. However, in the reconstructed networks as I summarise in table 3.2, it can be seen that there is no single individual giant vulnerable cluster. Instead there are many smaller vulnerable clusters in the reconstructed networks. In fact, the average size of the largest observed vulnerable cluster

is always smaller in the reconstructed networks compared to the real network. The abundance in the ER networks is very large, due to the homogeneous spread of links and edge weights, which creates a much larger frequency of contagion events [117].

Vulnerable cluster	Real	ER	Rewired	Fitness
Abundance	1	846	16	4
Maximum size	47	46	28	12

Table 3.2: Abundance and average maximum size of vulnerable clusters in real and reconstructed networks. There is no single giant vulnerable cluster in the reconstructed networks, but a much larger abundance of individual vulnerable clusters. This shows that the local concentration of risk is much larger in the real network. The values for the reconstructed networks are averaged over 10,000 simulations. Results shown for the first quarter of 2006.

The fact that there is no single GVC in the reconstructed networks shows that the null models cannot estimate the local concentration of risk in the real network. The multiple small vulnerable clusters in the reconstructed networks bear significantly less risk than the single giant vulnerable cluster in the real network.

3.7 Discussion and Future Developments

The main goal of this chapter is to study the required amount of information in order to reconstruct a real interbank network. Through a comparison of the required amount of information for each reconstructed network, the difference in their topologies can be investigated to discover the driving mechanisms of financial contagion. So far the literature has been divided in studies of the ability of null models to reconstruct networks and in studies of the contagion mechanisms in interbank networks. By combining these two research strands, I can judge the ability of Erdős-Rényi, directed configuration (rewired) and fitness null models to construct networks with levels of financial contagion similar to that of a real network.

My analysis across time shows that whilst the fitness and configuration models capture the general trends in financial contagion as measured by the Furfine algorithm for counterparty default risk, they do not accurately estimate the exact values of contagion observed in the real network. This is surprising as the previous literature has shown that the fitness model provides very accurate results. Previous studies have focused on the ability of the fitness model to reconstruct the topology of interbank networks and not their financial contagion via the Furfine algorithm [23, 123, 61, 41]. On the other hand, the null models, in particular the fitness model, measure the DebtRank relatively well, which is in agreement with the previous literature [41].

Whilst the topological properties of the fitness and rewired networks are similar to the Austrian network in terms of their assortativity and clustering, the abundance of triadic motifs is substantially different with respect to the real network. Counterparty default risk as measured by the Furfine algorithm depends strongly on the particular arrangement of the local network topology rather than global topological properties. The DebtRank algorithm on the other hand depends more on the global topological properties of the network and interbank leverage. This difference in the dependencies of the two algorithms explains why the reconstructed networks can estimate well the dynamics of the DebtRank but not those of the Furfine algorithm.

The good performance of the fitness model relative to the other two null models suggests that knowing the density of the network and the total interbank assets and liabilities of each bank might be sufficient to perform reliable tests to assess the stability of an interbank lending network, at least at the aggregate level. Moreover, I note that, although the fitness model provides the relatively best approximation of the absorbing state of the dynamics due to the Furfine algorithm, the contagion trajectories, which show the unfolding of contagion, are significantly different from those observed in the real network. I link this deviation to the fact that risky banks are more clustered amongst each other in the real network, suggesting a higher concentration of systemic risk. The adaptation of the vulnerable cluster to systemic risk proved useful in understanding this conclusion. The giant vulnerable cluster is much larger in the real relative to the reconstructed networks, which indicates a higher concentration of risk in the real network.

Future studies should focus their efforts on redefining the fitness model, particularly to better estimate triadic clustering and the vulnerable cluster via the distribution of vulnerable nodes. The estimation of the local topology may be enhanced by considering the reciprocated configuration model [124]. Vulnerable clusters are a convenient tool to connect the contagion patterns of a network to its topology. My analysis of vulnerable clusters highlights the importance of the local topology together with the edge weight distribution as the defining factors of the robust-yet-fragile nature of interbank networks. For this reason, the interplay of the local topology and edge weight distribution requires further investigation. The similarity of the contagion trajectories of the three null models despite their very different topologies, suggests that the edge weights have a strong influence on financial contagion. The edge weights were fitted using the RAS algorithm on each reconstructed network. An important question to be addressed is therefore whether the edge weight distribution determines the size of contagion trajectories, with the role of the local topology being reduced to providing a necessary structure for contagion to propagate. This would suggest that it is more important to control the equity

capital levels of banks, rather than to focus on the exact nature of their lending relationships in an interbank network.

Chapter 4

Topological Determinants of Control Energy

4.1 Introduction

4.1.1 Motivation

Many phenomena can be modelled using complex networks from neural networks, over interbank networks to social networks. In theoretical studies and applications of networks it is often of interest to change the state of nodes in a network. In neural networks for instance, one may aim to control epilepsy, the source of which are overfiring neurons [29]. Similarly, most research on interbank networks has the ultimate goal to make financial systems more resilient [36] and in social networks it is often desired to change a group's perception of a particular issue [136].

Control theory applied to networks stands to provide the solution to the problem of managing the state of networks. Control theory has been extensively studied in the field of engineering since the 1960s beginning with seminal studies by Kalman [80]. The central goal in control theory is to drive a dynamical system from an initial state to a desired final state by exerting an external influence upon the system. Beginning with a study by Liu et al in 2011 interest arose in applying control theory to complex networks [86]. Research has shown that control theory can provide insights into networks, for instance by identifying the essential nodes that are required to exert structural control in a network [47, 60, 88, 113]. The most important aspect of controlling networks, the cost of achieving the control target as measured via control energy, is not yet understood. Understanding the properties of control energy will show whether control problems are practically achievable in terms of their cost. Previous studies have indicated that control energy requirements may be very large and control targets therefore not achievable [137, 138]. These studies only investi-

gated the theoretical boundaries of control energy, but not how it is affected by the topological properties of networks. Indeed, most previous research has focused on defining the structural controllability of networks, which describes the topological constraints in controlling networks [78, 86, 88, 99, 104, 111, 116]. Much more research is therefore required to discover the feasibility of applying control theory to networks.

Ill-defined problem of control energy

Sparse networks provide the most relevant applications for control theory, because most real networks tend to be sparse, as I discuss for interbank networks in section 2.1.2. The controllability Gramian W is a matrix that provides the energy E_i required by each node i to be driven to its control target. The sum across all nodal energies E_i provides the total network control energy E . For sparse networks Sun and Motter (2013) show that W is liable to be ill-conditioned. The reciprocal condition number indicates the degree to which a matrix is ill-conditioned and describes the inverse of the relative numerical error when using W . A reciprocal condition number of 1 indicates that the matrix is well-conditioned and 0 that it is singular. The reciprocal condition number of W ranges from 10^{-4} for networks with $n = 50$ nodes and 10^{-18} for $n = 250$ [128]. Due to this computational uncertainty no complete analytical description of E_i exists.

In sparse networks the analytical problem of calculating the control energy reduces to finding the eigenvalues of sparse matrices. The nodal energies depend on the eigenvalues of W , which is based upon the adjacency matrix that defines the connections between nodes. Yan et al (2012) provide theoretical bounds on the control energy E using the Rayleigh-Ritz method to approximate the eigenvalues of W for directed and weighted scale-free networks. They find that the lower boundary is $E_{\min} = 1/\lambda_{\max}$ and the upper boundary is defined as $E_{\max} = 1/\lambda_{\min}$ [137]. With an increasingly heterogeneous degree distribution the bounds E_{\min} and E_{\max} diverge, due to the complexity of managing the different nodal dynamics. These bounds can therefore not provide a thorough understanding of the influence of network topology on control energy E .

4.1.2 Research objective and approach

In this project I investigate the impact of network connectivity, network size and the degree distribution on the properties of control energy. I devise a simple simulation to compute the control energy of Erdős-Rényi (ER) and scale-free (SF) networks. This simulation provides a first step in defining the topological determinants of

control energy. The simulation is set up in such a way as to avoid the numerical issues that I outline in (4.3) in the literature review. I achieve this by studying small networks with $n = 2,000$ and controlling all nodes $n_d = n$.

I additionally take a mean-field approach to provide an analytical solution to the problem of computing the control energy of ER networks. Mean-field theory is a standard method in statistical physics to study the dynamics of large complex systems. By averaging across nodal degrees and states, I solve a simpler optimisation problem than by computing the controllability Gramian W . This additionally avoids the computational issues arising through the use of W that I outlined in the previous section.

Contributions

In this project I make the following three main findings and contributions:

1. An increase in the network connectivity results in a decreasing control energy requirement due to a coordination advantage: As nodes become increasingly connected, they push one another to their common target state, resulting in a lower external input requirement.
2. I provide an analytical proof of the coordination advantage based on a mean-field approximation.
3. The coordination advantage diminishes with an increasing distance between the target states of neighbouring nodes. Nodes with different target states cannot push one another towards a common goal, which results in a larger control energy requirement.

4.2 Related Work

In this section I review the literature on control theory and its recent application to networks. I first provide a general discussion of the concepts of control theory. I then introduce the literature on control energy, which defines the cost of controlling a system. Lastly, I discuss the literature on the impact of network topology on control energy, which is the main focus of this project.

4.2.1 Control theory on networks

Control theory encompasses all necessary methods in defining the controllability of a (non-)linear system to drive it from its current state to a target state. To simplify

the analysis of control problems, system dynamics are typically linear and defined such that the relations between nodes do not change across time. These systems are referred to as linear time-invariant (LTI) systems [115]. I outline the set-up of LTI systems in section 4.3.1.

4.2.2 Defining the controllability of a dynamical system

In control problems the goal is to move a system from a given state defined by the state vector $\vec{x}(t_0)$ to a final state $\vec{x}(t_f)$, where the entries of $\vec{x}(t_0)$ are $x_1 = x_2 = \dots = x_n = 0$ and n is the number of nodes in the system. A system is controllable if it is possible to drive it to the desired final state $\vec{x}(t_f)$ in finite time t_f from the original state $\vec{x}(t_0)$.

The initial research that defined the controllability of dynamical systems was undertaken by Kalman in the 1960s [80]. To test the controllability of a LTI system the linear combination $[A, B]$ is studied, as I outline in detail in section 4.3.1. In traditional control problems A represents a set of linear functions, whereas in network control problems $G(A, B)$ represents a directed weighted network A with external controls B . The number of drivers, the nodes that exert the external influence upon the nodes, is referred to as n_d .

Kalman's rank condition is used to test the controllability of a LTI system [80]. The controllability matrix of size $n \times (n \times n_d)$ is defined as follows

$$\mathcal{C} \equiv [B, AB, A^2B, \dots, A^{n-1}B], \quad (4.1)$$

which is the linear combination of A and B . It's full rank is

$$\text{rank } \mathcal{C} = n. \quad (4.2)$$

If $\text{rank } \mathcal{C} < n$ then \mathcal{C} is rank deficient and the system is not fully controllable. In this case, the formal solution to drive the system to $\vec{x}(t_f)$ does not span the entire state space, as the vectors of \mathcal{C} are not linearly independent. In terms of network control, rank deficiency means that each node cannot be individually accessed by the controller, such that the network is not fully controllable [86].

4.2.3 Control energy

One of the most important aspects of control theory is the cost of driving a system to its target state. The required control energy is crucial in determining whether a control target is reachable in practical terms.

As I describe in the methodology section 4.3.4, the Gramian matrix W is crucial in defining the total network control energy E . Since the Gramian W grows in size

proportionally to n and since most networks are not complete, the chance of W being singular, i.e. ill-conditioned, increases as $n \rightarrow \infty$, especially when the number of links does not increase at the same rate as n . For similar reasons, the reciprocal condition number γ is also significantly reduced when the control time increases $t_f \rightarrow +\infty$ and when $n_d \rightarrow 1$ [86, 106, 128, 137].

Even though the traditional controllability criterion based on Kalman's rank condition would indicate that a network is controllable, control may not be practically feasible as the energy is not calculable. The cost of controlling a network is infinitely high when W cannot be inverted, which is the case when W is singular. Sun and Motter (2013) show that γ_W the reciprocal condition number of W , has a size in the order of 10^{-4} for $n = 50$ and for $n = 250$ in the range of 10^{-12} and 10^{-18} , depending on n_d [128]. A reciprocal condition number of 0 defines singular matrices. Sun and Motter (2013) find a critical threshold $\gamma_W \approx 10^{-10}$ below which most computations fail numerically. This threshold corresponds to $n_d/n \approx 20\%$ of nodes being controlled, such that when a lower fraction of nodes is controlled most computations fail. For SF networks this controllability transition is less sharp as their degree distribution is more heterogeneous, which results in a broader dispersion of the entries of W [128].

Theoretical bounds on the control energy

Theoretical bounds on E have been provided by Yan et al (2012), who use the Rayleigh-Ritz method to approximate the eigenvalues of W [137]. These approximations are provided for undirected and directed weighted ER and SF networks. When the adjacency matrix A is positive-definite the bounds are $E_{\min} = 1/\lambda_{\max}$ and $E_{\max} = 1/\lambda_{\min}$, if on the other hand A is negative-definite, the upper bound will vanish exponentially as $t_f \rightarrow \infty$. This is due to the largest eigenvalue λ_{\max} growing exponentially with t_f , as longer time intervals lead to a larger control energy requirement. When n_d is increased, the lower bound decreases whilst the upper bound increases. This follows from the fact that driver nodes tend to be low degree nodes, which are more difficult to access and therefore more costly to control [86, 72, 137]. Pasqualetti et al (2014) provide a similar analysis on the bounds of control energy [106]. Showing for undirected networks with constant n_d , that E increases exponentially with n . In turn, when n_d is a fixed fraction of the network size, then the energy grows independently of n .

To summarise, the bounds on the control energy E that have been found by the existing literature are very wide, especially when the degree distribution is heterogeneous. They therefore do not provide an understanding of the topological determinants of E . In the following cases significant reductions in the reciprocal

condition number of W are observed, such that the control energy E cannot be computed:

$$\begin{aligned}
 t_f &\rightarrow +\infty \text{ (control time);} \\
 n &\rightarrow +\infty \text{ (number of nodes);} \\
 n_d &\rightarrow 1 \text{ (number of driver nodes);} \\
 \lim_{e \rightarrow 1} \frac{e}{n(n-1)} &\text{ (density).}
 \end{aligned} \tag{4.3}$$

4.2.4 The impact of network topology on control energy

In this section I review the literature on the impact of the topology of networks on control energy. Since the main focus of this project is to study the influence of the topology on control energy, this literature review provides an essential background to my analysis.

A network consisting of well connected clusters in the form of strongly connected components (SCC), each controlled by an individual driver node, requires less energy than a network in which driver nodes are placed randomly [106]. Pasqualetti et al (2014) develop a decoupled control strategy by placing a driver node in each of the network's SCCs [106]. Importantly, the authors show that the strategy is scalable and bounded by the maximal energy requirements for clusters, rather than the entire network. However, the requirement of the existence of SCCs rules out most real world networks, in which SCCs are rare [102]. An additional major restriction of the decoupled control strategy is that a feedback mechanism in the form of reciprocal links must exist between the SCCs. Given the precise analysis of the network topology that underlies the decoupled control strategy, it is unsurprising that its energy requirements are lower compared to a strategy in which driver nodes are placed randomly.

Cowan et al (2012) discuss the importance of nodal dynamics for the controllability of networks [43]. Cowan et al suggest that in practical control problems inaccessible nodes may simply be ignored to significantly reduce the very large control energy requirement of isolated nodes. Future work should investigate the practicality of such an approach, by providing an understanding of the role of inaccessible nodes in real networks. I partly address this problem in the fourth project by showing that the removal of isolated nodes significantly reduces the external input. The removal of isolated nodes is sensible in interbank networks, as isolated nodes are a result of the way scholars construct interbank networks, rather than a reflection of truly isolated banks in the real system, see section 6.5. A further example of the effect of isolated nodes are gene regulatory networks. In gene regulatory networks a small number of nodes are sufficient to control the majority of a network [87]. The

small number of driving nodes required in gene regulatory networks is due to a high-degree of co-regulation in these networks. In this case the small part of the network that is not controlled by the driver nodes is frequently ignored by practitioners [87].

In section 4.2.3 I introduce the numerical controllability threshold for control energy. Research by Sun and Motter (2013) shows, that when a fraction of less than $n_d/n = N_D < 20\%$ of nodes is controlled in a network that the controllability Gramian is ill-conditioned, such that the energy cannot be computed. The threshold of 20% is much higher than suggested by the theoretical results of structural controllability [128]. In response, Nie et al (2016) investigated the effect of degree correlations on the numerical threshold [104]. The authors find that disassortative networks generally require a higher fraction of driver nodes than networks with neutral and assortative degree mixing. For instance, an ER network of size $n = 100$ with average degree $\langle k \rangle = 2$ requires $N_D = 60\%$ with neutral degree mixing $r = 0$, but with disassortative mixing $r = -0.7$ it requires a larger fraction of $N_D = 70\%$. This effect is stronger for SF networks, such that for a SF network of size $n = 100$ with a power-law exponent $\gamma = 3$ and $r = 0$ the fraction of drivers is $N_D = 80\%$ and for $r = -0.7$ it is $N_D = 95\%$. Additionally, the relative fraction of drivers increases with the network size when the density does not increase at the same rate. On the other hand, when the average degree $\langle k \rangle$ increases, the required fraction of drivers decreases as nodes can be reached more easily. This study by Nie et al shows the limitations of the current applicability of control theory to networks. Given the low connectivity of some nodes in real networks, it is often impractical to control all nodes in a network. This makes the computation of control energy impossible in many cases.

4.3 Methods

4.3.1 Control system set-up

A linear control system consists of a set of inputs, outputs and state variables whose behaviour is defined by a set of linear state equations and is expressed as a linear time-invariant (LTI) system [115]:

$$\vec{x}(t+1) = A\vec{x}(t) + B\vec{u}(t) \quad (4.4)$$

where:

- $\vec{x}(t) \in \mathbb{R}^n$ is the state vector, with n representing the number of nodes;

- $A \in \mathbb{R}^{n \times n}$ is the state matrix, representing the weighted directed adjacency matrix, describing the connections between the nodes;
- $B \in \mathbb{R}^{n \times m}$ is the input matrix, representing the matrix of the weights of the inputs upon the nodes, with m representing the number of inputs;
- $\vec{u}(t) \in \mathbb{R}^m$ is the input vector.

Since the system is time-invariant, A and B are constant in time. The dynamics in this chapter are described in discrete time.

Formally an LTI system is represented by a directed graph $G(A, B)$. The element a_{ij} of matrix A denotes the influence of node i upon node j . The element b_{ik} of the input matrix B reflects the strength of the control signal acting on node i .

4.3.2 Static method for scale-free network generation

I use the static method developed by Goh et al (2001) to generate directed scale-free networks [67]. An empty adjacency matrix A is set up of size $n \times n$. The n nodes are indexed with an integer $i \in \{1, 2, \dots, n\}$.

A control parameter z in the continuous range $[0, 1]$ defines the power-law exponent as follows $\gamma = (1 + z)/z$. The probability that node i has a link to node j is assigned as $p = i^{-z}$ and is normalised as $p = p / \sum_{i=1}^n p_i$.

To build the adjacency matrix A , the desired number of edges e is iterated through sequentially whilst picking an ingoing node i and an outgoing node j with probability p_i and p_j , respectively. An edge is added between them to A_{ij} , unless an edge already exists, whilst ensuring that the diagonal is empty $A_{ii} = 0$ to avoid self-loops.

4.3.3 System stability

This section describes the mathematical background to the stability properties of the solutions to linear time-invariant (LTI) systems.

Uniform exponential stability defines the boundedness and asymptotic behaviour of $\vec{x}(t)$ [115]. A system $\vec{x}(t_0+1) = A\vec{x}(t_0)$ is described as uniformly and exponentially stable when a finite constant β exists, such that

$$\sum_{t=0}^{\infty} \|A^t\|_2 \leq \beta. \quad (4.5)$$

This equation implies $|\lambda_1| \dots |\lambda_n| < 1$, because setting a real eigenvalue $\lambda \geq 1$

and \vec{v} the associated eigenvector, such that

$$A^t \vec{v} = \lambda^t \vec{v}, \quad t > 0, \quad (4.6)$$

and taking $\vec{x}(0) = \vec{v}$, then the solution $\vec{x}(t) = A^t \vec{v}$ grows unboundedly as $t \rightarrow \infty$. It follows that a necessary and sufficient condition for uniform exponential stability is that $\lim_{t \rightarrow \infty} A^t = 0$. This is guaranteed when $|\lambda_1| \dots |\lambda_n| < 1$. A result of this is that $\lim_{t \rightarrow \infty} \vec{x}(t) = 0$ for any $\vec{x}(t_0) > 0$.

4.3.4 Control energy

State transition matrix

The state vector $\vec{x}(t)$ at time t can be expressed in terms of the transition matrix $\Phi(k, t) = A^{k-t}$, for $k \geq t$, as

$$\vec{x}(t+1) = \Phi(k, t) \vec{x}(t). \quad (4.7)$$

Φ has the following transition property [115]:

$$\vec{x}(t_2) = \Phi(t_2, t_1) \vec{x}(t_1), \quad (4.8)$$

where

$$\vec{x}(t_1) = \Phi(t_1, t_0) \vec{x}(t_0). \quad (4.9)$$

This in turn implies the composition property:

$$\Phi(t, j) = \Phi(t, k) \Phi(k, j), \quad j \leq k \leq t. \quad (4.10)$$

Controllability Gramian

In order to compute $E(t_f)$ the total energy required to control the system, the controllability Gramian, which provides the eigenenergies E_i of the system, is defined as follows:

$$W(t_0, t_f) = \sum_{t=t_0}^{t_f-1} \Phi(t_f, t+1) B B^T \Phi(t_f, t+1)^T, \quad (4.11)$$

$t_f > t_0$. W must be positive definite (PD) for the system to be controllable within t_f time steps [115]. If it is not controllable in this period, it may be controllable at a later time $t > t_f$.

Control input and control energy

When the control energy is minimised it follows from (4.11) that the minimal control input $\vec{u}(t)$ required to reach the target state $\vec{x}(t_f)$ is

$$\vec{u}(t) = B^T \Phi(t_f, t + 1)^T W(t_0, t_f)^{-1} \vec{x}(t_f), \quad t \in [t_0, t_f]. \quad (4.12)$$

For a proof see Rugh [115, p. 483f]. The total energy required is

$$E(t_f) = \sum_{t=1}^{t_f} \|\vec{u}(t)\|_2^2, \quad (4.13)$$

where $\|\vec{v}\|_2 \forall \vec{v} \in \mathbb{R}^n$ denotes the Euclidean norm $\|\vec{v}\|_2 \equiv \sqrt{\vec{v}^T \vec{v}}$.

4.4 Topological Properties of Network Control Energy

In this section I describe the results of my simulations and analytical work on the topological determinants of network control energy.

4.4.1 Simulation set-up

The simulations in this project are set up as follows. Weighted and directed Erdős-Rényi (ER) networks A are generated of size $n = 2,000$ with an average degree $\langle k \rangle$ in the range of $[0.5, 100]$, see equation (3.18). Scale-free (SF) networks are generated using the static method described in section 4.3.2, with a power-law exponent $\gamma = 3$ and an average degree $\langle k \rangle$ in the same range as for ER networks.¹ The SF networks are relatively small and can therefore not technically be described as scale-free, however I refer to these networks as scale-free for convenience. The edge weights for both ER and SF networks are drawn from a normal distribution with $\mu = 0.009$ and $\sigma = 10^{-4}$. μ is chosen such that the system A is stable, i.e. $|\lambda_{\max}| < 1$.

Following section 4.3.3 any non-zero vector on the trajectory of the system is a viable target state, as the system naturally reverts to zero. The target state $\vec{x}(t_f)$ is therefore set as: $x_1(t_f) = x_2(t_f) = \dots = x_n(t_f) = 1$ and $\|\vec{x}(t_f)\|_2 = 1$ for the interval $[t_0, t_f]$, where $t_0 = 1$ and $t_f = 3$. This is a standard set-up in the control

¹I do not report results on $\langle k \rangle > 100$, because the average edge weight μ would need to be decreased in order for the necessary stability condition $|\lambda_{\max}| < 1$ to hold. An analysis across all possible $\langle k \rangle$ in the range $[0, 2,000]$ would therefore not compare the same networks. Note however that when both ER and SF networks are fully connected, their energy requirements are necessarily identical.

literature [115, 137, 138, 106]. A relatively short interval is chosen in order to provide a clearer intuition. I show in section 4.4.3 that the length of the chosen time interval has a small effect upon the results. I will show the effect of changing t_f and the distribution of $\vec{x}(t_f)$ in this chapter.

The number of drivers is set as $n_d = n$. The number of simulated networks for a given $\langle k \rangle$ is 5,000. In order for the control simulation to be successful the actual achieved final state $\tilde{x}(t_f)$ must be within a certain distance of the target state $\vec{x}(t_f)$. I find that in all simulations $\frac{|\tilde{x}(t_f) - \vec{x}(t_f)|}{\vec{x}(t_f)} < 10^{-6}$, which is the standard distance criterion used in the literature [128]. All computed controllability Gramian matrices W are positive definite.

4.4.2 Increasing network connectivity reduces control energy

It can be seen in fig. 4.1 that the total energy E required to control a network is a monotonous function of the average degree $\langle k \rangle$ in ER and SF networks. As the connectivity between nodes increases, nodes push one another to their common target state. This endogenous interaction results in a lower energy requirement, which I refer to as the coordination advantage. It can also be seen in fig. 4.1 that the energy requirement decreases more abruptly in SF relative to ER networks. This is due to the presence of a small number of hub nodes in SF networks. Hub nodes have a very large degree, which allows them to reach almost all nodes in a network. They are therefore able to coordinate the states of most nodes in a network, in contrast to ER networks in which hub nodes do not exist.

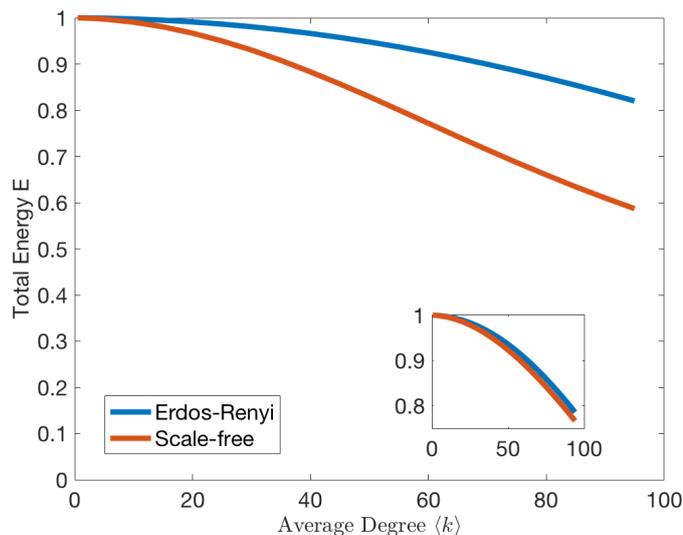


Figure 4.1: Average total network control energy E across simulations as a function of the average degree $\langle k \rangle$ for ER and SF networks. Inset: for $n = 100$, in which the difference between ER and SF networks is less strong.

I find that that the total energy requirement increases monotonically with the size of a network. Because I normed the target vectors $\|\vec{x}(t_f)\|_2 = 1$, I can compare the energy requirements across the different network sizes shown in fig. 4.1. In smaller compared to larger SF networks, hub nodes have a lower degree relative to the average node. The relative fraction of nodes whose state can be coordinated by hub nodes is therefore reduced in smaller compared to larger SF networks. Due to the diminishing effect of hub nodes on the total network energy, the energy requirement of smaller SF networks is therefore accordingly higher. It can clearly be seen in fig. 4.1 that the energy requirement in the smaller SF networks is much closer to ER networks when $n = 100$ shown in the inset, in contrast to the case of $n = 2,000$. Whilst the energy requirement at ≈ 0.8 is almost identical for ER networks of size $n = 100$ and $n = 2,000$, the energy requirement for SF networks with $n = 100$ is higher than for $n = 2,000$.

4.4.3 Analytical solution to optimal control

In this section I provide an analytical derivation of the control problem. Computing the energy requirements of networks analytically is computationally intensive and may fail if the controllability Gramian W (4.11) is ill-conditioned. I take the mean-field approach to derive a simpler analytical solution to the problem of finding the minimum control energy. For a general introduction to the mean-field approach see Sethna (2006) [120, p. 284].

Set-up of the problem

The essence of the mean-field approximation for ER networks is that the state and input of each node is approximated by the average state and input across all nodes:

$$\sum_{j=1}^n A_{ij} \approx \sum_{j=1}^n p_{ij} = pn = \langle k \rangle, \quad (4.14)$$

where p_{ij} is the probability of node i and j being connected.

So far I set $\vec{x}(t+1) = A\vec{x}(t) + B\vec{u}(t)$ according to the LTI system equation (4.4). Since $n_d = n$, B can be omitted as $B = I$ the identity matrix. Additionally, I set $\alpha = \langle k \rangle \mu$. In the following set of equations it can be seen that when each node can be assumed to have the same degree, which is the essence of the mean-field

assumption, that the LTI system equation (4.4) reduces to:

$$\begin{cases} x(1) &= u(1) \\ x(2) &= \alpha u(1) + u(2) \\ x(3) &= \alpha^2 u(1) + \alpha u(2) + u(3) \\ &\vdots \\ x(t_f) &= \alpha^{t_f-1} u(1) + \alpha^{t_f-2} u(2) + \dots + u(t_f) \\ &= \sum_{t=1}^{t_f} \alpha^{t_f-t} u(t). \end{cases} \quad (4.15)$$

The optimisation problem accordingly reduces to

$$\begin{aligned} \min_u \sum_{t=1}^{t_f} |u(t)|^2 \\ \text{s.t. } \sum_{t=1}^{t_f} \alpha^{t_f-t} u(t) &= x(t_f). \end{aligned} \quad (4.16)$$

Solution of the optimisation problem

To solve (4.16) I first state the problem that I seek to minimise as follows

$$\mathcal{L}(u, \lambda) = \sum_{t=1}^{t_f} u(t)^2 - \lambda \left(\sum_{t=1}^{t_f} \alpha^{t_f-t} u(t) - x(t_f) \right). \quad (4.17)$$

The first order conditions are

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial u(t)} &= 2u(t) - \lambda \alpha^{t_f-t} = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= \sum_{t=1}^{t_f} \alpha^{t_f-t} u(t) - x(t_f) = 0 \end{cases} \quad (4.18)$$

and I rearrange such that $u(t) = \frac{\lambda \alpha^{t_f-t}}{2}$ and insert $u(t)$ into $\frac{\partial \mathcal{L}}{\partial \lambda}$

$$\frac{\sum_{t=1}^{t_f} \alpha^{2(t_f-t)} \lambda}{2} = x(t_f). \quad (4.19)$$

The summation term represents a geometric series that can be simplified with an intermediate step as:

$$\frac{\sum_{t=1}^{t_f} \alpha^{2(t_f-t)} \lambda}{2} = x(t_f) \quad \Rightarrow \quad \frac{\alpha^{2t_f} \lambda \sum_{t=1}^{t_f} \alpha^{-2t}}{2} = x(t_f) \quad (4.20)$$

$$\frac{\alpha^{2t_f} \lambda \alpha^{-2t_f} (\alpha^{2t_f} - 1)}{2} = x(t_f) \quad \Rightarrow \quad \frac{\lambda \alpha^{2t_f} - 1}{2 \alpha^2 - 1} = x(t_f). \quad (4.21)$$

Equation (4.21) yields

$$\lambda = \frac{2x(t_f)(\alpha^2 - 1)}{\alpha^{2t_f} - 1}, \quad (4.22)$$

which I insert into $u(t)$, to arrive at the final expression for the input

$$u(t) = \frac{\alpha^{t_f-t} 2x(t_f)(\alpha^2 - 1)}{2(\alpha^{2t_f} - 1)} \Rightarrow \frac{\alpha^{t_f-t}(\alpha^2 - 1)x(t_f)}{\alpha^{2t_f} - 1}. \quad (4.23)$$

This allows the energy requirement to be expressed as

$$E = \frac{\sum_{t=1}^{t_f} \alpha^{2(t_f-t)} (\alpha^2 - 1)^2 x(t_f)^2}{(\alpha^{2t_f} - 1)^2} \quad (4.24)$$

$$= \frac{(\alpha^2 - 1)^2 x(t_f)^2 \alpha^{2t_f} \sum_{t=1}^{t_f} \alpha^{-2t}}{(\alpha^{2t_f} - 1)^2} \quad (4.25)$$

$$= \frac{(\alpha^2 - 1)^2 x(t_f)^2 \alpha^{2t_f}}{(\alpha^{2t_f} - 1)^2} \alpha^{-2t_f} \frac{\alpha^{2t_f} - 1}{\alpha^2 - 1} \quad (4.26)$$

$$= \frac{\alpha^2 - 1}{\alpha^{2t_f} - 1} x(t_f)^2, \quad (4.27)$$

where in (4.26) I again made use of the fact that the summation term represents a geometric series. I then finally arrive at the expression of the control energy E using another geometric series transformation

$$E = \frac{x(t_f)^2}{\sum_{t=0}^{t_f-1} \alpha^{2t}}. \quad (4.28)$$

Results of mean-field approximation and comparison with numerical simulations

This final result of the solution for E (4.28) confirms the coordination advantage. The minimum energy required reduces with an increasing average degree expressed by α . The energy computed from (4.28) together with the averaged results across 5,000 numerical simulations on ER networks with $n = 100$ are shown in fig. 4.2a. The figure shows an excellent match between the analytical solution and the numerical simulations.

The analytical expression of E (4.28) further suggests that the impact of increasing the control time t_f reduces the energy requirement, which decays exponentially as $t_f \rightarrow \infty$. E asymptotes when $t_f \approx 20$ as I show in fig. 4.2b.

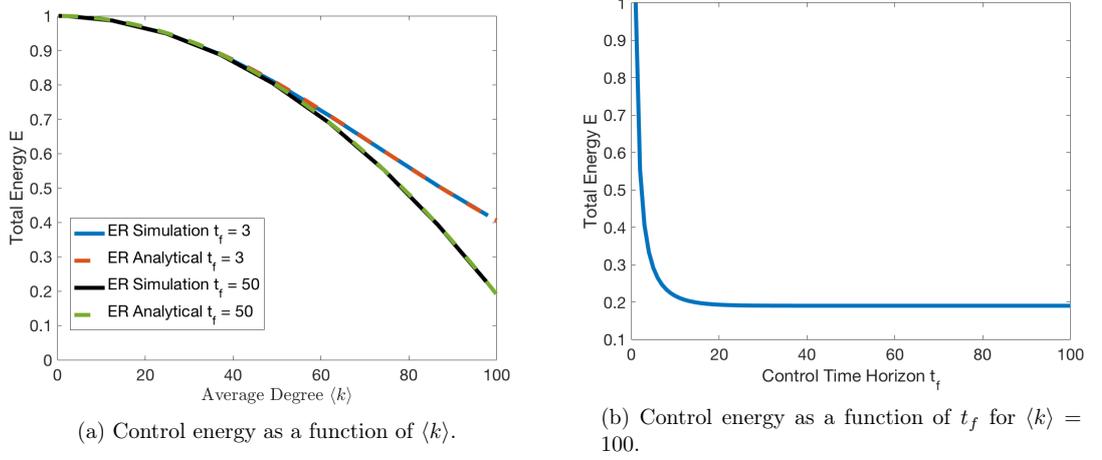


Figure 4.2: L.H.S.: Average across 5,000 ER network simulations (straight line) together with analytical solution (dashed line) from eq. (4.28) for $t_f = 3$ (blue and red respectively) and $t_f = 50$ (black and green respectively). The analytical and numerical solutions match perfectly. R.H.S.: Analytical solution as a function of t_f for $\langle k \rangle = 100$.

Distribution of network control energy

Figure 4.3 shows the control energy distribution for ER and SF networks. Whilst the energy distribution is homogeneous for ER networks, it has tail for SF networks. The networks in the tail of the SF energy distribution are characterised by a relatively lower abundance of small degree nodes and a relatively higher abundance of hubs. This suggests that hubs are able to reduce the energy requirement.

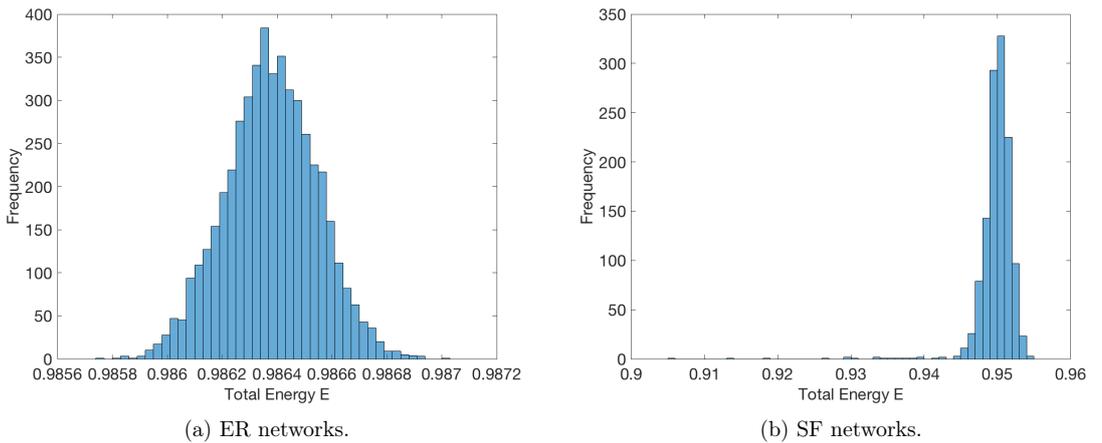


Figure 4.3: Distribution of total energy E for ER and SF networks, here shown for $\langle k \rangle = 25$, which is representative of all average degrees.

4.4.4 Topological determinants of control energy

Having presented the coordination advantage together with an analytical derivation, I present the results of investigations on the topological determinants of control energy in the remainder of this chapter.

I begin this section with an analysis of nodal energies E_i . Nodal energies as a fraction of the total network energy $\frac{E_i}{E}$ are much more varied for SF than for ER networks as shown in fig. 4.4. Each of the individual points in the scatter plot 4.4 represent $\frac{E_i}{E}$ for each node averaged over all simulations. The figure therefore depicts the expected energy of controlling a node with a given degree k_i . It can also be seen in fig. 4.4 that the relative fraction of energy required by hub nodes is several magnitudes larger than for an average node in SF networks.

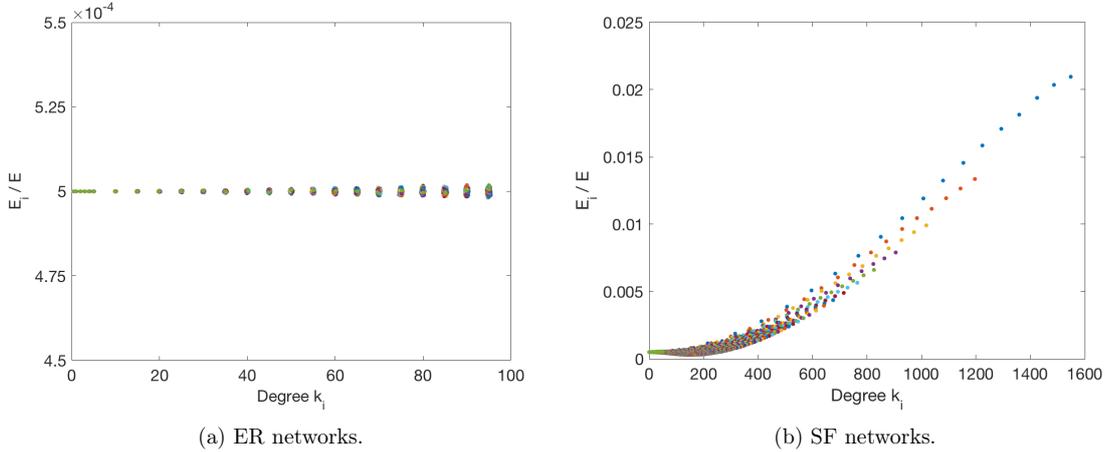


Figure 4.4: The figures represent the expected fractional nodal energy $\frac{E_i}{E}$ requirement for a given degree k_i averaged across all simulations for ER and SF networks, where the colours represent the different average degrees $\langle k \rangle$. The fraction is almost identical for nodes across all nodal degrees in ER networks, but highly heterogeneous for SF networks, with hub nodes requiring the largest fraction of the energy.

The inverse participation ratio (IPR) is an indicator of the concentration of observations, which allows an investigation of the variation of nodal energies. The IPR has a lower bound of 1 when one node requires all control energy and an upper bound n when each node requires an identical amount of energy. The IPR is calculated as

$$\text{IPR} = \frac{1}{\sum_{i=1}^n p_i^2}, \quad (4.29)$$

where $p_i = \frac{E_i}{E}$ for each node i . As it can be seen in fig. 4.5 the energy is concentrated in fewer nodes as the average degree $\langle k \rangle$ increases in SF networks, in which the IPR falls to ≈ 700 . In ER networks this concentration of nodal energies is essentially absent.

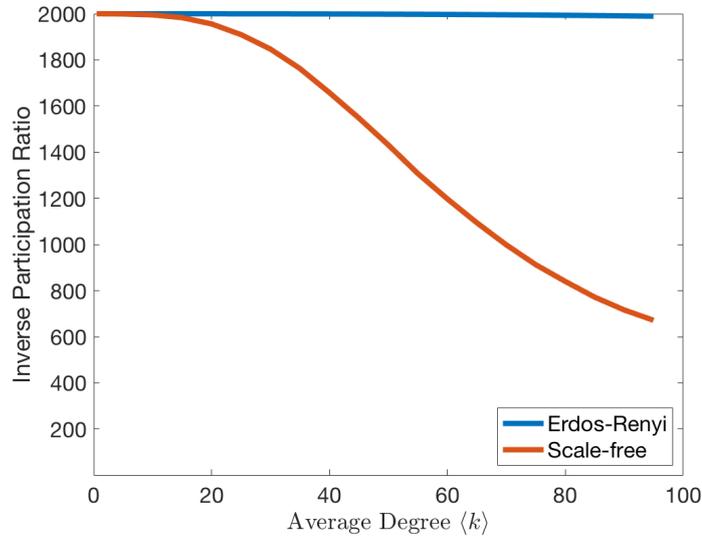


Figure 4.5: The inverse participation ratio (IPR) for ER (blue) and SF (red) networks. The concentration of nodal energies is much higher in SF relative to ER networks.

4.4.5 Absence of coordination advantage when nodal states deviate

I now investigate the impact of the distribution of the target state vector entries $x_i(t_f)$ on the properties of control energy. The existence of the coordination advantage is to be understood with regards to a specific goal in controlling a network. In the simulations presented so far, all nodes had the same target states. I will now look at the case in which each node has a different target state.

Figure 4.6a shows the total energy E as a function of the average degree $\langle k \rangle$ for various distributions of $\vec{x}(t_f)$ for ER and SF networks. In the figures the label *Identical* refers to the case when all nodes have identical target states as described in section 4.4.1 and used as the base case for the analysis throughout the previous sections. The label *Normal Distribution* refers to a target vector with a normal distribution $\mathcal{N}(\mu = 1, \sigma = 1)$. The label *Uniform Sphere* refers to a uniform sphere that is also normally distributed but around zero $\mathcal{N}(\mu = 0, \sigma = 1)$. In order to keep this analysis comparable I set all target state vectors to the same norm $\|\vec{x}(t_f)\|_2 = 1$. It can be seen in fig. 4.6a that the higher the variation amongst the nodal states, the more control energy is required.

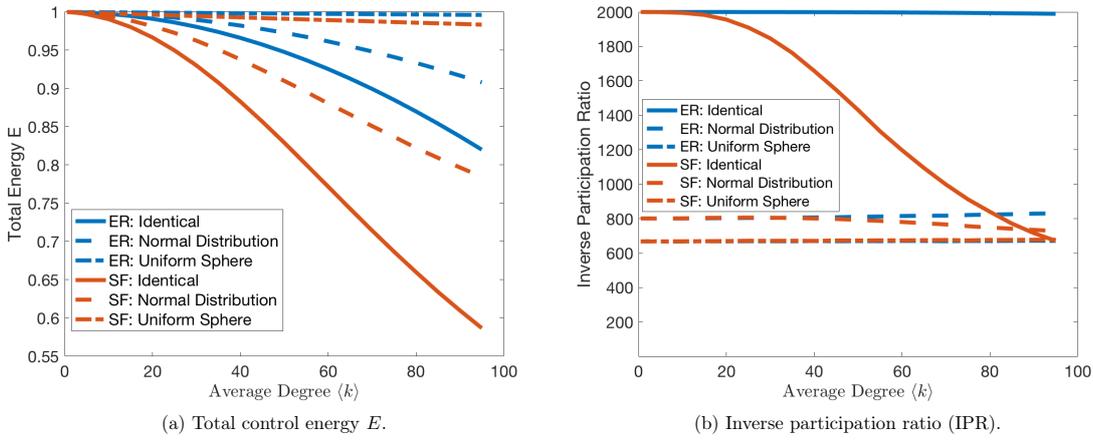


Figure 4.6: Total network control energy E and inverse participation ratio (IPR) as a function of the average network degree $\langle k \rangle$ for various target vector entry distributions. Blue lines reflect results for ER networks and red lines for SF networks. Straight lines reflect the case when each node has an identical target state, dashed lines when the target states are normally distributed around one and dashed-dot lines when target states are distributed in a uniform sphere around zero. When nodal target states deviate as shown by the normal and uniform sphere distributions, the energy requirement increases (L.H.S.) and the nodal energies are much more concentrated, in particular for low $\langle k \rangle$ (R.H.S.).

Figure 4.6b shows that the concentration of the fractional nodal energies $\frac{E_i}{E}$ for smaller $\langle k \rangle$ is much higher in the case of deviating relative to identical target states. This suggests that target states have a strong influence upon the distribution of nodal energies.

The variation in the fractional nodal energies $\frac{E_i}{E}$ can also be studied by plotting the values directly as a function of nodal degrees k_i across all $\langle k \rangle$. I have previously shown in fig. 4.4a that the nodal energies are essentially unchanged as a function of the degree of nodes in the case of identical targets for ER networks. This is also the case for target states that are distributed normally around one and zero for ER networks. In contrast, I showed for SF networks in fig. 4.4b, that nodal energies scale very strongly as a function of nodal degrees due to the presence of large hub nodes. This is also the case for normally distributed target vectors, albeit at a less strong scale, as the maximal nodal energy fraction of a hub node is $\approx 8 \times 10^{-3}$ as shown in fig. 4.7a and $\approx 4 \times 10^{-4}$ for the uniform sphere depicted in fig. 4.7b, compared to 0.02 in the case of identical target states shown in fig. 4.4b. The fractional energy requirement across nodes is therefore increasingly homogeneous when nodal target states deviate.

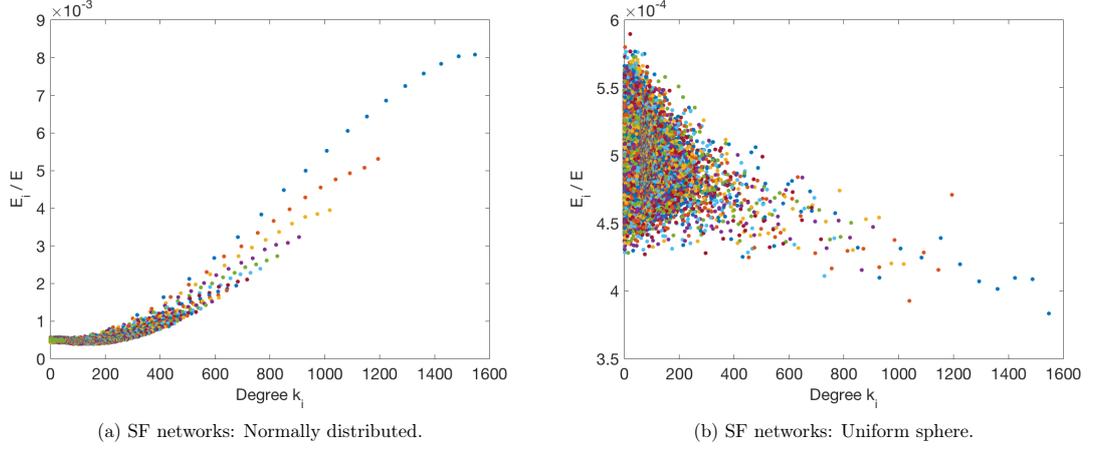


Figure 4.7: The figures represent the expected fractional nodal energy requirement $\frac{E_i}{E}$ for a given degree k_i averaged across all simulations for SF networks for normally distributed target vector entries and the uniform sphere. As opposed to the case of identical nodal target states, shown in fig. 4.4, the distribution of the energy requirements across nodes is less dependent upon k_i . In fact, in the case of the uniform sphere shown in panel (b) nodal energies are much more evenly split across nodal degrees. The colours represent the different average degrees $\langle k \rangle$. I do not show the respective figures for ER networks, as their nodal energies change very little as a function of k_i with respect to the different distributions, see fig. 4.4 as a reference.

The lower dependence of the energy of a particular node upon its degree when nodal target states deviate, can be understood when plotting $\frac{E_i}{E}$ as a function of the nodal target state $x_i(t_f)$. Nodes with a target state of zero require no external input, as nodes naturally move to their stable fixed point at zero. Nodes with a target state of larger magnitude $|x_i(t_f)|$ require more energy, independent of the sign of their target state. When plotting nodal energies as a function of the target states, a tick shape emerges in the case of a normally distributed target vector, as it can be seen in fig. 4.8a for a network with an average degree of $\langle k \rangle = 5$. The tick shape is due to a higher frequency of positive target states in the normally distributed target vector with $\mu = 1$ and $\|\vec{x}(t_f)\|_2 = 1$. As the average network degree increases, the tick shape disappears and transforms into a U-shape, which can be seen in fig. 4.8b for $\langle k \rangle = 95$. The reason for the change in the shape of the behaviour of nodal energies is that with an increasing interaction of nodes, the energy is more evenly spread throughout the network. The increasing homogeneity of nodal energies through the coordination of nodal activities is evidenced by the majority of the observations being in the centre of fig. 4.8b. These plots are very similar for ER and SF networks, for reasons of brevity I therefore only show the results for SF networks. This behaviour is also similar for the uniform sphere and the normal distribution, I therefore only present the results for the normal distribution.

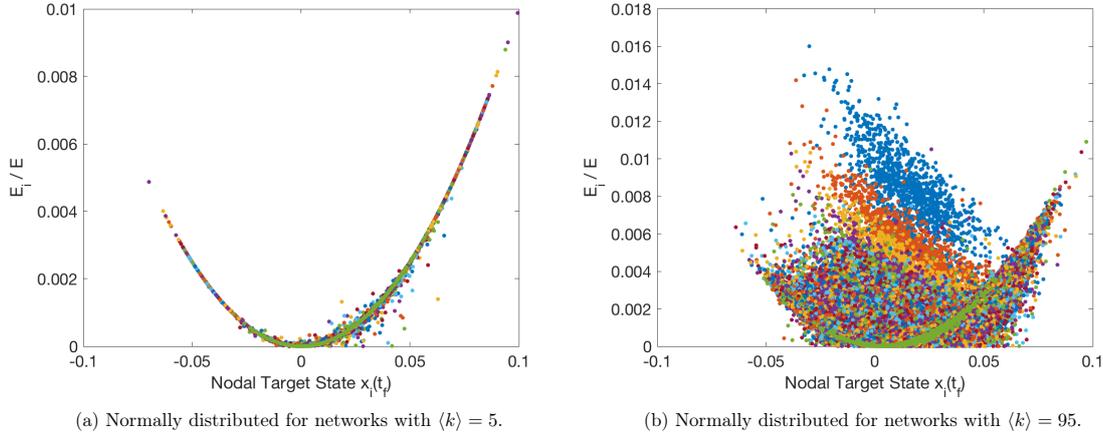


Figure 4.8: Fractional nodal energies $\frac{E_i}{E}$ for each nodal target state $x_i(t_f)$ averaged across all simulations. The figures represent the expected fractional nodal energy requirement for a given nodal target state $x_i(t_f)$ for normally distributed target vector entries. The colours represent the different average degrees $\langle k \rangle$. Figures shown refer to SF networks, the figures for ER networks are very similar.

The behaviour of the fractional nodal energies as a function of the individual nodal target states $x_i(t_f)$ explains the higher total control energy requirement that is observable in fig. 4.6a, when nodal target states deviate. With an increasing deviation of nodal target states the coordination advantage is reduced. When each node has a different target state, endogenous network effects cannot reduce the external network energy requirement, as each node needs to move into a different direction in the state space. In this case the control energy is concentrated in fewer nodes, which results in a significant decrease in the IPR, which can be observed in the dashed(-dot) lines in fig. 4.6b. The reason for this is that nodal energies are the highest for nodes with target states of the largest magnitudes and the lowest for nodes with target states close to their stable fixed point at zero.

4.5 Discussion and Future Developments

In this project I have shown that the cost of controlling a network, measured by control energy, decreases as a function of its average degree. In particular, I have shown that endogenous effects in networks reduce control energy. As the network connectivity increases, nodes are able to push each other towards their common target state. This coordination advantage results in a lower exogenous requirement for control energy, which is amplified in scale-free networks due to the presence of hub nodes. Hub nodes reach throughout the network and are thus able to coordinate the states of a large fraction of the nodes in a network.

It follows from this result that larger SF networks require less control energy than smaller SF networks with target states of the same norm. This is due to much

larger hub nodes in larger networks. The large degree of hub nodes is driven by the power-law nature of the degree distribution in scale-free networks, that emerges in a much stronger fashion in large networks. I note that due to the computational limitations of the control energy, the networks that I study in this project with 2,000 nodes are relatively small. The behaviour of the control energy in much larger scale-free networks may be more stable relative to smaller networks. In ER networks on the other hand, the energy requirement is relatively speaking independent of the network size. Moreover, I find that the coordination advantage is significantly diminished when nodal target states diverge. When each node has a different target state, nodal interactions cannot result in a coordination advantage, because each node is required to move into a different direction in the state space.

According to structural controllability the driver nodes of a network are low degree nodes [86]. In this project my focus on reducing control energy revealed that, especially in scale-free networks, larger degree nodes require the largest fraction of energy but reduce the overall energy requirement. This implies a trade-off between reducing the number of driver nodes, by controlling only nodes of low degree, and a reduction of control energy by controlling only nodes of large degree. This trade-off is of high relevance to applications of network control theory [43] and should therefore be investigated in future studies to identify an optimal control strategy.

Lastly, the analytical derivation that I presented confirms that control energy is reduced when nodal interactions increase. The mean-field analytical solution to the optimal control problem provides a more straight-forward solution to computing the control energy, as compared to the standard methodology based on the controllability Gramian matrix. However, the mean-field solution cannot provide detailed analytical solutions that are potentially achievable by the Gramian matrix. Such solutions include the control of a smaller number of nodes, instead of all nodes, which is a core problem for future work.

The contributions that I have made in this project may find several applications in real networks. In social networks for instance, a small fraction of nodes are frequently more sticky to their opinion than most other nodes; and in interbank networks a subset of nodes is much less resilient than others. For practitioners, for example organisations that desire to influence public opinion or financial regulators, this suggests that some agents, represented as hubs in their networks, should be given the highest focus as they can coordinate large parts of networks.

Chapter 5

Reverse Stress Testing Interbank Networks

5.1 Introduction

5.1.1 Motivation and research objective

In this third project of my thesis I use the insights of the previous two projects to develop a methodology to manage systemic risk. Specifically, I introduce a method to reverse stress test interbank networks. Using this method I reverse engineer the dynamics of stress propagation to identify the minimal exogenous shocks that need to be applied to each bank in the network to cause a specific loss of equity.

The focus of research carried out so far has been mainly to develop models to understand how exogenous shocks are amplified by the endogenous dynamics in an interbank network. I discuss this type of research on financial contagion in section 3.2.1. In contrast, in this chapter I study the reverse problem. I compute the time trajectories of shocks of minimal size that need to affect banks to produce a final loss of equity larger than a given threshold. The solution of this reverse problem is useful to identify stress scenarios that would lead to systemic events and with which it is therefore possible to identify the vulnerabilities of a financial system.

5.1.2 Approach

The reverse stress testing framework is inspired by (network) control theory, which is a methodology from engineering, that has recently been introduced to complex systems [86]. I study control theory in the second project to develop the methodological background for the reverse stress testing framework. In control theory the goal is to drive a system (in this case a network representing interbank lending between

banks) from an initial state to a desired target state (in this case to a minimum level of financial losses) with the least effort (in this case the smallest exogenous shocks to the balance sheets of banks).

After I introduce the methodology of the reverse stress testing framework in the next section, I apply the framework to real data in a case study, in which I reverse stress test a system composed of the 44 largest European banks that are the constituents of the STOXX Europe 600 Banks index. This index is the major equity benchmark of the most significant financial institutions in Europe. For each bank I collected from Bloomberg data on total interbank lending, total interbank borrowing and equity capital levels. I use the RAS algorithm to reconstruct the matrix of interbank exposures, which I outline in section 3.3.4. I then compute the trajectories of shocks of minimal size needed to reach a final level of distress within a given time horizon under a linear model of distress propagation, using the DebtRank algorithm that I describe in section 3.3.2. I choose this contagion algorithm because of its simplicity and because it can be considered as a first-order approximation for a more generic class of contagion algorithms [17, 19].

5.1.3 Contributions

The main contributions and findings of this project are:

- introduction of a framework to reverse stress test interbank networks;
- as the largest eigenvalue of the matrix of interbank leverages increases the optimal shocks become smaller and concentrate in a smaller set of banks;
- the distribution of optimal shock sizes is identified, thus providing a ranking of banks in terms of their systemic importance;
- this ranking can be used to make the system more robust through the implementation of targeted equity capital requirement policies.

5.2 Related Work

Reverse stress testing is a regulatory requirement in the United Kingdom (UK) and the European Union at the level of individual institutions. The Financial Conduct Authority (FCA), one of the UK's financial regulators, describes it as complementary exercise to general stress and scenario testing in its handbook (section SYSC 20). In standard stress testing a forward-looking methodology is employed, in which scenarios are selected to predict their potential impact on the financial health of

banks. Reverse stress testing on the other hand looks backward by identifying the scenarios that cause a specific loss to a bank. This way of identifying stress scenarios is the major advantage of reverse stress testing. Instead of relying on the judgement of experts to select scenarios, the most dangerous scenarios are identified via the methodology.

Previous work on reverse stress testing has focused on developing frameworks for the risk analyses of individual institutions and not of financial systems as a whole [69]. Some studies are dedicated to optimising scenario selection and defining probability distributions of the numerous intertwined driving variables across asset classes. For two recent reviews see Glasserman et al (2015) [64] as well as Flood and Korenko (2015) [55], and for a mathematical approach to worst case scenario selection see McNeil and Smith (2012) [92]. In brief, these papers find that a decrease in economic activity in a country and a steeper interest rate curve are the most risky scenarios for banks. Moreover, given the large data requirements of reverse stress tests, there is a large estimation risk in reverse stress testing [69].

5.2.1 Measuring the contribution of banks to systemic risk

The reverse stress testing framework that I develop in this chapter, is inspired by (network) control theory, I explain this connection in section 5.3.2. In this section I review recent work that applies control theory to interbank networks. In particular, I review the idiosyncratic contribution of each bank to the overall systemic risk in a network, which provides an essential background to the policy experiment that I develop in section 5.5.4, which is based on banks' individual systemic contributions. For a more general literature review on systemic risk and control theory, see section 2.1 and section 4.2, respectively.

Research concerned with controlling interbank networks has so far only considered their structural controllability, that defines the topological features that need to be present in order for a network to be controllable. A conceptual introduction to the application of control theory to interbank networks has been provided by Galbiati et al (2013) [60]. On a large payments network Galbiati et al find that there is no correlation between the feedback centrality of a bank, which measures the average centrality of the neighbours of a node, and whether it is a driver or not. A driver in network control theory is a node that receives an external control input, for instance an exogenous financial shock, and propagates this input to its neighbours in the network. Using data of the e-MID interbank market the authors found that the average degree of driving nodes is lower than the average degree of the network, in all periods under investigation from 1999 until 2009. Interestingly,

the share of the largest lenders who are drivers increases significantly in the run-up to the financial crisis. This suggests an increasing concentration of risk and a clustering of core banks, as found previously in a study by Squartini et al (2013) [125], that focused on the impact of clustering in interbank networks on systemic risk; and that I describe in section 3.2.2.

Delpini et al (2013) [47] find, in agreement with Galbiati et al (2013) [60], that driver nodes are neither the most interconnected hubs nor the largest lenders in interbank networks. The number of drivers as a fraction of the total network size N_D , is an efficiency and resilience indicator for control. N_D decreases with a higher time aggregation of bilateral exposures in the network. Monthly time aggregation is optimal in order to reduce N_D to $< 30\%$, whereas $N_D = 60\%$ when exposures are aggregated on a daily basis. Within smaller time windows relationships with many small borrowers appear, which tend to have a smaller connectivity and therefore increase the required number of drivers [47].

Contribution of banks to systemic risk in financial network theory literature

Independently of control theory, there are different methods that study the contribution of individual banks to the systemic risk of an interbank network. One such strand of the systemic risk literature was discussed in section 2.1.2, where I focused upon the robust-yet-fragile nature of networks. To briefly recap, it has been shown that the connectivity of Erdős-Rényi (ER) and scale-free (SF) networks are relatively stable with respect to a random removal of nodes [2]. However, when the nodes with the largest degrees are removed first, then SF networks disintegrate much faster relative to ER networks, due to the significant role of large degree nodes in coordinating SF networks. Roukny et al (2013) apply this insight to interbank networks by comparing the impact of random versus targeted shocks and varying other properties, such as the leverage of banks [114]. It was shown in the study that the effect of shocking the most connected banks in an interbank network is very similar in ER and SF networks when markets are liquid. In liquid markets banks can recover a large fraction of the value of non-performing, i.e. likely to default, debt by selling it in the market. However, when markets are illiquid, such that non-performing debt can only be sold with large losses –so called fire-sales, then the impact of targeted attacks is much larger in SF relative to ER networks [114].

Another strand of the systemic risk literature is based on a game theoretic analysis of the contribution of individual lenders to systemic risk. In these studies the contribution of a lender is measured by its relative weight of the overall interbank lending volume and the riskiness of its borrowers. For an introduction to this lit-

erature see Drehman and Tarashev (2013) [49]. Glasserman et al (2014) provide boundaries on the probability of the failure of a specific bank and estimate the impact of its failure upon the network [65]. By comparing the relative lending proportion of a bank in a network, which is the potential loss caused by its default to the equity of its neighbours, the probability of it defaulting other banks can be recursively measured. The authors confirm their method in an empirical study based on the data generated by the European Banking Authority 2011 Stress Test. Interestingly, Glasserman et al show that the boundaries on the probability of default hold independently of the network topology.

Lastly, Thurner et al (2013) build an agent-based model of banks that interact in an interbank network [131]. The authors study a normal mode, which replicates typical behaviour observed in real interbank networks, such as a heavy-tailed distribution of losses. A transparent mode is also studied, in which lenders have an incentive to extend loans only to low-risk banks. In this mode the DebtRank of each bank in the network is computed and published at each time step to the other banks. Lenders are incentivised to lend to banks in an increasing manner of their DebtRank. In this mode risk is significantly reduced, primarily because the risk is now much less concentrated in the network. Importantly, in the transparent mode the overall lending volume is identical to the normal mode, the only difference being that risks are now more homogeneously spread throughout the network. This study shows that systemic risk is significantly influenced by banks that have high levels of leverage.

5.3 Methods

5.3.1 Reverse stress testing problem set-up

In this section I outline the methodology that underlies the reverse stress testing framework. A system of n banks is studied, which interact through a network of mutual exposures (interbank assets and liabilities), via a dynamical setting in which the equity of banks is updated in discrete time-steps. It is assumed that a bank holds in its portfolio external assets that are not part of the modelled network in addition to interbank assets.

$A_{ij}(t)$ denotes the value of the exposure of bank i to bank j at time t , $E_i(t)$ is the equity of bank i at time t , $A_i^{\text{ext}}(t)$ the value of external (i.e. non-interbank) assets of bank i at time t , and finally L_i describes the liabilities of bank i , which are assumed to be constant over time. A further assumption is that banks do not rebalance their portfolio (i.e. the number of shares they own of an asset is assumed

to be constant), so that the changes in the balance sheet of a bank are only due to changes in the price of the bank's assets.

From the balance-sheet identity it follows that

$$E_i(t) = \sum_j^n A_{ij}(t) + A_i^{\text{ext}}(t) - L_i. \quad (5.1)$$

I consider a situation in which the value of external assets is subject to random market fluctuations, while the value of the interbank assets of a bank at time t depends on the equity of its counterparties at time $t - 1$. Following Battiston et al [23, 26] I assume that the relative devaluation of an interbank asset is proportional to the relative devaluation of the equity of the counterparty:

$$\frac{A_{ij}(t) - A_{ij}(0)}{A_{ij}(0)} = \beta \frac{E_j(t - 1) - E_j(0)}{E_j(0)}, \quad (5.2)$$

where β is a positive constant related to the strength of contagion between counterparties. Therefore the equity of bank i evolves in discrete time according to

$$E_i(t) = \beta \sum_j^n \frac{A_{ij}(0)E_j(t - 1)}{E_j(0)} + A_i^{\text{ext}}(t) - L_i. \quad (5.3)$$

Following Battiston et al [23] I define $h_i(t) = \frac{E_i(0) - E_i(t)}{E_i(0)}$ and $\Lambda_{ij} = \frac{A_{ij}(0)}{E_i(0)}$, so that

$$h_i(t) = \beta \sum_j^n \Lambda_{ij} h_j(t - 1) + \frac{A_i^{\text{ext}}(0) - A_i^{\text{ext}}(t)}{E_i(0)}. \quad (5.4)$$

The quantity Λ_{ij} represents the importance for bank i of its interbank asset associated with bank j measured in terms of i 's equity. For instance, if the value of the interbank asset drops by 1%, bank i would experience a loss of Λ_{ij} % of its equity. For this reason, Λ_{ij} is referred to as the matrix of interbank leverages [23].

I further define $u_i(t) = \frac{A_i^{\text{ext}}(0) - A_i^{\text{ext}}(t)}{E_i(0)}$, which represents the contribution to the relative equity loss of bank i due to shocks to its external assets between times 0 and t , so that

$$h_i(t) = \beta \sum_j^n \Lambda_{ij} h_j(t - 1) + u_i(t). \quad (5.5)$$

The interpretation of time t in this external asset shock scenario could for instance be days and the time horizon T could refer to the time of maturity of a liability or time scales in a regulatory framework.

I now imagine a situation in which I want to reverse stress test the system over

a time horizon T . In particular, I assume to be at time $t = 0$ and I look for the trajectories of shocks $\{\vec{u}(1), \vec{u}(2), \dots, \vec{u}(T)\}$ to external assets that can lead at time T to losses equal or greater than a given threshold, i.e. such that

$$h_i(T) = \sum_{t=1}^T \sum_{j=1}^n \beta^{T-t} (\Lambda^{T-t})_{ij} u_j(t) \geq \ell_i, \quad (5.6)$$

with $i \in \{1, 2, \dots, n\}$, and where ℓ_i is the threshold associated with the loss of bank i .

There are many possible trajectories that satisfy the constraints (5.6); here I am interested in identifying those that minimize fluctuations of relative losses on external assets over time, i.e. for which the following quantity is minimized:

$$K \equiv \sum_{i=1}^n \sum_{t=1}^T (u_i(t) - u_i(t-1))^2. \quad (5.7)$$

The cost function K can be interpreted as the aggregate size of the exogenous shock affecting the system. I do not make a distinction between positive or negative shocks. I select this cost function because I seek to identify the smallest exogenous shocks that should hit the system to cause a given final loss. I refer to these shocks as "worst case shocks". I note also that the underlying DebtRank algorithm is linear, such that a shock of 1% or 10% has a linear effect on the size of the resulting losses.

In summary, the following optimisation problem is to be solved to reverse stress test a system

$$\begin{aligned} \min & \left(\frac{1}{2} \sum_{i=1}^n \sum_{t=1}^T \Delta u_i(t)^2 \right), \\ \text{s.t.} & \sum_{t=1}^T \beta^{T-t} (\Lambda^{T-t})_{ij} u_j(t) \geq \ell_i \quad \forall i. \end{aligned} \quad (5.8)$$

where $\Delta u_i(t) = u_i(t) - u_i(t-1)$ and it is assumed that $u_i(0) = 0$ for all i . $\Delta u_i(t)$ represents the loss due to shocks on external assets experienced by bank i between times $t-1$ and t . The optimisation problem can be more conveniently written in terms of the cumulative variable Δu as

$$\begin{aligned} \min & \left(\frac{1}{2} \sum_{i=1}^n \sum_{t=1}^T \Delta u_i(t)^2 \right), \\ \text{s.t.} & \sum_{t=1}^T \beta^{T-t} \sum_{s=1}^t (\Lambda^{T-t})_{ij} \Delta u_j(s) \geq \ell_i \quad \forall i. \end{aligned} \quad (5.9)$$

Note that the reason for why I set up the optimisation as an inequality rather than an equality, is that I am interested in identifying the optimal, i.e. minimal, shock that leads to a given systemic loss. If I were to force each bank to reach a specific targeted loss, then some relatively more unstable banks that experience losses higher than the target losses would require negative losses, i.e. equity injections, to be stabilised to achieve the lower losses that were set as the target losses. This is clearly counter-intuitive in a stress-testing scenario, and I therefore set up the optimisation problem as an inequality.

5.3.2 Link to network control theory

The last expression (5.9) is of a form that can be solved using optimal control theory, which I study in depth in the previous chapter [115, 126]. In fact, the set-up of control problems is similar to the analytical approach to the framework (5.9) that I develop in section 5.6. The standard control dynamics are linear time-invariant (LTI) and evolve as follows:

$$\vec{x}(t+1) = A\vec{x}(t) + B\vec{u}(t). \quad (5.10)$$

As I outlined above, the input $u_i(t)$ of bank i at time t is the cumulative loss of external assets experienced by bank i up to that time. In order to see the similarity of the optimisation problem to that in control theory, note that in control problems $u_i(t)$ is the external input required to drive a dynamical system to its desired state ℓ_i at the end of the control period T . Specifically, $\vec{x} = h$ and $A = \Lambda$. Since all nodes receive external inputs in the reverse stress testing framework, the input matrix B is equal to the identity matrix I and can therefore be omitted. See section 4.2 for a discussion of control theory.

5.4 Analysis of Homogeneous System

In order to develop an intuition on the behaviour of the solutions of (5.9), I first consider the simple case of a homogeneous system in which all banks have the same interbank leverage c , i.e. the matrix Λ is such that $\sum_j \Lambda_{ij} = c$ for all i . In this case all banks are the same, in that they are subject to the same dynamic and to the same constraint. Therefore, by symmetry, also the trajectories of worst case shocks

are the same for all banks. The optimisation problem (5.9) reduces to

$$\begin{aligned} \min \quad & \left(\frac{1}{2} \sum_{t=1}^T \Delta u(t)^2 \right), \\ \text{s.t.} \quad & \sum_{s=1}^T \sum_{t=s}^T \lambda^{T-t} \Delta u(s) \geq \ell, \end{aligned} \quad (5.11)$$

where $\lambda = \beta c$ and $\Delta u(t) = u(t) - u(t-1)$, as before I set $u(0) = 0$ and I have exchanged the order of the sums over t and s ($\sum_{t=1}^T \sum_{s=1}^t = \sum_{s=1}^T \sum_{t=s}^T$).

This problem can be easily solved with the method of Lagrange multipliers, which brings

$$\Delta u(t) = \frac{\sum_{r=t}^T \lambda^{T-r}}{\sum_{s=1}^T \left(\sum_{r=s}^T \lambda^{T-r} \right)^2} \ell \quad (5.12)$$

and

$$K = \frac{\ell^2}{\sum_{s=1}^T \left(\sum_{t=s}^T \lambda^{T-t} \right)^2} \quad (5.13)$$

$$= \frac{(\lambda - 1)^3 (\lambda + 1) \ell^2}{T (\lambda^2 - 1) + \lambda (\lambda^T - 1) (\lambda^{T+1} - \lambda - 2)}. \quad (5.14)$$

From this formula it can be seen that, upon increasing the time horizon T over which stress propagates, the size of exogenous shocks needed to produce the sought final loss progressively reduces and goes to zero in the limit $T \rightarrow \infty$. This is expected, as shocks can reverberate over a longer time horizon, and eventually an infinite sequence of infinitesimal shocks can lead to the final loss ℓ .

However, the behaviour of the cost function for long time horizons shows the existence of two very distinct regimes: When $\lambda > 1$ the cost function approaches zero exponentially as $K \sim \lambda^{-2T} \ell^2$, while when $\lambda < 1$ the cost function decays to zero much more slowly, as $K \sim \frac{\ell^2 (\lambda - 1)^2}{T}$. The reason for this behaviour is that for $\lambda > 1$ shocks are exponentially amplified by the dynamics.

A similar behaviour can be observed for a general matrix of interbank leverages, where the largest eigenvalue λ_{\max} of $\beta \Lambda$ discriminates between the two regimes. I study this case in the following section.

5.5 Case Study Based on Largest European Banks

In the following sections I discuss an empirical application of the reverse stress test framework based on the optimisation problem (5.9) to an interbank system

representing the largest banks in the European Union. I explore the results of this problem of reverse engineering financial contagion via the following variables:

1. the largest eigenvalue λ_{\max} of the interbank leverage matrix, which determines the stability of the dynamics [17]. For simplicity of notation, I refer to λ_{\max} as the largest eigenvalue of the matrix $\beta\Lambda$ and I vary λ_{\max} by changing β . When $\lambda_{\max} < 1$ the dynamics are stable, but when $\lambda_{\max} > 1$ the dynamics are not stable;
2. the minimal financial loss ℓ_i , which is the target state of the optimised dynamics;
3. the time horizon T ;
4. the quantity $K_i = \sum_t^T \Delta u_i(t)^2$, which expresses the size of the exogenous shock experienced by bank i over the time horizon T .

5.5.1 Description of data

I collected data from Bloomberg of the 44 banks belonging to the STOXX Europe 600 Banks index, ticker symbol: SX7P.¹ In particular, for each bank I collected information on its equity, total interbank assets (advances and loans to banks) and total interbank liabilities (deposits due to other banks) for the entire year of 2015.

Figure 5.1 below shows the distribution of the equity capital, interbank assets and liabilities for each bank. The average equity capital is \approx EUR5.1bn and the average interbank assets \approx EUR3.2bn and for interbank liabilities \approx EUR3.7bn. As a reference, these numbers are much larger than those of the Austrian network discussed in the first project. The reason for this is that I here study the largest European banks, whereas in the Austrian network I investigated banks of all sizes.

¹Index constituents as of 11th November 2016.

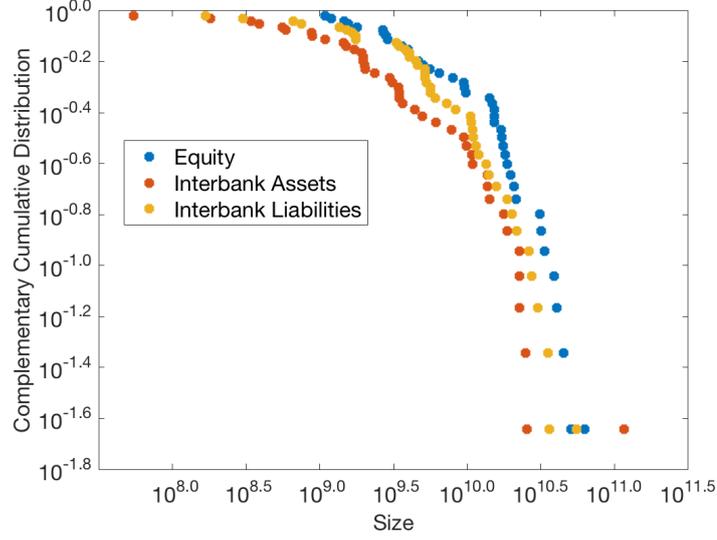


Figure 5.1: Distribution of equity capital, interbank assets and liabilities of each bank in 2015.

On the basis of this data, I used the RAS algorithm to construct a matrix of interbank exposures to represent an interbank lending network, see section 3.3.4 where I outline the algorithm. Starting from total interbank assets and liabilities of each bank, the RAS algorithm allows an allocation of interbank loans across counterparties [13]. If no further constraints are added, the outcome of the RAS algorithm is a complete weighted network of interbank claims. Although real interbank networks are far from complete [32, 54, 56, 125], here for simplicity I consider this limiting case that allows me to focus on the mechanics of reverse stress testing only, rather than on the interplay between network topology and financial contagion. In the next chapter, which describes the fourth project, I study the STOXX network with different densities.

5.5.2 Aggregate properties of optimal shocks

Figure 5.2a shows the behaviour of the cost function K as a function of λ_{\max} for different time horizons and for $\ell_i = 0.1$ for all i . It can be seen that the size of exogenous shocks decreases as a function of λ_{\max} across all T . As λ_{\max} increases the endogenous dynamics of the network lead to a larger amplification of the distress, such that a lower magnitude of exogenous shocks is required to reach the target loss ℓ_i .

For a similar reason larger T result in smaller K independently of λ_{\max} . The endogenous network dynamics propagate the distress of the previous time step, thereby reducing the exogenous shock requirement as T is increased. Because the iteration map (5.5) does not reach a fixed point when $\lambda_{\max} > 1$, even in the absence

of exogenous shocks beyond the first time step (i.e. $u(t) = 0$ for any $t > 1$), it is expected that the size of the exogenous shocks K will go towards zero exponentially fast in the limit $T \rightarrow \infty$. This is indeed the case, as shown in the inset of fig. 5.2a.

The behaviour is qualitatively similar for any value of final losses ℓ , as I show in fig. 5.2b, where I plot the cost function K as a function of λ_{\max} and ℓ , for $T = 20$. I choose this time horizon because the behaviour of K is similar to larger T and its computation is much less intensive. From this figure it can be seen that when λ_{\max} is large enough the shock needed to cause the sought final losses is relatively independent of ℓ , while it increases with ℓ when $\lambda_{\max} < 1$.

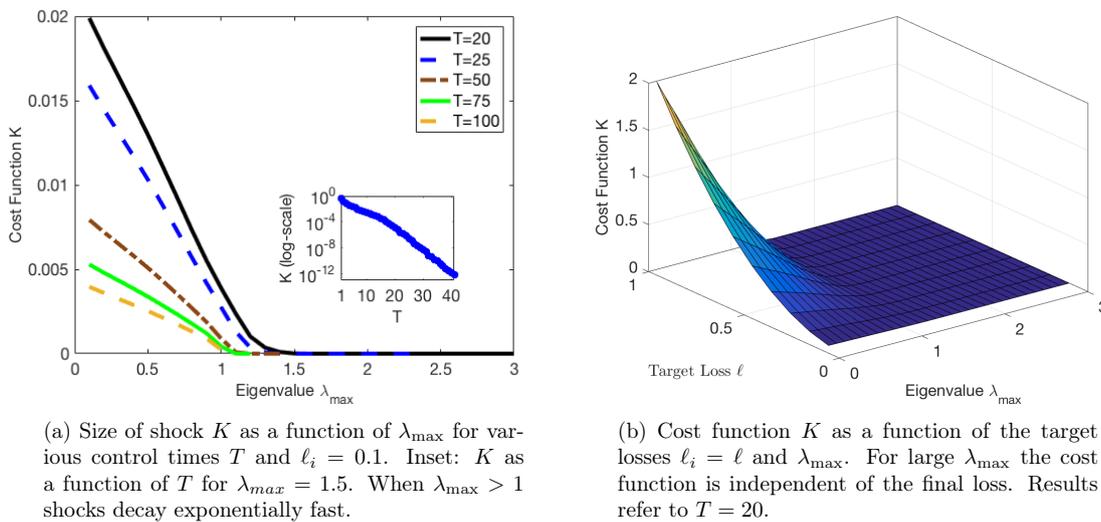


Figure 5.2: Size of shock K in STOXX network.

Figure 5.3 provides a similar picture of these results. It shows what I refer to as the loss fraction, which is the total amount of equity that is shocked externally, or lost, as a fraction of the total equity in the system to cause a specific loss ℓ . Since none of the edge weights (exposures) are negative in the network, this fraction always stays below one. The loss fraction indicates the extent to which the endogenous network dynamics amplify the exogenous shocks. This amplification increases as a function of λ_{\max} and thus results in lower required exogenous shocks K_i for each bank.

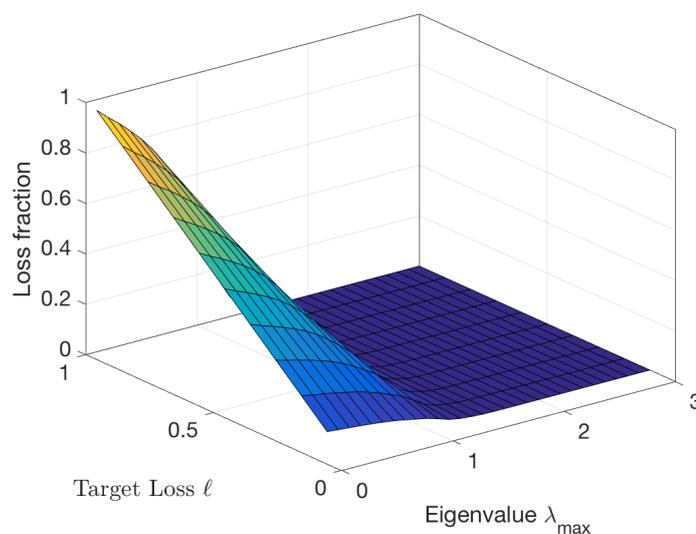


Figure 5.3: The loss fraction measures the total amount of equity externally shocked as a fraction of the total equity in system. Shown as a function of the target losses $\ell_i = \ell$ and λ_{\max} . The loss fraction always remains below one, indicating that the endogenous network dynamics propagate the losses.

5.5.3 Concentration of risk

I have so far looked at the properties of the aggregate shock, however the proposed methodology allows me to obtain the distribution of shocks across banks in the system. This information is useful as it enables me to rank banks in terms of their contribution to the aggregate shock, and to identify potential concentrations of vulnerability in the system. If the optimal aggregate shock is uniformly distributed across all banks, then I would expect the system to be more resilient with respect to idiosyncratic failures of individual banks. Although the system might be vulnerable with respect to common factors that affect the portfolios of banks. If the shock is instead highly concentrated in a few banks, the system is vulnerable with respect to the failure of these banks.

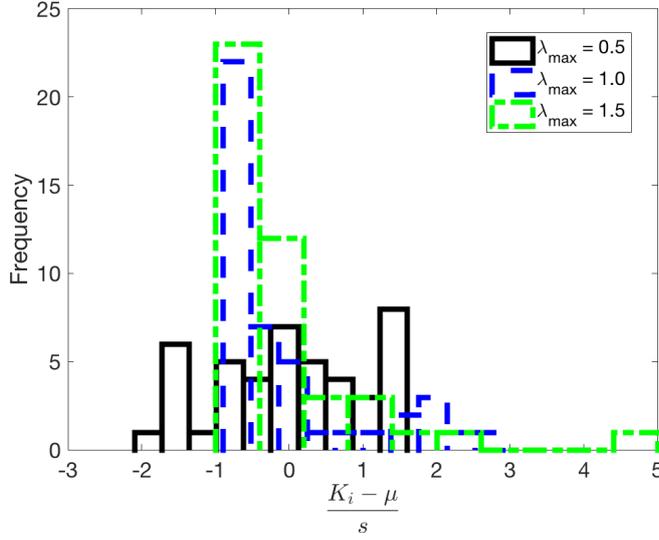


Figure 5.4: Distribution of the standardized shock of each node (z-score) for different values of λ_{\max} . Shown for $T = 20$ and for $\ell_i = 0.1$. s is the standard deviation and μ the mean of the respective distribution of nodal shocks K_i for each λ_{\max} . As λ_{\max} increases larger relative deviations with respect to the mean amplitude are observed in the right tail of the distribution.

Figure 5.4 shows the distribution of standardized shocks, i.e. the z-score of each nodal shock K_i , across all banks for three different values of λ_{\max} : 0.5, 1, and 1.5. The z-score is defined as $\frac{K_i - \mu}{s}$, where s is the standard deviation and μ the mean of the respective distribution of nodal shocks K_i . As it can be seen from the figure, the distribution of shocks is strongly affected by λ_{\max} . In particular, I observe that shocks appear to become more concentrated for higher values of λ_{\max} . Indeed, for higher values of λ_{\max} I observe a stronger correspondence of the z-score to the total interbank liabilities of each bank i , i.e. $\sum_{j=1}^n A_{ji}$. In fact, the largest z-score of 5 seen in fig. 5.4 corresponds to the bank with the largest interbank liabilities.

This concentration of systemic risk can be quantified by computing the inverse participation ratio (IPR), defined as

$$\text{IPR} = \frac{1}{\sum_{i=1}^n p_i^2}, \quad (5.15)$$

where $p_i = \frac{K_i}{K}$ for each node i . The IPR has a lower bound of 1 when the shock is concentrated in one node, and an upper bound of n when the shock is equally spread across all nodes.

As it can be seen in fig. 5.5a, the IPR is unaffected by ℓ , it decreases significantly as λ_{\max} approaches 1, and it becomes constant for $\lambda_{\max} > 1$. Figure 5.5b shows that the IPR is relatively unaffected by the time horizon T .

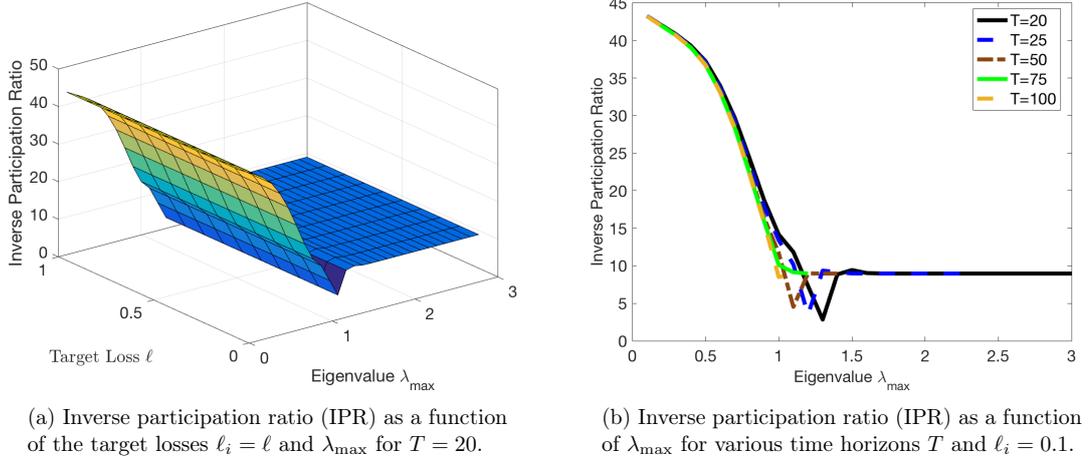


Figure 5.5: Inverse participation ratio (IPR) in the STOXX network. The nodal shocks become more concentrated as λ_{\max} increases.

I have shown that increasing λ_{\max} leads to a reduction of the aggregate size of the shock K needed to drive the system towards a certain loss and to a concentration of nodal shocks upon a smaller set of banks. I stress here that these two behaviours have different roots. The reduction of K is due to the fact that the system becomes more unstable as leverage increases. The concentration of risk is due to the heterogeneity of leverage across banks. In fact, in the homogeneous system considered in section 5.4 this concentration does not occur. The heterogeneity of banks' leverages, shown in fig. 5.6, suggests that some banks have lent significantly more than others relative to their equity capital. Some banks therefore contribute to the risk of the system much more than other banks, which leads to the concentration of shocks observed in fig. 5.5a. The interbank leverage of a bank is defined as the ratio between its total interbank assets and its equity. The interbank leverage of bank i is therefore calculated as $\sum_{j=1}^n \Lambda_{ij}$.

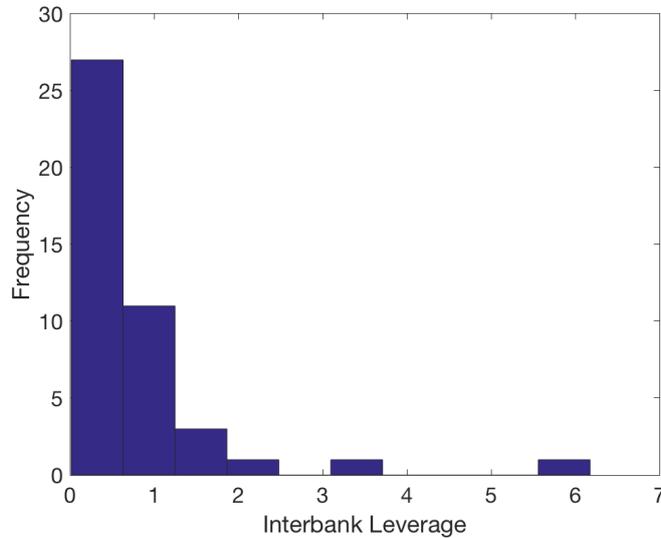


Figure 5.6: Histogram of interbank leverages of each bank in the STOXX network under investigation. It can be seen that lenders have heterogeneous strategies for their interbank leverage, defined as the ratio of the total amount lent (assets) over a bank’s equity capital.

The observed concentration of risk as measured by the IPR is one of the most significant findings of this project. It indicates that as the strength of the interactions between nodes becomes stronger, a decreasing number of nodes in the system can cause systemic events. In other words, while for $\lambda_{max} < 1$ a large fraction of nodes need to experience stress in order for the system to observe significant losses, when $\lambda_{max} > 1$ only a few nodes are sufficient to cause significant losses. In the next section I will use this insight to devise a policy experiment to reduce the observed system losses.

5.5.4 A simple policy experiment

In the following experiment I intend to show that the systemic vulnerabilities in the high λ_{max} regime can be associated with a small set of banks, in which the aggregate shock is concentrated. I run the dynamic (5.5) forward whilst applying the shocks that have been computed through the reverse stress test framework to a subset of the nodes (i.e. using as $u(t)$ the result of the optimisation (5.9)). I then compute the final loss R observed in the system divided by the final loss observed when all banks are stressed at once R_0 and plot this ratio as a function of the fraction of stressed banks.

Figure 5.7 shows the result of this experiment for different values of λ_{max} when the banks that are stressed are selected in a descending order of their nodal shock K_i . I also show in the circled lines the results for the benchmark case, in which the

stressed banks are randomly selected. As expected, deviations from the benchmark case become larger as the system becomes more unstable. The concentration of systemic risk in the system can be seen particularly when $\lambda_{\max} = 1.5$ (black line), in which case the exogenous shock of 5 banks can lead to roughly 70% of all observed final losses. Note that when banks are randomly selected then no such concentration is observed (circles in equivalent colours in the figure).

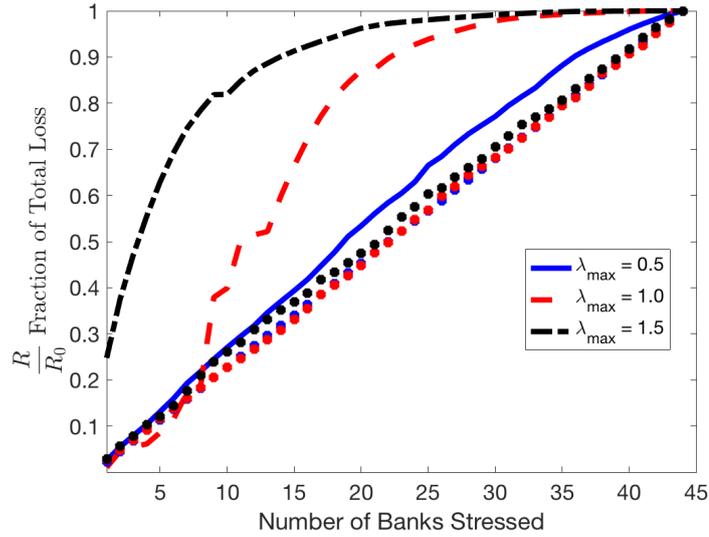


Figure 5.7: Incremental addition of $u_i(t)$ to banks in a decreasing order of the size of their individual shocks K_i for various λ_{\max} . The circles represent the average across 500 simulations in which banks were randomly selected (same colours as for the three lines corresponding to the three different λ_{\max} shown in the legend). For $\lambda_{\max} = 1.5$ the shock on the first five banks already accounts for 70% of final losses. Results shown for $T = 20$ and $\ell_i = 0.1$. As noted before, the behaviour is very similar for larger values of T .

Set-up of equity capital policy

The previous insights can be used to design a policy for the equity requirements of individual banks, that aims to reduce the observed financial losses under the scenario identified through the reverse stress test framework.

I use the contribution of each bank to the aggregate shock to rank banks in terms of their systemic impact. In particular, I consider a situation in which I increase the total capital in the system by 5% and a policy by which this additional capital is spread across banks proportionally to the size of the shock computed from the reverse stress test, i.e. bank i receives a proportion K_i/K of the total additional capital. I then compare this policy with a benchmark according to which the equity of each bank is increased by 5%. This benchmark mimics the case of a homogeneous (relative) increase in the capital requirement of banks. For both policies, I compute the total relative losses $R = \sum_i^n h_i(T)$ via DebtRank under the scenario identified

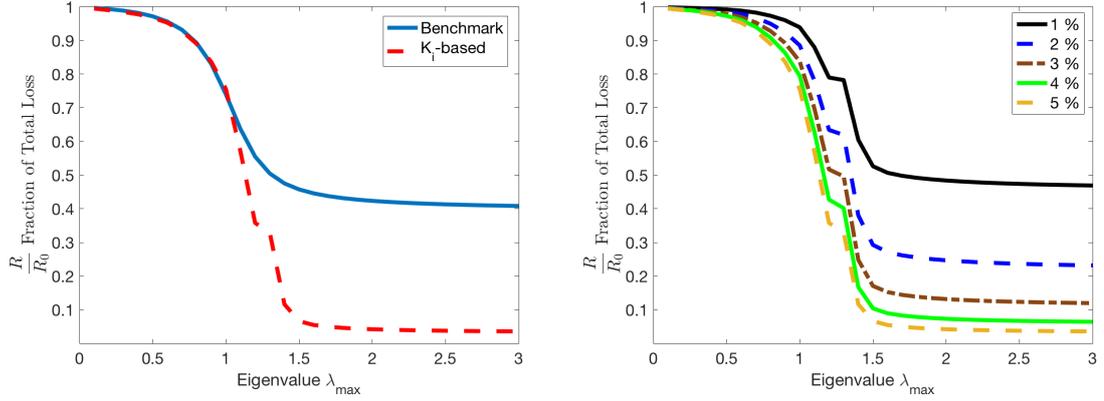
through the reverse stress test and compare it with the total losses observed in absence of policy intervention R_0 .

Currently, the world-wide Basel III capital regulations stipulate a minimum level of equity capital of 4.5% of a bank's total assets plus an additional 2.5% conservation buffer of a bank's total risk-weighted assets. Moreover, national regulators may require an additional capital increase of up to 2.5% if an increase in risk is observed across the system [22]. Put together these requirements are referred to as the Common Equity Tier 1 ratio. The definition of risk-weighted assets is thus a crucial factor in determining the level of equity capital. In brief, each asset held by a bank is given a factor based on its riskiness, which scales its weight in the aggregation of total risk-weighted assets. My policy approach of increasing the equity capital is thus different to current regulations. The focus of this section is to identify whether this difference is warranted, by testing the effectiveness of my policy in stabilising the system.

Note that my suggested 5% capital increase is not unrealistic. The European Banking Authority (EBA) reported after its 2016 stress test exercise of the 51 largest banking institutions in Europe, that the average Common Equity Tier 1 ratio increased from 8.9% at the end of 2010 to 13.2% at the end of 2015. This corresponds to an increase of $\approx 50\%$ of the equity capital across these banks [52]. Additionally, there is a large overlap between the banks involved in the stress tests of the EBA and the STOXX index used in this project. In fact, 10 out of the 15 European banks classified as global systemically important banks in 2015 are in the top 15 of my ranking of banks based on their systemic importance K_i [53].

Results of policy experiment

The results of this policy exercise are shown in fig. 5.8a. As it can be seen in fig. 5.8a when $\lambda_{\max} < 1$ the benchmark and K_i -based policies achieve a similar reduction of total losses, while the K_i -based policy becomes more effective when $\lambda_{\max} > 1$. The reason for this result is that the explosive dynamics when $\lambda_{\max} > 1$ lead to a concentration of systemic risk in a few banks on which an effective policy should concentrate.



(a) Comparison of two policies to reduce the observed total financial losses. In both policies an additional equity capital of 5% was allocated in different manners to each bank. When $\lambda_{\max} < 1$ the benchmark policy is slightly more effective than the K_i -based policy, however when $\lambda_{\max} > 1$ the K_i -based policy is significantly more effective.

(b) Results of financial loss reduction for the K_i -based policy as a function of λ_{\max} for different amounts of equity capital injected into the system. The percentage in the legend indicates the increase of the total equity capital of the system. The effectiveness of increasing capital injections has rapidly vanishing returns of scale.

Figure 5.8: Results of policy experiments. Losses R are recomputed after the equity was increased and expressed as a fraction of the original losses R_0 on the y-axis. These results are shown as a function of λ_{\max} along the x-axis. Results shown for $T = 20$ and $\ell_i = 0.1$.

Figure 5.8b shows that a larger increase in the additional capital to be allocated to the equity base of each bank results in a further reduction of losses. The figure shows this for the policy based on my ranking of nodes in terms of their systemic importance K_i . However, the impact of an increased equity allocation has decreasing returns of scale in terms of the reduction of financial losses. The impact of an increase from 1% to 2% is much larger than the impact of a 4% increase relative to that of 5%. This behaviour occurs for the benchmark policy as well, but in this case it is not as pronounced, for this reason I do not show the result separately. This is due to the fact that the K_i -based policy is much more effective in allocating the additional equity as compared to the benchmark, and therefore leads to a stronger reduction of losses that cannot be significantly improved with even larger equity capital injections.

The decreasing returns of scale of the policy are of high relevance to financial regulators. Higher capital requirements come at large costs to banks, not only due to the costs involved in raising the additional capital, but also in terms of lower returns. The fact that equity capital requirements beyond a certain level lead to drastically diminished marginal reductions in risk, points towards an optimal trade-off between the costs of increasing equity capital requirements and their effectiveness in reducing risk. In other words, reducing the leverage of the most risky lenders is sufficient to significantly increase the stability of an interbank market. This result corresponds to the insights provided by Poledna et al (2016), who show that a

systemic risk tax would force banks to lend money more homogeneously [109]. More homogeneous lending practices do not reduce the overall interbank volume, but avoid large cascades of bank failures by reducing the largest leverages. See also section 5.2.1 where I discuss similar studies.

5.5.5 Robustness of results

To test the robustness of the results presented in this chapter I collected data of the previous year (2014) and following year (2016) to compare to the simulations in this project, which are based on 2015 data. I chose these years as the financial conditions in the period from 2014 to 2016 are similar, and thus make the robustness test more straightforward. The results are very similar across the years. The total shock K and its properties such as the IPR are almost indistinguishable across the three years. Furthermore, the ranking of nodes is very robust across time, such that the results of the equity capital policy are very similar in 2014 and 2016 compared to the 2015 results presented in this chapter. In fig. 5.9 I show the relation of the nodal shocks for each bank across the three years in scatter plots. The correlation between nodal shocks across different years is always larger than 0.9. The strong similarity between the years is due to the stability of the equity, interbank assets and liabilities of the individual banks.

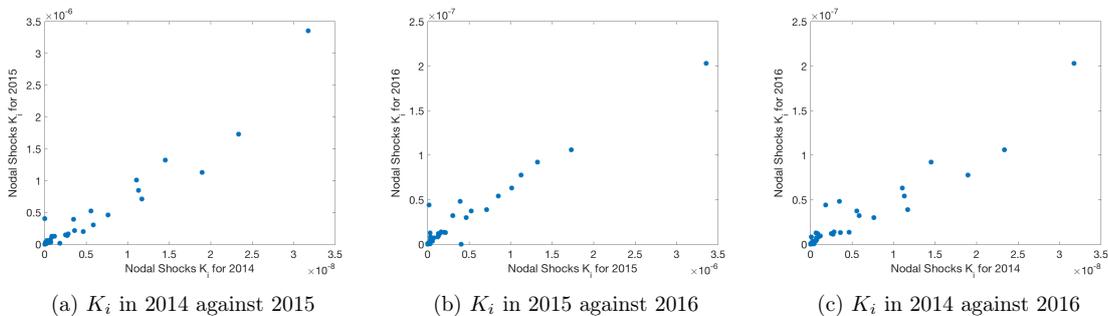


Figure 5.9: Scatter plots comparing nodal shocks K_i for each bank across three years 2014, 2015 and 2016. It can clearly be seen that the nodal shocks are very similar across the years. Results shown for $T = 20$ and $\ell_i = 0.1$.

5.6 Analytical Approach to the Optimisation Problem

To provide a deeper intuition for the reverse stress testing framework, I present an analytical derivation to the optimisation problem that underlies the framework.

I restate the main optimisation problem (5.9), as follows

$$\begin{aligned} & \min \left(\frac{1}{2} \sum_{i=1}^n \sum_{t=1}^T \Delta u_i(t)^2 \right), \\ \text{s.t. } & \sum_j \sum_{s=1}^T \sum_{t=s}^T \beta^{T-t} (\Lambda^{T-t})_{ij} \Delta u_j(s) \geq \ell_i, \quad \forall i, \end{aligned} \quad (5.16)$$

I exchanged the order of the sums over t and s ($\sum_{t=1}^T \sum_{s=1}^t = \sum_{s=1}^T \sum_{t=s}^T$) and made the aggregation over each node j explicit to simplify the following derivation.

I set $D(t) = (\beta\Lambda)^{T-t}$, the function to be minimised is then:

$$\mathcal{L} = \frac{1}{2} \sum_{i,t} \Delta u_i(t)^2 - \sum_i \eta_i \left(\sum_j \sum_{s=1}^T \sum_{t=s}^T D_{i,j}(t) \Delta u_j(s) - \ell_i \right) \quad (5.17)$$

where η_i refers to the Lagrange multipliers for each node i , used to enforce the constraints of the function \mathcal{L} . The partial derivatives are:

$$\begin{cases} \frac{\partial \mathcal{L}}{\partial \Delta u_k(r)} = \Delta u_k(r) - \sum_i \eta_i \sum_{t=r}^T D_{i,k}(t) = 0 \\ \frac{\partial \mathcal{L}}{\partial \eta_i} = \sum_j \sum_{s=1}^T \sum_{t=s}^T D_{i,j}(t) \Delta u_j(s) - \ell_i = 0. \end{cases} \quad (5.18)$$

Where

$$\Delta u_k(r) = \sum_i \eta_i \sum_{t=r}^T D_{i,k}(t). \quad (5.19)$$

Inserting $\Delta u_k(r)$ into $\frac{\partial \mathcal{L}}{\partial \eta_i}$ results in

$$\sum_j \sum_{s=1}^T \sum_{t=s}^T D_{i,j}(t) \sum_e \eta_e \sum_{r=s}^T D_{e,j}(r) = \ell_i. \quad (5.20)$$

To simplify this expression into matrix notation I set

$$\Phi_{i,j}(s) = \sum_{t=s}^T D_{i,j}(t), \quad (5.21)$$

such that I arrive at

$$\sum_e \sum_{s=1}^T \sum_j \Phi_{i,j}(s) \Phi_{e,j}(s)^+ \eta_e = \ell_i, \quad (5.22)$$

where $+$ denotes the transpose. I additionally set

$$W_{i,j} = \sum_{s=1}^T \sum_k \Phi_{i,k}(s) \Phi_{j,k}(s)^+. \quad (5.23)$$

$W_{i,j}$ in turn allows me to set

$$\eta_i = \sum_j^n W_{i,j}^{-1} \ell_j. \quad (5.24)$$

I thus finally arrive at the solution for the inputs

$$\Delta u_k(r) = \sum_i^n \Phi_{i,k}(r) \sum_j^n W_{i,j}^{-1} \ell_j \quad (5.25)$$

$$= \left(\Phi(r)^+ W^{-1} \vec{\ell} \right)_k. \quad (5.26)$$

It can be seen that the solution to the optimisation problem of the reverse stress testing framework can be neatly summarised in the W operator. The matrix W serves a similar function as the Gramian matrix in control problems, described in eq. (4.11).

Conditions for a successful solution to the optimisation problem In order for a solution to this linear inequality constrained optimisation problem to be found, the Karush-Kuhn-Tucker (KKT) conditions need to hold. For the simplified problem of

$$\begin{aligned} \min \left(\sum_i f(x_i) \right) & \quad (5.27) \\ \text{s.t.} \quad \sum_i g(x_i) & \geq c_i, \end{aligned}$$

where f is a generic function, g a generic constraint function and c a constraint, all used as a general representative case. The function to be minimised is

$$\mathcal{L}(x_i, \lambda) = f(x_i) - \sum_i \lambda_i (g(x_i) - c_i), \quad (5.28)$$

the KKT conditions are

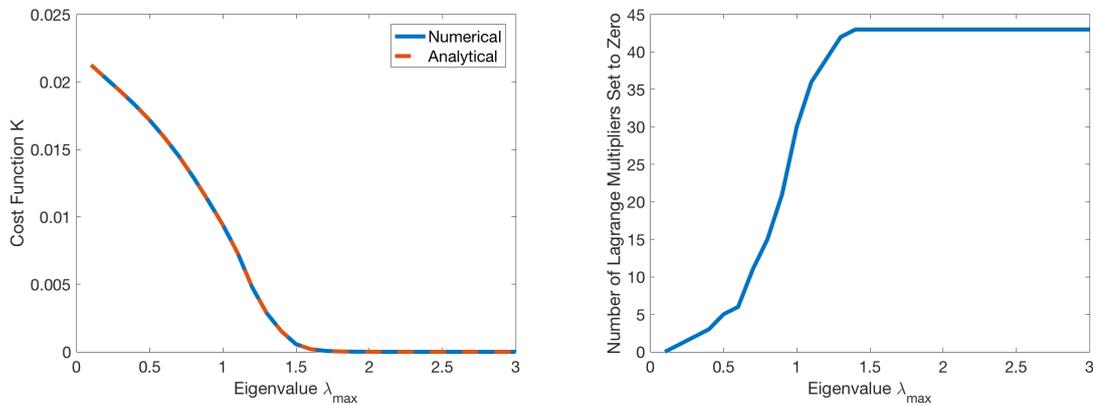
1. $\frac{\partial \mathcal{L}}{\partial x_i} = 0 \quad \forall i$;
2. $\lambda_j \geq 0 \quad \forall j$;
3. $g(x_i) \geq c_j$ or if $\lambda_j = 0$ then $\lambda(g(x_i) - c_j) = 0$.

5.6.1 Comparison of analytical and numerical results

In this section I present the results of a numerical solution to the optimisation problem (5.16), to cross check my analytical derivation (5.26). It can be seen in

fig. 5.10a that the result of the analytical solution for the cost function K agrees with the result of the numerical solution for the optimisation problem.

As λ_{max} increases the number of Lagrange multipliers η that are set to zero increases, indeed when $\lambda_{max} > 1.1$ all but one η are set to zero, as I show in fig. 5.10b. The reason for this is that when λ_{max} is large, the dynamics are very unstable, such that an infinitely small perturbation to the system causes most nodes to overshoot their target ℓ_i . In other words, the analytical solution asymptotes zero, because the constraint ℓ_i is non-binding (slack) for most nodes. Their Lagrange multipliers η therefore have to be set to zero according to the KKT conditions described above. η_i (5.24) in other words expresses the exogenous shock that is required to cause bank i a certain fraction of equity losses.



(a) Cost function K as a function of λ_{max} for the numerical (blue) and analytical solution (red). It can be seen that the analytical solution agrees with the numerical simulations.

(b) Number of Lagrange multipliers η set to zero as a function of λ_{max} . The nodes for which η_i is set to zero overshoot their target ℓ_i , for which the constraint is thus non-binding.

Figure 5.10: Results from analytical solution (5.26) for exogenous shocks K .

5.7 Discussion and Future Developments

In this chapter I introduced a framework to reverse engineer distress propagation in financial networks. I reversed the standard stress testing approach by setting a specific outcome, the loss of a certain fraction of the equity of each bank, and looking for the scenario with smallest shocks that could lead to such an outcome over a given time horizon.

I considered a system of interbank borrowing relationships based on 2015 annual data of the equity, interbank lending and borrowing of the 44 largest stock exchange listed European banks. My analysis of the preceding and following years shows that the insights resulting from this case study are robust across years. I found that at the aggregate level the size of the optimal exogenous shock decreases as the largest eigenvalue λ_{max} increases. At the same time the shock becomes concentrated in

a smaller number of banks. Based on the optimal pathways of shock propagation identified in my analysis, I ranked banks in terms of their systemic impact. Using this ranking I suggested a simple policy of capital allocations that significantly reduces the vulnerability of the system, with respect to the identified scenario in the regime of high endogenous amplification.

Furthermore, I derived an analytical solution of the reverse stress testing framework, which can be applied to other networks and be used to obtain analytical proofs of the behaviour of the total exogenous shocks. The main insight of the analytical solution is that the reverse stress test dynamic can be summarised in the W operator, whose inverse provides the nodal shocks of each bank.

5.7.1 Future developments

My analysis can be improved in several directions to widen the applicability of the reverse stress testing framework and to explore different channels of contagion.

First of all, I used a simple linear dynamical rule of distress propagation. Although common in the literature of financial contagion, this assumption can at best be considered an approximation of the true dynamics. In a more general case, it is still possible to write an optimisation problem analogous to (5.9) to perform the reverse stress test. The main difference with respect to the case considered here, would be the presence of a non-linear constraint, but the optimisation problem could still be solved numerically.

Secondly, I only considered the DebtRank algorithm as a means of stress propagation in interbank networks. Other important channels of contagion operate in interbank networks, without which a complete understanding of the propagation of risk in a network is not possible. I discuss these channels of contagion in section 3.2.

A third limitation of my analysis is the fact that I only considered direct long exposures between banks. Banks interact in many ways in the real world, and a more realistic scenario would consider a multilayer description of the network of interbank interactions. In this respect, my present analysis corresponds to an aggregation of the multilayer structure into a single layer [110, 71, 82]. However, it would be important to look also at the disaggregated multilayer structure to reflect different loan maturities, because the properties of aggregated and non-aggregated systems have been shown to differ in some cases [71].

Fourth, I considered the case of banks as passive investors. This is certainly a useful benchmark, but a more realistic scenario would also account for the reaction of banks to changing market conditions. For instance banks normally adjust their equity capital levels and interbank activities following the default of one of their

borrowers.

In spite of all the present limitations, this analysis suggests that reverse stress testing is a useful tool for the identification of vulnerabilities at the systemic level, and I believe this is an interesting avenue for future investigation with potentially relevant policy implications.

Chapter 6

Reverse Stress Testing Interbank Networks - The Influence of Topology

6.1 Introduction

6.1.1 Motivation and approach

In the previous chapter I introduced the reverse stress testing framework and focused on its dynamics rather than the influence of network topology on the framework. In this project I test the impact of network topology on the framework by investigating how the behaviour of the exogenous shock changes in interbank networks with different topologies. It is of particular interest to study the performance of the equity capital policy, which aims to reduce the observed system losses.

I study two different systems. Firstly, I apply the reverse stress testing framework to the STOXX network with different densities. The network describes the interbank exposures between the 44 largest European banks that are the constituents of the STOXX Europe 600 Banks index. Secondly, I apply the framework to the Austrian interbank network, which I also study in the first project of this thesis. The Austrian interbank network with 846 nodes is much larger than the STOXX network and also more heterogeneous. I have information on the bilateral exposures between individual banks in the Austrian network, whereas I construct the STOXX network using the RAS algorithm with information only of the total interbank assets and liabilities of each bank. This detailed information on the Austrian network supports the representativeness of my results.

With the insights of the comparison between the STOXX and Austrian networks in hand, I can explicitly study the ability of null models to reconstruct the reverse stress testing dynamics. In particular, I test three different network null models to estimate the size of the exogenous shock of the Austrian network. I consider Erdős-

Rényi, the directed configuration (rewired) and fitness null models. Within studies of systemic risk, network null models are used to reconstruct interbank networks from limited information. As I discuss in section 2.2, the reconstruction of networks is of interest as it is often difficult to obtain complete information on the interbank lending and borrowing activities between financial institutions. This lack of information can be due to privacy concerns and the lead time of financial reporting, amongst other reasons [97]. Practitioners thus regularly work with reconstructed networks to study financial contagion [7, 9, 23, 41, 48, 95, 98, 123].

In this project the reconstruction exercise allows me to investigate the ability of null models to estimate a new stress testing dynamic and to further study the influence of network topology on the reverse stress testing framework. I put a particular emphasis upon the fitness model, as it performed well in the first project and in experiments in the literature [23, 41, 61, 123]. Given the breadth of this experimental approach, this chapter presents a synthesis of all previous projects in this thesis.

6.2 Density Comparison across STOXX Networks

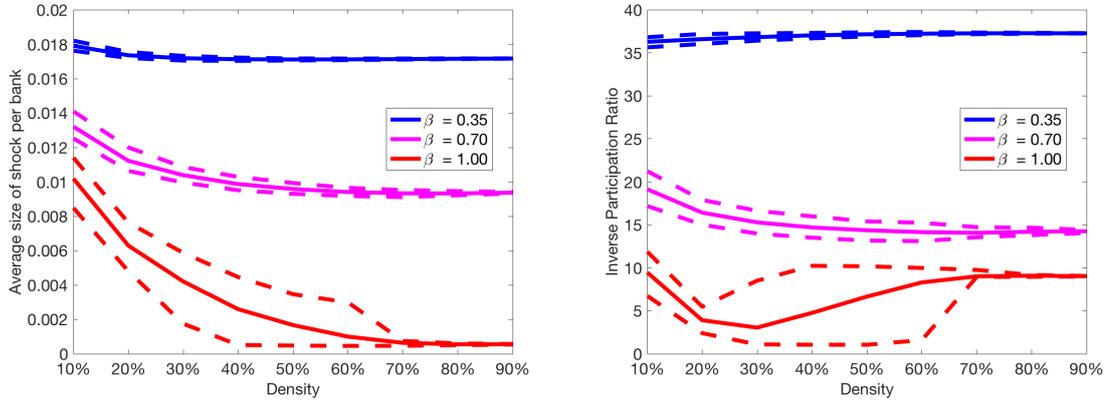
In this section I present the results of an analysis of the STOXX network with different densities. I generated the STOXX networks of varying densities by tuning the density parameter of the RAS algorithm based on the same data of the full year 2015 interbank assets and liabilities of the 44 largest European banks, which are the constituents of the STOXX Europe 600 Banks index. For each network density I created 100 simulations and I present the averaged results across these simulations. For an introduction to the RAS algorithm see section 3.3.4 and for a description of the STOXX network see section 5.5.1. I describe the methodology underlying the reverse stress test framework in section 5.3.

In order to keep the results comparable across different densities, I express the results as a function of β , which scales the strength of the interactions between nodes in the matrix of interbank leverages Λ , measured by $\beta\lambda_{max}$ as I describe in section 5.3. λ_{max} is the largest eigenvalue of Λ , which changes with the network topology.

6.2.1 Exogenous shock across densities

I find that the connectivity of a network, in this analysis expressed by the network density, has a strong effect on the properties of the exogenous shock K that is applied via the reverse stress test framework. I study the average exogenous shock $\sqrt{K/n}$

instead of the total exogenous shock, as it allows me to compare the size of the shock across different network sizes. As it can be seen in fig. 6.1a when β is large, the average exogenous shock decreases as the density of the STOXX network increases. The pace of the decrease of the exogenous shock accelerates as β increases. This indicates that networks of higher connectivity are more risky in terms of the reverse stress testing framework, as a smaller shock results in the same losses. There is an extensive literature on the effect of network connectivity on financial stability that I discuss in section 3.2.2.



(a) For larger β the average exogenous shock decreases as a function of the density, whereas it is almost unchanged for smaller values of β .

(b) Dip in the IPR for densities $< 50\%$ when β is large, i.e. when $\lambda_{max} > 1$, due to poorly connected nodes.

Figure 6.1: Average exogenous shock $\sqrt{K/n}$ and inverse participation ratio (IPR) as function of network density for various β . The blue lines refer to $\lambda_{max} \approx 0.5$, pink to $\lambda_{max} \approx 1.0$ and red to $\lambda_{max} \approx 1.5$. Shown for target loss $\ell_i = 0.1$ and time horizon $T = 20$. Dashed lines show the 10% and 90% quantiles, respectively. As it can be by the tight dashes lines, the results are stable across all simulations.

I find that a large fraction of the decrease of the exogenous shock K seen in fig. 6.1a is due to the disappearance of isolated nodes defined as $k_i^{in} = 0$, when the density increases. With a decreasing density the likelihood increases that nodes do not have an incoming link. For the purposes of reverse stress testing this means that these nodes cannot be influenced by the endogenous network dynamics. The nodal shock K_i of low degree nodes is thus relatively independent of β , such that low degree nodes require the largest K_i . I define $K_i = \sum_t^T \Delta u_i(t)^2$ as the size of the exogenous shock that affects bank i over the time horizon T . I show the dependence of the average exogenous shock on the number of isolated nodes in fig. 6.2a. The figure shows that in networks with a density of 10% and $\beta = 1.0$, the average exogenous shock is much larger when the network has a higher number of isolated nodes. In the case of networks with a density of 10%, all of the 100 simulated networks have at least four isolated nodes, in comparison in networks of density 50% only 9 networks have one isolated node.

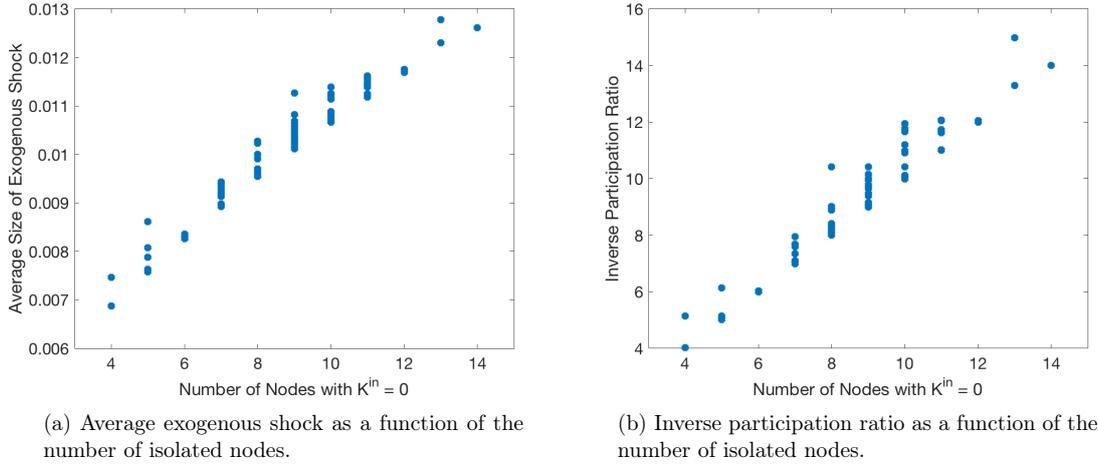


Figure 6.2: Average exogenous shock $\sqrt{K/n}$ and inverse participation ratio (IPR) for each of the 100 generated STOXX networks with density 10% and $\beta = 1.0$ ($\lambda_{max} = 1.5$). Isolated nodes ($k^{in} = 0$) are a core driver of the average exogenous shock and the IPR.

To understand the distribution of shocks across nodes I compute the inverse participation ratio (IPR). The IPR is an indicator of the concentration of nodal shocks and is computed as

$$\text{IPR} = \frac{1}{\sum_{i=1}^n p_i^2}, \quad (6.1)$$

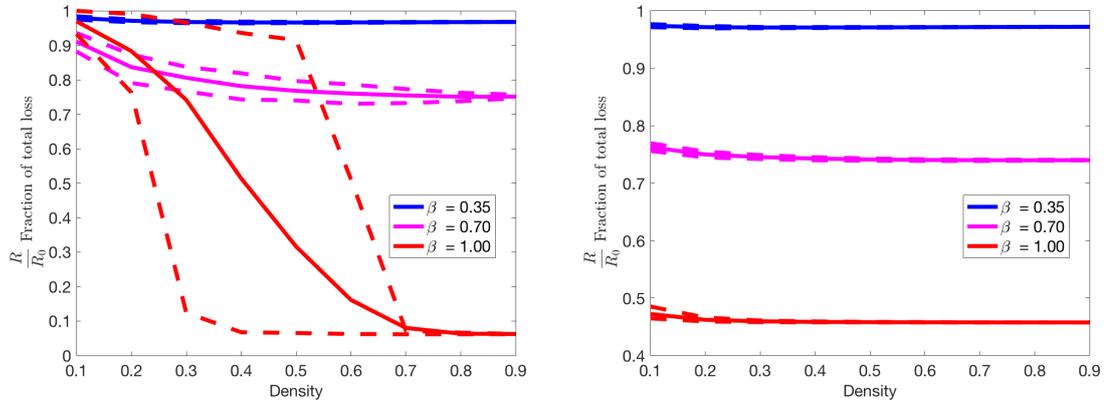
where $p_i = \frac{K_i}{K}$ for each node i . The IPR has a lower bound of 1 when the shock is concentrated in one node, and an upper bound of n when the shock is equally spread across all nodes. As with the exogenous shock K , the IPR is also unchanged as a function of the density for $\beta < 0.7$ ($\lambda_{max} < 1$) as it can be seen in fig. 6.1b. However, for $\beta > 0.7$ ($\lambda_{max} > 1$) a dip occurs for densities lower than 50%. This dip is due to poorly connected nodes taking up most of the nodal shocks, which drastically increases the concentration of shocks when the density is 10%. Once the network density increases, the IPR can return to its natural rate of concentration at ≈ 9 , as it can be observed in the full STOXX network in fig. 5.5a. Figure 6.2b shows that the IPR linearly increases as a function of the number of isolated nodes for $\beta = 1.0$ ($\lambda_{max} = 1.5$) in networks with a density of 10%.

6.2.2 Equity capital policy

I now use the concentration of nodal shocks in an equity capital policy to reduce the observed system losses. The equity capital policy, that I introduce in section 5.5.4, is based on a ranking of banks in order of their nodal shocks K_i . The core mechanism of the policy is that the total equity capital in the system is increased by 5% and the additional equity capital is allocated to banks in proportion of their nodal shock

K_i/K . After allocating the additional capital, I reapply the reverse stress test framework to recompute the total cumulative losses $R = \sum_i^n h_i(T)$ and compare this to the original losses R_0 before I increased the capital.

As it can be seen in fig. 6.3a the effectiveness of the K_i -based policy depends on the density of the network. When the density is larger than 50% the policy works well. Indeed, its performance is similar to that in the complete STOXX network, shown in fig. 5.8a. However, as the density tends towards zero the policy works increasingly less well. The reason for the worse performance of the policy is due to poorly connected nodes and in particular nodes with zero in-degree. As noted in the previous section, when the network density is low the likelihood increases that nodes are isolated. The effect of isolated nodes is particularly apparent when $\beta > 0.7$ ($\lambda_{max} > 1$), since most nodes only require a very small exogenous shock, whereas isolated nodes require a much larger shock. Isolated nodes are therefore allocated most of the additional equity, which is allotted according to the relative fraction of nodal shocks K_i/K . These isolated nodes also tend to have small outgoing links both in terms of number and weight, they can therefore not effectively propagate the additional equity, which leads to the decreased effectiveness of the policy.



(a) Comparison of the K_i -based policy as a function of the network density for selected β . The policy works less well for densities lower than 50%, as most of the additional equity is allocated to isolated nodes.

(b) Comparison of the benchmark (relative) policy as a function of the network density for selected β . Policy works well across all densities as each node receives an equal (relative) increase in its equity capital.

Figure 6.3: Comparison of policies across different densities of the STOXX network. Results shown for target loss $\ell_i = 0.1$ and time horizon $T = 20$. Dashed lines show the 10% and 90% quantiles, respectively. As it can be by the tight dashes lines, the results are stable across all simulations.

This point becomes clearer in comparison with the relative (benchmark) policy shown in fig. 6.3b, which performs equally well across all densities. In the case of the benchmark policy each node receives an equal relative increase of its equity capital of 5%, such that the well-connected nodes can propagate the additional equity to reduce the observed losses.

The effectiveness of any policy, as measured by the reduction of total losses, thus crucially depends on the ability of the system to reach each node in order to propagate the additional equity capital. Note however that the mere addition of incoming edges with minimal weight to those nodes with $k^{in} = 0$, does not improve the effectiveness of the K_i -based policy. In order for the additional equity to be propagated by these nodes, their edge weights need to be of a similar size to the average edge weight in the network.

6.3 Comparison to Austrian Network

In this section I describe the results of applying the reverse stress testing framework to the Austrian interbank network introduced in the first project. Note that I here study data from the first quarter of 2006. The Austrian network is much larger with $n = 846$ nodes compared to $n = 44$ in the STOXX network. Additionally, whilst I had to estimate the individual interbank assets and liabilities in the STOXX network, I obtained the complete interbank relationships between banks in the Austrian network.

6.3.1 Reverse stress testing the Austrian network

When comparing the behaviour of the average exogenous shock $\sqrt{K/n}$ and the IPR (6.1) in the Austrian network and the STOXX networks as a function of λ_{max} , a similar behaviour can be observed. In particular, K increases with the target losses ℓ and decreases with λ_{max} , as it can be seen in fig. 6.4a. Moreover the IPR is independent of ℓ and drops rapidly as λ_{max} increases, see fig. 6.4b. I do not show the comparison across time T , as the behaviours are very similar to those shown for the STOXX network in the previous chapter in fig. 5.2a and fig. 5.5b.

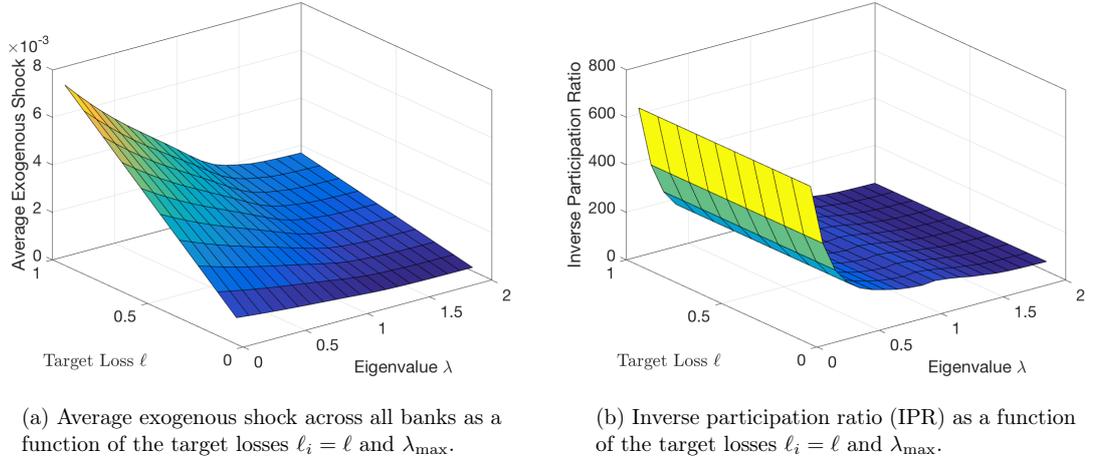


Figure 6.4: Average exogenous shock $\sqrt{K/n}$ and inverse participation ratio (IPR) for the Austrian network. Results shown for $T = 20$.

Despite these similarities between the Austrian and STOXX networks, it can be seen upon closer inspection that the distribution of exogenous nodal shocks shown in fig. 6.5, is more heterogeneous in the Austrian relative to the STOXX network, which is shown in fig. 5.4. A core reason for this difference is the larger underlying heterogeneity of the nodes in the Austrian network, which leads to a relatively higher concentration of exogenous shocks in the Austrian network. The larger heterogeneity in the Austrian network is due to the fact that in the STOXX network I have information only on the subset of the largest European banks, whereas in the Austrian interbank network I have information on banks of all sizes. I measure the difference in the concentration between the two networks with the relative decrease of the IPR rather than the absolute size of the IPR, as the Austrian network has much more nodes than the STOXX network. In the STOXX network when $\lambda_{\max} = 1.5$ the IPR is a fraction of 0.21 of the IPR measured when $\lambda_{\max} = 0.1$. In contrast for the Austrian network this fraction is lower at 0.09. Exogenous nodal shocks are thus relatively more concentrated in the Austrian than the STOXX network. Overall, it can also be seen that the average exogenous shock in the Austrian network shown in fig. 6.4a is smaller than that of the STOXX networks shown in fig. 6.1a. As a reference, the density of the Austrian network is 2%.

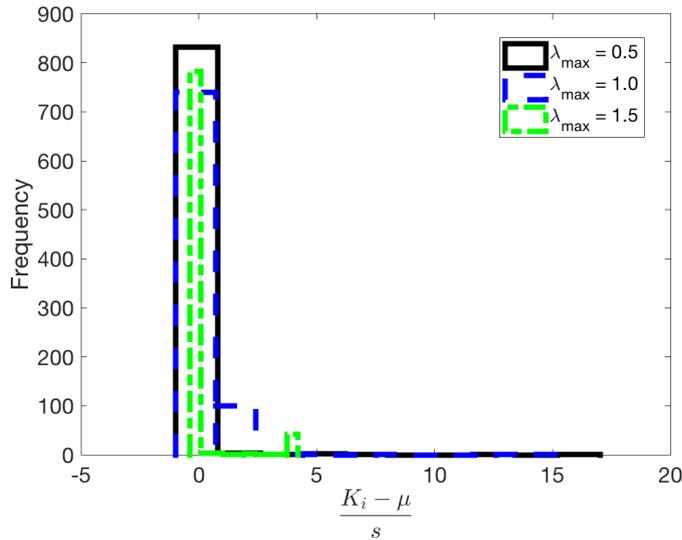


Figure 6.5: Distribution of the standardized shock of each node (z-score) for different values of λ_{\max} in the Austrian network. Results shown for $T = 20$ and for $\ell_i = 0.1$. s is the standard deviation and μ the mean of the respective distribution of nodal shocks K_i for each λ_{\max} . As λ_{\max} increases larger relative deviations with respect to the mean amplitude are observed in the right tail of the distribution.

6.3.2 Policy experiment

Using the distribution of nodal shocks K_i I create a ranking of banks in the identical manner as in section 5.5.4, that is build on the concentration of nodal shocks. I run the dynamics described in equation (5.5) forward and apply the shocks $u_i(t)$ that where computed via the reverse stress test framework (5.9) at each time step to banks in a decreasing order of K_i . At the end of each iteration I compute the final losses R observed in the system and divide these by the final losses observed when all banks are stressed at once R_0 .

Figure 6.6 shows the result of this experiment for different values of λ_{\max} . Similarly to the STOXX network shown in fig. 5.7, it can be seen that as λ_{\max} increases a large fraction of the observed total losses is achieved with relatively few nodal shocks added to the network. In comparison to the STOXX network this behaviour is more extreme in the Austrian network, as around 5 banks account for almost all of the final losses, whereas in the STOXX network the respective fraction is 70%. This difference is due to the higher concentration of nodal shocks in the much larger Austrian network, that I described at the end of the previous section and is shown in fig. 6.5.

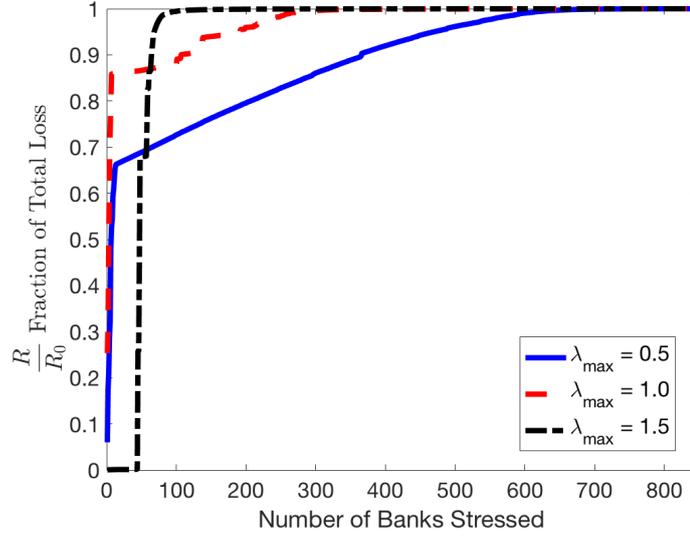
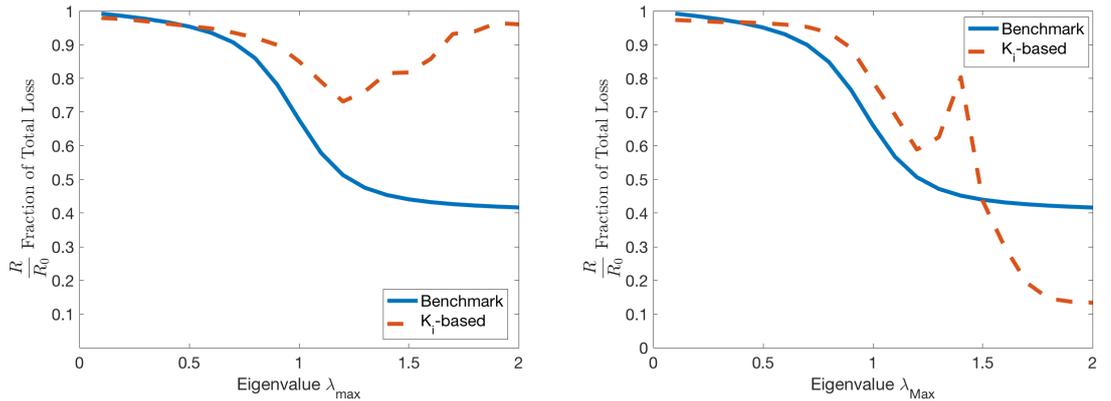


Figure 6.6: Incremental addition of $u_i(t)$ to banks in a decreasing order of the size of their individual shocks K_i for various λ_{max} . Results shown for $T = 20$ and $\ell_i = 0.1$.

Equity capital policy

Turning to the equity capital policy, it can be seen in fig. 6.7a that the relative (benchmark) policy, in which each bank receives a 5% increase of its equity capital, performs equally well across the STOXX and Austrian networks. In both networks the policy leads to a reduction of losses of approximately 60%.



(a) Comparison of different policies to reduce the observed total financial losses in the full Austrian interbank network.

(b) Comparison of different policies for the Austrian subnetwork with $n = 225$ in which no node has $k^{in} = 0$ and $k^{out} = 0$.

Figure 6.7: Equity capital policies applied to the Austrian network. The K_i -based policy in the subnetwork (R.H.S.) is much more effective than in the full network (L.H.S.).

The K_i -based policy in the Austrian network is faced with a similar problem as in the STOXX network. In the Austrian network with $n = 846$ the policy is much less effective relative to the benchmark policy and is indeed ineffective when

$\lambda_{max} > 1$. This, again, is due to the presence of isolated nodes with $k^{in} = 0$, which command almost all of the exogenous shocks and are therefore allocated almost all of the additional equity. These isolated nodes cannot effectively propagate the additional equity due to their small number of outgoing links and generally lower edge weights. This problem is solved when employing the policy on the Austrian subnetwork with $n = 225$, in which no node has $k^{in} \leq 1$ and $k^{out} \leq 1$. In this case the K_i -based policy works similarly well as in the STOXX network by reducing the observed total losses by 90%, as shown in fig. 6.7b.

Investigation of the effectiveness of the policy

Despite the increased effectiveness of the K_i -based policy in the subnetwork, a spike still occurs in the policy's ability to reduce the losses at $\lambda_{max} \approx 1.5$, as it can be seen in fig. 6.7b. To show that the effectiveness of the policy is related to connectivity of the network, I study the increasing k-cores of the Austrian subnetwork. The k-core is a subnetwork in which all nodes with degree smaller than k are removed. It can be seen in fig. 6.8 that the spike disappears when rerunning the policy experiment on increasing k-cores of the Austrian subnetwork. For the 5-core with $n = 161$ the spike decreased and shifted to a higher λ_{max} and for the 6-core with $n = 147$ the spike disappeared.

In summary, when the equity is not adequately allocated by the network, the effectiveness of the equity capital policy is drastically reduced. This analysis shows that having incomplete information on the local topology of a network can significantly impair the ability of policy-makers to reduce systemic losses.

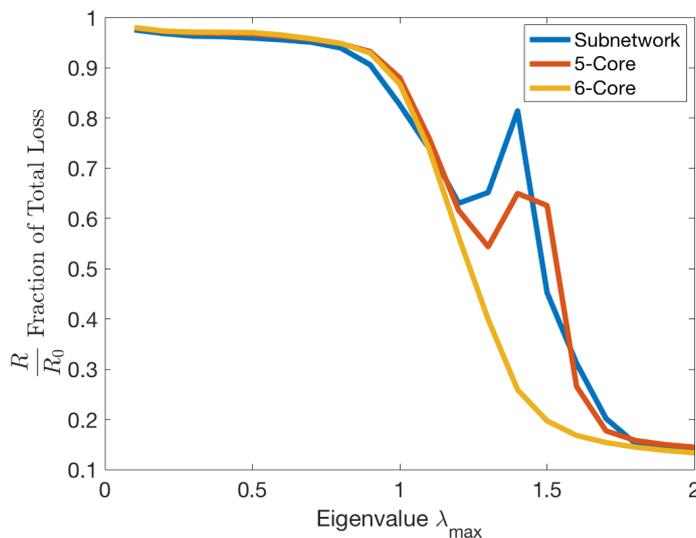


Figure 6.8: K_i -based policy for the subnetwork with $n = 225$ in comparison to the 5-Core with $n = 161$ and the 6-Core with $n = 147$. As the connectivity increases, indicated by the k-core, the policy becomes more effective.

6.3.3 Robustness tests

To test the robustness of the results presented in this section for the reverse stress testing exercise on the Austrian network of the first quarter of 2006, I performed the same tests based on data of the Austrian interbank network in the first quarter of 2007 and 2008. The results are very similar at the aggregate level of the total exogenous shock K and the policy. I show these results for the first quarter of 2007 in appendix A.3. I was not able to compare the results across the quarters at the nodal level. The information on the Austrian interbank networks that I obtained from the Austrian Central Bank (Österreichische Nationalbank) is anonymised, which makes it impossible to accurately identify individual banks across time. Note however that for the STOXX robustness tests, which I outline in section 5.5.5, I obtained public information from Bloomberg. This allowed me to compare nodal shocks K_i across time and to confirm that also the results at the nodal level are robust.

6.4 Reconstruction of Reverse Stress Testing Dynamics in Austrian Network

In the previous section I have shown that the properties of the reverse stress testing framework are robust in the STOXX network and the Austrian interbank network. Given the similarity of the results of the reverse stress testing framework in both networks, I now investigate whether it is possible to reconstruct the reverse stress test dynamics using network null models. I therefore reconstructed the Austrian network with the Erdős-Rényi (ER), the directed configuration (rewired) and the fitness null models, which I also study in the first project. The results presented in this section are based on 1,000 simulations for each null model.

6.4.1 Estimate of exogenous shock

Figure 6.9 shows the average exogenous shock as a function of β for each of the three null models, with the errorbars indicating the 5th and 95th percentiles, respectively. The performance of the ER networks is the worst of the three null models when $\beta < 0.24$ as they overestimate K , however the ER networks perform better than the fitness networks when $\beta > 0.24$. It can be seen in fig. 6.9 that the fitness networks drastically overestimate K as β increases. The rewired networks, whilst always slightly overestimating K , have the best performance of the three null models. The errorbars are relatively tight for the configuration and fitness model networks and their boundaries do not change the overall picture of the performances of the null

models. It is noteworthy that the errorbars of the ER networks are significantly larger relative to the other null model networks.

The main goal of using null models is to estimate the properties of the real network, including its largest eigenvalue λ_{max} . I therefore present the results in this section as a function of β . By multiplying the null model networks with the β of the real network, instead of supplying the null models with the desired λ_{max} , I can test the ability of the null models to reproduce the λ_{max} of the real network. This is particularly important because λ_{max} changes slightly with the network topology. For a translation of the β values to the corresponding λ_{max} values see table 6.1.

λ_{max}	0.1	0.5	1.0	1.5
β	0.02	0.12	0.24	0.35

Table 6.1: Conversion table for λ_{max} to β in the real network.

As a benchmark I have also plotted the maximum entropy version of the Austrian interbank network in fig. 6.9, which represents the full network version of the Austrian interbank network fitted with the RAS algorithm. It can be seen in the figure that when $\beta < 0.24$ the exogenous shock matches that of the real network. When $\beta > 0.24$ the exogenous shock K is slightly lower in the maximum entropy network. This is due to the higher "connectivity" of the maximum entropy network, which allows the shocks to propagate through the network more easily. This is a behaviour that is commonly observed in maximum entropy interbank networks [95].¹

¹Note that 39 banks in the real network have zero interbank liabilities and 3 banks have zero interbank assets. For this reason the RAS algorithm cannot fit any edges for these banks. I address this problem by running the algorithm only on rows and columns whose marginals are non-zero. The maximum entropy network thus has edges on all entries of the weighted and directed adjacency matrix A except for these nodes and the diagonal A_{ii} for all nodes i .

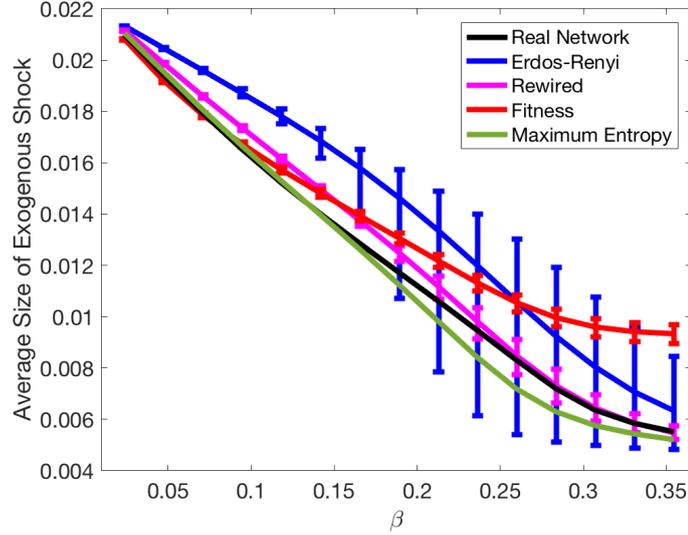


Figure 6.9: Average exogenous shock $\sqrt{K/n}$ for the three null models and the real network. The error bars indicate the 5th and 95th percentiles, respectively. Maximum entropy network plotted as a benchmark. The configuration model performs best out of the three null models, the ER networks perform worse than the fitness networks when $\beta < 0.24$, with this relation being reversed when $\beta > 0.24$. Results shown for $T = 20$ and $\ell_i = 0.1$.

Given the importance of the largest eigenvalue λ_{max} in determining the dynamics, it could reasonably be assumed that the ability of the null models to estimate λ_{max} drives the performance of the null models. To test this I compute the z-score, which measures the deviation of the distribution of λ_{max} in the null model simulations relative to the real network. Surprisingly, the z-score is the lowest for the ER networks at -0.004 , followed by the rewired networks at -0.846 and the fitness model networks at 1.215 . The performance of the null models can thus not be explained by their ability to measure the largest eigenvalue, as the ER networks do not perform best in measuring K , but have the lowest measuring error for λ_{max} .

6.4.2 Concentration of nodal shocks

I find that, a similar picture of the performance of the null models with respect to K is also found for the inverse participation ratio (IPR) (6.1), which is an indicator of the concentration of nodal shocks. It can be seen in fig. 6.10 that the rewired networks somewhat overestimate the IPR and when $\beta > 0.24$ they underestimate it. The fitness networks significantly underestimate the IPR when $\beta < 0.24$, but for $\beta > 0.24$ this reverses as the fitness networks overestimate the IPR. This relatively poor performance is noteworthy as the general behaviour of the IPR is to decrease, i.e. the exogenous shocks become more concentrated as β increases. However, instead the IPR as measured by the fitness networks displays the opposite behaviour

when $\beta > 0.24$.

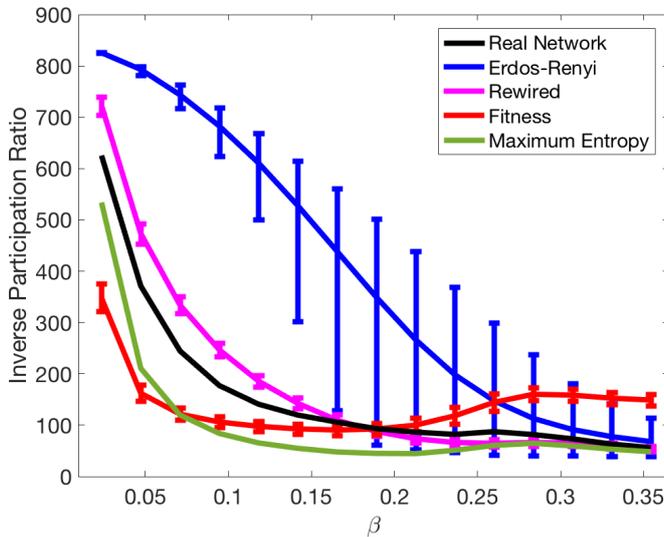


Figure 6.10: Inverse participation ratio (IPR) as a function of β for the three null models and the real network. The error bars indicate the 5th and 95th percentiles, respectively. Maximum entropy network plotted as a benchmark. Results shown for $T = 20$ and $\ell_i = 0.1$.

The poor performance of the fitness model is surprising given that it has achieved good results in the simulations in my first project and in the previous literature [23, 123, 61, 41]. I obtained further insights into the relatively poor performance of the fitness networks by investigating the nodal shocks K_i . The following two plots in fig. 6.11 show the average nodal shock for each bank across the three null models against the nodal shock in the real network for $\beta = 0.12$ and $\beta = 0.35$. It can clearly be seen that the fitness networks have the largest outliers. This is confirmed in table 6.2 below, which shows the root mean squared distance (RMSD) for each null model for three different β . The RMSD is the highest for the ER networks when $\beta < 0.24$, but for $\beta > 0.24$ it is the largest for the fitness networks. The configuration model always performs best across all β .

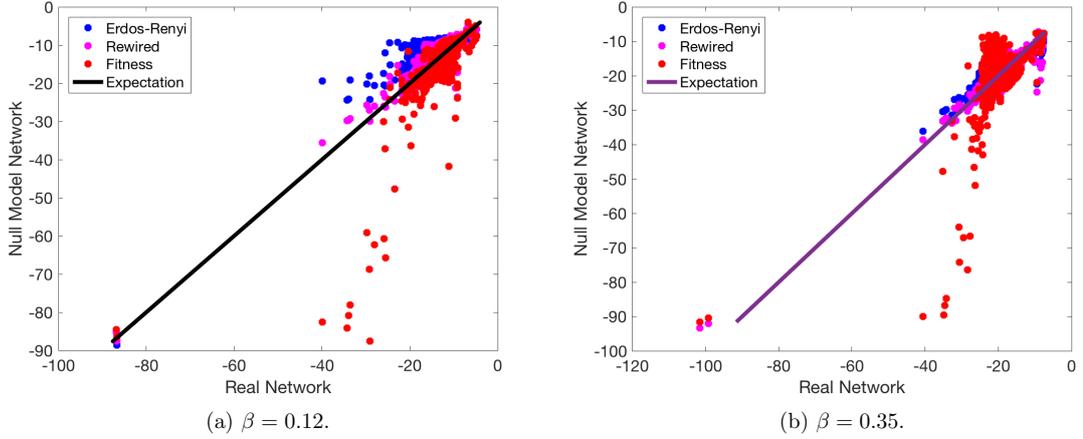


Figure 6.11: Nodal shocks K_i in log-scale for the three null model networks shown vis-à-vis the expectation from the real network. The fitness network has the largest outliers.

β	ER	Rewired	Fitness
0.12	0.214	0.119	0.130
0.24	0.186	0.085	0.118
0.35	0.052	0.020	0.109

Table 6.2: Root mean squared distance of nodal energies with respect to the real network. The configuration model always performs best and the fitness model has the largest error when β is large.

6.4.3 Role of network topology in reverse stress testing

An analysis of the isolated nodes can provide an explanation for the poor performance of the fitness model. The rewired and ER networks correctly estimate the number of isolated nodes. The real network has 39 such nodes. The configuration model performs well because it is given the complete degree distribution of the real network. The ER networks, despite having the largest measuring error with respect to nodal degrees, perform well as the information of isolated nodes is implied by the linking probability p , see equation (3.18). The fitness networks on the other hand have on average 140 such isolated nodes. Because isolated nodes cannot be affected by other nodes in the network and therefore require much more external influence, K is larger in the fitness networks.

To confirm this hypothesis I undertook the following experiment. In the fitness networks and the real network I removed all nodes that are isolated only in the fitness networks, whilst ensuring that during this removal process additional nodes that become isolated are also removed. I then recomputed the average exogenous

shock both for the reduced real and all 1,000 fitness network simulations. The result of this experiment is shown in fig. 6.12.

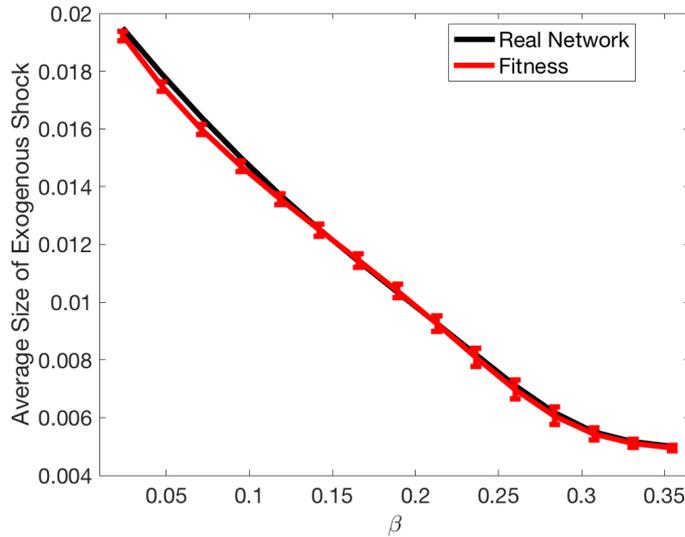


Figure 6.12: Average exogenous shock $\sqrt{K/n}$ as a function of β for the fitness networks and real network, in which the isolated nodes in the fitness networks have been removed from both networks. The match between the two networks is almost perfect, suggesting that the poor performance of the fitness networks is due to their overestimation of the abundance of isolated nodes.

It can be seen in the figure that the average total shock per node is now much closer to the real network, leading to an almost perfect match with the real network. The poor performance of the fitness model networks therefore results from the relatively high abundance of isolated nodes in the fitness networks relative to the real network. This insight also explains why in fig. 6.10 the IPR of the fitness networks increases as β increases, instead of decreasing as in the real network. As a network becomes increasingly unstable the isolated nodes command the highest fraction of exogenous shocks. Given the relatively high abundance of isolated nodes in the fitness networks, the IPR must therefore increase relative to the real network, as most other nodes require much less energy when the network is unstable.

6.4.4 Robustness tests

I undertook robustness tests of these results by rerunning the null model reconstruction exercise for the Austrian network in first quarter of 2007 and 2008. As before I generated 1,000 networks for each of the three null models and compared the results of the reverse stress test exercise to that of the real network. I find that the results are very robust across time and I show these for the first quarter of 2007 in appendix A.4. The behaviour of the null models is almost indistinguishable across

the three quarters at the aggregate level. Again as noted in section 6.3.3, due to the anonymised nature of the Austrian interbank data I am not in a position to compare the results at the nodal level.

6.5 Discussion and Future Developments

In this chapter I presented the results of applying the reverse stress testing framework to different networks. I tested the reverse stress test framework on the STOXX network with different densities, the Austrian interbank network and the networks generated by the Erdős-Rényi, directed configuration (rewired) and fitness null models. In particular, the goal of this chapter is to investigate the topological determinants of the aggregate size of the exogenous shock to cause a certain level of losses.

Across all networks I find that the main properties of the reverse stress test framework are robust. Specifically, I find that the size of the exogenous shock K decreases as a function of the largest eigenvalue of the leverage matrix λ_{max} , which acts as an amplifier of the dynamics. In the STOXX network K is not affected by the density when $\lambda_{max} < 1$, however when $\lambda_{max} > 1$ the exogenous shock decreases as a function of the density. I relate this behaviour of K to the connectivity of the network. When the connectivity of a network is increased, exogenous shocks can more easily propagate throughout the network, which results in a lower exogenous shock requirement.

The distribution of nodal shocks in the Austrian network is more heterogeneous relative to that of the STOXX network. The presence of relatively larger nodes in the Austrian compared to the STOXX network, leads to a higher concentration of nodal shocks in these larger nodes – which in turn increases the concentration of nodal shocks in the network.

The equity capital policy, which involves the targeted allocation of additional equity capital to reduce the observed losses, performs equally well across the different networks when the density is larger than 50%. For lower densities the presence of isolated nodes results in a reduced effectiveness of the policy. In these cases the connectivity of the network is not sufficient for the additional equity capital to be effectively spread throughout the network. I observe a similar effect in the Austrian network, in which the effectiveness of the equity policy is increased when nodes of very low degree are removed.

Isolated nodes also play an influential role in the ability of the null models to estimate the size of the exogenous shock in the Austrian network. I surprisingly find that the fitness model has the highest estimation error for the exogenous shock out of the three null models when $\lambda_{max} > 1$. The configuration model performs

best across all values of λ_{max} as it is given the entire degree distribution of the real network. The good performance of the configuration model relies on the pivotal role that the local topology plays for the exogenous shock to propagate the network effectively. Surprisingly, the ability of the null models to measure λ_{max} is not a core driver of their performance. The ER networks have the closest fit to the λ_{max} of the real network, yet they do not outperform the rewired networks. Instead, I find that the relatively poor performance of the fitness networks is due to their overestimation of isolated nodes. When the isolated nodes that are present in the fitness but not the real network are removed, the adjusted fitness networks have an almost perfect match with the real network.

The reduced effectiveness of the equity capital policy and the reverse stress testing framework as whole is an artefact of my experimental set-up, rather than a limitation of the framework itself. In real interbank systems true isolated banks do not exist for multiple reasons. The main source of liquidity and additional interest income for commercial banks is their interbank network, such that each bank has connections with a part of the interbank market in its country of residence. Furthermore, in the networks that I study here, a snapshot of the system is taken at a particular point in time. A small bank may not have been active in the interbank market in this particular time period. Isolated nodes should therefore not be a limitation of this framework in a real use-case.

In summary, the local topology has a significant impact on the size of the aggregate exogenous shock K . In particular when the largest eigenvalue λ_{max} of the leverage matrix is larger than one, a poorly connected network will have a much larger exogenous shock requirement in order to achieve the targeted losses. This insight is reflected in the ability of the three null models to reproduce the dynamics of the real network. A small error in the measurement of the local topology can lead to a large difference in the size of the exogenous shock, as is it can be seen in the fitness networks due to isolated nodes and in the ER networks through the large error bars of the exogenous shock estimate.

6.5.1 Future developments

In addition to the extensions and future work mentioned in the previous chapter in section 5.7, this project could be extended by testing different network topologies and different dynamics of shock propagation. An additional avenue for future research is to shock only some rather than all nodes in the network. This may present a more realistic application of the reverse stress testing framework. Studies of smaller networks using control theory show that the required fraction of drivers, to whom

the exogenous control input is applied, is relatively small at around 20% [128]. For larger networks however, which is particularly the case when applying control theory to real-world networks, the required number of drivers is much larger. For the Austrian network for instance, the computation of the control input is currently only possible when around 80% of nodes are also driver nodes. Moreover, studies on the structural controllability of interbank networks have shown that networks with smaller time aggregation are more difficult to control, I describe these studies in section 5.2.1. Shocking less than all banks in real interbank networks, may thus be an ambitious goal for future research.

Chapter 7

Conclusion and Discussion

7.1 Conclusions of Individual Projects

In this thesis I investigate the influence of the topology of interbank networks on financial contagion. I used the insights that I develop in this thesis to propose a methodology to manage systemic risk in interbank networks with different topologies. In the following four sections I summarise and discuss the contributions of each project. In the final section of this chapter I discuss common themes of my results across all projects and highlight avenues for future research.

7.1.1 First project: Systemic risk - Disentangling the influence of interbank network topology and dynamics

In the first project I study the required amount of information to reconstruct the financial contagion of a real interbank network. Previous studies investigated the ability of null models to capture the topology of interbank networks, but did not contrast their ability to estimate financial contagion.

I use a unique dataset of each bilateral exposure of all Austrian banks between 2006 and 2008. This dataset is particularly interesting because it captures the onset of the financial crisis. I find that the fitness model performs better in estimating financial contagion and network topology than Erdős-Rényi random networks and the directed configuration (rewired) model. The better performance of the fitness model in comparison to the configuration model is surprising as it is provided with information only of the total interbank assets and liabilities of each bank, rather than each bilateral exposure as in the configuration model. I therefore conclude that information on the aggregate interbank activity of a bank is sufficient to ascertain its individual exposures to other banks.

The null models measure the DebtRank of the Austrian network well, as Debt-

Rank depends on the global topological properties such as the density, which the null models reconstruct well. DebtRank furthermore depends on the interbank leverage, which is indirectly provided to the null models via the RAS algorithm that allocates the edge weights to each null model and is based on the total interbank assets and liabilities of each bank. In contrast, the null models are only able to capture general trends in financial contagion as measured by the Furfine counterparty default risk model, but cannot reconstruct the precise values of contagion observed in the real network. Through my investigations I show that the reason for the poor performance of the null models is due to their inability to accurately reproduce the local topology of the real interbank network. The local topology is crucial in defining the emergence and propagation of counterparty default risk as measured by the Furfine algorithm. I find that the local topology in the real network is characterised by a higher concentration of risk relative to the reconstructed networks. This higher concentration of risk is due to a larger vulnerable cluster in the real compared to the reconstructed networks. The vulnerable cluster is a set of banks, whose members default when only one of their counterparties defaults.

7.1.2 Second project: Topological determinants of control energy

I use optimal control theory to study the cost of managing the dynamics of a complex system. In particular, I investigate the required control energy of driving a network to a desired target state. The previous literature on optimal control theory in networks had only been able to compute the energy to control networks in highly limited settings. In this project I derive analytical solutions for the control energy in Erdős-Rényi networks and performed numerical simulations on Erdős-Rényi and scale-free networks. In this way I discover that the control energy depends crucially on network connectivity.

Networks of low connectivity are more costly to control as nodes cannot influence one another. As the network connectivity increases, nodes push and pull one another to their common target state and thereby reduce the need for an external control input. This coordination advantage reduces, when nodes have increasingly different target states, as they cannot help one another in reaching their target state, such that the required external energy increases. Moreover, the coordination advantage is more pronounced in scale-free networks. This is due to the relatively much larger degree of some nodes in scale-free networks, which can coordinate the state of a large subset of nodes.

7.1.3 Third project: Reverse stress testing interbank networks

Using the insights of the first and second projects I develop a framework to reverse stress test interbank networks. This method is a novel approach of applying reverse stress tests at the systemic level. It allows the identification of a set shocks that need to hit the equity capital of banks to result in a specific equity loss in the entire system. The framework is different to the standard forward-looking approach to financial contagion, which measures the outcome of an initial shock.

I study the reverse stress testing framework analytically and on a network representing the largest financial institutions in the European Union. I find that the size of the exogenous shock that needs to hit a system crucially depends on the stability of the system. When the system is stable, as indicated by the largest eigenvalue λ_{max} of the matrix of interbank leverages being smaller than 1, the required shock is much larger than the shock that needs to hit a system that is unstable, i.e. when $\lambda_{max} > 1$.

Moreover, I find that when the stability of the system increases, that the nodal exogenous shocks become increasingly concentrated in a few risky banks. I use the increasing concentration of risk in a subset of banks to create a ranking of banks in order of their systemic importance. Based on this ranking I implement a policy to provide targeted capital injections to each bank, that increase with the systemic importance of each bank. I show that using this policy the original losses of the system can be reduced by up to 95%.

7.1.4 Fourth project: Reverse stress testing interbank networks - The influence of topology

In the fourth and final project I test the impact of different network topologies on the reverse stress testing framework. I find that the framework is robust across different topologies, as the exogenous shock has similar properties to those found in the third project. Moreover, I find that the exogenous shock is independent of the network density when the network is stable. In contrast, when the network is unstable, which is indicated by the largest eigenvalue λ_{max} of the matrix of interbank leverages being larger than one, the exogenous shock decreases as a function of an increasing network density. When the network is well connected, as indicated by a high density, the most risky banks can more easily cause losses across the network, which results in a lower exogenous shock requirement. Equally, the equity capital policy is much more effective in reducing the overall system losses in networks with larger rather than smaller densities. The policy injects targeted amounts of additional equity capital

into banks based on their systemic importance.

I additionally test the reverse stress testing framework on reconstructed networks of the Austrian interbank network. I build the reconstructed networks using the Erdős-Rényi random networks, the directed configuration (rewired) and the fitness null models, which I introduce in the first project. In contrast to the first project, I find that the fitness model performs worst out of the three null models in estimating the size of the exogenous shock. The poor performance of the fitness model is due to its overestimation of the connectivity of large nodes and its underestimation of the connectivity of small nodes, which increases the exogenous shock requirement. The configuration model performs best in estimating the exogenous shock and the concentration of nodal shocks. I show that the good performance of the configuration model is due its specific knowledge of the topology of the real network, which is crucial for the nodal shocks to propagate throughout the network.

7.2 Conclusion and Discussion across Projects

Across the results of my four projects I identified some common themes, which I summarise and discuss in this section.

1. **Strong effect of the most connected nodes.** In the second project I find that the control energy required to drive a system to its target state scales as a function of the average degree of the network. This is an effect that I term the coordination advantage, as nodes with a larger degree require the largest fraction of control energy but also coordinate the states of many nodes in the network and thereby reduce the overall energy requirement. These nodes of large degree –hub nodes– are a distinctive feature of scale-free networks, resulting from the power-law nature of their degree distribution. The coordination advantage can also be observed in a similar manner in the interbank networks to which the reverse stress testing framework is applied. In this context I find that banks of larger degree require a larger fraction of the exogenous shock and that STOXX networks with a higher connectivity require a smaller total exogenous shock. Moreover, in the first project I find that highly leveraged hub nodes in the Austrian interbank network are the largest sources of potential financial instability via the extended vulnerable cluster.
2. **The difficulty of disentangling the effect of edge weights from the local topology.** In the first project I find despite the strong topological differences of the three null model networks, that the trajectories of contagion are very similar in all three networks. Since I fit the edge weights across all null

model networks with the RAS algorithm, I deduce that the size of contagion events across iterations is strongly influenced by edge weights, i.e. the size of the exposures between banks. The size of the edge weights is expressed as the leverage of each exposure of a lender in the matrix of interbank leverages in the third and fourth projects. In these projects I show that larger leverages significantly increase financial contagion, which confirms the importance of edge weights that I find in the first project.

3. **The influence of network diversification on financial contagion.** The different projects intermittently touch on the issue of diversification in interbank networks. As I discuss in the literature review, the impact of increased network connectivity upon financial stability is particularly evident when interbank leverage is high. In the first project on the Austrian network for instance, the larger concentration of risk in the Austrian relative to the reconstructed networks is shown through a larger vulnerable cluster, which contains highly leveraged banks and that bears a significant fraction of the counterparty default risk in the real network. Furthermore, in the third and fourth project the reverse stress testing framework shows that more densely connected nodes and higher interbank leverages lead to more fragile systems, as shocks can more easily propagate in the STOXX and Austrian interbank networks.
4. **The importance of the local topology in the propagation of financial contagion.** Interestingly, none of the null models are able to reconstruct the local topology of the Austrian network. The local topology is crucial in routing the dynamics of networks, which is a finding that was confirmed in all four projects. In the first project I find that specific triadic motifs, in particular the unreciprocated loop, results in higher financial contagion. This is particularly noticeable through the relatively strong changes in the contagion metrics measured in the Austrian network via the Furfine algorithm, which strongly depend on the existence, size and direction of bilateral exposures between banks. In the fourth project I show that knowledge of the specific bilateral connections between banks is essential for the equity capital policy to function effectively. Isolated nodes in particular absorb most of the additional equity capital, but cannot propagate it sufficiently well throughout the network to increase network stability.
5. **Policy implications resulting from the importance of the local topology.** An important implication of the results across my projects is that policy interventions in interbank markets need to be tailored to each bank. Highly leveraged banks with large exposures to many counterparties pose the most

substantial risks for the stability of interbank markets. Moreover, poorly diversified banks present a high risk not only to their own stability, i.e. by having poor access to liquidity, but also to the entire network as they hinder the diversification of risk. I show this with the equity capital policy in the fourth project, as isolated nodes cannot propagate the additional equity capital that they receive to increase the stability of the network. This is particularly important in the light of the marginal utility of higher equity capital requirements that I identify in the third project. The additional equity capital provided to the most systemically important banks, to enable these banks to reduce their leverage, accounts for the largest reduction in systemic risk. For this reason increasing banks' capital requirements has a marginal utility in reducing systemic risk, since even small but targeted equity capital injections address the most significant sources of financial instability.

7.3 Future Developments

Beyond the specific suggestions made in the individual projects, the following future steps could be taken to improve research on systemic risk.

In order to improve the applicability of the reverse stress testing framework, control theory, in particular the computation of control energy, should be developed for non-linear dynamics. A step towards this ambitious goal would be to consider time-varying instead of time-invariant adjacency matrices. Each time an input, or shock, is injected into an interbank network via the reverse stress testing framework, the entries of the leverage matrix on which the DebtRank algorithm evolves could in theory be readjusted to reflect the updated leverage of banks. This would result in a richer understanding of financial contagion.

I show in this thesis that the overall differences between types of networks, for instance their degree distribution and assortativity, accounts only for some of the differences in the dynamics of networks. If the drivers of financial contagion are to be understood, then the local topology of networks should be studied in future research. For instance, in this thesis I show the importance of the local topology using triadic motifs and the vulnerable cluster, which account for large parts of the observed dynamics. The local topology is also crucial in the measurement of the pace of financial contagion. For instance, to measure the unknown β parameter of the DebtRank algorithm in a real interbank network, which defines the strength of the interactions between nodes, an exceptionally detailed view at very small time scales of bilateral exposures is required. Much more research should be undertaken to understand how these very small building blocks of networks arise and influence

dynamics. However, many studies have shown that an investigation of triads proves difficult in real networks [10]. Many convoluted factors contribute to real interbank network dynamics. This complexity of financial contagion renders it practically impossible to isolate the effect of the local topology, let alone individual triadic motifs.

In response to the difficulty of isolating the driving factors of systemic risk, some studies, such as those by Adrian and Brunnermeier, investigate systemic risk only at the aggregate level [27, 1, 33]. These studies use econometric risk measures such as Conditional Value at Risk (CoVaR) and SRISK, that depend on aggregate factors such as the overall equity capital in a banking system and relative bank size in terms of total assets. By using aggregate instead of bilateral information on banks, these studies avoid the data availability problem in network studies due to slow reporting and privacy issues. Scholars argue that these econometric methods are less prone to underestimate systemic risk, because their measures depend on aggregate market prices, which represent the integration of the assessment of all risk factors by market participants.

In contrast, the network approach to systemic risk can only address a few channels of contagion at a time. This is due to the complexity of the models and the large data requirements at the nodal level. The network approach therefore leaves out some channels of financial contagion, which can result in an underestimation of risk. However, research on interbank networks, including my thesis, has shown that only network analyses of contagion can reveal the significant amplification of risk due to the endogenous dynamics in financial systems. My reverse stress testing framework may offer a way to combine the econometric and network theoretic approach to systemic risk. The exogenous shock modelled in my framework can be treated as an aggregate market variable that is propagated through the network. This would allow the endogenous risk that results from an aggregate exogenous shock, to be measured at the nodal level. The emergence of multilayer networks is important in this regard, as they allow the combination of different asset classes and channels of contagion in a single system.

Overall, it is important to recall that the total elimination of systemic risk is not the goal of policy-makers nor desirable for market participants. An economy is only free and open when it is characterised by natural cycles of "boom-and-bust". The famous instability hypothesis from Hyman Minsky suggests that stability is destabilizing: Optimistic market participants take more risks, which sow the seeds of future busts. In fact, the European Central Bank's Composite Stress Indicator was at its all time low before the financial crisis of 2008. The volatility of financial markets is again at all time lows in these days. It is at times as these that systemic

risk studies must prove their worth, by highlighting sources of endogenous risk and by making suggestions to inhibit the most risky activities in booms and thereby prevent the worst possible scenarios during busts.

Appendix

A.1 Properties of the Topology and Financial Contagion in the Austrian Network and its Reconstructed Networks in Different Quarters

A.1.1 Topological measures

In the main text in section 3.4 I described the topology of the Austrian network only for the first quarter of 2006. I also presented the results on the ability of the three null models to reconstruct the topology of the Austrian network only for the first quarter of 2006. I showed information only for one quarter for reasons of brevity, because the results across the quarters are very similar. I here prove this point by presenting the topological properties of the Austrian network and the reconstructed networks for the first quarter of 2007 in table A.1 and for the first quarter of 2008 in table A.2.

Metric		Rewired	Erdős-Rényi	Fitness Model	Real Network
Basic	Nodes (n)	834	834	834	834
	Edges (e)	13,740 (13,740 13,740)	13,740 (13,590 13,893)	13,740 (13,619 13,861)	13,740
	Density (d)	0.0198 (0.0198 0.0198)	0.0198 (0.0196 0.0200)	0.0198 (0.0196 0.0200)	0.0198
	Degree ($\langle k \rangle$)	32.9496 (32.9496 32.9496)	32.9504 (32.5899 33.3165)	32.9486 (32.6595 33.2398)	32.9496
	Assortativity	-0.55 (-0.55 -0.55)	0.0 (0.00 0.00)	-0.41 (-0.42 -0.41)	-0.5760
	Global Clustering	0.36 (0.36 0.37)	0.0058 (0.0056 0.0059)	0.32 (0.31 0.32)	0.3428
	Core Size	73 (73 73)	43 (43 44)	66 (65 67)	73
	Core-Periphery Error	-0.0392 (-0.0392 -0.0392)	0.9218 (0.9121 0.9315)	0.0692 (0.0651 0.0733)	-0.0392
	Cosine Similarity	0.664 (0.656 0.674)	0.0458 (0.0069 0.0870)	0.6714 (0.6669 0.6766)	NA
	RMSD	0.072 (0.072 0.073)	0.1610 (0.1391 0.1844)	0.0766 (0.0763 0.0769)	NA
Triads	# of nodes in Motif 1	764 (757 771)	834 (834 834)	748 (740 755)	389
	# of nodes in Motif 5	493 (481 505)	832 (831 834)	746 (738 754)	411
	# of nodes in Motif 9	142 (131 154)	777 (772 782)	467 (452 482)	124
	# of nodes in Motif 13	158 (154 162)	0 (0 0)	255 (244 265)	174

Table A.1: Summary statistics of topological properties of the real and reconstructed networks. For the reconstructed networks, the different metrics are averaged over 10,000 simulations and rounded to the nearest significant integer that can be described with confidence. The figures in the brackets represent the 10% and 90% quantiles, respectively. Results presented refer to the first quarter of 2007.

Metric		Rewired	Erdős-Rényi	Fitness Model	Real Network
Basic	Nodes (n)	825	825	825	825
	Edges (e)	13,944 (13,944 13,944)	13,893 (13,744 14,048)	13,943 (13,823 14,062)	13,944
	Density (d)	0.0205 (0.0205 0.0205)	0.0204 (0.0202 0.0207)	0.0205 (0.0203 0.0207)	0.0205
	Degree ($\langle k \rangle$)	33.8036 (33.8036 33.8036)	33.6809 (33.3188 34.0558)	33.8002 (33.5103 34.0897)	33.8036
	Assortativity	-0.56 (-0.56 -0.56)	0.0 (0.00 0.00)	-0.42 (-0.42 -0.41)	-0.5837
	Global Clustering	0.34 (0.34 0.34)	0.0060 (0.0058 0.0061)	0.32 (0.31 0.32)	0.3166
	Core Size	75 (75 75)	44 (43 45)	67 (66 68)	75
	Core-Periphery Error	-0.0351 (-0.0351 -0.0351)	0.9165 (0.9068 0.9263)	0.0618 (0.0578 0.0658)	-0.0351
	Cosine Similarity	0.664 (0.656 0.675)	0.0449 (0.0081 0.0795)	0.6313 (0.6266 0.6366)	NA
	RMSD	0.072 (0.071 0.073)	0.1529 (0.1345 0.1725)	0.0766 (0.0763 0.0769)	NA
Triads	# of nodes in Motif 1	761 (755 768)	825 (825 825)	744 (736 751)	399
	# of nodes in Motif 5	476 (464 487)	824 (823 825)	740 (731 748)	414
	# of nodes in Motif 9	151 (140 162)	769 (765 773)	463 (448 478)	114
	# of nodes in Motif 13	161 (157 165)	0 (0 0)	261 (250 271)	174

Table A.2: Summary statistics of topological properties of the real and reconstructed networks. For the reconstructed networks, the different metrics are averaged over 10,000 simulations and rounded to the nearest significant integer that can be described with confidence. The figures in the brackets represent the 10% and 90% quantiles, respectively. Results presented refer to the first quarter of 2008.

A.1.2 Trajectory profiles

In fig. 3.7a I showed the trajectory profiles of financial contagion due to counterparty default for the first quarter of 2006 and noted that the behaviour is similar in the other quarters. In the below fig. A.1 I present the same plot for the first quarter of 2007 and 2008, to show the similarity of the trajectories of the null models relative to the real network across the quarters. It can be seen that fig. 3.7a for 2006, looks very similar to fig. A.1a for 2007, in which the trajectories of the real network are more abrupt and impactful compared to the null models. The trajectories in the real network are also more impactful relative to the reconstructed networks in the first quarter of 2008, which I show in fig. A.1b. However, in this case the profiles in the real network are much larger in the second and third iterations relative to the first iteration. This behaviour is also observed in the second and third quarters of 2008 and is due to a much higher fragility of lenders at the onset of the global financial crisis of 2008. The inability of the null models to pick up on this change in the shape of the profiles, provides further evidence that the null models cannot capture the local topology of the real network.

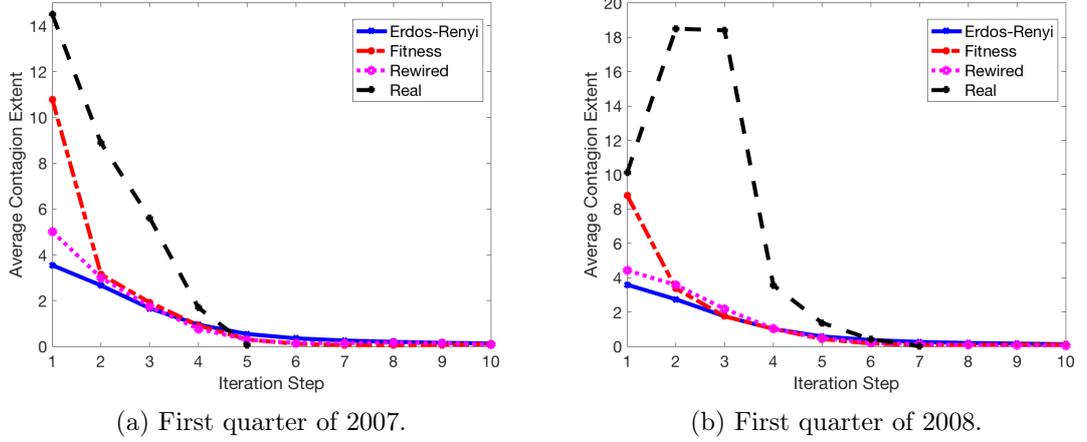


Figure A.1: Trajectories of financial contagion, measured as the total number of banks failed in each iteration of the dynamic scaled by the number of contagion events in the first iteration. Black line shows the measure of the real network. Blue refers to Erdős-Rényi, pink to configuration and red to fitness model networks. X-axis cut-off at 10 iterations.

A.2 Identification of Vulnerable Clusters

In this section I outline the algorithm that I developed to identify vulnerable clusters in real networks. Furthermore, I present the results of numerical simulations to prove the effectiveness of the algorithm, which I use in section 3.6.

A.2.1 Vulnerable cluster identification algorithm

In this section I introduce a method to identify the (extended) giant vulnerable cluster in interbank networks. As mentioned in the literature review in section 3.2.3, the fractional size of the extended giant vulnerable cluster approximates the probability of a global cascade, i.e. the contagion probability. When applying the analytical models of Watts [135] and Hurd [74] to real networks several problems arise that may result in a breakdown of the approximation of the fractional EGVC size of the contagion probability:

1. The analytical model assumes a pure branching structure, i.e. a network without loops. Loops frequently occur in real networks and distort the estimation of financial contagion, as they may increase systemic risk in a particularly dense part of the network by introducing a negative feedback loop between nodes [125].
2. In the analytical model the size of a global cascade is defined as a relative fraction of the network size, whereas in the Furfine model I assume a contagion event to occur when one bank fails beyond the seed, i.e. an absolute

threshold value defines a global cascade. I study the effect of this assumption in appendix A.2.2.

With these complications in mind, I now outline the vulnerable cluster identification method. The VCs were identified using a breadth-first algorithm, for details see Newman (2010) [103, p. 315ff]. This algorithm is in essence the counterparty risk model, described in section 3.3.1, adjusted by a weak-link condition: One borrower i must be sufficient to cause a single lender j to default, see section 3.6.1 for a definition. In the following I will refer to this algorithm as the VC identification algorithm.

When defaulting each bank in the network once as a seed, 375 unique banks fail following the Furfine algorithm in the Austrian network based on the first quarter of 2006. Out of these 375 failed banks, 353 are vulnerable with respect to one other bank, the remaining 22 banks are therefore classified as robust. Since I am interested in finding the largest cluster of vulnerable banks in a network, i.e. the giant vulnerable cluster (GVC), I run the VC identification algorithm on the subset of 353 vulnerable banks. As a result I find 47 unique vulnerable nodes that form a cluster with other vulnerable nodes, which I refer to as the GVC.

The EGVC consists of the set of nodes that constitute the GVC as well as the robust nodes that are immediately adjacent to it and can fail at least one GVC member. I therefore run the counterparty risk algorithm again with the set of GVC members as lenders and all robust banks as borrowers. After running this algorithm I identified an EGVC of size 116. The fractional size of the EGVC $116/846 \approx 0.14$ approximates well the contagion probability in the Austrian network of 0.14 in the first quarter of 2006. This good match reflects the effectiveness of the VC identification algorithm.

A.2.2 Simulations to confirm the efficacy of the VC identification algorithm

In this section I show that the VC identification algorithm developed in the previous section successfully identifies the (extended) giant vulnerable cluster in controlled experimental settings. I undertake this analysis to compare the effectiveness of my identification algorithm to the numerical simulations undertaken by Hurd (2016) [74].

Figure A.2a and fig. A.2b show the results of running the Furfine counterparty default risk algorithm on synthetic interbank networks with $n = 1,000$ and various average degrees $\langle k \rangle$ as shown on the x-axis, averaged over 1,000 simulations for each $\langle k \rangle$. The underlying degree distribution is created with the Erdős-Rényi approach

and the interbank balance sheets are homogeneously distributed as described in the following appendix A.2.3.

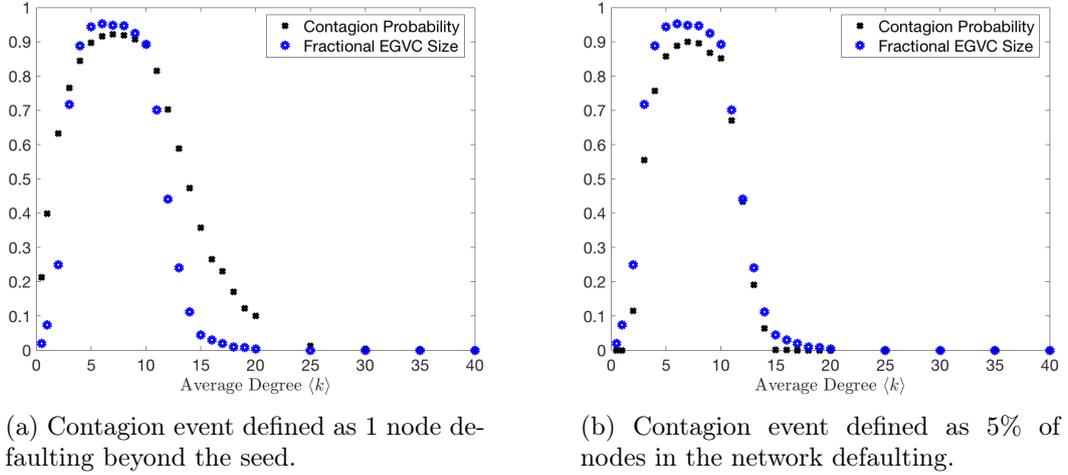


Figure A.2: Reproduction of experiments undertaken by Watts and Hurd, to prove the efficacy of the VC identification algorithm [135, 74]. Figures represent the averages over 1,000 simulations at each average degree $\langle k \rangle$ on the x-axis. The fractional EGVC size approximates the contagion probability almost perfectly. L.H.S.: In networks with $10 < \langle k \rangle < 20$ the fractional EGVC size does not match the contagion probability, as they are typically defined as $> 5\%$ of nodes failing.

In particular, fig. A.2a and fig. A.2b show the contagion probability computed using the Furfine algorithm and the fractional EGVC size for the assumption of a contagion event occurring when one bank defaults beyond the seed and when more than 5% of nodes default, respectively. The EGVC is measured using the VC identification algorithm that I developed in the previous section. As in the experiments of Watts and Hurd, the global cascade frequency (contagion probability) is approximated by the fractional EGVC size when a global cascade is defined at the 5% threshold [74, 135]. In networks with a larger average degree the typical contagion event is small, such that contagion events are more likely to occur and the approximation therefore breaks down when the contagion threshold is defined as 1 node beyond the seed, as it can be seen in fig. A.2a for $10 < \langle k \rangle < 20$.

This analysis shows that my VC identification algorithm works well in approximating the contagion probability with the fractional EGVC size and that it can reproduce the analytical results of Hurd and Watts, who introduce the theoretical model.

A.2.3 Creation of synthetic bank balance sheets

In appendix A.2.2 I undertake numerical experiments in which I measure the abundance of vulnerable clusters in a network composed of banks that are based on

synthetic balance sheets. In this section I describe the methods used to create these synthetic bank balance sheets.

I adapt the setting suggested by Gai and Kapadia (2010) to create homogeneous balance sheets [59]. The creation of interbank exposures presumes the prior creation of a binary directed adjacency matrix A of size n via the Erdős-Rényi method as per eq. (3.18). The interbank assets and liabilities as well as the equity capital are assigned as follows.

- Equity: 4 for all banks;
- Total interbank assets for bank i : $a_i = 20 \forall k_i^{out} > 0$, which each lender spreads uniformly amongst its counterparties;
- The individual loan amounts are assigned as follows to create the weighted directed liability matrix L_{ij} :

$$L_{ij} = \begin{cases} \frac{a_j}{k_j^{out}} & \text{if } A_{ij} = 1 \\ 0 & \text{if } A_{ij} = 0. \end{cases} \quad (\text{A.1})$$

- Total interbank liabilities for bank i are the sum across all its liabilities towards its lenders j :

$$l_i = \sum_{j=1}^n L_{ij}. \quad (\text{A.2})$$

The total interbank assets and the equity of each bank are exogenously set, the interbank liabilities of each bank are endogenously determined by the binary adjacency matrix A . The entire volume is held constant at $n * 20$ across all simulations, to ensure that interbank networks have an identical interbank volume independent of the density of the underlying binary adjacency matrix A .

A.3 Reverse Stress Testing the Austrian Interbank Network across Quarters

In section 6.3 I presented the results of the reverse stress testing framework applied to the Austrian framework based on data of the first quarter of 2006. Here I show the core results that I presented in the main text, for the first quarter of 2007 to prove the robustness of my results.

In fig. A.3a I show the average shock as a function of the target losses ℓ and the largest eigenvalue λ_{max} of the leverage matrix. Similarly, in fig. A.3b I show

the results for the inverse participation ratio. It can be seen that as for 2006 in the main text, that the average shock decreases as a function of λ_{max} in fig. A.3a and that nodal shocks become more concentrated as a function of λ_{max} in fig. A.3b.

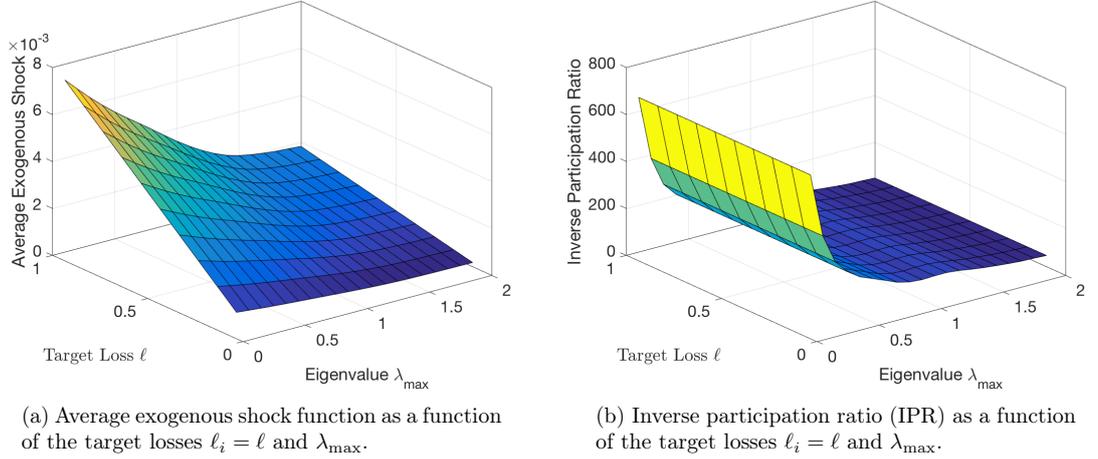
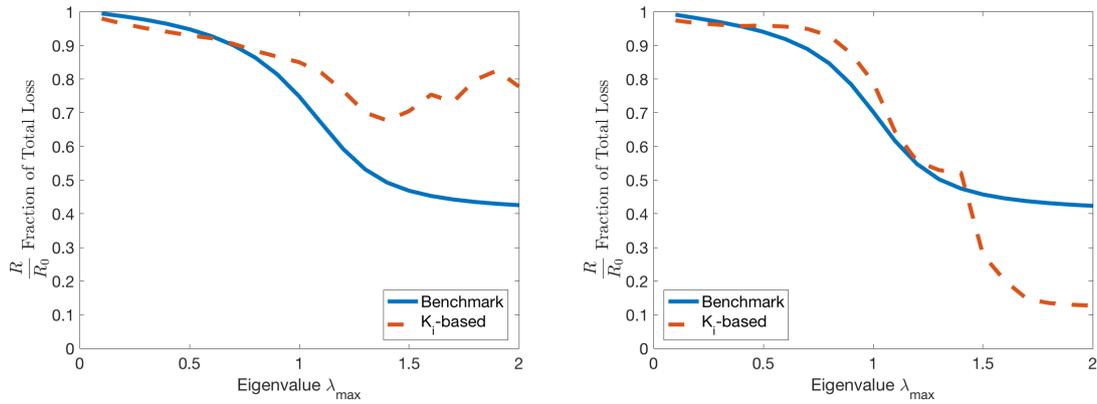


Figure A.3: Average exogenous shock $\sqrt{K/n}$ and inverse participation ratio (IPR) for the Austrian network based on data for the first quarter of 2007. Results shown for $T = 20$.

In fig. A.4a I show a comparison of the benchmark and the K_i -based policy that both aim to reduce the observed losses in the Austrian network, based on data of the first quarter of 2007. It can be seen that as for the 2006 data shown in fig. 6.7a, that the benchmark policy is more effective than the K_i -based policy. In fact, again, the K_i -based effectiveness decreases as λ_{max} increases beyond one. To show that this is due to isolated nodes defined as $k^{in} = 0$, I applied the two policies to a subnetwork of the Austrian network in which I have removed all nodes with $k^{in} \leq 1$ and $k^{out} \leq 1$. I show the result of this exercise for the subnetwork with $n = 219$ in fig. A.4b, which shows that the K_i -based policy is now more effective than the benchmark. The figure shows, as for the network based on 2006 data presented in fig. 6.7b, that the effectiveness of the K_i -based policy crucially depends on the abundance of poorly connected nodes.



(a) Comparison of policies to reduce the observed losses in the full Austrian interbank network.

(b) Policies for Austrian subnetwork with $n = 219$ in which no node has $k^{in} \leq 1$ and $k^{out} \leq 1$.

Figure A.4: Equity capital policies applied to the Austrian network based on the first quarter of 2007. The K_i -based policy is much more effective in the subnetwork (R.H.S.) than in the full network (L.H.S.).

A.4 Reverse Stress Testing the Austrian Interbank Network and its Reconstructions in Different Quarters

In section 6.4 I tested the ability of three null models to reconstruct the reverse stress testing dynamics in the Austrian network. In particular, I used the Erdős-Rényi (ER), the directed configuration (rewired) and the fitness null models based on data of the first quarter of 2006. In section 6.4.4 I described that I performed robustness tests of these experiments based on data of the first quarter of 2007 and 2008. I here show these results for the first quarter of 2007 that are based on 1,000 simulations for each null model.

Figure A.5 shows the average shock per node for the real network and the three null models. It can be seen as in the main text in fig. 6.9, that the configuration model always performs best and that the fitness model has the largest measuring error when $\beta > 0.24$ ($\lambda_{max} > 1$).

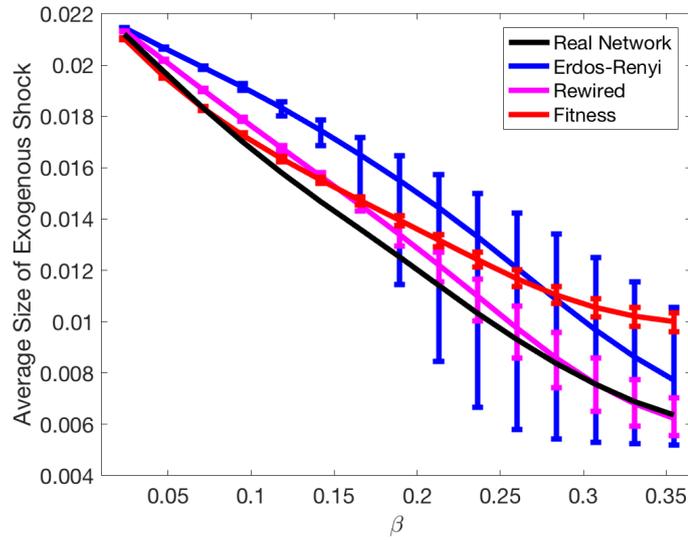


Figure A.5: Average exogenous shock $\sqrt{K/n}$ for the three null models and the real network as a function of β . The error bars indicate the 5th and 95th percentiles, respectively. Results shown for $T = 20$ and $\ell_i = 0.1$ based on data for the first quarter of 2007.

Figure A.6 shows the concentration of nodal shocks as measured by the inverse participation ratio for 2007. The behaviour is very similar to the same analysis shown for 2006 in fig. 6.10. The configuration model measures the nodal shock concentration best and the ER networks perform worst, with the exception of large values of β , when the fitness model performs worst. This is due to the large number of isolated nodes in the fitness networks relative to the real network, which decreases the concentration of shocks, as I explain in section 6.4.3.

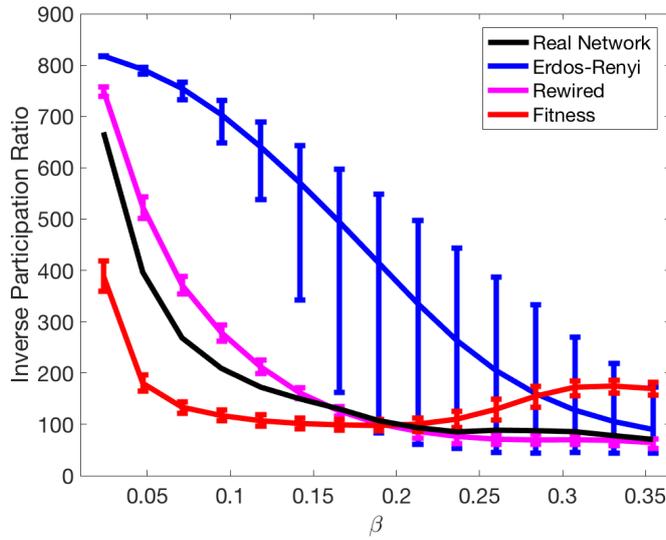


Figure A.6: Inverse participation ratio (IPR) as a function of β for the three null models and the real network. The error bars indicate the 5th and 95th percentiles, respectively. Results shown for $T = 20$ and $\ell_i = 0.1$ based on data of the first quarter of 2007.

Lastly, I study the nodal shocks estimated by the null models relative to those of the real network. For 2006 I show the results of this analysis in fig. 6.11. Below in fig. A.7 I show the same analysis for data based on the first quarter of 2007. The figure shows that as for 2006, the fitness model networks have the largest deviations from the real network both for $\beta = 0.12$ and for $\beta = 0.35$.

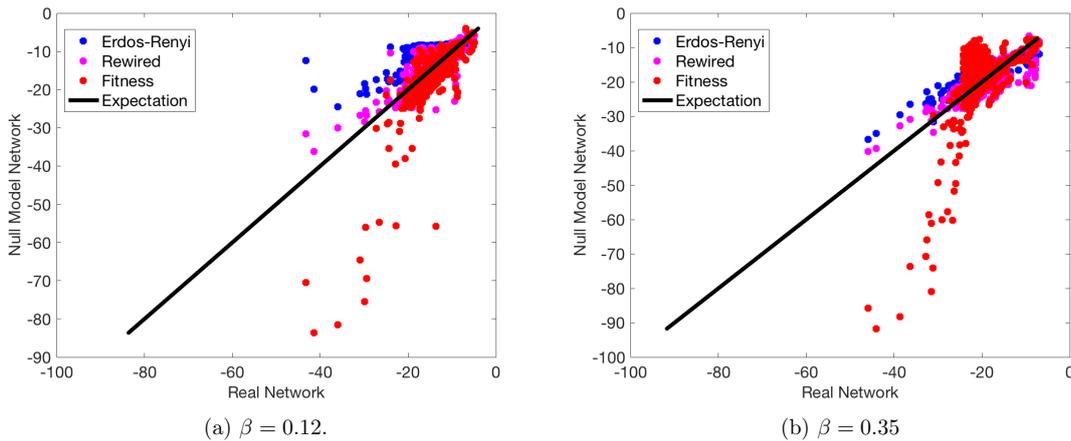


Figure A.7: Nodal shocks K_i in log-scale for the three null model networks shown vis-à-vis the expectation from the real network. The fitness networks have the largest outliers.

Bibliography

- [1] T. Adrian and M. K. Brunnermeier. CoVaR. *The American Economic Review*, 106(7):1705–1741, 2016.
- [2] R. Albert, H. Jeong, and A.-L. Barabasi. Error and attack tolerance of complex networks. *Nature*, 406:378–382, 2000.
- [3] I. Aldasoro, D. D. Gatti, and E. Faia. Bank Networks: Contagion, Systemic Risk and Prudential Policy. *CEPR Discussion Paper No. 10540*, (597), 2015.
- [4] F. Allen and D. Gale. Financial Contagion. *Journal of Political Economy*, 108(1):1, 2000.
- [5] H. Amini, R. Cont, and A. Minca. Stress Testing the Resilience of Financial Networks. *International Journal of Theoretical and Applied Finance*, 15(01):1250006, 2012.
- [6] H. Amini, R. Cont, and A. Minca. Resilience to contagion in financial networks. *Mathematical Finance*, 26(2):329–365, 2016.
- [7] K. Anand, B. Craig, and G. von Peter. Filling in the blanks: network structure and interbank contagion. *Quantitative Finance*, 15(4):625–636, 2015.
- [8] K. Anand, P. Gai, and M. Marsili. Rollover risk, network structure and systemic financial crises. *Journal of Economic Dynamics and Control*, 36(8):1088–1100, 2012.
- [9] K. Anand, I. van Lelyveld, Á. Banai, S. Friedrich, R. Garratt, G. Halaj, J. Figue, I. Hansen, S. Martinez-Jaramillo, H. Lee, J. L. Molina-Borboa, S. Nobili, D. Salakhova, T. Christiano Silva, L. Silvestri, and S. R. Stancato de Souza. The missing links: A global study on uncovering financial network structure from partial data. *Journal of Financial Markets*, (In Press), 2017.
- [10] M. T. Angulo, Y.-Y. Liu, and J.-J. Slotine. Network motifs emerge from interconnections that favour stability. *Nature Physics*, 11(10):848–852, 2015.

- [11] C. Aymanns. The dynamics of the leverage cycle. *SSRN Electronic Journal*, 2468800, 2014.
- [12] C. Aymanns, F. Caccioli, J. D. Farmer, and V. W. Tan. Taming the Basel leverage cycle. *Journal of Financial Stability*, 27:263–277, 2016.
- [13] M. Bacharach. Estimating Nonnegative Matrices from Marginal Data. *International Economic Review*, 6(3):294–310, 1965.
- [14] A. Bachem and B. Korte. On the RAS-Algorithm. *Computing*, 23:189–198, 1979.
- [15] O. Banwo, F. Caccioli, P. Harrald, and F. Medda. The effect of heterogeneity on financial contagion due to overlapping portfolios. *ArXiv e-prints*, 1704.06791, 2017.
- [16] A.-L. Barabasi and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286:509–512, 1999.
- [17] M. Bardoscia, S. Battiston, F. Caccioli, and G. Caldarelli. DebtRank: A microscopic foundation for shock propagation. *PLoS ONE*, 10(6):1–13, 2015.
- [18] M. Bardoscia, S. Battiston, F. Caccioli, and G. Caldarelli. Pathways towards instability in financial networks. *Nature Communications*, 8:1–7, 2016.
- [19] M. Bardoscia, F. Caccioli, J. I. Perotti, G. Vivaldo, and G. Caldarelli. Distress propagation in complex networks: The case of non-linear DebtRank. *PLoS ONE*, 11(10):1–12, 2016.
- [20] L. Bargigli, L. Infante, F. Lillo, and F. Pierobon. The multiplex structure of interbank networks. *Quantitative Finance*, 15(4):673–691, 2013.
- [21] P. Barucca and F. Lillo. Disentangling bipartite and core-periphery structure in financial networks. *Chaos, Solitons and Fractals*, 88:244–253, 2016.
- [22] Basel Committee On Banking Supervision. *Basel III: A global regulatory framework for more resilient banks and banking systems*. 2011.
- [23] S. Battiston, G. Caldarelli, M. D’Errico, and S. Gurciullo. Leveraging the network: a stress-test framework based on DebtRank. *ArXiv e-prints*, 1503.00621, 2015.
- [24] S. Battiston, M. D’Errico, and S. Gurciullo. DebtRank and the Network of Leverage. *Institutional Investor Journals*, 18(4):68–81, 2016.

- [25] S. Battiston, D. D. Gatti, M. Gallegati, B. Greenwald, and J. E. Stiglitz. Default cascades: When does risk diversification increase stability? *Journal of Financial Stability*, 8(3):138–149, 2012.
- [26] S. Battiston, M. Puliga, R. Kaushik, P. Tasca, and G. Caldarelli. DebtRank: too central to fail? Financial networks, the FED and systemic risk. *Scientific Reports*, 2(541):1–6, 2012.
- [27] N. Beale, D. G. Rand, H. Battey, K. Croxson, R. M. May, and M. A. Nowak. Individual versus systemic risk and the Regulator’s Dilemma. *Proceedings Of The National Academy Of Sciences Of The United States Of America-Biological Sciences*, 108(31):12647–12652, 2011.
- [28] S. Berardi and G. Tedeschi. From banks’ strategies to financial (in)stability. *International Review of Economics and Finance*, 47(November 2016):255–272, 2017.
- [29] A. Berenyi, M. Belluscio, D. Mao, and G. Buzsaki. Closed-Loop Control of Epilepsy by Transcranial Electrical Stimulation. *Science*, 337(6095):735–737, 2012.
- [30] D. Bisias, M. Flood, A. W. Lo, and S. Valavanis. A survey of systemic risk analytics. *Annual Review of Financial Economics*, 4(1):255–296, 2012.
- [31] S. P. Borgatti and M. G. Everett. Models of core/periphery structures. *Social Networks*, 21:375–395, 1999.
- [32] M. Boss, H. Elsinger, M. Summer, and S. Thurner. The Network Topology of the Interbank Market. *Quantitative Finance*, 4(6):677–684, 2004.
- [33] C. Brownlees and R. Engle. SRISK: A Conditional Capital Shortfall Index for Systemic Risk Measurement. *Review of Financial Studies*, 30(1):48–79, 2016.
- [34] M. K. Brunnermeier and L. H. Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22:2201–2238, 2009.
- [35] F. Caccioli, T. A. Catanach, and J. D. Farmer. Heterogeneity, correlations and financial contagion. *Advances in Complex Systems*, 15:1–15, 2012.
- [36] F. Caccioli, J. D. Farmer, N. Foti, and D. Rockmore. Overlapping portfolios, contagion, and financial stability. *Journal of Economic Dynamics and Control*, 51:50–63, 2015.

- [37] F. Caccioli, M. Shrestha, C. Moore, and J. D. Farmer. Stability analysis of financial contagion due to overlapping portfolios. *Journal of Banking & Finance*, 46:233–245, 2014.
- [38] G. Caldarelli, A. Capocci, P. De Los Rios, and M. Muñoz. Scale-Free Networks from Varying Vertex Intrinsic Fitness. *Physical Review Letters*, 89(25):258702–1–4, 2002.
- [39] X. Chen. Financial Contagion on Power Law Networks. *SSRN Electronic Journal*, 1751143, 2011.
- [40] R. Cifuentes, G. Ferrucci, and H. S. Shin. Liquidity Risk and Contagion. *Journal of the European Economic Association*, 3(2-3):556–566, 2005.
- [41] G. Cimini, T. Squartini, D. Garlaschelli, and A. Gabrielli. Systemic Risk Analysis on Reconstructed Economic and Financial Networks. *Scientific Reports*, 5:15758, 2015.
- [42] L. Clerc, A. Giovannini, S. Langfield, T. Peltonen, R. Portes, and M. Scheicher. Indirect contagion: the policy problem. *European Systemic Risk Board Occasional Paper Series*, January(9), 2016.
- [43] N. J. Cowan, E. J. Chastain, D. A. Vilhena, J. S. Freudenberg, and C. T. Bergstrom. Nodal dynamics, not degree distributions, determine the structural controllability of complex networks. *PLoS ONE*, 7(6), 2012.
- [44] B. Craig and G. von Peter. Interbank tiering and money center banks. *Journal of Financial Intermediation*, 23(3):322–347, 2014.
- [45] J. Danielsson, H. Song Shin, and J.-P. Zigrand. Endogenous and systemic risk. *NBER*, (January):73–94, 2012.
- [46] G. De Masi, G. Iori, and G. Caldarelli. Fitness model for the Italian interbank money market. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 74:1–5, 2006.
- [47] D. Delpini, S. Battiston, M. Riccaboni, G. Gabbi, F. Pammolli, and G. Caldarelli. Evolution of controllability in interbank networks. *Scientific Reports*, 3:1626, 2013.
- [48] T. Dette, S. Pauls, and D. N. Rockmore. Robustness and Contagion in the International Financial Network. *ArXiv e-prints*, 1104.4249, 2011.

- [49] M. Drehmann and N. Tarashev. Measuring the systemic importance of interconnected banks. *Journal of Financial Intermediation*, 22(4):586–607, 2013.
- [50] F. Duarte and T. M. Eisenbach. Fire-Sale Spillovers and Systemic Risk. *SSRN Electronic Journal*, (FRB of New York Staff Report No. 645):76, 2013.
- [51] P. Erdős and A. Rényi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.
- [52] European Banking Authority. 2016 EU-Wide Stress Test (29 July 2016).
- [53] Financial Stability Board. 2015 Update of List of Global Systemically Important Banks.
- [54] K. Finger, D. Fricke, and T. Lux. Network analysis of the e-MID overnight money market: the informational value of different aggregation levels for intrinsic dynamic processes. *Computational Management Science*, 10(2-3):187–211, 2013.
- [55] M. D. Flood and G. G. Korenko. Systematic scenario selection: stress testing and the nature of uncertainty. *Quantitative Finance*, 15(1):43–59, 2015.
- [56] D. Fricke, K. Finger, and T. Lux. On Assortative and Disassortative Mixing in Scale-Free Networks: The Case of Interbank Credit Networks. *Kiel Working Papers*, February(1830), 2013.
- [57] D. Fricke and T. Lux. Core-Periphery Structure in the Overnight Money Market: Evidence from the e-MID Trading Platform. *Computational Economics*, 45(3):359–395, 2014.
- [58] C. H. Furfine. Interbank Exposures: Quantifying the Risk of Contagion. *Journal of Money, Credit and Banking*, 35(1):111–128, 2003.
- [59] P. Gai and S. Kapadia. Contagion in financial networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 466(2120):2401–2423, 2010.
- [60] M. Galbiati, D. Delpini, and S. Battiston. The power to control. *Nature Physics*, 9(3):126–128, 2013.
- [61] D. Garlaschelli and M. I. Loffredo. Fitness-dependent topological properties of the world trade web. *Physical Review Letters*, 93:2–5, 2004.
- [62] J. Geanakoplos. The Leverage Cycle. *National Bureau of Economic Research*, 24:1–65, 2010.

- [63] C. Gietl and F. P. Reffel. Accumulation points of the iterative proportional fitting procedure. *Metrika*, 76:783–798, 2013.
- [64] P. Glasserman, C. Kang, and W. Kang. Stress scenario selection by empirical likelihood. *Quantitative Finance*, 15(1):25–41, 2015.
- [65] P. Glasserman and H. P. Young. How likely is contagion in financial networks? *Journal of Banking and Finance*, 50:383–399, 2014.
- [66] P. Glasserman and H. P. Young. Contagion in Financial Networks. *Journal of Economic Literature*, 54(3):779–831, 2016.
- [67] K. I. Goh, B. Kahng, and D. Kim. Universal behavior of load distribution in scale-free networks. *Physical Review Letters*, 87(27):278701, 2001.
- [68] R. Grilli, G. Tedeschi, and M. Gallegati. Markets connectivity and financial contagion. *Journal of Economic Interaction and Coordination*, 10(2):287–304, 2014.
- [69] P. Grundke and K. Pliszka. A macroeconomic reverse stress test. *Deutsche Bundesbank Discussion Paper*, 30, 2015.
- [70] S. Gualdi, G. Cimini, K. Primicerio, R. Di Clemente, and D. Challet. Statistically validated network of portfolio overlaps and systemic risk. *Scientific Reports*, 6(39467):1–14, 2016.
- [71] A. Hackett, D. Cellai, S. Gomez, A. Arenas, and J. P. Gleeson. Bond percolation on multiplex networks. *Physical Review X*, 6(2):1–8, 2016.
- [72] J. E. Hopcroft and R. M. Karp. An $O(n^{5/2})$ Algorithm for Maximum Matchings in Bipartite Graphs. *SIAM Journal on Computing*, 2(4):225–231, 1973.
- [73] X. Huang, I. Vodenska, S. Havlin, and H. E. Stanley. Cascading Failures in Bi-partite Graphs: Model for Systemic Risk Propagation. *Scientific Reports*, 3(1219):1–8, 2013.
- [74] T. Hurd. *Contagion! The Spread of Systemic Risk in Financial Networks*. Springer International Publishing, Hamilton, Canada, 2016.
- [75] G. Iori, G. De Masi, O. V. Precup, G. Gabbi, and G. Caldarelli. A network analysis of the Italian overnight money market. *Journal of Economic Dynamics and Control*, 32(1):259–278, 2008.

- [76] G. Iori, S. Jafarey, and F. G. Padilla. Systemic risk on the interbank market. *Journal of Economic Behavior and Organization*, 61(4):525–542, 2006.
- [77] G. Iori, M. Politi, G. Germano, and G. Gabbi. Banks’ Strategies and Cost of Money: Effects of the Financial Crisis on the European Electronic Overnight Interbank Market. *Journal of Financial Management, Markets and Institutions*, III(2/2015):179–202, 2015.
- [78] T. Jia, Y.-Y. Liu, E. Csoka, M. Posfai, J.-J. Slotine, and A.-L. Barabasi. Emergence of bimodality in controlling complex networks. *Nature Communications*, 4:2002, 2013.
- [79] Ò. Jordà, M. Schularick, and A. M. Taylor. Macrofinancial History and the New Business Cycle Facts. *NBER Macroeconomics Annual 2016*, 31:213–263, 2017.
- [80] R. E. Kalman. Mathematical Description of Linear Dynamical Systems. *Journal of the Society for Industrial and Applied Mathematics Series A Control*, 1(2):152–192, 1963.
- [81] C. Kok and M. Montagna. Multi-layered Interbank Model for Assessing Systemic Risk. *Kiel Working Papers*, 1873:1–42, 2013.
- [82] S. Langfield, Z. Liu, and T. Ota. Mapping the UK interbank system. *Journal of Banking and Finance*, 45(1):288–303, 2014.
- [83] I. V. Lelyveld. Finding the core: Network structure in interbank markets. *Journal of Banking & Finance*, 49:27–40, 2014.
- [84] S. Lenzu and G. Tedeschi. Systemic risk on different interbank network topologies. *Physica A: Statistical Mechanics and its Applications*, 391(18):4331–4341, 2012.
- [85] S. Z. W. Lip. A Fast Algorithm for the Discrete Core/Periphery Bipartitioning Problem. *ArXiv e-prints*, 1102.5511, 2011.
- [86] Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabási. Controllability of complex networks. *Nature*, 473(7346):167–173, 2011.
- [87] Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabasi. Few inputs can reprogram biological networks. *Nature*, 473:167–173, 2011.
- [88] Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabási. Control Centrality and Hierarchical Structure in Complex Networks. *Plos One*, 7(9):e44459, 2012.

- [89] S. Martinez-Jaramillo, B. Alexandrova-Kabadjova, B. Bravo-Benitez, and J. P. Solórzano-Margain. An empirical study of the Mexican banking system's network and its implications for systemic risk. *Journal of Economic Dynamics and Control*, 40:242–265, 2014.
- [90] S. Maslov and K. Sneppen. Specificity and stability in topology of protein networks. *Science (New York, N.Y.)*, 296:910–913, 2002.
- [91] I. Mastromatteo, E. Zarinelli, and M. Marsili. Reconstruction of financial networks for robust estimation of systemic risk. *Journal of Statistical Mechanics: Theory and Experiment*, 3:P03011, 2012.
- [92] A. J. McNeil and A. D. Smith. Multivariate stress scenarios and solvency. *Insurance: Mathematics and Economics*, 50(3):299–308, 2012.
- [93] S. Melnik, A. Hackett, M. a. Porter, P. J. Mucha, and J. P. Gleeson. The unreasonable effectiveness of tree-based theory for networks with clustering. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 83(3):1–12, 2011.
- [94] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network motifs: simple building blocks of complex networks. *Science (New York, N.Y.)*, 298(5594):824–827, 2002.
- [95] P. E. Mistrulli. Assessing financial contagion in the interbank market: Maximum entropy versus observed interbank lending patterns. *Journal of Banking and Finance*, 35(5):1114–1127, 2011.
- [96] M. Montagna and C. Kok. Multi-layered interbank model for assessing systemic risk. *ECB Working Paper Series - Macprudential Research Network*, 1944, 2016.
- [97] M. Montagna and T. Lux. Hubs and resilience: towards more realistic models of the interbank markets. *Kiel Working Papers*, 1826, 2013.
- [98] N. Musmeci, S. Battiston, G. Caldarelli, M. Puliga, and A. Gabrielli. Bootstrapping topology and systemic risk of complex network using the fitness model. *J. of Statistical Physics*, 151:720–734, 2012.
- [99] J. C. Nacher and T. Akutsu. Structurally robust control of complex networks. *Physical Review E*, 91(1):1–19, 2015.
- [100] M. Newman. Mixing patterns in networks. *Physical Review E*, 67(2):026126, 2003.

- [101] M. E. J. Newman. Assortative mixing in networks. *Physical Review Letters*, 89(4):208701, 2002.
- [102] M. E. J. Newman. The structure and function of complex networks. *Dialogues in clinical neuroscience*, 45:167–256, 2003.
- [103] M. E. J. Newman. *Networks - An Introduction*. Oxford University Press, Oxford, 2010.
- [104] S. Nie, X.-W. Wang, B. Wang, and L.-L. Jiang. Effect of correlations on controllability transition in network control. *Scientific Reports*, 6:23952, 2016.
- [105] E. Nier, J. Yang, T. Yorulmazer, and A. Alentorn. Network models and financial stability. *Journal of Economic Dynamics and Control*, 31(6):2033–2060, 2007.
- [106] F. Pasqualetti, S. Zampieri, and F. Bullo. Controllability metrics, limitations and algorithms for complex networks. *Proceedings of the American Control Conference*, 1(1):3287–3292, 2014.
- [107] R. Pastor-Satorras and A. Vespignani. Epidemic Spreading in Scale-Free Networks. *Physical Review Letters*, 86(14):3200–3203, apr 2001.
- [108] R. Pastor-Satorras and A. Vespignani. Immunization of complex networks. *Physical Review E*, 65(3):036104, feb 2002.
- [109] S. Poledna, O. Bochmann, and S. Thurner. Basel III capital surcharges for G-SIBs fail to control systemic risk and can cause pro-cyclical side effects. *arXiv preprint*, 1602.03505, 2016.
- [110] S. Poledna, J. L. Molina-Borboa, S. Martinez-Jaramillo, M. van der Leij, and S. Thurner. The multi-layer network nature of systemic risk and its implications for the costs of financial crises. *Journal of Financial Stability*, 20:70–81, 2015.
- [111] M. Pósfai, Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabási. Effect of correlations on network controllability. *Scientific Reports*, 3:1–7, 2013.
- [112] F. Pukelsheim. Biproportional scaling of matrices and the iterative proportional fitting procedure. *Annals of Operations Research*, 215(1):269–283, 2013.
- [113] G. Punzo, G. F. Young, M. Macdonald, and N. E. Leonard. Using Network Dynamical Influence to Drive Consensus. *Scientific Reports*, 6(26318), 2016.

- [114] T. Roukny, H. Bersini, H. Pirotte, G. Caldarelli, and S. Battiston. Default cascades in complex networks: topology and systemic risk. *Scientific Reports*, 3:2759, 2013.
- [115] W. J. Rugh. *Linear System Theory*. Prentice Hall, New Jersey, 2 edition, 1995.
- [116] J. Ruths and D. Ruths. Control profiles of complex networks. *Science (New York, N.Y.)*, 343(6177):1373–6, 2014.
- [117] A. Sachs. Completeness, Interconnectedness and Distribution of Interbank Exposures - A Parameterized Analysis of the Stability of Financial Networks. *Deutsche Bundesbank Discussion Paper Series*, 2(8), 2010.
- [118] A. Sergueiva. Multichannel Contagion vs Stabilisation in Multiple Interconnected Financial Markets. *ArXiv e-prints*, 1701.06975, 2016.
- [119] M. Á. Serrano and M. Boguná. Clustering in complex networks. II. Percolation properties. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 74(5):1–8, 2006.
- [120] J. P. Sethna. *Statistical Mechanics: Entropy, Order Parameters and Complexity*. Oxford University Press, Oxford, 2006.
- [121] H. A. Simon. The architecture of complexity. *Proc. Am. Phil. Soc.*, 107:467–482, 1962.
- [122] D. Sornette. Physics and Financial Economics (1776-2014): Puzzles, Ising and Agent-Based models. *Reports on Progress in Physics*, 77:062001, 2014.
- [123] T. Squartini, G. Cimini, A. Gabrielli, and D. Garlaschelli. Network reconstruction via density sampling. *ArXiv e-prints*, 1610.05494, 2016.
- [124] T. Squartini and D. Garlaschelli. Triadic motifs and dyadic self-organization in the World Trade Network. *Lec. Notes Comp. Science*, 7166(3):24–35, 2012.
- [125] T. Squartini, I. van Lelyveld, and D. Garlaschelli. Early-warning signals of topological collapse in interbank networks. *Scientific Reports*, 3:3357, 2013.
- [126] R. F. Stengel. *Optimal Control and Estimation*. Dover Publications, INC., New York, 1994.
- [127] J. E. Stiglitz. Risk and global economic architecture: Why full financial integration may be undesirable. *American Economic Review*, 100(2):388–392, 2010.

- [128] J. Sun and A. E. Motter. Controllability transition and nonlocality in network control. *Physical Review Letters*, 110(20):1–5, 2013.
- [129] P. Tasca and S. Battiston. Market procyclicality and systemic risk. *Quantitative Finance*, 16(8):1–15, 2016.
- [130] G. Tedeschi, A. Mazlounian, M. Gallegati, and D. Helbing. Bankruptcy Cascades in Interbank Markets. *PLoS ONE*, 7(12), 2012.
- [131] S. Thurner and S. Poledna. DebtRank-transparency: controlling systemic risk in financial networks. *Scientific Reports*, 3:1888, 2013.
- [132] C. Upper. Simulation methods to assess the danger of contagion in interbank markets. *Journal of Financial Stability*, 7(3):111–125, 2011.
- [133] C. Upper and A. Worms. Estimating bilateral exposures in the German interbank market: Is there a danger of contagion? *European Economic Review*, 48(4):827–849, 2004.
- [134] S. Vitali, S. Battiston, and M. Gallegati. Financial fragility and distress propagation in a network of regions. *Journal of Economic Dynamics and Control*, 62:56–75, 2016.
- [135] D. J. Watts. A simple model of global cascades on random networks. *PNAS*, 99(9):5766–5771, 2002.
- [136] D. J. Watts and P. S. Dodds. Influentials, networks, and public opinion formation. *Journal of Consumer Research: An Interdisciplinary Quarterly*, 34(4):441–458, 2007.
- [137] G. Yan, J. Ren, Y.-C. Lai, C.-H. Lai, and B. Li. Controlling Complex Networks: How Much Energy Is Needed? *Physical Review Letters*, 108(21):218703, 2012.
- [138] G. Yan, G. Tsekenis, B. Barzel, J.-J. Slotine, and Y.-Y. Liu. Spectrum of Controlling and Observing Complex Networks. *Nature Physics*, 11:1–22, 2015.