

Boston University

OpenBU

<http://open.bu.edu>

Boston University Theses & Dissertations

Boston University Theses & Dissertations

2018

Maximum entropy and network approaches to systemic risk and foreign exchange

<https://hdl.handle.net/2144/33267>

Downloaded from DSpace Repository, DSpace Institution's institutional repository

BOSTON UNIVERSITY
GRADUATE SCHOOL OF ARTS & SCIENCES

Dissertation

**MAXIMUM ENTROPY AND NETWORK APPROACHES
TO SYSTEMIC RISK AND FOREIGN EXCHANGE**

by

ALEXANDER PAUL BECKER

Diplom-Physiker, University of Duisburg-Essen, 2012

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2018

© 2018 by
ALEXANDER PAUL BECKER
All rights reserved

Approved by

First Reader

H. Eugene Stanley, PhD
Professor of Physics

Second Reader

Irena Vodenska, PhD
Associate Professor of Administrative Sciences

Acknowledgments

I cannot help but acknowledge the lucky coincidence that my advisor H. Eugene Stanley was the keynote speaker at the first physics conference I ever attended. Enthralled by his presentation, I wanted to become part of his academic family. I am forever grateful that Gene encouraged me to apply to Boston University, and that he has helped and advised me through these years. In so many ways, I would not be where I am today if it weren't for him. Thank you!

I would also like to thank my co-advisor, Irena Vodenska. She has been a wonderful mentor of science and life, and one of my best decisions while at Boston University was to reach out to her as I was beginning the research stage of my PhD. I have learnt so much from her, and I am most fortunate that I have the opportunity to continue working with her in the future.

I thank Marcel Wollschläger for a fruitful collaboration and for being a friend. Between watching the Red Sox and running around Boston, I thoroughly enjoyed working together on exciting research projects. A significant part of this thesis is rooted in our collaboration.

I further want to extend my thanks to the other members of my committee, Bill Klein, Kevin Black, and Manher Jariwala, for guiding me and helping me with their thoughtful questions and critiques in the final stages of my PhD. I would like to particularly thank Manher Jariwala for introducing me to the wonderful world of physics education research and for his encouragement. I am a better teacher today than I was at the start of my graduate career, and he has played no small role in that.

The physics department has been my academic home over the past five years, and I would like to thank the staff for supporting me however possible. I would like to particularly acknowledge Mirtha Cabello who has been an incredible help, even before I had officially started at Boston University. We are so fortunate to have you,

and you deserve more thanks than would fit in this section.

I have been fortunate to find wonderful friends in the department and at the university with whom I could celebrate successes and who would encourage me when I failed. Thank you Harshal Chaudhari, Clint Richardson, Joe Boales, Clover Su, John Quirk, Rajita Menon, Rashi Verma, Nahom Yirga, Nagendra Panduranga, Nabeel Akhtar, and many more. In particular I want to thank Ashish Bino George, for great advice, for sharing so many laughs, and for making our meager and windowless office a lively space.

I am grateful for the privilege to call two places home, and I want to thank the people that make that possible. A big shoutout goes to my roommates Bryant and Gil – you guys are a big reason I’ve been enjoying Boston so much. I want to thank my friends in Germany for making me a part of their life, even thousands of miles away. Matthias Ebke, Nina Preuß, and Nicolai ten Brinke: how you’ve welcomed me when I visited and your trips to Boston meant a lot to me.

I won’t even try to put into words how grateful I am towards my family. Thank you to my parents Dieter and Bärbel and my siblings Max and Anna. I am beyond fortunate to have your never-ending support, patience, encouragement, and love.

Alexander P. Becker

PhD Candidate

Department of Physics

MAXIMUM ENTROPY AND NETWORK APPROACHES TO SYSTEMIC RISK AND FOREIGN EXCHANGE

ALEXANDER PAUL BECKER

Boston University, Graduate School of Arts & Sciences, 2018

Major Professors: H. Eugene Stanley, PhD
Professor of Physics

Irena Vodenska, PhD
Associate Professor of Administrative Sciences

ABSTRACT

The global financial system is an intricate network of networks, and recent financial crises have laid bare our insufficient understanding of its complexity. In response, within the five chapters of this thesis we study how interconnectedness, interdependency and mutual influence impact financial markets and systemic risk.

In the first part, we investigate the community formation of global equity and currency markets. We find remarkable changes to correlation structure and lead-lag relationships in times of economic turmoil, implying significant risks to diversification based on historical data.

The second part focuses on banks as creators of credit. Bank portfolios generally share some overlap, and this may introduce systemic risk. We model this using European stress test data, finding that the system is stable across a broad range of asset liquidity and risk tolerance. However, there exists a phase transition: If banks become sufficiently risk averse, even small shocks may inflict great losses. Failure to address portfolio overlap thus may leave the banking system ill-prepared.

Complete knowledge of the financial network is prerequisite to such systemic risk analyses. When lacking this knowledge, maximum entropy methods allow a probabilistic reconstruction. In the third part of this thesis, we consider Japanese firm-bank data and find that reconstruction methods fail to generate a connected network. Deriving an analytical expression for connection probabilities, we show that this is a general problem of sparse graphs with inhomogeneous layers. Our results yield confidence intervals for the connectivity of a reconstruction.

The maximum entropy approach also proves useful for studying dependencies in financial markets: On its basis, we develop a new measure for the information content in foreign exchange rates in part four of this thesis and use it to study the impact of macroeconomic variables on the strength of currency co-movements.

While macroeconomic data and the law of supply and demand drive financial markets, foreign exchange rates are also subject to policy interventions. In part five, we classify the roles of currencies within the market with a clustering algorithm and study changes after political and monetary shocks. This methodology may further provide a quantitative underpinning to existing qualitative classifications.

Contents

1	Introduction	1
1.1	Outline	2
2	Community Analysis of Global Financial Markets	6
2.1	Introduction	7
2.2	Materials and methods	9
2.2.1	Data	9
2.2.2	Pearson correlation analysis	11
2.3	Results	12
2.3.1	Pearson correlation analysis	12
2.3.2	Summary statistics	12
2.3.3	Community formation and cluster analysis	16
2.4	Discussion	25
3	Overlapping Portfolios: a Source of Systemic Risk	28
3.1	Introduction	29
3.2	Literature review	30
3.3	Data	36
3.4	Simulation model	40
3.4.1	Modelling a bank's Exposure	41
3.4.2	Modeling sovereign debts and other assets	42
3.4.3	Iteration model	43
3.5	Simulation results	51

3.5.1	Shocking sovereign debt only	52
3.5.2	Allowing for spillover between asset classes	53
3.5.3	Shocking the equity of banks	57
3.5.4	Phase space analysis	61
3.6	Discussion and conclusion	67
4	Reconstruction of the Japanese Bank-Firm Network	70
4.1	Introduction	71
4.2	Methodology	72
4.2.1	Principle of maximum entropy	72
4.2.2	Maximum entropy for networks	73
4.2.3	The fitness ansatz	75
4.2.4	The challenge of sparse networks	77
4.3	Distribution of connection probabilities	79
4.3.1	Uniform distribution	80
4.3.2	Log-normal distribution	82
4.4	Hypothesis testing for connectivity	85
4.4.1	Analytical calculations	86
4.4.2	Numerical results	88
4.5	Data	98
4.6	Reconstruction of an empirical network	99
4.7	Conclusion	104
5	Measuring Currency Comovements with Relative Entropy	106
5.1	Introduction	107
5.2	Methodology	109
5.2.1	Symbolic performance	109
5.2.2	Maximum entropy estimation of probability distributions	113

5.2.3	Relative entropy and Jensen-Shannon divergence	115
5.2.4	Generalized linear models	117
5.2.5	Data	119
5.3	Results	123
5.3.1	Annual intervals	123
5.3.2	Quarterly intervals	126
5.4	Discussion	132
6	Political and Economic Effects on Currency Clustering Dynamics	135
6.1	Introduction	136
6.2	Data	140
6.3	Methodology	142
6.3.1	Exchange rates and returns	142
6.3.2	Symbolic performance	143
6.3.3	<i>k</i> -means++ clustering	145
6.4	Empirical results	149
6.4.1	Overall symbolic performance distributions	149
6.4.2	Dynamics of symbolic performance distributions	152
6.5	Conclusion	164
7	Summary and conclusions	166
A	Supplementary Material to Chapter 2	169
A.1	Table of countries in data set	169
B	Supplementary Material to Chapter 3	171
B.1	Interbank network results	171
C	Supplementary Material to Chapter 5	185
C.1	Macroeconomic data used for analysis	185

References	187
Curriculum Vitae	204

List of Tables

2.1	Statistics for the pre-crisis and the crisis periods.	15
3.1	Weighted-Average model for the risk-weighted Asset. The values for the exposure to different asset classes are for bank No.1, AT001. We present an example calculation for risk weights at the lower end of the range determined by our optimization procedure.	44
5.1	Trade volumes in billion US dollar between the different currency regions Australia (AUD), Canada (CAD), Switzerland (CHF), the eurozone (EUR) and the United Kingdom (GBP). The rows indicate the exporter and the columns indicate the importer.	121
5.2	Finite size error estimates for the KL divergence.	125
5.3	Result: Generalized linear model for currency comovements r_{KL} based on macroeconomic data and currency classifications.	131
5.4	Result: Generalized linear model for currency comovements r_{KL} based on macroeconomic data and conditioned on currencies.	132

6.1	We report how frequently (as a percentage) each currency appears in each cluster, enumerated from 1 to 6 as in Figure 6-2, as well as the average cluster value. The euro has the lowest average cluster component, followed by the US dollar, the two main reference currencies in the FX market. The Japanese yen, the New Zealand dollar, and the South African rand stand out as the currencies with the most tail-heavy symbolic performance distributions.	155
B.1	Number of banks that would fall below the Basel III threshold of 4.5% tier 1 capital ratio given a shock to different sectors in each of our ten scenarios after 50 simulation time steps (D_{50}) and after 100 simulation time steps (D_{100}). The initial shock is a sudden increase of the risk weights of the shocked sector of a particular country or region by 50%, causing an increase of the risk-weighted assets of banks. The numbers include the two banks which are below the threshold originally. . . .	172
B.2	Effect of a shock that originates in the GIIPS countries. There are a total of 43 banks in the GIIPS countries. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in the GIIPS countries (Loss GIIPS) as well as the average loss in tier 1 capital ratio for banks in other than the GIIPS countries (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for the GIIPS countries (D_{100} GIIPS) and for other than the GIIPS countries (D_{100} Other). .	173

B.3 Effect of a shock that originates in Eastern Europe. There are a total of 4 banks in Eastern Europe. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Eastern Europe (Loss EE) as well as the average loss in tier 1 capital ratio for banks in other countries than in Eastern Europe (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Eastern Europe (D_{100} EE) and for other than Eastern Europe countries (D_{100} Other). 174

B.4 Effect of a shock that originates in the Benelux countries. There are a total of 7 banks in the Benelux countries. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in the Benelux countries (Loss Benelux) as well as the average loss in tier 1 capital ratio for banks in other than the Benelux countries (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for the Benelux countries (D_{100} Benelux) and for other than the Benelux countries (D_{100} Other). . . . 175

B.5 Effect of a shock that originates in Greece. There are a total of 6 banks in Greece. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Greece (Loss Greece) as well as the average loss in tier 1 capital ratio for banks in other than Greece (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Greece (D_{100} Greece) and for other than Greece (D_{100} Other). 176

B.6 Effect of a shock that originates in Italy. There are a total of 6 banks in Italy. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Italy (Loss Italy) as well as the average loss in tier 1 capital ratio for banks in other than Italy (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Italy (D_{100} Italy) and for other than Italy (D_{100} Other). 177

B.7 Effect of a shock that originates in France. There are a total of 4 banks in France. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in France (Loss France) as well as the average loss in tier 1 capital ratio for banks in other than France (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for France (D_{100} France) and for other than France (D_{100} Other). 178

B.8 Effect of a shock that originates in Germany. There are a total of 12 banks in Germany. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Germany (Loss Germany) as well as the average loss in tier 1 capital ratio for banks in other than Germany (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Germany (D_{100} Germany) and for other than Germany (D_{100} Other). 179

B.9 Effect of a shock that originates in Spain. There are a total of 25 banks in Spain. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Spain (Loss Spain) as well as the average loss in tier 1 capital ratio for banks in other than Spain (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Spain (D_{100} Spain) and for other than Spain (D_{100} Other). 180

B.10 Effect of a shock that originates in the GIIPS countries. There are a total of 43 banks from these countries. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in GIIPS countries (Loss GIIPS) as well as the average loss in tier 1 capital ratio for banks in other than GIIPS countries (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for the GIIPS (D_{100} GIIPS) and for other than GIIPS (D_{100} Other). 181

B.11 Effect of a shock that originates in the Eastern European countries.	
There are a total of 4 banks from these countries. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Eastern Europe (Loss EE) as well as the average loss in tier 1 capital ratio for banks in other than Eastern European countries (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for the Eastern Europe (D_{100} EE) and for other than Eastern European countries (D_{100} Other).	181
B.12 Effect of a shock that originates in the Benelux countries. There are a total of 7 banks from these countries. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Benelux countries (Loss Benelux) as well as the average loss in tier 1 capital ratio for banks in other than Benelux countries (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for the Benelux countries (D_{100} Benelux) and for other than Benelux countries (D_{100} Other).	182

B.13 Effect of a shock that originates in Greece. There are a total of 6 banks in Greece. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Greece (Loss Greece) as well as the average loss in tier 1 capital ratio for banks in other than Greece (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Greece (D_{100} Greece) and for other than Greece (D_{100} Other). 182

B.14 Effect of a shock that originates in Italy. There are a total of 6 banks in Italy. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Italy (Loss Italy) as well as the average loss in tier 1 capital ratio for banks in other than Italy (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Italy (D_{100} Italy) and for other than Italy (D_{100} Other). 183

B.15 Effect of a shock that originates in France. There are a total of 4 banks in France. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in France (Loss France) as well as the average loss in tier 1 capital ratio for banks in other than France (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for France (D_{100} France) and for other than France (D_{100} Other).	183
B.16 Effect of a shock that originates in Germany. There are a total of 12 banks in Germany. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Germany (Loss Germany) as well as the average loss in tier 1 capital ratio for banks in other than Germany (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Germany (D_{100} Germany) and for other than Germany (D_{100} Other).	184

B.17 Effect of a shock that originates in Spain. There are a total of 25 banks in Spain. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Spain (Loss Spain) as well as the average loss in tier 1 capital ratio for banks in other than Spain (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Spain (D_{100} Spain) and for other than Spain (D_{100} Other). 184

C.1 Lookup table for the macroeconomic indicators of countries in our data set and classifications into commodity (Comm.) or reserve (Res.) currency. 186

List of Figures

2.1	Heat maps of the annual Pearson correlations for 2.1(a) stock markets and 2.1(b) foreign exchange markets logarithmic returns. For the stock markets, we considered all 56 countries; for the foreign exchange markets, we considered the 45 distinct currencies. The color shows the value of the correlations for different years, where red indicates strong positive correlation, blue indicates strong negative correlation, while green means that the correlation is weak.	13
2.2	(a) Annual mean correlation and (b) annual standard deviation for stock markets (x), foreign exchange markets (+), and between stock and foreign exchange markets (o). The stock markets were more correlated during the crisis period, whereas the foreign exchange market correlation were higher in the non-crisis period. The correlation between stock and foreign exchange markets was lowest during the crisis. The standard deviation for both, stock and foreign exchange markets remained fairly constant, while the standard deviation for the inter-layer correlations increased.	14

2·3	Planar maximally filtered graph (PMFG) for the stock markets obtained using synchronous correlations during (a) the economically calm period and (b) crisis period. The countries are denoted by their three-letter symbols and are color-coded according to their geographical locations: green for Asia, light blue for Europe and orange for the Americas. For the calm period, we detected five large clusters which are mostly geographically divided. The clusters significantly changed during the crisis period. Most notably the troubled Eurozone countries, Italy, Spain, Greece and Portugal, formed their own cluster. The Asian countries formed one larger cluster centered around Hong Kong and Singapore. We also observed a smaller cluster containing a diverse group of less connected countries.	18
2·4	Planar maximally filtered graph (PMFG) for the stock markets obtained using lagged correlations during (a) the economically calm period and (b) crisis period. The countries are denoted by their three-letter symbols and are color-coded according to their geographical locations: green for Asia, light blue for Europe and orange for the Americas. For the calm period, we observed four large clusters, while in the crisis period the number of communities changed to three, as the two European clusters merged.	20

2·6	<p>Planar maximally filtered graph (PMFG) for the foreign exchange markets obtained using lagged correlations during (a) the economically calm period and (b) crisis period. The currencies are denoted by their three-letter symbols and color-coded according to their geographical locations: green for Asia, light blue for Europe, orange for the Americas. For lagged correlations we used time t for the Americas and time $t + 1$ for Europe and Asia. For both periods, we detected three large clusters. In the non-crisis period we observed that the clusters were determined by the geographical currency location, except for the Japanese yen, New Zealand dollar and Australian dollar. Iceland appeared in the predominantly Asian cluster, while Malaysia, the Philippines and Indonesia were part of the American community. During the crisis period the interconnectedness between American and South East Asian currencies increased as they belonged to the same community. The purely Asian cluster comprised the Japanese Yen along with the currencies of the oil-exporting countries. The Russian ruble joined the European cluster.</p>	24
3·1	<p>Histograms of bank holdings where the x-axis is the percentage of a bank's portfolio an asset makes up. The dotted line indicates the mean value of the distribution.</p>	37

3·2	PMFG (Planar Maximally Filtered Graph) and MST (Minimum Spanning Tree) network showing portfolio overlap in sovereign debt of European banks, including British and Irish banks (bright red), French banks (blue), banks from the Benelux countries (orange), German banks (yellow), Austrian and Eastern European banks (gray), Spanish and Portuguese banks (light red), Italian banks (crimson red), Greek, Maltese and Cypriot banks (purple), and Scandinavian banks (light blue). The size of the node represents the value of all assets a bank. The largest banks are at the center of the network. We observe that banks from the same country cluster, indicating that their portfolios are similar due to home and regional bias.	39
3·3	Bipartite network of the banks β and sovereign debts s	40
3·4	Spreading parameter $Q(D_s)$ where D_s is the CDS spread in basis points for a given sovereign debt s . We choose a the transformation $Q(D_s) = 1 - 2^{-D_s/100}$ such that its value is between 0 and 1 and that a spread of 100 basis points corresponds to a value of 0.5. A CDS spread of 0 corresponds to a Q value of 0, and as the CDS spread grows, Q approaches 1. Note the similarity of the function $Q(D_s)$ with the default probability derived from a CDS spread for some maturity and recovery rate as outlined in Chan-Lau [2006].	43
3·5	The reduction factor $P(x)$ in Eq. (3.5). The dotted curve shows the linear case, and the solid curve shows the steep case.	48
3·6	Tier 1 capital ratios over time after a shock to the retail residential sector in GIIPS countries, given different bank response functions and both a small shock size and a small spreading parameter.	54

3·7	Tier 1 capital ratios over time after a shock to the retail residential sector in GIIPS countries, given different bank response functions and a large size and a small spreading parameter.	54
3·8	Tier 1 capital ratios over time after a shock to the retail residential sector in GIIPS countries, given different bank response functions and a small shock size and a larger spreading parameter.	55
3·9	Tier 1 capital ratios over time after a shock to sovereign debt in GIIPS countries, given different bank response functions and spreading parameters.	56
3·10	Evolution of the tier 1 capital ratio of Deutsche Bank (DE017) for a shock to various sectors in Germany, given a linear response function and a steep response function.	58
3·11	Evolution of the tier 1 capital ratio of Deutsche Bank (DE017) for a shock to corporate loans in various countries or regions, given a linear response function and a steep response function.	60
3·12	Phase space for different scenarios given a linear bank response function: We consider a variety of initial shock sizes to banks and asset classes as well as a variety of values for the spreading parameter. We observe no amplification of shocks. The size of the initial shock is indicated by the color, where from yellow (small shock) to dark purple (very large shock).	63

3.13	Phase space for different scenarios given a steep bank response function: We consider a variety of initial shock sizes to banks and asset classes as well as a variety of values for the spreading parameter. The size of the initial shock is indicated by the color, where from yellow (small shock) to dark purple (very large shock). We observe behavior reminiscent of a second order phase transition, where the critical parameter is the magnitude of the sector spread Q_a . The tipping point of the system appears to be around $Q_a = 0.55$	64
3.14	Stationary state that the system reaches in the simulation depending on the spreading parameter Q_a and the steepness parameter α . Green colors indicate a very small average deterioration of the tier 1 capital ratios of the banks in the system, whereas purple indicates the worst attainable stationary state. The bright line separating the two regimes indicates the point of the phase transition. The plot corresponds to a shock to loans in the French financial sector.	67
4.1	Distribution of the connection probabilities in a simulated sparse bipartite network of size 400×1000 with density $d = 0.1$. The probabilities for a single weak node in layer N are shown in blue, for a weak node in layer two in red, and the overall distribution of connection probabilities is shown in light gray. Note that a node in layer N has 1000 possible links, and a node in layer M has 400 possible links. We observe that the connection probabilities of the weak nodes are exclusively at the lower end of the overall distribution.	78

4.2	Distribution of probabilities of a node to be unconnected in any given realization of the network reconstruction. The nodes in layer N have at least degree one in practically all reconstructions, whereas there is a significant chance for a number of nodes in layer M to remain unconnected.	79
4.3	Fitness parameter z as a function of the network density d in a semilogarithmic representation based on Eq. (4.21) and from a simulation of a bipartite network with 300×100 nodes.	81
4.4	Cumulative probability distribution $F_Y(y)$ of the connection probabilities in the reconstructed network for different densities d . Left: less variance in the node strengths in the layers. Right: more variance. . .	84
4.5	Fitness parameter z as a function of the network density d in a semilogarithmic representation based on Eq. (4.25) and from a simulation of a bipartite network with 40×40 nodes.	85
4.6	Comparing the analytical solution of the connection probability for the node in the lowest quantile with a weak node results from a simulated network.	88
4.7	Average expected connection probability for node Q if its strength is at the quantile values $q = 1/400, 1/1000, 1/2500$ depending on the parameters of strength distributions. The network is reconstructed with density $d = 0.1$. The left plot shows $\mathbb{E}[Y X = x_q]$ if we vary the the parameter σ_N^2 of the log-normal distribution of the strengths of the nodes to which node Q can connect. The right plot shows $\mathbb{E}[Y X = x_q]$ if we vary the the parameter σ_M^2 of the log-normal distribution of the strengths in the layer of node Q. In both cases we keep the parameter σ^2 fixed to 1 for the other layer.	93

4·8	Average expected connection probability for node Q if its strength is at the quantile values $q = 1/400, 1/1000, 1/2500$ depending on the network density d . The left plot shows $\mathbb{E}[Y X = x_q]$ if σ^2 is bigger for the layer of the nodes to which Q can connect than for the layer in which Q is. The right plot shows $\mathbb{E}[Y X = x_q]$ if σ^2 in the opposite case.	94
4·9	Probability that node Q remains unconnected in a reconstruction of a network with density $d = 0.2$ as a function of the strength distributions. The strength of node Q is, from top to bottom rows, at quantile $q = 10^{-3}$, $q = 10^{-2}$, and $q = 10^{-1}$. From left to right, we calculate the cases $N = 100$, $N = 400$, and $N = 1000$	96
4·10	Probability that node Q remains unconnected in a reconstruction of a network where the layers have the same strength distributions. We consider different combinations of density d and quantile q to compute the probability. The parameter σ_2 of the strength distribution of the layers is, from top to bottom rows, $\sigma_2 = 1$, $\sigma_2 = 2$, and $\sigma_2 = 4$. From left to right, we calculate the cases $N = 100$, $N = 400$, and $N = 1000$	97
4·11	Scope of the Nikkei NEEDS database which includes over 100 banks and between 1400 and 2800 firms, depending on the year.	99
4·12	Left: Network density of the bipartite graph derived from banks in one layer and companies in the other. Right: Average degree of the nodes in each layer.	100
4·13	The Bipartite Configuration model assumes that there exists a relationship between the node strengths v_i , c_α and the degrees of the nodes. The empirical data is plotted as gray dots, and the theoretical prediction from the model in red.	101

4.14	Empirical and theoretical distribution of connection probabilities for the reconstruction of the network of Japanese banks and firms in the year 2010.	102
4.15	Empirical and theoretical distribution of connection probabilities for a small bank and a large bank.	103
4.16	Empirical and theoretical distribution of connection probabilities for a small company and a large company.	104
5.1	Left: foreign exchange market consisting of $K = 3$ currencies as a weighted directed network in which the links are the returns of the pairs. Right: binary directed network indicating the appreciation and depreciations for an arbitrary time step. In the example shown, the euro appreciates against the US dollar and the yen, and the US dollar appreciates against the yen only.	111
5.2	Kullback-Leibler divergences between the empirical bivariate distributions and the null model of independence for the years 2012 through 2017.	124
5.3	Kullback-Leibler divergence for currencies resolved on a quarterly basis. The value shown in this figure is the average KL divergence of a currency with all other currencies in the data set. As usual, a larger KL divergence indicates that the time series of the currency shown with other currencies contained more information than if we assumed independence.	128

6.1	Overall symbolic performance distributions $P(\zeta_i)$, $i = 1, \dots, 14$, for our data set. We distinguish between a few different shapes, characterized by their curvature. Some currencies exhibit rather small probabilities at the center and high probabilities in the tails. At the extremes we find the euro (EUR) with very concave curvature and the South African rand (ZAR) with very convex curvature.	150
6.2	Symbolic performance distributions $P(\zeta)$ corresponding to the cluster centers found by the k -means++ algorithm. For the sake of better illustration, we draw a line for each discrete probability distribution, keeping in mind that $P(\zeta)$ is only defined for odd integers between -13 and 13 . The violet distribution (cluster 1) corresponds to currencies which take central positions in the market, whereas the red distribution (cluster 6) corresponds to currencies which tend to take more extreme positions in the market.	153
6.3	Currency clustering dynamics according for the 14 currencies in our data set. Cluster 1 corresponds to reference-like behavior, whereas cluster 6 corresponds to heavy-tail behavior.	154
6.4	Cluster association of the Japanese yen, highlighting the time points and background of publicly announced central bank interventions. . .	157
6.5	Cluster association of the British pound, highlighting the Brexit. . . .	158
6.6	Cluster association of the Swiss franc, highlighting the time points of publicly announced central bank interventions.	161
6.7	Cluster association of the Mexican peso, highlighting the time points of publicly announced central bank interventions in red and political events in the United States in green.	162

6.8	Cluster association of the Singapore dollar, highlighting frequent publications of policy briefs by the Monetary Authority of Singapore. . .	163
-----	--	-----

List of Abbreviations

BIS	Bank for international Settlements
BOJ	Bank of Japan
BRIC	Brazil, Russia, India, China
CBI	Central Bank Intervention
CDO	Collateralized Debt Obligation
CDS	Credit Default Swap
EBA	European Banking Authority
ECB	European Central Bank
ETF	Exchange Traded Fund
FX	Foreign Exchange
GDP	Gross Domestic Product
GFC	Global Financial Crisis
GIIPS	Greece, Italy, Ireland, Portugal, Spain
GLM	Generalized Linear Model
IMF	International Monetary Fund
MAS	Monetary Authority of Singapore
MBS	Mortgage Backed Security
OTC	Over-the-counter
OLS	Ordinary Least Squares
PMFG	Planar Maximally Filtered Graph
QE	Quantitative Easing
RWA	Risk-weighted Assets
SD	Sovereign Debt
SDR	Special Drawing Rights
SME	Small and Medium-sized Enterprises
SNB	Swiss National Bank
SSM	Single Supervisory Mechanism

Chapter 1

Introduction

With his work “The Theory of Speculation” [Bachelier, 1900], the French mathematician Louis Bachelier laid the ground work for some of the most important and influential results in quantitative finance. In order to value stock options, that is, contracts that give the buyer the right to sell or purchase a security at a fixed price and date, Bachelier described the price movement of that security S_t as a random walk:

$$dS_t = a dt + b dW_t, \tag{1.1}$$

where t is time and W_t a stochastic process. Based upon this work, Black and Scholes [1973] developed the theory for derivative pricing which has become indispensable in modern finance. Eq. (1.1) is equivalent to Brownian motion with a drift, named after botanist Robert Brown and successfully explained by Albert Einstein¹.

“The same equations have the same solutions”, Richard Feynman reminded his students in his famous lectures [Feynman et al., 1971], and it is equations like Eq. (1.1) that offer physicists a natural connection to quantitative finance. As a result, equity and foreign exchange markets have historically been the main focus of econophysics. In addition to the mathematical connection, their abundance of time series data provided resources to study the distribution of asset returns, the properties of correlation matrices and the agent based modeling of trading activity.

¹ The movement of stock prices turns out to be better described by Geometric Brownian Motion (GBM) of the form $dS_t = a S_t dt + b S_t dW$. GBM has the desirable property that the relative change in stock price, or return, is independent of price. It also ensures positive prices.

In the past few years, this focus has shifted away from equity and foreign exchange markets, triggered in no small part by the global financial crisis of 2008 and the ensuing European sovereign debt crisis of 2011. The former chair of the Federal Reserve Board of Governors, Janet Yellen, summarized the causes of the global financial crisis five years after it started as follows²:

Losses arising from leveraged investments caused a few important, but perhaps not essential, financial institutions to fail. At first, the damage appeared to be contained, but the resulting stresses revealed extensive interconnections among traditional banks, investment houses, and the rapidly growing and less regulated shadow banking sector.

In other words, the high interconnectedness of the financial system had brought with it a build-up of a critical level of systemic risk. This systemic risk had not been properly monitored or regulated and caused much larger and more widespread damage than expected. The large societal cost of the crisis demonstrated the need to better understand the dynamics of the financial system, in particular its complexity. Therefore, in recent years more and more researchers have directed their attention to the question of financial stability, recognizing that the intricate topology of financial systems needs to be modeled as a complex network.

1.1 Outline

Chapters 2 through 6 contain the main contributions of this dissertation. Each chapter is written as a self-contained paper, providing its own introduction, overview of the methodology, analysis and conclusion. While these chapters can be read in any order, it is recommended to follow the order presented in this thesis as each chapter builds upon ideas and methodologies from the previous ones.

² <https://www.federalreserve.gov/newsevents/speech/Yellen20130104a.pdf>

Chapter 2 analyzes the community formation in global financial markets using methods from network science. The co-movement of financial indicators like stock prices or foreign exchange is typically measured through Pearson's correlation coefficient. Here, this approach is used to create and analyze a network from over 50 countries. This chapter is a lightly edited version of the paper "Community Analysis of Global Financial Markets", coauthored with Irena Vodenska, Di Zhou, Dror Y. Kenett, H. Eugene Stanley and Shlomo Havlin. It appeared in *Risks* **2016**, 4(2), 13 [Vodenska et al., 2016b]. The research shows that community structures and lead-lag relationships for correlation measures change remarkably in times of turmoil. This calls into question diversification and risk management based on historical data.

Chapter 3 focuses on banks as creators of credit and on portfolio overlap as a source of systemic risk. Even if banks are diversified, their portfolios generally overlap to a certain degree, and this may be a source of systemic risk for two reasons. One, a shock to the system may affect a large number of financial institutions at once. Two, in an effort to rebalance their portfolios, banks may inadvertently put more pressure on a distressed asset, potentially causing a fire sale and a downward spiral. Using data from the stress test performed by the European Banking Authority, the response of financial institutions to adverse financial conditions is modeled. The results indicate the financial system is stable across a broad range of asset liquidity and risk tolerance. There exists, however, a phase transition, and if banks turn risk averse, even small shocks have the potential to inflict great losses. While regulating banks individually is sufficient in favorable economic conditions, the research presented suggests that failure to address portfolio overlap and correlations leaves the banking systems ill-prepared for adverse conditions. At the time of this writing, the paper reprinted in this chapter is under review for the *Journal of Financial Stability*, following the *Second Conference on Network Models and Stress Testing for Financial Stability* in

Mexico City in September 2017.

Chapter 4 summarizes a forthcoming paper regarding the reconstruction of bipartite networks using methods of maximum entropy. Studies of systemic risk typically require complete knowledge of the network, which is typically unobtainable or only available for historic data. In this case, one can resort to maximum entropy methods like the Bipartite Configuration Model which reconstruct the network in a probabilistic sense from marginal information. The network of Japanese banks and firms is considered as an example to show the successes and limitations of maximum entropy reconstruction methods. The empirical study of this network and the analytical calculations show that sparse networks with heterogeneous layers may leave many nodes unconnected. This thesis provides confidence intervals for a successful reconstruction in terms of network structure and node properties. Originating from the principle of maximum entropy, the work described in this chapter forms a link between network science approaches to financial markets and stability in the first part of the thesis and the following chapters focusing on foreign exchange markets. Therein, the concept of maximum entropy in a slightly different context.

In Chapter 5, titled “Information Transmission Channels of the Foreign Exchange Network”, the foreign exchange market is interpreted as a network. This leads to the development of a new measure for the co-movement of assets with which the information content in a currency pair is quantified. The results indicate that this maximum entropy based measure is strongly impacted by macroeconomic data like trade volume or interest rates. Since the work described herein is at the time of this writing under preparation for publication, methods are described and developed in more detail than a research paper allows.

While macroeconomic data and the law of supply and demand generally drive the foreign exchange market, currency rates are subject to policy interventions by

central banks. Chapter 6 furthers the analysis of currencies by studying how this affects the FX market. With a clustering algorithm, roles of currencies in the market are classified, and it is studied how they change after political and monetary shocks. This methodology may further be useful to provide a quantitative underpinning to existing qualitative classifications in the FX market. The chapter is a slightly edited version of the paper titled “Political and Economic Effects on Currency Clustering Dynamics”, which is forthcoming in *Quantitative Finance*, co-authored with Marcel Wollschäger, Irena Vodenska, H. Eugene Stanley and Rudi Schäfer.

Finally, a summary of the previous chapters and concluding remarks are offered in Chapter 7.

Chapter 2

Community Analysis of Global Financial Markets

We analyze the daily returns of stock market indices and currencies of 56 countries over the period of 2002-2012. We build a network model consisting of two layers, one being the stock market indices and the other the foreign exchange markets. Synchronous and lagged correlations are used as measures of connectivity and causality among different parts of the global economic system for two different time intervals: non-crisis (2002-2006) and crisis (2007-2012) periods. We study community formations within the network to understand the influences and vulnerabilities of specific countries or groups of countries. We observe different behavior of the cross correlations and communities for crisis vs. non-crisis periods. For example, the overall correlation of stock markets increases during crisis while the overall correlation in the foreign exchange market and the correlation between stock and foreign exchange markets decrease, which leads to different community structures. We observe that the euro, while being central during the relatively calm period, loses its dominant role during crisis. Furthermore we discover that the troubled Eurozone countries, Portugal, Italy, Greece and Spain, form their own cluster during the crisis period.

2.1 Introduction

Financial crisis can cause substantial damages and economic losses not only locally, but also in other countries through trade relations, currency policies, financial contracts, and cross-country investments. Some examples of such crises are the 1997 Asian financial crisis, 1998 Russian bond crisis, 2001 dot-com bubble, 2007-08 global financial crisis, and 2010 EU sovereign debt crisis, all spilling over to various parts of the world. Similar to the transmission of a disease, small financial shocks initially affecting only a particular sector of the economy or geographic region can spread to other economic sectors and other countries with quite healthy economic outlook [Lin et al., 1994].

Many research studies have examined the connections among countries by exploring correlations of various financial time series data [Smith, 2009, Forbes and Rigobon, 2002, Flavin et al., 2002, Solnik et al., 1996, Ramchand and Susmel, 1998, Boyer et al., 2006, Bonanno et al., 2000, 2003, 2004, Sheedy, 1998, Meese, 1990, Longin and Solnik, 1995, Arshanapalli and Doukas, 1993, Kenett et al., 2010, 2011, 2012b, Onnela et al., 2003, Sandoval Jr and Franca, 2012, Sandoval Jr, 2014, Curme et al., 2014, Aste et al., 2010, Tumminello et al., 2007]. Moreover, many studies have analyzed relationships between stock and foreign exchange markets, given the significant increase in global capital flows in the last two decades [Dornbusch and Fischer, 1980, Dooley and Isard, 1982, Morley, 2002, Nieh and Lee, 2002, Bae et al., 2003, Gagnon and Karolyi, 2006, Cappiello and De Santis, 2005, Granger et al., 2000, Pan et al., 2007, Ning, 2010, Zhao, 2010, Katechos, 2011, Lin, 2012]. Other studies have focused on global stock market return predictability offering diverse findings across different regions and time periods [Cohen and Frazzini, 2008, Cochrane, 2008, Fama and French, 1988, 1989, Welch and Goyal, 2008, Ferson and Harvey, 1993, Dahlquist and Hasseltoft, 2013, Breen et al., 1989, Harvey, 1991, Bekaert et al., 2009, Rapach et al., 2013].

There have been dramatic advances in the field of complex networks in many research fields. The world-wide-web, the Internet, highway systems, and electric power grids are all examples of networks that can be modeled using coupled systems [Barabási and Albert, 1999, Buldyrev et al., 2010, Mantegna, 1999, Watts and Strogatz, 1998, Albert and Barabási, 2002, Amaral et al., 2000, Li et al., 2015, Kenett et al., 2012a], where the connectivity between network components is essential. Similarly, the economic system is composed of many agents, interacting at different levels. The agents in the system could be individual traders, firms, banks, financial markets, or countries, hence the global financial system can be well represented by using a complex network model. Recently, researchers have used network theory to study economic systems as well as systemic risk propagation through the financial network [Billio et al., 2012, Gai et al., 2011, Anand et al., 2012, Haldane and May, 2011, Battiston et al., 2012a, Huang et al., 2013, Schweitzer et al., 2009, Glasserman and Young, 2015, Acemoglu et al., 2015, 2012, Dehmamy et al., 2014b, Ellis et al., 2014]. In this chapter, a two-layer interdependent network is developed and analyzed, where each layer represents a different financial market and interactions exist not only within the same market, but also between the two layers. Because of these interdependences, failure in a certain network node can trigger global systemic risk and crisis propagation to other nodes in the network. This study comprises major global stock market indices and their corresponding currencies as the two layers in our coupled network model.

Stock markets are a common trade place for company shares thus reflecting companies' performances and investors' perceptions of company values. Moreover, stock markets are considered leading economic indicators and therefore useful as predictors of the economy. The foreign exchange market is the largest financial market in the world, with market participants actively involved in currency trading 24 hours a day

except weekends, with daily turnover of over 5 trillion US dollars, according to the Bank for International Settlements [BIS Monetary and Economic Department, 2013]. These two financial markets capture important aspects of a country’s economic status, and therefore, we use them as a centerpiece of our research. We use a complex network approach to model the interaction between stock and foreign exchange markets to capture the topology as well as the dynamics in this coupled financial system. We study 56 stock market indices and 45 distinct currencies since 12 of the countries in our data set use the euro as their official currency. The analysis reveals novel insights and interesting features of the interactions among global stock and currency markets. We divide the entire period of 2002 to 2012 into two time intervals, non-crisis (2002-2006) and crisis (2007-2012) periods. The research presented here finds that correlations exhibit different behavior during the crisis period such as higher stock market correlations and lower foreign exchange correlations when compared to the non-crisis period.

The rest of the chapter is organized as follows: The data set and the methods are described in Section 2.2. Section 2.3 presents the correlation-based community analysis results for two different sub-periods, the non-crisis period (2002-2006) and the crisis period (2007-2012); Section 2.4, provides a discussion of the findings.

2.2 Materials and methods

2.2.1 Data

We acquired the data from the Boston University Bloomberg terminal provided by Bloomberg L.P. for academic research. In our analysis, we used a time range of 11 years from January 1st, 2002, to December 31st, 2012, with daily frequency. We used the daily closing price. Weekends were excluded, and for holidays we repeated the closing price of the previous day. We selected 56 representative countries, which

included developed as well as emerging countries. The list of the countries in the data set can be found in the appendix in Table A.1.

When a country has more than one well-known major stock market index, the following criteria were used to select the most representative index: first we selected the stock market indices that are widely used in the financial industry, and from that subset, we selected the index that includes most of the companies listed on the respective exchange, covering mostly large capitalization stocks and in some instances including some mid or small capitalization stocks as well. Using this criteria, we have selected one single index for each country, such as the S&P 500 for the US, Nikkei 225 for Japan, STI for Singapore, etc.

For the foreign exchange rates, we used the closing mid price, which is the average of the closing bid and ask prices. The foreign exchange market convention is for the majority of the currencies to be expressed in terms of USD, except for the British pound, Australian dollar, New Zealand dollar, and the euro that are all expressed as USD per currency. For consistency, these four currencies were converted to be expressed in terms of USD. For any country in the eurozone that had adopted the euro at a later date than January 1st, 2002, we kept the separate currencies up to the date before the actual euro adoption, and replaced them with euro after the adoption day. This accounts for the somewhat independent movement of these currencies before the adoption of the euro [Slovenska, 2005, OECD, 2004]. Examples include the Maltese lira and Slovak koruna, which were replaced by the euro on January 1st, 2008, and January 1st, 2009, respectively.

In order to also include the US dollar as part of our analysis, we have expressed all currencies as currency over SDR. The SDR is an international reserve asset, created by the International Monetary Fund (IMF) in 1969 to supplement its member countries' official reserves. At the time of the analysis, its value was based on a basket of four key

international currencies, and SDRs could be exchanged for freely usable currencies. Though it is not a currency, *per se* the SDR can be used as a currency unit in our study [Frankel and Wei, 1993, Frankel et al., 1995]. The IMF fixes the value of one SDR in terms of US dollars daily.

2.2.2 Pearson correlation analysis

For all analyzed time-series, first the logarithmic returns were obtained as follows: $S(t)$ is the value of the time series at time t , where $t = 1, 2, \dots, N$, and $N = 2866$ days is the size of the time series. The logarithmic return of time series i is

$$R_i(t) = \ln S_i(t) - \ln S_i(t-1). \quad (2.1)$$

We then rescaled $R_i(t)$ to have zero mean and unit standard deviation,

$$r_i(t) = \frac{R_i(t) - \langle R_i \rangle}{\sigma_i}, \quad (2.2)$$

where $\langle R_i \rangle$ and σ_i are the mean value and standard deviation of time series i .

For stock market indices, the dimension of the logarithmic return time series matrices is 56×2865 . For foreign exchange markets, it is 45×2865 since all Eurozone countries are grouped together. We calculated the cross-correlation matrix C , and the value of each entry c_{ij} as follows,

$$c_{ij} = \langle r_i r_j \rangle, \quad (2.3)$$

where c_{ij} represents the Pearson correlation between rescaled logarithmic return time series for a pair of countries i and j .

2.3 Results

2.3.1 Pearson correlation analysis

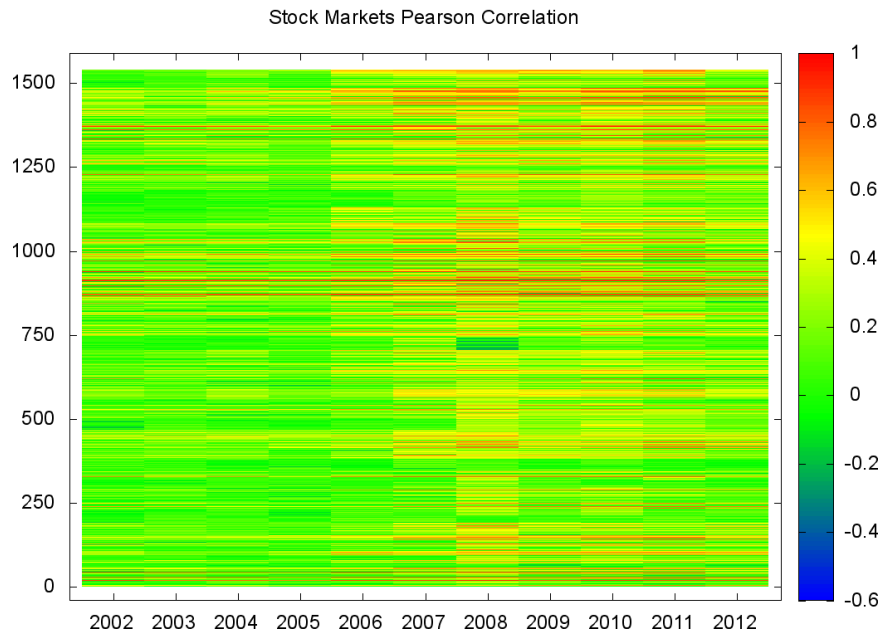
In order to see how the cross-correlation trends change with time, we divided the data into 11 annual periods. Figure 2·1(a) shows a heat map of the yearly stock market correlations. The x-axis represents the years, and the y-axis shows the 1540 unique pairwise correlations, excluding the diagonal of the correlation matrix. We used the color bar to show the magnitude of the correlations, where red means $c = 1$ and the two series are perfectly positively correlated, while blue means $c = -1$ and the two series were perfectly negatively correlated. Green means $c = 0$, or the two series were not correlated.

Figure 2·1(a) indicates that the overall correlation of the global stock markets trended upward starting in 2006, reached its peak in 2008, and stayed at a high level thereafter. These increasing correlation trends match the global financial crisis of 2007-2008 and could possibly be regarded as indicators of increased co-movements heading towards financial crisis.

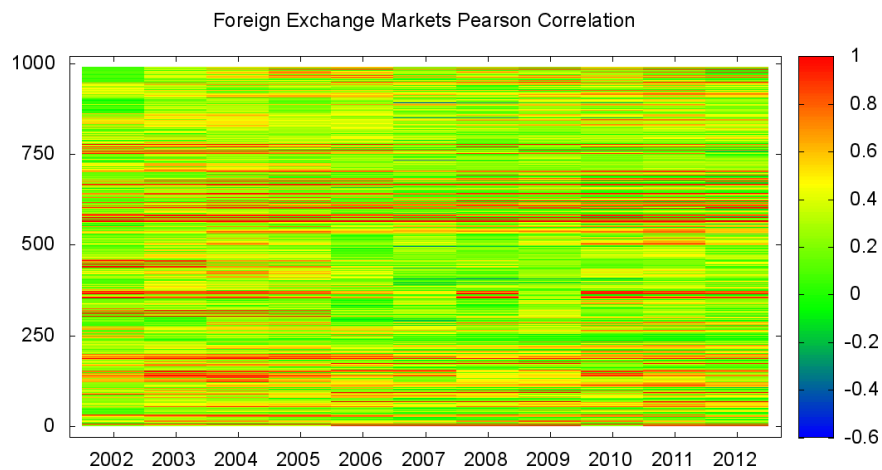
In Figure 2·1(b), the heat map for foreign exchange markets is plotted. It was generated in the same way as the heat map for the stock markets; however the number of entries is lower due to the lower number (45) of distinct currencies among the 56 countries. Hence, 990 distinct correlations make up the currency correlation matrix, excluding the diagonal. Generally foreign exchange markets showed stronger correlation compared to stock markets. It seems that the overall foreign exchange correlation fell during the financial crisis period, contrary to the stock market correlation trend.

2.3.2 Summary statistics

To study the statistical characteristics of the correlations, we calculated the mean and the standard deviation of any of the correlation distributions for the 11 years for



(a)



(b)

Figure 2.1: Heat maps of the annual Pearson correlations for 2.1(a) stock markets and 2.1(b) foreign exchange markets logarithmic returns. For the stock markets, we considered all 56 countries; for the foreign exchange markets, we considered the 45 distinct currencies. The color shows the value of the correlations for different years, where red indicates strong positive correlation, blue indicates strong negative correlation, while green means that the correlation is weak.




Figure 2-2: (a) Annual mean correlation and (b) annual standard deviation for stock markets (x), foreign exchange markets (+), and between stock and foreign exchange markets (o). The stock markets were more correlated during the crisis period, whereas the foreign exchange market correlation were higher in the non-crisis period. The correlation between stock and foreign exchange markets was lowest during the crisis. The standard deviation for both, stock and foreign exchange markets remained fairly constant, while the standard deviation for the interlayer correlations increased.

all 3 cases, including Pearson correlations for stock market layer, currency layer, as well as inter-layer correlations.

Figure 2-2 shows that the mean value of stock market correlations reached a peak during the crisis period (2007-2012). This finding suggests (similar to Fig. 2-1(a)) that during crisis the overall correlation increased and the stock markets tended to move together. This could be due to portfolio re-balancing, reducing equity market exposures and increasing allocations in the bond market. Such re-balancing in global portfolios occurs across different countries and produces declining stock market trends internationally.

On the contrary, the means of the foreign exchange market Pearson correlations were low during the crisis period, while they peaked during the non-crisis period. This finding suggests that in general the correlations among currencies are low during crises and high during non-crisis periods, contrary to the stock market behavior. For both,

	Stocks		Currency	
	2002–2006	2007–2012	2002–2006	2007–2012
Mean	0.175	0.315	0.370	0.318
Standard Deviation	0.208	0.249	0.260	0.279
Minimum	−0.320	−0.340	−0.155	−0.575
Maximum	0.965	0.966	1.000	1.000

Table 2.1: Statistics for the pre-crisis and the crisis periods.

however, the standard deviation of the correlations remained fairly constant, which suggests uniform increase (decrease) of the correlations in the stock markets (foreign exchange markets) throughout the entire time period.

In order to confirm this qualitative description that the means of the correlation were in fact different for the years in the calm period and for the years in the crisis period, we performed Student’s t-test. The null hypothesis was that the annual mean correlations were the same for both periods and that any differences come from the standard deviation. We applied the two-sample Student’s t-test. It rejected the null hypothesis that the two means were equal, and the p-value was 0.0002/ Thus the alternative hypothesis of unequal means was accepted.

The mean value of Pearson correlations between stock and foreign exchange markets exhibited a local minimum during the crisis period, which could be interpreted as positive stock market returns corresponding to currency appreciations. Overall the interlayer correlation did not change much throughout the entire period. The increase in standard deviation, however, indicates a larger spread in correlation values between the two layers during the crisis.

We distinguished between two different markets in our analysis, the network of stock market indices and the network of foreign exchange markets. Besides their obvious differences, the two markets also exhibited different correlation distributions, as visual inspection of Figures 2.1(a) and 2.1(b) hints. Performing the two-sample

Kolmogorov-Smirnov test (K-S test) confirmed this, with the null hypothesis being that the two sample data were from the same continuous distribution. The K-S test rejected this null hypothesis at a significance level $\alpha = 0.001$. Therefore we concluded that the correlation distributions of stock markets and of foreign exchange markets were different.

2.3.3 Community formation and cluster analysis

In our analysis we focused on two periods: a period of relative economic calmness (2002-2006) and a period of economic crisis (2007-2012). Using the correlation information, we depicted the network structure of the stock markets and of the foreign exchange markets for each of these time periods. Since countries have different peak trading times due to their respective geographical locations and different time zones, lagged correlations allow to infer information about regional influences. We separated the 56 countries and 45 currencies into 3 groups according to their geographical location. When we considered synchronous correlations, we used the returns at time t for every country. In that case, the Asian markets are the first to trade, then the European markets, and finally the American markets. A shock originating in the US, for example, would then not show its immediate effect on the other markets because the stock markets in other parts of the world were closed for most or all of the trading hours of the NYSE. When we considered lagged correlations, we used the returns at time t for the American countries and at time $t + 1$ for the other countries of the world. In that case, it is as if the American markets were first to trade, followed by the Asian and then the European markets. A shock from the US would then be very visible. For both, the stock and the foreign exchange markets, we first considered the synchronous correlations and then compared the results to those of the lagged correlations.

We used Planar Maximally Filtered Graphs (PMFG) [Tumminello et al., 2005,

Di Matteo et al., 2010] to study the properties of stock and foreign exchange market correlations. PMFG is useful for filtering meaningful correlations from the bulk of the 1540 correlation pairs, as it suppresses small correlations while maintaining the overall network structure. In order to build the graph, we ordered the correlations c_{ij} from largest to smallest. First, the pair with the largest correlation was connected. In subsequent steps we connected the countries i, j under the condition that a link between them maintained the planar structure of the graph; if it did not, the pair was skipped. This procedure resulted in an adjacency matrix with unweighted links from which the graph is plotted. We used *Wolfram Mathematica* to investigate the communities in the network, which were detected with respect to their modularity.

Stock markets

Synchronous correlations. Figure 2·3(a) shows the PMFG for the period from 2002 to 2006, where we identified 5 clusters which seemed to be organized by geographical locations. The cluster on the top left was led by Singapore, a financial hub, and Saudi Arabia, which connected to many other OPEC countries in the cluster. The second Asian cluster was centered around Hong Kong and Japan. The third cluster contained smaller European countries, organized around the Scandinavian countries. The fourth cluster contained the major European economies, except for Italy and Germany which were closely connected to the American countries in the fifth cluster, particularly through ties to the US and Canada. These four countries exhibited particularly strong connections within and, in case of Germany and Italy, to countries outside the cluster. The communities changed significantly during the crisis period, where Singapore and Hong Kong led a large Asian cluster, as illustrated in Figure 2·3(b). The American cluster formed during the non-crisis period became more mixed during the crisis period, as Norway, Iceland and Russia became part of this cluster. The change in the community comprised of mostly American countries

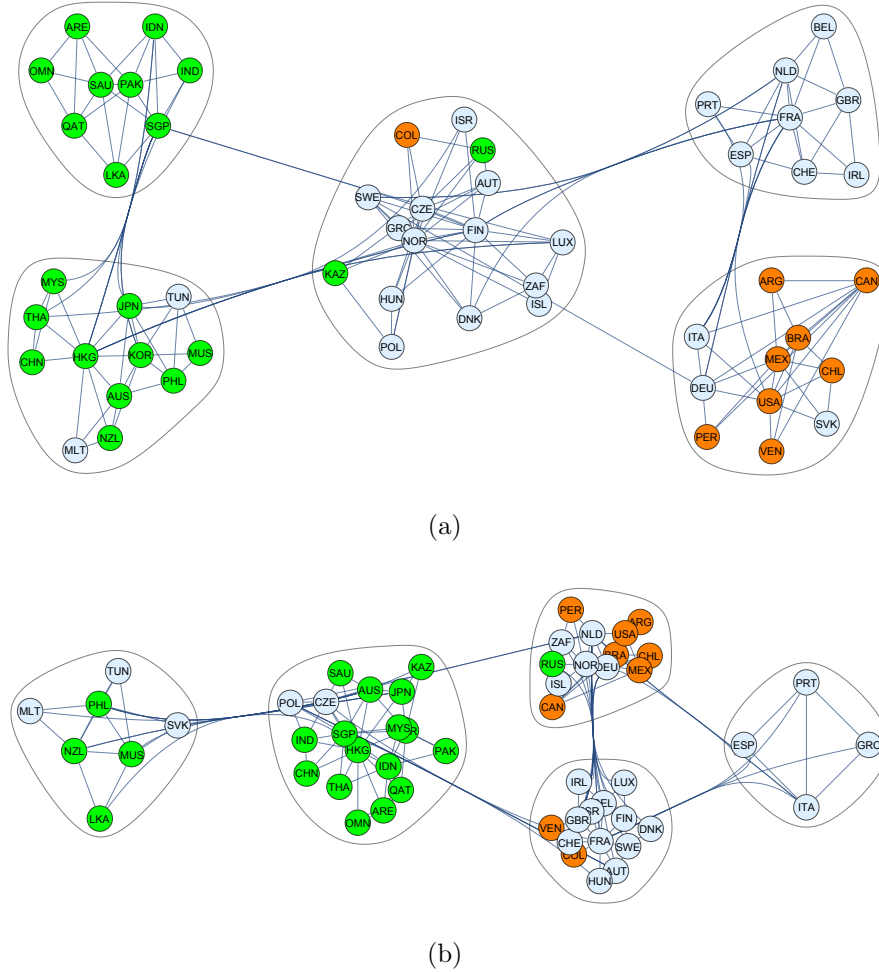


Figure 2-3: Planar maximally filtered graph (PMFG) for the stock markets obtained using synchronous correlations during (a) the economically calm period and (b) crisis period. The countries are denoted by their three-letter symbols and are color-coded according to their geographical locations: green for Asia, light blue for Europe and orange for the Americas. For the calm period, we detected five large clusters which are mostly geographically divided. The clusters significantly changed during the crisis period. Most notably the troubled Eurozone countries, Italy, Spain, Greece and Portugal, formed their own cluster. The Asian countries formed one larger cluster centered around Hong Kong and Singapore. We also observed a smaller cluster containing a diverse group of less connected countries.

suggests that Italy and Germany influenced the performance of the American markets in the non-crisis period, while during the crisis period, Norway and Russia seemed to increase their influence. The majority of the European countries, connected through France, formed another cluster, strongly linked to the countries in the American cluster. These observations are in line with the coordinated responses to the crisis from the US, large European economies, and the ECB; at the same time the Asian countries do not become more closely connected to them. Most notably, however, two new clusters appeared: one with the troubled Eurozone countries Portugal, Italy, Greece and Spain, and another consisting of rather less connected countries like Slovakia, Mauritius and Tunisia.

Lagged correlations. As pointed out before, lagged correlations are an important measure to study the influences of financial markets. They allow to consider effects originating in one region and spreading to another. We considered the correlation calculated for the returns for the Americas at time t , and for Asia and Europe at time $t + 1$. In other words, we focused on how the index movements in the Americas would affect the index movements in the rest of the world. Please note that such a procedure does not change the correlations within one geographical region, nor could the correlations between Asian and European countries change. However, due to the PMFG algorithm, larger correlations that appeared for countries in the Americas with countries in the rest of the world changed the structure of the network.

The European clusters remained mostly unchanged in the calm period, except that Germany and Italy no longer formed a community with American countries, when considering lagged correlations compared to synchronous correlations. Since here we considered the Americas at time t and Europe and Asia at time $t + 1$, it seems that the American countries, including the US, did not affect Italy and Germany significantly. The Netherlands, Italy, France, and the Scandinavian countries

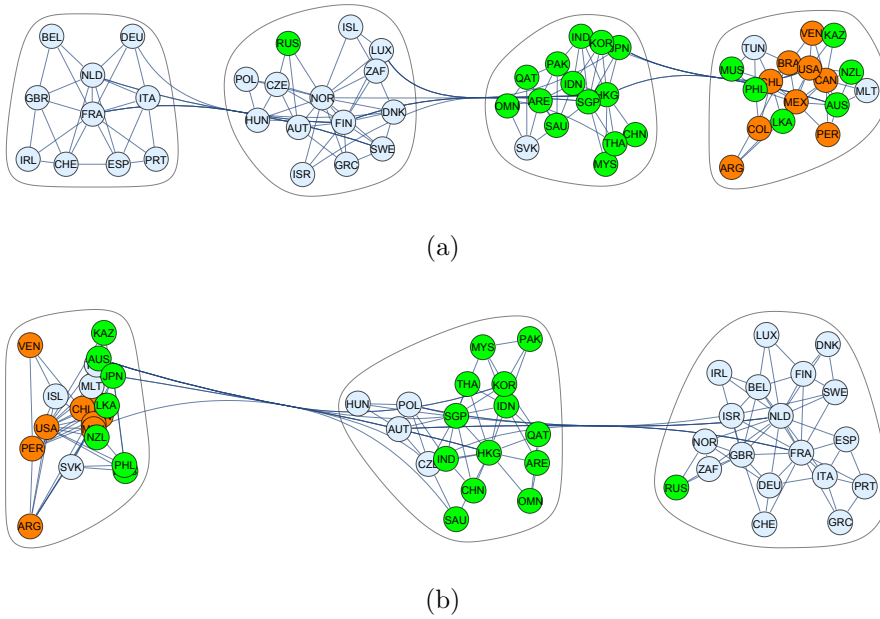


Figure 2-4: Planar maximally filtered graph (PMFG) for the stock markets obtained using lagged correlations during (a) the economically calm period and (b) crisis period. The countries are denoted by their three-letter symbols and are color-coded according to their geographical locations: green for Asia, light blue for Europe and orange for the Americas. For the calm period, we observed four large clusters, while in the crisis period the number of communities changed to three, as the two European clusters merged.

still manifested themselves as the most connected European countries. We did, however, observe large changes in the American cluster on the very right in Figure 2.4(a); Australia and New Zealand, economically close to the US, moved from the Asian community to connect more tightly to the American countries. Different responses by central banks, leading to higher interest rates particularly in Australia, can be considered a reason for this [Rodriguez, 2012]. They were joined by smaller countries, such as Sri Lanka and the Philippines, that we had previously identified as countries with weak links to other Asian countries. The Asian cluster was led by Singapore, Hong Kong and Japan, which displayed the most significant correlations to countries outside of their community. The community structure changed significantly during the years of economic turmoil in Figure 2.4(b). Most obviously, the number of communities reduced to three because the two European clusters from the non-crisis period mostly merged. Netherlands, United Kingdom and France were at the center of the European community in the crisis period. Portugal, Italy, Greece and Spain were found at the periphery of this cluster. Using synchronous correlations, Germany and Netherlands were tightly connected to the American cluster during the crisis period, while when considering lagged correlations, they belonged to the large European cluster. We observed that Japan, Australia and New Zealand, among others, formed a community with the American countries during the crisis, which suggests that the major financial markets in the Pacific followed the trends of the American stock markets. The importance of the US stock market is emphasized by the largest number of connections in this cluster. Australia was the second most connected country in this community.

Foreign exchange markets

While we analyzed a total of 56 global stock market indices, we only investigated 45 currencies because 12 countries in our data set use the euro.

Synchronous correlations. Figure 2·5(a) shows the dominant role that the euro played in Europe prior to the crisis. Together with the Danish krone, pegged to the euro via the European Exchange Rate Mechanism, the euro was connected with all European currencies, and it exhibited close ties to the Canadian dollar, the Australian dollar, and the New Zealand dollar. With 19 links each, the node for the euro and the node for the US dollar showed the highest interconnectedness among all currencies in the calm period. The US dollar is at the center of the cluster comprised of the majority of the American currencies and the oil-exporting countries. In Figure 2·5(b) we observed that the clear structure and hierarchy of the European community during the calm period fell apart during crisis. In addition, a fourth cluster appeared, comprised of European and South and Central American currencies, like the Brazilian real, the Mexican peso and the Chilean peso, which had been closely connected to the US dollar in the non-crisis period. As the financial crises unfolded, the Fed and its European counterparts employed “quantitative easing” as monetary policy, which in turn elicited strong criticism in the BRIC countries, as QE corresponded to currency devaluation [Blackden, 2012]. The US dollar maintained its strong ties with the currencies of oil-exporting countries. Its cluster was joined by the Japanese yen and the Chinese yuan. China moved to a managed floating regime during the crisis period [Reuters, 2012], whereas the Japanese government tried to stimulate its economy with policies similar to QE, known as “Abenomics” [BBC, 2011].

Lagged correlations. In Figure 2·6(a), using lagged correlations during the non-crisis period, we observed that the Euro was at the center of the cluster of European currencies. The Japanese yen joined this cluster. The US dollar and the currencies of the Asian (Middle Eastern) oil-exporting countries no longer shared close ties. Figure 2·6(b) reveals that during the crisis period there was a mixed cluster comprised of all the American currencies and a group of Asian currencies, including India, Singa-

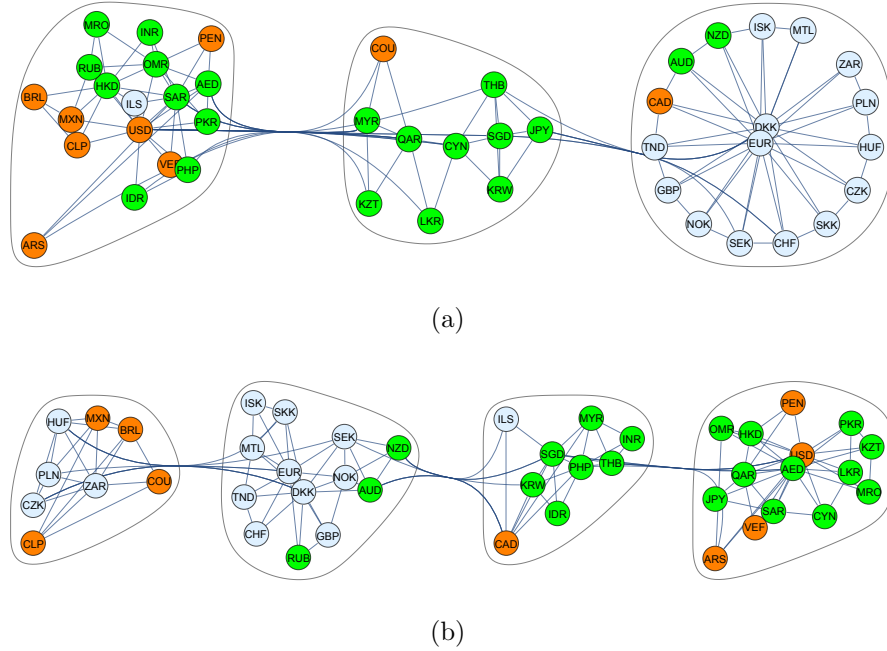


Figure 2-5: Planar maximally filtered graph (PFMG) for the foreign exchange markets obtained using synchronous correlations during (a) the economically calm period and (b) crisis period. The currencies are denoted by their three-letter symbols and color-coded according to their geographical locations: green for Asia, light blue for Europe, orange for the Americas. For the calm period, we detected three large clusters. Most distinctive was the European cluster shown on the right, with the euro and the Danish krone, which is pegged to the euro, at the center. The Asian countries split in two different clusters, with major oil-exporting countries being closely associated with the US dollar which were at the center of its community. The remaining Asian countries formed the third cluster. During the crisis, the hierarchy of the euro cluster collapsed. The community around the US dollar still contained the oil-exporting countries. It was joined by the Chinese yuan and the Japanese yen. The Brazilian real, Mexican peso and Chilean peso, however were no longer part of the US dollar-centered community during the crisis period.

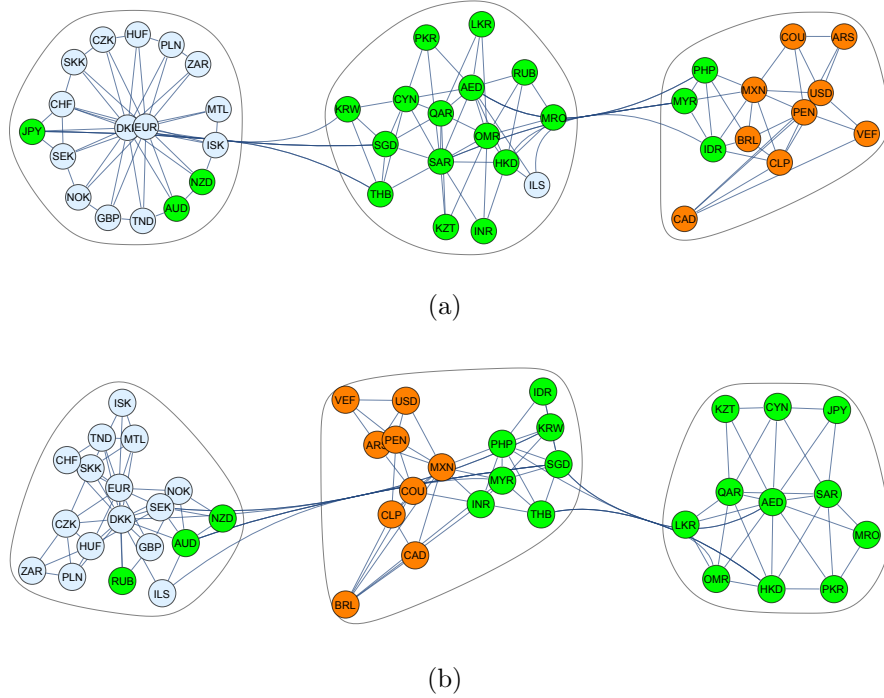


Figure 2-6: Planar maximally filtered graph (PMFG) for the foreign exchange markets obtained using lagged correlations during (a) the economically calm period and (b) crisis period. The currencies are denoted by their three-letter symbols and color-coded according to their geographical locations: green for Asia, light blue for Europe, orange for the Americas. For lagged correlations we used time t for the Americas and time $t+1$ for Europe and Asia. For both periods, we detected three large clusters. In the non-crisis period we observed that the clusters were determined by the geographical currency location, except for the Japanese yen, New Zealand dollar and Australian dollar. Iceland appeared in the predominantly Asian cluster, while Malaysia, the Philippines and Indonesia were part of the American community. During the crisis period the interconnectedness between American and South East Asian currencies increased as they belonged to the same community. The purely Asian cluster comprised the Japanese Yen along with the currencies of the oil-exporting countries. The Russian ruble joined the European cluster.

pore and Korea, while the South East Asian currencies formed their own community using synchronous correlations, as observed in Figure 2.5(b). The number of connections of the nodes comparing the two different periods remained stable; the majority of the currencies did not gain or lose more than two connections. Instead we find that different connections developed within the clusters. For example, Hong Kong lost the link with countries like India and Russia during the crisis, but became more closely connected to Singapore, the Philippines and Thailand.

Any currency is traded during every hour of the day, and therefore any sudden changes should be reflected in all the currencies on a time scale shorter than one full day. In fact, one would expect the strong correlations to correspond to shorter time scales. Hence the community structure of the currencies is depicted by synchronous correlations, while when using lagged correlations the true structure of the network disappears.

2.4 Discussion

In this study, we have investigated the daily logarithmic returns in the stock and foreign exchange markets of 56 countries and 45 currencies. We have used network theory and community analysis to understand the structure of the coupled financial network formed by global stock market indices and currencies. We defined weighted links within a network layer and between the two layers (stock markets on one hand and currencies on the other) using Pearson correlation. The overall correlations within stock markets increase during the crisis period, and the overall correlations within foreign exchange markets, as well as the correlations between stock and foreign exchange markets decreased during this period. We have investigated statistical properties of our results by presenting correlation summary statistics and performing the K-S test to closely study the characteristics of the correlation distributions. We applied the

PFMG method to discover distinct community formations and categorize the countries into clusters. We identified the importance of the countries according to their relative positions in the communities and the strength of their links with other member of the community as well as the strength of their external links with countries that do not belong to their cluster. In our analysis, we have distinguished between synchronous and lagged correlations. We divided the countries in three geographical regions, Asia, Europe and Americas. To study lagged correlations, we considered the returns at time t in the Americas, and the returns at time $t + 1$ in Asia and Europe. The comparison of the network and community structures allowed us to infer influences from one region to another, in both, crisis and non-crisis periods. Using synchronous correlations we have identified five clusters in the stock market during the non-crisis period. These clusters changed their structure during the crisis period, where, for example, four of the five troubled Eurozone countries, Portugal, Italy, Spain and Greece, formed their own community. In the case of lagged correlations, we observed four clusters in the stock market layer during the non-period period and only three communities during the crisis period. We observed that the American countries are most closely connected to Asian countries. The introduction of a time lag and the onset of the crisis caused the European countries to become more tightly connected. Unlike stock markets that are bound to different time zones, currencies trade around the clock. Any influences among currencies are immediately reflected in their returns, hence synchronous correlation define the true network structure of the foreign exchange markets. We observed that the Euro played a central role among the European currencies during the non-crisis period, while it lost its central position during the crisis-period. The US dollar was closely linked to the currencies of the major oil-exporting countries, both, during non-crisis and crisis periods.

These findings could have policy implications and could be helpful for central

bankers, policy makers and regulators, offering a tool for identifying tightly related communities of countries by stock market performance and by currency dynamics. If for instance a financial crisis originates in a specific country, the most vulnerable countries, where the crisis might spread first, are the countries with stronger links with the originating country, most likely belonging to the same community. If policy makers have knowledge of these communities, they might be able to focus needed bailout funds or implement temporary preventative measures in the most vulnerable countries, thus reducing the impact of inherent global financial crisis to the rest of the world and preventing the propagation of the crisis before severe damages cripple the entire economic system.

Chapter 3

Overlapping Portfolios: a Source of Systemic Risk

Stability of the banking system and macroprudential regulation are essential for healthy economic growth. In this chapter we study the European bank network and its vulnerability to stressing different bank assets. The importance of macroprudential policy is emphasized by the inherent vulnerability of the financial system, high level of leverage, interconnectivity of system's entities, similar risk exposure of financial institutions, and susceptibility for systemic crisis propagation through the system. Current stress tests conducted by the European Banking Authority do not take in consideration the connectivity of the banks and the potential of one bank vulnerability spilling over to the rest of the system. We create a bipartite network with bank nodes on one hand and asset nodes on the other with weighted links between the two layers based on the level of different countries' sovereign debt holdings by each bank. We propose a model for systemic risk propagation based on common bank exposures to specific asset classes. We introduce the similarity in asset distribution among the banks as a measure of bank closeness. We link the closeness of asset distributions to the likelihood that banks will experience a similar level and type of distress in a given adverse scenario. We analyze the dynamics of tier 1 capital ratio after stressing the bank network and find that while the system is able to withstand shocks for a wide range of parameters, we identify a critical threshold for asset risk beyond which the system transitions from stable to unstable.

3.1 Introduction

In today's interconnected world the financial system with increased global reach is becoming more vulnerable to international corporate, household, and sovereign exposures. Drawing national business boundaries for banks has become a difficult task. Even though in general banks still have significant exposures to domestic economies, the largest financial institutions have a substantial international presence. Systemic risk describes the vulnerability of a system, like the bank network, as a whole to exogenous and endogenous shocks. If not understood and closely monitored, these shocks can wipe out significant parts of a system. Therefore, monitoring the vulnerabilities of banks globally is highly important for policymakers. Stress tests are an important tool for determining the fragility of banks under adverse scenarios. While there are many factors that determine the health of financial institutions, (i) adequate level of capital and (ii) high quality liquid assets are two main aspects to consider when assessing banks' stability. Additionally, (iii) examining the interconnectivity among banks plays an important role in determining the macroprudential fragility of the network.

In this paper we address all three important factors from above and demonstrate that if banks are well capitalized, hold liquid high quality assets, and maintain well-balanced portfolios in relation to other banks, we can expect better overall stability of individual banks as well as the banking system as a whole. While the first two factors are well studied (Moyer [1990], Beltratti and Stulz [2012], Calice et al. [2013], Distinguin et al. [2013], Miles et al. [2013], Bank [2014], DeYoung and Jang [2016]), we offer a novel insight into the importance of the network structure and dynamics of the system of international financial institutions through the study of the third factor.

We propose a bipartite network model of banks on one hand and assets on the

other and test the stability of the network under specific adverse scenarios. The links between the banks are established indirectly through common exposures to asset classes such as corporate loans or commercial real estate loans, and the weights of the links are approximated based on bank exposures to sovereign debt of different countries. Our model has two parameters: (i) size of the initial shock to the banking system, and (ii) spreading or spillover parameter, which we identify as a measure of systemic risk. The initial shock reduces bank capital or increases the risk weights of bank assets. Both types of shocks cause a deterioration in the tier 1 capital ratio, which is used as benchmark to assess proper capitalization of financial institutions. This deterioration prompts a response from the affected banks which can further distress other parts of the system. We confirm, in accordance with Eisenberg and Noe [2001] and Glasserman and Young [2015], that any spillover effect in a linear system is not significant, and it is bounded by the initial shock to the system, i.e., no systemic effect is observed through network spreading dynamics. To investigate the conditions when systemic risk propagates through the banking network, we introduce non-linearity in our model. We simulate a non-linear bank response to a loss of equity or an increase of specific asset risk weights. In the non-linear scenario, we observe increased fragility of the banking system. The spreading parameter which describes asset vulnerability to risk is a critical parameter that separates the stable from the unstable regime of the system.

3.2 Literature review

Central banks often work in concert to steer the financial system towards stability. Just recently, researchers from 13 central banks joined efforts to study data from 25 different markets to understand how to reconstruct exposures in financial networks (Anand et al. [2017]). Financial crises such as the sub-prime credit crisis of 2007-2009

and the European sovereign debt crisis of 2010-2012 have revealed weaknesses in bank risk management practices as well as softness in the global regulatory framework. As a response to the economic system instabilities, regulators have focused on strengthening liquidity and capital requirement rules to increase the stability of the financial system and reduce the possibility of major negative impact on global economy (Basel Committee on Banking Supervision [2010], Congress [2010])

The Basel Committee on Bank Supervision has developed the first Basel Capital Accord in 1988 to address apprehensions arising from financial deregulation. The Basel III accord (endorsed by the G20 countries in November 2010) responds to the need for effective regulation to maintain the stability of the financial system in times of economic downturns, such as the financial crisis of 2007-2009. More specifically, Basel III calls for improvement in quantity and quality of capital, redefines tier 1 capital, weighs on banks' risk management, and calls for capital buffer requirements to increase the stability of the entire financial system (macroprudential regulation). Previous Basel regulations (I and II) have focused on the stability of the financial system's entities (microprudential regulation), disregarding the systemic risk and the vulnerability of the financial system to cascading failures. Similarly to Basel I and II, Basel III maintains the requirement that banks hold total capital of 8 percent of their risk-weighted assets (RWA). One of the main differentiating aspects of Basel III however is the introduction of a more stringent definition of tier 1 capital as a "going-concern" capital comprised mostly of common equity (common stock, common stock surplus, and retained earnings). Additionally, according to Basel III, 75 percent of the total banks' capital should consist of tier 1 capital. Common equity tier 1 capital should account for at least 4.5% of RWA of the bank.

The RWA are fundamental input for Basel III capital requirements and are determined by an internal rating-based (IRB) approach that assesses counterparty credit

risk (CCR). The majority of banks use external credit ratings attributed to their counterparties. Large banking institutions, however, may choose to use internal risk models to determine the capital needed to offset specific RWA, based on their estimates of exposures to loss or likelihood of loan defaults (King and Tarbert [2011], Padgett [2012]).

One of the challenges in macroprudential regulation is proper risk assessment of financial institutions and ensuring that, on a systemic level, the risk of network failure is minimized. Since the bank risk is based on exposures to different sectors of the main economy as well as the liquidity of the rest of the financial system, having an accurate assessment of asset risk factors as well as ensuring high quality tier 1 capital is essential. The Basel Committee on Banking Supervision introduced the Basel III reforms in 2010 to address the vulnerabilities in the financial system focusing on the liquidity requirements expressed through the Liquidity Coverage Ratio (LCR) and the Net Stable Funding Ratio (NSFR). Basel III stirred numerous arguments between supporters of the reforms who argue that Basel III will increase the stability of the financial system and opponents who claim that the reform will reduce credit availability and economic activity (Admati et al. [2011], Admati and Hellwig [2011], Allen et al. [2012], King [2013], Miles et al. [2013])

Given that certain banks are allowed to use internal models for calculating risk-weighted assets (RWA), and hence influence their tier 1 capital ratios, it is not always possible to state the proper capitalization of banks with great confidence. The banks can influence tier 1 capital ratios by holding assets weighted with a low risk that actually have higher risk. Additionally, there is a risk that asset risks might change over time. Moreover, correlations of banks' shared portfolios are prone to increase and contribute to higher risk in the banking system (Engle [2009], Acharya et al. [2014], Caccioli et al. [2014b], Brownlees and Engle [2015], Corsi et al. [2016]).

The establishment of the European Single Supervisory Mechanism (SSM) in November 2014 opened a new era of bank supervision in the euro zone. The intention of the SSM has been to harmonize key areas of bank supervision and to contribute significantly towards the safety and resilience of the European banking system. By conducting a comprehensive assessment of 130 banks in the euro area, the European Central Bank (ECB) has addressed several important objectives including: (i) strengthening bank balance sheets, (ii) enhancing transparency and improving quality of information regarding bank conditions, as well as (iii) building confidence by assuring banks' appropriate capitalization after completion of necessary corrective actions. The 130 banks, involved in the comprehensive assessment, account for €22 trillion total assets or over 80% of total assets in the SSM as of December 2013 (Authority [2016]). Similar efforts to regulate the financial system have been made in the US, most prominently in form of the Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010 (Congress [2010]). Just recently though, there have been significant efforts by the new US administration to revoke some of the regulations put in place by this bill (Rapperport [2017]). This shows that there is a significant political component to regulation, specifically regarding how much or how little regulation is appropriate. The focus, however, should not necessarily be on more or less regulation, but rather better regulation in order to curb banking crises or reduce systemic distress in the financial network.

This is all the more urgent as banking crises have been more frequent than expected in both developed and emerging market economies with an annual probability of a crisis of approximately 4-5 % (Walter [2010]).

Stability of financial institutions has attracted the attention of many economists who have offered insight into financial system fragility, measures of systemic risk as well the impact of regulation on systemic risk. In addition to central banks' stress

test methodologies, Acharya et al. [2014] have proposed an alternative approach to measuring systemic risk (SRISK) by offering an insight into how much capital a financial institutions will need to raise during economic crisis to bring their capital up to regulatory levels.

Comparison studies between regulatory stress tests and stress tests that use only public market data reveal that approaches used to weight assets by risk are not correlated with market measures of risk. The well capitalized banks also have not fared better than the rest of the European banks in light of the European sovereign debt crisis of 2011 (Acharya et al. [2012, 2014], Lucas [2014a]). On the other hand, Yan et al. [2012] find Basel III reforms to have significant long-term positive effects on the UK economy. Lucas argues that in addition to financial institutions, which are tightly interlinked with the main economy, the Government is also a significant source of systemic risk. While other factors such as lack of transparency of government actions and the scope of government's involvement in financial markets can contribute to overall systemic risk buildup in the economy, a notable systemic characteristic of the government is its enormous size as financial conglomerate. When considering the traditional credit programs, Fannie Mae, Freddie Mac, the Federal Home Loan Banks, deposit insurance, the Federal Reserve System, and the Pension Benefit Guarantee Corporation, the government becomes a \$20 trillion financial institution (Lucas [2014b]). The government, on the other hand, can also serve as an essential contributor to financial stability in times of crises undertaking appropriate actions carried out swiftly. Both aspects of government, carrying significant costs, may contribute to rethinking the notion that government debt is risk-free and hence reconsider sovereign debt risk weights used to determine banks' RWA.

To incorporate the complexity of the financial system, interdisciplinary approaches have been proposed. Understanding the interconnections among financial institutions

as well as interlinks between bank networks and the main economy have been a focus of many researchers during the past two decades starting with Eisenberg and Noe [2001] and Furfine [2003]. Using tools from network science, researchers have studied the likelihood of contagion (Upper [2011], Glasserman and Young [2015]), developed systemic risk measures (Battiston et al. [2012b]), and analyzed specific banking systems (Van Lelyveld and Liedorp [2006], Upper and Worms [2004]). Other approaches include the study of robustness, the cost of repair, and topological properties and their consequences for systemic risk (Garlaschelli and Loffredo [2004, 2005], Elsinger et al. [2006], Caldarelli [2007], Iori et al. [2008], Huang et al. [2011], Vitali et al. [2011], Battiston et al. [2012a], Dehmamy et al. [2014a,b], Joseph et al. [2014], Piškorec et al. [2014], Battiston et al. [2016], Majdandzic et al. [2016], Vodenska et al. [2016a]).

Following a shock, the natural response of financial institutions includes reducing losses by selling assets, which most likely trade at depressed prices: due to the distress, liquidity may dry out and a lack of buyers reduces market value. (Huang et al. [2013], Greenwood et al. [2015]). Previous studies have considered fire sale dynamics, in which banks attempt to unload their assets as efficiently as possible, leading to a downward spiral exacerbated by overlapping portfolios (Duarte and Eisenbach [2015], Cont and Schaanning [2017]). We expand upon this approach and introduce a behavioral component by including the risk tolerance of banks, modeling financial institutions' inclination to induce a fire sale. We suggest that by being able to monitor the dynamics of links in the bank network, regulators might be able to foresee specific bank network vulnerabilities in light of certain exogenous or endogenous economic shocks.

The rest of the chapter is organized as follows: In Section 3.3 we describe the data that we use for empirical analysis. In Section 3.4 we introduce the simulation model of systemic risk propagation through the bipartite network of banks and assets, while

in Section 3.5 we present the results of the simulation using the European Banking Authority stress test results data. Here we highlight the effect of inter-connectivity of the banks through shared portfolio networks. In Section 3.6 we discuss the regulatory implications of our results and offer our concluding remarks.

3.3 Data

We use the data from the 2011 European Banking Authority (EBA) stress test. The data from the EBA offers insight into asset portfolios of European banks. It shows the total exposure of 90 different banks in the following seven investment categories: sovereign debt, financial institutions, corporate, retail residential mortgage, retail revolving, retail small and medium-sized enterprises (retail SME), and commercial real estate (CRE). Figure 3·1 shows the structure of the banks' holdings. For each of the seven asset classes we plot the histogram of the percentage they make up in the banks' portfolios.

We observe that banks tend to hold large amounts of corporate loans and assets in the residential retail sector. Sovereign debt and loans to the financial sector play a smaller roll, while retail revolving, retail SME, and commercial real estate tend to make up the smallest part of the banks' portfolios. This is reflected by the means, indicated by dotted lines in the plot. An average bank portfolio is comprised of roughly 14% sovereign debt, 15% loans to financial institutions, 30% loans to corporate, 26% residential mortgage, 3% revolving retail, 5% retail SME and about 7% commercial real estate. The data set also details which kind of country's sovereign debt each bank holds. These countries include 30 European nations, the United States and Japan as well as a category "other".

Fig. 3·2 represents a network representation of the portfolio similarity of the banks. Since all banks exhibit a certain overlap in portfolio, we filter the network

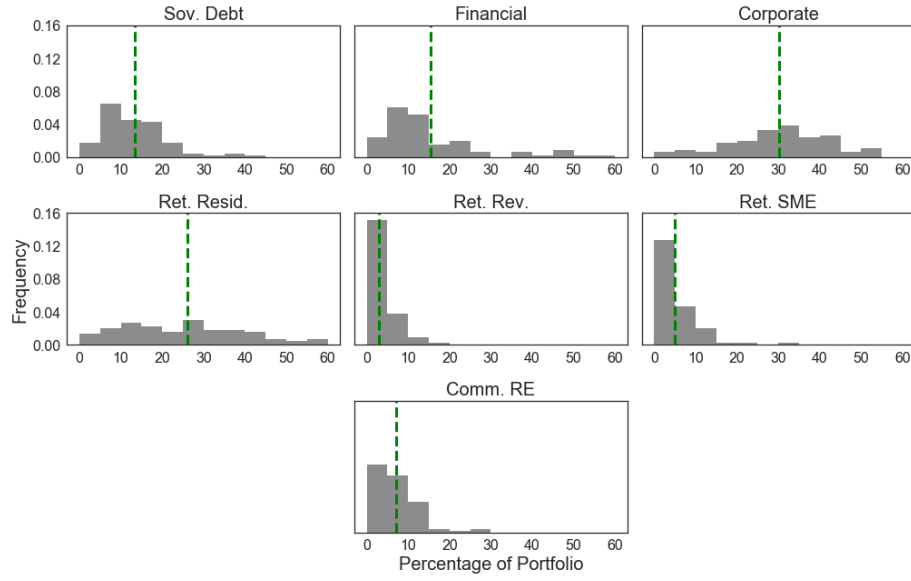


Figure 3-1: Histograms of bank holdings where the x-axis is the percentage of a bank’s portfolio an asset makes up. The dotted line indicates the mean value of the distribution.

using a method introduced by Tumminello et al. [2005]. Such a planar maximally filtered graph (PMFG) reduces the number of edges for information filtering and better graphical representation. A link, drawn in gray, between two banks indicates that there exists a significant overlap between their respective sovereign debt portfolios. The minimum spanning tree Kruskal [1956] of the graph is indicated by the thicker, solid edges. We define the overlap of two portfolios as their cosine similarity: given portfolios A and B with weights a_i, b_i distributed among asset classes $i = 1, \dots, N$ and $\sum_i^N a_i = \sum_i^N b_i = 1$, the cosine similarity is given by

$$\text{similarity} = \frac{\sum_{i=1}^N a_i b_i}{\sqrt{\sum_{i=1}^N a_i^2} \sqrt{\sum_{i=1}^N b_i^2}}, \quad (3.1)$$

which is the normalized inner product of the relative weights of portfolio A with the relative weights of portfolio B.

This product yields a number between 0 and 1 where a value of 0 means that

portfolio A does not contain any asset that is in portfolio B and vice versa, and a value 1 means that both portfolio holders allocate their money equally among the available assets. The size of the nodes and their color correspond to the value of all assets in the respective bank's portfolio and to the country or region in which it is headquartered.

The largest banks by asset value are positioned at the center of the minimum spanning tree. The French bank Société Générale, which is the 11th largest bank, has the most, 13, connections to non-domestic banks. The French bank BNP Paribas (2nd largest), the Belgian bank Dexia (17th largest), and the British bank HSBC (largest bank in the data set) follow in second place with 10 connections to banks from other countries. We can roughly separate the banks into two groups. On the one hand we have around 20 of the largest banks which show significant portfolio overlap with banks from other countries, indicating a broad and international portfolio of sovereign debt. On the other hand we find the remaining banks sharing a larger overlap with banks from their own country or neighboring countries. For example, sovereign debt holdings by Danish banks tend to be more similar to sovereign debt holdings of other Scandinavian banks as compared to the portfolios of banks from other European countries. These home and regional biases may be due to regulatory requirements such as Basel II, which have allowed to favor domestic sovereign debt¹, and probably reflect the focus of a bank's business operations.

The EBA data also details the tier 1 capital ratios of banks which is defined as the ratio of tier 1 equity and the sum of the risk-weighted assets of a bank. The risk-weighted assets of a bank are calculated by weighting the exposure to an asset by an estimate of their riskiness. Until 2013 Basel III has required banks to maintain a tier 1 capital ratio of at least 4.5 percent at all times. In line with the phase-in

¹ <http://www.bis.org/publ/bcbs128b.pdf>

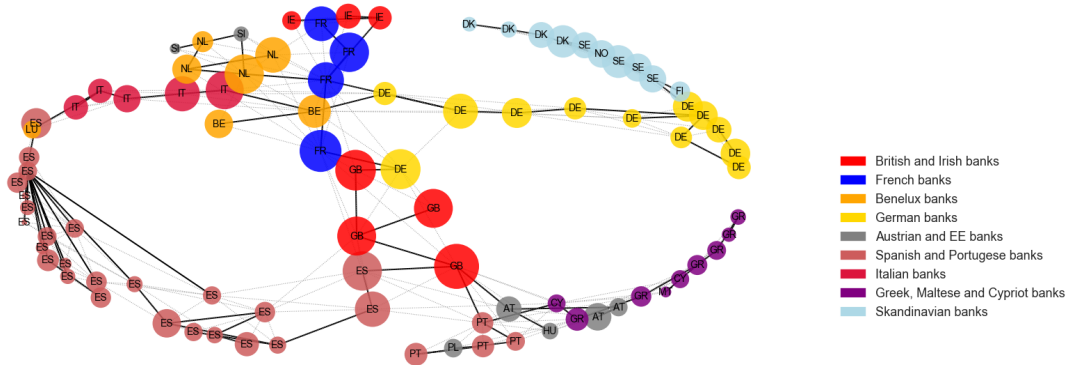


Figure 3-2: PMFG (Planar Maximally Filtered Graph) and MST (Minimum Spanning Tree) network showing portfolio overlap in sovereign debt of European banks, including British and Irish banks (bright red), French banks (blue), banks from the Benelux countries (orange), German banks (yellow), Austrian and Eastern European banks (gray), Spanish and Portuguese banks (light red), Italian banks (crimson red), Greek, Maltese and Cypriot banks (purple), and Scandinavian banks (light blue). The size of the node represents the value of all assets a bank. The largest banks are at the center of the network. We observe that banks from the same country cluster, indicating that their portfolios are similar due to home and regional bias.

arrangements of the Basel III framework², the requirements are currently higher than that; however we recognize that we need to compare the situation of the banks in our data set with the regulations at the time. The 2011 EBA report includes two banks that have a capital of less than 4.5 percent already, however, which we have to remember in the analysis of the results of our simulations. These banks are the Allied Irish Banks (tier 1 capital ratio of 3.7%) and the Spanish Caja de Ahorros (3.8%). Hypo Real Estate exhibits the largest value among all banks, with a tier 1 capital ratio of 28.4 percent. The mean and median capital ratio of the banks in the data set are 9.3 percent and 8.7 percent, respectively.

² http://www.bis.org/bcbs/basel3/basel3-phase.in_arrangements.pdf

3.4 Simulation model




Figure 3-3: Bipartite network of the banks β and sovereign debts s .

While the EBA data details the origin of sovereign debts in a bank’s portfolio, it does not identify the origin of other assets, that is, financial institutions, corporate, retail residential mortgage, retail revolving, retail SME, and commercial real estate. We make the assumption that the distribution of sovereign debts of a given bank is a good proxy for the entire portfolio of that bank. In other words, if half of a given bank’s sovereign debt exposure comes from German bonds, we assume that also half of its exposure to financial institutions, corporate, etc., comes from Germany. The static balance sheet and distribution of bank portfolio assumptions are based on the bank balance sheet characteristics, where the credit exposures that banks have on one hand (the asset side) and the financing, i.e. short and long term debt on the other (the liability side) of the balance sheet are usually quite stable over extended periods of time. Bank leverage (or the equity level of the bank) is also quite invariant and it is regulated by international regulatory frameworks such as Basel III.

This allows us to construct a complete bipartite network, in which banks and assets interact back and forth, as illustrated for the case of sovereign debts in Figure 3-3. Let us first define variables on the bank layer and the asset layer, and then the risk propagation procedure.

3.4.1 Modelling a bank's Exposure

A bank $\beta \in [1, 2, 3, \dots, N_B (= 90)]$ is invested into different asset subcategories $a \in [1, \dots, N_A (= 7)]$ from different countries $s \in [1, \dots, N_{SD} = 32]$. We define the following financial variables for bank β :

\mathbf{A}_β $N_A \times N_{SD}$ matrix detailing the exposures to different asset classes from different countries:

$$\mathbf{A}_\beta = \underbrace{\begin{pmatrix} S_{\beta,1} & \cdots & S_{\beta,N_{SD}} \\ A_{\beta,2,1} & A_{\beta,2,2} & \cdots \\ \vdots & \ddots & \vdots \\ A_{\beta,N_A,1} & \cdots & A_{\beta,N_A,N_{SD}} \end{pmatrix}}_{\text{Countries}} \left. \vphantom{\begin{pmatrix} S_{\beta,1} \\ A_{\beta,2,1} \\ \vdots \\ A_{\beta,N_A,1} \end{pmatrix}} \right\} \text{Asset Classes}$$

$S_{\beta,s}$ The first row of the matrix \mathbf{A}_β , contain the exposure to all sovereign debt SD held by this bank. An individual sovereign debt s is therefore $S_{\beta,s} = (\mathbf{A}_\beta)_{1,s}$.

C_β tier 1 capital.

W_β Sum of risk-weighted assets (RWA).

R_β tier 1 capital ratio, $R_\beta = C_\beta/W_\beta$.

The RWA of a bank β in the data are calculated according to Basel III. We instead make use of a weighted sum as our definition of RWA:

$$W_\beta = W_\beta^{\text{SD}} + W_\beta^{\text{other}} = \sum_{s=1}^{N_{SD}} r_s S_{\beta,s} + \sum_{s=1}^{N_{SD}} \sum_{a=2}^{N_A} w_{a,s} (\mathbf{A}_\beta)_{a,s}, \quad (3.2)$$

where r_s denotes the risk factor (weight) of the SD s and $w_{a,s}$ for $a \geq 2$ are the

risk factors of other assets from different countries. Cash is an important part of any bank's assets. It, however, presents no risk, and is therefore absent in this summation. The weights for the RWA are estimated using an optimization procedure against the reported aggregated holdings of risk-weighted assets. While debt with longer maturity poses less immediate risk, in this study we use a simple sum of all holdings, regardless of their maturity.

3.4.2 Modeling sovereign debts and other assets

The EBA data details the country of origin of the sovereign debt. We assign the following variables for sovereign debt s , where s is again the indicator for the country of origin:

- S_s Total exposure, owned by all the banks in our database, $S_s = \sum_{\beta=1}^{N_B} S_{\beta,s}$.
- r_s Risk factor, $r_s = w_{1,s}$. Before shock, $r_1 = r_2 = \dots = r$
- D_s Yearly average of the CDS spread.
- $Q(D_s)$ Spreading parameter; $Q(D_s) = 1 - 2^{-D_s/100}$, so that $Q(0) = 0$, $Q(100) = 0.5$, as shown in Figure 3.4.

We use the following variables for assets from the other subcategories:

- $A_{a,s}$ Total exposure to all assets from asset category a from country s , owned by the all banks in our database, $A_{a,s} = \sum_{\beta=1}^{N_B} (\mathbf{A}_\beta)_{a,s}$.
- $w_{a,s}$ Risk factor. Before shock, $w_{a,1} = w_{a,2} = \dots = w_a$

Q_a Spreading Parameter for a given asset class; constant across all countries.

As pointed out before, the data regarding the exposure to other asset classes than sovereign debt come in aggregate form. In other words, while we know $(\mathbf{A}_\beta)_a$, we infer $(\mathbf{A}_\beta)_{a,s}$ from the sovereign debt holdings.




Figure 3-4: Spreading parameter $Q(D_s)$ where D_s is the CDS spread in basis points for a given sovereign debt s . We choose a the transformation $Q(D_s) = 1 - 2^{-D_s/100}$ such that its value is between 0 and 1 and that a spread of 100 basis points corresponds to a value of 0.5. A CDS spread of 0 corresponds to a Q value of 0, and as the CDS spread grows, Q approaches 1. Note the similarity of the function $Q(D_s)$ with the default probability derived from a CDS spread for some maturity and recovery rate as outlined in Chan-Lau [2006].

3.4.3 Iteration model

We run simulations for different crisis scenarios. One, we shock the sovereign debts of a given country by increasing their risk factors. The only interaction occurs between banks and sovereign debts; the other asset categories are unaffected. Two, we shock any sector of a given country by increasing its risk factor, and we include the

a	Item	w_a range	A_a	w_a	$w_a A_a$
1	Sovereign Debt	[0.002, 0.1]	27,267	0.002	55
2	Financial institutions	[0.5, 1.0]	25,044	0.5	12522
3	Corporate	[0.5, 1.3]	61,237	0.5	42866
4	Retail: Residential Mortgages	[0.5, 0.8]	36,663	0.5	14665
5	Revolving	[0.8, 1.2]	23,153	0.8	18522
6	SME	[1.0, 1.3]	3,467	1.0	3467
7	Commercial real estate	[1, 2]	22,228	1.0	22228
Total RWA W, according to Eq. (3.2)					114325

Table 3.1: Weighted-Average model for the risk-weighted Asset. The values for the exposure to different asset classes are for bank No.1, AT001. We present an example calculation for risk weights at the lower end of the range determined by our optimization procedure.

possibility of spillover to other asset classes. The sector suffering the initial shock can be any of the seven asset categories, from sovereign debt to commercial real estate. Three, we shock the banks in a given country by decreasing their capital by a certain percentage.

Shocking sovereign debts

We simulate a crisis scenario which starts with sovereign debts. We consider ten different shock origins:

- (1) GIIPS (Greece, Italy, Ireland, Portugal, Spain)
- (2) Eastern Europe (Bulgaria, Czech Republic, Estonia, Hungary, Lithuania, Latvia, Poland, Romania, Slovakia, Slovenia)
- (3) Benelux (Belgium, Netherlands, Luxembourg)
- (4) Greece
- (5) Italy

- (6) France
- (7) Germany
- (8) Spain
- (9) US
- (10) Japan

The simulation process is as follows: Upon initialization, an SD or a set of SDs from the countries of origin above has their risk factor increased at $t = 0$.

1. At $t = 1$, these shocks on SDs are propagated to the banks who own those SDs. As a result of the increased risk factors, the RWA of the affected banks is increased and their tier 1 capital ratio decreased.
2. At $t = 2$, those shocks on the banks are propagated back to SDs. A bank that is affected by the shock in the previous step might be forced to take action, putting pressure on the SDs in its portfolio. The risk factors of all SDs owned by affected banks are therefore increased, depending on the decrease in the tier 1 capital ratio of those banks.
3. At $t = 3$, shocks on SDs are propagated to banks, in a manner much like the step $t = 1$. The propagation continues back and forth between SDs and banks for $t = 4, 5, \dots$

The capital C_β of bank β remains constant throughout the simulation run. Therefore the change tier 1 capital ratio is entirely determined by the change in the RWA.

As described above, we model the risk propagation from the banks to the SD s (at $t = 2, 4, 6, \dots$) through the increase in risk factors. Each SD s contributes to the risk-weighted assets W_β^{SD} of bank β , and thus each bank holding an affected SD will

see an increase in their RWA proportional to its exposure to this SD. The exposures from other sectors also contribute to the risk-weighted assets W_β^{other} of each bank. These, however, will remain unchanged:

$$r_s(t+1) = r_s(t) / \Omega_s(t), \quad (3.3)$$

$$w_{a,s}(t+1) = w_{a,s}(t), \quad (3.4)$$

where

$$\Omega_s(t) = 1 - Q(D_s) \left(1 - \frac{\sum_{\beta=1}^{N_B} S_{\beta,s} P\left(\frac{R_\beta(t)}{R_\beta(t-1)}\right)}{\sum_{\beta=1}^{N_B} S_{\beta,s}} \right) \quad (3.5)$$

with no changes on the bank side, that is, in C_β . The value for the risk weights is capped at $r_{\max} = 2$.

$P(x)$ is a function which allows us to model the bank response to distress. The argument x is the fraction of capital tier 1 ratio left after a shock with respect to the tier 1 capital ratio in the previous time step; accordingly $1 - x$ is the relative loss in one time step. Our analysis considers two broad cases: banks that are risk averse and react to a deterioration of their tier 1 capital ratio accordingly on one hand, and banks that are risk neutral on the other.

Risk attitudes correspond to different curvatures of the utility function. For pay-offs, a concave utility function describes to risk aversion (Kahneman and Tversky [1979], Shefrin and Statman [1985]); for losses, the utility function of a risk averse agent therefore is convex. Functions of the form $\text{Max}[1 - (1 - x)^\alpha, 0]$ are convex, linear, or concave on the interval 0 to 1 depending on the choice of α . For simplicity, we consider a linear approximation,

$$P_\alpha(x) = \text{Max}[1 - \alpha(1 - x), 0] \quad (3.6)$$

for $0 \leq x \leq 1$ as the possible bank response functions to a decrease in tier 1 capital ratio. In this case Eq. (3.5) simplifies to

$$\Omega_s(t) = 1 - \alpha Q(D_s) \left(1 - \frac{\sum_{\beta=1}^{N_B} S_{\beta,s} \frac{R_{\beta}(t)}{R_{\beta}(t-1)}}{\sum_{\beta=1}^{N_B} S_{\beta,s}} \right) \quad (3.7)$$

In our analysis we focus on the cases $\alpha = 1$ which we call the linear case (risk neutral attitude) and $\alpha = 2$ which we call the steep case (risk averse attitude). The linear case implies that a bank will find a 20% drop in tier 1 capital ratio twice as bad as a 10% drop in tier 1 capital ratio. In the steep case, losses would be perceived much more negatively, modeling a larger sensitivity to losses. Therefore the steep function causes a more drastic effect given the same decrease in tier 1 capital ratio. Alternatively, the impact of the parameter α can be interpreted as follows: α determines for which value x the function $P(x)$ reaches the minimum value of zero, the worst case scenario. The larger α , the smaller the decline in tier 1 capital ratio that corresponds to the worst case scenario.

Equations (3.5) and (3.7) incorporate the CDS spread in basis points D_s . As a bank suffers from an increase in their RWA and a decrease in tier 1 capital ratio, it is reasonable to assume that this will affect debts of different quality differently. The higher the spread is for a sovereign debt, the more its risk weight will increase for the next simulation step. If the SD has a low spreading parameter, it will not be very affected by the banks' financial condition. With this, we model that Greek debt, for example, will deteriorate more given additional market stress than German debt. We refer to $Q(D_s)$ as the spreading parameter.




Figure 3-5: The reduction factor $P(x)$ in Eq. (3.5). The dotted curve shows the linear case, and the solid curve shows the steep case.

The change in risk factors leads in Eq. (3.4) changes the RWA as follows:

$$W_{\beta}^{\text{SD}}(t+1) = \sum_{s=1}^{N_{\text{SD}}} S_{\beta,s} r_s(t) / \Omega_s(t), \quad (3.8)$$

$$W_{\beta}^{\text{other}}(t+1) = W_{\beta}^{\text{other}}(t). \quad (3.9)$$

Shocking other asset categories

In this crisis scenario, an asset category is being shocked, and there will be spill-over to other asset categories. The simulation process can then be described as follows. We initialize an asset categories in a country or a set of countries to have its risk factor increased at $t = 0$. We consider the same 10 scenarios as in the previous case.

1. At $t = 1$, these shocks on assets are propagated to the banks who own those assets. As a result of the increased risk factors, the RWA of the affected banks is increased and their tier 1 capital ratio decreased.
2. At $t = 2$, those shocks on the banks are propagated back to assets. A bank that is affected by the shock in the previous step might be forced to take action, putting pressure on all their assets in its portfolio. We emulate this by increasing

the risk factors of the assets that the affected bank holds, depending on how many of these assets the bank holds and how much its tier 1 capital ratio has decreased. Unlike in the previous scenario where the increase in risk factor was limited to sovereign debts, that is, the asset class in which the shock started, in this scenario the risk weights of all assets of affected banks will increase.

3. At $t = 3$, the increased risk factors across all asset classes lead to a new, higher value for the risk-weighted assets of the banks. Just like in step $t = 1$, this causes a further decrease in tier 1 capital ratio. The propagation continues back and forth until the system saturates.

The capital C_β of bank β remains constant throughout the simulation run.

As pointed out before, it is likely that a stress to a bank's portfolio originating from some of its exposures will lead to adjustments across all asset classes. In that case the shock propagation will not only change the risk factors for sovereign debts, but it will also affect the risk factors for the remaining asset classes. Again, we allow a maximum value of 2 for the risk weights. Extending the shock propagation, we rewrite Eq. (3.4) as follows:

$$\begin{aligned} r_s(t+1) &= r_s(t) / \Omega_s(t), \\ w_{a,s}(t+1) &= w_{a,s}(t) / \Omega_a(t), \end{aligned} \tag{3.10}$$

where $\Omega_a(t)$ is extended to the case of non-SDs:

$$\Omega_a(t) = 1 - \alpha Q_a \left(1 - \frac{\sum_{\beta=1}^{N_B} A_{\beta,a} \frac{R_\beta(t)}{R_\beta(t-1)}}{\sum_{\beta=1}^{N_B} A_{\beta,a}} \right). \tag{3.11}$$

Like its counterpart for sovereign debts, the shock parameter will depend on the

other asset classes through the asset-specific holdings of a bank $A_{\beta,a}$ and the equivalent of a spreading parameter, Q_a . Recall that we have chosen to set Q_a to be the same for all countries.

In the same way as $Q(D_s)$ models the susceptibility of sovereign debt s to deteriorating market conditions, Q_a describes how strongly asset class a is affected by a reduction in tier 1 Capital of banks exposed to it. If we set $Q_a = 0$, we recover Eq. (3.4) and we prohibit any spillover into other asset classes. For any other value $0 < Q_a \leq 1$, banks affected by an initial shock in the SD sector will cause an increase in risk factor for other asset classes, proportionally to their exposures.

Shocking the capital of banks

Previously we have kept the capital C_β of all banks constant. However, we can well imagine a scenario in which a shock to a bank or to a couple of banks stems from a sudden drop in equity. Such a sudden drop would leave a bank potentially over-leveraged, that is, their tier 1 capital ratio becomes too small. Our model allows us to study the effect of such a shock as well. We do this as follows, considering the same 10 scenarios as in the previous case.

1. At $t = 1$, the capital of a bank or a group of banks is reduced by a certain percentage. As a result of that, their tier 1 capital ratio decreases.
2. At $t = 2$, those shocks on the banks are propagated to assets, like in the previous scenarios, the only difference being that all assets start out with their original risk weights. The risk weights of all assets that are held by affected banks will increase.
3. At $t = 3$, the increased risk factors across all asset classes lead to a new, higher value for the risk-weighted assets for all banks holding them. In this step,

the crisis originating in a subset of banks spreads to other banks. Again this propagation continues back and forth until the system saturates.

The iteration process follows the same dynamics as in the scenarios in which we shocked sovereign debt and other asset classes, using Eqs. (3.10) and (3.11).

3.5 Simulation results

Using our model, we obtain simulation results for the three different scenarios described in the previous section, (a) shocking sovereign debt, (b) shocking any asset category allowing for spillover, and (c) shocking bank capital. We characterize the initial shock, market conditions, and the behavior of market participants by two parameters and a response function. As the *shock origin* we consider the ten scenarios as described in 3.4.3. The *shocked sector* refers to the part of the financial network in which the shock starts; this may be an asset class or the equity of a bank. When shocking assets, we increase their risk weights, increasing the risk-weighted assets and thus reducing banks' tier 1 capital ratio. The *shock size* ranges from increasing the risk weights by a factor of one-and-a-half to four³. We perturb the banks by simulating a sudden decrease of their equity and thus reducing the tier 1 capital ratio; this *shock size* ranges from a reduction by 10 to 90 percent. The *spreading parameter* describes the market conditions for sovereign debt Q_s and for other asset classes Q_a . The spreading parameter Q_s for the sovereign debt is a function of the CDS spread as shown in Figure 3-4. The spreading parameters for the other asset classes is a free parameter where for simplicity we set $Q_1 = Q_2 = \dots = Q_a$ for all asset classes. The *bank response function* $P(x)$ describes the behavior of market participants. We distinguish between the linear and the steep bank response function $P(x)$ as put forth in Eq. (3.6).

³ Recall that we limit the magnitude of a risk weight to $w_{\max} = 2$, which shall not be exceeded through the initial shock.

3.5.1 Shocking sovereign debt only

The first scenario in our simulation considers a shock to sovereign debt without any spillover to the other asset classes. Sovereign debt usually makes up a small fraction of a bank's assets, with exposure of most banks of about one tenth of the total as shown in 3.1 in the data section.

If we use risk weights (r_s) for sovereign debt at the lower end of the range described in Table 3.1, the impact of sovereign debt on a bank's risk-weighted assets becomes all but negligible. Therefore, even a major shock in any of the ten different shock origins will not have any significant effect on banks' risk-weighted assets for initial risk weights $r_s = 0.002$. In turn, the tier 1 capital ratio decrease is so minor that any shock propagation is stalled right away. This finding holds for both the linear and the steep bank response function. Therefore, the value we initially assume for the risk weights r_s becomes crucial.

If we use the risk weights at the upper end of the range in Table 3.1, $r_s = 0.1$, a major shock such as a fourfold increase in those risk weights for an affected country will have a measurable impact on banks holding a large chunk of this sovereign debt. However, the shock propagation stalls after at most a couple of steps for both the linear and the steep bank response function. If we consider a shock originating in a country or region with very low sovereign debt CDS spread, such as Germany or the Benelux countries, the increase in risk weights diminishes already in the step right after the initial shock. Considering the GIIPS countries, which exhibit the largest CDS spreads in our data set, the shock propagates for two steps; however it does not noticeably impact banks any more than the original shock.

3.5.2 Allowing for spillover between asset classes

We have observed that shocks to sovereign debt alone do not become systemic events within the framework of our model. This is mainly due to three reasons:

- (i) Many banks hold only a low amount of sovereign debt.
- (ii) The risks weights of sovereign debt are very low, signifying that it is considered practically riskless. The European sovereign debt crisis at least raises doubts regarding the risk-free nature of sovereign debt.
- (iii) We do not consider a spillover between asset classes, that is, a shock to sovereign debt will remain contained within the sovereign debt sector.

We therefore proceed to analyze the results when we include other asset classes with their larger risk weights as the source of shocks as well as when we allow for a spillover from one asset class to another, according to Eqs. (3.5) and (3.11).

We take 10 of the 90 banks as a sample to illustrate the dynamics and the outcome of a shock from various origins. The banks are Erste Bank Group (AT001), Deutsche Bank (DE017), ES065 (Banco de Sabadell), Credit Agricole (FR014), Barclays (GB090), National Bank of Greece (GR031), Bank of Ireland (IE038), Unicredit (IT041), ING Bank (NL047), and Banco Comercial Portugues (PT054). This selection aims to represent a good cross section of all European banks and to demonstrate one or more banks very vulnerable to the scenarios we consider. For example, we can expect Deutsche Bank to be initially more strongly effected by a crisis that originates in Germany, scenario (7).

The first scenario we consider is a shock which originates in the GIIPS countries, Greece, Italy, Ireland, Portugal and Spain. Figures 3-6 through 3-8 show the dynamics for the banks given an increase of the risk weights associated with residential mortgages from these countries. For most banks, retail residential mortgages are the most




Figure 3-6: Tier 1 capital ratios over time after a shock to the retail residential sector in GIIPS countries, given different bank response functions and both a small shock size and a small spreading parameter.

important or second most important asset class by exposure. We present results for different shock sizes, spreading parameters, and bank response function. If we consider a spreading parameter of $Q_a = 0.3$, the system is stable for both a linear and a steep bank response function, as Figure 3-6 indicates. As a matter of fact, a fourfold increase of the risk weights does not cause a systemic event either, see Figure 3-7. The sudden large increase in risk weights, however, causes a tremendous decline in the tier 1 capital ratio for banks heavily involved in mortgages in this region. Banks from other than GIIPS countries see an obvious downtick in their tier 1 capital ratios as well. As a consequence, a couple of banks fall below the Basel III threshold.




Figure 3-7: Tier 1 capital ratios over time after a shock to the retail residential sector in GIIPS countries, given different bank response functions and a large size and a small spreading parameter.

If we increase the spreading parameter to $Q_a = 0.6$ while still considering a linear bank response function, the system remains in the stable regime, and the initial shock does not cause a spillover. If we consider a spreading parameter of $Q_a = 0.6$ and assume a steep bank response function, however, the system exhibits instability. This can be seen in Figure 3-8, and it holds true for any shock size. While an increase of risk weights increase of $w_a \rightarrow 4w_a$ could not cause a systemic event for a lower spreading parameter or a linear bank response function, even the smallest increase of w_a will trigger a systemic event if the spreading parameter is sufficiently large and we use the steep bank response function. As an example, Figure 3-8 shows how $w_a \rightarrow 1.5w_a$ causes all banks to eventually fall below the Basel III threshold in our simulations.




Figure 3-8: Tier 1 capital ratios over time after a shock to the retail residential sector in GIIPS countries, given different bank response functions and a small shock size and a larger spreading parameter.

Figure 3-8 has demonstrated that even small shocks may eventually cause a systemic event, given an adverse set of parameters. While we concluded that sovereign debt in isolation cannot trigger a systemic event, we revisit this problem, now allowing for spillover to other asset classes. Specifically, we focus on the conditions that have facilitated a systemic event after an initial shock to residential mortgages. Figure 3-9 shows the dynamics for the banks given an increase of risk weights for sovereign debt from GIIPS countries from 0.002 to 0.008. We use a spreading parameter of $Q_a = 0.6$

and the steep bank response function. While the slope of the tier 1 capital ratio decline is much less than in Figure 3-8 and therefore the development of the systemic event is considerably slower, the dynamics become equivalent. Indeed, the final tier 1 capital values for the banks are very similar for both cases, not shown in the figure.




Figure 3-9: Tier 1 capital ratios over time after a shock to sovereign debt in GIIPS countries, given different bank response functions and spreading parameters.

Our observations for these two cases hold for a shock origin in other asset classes from the GIIPS countries as well. In absence of a systemic event, the final tier 1 capital ratio of a bank can be well approximated by the impact of the initial shock on its tier 1 capital ratio. In other words, the exposure to a specific asset class from a specific country determines the magnitude of tier 1 capital losses. If we observe a systemic event, however, banks start deteriorating uniformly as the initial shock has spread to all asset classes. In this case, the initial exposure to a specific asset class from a specific country influences only the initial rate of decline, but it does not affect the outcome.

We study how many banks fall below the Basel III threshold after 50 time steps and after 100 time steps. Table B.1 in Appendix B reports how many banks would be distressed in the ten scenarios we have proposed. Since only the steep bank response function and a high enough spreading parameter can trigger a systemic event, we consider shocks to all seven sectors given $P_{\text{quad}}(x)$ and $Q_a = 0.6$. As we can see by

comparing Figs. 3·8 and 3·9, shocks to different sectors and with different shock sizes permeate the system at different speeds.

In many cases, two thirds or more of the 90 banks can be considered distressed in our simulation after 100 time steps. This is true regardless of shock origin and the asset class which is initially shocked. The largest average distress across all asset classes can be found when shocking the GIIPS countries; 65.6 percent of all banks have fallen below the Basel III threshold after 100 time steps. Even after only 50 time steps, shocks to assets from the GIIPS country have large destructive capacity, distressing nearly half of all banks. A shock to assets from Spain leads to the second largest number of distressed banks across all sectors, with 64.9 percent of all banks having a tier 1 capital ratio of less than 4.5% after 100 time steps. As compared to a shock to the GIIPS countries of which Spain is part, a shock to assets from Spain tends to spread more slowly, leaving only about a quarter of banks distressed after 50 time steps. We observe that, while most shocks have large negative impacts, a shock to corporate loans damages the system very fast. In contrast, shocks to retail SME loans or commercial real estate tend to spread out more slowly, allowing for more time to intervene. We report these results in more detail in the appendix in Tables B.2 through B.9.

3.5.3 Shocking the equity of banks

So far we have only shocked the risk-weighted assets of the banks. Subsequently, we simulate the reduction of equity in a subset of banks. Since in this study we analyze the European banking system, we focus on scenarios (1) through (8). We report both the average loss in tier 1 capital ratio as well as the number of banks which fall below the Basel III threshold of 4.5 percent. The results are summarized in Appendix B in Tables B.10 through B.17.

It comes as no surprise that a shock to banks in the GIIPS countries would have

the largest effect on the entire system. Almost half of the banks in our data set are based in one of these five countries, and therefore the size of the initial shock would be quite significant. Furthermore, we recognize that a linear bank response function effectively stops spillover of the crisis. For example, in scenario (i), even given the largest shock size of 50 percent and a large spreading parameter of $Q_a = 0.9$, only five banks from other countries would fall below the Basel III requirements. The importance of the spreading parameter is again amplified in case of the steep response function: If the spreading parameter is large and we assume a steep bank response function, the initial shock size matters very little. Almost all banks would fall below the Basel III tier 1 capital ratio threshold given these parameters, regardless of where the shock originated.




Figure 3-10: Evolution of the tier 1 capital ratio of Deutsche Bank (DE017) for a shock to various sectors in Germany, given a linear response function and a steep response function.

In order to compare the dynamics of a shock to a bank's equity, that is, a direct shock, to the dynamics after an indirect shock through an asset class, we consider in more detail the outcomes for Deutsche Bank (DE017) as an example. Figure 3-10 illustrates the development of the tier 1 capital ratio of Deutsche Bank for a shock to

asset classes from Germany as well as the equity of all German banks. We initialize the shock by an increase of risk weights of the shocked asset class by 50 percent or decreasing the equity of German banks by 20 percent, respectively. We confirm that the linear response triggers now systemic event and we focus on the results for the steep response function. A shock to the equity of all German banks has both the largest initial impact on Deutsche Bank and leads to the lowest tier 1 capital ratio at the end of the simulation. In fact, the final tier 1 capital ratio is about one fifth lower than for shock to other asset classes. A shock to corporate loans or financial loans from Germany causes the fastest deterioration of the tier 1 capital ratio of Deutsche Bank, as this bank is strongly invested in these two asset classes. Shocks to other asset classes lead to slightly different dynamics: While the tier 1 capital ratio in the cases of initial shocks to corporate and financial loans is practically linear, a shock to other asset classes starts slowly, with a low downward slope, but picks up pace as time goes by. Apparently, an initial shock in these asset classes builds up slowly and it is not until other, more important asset classes to the bank become affected that the propagation causes bigger losses. Interestingly, for this reason, the shock that originates in the corporate loan sector and causes a tier 1 capital ratio decline the fastest leaves the bank at a slightly higher tier 1 capital ratio at the end of the simulation. The shock spreads at such a high rate that the maximum risk weight of $w_a = 2$ is achieved very quickly for this asset class, stalling the propagation before “infecting” other asset classes. It needs to be pointed out, though, that this difference is minor, and in all cases Deutsche Bank falls below the Basel III threshold in our simulations.

In Figure 3·10, we only consider shocks that originate in Germany, affecting German assets and German banks’ equity initially. Figure 3·11 focusses on the asset class of corporate loans, which we have identified as the shock affecting the bank the fastest,

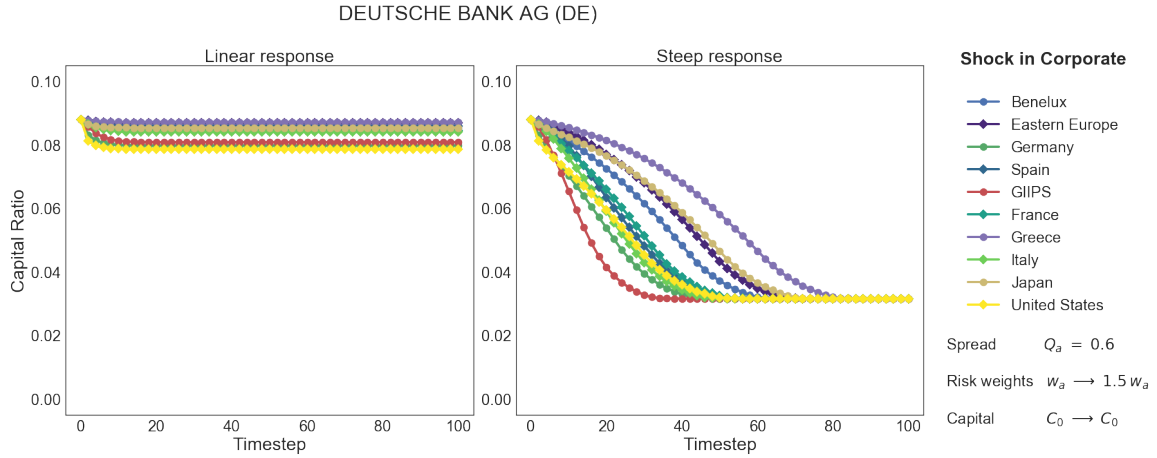


Figure 3-11: Evolution of the tier 1 capital ratio of Deutsche Bank (DE017) for a shock to corporate loans in various countries or regions, given a linear response function and a steep response function.

and we differentiate now by origin of shock. The shock size and spreading parameter are the same as in the previous discussion. Since we have used the sovereign debt holdings of Deutsche Bank to approximate its holdings in other classes, Figure 3-11 is an approximation to how affected the bank will be due to its holdings as well as the potential of a shock to spread throughout the market. Again, we only observe a decrease in tier 1 capital ratio through the initial shock using the linear bank response function, and we observe a systemic event using the steep bank response function. As we have seen before, a shock to assets from GIIPS countries leads to the fastest decline in tier 1 capital ratio for Deutsche Bank, which reflects the size of the original shock, followed by a shock to German assets which reflects the domestic bias of Deutsche Bank. Interestingly, a shock to US exhibits a very quick decline in tier 1 capital ratio as well. This is not too surprising however, because we would expect an global player like Deutsche Bank to be very active on the US market. It takes the longest time for shocks from Greece (due to the comparably low amount of Greek assets on the market), Japan and Eastern Europe (due to the limited amount of holdings in these

regions) to put Deutsche Bank below the Basel III threshold. However, as before the initial shocks all converge to roughly the same final value of tier 1 capital ratio for Deutsche Bank.

3.5.4 Phase space analysis

In our exploration of specific parameters, we have recognized that there appears to be a phase transition for the steep bank response function, depending on the value for the spreading parameter. In the following we expand our analysis to a wide range of parameters for both, a linear and a steep response function. We use a shock originating in the Benelux countries as an example to discuss the outcome for the banks in our sample given. We choose the Benelux countries because unlike a shock to the GIIPS countries which involves a very large amount of assets and banks or to Eastern Europe, for example, which involves much fewer assets and banks, this scenario corresponds to a mid-sized shock with respect to assets and number of initially affected banks. We want to point out that the following analysis and its findings are applicable to shocks from different origins.

Figure 3.12 shows the phase diagrams for a shock originating in the Benelux countries, using a linear bank response function. Each of the subplots captures a different crisis origin and explores the outcome for a wide range of parameters of shock size and spreading. We consider a crisis onset due to a sudden loss in tier 1 equity for all banks in the Benelux countries as well as crisis onsets due to a sudden increase in risk weights for assets from a given asset class in Portugal, Italy, Ireland, Greece, and Spain. We allow the shock between banks and assets to propagate back and forth for 200 time steps, that is, 100 propagations from assets to banks and 100 propagations back from banks to assets. On the y-axis we show the average decrease in tier 1 capital ratio of all banks. On the x-axis we show the sector spread, that is, the value of the parameter Q_a , which indicates how affected an asset from a given

asset class is when banks holding it experience distress. For the presentation of these results, we set $Q_2 = \dots = Q_a$ for all asset classes except the sovereign debt where we use the idiosyncratic values from the CDS swaps. Indicated by colors is the size of the original shock. We consider the following range of shock sizes: In the case of banks, we reduce the capital by 10 to 90 percent. In the case of the other assets, the shock is initiated by an increase of risk weights by 50 percent up to a fourfold increase.

Our main observation in Figure 3.12 is that, as expected, the impact on banks grows with the size of the spillover Q between sectors grows as well as with the size of the original shock. A shock scenario which begins with a sudden reduction in capital for banks in our dataset appears to have the largest negative effect on the system. Let us first consider the scenario in which all banks in the Benelux countries were to suffer from an abrupt loss of 10 percent of their equity and therefore suffer from a reduction of their tier 1 capital ratio by 10 percent. The systemic consequences then depend on the size of the spreading parameter, Q_a . In the case of weak spreading, $Q_a = 0.3$, banks from Benelux countries see a total reduction of their tier 1 capital ratio R of 11.0 percent, barely more than through the initial shock alone. Banks from all other countries will see a reduction in their R of 0.3 percent. This comes out to an average reduction in R of 1.1 percent, reflecting that there are far fewer banks in the Benelux countries than in the rest of Europe. In the case of strong spreading, $Q_a = 0.9$, the reduction in tier 1 capital ratio for banks from Benelux countries is more drastic, amounting to a total of 16.5 percent. We also observe a noticeable effect on banks from other European countries which see their R reduced by 4.1 percent. With larger initial shock size, these numbers become bigger, but not systemically. This is similar to our previous discussion in Section 3.5.3 and reflected in Table B.12.

The results are very different, as we have discussed before, if we consider a steep bank response function. We recognize two distinct regimes of the system in Fig-

Beginning crisis in Benelux
(Linear response)




Figure 3-12: Phase space for different scenarios given a linear bank response function: We consider a variety of initial shock sizes to banks and asset classes as well as a variety of values for the spreading parameter. We observe no amplification of shocks. The size of the initial shock is indicated by the color, where from yellow (small shock) to dark purple (very large shock).

Beginning crisis in Benelux
(Steep response)

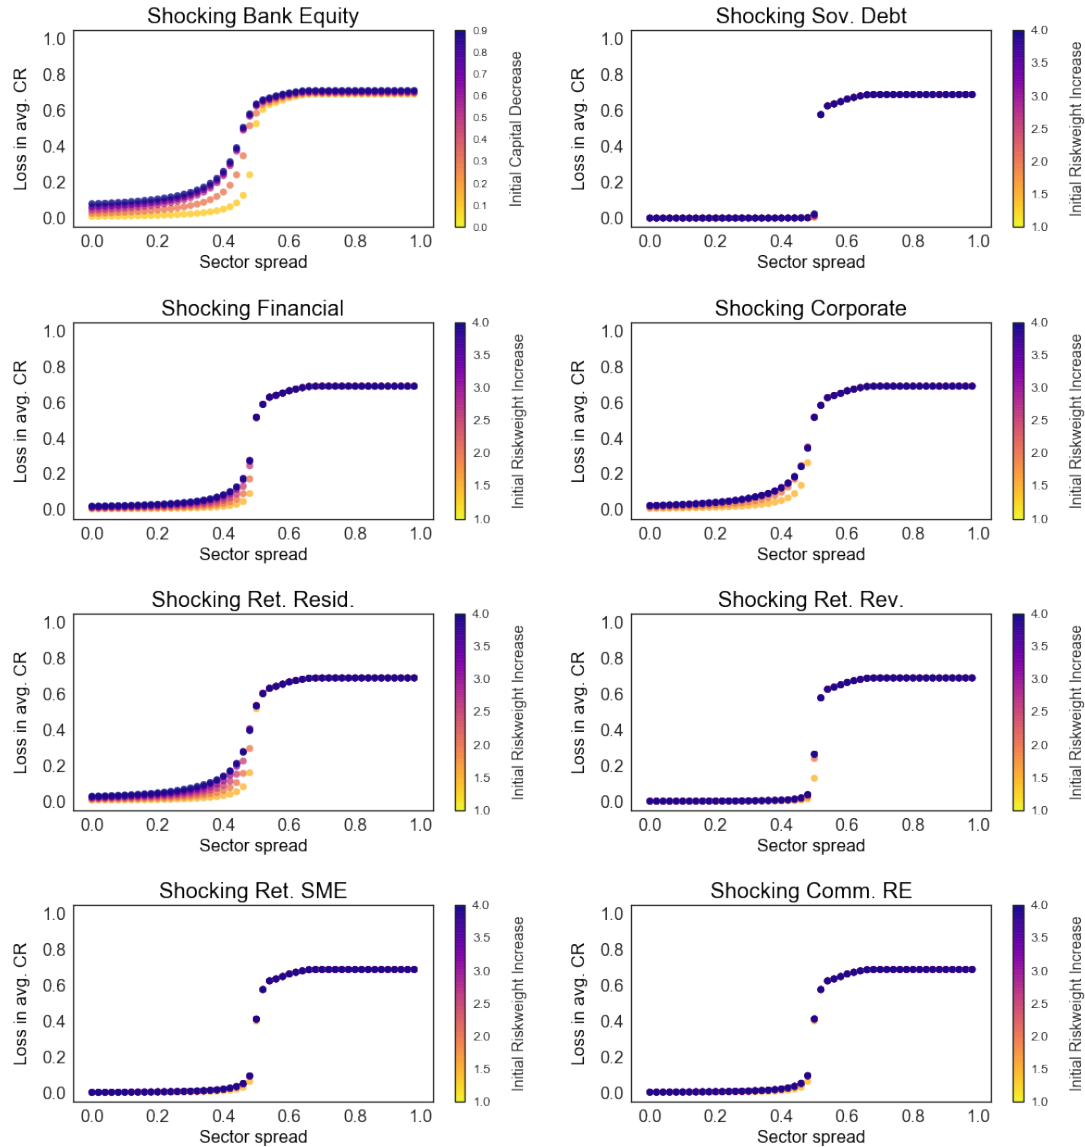


Figure 3.13: Phase space for different scenarios given a steep bank response function: We consider a variety of initial shock sizes to banks and asset classes as well as a variety of values for the spreading parameter. The size of the initial shock is indicated by the color, where from yellow (small shock) to dark purple (very large shock). We observe behavior reminiscent of a second order phase transition, where the critical parameter is the magnitude of the sector spread Q_a . The tipping point of the system appears to be around $Q_a = 0.55$.

ure 3-13, depending on the size of the spreading parameter. If $Q_a < 0.5$, then the loss in the tier 1 capital ratios of banks grows, but it grows slowly and steadily with increasing Q_a and with increasing initial shock size, very much in the same fashion as it has in Figure 3-12 for the linear bank response function. However, if $Q_a > 0.6$, banks lose, on average, about 75% of their tier 1 capital ratio. In other words, all banks with less than 18 percent tier 1 capital ratio at the onset of the simulation are in acute danger of failing the stress test in this scenario. Remarkably, this behavior is true for all asset classes and bank equity, with the exception of sovereign debt which we will discuss later. Furthermore, regardless of the asset class in which the shock originated and regardless of the size of the initial shock size, the final tier 1 capital ratio of banks will be the same.

We speculate that the system exhibits a phase transition around $Q_a = 0.55$, and as a consequence the banking network can be extremely fragile. Below this parameter, an initial shock can be quite damaging to banks, however it diminishes over time. Above this parameter, an initial shock does not diminish as quickly as it spreads to other sectors of the system. Instead it has the potency to affect all asset classes and thus all banks.

This sudden switching behavior can be better understood when considering the slope at the transition as well as the dependence on the initial shock size leading up to the transition point. We observe that the transition is smoothest for a shock to bank equity, followed by a shock to corporate loans and loans to financial institutions. Since corporate loans and loans to financial institutions tend to make up a very large part of the balance sheets of banks, especially larger banks, a shock to these asset classes is more impactful than a shock to other asset classes. Likewise, a shock to bank equity of 10 or 20 percent or more is a significant shock. In contrast, a shock to revolving retail or SME has only a minor initial impact on many banks. For the sake

of the argument, let's consider a fairly small shock to the risk weights of a given asset class. Obviously, no disruption to a specific asset class by itself by themselves cause a systemic event because banks are diversified across asset classes and countries, and therefore a small change to one risk weight will be a very small change to a bank's overall risk-weighted assets – and, in turn, to its tier 1 capital ratio. A systemic event gets underway if the initial shock can fester for long enough and spread to other countries and asset classes, increasing the associated risk weights. In the case of the bigger asset classes (corporate and financial) as well as equity, the initial shock will have had a measurable impact already, such that the transition looks less sharp. In the case of the smaller asset classes (ret. revolving and SME), however, the transition is very sudden: Below a certain spreading parameter, the shock remains contained and the possible damage to tier 1 capital ratios is bounded by the volume of these asset classes; but beyond that spreading parameter, the shock spreads to all parts of the system.

A notable exception is sovereign debt. Due the low risk weights and also the idiosyncratic spreading parameters derived from the CDS swaps, an initial shock within the scope of current risk weights is often just too small to cause a systemic event. Given a very large initial shock or very many time steps, a systemic event can, nevertheless, be triggered by a shock to sovereign debt. The value of Q_a for the phase transition, however, may be very different than that for the other asset classes. This again reflects the idiosyncratic spreading parameters.

Figure 3.14 shows the dependency of the outcome of the simulation with respect to the parameters Q_a and α which we had identified as critical parameters of the phase transition. We observe that there is a sharp boundary between two regimes, one corresponding to the resilient and one corresponding to the fragile phase. If we increase Q_a , the steepness parameter α necessary for a regime switch decreases




Figure 3-14: Stationary state that the system reaches in the simulation depending on the spreading parameter Q_a and the steepness parameter α . Green colors indicate a very small average deterioration of the tier 1 capital ratios of the banks in the system, whereas purple indicates the worst attainable stationary state. The bright line separating the two regimes indicates the point of the phase transition. The plot corresponds to a shock to loans in the French financial sector.

accordingly. Interpreting the steepness parameter as an indicator of risk aversion, this implies that if the market is fragile, small amounts of risk aversion among the banks may be sufficient to topple the system. Conversely, if Q_a , a measure related to the probability of default of the assets, is small, the banking system is particularly robust, even if banks are very risk averse. Figure 3-14 illustrates one specific shock scenario, that of a shock to financial institutions in France; other scenarios yield a similar result with a small shift to the regime boundary.

3.6 Discussion and conclusion

In this chapter we have analyzed the European Banking Authority stress test results. We have created a bipartite network with banks on one side and assets on the other, based on European banks' holdings of sovereign debt from 30 European nations, the United States, and Japan. We have proposed a systemic risk propagation in which

we take into consideration the interconnectivity between banks based on the overlap of their asset portfolios. We have analyzed the systemic impact of shocks to bank assets by increasing the risk weight of these assets as well as of shocks to the equity of banks. Both types of shocks lead to an initial decrease in the tier 1 capital ratio of the affected banks. In our propagation model a deterioration of the tier 1 capital ratio prompts a reaction of the banks, putting stress on assets in their portfolio and further enhancing the crisis. We have considered a linear response of banks to the shock and we confirm that banks are more affected by the initial shock than by subsequent spillovers. In accordance with Glasserman and Young [2015], we found no significant contagion, hence the effect of the shock is locally contained and can be explained by portfolio overlap and the size of the initial shock. If, however, the banks' response was described by a steeper function, the stability of the banking network becomes strongly depended on the spreading parameter. This spreading parameter has a critical value associated with a phase transition. Below the critical value, the system exhibits stability and is comparable to the linear case: We observe no spillover and losses depend on both the size of the initial shock and the origin of the shock. However, above the critical value the system breaks down, and we observe a very large deterioration of the tier 1 capital ratio of all banks; in fact most banks fall below the Basel III threshold. The critical value depends slightly on the origin of the shock and the asset class; however, the outcome for the banks beyond the critical value of the spreading parameter is then independent of the origin of the shock and of the size of the initial shock.

Our results show that even though the systemic risk propagation through the banking network is homogeneous once strong contagion is present, the trajectory of tier 1 capital ratio deterioration for a given bank depends on the origin of the shock, the size of the shock and its balance sheet. We suggest that our model is a

good complement to the current stress tests to capture the interconnectivity of banks due to their portfolio similarities, their business models, and their regional biases. Understanding these dynamics can be helpful to regulators as well as policy makers, and it may serve to inform the impact of interventions and the lack thereof in times of crisis.

Chapter 4

Reconstruction of the Japanese Bank-Firm Network

The study of systemic risk in complex networks requires both a good model and detailed knowledge of the network. Often, we cannot observe individual edges, and information is only available about the nodes on an aggregate level. If systemic risk is to be modeled as a shock propagation, an accurate estimate of the edges is imperative. Recently, a variety of network reconstruction methods has been put forth, of which the configuration model based on the principle of maximum entropy has emerged as a front runner. Inspired by reconstruction attempts on the Japanese firm-bank data set and their shortfall, we establish limits on the applicability of the configuration model for sparse networks.

4.1 Introduction

Knowing the network structure is a prerequisite to studying systemic risk. In Chapter 3 of this thesis, we introduced a model for systemic risk emanating from overlapping portfolios of financial institutions. The model takes into account the bank behavior as well as the liquidity. While the data set detailed the exposures of banks to sovereign debt, we had to reconstruct the portfolios in other asset categories from this partial information. Often, no such information is available due to confidentiality reasons or missing data. Typically, we can only rely on aggregate information. In the example of interbank lending, this means that we do not know the exposure individual banks have to one another; instead we know for each bank how much money they borrowed and lent in total¹. Much work has been devoted to reconstructing networks from such partial or marginal information with the goal of studying systemic risk, for example by Cimini et al. [2015] or Anand et al. [2015].

In a recent work, Ramadiah et al. [2017] performed a horse race of different reconstruction methods, including the enhanced capital asset-pricing (ECAPM) by Squartini et al. [2017], on the extensive Nikkei data set of Japanese banks and firms. Considering this credit network, the authors used a shock propagation model according to Huang et al. [2013] to test the network reconstruction methods on how well they describe the systemic risk. Surprisingly, Ramadiah et al. [2017] found that all network reconstruction methods underestimated the effect of a shock to the system. While the configuration model of Cimini et al. [2015] performed very well in a comparison of reconstruction methods studied by Anand et al. [2017], the extension for bipartite networks [Squartini et al., 2017] apparently was not successful in capturing critical parts of the Japanese bank-firm network.

In previous studies, the quality of network reconstructed was estimated through

¹ Of course, central banks or other regulatory bodies typically know the detailed exposures.

network topological and similarity measures, for example in the horse race in Anand et al. [2017]. We highlight another important feature to consider, the likelihood of connectivity. Due to their probabilistic nature, configuration models may fail to connect a subset of nodes in any given realization. We derive bounds for network properties like density and layer heterogeneity to guarantee connectivity at a given confidence level.

4.2 Methodology

4.2.1 Principle of maximum entropy

Given N possible events with probabilities $P(x_i), i = 1 \dots N$, the entropy H is defined as

$$H(X) = - \sum_{i=1}^N P(x_i) \log P(x_i), \quad (4.1)$$

with $x \log x = 0$ if $x = 0$. As originally defined by Shannon for a discrete random variable X with probability mass function $P(X)$, this is equivalent to the expected value

$$H(X) = \mathbb{E}[-\log P(X)]. \quad (4.2)$$

Generally the probability distribution is constrained to values F_k through some functions f_k :

$$\sum_{i=1}^N P(x_i) f_k(x_i) = F_k, \quad (4.3)$$

with $k = 1, \dots, m$. An obvious constraint for every probability distribution is

$$\sum_{i=1}^N P(x_i) = 1, \quad (4.4)$$

the normalization requirement. We can formulate the maximum entropy probability

distribution as

$$P(x_i) = \frac{\exp[\lambda_0 + \sum_{k=1}^m \lambda_k f_k(x_i)]}{\sum_{i=1}^N \exp[\lambda_0 + \sum_{k=1}^m \lambda_k f_k(x_i)]}, \quad (4.5)$$

where we recognize that the denominator is the partition function $Z(\lambda_0, \dots, \lambda_m)$. The parameter λ_0 is the Lagrange multiplier of the normalization, and $\lambda_1, \dots, \lambda_m$ are the Lagrange multipliers of the other m constraints, satisfying

$$F_k = \frac{\partial \log Z}{\partial \lambda_k}. \quad (4.6)$$

While we defined entropy in Eq. (4.1) for discrete random variables, the concept can easily be extended to continuous random variables by replacing the summation with integration. It can be shown that the exponential distribution and the normal distribution are the maximum entropy distributions given certain constraints. The exponential distribution is the maximum entropy distribution of the class of distributions over all positive real values with mean μ . The normal distribution is the maximum entropy distribution of the class of distributions over all real values with mean μ and variance σ^2 . That means that any other distribution over the respective domains assumes further knowledge over the data, like the value for higher moments.

4.2.2 Maximum entropy for networks

Following the lines of the seminal work of Park and Newman [2004] and Squartini and Garlaschelli [2011], we briefly explain how the principle of maximum entropy is applied to the reconstruction of complex networks. Assuming we know the degree sequence of a graph, that is, the degree of each node, denoted by k_i . Obviously, we can conceive many graphs with such a degree sequence, rendering the problem underdetermined. In their approach, Park and Newman [2004] consider a graph G within an ensemble \mathcal{G} of graphs. Each graph is assigned a probability $P(G)$ which

needs to be estimated. This is done by maximizing the entropy of the ensemble,

$$H = - \sum_{G \in \mathcal{G}} P(G) \log P(G), \quad (4.7)$$

subject to the constraints that the probability sums up to one and that the expected value of the degree sequence of the ensemble matches the observed degree sequence.

Using Eq. (4.5), we can write the result as

$$P(G) = \frac{e^{-\mathcal{H}(G)}}{Z}, \quad (4.8)$$

where $\mathcal{H}(G) = - \sum_i \lambda_i k_i(G)$ is the Hamiltonian with the Lagrange multipliers λ_i and degree sequence k_i for graph G , and Z is the partition function.

Let us consider the simple case of an unweighted, but directed network with the adjacency matrix \mathbf{A} with entries $a_{ij} = 1$ if i is linking to j , $a_{ji} = 1$ if j is linking to i , and $a_{ij} = 0$ if the nodes are not connected. We then find the in-degree of node i as $k_i^{in} = \sum_j a_{ji}$ and its out-degree $k_i^{out} = \sum_j a_{ij}$. Park and Newman [2004] proceeds to specify a coupling parameter Θ_{ij} for each edge, and the Hamiltonian becomes

$$\mathcal{H}(G) = - \sum_{i,j} \lambda_i a_{ij} = - \sum_{i<j} \Theta_{ij} a_{ij}, \quad (4.9)$$

and consequently the partition function is

$$Z = \sum_{\{a_{ij}\}} \exp \left(- \sum_{i<j} \Theta_{ij} a_{ij} \right) = \prod_{i<j} (1 + e^{-\Theta_{ij}}). \quad (4.10)$$

Following statistical mechanics, we calculate the free energy $F = - \log Z$ corresponding to the Hamiltonian to find the probability that there exists a link between node i and node j , specified by the coupling parameter Θ_{ij} :

$$p_{ij} = \frac{\partial F}{\partial \Theta_{ij}} = \frac{\exp[-\Theta_{ij}]}{1 + \exp[-\Theta_{ij}]}. \quad (4.11)$$

In our example of the unweighted directed network, the coupling parameter can be separated into a component describing the in-degree of node i and the out-degree of node j :

$$p_{ij} = \frac{\exp[-\Theta_{ij}]}{1 + \exp[-\Theta_{ij}]} = \frac{e^{-\lambda_i^{\text{out}} - \lambda_j^{\text{in}}}}{1 + e^{-\lambda_i^{\text{out}} - \lambda_j^{\text{in}}}} = \frac{x_i y_j}{1 + x_i y_j}, \quad (4.12)$$

where we have switched to the notation of Squartini and Garlaschelli [2011] in the last step. It is worth pointing out that this probability can also be calculated for bipartite networks. A bipartite network is a network in which we can divide the nodes in two subsets U, V , and all links exist only between a node in subset U and a node in subset V , but not within either subset. By treating all nodes in subset U to have an in-degree larger than zero and an out-degree of zero and all nodes in subset V to have an in-degree of zero and an out-degree larger than zero, we can apply Eq. (4.12) in a straightforward manner. Let's denote the nodes in subset U with the index i and the nodes in subset V with the index α , then the probability to find a link between these nodes is

$$p_{i\alpha} = \frac{x_i y_\alpha}{1 + x_i y_\alpha} \quad (4.13)$$

Since this thesis presents an empirical study of the bipartite Japanese firm and bank network, in the following we focus on the reconstruction of bipartite networks according to Squartini et al. [2017].

4.2.3 The fitness ansatz

The constrained maximum entropy approach yields an ensemble of graphs which satisfies the constraints in the sense of the ensemble average. In Section 4.2.2 we showed the calculation to find the probability for a link between two nodes, which is equivalent to the expected value of the adjacency matrix over the ensemble. This calculation requires knowledge of the degree sequence. Generally, we do not know the degree sequence. However, we may be able to estimate the density of the network or

equivalently the total number of links L . In this case the fitness *ansatz* has proven successful, for which we assume that there exist intrinsic node properties, or fitnesses, that are proportional to the x_i, y_α :

$$\begin{aligned} x_i &= \sqrt{z_V} v_i \\ y_\alpha &= \sqrt{z_C} c_\alpha. \end{aligned}$$

Defining $z = \sqrt{z_V z_C}$, Eq. (4.13) becomes

$$p_{i\alpha} = \frac{z v_i c_\alpha}{1 + z v_i c_\alpha}. \quad (4.14)$$

Since the number of links in a bipartite network is the sum over the degree of all nodes in both layers, the following relation has to hold:

$$L(G^*) = \langle L \rangle_{\mathcal{G}} = \sum_i \sum_\alpha \frac{z v_i c_\alpha}{1 + z v_i c_\alpha}, \quad (4.15)$$

where G^* is the graph we empirically observe. In order to find the probabilities $p_{i\alpha}$, we need to solve Eq. (4.15) numerically for the network we wish to reconstruct. This can be done with a simple optimization routine, minimizing the absolute difference of L and the double sum.

In the case of a weighted bipartite network, Squartini et al. [2017] reconstruct the weights of any given realization $G \in \mathcal{G}$ as

$$w_{i\alpha} = \frac{v_i c_\alpha}{W p_{i\alpha}} a_{i\alpha} = (z^{-1} + v_i c_\alpha) \frac{a_{i\alpha}}{W}, \quad (4.16)$$

where W is the sum of all weights and $a_{i\alpha}$ is the adjacency matrix of the realized graph G . Without the fitness parameter, the reconstructed network would be fully connected. As the last part Eq. (4.16) shows, however, the introduction of the fitness parameter z leads to a sparsification of the network by introducing a minimum weight

of z^{-1} . Since the reconstructed $w_{i\alpha}$ have to sum up to the total weight W , some links cannot be formed because an additional flat cost is incurred for each link. The fitness parameter z grows monotonously with the density of the network, and in the following section we show some results for different underlying distributions for the node properties v_i and c_α .

4.2.4 The challenge of sparse networks

While we are always able to reconstruct the network on the ensemble level as discussed in Sec. 4.2.2, there is one constraint that the methodology does not take into account: each node has to be part of the reconstructed network. If a certain node does not have a large value for its node strength, connection probability to a node in the other layer will be small. This is particularly problematic if the bipartite network is very imbalanced, meaning that one layer consists of a large number of nodes while the other layer has very few nodes in it. If we sparsify the network sufficiently, the expected degree of a weak node in the big layer will be less than one, and therefore most reconstructions in the ensemble leave this node unconnected. This would violate a constraint which we have ignored in the derivation of the reconstruction method, that is, we know of the presence of each node in the system, so it needs to be connected.

In fact, real life networks can be sparse with densities of $d \leq 0.1$. To illustrate this challenge, let's consider a bipartite network with 400 nodes in layer N and 1000 nodes in layer M. The strengths of the nodes in both layers follow a log-normal distributions, in layer N we have $V \sim \text{Lognormal}(0, 1)$ and in layer M we have $C \sim \text{Lognormal}(0, 2)$. By design, this network incorporates the two challenges for reconstruction: the layers are quite different in size, and the nodes on the bigger layer vary in strength more strongly. Using Eq. (4.15), we calculate the value for z that corresponds to a network density of $d = 0.1$, that is, a network with 20 000 links.

In Figure 4-1 we observe that weak nodes have very low connection probabilities.

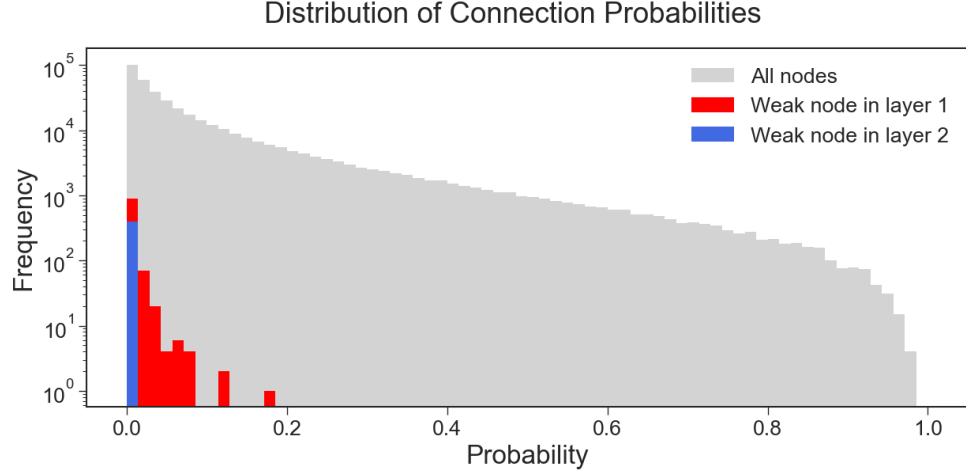


Figure 4-1: Distribution of the connection probabilities in a simulated sparse bipartite network of size 400×1000 with density $d = 0.1$. The probabilities for a single weak node in layer N are shown in blue, for a weak node in layer two in red, and the overall distribution of connection probabilities is shown in light gray. Note that a node in layer N has 1000 possible links, and a node in layer M has 400 possible links. We observe that the connection probabilities of the weak nodes are exclusively at the lower end of the overall distribution.

This is exacerbated in layer M where the strengths of the nodes have greater variance. layer N is more homogeneous with respect to the node strengths, and therefore higher connection probabilities are achieved even for the weak nodes in the layer.

Picking node i in layer N , we calculate its connection probability $p_{i\alpha}$ to all other nodes $\alpha = 1, \dots, 1000$ in layer M , and vice versa. The probability that node i or α remains disjoint from the network in any given realization of the reconstruction then is simply

$$P(\text{Node } i \text{ unconnected}) = \prod_{\alpha} (1 - p_{i\alpha}),$$

$$P(\text{Node } \alpha \text{ unconnected}) = \prod_i (1 - p_{i\alpha}).$$

Figure 4-2 shows the results of our simulation and illustrates that nodes in layer N ,




Figure 4.2: Distribution of probabilities of a node to be unconnected in any given realization of the network reconstruction. The nodes in layer N have at least degree one in practically all reconstructions, whereas there is a significant chance for a number of nodes in layer M to remain unconnected.

which are similar in strength and have many possible partners in layer M , have a probability of being unconnected of practically 0. In contrast, nodes in layer M , which are heterogenous in strength and have fewer possible partners in 1 , may have significant non-zero probabilities of remaining unconnected. In our example, like all nodes in its layer, the weak node in layer M has 400 possible links, and the largest connection probability among these links is just 0.45%. As a result, this node has a chance of not being connected in any given reconstruction of 67.5%. 5.2% of the nodes have a chance of more than 5% of remaining unconnected.

4.3 Distribution of connection probabilities

The node properties V, C and the density of the network to be reconstructed both determine the distribution of connection probabilities, $P(p_{i\alpha})$. Let us consider the limit of an infinitely large bipartite network, that is, a network with N nodes in layer N and M nodes in layer M where we take the limit $N, M \rightarrow \infty$ as we keep the degree of connectivity fixed. We show the distribution of connection probabilities

in terms of the density of the network d for two node property distributions: the uniform distribution and the log-normal distribution. While the uniform distribution is unrealistic for empirical data, it allows us to obtain analytical results to demonstrate the behavior of the fitness parameter z . Closest to empirical data in the reconstruction of financial networks is the log-normal distribution because the size loan portfolio of a set of banks can be approximated this way.

4.3.1 Uniform distribution

The continuous uniform distribution for a random variable with support between 0 and a real number a has the probability distribution

$$f(x) = \begin{cases} \frac{1}{a} & \text{for } 0 \leq x \leq a \\ 0 & \text{else.} \end{cases} \quad (4.17)$$

Let's assume both V and C are uniformly distributed between 0, 1. Note that the random variables V and C only appear as a product in Eq. (4.13), and therefore the behavior of $X = V \cdot C$ is important. If $V, C \sim \mathcal{U}(0, 1)$, then $X = V \cdot C$ follows the probability distribution

$$f_X(x) = -\log x, \quad 0 < x \leq 1.$$

Equation (4.13) transforms the input into a probability through the function $g(x) = zx/(1 + zx)$. Since $g(x)$ is strictly increasing for all allowed values and maps to the interval $[0, 1]$, it has an inverse function,

$$h(y) = g^{-1}(y) = \frac{y}{z(1 - y)}$$

with the derivative

$$h'(y) = \frac{1}{z(1 - y)^2}.$$

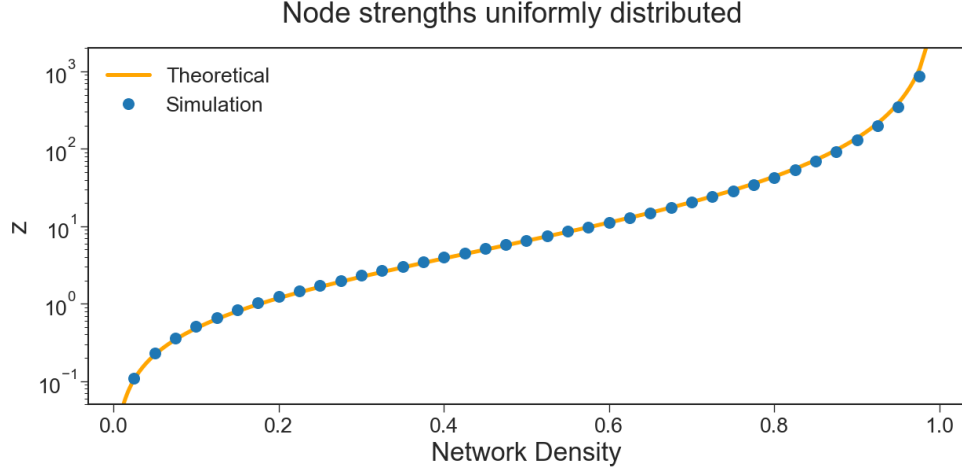


Figure 4.3: Fitness parameter z as a function of the network density d in a semilogarithmic representation based on Eq. (4.21) and from a simulation of a bipartite network with 300×100 nodes.

Assuming that the each V and C are uniformly distributed, we can calculate the distribution of connection probabilities $f_Y(y)$ from their distribution $f_X(x)$ in Eq. (4.18) using the transformation rule

$$f_Y(y) = f_X(h(y)) |h'(y)| = \frac{\log \frac{z(1-y)}{y}}{z(1-y)^2} \quad (4.18)$$

for $0 < y \leq \frac{z}{1+z}$ and 0 everywhere else.

Rewriting equation (4.15) in terms of a single random variable $X = V \cdot C$ yields

$$\langle L \rangle_{\mathcal{G}} = \sum_j \frac{zx_j}{1 + zx_j}. \quad (4.19)$$

The density of the network is the number of observed links divided by the number of possible links, and we can write

$$d = \frac{1}{N} \sum_j \frac{zx_j}{1 + zx_j} = \mathbb{E} \left[\frac{zX}{1 + zX} \right] = \mathbb{E}[Y]. \quad (4.20)$$

In the continuous limit we therefore find for $V, C \sim \mathcal{U}(0, 1)$ that density of the network

as a function of the fitness parameter z is

$$\begin{aligned}
 d(z) &= \mathbb{E}[Y] \\
 &= \int_0^{\frac{z}{1+z}} \frac{y}{z(1-y)^2} \log \frac{z(1-y)}{y} dy \\
 &= 1 - \frac{1}{z} \left[\frac{\pi^2}{6} + \log(z) \log(1+z) - \frac{1}{2} \log^2(1+z) - \text{Li}_2 \left(\frac{1}{1+z} \right) \right], \quad (4.21)
 \end{aligned}$$

where Li_2 is the polylogarithm of order 2. The inverse of this function $z(d)$ exists, but it has to be evaluated numerically. Figure 4.3 shows the behavior of z as a function of the network density and compares the theoretical result based on Eq. (4.21) with a simulated network of 300 nodes in one layer and 100 nodes in the other.

4.3.2 Log-normal distribution

If the random variable Y follows a normal distribution with mean μ and variance σ^2 , then $X = e^Y$ follows a log-normal distribution. The support of the log-normal distribution is the realm of positive real numbers, which makes it a natural candidate to describe the node properties or node strengths. Conveniently, the product of two independent log-normal distributions described by the parameters μ_1, σ_1^2 and μ_2, σ_2^2 , respectively, is again a log-normal distribution, with the parameters $\mu_1 + \mu_2$ and $\sigma_1^2 + \sigma_2^2$. If we therefore choose $V \sim \text{Lognormal}(0, \sigma^2/2)$ and $C \sim \text{Lognormal}(0, \sigma^2/2)$, the resulting random variable $X = V \cdot C$ is log-normally distributed:

$$f_X(x) = \frac{1}{x\sqrt{2\pi\sigma^2}} e^{-\frac{\log^2 x}{2\sigma^2}}.$$

However, it may be desirable to consider a network that consists of two inhomogeneous layers, particularly $\sigma_1^2 \neq \sigma_2^2$. To find the distribution of the connection probabilities,

we start with the cumulative distribution function (cdf)

$$P(Y \leq y) = P\left(Y \leq \frac{zX_1X_2}{1+zX_1X_2}\right) = P\left(X_2 \leq \frac{y}{z(1-y)} \frac{1}{X_1}\right), \quad (4.22)$$

which we can calculate as

$$\begin{aligned} F_Y(y) &= \frac{1}{\sqrt{4\pi^2\sigma_1^2\sigma_2^2}} \int_0^\infty \frac{dx_1}{x_1} \int_0^{\frac{y}{x_1z(1-y)}} \frac{dx_2}{x_2} \exp\left(-\frac{1}{2} \left[\frac{\log^2 x_1}{\sigma_1^2} + \frac{\log^2 x_2}{\sigma_2^2}\right]\right) \\ &= \frac{1}{\sqrt{8\pi\sigma_1^2}} \int_0^\infty \frac{dx_1}{x_1} e^{-\frac{\log^2 x_1}{2\sigma_1^2}} \left[1 + \operatorname{erf}\left(\frac{\log\left[\frac{y}{x_1z(1-y)}\right]}{\sqrt{2\sigma_2^2}}\right)\right]. \end{aligned} \quad (4.23)$$

This integral does not have an analytical solution. However, we can find the probability density function (pdf) where we can integrate out x_1 ,

$$\begin{aligned} f_Y(y) &= \frac{d}{dy} F_Y(y) \\ &= \frac{1}{2\pi\sigma_1\sigma_2 y(1-y)} \int_0^\infty dx_1 \exp\left(-\frac{1}{2} \left[\frac{\log^2 x_1}{\sigma_1^2} + \frac{\log^2 \frac{y}{x_1z(1-y)}}{\sigma_2^2}\right]\right) \\ &= \frac{1}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)}} \frac{1}{y(1-y)} \exp\left(-\frac{1}{2} \frac{\log^2 \frac{y}{z(1-y)}}{\sigma_1^2 + \sigma_2^2}\right). \end{aligned} \quad (4.24)$$

While we considered the case $\mu_1 = \mu_2 = 0$ for the derivation for better readability, including the parameters μ_1, μ_2 is straightforward:

$$f_Y(y) = \frac{\left(\frac{y}{z(1-y)}\right)^{\frac{\mu_1 + \mu_2}{\sigma_1^2 + \sigma_2^2}} \exp\left(-\frac{(\mu_1 + \mu_2)^2 + \log^2\left(\frac{y}{z(1-y)}\right)}{2(\sigma_1^2 + \sigma_2^2)}\right)}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)} y(y-1)}$$

In the following we continue with $\mu_1 = \mu_2 = 0$.

Since only the sum of the parameters σ^2 affects the distribution, let's use $\sigma_1^2 = \sigma_2^2 = \sigma^2/2$ to study the distribution and the implications for the fitness parameter. Figure 4.4 shows the behavior of the cdf $F_Y(y)$ for different σ^2 -parameters of the log-normal distribution and different desired network densities. If the desired network




Figure 4.4: Cumulative probability distribution $F_Y(y)$ of the connection probabilities in the reconstructed network for different densities d . Left: less variance in the node strengths in the layers. Right: more variance.

density d is small, a large number is assigned low probabilities. This effect is more pronounced if the parameter σ^2 is larger. Conversely, as the desired network density gets closer to 1, most node pairs are assigned a larger probability. Again this is more pronounced for larger σ^2 .

Let us now analyze the relationship between z and the network density d . Following Eq. (4.21), the density as a function of the fitness parameter in the log-normal case then is

$$d(z) = \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^1 \frac{e^{-\frac{\log^2 \frac{y}{z(1-y)}}{2\sigma^2}}}{(1-y)} dy. \quad (4.25)$$

This integral can only be solved numerically, but we observe that the fitness parameter z depends on the desired network density d and the combined variance of the layers in the network. For $d = 0.5$, the solution becomes independent of σ^2 , and we find $z = 1$. Figure 4.5 shows the behavior of z as a function of the network density for $\sigma^2 = 1$, comparing it to a simulation of a 40×40 bipartite network.

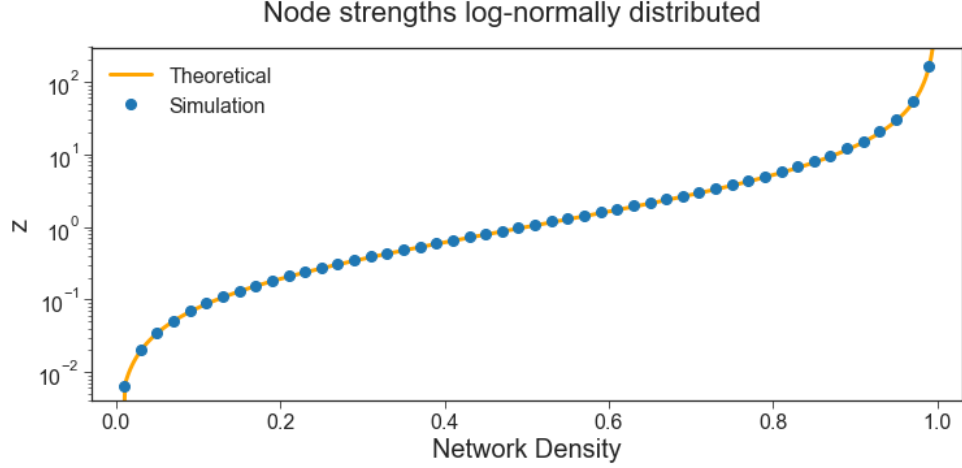


Figure 4-5: Fitness parameter z as a function of the network density d in a semilogarithmic representation based on Eq. (4.25) and from a simulation of a bipartite network with 40×40 nodes.

4.4 Hypothesis testing for connectivity

In Section 4.2.4 we used a simulated network to show that the reconstruction of sparse and lopsided bipartite networks may be unsuccessful in the sense that the reconstructed network is not connected. This is problematic because the study of systemic risk requires an accurate reconstruction of the network. If nodes remain unconnected in the reconstruction, we cannot account for their influence on systemic risk. In particular, we want to be able to check the following hypothesis:

\mathcal{H}_0^q : A node in layer i with strength q at the q -quantile of the distribution of strengths of nodes in the layer has at least degree 1.

In the following, we derive the equations to test this hypothesis for the case of log-normally distributed strengths, $X_1 \sim \text{Lognormal}(0, \sigma_1^2)$, $X_2 \sim \text{Lognormal}(0, \sigma_2^2)$. This allows us to estimate the limits of the reconstruction approach. Where we cannot solve the equations analytically, we provide numerical estimates when we have to reject hypothesis \mathcal{H}_0^q .

We formulated \mathcal{H}_0^q because we use as a proxy for the network connectivity an

individual node at a certain strength quantile q of the distribution of strength among the nodes in that layer. Typically, the likelihood that a weak node is connected in a reconstruction is of concern.

4.4.1 Analytical calculations

If we want to find the connection probability distribution of a node at the q -quantile of the distribution of strengths, we modify Eq. (4.23) in order to find the cumulative distribution function conditioned on the value of one node $P(Y \leq y \mid X_1 = x_q)$ with Bayes' rule:

$$P(Y \leq y \mid X_1 = x_q) = \frac{P(Y \leq y \cap X_1 = x_q)}{P(X_1 = x_q)}. \quad (4.26)$$

To compute the denominator, we need to find the value of the node strength pdf at the quantile. The quantiles of the log-normal distribution are calculated through

$$q = \int_0^{x_q} f_{X_1}(x_1) dx_1. \quad (4.27)$$

Thus, if a node is at the q -quantile of a Lognormal $(0, \sigma_1)$ distribution, then it has strength $x_q = \exp\left(-\sqrt{2\sigma_1^2} \operatorname{erfc}^{-1}(2q)\right)$, where erfc^{-1} denotes the inverse complementary error function. Obviously, $P(X_1 = x_q)$ depends on only σ_1 and q :

$$P(X_1 = x_q) = \frac{\exp\left(\operatorname{erfc}^{-1}(2q) \left[\sqrt{2\sigma_1^2} - \operatorname{erfc}^{-1}(2q)\right]\right)}{\sqrt{2\pi\sigma_1^2}} \quad (4.28)$$

In order to calculate the intersection, we need to modify Eq. (4.23) and introduce a Dirac delta to fix $X_1 = x_q$:

$$\begin{aligned}
P(Y \leq y \cap X_1 = x_q) = & \\
& \frac{1}{\sqrt{8\pi\sigma_1^2}} \int_0^\infty \frac{dx_1}{x_1} e^{-\frac{\log^2(x_1)}{2\sigma_1^2}} \left[1 + \operatorname{erf} \left(\frac{\log \left[\frac{y}{x_1 z (1-y)} \right]}{\sqrt{2\sigma_2^2}} \right) \right] \\
& \times \delta \left(x_1 - \exp \left(-\sqrt{2\sigma_1^2} \operatorname{erfc}^{-1}(2q) \right) \right). \tag{4.29}
\end{aligned}$$

The probability density function of the intersection is calculated as the derivative with respect to y as usual. Dividing by the pdf for $X = x_q$, we then get the conditional probability density function²:

$$f_{Y|X=x_q}(y) = \frac{1}{2\pi\sigma_2^2 y(1-y)} \exp \left(-\frac{\log^2 \frac{\exp(\sqrt{2\sigma_1^2} \operatorname{erfc}^{-1}(2q)) y}{z(1-y)}}{2\pi\sigma_2^2} \right) \tag{4.30}$$

The expected value μ_Y of $f_{Y|X=x_q}(y)$ can only be written in integral form,

$$\mathbb{E}[Y | X = x_q] = \int_0^1 dy y f_{Y|X=x_q}(y). \tag{4.31}$$

Likewise, the second moment only exists in integral form as well:

$$\mathbb{E}[Y^2 | X = x_q] = \int_0^1 dy y^2 f_{Y|X=x_q}(y). \tag{4.32}$$

² In case the marginal distributions have parameters μ_1 and μ_2 for the underlying normal distributions, this expression becomes

$$\begin{aligned}
f_{Y|X=x_q}(y) = & \frac{1}{2\pi\sigma_2^2 y(1-y)} \exp \left(-\frac{\mu_2^2 + \log^2 \frac{\exp(-\mu_1 + \sqrt{2\sigma_1^2} \operatorname{erfc}^{-1}(2q)) y}{z(1-y)}}{2\pi\sigma_2^2} \right) \\
& \times \left(\frac{\exp(-\mu_1 + \sqrt{2\sigma_1^2} \operatorname{erfc}^{-1}(2q)) y}{z(1-y)} \right)^{\mu_2/\sigma_2^2}.
\end{aligned}$$

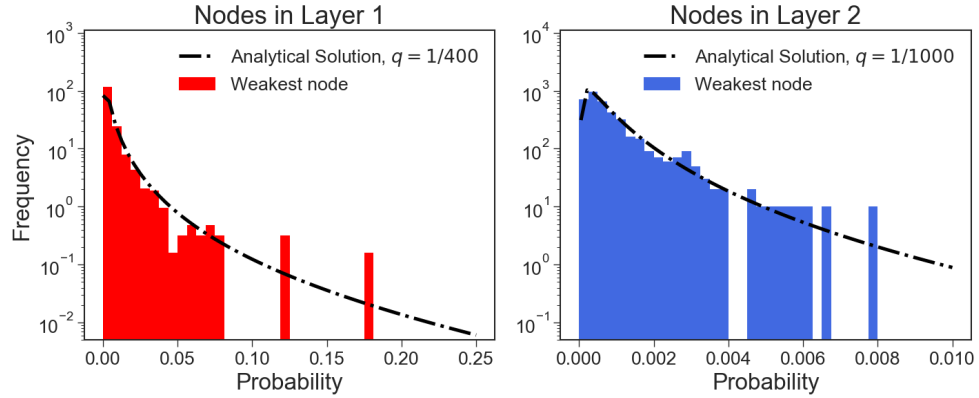


Figure 4-6: Comparing the analytical solution of the connection probability for the node in the lowest quantile with a weak node results from a simulated network.

It is integrable, meaning that the variance of Y is finite: $\sigma_Y^2 = \mathbb{E}[(Y - \mu_Y)^2] < \infty$.

Figure 4-6 shows that the analytical pdf is in good agreement with the connection probabilities for a weak nodes in our simulation. For the curve describing the analytical solution, we plug in the same parameters for σ_1 and σ_2 as in the simulation, and we choose the 1/400-quantile for the 400-node network and the 1/1000-quantile for the 1000-node network.

4.4.2 Numerical results

Equations (4.31) and (4.32) do not have closed form solutions. The integrals have to be solved numerically, and we explore the results for a large range of parameters.

Statistics primer

Before we delve deeper into the integrals and their implications, however, let us take a step back and consider a simple toy model for the connectivity of a node. We consider a node Q that can connect to N other nodes with some probability p . The degree of the node is the number of successful connections and can range from 0 to N . If we are concerned with connectivity, we want to study how likely it is that the node ends

up with degree zero.

From a statistical point of view, we are performing N independent Bernoulli trials with probability p and summing up their outcome. This sum K , which describes the degree of the node, follows a binomial distribution $B(N, p)$ with mean Np and variance $Np(1-p)$. In order to find the probability that the node is connected we simply have to calculate $P(K = 0) = (1-p)^N$.

So far, we have considered the special case that node Q can connect to each of the N nodes with the same probability. Let us relax this requirement such that the connection probability is different for each node $n = 1, \dots, N$, namely p_n . Then we are performing N independent Bernoulli trials with probabilities p_1, \dots, p_N . Again we calculate the sum K to find the degree of the node. K then follows a Poisson binomial distribution, a generalization of the binomial distribution. The distribution has mean $\sum_{n=1}^N p_n$ and variance $\sum_{n=1}^N (1-p_n)p_n$. Like in the case for identical probabilities, we can calculate the probability $P(K = 0)$ to find the probability that node Q has degree zero:

$$P(K = 0) = \prod_{n=1}^N (1 - p_n). \quad (4.33)$$

Until now, we have assumed that we know the connection probabilities p_1, \dots, p_N . If we consider a network in which the N nodes have some strength distribution from which we compute their connection probability to node Q, then these connection probabilities themselves are variates y_1, \dots, y_N of the random variables Y_1, \dots, Y_N . Node Q has degree zero with probability $P(K = 0 | Y_1 = y_1, \dots, Y_N = y_N)$, that is, the probability conditional on observing the values y_1, \dots, y_N for the N Bernoulli trials.

Using the law of total probability, we get

$$\begin{aligned}
 P(K = 0) &= \int_{y_1=0}^{\infty} \cdots \int_{y_N=0}^{\infty} P(K = 0 | Y_1 = y_1, \dots, Y_N = y_n) \\
 &\quad \times f_{Y_1}(y_1) \cdots f_{Y_N}(y_N) \, dy_1 \cdots dy_N,
 \end{aligned} \tag{4.34}$$

where we have used the fact that the Y_n are independent random variables, and $f_{Y_n}(y_n)$ are their corresponding probability density functions. The conditional probability for $K = 0$ has the same form as the probability in Eq. (4.33), and therefore the integrals factorize, yielding

$$P(K = 0) = \left(\int_0^{\infty} (1 - y) f_Y(y) \, dy \right)^N. \tag{4.35}$$

We can further simplify this expression using the expected value,

$$P(K = 0) = (1 - \mathbb{E}[Y])^N, \tag{4.36}$$

showing that the connectivity of the node depends on the number of nodes to connect to and the probability distribution from which we draw the connection probabilities.

Equation (4.36) calculates the probability of degree zero over all possible combinations for the probabilities y_1, \dots, y_N underlying the Bernoulli trials. This allows us to test the null hypothesis \mathcal{H}_0^q in terms of the expected value of the connection probability of node Q.

Connectivity of a node with $f_{Y|X=x_q}(y)$

Let us now return to our example of the bipartite network in which one layer has N nodes, and their strengths are distributed log-normally with parameters $\mu = 0$ and $\sigma^2 = \sigma_N^2$. We also refer to this layer as the in-layer. Similarly, the second layer has M nodes, and their strengths are also distributed log-normally with parameters $\mu = 0$,

$\sigma^2 = \sigma_M^2$. Accordingly, we refer to the second layer as the out-layer (because a node wants to connect out of it). We consider node Q from our toy model be one of the M nodes in this second layer. We assume node Q to have strength x_q which is the value at the q -quantile of the strength distribution across the nodes in its layer. In Eq. (4.30), we calculated the distribution of connection probabilities of such a node, denoted by $f_{Y|X=x_q}$. We recall the expected value of Y conditioned on $X = x_q$,

$$\mathbb{E}[Y | X = x_q] = \frac{1}{2\pi\sigma_M^2} \int_0^1 \frac{dy}{1-y} \exp\left(-\frac{\log^2 \frac{\exp(\sqrt{2\sigma_N^2} \operatorname{erfc}^{-1}(2q)) y}{z(1-y)}}{2\pi\sigma_M^2}\right). \quad (4.37)$$

For readability, let μ_Q denote the expected value $\mathbb{E}[Y | X = x_q]$ for node Q. Its probability of being disconnected in a reconstruction then is $P(K = 0) = (1 - \mu_Q)^N$. Obviously, this probability becomes negligible if μ_Q or N are sufficiently large. If the connection probabilities for node Q are small, however, then μ_Q is small as well. In this case the probability that Q is disconnected in a realization the network reconstruction may be significantly larger than 0. Considering $\mu_Q < 0.1$, we set a threshold $10^{-\theta}$ and take the logarithm of the probability for degree equal to zero. This yields the condition $N \log(1 - \mu_Q) \approx N\mu_Q \leq \theta \log 10$, an easy criterion to determine whether a combination of N and μ_Q exceeds a probability threshold.

An empirical network which we want to reconnect typically has N between 100 and 2000. A 99.9% success rate in reconstruction (the node Q has degree of zero in only 0.1% of the cases) implies $\theta = 3$. The critical value of μ_Q then is in the range of 0.0035 for $N = 2000$ and 0.069 for $N = 100$. Put another way, if the expected value of $Y | X = x_q$ is larger than these values for a given N , node Q is connected in practically all network reconstructions.

Figures 4.7 and 4.8 show the expected value for a broad range of network parameters. We plot the results for different quantiles of node strength for node Q. Since

the quantile provides a relative strength estimate, this is equivalent to considering different values M for the size of the layer in which Q is.

Figure 4-7 varies the σ^2 parameter of the strength distribution in one of the layers while holding it constant for the other layer. The network is reconstructed with a density of $d = 0.1$. In the left panel we see that the expected value of connection probabilities grows for node Q if the parameter σ^2 in the layer of nodes Q can connect to becomes larger. The intuition behind this is as follows: If the network is sparse, as is the case in this example, a node with low strength in one layer is likely only able to connect to nodes with high strength in the other layer. If the parameter σ^2 is larger in that other layer, there likely exist a few nodes with very high strength. Then node Q , even if its strength is low, has a reasonable chance of forming at least one link. In the right panel we see that the expected value of connection probabilities decreases for node Q if the parameter σ^2 in the layer in which Q is becomes larger. If the strength distribution of this layer becomes more heterogeneous, then node Q at a low quantile becomes weaker in relation to most other nodes in this layer. For example, a node at quantile $q = 0.95$ is stronger than a node at quantile $q = 0.05$ by factor of roughly $e^{3.29\sigma^2}$ if the strengths are log-normally distributed with parameters $\mu = 0$ and σ^2 . As a result, the nodes with large strength have such high probabilities in the reconstruction that most links in a sparse network are assigned to them, leaving node Q more likely to be unconnected. Note that we have not considered yet the number of nodes N to which node Q can connect. The larger N is, the lower is the threshold for the expected connection probability for Q that yields at least one connection in the reconstruction.

Figure 4-8 varies the density of the network which we want to reconstruct for fixed parameters σ^2 for the two layers of the network. Unsurprisingly, the expected value of the connection probability for Q grows with the density of the network, and it




Figure 4-7: Average expected connection probability for node Q if its strength is at the quantile values $q = 1/400, 1/1000, 1/2500$ depending on the parameters of strength distributions. The network is reconstructed with density $d = 0.1$. The left plot shows $\mathbb{E}[Y | X = x_q]$ if we vary the the parameter σ_N^2 of the log-normal distribution of the strengths of the nodes to which node Q can connect. The right plot shows $\mathbb{E}[Y | X = x_q]$ if we vary the the parameter σ_M^2 of the log-normal distribution of the strengths in the layer of node Q. In both cases we keep the parameter σ^2 fixed to 1 for the other layer.




Figure 4-8: Average expected connection probability for node Q if its strength is at the quantile values $q = 1/400, 1/1000, 1/2500$ depending on the network density d . The left plot shows $\mathbb{E}[Y | X = x_q]$ if σ^2 is bigger for the layer of the nodes to which Q can connect than for the layer in which Q is. The right plot shows $\mathbb{E}[Y | X = x_q]$ if σ^2 in the opposite case.

approaches one for a fully connected network. To get the probability that node Q remains unconnected, we again need to consider what N is. Assume $N = 100$ and the more favorable case of $\sigma_N^2 = 2, \sigma_M^2 = 1$ like in the left panel of Figure 4-8. In order to have a probability of less than 10^{-3} that node Q is unconnected, the critical value for the expected value is roughly 0.069, as we calculated earlier. If $q = 1/400$, this value is reached at network density $d \approx 0.41$. In relation to empirical networks, this is a large value. If $q = 1/1000$ or $q = 1/2500$, d needs to be roughly 0.44 and 0.50, respectively. In the less favorable case of $\sigma_N^2 = 1, \sigma_M^2 = 2$ like in the right panel, d needs to be 0.66, 0.72, and 0.76 for the different quantiles. Most empirical networks are not this highly connected, and the network reconstruction would fail for node Q for the parameters used in this example.

So far, we have focused on the expected value $\mathbb{E}[Y | X = x_q] = \mu$. We continue the numerical analysis by calculating the probability that node Q remains unconnected, which is given by $(1 - \mu_Y)^N$, for different network densities, strength distributions,

and values of $N = 100, 400, 1000$.

Figure 4.9 summarizes the probability that node Q remains unconnected for a network of density $d = 0.20$ if its strength is at different quantiles of the strength distribution. The color shows the probability of $P(K = 0)$, and the x and y-axes correspond to different strength distributions in the two layers of the bipartite network. The quantile q is an estimate for its strength in relation to the other nodes in its layer. Naturally, the smaller the quantile value, the more likely it is that node Q is not connected to any of the N nodes it could connect to. An alternative interpretation of the quantile value is the number of nodes node Q needs to compete against: if q is very small, this is equivalent to node Q competing against a large number of nodes to connect to the N nodes, and vice versa. Increasing the number of potential partners N improves the connection probability considerably. A large value of N is particularly important if the parameter σ_M^2 , which describes the heterogeneity of node Q's layer, is large.

Let us briefly discuss the two most extreme cases. The top left panel corresponds to a small network in which the layer of node Q has very many other nodes in it, while the number of connection partners is rather low. Then for most strength distributions node Q remains unconnected. The bottom right panel corresponds to the opposite case; there are few nodes in the layer of node Q, and they have many options to connect to. Then for practically all strength distributions, node Q is connected in the reconstructed network.

Let us now consider the case of two layers which have the same strength distribution, that is, the parameter σ^2 is identical. Then the only difference between the two layers is the number of nodes in them, measured by N and $1/q$. Figure 4.10 shows, from top to bottom, that networks with more homogeneous strength distributions are easier to reconstruct in the sense that node Q is less likely to be disconnected. On

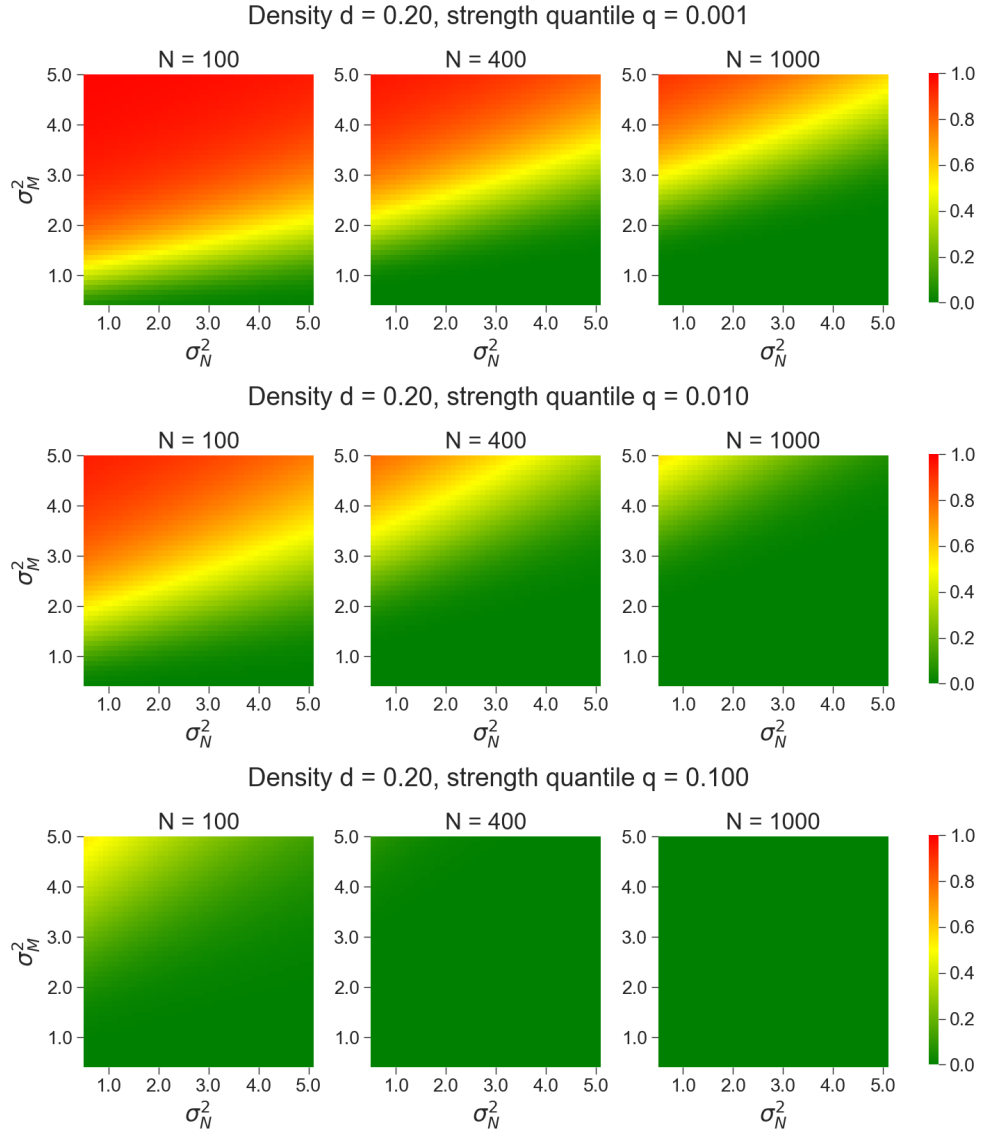


Figure 4-9: Probability that node Q remains unconnected in a reconstruction of a network with density $d = 0.2$ as a function of the strength distributions. The strength of node Q is, from top to bottom rows, at quantile $q = 10^{-3}$, $q = 10^{-2}$, and $q = 10^{-1}$. From left to right, we calculate the cases $N = 100$, $N = 400$, and $N = 1000$.




Figure 4-10: Probability that node Q remains unconnected in a reconstruction of a network where the layers have the same strength distributions. We consider different combinations of density d and quantile q to compute the probability. The parameter σ_2 of the strength distribution of the layers is, from top to bottom rows, $\sigma_2 = 1$, $\sigma_2 = 2$, and $\sigma_2 = 4$. From left to right, we calculate the cases $N = 100$, $N = 400$, and $N = 1000$.

the y-axis of the heat maps we mark the density of the network, and on the x-axis we show the strength quantile q , equivalent to the inverse of the layer size M . If Q has a lot of competition to connect, that is, for small q or larger M , the reconstruction is difficult for a broad range of densities, and it gets more and more difficult with growing σ^2 , the layer inhomogeneity. For sufficiently large q or small M , node Q is likely to be connected in the reconstructed network for all but very sparse networks or low N .

In the following we supplement these theoretical calculations with a case study, the reconstruction of the Japanese bank-firm network.

4.5 Data

The Nikkei Economic Electronic Databank System (NEEDS) contains large amounts of financial data of Japanese banks and firms. It details the lending relationship between them, including the outstanding amount in Japanese yen as well as a distinction between long-term and short-term debt. The data set further includes corporate balance sheets, making it an ideal source to study the effect of financial shocks and their effect on the bank-firm network. While the database captures bank-firm relationships since the 1980s for the biggest publicly listed companies, NEEDS underwent some changes in the mid-1990s when its scope was expanded. Figure 4.11 illustrates this change; since 1996 the database includes companies traded in over-the-counter market, leading to an increase in the number of firms covered. These companies are generally smaller than its previously covered counterparts. The number of banks in the data set has slightly declined over time.

Due to the fundamental change in the composition of the database, the research presented in this thesis covers the years 2000 through 2012. The significant change in the pool of companies yielded consequences for the structure of the bank-firm net-

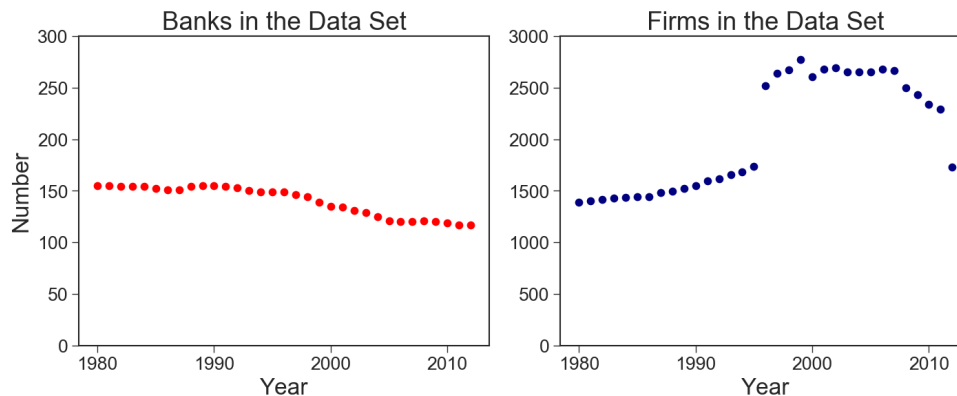


Figure 4-11: Scope of the Nikkei NEEDS database which includes over 100 banks and between 1400 and 2800 firms, depending on the year.

work. The average degree of companies decreased after the expansion in the database, presumably due to the inclusion of smaller firms with less need for large financing. Roughly since the year 2000, however, the density of the network (which is the number of empirical links divided by the number of possible links), the average degree of banks, and the average degree of companies has remained fairly constant.

4.6 Reconstruction of an empirical network

Since the data in NEEDS contains firm and bank specific information as well as the exact lending relationship for each year, it allows to test the reconstruction obtained through the configuration model with its fitness *ansatz*, introduced in Section 4.2.3. At the heart of the configuration model is the assumption that there exists a connection between a specific, intrinsic node property called strength and its degree k . Recalling Eq. (4.14), the probability for a link between a node i in one layer and a node α in the other layer is

$$p_{i\alpha} = \frac{z v_i c_\alpha}{1 + z v_i c_\alpha},$$

where z is the fitness parameter and v_i and c_α are the node strengths.




Figure 4.12: Left: Network density of the bipartite graph derived from banks in one layer and companies in the other. Right: Average degree of the nodes in each layer.

A link in the original network indicates an outstanding loan between bank and their client, a company, and its weight corresponds to the size of the loan. Mathematically, this is described by the weighted adjacency matrix \mathbf{W} with entries $w_{i\alpha}$. The marginals then are the total outstanding loans for both banks and customers. For banks, this number corresponds to their exposure to the market, and for companies, this number corresponds to the amount of money they owe:

$$v_i = \sum_{\alpha} w_{i\alpha},$$

$$c_{\alpha} = \sum_i w_{i\alpha}.$$

Using Eq. (4.15), we calculated the marginals of all nodes and found the fitness parameter z that reconstructed the density of the network. Since z is the only free parameter of this reconstruction method, finding its value directly yielded a connection probability for each node pair. Furthermore, a given value of z implies an expected degree $\langle k \rangle$ for a node of strength v_i or c_{α} . For more information regarding the calculation, see Squartini et al. [2017]. Figure 4.13 confirms that the strength of the node as given by the total of outstanding loans is a reasonable good predictor for

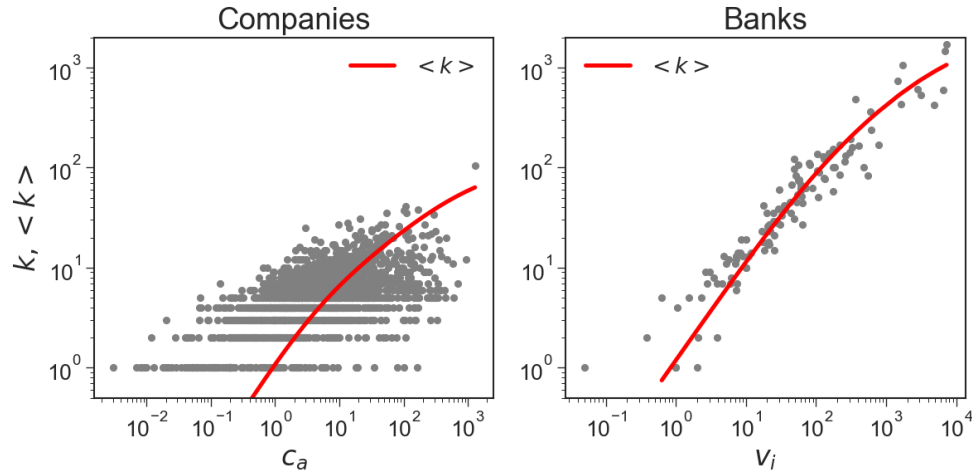


Figure 4-13: The Bipartite Configuration model assumes that there exists a relationship between the node strengths v_i , c_α and the degrees of the nodes. The empirical data is plotted as gray dots, and the theoretical prediction from the model in red.

the degree of the corresponding bank or firm. It is noteworthy that the variance of firm size given a certain degree is very large, while the model expectation is close to the empirical data for banks.

We tested whether the marginals followed a log-normal distribution by using the Shapiro-test on the logarithm of the strength values. As a null hypothesis, we assumed that the log-transformed data were normally distributed. The alternative hypothesis was that they were not normally distributed. We rejected the null hypothesis for the banks; however the sample was too small to form another hypothesis. For most years we fail to reject the null hypothesis for the companies at the 95% confidence level, implying that the strengths of firm nodes is log-normal. Coincidentally, we rejected the null hypothesis in the years 2000, 2001 and 2008, coinciding with major financial crisis.

We present the network reconstruction exemplarily for the year 2010. While the distribution of the bank strengths is unknown, we determined the distribution of firm strengths to be log-normal. The means and standard deviations of the strength data




Figure 4-14: Empirical and theoretical distribution of connection probabilities for the reconstruction of the network of Japanese banks and firms in the year 2010.

computed to $\mu_B = 1.401$, $\sigma_B = 1.788$ and $\mu_C = 3.665$, $\sigma_C = 2.134$ for banks and companies, respectively. Figure 4-14 shows the distribution of connection probabilities, comparing it to the analytical solution for $f_Y(y)$, assuming log-normal marginals with the parameters above and the empirical network density which is known from the data set. Despite the deviation from the model in case of the banks, the agreement is remarkable. Due to the low density of the network, $d \approx 0.051$, most probabilities are very small.

We concluded from our analytical calculations and numerical results that the reconstruction of low sparsity networks proves particularly difficult for the nodes in the much larger layer; in the case of the Japanese bank-firm network, this is the layer with the companies. In the data set, they outnumber banks by a factor of 10 to 25, depending on the year. For the year 2010, which we focused our analysis on, the data set contains 119 banks and 2339 companies. We thus hypothesized that the ensemble of reconstructions would contain many networks in which a significant number of companies have degree zero, that is, they are excluded from the reconstructed network. We generated 1000 realizations of the network, and we checked




Figure 4-15: Empirical and theoretical distribution of connection probabilities for a small bank and a large bank.

how many banks and firms remained unconnected. Of the 119 banks, on average 3.47 (2.92%) were unconnected with a standard deviation of 1.28, and the median was 3. The minimum number of unconnected banks was 0 and the maximum 8. Of the 2339 companies, on average 373.4 (16.0%) were unconnected with a standard deviation of 13.0, and the median was 374. The minimum of unconnected companies was 334 and the maximum 415. In the 1000 realizations, one bank and 20 companies remained unconnected every single time.

Exemplarily, we show the connection probabilities for a bank with small outstanding loans and bank with large outstanding loans in Figure 4-15. We picked the bank with the second smallest and the bank with the second largest portfolio and assumed their strengths to be at the quantile values $q = 2/119$ and $118/119$. To calculate the conditional distributions for their connection probabilities, we estimated the parameters for mean and standard deviation for the log of their strength values v_i and c_α . We assumed the marginal distributions to be log-normal with the estimated parameters. Despite these approximations, the fit is remarkable for both a bank at a low and a bank at a high quantile. Similarly in Figure 4-16 we show the connection probabili-




Figure 4.16: Empirical and theoretical distribution of connection probabilities for a small company and a large company.

ties for a small company in terms of outstanding loans and a large company. Since $N = 119$ possible connection yields a much smaller sample, the agreement of empirical probabilities and the theoretical curve is not as good as for the banks. However, it describes the trend well.

4.7 Conclusion

When detailed network information is not available, we have to resort to reconstruction methods. For bipartite networks, the enhanced capital-asset pricing model (ECAPM), a configuration model, has proven successful in recent studies. Using a maximum entropy approach, ECAPM generates an ensemble of networks which fulfills known constraints of the original network. To generate the ensemble, the algorithm relies on knowledge about the density of the network and some intrinsic node parameters which are typically called strengths. If the network that needs to be reconstructed is very sparse, like the Japanese bank-firm network which we studied, ECAPM may fail to connect a significant portion of nodes in many graphs of the ensemble. Analytically and numerically, we have determined ranges of validity

for the bipartite configuration model. Using this chapter as a blueprint, our calculations can easily be transferred to other reconstruction methods, like the universally renormalizable model. to establish similar bounds.

Assuming log-normal distributions for node strengths in the layers of the network, we derived the probability distributed of the connection probabilities according to which links are formed in the reconstruction. We explored the results numerically to determine for which parameters, like network density, variance of the marginal distributions, and network size, the chance of not connecting nodes at the lower quantiles of the strength distribution is significantly different from zero. Using empirical data from the Japanese bank-firm network, we illustrated that the analytical results provide valuable insights into the success or failure of network reconstruction even if the assumptions of log-normal marginals are violated. So far, reconstructions have been evaluated by considering how well the method recovers topological features like degrees, nearest-neighbor degrees or clusters. Our work may be used to add to this toolkit to determine the reliability of network reconstructions upon on which systemic risks models are built.

A couple of avenues for future research offer themselves. Firstly, we may reconstruct networks imposing a minimum degree of one for each node. In existing reconstruction models, either the degree of nodes is imposed, requiring extensive knowledge of the system, or the total number of links. An improvement of ECAPM, which is of the latter type, needs to enforce degree of at least one for each node. Secondly, renormalizable models may mitigate the limitations of configuration models as they allow to group nodes and reconstruct on an aggregate level.

Chapter 5

Measuring Currency Comovements with Relative Entropy

We develop a new tool to measure the correlation of currencies. Using a network interpretation of the foreign exchange market, we compute the information contained in the edges, that is, in individual currency pairs. We identify currency pairs with large deviations from the null model of no shared information. Studying the comovement on an annual and quarterly basis, we observe relative stability of the foreign exchange market. Finally, we relate the new measure to macroeconomic data such as government bond yields, bilateral trade volumes, currency classifications, and daily turnover. We build a generalized linear model in which we incorporate both the macroeconomic data and idiosyncratic properties of currencies. We find that the bilateral trade volume between two countries or regions is a significant predictor of their currencies' comovement. However, stock index correlation or share in daily turnover cannot be confirmed as statistically significant influencers of our comovement measure.

5.1 Introduction

The foreign exchange market is the largest financial market in the world by daily turnover, and it is more liquid than other well-studied financial markets such as equity or bond markets. While currency exchange is essential for world trade of goods and services, participants in the foreign exchange market also include investors who wish to invest in non-domestic assets like international stocks or bonds and need to exchange currency, money managers who wish to hedge their positions in global equity markets, or speculators. As for any traded asset, understanding correlations and dependencies between different currencies is crucial to a successful trading strategy. Generally, traders analyze the correlations between different pairs to guide their actions and tune their strategies. While useful for practical application, these correlations between currency pairs offer only limited insight into the underlying structure of the foreign exchange market.

Taking a step back, we interpret the foreign exchange market as a network in which the nodes are the currencies like the US dollar or the euro. The nodes are connected through weighted links which reflect the exchange rate of a currency pair. For example, the link between the US dollar node and the euro node is the EUR/USD exchange rate, and as such it is a weighted link. Since the links are determined by exchange rates, we refer to this network as the exchange rate network.

While the underlying macroeconomic fundamentals such as interest rates, inflation or GDP affect the nodes of the exchange rate network, trading happens on the edges. Changes in the characteristics of nodes happen over long time periods, that is, quarters, years, or even decades. In contrast, the weights of the edges fluctuate on very short time scales as the laws of supply and demand constantly incorporate information about the nodes and future expectations of existing and new market participants. To understand the dynamics of the foreign exchange market better, the

task is then to relate the information contained in the nodes to the exchange rate fluctuations purported by the links.

In this chapter, we identify channels of information transmission in the foreign exchange network, that is, edges which carry significant information about the comovement of the currencies which they connect. We achieve this by introducing a new measure rooted in information theory and with close connection to statistical physics, the principle of maximum entropy. Considering a pair of currencies, we compute the bivariate distribution of a network observable. We measure the amount of information encoded in this empirical distribution and compare it to the null model of independent currencies. This analysis provides novel insight into the correlation structure of the foreign exchange market in two ways. Firstly, our study extends beyond considering the correlation of currency pairs in isolation of the rest of the market. Secondly, we are able to quantify the comovement of currencies regardless of base currency. This removes the bias that is inherent to calculating, e.g., Pearson's correlation or copulas using a specific base currency. Papell and Theodoridis [2001] and Hovanov et al. [2004] constitute only two examples of many in the literature showing the impact of this bias.

The network observable that we study is called the degree imbalance and describes the difference between in-degree and out-degree of a node [Mubayi et al., 2001]. This thesis offers a network interpretation of the symbolic performance, introduced by Wollschläger et al. [2018]. The symbolic performance transforms the exchange rate data into a ranking of appreciation and depreciation of individual currencies against other currencies in the market. Thus it both discretizes the time series and removes the effect of market volatility. The symbolic performance has proven useful to identify the hierarchy in the foreign exchange market and to study the effect of central bank interventions. It is discussed in more detail in Chapter 6 of this dissertation.

The rest of this chapter is organized as follows: We cover the necessary concepts from network science and information theory in Section 5.2.1. We briefly present the rationale behind the symbolic performance (Wollschläger et al. [2018] or Chapter 6). Then we further show how to calculate the degree imbalance in the foreign exchange network in order to recover the symbolic performance. We establish the maximum entropy approach to study the null model of independence among currencies, and we introduce the Kullback-Leibler divergence to measure the deviations of the empirical data from the null model. In Section 5.2.5 we describe the foreign exchange data and the macroeconomic data which we have downloaded from various sources. Section 5.3 introduces a generalized linear model with macroeconomic indicators as predictors for the Kullback-Leibler divergence. We discuss these results in Section 5.4.

5.2 Methodology

5.2.1 Symbolic performance

Following Wollschläger et al. [2018], we study a financial market made up of K currencies. $K - 1$ of these assets are quoted in a base currency, in our case that is the euro. This gives us the exchange rate vector

$$\vec{S}(t) = (S_1(t), S_2(t), \dots, S_K(t)) = (1, S_2(t), \dots, S_K(t)), \quad (5.1)$$

where we have set $S_1 = S_{\text{EUR}/\text{EUR}} = 1$ as a placeholder for the base currency. If the market is sufficiently liquid, there should be no opportunity for arbitrage, and then we can infer all exchange rates between currency i and j , with $i, j = 2, \dots, K$ from this vector. In other words, considering $\vec{S}(t)$ is sufficient to capture the entire market. We proceed to calculate the time series of the log-returns,

$$R_i(t + \Delta t) = \log \frac{S_i(t + \Delta t)}{S_i(t)}. \quad (5.2)$$

Log-returns have the desirable probability that they are additive across time, which is not the case for returns calculated as the relative change. Especially in foreign exchange markets, however, where the returns are very small for time steps of the order of days or less, both measures practically yield the same numerical result.

The symbolic performance of currency i is defined as the number of currencies j it rises against subtracted by the number of currencies $K - j - 1$ it falls against. The symbolic performance $\zeta_i(t)$ of currency i at time t can then be found from ranking the returns of all currencies and identifying its ordinal number:

$$\zeta_i(t) = 2 \text{rank}[R_i(t)] - (K + 1). \quad (5.3)$$

The values for $\zeta_i(t)$ range from $-(K - 1)$ to $(K - 1)$.

Alternatively, let us consider a network of K nodes which correspond to K currencies. If the links correspond to foreign exchange rate information, then the network is fully connected, and thus it has $K(K - 1)$ links. Each of these links is bidirectional: connecting currency 1 with currency 2 in one direction represents the value of currency 1 expressed in currency 2, and vice versa. Then the weighted adjacency matrix of the network is the matrix of foreign exchange rates,

$$\mathbf{S}(t) = \begin{pmatrix} 0 & S_{1,2}(t) & \cdots & S_{1,K}(t) \\ S_{2,1}(t) & 0 & \cdots & S_{2,K}(t) \\ \vdots & \vdots & \ddots & \vdots \\ S_{K,1}(t) & S_{K,2}(t) & \cdots & 0 \end{pmatrix}, \quad (5.4)$$

where the relationship $S_{i,j}(t) = 1/S_{j,i}(t)$ holds. If $S_{i,j}(t + \Delta t) > S_{i,j}(t)$, then the base currency i appreciates against the quote currency j , and vice versa. Figure 5.1 shows the foreign exchange network for the case $K = 3$ with the currencies euro (EUR), US dollar (USD), and Japanese yen (JPY). At any given time, one of the

three currencies appreciates with respect to the other two, one of the three currencies appreciates with respect to one and depreciates with respect to one, and one currency depreciates with respect to the other two. Modifying the network illustrated in the left panel of Figure 5.1 such that a link pointing from node i to node j indicates the appreciation of currency i , we get a simple binary directed network. Let's assume a time step from t to $t + \Delta t$ in which the euro appreciates against the US dollar and the Japanese yen, and the US dollar appreciates against the yen. The resulting directed network is shown in the right panel of Figure 5.1.

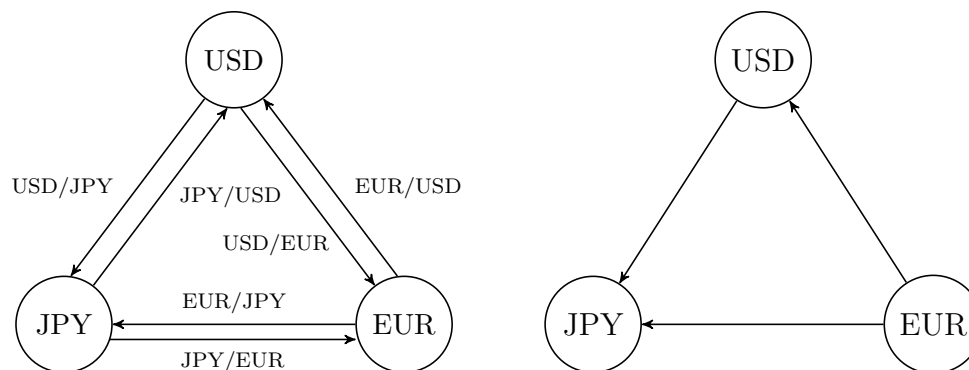


Figure 5.1: Left: foreign exchange market consisting of $K = 3$ currencies as a weighted directed network in which the links are the returns of the pairs. Right: binary directed network indicating the appreciation and depreciations for an arbitrary time step. In the example shown, the euro appreciates against the US dollar and the yen, and the US dollar appreciates against the yen only.

By construction, the symbolic performance is independent of base currency since the ranking is invariant. This helps us in two ways. For one thing, we will be able to include the base currency in our analysis and do not have to forgo its information. This is one critical limitation of typical correlation analyses. For another thing, since the time series $\zeta_i(t)$ is identical for any base currency, we are able to base our analysis on the currency for which data are most readily available.

While correlation analyses in the classical sense always hinge on the choice of pairs

or a base currency, the symbolic performance allows us to observe the comovement of currencies against the background of the entire market. At any given time, we can compare the positions of two currencies, albeit in broader terms than the specific returns. Naturally the symbolic performance time series does not offer itself to be investigated using Pearson's correlation coefficient for a couple of reasons. One of them is that two currencies can never have the same value at any given time since we are strictly ranking returns, that is $\zeta_i(t) \neq \zeta_j(t)$ if $i \neq j$. An analysis with Pearson's correlation coefficient would therefore be underestimating the correlation. This problem with ordinal data was pointed out, for example, by Aitchison [1982].

Instead we consider the empirical joint distribution of a pair of two currencies. At any given time, a pair of two currencies can be in a total of $K(K - 1)$ states, where each state is a pair of two ranks. For our data set, this amounts to 132 possible states. The likelihood of each combination to occur is dominated by two factors:

- (i) The marginal distributions of each currency set the boundaries for the shape of the joint distribution. If both currencies have neither very high or very low rank, then the states that describe the appearance of both currencies in one of those ranks at the same time is very unlikely. This is a simple consequence of Bayes' theorem.
- (ii) Correlations and anti-correlations determine the shape of the joint distributions within these boundaries.

It is entirely possible to consider two different pairs of currencies with identical marginal distributions, but entirely different joint distributions. This difference would then arise from comovements, for example.

In either case, we start with the null model of independent currencies, that is, currencies whose bivariate distribution is determined entirely by the marginals according to (i). In order to quantify the deviation of an empirical distribution from

the null model, we first need to determine the distributions that arises from the marginal distributions of the two currencies we wish to analyze. The most general joint distribution of the two currencies would be the one that assumes, firstly, no further knowledge than their marginal distributions and, secondly, that the two currencies cannot have the same symbolic performance at any given time. Using the method of the the principle of maximum entropy, which we explain further in Section 5.2.2, we are able to compute the null distribution. In brief, the principle of maximum entropy yields the distribution which under the given constraints maximizes our ignorance. Consequently, any deviation of an empirical joint distribution from the MaxEnt distribution encodes some kind of information.

We hypothesize that this information arises from the correlation structure of the foreign exchange market. In Section 5.2.3 we proceed to explain how we can measure the difference between two bivariate distributions in a way that is naturally linked to the amount of information encoded in the empirical distribution.

5.2.2 Maximum entropy estimation of probability distributions

For any given pair of currencies, we can extract the empirical bivariate distribution of their symbolic performances. This distribution indicates which positions the two currencies take in the market simultaneously. This offers insight into whether they take similar or opposite market positions over time.

Furthermore we compute the bivariate distribution that arises from the marginal distributions of the currencies if we assume independence. We then use distance measures in probability space to quantify the degree to which the empirical bivariate distribution of two currencies is similar or dissimilar to that of two independently moving currencies that otherwise share the same characteristics. This enables us to measure the degree of dependence between two currencies and their roles in the entire foreign exchange market. Both the estimate of the bivariate distribution of

independent currencies and the distance measure in probability space require us to introduce the concept of entropy.

Researchers and traders have been successful in applying entropy methods to finance, based on their connection to statistics and physics. Zhou et al. [2013] offers a review of a variety of uses. Some applications include portfolio selection and optimization as well as asset pricing and risk management (see Philippatos and Wilson [1972], White [1974], Avellaneda [1998], Gulko [2002]).

However, entropy has first been used in the context of statistical physics and thermodynamics. It can be interpreted as a measure of order in a physical system. More precisely, it measures the number of available microstates, that is, specific configurations of the system which the system can attain. Systems with a few but very well distinguishable microstates carry a low amount of entropy. We can say that the potential for disorder is rather low. Conversely, systems with a large number of indistinguishable microstates exhibit a large amount of entropy.

One of the fundamental laws of thermodynamics specifies that the entropy of a system increases over time unless we put work into the system to maintain order. Information theory takes a slightly different twist on the matter and defines that entropy quantifies the information content, e.g., of a message. In the physics picture, a system with few available microstates carries a large amount of information because we can easily identify the possible configurations. Likewise, a system with a large number of indistinguishable microstates contains little information because we cannot identify in which one of the many available microstates the system is. In other words, even in physics systems with low entropy carry a lot of information, and systems with high entropy exhibit a lower amount of information content. This leads to the following conclusion: When we observe a system that maintains a state that is not the state with the highest possible entropy, there is a mechanism which creates some

kind of order in the system. This may be the sun providing energy to earth such that complex, more ordered organisms could form, or this may be a writer assembling letters in a particular order to create readable content, i.e., provide information.

Using these insights, the principle of maximum entropy was formulated. Let us assume we are given some prior information which constrain the probability distribution. The most general distribution to describe the data is the distribution that exhibits the largest entropy while satisfying the constraints. Any distribution with lower entropy than that of the maximum entropy distribution would assume information that we do not have. Jaynes [1982] stated the principle of maximum entropy, or MaxEnt principle, and asserted: “When we make inferences based on incomplete information, we should draw them from that probability distribution that has the maximum entropy permitted by the information we do have.”

There exists a close connection between Bayes’ theorem and the principle of maximum entropy. In both cases, we use information to update our prior beliefs. If the information comes strictly in the form of data we gather, we use Bayes’ theorem to update our prior distribution. If, however, the information comes in the form of constraints, that is, testable information, the principle of maximum entropy is used instead. In fact, as Caticha et al. [2004] pointed out, if the prior knowledge we possess can be interpreted both as data and as a constraint, Bayes’ theorem and the principle of maximum entropy are equivalent.

5.2.3 Relative entropy and Jensen-Shannon divergence

Much in the same way as entropy allows us to quantify the information content of a distribution, it can be used to measure distances in distribution space. Consider an underlying true distribution $P(x)$. Then any approximation of this distribution by $Q(x)$ will result in a loss of information. The Kullback-Leibler divergence (KL divergence) was introduced in information theory to quantify this information loss.

It can be extended to a metric in the space of probability distributions, and it is defined as

$$D_{\text{KL}}(P||Q) = \sum_{i=1}^N P(x_i) \log \frac{P(x_i)}{Q(x_i)}, \quad (5.5)$$

is a generalization of the entropy given in Eq. (4.1) and also called “relative entropy.” It measures the information difference between the distributions $P(x)$ and $Q(x)$. In that light, entropy is a special case of relative entropy given uniform Q , modulo a constant.

Let us assume we found the proper MaxEnt distribution for the given marginals. In the process, we have assumed that nothing is known about the interaction and comovement of two currencies. However, we want to find out how much information is encoded in the empirical distribution. Then we have to compare this empirical distribution to the joint distribution that arises from the constraints of our problem. Here the KL divergence emerges as a natural choice.

The Kullback-Leibler divergence can also be understood as a measure of how far apart two distributions are in distribution space. However, it is not a metric. The KL divergence is, for example, not symmetric under the exchange $P \leftrightarrow Q$. However, it can be generalized to fulfill the requirements of a metric. Using the arithmetic mean $M = \frac{1}{2}(P + Q)$, one can define the so-called Jensen-Shannon divergence (JS divergence) as follows [Fuglede and Topsoe, 2004]:

$$D_{\text{JS}}(P||Q) = \frac{1}{2}D_{\text{KL}}(P||M) + \frac{1}{2}D_{\text{KL}}(Q||M). \quad (5.6)$$

Then $d(P, Q) = \sqrt{D_{\text{JS}}(P||Q)}$ is a metric in distribution space.

For our purposes the metric property is not as important. Instead it may be useful to rescale the Kullback-Leibler divergence to a range from 0 to 1, making it

comparable to an absolute correlation measure:

$$D_{\text{KL}} \rightarrow 1 - \exp[-D_{\text{KL}}] \equiv r_{\text{KL}}. \quad (5.7)$$

Note that this transformation reduces to the original expression in a first order approximation. This transformation finds applications in many fields of finance in order to rescale a variable to the interval from zero to one. One example we already saw in Chapter 3 is the calculation of the default probability of a credit instrument from the spread of the underlying credit default swap (compare Chan-Lau [2006]).

5.2.4 Generalized linear models

To explain a variable y that can take any value on the real axis in terms of one or more independent variables x_n , the standard tool of statistics is linear regression. Given a data set with $i = 1, \dots, M$ observations and $n = 1, \dots, N$ independent variables, one builds a model of the form

$$y^{(i)} = \beta_0 + \beta_1 x_1^{(i)} + \dots + \beta_N x_N^{(i)} + \epsilon^{(i)}, \quad (5.8)$$

where the superscript indicates an individual observation and the subscript denotes the independent variables and their coefficients. $\epsilon^{(i)}$ is an error term, accounting for noise and deviations between the model and observations not explained by the independent variables. In the case of ordinary least squares (OLS) regression, the coefficients β_0, \dots, β_N are determined such that the model minimizes the squared error of the observations $y^{(i)}$ and the model prediction $\hat{y}^{(i)}$. In matrix form, the model equation becomes

$$\mathbf{y} = \mathbf{X}\beta + \epsilon. \quad (5.9)$$

The solution for the coefficient vector β can be found analytically. This, however, involves matrix inversion; in many situations it is advantageous to consider numerical

methods like gradient descent that are computationally less expensive, for example when the matrix \mathbf{X} becomes large.

OLS regression assumes that the entries in the error term ϵ are distributed normally with zero mean and constant variance. Often, this assumption is violated to some degree, and in that case different regression approaches need to be considered. Additionally, the transformed Kullback-Leibler divergence r_{KL} is constrained to the interval from zero to one, per Eq. (5.7). Mitigating the shortcomings of OLS regression, we can turn to generalized linear models (GLM) which allow the dependent variable \mathbf{y} to be distributed according to a member of the exponential family, relaxing the normality assumption in OLS regression. Introduced by Nelder and Wedderburn [1972], generalized linear models use a link function to describe the relationship of the expected value of the data and the mean of a distribution function, $\eta = g(\mathbb{E}[y^{(i)}]) = \mathbf{X}\beta$. Such a generalized linear model then takes the following form:

$$\eta^{(i)} = \beta_0 + \beta_1 x_1^{(i)} + \cdots + \beta_N x_N^{(i)} + \epsilon^{(i)}, \quad (5.10)$$

in which the link function $\eta = g(\mathbb{E}[Y])$ is supplemented by a variance function $\text{Var}(Y) = \phi V(\mathbb{E}[Y])$, where ϕ is a constant. Possible choices for the link function include the identity, used to recover a normal distribution, or the logarithm, typically used in connection with a Poisson distribution, that is $\log(\mathbb{E}[Y]) = \mathbf{X}\beta$. The data we describe in Section 5.2.5 includes categorical data like currency classifications and suggests the use of a GLM based on the Poisson distribution.

5.2.5 Data

Foreign exchange data

We used the trading platform *Oanda* to download intraday foreign exchange data. We employed *Oanda's* open access API¹. As outlined in Section 5.2.1, the time series of one base currency is sufficient, and we could infer all information from the exchange rates of the other currencies with that base currency. We chose the euro as the base currency since it offers the largest and most complete dataset on *Oanda*. We downloaded data from January 2005 to May 2018 in ten minute intervals. Our dataset contained the following currencies, where in brackets we provide the three-letter code which we will often use to refer to the currencies: Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), euro (EUR), British pound (GBP), Japanese yen (JPY), Mexican peso (MXN), New Zealand dollar (NZD), Norwegian krona (NOK), Swedish krona (SEK), US dollar (USD), and South African rand (ZAR).

This gave us $K = 12$ distinct exchange rate time series $S_i(t)$, $i = 1, \dots, 12$, $t = 1, \dots, T$, where we included the euro time series consisting of only ones, $S_1(t) = S_{\text{EUR/EUR}}(t) \equiv 1$. Analogously, we had $K = 12$ distinct return time series $R_i(t)$, $i = 1, \dots, 12$, $t = 2, \dots, T$, where, again, we included the euro time series, now consisting of only zeros, $R_1(t) = R_{\text{EUR/EUR}}(t) \equiv 0$. If any exchange rate data were missing for a time point, we excluded the time point from our analysis. If two consecutive exchange rates did not correspond to a ten minute interval, we excluded the return corresponding between these two time points.

Macroeconomic data

Since foreign exchange rates are affected by macroeconomic indicators like interest rates, trade volume, commodity prices, stock index correlations and others, we down-

¹ <http://developer.oanda.com>

loaded the corresponding data. In the following we describe the data sources as well as the necessary data cleaning.

Trade volume. The Comtrade Database provided by the United Nations (UN) is a “repository of official international trade statistics”². It is compiled by the Trade Statistics Branch of the Statistics Division, Department of Economic and Social Affairs of the UN. While it contains a detailed breakdown of goods and services traded between more than 170 nations, we downloaded only the gross total of annually imported and exported commodities. We selected the countries according to the twelve currencies in our data set: Australia, Canada, Switzerland, the countries of the eurozone, the United Kingdom, Japan, Mexico, New Zealand, Norway, Sweden, the United States and South Africa. Each entity reports their trade statistics with their trading partners to the UN which compiles the data into a standard format as outlined in the Comtrade user manual³: “All commodity values are converted from national currency into US dollars using exchange rates supplied by the reporter countries, or derived from monthly market rates and volume of trade.” Using the user interface provided in UN Comtrade database, we downloaded the trade volume reported by each country in our data set with annual frequency from 2012 to 2017. Table 5.1 shows the result for a subset of the countries; entries are in billions of US dollars.

Treasury bond yields. A treasury bond is a bond issued by the United States Department of the Treasury. In order to finance the operations of government, the Treasury auctions debt securities like the treasury bond (T-bond). While a variety of such debt instruments exists, the 10-year T-bond stands out in importance. It is used to price mortgage rates, and it is very actively traded in secondary markets. When the T-bond is issued, the coupon payments through which the interest on the bond

² <https://comtrade.un.org>

³ <https://unstats.un.org/unsd/tradekb/Knowledgebase/50075/What-is-UN-Comtrade>

	Imp.	AUD	CAD	CHF	EUR	GBP
Exp.						
AUD		—	1.414	0.8	6.044	7.3
CAD		1.5	—	1.0	16.5	13.1
CHF		2.4	3.5	—	104.9	29.6
EUR		26.7	31.6	126.4	—	285.6
GBP		5.3	6.2	19.4	163.0	—

Table 5.1: Trade volumes in billion US dollar between the different currency regions Australia (AUD), Canada (CAD), Switzerland (CHF), the eurozone (EUR) and the United Kingdom (GBP). The rows indicate the exporter and the columns indicate the importer.

is paid out, is fixed. In the secondary market, traders react to changes in economic policy like interest rate changes and trade the bond at a discount or at a premium. If the Federal Reserve (Fed) changes its target federal funds rate, colloquially just called “interest rate”, this has an impact on T-bond prices. If the Fed increases its target rate, for example, an older T-bond become less desirable to own since it was priced with a lower target rate in mind. Consequently, such a T-bond is sold at a discount in the secondary market. As a result, the yield of the T-bond increases according to

$$\text{Yield} = \frac{\text{Annual Dollar Interest Paid}}{\text{Market Price}} \times 100\%. \quad (5.11)$$

Since 10-year Treasury bonds are so actively traded and have an equivalent in most economies, we used them as a proxy for monetary policy such as changes in interest rates. Using the provided API, we downloaded the data from the database of the Federal Reserve Bank of St. Louis (FRED) which uses the Organization for Economic Co-operation and Development (OECD) as its source [OECD, 2018]. We downloaded the quarterly yield of the 10-year governments bond from 2012 to 2017. The appendix (C.1) offers an overview of the respective government debt instruments for the countries in our data.

We used the treasury bond yields as a proxy for the nominal interest rates to consider the so-called international Fisher effect. Observing, for example, the euro and the US dollar and their respective nominal interest rates i_e and $i_{\$}$, the international Fisher effect anticipates a change $E(e)$ in the exchange rate of

$$E(e) = \frac{1 + i_e}{1 + i_{\$}} - 1. \quad (5.12)$$

For typical interest rates or bond yields in our case in the range of -2% to 10% , this expression is well approximated by $E(e) \approx \log\left(\frac{1+i_e}{1+i_{\$}}\right)$.

Stock indices. A major indicator of economic health is the stock market. We approximated the stock market performance of each country or economic region by considering the iShares exchange-traded funds (ETF) for the countries in our data set. For the eurozone, for example, we used the iShares ETF which tracks large and mid-sized publicly listed companies from countries using the euro as their official currency. A detailed list of the ETFs used in this thesis is provided in the appendix, C.1. While tracking the performance of an index (or any other basket of assets), an ETF trades like a security itself. In particular, we could choose ETFs that are traded on the same stock exchange. We downloaded the data from Morningstar in daily frequency for the years 2012 through 2017. Using Pearson's correlation coefficient for two time series $X = x_1, \dots, x_n, Y = y_1, \dots, y_n$,

$$r = \frac{\sum_{i=1}^n (x_i - \langle x \rangle)(y_i - \langle y \rangle)}{\sqrt{\sum_{i=1}^n (x_i - \langle x \rangle)^2} \sqrt{\sum_{i=1}^n (y_i - \langle y \rangle)^2}}, \quad (5.13)$$

with mean $\langle x \rangle$ and $\langle y \rangle$, respectively, we calculated the quarterly and the annual correlation coefficient of the ETFs corresponding to the stock markets in our data set.

Classification of currencies. Our regression analysis included dummy variables to study whether the classification of a currency as a reserve currency or commodity currency had an effect on its behavior in the foreign exchange network. While these classifications are more qualitative in nature, they have been well established for practitioners as well as in the literature. Compare, for example Chen et al. [2010] or Hossfeld and MacDonald [2015]. We again refer to the appendix (C.1) for a list of the classifications we have used.

5.3 Results

5.3.1 Annual intervals

We observed significant deviations from the null model distribution which were persistent for many years. These deviations were particularly pronounced for some currency pairs such as the euro and the Swiss franc or the Australian dollar and the New Zealand dollar. Other currency pairs indicated codependence in some years while not in others, such as the US dollar and the Japanese yen, or the US dollar and the Mexican peso. Some currencies exhibited overall larger deviations from the null model across a broad range of currencies, like the euro or the US dollar, while others like the British pound or the Canadian dollar appeared to move independently from other currencies. Figure 5.2 shows the Kullback-Leibler divergence of the bivariate distributions for all possible combinations of two currencies and the null model of independence for the years 2012 through 2017.

Note that the matrices in Figure 5.2 are symmetric. Furthermore, the diagonal elements in each matrix are not defined since the null model of independence does not exist. All other entries, however, are non-zero: due to the finiteness of the data, the Kullback-Leibler divergence like any entropic measure yields a finite result, bounded by the system being studied.




Figure 5.2: Kullback-Leibler divergences between the empirical bivariate distributions and the null model of independence for the years 2012 through 2017.

The finiteness of the data also induces some noise which results in an overestimation of KL divergence. Steuer et al. [2002] provides an estimate of the finite size effect for the mutual information $I(X;Y)$ which is closely related to the Kullback-Leibler divergence:

$$\Delta I(X;Y) = \frac{M_{x,y} - M_x - M_y + 1}{2N} \quad (5.14)$$

where $M_{x,y}$, M_x , M_y denote the number of discrete states with nonzero probability. In our case, $M_x = M_y = M = 12$, the number of currencies, and $M_{x,y} = M(M-1) = 132$ to account for the diagonal. N is the number of observations. We used this estimate for our measure of the Kullback-Leibler divergence and calculated the systematic error through finite size effects as $\Delta D_{KL}(P||Q) = 109/(2N)$. This yields the error estimates shown in Table 5.2.

time	N	ΔD_{KL}
2012	23481	0.00232
2013	34135	0.00160
2014	33043	0.00165
2015	33287	0.00164
2016	35025	0.00156
2017	34762	0.00157

Table 5.2: Finite size error estimates for the KL divergence.

While some values in Figure 5.2 were small, each calculated value for the Kullback-Leibler divergence exceeded the threshold from the finite size effect calculated in Table 5.2. Nevertheless we observed that the majority of currency pairs, on average, only differed very slightly from the behavior we would expect of two independently moving currencies. Yet we observed some significant outliers.

Especially the euro and the US dollar showed overall higher values. Some specific currency pairs, such as the NOK and SEK or the AUD and NZD, also tended to show

larger comovement than two currencies moving independently of each other. Overall, the euro showed the highest average value for the JS divergence, closely followed by the US dollar. This confirms that these currencies are indeed major entities in the foreign exchange market in the sense that their course movements strongly affect other currencies. Additionally, the difference between the empirical joint distribution of the symbolic performances of euro and US dollar was significant. Further data analysis indicated that euro and US dollar, however, tended to stay on opposite sides of the market. In other words, if one currency appreciated with respect to a majority of the remaining currencies, the other currency was more likely to depreciate overall than just by chance. This hints that these two currencies served as the pole markers of the foreign exchange market for the time period we studied.

It is not surprising that the South African rand, as the smallest currency of our set and the economically least intertwined, showed a very low KL divergence across most pairs. This indicates that it moved like an independent and, on average, uncorrelated currency within the market. The South African rand showed only some dependence on the euro and US dollar, which is in line with the assumption that these two currencies dominated the overall market development.

More striking, however, was the apparent independence of the British pound. Its average KL divergence was even smaller than that of the South African rand, and it was even less affected by EUR and USD than the South African rand.

5.3.2 Quarterly intervals

For a more detailed analysis of the development over time, we extended the observation period and study the KL divergence of one currency with all other currencies at a finer time-resolution, considering the data quarter by quarter. Given that one time interval is 10 minutes long, one quarter contained approximately 8000 to 9000 observations. Only in the year 2012, some quarters contained less data points, re-

flecting the poorer data quality during that time. Corresponding to the analysis in Table 5.2, this corresponded to a lower bound to our ability to resolve the KL divergence to 0.6×10^{-2} for most quarters and 1.1×10^{-2} for the affected quarters in 2012. Again, we determined that the KL divergence we observed exceeded this lower bound for every currency pair.

Average comovements over time

The results are presented in Figure 5.3 and show that the KL divergence was a fairly constant quantity for most currencies. Notable exceptions included the Swiss franc (CHF), euro (EUR) and the US dollar (USD). The pegging and unpegging of the Swiss franc to the euro altered the relationship of the foreign exchange market in that the pegging increased the information contained in the time series on one hand and altered possible diversification strategies on the other. The first peak in the panel for the Swiss franc corresponds to the time immediately after the peg to the euro, and the second peak corresponds to the unpegging. A result of the unpegging was that the Swiss franc moved rather independently from the rest of the market, as indicated by a very low KL divergence value. Similar results, albeit weaker, were found for the euro. Since it has been exhibiting large comovements with other European currencies, like the Swedish and Norwegian krona, the effect of the peg and unpeg was not as pronounced for the euro. The third peak for the Swiss franc and the euro also appears in the graphs for most other currencies, indicating a global event in the first quarter of 2016. It was observed for the Australian and Canadian dollar, the Japanese yen, the Mexican peso, and the New Zealand dollar. Potentially also the Swedish krona and the South African rand exhibited a small spike in this quarter. Notably, the increase in average KL divergence for these currencies could not be attributed to the behavior of two or three currencies only; instead a majority of the currencies in the market exhibited a larger comovement. Interestingly, the US dollar and the Norwegian krona




Figure 5-3: Kullback-Leibler divergence for currencies resolved on a quarterly basis. The value shown in this figure is the average KL divergence of a currency with all other currencies in the data set. As usual, a larger KL divergence indicates that the time series of the currency shown with other currencies contained more information than if we assumed independence.

exhibited a decrease in KL divergence at the same time, both quarter over quarter and year over year for the first quarter of 2016. Future research into the causes is necessary, however, one event comes to mind: in January 2016, the sanctions on Iran were lifted, changing the landscape of the oil market – incidentally, the two currencies with decrease in KL divergence belong to two major players in this market. Another observation in Figure 5-3 supports this hypothesis: in early 2015, both the US dollar and the Norwegian krona exhibited a spike in their KL divergence. This coincided with large swings in oil price: in Q4 of 2014 the oil price dropped 40 percent, and in Q1 of 2015 it dropped another 10 percent, only to recover by 25 percent in the following quarter. Our hypothesis is further motivated by the findings in Chapter 2, where we found that currencies of oil-exporting nations and of the US were positively correlated for many years before 2012 already.

GLM with macroeconomic data

In Section 5.2.5 we described macroeconomic data which we identified as potential predictors for the KL divergence of a currency pair. Using a generalized linear model, we studied the impact of stock index covariance, bilateral trade volume, bilateral trade imbalances, treasury bonds yields, share of the currency pair in daily FX market turnover, and geographic proximity. Additionally the model included dummy variables for currency classifications, that is, indicators whether the two currencies were both considered as reserve currencies, commodity currencies, a mix of both, or none of the above. This model based on classification is abbreviated “class”. We then considered a second model in which we included dummy variables for each currency to capture idiosyncratic characteristic properties which may not be explained otherwise instead of classifications. This model based on idiosyncratic currency characteristics is abbreviated “idio”. Features were then selected based on their impact on the root-mean-square error which we computed after splitting the data in a training set (80%)

and a test set (20%). If a feature neither reduced the RMSE nor exhibited statistical significance at the 95 percent level, we excluded the feature. As a result we were left with the following independent variables: stock market covariance, log of mean trade volume, log of the ratio in Eq. (5.12) using the treasury bond yields and a dummy variable for the currency cap between the Swiss franc and the euro until early 2015.

$$\begin{aligned}
\eta^{(class)} = & \beta_0 + X_{IndexCov} \beta_1 + X_{\log(MeanTradeVol)} \beta_2 + X_{\log(Fisher)} \beta_3 \\
& + D_{Cap} \beta_4 + D_{CommComm} \beta_5 + D_{CommRes} \beta_6 \\
& + D_{ResRes} \beta_7 + D_{CurrOther} \beta_8.
\end{aligned} \tag{5.15}$$

Similarly, the second model was described by

$$\begin{aligned}
\eta^{(idio)} = & \tilde{\beta}_0 + X_{IndexCov} \tilde{\beta}_1 + X_{\log(MeanTradeVol)} \tilde{\beta}_2 + X_{\log(Fisher)} \tilde{\beta}_3 + D_{Cap} \tilde{\beta}_4 \\
& + D_{AUD} \tilde{\beta}_5 + D_{CAD} \tilde{\beta}_6 + D_{CHF} \tilde{\beta}_7 + D_{EUR} \tilde{\beta}_8 \\
& + D_{GBP} \tilde{\beta}_9 + D_{JPY} \tilde{\beta}_{10} + D_{MXN} \tilde{\beta}_{11} + D_{NOK} \tilde{\beta}_{12} \\
& + D_{NZD} \tilde{\beta}_{13} + D_{SEK} \tilde{\beta}_{14} + D_{USD} \tilde{\beta}_{15} + D_{ZAR} \tilde{\beta}_{16}
\end{aligned} \tag{5.16}$$

In both equations, X indicates a real-valued variable, while D denotes a dummy variable for categorical data. We used β and $\tilde{\beta}$ for the regression coefficients in the two different models, respectively, to avoid confusion. The results for the generalized linear model based the classification predictor in Eq. (5.15) are shown in Table 5.3 where we report the coefficients and their significance levels. We report results for the generalized linear model in Eq. (5.16) in Table 5.4.

In the model based on currency classification we observed that the trade volume, the currency cap indicator for the Swiss franc and the euro, and currency classifications were significant predictors of the currency comovement as measured through r_{KL} . Positive coefficients contribute to higher comovement, that is, a larger deviation from the null hypothesis of independence. Large bilateral trade volume and the cap

Model:	GLM	AIC:	366.3	
Link Function:	log	BIC:	-11589.4	
Dependent Variable:	r_{KL}	Log-Likelihood:	-175.2	
Method:	IRLS	LL-Null:	-183.5	
No. Observations:	1584	Deviance:	22.2	
Df Model:	7	Pearson χ^2 :	27.2	
Df Residuals:	1576	Scale:	1.0	
	Coef.	Std. Err.	z	$P > z $
β_0	-6.175 ***	1.672	-3.694	0.000
IndexCov	2516.	1877.	1.341	0.180
log(MeanTradeVol)	0.413 *	0.208	1.985	0.047
log(Fisher)	-11.57	8.759	-1.321	0.187
Cap	1.755 *	0.688	2.550	0.011
CommComm	-1.557 **	0.591	-2.636	0.008
CommRes	-1.659 **	0.548	-3.008	0.003
ResRes	-1.812 **	0.694	-2.611	0.009
CurrOther	-1.157 **	0.438	-2.639	0.008

Table 5.3: Result: Generalized linear model for currency comovements r_{KL} based on macroeconomic data and currency classifications.

contributed to higher comovement. The more two currency regions traded bilaterally, measured as the logarithm of the sum of the imports and exports to one another, the bigger r_{KL} . Since the cap enforced by the Swiss National Bank against the euro linked the value of the currencies closer together, it is not surprising that we observed it to have a positive effect on their comovement. The dummy variables for classification all have negative coefficients, albeit to varying degrees. The comovement of two reserve currencies, or the amount of information in their link in the foreign exchange network, was most strongly negatively affected. This is explicable in part due to the influence of the Swiss franc-euro cap on the one hand and the fact that reserve currencies in our data set belong to countries and regions with a large trade volume, e.g. the EU and the United States on the other. The stock index covariance appeared to contribute positively to the comovement of currencies, however, the result was not significant at the 95% level. The impact of bond yields was determined to be negative, albeit not at the significance level we considered.

Table 5.4 describes the model focusing on idiosyncratic properties of currencies as

Model:	GLM	AIC:	374.3	
Link Function:	log	BIC:	-11538.4	
Dependent Variable:	r_{KL}	Log-Likelihood:	-171.2	
Method:	IRLS	LL-Null:	-183.5	
No. Observations:	1584	Deviance:	14.1	
Df Model:	15	Pearson χ^2 :	17.2	
Df Residuals:	1568	Scale:	1.0	
	Coef.	Std. Err.	z	$P > z $
β_0	-7.952 ***	2.287	-3.477	0.001
IndexCov	2174.	2075.	1.048	0.295
log(MeanTradeVol)	0.735 *	0.301	2.373	0.018
log(Fisher)	-18.65	13.60	-1.371	0.170
Cap	0.936	0.785	1.193	0.233
AUD	-1.114 *	0.511	-2.182	0.029
CAD	-1.664 **	0.567	-2.935	0.003
CHF	-1.073 *	0.502	-2.135	0.033
EUR	-1.562 *	0.746	-2.095	0.036
GBP	-2.630 ***	0.742	-3.545	0.000
JPY	-1.653 **	0.616	-2.684	0.007
MXN	-0.972	0.567	-1.716	0.086
NOK	-1.124 *	0.493	-2.376	0.018
NZD	-0.603	0.408	-1.480	0.139
SEK	-0.989 *	0.497	-2.071	0.038
USD	-1.868 *	0.769	-2.430	0.015
ZAR	-0.653	0.668	-0.978	0.328

Table 5.4: Result: Generalized linear model for currency comovements r_{KL} based on macroeconomic data and conditioned on currencies.

opposed to qualitative currency classifications like in Table 5.3. Of the macroeconomic indicators, only the trade emerged as significant at the 95% level. The dummy variables for the currencies are significant, except for the Mexican peso, the New Zealand dollar, and the South African rand, which coincidentally are the currencies with the smallest daily turnover in our data set. The coefficients for stock index covariance and Fisher effect appeared with the same sign as for the classification model.

5.4 Discussion

We presented a network interpretation of the symbolic performance, which can be calculated as the degree imbalance in a network in which the nodes are currencies and directed links indicate appreciation and depreciation relationships. Using tools

from information theory, we introduced a new measure for comovement of currencies. For this new measure we calculate the Kullback-Leibler divergence of distribution of degree imbalances of currencies. A large value indicates a strong deviation of the null hypothesis that the degree imbalance of currencies, and thus their appreciations and depreciations, are independent and due to noise.

Our results showed the major roles of the euro, buoyed by the interaction with other European currencies in our data set. As expected, the smallest currency in the data set, the South African rand, shared little information with other currencies. More surprisingly, the British pound appreciated and depreciated practically independently from all other currencies, including the euro and the US dollar, the currencies of two major trading partners. Interestingly, this did not change after the Brexit vote in 2016. Our analysis furthermore recovered the effect of the Swiss franc peg to the euro and its effect on the market. It also showed an interesting emergence of comovement of many currencies in our data set in early 2016, and we hypothesize that this can be linked to the lifting of sanctions on Iran. Our hypothesis is supported by the fact that the comovement excluded the US dollar and the Norwegian krona, two important currencies in the oil market.

Building a generalized linear model, we studied macroeconomic indicators and their influence on the comovement of currencies. We found a significant positive impact from the bilateral trading volume. The stock index covariance as well as the international Fisher effect had positive and negative coefficients, respectively; however, they were not statistically significant at the 95% level. While currency specific characteristics had a significant impact on their comovement, we showed that characterizing them into reserve currencies, commodity currencies, and other was a sufficient substitute.

For future research, the investigation of lead-lag relationships in the comovement

is an obvious next step. Identifying the empirical and theoretical bivariate distributions of the symbolic performance, our new comovement measure allows to study which pairs currencies are leading indicators. While this measure was undirected for synchronous appreciations and depreciations, lead-lag relationships are directed, allowing us to also consider directed macroeconomic indicators, for example on bilateral trade or differences in employment or GDP.

Chapter 6

Political and Economic Effects on Currency Clustering Dynamics

We propose a new measure named the *symbolic performance* to better understand the structure of foreign exchange markets. Instead of considering currency pairs, we isolate a quantity that describes each currency's position in the market, independent of a base currency. We apply the k -means++ clustering algorithm to analyze how the roles of currencies change over time, from reference status or minimal appreciations and depreciations with respect to other currencies to large appreciations and depreciations. We show how different central bank interventions and economic and political developments, such as the cap on the Swiss franc to the euro enforced by the Swiss National Bank or the Brexit vote, affect the position of a currency in the global foreign exchange market.

6.1 Introduction

Similar to other financial markets, exchange rates between different currencies are determined by the laws of supply and demand in the forex market. Additionally, market participants (financial institutions, traders, and investors) consider macroeconomic factors such as interest rates and inflation to assess the value of a currency. Central banks may participate in the forex market as well, when pursuing their fiscal and monetary policy goals. The degree of central bank interventions in the market determines the regime in which a currency trades. Some central banks peg their currency to another currency, using their assets in the market to accomplish a fixed exchange rate. If a central bank does not intervene, its currency is considered free-floating, meaning that the exchange rate is mostly determined by market forces. Some central banks allow their currency to float freely within a certain range, in a so-called managed float regime.

The value of any given currency is expressed with respect to the rest of the market through pairwise exchange rates. Currency quotes thus exhibit an important difference to equity, fixed income or commodity markets where prices of these assets are quoted in one specific currency. A consequence of this is that the appreciation of one currency implies the depreciation of the currency against which it is traded. This structural property of the market in combination with the strong influence by macroeconomic fundamentals and central banks as market participants have led to specific characteristic behaviors for currencies.

Qualitatively, we can distinguish between hard and soft currencies; hard currencies are considered a store of value due to their stability even in adverse global economic conditions, and soft currencies are more volatile, for example due to deteriorating economic conditions in respective countries. Examples of hard currencies are typically the US dollar, the euro or the Japanese yen [Hossfeld and MacDonald, 2015]. The

Venezuelan bolivar, on the other hand, and its continued devaluation over the last decade is an example of a soft currency.

Alternatively, we can distinguish between reserve currencies, funding currencies, and commodity currencies. Reserve currencies are currencies which central banks typically hold as foreign exchange reserves, for which they prefer hard currencies [Habib and Stracca, 2012]. Most of the world's currency reserves are held in US dollar or euro, and to a lesser extent in British pound and Japanese yen. As a result, until the recent inclusion of the Chinese yuan, these four currencies also comprised the currency basket used for accounting purposes at the International Monetary Fund (IMF)¹. Funding currencies are currencies which can be borrowed at low interest rates. Historically, the Swiss franc and the Japanese yen have been used to fund purchases of currencies with higher interest rates, for example. Commodity currencies are currencies of countries whose economic output strongly depends on the price of one or more commodities. Examples of commodity currencies include the Norwegian krona due to Norway's significant oil exports or the Australian dollar due to Australia's significant exports of metals and coal.

It becomes clear that the exchange rate of two currencies is therefore determined by their own idiosyncratic behaviors and economic factors as well as by their relationships with other currencies in the market. The structure and the characteristics of the foreign exchange market pose an extraordinary challenge to traders and researchers. The choice of the transaction or base currency influences the results, which has been observed in the literature in several works and in many different contexts. Papell and Theodoridis [2001] analyze efforts to calculate the purchasing power parity (PPP) via panel tests. They show that the choice of base currency influences the outcomes of PPP-tests. Recognizing that the choice of base currency affects the correlation

¹ The accounting currency of the IMF are the so-called Special Drawing Rights, and as of 2018 their value is determined through a weighted basket of the US dollar, euro, Japanese yen, British pound and Chinese yuan.

between different currencies, Hovanov et al. [2004] create a currency index to determine the value of an individual currency within the global forex market, independent of base currency. Many network approaches to understanding the foreign exchange markets rely on correlation measures. Górski et al. [2008], Kwapien et al. [2009] investigate the effect of base currency on interrelationships between currencies when studied through a network lens. They provide evidence that the topology of the foreign exchange network and the structure of its minimum spanning tree is different for different base currencies.

In this chapter, we present a novel approach to address these issues. Instead of considering currency pairs, we treat each currency i in a market of K currencies as an individual entity with assigned symbolic performance ζ_i . This approach introduces a measure independent of base currency to investigate the hierarchy and the dynamics of the FX market. We aim to encode the relationship of each currency with the remaining currencies in the market into one quantity. Instead of considering all currency pairs for a given currency we compress information of its pairwise exchange rates into one quantity for each currency. We do this by measuring the relative performance of a currency in relation to other currencies, and we call this measure the symbolic performance.

Using the symbolic performance, we investigate how a currency's role evolves within the market in the wake of changing economic conditions. As exchange rates are affected by monetary policy, we especially consider central banks' currency interventions that may be conducted directly, for example, if a central bank purchases or sells the domestic currency. In more extreme cases central banks may introduce a cap or a peg of its currency, backing this policy by currency transactions. Data for interventions publicly accessible for the Swiss franc, the Mexican peso, the Singapore dollar and Japanese yen among the currencies considered in this chapter.

In the literature, the effects of central bank interventions (CBIs) on foreign exchange rates have been studied by various techniques, particularly focusing on volatility of exchange rates. These techniques include GARCH type models [Almekinders and Eijffinger, 1996, Baillie and Osterberg, 1997a,b, Dominguez, 1998, Beine et al., 2002], implied volatility estimation of currency options [Bonser-Neal and Tanner, 1996, Dominguez, 1998], regime-switching analysis of mean and variance of exchange rates [Beine et al., 2003], realized volatility estimation [Dominguez, 2006, Beine et al., 2009, ru Cheng et al., 2013], time series study of news reports [Fatum and Hutchison, 2002], and event study of CBIs [Fatum and Hutchison, 2002, 2003, Fatum, 2008].

Most of the works quoted consider only three currencies – the German deutschmark (euro), Japanese yen and US dollar – and study the respective CBIs of the German Bundesbank (European Central Bank), Bank of Japan and Federal Reserve System on their domestic currency. Our approach, however, explicitly incorporates information of currencies whose central banks did not intervene in the time period being analyzed. This methodological distinction allows us to examine not only effects of CBIs on the domestic currency, but on the currency embedded in the FX market.

The rest of this chapter is structured as follows. We present our foreign exchange data set in Section 6.2. We lay out the framework and the methodology of obtaining the symbolic performance measure in Section 6.3. In Section 6.4, we study the statistics of symbolic performances for the entire time horizon as well as for specific subintervals. In particular, we present the results of our cluster analysis revealing the temporal evolution of the symbolic performances and identifying different roles currencies play within the FX market. We link changes in roles and behaviors of currencies to central bank interventions as well as economic shocks. We offer our conclusion in Section 6.5.

6.2 Data

We downloaded currency exchange rate time series from *Oanda* via an open access API². Our data comprised the following currencies, which are all quoted in terms of euro (EUR) and listed in alphabetical order of their ISO currency code: Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), British pound (GBP), Hong Kong dollar (HKD), Japanese yen (JPY), Mexican peso (MXN), Norwegian krone (NOK), New Zealand dollar (NZD), Swedish krona (SEK), Singapore dollar (SGD), US dollar (USD), and South African rand (ZAR). These currencies cover 14 of the 20 globally most traded currencies and each of them accounted for a share of at least 1% of average daily turnover in April 2016 according to BIS Monetary and Economic Department [2016]. Since *Oanda* does not provide sufficiently complete data for the remaining six currencies, we omitted them in our analysis. This selection of currencies yielded 14 distinct exchange rate time series after we included a dummy euro time series which consisted of only ones. The data set spans from January 2, 2005, to May 9, 2017, a period of more than 12 years. The start date was chosen such that we could observe a time window as long as possible while at the same time maintaining the quality of the data. There was no trading on certain holidays as well as from Friday night to Sunday night. We downloaded exchange rates in 10-minute time intervals. Note that our results are independent of the time resolution of the underlying data set. When studying currency dynamics, however, a large amount of data is a prerequisite, and therefore we selected the 10-minute time intervals.

We treated missing data as follows: If no euro pair had been traded at all in a given 10-minute interval, the *Oanda* platform did not report any exchange rates for this time point. We excluded these time points from our analysis. If one particular pair was not traded in a given 10-minute interval, the *Oanda* platform recorded a

² <http://developer.oanda.com>

return of zero. If more than two euro pairs had not been traded, we discarded this time point. If at most two euro pairs had not been traded, we imputed data by drawing values from a normal distribution. We estimated mean and variance of that normal distribution as sample mean and sample variance of empirical log-returns of the missing currency within the past 24 hours. If data were missing and we could not estimate sample mean and variance over the last 24 hours, we excluded the time point from our analysis.

Overall our data set contained 444 360 intervals, and after we accounted for missing data according to the aforementioned procedure, approximately 413 000 intervals per currency remained. Since we analyzed 14 currencies, this yielded about 5.8 million exchange rates which we use for our study.

While we used the intraday foreign exchange data, we identified the roles which currencies play with a resolution of two weeks. In visualizations and in alignment with other data sets, these non-overlapping windows of two weeks were represented by their midpoint.

We gathered information on central bank interventions from the websites of central banks who publish their activity³ as well as news reports from sources like Reuters⁴. This information is of daily frequency, allowing us to specify on which day an intervention had taken place and to associate it with a two-week window describing the currency roles.

³ E.g. <http://www.banxico.org.mx/sistema-financiero/estadisticas/mercado-cambiario/banco-mexico-s-foreign-exchan.html>

⁴ E.g. <http://www.reuters.com/article/global-forex-boj-idUSL8N15Q4BY>

6.3 Methodology

6.3.1 Exchange rates and returns

Let us consider a financial market consisting of K assets. In the foreign exchange market, each asset $i = 1, \dots, K$, is a currency and linked to the remaining $K - 1$ currencies through their exchange rate $S_{i,j}(t)$ at time t , where $j = 1, \dots, K, j \neq i$. Obviously $S_{i,i} \equiv 1$. In matrix form, this becomes

$$\mathbf{S}(t) = \begin{pmatrix} 1 & S_{1,2}(t) & \cdots & S_{1,K}(t) \\ S_{2,1}(t) & 1 & \cdots & S_{2,K}(t) \\ \vdots & \vdots & \ddots & \vdots \\ S_{K,1}(t) & S_{K,2}(t) & \cdots & 1 \end{pmatrix}. \quad (6.1)$$

Each row of this matrix describes exchange rates given a base currency, while each column of the matrix describes exchange rates given a counter currency. In the absence of a bid-ask spread, the reciprocal relationship implies $S_{i,j}(t) = 1/S_{j,i}(t)$.

Using the structure of the FX market, we can construct this matrix from the exchange rates of just one currency with all other currencies at time t . Under the assumption of no arbitrage, the exchange rates $S_{i,j}(t)$ and $S_{i,k}(t)$ imply the exchange rate $S_{j,k}(t)$ via $S_{j,k} = S_{j,i}(t) S_{i,k}(t)$. It then suffices to consider the exchange rate vector $\mathbf{S}_i(t)$ at time t which comprises of the exchange rates of currency i , acting as a base currency, with all other currencies:

$$\mathbf{S}_i(t) = (S_{i,1}(t), S_{i,2}(t), \dots, S_{i,K}(t)). \quad (6.2)$$

It carries complete information of the system, in that the outer product of the vector

and its inverse yields the exchange rate matrix in equation (6.1),

$$\mathbf{S}(t) = \begin{pmatrix} S_{1,i}(t) \\ S_{2,i}(t) \\ \vdots \\ S_{K,i}(t) \end{pmatrix} \cdot (S_{i,1}(t), S_{i,2}(t), \dots, S_{i,K}(t)). \quad (6.3)$$

In this study, we calculate logarithmic returns between consecutive exchange rates in a time interval $\Delta t = 10$ minutes:

$$R_{i,j}(t) = \log S_{i,j}(t) - \log S_{i,j}(t - \Delta t). \quad (6.4)$$

Analogous to equation (6.1), we can write the exchange rate returns in matrix form:

$$\mathbf{R}(t) = \begin{pmatrix} 0 & R_{1,2}(t) & \cdots & R_{1,K}(t) \\ R_{2,1}(t) & 0 & \cdots & R_{2,K}(t) \\ \vdots & \vdots & \ddots & \vdots \\ R_{K,1}(t) & R_{K,2}(t) & \cdots & 0 \end{pmatrix}. \quad (6.5)$$

Here the reciprocal relationship becomes $R_{i,j}(t) = -R_{j,i}(t)$.

6.3.2 Symbolic performance

Considering the foreign exchange return matrix in equation (6.5), it is intuitively clear that a base currency with a large number of positive returns appreciates overall, whereas a base currency with a large number of negative returns depreciates overall. This information for a given base currency i is stored in the i th row of the return matrix (6.5).

We introduce the symbolic performance $\zeta_i(t)$ which we define as the difference

between the number of positive and negative returns currency i , acting as base currency, has at time t . This is achieved by applying the sign-function to returns $R_{i,j}(t)$ in the return matrix, yielding $+1$ for positive returns and -1 for negative returns, and summing row-wise:

$$\zeta_i(t) = \sum_{j=1}^K \text{sgn } R_{i,j}(t). \quad (6.6)$$

We know that $\text{sgn } R_{i,j}(t) = -\text{sgn } R_{j,i}(t)$. In other words, if the currency i rises with respect to currency j , then currency j falls with respect to currency i . The symbolic performance vector then lists the symbolic performances $\zeta_i(t)$ of all currencies at time t ,

$$\zeta(t) = (\zeta_1(t), \zeta_2(t), \dots, \zeta_K(t)). \quad (6.7)$$

By construction, the symbolic performance can take values from $-K + 1$ to $K - 1$ in steps of 2, and for a given time t each value appears exactly once; this means $\sum_i \zeta_i(t) = 0$. As a consequence, if currency j has a larger return than currency k with respect to the base currency i , then currency j will appreciate with respect to currency k . It is important to point out that the symbolic performance is a measure for each asset itself: We get K values $\zeta_i(t)$, $i = 1, \dots, K$, which describe the behavior of the K currencies individually, while containing the correlation structure of the FX market.

In a simple model, we can decompose the variance of a given return time series into one part which corresponds to the state of the market, i.e. the sum of the variances of all returns, and into another part which corresponds the idiosyncratic variance. While the symbolic performance of a given currency contains information about all other currencies, it does not provide information about the magnitude of an average return at a given time. In other words, it disregards the variance of the return time series due to the state of the market. However, it retains information on the idiosyncratic variance of the return time series of currencies. It is easy to see that equation (6.6)

is equivalent to a ranking of returns given a fixed base currency i where the largest positive (negative) return is assigned the largest positive (negative) value for $\zeta_i(t)$:

$$\begin{aligned}\zeta_i(t) &= \sum_{j=1}^K \operatorname{sgn} R_{i,j}(t) \\ &= \sum_{j=1}^K [2\mathbb{H}(R_{i,j}(t)) - 1] = 2 \sum_{j=1}^K \mathbb{1}(R_{j,1}(t) \leq R_{i,1}(t)) - 1 - K \\ &= 2 \operatorname{rank} R_{i,1}(t) - (K + 1),\end{aligned}\tag{6.8}$$

where \mathbb{H} denotes the Heaviside step function and $\mathbb{1}$ the indicator function. Note that this is true regardless of our choice of base currency i , as pointed out earlier.

Currencies with larger variance relative to the other currencies are more likely to exhibit large swings and thus higher magnitudes of ζ , and vice versa. As a result, currencies that tend to be more volatile regardless of market state tend to have more fat-tailed symbolic performance distributions $P(\zeta_i)$. Currencies that take a position at the center of the market in comparison to the remaining currencies tend to have a symbolic performance distribution $P(\zeta_i)$ that is reminiscent of a Gaussian. For the purpose of finding the position of a currency, it is then sufficient to consider the distribution of the absolute values of the symbolic performance, $P(|\zeta_i|)$, a discrete probability distribution which describes how often a currency takes the symbolic performance $-(K - 1)$ or $K - 1$, $-(K - 3)$ or $K - 3$, etc.

6.3.3 k -means++ clustering

The movement of currencies is closely linked to the macroeconomic developments in the corresponding countries. Therefore the more qualitative description of currencies as, for example, G4 currencies, G10 currencies or commodity currencies, is valid on long time scales. While this classification helps describe the roles and behaviors of currencies in general, economic shifts as well as political or monetary shocks can alter

the position of a currency within the market on varying time scales, and the symbolic performance is able to capture these shifts. To investigate these currency dynamics, we evaluate the symbolic performance distributions $P_t(\zeta_i)$ on reasonably short time scales, from t to $t + \Delta T$, and classify them according to different currency behavior.

We use k -means++ clustering for this task. Cluster analyses find application in many fields, such as, biology and bioinformatics, business and marketing, medicine, social networks, computer science, climate and weather research, and data mining, where they are used to discover similarities and patterns in large amounts of data. The k -means++ clustering algorithm takes as input a set of vectors as well as a number of clusters or classifications k . It then determines k points which serve as cluster centroids such that the overall distance between all data points and their respective cluster center is minimized. This algorithm itself does not require any knowledge of the data set to operate, and it will find a solution even if the data is not properly clustered into k centers. Therefore, the choice of the number of clusters k becomes critical, as we explain later in this section.

The algorithm schematically works as follows: Given N different points in d -dimensional space, k random points are selected as the initial cluster centers μ_i , $i = 1, \dots, k$. Each of the $N - k$ remaining points are then associated with the cluster center to which they are closest. After assigning all points to clusters, the cluster centers, defined as the mean of all points in the cluster, changes. The cluster centers are thus re-evaluated before the next iteration, and all points are re-classified to be closest to the new cluster centers, again. This procedure is reiterated until convergence is reached, that is, all points \mathbf{x} are distributed in clusters C_i , $i = 1, \dots, k$, with centers μ_i such that

$$\sum_{i=1}^k \sum_{\mathbf{x} \in C_i} \|\mathbf{x} - \mu_i\|^2 \quad (6.9)$$

is minimized.

The k -means++ algorithm uses the squared Euclidean metric as its distance measure. Recall that we are interested in the classification of symbolic performance distributions $P_t(\zeta_i)$, estimated on time intervals $[t, t + \Delta T]$. To this end we arrange the relative frequencies of $P_t(\zeta_i)$, in a vector,

$$\mathbf{P}_t(|\zeta_i|) = \begin{pmatrix} P_t(\zeta_i = \pm(K-1)) \\ P_t(\zeta_i = \pm(K-3)) \\ \vdots \\ P_t(\zeta_i = \pm 1) \end{pmatrix}. \quad (6.10)$$

Each vector $\mathbf{P}_t(|\zeta_i|)$ is a composition. This has implications regarding the distance measure. For compositional data the use of Aitchison’s distance, introduced in Aitchison [1982], is much more appropriate. In contrast, the use of Euclidean distance may be problematic, as Martín-Fernández et al. [1998] and Aitchison et al. [2000] point out. Alternatively, we can apply an isometric log-ratio (ilr) transformation to $\mathbf{P}_t(\zeta_i)$. The so-transformed data is then accessible to analytical methods that use the Euclidean metric. Since the scikit-learn package which we use to classify our data requires the Euclidean distance as the distance metric for k -means++, we apply an ilr transformation and cluster the results. In the last step, we perform the inverse ilr transformation to express the cluster centers in terms of relative frequencies again⁵.

In order to be able to provide a meaningful interpretation of the symbolic performance distribution, ΔT needs to be carefully chosen considering three criteria. One, there is no trading on the *Oanda* platform for a few hours every weekend. In order to avoid capturing beginning- or end-of-the-week effects, it is sensible that each time window is one week long or multiples of one week. Two, if we choose the length of the

⁵ By using the absolute values of the symbolic performances, we reduce the dimensionality from K to $K/2$, and by applying the ilr transformation, we further reduce it to $K/2 - 1$. This mitigates the problem of sparsity of high-dimensional spaces significantly.

window to be too small, we undersample the distribution. Three, if ΔT is too large, we may miss local trends and fluctuations, and we may limit our ability to resolve the effects of political or economic shocks whose impact can last anywhere from a few days to several months or years. Given our data resolution of 10 minutes, $\Delta T = 2$ weeks appears to be the appropriate choice to satisfy our criteria.

It is worth pointing out that we perform the cluster analysis on the pool of all symbolic performance distributions of all currencies, that is, on the entire FX market. As a result, information about each currency as well as each point in time informs the classification process. In other words, if two different currencies are classified to belong to the same cluster at different times, their respective behavior or role in the market is very comparable.

As pointed out previously, we also have to decide on a reasonable number of clusters k , based on the data set. We want to choose as few clusters as possible with as large explanatory power as possible. The gap statistic, introduced by Tibshirani et al. [2001], is a useful metric to determine the appropriate value for k as suggested by the structure of the data. Since k -means++ does not penalize model complexity, that is, the number of clusters to be found, increasing k does not increase the value of the cost function. This is true whether the data is clearly clustered or not. The gap statistic contrasts the benefit of adding one more cluster to a structured data set to the benefit of adding one more cluster to a comparable but random and unclustered data set. The gap is the difference in the cost function of clustering these two data sets. We increase k , starting with just one cluster, until the size of the gap reaches its first local maximum; at this point adding the $k + 1^{\text{th}}$ cluster provides diminishing benefits. After performing this analysis on our data set, we conclude that separating our data in six clusters is the most sensible choice.

6.4 Empirical results

6.4.1 Overall symbolic performance distributions

Based on the foreign exchange data introduced in Section 6.2 and the symbolic performances calculated according to equation (6.6), we analyzed the overall symbolic performance distributions $P(\zeta_i)$, of all currencies $i = 1, \dots, 14$, as shown in Figure 6.1. These discrete probability functions indicated how often each currency i took each value ζ_i from -13 to $+13$ during the entire time period from 2005 to 2017.

We observed that the distributions exhibit characteristic shapes for different currencies. Some are clearly convex, like the distributions for the Japanese yen, the New Zealand dollar, and the South African rand. Others are clearly concave, like the distributions for the euro, the Hong Kong dollar, the Singapore dollar, and the US dollar. This reflects the degree to which extent currencies tended to exhibit large swings, that is, had large volatility relative to the remaining currencies regardless of market conditions on one hand, or to which currencies tended to stay in the center of the market on the other. Irrespective of their shapes, however, all distributions are quite symmetric around zero⁶.

We describe the distributions based on the appearance of maxima in the center or at the tails; in other words we can distinguish the currencies' performance distributions by their curvature:

1. Strongly concave: EUR, HKD, SGD, USD.
2. Slightly concave: CHF, GBP.
3. Fairly flat: CAD, SEK.
4. Slightly convex: AUD, MXN, NOK.

⁶ We make use of this feature in our cluster analysis by considering only the absolute value of the symbolic performance and its distribution.

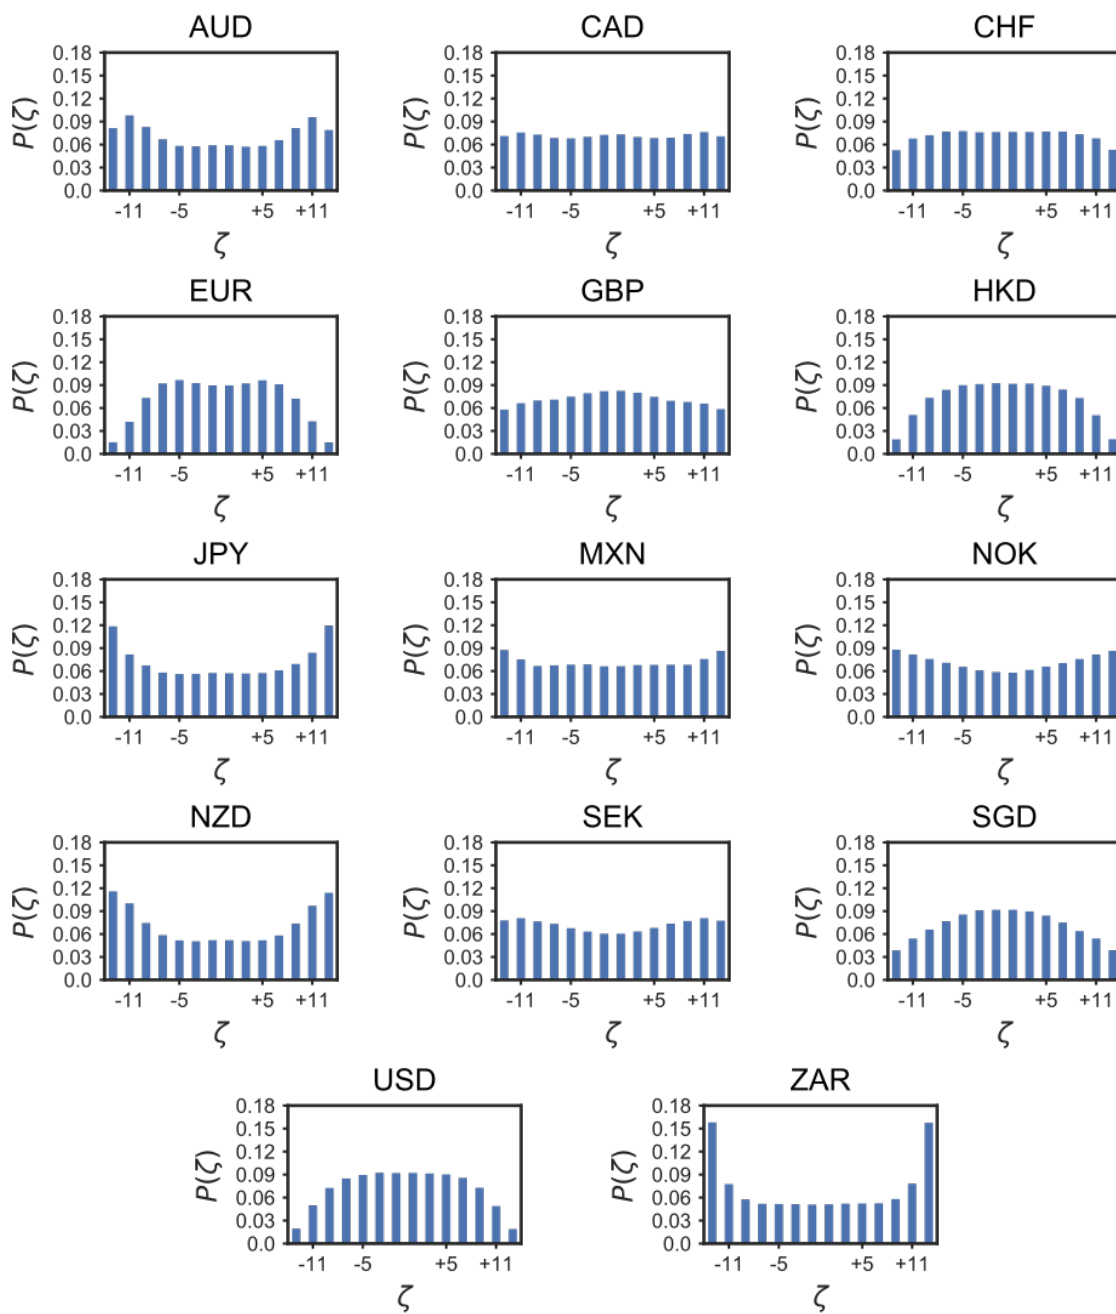


Figure 6-1: Overall symbolic performance distributions $P(\zeta_i)$, $i = 1, \dots, 14$, for our data set. We distinguish between a few different shapes, characterized by their curvature. Some currencies exhibit rather small probabilities at the center and high probabilities in the tails. At the extremes we find the euro (EUR) with very concave curvature and the South African rand (ZAR) with very convex curvature.

5. Strongly convex: JPY, NZD, ZAR.

Let us first consider the tail-heavy distributions. Currencies like the ZAR with an overall convex symbolic performance distribution were most likely to exhibit returns that are larger in magnitude than all other returns. These currencies tended to outperform or underperform the rest of the market. This does not, however, imply that these currencies necessarily were currencies with large exchange rate volatility. We could easily change the base pair to change the exchange rate volatility of a currency. When measured in euro, the Swiss franc, for example, was the currency with the lowest exchange rate volatility, but when measured in US dollar, the Swiss franc exhibited higher volatility. This observation underscores the usefulness of the symbolic performance as a measure of relative volatility since it takes into consideration all currency pairs at once.

Currencies like the EUR or USD with strongly concave symbolic performance distributions tended to maintain a position at the center of the market. This means that they were unlikely to consistently strongly appreciate or depreciate against other currencies.

Instead, they served as some form of reference against which other currencies are held. Consider an imaginary currency with true reference status. The size of its returns would be distributed like a normal distribution centered around zero with a very small standard deviation due to minor fluctuations. As a consequence, large fluctuations would be very unlikely, and likewise the chance for being at the extreme positions of the market would be very small as compared to other currencies.

We also identified currencies which appeared to take all symbolic performance values and therefore roles in the market with similar frequency. This could be part of the behavior of the currencies, or it could be a mixture of periods in which their symbolic performance distributions exhibit concave curvature and periods in which

they exhibit convex curvature.

6.4.2 Dynamics of symbolic performance distributions

According to the 2016 Triennial Central Bank Survey of FX and over-the-counter (OTC) derivatives markets, the daily trading volume of the foreign exchange market exceeded 5 trillion USD.⁷ This enormous volume results in high liquidity and rich dynamics, which are enhanced by central banks interventions as well as major political developments with macroeconomic consequences. The clustering procedure introduced and explained in Section 6.3.3 is able to describe the impact of such events. The results of the cluster analysis are presented in Figures 6.2 and 6.3. Training the algorithm on the symbolic performance distributions $P_t(\zeta_i)$ estimated on time intervals $[t, t + \Delta T]$ with $\Delta T = 2$ weeks for all t , the k -means++ algorithm identifies six cluster centers corresponding to six distinct distributions $P(\zeta)$. Figure 6.2 shows these characteristic symbolic performance distributions found by the k -means++ algorithm. Figure 6.3 exhibits how the currencies evolve through the clusters over time, that is, which symbolic performance distribution corresponding to certain currency behaviors describes them.

Table 6.1 offers a summary of how often each currency was classified into each cluster during the 12 year observation period. The euro and the US dollar occupied clusters 1 and 2 most often. The Japanese yen, the New Zealand dollar, and the South African rand could be found at the opposite end, as they were most often classified into clusters 5 and 6.

We can distinguish two different groups of currencies: Currencies that maintained their role in the foreign exchange market over a long time and spent the majority of the observation period in one or two clusters, and currencies that changed their role significantly over time, either in the form of short bursts or long-term shifts. A

⁷ Bank for International Settlements, <http://www.bis.org/publ/rpfx16.htm>




Figure 6.2: Symbolic performance distributions $P(\zeta)$ corresponding to the cluster centers found by the k -means++ algorithm. For the sake of better illustration, we draw a line for each discrete probability distribution, keeping in mind that $P(\zeta)$ is only defined for odd integers between -13 and 13 . The violet distribution (cluster 1) corresponds to currencies which take central positions in the market, whereas the red distribution (cluster 6) corresponds to currencies which tend to take more extreme positions in the market.

typical example of the first group was the euro which maintained its reference role almost throughout the entire time period. Typical examples for the second group were the British pound with the sudden change in clusters after the Brexit vote in June 2016, and the Swiss franc with drastic changes in September 2011 (pegging to the euro) and January 2015 (unpegging from the euro).

In the following we further investigate a selection of currencies which serve as an example to demonstrate the ability of our methodology to resolve currency dynamics in response to economic and political interventions have occurred. These are in descending order of their average daily turnover in April 2016 according to BIS Monetary and Economic Department [2016]: US dollar, euro, Japanese yen, British pound, Canadian dollar, Swiss franc, Mexican peso, Singapore dollar, and Hong Kong dollar.

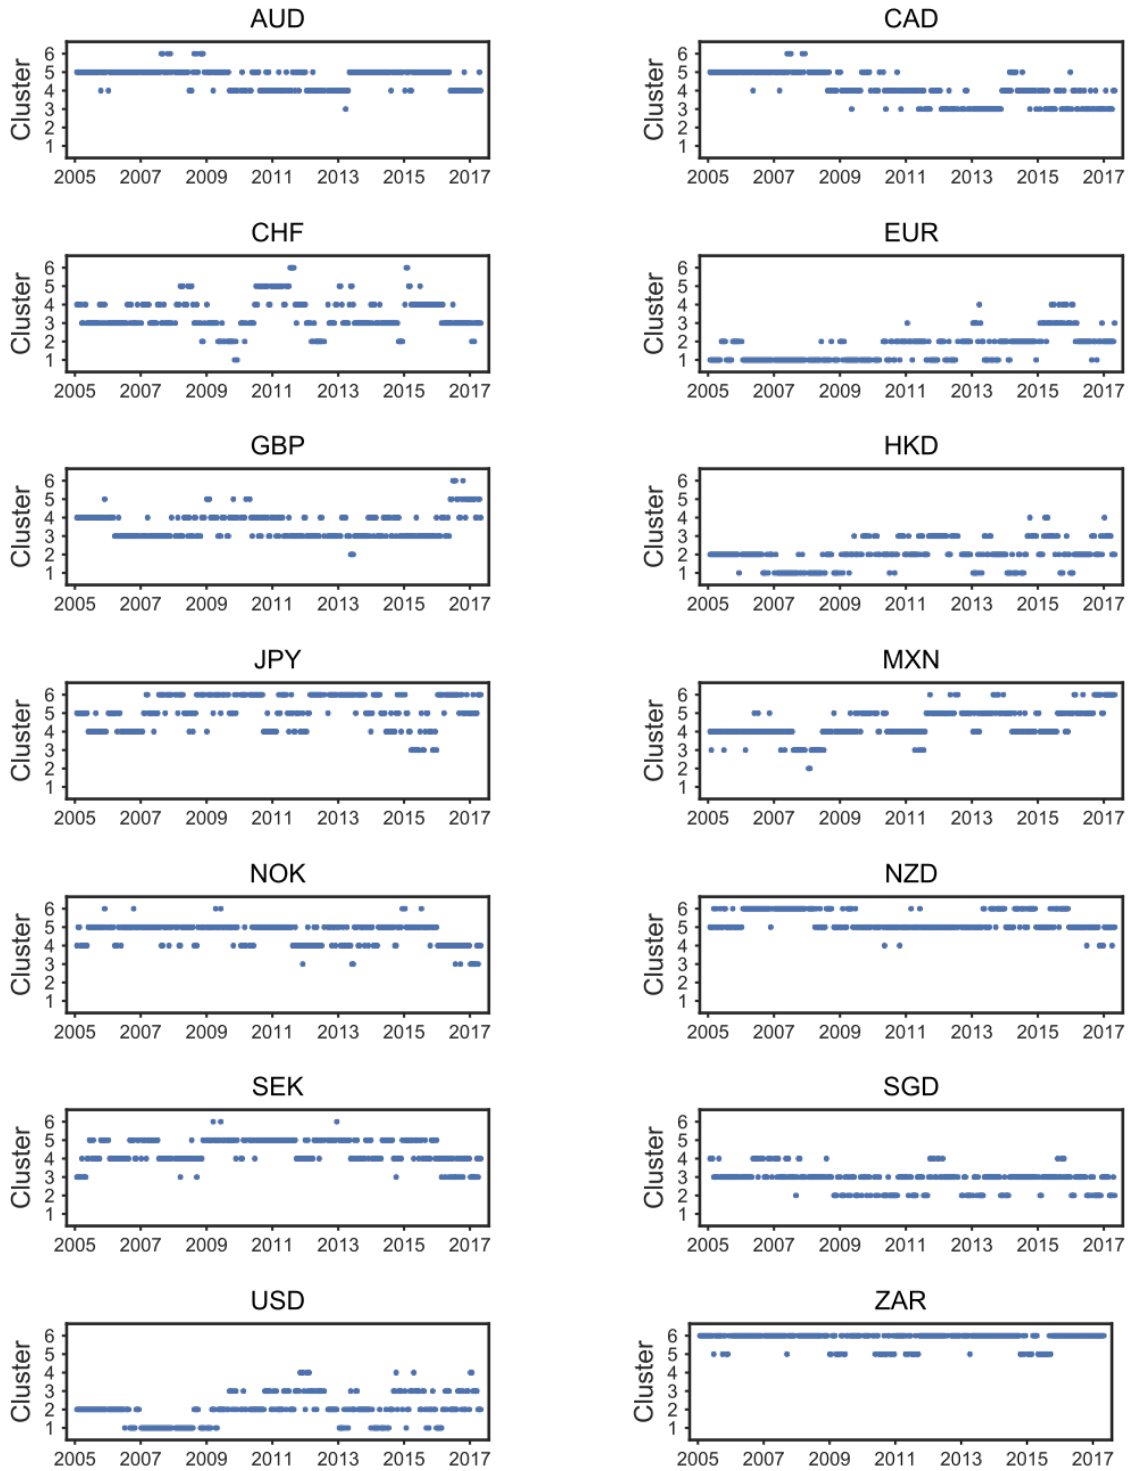


Figure 6-3: Currency clustering dynamics according for the 14 currencies in our data set. Cluster 1 corresponds to reference-like behavior, whereas cluster 6 corresponds to heavy-tail behavior.

Table 6.1: We report how frequently (as a percentage) each currency appears in each cluster, enumerated from 1 to 6 as in Figure 6.2, as well as the average cluster value. The euro has the lowest average cluster component, followed by the US dollar, the two main reference currencies in the FX market. The Japanese yen, the New Zealand dollar, and the South African rand stand out as the currencies with the most tail-heavy symbolic performance distributions.

cluster	AUD	CAD	CHF	EUR	GBP	HKD	JPY	MXN	NOK	NZD	SEK	SGD	USD	ZAR
1	0	0	1	49	0	24	0	0	0	0	0	0	27	0
2	0	0	10	39	1	53	0	1	0	0	0	24	50	0
3	0	30	49	9	53	22	4	10	4	0	8	65	21	0
4	32	35	27	3	37	1	23	47	30	3	42	10	3	0
5	65	33	12	0	7	0	29	34	64	60	49	0	0	16
6	3	2	2	0	1	0	44	9	2	37	1	0	0	84
average	4.7	4.1	3.5	1.7	3.5	2	5.1	4.4	4.6	5.3	4.4	2.9	2	5.8

USD For the majority of the time period the US dollar was associated with clusters 1 and 2 and appeared to be a reference currency. In particular, the US dollar shared this role with the euro in cluster 1 during the year 2007, indicating very little movement with respect to the overall market. However, with the onset of the global financial crisis of 2008 the reference role of the US dollar within the foreign exchange market appeared to weaken, as bigger swings and movements with respect to the remaining 13 currencies were more prevalent. This phase in which the dollar spent most time in cluster 3 lasts until late 2012. Hereafter the US dollar mostly stayed in cluster 2 with some appearances in cluster 1 in 2014, reflecting the upswing of the US economy in the wake of the persisting European sovereign debt crisis and fear of deflation.

EUR Our analysis confirms the euro as a major reference. This is detailed by its continued presence in cluster 1, despite global and local financial turmoil. During the European sovereign debt crisis, though, the euro left cluster 1 to move to cluster 2 more frequently. On July 26, 2012, Mario Draghi announced that the ECB would

do ‘whatever it takes’ to protect the euro⁸. This led to a large appreciation of the euro from roughly 1.23 USD to 1.38 USD per euro in the following nine months. As a result of this, the euro left cluster 1 for almost a year, and it even briefly moved to cluster 3 and cluster 4, indicating a more volatile period, before returning to cluster 2 and then 1. During most of 2015, the euro was in cluster 3 again, with significant time in cluster 4. Towards the end of the observation interval, no currency belonged to cluster 1 for more than a few months at a time. We can interpret this as a lack of reference in the FX market. Even after the return to cluster 2 in early 2016, the euro occupied cluster 1 only for brief periods of time.

JPY Despite its role as one of the IMF reference currencies, the Japanese yen was found at the tails of the market which correspond to relatively large movements. Unlike most other currencies, which tended to stay in the same cluster for extended periods of time and moved slowly in Figure 6-3, the Japanese yen switched clusters rather frequently, moving mostly between cluster 5 and 6 starting in the year 2007.

It is noteworthy to highlight that this was not an artifact of the size of the interval we had selected to estimate the symbolic performance distributions, as this observation would hold similarly for longer and shorter time intervals. Instead the symbolic performance distribution changed frequently without deviating much from the cluster centers. We observed, however, that the central bank interventions triggered the move of the Japanese yen from cluster 6 to cluster 4, where it remained for a while. Likewise, it took an outside shock for the yen to change its cluster association again and return to cluster 6, as illustrated in Figure 6-4. The first line marks an intervention on September 15, 2010, after reaching a 15-year-low of the US dollar with respect to the yen, and it put the yen in cluster 4 where it remained fairly steadily. The second intervention happened in the light of the Fukushima nuclear disaster. It

⁸ <https://www.ecb.europa.eu/press/key/date/2012/html/sp120726.en.html>

is worth mentioning that this intervention was a joint effort by the G7, and thus it is not surprising that the yen returned to cluster 6.

The yen was mostly in cluster 3 in 2015; one may suspect that these changes are due to central bank actions. Given the USD/JPY rate development in 2015 and 2016, the Bank of Japan (BOJ) and government officials issued statements indicating the possibility of intervention. Our methodology indicates that these talks and any possible covert interventions have been successful, as the yen stabilized in its standard role in clusters 3 and 4 on the foreign exchange market in 2016. Recently, however, it returned to clusters 5 and 6. Since Figure 6-4 shows that the yen tends to eventually return to higher clusters a few weeks or months after interventions, we suspect that the BOJ has been intervening less actively or forcefully in the recent couple of years.





Figure 6-4: Cluster association of the Japanese yen, highlighting the time points and background of publicly announced central bank interventions.

GBP Being in cluster 3 for the majority of the time horizon, the British pound played the role of a rather hard currency. On June 23, 2016, the British people voted to leave the European Union, causing a major shock to the world economy. This was reflected by a move of the British pound to cluster 4 already at announcement of the referendum date, as we observe in Figure 5 in more detail, indicating higher relative volatility. In the weeks up to the referendum, uncertainty increased further

and the pound was classified into cluster 5. The surprising outcome of the vote caused large movements in exchange rates around the globe, especially involving the British pound. With our methodology we can monitor if the British pound is going return to its role or if this kind of political intervention has changed the hierarchy in the foreign exchange market for the long run. It is going to be interesting to analyze whether the outcome of the British referendum to leave the European Union will have a profound and lasting effect on the role and importance of the British pound or whether it will be a short-term shock to the currency. In the case of the Swiss franc we observed that its pegging and unpegging has not permanently changed its position in the market. (See the paragraph regarding the CHF.) Instead, it was a temporary change, and within one year of unpegging, the Swiss franc returned to cluster 3 where it was before the major interventions. However, the uncertainty induced by the pending Brexit negotiations with the European Union has seemed to leave this matter unresolved. While the pound spent only a total of eight weeks in cluster 6, the cluster indicating the largest relative volatility, it since has remained in cluster 5 most of the time.



cluster 5. Around late 2008, with the onset of the global financial crisis, the heavy tails became weaker, as shown by its presence in cluster 4. From mid 2011 until late 2013, the Canadian dollar behaved more like the Swiss franc or the British pound, as shown by its association to cluster 3. In 2014 this development appeared to revert itself, briefly. Eventually, however, the Canadian dollar maintained clusters 3 and 4. It is crucial to emphasize that the symbolic performance distributions associated with the clusters provide relative information. Therefore, we cannot automatically infer whether the Canadian dollar itself has changed its role or the changes pertain to the rest of the market; instead we have to consider the development of other currencies during these time periods.

CHF In figure 6-6 we zoom into the cluster association of the Swiss franc, drawing lines for each publicly announced central bank intervention of the Swiss National Bank (SNB). The Swiss franc exhibited particularly stable behavior prior to the onset of the global financial crisis in 2008, being in cluster 3 most of the time. While the franc has long been considered a hard currency given the financial and political stability of Switzerland, its use as a funding currency may frequently induce some additional volatility in the foreign exchange rates which include the franc. As a result its behavior is not quite like that of US dollar and euro. Instead our analysis determined its behavior to be closer to the behavior of the British pound and the Singapore dollar. With the onset of the financial crisis in 2008, the Swiss franc exhibited larger swings than other currencies for a while, but soon returned to cluster 3.

Starting early 2009, the Swiss franc began an extended period in cluster 2 which is usually only occupied by reference currencies. Additionally, currencies in a managed float regime with respect to a reference currency or a large basket of currencies are typically associated with cluster 2. We explain this in more detail in our discussion of the Hong Kong dollar. In case of the Swiss franc, we can link this behavior

to aggressive action of the SNB which acted to maintain the stability of the franc and curb its appreciation against other major currencies by buying euros and other currencies, effectively managing the exchange rate. These public interventions started in early 2009, and we consider them successful in that they moved the CHF into cluster 1, where we refer to currencies as reference currencies. In the wake of the European sovereign debt crisis, the SNB intervened more frequently, as indicated by the numerous lines in early 2010 in Figure 6-6. During the time of these interventions the Swiss franc stayed in clusters 2 and 3 despite significant market pressure. However, eventually the SNB abandoned its policy to purchase euros and other currencies⁹ In the following weeks, the impact that this policy had became obvious as the Swiss franc was mostly in cluster 5, corresponding to above-average volatility. A consequence of the end of central bank interventions was an overall appreciation of the Swiss franc in relation to the euro and the US dollar. These appreciations eventually shifted the Swiss franc to cluster 6 which it had never occupied before.

These large currency movements sparked further central bank interventions, resulting in the enforcement of a cap on the exchange rate of the Swiss franc to the euro on September 6, 2011, by the Swiss National Bank. Our analysis indicates that after a couple of months the market accepted the SNB's controlled cap, and the Swiss franc moved back to its more pre-crisis role, appearing in cluster 3 most of the time. Another intervention on January 15, 2015, corresponded to the uncapping of the Swiss franc with respect to the euro. This SNB intervention pushed the Swiss franc in clusters 5 and 6 again. For roughly two months the Swiss franc exhibited heavy tails in the symbolic performance distribution. Over the time scale of a year, it returned to cluster 3, its standard position. Further intervention is suggested to have taken place in early February 2017, inferred from rising sight deposits¹⁰, and we

⁹ Speech by Fritz Zurbrugg, Member of the Governing Board of the Swiss National Bank, <https://www.bis.org/review/r150330c.png>

¹⁰ <http://uk.reuters.com/article/uk-swiss-snb-idUKKBN15S1NU>

indeed observe a drop in cluster state around that period.




Figure 6-6: Cluster association of the Swiss franc, highlighting the time points of publicly announced central bank interventions.

MXN The Mexican peso was, until late 2011, mostly classified into cluster 4 with some extended stints in clusters 3 and 5. During this time period the peso therefore could be characterized as a currency of average relative volatility when compared to the remaining currencies in our data set. Given Mexico’s status as an emerging market, this is not surprising. Starting in late 2011, the peso shifted to cluster 5 and remained there for most of a time for a couple of years, reflecting a higher relative volatility. This coincided with growth rates which remained below expectations.

Trade with the United States has long had a great impact on the Mexican economy since the overwhelming majority of Mexican exports go to the US. Generally the Mexican peso considered a satellite currency of the US dollar, and more than 90% of all peso transactions are trades on the USD/MXN pair. As a result of this close relationship and economic interconnectedness, the Mexican peso has been susceptible to political and economic developments that potentially affect the trade relationships with the US. As of 2018, the North American Free Trade Agreement (NAFTA) is being renegotiated. The market seemed to have anticipated this risk after the surge of Donald Trump in the primary polls of the Republican party in mid to late July 2016, corresponding to a change in cluster of the Mexican peso that moved from cluster

4 to 5 and eventually to 6. It only left this cluster to return to cluster 5 following a surprising intervention in February 2016 when Banxico intervened with a ‘shock and awe’ move by selling US dollars and simultaneously increasing the interest rate¹¹. Following the inauguration of the new president in the US in early 2017, the peso returned to cluster 6 where it remained for the rest of the observation period. In summary, the peso behavior was characteristic of commodity currencies which are closely correlated to world market prices of commodities which they export. The higher relative volatility of such currencies reflects the vulnerability of the underlying economies to major market shifts. Our results suggest that the market has been considering the importance of Mexico’s access to North American markets and its implications for the peso to be similar to that of a commodity exporter and its exposure to the price of that commodity.




Figure 6-7: Cluster association of the Mexican peso, highlighting the time points of publicly announced central bank interventions in red and political events in the United States in green.

SGD Singapore’s central bank is the Monetary Authority of Singapore (MAS). It is the mission of the MAS to hold the Singapore dollar in a managed float regime to control inflation in Singapore, which is detailed in biannual policy publications.

¹¹ <https://www.ft.com/content/7129d13c-d364-11e6-9341-7393bb2e1b51>

Through direct intervention the MAS holds the Singapore dollar within a fixed band against an undisclosed trade-weighted basket of currencies. This method of controlling inflation stands in contrast to what most other central banks in our data set do, which is choosing to adjust interest rates to control inflation. Our clustering analyses show the results we would expect given the MAS policies and procedures: First, the Singapore dollar has occupied cluster 3 for most of the time, interrupted by moves to cluster 2 and sometimes cluster 4. As the MAS enforces the policy band, bigger overall swings become less likely. This is typical of currencies in clusters 2 and 3. Furthermore, the main trading partners of Singapore comprise Hong Kong and China, whose currencies have been closely linked to the US dollar, and the United States. This implies that the cluster of the Singapore dollar tends to be similar to that of the US dollar. Figure 6-8 indicates with red vertical lines the times when the MAS published its policy updates, and we observe that changes in clusters for the Singapore dollar were often in accordance to these updates.




Figure 6-8: Cluster association of the Singapore dollar, highlighting frequent publications of policy briefs by the Monetary Authority of Singapore.

HKD When a currency is in a managed float regime, the central bank participates in the foreign exchange market. It does so by buying and selling currency to achieve the target rate within a certain band in which the currency is allowed to float. This

reduces the relative volatility of a currency in two ways. One, since central banks limit the amount of appreciation and depreciation through their market actions, the volatility of the exchange rates of such a currency is suppressed. Two, when a central bank manages its currency in such a way, it usually does so with respect to a reference currency or a basket of currencies. A reference currency has lower relative volatility than most currencies in the market, as evidenced by the clustering dynamics of the US dollar or the euro. Likewise, the relative volatility of a basket of currencies generally is smaller. As a result the relative volatility of the managed currency is lower as well. In case of the Hong Kong dollar, the currency has been managed with respect to the US dollar which explains the similarity of their clustering dynamics. We observed only minor deviations at the end of 2007, where the Hong Kong dollar was in cluster 2 while the US dollar maintained cluster 1, and around 2012, where the Hong Kong dollar stayed in cluster 3 while the US dollar moved to cluster 4.

6.5 Conclusion

We have introduced a symbolic performance measure to quantify the roles of currencies in the foreign exchange market. Instead of investigating currency pairs, this novel approach allows for analysis of currencies individually, embedded within the entire system. We were able to identify the central (or reference) currencies in the foreign exchange market, and unsurprisingly the euro has been playing a major role, much like the US dollar and the British pound. Our methodology has also proven effective in observing the dynamics of currency roles over time. We found that, in general, currencies maintain certain positions in the FX market for extended periods of time.

The independence from the choice of base currency in our methodology and its design enabled us to closely investigate the characteristics of individual currencies with respect to the rest of the currencies in the network. We used our methodology

to analyze the effects of central bank interventions, and we could hypothesize when a covert intervention by a central bank seems to have occurred. Furthermore we have been able to examine the impact of specific shocks to the system, such as the pegging of the Swiss franc with respect to the euro, major interventions of the Bank of Japan, and the vote for Brexit.

We suggest that our methodology can be used to trace movements in the foreign exchange market to detect currency interventions, political developments and economic downturns or booms. One natural expansion of this work would be to include analysis of real-time data, monitoring the distribution of symbolic performances on shorter time scales and using observed turbulences in the foreign exchange market for currency forecasting.

Chapter 7

Summary and conclusions

This thesis has studied the effects of interconnectedness and interdependence of financial systems across a range of scales. We began by analyzing the community structure and lead-lag relationships on global equity and currency markets in Chapter 2. Our results indicated significant differences depending on the state of the global economy. Using network filtering techniques, we showed that, even in the today's age of extensive global trade, geography has been a good predictor of community formation; in times of crisis, however, the communities were more mixed. We observed that the euro was at the center of a community of European currencies pre-crisis, while this hierarchy vanished during the crisis. Remarkably, we identified a cluster consisting of the equity markets of the troubled GIPS countries – Greece, Ireland, Portugal, and Spain – during the crisis period. Different lead-lag relationships emerged for the time of the global financial crisis, and therefore our results caution against the reliance on historical data for a global macro-hedging strategy.

Financial institutions and their holdings of mispriced assets took the center stage during the global financial crisis, which showed the need for smarter regulation of the banking system. Before 2008, regulators had put their focus on the solvency and resilience of individual banks, also called microprudential regulation. This kind of regulation proved too soft as it disregarded the complex and interwoven nature of the financial system. In Chapter 3 we proposed a model to study the vulnerability of a system of banks with portfolio overlap. Using empirical data from the European

banking authority, we studied the effect of shocks for a broad range of risk aversion and liquidity parameters. Precariously, microprudential regulation proves sufficient for times of high confidence and asset liquidity, potentially providing a false sense of security to both regulators and banks. Under adverse conditions, however, the financial network exhibited a large degree of fragility, especially once we accounted for the regulatory requirements imposed on banks. In fact, the more adverse conditions became, e.g., higher risk-aversion or default probability of assets, the more sensitive the system became to small shocks.

Since network studies on systemic risk – such as ours – typically include a shock propagation mechanism, they require accurate knowledge of the system. Often, such information is confidential, and the network needs to be reconstructed from incomplete information. Probabilistic methods have emerged as the front runner, creating an ensemble of possible reconstructions based on the principle of maximum entropy. Considering the Nikkei data set, which includes detailed information on Japanese banks and firms for more than three decades, we tested reconstruction method, the enhanced capital asset-pricing model (ECAPM), in Chapter 4. ECAPM is an extension of the popular bipartite configuration model (BiCM). Both models assume a node property called strength to be a predictor of connectivity. Our analysis showed that a significant number of nodes in the network would remain unconnected in most reconstructions of the Japanese credit network, in line with recent findings that systemic risk would be underestimated in this system. We derived equations to study the range of validity of the reconstruction, depending on the distribution of node strengths, the network density, and the size of the network. With numerical calculations, we demonstrated bounds and showed that sparse networks with inhomogeneous layers are particularly hard to reconstruct. We propose that our results complement typical tests of reconstruction quality which put a focus on network topology mea-

asures, like degree distributions or the nearest-neighbor degrees.

In Chapter 5 we used the principle of maximum entropy to develop a novel measure of co-movement of assets, applying it to the foreign exchange market. Unlike bonds or equity, currencies are traded in a decentralized fashion, and exchange rates are quoted pairwise. Interpreting the FX market as a fully connected network, we measured the amount of information contained in each traded pair. We tested the empirical networks against the null model of no information, meaning complete independence of currencies. We found significant deviations from the null model for some currency pairs, and with a generalized linear model we linked these findings to macroeconomic observables like interest rates and trade volume. Our analysis showed that both, macroeconomic indicators as well as idiosyncratic properties of currencies contribute to the deviation from the null model.

Therefore, while macroeconomic data drives foreign exchange markets, other aspects like the intervention of central banks may strongly influence currency movements. In Chapter 6 we introduced a clustering approach to determine the roles of currencies within the market. We demonstrated that this methodology is sensitive to detecting political and monetary interventions and their impact on the structure of the foreign exchange market. We believe that our approach can be useful in the future to provide an underpinning of qualitative concept in foreign exchange markets.

Appendix A

Supplementary Material to Chapter 2

A.1 Table of countries in data set

Country	Currency	Index Label	Currency Label
Argentina	Argentine peso	ARG	ARS
Australia	Australian dollar	AUS	AUD
Austria	euro	AUT	EUR
Belgium	euro	BEL	EUR
Brazil	Brazilian real	BRA	BRL
Canada	Canadian dollar	CAN	CAD
Chile	Chilean peso	CHL	CLP
China	Chinese yuan	CHN	CNY
Colombia	Colombian peso	COL	COU
Czechia	Czech koruna	CZE	CZK
Denmark	Danish krone	DNK	DKK
Finland	euro	FIN	EUR
France	euro	FRA	EUR
Germany	euro	DEU	EUR
Greece	euro	GRC	EUR
Hong Kong	Hong Kong dollar	HKG	HKD
Hungary	Hungarian forint	HUN	HUF
Iceland	Icelandic krona	ISL	ISK
India	Indian rupee	IND	INR
Indonesia	Indonesian rupiah	IDN	IDR
Ireland	euro	IRL	EUR
Israel	Israel new shekel	ISR	ILS
Italy	euro	ITA	EUR

Japan	Japanese yen	JPN	JPY
Kazakhstan	Kazakhstani tenge	KAZ	KZT
Korea	South Korean won	KOR	KRW
Luxembourg	euro	LUX	EUR
Malaysia	Malaysian ringgit	MYS	MYR
Malta	Maltese lira	MLT	MTL/EUR
Mauritius	Mauritian rupee	MUS	MRO
Mexico	Mexican peso	MEX	MXN
Netherlands	euro	NLD	EUR
New Zealand	New Zealand dollar	NZL	NZD
Norway	Norwegian krone	NOR	NOK
Oman	Omani rial	OMN	OMR
Pakistan	Pakistani rupee	PAK	PKR
Peru	Peruvian nuevo sol	PER	PEN
Philippines	Philippine peso	PHL	PHP
Poland	Polish zloty	POL	PLN
Portugal	euro	PRT	EUR
Qatar	Qatari riyal	QAT	QAR
Russia	Russian ruble	RUS	RUB
Saudi Arabia	Saudi riyal	SAU	SAR
Singapore	Singapore dollar	SGP	SGD
Slovakia	Slovak koruna	SVK	SKK/EUR
South Africa	South African rand	ZAF	ZAR
Spain	euro	ESP	EUR
Sri Lanka	Sri Lankan rupee	LKA	LKR
Sweden	Swedish krona	SWE	SEK
Switzerland	Swiss franc	CHE	CHF
Thailand	Thai baht	THA	THB
Tunisia	Tunisian dinar	TUN	TND
Utd. Arab Emirates	UAE dirham	ARE	AED
United Kingdom	British pound	GBR	GBP
United States	US dollar	USA	USD
Venezuela	Venezuelan bolivar	VEN	VEF

Appendix B

Supplementary Material to Chapter 3

B.1 Interbank network results

Sector	Financial		Corporate		Ret. Resid.		Ret. Rev.		Ret. SME		Comm. RE		Average	
	D_{50}	D_{100}	D_{50}	D_{100}	D_{50}	D_{100}	D_{50}	D_{100}	D_{50}	D_{100}	D_{50}	D_{100}	D_{50}	D_{100}
(1) GIIPS	72	77	71	72	78	81	12	75	48	76	55	76	48.6	65.6
(2) EE	5	75	35	76	5	74	3	27	4	56	5	67	8.4	53.9
(3) Benelux	7	75	57	76	19	76	3	8	5	69	5	69	14.0	53.6
(4) Greece	3	56	15	76	4	59	3	15	3	18	3	48	4.7	39.1
(5) Italy	27	77	69	75	20	76	3	20	16	76	17	76	22.0	57.4
(6) France	8	76	68	75	22	76	3	44	8	75	5	71	16.6	59.9
(7) Germany	72	77	71	75	21	78	3	22	5	69	25	76	28.4	57.0
(8) Spain	19	78	68	75	60	80	5	69	12	75	12	75	25.7	64.9
(9) US	12	77	69	75	25	76	5	69	4	64	10	75	18.3	62.6
(10) Japan	3	62	22	76	5	69	3	15	3	53	3	55	5.6	47.6
Average	22.9	73.1	54.5	75.1	25.9	74.5	4.4	36.4	11.0	63.1	14.1	68.8	19.3	56.1

Table B.1: Number of banks that would fall below the Basel III threshold of 4.5% tier 1 capital ratio given a shock to different sectors in each of our ten scenarios after 50 simulation time steps (D_{50}) and after 100 simulation time steps (D_{100}). The initial shock is a sudden increase of the risk weights of the shocked sector of a particular country or region by 50%, causing an increase of the risk-weighted assets of banks. The numbers include the two banks which are below the threshold originally.

Origin	Sector	Shock		Loss Avg	Loss GIIPS	Loss Other	D_{100} Total	D_{100} GIIPS	D_{100} Other		
		Size (%)	Q_A $P(x)$								
GIIPS	Sov. Debt	150	lin	0.3	0.0	0.0	0.0	2	2	0	
			steep	0.3	0.0	0.0	0.0	2	2	0	
		400	lin	0.9	0.0	0.0	0.0	2	2	0	
			steep	0.9	69.0	69.0	69.1	87	41	46	
		Financial	150	lin	0.3	3.3	4.7	2.0	2	2	0
				steep	0.3	4.9	6.6	3.2	3	2	1
	400		lin	0.9	10.5	12.8	8.4	3	2	1	
			steep	0.9	69.0	69.0	69.1	87	41	46	
	Corporate		150	lin	0.3	10.1	17.2	3.6	4	3	1
				steep	0.3	14.3	22.6	6.8	8	7	1
		400	lin	0.9	30.3	39.9	21.6	20	17	3	
			steep	0.9	69.0	69.0	69.1	87	41	46	
		Ret. Resid.	150	lin	0.3	23.7	39.0	9.7	21	19	2
				steep	0.3	27.2	41.5	14.1	22	19	3
	400		lin	0.9	38.5	49.5	28.5	33	29	4	
			steep	0.9	69.0	69.0	69.1	87	41	46	
	Ret. Rev.		150	lin	0.3	6.2	11.5	1.3	2	2	0
				steep	0.3	8.4	14.7	2.6	3	3	0
		400	lin	0.9	16.3	24.2	9.0	7	6	1	
			steep	0.9	69.0	69.0	69.1	87	41	46	
		Ret. SME	150	lin	0.3	24.1	42.9	6.8	23	23	0
				steep	0.3	29.5	48.5	12.1	30	29	1
	400		lin	0.9	42.0	56.9	28.4	37	34	3	
			steep	0.9	69.0	69.0	69.1	87	41	46	
	Comm. RE		150	lin	0.3	1.5	3.0	0.2	2	2	0
				steep	0.3	2.1	3.8	0.5	2	2	0
		400	lin	0.9	3.7	5.9	1.7	2	2	0	
			steep	0.9	69.0	69.0	69.1	87	41	46	
		Other	150	lin	0.3	4.1	7.9	0.7	2	2	0
				steep	0.3	5.2	9.4	1.3	2	2	0
	400		lin	0.9	8.6	13.5	4.2	4	3	1	
			steep	0.9	69.0	69.0	69.1	87	41	46	

Table B.2: Effect of a shock that originates in the GIIPS countries. There are a total of 43 banks in the GIIPS countries. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in the GIIPS countries (Loss GIIPS) as well as the average loss in tier 1 capital ratio for banks in other than the GIIPS countries (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for the GIIPS countries (D_{100} GIIPS) and for other than the GIIPS countries (D_{100} Other).

Origin	Sector	Shock Size (%)	Q_A	$P(x)$	Loss Avg	Loss EE	Loss Other	D_{100} Total	D_{100} EE	D_{100} Other	
Eastern Europe	Sec. Debt	150	0.3	lin	0.0	0.0	0.0	2	0	2	
			steep	0.0	0.0	0.0	2	0	2		
		400	0.3	lin	0.0	0.2	0.0	2	0	2	
			steep	0.0	0.2	0.0	2	0	2		
		Financial	150	0.3	lin	0.6	4.1	0.4	2	0	2
				steep	0.9	4.8	0.7	2	0	2	
	400		0.3	lin	1.9	6.5	1.7	2	0	2	
			steep	69.0	67.9	69.1	87	4	83		
	Corporate		150	0.3	lin	1.8	16.4	1.1	2	0	2
				steep	2.5	18.4	1.8	2	0	2	
		400	0.3	lin	5.4	23.0	4.6	4	1	3	
			steep	69.0	67.9	69.1	87	4	83		
		Ret. Resid.	150	0.3	lin	0.6	7.3	0.3	2	0	2
				steep	0.9	8.1	0.6	2	0	2	
	400		0.3	lin	1.8	9.8	1.5	2	0	2	
			steep	69.0	67.9	69.1	87	4	83		
	Ret. Rev.		150	0.3	lin	3.1	28.6	1.9	2	0	2
				steep	4.4	31.0	3.1	2	0	2	
		400	0.3	lin	9.7	37.3	8.4	4	1	3	
			steep	69.0	67.9	69.1	87	4	83		
		Ret. SME	150	0.3	lin	0.4	5.0	0.1	2	0	2
				steep	0.5	5.7	0.2	2	0	2	
	400		0.3	lin	0.8	6.7	0.5	2	0	2	
			steep	69.0	67.9	69.1	87	4	83		
	Comm. RE		150	0.3	lin	0.9	12.7	0.4	2	0	2
				steep	1.1	13.4	0.6	2	0	2	
		400	0.3	lin	1.9	14.8	1.3	2	0	2	
			steep	69.0	67.9	69.1	87	4	83		
		Ret. Rev.	150	0.3	lin	0.4	5.3	0.2	2	0	2
				steep	0.6	5.9	0.3	2	0	2	
	400		0.3	lin	1.1	7.1	0.8	2	0	2	
			steep	69.0	67.9	69.1	87	4	83		
	Comm. RE		150	0.3	lin	0.8	9.8	0.4	2	0	2
				steep	1.1	10.5	0.6	2	0	2	
		400	0.3	lin	2.0	12.1	1.5	2	0	2	
			steep	69.0	67.9	69.1	87	4	83		
		Comm. RE	150	0.3	lin	0.3	1.1	0.3	2	0	2
				steep	0.5	1.5	0.5	2	0	2	
	400		0.3	lin	1.2	2.5	1.1	2	0	2	
			steep	69.0	67.9	69.1	87	4	83		
	Comm. RE		150	0.3	lin	0.7	2.1	0.6	2	0	2
				steep	1.0	2.7	0.9	2	0	2	
		400	0.3	lin	2.1	4.4	2.0	2	0	2	
			steep	69.0	67.9	69.1	87	4	83		

Table B.3: Effect of a shock that originates in Eastern Europe. There are a total of 4 banks in Eastern Europe. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Eastern Europe (Loss EE) as well as the average loss in tier 1 capital ratio for banks in other countries than in Eastern Europe (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Eastern Europe (D_{100} EE) and for other than Eastern Europe countries (D_{100} Other).

Origin	Sector	Shock		Loss Avg	Loss Benelux	Loss Other	D_{100} Total	D_{100} Benelux	D_{100} Other	
		Size (%)	Q_A $P(x)$							
Benelux	Sov. Debt	150	lin	0.0	0.0	0.0	2	0	2	
			steep	0.0	0.0	0.0	2	0	2	
		400	lin	0.0	0.0	0.0	2	0	2	
			steep	0.0	0.0	0.0	2	0	2	
		Financial	150	lin	0.5	3.6	0.2	2	0	2
				steep	0.7	4.1	0.5	2	0	2
	400		lin	1.8	5.5	1.5	3	0	3	
			steep	69.0	71.9	68.8	87	7	80	
	Corporate		150	lin	2.3	15.7	1.2	3	0	3
				steep	3.5	17.1	2.4	3	0	3
		400	lin	8.3	21.5	7.2	3	0	3	
			steep	69.0	71.9	68.8	87	7	80	
		Ret. Resid.	150	lin	1.0	5.7	0.6	2	0	2
				steep	1.8	7.0	1.4	2	0	2
	400		lin	5.3	11.2	4.8	3	0	3	
			steep	69.0	71.9	68.8	87	7	80	
	Ret. Rev.		150	lin	3.2	16.8	2.1	4	0	4
				steep	5.2	19.0	4.0	4	0	4
		400	lin	12.7	25.5	11.6	4	0	4	
			steep	69.0	71.9	68.8	87	7	80	
		Ret. SME	150	lin	0.8	6.8	0.3	2	0	2
				steep	1.2	7.7	0.7	2	0	2
	400		lin	3.2	10.3	2.6	3	0	3	
			steep	69.0	71.9	68.8	87	7	80	
	Comm. RE		150	lin	3.7	28.1	1.7	3	1	2
				steep	5.9	31.5	3.7	4	1	3
		400	lin	14.5	38.6	12.5	5	1	4	
			steep	69.0	71.9	68.8	87	7	80	
		Ret. Rev.	150	lin	0.0	0.1	0.0	2	0	2
				steep	0.1	0.2	0.1	2	0	2
	400		lin	0.2	0.3	0.2	2	0	2	
			steep	69.0	71.9	68.8	87	7	80	
	Ret. Rev.		150	lin	0.2	0.4	0.1	2	0	2
				steep	0.3	0.5	0.2	2	0	2
		400	lin	0.7	1.0	0.7	2	0	2	
			steep	69.0	71.9	68.8	87	7	80	
		Ret. SME	150	lin	0.3	1.9	0.1	2	0	2
				steep	0.4	2.3	0.3	2	0	2
	400		lin	1.2	3.3	1.0	2	0	2	
			steep	69.0	71.9	68.8	87	7	80	
	Comm. RE		150	lin	0.5	3.8	0.3	2	0	2
				steep	0.9	4.3	0.6	2	0	2
400		lin	2.2	5.9	1.9	2	0	2		
		steep	69.0	71.9	68.8	87	7	80		
Comm. RE		150	lin	0.3	2.4	0.1	2	0	2	
			steep	0.5	2.7	0.3	2	0	2	
	400	lin	1.2	3.7	1.0	2	0	2		
		steep	69.0	71.9	68.8	87	7	80		
	Comm. RE	150	lin	0.6	4.4	0.3	2	0	2	
			steep	0.9	4.9	0.6	2	0	2	
400		lin	2.3	6.4	1.9	3	0	3		
		steep	69.0	71.9	68.8	87	7	80		

Table B.4: Effect of a shock that originates in the Benelux countries. There are a total of 7 banks in the Benelux countries. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in the Benelux countries (Loss Benelux) as well as the average loss in tier 1 capital ratio for banks in other than the Benelux countries (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for the Benelux countries (D_{100} Benelux) and for other than the Benelux countries (D_{100} Other).

Origin	Sector	Shock Size (%)	Q_A	$P(x)$	Loss	Loss	Loss	D_{100}	D_{100}	D_{100}		
					Avg	Greece	Other	Total	Greece	Other		
Greece	Sv. Debt	150	0.3	lin	0.0	0.0	0.0	2	0	2		
				steep	0.0	0.1	0.0	2	0	2		
			0.9	lin	0.0	0.1	0.0	2	0	2		
				steep	69.0	68.8	69.0	87	5	82		
			400	0.3	lin	0.0	0.4	0.0	2	0	2	
					steep	0.0	0.4	0.0	2	0	2	
		0.9		lin	0.0	0.5	0.0	2	0	2		
				steep	69.0	69.0	69.0	87	5	82		
		Financial		150	0.3	lin	0.5	4.2	0.3	2	0	2
						steep	0.7	5.1	0.4	2	0	2
			0.9		lin	1.3	6.8	0.9	2	0	2	
					steep	69.0	68.8	69.0	87	5	82	
	400		0.3		lin	2.7	19.8	1.4	2	0	2	
					steep	3.3	22.1	2.0	3	1	2	
			0.9	lin	5.8	26.9	4.3	4	1	3		
				steep	69.0	68.8	69.0	87	5	82		
			Corporate	150	0.3	lin	1.7	16.0	0.7	2	0	2
						steep	2.2	19.0	1.0	2	0	2
	0.9				lin	4.1	25.4	2.6	2	0	2	
					steep	69.0	68.8	69.0	87	5	82	
	400	0.3			lin	4.3	38.1	1.9	5	2	3	
					steep	5.1	39.7	2.6	5	2	3	
		0.9		lin	8.2	44.1	5.6	6	2	4		
				steep	69.0	68.8	69.0	87	5	82		
		Ret. Resid.		150	0.3	lin	0.8	10.0	0.2	2	0	2
						steep	1.1	11.4	0.3	2	0	2
	0.9				lin	1.7	14.1	0.8	2	0	2	
					steep	69.0	68.8	69.0	87	5	82	
	400		0.3		lin	3.5	38.8	1.0	3	1	2	
					steep	4.4	42.5	1.7	3	1	2	
			0.9	lin	7.4	48.8	4.4	4	1	3		
				steep	69.0	68.8	69.0	87	5	82		
			Ret. Rev.	150	0.3	lin	0.5	6.4	0.0	2	0	2
						steep	0.6	7.4	0.1	2	0	2
	0.9				lin	0.8	8.9	0.3	2	0	2	
					steep	69.0	68.8	69.0	87	5	82	
	400	0.3			lin	1.2	15.9	0.1	2	0	2	
					steep	1.4	17.1	0.2	2	0	2	
		0.9		lin	2.0	19.2	0.7	2	0	2		
				steep	69.0	68.8	69.0	87	5	82		
		Ret. SME		150	0.3	lin	0.3	2.9	0.1	2	0	2
						steep	0.4	3.6	0.2	2	0	2
	0.9				lin	0.7	5.0	0.4	2	0	2	
					steep	69.0	68.8	69.0	87	5	82	
	400		0.3		lin	0.5	5.2	0.2	2	0	2	
					steep	0.7	6.2	0.3	2	0	2	
			0.9	lin	1.2	8.0	0.7	2	0	2		
				steep	69.0	68.8	69.0	87	5	82		
Comm. RE			150	0.3	lin	0.6	6.0	0.2	2	0	2	
					steep	0.8	7.1	0.3	2	0	2	
	0.9			lin	1.2	9.3	0.7	2	0	2		
				steep	69.0	68.8	69.0	87	5	82		
	400	0.3		lin	1.0	10.7	0.4	2	0	2		
				steep	1.3	12.1	0.5	2	0	2		
		0.9	lin	2.0	14.6	1.1	2	0	2			
			steep	69.0	68.8	69.0	87	5	82			

Table B.5: Effect of a shock that originates in Greece. There are a total of 6 banks in Greece. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Greece (Loss Greece) as well as the average loss in tier 1 capital ratio for banks in other than Greece (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Greece (D_{100} Greece) and for other than Greece (D_{100} Other).

Origin	Sector	Shock Size (%)	Q_A	$P(x)$	Loss Avg	Loss Italy	Loss Other	D_{100} Total	D_{100} Italy	D_{100} Other	
Italy	Sv. Debt	150	0.3	lin	0.0	0.0	0.0	2	0	2	
			steep	0.0	0.0	0.0	2	0	2		
		400	0.3	lin	0.0	0.1	0.0	2	0	2	
			steep	0.0	0.1	0.0	2	0	2		
		150	0.9	lin	0.0	0.0	0.0	2	0	2	
			steep	0.0	0.0	0.0	2	0	2		
		400	0.9	lin	0.0	0.2	0.0	2	0	2	
			steep	69.0	66.1	69.2	87	5	82		
		Financial	150	0.3	lin	0.9	5.2	0.7	2	0	2
				steep	1.6	6.7	1.3	2	0	2	
			400	0.3	lin	4.1	10.5	3.8	3	0	3
				steep	69.0	66.1	69.2	87	5	82	
	150		0.9	lin	4.4	22.9	3.3	3	0	3	
			steep	6.5	25.9	5.3	3	0	3		
	400		0.9	lin	14.7	33.7	13.6	7	3	4	
			steep	69.0	66.1	69.2	87	5	82		
	Corporate		150	0.3	lin	2.2	18.8	1.2	2	0	2
				steep	3.8	22.4	2.7	5	2	3	
			400	0.3	lin	11.2	32.3	10.0	5	2	3
				steep	69.0	66.1	69.2	87	5	82	
		150	0.9	lin	5.9	43.1	3.7	7	4	3	
			steep	8.6	44.5	6.5	7	4	3		
		400	0.9	lin	18.7	48.8	16.9	9	4	5	
			steep	69.0	66.1	69.2	87	5	82		
		Ret. Resid.	150	0.3	lin	0.8	5.3	0.5	2	0	2
				steep	1.3	6.5	1.0	2	0	2	
			400	0.3	lin	3.3	9.5	3.0	3	0	3
				steep	69.0	66.1	69.2	87	5	82	
	150		0.9	lin	4.1	25.0	2.9	3	1	2	
			steep	6.6	29.2	5.3	5	2	3		
	400		0.9	lin	16.7	39.7	15.3	8	3	5	
			steep	69.0	66.1	69.2	87	5	82		
	Ret. Rev.		150	0.3	lin	0.1	0.2	0.0	2	0	2
				steep	0.1	0.3	0.1	2	0	2	
			400	0.3	lin	0.4	0.6	0.3	2	0	2
				steep	69.0	66.1	69.2	87	5	82	
		150	0.9	lin	0.2	0.7	0.2	2	0	2	
			steep	0.4	1.0	0.4	2	0	2		
		400	0.9	lin	1.1	1.9	1.1	2	0	2	
			steep	69.0	66.1	69.2	87	5	82		
		Ret. SME	150	0.3	lin	0.6	6.7	0.3	2	0	2
				steep	1.1	8.0	0.7	2	0	2	
			400	0.3	lin	2.9	11.2	2.4	3	0	3
				steep	69.0	66.1	69.2	87	5	82	
	150		0.9	lin	1.2	12.1	0.5	2	0	2	
			steep	1.9	13.5	1.2	2	0	2		
	400		0.9	lin	4.5	16.8	3.8	3	0	3	
			steep	69.0	66.1	69.2	87	5	82		
Comm. RE	150		0.3	lin	0.8	7.4	0.4	2	0	2	
			steep	1.3	8.7	0.8	2	0	2		
	400		0.3	lin	3.2	12.0	2.7	3	0	3	
			steep	69.0	66.1	69.2	87	5	82		
	150	0.9	lin	1.5	13.1	0.8	2	0	2		
		steep	2.2	14.5	1.4	2	0	2			
	400	0.9	lin	5.0	17.8	4.2	3	0	3		
		steep	69.0	66.1	69.2	87	5	82			

Table B.6: Effect of a shock that originates in Italy. There are a total of 6 banks in Italy. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Italy (Loss Italy) as well as the average loss in tier 1 capital ratio for banks in other than Italy (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Italy (D_{100} Italy) and for other than Italy (D_{100} Other).

Origin	Sector	Shock Size (%)	Q_A	$P(x)$	Loss Avg	Loss France	Loss Other	D_{100} Total	D_{100} France	D_{100} Other		
France	Sv. Debt	150	0.3	lin	0.0	0.0	0.0	2	0	2		
			steep	0.0	0.0	0.0	2	0	2			
		400	0.3	lin	0.0	0.0	0.0	2	0	2		
			steep	0.0	0.1	0.0	2	0	2			
		Financial	150	0.3	lin	0.3	1.6	0.3	2	0	2	
				steep	0.6	2.0	0.6	2	0	2		
			400	0.3	lin	1.8	3.4	1.7	2	0	2	
				steep	69.0	68.1	69.1	87	4	83		
			Corporate	150	0.3	lin	1.9	8.8	1.5	2	0	2
					steep	3.3	10.6	3.0	3	0	3	
				400	0.3	lin	8.8	16.1	8.5	3	0	3
					steep	69.0	68.1	69.1	87	4	83	
	Ret. Resid.			150	0.3	lin	1.2	8.8	0.9	2	0	2
					steep	2.4	10.6	2.0	2	0	2	
				400	0.3	lin	7.6	16.7	7.2	3	0	3
					steep	69.0	68.1	69.1	87	4	83	
		Ret. Rev.		150	0.3	lin	3.7	23.8	2.8	2	0	2
					steep	6.2	25.9	5.3	4	0	4	
				400	0.3	lin	15.7	32.3	14.9	5	1	4
					steep	69.0	68.1	69.1	87	4	83	
			Ret. SME	150	0.3	lin	0.6	3.3	0.5	2	0	2
					steep	1.1	4.1	1.0	2	0	2	
				400	0.3	lin	3.2	6.7	3.1	3	0	3
					steep	69.0	68.1	69.1	87	4	83	
	Comm. RE			150	0.3	lin	3.3	15.9	2.7	2	0	2
					steep	5.7	19.5	5.1	3	0	3	
				400	0.3	lin	16.1	29.7	15.4	5	1	4
					steep	69.0	68.1	69.1	87	4	83	
		Ret. Rev.		150	0.3	lin	0.1	0.5	0.1	2	0	2
					steep	0.2	0.6	0.1	2	0	2	
				400	0.3	lin	0.5	1.0	0.5	2	0	2
					steep	69.0	68.1	69.1	87	4	83	
			Ret. Rev.	150	0.3	lin	0.3	1.4	0.2	2	0	2
					steep	0.5	1.8	0.4	2	0	2	
				400	0.3	lin	1.5	3.0	1.5	2	0	2
					steep	69.0	68.1	69.1	87	4	83	
	Ret. SME			150	0.3	lin	0.3	2.6	0.2	2	0	2
					steep	0.6	3.2	0.5	2	0	2	
				400	0.3	lin	1.9	4.8	1.8	2	0	2
					steep	69.0	68.1	69.1	87	4	83	
		Comm. RE		150	0.3	lin	0.6	4.9	0.4	2	0	2
					steep	1.1	5.7	0.9	2	0	2	
				400	0.3	lin	3.2	8.0	3.0	3	0	3
					steep	69.0	68.1	69.1	87	4	83	
			Comm. RE	150	0.3	lin	0.3	1.0	0.2	2	0	2
					steep	0.5	1.3	0.4	2	0	2	
				400	0.3	lin	1.4	2.4	1.3	2	0	2
					steep	69.0	68.1	69.1	87	4	83	
Comm. RE	150			0.3	lin	0.5	2.0	0.4	2	0	2	
				steep	0.9	2.5	0.8	2	0	2		
	400			0.3	lin	2.5	4.4	2.5	3	0	3	
				steep	69.0	68.1	69.1	87	4	83		

Table B.7: Effect of a shock that originates in France. There are a total of 4 banks in France. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in France (Loss France) as well as the average loss in tier 1 capital ratio for banks in other than France (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for France (D_{100} France) and for other than France (D_{100} Other).

Origin	Sector	Shock Size (%)	Q_A	$P(x)$	Loss Avg	Loss Germany	Loss Other	D_{100} Total	D_{100} Germany	D_{100} Other	
Germany	Sec. Debt	150	0.3	lin	0.0	0.0	0.0	2	0	2	
			steep	0.0	0.0	0.0	2	0	2		
		0.9	lin	0.0	0.0	0.0	2	0	2		
			steep	0.0	0.1	0.0	2	0	2		
		400	0.3	lin	0.0	0.2	0.0	2	0	2	
			steep	0.0	0.3	0.0	2	0	2		
		0.9	lin	0.1	0.3	0.1	2	0	2		
			steep	69.0	69.0	69.0	87	11	76		
		Financial	150	0.3	lin	2.5	14.0	0.7	3	1	2
				steep	3.7	16.6	1.7	3	1	2	
			0.9	lin	9.2	24.0	6.9	3	1	2	
				steep	69.0	68.9	69.0	87	11	76	
	400		0.3	lin	8.5	44.5	3.0	6	4	2	
			steep	10.9	45.8	5.6	6	4	2		
	0.9		lin	21.2	50.4	16.7	8	5	3		
			steep	69.0	68.9	69.0	87	11	76		
	Corporate		150	0.3	lin	3.0	13.5	1.4	3	1	2
				steep	4.9	16.5	3.1	3	1	2	
			0.9	lin	12.8	25.8	10.8	3	1	2	
				steep	69.0	68.9	69.0	87	11	76	
		400	0.3	lin	8.1	33.4	4.3	3	1	2	
			steep	11.3	35.7	7.5	3	1	2		
		0.9	lin	22.4	42.5	19.3	6	2	4		
			steep	69.0	68.9	69.0	87	11	76		
		Ret. Resid.	150	0.3	lin	0.8	1.6	0.6	2	0	2
				steep	1.3	2.3	1.1	2	0	2	
			0.9	lin	3.5	4.8	3.2	3	1	2	
				steep	69.0	68.9	69.0	87	11	76	
	400		0.3	lin	4.2	8.6	3.5	3	1	2	
			steep	6.9	12.2	6.1	3	1	2		
	0.9		lin	18.1	24.4	17.1	4	1	3		
			steep	69.0	68.9	69.0	87	11	76		
	Ret. Rev.		150	0.3	lin	0.1	0.2	0.1	2	0	2
				steep	0.1	0.3	0.1	2	0	2	
			0.9	lin	0.4	0.6	0.4	2	0	2	
				steep	69.0	68.9	69.0	87	11	76	
		400	0.3	lin	0.3	0.6	0.2	2	0	2	
			steep	0.5	0.9	0.4	2	0	2		
		0.9	lin	1.3	1.7	1.2	2	0	2		
			steep	69.0	68.9	69.0	87	11	76		
		Ret. SME	150	0.3	lin	0.2	0.3	0.2	2	0	2
				steep	0.4	0.5	0.4	2	0	2	
			0.9	lin	1.2	1.4	1.1	2	0	2	
				steep	69.0	68.9	69.0	87	11	76	
	400		0.3	lin	0.4	0.5	0.4	2	0	2	
			steep	0.8	1.0	0.8	2	0	2		
	0.9		lin	2.2	2.6	2.1	3	1	2		
			steep	69.0	68.9	69.0	87	11	76		
Comm. RE	150		0.3	lin	1.2	5.7	0.5	3	1	2	
			steep	1.8	7.0	1.0	3	1	2		
	0.9		lin	4.5	10.5	3.6	3	1	2		
			steep	69.0	68.9	69.0	87	11	76		
	400	0.3	lin	2.1	10.4	0.9	3	1	2		
		steep	3.1	11.8	1.8	3	1	2			
	0.9	lin	6.8	15.7	5.5	3	1	2			
		steep	69.0	68.9	69.0	87	11	76			

Table B.8: Effect of a shock that originates in Germany. There are a total of 12 banks in Germany. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Germany (Loss Germany) as well as the average loss in tier 1 capital ratio for banks in other than Germany (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Germany (D_{100} Germany) and for other than Germany (D_{100} Other).

Origin	Sector	Shock Size (%)	Q_A	$P(x)$	Loss Avg	Loss Spain	Loss Other	D_{100} Total	D_{100} Spain	D_{100} Other		
Spain	Sv. Debt	150	lin		0.0	0.0	0.0	2	1	1		
			steep		0.0	0.0	0.0	2	1	1		
		400	lin		0.0	0.1	0.0	2	1	1		
			steep		0.0	0.2	0.0	2	1	1		
		150	lin	0.3		0.1	0.2	0.0	2	1	1	
			steep	0.3		69.0	69.3	68.9	87	24	63	
		Financial	150	lin			1.4	4.2	0.3	2	1	1
				steep			2.0	5.7	0.6	2	1	1
			400	lin	0.3		4.2	9.4	2.1	3	1	2
				steep	0.3		69.0	69.3	68.9	87	24	63
			150	lin	0.9		6.6	19.3	1.7	3	2	1
				steep	0.9		8.5	22.9	2.9	5	3	2
	400		lin	0.3		15.0	31.7	8.6	9	7	2	
			steep	0.3		69.0	69.3	68.9	87	24	63	
	Corporate		150	lin			5.2	17.1	0.6	3	2	1
				steep			7.3	22.1	1.6	5	4	1
			400	lin	0.3		15.4	35.5	7.7	11	9	2
				steep	0.3		69.0	69.3	68.9	87	24	63
		150	lin	0.9		12.2	39.0	1.9	11	10	1	
			steep	0.9		14.1	41.3	3.6	12	10	2	
		400	lin	0.3		21.7	48.1	11.5	17	15	2	
			steep	0.3		69.0	69.3	68.9	87	24	63	
		Ret. Resid.	150	lin			3.8	12.9	0.3	2	1	1
				steep			5.1	16.2	0.8	3	2	1
			400	lin	0.3		9.7	24.7	3.9	5	3	2
				steep	0.3		69.0	69.3	68.9	87	24	63
	150		lin	0.9		14.0	46.7	1.4	16	15	1	
			steep	0.9		16.9	52.0	3.5	21	19	2	
	400		lin	0.3		25.3	58.3	12.6	22	20	2	
			steep	0.3		69.0	69.3	68.9	87	24	63	
	Ret. Rev.		150	lin			0.9	3.2	0.0	2	1	1
				steep			1.2	4.1	0.2	2	1	1
			400	lin	0.3		2.2	6.0	0.7	2	1	1
				steep	0.3		69.0	69.3	68.9	87	24	63
		150	lin	0.9		2.5	8.7	0.1	2	1	1	
			steep	0.9		3.2	10.4	0.4	2	1	1	
		400	lin	0.3		5.3	14.1	1.9	3	2	1	
			steep	0.3		69.0	69.3	68.9	87	24	63	
		Ret. SME	150	lin			1.8	6.2	0.1	2	1	1
				steep			2.3	7.6	0.3	2	1	1
			400	lin	0.3		4.1	11.1	1.4	3	2	1
				steep	0.3		69.0	69.3	68.9	87	24	63
	150		lin	0.9		3.2	11.0	0.2	3	2	1	
			steep	0.9		3.9	12.7	0.5	3	2	1	
	400		lin	0.3		6.1	16.3	2.2	3	2	1	
			steep	0.3		69.0	69.3	68.9	87	24	63	
	Comm. RE		150	lin			1.9	6.4	0.2	2	1	1
				steep			2.5	7.9	0.4	2	1	1
400			lin	0.3		4.3	11.5	1.6	2	1	1	
			steep	0.3		69.0	69.3	68.9	87	24	63	
150		lin	0.9		3.4	11.5	0.4	2	1	1		
		steep	0.9		4.2	13.1	0.7	3	2	1		
400		lin	0.3		6.5	16.8	2.5	4	2	2		
		steep	0.3		69.0	69.3	68.9	87	24	63		

Table B.9: Effect of a shock that originates in Spain. There are a total of 25 banks in Spain. We simulate a sudden increase of risk weights to 150 percent and to 400 percent of their original value. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Spain (Loss Spain) as well as the average loss in tier 1 capital ratio for banks in other than Spain (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Spain (D_{100} Spain) and for other than Spain (D_{100} Other).

Origin	Shock Size	Q_A	$P(x)$	Loss Avg	Loss GI-IPS	Loss Other	D_{100} Total	D_{100} GI-IPS	D_{100} Other
GIIPS	20%	0.3	lin	12.1	24.0	1.2	5	5	0
			steep	16.1	29.5	3.8	11	10	1
		0.9	lin	31.9	46.9	18.2	27	25	2
			steep	72.0	75.2	69.1	88	42	46
	50%	0.3	lin	28.4	56.1	3.0	35	34	1
			steep	33.7	61.8	7.9	38	37	1
0.9		lin	57.6	79.4	37.7	47	42	5	
		steep	76.4	84.5	69.1	89	43	46	

Table B.10: Effect of a shock that originates in the GIIPS countries. There are a total of 43 banks from these countries. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in GIIPS countries (Loss GIIPS) as well as the average loss in tier 1 capital ratio for banks in other than GIIPS countries (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for the GIIPS (D_{100} GIIPS) and for other than GIIPS (D_{100} Other).

Origin	Shock Size	Q_A	$P(x)$	Loss Avg	Loss EE	Loss Other	D_{100} Total	D_{100} EE	D_{100} Other
Eastern Europe	20%	0.3	lin	1.0	21.2	0.0	3	1	2
			steep	1.1	22.5	0.1	3	1	2
		0.9	lin	1.4	24.8	0.3	3	1	2
			steep	69.3	74.3	69.1	87	4	83
	50%	0.3	lin	2.4	52.0	0.0	4	2	2
			steep	2.5	53.2	0.1	4	2	2
		0.9	lin	3.4	57.7	0.9	4	2	2
			steep	69.7	84.0	69.1	87	4	83

Table B.11: Effect of a shock that originates in the Eastern European countries. There are a total of 4 banks from these countries. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Eastern Europe (Loss EE) as well as the average loss in tier 1 capital ratio for banks in other than Eastern European countries (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for the Eastern Europe (D_{100} EE) and for other than Eastern European countries (D_{100} Other).

Origin	Shock Size	Q_A	$P(x)$	Loss Avg	Loss Benelux	Loss Other	D_{100} Total	D_{100} Benelux	D_{100} Other
Benelux	20%	0.3	lin	2.2	21.8	0.6	2	0	2
			steep	3.5	24.1	1.8	3	0	3
		0.9	lin	10.1	31.4	8.3	3	0	3
			steep	69.5	77.5	68.8	87	7	80
	50%	0.3	lin	5.4	52.8	1.4	3	1	2
			steep	7.8	55.3	3.8	6	3	3
		0.9	lin	24.3	67.3	20.7	14	7	7
			steep	70.1	85.9	68.8	87	7	80

Table B.12: Effect of a shock that originates in the Benelux countries. There are a total of 7 banks from these countries. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Benelux countries (Loss Benelux) as well as the average loss in tier 1 capital ratio for banks in other than Benelux countries (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for the Benelux countries (D_{100} Benelux) and for other than Benelux countries (D_{100} Other).

Origin	Shock Size	Q_A	$P(x)$	Loss Avg	Loss Greece	Loss Other	D_{100} Total	D_{100} Greece	D_{100} Other
Greece	20%	0.3	lin	1.6	22.5	0.1	2	0	2
			steep	2.0	25.3	0.3	2	0	2
		0.9	lin	3.2	31.2	1.2	3	1	2
			steep	69.4	75.0	69.0	88	6	82
	50%	0.3	lin	3.8	53.8	0.2	5	3	2
			steep	4.3	56.8	0.6	5	3	2
		0.9	lin	7.6	67.1	3.4	8	5	3
			steep	70.1	84.4	69.0	88	6	82

Table B.13: Effect of a shock that originates in Greece. There are a total of 6 banks in Greece. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Greece (Loss Greece) as well as the average loss in tier 1 capital ratio for banks in other than Greece (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Greece (D_{100} Greece) and for other than Greece (D_{100} Other).

Origin	Shock Size	Q_A	$P(x)$	Loss Avg	Loss Italy	Loss Other	D_{100} Total	D_{100} Italy	D_{100} Other
Italy	20%	0.3	lin	1.6	22.7	0.4	4	2	2
			steep	2.7	26.0	1.4	4	2	2
		0.9	lin	8.3	34.9	6.8	6	3	3
			steep	69.4	72.9	69.2	87	5	82
	50%	0.3	lin	4.0	54.2	1.1	7	5	2
			steep	5.9	57.7	2.9	8	5	3
		0.9	lin	20.9	71.7	18.0	10	5	5
			steep	70.0	83.1	69.2	87	5	82

Table B.14: Effect of a shock that originates in Italy. There are a total of 6 banks in Italy. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Italy (Loss Italy) as well as the average loss in tier 1 capital ratio for banks in other than Italy (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Italy (D_{100} Italy) and for other than Italy (D_{100} Other).

Origin	Shock Size	Q_A	$P(x)$	Loss Avg	Loss France	Loss Other	D_{100} Total	D_{100} France	D_{100} Other
France	20%	0.3	lin	1.7	21.7	0.8	2	0	2
			steep	3.4	24.1	2.4	3	0	3
		0.9	lin	12.0	32.8	11.0	4	0	4
			steep	69.3	74.5	69.1	87	4	83
	50%	0.3	lin	4.2	52.6	1.9	6	4	2
			steep	7.3	55.3	5.0	7	4	3
		0.9	lin	28.3	68.6	26.4	12	4	8
			steep	69.7	84.0	69.1	87	4	83

Table B.15: Effect of a shock that originates in France. There are a total of 4 banks in France. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in France (Loss France) as well as the average loss in tier 1 capital ratio for banks in other than France (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for France (D_{100} France) and for other than France (D_{100} Other).

Origin	Shock Size	Q_A	$P(x)$	Loss Avg	Loss Ger-many	Loss Other	D_{100} Total	D_{100} Ger-many	D_{100} Other
Germany	20%	0.3	lin	3.6	22.5	0.6	3	1	2
			steep	5.3	25.8	2.1	3	1	2
		0.9	lin	13.6	35.8	10.2	3	1	2
			steep	69.9	75.2	69.0	87	11	76
	50%	0.3	lin	8.6	53.9	1.6	8	6	2
			steep	11.5	57.4	4.4	9	7	2
	0.9	lin	31.8	72.8	25.5	18	11	7	
		steep	71.1	84.5	69.0	87	11	76	

Table B.16: Effect of a shock that originates in Germany. There are a total of 12 banks in Germany. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Germany (Loss Germany) as well as the average loss in tier 1 capital ratio for banks in other than Germany (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Germany (D_{100} Germany) and for other than Germany (D_{100} Other).

Origin	Shock Size	Q_A	$P(x)$	Loss Avg	Loss Spain	Loss Other	D_{100} Total	D_{100} Spain	D_{100} Other
Spain	20%	0.3	lin	7.0	24.3	0.4	3	2	1
			steep	9.3	30.0	1.4	7	6	1
		0.9	lin	18.8	45.7	8.4	17	15	2
			steep	70.7	75.4	68.9	87	24	63
	50%	0.3	lin	16.4	56.7	0.9	21	20	1
			steep	19.5	62.5	3.0	23	21	2
	0.9	lin	37.4	80.2	21.0	29	24	5	
		steep	73.3	84.6	68.9	88	25	63	

Table B.17: Effect of a shock that originates in Spain. There are a total of 25 banks in Spain. We simulate a sudden drop of bank equity by 20% and by 50%. We show the average loss in tier 1 capital ratio across the entire system (Loss Avg), the average loss in tier 1 capital ratio for banks in Spain (Loss Spain) as well as the average loss in tier 1 capital ratio for banks in other than Spain (Loss Other). We also show the number of banks failing the stress test by having their tier 1 capital ratio drop below 4.5 percent after 100 time steps for the entire system (D_{100} Total), for Spain (D_{100} Spain) and for other than Spain (D_{100} Other).

Appendix C

Supplementary Material to Chapter 5

C.1 Macroeconomic data used for analysis

Country	ISO code	FX symbol	Gov't bond	ETF ticker	Timezone	Classification
Australia	AUS	AUD	AU	EWA	Pacific	Comm.
Austria	AUT	EUR	EZ	EZU	Europe	Res.
Belgium	BEL	EUR	EZ	EZU	Europe	Res.
Canada	CAN	CAD	CA	EWC	Americas	Comm.
Cyprus	CYP	EUR	EZ	EZU	Europe	Res.
Estonia	EST	EUR	EZ	EZU	Europe	Res.
Finland	FIN	EUR	EZ	EZU	Europe	Res.
France	FRA	EUR	EZ	EZU	Europe	Res.
Germany	DEU	EUR	EZ	EZU	Europe	Res.
Greece	GRC	EUR	EZ	EZU	Europe	Res.
Ireland	IRL	EUR	EZ	EZU	Europe	Res.
Italy	ITA	EUR	EZ	EZU	Europe	Res.
Japan	JPN	JPY	JP	EWJ	Asia	Res.
Latvia	LVA	EUR	EZ	EZU	Europe	Res.
Lithuania	LTU	EUR	EZ	EZU	Europe	Res.
Luxembourg	LUX	EUR	EZ	EZU	Europe	Res.
Malta	MLT	EUR	EZ	EZU	Europe	Res.
Mexico	MEX	MXN	MX	EWX	Americas	—
Netherlands	NLD	EUR	EZ	EZU	Europe	Res.
New Zealand	NZL	NZD	NZ	ENZL	Pacific	—
Norway	NOR	NOK	NO	ENOR	Europe	Comm.
Portugal	PRT	EUR	EZ	EZU	Europe	Res.
Slovakia	SVK	EUR	EZ	EZU	Europe	Res.
Slovenia	SVN	EUR	EZ	EZU	Europe	Res.
South Africa	ZAF	ZAR	ZA	EZA	Africa	Comm.
Spain	ESP	EUR	EZ	EZU	Europe	Res.
Sweden	SWE	SEK	SE	EWD	Europe	—
Switzerland	CHE	CHF	CH	EWL	Europe	Res.
UK	GBR	GBP	GB	EWU	Europe	Res.
USA	USA	USD	US	SPY	Americas	Res.

Table C.1: Lookup table for the macroeconomic indicators of countries in our data set and classifications into commodity (Comm.) or reserve (Res.) currency.

References

- Daron Acemoglu, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012.
- Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *The American Economic Review*, 105(2):564–608, 2015. doi: 10.1257/aer.20130456.
- Viral Acharya, Robert Engle, and Matthew Richardson. Capital shortfall: A new approach to ranking and regulating systemic risks. *The American Economic Review*, 102(3):59–64, 2012. doi: 10.1257/aer.102.3.59.
- Viral Acharya, Robert Engle, and Diane Pierret. Testing macroprudential stress tests: The risk of regulatory risk weights. *Journal of Monetary Economics*, 65: 36–53, 2014. doi: 10.1016/j.jmoneco.2014.04.014.
- Anat Admati, Peter DeMarzo, Martin Hellwig, and Paul Pfleiderer. Fallacies. *Irrelevant Facts, and Myths in the Discussion of Capital Regulation: Why Bank Equity is Not Expensive*, 2011.
- Anat R Admati and Martin F Hellwig. Good Banking Regulation Needs Clear Focus, Sensible Tools, and Political Will. *International Centre for Financial Regulation Research Paper*, 2011.
- JA Aitchison. The statistical analysis of compositional data. *Journal of the Royal Statistical Society, Series B. (Statistical Methodology)*, pages 139–177, 1982. doi: 10.1007/978-94-009-4109-0.
- JA Aitchison, Carles Barceló-Vidal, Josep A Martín-Fernández, and Vera Pawlowsky-Glahn. Logratio analysis and compositional distance. *Mathematical Geology*, 32(3):271–275, 2000.
- Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47, 2002. doi: 10.1103/revmodphys.74.47.
- Bill Allen, Ka Kei Chan, Alistair Milne, and Steve Thomas. Basel III: Is the cure worse than the disease? *International Review of Financial Analysis*, 25:159–166, 2012. doi: 10.1016/j.irfa.2012.08.004.

- Richard T Baillie and William P Osterberg. Why do central banks intervene? *Journal of International Money and Finance*, 16(6):909–919, 1997b. doi: 10.1016/s0261-5606(97)00012-0.
- European Central Bank. Aggregate Report on the Comprehensive Assessment. *Frankfurt am Main*, 2014.
- Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999. doi: 10.1515/9781400841356.349.
- Basel Committee on Banking Supervision. Basel III: A global regulatory framework for more resilient banks and banking systems. Technical report, Bank for International Settlements, 2010.
- Stefano Battiston, Domenico Delli Gatti, Mauro Gallegati, Bruce Greenwald, and Joseph E Stiglitz. Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. *Journal of Economic Dynamics and Control*, 36(8):1121–1141, 2012a. doi: 10.1016/j.jedc.2012.04.001.
- Stefano Battiston, Michelangelo Puliga, Rahul Kaushik, Paolo Tasca, and Guido Caldarelli. DebtRank: Too Central to Fail? Financial Networks, the FED and Systemic Risk. *Scientific Reports*, 2, 2012b. doi: 10.1038/srep00541.
- Stefano Battiston, J. Doyne Farmer, Andreas Flache, Diego Garlaschelli, Andrew G Haldane, Hans Heesterbeek, Cars Hommes, Carlo Jaeger, Robert May, and Marten Scheffer. Complexity theory and financial regulation. *Science*, 351(6275):818–819, 2016. doi: 10.1126/science.aad0299.
- BBC. <http://www.bbc.co.uk/news/business-14398392>, Apr. 2011.
- Michel Beine, Agnes Bénassy-Quéré, and Christelle Lecourt. Central bank intervention and foreign exchange rates: new evidence from FIGARCH. estimations. *Journal of International Money and Finance*, 21:115–144, 2002. doi: 10.1016/s0261-5606(01)00040-7.
- Michel Beine, Sébastien Laurent, and Christelle Lecourt. Official central bank interventions and exchange rate volatility: Evidence from a regime-switching analysis. *European Economic Review*, 47:891–911, 2003. doi: 10.1016/s0014-2921(02)00306-9.
- Michel Beine, Sébastien Laurent, and FC Palm. Central bank FOREX. interventions assessed using realized moments. *Journal of International Financial Markets, Institutions & Money*, 19:112–127, 2009. doi: 10.1016/j.intfin.2007.09.001.
- Geert Bekaert, Robert J Hodrick, and Xiaoyan Zhang. International stock return comovements. *The Journal of Finance*, 64(6):2591–2626, 2009. doi: 10.1111/j.1540-6261.2009.01512.x.

- Andrea Beltratti and René M Stulz. The credit crisis around the globe: Why did some banks perform better? *Journal of Financial Economics*, 105(1):1–17, 2012. doi: 10.1016/j.jfineco.2011.12.005.
- Monica Billio, Mila Getmansky, Andrew W Lo, and Lioriana Pelizzon. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559, 2012. doi: 10.1016/j.jfineco.2011.12.010.
- BIS Monetary and Economic Department. Triennial Central Bank Survey of foreign exchange turnover in April 2013. Technical report, Bank for International Settlements, 2013.
- BIS Monetary and Economic Department. Triennial Central Bank Survey of foreign exchange turnover in April 2016. Technical report, Bank for International Settlements, 2016.
- Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/1831029>.
- Richard Blackden. BRICs attack QE and urge Western leaders to be 'responsible'. <http://www.telegraph.co.uk/finance/economics/9174292/BRICs-attack-QE-and-urge-Western-leaders-to-be-responsible.html>, Mar. 2012.
- Giovanni Bonanno, Nicolas Vandewalle, and Rosario N Mantegna. Taxonomy of stock market indices. *Physical Review E*, 62(6):R7615, 2000. doi: 10.1103/physreve.62.r7615.
- Giovanni Bonanno, Guido Caldarelli, Fabrizio Lillo, and Rosario N Mantegna. Topology of correlation-based minimal spanning trees in real and model markets. *Physical Review E*, 68(4):046130, 2003. doi: 10.1103/physreve.68.046130.
- Giovanni Bonanno, Guido Caldarelli, Fabrizio Lillo, F Miccichè, Nicolas Vandewalle, and Rosario N Mantegna. Networks of equities in financial markets. *The European Physical Journal B*, 38(2):363–371, 2004. doi: 10.1140/epjb/e2004-00129-6.
- C. Bonser-Neal and G. Tanner. Central bank intervention and the volatility of foreign exchange rates: evidence from the options market. *Journal of International Money and Finance*, 15(6):853–878, 1996. doi: 10.1016/s0261-5606(96)00033-2.
- Brian H Boyer, Tomomi Kumagai, and Kathy Yuan. How do crises spread? Evidence from accessible and inaccessible stock indices. *The journal of finance*, 61(2):957–1003, 2006. doi: 10.1111/j.1540-6261.2006.00860.x.

- William Breen, Lawrence R Glosten, and Ravi Jagannathan. Economic significance of predictable variations in stock index returns. *The Journal of Finance*, 44(5): 1177–1189, 1989. doi: 10.2307/2328638.
- TC Brownlees and Robert F Engle. SRISK: A conditional capital shortfall index for systemic risk measurement. *Department of Finance, New York University (Draft on Jan. 2015)*, 2015.
- Sergey V Buldyrev, Roni Parshani, Gerald Paul, H Eugene Stanley, and Shlomo Havlin. Catastrophic cascade of failures in interdependent networks. *Nature*, 464 (7291):1025–1028, 2010. doi: 10.1038/nature08932.
- Fabio Caccioli, Munik Shrestha, Christopher Moore, and J Doyne Farmer. Stability analysis of financial contagion due to overlapping portfolios. *Journal of Banking & Finance*, 46:233–245, 2014a. doi: 10.1016/j.jbankfin.2014.05.021.
- Fabio Caccioli, Munik Shrestha, Cristopher Moore, and J Doyne Farmer. Stability analysis of financial contagion due to overlapping portfolios. *Journal of Banking & Finance*, 46:233–245, 2014b.
- Guido Caldarelli. *Scale-free networks: complex webs in nature and technology*. Oxford University Press, 2007.
- Giovanni Calice, Jing Chen, and Julian Williams. Liquidity spillovers in sovereign bond and CDS. markets: An analysis of the Eurozone sovereign debt crisis. *Journal of Economic Behavior & Organization*, 85:122–143, 2013. doi: 10.1016/j.jebo.2011.10.013.
- Lorenzo Cappiello and Roberto A De Santis. *Explaining exchange rate dynamics: The uncovered equity return parity condition*. Citeseer, 2005.
- Ariel Caticha, Gary J Erickson, and Yuxiang Zhai. Relative entropy and inductive inference. In *AIP. Conference Proceedings*, volume 707, pages 75–96. AIP, 2004. doi: 10.1063/1.1751358.
- Jorge A Chan-Lau. Market-Based Estimation of Default Probabilities and Its Application to Financial Market Surveillance. Technical report, IMF Working Paper, Monetary and Financial Systems Department, 2006.
- Yu-Chin Chen, Kenneth S. Rogoff, and Barbara Rossi. Can exchange rates forecast commodity prices?*. *The Quarterly Journal of Economics*, 125(3):1145–1194, 2010. doi: 10.1162/qjec.2010.125.3.1145. URL <http://dx.doi.org/10.1162/qjec.2010.125.3.1145>.
- Guilio Cimini, Tiziano Squartini, Diego Garlaschelli, and Andrea Gabrielli. Systemic risk analysis on reconstructed economic and financial networks. *Scientific Reports*, 5(15758), 2015.

- John H Cochrane. The dog that did not bark: A defense of return predictability. *Review of Financial Studies*, 21(4):1533–1575, 2008.
- Lauren Cohen and Andrea Frazzini. Economic links and predictable returns. *The Journal of Finance*, 63(4):1977–2011, 2008. doi: 10.1111/j.1540-6261.2008.01379.x.
- One Hundred Tenth Congress. This Act has been amended by the: DODD-FRANK. WALL. STREET. REFORM and. CONSUMER. PROTECTION. ACT. [PUBLIC. LAW. 111–203, JULY. 21, 2010] This PDF. will be replaced in a future Update to reflect changes. *Public Law*, 111:203, 2010.
- Rama Cont and Eric Schaanning. Fire sales, indirect contagion and systemic stress testing, 2017. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2541114.
- Fulvio Corsi, Stefano Marmi, and Fabrizio Lillo. When micro prudence increases macro risk: The destabilizing effects of financial innovation, leverage, and diversification. *Operations Research*, 64(5):1073–1088, 2016.
- Chester Curme, Michele Tumminello, Rosario N Mantegna, H Eugene Stanley, and Dror Y Kenett. Emergence of statistically validated financial intraday lead-lag relationships. *SSRN. 2373787*, 2014. doi: 10.2139/ssrn.2373787.
- Magnus Dahlquist and Henrik Hasseltoft. International bond risk premia. *Journal of International Economics*, 90(1):17–32, 2013. doi: 10.1016/j.jinteco.2012.11.008.
- Nima Dehmamy, Sergey V Buldyrev, Shlomo Havlin, H Eugene Stanley, and Irena Vodenska. Classical mechanics of economic networks. *arXiv preprint arXiv:1410.0104*, 2014a.
- Nima Dehmamy, Sergey V Buldyrev, Shlomo Havlin, H Eugene Stanley, and Irena Vodenska. A systemic stress test model in bank-asset networks. *arXiv preprint*, 2014b.
- Robert DeYoung and Karen Y Jang. Do banks actively manage their liquidity? *Journal of Banking & Finance*, 66:143–161, 2016. doi: 10.1016/j.jbankfin.2015.11.013.
- Tiziana Di Matteo, F Pozzi, and Tomaso Aste. The use of dynamical networks to detect the hierarchical organization of financial market sectors. *The European Physical Journal B*, 73(1):3–11, 2010.
- Isabelle Distinguin, Caroline Roulet, and Amine Tarazi. Bank regulatory capital and liquidity: Evidence from US and European publicly traded banks. *Journal of Banking & Finance*, 37(9):3295–3317, 2013. doi: 10.1016/j.jbankfin.2013.04.027.

- KME Dominguez. Central bank intervention and exchange rate volatility. *Journal of International Money and Finance*, 17:161–190, 1998. doi: 10.1016/s0261-5606(97)98055-4.
- KME Dominguez. When do central bank interventions influence intra-daily and longer-term exchange rate movements? *Journal of International Money and Finance*, 25:1051–1071, 2006. doi: 10.1016/j.jimonfin.2006.08.009.
- Michael Dooley and Peter Isard. A portfolio-balance rational-expectations model of the dollar-mark exchange rate. *Journal of International Economics*, 12(3):257–276, 1982. doi: 10.1016/0022-1996(82)90039-3.
- Rudiger Dornbusch and Stanley Fischer. Exchange rates and the current account. *The American Economic Review*, 70(5):960–971, 1980.
- Fernando Duarte and Thomas Eisenbach. Fire-sale spillovers and systemic risk. Technical report, FRB of New York Staff Report N0 645, 2015.
- Larry Eisenberg and Thomas H Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001. doi: 10.1287/mnsc.47.2.236.9835.
- Luci Ellis, Andy Haldane, and Fariborz Moshirian. Systemic risk, governance and global financial stability. *Journal of Banking & Finance*, 45:175–181, 2014. doi: 10.1016/j.jbankfin.2014.04.012.
- Helmut Elsinger, Alfred Lehar, and Martin Summer. Risk assessment for banking systems. *Management science*, 52(9):1301–1314, 2006. doi: 10.1287/mnsc.1060.0531.
- Robert Engle. *Anticipating correlations: a new paradigm for risk management*. Princeton University Press, 2009. doi: 10.1515/9781400830190.
- Eugene F Fama and Kenneth R French. Dividend yields and expected stock returns. *Journal of Financial Economics*, 22(1):3–25, 1988. doi: 10.1016/0304-405x(88)90020-7.
- Eugene F Fama and Kenneth R French. Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics*, 25(1):23–49, 1989. doi: 10.1016/0304-405x(89)90095-0.
- Rasmus Fatum. Daily effects of foreign exchange intervention: Evidence from official Bank of Canada data. *Journal of International Money and Finance*, 27:438–454, 2008. doi: 10.1016/j.jimonfin.2008.01.003.
- Rasmus Fatum and Michael M Hutchison. ECB. Foreign Exchange Intervention and the EURO: Institutional Framework, News, and Intervention. *Open economies review*, 13:413–425, 2002. doi: 10.2139/ssrn.304519.

- Rasmus Fatum and Michael M Hutchison. Is sterilised foreign exchange intervention effective after all? An event study approach. *The Economic Journal*, 113:390–411, 2003. doi: 10.1111/1468-0297.00122.
- Federal Reserve Bank of St. Louis. FRED, Gross Domestic Product by Industry. <https://fred.stlouisfed.org/release/tables?rid=331&eid=211>, April 2018.
- Wayne E Ferson and Campbell R Harvey. The risk and predictability of international equity returns. *Review of Financial Studies*, 6(3):527–566, 1993. doi: 10.1093/rfs/6.3.527.
- Richard P Feynman, Robert B Leighton, and Matthew Linzee Sands. *The Feynman lectures on physics*, volume II of *Addison-Wesley World Student Series*. Addison-Wesley, Reading, MA, 1971.
- Thomas J Flavin, Margaret J Hurley, and Fabrice Rousseau. Explaining stock market correlation: A gravity model approach. *The Manchester School*, 70(S1):87–106, 2002. doi: 10.1111/1467-9957.70.s1.5.
- Kristin J Forbes and Roberto Rigobon. No contagion, only interdependence: measuring stock market comovements. *The Journal of Finance*, 57(5):2223–2261, 2002. doi: 10.1111/0022-1082.00494.
- Jeffrey A Frankel and Shang-Jin Wei. Is there a currency bloc in the Pacific? In *The Exchange Rate, International Trade and the Balance of Payments*, pages 275–307. Reserve Bank of Australia, 1993.
- Jeffrey A Frankel, Shang-Jin Wei, Matthew Canzoneri, and Morris Goldstein. *Emerging currency blocks*. Springer, 1995. doi: 10.1007/978-3-642-79681-4.6.
- B. Fuglede and F. Topsøe. Jensen-shannon divergence and hilbert space embedding. In *International Symposium on Information Theory, 2004*, pages 31–, June 2004. doi: 10.1109/ISIT.2004.1365067.
- Craig Furfine. Interbank exposures: Quantifying the risk of contagion. *Journal of Money, Credit, and Banking*, 35(1):111–128, 2003. doi: 10.1353/mcb.2003.0004.
- Louis Gagnon and G Andrew Karolyi. Price and volatility transmission across borders. *Financial Markets, Institutions & Instruments*, 15(3):107–158, 2006. doi: 10.1111/j.1468-0416.2006.00115.x.
- Prasanna Gai, Andrew Haldane, and Sujit Kapadia. Complexity, concentration and contagion. *Journal of Monetary Economics*, 58(5):453–470, 2011. doi: 10.1016/j.jmoneco.2011.05.005.

- Diego Garlaschelli and Maria I Loffredo. Fitness-dependent topological properties of the world trade web. *Physical Review Letters*, 93(18):188701, 2004. doi: 10.1103/physrevlett.93.188701.
- Diego Garlaschelli and Maria I Loffredo. Structure and evolution of the world trade network. *Physica A: Statistical Mechanics and its Applications*, 355(1):138–144, 2005. doi: 10.1016/j.physa.2005.02.075.
- Paul Glasserman and H Peyton Young. How likely is contagion in financial networks? *Journal of Banking & Finance*, 50:383–399, 2015. doi: 10.1016/j.jbankfin.2014.02.006.
- Parameswaran Gopikrishnan, Martin Meyer, Luis A. Nunes Amaral, and H Eugene Stanley. Inverse cubic law for the distribution of stock price variations. *The European Physical Journal B: Condensed Matter and Complex Systems*, 3(2):139–140, 1998. URL <https://EconPapers.repec.org/RePEc:spr:eurphb:v:3:y:1998:i:2:p:139-140>.
- Andrzej Z Górski, Stanislaw Drozd, and Jaroslaw Kwapien. Scale free effects in world currency exchange network. *The European Physical Journal B*, pages 91–96, 2008. doi: 10.1140/epjb/e2008-00376-5.
- Clive WJ Granger, Bwo-Nung Huangb, and Chin-Wei Yang. A bivariate causality between stock prices and exchange rates: evidence from recent Asianflu. *The Quarterly Review of Economics and Finance*, 40(3):337–354, 2000. doi: 10.1016/S1062-9769(00)00042-9.
- Robin Greenwood, Augustin Landier, and David Thesmar. Vulnerable banks. *Journal of Financial Economics*, 115(3):471–485, 2015.
- Les Gulko. The entropy theory of bond option pricing. *International Journal of Theoretical and Applied Finance*, 5(04):355–383, 2002. doi: 10.1142/S021902490200147x.
- Maurizio M Habib and Livio Stracca. Getting beyond carry trade: What makes a safe haven currency? *Journal of International Economics*, 87(1):50–64, 2012. doi: 10.1016/j.jinteco.2011.12.005.
- Andrew G Haldane and Robert M May. Systemic risk in banking ecosystems. *Nature*, 469(7330):351–355, 2011. doi: 10.1038/nature09659.
- Campbell R Harvey. The world price of covariance risk. *The Journal of Finance*, 46(1):111–157, 1991. doi: 10.2307/2328691.
- Oliver Hossfeld and Ronald MacDonald. Carry funding and safe haven currencies: A threshold regression approach. *Journal of International Money and Finance*, 59:185–202, 2015. doi: 10.1016/j.jimonfin.2015.07.005.

- Nikolai V Hovanov, James W Kolari, and Mikhail V Sokolov. Computing currency invariant indices with an application to minimum variance currency baskets. *Journal of Economic Dynamics and Control*, 28:1481–1504, 2004. doi: 10.1016/s0165-1889(02)00087-8.
- Xuqing Huang, Irena Vodenska, Fengzhong Wang, Shlomo Havlin, and H Eugene Stanley. Identifying influential directors in the United States corporate governance network. *Physical Review E*, 84(4):046101, 2011. doi: 10.1103/physreve.84.046101.
- Xuqing Huang, Irena Vodenska, Shlomo Havlin, and H Eugene Stanley. Cascading Failures in Bi-partite Graphs: Model for Systemic Risk Propagation. *Scientific Reports*, 3, 2013. doi: 10.1038/srep01219.
- Giulia Iori, Giulia De Masi, Ovidiu Vasile Precup, Giampaolo Gabbi, and Guido Caldarelli. A network analysis of the Italian overnight money market. *Journal of Economic Dynamics and Control*, 32(1):259–278, 2008. doi: 10.1016/j.jedc.2007.01.032.
- ET Jaynes. On the rationale of maximum-entropy methods. *Proceedings of the IEEE*, 70(9):939–952, Sept. 1982. ISSN 0018-9219. doi: 10.1109/PROC.1982.12425.
- Andreas Joseph, Irena Vodenska, Eugene Stanley, and Guanrong Chen. Netconomics: Novel Forecasting Techniques from the Combination of Big Data, Network Science and Economics. *arXiv preprint arXiv:1403.0848*, 2014.
- Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the econometric society*, pages 263–291, 1979. doi: 10.1142/9789814417358_0006.
- Georgios Katechos. On the relationship between exchange rates and equity returns: A new approach. *Journal of International Financial Markets, Institutions and Money*, 21(4):550–559, 2011. doi: 10.1016/j.intfin.2011.03.001.
- Dror Y Kenett, Yoash Shapira, A Madi, S Bransburg-Zabary, G Gur-Gershgoren, and Eshel Ben-Jacob. Dynamics of stock market correlations. *AUCO. Czech Economics Review*, 4(3):330–341, 2010.
- Dror Y Kenett, Yoash Shapira, A Madi, S Bransburg-Zabary, G Gur-Gershgoren, and Eshel Ben-Jacob. Index cohesive force analysis reveals that the US market became prone to systemic collapses since 2002. *PLoS ONE*, 6(4):e19378, 2011. doi: 10.1371/journal.pone.0019378.

- Dror Y Kenett, Tobias Preis, G Gur-Gershgoren, and Eshel Ben-Jacob. Dependency network and node influence: application to the study of financial markets. *International Journal of Bifurcation and Chaos*, 22(07), 2012a. doi: 10.1142/s0218127412501817.
- Dror Y Kenett, Matthias Raddant, Thomas Lux, and Eshel Ben-Jacob. Evolvement of uniformity and volatility in the stressed global financial village,. *PLoS ONE*, 7(2):e31144, 2012b. doi: 10.1371/journal.pone.0031144.
- Michael R King. The Basel III net stable funding ratio and bank net interest margins. *Journal of Banking & Finance*, 37(11):4144–4156, 2013. doi: 10.1016/j.jbankfin.2013.07.017.
- Peter King and Heath Tarbert. Basel III: an overview. *Banking & Financial Services Policy Report*, 30(5):1–18, 2011.
- Joseph B Kruskal. On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem. *Proceedings of the American Mathematical Society*, 7(1):48–50, 1956. ISSN 00029939, 10886826. doi: 10.2307/2033241.
- Jaroslav Kwapien, Sylwia Gworek, Stanislaw Drozd, and Andrzej Górski. Analysis of a network structure of the foreign currency exchange market. *Journal of Economic Interaction and Coordination*, 4:55–72, 2009. doi: 10.1007/s11403-009-0047-9.
- Laurent Laloux, Pierre Cizeau, Jean-Philippe Bouchaud, and Marc Potters. Noise dressing of financial correlation matrices. *Physical Review Letters*, 83(7):1467, 1999. doi: 10.1103/physrevlett.83.1467.
- Daqing Li, Bowen Fu, Yunpeng Wang, Guangquan Lu, Yehiel Berezin, H Eugene Stanley, and Shlomo Havlin. Percolation transition in dynamical traffic network with evolving critical bottlenecks. *Proceedings of the National Academy of Sciences*, 112(3):669–672, 2015.
- Chien-Hsiu Lin. The comovement between exchange rates and stock prices in the Asian emerging markets. *International Review of Economics & Finance*, 22(1):161–172, 2012. doi: 10.1016/j.iref.2011.09.006.
- Wen-Ling Lin, Robert F Engle, and Takatoshi Ito. Do Bulls and Bears Move Across Borders? International Transmission of Stock Returns and Volatility as the World Turns. *The Review of Financial Studies*, 7(3):507–538, 1994.
- Francois Longin and Bruno Solnik. Is the correlation in international equity returns constant: 1960–1990. *Journal of international money and finance*, 14(1):3–26, 1995. doi: 10.1016/0261-5606(94)00001-h.

- Roger Lowenstein. *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. Random House, 2000.
- Deborah Lucas. Comment on: “Testing macroprudential stress tests: The risk of regulatory risk weights” by Viral Acharya, Robert Engle, and Diane Pierret. *Journal of Monetary Economics*, 65:54–56, 2014a. doi: 10.1016/j.jmoneco.2014.05.003.
- Deborah Lucas. Evaluating the government as a source of systemic risk. *Journal of Financial Perspectives*, 2(3), 2014b.
- Antonio Majdandzic, Lidia A Braunstein, Chester Curme, Irena Vodenska, Sary Levy-Carciente, H Eugene Stanley, and Shlomo Havlin. Multiple tipping points and optimal repairing in interacting networks. *Nature Communications*, 7, 2016. doi: 10.1038/ncomms10850.
- Benoit Mandelbrot. The variation of certain speculative prices. *The Journal of Business*, 36(4):394–419, 1963. ISSN 00219398, 15375374. URL <http://www.jstor.org/stable/2350970>.
- Rosario N Mantegna. Hierarchical structure in financial markets. *The European Physical Journal B*, 11(1):193–197, 1999. doi: 10.1007/s100510050929.
- J. A Martín-Fernández, C Barceló-Vidal, and V Pawłowsky-Glahn. Measures of difference for compositional data and hierarchical clustering methods. In *Proceedings of IAMG*, volume 98, pages 526–531, 1998.
- Richard Meese. Currency fluctuations in the post-Bretton Woods era. *The Journal of Economic Perspectives*, 4(1):117–134, 1990. doi: 10.1257/jep.4.1.117.
- David Miles, Jing Yang, and Gilberto Marcheggiano. Optimal bank capital. *The Economic Journal*, 123(567):1–37, 2013.
- Bruce Morley. Exchange rates and stock prices: implications for European convergence. *Journal of Policy Modeling*, 24(5):523–526, 2002. doi: 10.1016/s0161-8938(02)00126-6.
- Susan E Moyer. Capital adequacy ratio regulations and accounting choices in commercial banks. *Journal of Accounting and Economics*, 13(2):123–154, 1990. doi: 10.1016/0165-4101(90)90027-2.
- Dhruv Mubayi, Todd G. Will, and Douglas B. West. Realizing degree imbalances in directed graphs. *Discrete Mathematics*, 239(1):147 – 153, 2001. ISSN 0012-365X. doi: [https://doi.org/10.1016/S0012-365X\(01\)00048-6](https://doi.org/10.1016/S0012-365X(01)00048-6). URL <http://www.sciencedirect.com/science/article/pii/S0012365X01000486>.

- J. A. Nelder and R. W. M. Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society. Series A (General)*, 135(3):370–384, 1972. ISSN 00359238. URL <http://www.jstor.org/stable/2344614>.
- Chien-Chung Nieh and Cheng-Few Lee. Dynamic relationship between stock prices and exchange rates for G-7 countries. *The Quarterly Review of Economics and Finance*, 41(4):477–490, 2002.
- Cathy Ning. Dependence structure between the equity market and the foreign exchange market—a copula approach. *Journal of International Money and Finance*, 29(5):743–759, 2010. doi: 10.1016/j.jimonfin.2009.12.002.
- OECD. *Euro Area*, volume 2004/5 of *Economic Surveys*. OECD Publishing, 2004. ISBN 9789264106567.
- OECD. Main economic indicators - complete database, 2018.
- J-P Onnela, Anirban Chakraborti, Kimmo Kaski, Janos Kertesz, and Antti Kanto. Dynamics of market correlations: Taxonomy and portfolio analysis. *Reviews of modern physics*, 68(5):056110, 2003. doi: 10.1103/physreve.68.056110.
- Sarah Padgett. The negative impact of Basel III. on small business financing. *Ohio St. Entrepren. Bus. LJ*, 8:183, 2012.
- Ming-Shiun Pan, Robert Chi-Wing Fok, and Y. Angela Liu. Dynamic linkages between exchange rates and stock prices: Evidence from East Asian markets. *International Review of Economics & Finance*, 16(4):503–520, 2007. doi: 10.1016/j.iref.2005.09.003.
- David H Papell and Hristos Theodoridis. The choice of numeraire currency in panel tests of purchasing power parity. *Journal of Money, Credit and Banking*, 33(3):790–803, 2001. doi: 10.2307/2673894.
- Juyong Park and MEJ Newman. Statistical mechanics of networks. *Phys. Rev. E*, 70:066117, Dec 2004. doi: 10.1103/PhysRevE.70.066117. URL <https://link.aps.org/doi/10.1103/PhysRevE.70.066117>.
- George C Philippatos and Charles J Wilson. Entropy, market risk, and the selection of efficient portfolios. *Applied Economics*, 4(3):209–220, 1972. doi: 10.1080/00036847200000017.
- Matija Piškorec, Nino Antulov-Fantulin, Petra Kralj Novak, Igor Mozetič, Miha Grčar, Irena Vodenska, and Tomislav Šmuc. Cohesiveness in financial news and its relation to market volatility. *Scientific Reports*, 4:5038, 2014. doi: 10.1038/srep05038.

- Vasiliki Plerou, Parameswaran Gopikrishnan, Bern Rosenow, Luis A. Nunes Amaral, Thomas Guhr, and H Eugene Stanley. Random matrix approach to cross correlations in financial data. *Physical Reviews E*, 65(6):066126, 2002. doi: 10.1103/PhysRevE.65.066126.
- Amanah Ramadiah, Fabio Caccioli, and Daniel Fricke. Reconstructing and stress testing credit networks, 2017. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3084543.
- Latha Ramchand and Raul Susmel. Volatility and cross correlation across major stock markets. *Journal of Empirical Finance*, 5(4):397–416, 1998. doi: 10.2139/ssrn.57948.
- David E Rapach, Jack K Strauss, and Guofu Zhou. International stock return predictability: what is the role of the United States? *The Journal of Finance*, 68(4):1633–1662, 2013. doi: 10.1111/jofi.12041.
- Alan Rapperport. Bill to Erase Some Dodd-Frank Banking Rules Passes in House. *The New York Times*, 2017.
- Reuters. Timeline: China’s reforms of yuan exchange rate. <http://www.reuters.com/article/2012/04/14/us-china-yuan-timeline-idUSBRE83D03820120414>, Apr. 2012.
- David Rodriguez. Why is the Australian Dollar Correlated to the US. S&P. 500? https://www.dailyfx.com/forex/technical/article/forex_correlations/2012/03/27/forex_correlations_australian_dollar_us_dollar.html, Mar. 2012.
- Ai ru Cheng, Kuntal Das, and Takeshi Shimatani. Central bank intervention and exchange rate volatility: Evidence from Japan using realized volatility. *Journal of Asian Economics*, 28:87–98, 2013. doi: 10.1016/j.asieco.2013.05.001.
- Leonidas Sandoval Jr. To lag or not to lag? How to compare indices of stock markets that operate on different times. *Physica A: Statistical Mechanics and its Applications*, 403:227–243, 2014. doi: 10.1016/j.physa.2014.02.039.
- Leonidas Sandoval Jr and Italo De Paula Franca. Correlation of financial markets in times of crisis. *Physica A: Statistical Mechanics and its Applications*, 391(1):187–208, 2012. doi: 10.1016/j.physa.2011.07.023.
- Frank Schweitzer, Giorgio Fagiolo, Didier Sornette, Fernando Vega-Redondo, and Douglas R White. Economic Networks: What do we know and what do we need to know? *Advances in Complex Systems*, 12(04n05):407–422, 2009. doi: 10.1142/s0219525909002337.

- Elizabeth Sheedy. Correlation in currency markets a risk-adjusted perspective. *Journal of International Financial Markets, Institutions and Money*, 8(1):59–82, 1998. doi: 10.1016/s1042-4431(98)00024-9.
- Hersh Shefrin and Meir Statman. The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of Finance*, 40(3):777–790, 1985. doi: 10.1111/j.1540-6261.1985.tb05002.x.
- Národnej Banky Slovenska. Slovak Koruna Included in the ERM II. <http://web.archive.org/web/20061002212008/http://www.nbs.sk/PRESS/PR051128.HTM>, Nov. 2005.
- Reginald D Smith. The spread of the credit crisis: view from a stock correlation network. *Journal of the Korean Physical Society*, 54(6):2460–2463, 2009. doi: 10.3938/jkps.54.2460.
- Bruno Solnik, Cyril Boucrelle, and Yann Le Fur. International market correlation and volatility. *Financial Analysts Journal*, pages 17–34, 1996. doi: 10.2469/faj.v52.n5.2021.
- Tiziano Squartini and Diego Garlaschelli. Analytical maximum-likelihood method to detect patterns in real networks. *New Journal of Physics*, 13(8):083001, 2011. URL <http://stacks.iop.org/1367-2630/13/i=8/a=083001>.
- Tiziano Squartini, Assaf Almog, Guido Caldarelli, Iman van Lelyveld, Diego Garlaschelli, and Giulio Cimini. Enhanced capital-asset pricing model for the reconstruction of bipartite financial networks. *Phys. Rev. E*, 96:032315, Sep 2017. doi: 10.1103/PhysRevE.96.032315. URL <https://link.aps.org/doi/10.1103/PhysRevE.96.032315>.
- Ralf Steuer, Jürgen Kurths, Carsten O Daub, Janko Weise, and Joachim Selbig. The mutual information: detecting and evaluating dependencies between variables. *Bioinformatics*, 18(suppl_2):S231–S240, 2002.
- Robert Tibshirani, Guenther Walther, and Trevor Hastie. Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2):411–423, 2001.
- Michele Tumminello, Tomaso Aste, Tiziana Di Matteo, and Rosario N Mantegna. A tool for filtering information in complex systems. *Proceedings of the National Academy of Sciences*, 102(30):10421–10426, 2005. doi: 10.1073/pnas.0500298102.
- Michele Tumminello, Tiziana Di Matteo, Tomaso Aste, and Rosario N Mantegna. Correlation based networks of equity returns sampled at different time horizons. *The European Physical Journal B*, 55(2):209–217, 2007.

- Christian Upper. Simulation methods to assess the danger of contagion in interbank markets. *Journal of Financial Stability*, 7(3):111–125, 2011. doi: 10.1016/j.jfs.2010.12.001.
- Christian Upper and Andreas Worms. Estimating bilateral exposures in the German interbank market: Is there a danger of contagion? *European Economic Review*, 48(4):827–849, 2004. doi: 10.1016/j.euroecorev.2003.12.009.
- Iman Van Lelyveld and Franka R Liedorp. Interbank Contagion in the Dutch Banking Sector: A Sensitivity Analysis. *International Journal of Central Banking*, 2(2), 2006.
- Stefania Vitali, James B Glattfelder, and Stefano Battiston. The network of global corporate control. *PLoS ONE*, 6(10):e25995, 2011. doi: 10.1371/journal.pone.0025995.
- Irena Vodenska, Hideaki Aoyama, Yoshi Fujiwara, Hiroshi Iyetomi, and Yuta Arai. Interdependencies and causalities in coupled financial networks. *PLoS ONE*, 11(3):e0150994, 2016a. doi: 10.1371/journal.pone.0150994.
- Irena Vodenska, Alexander P Becker, Di Zhou, Dror Y Kenett, H Eugene Stanley, and Shlomo Havlin. Community Analysis of Global Financial Markets. *Risks*, 4(2):13, 2016b. doi: 10.3390/risks4020013.
- Stefan Walter. Basel III and Financial Stability, Speech at the 5th Biennial Conference on Risk Management and Supervision. *Financial Stability Institute, Bank for International Settlements, Basel, 2010*, <http://www.bis.org/speeches/sp101109a.htm>, 2010.
- Duncan J Watts and Steven H Strogatz. Collective dynamics of ‘small-world’ networks. *nature*, 393(6684):440–442, 1998. doi: 10.1515/9781400841356.301.
- Ivo Welch and Amit Goyal. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508, 2008.
- DJ White. Entropy, market risk and the selection of efficient portfolios: comment. *Applied Economics*, 6(1):73–76, 1974. doi: 10.1080/00036847400000014.
- Marcel Wollschläger, Alexander P Becker, Irena Vodenska, H Eugene Stanley, and Rudi Schäfer. Economic and Political Effects on Currency Clustering Dynamics. *Quantitative Finance* (under review), 2018.
- WTO. *World Trade Statistical Review 2017*. Bernan Distribution, 2017. p. 109.
- Meilan Yan, Maximilian JB Hall, and Paul Turner. A cost–benefit analysis of Basel III: Some evidence from the UK. *International Review of Financial Analysis*, 25: 73–82, 2012. doi: 10.1016/j.irfa.2012.06.009.

Hua Zhao. Dynamic relationship between exchange rate and stock price: Evidence from China. *Research in International Business and Finance*, 24(2):103–112, 2010. doi: 10.1016/j.ribaf.2009.09.001.

Rongxi Zhou, Ru Cai, and Guanqun Tong. Applications of Entropy in Finance: A Review. *Entropy*, 15(11):4909–4931, 2013. ISSN 1099-4300. doi: 10.3390/e15114909.

CURRICULUM VITAE

