Coping with the Complexity of Financial Markets

Inauguraldissertation zur Erlangung des akademischen Grades
eines Doktors der Wirtschafts- und Sozialwissenschaften
der Wirtschafts- und Sozialwissenschaftlichen Fakultät
der Christian-Albrechts-Universität zu Kiel

vorgelegt von
Diplom-Volkswirt Daniel Fricke
aus Hildesheim
Kiel, Januar 2013

Gedruckt mit Genehmigung der Wirtschafts- und Sozialwissenschaftlichen
Fakultät der Christian-Albrechts-Universität zu Kiel
Dekan: Professor Dr. Horst Raff
Erstberichterstattender: Professor Dr. Thomas Lux
Zweitberichterstattender: Professor Dr. Stefan Reitz
Tag der Abgabe der Arbeit: 09.01.2013
Tag der mündlichen Prüfung: 01.03.2013

Coping with the Complexity of Financial Markets

Inauguraldissertation zur Erlangung des akademischen Grades
eines Doktors der Wirtschafts- und Sozialwissenschaften
der Wirtschafts- und Sozialwissenschaftlichen Fakultät
der Christian-Albrechts-Universität zu Kiel

vorgelegt von
Diplom-Volkswirt Daniel Fricke
aus Hildesheim

To Lorena

Acknowledgements

I owe my gratitude to all people who made this dissertation possible. First and foremost, my deepest gratitude is to my advisor and colleague, Prof. Thomas Lux. I consider myself extremely fortunate to have an advisor who both gives me all the freedom to carry out research on my own and provides perfect guidance whenever I need it. His patience and support helped me overcome all emerging difficulties during my dissertation. I hope that one day I will be able to become as good an advisor to my students as Prof. Lux has been to me.

Second, I thank my co-advisor Prof. Stefan Reitz for his patience and his continuous listening to my worries. I am deeply grateful to have such a fantastic second supervisor.

I also want thank my colleagues both from the CAU and the IfW for extensive discussions and helpful comments in several staff seminars. Particular thanks go those with whom I had epic table tennis battles.

Similarly, I am indebted to participants of numerous conferences and workshops, for helpful comments and for giving me the possibility to reflect upon the relevance of my work. Many thanks also go to Prof. Tony He, who arranged a wonderful research visit in October 2012 at the University of Technology in Sydney.

Both my family and many friends have helped me navigate even the most troubled waters during these exciting years. Special thanks go to Jonas, Karl, Martin, Nikolai, Sanna, and Tillmann.

Finally, I would like to thank Lorena, to whom this thesis is dedicated. You know why.

CONTENTS

Contents i

List of Figures iv

List of Tables viii

1. General Introduction 1
1.1 Motivation 2
1.2 Background and Overview 5
1.2.1 Financial Transaction Tax and Agent-Based Modelling 5
1.2.2 Interbank Markets and Network Analysis 9
1.2.3 Contributions 14

Part I Financial Transaction Tax and Agent-Based Modelling 16

2. Financial Transaction Tax 17
2.1 Introduction and Existing Literature 18
2.2 Model 22
2.2.1 CDA and Information 22
2.2.2 Trader Types 25
2.2.3 Asset Demand 32
2.2.4 Cancellation and Trading Process 34
2.2.5 GA Learning 36
2.3 Pseudo-Empirical Results 38
2.3.1 Properties 40
2.3.2 Effects of FTT 53
2.4 Conclusions 60

Part II Interbank Markets and Network Analysis 63

3. Network Analysis of the e-MID 64

3.1 Introduction and Existing Literature . . . 65
3.2 Networks . . . 66
3.3 The Italian Interbank Market e-MID . . . 67
3.4 Results . . . 68
3.4.1 General Features . . . 68
3.4.2 Density . . . 73
3.4.3 What is a Sensible Aggregation Period? . . . 75
3.4.4 Transitivity . . . 80
3.4.5 Small-World Property . . . 82
3.4.6 Effects of the Global Financial Crisis . . . 84
3.5 Conclusions and Outlook . . . 86

4. Distribution of Links in the Interbank Market . . . 89
4.1 Introduction and Existing Literature . . . 90
4.2 Networks . . . 94
4.3 The Italian Interbank Market (e-MID) . . . 95
4.4 Results . . . 95
4.4.1 Dynamics of the Degrees and Number of Transactions . 96
4.4.2 The Degree Distributions . . . 99
4.4.3 The Distribution of the Number of Transactions . . . 112
4.5 Conclusions . . . 116
4.6 Appendix . . . 118
4.6.1 Truncated Distributions and Maximum Likelihood . . . 118
4.6.2 Discrete Power-laws and Parameter Estimation . . . 120
4.6.3 Goodness-of-Fit Test for the Estimated Distributions . 121
4.6.4 Distributional Properties of Transaction Volumes . . . 122

5. Core-Periphery Structure in the Interbank Market . . . 126
5.1 Introduction and Existing Literature . . . 127
5.2 Networks . . . 129
5.3 Dataset . . . 130
5.4 Models . . . 134
5.4.1 The Discrete Model . . . 134
5.4.2 The Continuous Model . . . 137
5.5 Results . . . 140
5.5.1 Model Similarity . . . 140
5.5.2 Discrete and Tiering Model . . . 143
5.5.3 Continuous Model . . . 153
5.5.4 What Defines a Core Bank? . . . 157
5.5.5 What Happened During the GFC? . . . 160
5.6 Discussion . . . 163

5.7 Conclusions . . . 167
5.8 Appendix . . . 169
5.8.1 Genetic Algorithm . . . 169
5.8.2 Discrete Model: Illustration . . . 170
5.8.3 Empirical Estimation of the Asymmetric Continuous (AC) Model . . . 171
5.8.4 Model Fit: AC Model . . . 174
5.8.5 What Defines a Core Bank (in the Continuous Model)? 174
5.8.6 Changing the Aggregation Period . . . 177
5.8.7 Including Foreign Banks . . . 179
5.8.8 Continuous Model Using the Number of Transactions . 180
5.8.9 Further Robustness Checks . . . 180

6. Interbank Trading Strategies . . . 182
6.1 Introduction and Existing Literature . . . 183
6.2 Dataset . . . 184
6.3 Empirical Analysis of Trading Strategies . . . 186
6.3.1 Measuring Correlations Between Strategies . . . 186
6.3.2 Structure in the Correlation Matrices . . . 190
6.3.3 Clustering of Trading Behavior . . . 198
6.3.4 Time-Persistence of Clusters . . . 199
6.3.5 A Closer Look at the Clusters . . . 202
6.4 Conclusions . . . 205

References . . . 207

Bibliography . . . 207

LIST OF FIGURES

2.1 Time-horizons in the model. 22
2.2 Example: Development of expected price for different forward trend horizons. 28
2.3 Example: Limit price determination for a sell order. 31
2.4 Simulation results: Single run, $\theta = 0.2$, baseline scenario (GA off). 42
2.5 Simulation results: Single run, baseline scenario (GA off). Relative wealth. 42
2.6 Simulation results: Single run, $\theta = 0$ (GA off). 44
2.7 Simulation results: Single run, $\theta = 0.5$ (GA off). 45
2.8 Simulation results: Monte-Carlo simulations. Dependence on θ (GA off). 46
2.9 Simulation results: Single run, baseline scenario (GA on). . . . 48
2.10 Simulation results: Single run, baseline scenario (GA off). $\langle H^w/H^t \rangle = 1$. 49
2.11 Simulation results: Single run, baseline scenario (GA off). $\langle H^w/H^t \rangle = 0.2$. 50
2.12 Simulation results: Single run, baseline scenario (GA off). $\langle H^w/H^t \rangle = 1.8$. 51
2.13 Simulation results: Monte-Carlo simulations. Dependence on $\langle H^w/H^t \rangle$ (GA off). 52
2.14 Simulation results: Monte-Carlo simulations. Dependence on χ (GA off). 54
2.15 Simulation results: Monte-Carlo simulations. Dependence on χ (GA on). 56
2.16 Single run, baseline scenario (GA off) with three groups of informed traders without a tax. 58
2.17 Single run, baseline scenario (GA off) with three groups of informed traders and a small tax. 59
2.18 Single run, baseline scenario (GA off) with three groups of informed traders and a large tax. 60

2.19 Simulation results: Monte-Carlo simulations. Fraction of order submissions. 61
2.20 Simulation results: Monte-Carlo simulations. Dependence on χ (GA off) with different groups. 62

3.1 Number of active banks and traded volume over time. 69
3.2 Number of trades per quarter. 70
3.3 Fraction of trades and traded volume between banks from different countries. 71
3.4 The banking network in the 4th of quarter 2010. 72
3.5 Network density. 73
3.6 Density for aggregated networks. 75
3.7 Jaccard Index 77
3.8 Reciprocity for aggregated networks. 79
3.9 Transitivity. 80
3.10 Clustering coefficients. 82
3.11 Clustering coefficients for aggregated networks. 83

4.1 Mean and median degree over time. 97
4.2 In- vs. out-degree (Italian banks). 98
4.3 Mean and median number of transactions over time. 98
4.4 Quarterly data, degree. Histograms for the degrees (Period 1, 2, and 3). 101
4.5 Quarterly data, degree. Histograms for the degrees (Period 1-3). 101
4.6 Quarterly data, degree. Complementary cumulative distribution functions for the degrees on a log-log scale. 102
4.7 Daily data, degree. Complementary cumulative distribution functions for the degrees on a log-log scale. 102
4.8 Daily data, degree. Maximum degrees. 106
4.9 Daily data, degree. Histograms for the power-law exponents of the complete distributions. 106
4.10 Daily data, degree. Histograms for the power-law exponents of the tail observations. 107
4.11 Daily data, degree. Number of observations in total (complete) and tail. 107
4.12 Quarterly data, ntrans. Complementary cumulative distribution functions for the number of transactions on a log-log scale. 114
4.13 Daily data, ntrans. Complementary cumulative distribution functions for the number of transactions on a log-log scale. . . 114
4.14 Quarterly data, tvol. Complementary cumulative distribution functions for the transaction volumes on a log-log scale. 122

4.15 Daily data, tvol. Complementary cumulative distribution functions for the transaction volumes on a log-log scale. 123

5.1 Number of active banks and traded volume over time. 132
5.2 Density of the network over time. 133
5.3 Time-varying correlations between different coreness vectors. . 142
5.4 Absolute size of the core over time. 143
5.5 Relative size of the core over time. 145
5.6 Density of the blocks. 146
5.7 Fraction of intermediaries, lenders and borrowers over time. . . 147
5.8 Structure of the core and periphery in the discrete model. . . . 148
5.9 Transition probabilities over time, discrete model. 149
5.10 Jaccard Index for the different blocks over time. 151
5.11 Error score by block. 152
5.12 Error scores and core sizes in discrete model. Actual and random graphs. 154
5.13 In-coreness vs. Out-coreness for all observations, by core and periphery, as indicated by the discrete model. 155
5.14 Persistence of coreness vectors. 155
5.15 PRE for the SC and the AC model, actual and random graphs. 156
5.16 Time-varying correlation between discrete coreness and degree measures. 158
5.17 Time-varying correlation between discrete coreness and size measures. 159
5.18 Time-varying correlation between discrete coreness and activity measures. 159
5.19 In- and out-degrees of bank IT0278, by core and periphery. . . 162
5.20 Transaction volumes of bank IT0278. 163
5.21 Example for the core-periphery structure of the Italian interbank network, 1999Q1. 170
5.22 Example: Data matrix and approximation based on the AC model for 2000 Q3. 175
5.23 Time-varying correlation between in-/out-coreness and degree measures. 176
5.24 Time-varying correlation between in-/out-coreness and size measures. 176
5.25 Time-varying correlation between in-/out-coreness and activity measures. 177
5.26 Comparison of the time-varying relative core sizes for different aggregation periods. 178
5.27 Comparison of the error-scores for different aggregation periods. 179

6.1 Number of active banks and traded volume over time. 185
6.2 Activity by time of the day. 187
6.3 Number of sample banks over time. 187
6.4 Examples of banks' trading strategies in the Italian interbank market. 188
6.5 Time evolution of $Q = T/N$. 189
6.6 Example of a correlation matrix after applying the clustering approach for the first half of 2006. 190
6.7 Fraction of significant correlations over time. 191
6.8 Absolute and relative size of the first 5 Eigenvalues of $\mathbf{C}$ over time. 193
6.9 Comparison of observed and predicted eigenvalues from the MP distribution, for semester 9. 195
6.10 Relative size of the largest eigenvalue. 196
6.11 Inverse participation ratio (logarithm) for the first and last eigenvector over time. 198
6.12 Example of a correlation matrix after applying the clustering approach for the first half of 1999. 200
6.13 Frequency matrix $\mathbf{F}$. 201
6.14 Transition probabilities over time. 202
6.15 Most anti-correlated (signed) trading strategies in semester 3. 203
6.16 Most anti-correlated (signed) trading strategies in semester 10. 204
6.17 Relation between clusters and core-periphery model. 205

LIST OF TABLES

2.1 Example for a LOB at a certain point in time. 24
2.2 Baseline parameter setting for the simulations. 39

3.1 Jaccard Index for daily, monthly, quarterly and yearly networks. 77
3.2 Reciprocity of the Italian Banking network. 78
3.3 The average CC and ASPL for the observed and random networks. 84
3.4 Summary statistics for quarters 36-45. 85

4.1 Daily data, degree. KS statistic for the candidate distributions (complete). 109
4.2 Power-law parameters with standard deviations for daily and quarterly data, degree. 110
4.3 Daily data, degree. KS statistic for the candidate distributions (tail). 110
4.4 Quarterly data, degree. KS statistic for the candidate distributions (complete). 111
4.5 Quarterly data, degree. KS statistic for the candidate distributions (tail). 111
4.6 Power-law parameters with standard deviations for quarterly data, ntrans. 113
4.7 Quarterly data, ntrans. KS statistic for the candidate distributions (complete). 115
4.8 Quarterly data, ntrans. KS statistic for the candidate distributions (tail). 115
4.9 Daily data, tvol. KS statistic for the candidate distributions (complete). 123
4.10 Power-law parameters with standard deviations for daily and quarterly data, tvol. 124
4.11 Daily data, tvol. KS statistic for the candidate distributions (tail). 124
4.12 Quarterly data, tvol. KS statistic for the candidate distributions (complete). 125

4.13 Quarterly data, tvol. KS statistic for the candidate distributions (tail). 125

5.1 Correlations between individual coreness vectors of different models. 142
5.2 Transition matrix: trading strategies. 147
5.3 Transition matrix: discrete model. 148

6.1 Transition matrix: clusters. 201

1. GENERAL INTRODUCTION

This thesis can be split into two parts, containing the following papers:

Part I — Fricke, D., Lux, T. (2012), The Effects of a Financial Transaction Tax in an Artificial Financial Market, unpublished manuscript

Part II — Finger, K., Fricke, D., Lux, T. (2012), Network Analysis of the e-MID Overnight Money Market: The Informational Value of Different Aggregation Levels for Intrinsic Dynamic Processes, Kiel Working Paper 1782, Kiel Institute for the World Economy

Fricke, D., Lux, T. (2013), On the Distribution of Links in the Interbank Network: Evidence from the e-MID Overnight Money Market, unpublished manuscript

Fricke, D., Lux, T. (2012), Core-Periphery Structure of the Overnight Money Market: Evidence from the e-MID Trading Platform, Kiel Working Paper 1759, Kiel Institute for the World Economy

Fricke, D. (2012), Trading Strategies in the Overnight Money Market: Correlations and Clustering on the e-MID Trading Platform, Physica A, 391 (24), pp. 6528-6542

Section 1.2.3 contains a description of my contribution to each of those papers.

In the first part (chapter 2) I present a detailed artificial financial market with a large number of heterogeneous interacting agents. The model is used to assess the effects of a financial transaction tax (FTT). The second part (chapters 3-6) is concerned with the properties and the topological structure of the Italian interbank market. The contents of both parts are, at least on first sight, not connected to each other, as they provide contributions to very different areas of the literature. In this introduction, I mainly aim at motivating the work in the two parts and providing some background. I will also point out that the two parts are not as disconnected as it may seem, since both essentially aim at tackling the complexity of financial markets by taking into account the heterogeneity of agents and their interactions.

1.1 Motivation

Financial markets are prototypes of complex systems[1]: they consist of a large number of heterogenous components (e.g. households, firms, and financial institutions), with each component interacting with a subset of all

[1] I should stress that there exists no workable rigorous definition of complex systems so far. Similarly, the term financial market covers a wide range of different segments of the global financial system.

components. These interactions are not limited to trading in the marketplace, but also include connections between individual components. Thus, individual components do not act autonomously in general and their actions may induce feedbacks by others. At the macro-level, these interactions lead to collective market dynamics which are, as usual in complex systems, very hard to predict. This is even more so, since financial markets are only a subsystem, albeit one with a large and increasing importance, of the global economy. There is no doubt that well-functioning financial markets are crucial for the working of the economy, and their instability has significant effects on the 'real' sector.

It has been shown that components following rather simple rules may, by their very interactions, be the source of complicated aggregate dynamics.[2] Applying this insight to financial markets makes clear that higher-level aggregates e.g. market prices, cannot be understood on the basis of their constituent units.[3] The aggregate is not simply the sum of the individual parts and it is crucial to understand the links between the micro- and macro-level. For the economy in general, and financial markets in particular, this raises the necessity of understanding and modelling the system from bottom-up. This thesis aims at describing and tackling the complexity of certain financial market segments.

In order to frame the work in this thesis, at this point it seems necessary to briefly explain how mainstream economists, in particular from macroeconomics and finance, tend to cope with the complexity of the economy.[4] Modern mainstream macroeconomics, most prominently the so-called dynamic stochastic general equilibrium (DSGE) models, tells us that the macroeconomic equilibrium can be found by maximizing the present value of the expected utility of a rational, representative individual ('homo oeconomicus'). In this view, deriving the optimal behavior of the representative agent is already sufficient for understanding the dynamics of the whole economy. However, as Sonnenschein (1972), Mantel (1974), and Debreu (1974) have shown 40 years ago, aggregate demand is not as well behaved as individual demand. Obviously, the linear aggregation of individual behavior misses the interactions between individuals, which are crucial for understanding the system's dynamics. The aggregate behavior is fundamentally different from that of the 'average' individual.[5] Related thoughts can be found in finance:

[2] See for example Schelling's (1969) model of racial segregation.

[3] See Lux (2011). This has important implications for the regulation of financial institutions, as I will explain below.

[4] See e.g. Kirman (1992), Bouchaud (2008), Colander *et al.* (2009), and Lux and Westerhoff (2009) for detailed discussions of the topic.

[5] See Anderson (1972).

the efficient market hypothesis (EMH) states that non-rational traders can be neglected, since they would have to leave markets after continuously losing to rational traders.[6] From this viewpoint, it again suffices to consider a representative, rational individual for analyzing the complete system. Hence, standard macro and finance models fail to account for the complexity of the economy. In addition, many of the very basic assumptions, such as perfect rationality, have been rejected spectacularly.[7]

It is easy to imagine the potential dangers of predictions from such overly simplistic models. For example, perfect rationality implies efficient, and therefore incredibly stable, financial markets. Prices will always equal fundamental values, and the only source for financial crises is due to exogenous (fundamental) shocks. This explains why financial markets are practically absent from modern macro-models: there is simply no reason why a perfectly functioning market should have substantial impact on the economy. For a large part, the opacity of many market segments, among them the absence of real-time data for the interbank linkages, is to some extent favored by the broad disinterest in the structure of the financial system. Once again, the global financial crisis (GFC) of 2007-08 made painfully clear that such approaches are inadequate. Financial markets profoundly affect the stability of the entire economy. It is likely that the general trend of 'financialization',[8] involving a larger share of financial institutions profits' among all corporate profits[9] and an exponential growth of financial markets around the world,[10] significantly increases the impact of financial markets on the stability of the global economy.[11] The sheer size of financial markets is worrying by itself, but even more threatening appears the increasing level of interconnectedness and concentration in the global financial system.[12] In the end, a substantial

[6] See e.g. Fama (1965, 1970).

[7] See e.g. Kahneman and Tversky (1973) and Tversky and Kahneman (1974). The booming research fields of behavioral and experimental economics continuously add new evidence against the rationality of agents, see e.g. Shiller (2003).

[8] See Dore (2008).

[9] For the US for example, the fraction of financial institutions' profits among all corporate profits was roughly 35% at the beginning of 2012. In the 1950s (2002) the same figure read 9.5% (45%). See section 6 of the National Income and Product Accounts Tables, available online: http://www.bea.gov/iTable/index_nipa.cfm.

[10] According to statistics from the Bank of International Settlements (2011), the notional amount of outstanding over-the-counter derivative contracts was 648 trillion US Dollars in December 2011. Compare this with the world domestic product of roughly 70 trillion US Dollars in 2011, i.e roughly one tenth of the derivative contracts and these are just estimates for the OTC derivative market.

[11] Of course, part of that growth can be attributed to globalization and the growth of world trade.

[12] See Haldane (2009) and Kubelec and Sa (2010).

fraction of financial market growth is likely to be favored by moral hazard: when losses are socialized, while gains are not, financial institutions have incentives to become 'systemically important' by increasing their risk exposures. Unfortunately, it took the recent crisis for policymakers to realize that the absence of very basic real-time information, such as maps of interbank exposures and a quantitative definition of systemic importance, creates additional uncertainty.[13] In hindsight, it seems unlikely that Lehman Brothers would have been allowed to fail in September 2008, if policymakers had been aware of the scale of Lehman's systemic importance. However, not only the lack of data, but also the lack of appropriate models for assessing the systemic importance of individual institutions, is striking.

The two parts of this thesis cope with some of these issues in detail, always acknowledging the complexity of the system. In the first part, I present a detailed agent-based model of an asset market, where a large number of heterogenous agents compete against each other. The model is used to analyze the effects of a financial transaction tax (FTT), which aims at reducing financial market fluctuations and generating tax revenues by imposing a small transaction tax. In the second part, I present basic work on the interconnectedness of the Italian interbank network, i.e. a network of banks which are connected to each other through overnight loans. Haldane (2009) argues that the first step for efficient regulation and crisis prevention in financial systems consists of mapping the network. Only then, after understanding its structure and dynamics, can we move on to improve regulations and possibly restructure the system. In a series of papers, my co-authors and I follow this proposal by analyzing the (dynamical) structure of this particular network. It is of utmost importance to take the next steps in the future.

In the next section, I will present a brief overview of the individual chapters. I will give a separate motivation for both parts and outline possible areas of future research.

1.2 Background and Overview

1.2.1 Financial Transaction Tax and Agent-Based Modelling

The original idea that a general FTT might curb financial market volatility goes back to John Maynard Keynes (1936), ch. 12. James Tobin (1978) transferred this idea to currency markets (Tobin-tax). Only recently, given the surging government deficits from responses to the crisis, policymakers started realizing that such a tax might actually generate substantial mone-

[13] See e.g. Trichet (2011).

tary revenues.[14]

FTTs theoretical appeal stems from its potential to limit short-term speculative behavior, and thus transaction volumes, on financial markets due to the rise in transaction costs. The hypothesis is that there is a U-shaped relationship between the tax rate and volatility: for small tax rates, volatility decreases, but increases for larger values. In the light of the previously mentioned explosion of financial market activity, and the corresponding decoupling from 'real' activity, restricting the turnover appears to be a reasonable aim. Note also that the increasing level of activity goes along with a higher level of liquidity.[15] From a naive point of view, one would expect more liquid markets to be more resilient. However, there is some evidence that higher liquidity coincided with higher fragility in the sense that financial crises, i.e. the bursting of speculative bubbles, became more likely.[16] FTTs, by harming short-term investments overproportionally, should have the potential to reunite financial market and real activity. Additionally, besides generating tax revenues, it frees-up economic resources from the financial sector for more productive uses.[17] However, empirical evidence on the expected U-shaped relationship between FTTs and volatility is rather mixed: some studies find that volatility decreases, increases or does not react at all in response to a tax increase.[18] Given these contradicting empirical results, it becomes necessary to analyze the effects of the FTT in a different way. In the light of the complexity of financial markets, agent-based models (ABM) are a very promising approach, which allow to non-invasively evaluate the effects of regulatory measures in general.

More detailed (realistic) models are usually hard to tackle analytically, so numerical simulations are needed. ABMs are such computerized simulations, containing a number of components (agents) interacting with each other through prescribed rules.[19] As such, they take all the necessary ingredients for modelling complex systems (heterogeneity and interactions) into

[14] These revenues are estimated to range between 1 and 3% of national GDPs. See, e.g. Pollin *et al.* (2003).

[15] Liquidity is the ability to trade large size quickly, at low costs, see Harris (2003).

[16] See Shiller (1981) and Bordo *et al.* (2001). Additionally, the arrival of high-frequency trading is likely to have increased the speed to the market dynamics, see e.g. Sornette and der Becke (2009), which increases the probability of micro-crashes as seen in the Flash crash of 2010.

[17] For example, Dore (2008) reports that in 2006 22% of the British labor force was employed in the financial industry. Among them you will find a large number physicists, mathematicians, and engineers, designing clever trading algorithms.

[18] See, e.g. Jones and Seguin (1997), Hau (2006) and Roll (1989), respectively.

[19] See LeBaron and Tesfatsion (2008), Farmer and Foley (2009), and Helbing (2012), ch. 2.

account.[20] ABMs have guided policymakers in other areas such as epidemiology and traffic control, and there is no reason why economic systems could not be modeled similarly. A major challenge consists of specifying how agents behave and interact, due to the infinite number of ways to model deviations from perfect rationality. The degrees of freedom can be partly reduced using results from behavioral and experimental economics, serving as a guideline for agent behavior in many cases. Obviously, attempting to model all the details of a realistic problem can rapidly lead to a 'black box' where it is difficult to determine what causes what. Without empirical guidelines, it is clearly necessary to test different behavioral strategies and evaluate their degree of realism. Validation is commonly done with regard to the replication of certain well-known 'stylized facts'. A model that replicates more of these, is likely to be more favorable compared to others.

There is a long list of ABMs of financial markets, usually within the chartist-fundamentalist framework as in e.g. Kirman (1991, 1993) and Lux (1995), that are able to replicate many of the stylized facts apparent in financial market data.[21] Important insights can be gained from many of such models, but when it comes to evaluating regulatory policies, crucial (oversimplifying) assumptions concerning agents' behavior and market microstructure are likely to affect the results. For example, many authors assume that a market-maker provides infinite liquidity, in which case FTTs are potentially stabilizing for small tax rates.[22] More realistic approaches take the endogeneity of liquidity into account, usually based on order-driven continuous double auctions (CDA). However, many of these so-called 'econophysics' approaches model the case of zero-intelligence traders,[23] i.e. the exact opposite of the 'homo oeconomicus' in standard economics.[24] Psychological factors, such as imitation and herding, are assumed to be absent in such models working

[20] It is beyond the scope of this introduction to name all the advantages and disadvantages of ABMs. I refer the interested reader to Helbing (2012), ch. 2.

[21] See for example Beja and Goldman (1980), LeBaron *et al.* (1999), and Lux and Marchesi (1999, 2000).

[22] For a single asset market, see Ehrenstein (2002) and Westerhoff (2003, 2004a). However, Giardina and Bouchaud (2003b) find that only substantial trading costs will actually stabilize the market, while a small tax (of the order of a few basis points) would have no real effect. For two ex-ante identical markets, with one country unilaterally introducing the tax, Westerhoff (2004b) finds that the taxed market is stabilized while volatility in the tax haven strongly increases. For markets of different size, Hanke *et al.* (2010) conclude that if the tax is introduced in the large (small) market, volatility decreases (increases) there.

[23] See Cliff and Bruten (1997).

[24] See for example Bak *et al.* (1997), Maslov (2000), and Challet and Stinchcombe (2001).

at very short time-scales. To date, few attempts have been made to model the limit order book based on detailed interactions between many boundedly rational agents and it is not clear how strategic interactions would alter the findings as compared to zero intelligence models. In particular, most existing studies from the 'econophysics' literature appear to be reluctant to model the market structure in a realistic way, e.g. incorporating budget constraints, mostly in order to avoid making behavioral assumptions about the agents.[25]

The model presented in chapter 2 aims at bridging the gap between short and long time-scales. To quantify the effects of the FTT, we take into account adjustments in trading strategies (time horizons) and their effects on liquidity. The ABM is based on a realistic CDA mechanism and allows for a continuum of investment strategies within the chartist/fundamentalist framework. Successful strategies spread through the population via a Genetic Algorithm (GA), i.e. agents can adjust their strategies in response to the FTT. The main conclusions are as follows: First, the model is able to replicate certain stylized facts of real financial time-series for several parameter combinations, e.g. the model replicates the building up and bursting of price bubbles. Second, for both fixed and flexible behavioral strategies, we find the usual trade-off between monetary revenues (a kind of Laffer curve) and stability, as higher tax revenues come along with higher volatility. This finding is in line with the results from the existing literature. However, we find somewhat different results for very small and large tax rates, indicating that the effects of the tax may not be entirely negative. In any case, the tax allows to generate substantial tax revenues, which could be used for a number of productive purposes. In reality, it appears likely that social welfare could be enhanced by the imposition of such a tax, e.g. through directing money and resources from the financial sector towards more productive purposes. Unfortunately, given the political power of the financial sector, it is still an open question whether policymakers indeed decide to impose such a tax.

We discuss numerous avenues of possible future research are discussed in the paper. Here I just mention that we, that is Karl Finger and myself, are working on a simpler version of the model. Our main aim is to extend such a model to the case of two markets, i.e. allowing for a tax haven. Furthermore, we want to compare the effects of a FTT with that of a financial activity tax, which taxes the profits from trading rather than all transactions.

[25] To our knowledge, Chiarella and Iori (2002) were the first to incorporate trading strategies into a CDA setup. The authors state that without these strategies, it is impossible to generate realistic time-series. See also Chiarella *et al.* (2009).

1.2.2 Interbank Markets and Network Analysis

Interbank markets allow banks to exchange central bank money in order to share liquidity risks.[26] At the macro level, however, a high number of bank connections could give rise to systemic risk.[27] Since it is well known that the structure of a network is important for its resilience,[28] policymakers need information on the actual topology of the interbank network. This is even more so, since market-based indices for systemic risk fail in serving as early-warning indicators.[29] In order to move from micro- to macro-prudential regulations, it is crucial to map and understand the dynamics of the interbank network. Moreover, the identification of meaningful systemic risk indices has to take the actual structure of the interbank network into account.

Recent research in the natural sciences has significantly advanced our understanding of the structure and functioning of complex networks. Network ideas have been applied to very diverse areas and data sets such as the internet, epidemiology, ecosystems, scientific collaboration and financial markets, to name a few. Quite interestingly, despite the heterogeneity of the underlying systems, many real-world network share a number of common features.[30]

The experiences of the last few years have made policymakers aware of the necessity of gathering information on the structure of the financial network in general and the interbank market in particular.[31] Clearly it is of overwhelming importance, to set up a dynamic ABM of the financial system in the future, where the interbank network is only one of many segments. One reason for the previous scarcity of research on the connections between financial institutions is certainly the limitation of available data,[32] the other reason being the neglect of the internal structure of the financial system in standard macroeconomic models as explained before. Thus, it comes as no surprise that most previous studies on the topology of interbank markets have been conducted by physicists applying measures from the natural sciences to a network formed by interbank liabilities.[33] It has been found that

[26] See Ho and Saunders (1985), Freixas *et al.* (2000) and Allen and Gale (2000).

[27] Systemic risk is closely related to financial contagion, see de Bandt and Hartmann (2000), and implies that an idiosyncratic shock causing the failure of one or few institutions may destabilize the entire system.

[28] See also Allen and Gale (2000).

[29] See Markose et al. (2010).

[30] One of the most prominent features is the 'small-world' phenomenon. This means that most nodes can be reached by every other node by a very small number of steps. See Watts and Strogatz (1998).

[31] See Haldane (2009), Haldane and May (2011) and Trichet (2011).

[32] See Mistrulli (2007).

[33] See for example Boss *et al.* (2004), Inaoka *et al.* (2004), Soramäki *et al.* (2007), Bech

interbank networks appear to share several common properties. For example, such networks are

1. sparse (only few links exist at any point in time),
2. scale-free (power-law degree distribution, with the degree being the number of active links per bank),
3. clustered (non-random network structure),
4. disassortative (high-degree nodes tend to trade with low-degree nodes and vice versa).

The second part of this thesis contains basic analyses of a particular interbank network. The overarching themes are the identification of new (stylized?) facts and the challenge of existing ones. For that purpose, we investigated the Italian interbank network, based on a dataset covering all transactions on the e-MID trading platform during the period 1999 to 2010. This dataset is used simply due to the fact that the e-MID data are the only interbank data which can be purchased freely without any restrictions. In contrast, getting access to similar datasets for other markets is usually far more complicated, but will be necessary to make statements about the universality of our findings.

Chapter 3 contains a basic analysis of some network properties of the e-MID data. The basic motivation is to match the seemingly incompatible findings of random daily networks,[34], and the existence of preferential lending relationships at other aggregation levels, i.e. significant non-random structure.[35] In this paper, we show that the aggregation period indeed has an effect on the informational value of the underlying networks. The main finding is that daily networks indeed feature a substantial amount of randomness and cannot be considered as being sufficiently informative for the underlying 'latent' network. This is illustrated on the basis of a number of network statistics which are compared to those of random networks. Furthermore, we find a substantial amount of asymmetry in the network. Last but not least, we find that the GFC can be identified as a significant structural break for many network measures.

Essentially, these results show that it is far from trivial to map a given data structure into a 'network'. While daily records of the interbank trading system can be arranged in an adjacency matrix and treated with all types of network statistics, they provide probably only a very small sample of realizations from a richer structure of relationships. Just like daily contacts of

and Atalay (2010), De Masi *et al.* (2006), and Iori *et al.* (2008).

[34] See De Masi *et al.* (2006) and Iori *et al.* (2008).

[35] See e.g. Cocco *et al.* (2009) for the Portuguese interbank market.

humans provide very incomplete information of networks of friendship and acquaintances, the daily interbank data might only provide a small selection of existing, dormant established trading channels. Hence, inference based on such high-frequency data may be misleading while a higher level of time aggregation might provide a more complete view on the interbank market. What level of aggregation is sufficient for certain purposes is an empirical question depending on the research questions at hand. However, saturation of certain measures may be a good indicator that most dormant links have been activated at least once over a certain time horizon. At the same time, such dependence of statistics on the time horizon serves to sort out a number of simple generating mechanisms (i.e. completely randomly determined networks in every period) and reveal interesting dynamic structure.

Chapter 4 is closely related to the analysis in chapter 3, explicitly focusing on the distribution of in- and out-degrees in the Italian interbank network. The analysis challenges the somewhat universal finding of scale-free interbank networks. Quite interestingly, many interbank networks have been reported to be scale-free networks. For example, using daily e-MID data over the period 1999-2002, De Masi *et al.* (2006) reported power-laws for the distribution of in- and out-degrees, with tail parameters 2.7 and 2.15, respectively. Besides the existence of several well-known generating mechanisms for scale-free networks,[36] this finding carries important policy implications. One important feature of scale-free networks is that they can be described as robust-yet-fragile,[37] indicating that random disturbances are easily absorbed (robust) whereas targeted attacks on the most central nodes may lead to a breakdown of the entire network (fragile). If the network of credit relationships had such a structure, this would carry important policy implications. For instance, such a network might experience long stable periods, during which disruptions are confined to peripherical banks and can be absorbed easily within the entire system. However, such periods could be a misleading indicator of the overall stability of the system as problems affecting the most central nodes could suddenly cause a breakdown of the entire network, cf. Haldane (2009). Given the substantial level of randomness in the daily networks, cf. chapter 3, we will also look at the distributional properties at the quarterly level. Due to the aggregation properties of power-laws, one would expect any 'true' power-laws at the daily level to carry over to other frequencies as well.

Quite surprisingly in view of the previous literature, we find hardly any support in favor of power-laws: at the daily level the degrees are usually fit

[36] See for example Barabasi and Albert (1999).
[37] See Albert *et al.* (2000).

best by negative Binomial distributions, while the power-law may provide the best fit for the tail data. However, we typically find very large power-law exponents (with values as large as 7), i.e. levels where the power-law is virtually indistinguishable from exponential decay. At the quarterly level, Weibull, Gamma, and Exponential distributions tend to provide comparable fits for the complete degree distribution, while the tails again tend to display exponential decay. We find comparable results when investigating the distribution of the number of transactions, even though in this case the tails of the quarterly variables are somewhat fatter. However, the Log-normal distribution typically outperforms the power-law. Overall these findings indicate that the power-law is typically a poor description of the data, implying that preferential attachment and other generating mechanisms for scale-free networks are unsuitable explanatory mechanisms for the structure of the Italian interbank network. Moreover, the networks contain a substantial level of asymmetry, due to the low correlation between in- and out-degrees. Additionally, we find that the two variables do not follow identical distributions in general.

Chapter 5 sheds light on the hierarchical structure of the interbank network. The hypothesis of the existence of a hierarchical structure can be drawn from the disassortative mixing patterns, which imply that small banks tend to trade with large banks, but rarely among themselves. In passing, many authors have indeed remarked that there seemed to be some kind of community structure in the interbank network they analyzed. For example, Boss *et al.* (2004) note that the Austrian interbank network shows a hierarchical community structure that mirrors the regional and sectoral organization of the Austrian banking system. Soramäki *et al.* (2007) show that the network includes a tightly connected core of money-center banks to which all other banks connect. Thus there is some form of tiering in the interbank market. The empirical findings of Cocco *et al.* (2009) also show that relationships between banks are important factors to explain differences in interest rates. Identifying communities in networks is an important aspect and in this paper we are concerned with the identification of the set of arguably systemically important (core) banks. In order to do so, we estimate various versions of core-periphery models in the spirit of Borgatti and Everett (2000). The basic idea is that a network can be divided into subgroups of core and periphery members, where the core (periphery) banks are maximally (minimally) connected to each other. To our knowledge, Craig and von Peter (2010) is the first and so far only contribution applying a core-periphery structure to the German interbank market. In this paper, we applied the (unrestricted) discrete core-periphery model, the (restricted) tiering model due to Craig and von Peter (2010) as well as symmetric and asymmetric versions of a

continuous core-periphery model (hitherto not applied to interbank data) to the Italian interbank data. We find that a core-periphery structure provides a better fit for these interbank data than alternative network models. The identified core shows a high degree of persistence over time, consisting of roughly 28% of all banks before the global financial crisis and 23% afterwards. Quite interestingly, other standard community finding approaches are unable to identify other community structures. We can classify the majority of core banks as intermediaries, i.e. as banks both borrowing and lending money in the market. Furthermore, allowing for asymmetric 'coreness' with respect to lending and borrowing activity considerably improves the fit, and reveals more concentration in borrowing than lending activity of money center banks. We also shed light on the development during the financial crisis of 2008, finding that the reduction of interbank lending was mainly due to core banks' reducing their numbers of active outgoing links.

The general conclusion is that preferential lending relationships at the micro-level lead to hierarchical structure at the macro-level. The findings carry important policy implications, e.g. is has been shown that a disassortative core-periphery framework might be more robust in 'normal' times, but more fragile under exceptional circumstances when key nodes are under stress or withdraw from the market. In this sense, the 'coreness' translates to a certain extent into 'systemic relevance' of individual institutions.[38] Additionally, we argue that such network approaches allow to reduce the complexity of interbank data sets in a number of ways. For example, the goodness-of-fit of the core-periphery framework can be seen as an indicator of money market stress.

Chapter 6 moves on from the finding that it is impossible to identify other community structures than the core-periphery partition. Here we aimed at analyzing the trading strategies of individual institutions in more detail. Splitting the dataset into half-yearly subsamples, we used the (intra-) daily net trading volumes of the individual institutions as the trading strategies and analyzed the correlations between them.[39] We find evidence for significant and persistent bilateral correlations between institutions' trading strategies. In most semesters there are two anti-correlated clusters of trading strategies, indicating that banks tend to trade preferentially with banks from the other cluster.[40] This result is not only in line with the existence of preferential relationships, but also with herding phenomena in banks' trading strategies. Interestingly, the information whether a bank is a continuous net

[38] See also Markose et al. (2010).

[39] See Kyriakopoulos *et al.* (2009) for a similar approach.

[40] Zovko and Farmer (2007) performed a similar analysis for the trading of members in the London Stock Exchange.

buyer or net seller appears to be the only defining characteristic for its group membership in this anonymous dataset. This is in contrast to the findings of Iori *et al.* (2007), who performed a very similar analysis for the sample period 1999-2002, and found the two clusters containing mostly large and small banks, respectively. Overall, the findings are quite persistent over time and appear to be largely unaffected by the GFC.

These findings add further evidence on the fact that preferential lending relationships on the micro-level lead to community structure on the macro-level. In this sense, trading on the e-MID platform appears to be a relatively structured process in terms of trading strategies. Given that each trade involves two counterparties, i.e. a buying and a selling side, it may appear trivial that we identify the two clusters. However, the high level of persistence shows that most banks tend to have rather stable trading strategies over time, since they tend to appear in the same cluster over time, trading with the same counterparties.

All of the presented work is just a first step in describing and understanding the structure of the interbank network in more detail. We are still at the beginning of understanding individual bank behavior, which is crucial for the complex aggregate dynamics. Building realistic ABMs of the financial system, however, involves more detailed knowledge about the universality of our findings. As argued above, financial institutions are exposed to each other in a variety of markets, of which the interbank market is just one segment. Future research should aim at mapping the complete financial system. Only after understanding the structure and dynamics, it will be possible to make reliable policy recommendations for regulating (and possibly restructuring) the financial system. I plan to contribute to this proposal in a number of future research projects.

1.2.3 Contributions

Here I briefly describe my contribution to the individual chapters.

- Chapter 2: This paper is joint work with Thomas Lux and the result of extensive discussions between the two of us. Professor Lux came up with the idea to model informed traders in the spirit of Youssefmir *et al.* (1998). The technical implementation, simulation, analysis, and most of the writing, were done by myself.

- Chapter 3: This paper is joint work with Karl Finger and Thomas Lux. The technical implementation and the writing, usually resulting from extensive discussions with Professor Lux, were mostly done by Karl and myself.

- Chapter 4: This paper is joint work with Thomas Lux. The technical implementation and the writing, again based on extensive discussions with Professor Lux, were done by myself.

- Chapter 5: This paper is joint work with Thomas Lux. He provided the basic idea of applying the core-periphery framework to the dataset. I came up with the idea to estimate both the restricted and unrestricted discrete models and the continuous models, which has not been done so far. The technical part was all done by myself, writing and editing was shared roughly equally. Mattia Montagna helped creating one of the network pictures.

- Chapter 6: All the work in this paper was done by myself.

Part I

FINANCIAL TRANSACTION TAX AND AGENT-BASED MODELLING

2. THE EFFECTS OF A FINANCIAL TRANSACTION TAX IN AN ARTIFICIAL FINANCIAL MARKET

2.1 Introduction and Existing Literature

For several decades the financial transaction tax (FTT) has been discussed as an instrument to curb financial market volatility.[1] Only recently -given the surging government deficits from responses to the global financial crisis- the focus has shifted to the FTT's large potential monetary revenues.[2] In this paper we investigate the effects of a FTT in an agent-based artificial financial market.

The FTT's appeal stems from its potential to limit short-term speculative behavior, and thus transaction volumes, on financial markets. This seems a reasonable aim given the divergence of financial market and 'real' activity during the last decades, when increases in financial market transaction volumes continuously exceeded those of the real economy. The exponential growth of financial transaction volumes was fueled by a continuous fall in transaction costs for many assets due to the technological progress in computer-based trading and an increased competition between stock exchanges. One result of this development is the increased presence of so-called high-frequency trading (HFT), which is predominantly employed by large hedge funds and essentially represents the modern version of market making.[3] Unfortunately, higher liquidity[4] seems to have come along with higher fragility in the sense that financial crises, i.e. the bursting of speculative bubbles, became more likely.[5] In this way, a FTT that favors longer-term investments could have the effect of reducing the decoupling of financial market from real activity and additionally free resources from the financial sector for more productive uses.

Critics of the FTT, most importantly from the financial industry, usually bring forward the following arguments: (1) market liquidity will dry up, (2) volatility may thereby in fact increase, (3) banks will pass on the tax burden to firms and other bank customers, raising capital costs in general, and (4)

[1] The first proposal for a general FTT was made by John Maynard Keynes (see Keynes (1936), Chapter 12). James Tobin's proposal of what has hitherto been christianed Tobin-tax, referred to currency markets, see Tobin (1978).

[2] These revenues are estimated to range between 1 and 3% of national GDPs. See, e.g. Pollin *et al.* (2003).

[3] We should note, however, that many economists actually favor HFT, since the increased liquidity should make markets more resilient. See Brogaard (2010) for an empirical assessment.

[4] Liquidity is the ability to trade large size quickly, at low costs, see Harris (2003).

[5] Bordo *et al.* (2001) find that the frequency of crises since 1973 has been double that of the Bretton Woods and classical gold standard periods. Two important explanatory factors are financial globalization and safety nets encouraging financial institutions to take on higher risks.

there is a danger of capital flights from a taxed market towards untaxed markets. In this paper we are concerned with the first two (interrelated) points.[6]

High liquidity, i.e. small transaction costs, fuels excess volatility[7] (compared to 'fundamentals') as it makes round-trips relatively cheap. Empirical evidence verifies that FTTs, despite applying for all market participants, harm short-term speculators disproportionately more. For example, it has been found that an increase of a transaction tax has increased asset holding periods, while transaction volumes have decreased.[8] However, this does not imply that volatility will decrease as well. In theory, there should be a U-shaped relationship: for small tax rates volatility should decrease, since (destabilizing) short-term oriented speculation becomes unprofitable. However, larger tax rates will affect (stabilizing) longer-term strategies as well, thereby reducing liquidity and potentially increasing volatility. Empirical evidence on point (2) is therefore rather mixed: some studies find that volatility decreases, increases or does not react at all in response to a tax increase.[9]

Given these contradicting results, simulations of artificial financial markets are a promising way to non-invasively evaluate the effects of regulatory measures in general. Numerous agent-based models, usually within the chartist-fundamentalist framework as in e.g. Kirman (1991, 1993) and Lux (1995), are able to replicate many of the stylized facts characterizing financial market data.[10] When used to evaluate regulatory policies, however, crucial (simplifying) assumptions concerning agents' behavior and market microstructure should affect the results. For example, many authors assume that a market-maker provides infinite liquidity, in which case FTTs are potentially stabilizing for small tax rates. For a single asset market, see Ehrenstein (2002) and Westerhoff (2003, 2004a). However, Giardina and Bouchaud (2003b) find that only substantial trading costs will actually stabilize the market, while a small tax (of the order of a few basis points) would have no

[6] In future research we are planning to deal with point (4) as well. This could be done by having two markets and imposing the tax only on one of them.

[7] See Shiller (1981).

[8] See Jackson and O'Donnell (1985) and Baltagi *et al.* (2006). For example Baltagi *et al.* (2006) find that a tax increase from 0.3 to 0.5% reduced trading volume in China by roughly 1/3.

[9] See, e.g. Jones and Seguin (1997), Hau (2006) and Roll (1989), respectively.

[10] An early example is Beja and Goldman (1980). See also LeBaron *et al.* (1999), Challet and Zhang (1997), Chiarella and Iori (2002), Lux and Marchesi (1999, 2000), Lux and Schornstein (2005), Raberto *et al.* (2003) and Chiarella *et al.* (2009). See e.g. Allen and Taylor (1990) and Menkhoff (1998) for empirical evidence on the use of chartist and fundamentalist strategies, respectively.

real effect.[11] However, since liquidity is a major determinant of volatility in real markets,[12] it is crucial when discussing the effects of FTTs. Therefore we explicitly take the microstructure of real markets and provision of liquidity into account by simulating an order-driven continuous double-auction (CDA).

While the models based on the marker-maker setting usually incorporate important psychological factors that drive the system's properties, e.g. through herding and imitation, CDA models usually work at shorter time-scales where psychological factors are either hard to model or simply assumed to be absent. For example, Bak *et al.* (1997)[13] treat the limit order book (LOB) as a system of particles with each particle (order) having a mass (order size) and a price (spatial position). Price variations stem from diffusion and annihilation of particles, well-known processes in physics, which allows to obtain analytical results. Even though many important insights can be gained from such approaches, the usual criticism is that these models operate within a zero intelligence framework.[14] This is the exact opposite to 'homo oeconomicus' in mainstream economics, but traders are unlikely to be either fully rational or plainly stupid. Another problem is that, by avoiding detailed behavioral assumptions, these models typically ignore budget constraints and wealth dynamics. Nevertheless, since these models are able to replicate certain stylized facts of LOB data, the structure of the trading protocol is likely to have a significant effect on the data-generating process. To date, few attempts have been made to model the LOB based on detailed strategic interactions between many boundedly rational agents, while incorporating economic constraints.[15] Our model aims at bridging the gap between models with short and long time-scales.

To our knowledge, only few studies dealt with FTTs in detailed order-driven markets. Two examples are Mannaro *et al.* (2008) and Pellizzari and Westerhoff (2009). Mannaro *et al.* (2008) use a zero intelligence frame-

[11] For two ex-ante identical markets, with one country unilaterally introducing the tax, Westerhoff (2004b) finds that the taxed market is stabilized while volatility in the tax haven strongly increases. For markets of different size, Hanke *et al.* (2010) conclude that if the tax is introduced in the large (small) market, volatility decreases (increases) there.

[12] See Mike and Farmer (2008).

[13] See also Maslov (2000) and Challet and Stinchcombe (2001).

[14] See Cliff and Bruten (1997). Interestingly, Gode and Sunder (1993) were the first to introduce the zero intelligence framework in a trading setup. The authors show that the double auction mechanism ensures allocative efficiency irrespective of the level of rationality of the agents.

[15] To our knowledge, Chiarella and Iori (2002) were the first to incorporate trading strategies into a CDA setup. The authors state that without these strategies, it is impossible to generate realistic time-series. See also Chiarella *et al.* (2009).

work combined with a once-a-day supply-demand based market-clearing rule, deleting all orders not executed during the clearing session and thus strongly limiting their impact. In this setting the FTT is found to be destabilizing. Pellizzari and Westerhoff (2009) compare the effects of the FTT in different market settings. The main finding is that the FTT destabilizes a CDA market, while it stabilizes a dealership market where specialists provide abundant liquidity. However, these studies suffer from the assumption that all agents act with equal probability, i.e. they neglect the importance of investment horizons.[16] In this way, the FTT's effect of more severe taxation of short-term speculation is missed.[17] Furthermore, possible adjustments of agents' trading strategies in response to the FTT are often missing or at least strongly limited. We incorporate learning of agents as well. Another novelty compared to the existing literature is that the limit orders in our model emerge from a rule-based decision process, rather than from a pure zero intelligence framework.

In our model two groups of agents compete in the market: Noise traders act as liquidity providers, by posting random orders. Informed traders use information about past prices and the fundamental value when forming their price expectations. As in Youssefmir *et al.* (1998), their price expectations depend on three different time horizons (see Figure 2.1): the investment horizon (denoted by H^w) basically defines how often a particular agent acts and how long his planning horizon is when making investment decisions. Two different trend horizons model the trend chasing behavior of agents: the backward trend horizon (H^b) defines how many past price observations are relevant when calculating the trend. The forward trend horizon (H^t) defines how long the agent expects his calculated trend to last before the price will start returning to the fundamental value. This setting is very flexible concerning the strategies and we essentially allow for all combinations of time horizons within a certain set. Most importantly, the relative size of forward trend and investment horizon defines whether an agent is a chartist, a fundamentalist or something in between.[18]

In order to endogenize agents' strategies, we employ a Genetic Algorithm (GA). GAs constitute a class of search, adaptation and optimization techniques based on the principles of natural evolution. In this way, switches be-

[16] The importance of time-scales is a relatively recent research topic. See for example Zumbach and Lynch (2001) and Borland and Bouchaud (2005).

[17] See Anufriev and Bottazzi (2004) for the importance of investment horizons. In an infinite-liquidity model, Demary (2010) incorporates investment horizons and finds that investment horizons increase for small tax rates.

[18] Below, we will impose symmetry between the trend horizons, i.e. set $H^b = H^t$.

tween strategies depend upon the relative strength of each strategy.[19] Consequently, the GA allows us to characterize the 'optimal' (time-varying) strategies for different tax rates, where we are mainly interested in the adjustments in investment horizons.

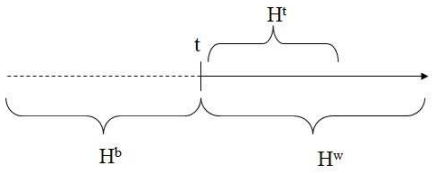

Fig. 2.1: Time-horizons in the model. In order to reduce the complexity of the model, we will set $H^b = H^t$ in the following. More details can be found below.

Our main conclusions can be summarized as follows: First, the model is able to replicate certain stylized facts of real financial time-series for several parameter combinations, e.g. the model replicates the building up and bursting of price bubbles. Second, for both fixed and flexible behavioral strategies, we find the usual trade-off between monetary revenues (a kind of Laffer curve) and stability, as higher tax revenues come along with higher volatility. This finding is in line with the results from the existing literature. However, we find somewhat different results for very small and large tax rates, indicating that the effects of the tax may not be entirely negative. In any case, the tax allows to generate substantial tax revenues, which could be used for a number of productive purposes.

The remainder of this paper is structured as follows: Section 2.2 introduces the structure of the model. Section 2.3 presents pseudo-empirical results and Section 2.4 concludes.

2.2 Model

2.2.1 CDA and Information

The basic model structure is as follows: the financial market consists of N heterogenous agents trading one asset (which pays no dividend and has

[19] See Lux and Schornstein (2005). Our usage of the GA allows to analyze the evolution of investment strategies on a micro-level. See LeBaron *et al.* (1999), Lux and Marchesi (1999, 2000) as well.

fixed supply) against cash. Cash earns zero interest, so there are no interest payments (or they are spent elsewhere). The market is order-driven and the quoted price p_t (midprice) is the average of the best ask (a_1) and best bid (b_1) in the limit-order book (LOB), while $a_1 - b_1 > 0$ is the bid-ask spread. In case there are no orders in the LOB, the quoted price is simply the last quoted price. Prices are discrete and can only be submitted on a specified grid, defined by the tick-size Δ.[20]

We simulate a CDA market, where two types of orders exist: a market order specifies to buy or sell a certain amount of the asset at the best available price. A limit order additionally specifies a limit price at which the agent is still willing to trade. In general, market orders are guaranteed execution but not price, since with a market order a trader is assured that it will be executed against the best price in the LOB within a short amount of time. Limit orders, on the other hand are guaranteed price but not execution as they will only be executed at, or below (above) for buy (sell) orders, the specified price which may never happen if no matching order is found. Each transaction involves a market order transacting against an existing limit order. Generally speaking, patient investors are more likely to place limit orders, while impatient investors place market orders. The price of immediacy is simply the bid-ask spread. Thus, choosing a limit price is a strategic decision that induces a trade-off between patience and (expected) profit.[21] The price dynamics within the LOB are therefore driven by three forces: limit order arrivals, market order arrivals (i.e. trades) and cancellations of limit orders.[22]

Table 2.1 illustrates the structure of the CDA: Buy orders are stored on the bid side (left), while sell orders are on the ask side (right). The two relevant features are price-priority and time-priority. Price-priority means that the best orders are placed on top of the book, i.e. the order with the highest bid price (best bid) and the order with the lowest ask price (best ask). Obviously, orders stored in the LOB cannot be executed immediately: in the example, the best bid (100.50) is smaller than the best ask (101.50) such that currently no trade is possible. Time-priority means that, after providing price-priority, orders with the same limit price are sorted according to submission date.[23] Therefore the best bid is placed above the second best

[20] In the future, it might be interesting to look at the effects of tick size changes in the market. It has been shown that the effects of tick size changes are more subtle than previously thought, see e.g. Ahn *et al.* (1996) and La Spada *et al.* (2011).

[21] See Harris and Hasbrouck (1996).

[22] There is a growing literature on the stylized facts of LOB data, see e.g. Bouchaud *et al.* (2002), Farmer *et al.* (2004), and Mike and Farmer (2008).

[23] Another alternative would be price-size priority, where large orders enjoy priority.

Bid			Ask		
Price	Quantity	Time	Price	Quantity	Time
100.50	20	12:38:39	101.50	10	12:15:01
100.50	10	12:42:08	105.00	5	12:28:40
95.60	8	12:10:52	110.50	10	09:01:05
87.90	5	10:15:23	125.50	8	12:40:18

Tab. 2.1: Example for a LOB at a certain point in time.

bid (with the same limit price), since it was submitted earlier to the LOB. In the example the quoted price would be $\frac{a_1+b_1}{2} = 101.00$. Note however, that this quoted price is just a proxy for the price of an immediate transaction: For example, assume there arrives a new sell order with a quantity of 25 and limit price 100.00. In this case the order is marketable, such that the offered 25 assets are sold at a price of 100.00, which differs from the quoted price of 101.00.[24]

Despite disregarding dividends,[25] we assume that the fundamental value p_f of the asset is positive and follows a geometric Brownian motion of the form

$$p^f_{t+1} = p^f_t \, \mathrm{e}(\mu_{t+1}) \tag{2.1}$$

where the noise term μ_t is iiN$(0, \sigma_\mu)$, $\sigma_\mu \geq 0$ is the corresponding volatility and $\mathrm{e}(\cdot)$ is the exponential function. In the following, we will set $\sigma_\mu = 0$ for simplicity and thus assume a constant fundamental value.[26] Only informed agents know the fundamental value, whereas the state of the LOB and the

Although no major exchange currently uses price-size priority, Nasdaq currently conducts research in this direction.

[24] Note how time- and price-priority favor the buying agent, i.e. the trade initiator, in the Example: He submitted a limit price of 100.50 but only pays 100.00.

[25] Dividend payments are negligible on a short-term basis, since they are only paid once a year and usually only have a small effect on wealth. Ignoring dividend payments simplifies the analysis, since (without the FTT and with cash earning zero interest) the total wealth is constant. With a positive tax rate, however, the total amount of cash will be decreasing over time. Therefore, we redistribute tax revenues equally among all agents for positive tax rates to keep the total wealth constant.

[26] We model trading dynamics on very short time-scales where the fundamental value is unlikely to change significantly. Furthermore, this assumption makes it possible to ignore adverse selection problems due to news arrival. As long as the fundamental volatility is relatively small (compared to the volatility of noise traders expectations), this does not affect the qualitative results. In the future it would be interesting to incorporate discrete jumps of the fundamental value to make the model more realistic.

history of the quoted prices are public information.[27] Each agent is initially endowed with $S_0^i = N_s$ assets and $C_0^i = N_s p_0$ units of cash. We impose short-selling and capital constraints, such that $S_t^i, C_t^i \geq 0$ for all t, i.[28]

2.2.2 Trader Types

Two groups of traders, differing in the way they form their price expectations and choose their limit prices, compete in the market: there is a fraction $\theta \in [0, 1]$ of informed traders and a fraction $(1 - \theta)$ of noise traders. Hence there are $N\theta = N_\theta$ informed and $(N - N_\theta)$ noise traders.[29] Each agent acts once every $H^{w,i}$ time-steps on average, which we model based on Poisson waiting times with expected value $H^{w,i}$ individually for each agent.[30] Submitted orders will remain at most $H^{w,i}$ time-steps in the LOB, after which the orders will be removed. In section 2.2.4 we will introduce another channel for the cancellation of the noise traders' orders.

One general point worth mentioning is that all agents act strategically in terms of their order submission. This means that all agents create two orders, where the first order (with corresponding limit price $p^{(1)}$) is always being sent to the LOB, while the submission of the second order (limit price $p^{(2)}$) depends on whether the first order was fully executed.[31] For example, an agent buying the asset today will try to sell it again at a higher price later on. Quite interestingly, while this 'buy-low/sell-high' framework is straightforward and incorporates the aim of wealth/utility maximization, we do not know of a single study dealing with these conditional (take-profit) orders within a simulation model.

Noise Traders

Noise traders are usually assumed to trade for personal but not for speculative reasons. In our setting, noise traders have random price expectations and act as liquidity providers in order to preclude no-trade situations. Noise traders' limit prices are chosen randomly around the current best bids and

[27] Of course, 'true' fundamentals are unobservable in reality. Another interesting feature would be to model costly acquisition of the fundamental value.

[28] This assumption obviously affects the admissible range of order sizes, see Section 2.2.3.

[29] Liquidity providers would be another possible label for the group of noise traders. The term noise traders however, emphasizes the random nature of those agents random limit price determination.

[30] Thus the probability of a particular agent being chosen equals $(H^{w,i})^{-1}$. Agents with relatively small investment horizons are thus acting more frequently than those with longer horizons.

[31] In real markets the second order corresponds to a take-profit order.

asks according to

$$p^{(k)} = p_t \, \mathrm{e}(\epsilon_t^k), \tag{2.2}$$

with $k = 1, 2$ and ϵ denoting iiN$(0, \sigma_\epsilon)$ random numbers, with $\sigma_\epsilon > 0$, independent of μ (the fundamental increments).[32] With this approach, we directly obtain two limit prices, $p^{(1)}$ for the first (unconditional) order and and $p^{(2)}$ for the second (conditional) take-profit order. Based on these limit prices, we identify the market side that the trader acts upon by comparing the limit prices, as we impose that their orders should not create a sure loss, i.e. the agent buys first and sells later if $p^{(2)} > p^{(1)}$, or the agent sells first and buys later if $p^{(2)} < p^{(1)}$.[33] Note that we could make the distribution of ϵ explicitly dependent on other variables, for example positively related to the historical volatility of returns. This type of volatility feedback, i.e. having noise traders choose their limit prices from a broader distribution when historical volatility is large, would then for example generate volatility clustering mechanically (see Section 2.3.1).[34] However, given that we are more interested in the behavior of informed traders and their effect on the system's properties, we are reluctant to impose this feedback and hold σ_ϵ constant in the following.[35]

We would like to stress here that we incorporate noise traders as liquidity providers in our model, since, as will become clear in the next section, it is possible that many of the informed traders appear on the same market side. Thus, noise traders provide liquidity when the informed agents are not willing to do so or at least not sufficiently to generate trades (and hence price changes).[36] Given the random structure of their limit prices, noise traders

[32] Furthermore, we assume that $\sigma_\epsilon > \sigma_\mu$. This assumption ensures some form of excess volatility, where σ_ϵ is essentially a scaling parameter for the return distribution.

[33] In everything that follows, we treat market orders as limit orders with limit prices equal to the best opposite price (marketable limit orders). For example, an agent will submit a buy limit price at most equal to the best ask. This assumption simplifies taking into account budget constraints at order submission, since otherwise the trading price would be uncertain.

[34] For example, the noise traders in Raberto *et al.* (2003) are constructed exactly in such a way, i.e. in their model informed traders are not necessary to reproduce volatility clustering and excess kurtosis. However, this is a very 'direct' way to guarantee volatility clustering in a model. It is not clear why agents should behave like this and there is in fact some evidence that past price volatility tends to lead the arrival of limit orders, see Zovko and Farmer (2002).

[35] Note that since the width of the distribution is fixed, noise traders are more likely to submit market orders when the spread is small, while they are more likely to submit limit orders when the spread is large. This is in line with empirical findings, e.g. Biais *et al.* (1995), Bae *et al.*(2003), and Foucault *et al.* (2005).

[36] Note that, in contrast to other studies, we find that many of the stylized facts cannot be easily replicated using noise traders alone.

tend to lose money to the informed traders on average, in particular when there are pronounced trends. A relatively small number of noise traders is already sufficient for the model to work. However, particularly in case that the GA is deactivated, with such a small number of noise traders the generated bubbles appear relatively smooth, such that prices and returns would be autocorrelated. Therefore, θ will not be too large during the simulations.

Informed Traders: Chartists and Fundamentalists

Price Expectations. Whether an informed agent buys or sells the asset depends on his expectation of the asset's future price at the end of his investment horizon. When forming price expectations, informed traders use information about past prices and fundamental values. Expectations evolve, following Youssefmir *et al.* (1998), as

$$\frac{d\hat{p}^i_{t+\tau}}{d\tau} = -\frac{\hat{p}^i_{t+\tau-1} - p^f_t}{H^{t,i}} + \left(T^i_t + \frac{p_t - p^f_t}{H^{t,i}} \right) \mathrm{e}\left(-\frac{\tau}{H^{t,i}} \right), \tag{2.3}$$

where $\hat{p}^i_t$ is agent i's expected price at t, T^i_t is the calculated trend and $H^{t,i}$ is the forward trend horizon over which agent i expects the trend to last. The trend itself is an weighted average rate of price changes over a backward horizon $H^{b,i}$ of the form

$$T^i_t = \frac{1}{H^{b,i}} \int_{t_0}^{t} \frac{dp}{d\tau} \mathrm{e}\left(-\frac{t-\tau}{H^{b,i}} \right) d\tau, \tag{2.4}$$

where dp is the price change between $(t-\tau)$ and t and $H^{b,i}$ is the backward trend horizon of i. As noted by Youssefmir *et al.* (1998), Eq. (2.4) can be integrated as $T^i_t = \frac{p_t - \langle p_t \rangle_{H^{b,i}}}{H^{b,i}}$, where $\langle p_t \rangle_{H^{b,i}}$ is the exponential average price over the horizon $H^{b,i}$. Thus the trend measures the deviation from the moving average of prices, which is a popular approach among technical analysts.

The evolution of trends can be obtained by taking the time derivative of Eq. (2.4) which yields

$$\frac{dT^i_t}{dt} = \frac{1}{H^{b,i}} \left(\frac{dp}{dt} - T^i_t \right). \tag{2.5}$$

Subject to the boundary condition $\hat{p}^i_t = p_t$, each agent formulates his expected price development over the next $H^{w,i}$ time-steps via Eq. (2.3) using the calculated trend from Eq. (2.5). This system incorporates, depending on the corresponding horizons, chartist and fundamentalist components. In

principle, all agents are fundamentalists in the sense that for $H^{w,i} \to \infty$ (given $H^{t,i}$) the expected price will collapse towards the fundamental value. As agents do not have infinite investment horizons in general, the relative magnitude of $H^{t,i}$ to $H^{w,i}$ matters: agents with a small value of $H^{t,i}/H^{w,i}$ can be considered as fundamentalists, a large value indicates a more chartist strategy, and intermediate values are a combination of both. We should stress already here that we found heterogeneity in H^b to be of minor importance compared to H^w and H^t. Therefore, in everything that follows, we simply take $H^{b,i} = H^{t,i}$ for all agents. Note that this substantially reduces the total number of strategies.

Technically, the nonlinear price expectations are influenced by three terms: first, agents expect the observed trend and the difference between price and fundamental value to continue in the near term (this is the second part on the right-hand side of Eq. (2.3)). However, the influence of this term decreases for increasing τ (and hence for large $H^{w,i}$). Second, via the decreasing impact of the calculated trend, the agent expects the price to eventually relax towards the fundamental value at a rate of $-\frac{\hat{p}_t^i - p_t^f}{H^{t,i}}$. Note that fundamentalism is defined in terms of the expected price at the end of the investment horizon, but a fundamentalist may nevertheless try to make a profit based on short-term trends.

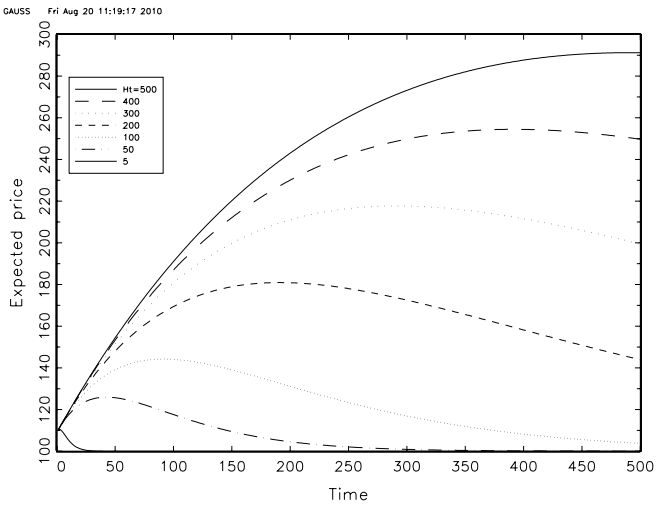

Fig. 2.2: Example: Development of expected price for different forward trend horizons, $T_t^i = 1$, $H^{w,i} = 500$, $p_t = 110$ and $p_t^f = 100$.

In the following, a discretized version (i.e. $dt = d\tau = 1$) of Eqs. (2.3)

and (2.5) will be employed. This produces a time-series of the expected price development for each agent. As an example, consider Figure 2.2 which plots, for different forward trend horizons, the expected price development of an agent with $T_t^i = 1$ and an investment horizon of 500. The current price is 110 and the fundamental value equals 100. Obviously, for relatively small trend horizons the agent expects the price to revert towards the fundamental value soon. For larger trend horizons, the agent expects the trend to last in the near term but the price to revert towards the fundamental value at the end of his planning period. For very large trend horizons, this no longer holds. Consequently there is a low level of speculation for small trend horizons, in which case the dynamics are dominated by the first term in Eq. (2.3).

Note that this way of modelling informed agents implies that all of them are trend-followers, at least to some extent. This clearly affects the autocorrelation of prices and returns and makes the (in-)efficiency of the market a serious issue, at least when the fraction of informed traders is not too small.[37] We will comment on this in more detail below.

Limit Prices. The limit price determination of informed agents can be split into three parts: In the first part the agent uses Eqs. (2.3) and (2.5) to forecast the evolution of the midprice between t and $t + H^{w,i}$. In case the agent expects the price to be higher (lower) than the current price he will submit a buy (sell) order. The limit prices of this order and the corresponding conditional order depend on the expected development of the price between t and $t + H^{w,i}$.

In the second step, the agent uses information about the midprice and a proxy for the expected price volatility to form more detailed expectations about the best bid and ask over time. As proxy for the expected price volatility, we use the average distance between the current best and second best prices on the two market sides.[38] This is calculated as

$$\hat{\sigma}_t^i = \frac{((a_2 - a_1) + (b_1 - b_2))}{2}, \tag{2.6}$$

where a_2 and b_2 denote the second best ask and bid prices, respectively. In

[37] In an earlier version of the paper, we mitigated this problem, by allowing informed traders to employ a contrarian strategy, i.e. to use $-T_t^i$ rather than T_t^i in Eq. (2.3) with some probability (see Giardina and Bouchaud (2003a) as well). As it turned out, the structure of the system made contrarian behavior very profitable, so that the price eventually collapsed towards the fundamental value and remained there during the rest of the simulation.

[38] The motivation for using the average gap is based on the finding that large price changes are in fact due to gaps in the LOB, see e.g. Farmer *et al.* (2004) and Farmer and Lillo (2004).

this view, the average price change due to immediate market orders wiping out the best prices on either market side is being calculated. Note that this is the only channel in our model where (informed) traders use higher-order information on limit prices. We leave it to future research to model the information usage of LOB data in more detail.

Equipped with expectations about the midprice, the corresponding price volatility, and the decision to buy or sell from the first step, the agent then chooses his limit prices as follows: Defining

$$\begin{aligned} \hat{p}^i_{\max} &= \max\{\hat{p}_{t:\tau:t+H^{w,i}}\} \\ \hat{p}^i_{\min} &= \min\{\hat{p}_{t:\tau:t+H^{w,i}}\} \end{aligned} \tag{2.7}$$

gives the maximum and minimum of the expected midprice over the investment horizon of agent i. For an initial buy order, the agent then has to decide between buying the asset right away using a market order with price a_1 or setting a limit order at the minimum expected ask, i.e. $\hat{p}^i_{\min} + 0.5\hat{\sigma}$. Quite intuitively, he will choose the minimum of the two, in order to take favorable future developments into account. Therefore, if he expects the best ask to drop significantly below the current value in the near future, he will submit a limit order with a price below the current best ask. For the limit price of the conditional sell order the agent has to decide between the maximum expected bid, i.e. $\hat{p}^i_{\max} - 0.5\hat{\sigma}^i_t$ and the expected bid at the end of his investment horizon, and naturally takes the maximum of the two. Similar arguments can be used for the case of an initial sell order. More formally, the strategies of finding limit prices can be written as follows:

Definition 1. *Buy limit prices:* $\hat{p}_{t+H^{w,i}} > p_t$.

$$\begin{aligned} p^{(1)} &= b^i = \min(a_1, \hat{p}^i_{\min} + 0.5\hat{\sigma}^i_t) \\ p^{(2)} &= a^i = \max(\hat{p}^i_{\max} - 0.5\hat{\sigma}^i_t, \hat{p}_{t+H^{w,i}} - 0.5\hat{\sigma}^i_t). \end{aligned} \tag{2.8}$$

Definition 2. *Sell limit prices:* $\hat{p}_{t+H^{w,i}} < p_t$.

$$\begin{aligned} p^{(1)} &= a^i = \max(b_1, \hat{p}^i_{\max} - 0.5\hat{\sigma}^i_t) \\ p^{(2)} &= b^i = \min(\hat{p}^i_{\min} + 0.5\hat{\sigma}^i_t, \hat{p}_{t+H^{w,i}} + 0.5\hat{\sigma}^i_t). \end{aligned} \tag{2.9}$$

With this definition it may happen, in particular for large spreads, small trends and/or small deviations from the fundamental value, that the two limit prices are not in line with the agents' price expectations. For example, an agent with an expected price increase might end up with limit prices $p^{(1)} > p^{(2)}$. To ensure consistency, such orders will not be submitted to the LOB.

Figure 2.3 illustrates the concept for a sell order: Again the current price is equal to 110 and the fundamental value equals 100. The investment horizon is $H^{w,i} = 500$. The expected price at the end of the investment horizon is below the current price, therefore the agent will first sell the asset and try to buy it back at a lower price. Since the agent expects a positive trend to continue in the near term, $\hat{p}^i_{\max}$ exceeds the current price, while $\hat{p}^i_{\min}$ coincides with $\hat{p}_{t+H^{w,i}}$. Given his risk aversion, he places a sell order with limit price $a^i = \hat{p}^i_{\max} - 0.5\sigma^i_t$ and a conditional buy order with limit price $b^i = \hat{p}_{t+H^{w,i}} + 0.5\sigma^i_t$. The expected return (after tax) of the agent equals $r^e = |\ln(p^{(2)}/p^{(1)})| - \chi$, with χ denoting the two-sided tax rate.

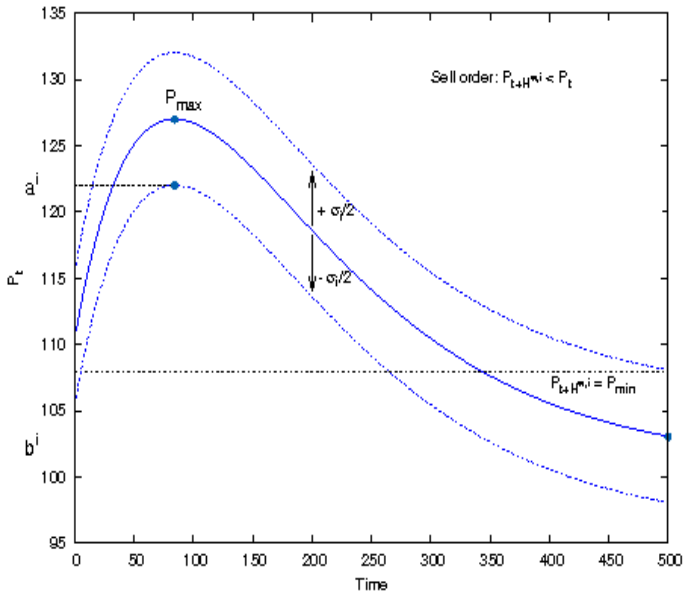

Fig. 2.3: Example: Limit price determination for a sell order.

Price Bubbles and FTT. At this point, a brief explanation on the building up and bursting of price bubbles, and their relation to the FTT, is in order:[39] At the beginning of the simulation, random strategies are assigned to the informed traders. Thus, approximately half of the population is willing to buy the asset while the other half is willing to sell. In any case, the price will move either upwards or downwards. Suppose it moves upwards, then the trend signals of informed agents with high-frequency strategies will turn positive. This

[39] See Giardina and Bouchaud (2003b) for a similar explanation.

can be the beginning of the bubble, where exactly those traders induce additional positive price changes through their positive demand, which in turn affects the calculated trend variables of lower-frequency traders. Thus, the buy pressure increases even more, where the sellers of the asset are informed agents using fundamentalist strategies or noise traders. The larger the distortion between price and fundamental value, the more fragile the bubble becomes: first, the buying power of potential buyers decreases due to the higher price and lower amount of cash available. This implies that the trend signals of high-frequency traders become smaller, possibly even turning negative. A random negative shock, induced by a fundamentalist trader or a noise trader, raises the possibility of a sudden downturn, where again the trend signals of lower-frequency traders follow those of the high-frequency traders. In this case a negative bubble may appear and the described process starts again, however now with negative signs.

This description of a bubble process also clarifies the arguments favoring a FTT: High-frequency oriented chartists have an incentive to create bubbles, since trend-following strategies are only profitable when there is some trend to follow. These traders are also the first to leave the sinking ship when the bubble bursts, again amplifying the negative trend. This is particularly true for high-frequency strategies which would be very costly when the buy/sell signal is constantly wrong. Thus, with a large proportion of fundamental traders in the population, the initial bubble might not even build up. The FTT should be one way to increase the proportion of fundamentalists, as it reduces the profitability of marginally profitable trades in the first place. In this way, it should lower the frequency of 'false' signals appearing in the system. In Section 2.3 we will discuss the effects of such a tax within our model.

2.2.3 Asset Demand

Any (limit) order is a commitment to trade (at most) a certain quantity at the specified limit price. In this part, we focus on the determination of the order sizes. While defining a strategy that maps price expectations into order sizes appears to be a trivial task, our imposed short-selling and budget constraints complicate things considerably.[40] In our setting all agents are initially endowed with the same level and composition of wealth. The wealth of agent i at time t is simply $W_t^i = p_t S_t^i + C_t^i$ and his wealth at the next date is $W_{t+1}^i = p_{t+1} S_t^i + C_t^i = W_t^i + dp_t S_t^i$. Therefore, given the price expectations of the agent, the trading behavior reduces to an optimization problem with

[40] See Franke and Asada (2008).

respect to the asset holdings S_t^i.

As future price developments are uncertain, we assume agents to be risk-averse. To some extent, this risk-aversion is reflected in the determination of the informed agents' limit prices in Eqs. (2.8-2.9), but should also be present in the choice of the order sizes. The usual approach in the literature is to either use some form of utility maximization (often CARA or CRRA utility functions)[41], to use random order sizes[42], or use rules-of-thumb, most importantly unit orders with constant size equal to one.[43]. There are problems with all these approaches: First, much of the literature favors the CARA approach, mainly due to the fact that under this approach the desired asset holdings are independent of wealth (for Gaussian returns). However, when taking accumulated asset positions into account, the actual order size (desired minus actual holdings) is by definition not independent of wealth.[44] Imposing short-selling and credit constraints will then yield order sizes that are not in line with economic principles. From an economic viewpoint, random and unit order sizes are not very appealing as well, as they imply that order sizes are independent of current wealth and behavioral parameters. We overcome these problems by combining economic variables with rule-of thumb behavior.[45]

The order size depends on three crucial variables: the agent's aggressiveness, his available resources, and, in case of a market order, the liquidity available at the best opposite price. We specify the demand function as

$$d_t^i = \begin{cases} \left[\alpha^i(\cdot)\frac{C_t^i}{(1+\chi/2)b^i}\right] & \text{if buy order} \\ [\alpha^i(\cdot)S_t^i] & \text{if sell order,} \end{cases} \tag{2.10}$$

where $[x]$ denotes rounding towards minus infinity and $\alpha^i(\cdot) \in (0, 1]$ is an aggressiveness parameter determining the proportion of cash/assets an agent actually wants to use for investment. Note that only half the tax rate is taken into account since buyer and seller will share the tax burden equally.

In principle, $\alpha^i(\cdot)$ could take any functional form, which is why we left the arguments unspecified. Following Giardina and Bouchaud (2003b) and Martinez-Jaramillo (2007), we simply set

$$\alpha^i(\cdot) = \bar{\alpha}, \tag{2.11}$$

[41] See e.g. Chiarella *et al.* (2009) and Bottazzi *et al.* (2005) for an application of CARA asset demand in an order-driven market.

[42] See e.g. the random traders in Mannaro *et al.* (2008).

[43] See e.g. Pellizzari and Westerhoff (2009).

[44] See Franke (2008).

[45] Note that rule-of-thumb behavior, although having the weakness of being 'ad-hoc', is more realistic in terms of how actual people make decisions, see Gigerenzer (2008). In this regard, we are planning to implement the findings of Lachapelle and Challet (2010).

i.e. a fixed parameter identical for all agents.[46] Note that, by construction, the agent's budget constraint is never binding since he willingly only uses a fraction of his wealth to invest in the risky asset.[47] Thus, agents are reluctant to submit very large orders which are likely to have strong market impact.[48]

Additionally, when submitting market orders, agents will at most trade the amount available at the opposite best price, denoted as d_1, i.e.

$$\bar{d}_t^i = \min(d_1, d_t^i). \tag{2.12}$$

This takes into account the empirical fact that orders removing more than the volume available at the opposite best quote are rare.[49]

A brief note on the effects of the FTT is in order here. The tax affects transaction volumes negatively in two ways: first via Eq. (2.10) single order sizes necessarily become smaller due to the negative impact of the tax. Second, agents will only post orders with an expected return larger than the tax rate. Without the tax, the only requirement to post an order is that the two limit prices differ by at least one tick. Depending on the tax rate, orders with rather small expected returns will not be posted anymore. In this way, possible liquidity reductions for higher tax rates may be both due to noise and informed traders.[50]

2.2.4 Cancellation and Trading Process

Cancellation. The price dynamics of the LOB are driven by the non-trivial interplay between liquidity takers and liquidity providers.[51] Prices may change due to the arrival of market and limit orders, and the cancellation of existing limit orders. Limit orders can disappear from the LOB in four different ways: 1) a newly arriving market order executes against the existing

[46] As an alternative, we could make agents' aggressiveness explicitly dependent on economic variables (such as volatility) or on the relative weight of chartism and fundamentalism. In such a setting, the aggressiveness of chartists would be higher, since chartists are usually found to be less risk-averse than fundamentalists, see e.g. Menkhoff and Schmidt (2005). In order to reduce the complexity of the current model, we leave that for future research.

[47] This ensures that agents do not run out of assets/cash.

[48] See Harris (2003), Ch. 15.

[49] Farmer and Lillo (2004) have shown that roughly 87% of the market orders creating an immediate price change have a volume equal to the volume at the opposite best, while 97% of the market orders creating an immediate price change have a volume at most of the sum of volumes available at the two best opposite prices.

[50] As a robustness check, we could also force noise traders to post orders with negative returns. In this way, we could investigate the effect of the tax in the case that market makers do not have to pay the tax.

[51] See Bouchaud *et al.* (2003), Bouchaud *et al.* (2004) and Toth *et al.* (2011).

limit order, 2) a limit order can remain at most H^w intraday time-steps in the LOB, afterwards it will be deleted automatically, 3) an agent being chosen to act again will cancel any outstanding orders, 4) an agent cancels his outstanding orders autonomously (possibly even at random), thus the order 'evaporates'.

While the first three channels are obvious, the fourth channel has usually attracted not as much attention in agent-based modelling. However, Farmer *et al.* (2004) find that cancellation occurs roughly 32% of the time at the best price and 68% of the time inside the book. Challet and Stinchcombe (2001) find that typically 80% (20%) of all orders are cancelled (executed).[52] The main argument for these huge numbers is that placing and cancelling limit orders is usually free of charge and therefore a strategic opportunity for all types of traders.[53] Thus, when modelling the LOB, cancellation of orders cannot be neglected.[54]

There may well be an important link between investment horizons and the average order lifetime: It is widely believed that a power-law in the distribution of investment horizons may be the driving force behind the power-law tail of price changes.[55] However, when investment horizons are indeed fat tailed, the same should be true for order lifetimes. And indeed, the lifetime of an order increases as one moves away from the best bid/ask.[56] Patient investors are therefore less likely to cancel their orders, as found in Potters and Bouchaud (2003).

For the sake of simplicity, we assume a Poisson process of order cancellation for the noise traders as in Daniels *et al.* (2001). At the beginning of each time-step, each noise trader cancels his outstanding orders with probability π_{canc}. For informed traders, we neglect this channel of order cancellation, since this would inject a significant amount of (additional) randomness to their strategies.

Trading Process. Time evolves continuously within our CDA trading framework. For simulation purposes, we discretize time such that between two

[52] Similar values were reported in Lachapelle and Challet (2010). See Bouchaud *et al.* (2002) as well.

[53] See e.g. Cao *et al.* (2008). Note that fundamental traders will post limit orders with prices far away from the best quotes. If the agent is not patient enough, he will cancel his order prematurely.

[54] See also Challet and Stinchcombe (2001).

[55] There is some indirect evidence of a power-law distribution in time-scales, see e.g. Lynch and Zumbach (2003). For a theoretical argument, see Lillo (2007). There, heterogeneity in the time horizons is identified as the most likely explanation of the fat-tailed distribution of limit-order prices.

[56] See also Challet and Stinchcombe (2001) and Farmer *et al.* (2004).

points in time the algorithm does the following:

1. Update fundamental value.

2. Update outstanding orders (reduce order lifetime). Delete expired orders. Cancel outstanding orders of noise traders with probability π_{canc}.

3. Agents decide whether they will be active during this time-step, and approach the market in a random order. During these (intraday) time-steps[57] the following happens:

   - The active agent i deletes his outstanding orders and (possibly) generates his two orders. The first order is submitted to the LOB.
   - Execute all possible trades (sequentially) taking into account price-time priority and send conditional orders to the LOB. Repeat until no trades possible anymore. Update wealth continuously.

4. Save midprice, update trends and wealth. At times, redistribute the tax revenues across all agents.

5. Optional: carry out the GA (and delete informed traders' outstanding orders).

2.2.5 GA Learning

A major aim of this paper is to investigate informed traders' optimal strategies in response to the FTT, where a strategy can be fully described by the agents' time horizons. Strategies evolve over time through a 'social learning' GA,[58] where agents tend to imitate and recombine successful ideas of other agents.[59]

Given a well-defined search space, the proposed solutions are encoded in strings (chromosomes) of length L using binary alphabet, i.e. the maximum number of possible solutions equals 2^L for each of the horizons. In order to reduce the search space, we fix H^t for all agents (equal to the average value in the admissible range) in the following and only optimize over H^w.[60] In order to increase the width of the search space, the parameter $\tilde{H}$ defines the minimum/maximum horizons and the number of time steps between adjacent

[57] Note that the number of intraday time-steps depends on the average investment horizon and the number of agents. Having a large number of agents and small average investment horizons obviously slows down the simulation significantly.

[58] Holland (1975) introduced GAs as stochastic search algorithms for numerical optimization.

[59] See Vriend (2000) for a comparison of 'social' and 'individual learning' GAs.

[60] Recall that we replaced the backward trend horizon by the forward trend horizon.

strategies. Using the binary matrix $B^w_{\{N_\theta \times L\}}$, we can represent informed traders' strategies as

$$H^{w,i} = \tilde{H}(1 + \sum_{j}^{L} 2^{(j-1)} B_j^{w,i}), \tag{2.13}$$

with $B_j^{w,i}$ denoting the jth bit of i's strategy,

The GA maintains a population of N_θ solution candidates and evaluates the quality of each according to a problem-specific fitness function. The genetic operators are the following:[61]

1. **Reproduction:** From the pool of all informed traders, N_θ we randomly select copies (with replacement) with probabilities depending on each strategy's relative fitness, i.e. prob(reproduce agent i's strategy) $= f_t^i / \sum f_t^i$ with f_t^i being the fitness of i's strategy (tournament selection).

2. **Crossover:** For a fraction π_{cross} of the population we randomly couple pairs of individuals and swap bits between them. The cutoff point is a random integer between 1 and $(L-1)$. The offsprings are constructed by combining the left of this position from parent one with that from the right-hand part of parent two and vice versa.

3. **Mutation:** With probability π_{mut}, each position within a string is altered to the other value.

Thus, a binary GA has the two parameters π_{cross} and π_{mut}.[62] The two parameters L and $\tilde{H}$ define the search space.

Obviously, the fitness function is the crucial element of the GA. In our setting, agents try to maximize wealth by optimizing their trading strategies. To keep things simple, we relate the fitness of a strategy to agent i's gross return on wealth since the last activation of the GA

$$f_t^i = \left(\frac{W_t^i}{W_{t-\omega^{\text{GA}}}^i} \right). \tag{2.14}$$

We calculate the values for all informed traders and compare them in the reproduction step.[63] In principle, it would be more useful to use a smoothed

[61] Note that we neglect the so-called 'election' operator, which ensures that the overall fitness of the new population actually exceeds that of the previous population. Given that the calculation of the potential fitness of new strategies is rather involved, we neglect this operator in the following.

[62] We also experimented with real-coded GAs, see Herrera *et al.* (1998), but found very similar results.

[63] In order to ensure that the selection pressure is not too large, we normalize the fitness values to the unit interval, where the best and worst performing strategies receive the values 1 and 0, respectively. In this way, poorly performing strategies are not necessarily being abandoned right away.

fitness measure rather than the raw returns on wealth as in Eq. (5.12). However, since the agents' strategies will change when the GA is being activated, such a smoothed value would not properly reflect whether a particular strategy was successful in the past. Additionally, agents with different time horizons would smoothen their fitness values differently, with high-frequency traders focusing more on short-term gains. We avoid both problems by using the above fitness measure and activating the GA every ω^{GA} time-periods on average. Note that we will not activate the GA too often, in order to ensure the meaningfulness of the fitness function.[64]

2.3 Pseudo-Empirical Results

In this section we present pseudo-empirical results from the model simulations. If not stated otherwise, the reported results are the outcome of Monte-Carlo simulations of $22,500$ time-steps (where we disregard an initial period of 2,500 time-steps), each of which repeated 20 times with different random seeds.[65] In the following, we always differentiate between two cases in our simulations, namely when the GA is switched off and when it is switched on. In the latter case it will be activated every ω^{GA} periods.

In order to get a feeling for the model's properties, we will first present time-series of single simulation runs for both cases. For such single runs, we will always present time-series for the most interesting variables (e.g. price and log-returns) and comment on certain statistical properties. Afterwards, we will investigate the effects of a FTT based on the Monte-Carlo approach. The basic parameter values used in our simulations and brief descriptions for all parameters are given in Table 2.2. In this baseline scenario, agents' time horizons are picked uniformly from the admissible range. However, we will also highlight the dependence of the system's properties on the choice of the horizons.

At this point, we should mention one of the main drawbacks of agent-based modelling, namely the large number of degrees of freedom in the choice of the parameter values. This holds even more, when modelling very complex decision-making of agents, as in our case. While one should employ empirical estimates whenever possible, in case there is no (and perhaps never will be one) empirical estimate, the modeler has to decide about this value. Modelers often denote this as 'calibration', which is a neat description for something which can be dangerously misleading. Obviously, a model cannot be robust to

[64] As a robustness check we use different values of ω^{GA} below.

[65] We will see below that the error bars usually tend to be quite small, so this small number of runs is indeed already sufficient.

Parameter		Value	Description
General			
$\tilde{H}$	=	20	Minimum horizon and distance between adjacent strategies
L	=	5	Parameter related to the number of strategies
N	=	500	Total number of agents
N_s	=	100	Parameter for initial endowment
p_0, p_0^f	=	100	Starting value: price/fundamental value
$\bar{\alpha}$	=	0.10	Order aggressiveness
Δ	=	10^{-3}	Tick size for the price
θ	=	0.20	Fraction of informed traders
π_{canc}	=	0.01	Cancellation probability (noise traders)
σ_ϵ	=	10^{-3}	Volatility for noise traders' expectations
σ_μ	=	0	Volatility of fundamental noise
χ	=	0	Tax rate
GA			
$\{H^w\}$			Set of horizons relevant for the GA
π_{cross}	=	0.2	Crossover probability
π_{mut}	=	10^{-3}	Mutation probability
ω^{GA}	=	1,000	Activation frequency

Tab. 2.2: Baseline parameter setting for the simulations. Note that we differentiate between the case with and without the GA.

changes in all parameters, but a model can only be considered relevant if (1) it is able to produce realistic dynamics for (economically) plausible parameters values, and (2) is robust with respect to changing certain parameters.

To stress this point, consider for example the parameter θ: What would be a reasonable value for the fraction of informed traders? A priori we should expect a relatively large number of agents to use information about past prices and the fundamental value when forming price expectations in real markets. If so, how may of those agents will be chartists and fundamentalists, respectively? Will it matter whether the GA is being activated? Within our model, we found θ to be a very important parameter for the time-series properties. Consequently, while there is always scope for fine-tuning of the parameters in order to obtain more 'realistic' time-series we found the qualitative results to be rather robust with respect to parameter changes as compared to our baseline scenario in Table 2.2.[66] For the effects of the FTT, we found the results to be quite robust as well.

[66] In principle, we would be happy to use an approach similar to Franke and Westerhoff (2012), where the 'optimal' parameters (with respect to the stylized facts) are estimated via moment-matching. However, given the complexity of our model, and the related high-number of degrees-of-freedom, parameter estimation would be prohibitive.

2.3.1 Properties

One obvious question is whether our model is able to replicate some of the stylized facts of empirical financial time-series. Without going into the details, the most basic stylized facts of asset prices and returns can be summarized as follows:[67]

- Martingale property (unit root) of prices: Price dynamics close to a random walk. Zero expected return, with only the very first lags positively autocorrelated (at least for high-frequency data).
- Fat-tailed return distribution: Positive excess kurtosis and power-law tails both imply more probability mass in the center and the tails of the return distribution (compared to Gaussian). Tail exponent around 3.
- Volatility clustering: Autoregressive dependence with very slow (hyperbolic) decay in various measures of volatility.

Recently, many more stylized facts of order-book data have been identified which we will not comment on in the following.[68]

Here, we do not aim to test quantitatively whether all of these stylized facts are present in the model. Rather we present several basic properties, which are illustrated by individual representative time-series from the model. After explaining the baseline scenario in detail, we move on to comment on certain statistical properties, when changing a set of selected parameters. In the next section we then move on to the introduction of a FTT, which can be seen as a simple parameter variation of χ.

Baseline Scenario and Dependence on θ

Without GA. Figure 2.6 shows the results for a single simulation run of the baseline scenario, with the GA deactivated. The top left panel shows the price and the fundamental value. We see that the price fluctuates around the fundamental value over time, and there is a continuous building up and bursting of bubbles. On the top right panel we see the corresponding log-returns, with kurtosis and skewness being 12.28 and -.04, respectively. Thus, the return distribution is highly non-normal but roughly symmetric around zero. In this regard, the center left panel further quantifies the fat-tail of the return distribution by means of the complementary cumulative distribution (ccdf) of the return time-series on a log-log scale.[69] The tail region shows

[67] See e.g. Lux (2008).

[68] See Bouchaud *et al.* (2008) for an extensive overview.

[69] Here we show a rank-frequency plot of the ccdf, which is one minus the cdf, on log-log scale.

a somewhat linear decay, with an estimated tail exponent of 3.50.[70] This value lies within the range observed in real markets. The center right panel shows the autocorrelation function (ACF) of the raw and absolute returns, including the 95% confidence intervals. For the raw returns (blue), the bid-ask bounce is the major source for the first few negative lags. However, due to the trend-chasing behavior of informed traders, larger values of θ lead to more autocorrelated prices and returns (see below), which wash-out the bid-ask bounce. For higher lags, the autocorrelations are marginally insignificant. For the absolute returns, the first lags are significantly positively autocorrelated, i.e. there is a small level of volatility clustering. However, the decay of the autocorrelation function is much faster than empirically observed. As we will see below, the trend-chasing behavior of informed traders tends to induce a somewhat larger level of autocorrelation in the raw returns, as compared to the absolute returns. The bottom left panel shows the bid-ask spread (in ticks). We should first note that the average spread is rather large, here with an average value of 6.09 ticks. We see some persistence in the spread, so large (small) spreads tend to be followed by large (small) spreads. The bottom right panel shows the transaction volumes, i.e. the number of stocks traded per time step. Quite interestingly, the Figure shows a smaller (if any) level of persistence in the transaction volumes, such that large price changes appear to be more or less unrelated to large volumes traded.

We checked that large price changes are in fact driven by gaps in the LOB (unreported), as proposed by Farmer and Lillo (2004). Furthermore we checked that, as long as there are some trends to follow, informed traders tend to gain, while noise traders tend to lose on average. Figure 2.5 illustrates this for the simulation run in Figure 2.4, where we show the average wealth difference between informed and noise traders over time, which steadily increases over time. Even though the interpretation is different, this finding is consistent with the results of Kyle (1985), where more informed traders tend to gain from noise traders. However, we should stress that these are only average values. There are also informed traders that lose substantial amounts of their wealth, in part to other informed traders, but also to some of the noise traders. In section 2.3.2 below we look at this issue in more detail, with a particular focus on the effects of the FTT on different subsets (in terms of trading strategies) of the population of informed traders.

As a next step, we briefly compare the findings for different values of θ. Figures 2.6 and 2.7 show the results for a single simulation run for $\theta = 0$ and

[70] For the estimation of the tail parameter, we used the usual Hill (1975) estimator based on the top 15% observations of the absolute returns, i.e. ignoring signs. The values are not affected by focusing on positive or negative returns only.

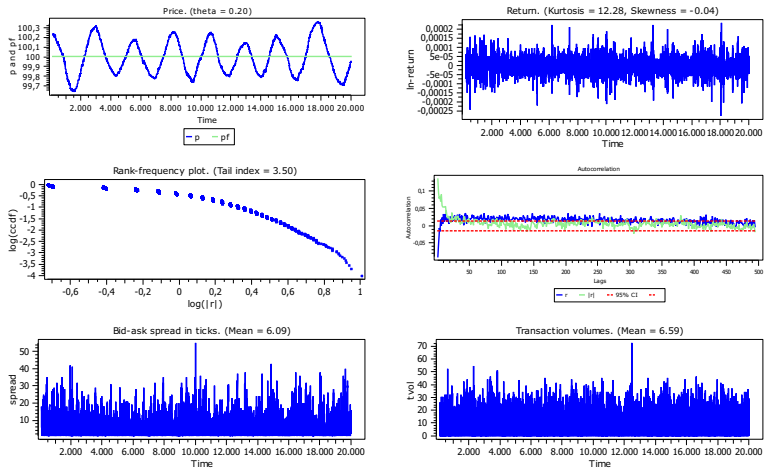

Fig. 2.4: Simulation results: Single run, $\theta = 0.2$, baseline scenario (GA off). Top left: price (blue) and fundamental value (green). Top right: log-returns. Center left: rank-frequency plot (ccdf, log-log-scale) of log-returns. Center right: Autocorrelation of raw (blue), absolute returns (green), and 95% confidence interval. Bottom left: bid-ask spread (in ticks). Bottom right: transaction volumes.

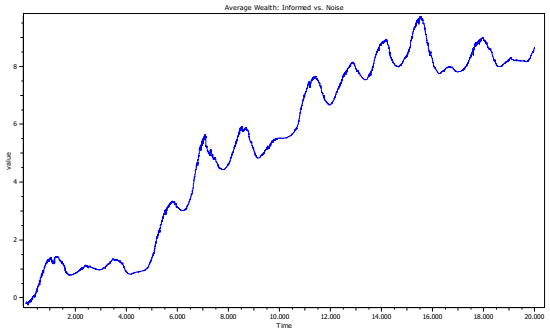

Fig. 2.5: Simulation results: Single run, baseline scenario (GA off). Dynamics of relative wealth of average informed trader vs. average noise trader.

.5, respectively. For $\theta = 0$, the price is completely unrelated to the fundamental value: here it continuously exceeds the fundamental value. Without informed traders there is no force that would push the price towards the fundamental value. Quite interestingly, the return distribution has a fat-tail already in the case with noise traders only (kurtosis 8.29, tail parameter 3.24) and the first lags of the ACFs are again significant. In contrast, there is no fundamental persistence in the returns, the bid-ask spread, and the transaction volumes, whatsoever. Also the bid-ask spread is substantially larger (19.88 ticks on average), while transaction volumes are smaller. This changes for large values of θ, where we see very smooth bubbles, which translate into highly autocorrelated raw and absolute returns, with the raw returns having a substantially larger positive autocorrelation at all lags, as compared to the absolute returns.[71] Choosing θ such that the ACF of the absolute returns decays hyperbolically results in a very inefficient market, with the ACF of the raw returns decaying similarly slowly. In part, we already saw this for the simulation of the baseline scenario. Still, our choice of θ in the baseline scenario takes this trade-off into account.

An important reason for the small level of volatility clustering is probably the absence of the leverage effect in our model.[72] The basic idea is that leverage would increase the necessity of portfolio adjustments after large negative price changes, i.e. large downward spirals.[73] We leave this model extension for future research.

Moving to the Monte-Carlo simulations for different values of θ, Figure 2.8 summarizes some of the above findings, where we show the average values (including plus and minus one standard deviation) of many runs. We find non-linear relationships for many of the variables. For example, the top left and top right panels show the kurtosis and the tail parameter. We see that for small θ a higher number of informed traders makes the tails somewhat fatter, while for intermediate values, informed traders tend to smoothen the price dynamics and thus lead to somewhat thinner tails. For very large values of θ the tails are the fattest. The center left panel shows a U-shaped relationship between the bid-ask spread and θ. To a much lesser extent this is also true for the distortion (i.e. the absolute percentage difference between the price and the fundamental value) and the volatility (i.e. the absolute price change). In contrast, the center right panel shows that the trading volumes tend to

[71] The simulated prices are qualitatively very similar to the results for the unstable case in Youssefmir *et al.* (1998). There, the analysis is based on a much simpler version of the model, solved based on a mean-field approximation.

[72] In simple terms, the leverage effect corresponds to a negative correlation between past returns and future volatility, see e.g. Bouchaud *et al.* (2001).

[73] See e.g. Thurner *et al.* (2010).

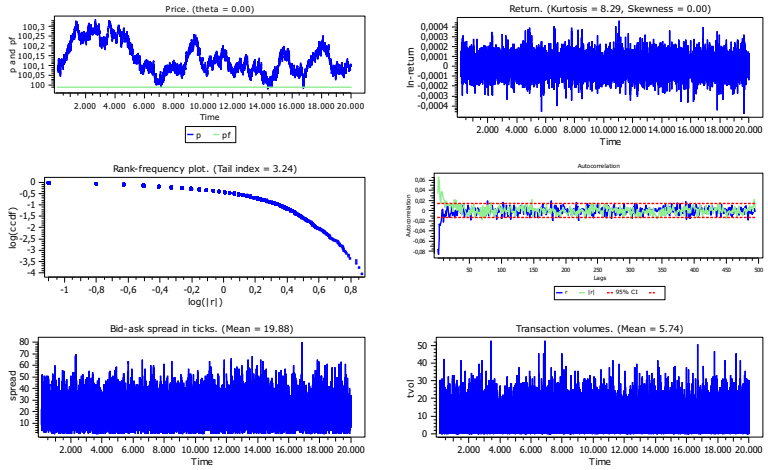

Fig. 2.6: Simulation results: Single run, $\theta = 0$ (GA off). Top left: price (blue) and fundamental value (green). Top right: log-returns. Center left: rank-frequency plot (ccdf, log-log-scale) of log-returns. Center right: Autocorrelation of raw (blue), absolute returns (green), and 95% confidence interval. Bottom left: bid-ask spread (in ticks). Bottom right: transaction volumes.

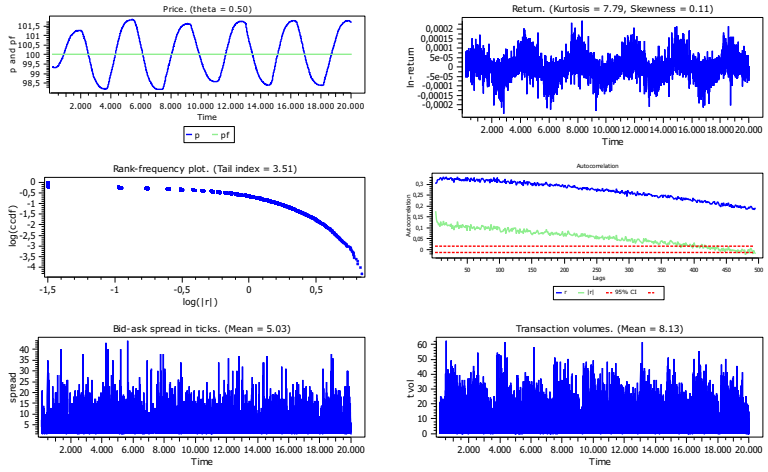

Fig. 2.7: Simulation results: Single run, $\theta = 0.5$ (GA off). Top left: price (blue) and fundamental value (green). Top right: log-returns. Center left: rank-frequency plot (ccdf, log-log-scale) of log-returns. Center right: Autocorrelation of raw (blue), absolute returns (green), and 95% confidence interval. Bottom left: bid-ask spread (in ticks). Bottom right: transaction volumes.

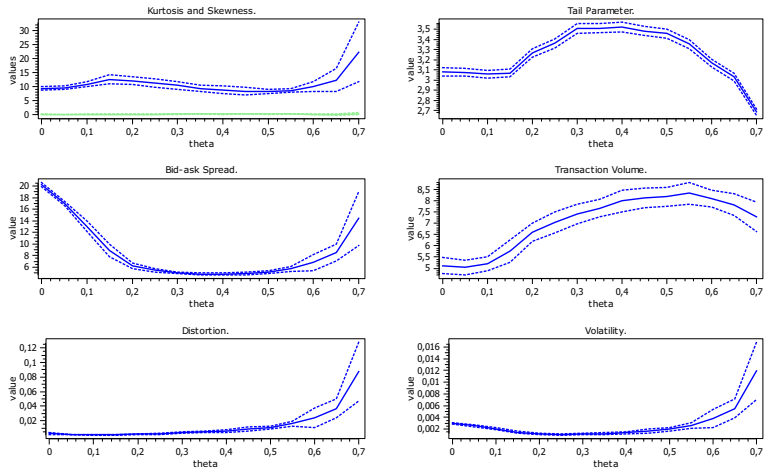

Fig. 2.8: Simulation results: Monte-Carlo simulations. Dependence on θ (GA off). Top left: kurtosis (blue) and skewness (green). Top right: tail exponent (top 15% observations). Center left: bid-ask spread. Center right: transaction volumes. Bottom left: distortion ($|\ln(p/pf)|$). Bottom right: volatility (absolute price change). Plotted are mean values (solid lines), plus and minus one standard deviation (dashed lines).

increase in θ, but become smaller for very large values.

With GA. Figure 2.9 illustrates the results of a single simulation run with the GA activated. In this case, the bottom right panel shows that both the average (blue) and best (green) investment horizon tends to increase slightly.[74] However, the dynamics of the investment horizons are far from being cyclical (which might be a potential source of interesting dynamics). This is true for practically all parameter settings and variants of the GA. In most cases, the average investment horizons remain close to or slightly above the midpoint of the admissible range. We show in section 2.3.2 that this 'convergence' can be partly explained by the bad performance of strongly chartist strategies and the good performance of intermediate and fundamentalist strategies. However, as soon as most agents switch to these strategies, they become less profitable. This self-referential nature of the optimization problem shows that extracting valuable information, in terms of an 'optimal' trading strategy, is a complicated task. In that regard, even focusing

[74] The best investment horizon is the H^w of the fittest strategy at GA activation.

on H^w only and restricting the size of the search space was not sufficient to reduce the complexity of this task. Furthermore, the election operator usually makes sure that the fitness of new strategies is higher than the previous ones. However, in a highly complex system as the present one, it is very hard (impossible?) to calculate the potential fitness of a new strategy. We plan to work on a simpler version of the model in the future, in order to see whether the dynamics of the GA might become more cyclical.

We also performed a number of checks on the working of the GA (e.g. different GA parameters, different fitness functions, and the real GA), but found similar results. Additionally, we checked the robustness of the results using two very simple algorithms, based on selection pressure and imitation only.[75] Again, the results are very similar to those of the baseline GA. Thus, if anything, the GA leads to an increase of the average investment horizons towards intermediate strategies, i.e. strong chartism appears not to pay. In most of the following, we will therefore show the results without the GA, except for the case focusing on the FTT.

Dependence on $\langle H^w/H^t \rangle$

In the baseline scenario, both the informed and noise traders' investment horizons are uniformly distributed, such that the average value of investment over (forward) trend horizon, denoted as $\langle H^w/H^t \rangle$, equals 1. Here we show that the different values indeed matter for the system dynamics. For this purpose, we draw the informed traders' investment horizons from Normal distributions with different mean values, while keeping H^t fixed as the midpoint of the admissible range for all informed traders.[76] Obviously, a smaller average investment horizon for the informed traders should lead to more activity and thus to higher volatility. Additionally, given that we fix H^t, smaller (larger) average investment horizons corresponds to a situation with more chartism (fundamentalism). This is also the basic justification of letting only H^w fluctuate in the GA, as it allows us to easily identify the different states of the system.

Figure 2.10 shows a simulation run for $\langle H^w/H^t \rangle = 1$. As expected, the results are very similar to those of the baseline scenario presented in Figure 2.4. The major difference is the significantly smaller level of autocorrelation

[75] These two algorithms work as follows: (1) Selection pressure. The strategies of the least successful individuals are being replaced by random strategies. (2) Imitation. Two agents 'meet' randomly and less successful agent adopts the more profitable strategy.

[76] Note that we do not change the standard deviation, but only shift the mean of the distribution. In order to identify the effects of these changes, we leave the noise traders' investment horizons uniformly distributed.

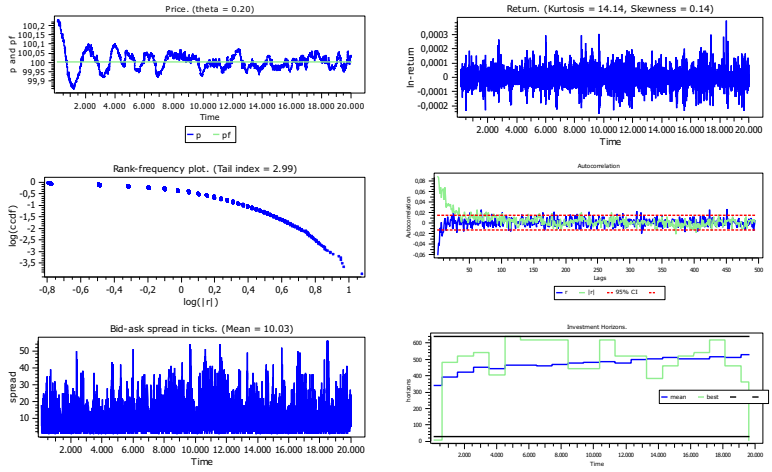

Fig. 2.9: Simulation results: Single run, baseline scenario (GA on). Top left: price (blue) and fundamental value (green). Top right: log-returns. Center left: rank-frequency plot (ccdf, log-log-scale) of log-returns. Center right: Autocorrelation of raw (blue), absolute returns (green), and 95% confidence interval. Bottom left: bid-ask spread (in ticks). Bottom right: average (blue), best (green), and maximum/minimum allowed (black) investment horizons.

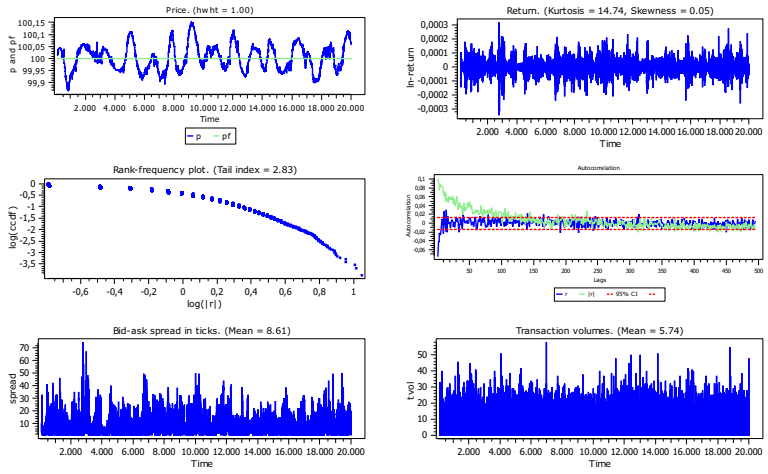

Fig. 2.10: Simulation results: Single run, baseline scenario (GA off). $\langle H^w/H^t \rangle = 1$. Top left: price (blue) and fundamental value (green). Top right: log-returns. Center left: rank-frequency plot (ccdf, log-log-scale) of log-returns. Center right: Autocorrelation of raw (blue), absolute returns (green), and 95% confidence interval. Bottom left: bid-ask spread (in ticks). Bottom right: transaction volumes.

in raw returns, making the resulting price time-series appear more realistic. In contrast, Figure 2.11 show a simulation run for $\langle H^w/H^t \rangle = 0.2$. This corresponds to a situation with pronounced chartism. Quite interestingly, in this case the system's properties are very similar to the case of large θ in the previous section. The major difference here is very small level of volatility clustering. For the sake of completeness, Figure 2.12 shows a simulation run for $\langle H^w/H^t \rangle = 1.8$. This case corresponds to a situation with pronounced fundamental trading. As we can see, there are no clear price bubbles, such that the price tends to track the fundamental value closely. The return distribution is again fat-tailed and there is still some volatility clustering.

Figure 2.13 summarizes the results of the Monte-Carlo simulations for different values of $\langle H^w/H^t \rangle$. Again we see a non-linear relationship for several summary statistics. Quite interestingly, the kurtosis has a peak at the value of 1 and similarly the tail parameter has a minimum around 1.2 (top left and right, respectively). Thus, the tails are fattest around 1, indicating that this value might be a critical value for the system dynamics. Given the smaller level of activity for higher values of $\langle H^w/H^t \rangle$, it is not surprising that the bid-ask spread tends to increase (center left) and trading volumes tend to decrease (bottom right). This is clearly driven by a substantially smaller

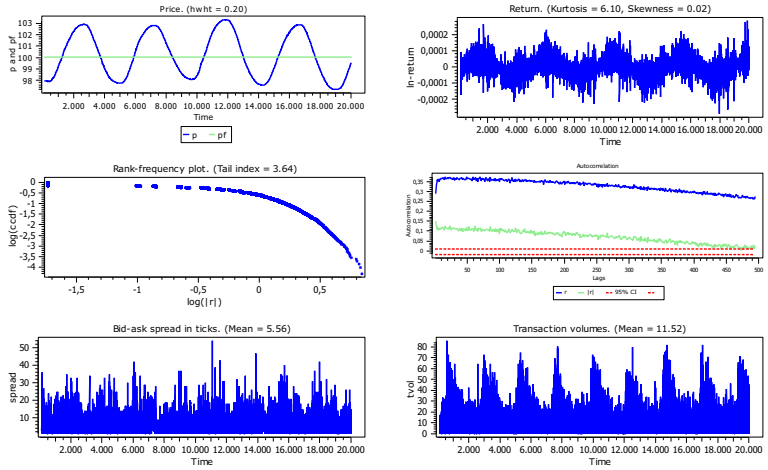

Fig. 2.11: Simulation results: Single run, baseline scenario (GA off). $\langle H^w/H^t \rangle = 0.2$. Top left: price (blue) and fundamental value (green). Top right: log-returns. Center left: rank-frequency plot (ccdf, log-log-scale) of log-returns. Center right: Autocorrelation of raw (blue), absolute returns (green), and 95% confidence interval. Bottom left: bid-ask spread (in ticks). Bottom right: transaction volumes.

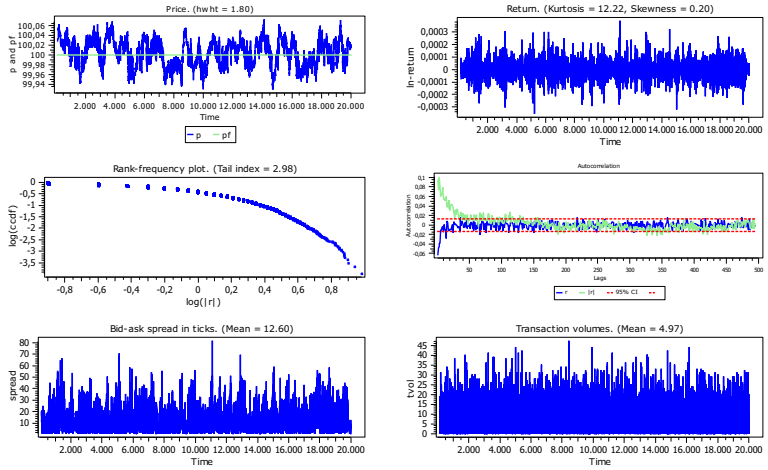

Fig. 2.12: Simulation results: Single run, baseline scenario (GA off). $\langle H^w/H^t \rangle = 1.8$. Top left: price (blue) and fundamental value (green). Top right: log-returns. Center left: rank-frequency plot (ccdf, log-log-scale) of log-returns. Center right: Autocorrelation of raw (blue), absolute returns (green), and 95% confidence interval. Bottom left: bid-ask spread (in ticks). Bottom right: transaction volumes.

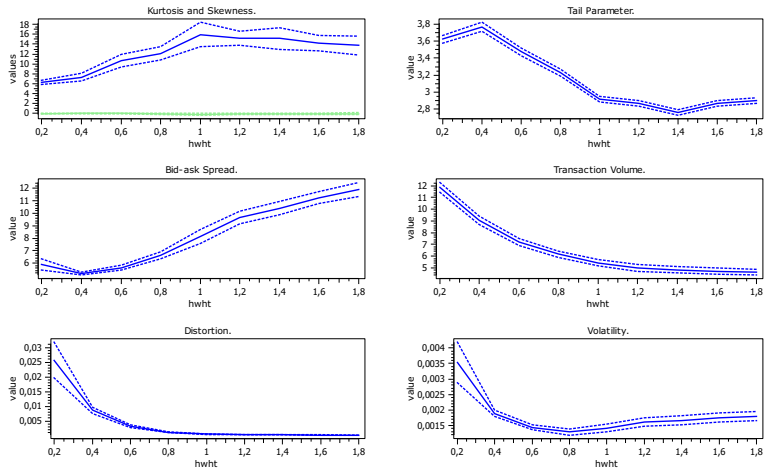

Fig. 2.13: Simulation results: Monte-Carlo simulations. Dependence on $\langle H^w/H^t \rangle$ (GA off). Top left: kurtosis (blue) and skewness (green). Top right: tail exponent (top 15% observations). Center left: bid-ask spread. Center right: transaction volumes. Bottom left: distortion ($|\ln(p/p^f)|$). Bottom right: volatility (absolute price change). Plotted are mean values (solid lines), plus and minus one standard deviation (dashed lines).

number of outstanding orders of informed traders (unreported). Concerning the system's stability, we see that the distortion is clearly negatively related to $\langle H^w/H^t \rangle$, i.e. the average distortion is largest in the highly chartist regime (bottom left). Thus, more chartists indeed drive prices away from the fundamental value. Interestingly, this does not necessarily translate itself into smaller volatility, as can be seen from the bottom right panel. We see that the volatility decreases up to a value $\langle H^w/H^t \rangle = 0.8$ and then starts to increase slightly again. This suggests that the chartist regime is the most unstable, followed by the completely fundamental regime.[77] The least volatile regime appears to be one with both chartists and fundamentalists, even though this might lead to fatter tails. This does make sense, given that the tail contains only the most extreme observations, which tend to have a small effect on the average level of volatility.

[77] Note that the higher level of volatility in the fundamental regime, as compared to the chartist/fundamentalist setting, is driven by noise traders' activities. In this case, mostly noise traders are responsible for the dynamics, which can involve substantial price fluctuations. The (fundamentalist) informed traders then simply take care that the price does not diverge too much from the fundamental value, but remain inactive otherwise.

2.3.2 Effects of FTT

Now we turn to the effects of a (two-sided) FTT. We will present aggregate results from Monte-Carlo simulations for 18 different tax rates for both cases, i.e. with and without GA.[78] The parameters used are summarized in Table 2.2. We should note that the case without GA corresponds to fixed behavioral rules, which is similar to what is often done in the literature on FTTs in order-driven artificial financial markets. The novelty of our approach lies in the endogeneity of the agents' investment horizons.

GA off

Figure 2.14 summarizes the most important model properties in dependence of the tax rate. In the following, we present results from the Monte-Carlo simulations for the case without the GA. The top left panel shows the tax revenues in dependence on the tax rate. We find the usual Laffer curve relationship, i.e. the tax revenues tend to increase for small tax rates until they start decreasing again for larger ones. The top right panel shows that the tails become fatter for larger tax rates. This indicates that the variance of the most extreme observations becomes substantially larger. Note that the tail exponent is usually larger than 2, except for very large tax rates where we might end up in the Levy-stable regime. The center left panel shows that the tax leads to a significant increase of the bid-ask spread. Thus, as expected, the liquidity is significantly reduced.[79] The center right panel shows that transaction volumes tend to decrease in the tax rate, which is not surprising given that all trades become less profitable. Most interestingly, the bottom panels show the distortion (left) and volatility (right), respectively. Quite surprisingly, the distortion tends to decrease for very small tax rates, increases later on and reaches the initial level without tax only for relatively large tax rates. In contrast, the volatility tends to increase for very small to intermediate tax rates, but decreases later on.[80] We will comment on the

[78] We used the following tax rates (in percent): 0 (baseline scenario), 0.01, 0.02, 0.03, $\cdots$, 0.18%. We use 0.18% as the maximum value, since the LOB may be empty at times for larger tax rates and we require at least one order to be on each side of the book to be existent at any point in time.

[79] Note the relatively large change for very small tax rates (also present in the transaction volumes and the distortion). This effect is mostly driven by liquidity reductions from very short-term oriented informed traders, whose trades become unprofitable even for these very small tax rates. For larger tax rates, the amount of liquidity provided both by noise traders and informed traders tend to increase the distortion again.

[80] For larger tax rates, the volatility increases substantially. We do not show the results, since the LOB might become very sparse, with only few or no orders present at certain points in time.

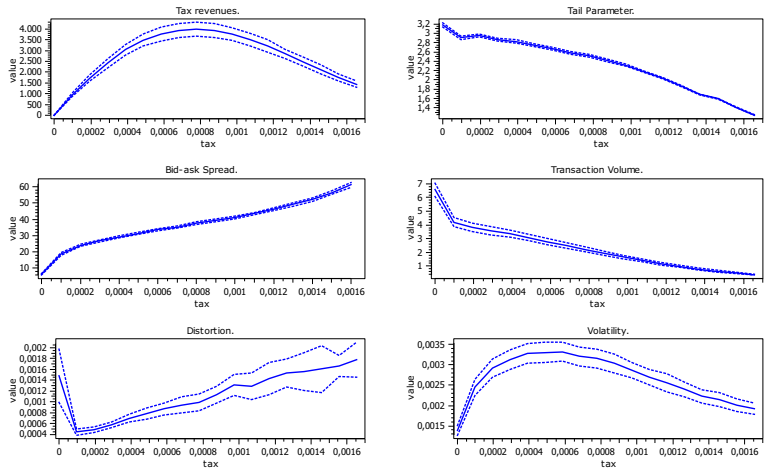

Fig. 2.14: Simulation results: Monte-Carlo simulations. Dependence on χ (GA off). Top left: tax revenues. Top right: tail exponent (top 15% observations). Center left: bid-ask spread. Center right: transaction volumes. Bottom left: distortion ($|\ln(p/pf)|$). Bottom right: volatility (absolute price change). Plotted are mean values (solid lines), plus and minus one standard deviation (dashed lines).

'kinks', i.e. the strong changes in some of the variables for very small tax rates in section 2.3.2. The findings for the distortion and volatility can be explained by the wedge that the tax drives between agents' two limit prices. Only those limit prices corresponding to a positive expected return (post-tax) are submitted to the LOB. However, these orders are unlikely to be executed. For example, in case of an initial market order, the limit price of the second order needs to be substantially far away from the current price. However, this second order will probably never be executed. In this way, price fluctuations are relatively rare (volatility decreases), but if the price changes it does so substantially (tail index smaller). Additionally, the informed traders make sure that prices do not depart too much from the fundamental value. We should also note that the tax rate with the maximum volatility level is close to the tax rate that maximizes tax revenues, representing the usual trade-off between stability and tax revenues.

Summing up, the results for small tax rates are roughly in line with those from the literature, except for the strong decrease in distortion for very small tax rates.[81] This suggests that a small tax rate should reduce liquidity and

[81] See Mannaro *et al.* (2008) and Pellizzari and Westerhoff (2009) for the CDA case.

transaction volumes, but might actually bring prices closer to the fundamental values and only marginally increasing volatility. Additionally, our findings also show that larger tax rates may not create entirely negative effects. In this case, comparable values for distortion and volatility as in the no-tax case come along with substantial tax revenues.

GA on

Figure 2.15 shows the results for the case with the GA activated. The findings are very similar to the case without the GA: the tax revenues shows the Laffer curve relationship, the tail parameter decreases in the tax rate, and the bid-ask spread increases in the tax rate. Similarly, the distortion drops for small tax rates and decreases for larger ones, whereas the volatility increases for small tax rates and decreases later on. In the center right panel we also plot the average (blue) and best (green) investment horizons for the different tax rates. Clearly, there is no convergence in the time horizons, and the average values are very close to (or slightly above) the initial mean. We also see that the standard deviation does not change for the different tax rates. In the end, the conclusions from the case without GA remain largely unaffected.

Investment Horizons, Performance, and FTT

In the previous section, we analyzed the effects of the FTT in terms of aggregate system properties both with and without GA. Here, given that we expect the tax to affect the behavior of certain groups of traders in different ways, we take a closer look at different subgroups of the population of the informed traders. In everything that follows, we use the parameters from the baseline scenario (GA off) and a set of rather small tax rates compared to the previous sections.[82] The main reason for using these smaller rates is the relatively high sensitivity of informed traders' trading decisions, in particular for very small investment horizons, to minor changes in the tax rate. Here we mainly aim at explaining the 'kinks', for example in distortion, in Figure 2.14 for very small tax rates, but focusing on the behavior of certain groups of informed traders. This analysis also allows to make further statements about the dynamics of the GA. For this purpose, we do not randomly assign investment horizons to the informed traders within the admissible range, but rather divide them into three different groups: All agents in group 1 use a small investment horizon, those in group 2 an intermediate value, and those in group 3 a large value. To be precise, for the three groups H^w corresponds

[82] The maximum tax rate used here is only 0.0035%.

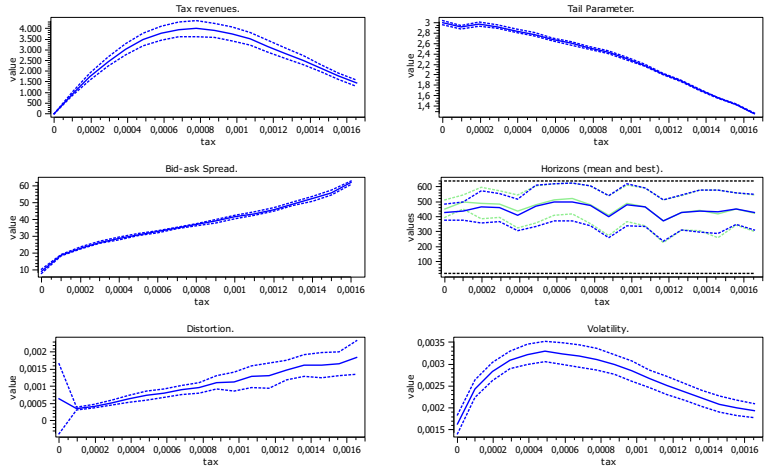

Fig. 2.15: Simulation results: Monte-Carlo simulations. Dependence on χ (GA on). Top left: tax revenues. Top right: tail exponent (top 15% observations). Center left: bid-ask spread. Center right: investment horizons (mean and best). Bottom left: distortion ($|\ln(p/pf)|$). Bottom right: volatility (absolute price change). Plotted are mean values (solid lines), plus and minus one standard deviation (dashed lines).

to the minimum, the midpoint, and the maximum investment horizons in the admissible range, respectively. Thus, group 1 uses a chartist strategy, group 3 a fundamentalist strategy, and group 2 something in between.

We see that informed traders tend to gain, while noise traders tend to lose on average (cf. Figure 2.5). In the following, by analyzing the wealth dynamics of the three groups (relative to the performance of noise traders), we will show that the distribution of gains and losses of informed traders is far from uniform. This analysis also helps us in understanding the dynamics of the investment horizons with the GA activated. Most importantly, we find that the increase in investment horizons shown before is mainly due to the very bad performance of group 1, i.e. the chartists. This can be seen from Figures 2.16-2.18, where we plot the average wealth difference (in absolute terms) between the three groups and the group of noise traders over time for single simulation runs and different tax rates.

Figure 2.16 shows that, without a tax group 2 (green line) performs best, i.e. the intermediate strategy is the most lucrative, as it tends to strongly gain in wealth over time. Similarly, group 3 tends to gain as well, but on a smaller scale, so the fundamental strategy is also profitable. Somewhat surprisingly, the worst performers are in group 1, whose average wealth tends to be comparable to the wealth of noise traders (at times even smaller as the wealth difference may be negative and highly cyclical). Thus, chartists tend to perform very poorly on average. An explanation for the good performance of group 2 is that group 1 is the major source of predictable bubbles (high-frequency traders), while group 3 tries to drive prices back towards the fundamental level (low-frequency traders). The intermediate strategy works on a higher frequency than the fundamental traders, thereby leading them to follow the trend at times or expecting reversal towards the fundamental value at other times.

From this viewpoint the dynamics of the average investment horizons with the GA, as described in section 2.3.1, are sensible; the GA spreads the best performing strategies, replacing less successful ones. Thus, the average investment horizon tends to move quickly towards intermediate investment horizons. However, we should stress that we only show the average wealth dynamics for the three groups, i.e. neglecting within-group variations. These variations, however, may crucially affect the working of the GA.[83] We should also stress that, in order to avoid strong path dependence, our fitness function (cf. Equation 5.12) may at times suggest that chartists performed rather well;

[83] For example, we checked the dynamics of the average investment horizons for varying (average) initial values of $\langle H^w/H^t \rangle$ (cf. section 2.3.1). It turns out that the average horizons with the GA typically end up close to the initial average. This may to a large extent be driven by a significant amount of noise in the profitability of individual strategies.

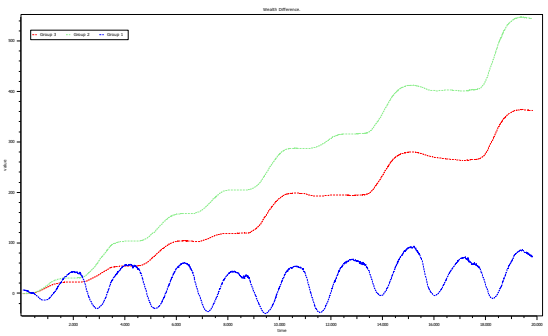

Fig. 2.16: Single run, baseline scenario (GA off) with three groups of informed traders with small, intermediate, and large H^w. Average wealth differences between the three groups and noise traders. Group 1 (3) are the chartists (fundamentalists), while group 2 uses an intermediate strategy.

even though they might lose money on average, the strong cyclicality of their fitness values can lead to very large fitness values of individuals. Clearly, the GA has difficulties extracting this information.

Figure 2.17 shows the wealth dynamics for a very small tax rate. We see that the ordering is conserved, since group 2 still performs better on average than the other groups. As before, chartists tend to perform very badly, in this example even losing money over time, while fundamentalists and those with an intermediate strategy tend to gain. Thus, for very small tax rates, chartists continue posting very unprofitable orders. This changes for very large tax rates, as we can see from Figure 2.18: here the average wealth of chartists now exceeds that of the noise traders. Large tax rates reduce the amount of strong chartist orders, so chartists post fewer orders (i.e. effectively work on lower frequencies) and perform better. Also we see that their wealth dynamics are significantly less cyclical, so reducing their posting of high-frequency orders tends to reduce the occurrence of bubbles and bursts. Interestingly, we see that for the fundamental strategy becomes most profitable for larger tax rates. Note that in this case there are practically no cyclical fluctuations, so the calculated trends (necessary for calculating the expected prices) tend to be very small. This implies that the trending component for all strategies becomes negligible, so practically all traders tend to expect the price to revert towards the fundamental value. From this viewpoint, it is clear that fundamental traders tend to be the most successful, since their orders have the longest lifetime, i.e. the highest probability that their two limit prices are being hit by the noise traders. The longer

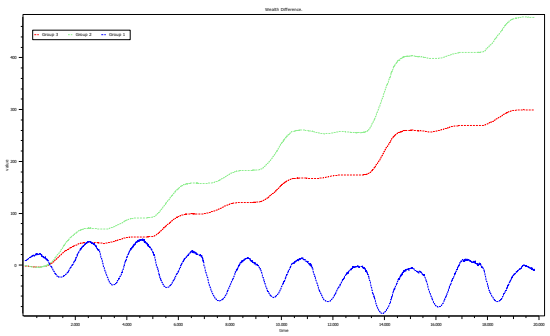

Fig. 2.17: Single run, baseline scenario (GA off) with small tax rate and three groups of informed traders with small, intermediate, and large H^w. Average wealth differences between the three groups and noise traders. Group 1 (3) are the chartists (fundamentalists), while group 2 uses an intermediate strategy.

investment horizon also implies that the difference between the current price and the expected price can be larger, such that the majority of orders with expected profits (after tax) are posted by the group of fundamental traders.

Figure 2.19 illustrates the effects of different tax rates on the order submission process of the three groups. There we show the average fraction of order submissions relative to the frequency of being active for all groups. First, we see that noise traders (black) are practically unaffected by these relatively small tax rates, as they post orders practically every time when they become active (i.e. the fraction of order submissions is close to 1). In contrast, we see that group 1 posts orders roughly 60% of the time without a tax, but this fraction approaches zero relatively quickly. We see a stepwise relationship, which is due to the fact that prices (and thus the taxes to be paid) are not continuous, but depend on the tick size. For groups 2 and 3 we observe similar relationships, but the values are on a significantly higher level for these relatively small tax rates. Quite interestingly, at the highest tax rate in this analysis, group 3 continues posting orders 80% of the time. The kinks in the previous section are therefore mainly driven by the inactivity of chartists already for relatively small tax rates (cf. Figure 2.20). As expected, the tax affects strategies with higher trading frequency significantly more. Note that for such small tax rates, the other groups continue posting orders, so the absence of strong speculative bubbles reduces the average distortion. The fraction of order submissions for group 2 and 3 also tend to approach zero for larger tax rates, with group 3 being active also

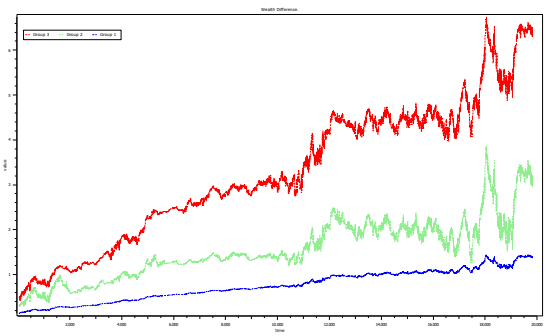

Fig. 2.18: Single run, baseline scenario (GA off) with larger tax rate and three groups of informed traders with small, intermediate, and large H^w. Average wealth differences between the three groups and noise traders. Group 1 (3) are the chartists (fundamentalists), while group 2 uses an intermediate strategy.

for larger tax rates (unreported). For very large tax rates, i.e. those used in the previous sections, even the fundamental traders act rarely, such that the average distortion starts increasing again.

2.4 Conclusions

In this paper we presented a detailed artificial financial market, where agents compete against each other within a CDA mechanism. While incorporating the usual chartist/fundamentalist/noise trader framework, our model has the advantage of explicitly accounting for the importance of time horizons in financial markets. We showed evidence that the model is able to replicate certain well-known stylized fact of financial time-series, among them the martingale property, fat-tailed return distributions, and, albeit to a lesser extent, volatility clustering. Moreover, for certain parameter combinations the model is able to generate the building-up and bursting of asset price bubbles. We also discussed the results for a number of parameter variations.

The main focus of this paper was on the effects of a FTT in the artificial financial market. In this regard, the two cases, i.e. with fixed and flexible behavioral strategies, yield very similar results, due to typical convergence to the intermediate investment strategy due to the GA. We find the usual trade-off between monetary revenues (Laffer curve) and stability, as higher tax revenues come along with higher volatility. The results for small tax rates are roughly in line with those from the literature, except for the strong

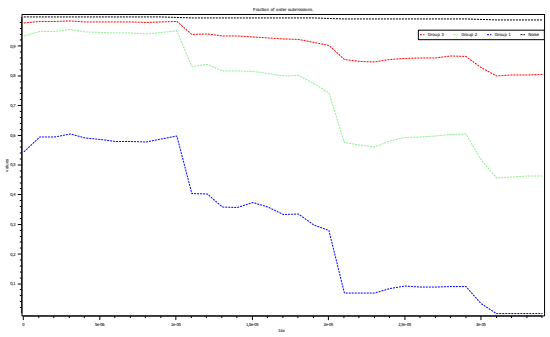

Fig. 2.19: Simulation results: Monte-Carlo simulations. Fraction of order submissions relative to the activity frequency for noise traders (black), and group 1-3.

decrease in distortion for very small tax rates.[84] This suggests that a small tax rate should reduce liquidity and transaction volumes, but might actually bring prices closer to the fundamental values and only marginally increasing volatility. Additionally, we also show that larger tax rates may not create entirely negative effects. In this case, comparable values for distortion and volatility as in the no-tax case come along with substantial tax revenues. These revenues could be used for a number of productive purposes. Additionally, the reduced market activity also frees-up resources, both in terms of financial and human capital, that could be directed to other parts of the economy.

In reality, we would need to weigh the different effects in order to come up with a welfare-optimizing solution. For example, here we simply redistributed the tax revenues among the traders to keep total wealth constant. In reality one could think of transferring the money to those who are in need. Given the extraordinary high transaction volumes in real markets, partly driven by the arrival of HFTs, it appears promising to introduce a small tax to reduce the possibly distorting effects of their activities and generate large tax revenues at the same time. In part, our results suggest that imposing a very small tax would make HFT strategies highly unprofitable. However, since their trading algorithms are usually not meant to follow trends or drive prices towards some fundamental value, their effects on the macro-properties of the system are still under debate. In the end, even a tiny FTT would lead to a shrinking of the financial sector, allowing to extract highly productive resources (e.g. human capital) for other purposes.

[84] See Mannaro *et al.* (2008) and Pellizzari and Westerhoff (2009) for the CDA case.

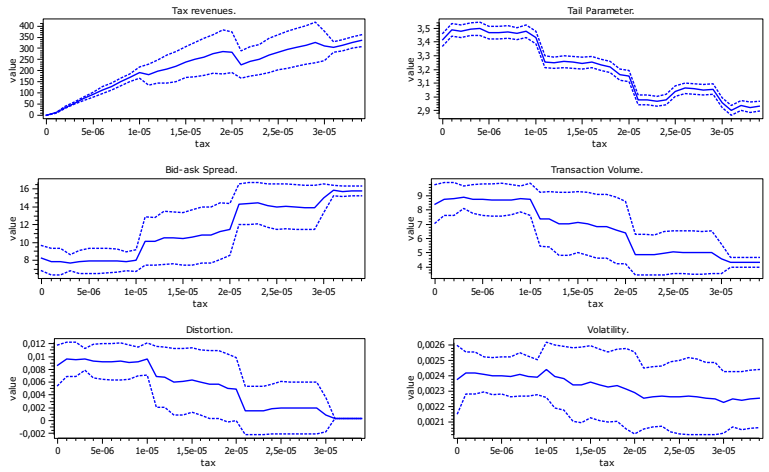

Fig. 2.20: Simulation results: Monte-Carlo simulations. Dependence on χ (GA on) for relatively small tax rates with three groups of informed traders. Top left: tax revenues. Top right: tail exponent (top 15% observations). Center left: bid-ask spread. Center right: transaction volumes. Bottom left: distortion ($|\ln(p/pf)|$). Bottom right: volatility (absolute price change). Plotted are mean values (solid lines), plus and minus one standard deviation (dashed lines).

The presented model is very flexible and serves as a good illustration of the complexity of the optimization task in real markets. The self-referential nature of this task makes the extraction of valuable information quite difficult. In this regard, it is worth noting that the order generation process of individuals is still poorly understood, in contrast to the aggregate order-book dynamics for which a number of scaling-laws have been identified.[85] We hope that future research, for example by means of laboratory experiments, may help us in deciphering these processes. It is crucial to understand the agents' individual behavior at the micro-level to generate more realistic dynamics at the macro-level. In the end, we believe that our model is an ambitious first step towards more realistic 'wind-channels' for testing regulatory policies. In future research, we plan to tackle several of the issues mentioned throughout the text.

[85] See e.g. Zovko and Farmer (2002).

Part II

INTERBANK MARKETS AND NETWORK ANALYSIS

3. NETWORK ANALYSIS OF THE E-MID OVERNIGHT MONEY MARKET: THE INFORMATIONAL VALUE OF DIFFERENT AGGREGATION LEVELS FOR INTRINSIC DYNAMIC PROCESSES

3.1 Introduction and Existing Literature

Interbank markets are crucial for the functioning of the economy. However, as painfully illustrated by the global financial crisis (GFC) in 2007/08, creating links at the micro-level may generate systemic risk at the macro-level. Thus, the structure of the interbank network, with the banks having connections in terms of credit relationships, is important for its stability. The economy depends on stable interbank markets, since short-term money market rates affect those of longer maturities and thus the real economy. From this viewpoint, it appears quite surprising that the economic profession has not been concerned much with the functioning of interbank markets until recently. The usual focus is on the overnight segment of the interbank deposit market, since it tends to be the largest spot segment of money markets. In this paper, we analyze the network properties of the Italian e-MID (electronic market for interbank deposits) data based on overnight loans during the period 1999-2010.

Most existing studies on the structure of real interbank markets have been conducted by physicists trying to get an idea of the topology of different interbank markets. Examples include Boss *et al.* (2004) for the Austrian interbank market, Inaoka *et al.* (2004) for the Japanese BOJ-Net, Soramäki *et al.* (2007) for the US Fedwire network, Bech and Atalay (2010) for the US Federal funds market, and De Masi *et al.* (2006) and Iori *et al.* (2008) for the Italian e-MID (electronic market for interbank deposits). The most important findings reported in this literature are: (1) most interbank networks are quite large (e.g. more than 5000 banks in the Fedwire network), (2) interbank networks are sparse, meaning that only a minority of all possible links do actually exist, (3) degree distributions appear to be scale-free (with coefficients between 2-3), (4) transaction volumes appear to follow scale-free distributions as well, (5) clustering coefficients are usually quite small, (6) interbank networks are small worlds and (7) the networks show disassortative mixing with respect to the bank size, so small banks tend to trade with large banks and vice versa.

Most relevant for our study are the two papers on the e-MID. We should stress here that the e-MID data are the only interbank data which can be purchased freely without any restrictions. In contrast, getting access to similar datasets for other markets is usually far more complicated. The authors analyze daily networks from 1999-2002 and find intradaily and intramonthly seasonalities. The authors conclude that the networks appear to be random at the daily level. This finding is in stark contrast with the findings of preferential lending relationships in the Portuguese interbank market by Cocco *et al.* (2009). In this paper, we are mostly concerned with matching these

seemingly incompatible findings, by showing that the aggregation period has an effect on the informational value of the underlying networks. The main finding is that daily networks indeed feature a substantial amount of randomness and cannot be considered as being representative for the underlying 'latent' network. This is illustrated on the basis of a number of network statistics which are compared to those of random networks. Furthermore, we find a substantial amount of asymmetry in the network. Last but not least, we find that the GFC can be identified as a significant structural break for many network measures.[1]

The remainder of this paper is structured as follows: Section 3.2 gives a brief introduction into (interbank) networks, section 3.3 introduces the Italian e-MID trading system and gives an overview of the data set we have access to. Section 3.4 describes our findings and section 3.5 concludes and discusses the relevance of these findings for future research.

3.2 Networks

A network consists of a set of N nodes that are connected by M edges (links). Taking each bank as a node and the interbank positions between them as links, the interbank network can be represented as a square matrix of dimension $N \times N$ (data matrix, denoted $\mathbf{D}$).[2] An element d_{ij} of this matrix represents a gross interbank claim, the total value of credit extended by bank i to bank j within a certain period. The size of d_{ij} can thus be seen as a measure of link intensity. Row (column) i shows bank i's interbank claims (liabilities) towards all other banks. The diagonal elements d_{ii} are zero, since a bank will not trade with itself.[3] Off-diagonal elements are positive in the presence of a link and zero otherwise.

Interbank data usually give rise to directed, sparse and valued networks.[4] However, much of the extant network research ignores the last aspect by focusing on binary adjacency matrices only. An adjacency matrix $\mathbf{A}$ contains elements a_{ij} equal to 1, if there is a directed link from bank i to j and 0

[1] In a companion paper, we focus explicitly on fitting the degree distribution, see Fricke *et al.* (2013). The main findings are (1) The degree distributions are unlikely to be scale-free, and (2) the in- and out-degrees do not follow the same distribution.

[2] In the following, matrices will be written in bold, capital letters. Vectors and scalars will be written as lower-case letters.

[3] This is true when we think of individual banks as consolidated entities.

[4] Directed means that $d_{i,j} \neq d_{j,i}$ in general. Sparse means that at any point in time the number of links is only a small fraction of the $N(N-1)$ possible links. Valued means that interbank claims are reported in monetary values as opposed to 1 or 0 in the presence or absence of a claim, respectively.

otherwise. Since the network is directed, both $\mathbf{A}$ and $\mathbf{D}$ are asymmetric in general. In this paper, we also take into account valued information by using both the raw data matrix as well as a matrix containing the number of trades between banks, denoted as $\mathbf{T}$. In some cases it is also useful to work with the undirected version of the adjacency matrices, $\mathbf{A}^u$, where $a^u_{ij} = \max(a_{ij}, a_{ji})$.

As usual, some data aggregation is necessary to represent the system as a network. In the following, we use quarterly networks.

3.3 The Italian Interbank Market e-MID

The Italian electronic market for interbank deposits (e-MID) is a screen-based platform for trading of unsecured money-market deposits in Euros, US-Dollars, Pound Sterling, and Zloty operating in Milan through e-MID SpA.[5] The market is fully centralized and very liquid; in 2006 e-MID accounted for 17% of total turnover in the unsecured money market in the Euro area. Average daily trading volumes were 24.2 bn Euro in 2006, 22.4 bn Euro in 2007 and only 14 bn Euro in 2008.

Available maturities range from overnight up to one year. Most of the transactions are overnight. While the fraction was roughly 80% of all trades in 1999, this figure has been continuously increasing over time with a value of more than 90% in 2010.[6] As of August 2011, e-MID had 192 members from EU countries and the US. Members were 29 central banks acting as market observers, 1 ministry of finance, 101 domestic banks and 61 international banks. We will see below that the composition of the active market participants has been changing substantially over time. Trades are bilateral and are executed within the limits of the credit lines agreed upon directly between participants. Contracts are automatically settled through the TARGET2 system.

The trading mechanism follows a quote-driven market and is similar to a limit-order-book in a stock market, but without consolidation. The market is transparent in the sense that the quoting banks' IDs are visible to all other banks. Quotes contain the market side (buy or sell money), the volume, the interest rate and the maturity. Trades are registered when a bank (aggressor) actively chooses a quoted order. The platform allows for credit line checking before a transaction will be carried out, so trades have to be confirmed by

[5] The vast majority of trades (roughly 95%) is conducted in Euro.

[6] This development is driven by the fact that the market is unsecured. The recent financial crisis made unsecured loans in general less attractive, with stronger impact for longer maturities. See below. It should be noted, that there is also a market for secured loans called e-MIDER.

both counterparties. The market also allows direct bilateral trades between counterparties.

The minimum quote size is 1.5 million Euros, whereas the minimum trade size is only 50,000 Euros. Thus, aggressors do not have to trade the entire amount quoted.[7] Additional participant requirements, for example a certain amount of total assets, may pose an upward bias on the size of the participating banks. In any case, e-MID covers essentially the entire domestic overnight deposit market in Italy.[8]

We have access to all registered trades in Euro in the period from January 1999 to December 2010. For each trade we know the two banks' ID numbers (not the names), their relative position (aggressor and quoter), the maturity and the transaction type (buy or sell). As mentioned above, the majority of trades is conducted overnight and due to the global financial crisis (GFC) markets for longer maturities essentially dried up. We will focus on all overnight trades conducted on the platform, leaving a total number of 1,317,679 trades. The large sample size of 12 years allows us to analyze the network evolution over time. Here we focus on the quarterly aggregates, leaving us with 48 snapshots of the network.

3.4 Results

In this section, we look at the network structures formed by interbank lending over various horizons of time aggregation of the underlying data. We will see that comparing various network measures at different levels of time aggregation reveals interesting features suggestive of underlying behavioral regularities. Given that most studies focus on overnight data, it has become quite standard to focus on networks constructed from daily data. Here we find, that, at least for the Italian interbank network, it may be more sensible to focus on longer aggregation periods, namely monthly or quarterly data. We also discuss in how far the network structure has changed (and, in how far it has remained intact) after the default of Lehman Brothers in September 2008.

3.4.1 General Features

In total, 350 banks (255 Italian and 95 foreign) were active at least once during the sample period. However the number of active banks changes

[7] The minimum quote size could pose an upward bias for participating banks. It would be interesting to check who are the quoting banks and who are the aggressors. Furthermore it would be interesting to look at quote data, as we only have access to actual trades.

[8] More details can be found on the e-MID website, see http://www.e-mid.it/.

substantially over time as can be seen from the left panel of Figure 6.1.[9] We see a clear downward trend in the number of active Italian banks over time, whereas the additional large drop after the onset of the GFC is mainly due to the exit of foreign banks. The right panel shows that the decline of the number of active Italian banks went along with a relatively constant trading volume in this segment until 2008. This suggests that the decline of active Italian banks was mainly due to mergers and acquisitions within the Italian banking sector. Given the anonymity of the data set, it is impossible to shed more light on this interesting issue. The overall upward trend of trading volumes was due to the increase of active foreign banks until 2008, while their activities in this market virtually faded away after the onset of the crisis. Interestingly, the average volume per trade tends to increase over time, as can be seen from the strong negative trend in the total number of trades in Figure 3.2, at least for the Italian banks.

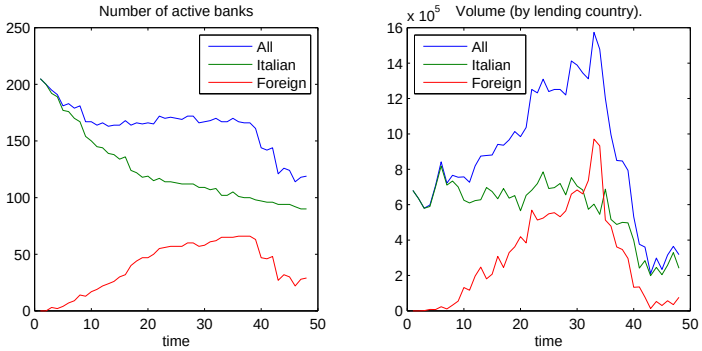

Fig. 3.1: Number of active banks (left) and traded volume (right) over time. We also split the traded volume into money lent by Italian and foreign banks, respectively.

An interesting question in this regard is, who trades with whom. Figure 3.3 illustrates this for the number of trades (top) and the transacted volume (bottom) by country. For example, the green lines show the total number of trades (traded volumes) of foreign banks lending money to Italian banks, relative to all outgoing trades of foreign banks. Similarly, the blue lines show the total number of trades (traded volumes) of money flowing between Italian banks, as a fraction of all outgoing trades of Italian banks. The general patterns are the same for both Figures: Italian banks lend most of the time to other Italian banks (99.31% on average) and only a negligible amount

[9] Similar developments are reported by Bech and Atalay (2010) for the federal funds market.

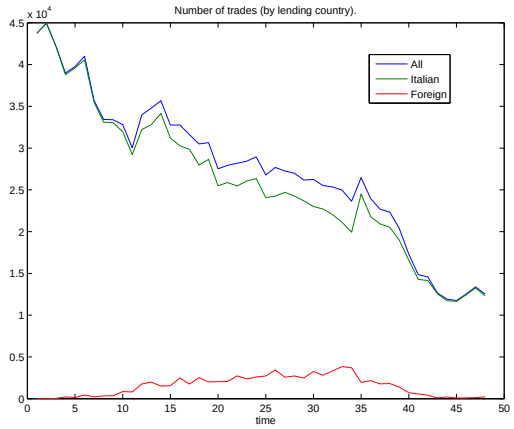

Fig. 3.2: Number of trades per quarter.

to foreign banks (0.61% on average). This pattern is remarkably stable over time. In contrast, at the beginning of the sample period, foreign banks mostly used the market in the absence of (many) other foreign counterparties to lend money to Italian banks. This has changed over time and foreign banks mostly later on used the platform to trade with other foreign banks. It is not quite clear why this is the case, the underlying trend seems to point towards structural changes altering the (foreign) banks' behavior. For many research questions, one should therefore only use the subsample of Italian banks. In most of what follows, we stick to this choice.

This leads us to a first glance at the network structure. Figure 3.4 shows the banking network formed by the 119 active banks (89 Italian) in the last quarter of 2010.[10] The network consists mainly of two components: The very dense part formed by the Italian banks (circles) on the right-hand side and the far less interconnected foreign banks (triangles) on the left-hand side. The higher activity of the Italian banks is not represented in terms of the volume traded. We use total outgoing volume as a proxy for banks size and group the banks into 4 classes according to which percentile (30th, 60th, 90th or above) they belong to. This attribute is shown in the Figure as the size and the brightness of the nodes. We should note that 3 out of 12 banks of group 4 are foreign banks which is in line with their fraction of the total banks (30 out of 119). Hence, foreign banks trade less on average (both in terms of volume and number of trades), however, the volume per trade is higher.

[10] The Figure was produced using visone, http://www.visone.info/, by Brandes and Wagner (2004).

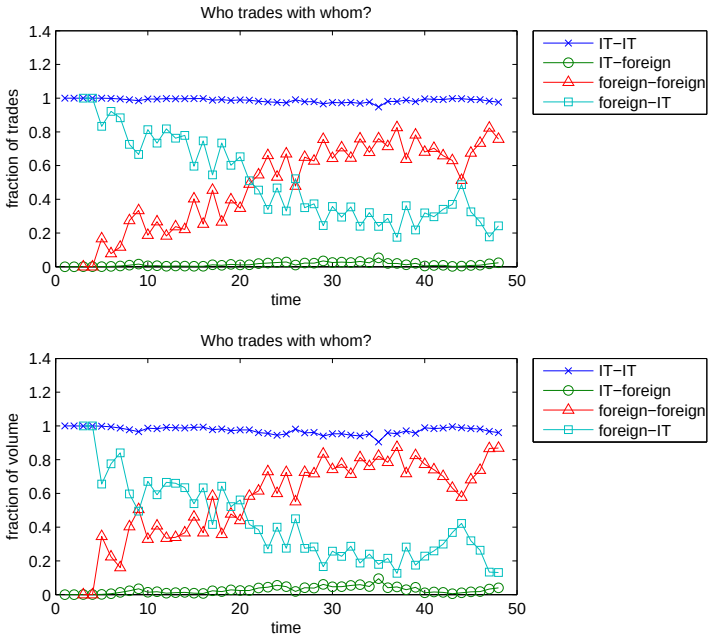

Fig. 3.3: Fraction of trades (top) and traded volume (bottom) between banks from different countries.

It is also interesting to highlight some specific features of the trading behavior of individual banks in this particular quarter, since not all banks use the market in the same way:[11] There are 14 banks with zero in-degree and 29 banks with zero out-degree. Surprisingly these banks are quite heterogeneous and not, as one might expect, just small banks. As an example the highest overall transaction volume of 58.6 bn Euro for a single bank, and therefore roughly 9.3% of the total trading volume, was traded by a German bank borrowing this sum in 90 trades from 8 counterparties. Another interesting case is an Italian bank trading only with one counterparty, lending this other bank 5.02 billion Euro in 76 trades, whereas borrowing just 0.03 billion in 3 trades. Even though these special relationships are quite interesting, the anonymity of the data set makes it impossible for us to say more on the particular relationships that might lead to these interesting outcomes. After this broad overview of the market and the ongoing interactions, we turn to the question of a sensible aggregation period. As should be clear from the discussion above, we will mostly focus on the (sub)network formed by Italian banks only.

[11] For a detailed analysis of the trading strategies in the e-MID, see Fricke (2012).

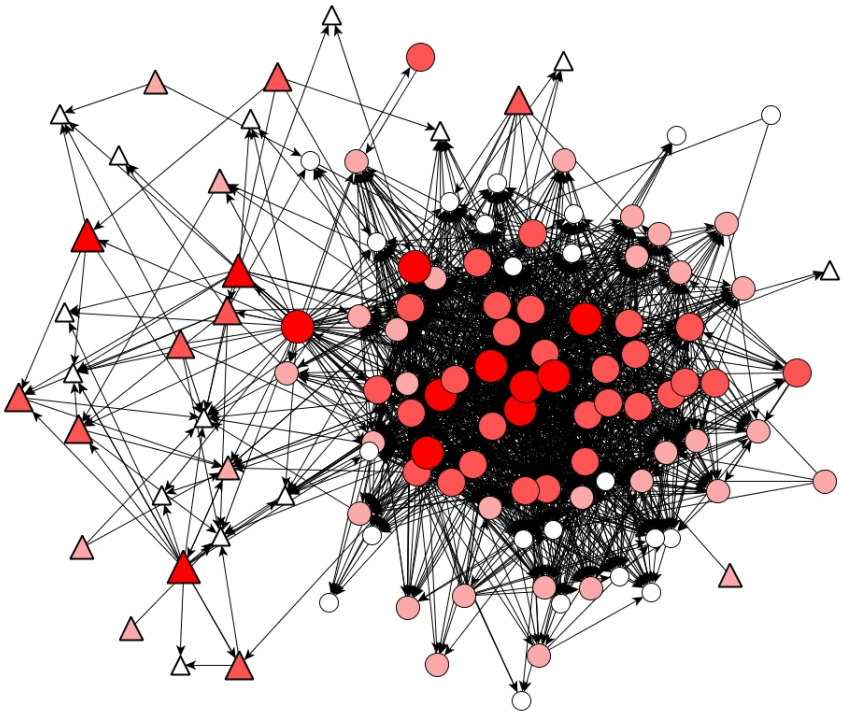

Fig. 3.4: The banking network in the 4th of quarter 2010: triangles are foreign banks. The size of the node as well as the brightness of the red color indicate the size in terms of volume lent.

3.4.2 Density

The density ρ of a network is defined as the number of existing links (M) relative to the maximum possible number of links. It can be calculated as

$$\rho = \frac{M}{N(N-1)}. \tag{3.1}$$

Figure 5.2 illustrates the evolution of the density for four different aggregation periods (day, month, quarter, year). Except for the daily networks the density is quite stable over time and slightly increases until the GFC, which was a significant structural break for the monthly and quarterly network. We should note that the breakpoint (quarter 39), coincides with the quarter during which Lehman Brothers collapsed. The daily density fluctuates much more strongly, but overall increases throughout the sample.[12]

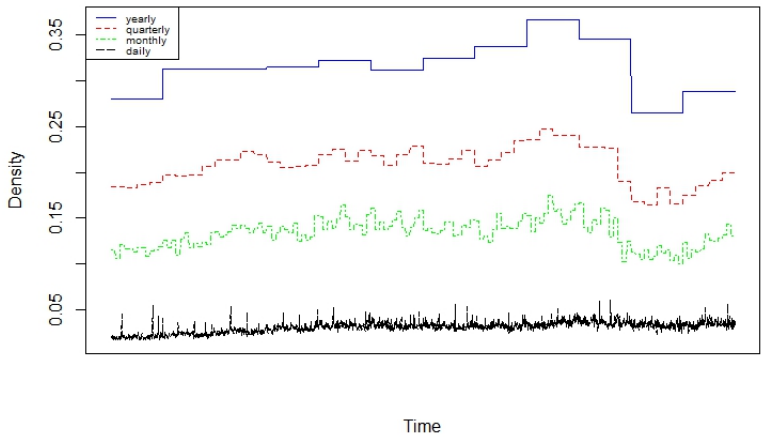

Fig. 3.5: The density for yearly (blue), quarterly (red), monthly (green) and daily (black) aggregated networks. A Chow-test and an additional CUSUM-test indicate a structural break for quarter 39 (month 117) at the 1% significance level, but not for the yearly or daily networks.

Compared to the findings for other interbank networks, the density of the Italian interbank network is quite high. For example, Bech and Atalay (2010)

[12] The density for the total network, including the foreign banks, seems to steadily decline over the sample period. This illustrates the fact that the increasing fraction of foreign banks are less interconnected with the (smaller) Italian banks.

and Soramäki *et al.* (2007) report an average density of below 1% in daily interbank networks, compared to an average density of roughly 3.1% in our case. The main reasons for the higher density are most likely the relatively small number of participating banks in the market and the transparent market structure which easily allows each bank to trade with any other bank in the market. For comparison, the Fedwire network investigated by Soramäki *et al.* (2007) contains 5,086 institutions.[13] The means of 20.8% for quarterly aggregated networks (13.4% monthly) reveal much higher figures. Obviously the network density is positively related to the aggregation period, but to our knowledge the structure of this relation has not been investigated for interbank networks so far.

For this reason, we compare the aggregation properties of the empirical networks with those of random networks. Here we use Erdös-Renyi networks, i.e. completely random networks, and random scale-free networks, where the out-degrees follow a power-law distribution with scaling parameter 2.3.[14] The experiments work as follows: For each year, we aggregate the daily networks and plot the resulting density in dependence of the aggregation period, from one day up to one year (roughly 250 days). For the random networks, we aggregate artificial Erdös-Renyi and scale-free networks for each day with the same number of active banks and density as the observed daily network. The results are the average values for 100 runs for the Erdös-Renyi and the scale-free networks.[15] We find very similar qualitative results for all 12 years.

As an example, Figure 3.6 illustrates the results for 1999. For all three networks, there appears to be a saturation level for the density, however at different levels. The Erdös-Renyi networks always show the highest density

[13] Additionally, the electronic nature of the trading platform might make links between any two institutions more likely.

[14] The power-law distribution with tail exponent 2.3 is a common finding in many interbank markets, see e.g. Boss *et al.* (2004). The resulting sequences of the out-degrees are attributed to the nodes by ranking those according to the observed out-degrees, considering only active banks during the particular day. Note that if we did not account for the ordering of the observed degree sequences, we would end up with very similar aggregation properties as the Erdös-Renyi case. The in-degrees are distributed in a random uniform way, ruling out self-links and counting each link at most once. For a detailed analysis of the degree distributions for this data set see Fricke *et al.* (2013).

[15] Note that the density of aggregated Erdös-Renyi networks can be written as

$$\rho_T^r = 1 - \prod_{t=1}^{T} (1 - \rho_t^r),$$

where $(1-\rho_t^r)$ is the probability of observing no links in the network at time t whatsoever, but since we adjust the number of active banks on a daily basis using tese probabilities would not constitute a completely satisfactory approach in our case.

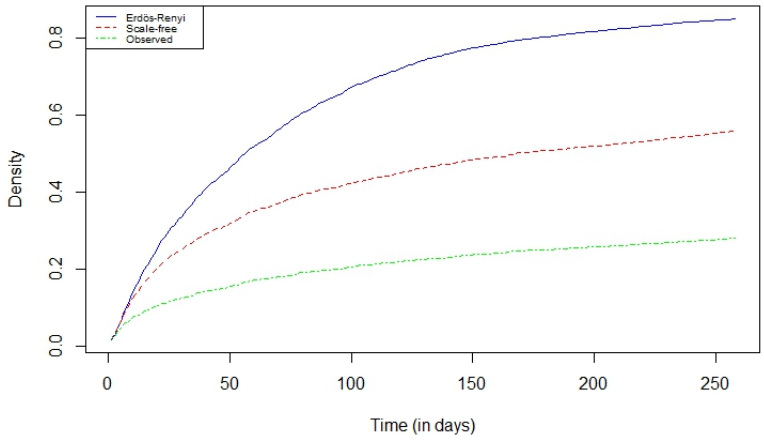

Fig. 3.6: Data for 1999. Density for the aggregated Erdös-Reny (blue), Scale-Free (red) with $\alpha = 2.3$, and observed networks (green). Aggregation period in days. Note: we do not plot standard deviations, since these are negligible.

(up to .851), followed by the Scale-Free networks (up to .559) and the observed networks (only up to .280). Apparently, it is much more likely for the empirical data that the same link gets activated several times than for the randomized data of Erdös-Renyi and scale-free networks (where the overall number of links is the same by construction). This is supported by the fact that in the observed networks a total number of $2,757$ links are observed only once in the year, while for the scale-free and random networks these values are $5,040$ and $5,746$, respectively. Hence, these results indicate the existence of lasting (preferential) lending relationships in the actual banking network.

3.4.3 What is a Sensible Aggregation Period?

After showing that longer than daily aggregation tends to reveal non-random structures for the Italian banking network, we are concerned with determining the 'correct' aggregation period in more detail in this section. This question is crucial for extracting relevant information, since the banking network cannot be observed at a given point in time, but always has to be approximated by aggregating trades over a certain period.[16] Most studies use

[16] The literature on interbank networks is surprisingly silent about the choice of the aggregation period. We are aware of only one paper (Kyriakopoulos *et al.* (2009)) inves-

daily aggregates (daily networks), which seems justified by the fact that the underlying loans are (mostly) overnight. Economically, however, overnight loans can be seen as longer-term loans, where the lender can decide every day whether to prolong the loan or not. Aggregating over a longer period is only preferable, if it can reveal a non-random structure of the banking network. The existence of preferential relationships would imply that daily transactions are not determined myopically, but that a virtual network of longer lasting relationships exists. Daily transactions would then be akin to random draws from this underlying network with the realizations depending on current liquidity needs and liquidity overhang. Aggregation over a sufficiently long time horizon might reveal more and more of the hidden links, rather than adding up purely random draws from all possible links. The relatively fast saturation of the empirical density that we observe in figure 3.6 is consistent with this interpretation. To shed more light on this issue we consider the consistency of yearly, quarterly, monthly and daily aggregated networks. The main finding is that we observe a much higher degree of structural stability for monthly or quarterly networks, depending on the application, rather than daily networks.

The use of a 'sensible' aggregation period should ensure that we extract stable features (if they exist) of the banking network rather than noisy trading patterns at different points in time. In this regard, it is important to investigate the stability of the link structure, in order to assess whether subsequent occurrences of the network share many common links. In order to do this, we rely on the Jaccard Index (JI),[17] which can be used to quantify the similarity of two sample sets in general. Here it is defined as

$$J = \frac{S_{11}}{S_{01} + S_{10} + S_{11}}, \tag{3.2}$$

where S_{xy} counts the number of relations having status $a_{ij} = x$ at the first instance and $a_{ij} = y$ at the second. The JI measures links which survive as a fraction of links which are established at any of the two points in time. Hence, it also takes into account those banks which are active in only one of the two periods. Figure 3.7 shows that the JI is very stable over time for longer aggregation periods, but not for the daily level. As expected, the JI tends to be higher for longer aggregation intervals. The daily measures are much more unstable and increase substantially until the GFC. More problematic than the smaller average level are however the extreme outliers on the downside.

tigating this issue.

[17] The so-called graph correlation, see e.g. Butts and Carley (2001), shows qualitatively very similar results, but is not able to cope with banks entering or exiting market. The correlation of both measures is always above .9 irrespective of the aggregation period.

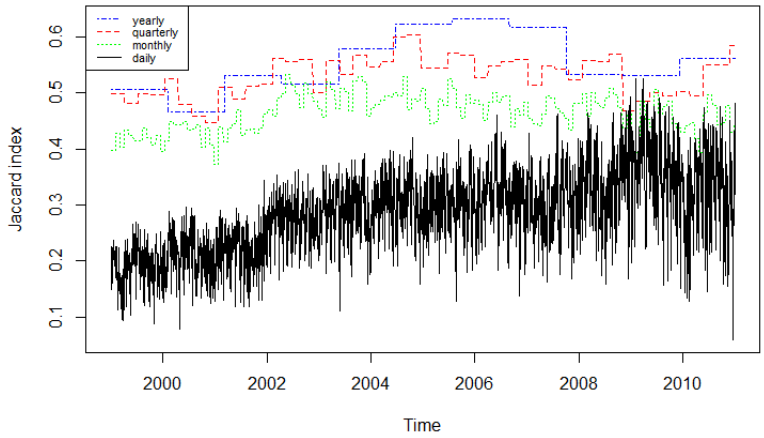

Fig. 3.7: Jaccard Index for daily (black), monthly (green), quarterly (red) and yearly (blue) networks.

As a rule-of-thumb, in social network analysis one considers networks with JI values above .3 as substantially stable.[18]

Table 3.1 shows the mean, minimum, 10th percentile and standard deviation of the JIs for different aggregation periods. Again, the most evident observation is that the daily networks are rather special: the minimum and the 10th percentile of the JI are significantly smaller, indicating that we observe values below .2 in at least 10% of the sample, which is not a rare event.[19]

[18] See Snijders *et al.* (2009).

[19] We should note that, as apparent from Figure 3.7, the reason for the 10th percentile to be below the .2 threshold is not the GFC.

Jaccard Index	Year	Quarter	Month	Day
mean	.5543	.5302	.4638	.2861
standard deviation	.0535	.0368	.0333	.0740
min	.4652	.4479	.3735	.0603
10th percentile	.5072	.4835	.4183	.1904

Tab. 3.1: Jaccard Index for daily, monthly, quarterly and yearly networks. Calculations were carried out for all subsequent networks at the different aggregation periods. Standard deviations based on all observations.

Reciprocity	Year	Quarter	Month	Day
mean	.4264	.2085	.0829	.0042
standard deviation	.0580	.0423	.0244	.0060

Tab. 3.2: Reciprocity of the Italian Banking network. Calculations were carried out for all networks at the different aggregation periods. Standard deviations based on all observations.

These results suggest a high degree of randomness in the daily networks.

Obviously, higher values of the JI are no guarantee that we are closer to the 'real' network per se. Note that in a network with randomly drawn connections, the index should be positively related to the length of the aggregation period. Thus, it is important to show that other network measures also take on values significantly different from random networks for longer aggregation periods. In the following, we will therefore have a closer look at the reciprocity of the network.

Reciprocity is a global concept for directed networks that measures how many of the existing links are mutual. It can be calculated by adding up all loops of length two, i.e. reciprocal links, and dividing them by the total number of links.

Table 3.2 shows higher levels of reciprocity for longer aggregation periods.[20] In the case of daily networks we observe very few mutual links, as Iori *et al.* (2008) stated this is a very plausible finding, since banks rarely borrow and lend money from the same bank within a particular day. However, the values for longer aggregation periods show that the banking network is not one-sided, supporting the evidence on the inability of daily networks to represent the 'true' underlying (directed) banking network. The left panel of Figure 3.8 illustrates the results for 1999, where we perform a similar analysis as for the density above, by comparing the observed network reciprocity to those of Erdös-Renyi and scale-free random networks.[21] The actual values are, again, always the lowest. The right panel of Figure 3.8 shows that the reciprocity (after 19 days for the real network) exceeds the density for all three networks. Note that for random networks one would expect the two measures to be almost identical. However, different banks are active at different days and if many banks are often simultaneously active the chance

[20] A structural break (after the GFC) is detected by a Chow-test as well as an additional CUSUM test for the 10th year, the 39th quarter and the 117th month respectively, but not for daily networks. For the yearly networks only the Chow-test indicates a structural break. Note that the yearly analysis involves only 12 data points.

[21] Again the results are qualitatively very similar for the other years as well.

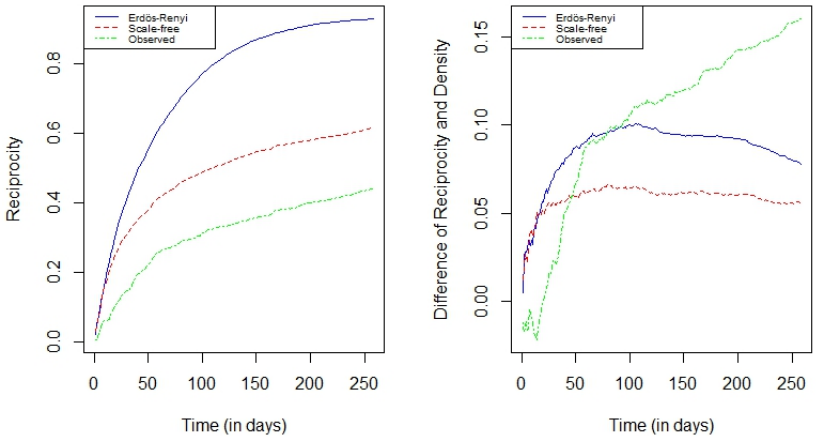

Fig. 3.8: Data for 1999. Left: reciprocity for the Erdös-Renyi (blue), Scale-Free (red) with $\alpha = 2.3$, and observed networks (green). Right: difference between reciprocity and density for the respective networks. Aggregation period in days.

of forming a reciprocal link is higher (remember that we used the actually active banks of each day in the Monte Carlo exercise). However, more important for our analysis is that for the real network the difference between reciprocity and density increases steadily and exceeds the difference of the random networks. Hence, using longer than daily aggregation is not only capable of taking mutual credit relationships into account, but even indicates a preference of banks to form them. On the other hand, the saturation of the reciprocity indicates that banks will have mutual credit relationships with most of their counterparties.[22] Thus, the noise level of networks with longer aggregation periods is smaller and the directed version of the networks contains a substantial amount of information.

On the base of the Jaccard index, monthly and quarterly networks appear most stable as they have a high index with a very low standard deviation, i.e., the highest degree of structural stability of lending relationships over time. Yearly aggregation levels, in contrast, have somewhat higher variation in their JI and might be problematic, because the banking network (in particular during unstable times) is likely to evolve much faster. Somewhat related is the change in the composition of banks, i.e. banks leaving and entering the

[22] This may of course be affected by the time-varying composition of active banks. See below.

market, since we consider a bank as active for the whole year even if it leaves the market after the first trading day.[23] Concerning the monthly level, Iori *et al.* (2008) discovered intradaily and -monthly seasonalities which may affect our results. In everything that follows, we will therefore mostly focus on the quarterly networks.

3.4.4 Transitivity

Here we are interested in transitive relations between three banks. The concept of transitivity states that a specific relationship $\diamond$ is transitive if from $i \diamond j$ and $j \diamond k$ it follows that $i \diamond k$ holds. Equality is a transitive relation, but inequality is not. From $i = j = k$ follows $i = k$, yet $i \neq j \neq k$ does not imply $i \neq k$.

The relation we are interested in is i has a link to j or $a_{ij} = 1$. The relationship is obviously not transitive, since i has a link to j and j has a link to k does not strictly imply that k also has a link to i. However, it is interesting to investigate how many such closed triplets occur. More generally speaking, transitivity measures whether the existence of certain links depend both on the relation between the two counterparties and on the existence of links with a third party. The measure most prominently used for this

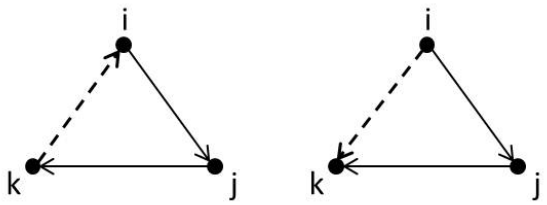

Fig. 3.9: The two possibilities how the directed path of length two (solid lines) between i and k can be closed. On the left hand side the path is closed into a loop of length three (CC^1). On the right hand side the triplet is interconnected but not in the same single direction (CC^2).

purpose is the (directed) clustering coefficient[24], which, despite its name, has no relation to cluster identification whatsoever.[25] It measures the number of

[23] This problem occurs for each aggregation period, but is likely to become more pronounced for longer frequencies.

[24] For more detailed definitions of clustering coefficients see Zhou (2002).

[25] See Fricke and Lux (2012) and Fricke (2012) for detailed approaches of cluster identification in the e-MID market.

(directed) paths of length two in the network and takes the fraction of these which are closed.[26] Figure 3.9 illustrates two ways to close the triplet i, j, k in a directed network. First, the directed path may be closed into a loop as shown on the left hand side of the figure. The function of such closure is given by the coefficient CC^1:

$$CC^1 = \frac{\sum_{j \neq i \neq k} a_{ij} a_{jk} a_{ki}}{\sum_{j \neq i \neq k} a_{ij} a_{jk}} \tag{3.3}$$

Second, the link from i to k may be reversed. The function of such closure is given by the coefficient:

$$CC^2 = \frac{\sum_{j \neq i \neq k} a_{ij} a_{jk} a_{ik}}{\sum_{j \neq i \neq k} a_{ij} a_{jk}} \tag{3.4}$$

An important distinction is that the nodes in the case of CC^1, as apparent from Figure 3.9, have each one in- and one outgoing link and therefore show no hierarchical ordering. Figure 3.10 shows that the results for the two coefficients are very different. The mean for CC^1 is .164 and .571 for CC^2. This is further evidence for the non-random character of the banking network, since the probability of an 'average' link to exist is just equal to the density of .208. Hence, the existence of a path of length two between i and k via j makes it 2.75 times more likely that the link from i to k exists compared to a random link[27], but reduces the probability that there is a link from k to i by 21%. The huge difference indicates that the banking network has a hierarchical ordering on the triadic level.

Figure 3.11 illustrates that for the Erdös-Renyi networks the evolution of the clustering coefficients is almost identical (correlation above .999). The exact numbers of the clustering coefficients for the observed networks change with the aggregation, but CC^1 is always much smaller than CC^2 for all aggregation levels as shown by figure. CC^2 is at the beginning even higher for the observed network than for the random networks, but saturates after a steep increase relatively quickly on a much lower level (up to .624). CC^1 on the other hand is almost zero at the beginning. Note that a loop of length three at a single day implies that each involved bank would get back some of its own lending via an intermediate bank, which appears very unlikely. However, CC^1 increases (up to .386) for longer aggregation periods showing that such relations do exist. Note that the daily or undirected networks are not capable of taking this into account.

[26] Any connection along directed links between two nodes i and j is called a path and the length of the path is defined as the number of edges crossed. There are no restrictions on visiting a node or link more than once alongside a path.

[27] The probability of random link is exactly the density of .208.

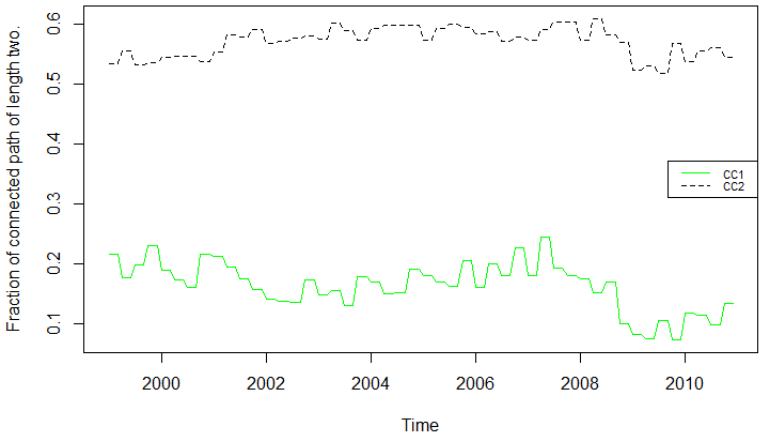

Fig. 3.10: The two clustering coefficients of quarterly networks: CC^1 (green) is the fraction of path of length two which are closed into a loop by a third link. CC^2 characterizes links in which the triangle is closed in a hierarchical way.

3.4.5 Small-World Property

Another very prominent measure in the network literature is the average shortest path length (ASPL). Interest in this measure stems from the remarkable finding that in many 'real world' networks the ASPL is quite small, also known as the small-world phenomenon.[28] Here we focus on the undirected version of the network.

Watts and Strogatz (1998) show that completely random networks already have a very small ASPL, but at the same time a relatively low (undirected) clustering coefficient (CC) equal to its density. On the other hand, a regular network, where all nodes have connections to their s nearest neighbours, has a CC of 1, but a very high ASPL. Interestingly, already the introduction of a few 'random' links reduces the ASPL significantly, because these 'long range' links connect different clusters of the network. The authors argue that the interesting 'real world' networks are neither purely regular nor random. Therefore Watts and Strogatz (1998) define small-world networks

[28] Small ASPLs have been detected for social, information, technological and biological networks. The first empirical finding dates back to the chain letter experiments conducted by Milgram (1967). His finding that on average only six acquaintances are needed to form a link between two random selected persons led to the famous phrase of 'six degrees of separation'.

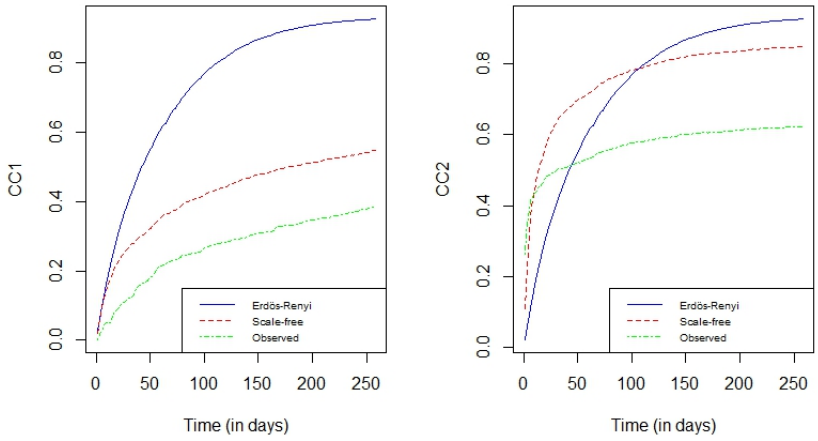

Fig. 3.11: Left: CC^1 for the Erdös-Renyi (blue), Scale-free (red) with $\alpha = 2.3$, and observed networks (green), while the are aggregated on a daily basis up to one year(1999). Right: CC^2.

to additionally have a higher CC than random networks. Hence, to further investigate if the banking network exhibits the small-world property, we calculate the ASPL and the CC and compare the values to random networks.

The (symmetric) matrix $\mathbf{G}$ of dimension $N \times N$ contains the geodesic distances between any two nodes, i.e. each element g_{ij} is the length of the geodesic path between node i and j. The ASPL is calculated by dividing the sum of all (existing) geodesic path lengths by the total number of (existing) geodesic paths.[29] The CC is calculated similar to the directed case (CC^1 and CC^2), but ignores the directedness of the links.

In this case a comparison of the networks aggregated day by day is not suitable, since the density of both random networks as visible in figure 3.6 after a quarter (around 63 days) is much higher. Therefore the simulated random networks correspond to the aggregated quarterly networks with respect to their density (and their out-degree distribution for scale-free networks). Table 3.3 summarizes the results: The ASPL for all networks is small and

[29] It is not necessary for two nodes to have a shortest path, since there might be no link leading from one to the other. In this case the two nodes lie in different components of the network and by convention the length of these non existing geodesic paths are set to infinity. The undirected banking network consists always of only one component and the same is true for the investigated random networks, because of the high density.

	CC	ASPL
Erdös-Renyi	.3736	1.6470
	(.0058)	(.0073)
Scale-free	.4131	1.6259
	(.0103)	(.0134)
Observed	.5422	1.6486
	(.0434)	(.0373)

Tab. 3.3: The average CC and ASPL for the observed network and for 100 Erdös-Renyi and scale-free (with $\alpha = 2.3$) random networks. Standard deviations in brackets.

almost identical, in which the very low value is due to the high density. The CC for the Erdös-Renyi networks is by construction close to the density, since all links have per quarter exactly this probability to occur. The CC for the scale-free is with .413 higher, but is exceeded by the observed network (.542). This indicates the higher regularity in the link structure. Hence, the banking network lies midway between regular and completely random graphs. As can be seen from the last column of Table 3, the ASPL does not provide much scope for distinguishing between the benchmark Erdös-Renyi and scale-free networks and the empirical ones. The reason might be that the relatively high density leads to relatively short path lengths anyway.[30]

3.4.6 Effects of the Global Financial Crisis

Finally, we take a closer look at the effects of the GFC on the banking network.[31] The start of the GFC is not easy to determine, but we have seen that the collapse of Lehman Brothers in quarter 39 (2008 Q4) was a major shock for the global financial market in general and the Italian interbank market as well.[32] The effects of this event were twofold: first, the counterparties of Lehman Brothers realized huge loses. Second, the perception of risks changed, since Lehman had been considered to be 'too-big-too-fail' before. The resulting dramatic increase of perceived counterparty risk reduced

[30] The results are qualitatively the same if we consider the directed version of the network. The ASPL for the quarterly networks is 1.912 against 1.802 for the random networks. The CC^2 (.571) is significantly higher than for Erdös-Renyi networks with .208, whereas CC^1 (.164) is even smaller. However, in this case the network consists not only of one giant component which makes the interpretation of the ASPL more difficult.

[31] See also Fricke and Lux (2012).

[32] See Brunnermeier (2008).

the willingness of banks to lend to each other, which ultimately affected the real economy due to tighter lending restrictions. The monetary authorities and governments around the world injected substantial amounts of capital into the financial system to prevent interbank markets from freezing in the following weeks. We have seen, that important network measures such as density and reciprocity, were significantly affected by these events as well.

Quarter	36	37	38	39	40
Banks	101	100	100	98	97
Volume	445,991	409,340	435,338	404,353	385,819
Trades	20,984	20,078	19,963	18,160	16,477
Links	2,425	2,253	2,249	2,153	1,768
Trades per Link	8.65	8.91	8.88	8.43	9.32
Quarter	41	42	43	44	45
Banks	96	96	94	94	94
Volume	234,102	267,057	197,021	227,076	196,503
Trades	14,184	13,981	12,525	11,636	11,577
Links	1,533	1,506	1,599	1,449	1,530
Trades per Link	9.25	9.28	7.83	8.03	7.57

Tab. 3.4: The table summarizes the number of Italian banks, total volume (million Euros), number of trades, number of links and links per trade for the quarters 36-45, while the Lehman collapse has been in quarter 39.

Here we investigate the change in banks' behavior during and after the breakdown of Lehman Brothers. To begin with, Table 3.4 contains several basic network statistics for ten quarters around the breakpoint (quarter 39). Interestingly, the number of active Italian banks remained relatively stable during this period, and in fact 86 banks were active in all of the ten quarters. The stability of this composition is important, since under these circumstances changes in the behavior of the banks on the aggregated or individual level should be mainly driven by their response to this exogenous shocks.

We also see that the total trading volume, the number of trades and the number of links are all decreasing over this period, but the exact patterns are distinct. Surprisingly, the volume is quite stable until the 40th quarter, but drops by 39.3 percent in the 41st quarter. In contrast, the number of trades starts to fall in the 39th quarter and decreases further until quarter 44. The total number of links already starts to decrease in quarter 37, but overall

tends to develop in a very similar way to the number of trades.[33] In the end, the most immediate reaction to the crisis was that banks traded similar total volumes but in fewer trades and with a smaller number of counterparties in order to minimize their (perceived) counterparty risks.[34]

Interestingly, the trades per link are the highest for the three quarters after the Lehman collapse. This indicates that the banks relied stronger on their preferred counterparties. Preferred in this context might simply means that the banks had more reliable information about these banks, which however should coincide with former trading relationships. Eventually, we conclude that the breakdown of Lehman Brothers significantly affected the behavior of individual banks and thus had a clear impact on the structure of the network. Quite surprisingly, we find that the link structure of subsequent networks remained rather stable during this period, since no structural break is detected for the Jaccard Index. Furthermore, despite the significant impact of the GFC, we do not find evidence for a complete drying up of the e-MID market, even at the daily level.[35]

3.5 Conclusions and Outlook

In this paper, we have investigated the interbank lending activity as documented in the e-MID data from 1999 until the end of 2010 from a network perspective. Our main finding is that daily networks feature too much randomness to be considered a representative statistic of some underlying latent network. The JI shows the higher consistency over time for longer aggregation periods and the very low density compared to (aggregated) random networks indicates the existence of preferred trading relations. In general the evolution of all global network measures for longer aggregation periods (month, quarter, year) is very similar in their deviation from the Erdös-Enyi and scale-free benchmarks. Moreover, the monthly and quarterly networks are characterized by a significantly higher than random clustering coefficient, and thus reveal some regularity in the link structure. The (almost) zero reciprocity and CC^1 of daily networks proves the inability of this aggregation level to reveal information on such structural elements. However, quarterly networks consistently exhibit a non-random structure and allow us to consider the mutuality of the relations and are therefore a preferable subject of study, especially if one is interested in the evolution of the network over time.

[33] The correlation between both is .963 for this period and .957 over the complete sample.

[34] Fricke and Lux (2012) show that most of these changes were in fact driven by behavioral changes of core banks.

[35] As noted above, this is not true for loans longer than overnight. These markets essentially collapsed completely, which is not surprising given that the loans are unsecured.

Essentially, these results show that it is far from trivial to map a given data structure into a 'network'. While daily records of the interbank trading system can be arranged in an adjacency matrix and treated with all types of network statistics, they provide probably only a very small sample of realizations from a richer structure of relationships. Just like daily contacts of humans provide very incomplete information of networks of friendship and acquaintances, the daily interbank data might only provide a small selection of existing, dormant established trading channels. Hence, inference based on such high-frequency data may be misleading while a higher level of time aggregation might provide a more complete view on the interbank market. What level of aggregation is sufficient for certain purposes is an empirical question depending on the research questions at hand. However, saturation of certain measures may be a good indicator that most dormant links have been activated at least once over a certain time horizon. At the same time, such dependence of statistics on the time horizon serves to sort out a number of simple generating mechanisms (i.e. completely randomly determined networks in every period) and reveal interesting dynamic structure.

Another interesting result is that the network is asymmetrical in many respects. For the quarterly network the fraction of reciprocal links is very similar to the density. Furthermore, the two directed clustering coefficients are very different. The probability for path of length two to be closed into a loop is 3.48 times smaller than the other way. Additionally, the correlation between in- and out-degree is merely .12 for the complete sample. Therefore, the information that a bank has a a large number of incoming links is a surprisingly poor indicator of how many outgoing links the bank has.

Moreover, for many measures the GFC could be identified as a structural break and also the decreasing number of volume, trades and links support that the GFC heavily affected the Italian interbank market. However, the network overall remained surprisingly stable and despite the decrease of its volume (in the beginning of 2010) the e-MID market was never close to drying up completely.

In the future more attention should be given to the analysis of directed banking networks using longer aggregation periods to identify structural commonalities. This has important consequences for the regulation of credit institutions, since at the daily level it is difficult to detect the systemically important institutions. For policymakers and regulators, it would be potentially (dangerously) misleading to focus on the noisy daily networks, even more since the low level of connectivity suggests a low-degree of systemic risk at any point in time. More important, in our view, is to get a better idea on the wider pool of counterparties of all credit institutions, in order to detect possible behavioral changes among the set of relatively active banks. Such

changes might then serve as an indicator for funding problems of individual institutions.[36] In the end, it would be important to extend our phenomenological analysis in order to test hypotheses about the behavior of banks at the micro-level that drives the system's properties.

[36] See Fricke and Lux (2012).

4. ON THE DISTRIBUTION OF LINKS IN THE INTERBANK NETWORK: EVIDENCE FROM THE E-MID OVERNIGHT MONEY MARKET

4.1 Introduction and Existing Literature

Since the onset of the global financial crisis (GFC) in 2007/08, the analysis of network structures formed by interbank liabilities has received increasing attention. Considering an ensemble of financial institutions, individual banks are connected to each other through some of their activities (usually credit flows) and the bilateral exposures can be mapped into a credit network. Such a perspective is useful in order to study the knock-on effects on other banks due to disruptions of the system caused by the failure of individual nodes (e.g. insolvency of one bank). A new strand of literature has started to construct financial networks based on empirical data available at supervisory authorities or hypothetical network structures to investigate the contagious effects of failures of single banks.[1] A basic finding of network theory is that the topology of a network is important for its stability, with the interbank network obviously being no exception.[2] In this regard, the understanding of the structure and functioning of complex networks has advanced significantly in recent years.

In this paper, we focus on one of the most prominent network characteristics, namely the degree distribution, where the degree is the number of (incoming/outgoing) connections per node. Even though the degree distribution does not provide sufficient information for all facets of the structure of the network (Alderson and Li, 2007), it is often considered as one of the defining characteristics of different network types. For example, networks with random link formation (Erdös and Renyi, 1959, or ER random networks) display Poisson degree distributions, i.e. most nodes have degrees within a relatively narrow range. In contrast, many real-world networks have been reported to display fat-tailed degree distributions: most nodes have a very small degree, but the tail contains nodes with substantially larger degrees (cf. Clauset *et al.*, 2009). This feature is shared by the important class known as scale-free (SF) networks, in which the fraction of nodes with degree k is proportional to $k^{-\alpha}$, where α is the so-called scaling parameter. The term scale-free indicates that there is no typical scale of the degrees, i.e. the mean may not be representative. These networks received considerable attention in the literature due to a number of interesting properties (cf. Caldarelli, 2007). One important feature of scale-free networks is that they can be described as robust-yet-fragile,[3] indicating that random disturbances are easily absorbed (robust) whereas targeted attacks on the most central nodes may lead to a breakdown of the entire network (fragile). Quite interestingly,

[1] See e.g. Upper and Worms (2004), Nier *et al.* (2007), and Gai *et al.* (2011).
[2] See Haldane and May (2011) and Albert *et al.* (2000).
[3] See Albert *et al.* (2000).

many interbank networks have been reported to resemble scale-free networks (cf. Boss *et al.*, 2004, Soramäki *et al.*, 2007, De Masi *et al.*, 2006, and Iori *et al.*, 2008). If the network of credit relationships had such a structure, this would carry important policy implications. For instance, such a network might experience long stable periods, during which disruptions are confined to peripherical banks and can be absorbed easily within the entire system. However, such periods could be a misleading indicator of the overall stability of the system as problems affecting the most central nodes could suddenly cause a breakdown of the entire network, cf. Haldane (2009).

The distribution of network degrees is just one example among many phenomena in the natural sciences as well as from the socio-economic sphere that have been claimed to follow a scaling law (power-law or Pareto-law). Other well-known examples include: Zipf's law for the city size distribution (Gabaix, 1999), the distribution of firm sizes (Axtell, 2001), the size distribution of innovations (Silverberg and Verspagen, 2007) or the distribution of large asset returns (Mandelbrot, 1963, Lau *et al.*, 1990, and Jansen and de Vries, 1991). While these examples appear to be supported by empirical evidence and meanwhile count as stylized facts, a variety of other findings of power-laws seem more questionable. It appears from a number of recent reviews of power-law methodology and power-law findings (cf. Avnir *et al.*, 1998, Stumpf and Porter, 2012) that there had been an over-emphasis on scaling laws and often too optimistic interpretation of statistical findings in the literature of the natural sciences. For instance, in a meta-study of power-laws reported in publications in the main physics outlet *Physical Review* between 1990 and 1996, Avnir *et al.* (1998) found that most claims of power-laws (aka scaling or fractal behavior) had a very modest statistical footing. As they say '... the scaling range of experimentally declared fractality is extremely limited, centered around 1.3 orders of magnitude.' In terms of statistics jargon this means that the more typical declaration of a power-law in these publications is based on a partially linear slope in a relatively small intermediate range of the empirical cumulative distribution of some observable.

The power-law exponent (like the ones reported for the degree distribution) is typically obtained by a linear regression in a log-log plot of the cumulative distribution. Obviously, this approach suffers from a number of shortcomings: (i) even if the hypothetical data-generating process is a Pareto distribution, this log-log fit would not be an efficient way to extract the parameter of the underlying distribution.[4] It is actually a method that is definitely inferior to maximum likelihood (which is easy to implement), and

[4] See Goldstein *et al.* (2004). Gabaix and Ibragimov (2011) improve the regression method by shifting the rank observations.

results are hard to interpret as, due to the dependency of observations in the log-log plot of the *cumulative* distribution, the statistical properties of this estimator are not straightforward, (ii) the implicit censoring of the data that is exerted by selecting a *scaling* range makes it easy to deceive oneself. Many distributions might actually have some intermediate range in their 'shoulders' where their cdf looks appropriately linear. But their remaining support (small and large realizations) might display a completely different behavior. Since power-laws in the natural sciences are thought to be interesting if they extent over several orders of magnitude, it is unclear what the interpretation of such an intermediate power-law approximation would be.

Statistical extreme value theory (EVT) provides yet another perspective on power-law behavior. The basic result of this branch of statistics is a complete characterization of the limiting distributions of extremes (maximum or minimum) of time series of iid observations (where results for the iid case have been generalized for dependent processes under relatively mild conditions, cf. Leadbetter, 1983, and Reiss and Thomas, 2007, for details). According to EVT, the appropriately scaled minimum or maximum of a series of observations converges in distribution to one of only three functional forms: the Fréchét, Gumbel or inverted Weibull distribution. Since extremes are by definition very rare, it is often even more relevant, that the tail of a distribution converges in distribution in a similar way to one of three adjoint functional forms. Namely, the outer part approaches either a power-law decay, an exponential decay or a decay towards a fixed endpoint for the three types of extremal behavior, respectively. Power-law behavior is, therefore, a very general form of limiting behavior for the large realizations of a stochastic process. EVT has originally been developed for continuous distribution function. Since degree distributions are discrete (degrees being integer numbers), its is worthwhile to note that corresponding limit laws for discrete variables are available as well, cf. Anderson (1970). In our context this might imply that very large realizations of the degree distribution could still decay like a power-law even if the bulk of the distribution does not appear to follow such a distribution (and the implications for the fragility of the system might be similar as for 'true' scale-free networks). It is important to emphasize that both the limiting behavior of extremes and tails are stable under aggregation. Hence, data at different levels of (time-) aggregation should obey the same extreme value and tail behavior.

One reason for the 'popularity' of power-laws in the natural sciences is that they are often the signatures of relatively simple and robust generating mechanisms that might apply to a variety of phenomena. In the case of networks, a power-law distribution of degrees is the imprint of so-called scale-free networks. Reported power-laws for interbank networks have been

within a relatively narrow range around 2.3 both for the in- and out-degree distributions (see e.g. Boss *et al.*, 2004, Soramäki *et al.*, 2007, and De Masi *et al.*, 2006), even though most papers lack a thorough statistical analysis of the issue, with Bech and Atalay (2010) being a notable exception. If these findings were robust, the known generating mechanisms for scale-free networks would be strong candidates as mechanisms for the formation of interbank links. Furthermore, the well-known reactions of scale-free networks to disturbances would be of immediate concern for macro-prudential regulation. Thus, taking into account the relevance of such topological features, and the documented over-emphasis on power-law behavior, a more rigorous statistical analysis of the distributional properties of interbank network data should be worthwhile. Similar approaches have revealed that numerous previous claims of power-law behavior were not supported by the data (Stumpf and Porter, 2012).

In this paper, we consider interbank networks based on the Italian e-MID (electronic market for interbank deposits) data for overnight loans during the period 1999-2010. Our main focus is to fit a set of different candidate distributions to the degrees for different time horizons. Using daily data over the period 1999-2002, De Masi *et al.* (2006) reported power-laws for the distribution of in- and out-degrees, with tail parameters 2.7 and 2.15, respectively. Finger *et al.* (2012) have shown recently that the networks' properties depend on the aggregation period.[5] We will, therefore, not confine our analysis to daily data (the basic frequency of our data set), but also look at the distribution of in- and out-degrees for networks constructed on the base of aggregated data over longer horizons. Quite surprisingly in view of the previous literature, we find hardly any support in favor of previously reported power-laws: at the daily level the degrees are usually fit best by negative Binomial distributions, while the power-law may provide the best fit for the tail data. However, we typically find very large power-law exponents (with values as large as 7), i.e. levels where the power-law is virtually indistinguishable from exponential decay. At the quarterly level, Weibull, Gamma, and Exponential distributions tend to provide comparable fits for the complete degree distribution, while the tails again tend to display exponential decay. We find comparable results when investigating the distribution of the number of transactions, even though in this case the tails of the quarterly variables are somewhat fatter. However, the Log-normal distribution typically outperforms the power-law. Overall these findings indicate that the

[5] Since we cannot easily observe the state of a hypothesized network of interbank links at a given point in time, some data aggregation is necessary. Usually, for time-aggregated data a link is assumed to exist between two banks, if there has been a trade at any time during the aggregation period.

power-law is typically a poor description of the data, implying that preferential attachment and other generating mechanisms for scale-free networks are unsuitable explanatory mechanisms for the structure of the Italian interbank network. Moreover, the networks contain a substantial level of asymmetry, due to the low correlation between in- and out-degrees. Additionally, we find that the two variables do not follow identical distributions in general.

The remainder of this paper is structured as follows: Section 4.2 gives a short introduction into (interbank) networks, section 4.3 briefly introduces the Italian e-MID trading system and gives an overview of the data set we have access to. Section 4.4 describes our findings and section 4.5 concludes and discusses the relevance of these findings for future research.

4.2 Networks

A network consists of a set of N nodes that are connected by M edges (links). Taking each bank as a node and the interbank positions between them as links, the interbank network can be represented as a square matrix of dimension $N \times N$ (data matrix, denoted $\mathbf{D}$). An element d_{ij} of this matrix represents a gross interbank claim, the total value of credit extended by bank i to bank j within a certain period. The size of d_{ij} can thus be seen as a measure of link intensity. Row (column) i shows bank i's interbank claims (liabilities) towards all other banks. The diagonal elements d_{ii} are zero, since a bank will not trade with itself.[6] Off-diagonal elements are positive in the presence of a link and zero otherwise.

Interbank data usually give rise to directed, sparse and valued networks.[7] However, much of the extant network research ignores the last aspect by focusing on binary adjacency matrices only. An adjacency matrix $\mathbf{A}$ contains elements a_{ij} equal to 1, if there is a directed link from bank i to j and 0 otherwise. Since the network is directed, both $\mathbf{A}$ and $\mathbf{D}$ are asymmetric in general. In this paper, we also take into account valued information by using both the raw data matrix as well as a matrix containing the number of trades between banks, denoted as $\mathbf{T}$. In some cases it is also useful to work with the undirected version of the adjacency matrices, $\mathbf{A}^u$, where $a^u_{ij} = \max(a_{ij}, a_{ji})$.

As usual, some data aggregation is necessary to represent the system as a network. In the following, we define interbank networks by aggregating over daily as well as quarterly lending activity.

[6] This is of course only true when taking banks as consolidated entities.

[7] Directed means that $d_{i,j} \neq d_{j,i}$ in general. Sparse means that at any point in time the number of links is only a small fraction of the $N(N-1)$ possible links. Valued means that interbank claims are reported in monetary values as opposed to 1 or 0 in the presence or absence of a claim, respectively.

4.3 The Italian Interbank Market (e-MID)

The Italian electronic market for interbank deposits (e-MID) is a screen-based platform for trading of unsecured money-market deposits in Euros, US-Dollars, Pound Sterling, and Zloty operating in Milan through e-MID SpA.[8] The market is fully centralized and very liquid; in 2006 e-MID accounted for 17% of total turnover in the unsecured money market in the Euro area, see European Central Bank (2007). Average daily trading volumes were 24.2 bn Euro in 2006, 22.4 bn Euro in 2007 and only 14 bn Euro in 2008.

Detailed descriptions of the market and the corresponding network properties can be found in Finger *et al.* (2012).[9] In this paper we used all registered trades in Euro in the period from January 1999 to December 2010. For each trade we know the banks' ID numbers (not the names), their relative position (aggressor and quoter), the maturity and the transaction type (buy or sell). The majority of trades is conducted overnight and due to the global financial crisis (GFC) markets for longer maturities essentially dried up. We will focus on all overnight trades conducted on the platform, leaving a total number of 1,317,679 trades. If not stated otherwise, the reported results are based on trades conducted between Italian banks only, reducing the total number of trades to 1,215,759.

4.4 Results

In this section we present empirical results on the dynamics and distribution of the number of links (degrees) and the number of transactions (ntrans) of individual institutions. The degree of a node gives the total number of links that a bank has with all other banks and can thus be seen as a measure for the importance of individual nodes. Undirected networks imply symmetric adjacency matrices. In this case bank i's total degree k_i is simply the number of relationships bank i has with other banks, i.e.

$$k_i^{total} = \sum_{j \neq i} a_{ij}^u. \tag{4.1}$$

For directed networks, we differentiate between incoming links (bank i borrows money from other banks) and outgoing links (i lends money to other

[8] The vast majority of trades (roughly 95%) is conducted in Euro.
[9] See also the e-MID website http://www.e-mid.it/.

banks), and define the in- and out-degree of i (k_i^{in} and k_i^{out}) as

$$\begin{aligned} k_i^{in} &= \sum_{j \neq i} a_{ji} \\ k_i^{out} &= \sum_{j \neq i} a_{ij}, \end{aligned} \tag{4.2}$$

respectively. Note that our networks contain only banks with at least one (directed) link. In this way, the total degree of a sample bank is always at least equal to one, while it may be the case that either the in- or out-degree equals zero for a particular bank. Since we ignore zero values in the distribution fitting approach, this affects the number of observations for the different variables.

For the number of transactions, we use similar definitions based on the $\mathbf{T}$ matrix, with each element $t_{i,j}$ giving the number of trades with credit extended from bank i to bank j. To be precise, we calculate the number of in-/out-transactions as

$$\begin{aligned} n_i^{in} &= \sum_{j \neq i} t_{ji} \\ n_i^{out} &= \sum_{j \neq i} t_{ij}. \end{aligned} \tag{4.3}$$

Additionally, we analyze the total number of transactions, for simplicity defined as the sum of in- and out-transactions

$$n_i^{total} = n_i^{in} + n_i^{out}. \tag{4.4}$$

4.4.1 Dynamics of the Degrees and Number of Transactions

Before investigating the distribution of the variables under study, we provide a brief overview of their dynamics over time, restricting ourselves to quarterly data here. Figure 4.1 shows the in-/out-degrees (left) from the directed networks and the total degrees from the undirected networks (right). The upper left panel shows the mean and median in-degree and out-degrees over time.[10] Clearly, the mean values are decreasing over time, and so does the median in-degree which is mostly very close to the mean value. For both series we find a significant structural break after quarter 10. In contrast, the median out-degree fluctuated around an average value of roughly 17 over most of the sample period, but with a significant structural break after quarter 39 due to the GFC. These values are considerably smaller than the values

[10] Note that the mean in- and out-degree are identical by definition.

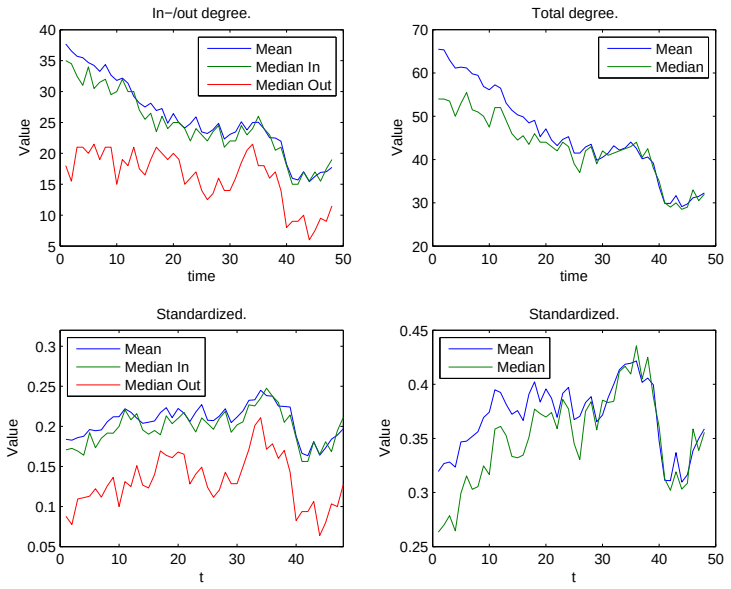

Fig. 4.1: Mean and median degree over time. Left: in- and out-degree. Right: total degree. Top: absolute levels. Bottom: standardized values (divided by the number of active banks per quarter).

for the in-degree, pointing towards a substantial level of skewness in the out-degree distribution. Thus, the distributions of in- and out-degrees are likely to be not identical. The lower left panel shows the relative mean and median degree over time, i.e. the values in the upper panel standardized by the number of nodes active in each quarter. We see that the negative trend in the upper panel is mostly driven by the negative trend in the number of active banks. Thus, the standardization appears to make the in-degrees of different quarters comparable. This is less so for the median out-degree, which is far more volatile over the sample period.[11] For the sake of completeness, the corresponding values for the degrees from the undirected networks are shown on the right-hand side. Both for the absolute and relative values the mean and median values are very similar, except for the beginning of the sample period. This is driven by the high level of asymmetry in the out-degree distribution for the first half of the sample, which appears to decrease later on.

What does the evidence on the differences between the in- and out-degree

[11] Interestingly, after standardizing the degrees, we find structural breaks in all three time series close to quarter 39, i.e. around the GFC.

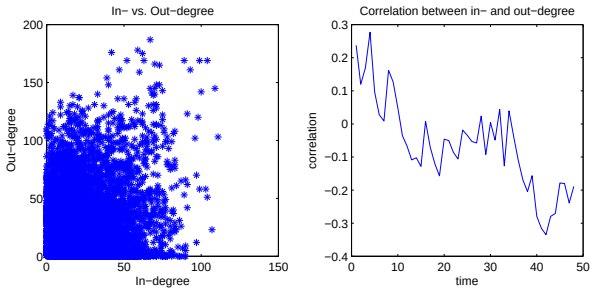

Fig. 4.2: Left: Scatter plot of in- vs. out-degree. Correlation: .0899. Right: Correlation between individual banks' in- and out-degree over time. Italian banks.

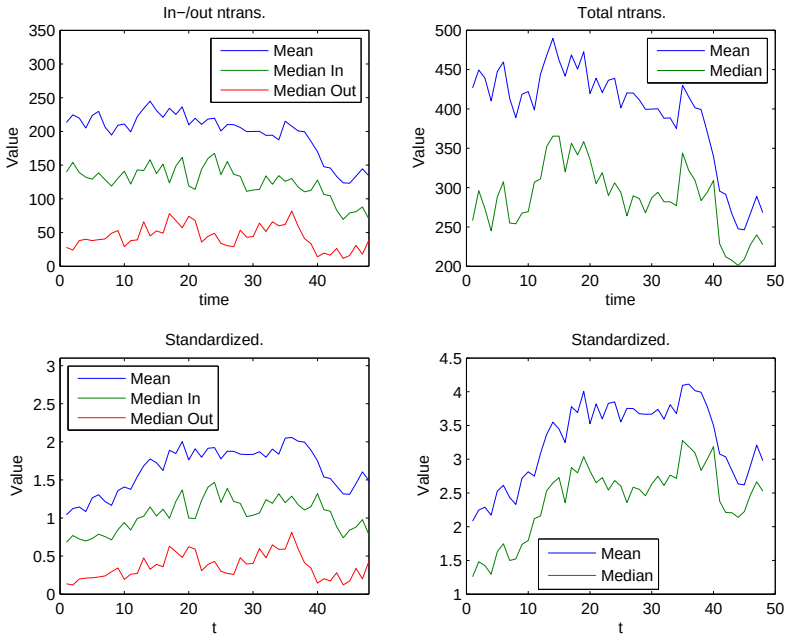

Fig. 4.3: Mean and median number of transactions over time. Left: Directed Network. Right: Undirected Network. Top: absolute levels. Bottom: standardized values (divided by the number of active banks per quarter).

distributions imply? Given that many studies on interbank markets work with undirected networks, these studies entail the implicit assumption of a high correlation between in- and out-degrees of individual banks. The left panel of Figure 4.2 shows a scatter-plot of in-degree against out-degree for Italian banks, showing a small correlation of .0899 for all observations. For single quarters, we find that the correlation between these measures may be very small, at times even negative. Thus, banks with a high in-degree do not necessarily have a high out-degree and vice versa. The directed version of the network contains a considerable amount of information. The right panel of Figure 4.2 indeed shows a relatively monotonic decline of the correlation over time. This implies that banks have become more 'specialized', i.e. in any quarter they appear to enter the market predominantly as lenders or borrowers.

For the number of transactions, Figure 4.3 shows the dynamics of the mean and median in-/out-ntrans (left) and the total ntrans (right). The upper left panel shows that the average number of transactions per bank is close to 200 during most quarters, but significantly decreases during and after the GFC. For both variables, the median values are substantially smaller than the mean, which hints towards a high level of skewness. Again, substantial differences in the median values indicate that the in- and out-variables are unlikely to follow identical distributions. The bottom left panel shows the standardized mean and median values. Quite interestingly, the somewhat negative trend of the variables vanishes, except for the GFC period. The same observation applies to the total number of transactions on the right panels. The results concerning the correlation between in- and out-transactions are comparable to those for the degrees (not reported).

4.4.2 The Degree Distributions

Due to the change in the size of the Italian interbank network, and the detection of two candidates for significant structural breaks during our sample period, we split the data set into three periods: Period 1 covers quarters 1-10, period 2 covers quarters 11-39, and period 3 covers the remaining quarters 40-48.[12] Assuming that the realizations of single days (quarters) are iid draws (or weakly dependent ones) from the same underlying data generating process, allows us to pool the data of the three subperiods into larger samples for the in-, out-, and total degrees (ntrans) of active banks, respectively. We use both daily and quarterly aggregates, i.e. construct variables that count

[12] Note that the first subsample roughly coincides with the data set used by De Masi *et al.* (2006).

the number of unique counterparties (degree) and total number of transactions (ntrans) for each bank within each day and quarter, respectively.[13] For the daily (quarterly) data this amounts to a total of 96,892 (1,780), 188,582 (3,369), and 41,775 (843) pooled observations for the three periods, respectively. For the sake of completeness, we also show the results when pooling all observations for the three time periods (1-3) for each degree measure. We should stress that pooling observations from several periods is crucially necessary in order to obtain reliable parameter estimates, in particular for daily data. We will elaborate on this issue in more detail in the next section.

As a first step, we compare the in- and out-degree distributions and check whether they could be realizations from the same underlying distribution. Figure 4.4 shows the histograms of the in-, out-, and total degrees for the different time-periods using quarterly data. We see that the histograms look very different when comparing in- and out-degrees for each sample period. We should note that a substantial fraction of observations equals zero, both for in- and out-degrees. While the in-degree histograms appear to have a certain hump-shape, the out-degrees look more like a slowly decaying function with monotonic decline of probability from left to right. Furthermore, the L-shaped form of the out-degree distributions appears to be more stable over time, even though the scale on the x-axis changes substantially. Individual Kolmogorov-Smirnov (KS) tests provide further evidence against the equality of in- and out-degree distributions for all sample periods. The KS test allows to check whether two variables follow the same probability distribution, but also whether one variable follows a certain specific distribution. In our case, the KS test statistic is calculated as

$$\mathrm{KS}_n = \sup_x |F_{1,n}(x) - F_{2,n}(x)|, \tag{4.5}$$

where $\sup_x$ denotes the supremum of all possible values, while $F_{1,n}(\cdot)$ and $F_{2,n}(\cdot)$ are the empirical distribution functions of the sample of in-degrees and out-degrees, respectively. At all sensible significance levels, we have to reject the null hypothesis of the equality of both distributions. Similar observations can be made when pooling all observations across the three subperiods, see Figure 4.5.

Figure 4.6 shows the complementary cumulative distribution functions (ccdf) for the quarterly degree measures for all sample periods on a log-log scale, the typical way to represent data when suspecting power-law decay. Note that for a power-law, these ccdfs would be straight lines, which upon inspection seems unlikely to provide a good approximation to any of our

[13] In Appendix 4.6.4 we present a similar analysis for the distribution of transaction volumes of individual institutions.

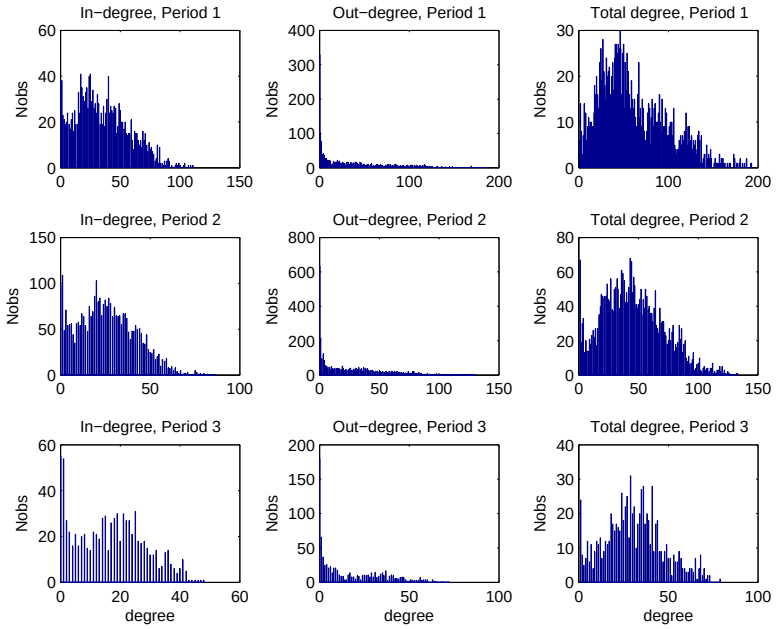

Fig. 4.4: Quarterly data, degree. Histograms for in-degree (left), out-degree (center), and total degree (right) for Period 1, 2, and 3.

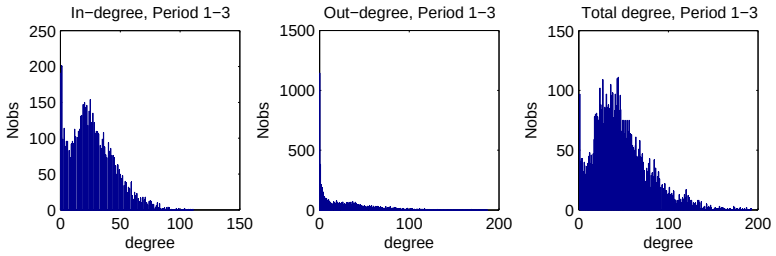

Fig. 4.5: Quarterly data, degree. Histograms for in-degree (left), out-degree (center), and total degree (right) using all observations.

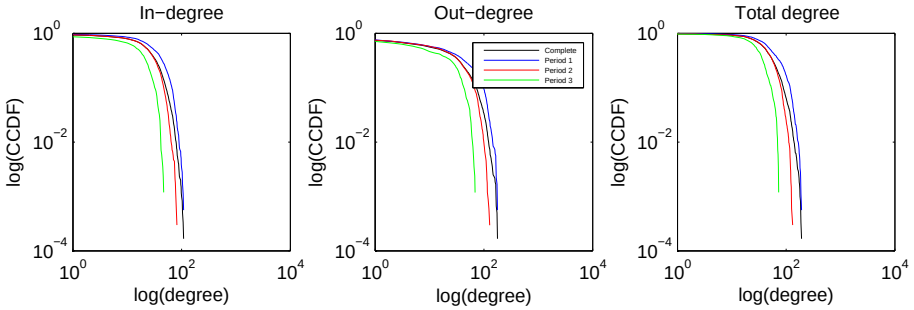

Fig. 4.6: Quarterly data, degree. Complementary cumulative distribution functions (ccdf) in-degree (top), out-degree (center), and total degree (bottom) for all time periods on a log-log scale.

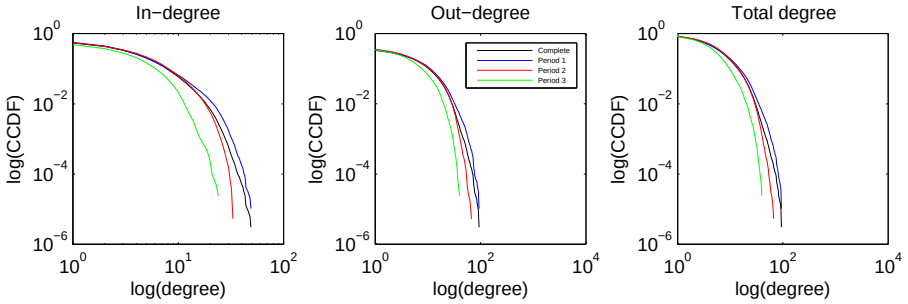

Fig. 4.7: Daily data, degree. Complementary cumulative distribution functions (ccdf) in-degree (top), out-degree (center), and total degree (bottom) for all time periods on a log-log scale.

subsamples, even for the tail regions. Again the distributions of in- and out-degrees look quite different in general, even though the shapes of the tail regions appear to be more homogeneous than what one might have expected after inspection of the raw data in Figures 4.4 and 4.5. Similar arguments hold for the distribution of total degrees, which has a somewhat similar shape as the in-degree distribution. For this reason, we will mostly restrict ourselves to comment on the results for the in- and out-degrees, respectively. We also show the ccdfs for the daily observations in Figure 4.7. Again, it is hard to detect a linear decay for most samples, at least not over several orders of magnitude.

Distribution Fitting Approach

Our basic approach is to fit a number of candidate distributions in order to investigate which distribution describes the data 'best' in a statistical sense. We should note that, similar to the approach in Stumpf and Ingram (2005), we use both discrete and continuous candidate distributions, implying that for the latter we treat the degrees as continuous variables. The candidate distributions, always fitted using maximum likelihood (ML), are:

- the Exponential distribution, with parameter $\lambda > 0$ (rate),
- the Gamma distribution, with parameters $k > 0$ (shape) and $\theta > 0$ (scale),
- the Geometric distribution, with probability parameter $p \in [0, 1]$,
- the Log-normal distribution, with parameters μ (scale) and $\sigma > 0$ (shape),
- the negative Binomial distribution, with parameters $r > 0$ (number of failures) and $p \in [0, 1]$ (success probability),
- the Poisson distribution, with parameter $\lambda > 0$,
- the discrete power-law or Pareto distribution, with parameters $x_m > 0$ (scale) and $\alpha > 0$ (shape),
- the Weibull or stretched exponential distribution, with parameters $\lambda > 0$ (scale) and $k > 0$ (shape).

We should note that a large part of the literature focuses on fitting the power-law only, in particular when the ccdfs have an apparently linear shape. Given that this is not the case here, we test a number of alternative distributions to find the distributions that fit the data best. Nevertheless, even though the power-law might not be a good description of the complete distribution, it could still provide of good fit of the (upper) tail region. Therefore, we conducted two sets of estimations of the above distributions for each sample: first, we fitted the complete distribution using all entries of our samples. Here we should stress, that several of the distributions have strictly positive support, while the others also allow for the occurrence of zero links. For the sake of consistency we will therefore only use non-zero values for the degree

and ntrans variables in the following.[14] This means, for some distribution functions, we are using truncated variables in general (both for the complete and tail observations) and need to adjust the ML estimators for these distributions accordingly, cf. Appendix 4.6.1. In a second step, we explicitly fitted three of the eight candidate distributions, namely the Exponential, the Log-normal, and the power-law, to a certain upper tail region for each period and variable (the other candidates would obviously make little sense as tail distributions). There are different possibilities to identify the 'optimal' tail region and here we employ the approach of Clauset *et al.* (2009), which has been demonstrated to yield reliable estimates of both power-law parameters for certain distributions converging to Paretian tail behavior. The basic idea of this approach is to find the optimal tail parameter for all possible cutoff points using maximum likelihood, where the optimal x_m is the one corresponding to the lowest KS statistic. Details can be found in Appendix 4.6.2.[15] The tail region is then defined by the scale parameter x_m, and the other distributions are fitted to all observations where $x \geq x_m$. Note that this approach gives an obvious advantage to the fit of the power-law in the 'tail' region. Quite surprisingly, however, in many cases the power-law is not the best description of the data tailored in this way as we will see below.

In these goodness-of-fits (GOF) experiments, we first estimate the parameters for each candidate distribution, both for the complete data set and the upper tail region, respectively, using ML. Using these parameters, we calculate the KS test-statistic for each candidate distribution and take the one with the lowest value as the 'best' fit of the respective data.[16] As a last step, we evaluate the GOF of this candidate distribution based on the KS test statistic. Given that the critical values of the KS distribution are only valid for known distributions (i.e. without estimating parameters), we have

[14] This is important, since we cannot replicate the large number of zero values based on these distributions that we observe in the empirical data. Ignoring zeros reduces the number of quarterly observations to 1,742, 3,271, and 788 for the in-variables, and 1,450, 2,733, and 663 for the out-variables, respectively. For the daily data this leaves 70,584, 133,280, and 28,093 for the in-variables, and 39,619, 83,723, and 17,961 for the out-variables, respectively. The number of observations for the total degree and ntrans variables remain unaffected, since only active banks are in the sample.

[15] There exist a number of alternative approaches in statistical extreme value theory for determining the optimal tail size. The approaches by Danielsson *et al.* (2001) and Drees and Kaufmann (1998) yielded results very similar to those reported in the text. We also checked certain fixed thresholds for identifying the tail region. The results remain qualitatively the same as long as the chosen upper quantile is reasonably large.

[16] In principle, we could also use likelihood-based criteria, e.g. AIC or BIC. However, Clauset *et al.* (2009) provide some evidence that the KS statistic is preferable as it is more robust to statistical fluctuations.

to perform individual Monte-Carlo exercises.[17] In these exercises, we randomly sample many degree sequences from the best fitting distributions with their estimated parameter values and then calculate the KS test statistic of these synthetic data sets. The reported p-values count the relative fraction of observations larger than the observed ones, such that low p-values (say 5%) indicate that the pertinent distribution can be rejected. We should stress that we carry out this analysis only for the best fitting distribution, since the remaining ones have already been found to be inferior under the KS criterion. Details on the Monte-Carlo design can be found in Appendix 4.6.3.

In the following we will use this approach to investigate the distribution of degrees and number of transactions for both daily and quarterly aggregates. Already at this point we should stress that the GOF tests mostly indicate that the distributions have to be rejected at traditional levels of significance for the complete samples, while the fits to the tail tend to perform better. This finding is, however, strongly driven by the significantly smaller number of observations for the tail data, which yields relatively large and more volatile KS statistics compared to the complete distributions.

Daily Data

We start our analysis with the daily degree data for which earlier studies have reported power-laws (De Masi *et al.*, 2006, and Iori *et al.*, 2008). Before turning to the results, we need to stress several complicating issues arising from network data in general, and our data in particular. For example, Stumpf and Porter (2012) note that '[a]s a rule of thumb, a candidate power-law should exhibit an approximately linear relationship on a log-log plot over at least two orders of magnitude in both the x and y axes. This criterion rules out many data sets, including just about all biological networks'. In this sense, finite and possibly very small network sizes make it hard to provide evidence for scale-free networks (Avnir *et al.*, 1998, and Clauset *et al.*, 2009).

For our data, Figure 4.8 shows the maximum in- and out-degrees for the individual days over time. We see that the criterion of Stumpf and Porter (2012) is typically violated. Thus, it should be hard to find evidence in favor of the power-law hypothesis for the complete distributions. Additionally, the number of observations in the 'tail' of the data for a single day becomes very small leading to large fluctuations of estimates across days and large error bands of single estimates. These issues highlight the importance of applying rigorous statistical methods to identify the best fitting distributions, i.e. simply identifying a linear slope of the ccdf on a log-log scale might easily

[17] See Clauset *et al.* (2009) and Stumpf *et al.* (2005) for similar approaches.

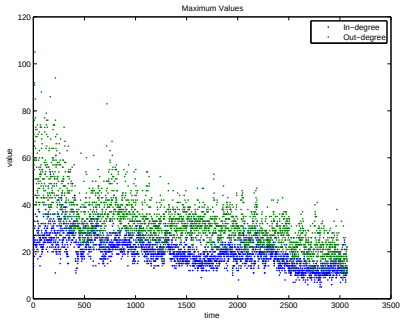

Fig. 4.8: Daily data. Maximum in- and out-degrees over time.

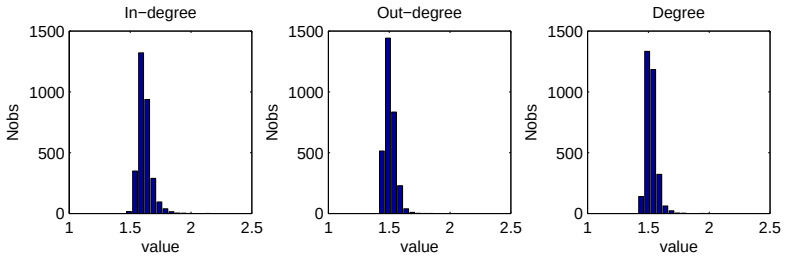

Fig. 4.9: Daily data, degree. Histograms for the power-law exponents for the complete distributions, in-, out- and total degree, respectively.

be misleading. Similar remarks also apply for the daily ntrans variables (see below), while quarterly data are typically slightly less problematic.

To highlight our previous comments, Figures 4.9 and 4.10 show the distribution of the estimated daily power-law parameters for the complete and tail observations, respectively, for all sample days. For the complete daily samples, the results are very stable over time and across types of degrees, cf. Figure 4.9. In fact, we will see that this stability tends to carry over to the complete distributions of the aggregated data as well. In contrast, there is a substantial level of heterogeneity for the power-law exponent of the tail observations for the individual days, cf. Figure 4.10. Thus, we cannot confirm previously reported findings of 'typical' tail parameters between 2 and 3 for any of the degree variables.[18] While numerous observations lie within this range, for many days we find substantially larger values, at times as large as

[18] The results are very similar when focusing on the individual period 1-3 as defined before.

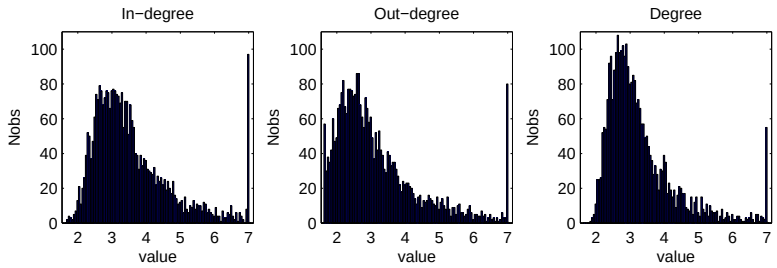

Fig. 4.10: Daily data, degree. Histograms for the power-law exponents for the tail observations, in-, out- and total degree, respectively.

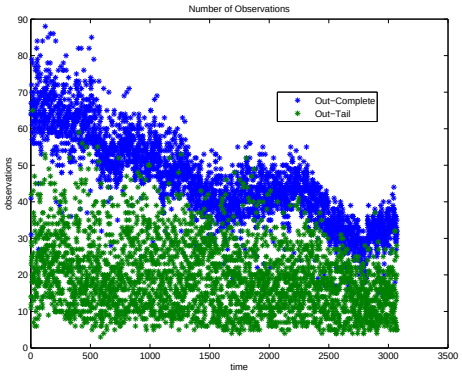

Fig. 4.11: Daily data, degree. Total number of observations (complete) and number of tail observations for out-degree.

7.[19] Apparently, the daily tail data are too noisy to identify a 'typical' tail parameter, cf. Figure 4.11.[20] The mismatch between the narrow range of values obtained for the complete data set of single days and the broad range of estimates for their tail might also indicate that the former are mainly determined by the more central part of the distribution.

Since data for single days are too scarce to allow reliable parameter estimation, pooling observations over longer horizons might be advisable to obtain better estimates. This, of course, requires the assumption of daily data being drawn independently from the same underlying distribution, or only with weak dependence of adjacent observations. While it is not straightforward to check this assumption for complete daily ensembles (as opposed to a time-series of univariate daily data), we have made some attempt at checking for statistical breaks for averages of degree statistics and have cut our complete sample into subsamples accordingly. Note also that any analysis of a network structure would be more or less futile, if we could not assume some stationarity of the structural characteristics of the network. Fricke and Lux (2012) demonstrate that the e-MID network is indeed structurally stable along many dimensions.

We report our estimation results for the pooled daily data in Tables 4.1-4.3. Our main finding is that the negative Binomial distribution provides the best fits (in bold) for all daily degree measures and for all samples (i.e. the complete samples and the three subsamples identified via tests for structural breaks), cf. Table 4.1. The results from the GOF experiments indicate, however, that the best fitting candidate distributions have to be rejected. Therefore, even the winner among the candidate distributions appears to be an unlikely description of the data. We should also stress that the fit of the power-law is usually rather poor, competing with the Poisson distribution for the worst description of the data. Similar to the findings for the individual days, the estimated tail parameters are between 1.5 and 1.6, cf. Table 4.2 (top, complete). Figure 4.7 together with the relatively poor KS statistics for estimated power-laws suggests that estimates in the scaling range 1-2 are

[19] We have set 7 as the upper bound of the power-law parameter in our numerical ML implementation. For larger values the evaluation of the zeta function appearing in the discrete Pareto law, cf. Appendix 4.6.2, is not accurate enough to obtain reliable estimates. The fact that the estimated values hit the upper bound quite frequently indicates that the estimated values may become even larger when increasing the upper bound.

[20] We also generated synthetic power-law distributed random draws and estimated their scaling parameters based on the algorithm for the selection of the tail region detailed above (not reported). For the small sample sizes of the typical daily data, the tail parameter of these synthetic data is highly volatile as well, even though the very large values observed for the actual data are very rare. As usual, however, increasing the number of observations (say more than 500), typically yields estimates very close to the true parameters.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Exponential	.0465	.0627	.0448	.0498	.0789	.0911	.0777	.0447	.0488	.0764	.0374	.0406
Gamma	.0627	.0670	.0661	.0870	.0515	.0511	.0536	.0559	.0562	.0637	.0512	.0592
Geometric	.0132	.0250	.0129	.0299	.0608	.0759	.0595	.0224	.0214	.0510	.0127	.0289
Log-normal	.0814	.0748	.0816	.1001	.0725	.0722	.0736	.0746	.0631	.0605	.0641	.0701
Neg. Bin.	**.0063**	**.0177**	**.0082**	**.0224**	**.0114**	**.0160**	**.0138**	**.0105**	**.0153**	**.0208**	**.0115**	**.0103**
Poisson	.2313	.2409	.2347	.1715	.3500	.3774	.3476	.2678	.2973	.3318	.2892	.2087
Power-law	.2099	.2151	.2107	.1985	.2077	.2024	.2079	.2140	.2366	.2219	.2427	.2373
Weibull	.0591	.0630	.0646	.0872	.0547	.0552	.0574	.0555	.0522	.0575	.0481	.0581

Tab. 4.1: Daily data, degree. KS statistic for the candidate distributions (complete). Minimum values in bold indicate the best fitting distribution. Asterisks would indicate non-rejection of this distribution at the 5% confidence level, where the critical values were obtained from a Monte-Carlo exercise as described in the main text. There is, however, no such case in this table.

obtained as very inaccurate straight lines fitted to a strongly curved distributional shape. Moving to the tail observations, we find that exponential and power-law distributions tend to provide the best fit for all variables, cf. Table 4.3. Thus, it appears that the power-law is a better description of the tail observations - a usual finding for many data sets. In contrast to the complete distributions, the GOF experiments suggest that the estimated distributions are mostly not rejected for the tail observations.[21] Upon closer inspection, however, we see the KS statistics of the exponential and the power-law are typically close to each other, in particular when the tail exponents are very large, cf. Table 4.2 (top, complete). Even though the power-law appears to provide the best fit for some of the tail data, the very large parameter values (larger than 4, often close to 7) are in a range where the power-law becomes almost undistinguishable from exponential decay. Often such high values would be obtained spuriously from distributions with an exponential decline as semi-parametric estimators of the tail index would not be able to 'identify' the limit of $\alpha \to \infty$. The huge difference in estimated power-law parameters for the complete sample compared to the tail also indicates that the empirical distribution shows pronounced curvature (actually confirming the visual inspection of absence of a linear slope over the complete support and very fast decline at the end in Figure 4.7). On the other hand, it is also interesting to remark that the estimated coefficients are relatively uniform for both the complete sample and the tail, respectively, across periods and

[21] This result is driven by the higher noise level in the tail data due to a smaller number of observations compared to the complete distributions.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Daily												
Complete	1.61	1.61	1.60	1.67	1.50	1.48	1.50	1.54	1.52	1.51	1.51	1.57
	(.001)	(.003)	(.002)	(.005)	(.002)	(.003)	(.002)	(.005)	(.001)	(.002)	(.001)	(.003)
Tail	7.00	7.00	7.00	7.00	5.93	4.43	7.00	5.53	6.03	4.70	7.00	7.00
	(.175)	(.300)	(.160)	(.260)	(.110)	(.078)	(.170)	(.161)	(.100)	(.071)	(.146)	(.393)
Quarterly												
Complete	1.28	1.26	1.28	1.33	1.29	1.28	1.29	1.33	1.24	1.23	1.24	1.27
	(.004)	(.006)	(.005)	(.001)	(.004)	(.008)	(.006)	(.001)	(.003)	(.006)	(.004)	(.009)
Tail	5.13	7.00	7.00	4.63	7.00	6.90	7.00	4.82	5.20	6.90	7.00	5.01
	(.134)	(.460)	(.325)	(.233)	(.482)	(.532)	(.412)	(.306)	(.145)	(.421)	(.330)	(.261)

Tab. 4.2: Power-law parameters and standard deviations, degree. Values obtained via numerical maximization of the log-likelihood for discrete data. Standard deviations (in parentheses) approximated as $(\alpha-1)/\sqrt{T}$, with T being the number of observations. Top: daily data, bottom: quarterly data.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Exponential	**.0357***	**.0354***	**.0355***	.0954	.0642	.0580	**.0353***	**.0388***	.0685	.0637	**.0300***	.0457
Log-normal	.0664	.0639	.0762	.1036	.0479	.0479	.0541	.0771	.0484	.0455	.0521	.0927
Power-law	.0372	.0376	.0400	**.0203***	**.0129***	**.0305***	.0392	.0455	**.0114***	**.0192***	.0352	**.0382***

Tab. 4.3: Daily data, degree. KS statistic for the candidate distributions (tail). Minimum values in bold indicate the best fitting distribution. Asterisks indicate non-rejection of this distribution at the 5% confidence level, where the critical values were obtained from a Monte-Carlo exercise as described in the main text.

for all the measures of degree. This speaks of relatively uniform shapes of the distributions, at least in view of this simple statistic. Summing up, the power-law distribution appears to be a poor description of the data, both for the complete distribution and the tail observations (where it more or less coincides with an exponential for the high estimates of the tail index). We also need to stress that the identified power-law exponents, both for individual days and pooled observations, are far off from those reported in earlier studies. It is not clear how these estimates were obtained.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Exponential	.1474	.1544	.1661	.1520	.0797	.0932	**.0710**	.0925	.1740	.1675	.1887	.2171
Gamma	.0573	.0543	.0723	.0942	**.0595**	**.0514**	.0738	.0961	.0414	**.0284***	.0673	.0943
Geometric	.1533	.1586	.1723	.1619	.0778	.0918	.0728	.1009	.1771	.1696	.1920	.2223
Log-normal	.1141	.0972	.1274	.1377	.1063	.0984	.1164	.1226	.0972	.0760	.1185	.1453
Neg. Bin.	.0601	.0580	.0729	.1025	.0708	.0615	.0836	.1081	.0395	.0318	.0627	.0881
Poisson	.3117	.3462	.2753	.2561	.4367	.4707	.4116	.4183	.3601	.4115	.3183	.2489
Power-law	.3828	.4023	.3849	.3522	.2727	.2728	.2842	.2608	.4376	.4546	.4387	.4291
Weibull	**.0380**	**.0342***	**.0456**	**.0689**	.0624	.0609	.0736	**.0912**	**.0246**	.0361	**.0325**	**.0495***

Tab. 4.4: Quarterly data, degree. KS statistic for the candidate distributions (complete). Minimum values in bold indicate the best fitting distribution. Asterisks indicate non-rejection of this distribution at the 5% confidence level, where the critical values were obtained from a Monte-Carlo exercise as described in the main text.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Exponential	**.0248***	.0466	**.0325***	**.0628***	**.0395***	.0394	**.0352***	**.0459***	**.0331***	**.0315***	**.0437***	**.0530***
Log-normal	.0379	**.0441***	.0431	.0766	.0526	.0533	.0663	.0756	.0451	.0494	.0769	.0794
Power-law	.0651	.0748	.0471	.0949	.0502	**.0384***	.0559	.0918	.0515	.0405	.0602	.0778

Tab. 4.5: Quarterly data, degree. KS statistic for the candidate distributions (tail). Minimum values in bold indicate the best fitting distribution. Asterisks indicate non-rejection of this distribution at the 5% confidence level, where the critical values were obtained from a Monte-Carlo exercise as described in the main text.

Quarterly Data

The results for the quarterly data are shown in Tables 4.4 and 4.5. Weibull distributions typically provide the best fits for the in- and total degrees, while Exponential and Gamma distributions yield comparable fits as the Weibull for the out-degrees. Similar to the complete distributions for the daily data, the optimal fits are insignificant, except for three cases. Again, the best fits appear to be unlikely descriptions of the observed data. The power-law exponents for the complete sample are again quite small, typically between 1.25 and 1.3, cf. Table 4.2 (bottom, complete). Turning to the tail observations, we find that Exponential distributions provide the best fits in all but two cases (in-degree in period 1, out-degree period 1). Similar to the daily observations, for the best fits of the tail data, the pertinent distributions cannot be rejected as the 'true' data-generating process at the

95 percent significance level. The poor fit of the power-law again comes along with relatively large tail exponents, cf. Table 4.2 (bottom, tail). In summary, similar to the daily data, we do not find evidence in favor of scale-free networks.

Robustness and Discussion

A reason for not finding evidence for power-law distributions may be the fact that we focus on the subnetwork formed by Italian banks only. Stumpf *et al.* (2005) have shown that (randomly chosen) subnetworks of scale-free networks are in fact not scale-free. Therefore we also checked the distributions including foreign banks as well, similar to the existing papers using the e-MID data. We found that the results (including the tail parameters) remain qualitatively unaffected (not reported). In terms of Stumpf *et al.* (2005), these findings indicate again that the networks including all banks are unlikely to be scale-free, and that our previous findings for Italian banks alone are not biased due to random sampling from a larger scale-free network (indeed, it seems very unlikely that the Italian banks should constitute a random set from the overall sample of all banks). Then it comes as no surprise that the subnetworks formed by Italian banks only are also not scale-free. In fact, it is remarkable that there appears to be no significant qualitative effect of incorporating foreign banks or not.

In Finger *et al.* (2012) it has been shown that the quarterly e-MID networks are more complete representations of the underlying 'latent' network structure, whereas daily networks might be seen as random activations of parts of the more complex, hidden structure. Under this perspective, the lack of coincidence of the fitted distributions for different levels of time aggregation might not be too surprising.

Summing up, our results indicate that the power-law hypothesis needs to be tested more thoroughly for other networks in general and the interbank network in particular,[22] with the power-law being one of many candidate distributions. The findings are in line with other studies casting doubts on certain claims of power-law and scaling behavior in a broad range of empirical studies (cf. Avnir *et al.*, 1998).

4.4.3 The Distribution of the Number of Transactions

Note that the quarterly degree of a given bank is not the simple sum of its daily degrees, since a link that has been activated many times over a quarter, is counted only as one link on the quarterly level of network activity.

[22] See Stumpf and Porter (2012).

If we consider the number of links in the daily data as (possibly power-law distributed) random variables, the number of transactions over a longer time horizon is, in fact, what we obtain from simple aggregation of the daily degrees observed for any bank i. Assuming that the degrees of all banks are drawn from the same distribution, we obtain in this way a sample of sums of random variables following the same underlying distribution. Note that we would expect a power-law at the daily level to survive in the aggregation process for an iid random process of link formation as well as for various extensions allowing for 'weak' dependency.[23] The extremal behavior of the distribution of degrees should, therefore, be preserved in the distribution of the number of transactions over longer horizons. We turn to the analysis of this quantity in this subsection.

Note also that the finite size of the network might pose a problem due to the effective imposition of an upper limit on the observable degrees. It might, therefore, be the case that a scale-free distribution is just hard to verify because of the small number of observations. In contrast, the aggregated ntrans variables have the advantage that they have no obvious upper bound, so testing the power-law hypothesis might be more sensible in this case.

Figures 4.12 and 4.13 show the ccdfs of the quarterly and daily ntrans variables. Again, linear decay over several orders of magnitude is hard to detect visually. However, at least for the quarterly data we see that the variables under study span several orders of magnitude, making the data more useful candidates for our distribution fitting approach.

	In				**Out**				**Total**			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Quarterly												
Complete	1.19	1.19	1.19	1.21	1.21	1.22	1.21	1.23	1.16	1.16	1.16	1.17
	(.003)	(.005)	(.003)	(.008)	(.003)	(.006)	(.004)	(.009)	(.002)	(.004)	(.003)	(.006)
Tail	3.48	2.78	3.54	3.06	3.16	2.99	2.21	2.95	2.77	3.03	2.76	3.68
	(.091)	(.082)	(.116)	(.136)	(.094)	(.143)	(.044)	(.180)	(.040)	(.106)	(.050)	(.200)

Tab. 4.6: Power-law parameters and standard deviations, ntrans. Values obtained via numerical maximization of the log-likelihood for discrete data. Standard deviations (in parentheses) approximated as $(\alpha-1)/\sqrt{T}$, with T being the number of observations. Quarterly data.

For daily data, the range of the observed variables remains rather limited,

[23] The stability under aggregation of power-laws characterizing the tails of iid random variables is one of the basic tenets of the statistical theory of extremes, cf. Reiss and Thomas (2007). In this sense, summing up daily power-law networks should preserve the tail index for different frequencies.

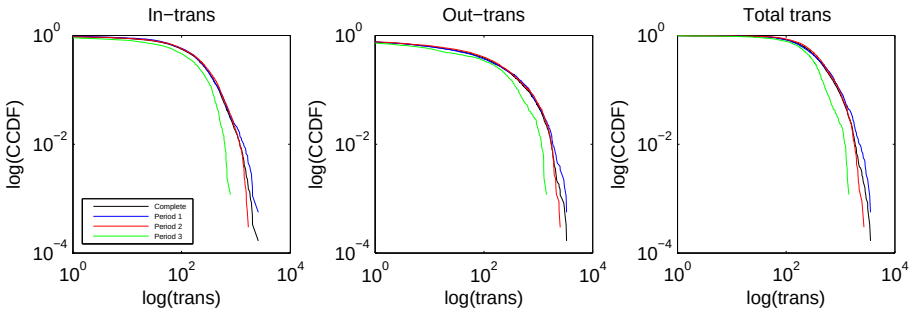

Fig. 4.12: Quarterly data, ntrans. Complementary cumulative distribution functions (ccdf) in-trans (top), out-trans (center), and total trans (bottom) for all time periods on a log-log scale.

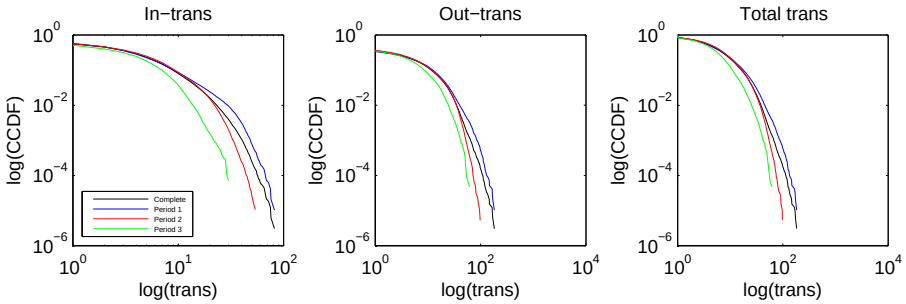

Fig. 4.13: Daily data, ntrans. Complementary cumulative distribution functions (ccdf) in-trans (top), out-trans (center), and total trans (bottom) for all time periods on a log-log scale.

even though the maximum value is roughly twice the one for the degrees (not reported). Since for daily realizations of the number of transactions we find virtually identical results to those of the daily degrees, we abstain from presenting these here, and immediately turn to quarterly aggregated observations.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Exponential	.0452	.0516	.0467	.0748	.2334	.2598	.2131	.2716	.0612	.0524	.0761	.0955
Gamma	**.0153**	.0287	.0199	**.0388**	**.0167**	**.0165**	**.0185**	**.0498**	.0356	.0485	.0325	.0549
Geometric	.0450	.0511	.0465	.0744	.2331	.2595	.2129	.2713	.0614	.0525	.0763	.0960
Log-normal	.0801	.0786	.0843	.0953	.0712	.0610	.0725	.1013	.0701	.0508	.0738	.1207
Neg. Bin.	.0205	.0280	.0249	.0419	.0205	.0188	.0223	.0586	**.0353**	**.0477**	**.0321**	.0543
Poisson	.5744	.5859	.5775	.5527	.6706	.6860	.6710	.6221	.5896	.6164	.5813	.5216
Power-law	.3610	.3865	.3658	.3384	.2381	.2397	.2607	.2209	.4497	.4488	.4648	.4236
Weibull	.0185	**.0244**	**.0186**	.0420	.0383	.0389	.0392	.0668	.0391	.0481	.0386	**.0437***

Tab. 4.7: Quarterly data, ntrans. KS statistic for the candidate distributions (complete). Minimum values in bold indicate the best fitting distribution. Asterisks indicate non-rejection of this distribution at the 5% confidence level, where the critical values were obtained from a Monte-Carlo exercise as described in the main text.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Exponential	.0684	.0874	.0709	.0687	**.0393***	.0805	.0532	.0751	.0528	.0616	.0492	.0948
Log-normal	**.0258***	**.0202***	**.0441***	**.0560***	.0408	**.0594***	**.0495**	**.0700***	**.0274***	**.0336***	**.0362***	.0579
Power-law	.0476	.0570	.0520	.1219	.0988	.0741	.1024	.0718	.0676	.0604	.0683	**.0426***

Tab. 4.8: Quarterly data, ntrans. KS statistic for the candidate distributions (tail). Minimum values in bold indicate the best fitting distribution. Asterisks indicate non-rejection of this distribution at the 5% confidence level, where the critical values were obtained from a Monte-Carlo exercise as described in the main text.

We show the results in Tables 4.7 and 4.8, finding that negative Binomial, Gamma and Weibull distributions appear among the best fits, depending on the concept (in-, out-, or total transactions) and the period considered. However, their KS statistics are typically at a comparable level. The results from the GOF experiments show that the best fitting distributions are nevertheless rejected as data-generating processes (exception: total ntrans in period 3). Again, the fit of the power-law is very bad in general, with tail parameters around 1.20, cf. Table 4.6, and KS statistics that consistently come in

second to last (with the Poisson distribution performing worst). Moving to the quarterly tail data, we find that in most cases the Log-normal provides the best fit (exceptions: out-degree for the complete sample and total degree in period 3). This is quite surprising, given that the scaling parameters now lie in the 'typical' range for power-laws, here between 2.21 and 3.68. Therefore, even though the power-law estimates appear more sensible, it is inferior by some margin in fitting the tail data (with a cut-off determined by the best-fitting Pareto law) to the Log-normal, and sometimes also to the Exponential. As with the previous cases, the results from the GOF experiments indicate that the best fitting tail distributions usually cannot be rejected via KS tests with Monte-Carlo distributions. While it is well-known that it is hard to distinguish Log-normal from power-law tails, these findings raise doubts on the universality of power-law tails and highlight the need for thorough statistical approaches of testing the power-law hypothesis.

As another robustness check, we investigated the distribution of transaction volumes (tvol), cf. Appendix 4.6.4, again differentiating between in-tvol, out-tvol and their sum (total tvol), respectively. While the tails of the tvol variables are typically much fatter compared to the degree and ntrans variables, the power-law remains a poor description both for daily and quarterly data.

4.5 Conclusions

In this paper, we have revisited the distributional properties of interbank loans for the Italian interbank network during the years 1999-2010. Using both the degrees and the number of transactions, we fitted a set of different candidate distributions to these data for daily and quarterly aggregates, respectively. Given that the daily networks have previously been claimed to be scale-free (De Masi *et al.*, 2006), it comes as a surprise that we find no evidence in favor of the power-law hypothesis: at the daily level the degrees are usually fit best by negative Binomial distributions, while the tails tend to decay exponentially, i.e. the fitted power-laws display very large tail parameters. At the quarterly level, Weibull, Gamma, and Exponential distributions tend to provide comparable fits for the complete degree distribution, while the tails again tend to display exponential decay. For the number of transactions we find comparable results, even though the tails of the quarterly data appear to be fatter. However, in this case the Log-normal distribution usually outperforms the power-law. Moreover, we found that the networks are characterized by a substantial level of asymmetry, as exemplified by the low correlation between in- and out-degrees. We also find that the two variables

do not follow identical distributions in general.

Overall these findings indicate that the power-law is typically a poor description of the data. This implies that preferential attachment and related mechanisms (see e.g. De Masi *et al.*, 2006), are unlikely explanations for the formation of the Italian interbank network. Note that these findings are also not in line with a large part of the empirical (interbank) network literature for other data sets, putting doubts on the universality of scale-free interbank networks. Our results also indicate that the power-law hypothesis needs to be tested more thoroughly for other networks in general and the interbank network in particular. The findings are related to other studies casting doubts on certain claims of power-law and scaling behavior in a broad range of empirical studies (cf. Avnir *et al.*, 1998, and Stumpf and Porter, 2012), and it seems possible that claims of scale-free behavior of interbank lending activity may not survive under closer statistical scrutiny.

4.6 Appendix

4.6.1 Truncated Distributions and Maximum Likelihood

The distribution fitting approach described in the main text involves fitting a set of candidate distributions with possibly differing support. For example, some distributions have support at zero, while others do not. Similarly, when focusing on the tail observations we have to get rid of the probability mass below the cutoff point in order to accurately calculate the statistics. Therefore, we describe the use of truncated distributions and ML fitting in this Appendix in more detail.

Normalization

When working with truncated variables, we need to make sure to use the correct pdfs and cdfs, since the ML estimation and the evaluation of the fit (KS statistic) depend on them. In order to illustrate this issue, let variable x have the pdf $p(x)$ with support $[0, \infty]$. As usual, the cdf is defined as

$$P(a) = P(X \leq a) = \int_0^a p(x)dx. \tag{4.6}$$

Now, suppose the data are (left-)truncated at some value x_m, i.e. the variable $\tilde{x}$ follows the same distribution as x, but the pdf has limited support $[x_m, \infty]$ with minimum value $x_m > 0$. For our purposes, it is therefore useful to define the quantity

$$P_<(a) = P(X < a) = 1 - \int_a^\infty p(x)dx, \tag{4.7}$$

or more compactly

$$P_<(a) = P(a) - p(a). \tag{4.8}$$

We can properly construct the pdf of $\tilde{x}$, say $\tilde{p}$, as

$$\tilde{p}(x) = \begin{cases} \frac{p(x)}{1-P_<(x_m)}, & \text{if } x \geq x_m \\ 0, & \text{else.} \end{cases} \tag{4.9}$$

where the denominator distributes the probability mass of $p(x)$ among the support of $\tilde{x}$.

For the calculation of the KS statistics, we also need the adjusted cdf. For the supported values of $\tilde{x}$ it takes the form

$$\tilde{P}(x) = \int_{x_m}^x \frac{p(x)}{1 - P_<(x_m)}dx = \frac{1}{1 - P_<(x_m)} \int_{x_m}^x p(x)dx, \tag{4.10}$$

or

$$\tilde{P}(x) = \frac{P(x) - P_<(x_m)}{1 - P_<(x_m)}, \tag{4.11}$$

which can be easily evaluated.

Maximum Likelihood for Truncated Variables

Using the previous definitions, we can show that the ML estimator for left-truncated variables does not coincide with the standard estimator. The standard ML estimator, i.e. using a sample of n observations of x and denoting by θ the vector of parameters, can be written as

$$L(\theta|x_1, \cdots, x_n) = p(x_1, \cdots, x_n|\theta) = \prod_i^n p(x_i|\theta), \tag{4.12}$$

or in logs

$$\ln(L) = \sum_i^n \ln[p(x_i|\theta)]. \tag{4.13}$$

Using the definitions from above, we can show that the ML estimator for left-truncated variables differs from the one in Eq. (4.13). Using Eq. (4.9), we can write the likelihood as

$$L = \prod_i^{\tilde{n}} \tilde{p}(x_i|\theta) = \prod_i^{\tilde{n}} \frac{p(x|\theta)}{1 - P_<(x_m|\theta)}, \tag{4.14}$$

where x ignores those observations smaller than x_m and the total number of observations is $\tilde{n}$ instead of n. Taking logarithms we obtain

$$\ln(L) = \sum_i^{\tilde{n}} \ln\left[\frac{p(x_i|\theta)}{1 - P_<(x_m|\theta)}\right] = \sum_i^{\tilde{n}} \ln[p(x_i|\theta)] - \sum_i^{\tilde{n}} \ln[1 - P_<(x_m|\theta)], \tag{4.15}$$

which can be written as

$$\ln(L) = -\tilde{n} \ln[1 - P_<(x_m|\theta)] + \sum_i^{\tilde{n}} \ln[p(x_i|\theta)]. \tag{4.16}$$

The second part of this Eq. looks familiar, as it corresponds to Eq. (4.13) for the $\tilde{n}$ observations with values $\geq x_m$. However, the normalization term on the left does not vanish (as it depends on the parameter vector) and affects the location of the maximum likelihood estimator. Therefore, we need to find the θ that maximizes Eq. (4.16). The standard ML estimator would not be efficient.

4.6.2 Discrete Power-laws and Parameter Estimation

This presentation is mostly based on Clauset *et al.* (2009).

Discrete Power-laws

A power-law distributed variable x obeys the pdf

$$p(x) \propto x^{-\alpha}, \tag{4.17}$$

where $\alpha > 0$ is the tail exponent with 'typical' interesting values in the range between 1 and 3. In many cases, however, the power-law only applies for some (upper) tail region, defined by the minimum value x_m. While it is common to approximate discrete power-laws by the (simpler) continuous version, for our (integer-valued) data, we employ the more accurate discrete version in the paper.[24]

In the discrete case, the cdf of the power-law can be written as

$$P(x) = \frac{\zeta(\alpha, x)}{\zeta(\alpha, x_m)}, \tag{4.18}$$

where

$$\zeta(\alpha, x_m) = \sum_{n=0}^{\infty} (n + x_m)^{-\alpha} \tag{4.19}$$

is the generalized or Hurwitz zeta function.

Estimation of α and x_m

For a given lower bound x_m, the ML estimator of α can be found by direct numerical maximization of the log-likelihood function

$$\mathcal{L}(\alpha) = -n \ln[\zeta(\alpha, x_m)] - \alpha \sum_{i=1}^{n} \ln[x_i], \tag{4.20}$$

where n is the number of observations.[25] For simplicity, we approximate the standard error of the estimated $\hat{\alpha}$ (for $\hat{\alpha} > 1$) using the closed-form

[24] Clauset *et al.* (2007) show that this is necessary for data sets from the social sciences, where the maximum value is usually only a few orders of magnitude larger than the minimum, i.e. the tail is heavy but rather short. In such cases the estimated exponents can be biased severely when using the continuous approximation.

[25] Using a quadratic approximation of the log-likelihood at its maximum, Clauset *et al.* (2009) also derive an approximate closed-form solution for the estimate of $\alpha \simeq 1 + n/\left(\sum_{i=1}^{n} \ln\left[\frac{x_i}{x_m - 0.5}\right]\right)$. This can be seen as an adjusted Hill-estimator, see Hill (1975). While we always report the exact ML estimator, we checked that the approximation is typically not too bad.

solution based on continuous data.[26] Neglecting higher-order terms, this can be calculated as

$$\sigma = \frac{\hat{\alpha} - 1}{\sqrt{n}}. \tag{4.21}$$

However, the equations assume that x_m is known in order to obtain an accurate estimate of α.[27] When the data span only a few orders of magnitude, as usual in many social or complex systems, an underpopulated tail would come along with little statistical power. Therefore, we employ the numerical method proposed by Clauset *et al.* (2007) for selecting the x_m that yields the best power-law model for the data. To be precise, for each x_m over some reasonable range, we first estimate the scaling parameter using Eq. (4.20) and calculate the corresponding KS statistic between the fitted data and the theoretical distribution with the estimated parameters. The reported x_m and α are those that minimize the KS statistic, i.e. minimize the distance between the observed and fitted probability distribution. According to Clauset *et al.* (2007; 2009), minimizing the KS statistic is generally superior to other distance measures, e.g. likelihood-based measures such as AIC or BIC.

4.6.3 Goodness-of-Fit Test for the Estimated Distributions

Since the distribution of the KS statistics is unknown for the comparison between an empirical subsample and a hypothetical distribution with estimated parameters, we carry out a Monte Carlo approach. We sample synthetic data sets from the estimated distribution, compute the distribution of KS statistics and compare the results to the observed value for the original data set. If the KS statistic of the empirical data set is beyond the α percent quantile of the Monte Carlo distribution of KS values, we reject the pertinent distribution at the $1 - \alpha$ level of significance. In our results, we indicate significant fits at the 5% confidence level using asterisks. We should stress that we carry out this (very time-consuming) GOF experiment only for the distribution with the minimum KS statistic for each sample and variable, respectively. This can be justified by the fact that, even though other candidate distributions may not be rejected as well, they are clearly inferior to the optimal distribution in terms of the KS statistic.

[26] Clauset *et al.* (2009) also derive an (approximate) estimator for the standard error based on discrete data, which is, however, much harder to evaluate as it involves derivatives of the generalized zeta function.

[27] See Clauset *et al.* (2007; 2009) for an extensive discussion.

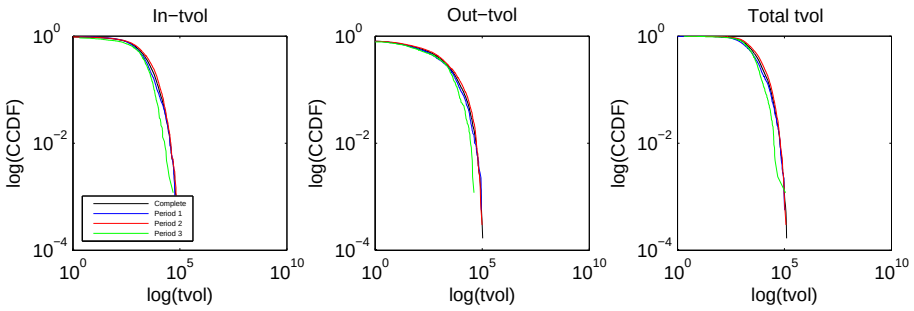

Fig. 4.14: Quarterly data, tvol. Complementary cumulative distribution functions (ccdf) in-tvol (top), out-tvol (center), and total tvol (bottom) for all time periods on a log-log scale.

4.6.4 Distributional Properties of Transaction Volumes

Here we report the results using another important measure for interbank networks, namely the transaction volumes (tvol). We use the same distribution fitting approach as before, differentiating between in-tvol, out-tvol and their sum (total tvol), respectively. Figures 4.14 and 4.15 show the ccdfs on a for the quarterly and daily variables on a log-log scale. We should stress that the minimum trade size on the e-MID market is 50,000 Euros. In order to run our estimation procedure in a reasonable amount of time, we rescale the tvol variables by a factor of 10^{-6} such that a transaction size of 50,000 is represented by a value of 0.05.[28] We then round the tvol variable towards the nearest integer (otherwise the discrete candidate distributions could not be accurately evaluated), again ignoring zero values. In this way, we restrict our samples to relatively large transaction volumes with at least 500,000 Euros, represented by positive integer values. Note that, besides the upward bias of the data and the fact that the data now span several orders of magnitude, it is again hard to visually detect linear decay over several orders of magnitude in the ccdfs. We should also stress that we did not perform the GOF exercise for the tvol variables, since it is too time-consuming in this case.

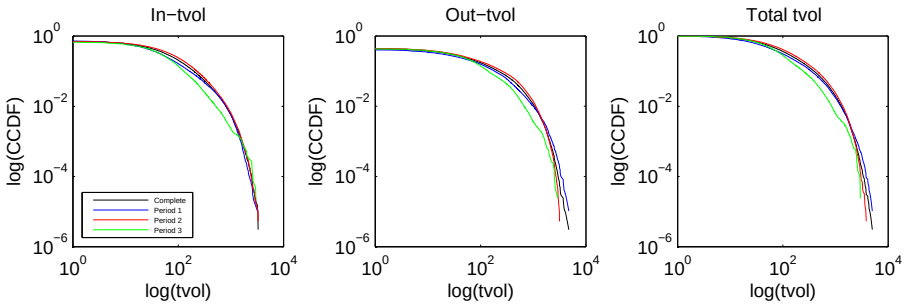

Fig. 4.15: Daily data, tvol. Complementary cumulative distribution functions (ccdf) in-tvol (top), out-tvol (center), and total tvol (bottom) for all time periods on a log-log scale.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Exponential	.1485	.1868	.1252	.1141	.1926	.1981	.1986	.1294	.1737	.2118	.1578	.1242
Gamma	.0765	.0972	.0675	.0821	.0685	.0566	.0690	.0834	.0833	.0876	.0766	.0942
Geometric	.1471	.1851	.1241	.1120	.1919	.1975	.1980	.1281	.1729	.2108	.1570	.1226
Log-normal	**.0274**	**.0253**	**.0276**	**.0340**	**.0319**	**.0344**	**.0310**	**.0357**	**.0231**	**.0248**	**.0224**	**.0259**
Neg. Bin.	.0686	.0876	.0604	.0753	.0599	.0489	.0598	.0766	.0751	.0779	.0685	.0880
Poisson	.6552	.6878	.6423	.6253	.6780	.6785	.6728	.6618	.6803	.6971	.6715	.6527
Power-law	.3587	.3324	.3643	.3836	.3416	.3320	.3490	.3774	.3678	.3406	.3693	.3897
Weibull	.0499	.0648	.0442	.0602	.0427	.0355	.0482	.0590	.0536	.0561	.0496	.0698

Tab. 4.9: Daily data, tvol. KS statistic for the candidate distributions (complete). Minimum values in bold. Significance tests not carried out.

Daily Data

Tables 4.9-4.11 show the results for the daily data. The complete distributions are now usually fitted best by Log-normal distributions, whereas the fit of the power-law is very poor in general. The power-law parameters are again very small, with typical values around 1.22, cf. Table 4.10 (top, complete). For the tail observations, the best fit again is always provided by Log-normal distributions, cf. Table 4.11. Interestingly, the tail exponents of the daily data are within the typical range of meaningful power-laws, cf. Table 4.10 (top, tail), but the power-law is still not the best description of

[28] Note that the maximum daily (quarterly) transaction volumes were 3.75bn (113.46bn) Euros for in-tvol, 4.96bn (111.93bn) Euros for out-tvol and 5.32bn (146.06bn) Euros for total tvol, respectively. For such huge numbers, the estimation procedure, in the numerical optimization for the power-law parameters, tends to take a very long computation time. Therefore, the results in this section should be treated with care, since the rescaling might affect our statistical analysis.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Daily												
Complete	1.23	1.24	1.22	1.24	1.21	1.22	1.21	1.22	1.22	1.23	1.21	1.22
	(.001)	(.001)	(.001)	(.001)	(.001)	(.001)	(.001)	(.002)	(.001)	(.001)	(.001)	(.001)
Tail	2.66	2.15	2.72	2.86	3.23	2.63	4.88	3.39	3.33	3.27	3.35	3.38
	(.010)	(.009)	(.013)	(.042)	(.022)	(.023)	(.239)	(.079)	(.019)	(.036)	(.023)	(.069)
Quarterly												
Complete	1.13	1.15	1.13	1.16	1.14	1.15	1.14	1.15	1.12	1.12	1.14	1.14
	(.002)	(.004)	(.002)	(.006)	(.002)	(.004)	(.003)	(.006)	(.002)	(.003)	(.002)	(.005)
Tail	2.53	2.31	2.55	2.81	3.37	1.97	3.43	2.11	2.02	2.02	3.36	2.46
	(.050)	(.074)	(.060)	(.155)	(.346)	(.050)	(.284)	(.077)	(.020)	(.040)	(.315)	(.091)

Tab. 4.10: Power-law parameters and standard deviations, tvol. Values obtained via numerical maximization of the log-likelihood for discrete data. Standard deviations (in parentheses) approximated as $(\alpha - 1)/\sqrt{T}$, with T being the number of observations. Top: daily data, bottom: quarterly data.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Exponential	.0752	.1339	.0720	.1270	.0306	.0851	.0442	.0887	.0315	.0547	.0201	.0782
Log-normal	**.0194**	**.0235**	**.0175**	**.0169**	**.0141**	**.0170**	**.0425**	**.0245**	**.0133**	**.0162**	**.0170**	**.0279**
Power-law	.0516	.0536	.0526	.0226	.0587	.0466	.0627	.0301	.0557	.0455	.0645	.0351

Tab. 4.11: Daily data, tvol. KS statistic for the candidate distributions (tail). Minimum values in bold. Significance tests not carried out.

the data. In the end, for the transaction volumes we find no evidence in favor of power-laws.

Quarterly Data

Tables 4.12 and 4.13 show the results for the quarterly data. The complete in-, out-, and total degree distributions are now fit best by Weibull, negative Binomial, and Log-normal distributions, respectively. In many cases, these distributions yield comparable KS statistics, but the clear advantage of the Log-normal distribution for the daily data does not carry over to the quarterly level in all cases. Similar to the daily estimates, the power-law parameters are within the usual range of empirical power-laws. As before, however, the tails are best described by Log-normal distributions. Therefore, while the tails of the tvol variables are somewhat fatter compared to the degree and ntrans variables, the power-law remains a poor description of the data.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Exponential	.1741	.1940	.1722	.1258	.3419	.3656	.3383	.3139	.1520	.1808	.1390	.0887
Gamma	.0511	.0713	.0429	**.0313**	.0342	.0370	.0465	.0449	.0722	.0881	.0631	.0651
Geometric	.1741	.1940	.1722	.1258	.3418	.3655	.3383	.3139	.1519	.1808	.1390	.0886
Log-normal	.0619	.0486	.0626	.1075	.0610	.0635	.0566	.0969	**.0332**	**.0212**	**.0316**	.0810
Neg. Bin.	.0491	.0701	.0404	.0314	**.0298**	**.0319**	.0417	**.0414**	.0712	.0867	.0626	.0649
Poisson	.7128	.7366	.6989	.6848	.7579	.7574	.7516	.7264	.7139	.7277	.6969	.6907
Power-law	.4030	.3888	.4106	.3765	.2730	.2752	.2910	.2930	.4733	.4614	.4545	.4588
Weibull	**.0202**	**.0378**	**.0165**	.0462	.0308	.0334	**.0260**	.0585	.0447	.0529	.0430	**.0566**

Tab. 4.12: Quarterly data, tvol. KS statistic for the candidate distributions (complete). Minimum values in bold. Significance tests not carried out.

	In				Out				Total			
Period	1-3	1	2	3	1-3	1	2	3	1-3	1	2	3
Exponential	.0855	.1380	.0759	.1476	.0961	.1529	.0567	.1563	.1071	.1434	.0721	.1508
Log-normal	**.0271**	**.0542**	**.0267**	**.0521**	**.0674**	**.0407**	**.0475**	**.0459**	**.0353**	**.0483**	**.0695**	**.0394**
Power-law	.0642	.0588	.0672	.0703	.0785	.0733	.0813	.0790	.0788	.0778	.0812	.0484

Tab. 4.13: Quarterly data, tvol. KS statistic for the candidate distributions (tail). Minimum values in bold. Significance tests not carried out.

5. CORE-PERIPHERY STRUCTURE OF THE OVERNIGHT MONEY MARKET: EVIDENCE FROM THE E-MID TRADING PLATFORM

5.1 Introduction and Existing Literature

Interbank markets allow banks to exchange central bank money in order to share liquidity risks.[1] At the macro level, however, a high number of bank connections could give rise to systemic risk.[2] Since it is well known that the structure of a network is important for its resilience,[3] policymakers need information on the actual topology of the interbank network.

The experiences of the last few years have made policymakers aware of the necessity of gathering information on the structure of the financial network in general and the interbank market in particular.[4] One reason for the previous scarcity of research on the connections between financial institutions is certainly the limitation of available data,[5] the other reason being the neglect of the internal structure of the financial system by the dominating paradigm in macroeconomics during the last quarter of a century.[6]

Recent research in the natural sciences has significantly advanced our understanding of the structure and functioning of complex networks. Network ideas have been applied to very diverse areas and data sets such as the internet, epidemiology, ecosystems, scientific collaboration and financial markets, to name a few.

Most previous studies on the topology of interbank markets have been conducted by physicists applying measures from the natural sciences to a network formed by interbank liabilities. Examples include Boss *et al.* (2004) for the Austrian interbank market, Inaoka *et al.* (2004) for the Japanese BOJ-Net, Soramäki *et al.* (2007) for the US Fedwire network, Bech and Atalay (2010) for the US Federal funds market, and De Masi *et al.* (2006) and Iori *et al.* (2008) for the Italian e-MID (electronic market for interbank deposit). Overall, the most important findings of this literature are: (1) interbank networks are sparse, i.e. their density is relatively low,[7] (2) degree distributions appear to be scale-free (with coefficients between 2-3),[8] (3)

[1] See Ho and Saunders (1985), Freixas *et al.* (2000) and Allen and Gale (2000).

[2] Systemic risk is closely related to financial contagion, see de Bandt and Hartmann (2000), and implies that an idiosyncratic shock causing the failure of one or few institutions may destabilize the entire system.

[3] See also Allen and Gale (2000).

[4] See Haldane (2009), Haldane and May (2011) and Trichet (2011).

[5] See Mistrulli (2007).

[6] See Colander *et al.* (2009) for a more general critique.

[7] The density of a network is simply the fraction of existing links, relative to the maximum possible number of links. Ignoring the diagonal elements, the density can be calculated as $M/(N^2 - N)$, with M being the number of observed links and N the number of active nodes (banks).

[8] The in-degree is the number of incoming links, while the out-degree is the number of

transaction volumes appear to follow scale-free distributions as well, (4) clustering coefficients are usually quite small, (5) interbank networks are close to 'small world' structures, and (6) the networks show disassortative mixing, i.e. high-degree nodes tend to trade with low-degree nodes, and vice versa.[9] This indicates that small banks tend to trade with large banks, but rarely among themselves. Thus, we might expect the interbank network to display some sort of hierarchical community structure.

In passing, many authors have indeed remarked that there seemed to be some kind of community structure in the interbank network they analyzed. For example, Boss *et al.* (2004) note that the Austrian interbank network shows a hierarchical community structure that mirrors the regional and sectoral organization of the Austrian banking system. Soramäki *et al.* (2007) show that the network includes a tightly connected core of money-center banks to which all other banks connect. Thus there is some form of tiering in the interbank market. The empirical findings of Cocco *et al.* (2009) also show that relationships between banks are important factors to explain differences in interest rates.

Identifying communities in networks is an important aspect and in this paper we are concerned with the identification of the set of arguably systemically important (core) banks. In order to do so, we estimate various versions of core-periphery models in the spirit of Borgatti and Everett (2000).[10] Similar to De Masi *et al.* (2006) and Iori *et al.* (2008) we use data from the Italian e-MID trading platform, which is a market for unsecured deposits virtually covering the entire domestic overnight deposit market in Italy. Core-periphery models have been applied in a number of interesting fields before, for example to identify the spreaders of sexually transmitted diseases (see Christley *et al.* (2005)), in protein interaction networks (see Luo *et al.* (2009)), and to identify opinion leaders in economic survey data (see Stolzenburg and Lux (2011)). To our knowledge, Craig and von Peter (2010) is the first and so far only contribution applying a core-periphery structure to an interbank market. Applying this core-periphery framework to a data set of credit relationships between German banks,[11] their results speak in favor of

outgoing links per bank.

[9] Quite interestingly, the conventional explanation of the scale-free degree distribution is that of preferential attachment. Note that this is rather the opposite of disassortative mixing.

[10] Another interesting approach in using network-based measures for financial regulation is presented in Markose et al. (2010). The authors construct a so-called super-spreader tax based on eigenvector centrality.

[11] The authors use comprehensive statistics from the so-called 'Gross- und Millionenkreditstatistik' (statistics on large loans and concentrated exposures) from the Deutsche Bundesbank. In Germany, financial institutions have to report (on a quarterly basis) their

a very stable set of core banks. Furthermore, they show that core membership can be predicted using bank-specific features such as balance sheet size.[12] In this paper we will apply the (unrestricted) discrete core-periphery model, the (restricted) tiering model due to Craig and von Peter (2010) as well as symmetric and asymmetric versions of a continuous core-periphery model (hitherto not applied to interbank data) to a different set of interbank market data. Using a detailed dataset containing all overnight interbank transactions in the Italian interbank market from January 1999 to December 2010, we find that a core-periphery structure provides a better fit for these interbank data than alternative network models. The identified core shows a high degree of persistence over time, consisting of roughly 28% of all banks before the global financial crisis and 23% afterwards. We can classify the majority of core banks as intermediaries, i.e. as banks both borrowing and lending money in the market. Furthermore, allowing for asymmetric 'coreness' with respect to lending and borrowing activity considerably improves the fit, and reveals more concentration in borrowing than lending activity of money center banks. We also shed light on the development during the financial crisis of 2008, finding that the reduction of interbank lending was mainly due to core banks' reducing their numbers of active outgoing links.

The remainder of this paper is structured as follows: section 5.2 gives a brief introduction into necessary terminology for the formalisation of (interbank) networks, section 5.3 introduces the Italian e-MID interbank data and highlights some of its important properties. Section 5.4 introduces different variants of the core-periphery model. Section 5.5 presents the results and different robustness checks. Section 5.6 discusses the findings and section 5.7 concludes. A set of appendices provides more technical details as well as further robustness checks.

5.2 Networks

A network consists of a set of N nodes that are connected by M edges (links). Taking each bank as a node and the interbank positions between them as links, the interbank network can be represented as a square matrix of dimension $N \times N$ (data matrix, denoted $\mathbf{D}$).[13] An element d_{ij} of this matrix represents a gross interbank claim, the total value of credit extended

total exposure to each counterparty to whom they have extended credit of at least 1.5 million Euros or 10% of their liable capital to the Bundesbank. These reports include outstanding claims of any maturity.

[12] We cannot carry out such an analysis since we do not observe bank IDs, see below.

[13] In the following, matrices will be written in bold, capital letters. Vectors and scalars will be written as lower-case letters.

by bank i to bank j within a certain period. The size of d_{ij} can thus be seen as a measure of link intensity. Row (column) i shows bank i's interbank claims (liabilities) towards all other banks. The diagonal elements d_{ii} are zero, since a bank will not trade with itself.[14] Off-diagonal elements are positive in the presence of a link and zero otherwise.

Interbank data usually give rise to directed, sparse and valued networks.[15] However, much of the extant network research ignores the last aspect by focusing on binary adjacency matrices only. An adjacency matrix $\mathbf{A}$ contains elements a_{ij} equal to 1, if there is a directed link from bank i to j and 0 otherwise. Since the network is directed, both $\mathbf{A}$ and $\mathbf{D}$ are asymmetric in general. In this paper, we also take into account valued information by using both the raw data matrix as well as a matrix containing the number of trades between banks, denoted as $\mathbf{T}$. In some cases it is also useful to work with the undirected version of the adjacency matrices, $\mathbf{A}^u$, where $a^u_{ij} = \max(a_{ij}, a_{ji})$.

As usual, some data aggregation is necessary to represent the system as a network. In the following, we use quarterly networks. The next section summarizes the most important properties of our data, more detailed information can be found in Finger *et al.* (2012).

5.3 Dataset

The Italian electronic market for interbank deposits (e-MID) is a screen-based platform for trading of unsecured money-market deposits in Euros, US-Dollars, Pound Sterling, and Zloty operating in Milan through e-MID SpA.[16] The market is fully centralized and very liquid; in 2006 e-MID accounted for 17% of total turnover in the unsecured money market in the Euro area. Average daily trading volumes were 24.2 bn Euro in 2006, 22.4 bn Euro in 2007 and only 14 bn Euro in 2008.

Available maturities range from overnight up to one year. Most of the transactions are overnight. While the fraction was roughly 80% of all trades in 1999, this figure has been continuously increasing over time with a value of

[14] This is of course only true when taking banks as consolidated entities. There are, however, important examples of self-referential networks: the typical node in a connection matrix of the brain represents a group of neurons; in citation networks authors cite articles appearing in the same journal. See Boyd *et al.* (2010) for a discussion.

[15] Directed means that $d_{i,j} \neq d_{j,i}$ in general. Sparse means that at any point in time the number of links is only a small fraction of the $N(N-1)$ possible links. Valued means that interbank claims are reported in monetary values as opposed to 1 or 0 in the presence or absence of a claim, respectively.

[16] The vast majority of trades (roughly 95%) is conducted in Euro.

more than 90% in 2010.[17] As of August 2011, e-MID had 192 members from EU countries and the US. Members were 29 central banks acting as market observers, 1 ministry of finance, 101 domestic banks and 61 international banks. We will see below that the composition of the active market participants has been changing substantially over time. Trades are bilateral and are executed within the limits of the credit lines agreed upon directly between participants. Contracts are automatically settled through the TARGET2 system.

The trading mechanism follows a quote-driven market and is similar to a limit-order-book in a stock market, but without consolidation. The market is transparent in the sense that the quoting banks' IDs are visible to all other banks. Quotes contain the market side (buy or sell money), the volume, the interest rate and the maturity. Trades are registered when a bank (aggressor) actively chooses a quoted order. The platform allows for credit line checking before a transaction will be carried out, so trades have to be confirmed by both counterparties. The market also allows direct bilateral trades between counterparties.

The minimum quote size is 1.5 million Euros, whereas the minimum trade size is only 50,000 Euros. Thus, aggressors do not have to trade the entire amount quoted.[18] Additional participant requirements, for example a certain amount of total assets, may pose an upward bias on the size of the participating banks. In any case, e-MID covers essentially the entire domestic overnight deposit market in Italy.[19]

We have access to all registered trades in Euro in the period from January 1999 to December 2010. For each trade we know the two banks' ID numbers (not the names), their relative position (aggressor and quoter), the maturity and the transaction type (buy or sell). As mentioned above, the majority of trades is conducted overnight and due to the global financial crisis (GFC) markets for longer maturities essentially dried up. We will focus on all overnight trades conducted on the platform, leaving a total number of 1,317,679 trades. The large sample size of 12 years allows us to analyze the network evolution over time. Here we focus on the quarterly aggregates, leaving us with 48 snapshots of the network.

[17] This development is driven by the fact that the market is unsecured. The recent financial crisis made unsecured loans in general less attractive, with stronger impact for longer maturities. See below. It should be noted, that there is also a market for secured loans called e-MIDER.

[18] The minimum quote size could pose an upward bias for participating banks. It would be interesting to check who are the quoting banks and who are the aggressors. Furthermore it would be interesting to look at quote data, as we only have access to actual trades.

[19] More details can be found on the e-MID website, see http://www.e-mid.it/.

The left panel of Figure 6.1 shows the development of the number of active banks over time. We see a clear downward trend in the number of active Italian banks over time (green line), whereas the additional large drop after the onset of the GFC is mainly due to the exit of foreign banks. The right panel shows that the decline of the number of active Italian banks went along with a relatively constant trading volume in this segment until 2008. This suggests that the decline of active Italian banks was mainly due to mergers and acquisitions within the Italian banking sector. The overall upward trend of trading volumes was due to the increase of active foreign banks until 2008, while their activities in this market virtually faded away after the onset of the crisis.

The data show a trivial community structure in that foreign banks tend to trade with each other preferentially, and so do Italian banks. Due to the limited extent of trading between both components, and the smaller number of foreign banks, we will focus on Italian banks only in our subsequent analysis. This leaves a total number of 1,215,759 trades for the analysis.

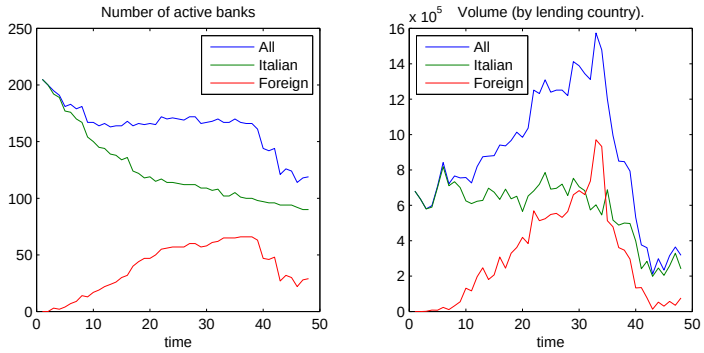

Fig. 5.1: Number of active banks (left) and traded volume (right) over time. We also split the traded volume into money lent by Italian and foreign banks, respectively.

Other important findings are:

- The e-MID network has a relatively high density compared to other interbank networks investigated in the literature.[20] See Figures 6.1

[20] Note that the density in the German interbank network is smaller for two reasons: first, the number of active banks is much larger, so it is more likely to observe missing links. Second, in our analysis we focus on overnight trades only, while Craig and von Peter (2010) use aggregate credit volumes of all maturities (probably only with a small fraction of overnight trades). It seems plausible that the probability of observing a link between any two banks should be inversely related to the maturity of the loan.

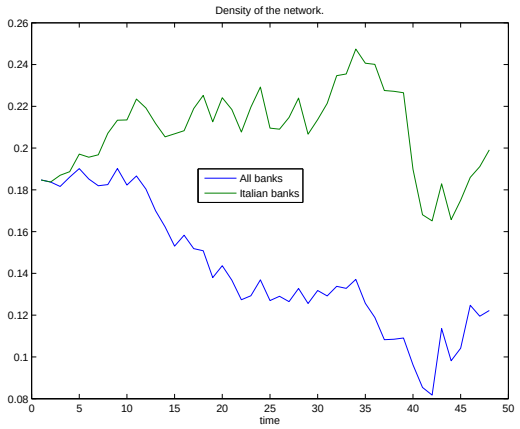

Fig. 5.2: Density of the network over time, calculated as $M_t/(N_t^2 - N_t)$, with M_t being the number of observed links and N_t the number of active banks in the respective quarter. A Chow-test indicates that there is a structural break after quarter 39 at all sensible significance levels for the Italian banks. A CUSUM-test also indicates a structural break, however, the time series seems to revert towards its pre-GFC level.

and 5.2. For the density of the network formed by Italian banks, a Chow-test and a CUSUM-test both indicate that there is a structural break after quarter 39 (i.e. at the onset of the financial crisis). Later on, we will see that the core-periphery structure was also influenced by the GFC.

- The aggregation period is important for economic applications as the network structure is less volatile with longer aggregation periods. Since the network is sparse, short periods will only give an incomplete image of existing linkages, where many links between otherwise frequent trading partners may be dormant. In order to obtain a more comprehensive and less random picture of existing links, a larger aggregation period is required. We will, therefore, use quarterly data in the following (but results are robust to somewhat shorter or larger aggregation periods).

- There is very small (at times even negative) correlation between the banks' in- and out-degrees. Hence, the directed version of the network might contain important additional information.

- The underlying distributions of in- and out-degrees are apparently not scale-free at any aggregation level (including the daily level), cf. Finger

et al. (2012). The same holds for transaction volumes.

- The network shows disassortative mixing patterns, so nodes with high overall degree (number of connections) tend to connect with low-degree nodes. We find similar assortativity coefficients for the relation between in- and out-degree, so high in-degree (out-degree) nodes tend to connect to low out-degree (in-degree) nodes.

In the next section, we will describe the different versions of the core-periphery model in detail.

5.4 Models

Core-periphery network models have been proposed first by Borgatti and Everett (2000). The basic idea is that a network can be divided into subgroups of core and periphery members. The discrete model partitions banks such that core (periphery) banks are maximally (minimally) connected to each other. The concept of discrete group membership can be extended by considering the core and periphery as opposite ends of a continuum. The continuous model overcomes the excessive simplicity of the discrete partitioning, by assigning a 'coreness' level to each bank. In the following we will first present the discrete model, with the tiering model proposed by Craig and von Peter (2010) as a special case, and then move on to the asymmetric continuous model for directed networks due to Boyd *et al.* (2010). Throughout the following we assume that a network cannot have more than one core.[21]

5.4.1 The Discrete Model

Formalisation

To identify the N_c core members among our sample of N banks, we aim at sorting the binary adjacency matrix such that we have the core-core region as a 1-block in the upper left part (of dimension $N_c \times N_c$) and the periphery-periphery region as a 0-block in the lower right part (of dimension $(N - N_c) \times (N - N_c)$). The idealized pattern matrix ($\mathbf{P}_I$) for a 'pure' core-periphery segmentation, then, looks as follows:[22]

$$\mathbf{P}_I = \begin{pmatrix} \mathbf{CC} & \mathbf{CP} \\ \mathbf{PC} & \mathbf{PP} \end{pmatrix} = \begin{pmatrix} \mathbf{1} & \mathbf{CP} \\ \mathbf{PC} & \mathbf{0} \end{pmatrix}, \tag{5.1}$$

[21] Everett and Borgatti (2000) include the possibility of multiple cores.

[22] The diagonal elements will be ignored in all that follows, since the network is not self-referential.

where **1** and **0** denote submatrices of ones and zeros.

The **CC**-block contains the top-tier banks, while the **PP**-block contains the periphery. Note that the off-diagonal blocks may be 1-blocks (each core member connected to all periphery-nodes), 0-blocks (no connection between core and periphery members) or something in between, depending on the problem. Borgatti and Everett (2000) claim that only the diagonal blocks are characteristic of CP structures and are thus the defining property. We will denote this version, without any restrictions on the off-diagonal blocks, as the discrete model.

In some cases however, the underlying model explicitly dictates requirements on the **CP** and **PC** blocks. For instance, Craig and von Peter (2010) propose a more strictly tiered interbank market than the benchmark discrete structure. In this model, a key characteristic of core banks (top tier) is that they intermediate between periphery banks. If at least a minimum level of intermediation activity is required of a 'core' bank, this means that **CP** and **PC** have to be row- and column-regular,[23] respectively, i.e. at least one entry has to be non-zero in each row of **CP** and in each column of **PC**.

Optimization Problem

The discrete core-periphery framework amounts to assigning to each bank the property of membership in the core or the periphery. This classification can be summarized in a vector c of zeros and ones of length N (the total number of banks). The usual approach to find the optimal coreness vector, c, referred to as the minimum residual (MINRES) approach, is to fit a pattern matrix $\mathbf{P} = cc'$, which should be as close as possible to the observed network matrix $\mathbf{A}$. This requires to identify the core banks, which are unknown a priori.

We start by defining a coreness vector, ordering the core banks first and writing the set of core members as $\mathcal{C} = \{1, \cdots, N_c\}$.[24] Then we can measure the 'fit' of the corresponding core-periphery structure as the total number of inconsistencies between the observed network and the idealized pattern matrix $\mathbf{P}_I$ of the same dimension. Depending on the problem, the distance involves certain restrictions on the off-diagonal blocks, **CP** and **PC**. The optimal partition $\mathcal{C}^*$ thus minimizes the residuals and gives the optimal set of core banks.

Residuals are obtained by simply counting the errors in each of the four

[23] See Doreian *et al.* (2005).

[24] Note that in order to have a core, N_c has to be ≥ 2. Also note the difference between $\mathcal{C}$ and c: $\mathcal{C}$ is the set of core banks and thus is a vector of dimension N_c, while c is a vector of zeros and ones. Of course, both $\mathcal{C}$ and c carry the same information.

blocks of Eq. (5.1) and aggregating over the blocks. The core-core block should be a complete **1**-block of dimension N_c, so any missing link represents an inconsistency (residual) with respect to the model.[25] Likewise any link between two periphery banks constitutes an error relative to the benchmark. Obviously, we can introduce any constraints on the off-diagonal blocks, so the tiering model can be easily implemented here as well: errors in the off-diagonal blocks penalize zero rows and columns, because these are inconsistent with row- and column-regularity, respectively. For example, a zero column could be penalized by as many errors as there are banks in the periphery ($N - N_c$).

For the general version of the discrete model with arbitrary off-diagonal blocks, the aggregate errors in the individual blocks can be written as

$$\mathbf{E}(\mathcal{C}) = \begin{pmatrix} E_{CC} & E_{CP} \\ E_{PC} & E_{PP} \end{pmatrix} = \begin{pmatrix} N_c(N_c-1) - \sum_{i,j\in\mathcal{C}} a_{ij} & 0 \\ 0 & \sum_{i,j\notin\mathcal{C}} a_{ij} \end{pmatrix}. \tag{5.2}$$

The total error score (e) then simply aggregates the errors across the relevant blocks, normalized by the total number of links in the network.[26] Formally this can be written as

$$e(\mathcal{C}) = \frac{E_{CC} + E_{CP} + E_{PC} + E_{PP}}{M} = \frac{E_{CC} + E_{PP}}{M}, \tag{5.3}$$

with $e(\cdot)$ being a function of $\mathcal{C}$ since every possible partition is associated with a particular value of e.

For the tiering model proposed by Craig and von Peter (2010), the aggregate errors in the off-diagonal blocks can be calculated as

$$E_{CP} = (N - N_c) \sum_{i\in\mathcal{C}} \max(0, 1 - \sum_{j\notin\mathcal{C}} a_{ij}) \tag{5.4}$$

and

$$E_{PC} = (N - N_c) \sum_{j\in\mathcal{C}} \max(0, 1 - \sum_{i\notin\mathcal{C}} a_{ij}), \tag{5.5}$$

respectively, leading to additional non-zero entries in $e(\mathcal{C})$.

The optimal partition $\mathcal{C}^*$ is the set of core banks producing the smallest distance to an idealized pattern matrix of the same dimension, i.e.

$$\mathcal{C}^* = \arg\min e(\mathcal{C}) = \{\mathcal{C} \in \Omega | e(\mathcal{C}) \leq e(C) \forall C \in \Omega\}, \tag{5.6}$$

[25] The maximum number of possible inconsistencies in this block would be $N_c(N_c - 1)$ since the main diagonal is ignored. This upper bound is obviously never reached since otherwise there would be no core-periphery structure.

[26] Note that M is the maximum error possible in a network consisting only of a periphery.

where Ω denotes all strict and non-empty subsets of the population $\{1, \cdots, N\}$. It should be noted, however, that the discrete approach implicitly assumes symmetry of the underlying structure (or irrelevance of the direction of links). Therefore, in section 5.4.2 we will turn to a continuous core-periphery model, which explicitly takes the directed nature of the network into account, characterizing coreness by two vectors rather than one.

Implementation

Fitting the discrete and the tiering model to a real-world network is a large scale problem in combinatorial optimization. Exhaustive search becomes impractical for large matrices, since the number of possible labeled bipartitions increases exponentially with the dimension of the matrix. More precisely, the number of nontrivial bipartitions (with both the core and the periphery having at least two members) is $2^N - 2N - 2$. The term 2^N corresponds to the number of all possible subsets, while the negative terms exclude partitions with only core or periphery banks. For example, with $N = 10$ banks there are 1002 nontrivial possible bipartitions. For a system with $N = 100$ banks there are already roughly 10^{30} partitions.

A number of algorithms have been applied to tackle such problems. We will use a Genetic Algorithm (GA) to fit both the discrete and the tiering model.[27] A GA uses operations similar to genetic processes of biological organisms to develop better solutions of an optimization problem from an existing population of (randomly initiated) candidate solutions. Typically the proposed solutions are encoded in strings (chromosomes) mostly using a binary alphabet, i.e. in our setting the strings have length N and consist of ones and zeros, depending on whether a bank is in the core or periphery. We use the rate of correct classifications (in terms of the error score) by a string l, $f_l = 1 - e(\mathcal{C}_l)$ as a fitness function that drives the evolutionary search. Details are explained in Appendix 5.8.1.

5.4.2 The Continuous Model

Basic Structure

One limitation of the partition-based approach presented above is the excessive simplicity of defining just two homogeneous classes of nodes: core

[27] We cross-checked the results using the sequential algorithm applied in Craig and von Peter (2010). Alternatives would be the Kerninghan-Lin Algorithm (Kernighan and Lin (1970)), see Boyd *et al.* (2006) for an application, and Branch-and-Bound Programming, see Brusco (2011).

and periphery. Assuming that the network data consist of continuous values representing strengths or capacities of relationships (for banking data: credit volumes or number of transactions), it seems sensible to also consider a continuous model in which each node is assigned a measure of 'coreness'. Since a continuous measure of coreness allows for more flexibility in capturing the role of an institution, we apply this model to the valued matrix $\mathbf{D}$ of interbank liabilities rather than the binary adjacency matrix $\mathbf{A}$.

The usual approach in the symmetric continuous (SC) model is to find a coreness vector c, where $1 \geq c_i \geq 0 \; \forall i$, with pattern matrix $\mathbf{P} = cc'$ that approximates the observed data matrix as closely as possible. Similar to the presentation of the discrete model, the optimal coreness vector in the symmetric continuous (SC) model can be found using the MINRES approach.[28] Again however, this method imposes a symmetric pattern matrix, i.e. $p_{ij} = p_{ji} \; \forall i, j$. Thus, it is assumed that the strength of the relation from i to j is the same as that from j to i. To overcome this restriction, we also estimate an asymmetric continuous (AC) core-periphery model, as introduced by Boyd *et al.* (2010). This formulation involves two vectors, representing the degrees of outgoing and incoming centrality for each node. For networks of international trade, for example, the two vectors would correspond to exports and imports, respectively. In our setting, the two vectors correspond to out- and in-coreness. Note that both the SC and AC model can be applied to valued matrices, with binary adjacency matrices being just a special case. Thus the continuous models might allow us to extract important additional information from the directed, valued networks. However, a disadvantage of the continuous models is that restrictions, such as the tiering model, cannot be implemented. In the following, we will briefly introduce both model versions. More details on the AC model can be found in Appendix 5.8.3.

The Symmetric Continuous (SC) Model

The SC model will again be estimated by minimization of residuals. MINRES seeks a column vector c such that the square matrix $\mathbf{D}$ is approximated by the pattern matrix $\mathbf{P} = cc'$. Ignoring the diagonal elements, this amounts to minimizing the sum of squared differences of the off-diagonal elements, or

$$\arg\min_{c} \sum_{i} \sum_{j \neq i} (d_{ij} - c_i c_j)^2. \tag{5.7}$$

In the same spirit as with our optimization algorithm in the discrete case, we use the proportional reduction of error (PRE) as our measure of fit. PRE

[28] An interesting alternative approach, based on the Kullback-Leiber distance, can be found in Muniz and Carvajal (2006) and Muñiz *et al.* (2011).

is defined as

$$\text{PRE}(cc'|\langle D\rangle) = 1 - \frac{SS(\mathbf{D} - cc')}{SS(\mathbf{D} - \langle D\rangle)}, \tag{5.8}$$

with $\langle D\rangle$ being the global average (across all elements, excluding the diagonal) of $\mathbf{D}$ and $SS(\cdot)$ is the sum of squared deviations of the off-diagonal elements of the input matrix. Thus, maximizing the PRE is equivalent to minimizing $SS(\mathbf{D} - cc')$. Boyd *et al.* (2010) argue that the continuous core/periphery model makes a reasonable contribution towards explaining empirical structures if the PRE significantly exceeds 0.5. Note that the reported coreness vectors in both the SC and AC model will be standardized by the Euclidean norm of the optimal solution vectors.

The Asymmetric Continuous (AC) Model

The idea of the asymmetric continuous (AC) model is to decompose overall 'coreness' into 'out-coreness' and 'in-coreness' (denoted by u_i and v_i in the following), respectively. Applying this distinction allows us to write the objective function for the AC model as

$$\arg\min_{u,v} \sum_i \sum_{j \neq i} (d_{ij} - u_i v_j)^2. \tag{5.9}$$

The optimal coreness vectors can be determined by finding the roots of the first-order conditions of Eq. (5.9).[29] The PRE of the AC model can be defined similarly as in Eq. (5.8) as

$$\text{PRE}(cc'|\langle D\rangle) = 1 - \frac{SS(\mathbf{D} - uv')}{SS(\mathbf{D} - \langle D\rangle)}. \tag{5.10}$$

For both the SC and the AC models, we will, in order to adjust for the skewness of the network matrices, log-transform the data matrix in the form $\log(1 + \mathbf{D})$, where the factor 1 makes sure that zeros in the original matrix remain zeros in the transformed matrix.[30] Note that the split into in- and out-coreness is germane to a singular value decomposition of our matrix $\mathbf{D}$ of interbank liabilities. This similarity is exploited in the empirical estimation of the vectors u and v. Our numerical approach for estimating these two coreness vectors follows Boyd *et al.* (2010) and is detailed in Appendix 5.8.3.

[29] This could be implemented by using standard algorithms for numerical optimization. Here we used a trust-region algorithm.

[30] We also tried to fit the core-periphery models to the raw network matrices, however, the high level of skewness in the data results in a very poor fit in general. These results are hardly comparable to those presented below, see Appendix 5.8.9.

5.5 Results

This section presents and discusses the results from the different versions of the core-periphery framework. In the following, as noted above, we focus on the quarterly networks formed by Italian banks only. Robustness checks, using different aggregation periods and sample banks can be found in the Appendix.[31] Recall that the discrete and tiering model use the (binary) adjacency matrices $\mathbf{A}$, while the continuous model uses the valued matrix of transaction volumes $\mathbf{D}$, as defined in section 5.2.[32]

As a first step, we compare the coreness vectors between the different models. It will become clear that the discrete and tiering model are almost identical throughout. Later on, we show that the AC model contains important information from the asymmetric nature of the network, since the in- and out-coreness vectors are far from being perfectly correlated. Secondly, we investigate the properties of the core/periphery banks. We find that the core is large compared to the findings in Craig and von Peter (2010), but also very persistent over time. Due to the high network density, we find that the error scores are also much higher compared to the German market. In particular, the model fit deteriorates over time due to the GFC. Formal tests suggest a significant worsening of the fit of the core-periphery model after the GFC, pointing towards the breakdown some part of the core-periphery structure. As a last step, we investigate the significance of the results by comparing the identified cores and the corresponding error scores to the cores obtained from random and scale-free networks, calibrated to share similar properties as the observed ones along certain dimensions. Here we find that the identified cores are significant, i.e. the identified core-periphery structure is not a spurious network property.

5.5.1 Model Similarity

Table 5.1 presents selected correlations between the identified coreness vectors of the different model versions. For each combination, we compute the correlation between the (stacked) coreness vectors for the complete sample period. Note that the discrete and tiering coreness vectors contain only binary values, while the in- and out-coreness vectors contain real numbers. Obviously the correlation between the cores in the discrete and tiering model

[31] Appendix 5.8.6 discusses the findings for other aggregation periods, most importantly for monthly and yearly networks. Appendix 5.8.7 discusses the results when including foreign banks to the analysis.

[32] Appendix 5.8.8 discusses the results for the continuous model using the matrix containing the number of transactions $\mathbf{T}$. Appendix 5.8.9 discusses further robustness checks.

is very high with a value of around .95.[33] The same is true for the discrete core and the out-coreness with a value of .73, whereas the correlation between the discrete core and the in-coreness is much smaller with a value of .26.[34] Core banks from the discrete model are therefore more likely to be in the out-core of the continuous model as well, but not necessarily in the in-core. This result seems rather surprising at first, since for example the results from Cocco *et al.* (2009) suggest that small (periphery) banks are net lenders, which offer their excess liquidity to a preferred set of large (core) banks. Our analysis shows that at least in the present data set, the pattern of interbank linkages is more complex: again, periphery banks lend money to a relatively small set of selected core banks, but the core banks in turn tend to redistribute this liquidity not only among the other core banks, but also among a larger part of the periphery. Technically, we find that the density in the CP-block is on average three times higher than the density in the PC-block (see Figure 5.6 below), so for most core banks the out-degree clearly exceeds the in-degree.[35] Therefore, it is not surprising that the correlation between the discrete and the out-coreness is higher than the correlation with the in-coreness. This shows that there is a considerable amount of asymmetry in the network, also captured by the negative correlation of -.08 between the in- and out-coreness vectors, cf. Figure 5.3. We see that these relations are rather stable. Interestingly, the correlation between in- and out-coreness was always the smallest of these combinations, turning negative after 12 quarters and remaining so for the rest of the sample period. This hints towards the existence of different subgroups in the core.

In the following we present more detailed results for the discrete and tiering model, then moving on to the continuous model.

[33] Therefore, the correlations between the tiering core and the in-/out-coreness are not presented here since they are very similar to those from the general discrete model.

[34] Interestingly, we see that the correlation between the coreness vectors from the (symmetric) discrete and the SC model is only .7578. One might expect that this is partly driven by the fact that the input matrix is valued, rather than binary in the continuous case. Estimating the continuous model with binary network matrices, however, yields very similar results, see Appendix 5.8.9, with a correlation of .7635. Thus, the main reason for the low correlation between the two vectors lies in the objective function: the continuous models approximate the complete matrix, while the discrete model focuses on the diagonal blocks.

[35] This also explains the small (at times even negative) correlation between individual banks' in- and out-degree.

Models		Correlation
Discrete	Tiering	.9526
Discrete	Out-coreness	.7267
Discrete	In-coreness	.2567
Discrete	Sym. coreness	.7578
In-coreness	Out-coreness	-.0809

Tab. 5.1: Correlations between individual coreness vectors of different models. For each model, we stack the coreness vectors over the entire sample period in a single vector. Then we compute the correlations between each combination. Note that the discrete and tiering coreness consists of binary values, while the in-, out-, and symmetric coreness vectors contain real numbers.

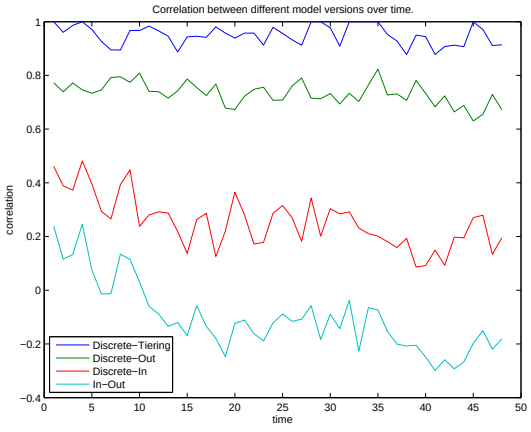

Fig. 5.3: Time-varying correlations between different coreness vectors.

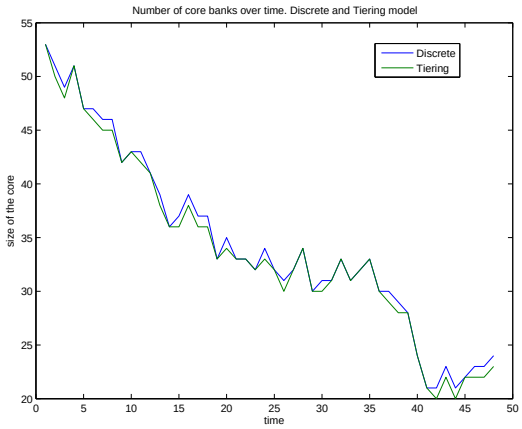

Fig. 5.4: Absolute size of the core over time. A Chow-test indicates that there is a structural break for the detrended time series after quarter 10, while there is no evidence for a significant structural break after quarter 39. An additional CUSUM test indicates that this break is significant at all sensible confidence levels. We also note a significant level of autocorrelation in the detrended time series, while the first difference of the original time series is stationary.

5.5.2 Discrete and Tiering Model

The Size of the Core and Periphery

We saw that the identified cores in the discrete and tiering model are highly correlated. In fact, Figure 5.4 shows that the sizes of their cores are very similar over time. Note that the core in the discrete model is always at least as large as the core in the tiering model. The reason lies in the requirement that all core banks in the tiering model act as intermediaries, which is not necessarily true for the discrete model, even though again the vast majority of core banks acts as intermediaries in this case. Overall, the differences between the two model versions consist of a few borderline cases.[36]

Note also the negative trend in the absolute size of the cores over time. This is not surprising given that the number of active Italian banks has been decreasing over time. Interestingly, a Chow-test indicates the existence of a structural break in the (detrended) core sizes after quarter 10, with the

[36] In cases where the row- and column-regularity constraints are binding, it may also happen that core banks from the discrete model are part of the periphery in the tiering model.

trend going back towards its initial level in the post-GFC period.[37] The economic significance of this result is, however, questionable as we see in Figure 5.4 that a linear negative trend might fit the entire sample period quite well, and we know that the sharp drop after quarter 39 was due to the GFC. Given the overall trend in the number of active banks, it seems more interesting to consider the relative size of the core compared to the complete interbank network. Figure 5.5 shows that the relative size of the core is rather stable over time, fluctuating around 28% before the GFC, and around 23% afterwards. A Chow-test indicates that there is a structural break after quarter 39. However, under a CUSUM test this break is only marginally significant at the 5% level for the discrete model, and insignificant for the tiering model. Thus, there is some evidence that the GFC has led to a structural break in the formerly relatively stable structure of intermediation in the interbank market. However, we also see a positive trend in the core sizes for the last 3 quarters of the sample period, so that the relative core size seemed to revert to its pre-GFC level. Not surprisingly, the size of the core is highly correlated with the density of the network (cf. Figure 5.2). We should note that relative core sizes are very high compared to the value of 3% found for the German interbank market by Craig and von Peter (2010). This is driven by the very high overall network density of above 20%, compared to only 0.61% for the German market.[38]

The left panel of Figure 5.6 shows the densities of the complete network and the core-core and periphery-periphery subnetworks over time. Since results are virtually the same for both models, we only display those of the baseline discrete model[39] with rather stable values for the pre-GFC period, but again with a structural break after quarter 39 for all time series in the Figure. The density in the CC-block is at least 2.5 times that of the entire network and at least 6 times that of the PP-block. The right panel of Figure 5.6 shows the densities in the off-diagonal blocks. As already mentioned, the

[37] Iori *et al.* (2007) also mention this structural break in the Italian interbank network in quarter 10, however, without conducting formal tests. They relate this breakpoint to two events: (1) official and market interest rates changed their trend from positive to negative, (2) the ECB tried to support economic growth by increasing the amount of liquidity provided.

[38] Recall that the number of banks in the German market is roughly 1800, so the network is at least 10 times larger than the Italian network. Thus it is not surprising, that the density is much higher in the Italian case. Since the e-MID sample presumably contains mainly large banks, our core might be the core of the overall banking network. See Figure 5.21 in Appendix 5.8.2 for a network illustration for one particular quarter.

[39] Results from the tiering model are available upon request. We checked that the results from the tiering model are statistically not distinguishable from the results of the discrete model.

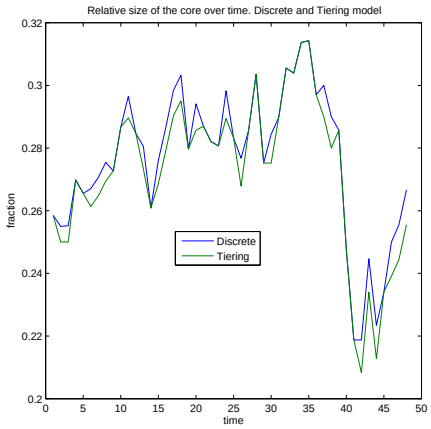

Fig. 5.5: Relative size of the core over time. A Chow-test indicates that there is a structural break after quarter 39 at all sensible significance levels. An additional CUSUM test indicates that this break is marginally significant at the 5% level.

density in the CP-block is three times higher than the corresponding density in the PC-block. These values are very stable over time, and we do not find evidence for a structural break.

The Structure of the Core and Periphery

To gain more insights into the structure of the network, Figure 5.7 shows the fraction of intermediaries, lenders and borrowers in the complete network over time. Here we define borrowers as banks with an out-degree of zero but positive in-degree in a given quarter, whereas the reverse holds for lenders. The remaining banks, with both positive in- and out-degree, are thus intermediaries. We see that these fractions are relatively stable over time: most of the banks (roughly 75%) act as intermediaries, a smaller fraction acts as lenders (20%) and the remainder consists of sole borrowers. Interestingly, the fraction of sole borrowers seems to increase significantly after the GFC, since we find a structural break after quarter 39. This may hint towards the entry of banks who use the market only to attract funds. In contrast, there is no significant structural break for the fraction of intermediaries and lenders. Table 5.2 shows the transition probabilities for each strategy, with I_t denoting that a bank is an intermediary in t. L, B and E stand for lending, borrowing and exit, respectively. The matrix shows, for example, that with a probability of 88.45% an intermediating bank in $t-1$ will also be an intermediary in t. Note that the diagonal elements are largest, even though

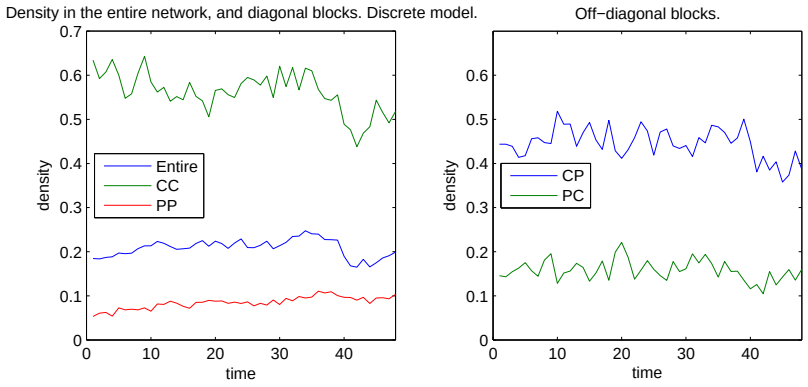

Fig. 5.6: Density of the entire network, CC/PP blocks (left), and off-diagonal blocks (right). Individual Chow-tests indicate that there is a structural break in the time-series in the left panel after quarter 39 at all sensible significance levels (see also Figure 5.2). Additional CUSUM-tests indicate that the structural breaks are significant at all sensible significance levels, with the PP-density apparently containing an additional structural break around quarter 10. In contrast, we cannot reject the hypothesis of no structural break in the time-series of the right panel.

	I_t	L_t	B_t	E_t
I_{t-1}	.8845	.0752	.0190	.0214
L_{t-1}	.2971	.6508	.0009	.0513
B_{t-1}	.3661	.0164	.4590	.1585
E_{t-1}	.0049	.0038	.0028	.9885

Tab. 5.2: Transition matrix: trading strategies. I, L, B, and E denote intermediary, lender, borrower, and exit, respectively.

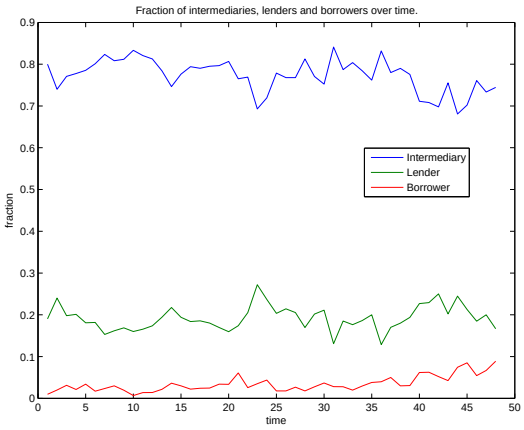

Fig. 5.7: Fraction of intermediaries, lenders and borrowers over time. Individual Chow-tests point towards the existence of a structural break after quarter 39 in all time-series. Additional CUSUM-tests, however, indicate that this structural break is only significant for the fraction of borrowers.

the borrowing strategy is less persistent over time compared to the other strategies. This is in line with the observation of a more intense entry of sole borrowers during and after the GFC.

Figure 5.8 shows the fractions of intermediaries, lenders, and borrowers in the core and periphery of the general discrete model. Again these results are very similar to those of the tiering model: the fraction of intermediaries in the core is highest, while the fraction of intermediaries in the periphery is second highest. As expected, only very small fractions of borrowers and lenders are found in the core (none in the tiering model), while banks that appear only as borrowers are a significant fraction (about 30 percent) of the periphery.

To elucidate the stability of these structural properties, consider Table

Fig. 5.8: Structure of the core and periphery in the discrete model. Fractions of intermediaries, borrowers and lenders, in the core and periphery, respectively. *Note:* ICore=intermediaries in the core, BCore=borrowers in the core, LCore=lenders in the core. Similarly for the periphery.

	C_t	P_t	E_t
C_{t-1}	.8324	.1565	.0110
P_{t-1}	.0555	.9055	.0391
E_{t-1}	.0012	.0104	.9885

Tab. 5.3: Transition matrix: discrete model. C, P and E stand for core, periphery and exit, respectively.

5.3 containing transition probabilities of the state of a bank for the discrete model. For example, the first rows show the probabilities of a core bank in $t-1$ (C_{t-1}) being a core member in t, switching to the periphery (P_t) or exiting the market (E_t). There is some asymmetry in the Table, for example, the probability of switching from the core to the periphery is roughly 15.6%, while the reverse probability is only 5.5%. In particular, the diagonal entries are very high with values above 80%, such that there is significant persistence (autocorrelation) in the group memberships.[40]

The above transition probabilities are aggregate values over the entire sample period. To investigate the inherent structural stability, the values in this matrix should be roughly constant over time. Figure 5.9 shows the time

[40] Note that the structure is very stable despite the existence of a structural break due to the GFC after quarter 39, cf. section 5.5.5.

evolution of these values for the discrete model. We see that the elements on the main diagonal are quite stable over time and very large in general. However $P(C|C)$ becomes smaller due to the GFC simply because a number of core banks become part of the periphery, which can be seen by the increase in $P(P|C)$ to more than 20%. Again we emphasize that we do not observe banks' names, so we are unable to track for example bank mergers and acquisitions.

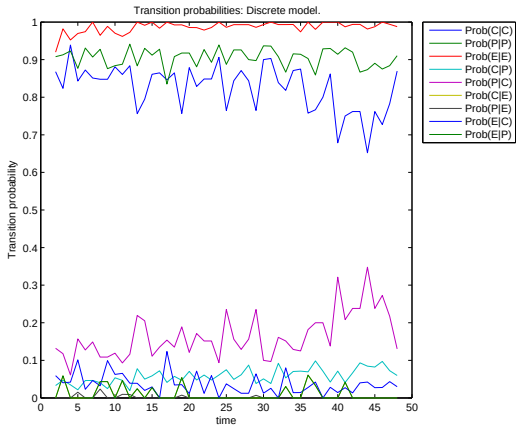

Fig. 5.9: Transition probabilities over time, discrete model. $P(y|x)$ is the probability of going from state x to state y.

Besides the overall structural stability, one might also be interested in the stability of the system at the micro-level of bilateral connections. In order to assess the stability of the link structure in the different blocks, we use the so-called Jaccard Index. This is defined as

$$J = \frac{M_{11}}{M_{01} + M_{10} + M_{11}}, \tag{5.11}$$

where M_{xy} is the number of relations with status x in period $t-1$ and with status y in the next period. It thus measures the similarity between subsequent graphs, taking only links into account which were present in at least one period. Social networks are usually considered to be sufficiently stable for values of J larger than .3, in which case the network is likely to display recognizable structure.[41] For the complete Italian interbank network we observe an average Jaccard Index of .5302 (std. dev.: .0368).[42] When calculating the Jaccard Index for the different blocks, we restrict ourselves

[41] See Snijders *et al.* (2009).
[42] See Finger *et al.* (2012).

to those banks having the same status of being a core/periphery bank in the two adjacent quarters.[43] Figure 5.10 shows the results: the Jaccard Index is largest for the CC- and the CP-blocks with average values of .6273 and .6565 (std. devs: .0366 and .0380), respectively, i.e. two thirds of all links are maintained over adjacent quarters. These values are roughly 1.5 times larger than those in the PC- and PP-blocks, with average values of .4261 and .4241 (std. devs.: .0471 and .0622), respectively. Interestingly, we do not find significant evidence of a structural break due to the GFC in any of the time series. Even though the values dropped for most of the time-series after quarter 39 (except for the CP-block), the values tend to stabilize later around the pre-GFC levels. This might indicate that many interbank relationships tended to survive through the GFC.[44]

Overall, our calculations show that the outgoing links of core banks are highly persistent, both with respect to the core and the periphery. Outgoing links from the periphery are persistent as well, but to a significantly lower degree. Given that the Jaccard Index is independent of the density of the network (non-existing links are ignored), these findings indicate that core banks generally lend towards a large set of core and periphery banks. In contrast, periphery banks are not only reluctant to create links among themselves, but also, given the relatively small density in the PC-block, trade with a small set of core banks, which is not necessarily the same set in each quarter. This finding is interesting, since the persistence in the PC-block should be much larger, if periphery banks would have a preferred partner among the core banks. These findings may, however, be driven to some extent by the relatively small trading volumes of periphery banks (see below). In any case, the asymmetry between the CP- and PC-blocks is remarkable and will be discussed in more detail below.

Model Fit and Significance

In this section we turn to a quantitative analysis of the error scores and their significance. When investigating the significance of our results, we compare the core sizes and error scores of the empirical networks with those of network structures sharing similar properties along certain dimensions. This analysis helps us evaluating whether the core-periphery structure offers a meaningful characterization of our data or whether the data rather generate a 'spurious' core by chance.

[43] Given that the coreness vectors are highly autocorrelated, this is not a very restrictive assumption, but it is likely to reduce some noise in the calculated numbers.

[44] Cf. Affinito (2011) and Braeuning (2011) for related evidence on the robustness of lending relationships over the crisis.

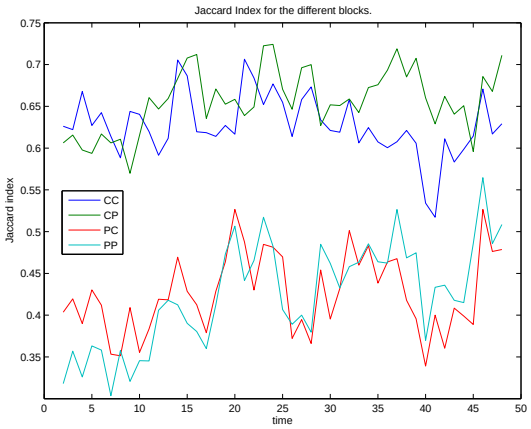

Fig. 5.10: Jaccard Index for the CC-, CP-, PC-, and PP-blocks over time. Coreness is taken from the discrete model. Individual Chow- and CUSUM-tests show no evidence of a structural break due to the GFC in any of the time series. However, the CP- and PP-blocks appear to contain a structural break after quarter 10.

The left panel of Figure 5.11 shows that the error scores (fractions of residuals) are on average roughly 42%, which is rather high compared to the maximum value of 12% for the German interbank market reported by Craig and von Peter (2010). These values are, however, way below unity, so the core-periphery model is indeed a better explanation of the data than an unstructured alternative consisting only of a periphery. We also see that the GFC made the fit somewhat worse, yielding an error score that is roughly 1.3 times the average score before the GFC, albeit with a declining trend. A Chow-test and a CUSUM-test again indicate the existence of a structural break after quarter 39 at all sensible significance levels. The right panel of Figure 5.11 shows that this structural break is mainly due to the increase in the error score in the PP-block. In contrast, we find no evidence for a structural break in the error score of the CC-block after quarter 39, but after quarter 10. Given that the relative core size has been significantly smaller, the overall picture is thus that some previous core banks have reduced their interbank activities so strongly that they are assigned to the periphery after quarter 39. We will investigate the effect of the GFC in more detail in section 5.5.5.

In order to shed light on the significance of the observed error scores, we compute the average core size and error scores by generating 100 random samples of particular network structures (see below) and compare the results

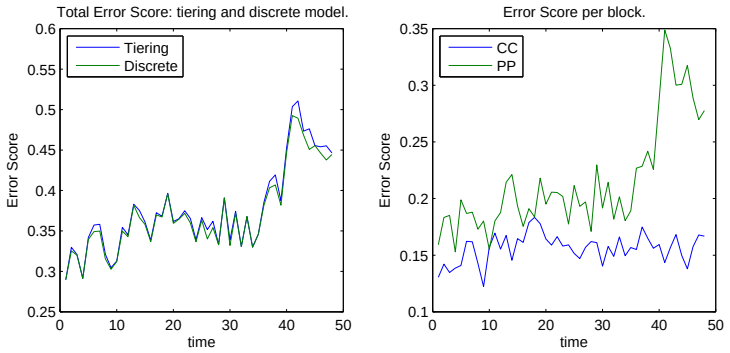

Fig. 5.11: Left: Error score in tiering (blue) and discrete (green) model over time. A Chow-test indicates that there is a structural break after quarter 39 at all sensible significance levels. The results from an additional CUSUM-test are also in favor of the existence of a structural break. Right: Error score for the CC- and the PP-block in the discrete model. For the CC-block there is a significant structural break after quarter 10, while the PP-block contains a significant structural break after quarter 39.

to our findings above. The analyzed networks are:

- Erdös-Renyi (ER) random graphs, where a link is formed with probability p. The value of p will be set equal to the observed density of the network. This network is completely random and we do not expect to find a convincing core-periphery structure in this case. The error scores should be relatively high, since identified cores would be completely spurious. Note that this is tantamount to a bootstrap test for the significance of our identified core-periphery structures, as the random graphs could also be generated by random resampling of the empirical links (with replacement). If the error scores of the core-periphery model are below a certain percentage boundary of those obtained for the sample of random networks, we could exclude with a significance level equal to the inverse of that probability, that our results are spuriously obtained from a completely random system of interbank liabilities. As it turns out, all error scores are always way below the minimum obtained for the random networks.

- Scale-free random graphs, with scaling parameter 2.3.[45] Even though we found the degree distribution not to be scale-free, see Finger *et al.*

[45] In actual interbank networks, the observed scaling parameters vary between 2 and 3.

(2012), most interbank markets appear to have a certain resemblance of their degree distributions to a scale-free distribution. Reported scaling parameters vary between 2 and 3, but are roughly similar for in- and out-degree. We generate these networks using the approach of Goh *et al.* (2001). Note that scale-free networks are assortative by definition, since high-degree nodes tend to connect with each other. Therefore we expect the scale-free network to have a much tighter core and significantly lower error scores.[46]

In the following we only discuss results from the discrete model to save space.[47] Most of the results were expected: all models show a structural break due to the GFC. The actual error scores lie between those from a completely random network (ER) and those of a scale-free network as can be seen from the left panel of Figure 5.12, where we plot the actual error score and the average error scores of the ER and scale-free networks (including plus and minus one standard error for the simulated models). Not surprisingly, the actual network is closer to a scale-free network even though the distance seems to increase with the GFC.

The right panel of Figure 5.12 shows the core sizes for the actual and random networks (again including one standard error for the simulated models). We see that the observed core is significantly larger than both the core of the scale-free and ER networks. For the ER network, the core is spurious, while it would capture the most highly connected nodes in the scale-free network (although the core-periphery model would be a misspecification of the overall structure of such a network).

Overall, we find that we can reject the hypothesis that our findings are just artifacts of applying the core-periphery algorithms to random data, and we also find that the popular scale-free model could not have generated our particular sets of identified cores and fit of the model (error scores).

5.5.3 Continuous Model

We now move to the results from the continuous framework, mostly concentrating on the added explanatory power of the asymmetric version. We have seen in Table 5.1 and Figure 5.3 that the in- and out coreness vectors are mostly negatively correlated. Figure 5.13 shows a scatter-plot of the two variables, explicitly linking the findings to the results of the discrete model.[48]

Here we take the value found by De Masi *et al.* (2006).

[46] Interestingly, Craig and von Peter (2010) found that the error scores for the German interbank market are significantly smaller than for SF networks.

[47] Again we note that the results are almost identical to those from the tiering model.

[48] Recall that the coreness values from the continuous model are standardized values.

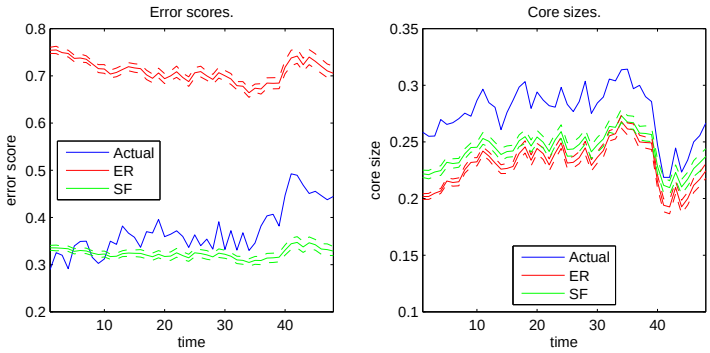

Fig. 5.12: Error scores (left) and core sizes (right) in discrete model. Actual and random graphs. For the SF networks we used a scaling parameter of $\alpha = 2.3$.

Obviously, core banks have on average a higher in- and out-coreness. Indeed, we see a relatively sharp distinction between core and periphery banks. Core banks (red) are typically characterized by a sum of their in- and out-coreness above .2, while this sum is lower for banks assigned to the periphery. For both categories, there might be a dominance of lending and borrowing or a more balanced composition of their transactions. The systemic importance of a bank, in terms of its in- and out-coreness, is therefore not identical in general.[49]

In Figure 5.14, we show the time-varying autocorrelations of the two coreness vectors. The autocorrelations were calculated as the correlation between two subsequent coreness vectors, using only banks that were active in both periods. We see that both the in- and out-coreness vectors are highly autocorrelated (average values: .8474 and .9186, respectively). We also calculated cross-correlations between the two vectors, where In-Out (Out-In) is the correlation between in-coreness in $t-1$ (t) and out-coreness in t ($t-1$). These cross-correlations are significantly lower with slightly negative average values of -.0698 and -.0764, respectively. Thus, lagged values of one coreness vector are not very informative for the expected value of the other coreness vector in the next period.

An important question is by how much the fit of the model improves by using the AC model rather than the SC model. As a rule of thumb, Boyd *et al.* (2010) argue that the PRE of the SC model should be at least .5 in order to have a superior fit to an unstructured distribution of activity. Here we find values around .2 for the SC model, but higher values of around .58 for

[49] An example of a fitted network matrix is shown in Appendix 5.8.4.

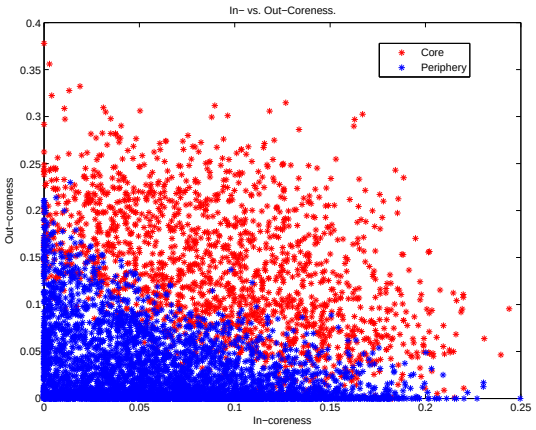

Fig. 5.13: In-coreness vs. Out-coreness for all observations, by core and periphery, as indicated by the discrete model.

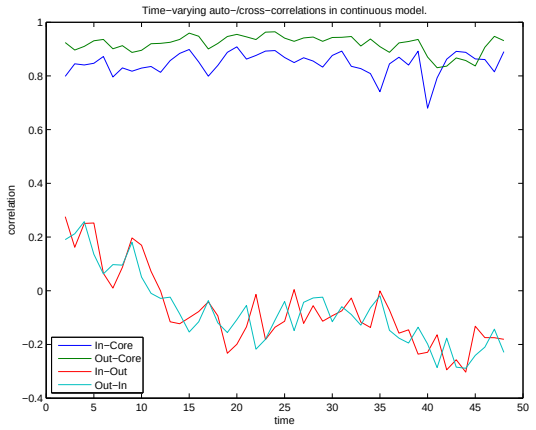

Fig. 5.14: Persistence of coreness vectors. The plot shows the auto-correlations and cross-correlations of the two vectors over time. The autocorrelation is simply the correlation of the coreness vector in t with the one in $t-1$, using only the banks active in both periods. The cross-correlations are the correlations between in-coreness in $t-1$ and out-coreness in t (In-Out), and vice versa (Out-In).

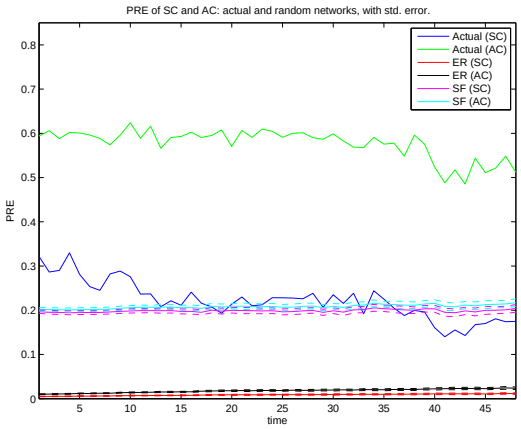

Fig. 5.15: PRE for the SC and the AC model, actual and random graphs. A Chow-test indicates that there is a structural break after quarter 39 at all sensible significance levels for the PRE of the AC model. The results from an additional CUSUM-test are also in favor of the existence of a structural break. The PRE of the SC model appears to display an additional structural break after quarter 10. For the SF networks we used a scaling parameter of $\alpha = 2.3$.

the AC model (cf. Figure 5.15).[50] Similar to the discrete and tiering models above, the fit of the model deteriorates somewhat with the GFC, with lower average values afterwards. In line with the previous findings for the discrete and tiering model, the PRE of the AC model displays a structural break after quarter 39 (based on a Chow-test and a CUSUM test), but not in the SC model.

In order to check the significance of the PREs, we use a similar approach as in the previous section on the discrete and tiering model, however, here we use the (valued) network of interbank liabilities.[51] Figure 5.15 compares the PREs of the actual networks with the mean values from 100 realizations of random ER and SF networks (again with scaling parameter 2.3) minus and plus one standard deviation. As expected, the actual PREs of the SC

[50] Obviously the fit has to be better in the AC model, since we have twice as many parameters. Interestingly, the fit is mostly more than twice as good as the fit of the SC model.

[51] In this approach, we generated random ER and SF networks as explained above. Then, we randomly assigned observed transaction volumes from the actual networks (log-transformed) to the random ones. The results are essentially identical with and without replacement. Here we present the results without replacement.

and AC models significantly exceed those from the ER networks, which are very low in general. In contrast, for the SF networks, the PREs of the SC model are close to the actual ones, while this is not true for the AC model.[52] This finding underscores the observed asymmetries in the network, which are absent from scale-free networks, where in- and out-degrees of individual banks are highly correlated by construction.

In comparison with the closeness of the error scores of the empirical data and their scale-free resamples in Figure 5.12, the consideration of the asymmetries of the concentration of incoming and outgoing links shows the limitations of the scale-free networks. While it appears reasonably similar to a symmetric core-periphery framework, it falls back behind the asymmetric continuous CP model at all levels of significance. Since the fit of the two-dimensional continuous approach (AC) is way better than that of the one-dimensional continuous approach (SC), we conclude that the directed version of the model contains important information about the structure of the interbank market.

5.5.4 What Defines a Core Bank?

In the following we will focus on the results from the discrete model.[53] As a first step, we calculate the correlations between the coreness vectors and different observable variables (degree, size, and trading activity).[54] It would also be very interesting to forecast the coreness vectors based on non-network-related observable variables, e.g. balance sheet size. Due to the anonymity of the data, such an analysis is, unfortunately, not possible.

Figure 5.16 shows the correlation between the discrete coreness vectors and the in-, out-, and total degree (total degree is the degree from the undirected version of the network), respectively. The correlation is far higher for

[52] Note that the PREs of the AC model are always larger than those from the SC model, both for the actual and the random networks (even though for the random networks not always significantly). This is driven by the higher number of parameters (degrees of freedom) in the AC model.

[53] Again the results for the tiering model are very similar and available upon request.

[54] It would be interesting to analyze the interest rates charged in the different blocks in more detail. This is, however, beyond the scope of this paper. Here we just note that the (volume-weighted) average interest rate charged in the CC-block exceeds that of all other blocks (average value of 2.72% for the complete sample). Thus, it seems that core banks price in the systemic importance of other core banks, while giving more favorable prices to periphery banks (average value of 2.70%). Furthermore, the average interest rate charged between periphery banks tends to be quite small as well (average value of 2.71%). Thus, two periphery banks may grant each other more favorable prices as soon as they trade with each other on a regular basis. Note that, due to the non-stationarity of the interest rates, we checked the significance by comparing the first differences.

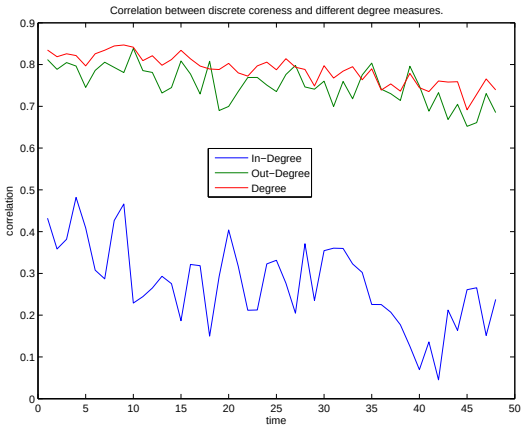

Fig. 5.16: Time-varying correlation between discrete coreness and in-degree/out-degree/total degree. Total degree is the degree we would obtain from transforming the directed network into an undirected network.

out-degree compared to in-degree, and the former has practically the same correlation with coreness as the total degree. Hence, it is the distribution of liquidity rather than its absorption, that identifies the core banks in our sample.

We constructed similar measures for the individual sizes and the number of transactions per bank, see Figures 5.17 and 5.18, respectively. We proxy the bank size by the transaction volumes in a particular quarter.[55] Here, in-size contains the total volume of borrowing transactions (per quarter and per bank), out-size the total volume of lending transactions and total size the sum of in-size and out-size. Similarly, the number of in-transactions (out-transactions) is the number of borrowing (lending) transactions per bank. The total number of transactions is the sum of the two. We find that the core banks are significantly larger and more active than periphery banks (unreported). However, the size measure appears to be a less reliable indicator than the simple number of transactions, since it is far more volatile. Both measures, however, confirm again the dominant aspect of the out-direction (lending activity) for the core membership of a bank.

We also constructed the same figures for the continuous model, see Appendix 5.8.5. As expected, the two coreness vectors can be better explained based on the directed version of the network. Most importantly, the correlation with the total degree is smaller compared to the correlation of in-coreness

[55] See De Masi *et al.* (2006).

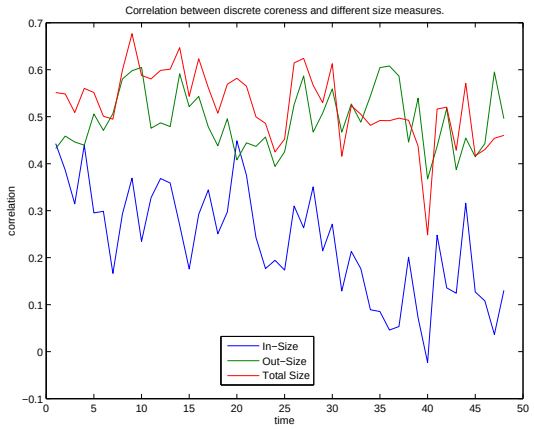

Fig. 5.17: Time-varying correlation between discrete coreness and in-size/out-size/total size as defined in the text.

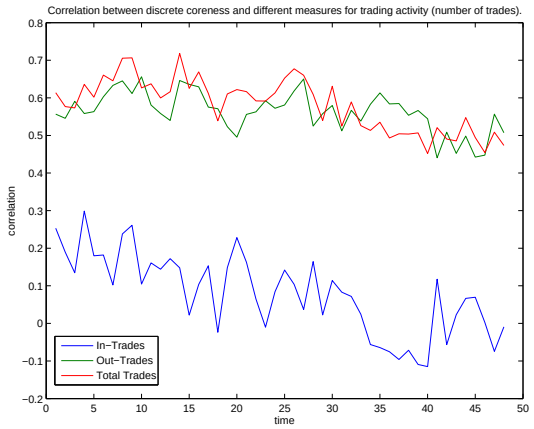

Fig. 5.18: Time-varying correlation between discrete coreness and number of in- and out transactions. Total number of transactions is the sum of the two.

with in-degree and out-coreness with out-degree, respectively. Again, the correlations with the size measures are highly volatile.

We conclude that all measures point towards the lending activity as the more relevant aspect of core banks' participation in the market. The much lower relevance of their borrowing activity, then, explains why in- and out-coreness vectors in the asymmetric model are virtually uncorrelated.

5.5.5 What Happened During the GFC?

In this section, we provide a more detailed analysis of the effects of a major shock to the interbank network, namely the collapse of Lehman Brothers in quarter 39. So far, our analysis shows that the GFC indeed had a substantial impact on the network along many dimensions, in particular in terms of the goodness-of-fit of the core-periphery models. To investigate the effects of the structural break in more detail, we split our sample into a short pre-crisis period (quarters 37 and 38) and a post-crisis period (40 and 41).[56] Interestingly, despite the clear negative trend in the number of active banks during the complete sample period (cf. Figure 6.1), the actual number during the analyzed subperiod is relatively stable with an average value of 98 banks. Thus the network sizes during this particular period are comparable, which allows to compare different network-related measures. As a first step, we will investigate network-related variables from a macro perspective. Then we take a closer look at the behavior of one particular exemplary core bank around the breakpoint.

As we have seen (cf. Figure 5.6), the GFC affected the block-structures of the discrete and the tiering model: Core banks trade significantly less with each other (density in the CC-block smaller), and so do periphery banks (density in the PP-block smaller). In contrast, there is no evidence for a significant structural break in the densities of the off-diagonal blocks. Core banks also tend to lend less money to the periphery (density in the CP-block smaller), while there is no clear trend in the amount that peripheral banks lend to the core, thus the periphery tends to maintain their links to the core during and after the crisis. Given that the GFC, and the resulting tensions in money markets, can be seen as the result of a crisis of confidence, it comes as no surprise that core banks tend to reduce their risk exposure by cutting down the number of links going both to core and periphery banks.[57] Concerning

[56] Of course one could argue that the pre-GFC period should be further away from the breakpoint, however, here we are particularly interested in the network changes right at the phase transition.

[57] Interestingly, the number of reciprocal links, i.e. the fraction of links pointing in both directions, goes down due to the GFC. This is somewhat surprising, since we would expect

the market activity, we find that the total trading volumes (and also the total number of trades) in the CC- and the CP-blocks dropped substantially during the crisis, while it actually increased in the PP-block immediately after the GFC but then dropped substantially. In contrast, after a sharp drop of market activity in the PC-block right before the GFC, the total amount of credit flowing from the periphery to the core actually increased after the GFC.[58] Thus it seems, that the crisis mainly affected the behavior of core banks, which rather hoarded their liquidity than providing it to a large number of other counterparties.[59] In contrast, periphery banks tend to keep (at times even expand) the number of outgoing links with core banks, while reducing the exposure to the periphery. The findings on the Jaccard Index of the PC-block (cf. Figure 5.10), however, indicate that periphery banks do not necessarily lend money to the same core banks over time. Overall, from the relatively stable Jaccard indices it appears that no major disruption of the network pattern occurred (cf. section 5.5.2), but that the aggregate volume of lending by core banks has declined substantially. Hence, most of the network structure remained intact, but continued its operations at a much lower level of activity. This finding speaks in factor of a positive effect of relationship lending that helped to prevent a complete collapse of the interbank market after the onset of the financial crisis (as suggested by Affinito (2011), and Braeuning (2011)).

To illustrate the generally observed tendencies, we picked the (core) bank with the highest aggregate trading volume.[60] During this period, the particular bank had an average in-degree of 30, while its average out-degree was substantially higher with 64.[61] These mean values, however, hide the dynamic development, since there was a sharp drop in the banks' out-degree during the GFC (the maximum level pre-GFC was 80, the minimum level at the end of the period is merely 32), while the in-degree actually increased during the crisis (the minimum level pre-GFC was 15, the maximum level at the end of the period is 45). In Figure 5.19, we split up the bank's links into outgoing links to core and periphery banks, respectively, and the same for the incoming links during the period under study. We see that the bank had

that bilateral relationship become closer in crisis times.

[58] The increase in the number of trades in the PC-block after the GFC is even more impressive, ending up above the pre-GFC level.

[59] Interestingly, core banks lend more money than they borrow from the periphery, thus the core is a net lender to the periphery.

[60] In fact, this bank (ID number 'IT0278') was in the core during the complete sample period.

[61] These numbers just underline the observed asymmetry between the CP- and PC-blocks.

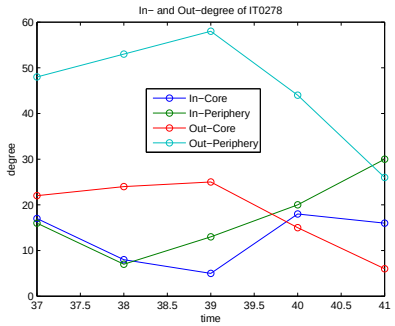

Fig. 5.19: In- and out-degrees of bank IT0278, by core and periphery. In-core gives the number of incoming links from other core banks, In-periphery the number of incoming links from periphery banks. Out-core gives the number of outgoing links to other core banks, Out-periphery the number of outgoing links to periphery banks.

reduced the number of outgoing links, both with core and periphery banks, but it had increased the number of incoming links, in line with the overall tendencies.[62] Interestingly, while the bank was a net-lender during most of the sample period, we see that the bank actually reversed its strategy during the GFC, since it became a net-borrower afterwards (see Figure 5.20). Thus, the bank tried to attract liquidity, mainly from periphery banks, since core banks became reluctant to trade with other core banks.

Summing up, we conclude that the GFC both affected the behavior of core and periphery banks: Periphery banks seem to have increased their lending to the core, both in terms of the number of links and trading volumes. In contrast, core banks have reduced their lending, not only to other core banks, but also to the periphery. The decline in goodness-of-fit of the core-periphery structure is therefore mostly due to a loosening of the core. Core banks activated a smaller number of their previous outgoing links. Hence they started to hoard liquidity rather than distributing it in the system. Therefore, it seems that core banks tend to rely on the liquidity of periphery banks during times of distress, while in 'normal' times they would more freely redistribute liquidity in the complete system.

[62] It would be interesting to see the quote data, rather than the transaction data. We suspect, that many quotes are simply never executed during the GFC.

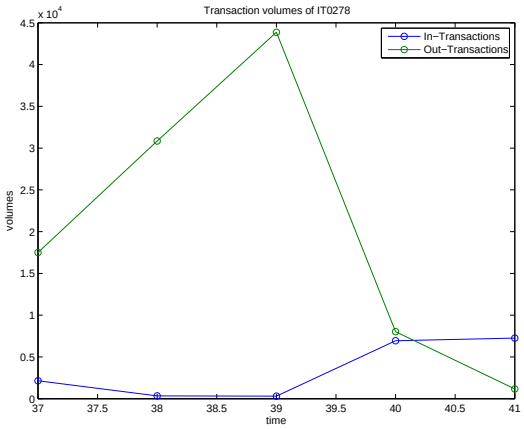

Fig. 5.20: Transaction volumes of bank IT0278. In-transactions gives the total amount of credit borrowed by the bank, while Out-transactions gives the total amount of credit lend by the bank to other banks.

5.6 Discussion

The majority of studies on the structure of interbank networks has hitherto concentrated on the distribution of degrees. Many authors mention the finding of some form of community structure in the interbank market, suggesting a tightly connected core of money-center banks.[63] The finding of a core-periphery structure in the Italian interbank market can be seen as a special case of community structure,[64] where the core is a tightly connected part of the network, and the periphery is the loosely connected component.[65] Even though we only know of only one other study in this regard, it may well be that the finding of a core-periphery structure could be seen as a new 'stylized fact' of modern banking systems. As far as data are available, it would be important to test this hypothesis in other interbank networks.

An important question is of course why we find a core-periphery structure in the interbank market. In the literature on social network analysis, two main explanations for the emergence of a core-periphery structure exist: (1)

[63] See Iori *et al.* (2006) and Soramäki *et al.* (2007).

[64] Note that communities are usually defined as very dense subgraphs, with few connections between them. The periphery is thus more of an anti-community.

[65] We also checked several standard community detecting algorithms for the Italian interbank network. The main finding is that, for the entire market, we find two separate communities consisting of foreign and Italian banks, respectively. Interestingly, it is impossible to split these communities further into smaller subcommunities. Thus it seems even more remarkable that we find a core-periphery structure in this market.

'Superior' core members possess an intrinsic advantage over the 'inferior' periphery members, such that the core exerts power over the periphery.[66] In order for a core-periphery structure to emerge, the advantage of the core members must be reflected in attributes affecting the linking behavior of all agents.[67] Then core agents would be able to translate their advantage into a positional advantage in the social network.[68] Transferring this idea to banking networks, one encounters several problems. First, it is not clear a priori which attributes might make core banks 'superior' to the periphery. We would also need to come up with an explanation why core banks share attributes that periphery banks do not have. Note also that this definition implies that it should be preferable for all banks to be part of the core, which is not very plausible. For example, a small bank (in terms of its balance sheet) would find it hard to intermediate between other core banks, simply because it does not command a sufficient amount of funds to do so. Therefore, this bank will always prefer being in the 'inferior' periphery, where it still might intermediate between other small banks. Furthermore, the general finding of disassortative mixing patterns in banking networks[69] is not in line with the power-based explanation, since core banks would then be reluctant to create links with periphery agents. Nevertheless, if we define power as the ability of influencing the market, it may well be that it is an important driver for the emergence of a core-periphery structure in the banking network.

(2) Core members have a comparative advantage in gathering (and spreading) information about other members of the network.[70] Thus, information costs are higher for periphery-periphery relationships compared to core-periphery relationships (in both directions). For the banking network, this would mean that periphery banks have an incentive in cutting down the number of links to other periphery banks, maintaining only a few links to core banks. Core banks on the other hand connect among themselves and to periphery banks.[71] This explanation would not only be in line with the disas-

[66] See Persitz (2009).

[67] For example, in a scientific network, the core agents are the highly productive agents being cited by many others. See Mullins *et al.* (1977).

[68] Persitz (2009) provides a formal model for a power-based core-periphery network. The basic idea is that linking preferences are such that all agents prefer establishing links to 'superior' agents relative to 'inferior' agents.

[69] See Finger *et al.* (2012).

[70] For banks, the comparative advantage may stem from economies of scope and scale, but also from very frequent interactions on the market which small periphery banks usually do not have.

[71] Note that, despite the overall disassortative mixing patterns, the core-periphery structure indicates that we should actually differentiate between these patterns in the core and the periphery: the periphery mostly shows disassortative mixing within itself, while the

sortative mixing patterns, but also with the evidence in Cocco *et al.* (2009): small banks, with limited access to international capital markets and possibly limited investment/financing opportunities due to their more locally oriented business model, tend to rely on preferential relationships with (large) core banks. Thus, core banks act as intermediaries between different parts of the periphery of the domestic banking system, resulting in indirect relationships between peripheral banks. Note that this explanation is also in line with the observed asymmetry between the densities in the CP- and the PC-blocks, since they imply that periphery banks cut down their credit risk by focusing on a few selected core banks, while they are prepared to borrow money from a larger set of core banks.

Finally, we would like to focus on the potential implications of our findings for regulators. It is well known that the structure of a network is important for its resilience, hence policymakers should be interested in the actual topology of the interbank network. For stress-testing exercises, it would, therefore, be crucial to use a topological description of the connections within the banking sector that is both realistic and computationally tractable. Most stress-testing scenarios have actually adapted an entropy-maximization approach for filling the unknown matrices of interbank liabilities.[72] This means, that given some overall statistics for the whole system, interbank credit is spread as evenly as possible across the system[73] An idealized core-periphery structure amounts to pretty much the opposite in terms of concentration of interbank liabilities. If the data were closer to the latter type of structure, the entropy-based approach could give misleading results for the expected aftereffects to shocks affecting single institutions. If, as we believe, the core-periphery structure turns out to be a stylized fact of the interbank market, stress-tests should take this particular topology into account. Unknown amounts of interbank liabilities could then be calibrated along the structural features of typical core-periphery models for available data (like those of the present paper and Craig and von Peter, 2011). As our results show, it might also be important to take into account asymmetries in the borrowing and lending attitudes of core banks. Even when comparing the effects of shocks between different network models with some tendency of concentration of links, important differences might exist. For example, networks with scale-free degree distributions are known to be robust with respect to random failures, but fragile with respect to targeted attacks on the most central

core shows more of assortative mixing among its members, since core banks tend to connect among themselves.

[72] See Sheldon and Maurer (1998), and Upper and Worms (2004).

[73] Note that this is equivalent to the benchmark against which the error reduction by the continuous core-periphery model is measured.

nodes.[74] The usual mechanism to construct scale-free degree distributions is that of preferential attachment, see Barabasi and Albert (1999), so high-degree nodes tend to attract more links than low-degree nodes over time. As described above, we did not find evidence for scale-free degree distributions in the e-MID data and also find disassortative rather than assortative mixing patterns. Comparing assortative to disassortative networks, Newman (2002) shows that, for the same degree distribution, assortative ones are more robust to targeted attacks compared to disassortative ones. Since an assortative network possesses a whole set of nodes with large in- and out-degrees, i.e. many connections across the entire network, the system is characterized by a certain degree of redundancy that makes it more robust under attacks on single highly connected nodes. In contrast, the disassortative network is more susceptible to removal of high-degree nodes, which are not as tightly connected as in the assortative case. Thus, removing high-degree nodes allows to attack different parts of the network.

In a somewhat related strand of research, Brede and de Vries (2009) show that core-periphery structures might emerge from an evolutionary process as a compromise between resilience (concentration makes the network more vulnerable) and efficiency of a network (concentration creates short average path lengths).[75] From an economic point of view, the question would be whether the self-organization of the interbank network into a core-periphery structure creates important externalities so that policymakers should attempt to shift the balance towards higher resilience and somewhat lower efficiency.

Of course, regulators should also be interested in the dynamics of a network, when the breakdown of one node has knock-on effects on other nodes. This contagion effect is for example investigated by Caccioli *et al.* (2011) for different network structures. The authors analyze the extent of contagion in artificial banking systems after the random failure of individual institutions. Their main finding is that the likelihood of contagion, i.e. the breakdown of the entire system, is smaller for disassortative networks. Since in the latter, high-degree nodes tend to connect with low-degree nodes, the failure of a random node is unlikely to spread through the entire system. Conversely, the random breakdown of a high-degree node will severely affect other high-degree nodes in assortative networks. Note that this is different from the aspect of vulnerability under targeted attacks. As a consequence, a disassortative core-periphery framework might be more robust in 'normal' times, but more fragile under exceptional circumstances when key nodes are under

[74] See Albert *et al.* (2000).

[75] Note that the highest efficiency is realized in star-like configurations, while the highest resilience is related to the avoidance of short loops and degree homogeneity. See also Netotea and Pongor (2006).

stress or withdraw from the market. Hence, the 'coreness' translates to a certain extent into 'systemic relevance' of certain institutions.[76] The GFC seems to have been a major shock to the interbank network, as tests for structural breaks indicate. The observation that the fit of the core-periphery models significantly worsened with the GFC, might provide important information per se on the endogenous reaction of the system to stress which could be incorporated in stress-test scenarios. Furthermore the goodness-of-fit of the core-periphery framework might be seen as an indicator of tensions in the interbank market, so that various statistics based upon such a framework could be used as early warning signals of impending crises.

5.7 Conclusions

The main findings of our paper are the following: we find a significant core-periphery structure in the Italian interbank network for a sample period from January 1999 to December 2010. The identified core is quite persistent over time, consisting of roughly 28% of sample banks before the GFC and 23% afterwards (discrete model). Given the substantial differences in the German and Italian interbank market data investigated by Craig and von Peter (2010) and the present paper, e.g. in the underlying region and the maturity structure of the credit relationships, the finding of a core-periphery structure is unlikely to be a coincidence. We expect that other interbank markets display a similar hierarchical structure, which might be classified as a new 'stylized fact' of modern interbank networks and actually concretizes on a system level the role of money center banks. Going beyond the analysis of Craig and von Peter (2010), we also investigate the continuous and asymmetric versions of core-periphery models and find evidence for strong asymmetries. In particular, overall coreness is mainly driven by the function of provision of liquidity to large parts of the banking system by the core members. Overall coreness is, therefore, largely identical to out-coreness, while its connection to in-coreness is very weak. Regulators should be aware of the fact that a bank which is part of the in-core but not of the out-core, may play a completely different role in the system than a bank with the reverse strategy.

Formal tests favor the existence of a structural break in the last quarter of 2008, the time when Lehman Brothers collapsed. We investigated this time period in more detail and found that the deteriorating fit of the core-periphery structure in the post-GFC period is mainly due to the loosening of connections in the core, particularly on the lending side. Furthermore,

[76] See also Markose et al. (2010).

it seems that during times of distress, core banks tend to rely on periphery banks as an important source of funding, since other core banks are reluctant to provide as much liquidity to other banks as in normal times.

Our findings provide some support for the view that the network structure is non-random due to the existence of preferential lending relationships. This is in line with the results of Cocco *et al.* (2009), Affinito (2011), and Braeuning (2011). Further evidence in this regard is provided by Finger and Lux (2013), who analyze the evolution of the banking network using the actor-oriented approach by Snijders (1996, 2001). The general conclusion is that preferential lending relationships at the micro-level lead to hierarchical structure at the macro-level. An open question is why the interbank network shows such a hierarchy. We argue that the comparative advantage of core banks in gathering and distributing information about their counterparties is likely to be a crucial factor.

In the future we plan to apply the model to other interbank data, in order to evaluate whether the core-periphery structure is indeed a new stylized fact of banking systems. Furthermore, it would be interesting to relate the results to bank-specific variables, such as individual balance-sheet data. In any case, this approach can be seen as a contribution to identifying the systemically important banks in a quantitative way. We also believe that the methods presented here could be an important tool for regulators since they allow to reduce the complexity of large-scale network data, and to represent the salient structural features of the complicated web of dispersed activity in the interbank market in a compact way.

5.8 Appendix

5.8.1 Genetic Algorithm

The GA maintains a population of L solution candidates and evaluates the quality of each solution candidate according to a problem-specific fitness function, which defines the environment for the evolution. New solution candidates are created by selecting relatively fit strategies which are then recombined through various genetic operators:

1. **Reproduction:** From the pool of current solutions, L copies are selected (with replacement) randomly with probabilities depending on a solution's relative fitness, i.e. prob(reproduce solution j) $= f_j / \sum_i f_i$ with f being the fitness function as defined below (tournament selection).

2. **Crossover:** Copies from the reproduction step are randomly paired in a mating process with each couple producing two offspring via exchange of genetic material. The simplest way is to select two copies randomly and swap bits between both of them. Here we randomly select an integer in the range of $[1, L-1]$ and construct offsprings by combining the left of this position from parent one with that from the right-hand part of parent two and vice versa. The cross-over operation is carried out with probability π_{cross}, while with probability $(1 - \pi_{\text{cross}})$ the offspring are unchanged copies of their parents.

3. **Mutation:** After exchanging genetic material, slightly different solutions are formed by altering each position within a string with probability π_{mut} to the other value of the binary alphabet.

4. **Election:** The election operator avoids an overall decrease of fitness in the population by allowing only offspring with a higher fitness than their parents to the new generation.

We also add an operator that we call 'genetic engineering' (GE) to accelerate the convergence of the algorithm.[77] GE selects the binary string with highest fitness value of the current population so far and tries to improve its fitness value as follows: we search the core member with the lowest connection to the core in this partition (i) and the periphery member with the highest connection to the core (j).[78] Then we compare the fitness of the current partition with three alternatives: (1) flip i's core membership from 1

[77] See also Stolzenburg and Lux (2011).

[78] This is done by computing the average connections with all core members.

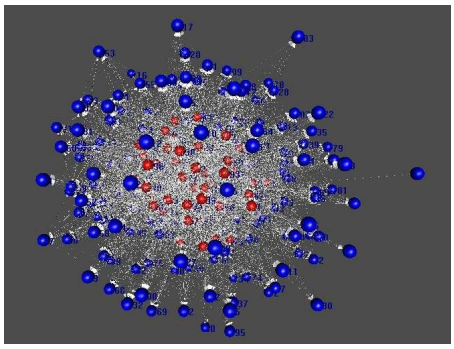

Fig. 5.21: Example for the core-periphery structure of the Italian interbank network, 1999Q1. Red dots show core, blue dots periphery banks. Dotted lines show directed edges from one bank to another.

to 0, and j's core membership from 0 to 1, (2) only flip i's core membership, (3) only flip j's core membership. Using the best of these alternatives, we put the resulting string back into the evolving population. By using this additional operation we can manipulate our population target-oriented and do not have to rely solely on blind exploration of the search space by the random evolutionary steps of 'mutation' and 'crossover'.

Summing up, the binary-coded GA has only two parameters π_{cross} and π_{mut}. Obviously, the fitness function is the crucial element of the GA. In our setting, the fitness of each solution depends on the corresponding error score e. Since the optimal solution would have zero errors, we simply take the rate of 'correct' classifications as the fitness function, i.e.

$$f_l = 1 - e(\mathcal{C}), \tag{5.12}$$

with the error score defined in Eq. (5.3).

5.8.2 Discrete Model: Illustration

Figure 5.21 illustrates the outcome of the estimation of the discrete core-periphery model for the first quarter of 1999. We clearly see that the core banks (red dots) form the cluster of the most central nodes, with periphery banks (blue dots) connecting to parts of this cluster.[79]

[79] Note the high network density and the existence of only a single network component.

5.8.3 Empirical Estimation of the Asymmetric Continuous (AC) Model

This section summarizes Boyd *et al.*'s (2010) approach for estimation of the AC model.

Problem Formulation

We first note the similarity between the optimization problem of Eq. 5.9 in the main text, and a Singular Value Decomposition (SVD) of a matrix. SVD allows to decompose a matrix $\mathbf{D}_{\{M\times N\}}$ of rank r into

$$\mathbf{D} = \mathbf{USV}', \tag{5.13}$$

where $\mathbf{U}_{\{M\times r\}}$, $\mathbf{V}_{\{N\times r\}}$ are real matrices with orthonormal columns, with the columns corresponding to the singular vectors, and $\mathbf{S}_{\{r\times r\}}$ is a diagonal matrix containing the singular values (ordered) on the main diagonal.[80]

For any $k \leq r$, SVD gives the best (least-squares) rank-k approximation of $\mathbf{D}$, i.e.

$$\mathbf{D}_{(k)} = \mathbf{U}_{(k)}\mathbf{S}_{(k)}\mathbf{V}'_{(k)}, \tag{5.14}$$

with $\mathbf{U}_{(k)}$ and $\mathbf{V}_{(k)}$ as the first k columns of $\mathbf{U}$ and $\mathbf{V}$ respectively, and $\mathbf{S}_{(k)}$ as the diagonal matrix formed by the first k singular values.

Singular Value Decomposition (SVD) is defined for rectangular matrices, for which symmetry is not an issue, but it can also handle square matrices, whether symmetric or asymmetric, as a special case.[81] However, it does require the presence of the diagonal elements of a matrix. Thus, by definition, SVD can handle asymmetric data matrices.

Estimation of the AC model is again performed via minimization of residuals (MINRES) taking stock of the proximity of the problem to a SVD. The basic idea of MINRES/SVD is to use a rank-1 approximation of $\mathbf{D}$, using the first singular value s and the first singular vectors u and v: $\mathbf{D}_{(1)} = usv'$.[82] Since SVD requires diagonal elements, we use a SVD of rank 1 and, similar to the MINRES approach, exclude the diagonal elements in the analysis. Thus, our objective function for the AC model looks as follows

$$\arg\min_{u,v} \sum_i \sum_{j\neq i} (d_{ij} - u_i s v_j)^2. \tag{5.15}$$

[80] $\mathbf{U}$ contains the eigenvectors of $\mathbf{D}\ \mathbf{D}'$ while $\mathbf{V}$ contains the eigenvectors of $\mathbf{D}'\mathbf{D}$. The diagonal elements of $\mathbf{S}$ are the square roots of the non-zero eigenvalues of $\mathbf{D}\ \mathbf{D}'$.

[81] See Stewart (1993) for an overview.

[82] In the future it would be interesting to look at higher dimensional approximations. This would allow splitting up the core and periphery even further.

The normality constraints on u and v can be eliminated, such that we can neglect the singular value s by absorbing it into the unconstrained vectors u and v. Obviously, the solution is not unique, but without s the model is even simpler, since we approximate $\mathbf{D}$ using only uv', leaving us with the objective function

$$\arg\min_{u,v} \sum_i \sum_{j \neq i} (d_{ij} - u_i v_j)^2. \tag{5.16}$$

The optimal vectors can be determined by finding the roots of the first-order conditions of Eq. (5.16). The original u, s and v can be obtained by defining $s = ||u||||v||$, with $||u|| = \sqrt{\sum_i u_i^2}$ being the Euclidian norm of u, and then dividing u and v by their norms. The reported coreness vectors are normalized.

Optimization Problem for MINRES/SVD

We could solve the problem of finding the vectors $u_{\{N\times 1\}}$ and $v_{\{N\times 1\}}$ numerically by using standard optimization procedures. However, following Boyd *et al.* (2010), the problem can be solved easier by setting the first derivative of Eq. (5.9) with respect to u_i and v_j equal to zero and solving the resulting equations numerically.
More formally, this amounts to

$$\frac{\partial L}{\partial u_i} = \sum_{j \neq i}^{N} (d_{ij} v_j - u_i v_j^2) = 0. \tag{5.17}$$

Remembering that the diagonal elements in $\mathbf{A}$ equal zero, we can write this as

$$\sum_{j=1}^{N} d_{ij} v_j = -u_i v_i^2 + \sum_{j=1}^{N} u_i v_j^2. \tag{5.18}$$

For each row i this equation has to hold, so we can write the set of equations in matrix notation as

$$\mathbf{D}v = u.(-v.^2 + v'v), \tag{5.19}$$

where a dot indicates elementwise multiplication. The other set of equations can be calculated in a similar fashion

$$\frac{\partial L}{\partial v_j} = \sum_{i \neq j}^{N} (d_{ij} u_i - u_i^2 v_j) = 0, \tag{5.20}$$

leading to

$$\sum_{i=1}^{N} d_{ij} u_i = -u_j^2 v_j + \sum_{j=1}^{N} u_i^2 v_j. \tag{5.21}$$

Now, this equation has to hold for each column j, which can be written in compact form as

$$u'\mathbf{D} = v'.(-u.^2 + u'u)'. \tag{5.22}$$

Consequently, the optimal vectors u and v can be obtained by solving Eqs. (5.19) and (5.22) simultaneously. Using an appropriate set of initial values (see below), the optimization is much faster than solving Eq. (5.15) directly.[83]

We should also note that we checked an alternative approach proposed by Boyd *et al.* (2010), where we impute values on the diagonal and apply the usual one-dimensional SVD to this matrix. The results from this approach cannot be distinguished from those presented in the following, so this approach is likely to be more efficient when working with very large networks.

Initial values for MINRES/SVD

The choice of the initial values is important in many numerical problems, most importantly with respect to computation time. Here we follow the approach in Boyd *et al.* (2010) and impute diagonals to the data matrix first, then using the first step in the reciprocal averaging method for computing the SVD. The algorithm will then work on the original data matrix, without diagonal elements.

Let c_i, r_j and t be the column, row and total sums of the matrix $\mathbf{D}$, respectively, excluding the diagonal elements.[84] A single missing value at the position d_{ij} could then be imputed by assuming independence. This leads us to $(r_i + d_{ij})(c_j + d_{ij}) = d_{ij}(t + d_{ij})$ or solving for the missing entry

$$d_{ij} = \frac{r_i c_j}{t - r_i - c_j}. \tag{5.23}$$

If all of the diagonal elements were missing, one could use this formula to estimate each of the diagonal elements. However, this neglects the contribution of the other $N - 1$ diagonal elements to the total sum t. So a better approximation would be to estimate the sum of all the matrix elements by adding to t an estimate for the other $N - 1$ diagonal elements, the average value of the off-diagonal elements, $t/(N^2 - N)$. After canceling the factor $N - 1$, the independence model for estimating the diagonal elements appears as $(r_k + d_{kk})(c_k + d_{kk}) = d_{kk}(t + t/N + d_{kk})$ for each element d_{kk}.[85] This leads

[83] We have used a very similar approach for the SC model, where we can also speed up the estimation by taking the first derivative of Eq. (5.7) with respect to c and solving the resulting system of equations numerically. See Boyd *et al.* (2010).

[84] For simplicity, if these were zero the sums would remain unaffected.

[85] There is a slight error in the version by Boyd *et al.* (2010), since they missed the d_{kk}^2 term on the right-hand side of the Equation.

to

$$d_{kk} = \frac{r_k c_k}{t + t/N - r_k - c_k}. \tag{5.24}$$

The reciprocal averaging method is analogous to the power method for computing eigenvectors. It works as follows: choose initial vectors x_0 and y_0 to be $x_0(i) = y_0(i) = 1, i = 1, \cdots, N$. Then the iterative formulas

$$\begin{aligned} \tilde{x}_k &= \mathbf{D} y_{k-1}, \;\; x_k = \tilde{x}_k / ||\tilde{x}_k|| \\ \tilde{y}_k &= x_{k-1} \mathbf{D}, \;\; y_k = \tilde{y}_k / ||\tilde{y}_k||. \end{aligned} \tag{5.25}$$

give a sequence of vectors such that x_k and y_k converge to the first singular vectors u and d, respectively. A good approximation for usv' would be $x_2 \mathbf{D} y_2'$. However, we do not specify the singular value but absorb it into the vectors u and d, which are now not normalized. Thus the initial vectors u_0 and v_0 are $\sqrt{s}x_2$ and $\sqrt{s}y_2$, where a good approximation for d is

$$d_0 = \frac{r}{||r||} \mathbf{D} \frac{c'}{||c||}. \tag{5.26}$$

5.8.4 Model Fit: AC Model

Figure 5.22 shows an example of a fitted network matrix, where the matrix has been sorted according to the individual sums of the in- and out-coreness values. The Figure illustrates the typical asymmetry of the in- and out-cores, since most of the core banks lend money to the periphery (out-core), while there are fewer connections from the periphery to the core.

5.8.5 What Defines a Core Bank (in the Continuous Model)?

Figures 5.23-5.25 show the results for the continuous model. Obviously, the two coreness vectors can be better explained based on the directed version of the network. Most importantly, the correlation with the total degree is smaller compared to the correlation of in-coreness with in-degree and out-coreness with out-degree, respectively. Again, the correlations with the size measures are highly volatile.

These results suggest that, in contrast to the findings of Craig and von Peter (2010), total balance sheet size may not be as informative for explaining banks' coreness as we might expect. When splitting the banking network into in-core, out-core and periphery, we should rather focus on asymmetric figures. Examples might be loans granted on the asset side of the balance sheet and the size of the debt on the liability side. It would be interesting to investigate this in other interbank markets in the future.

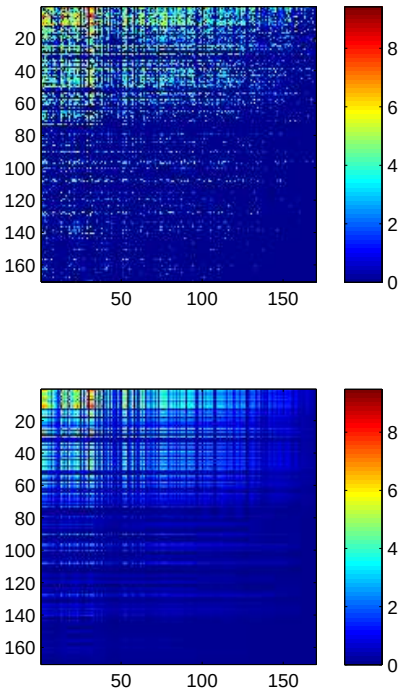

Fig. 5.22: Example: Data matrix and approximation based on the AC model for 2000 Q3, with warm colors indicating high transaction values between individual banks. The upper panel shows the log-transformed data matrix **D** after applying the MINRES/SVD approach and sorting the network according to the coreness vectors. The lower panel shows the MINRES/SVD approximation based on the AC model, with the same ordering of the network nodes.

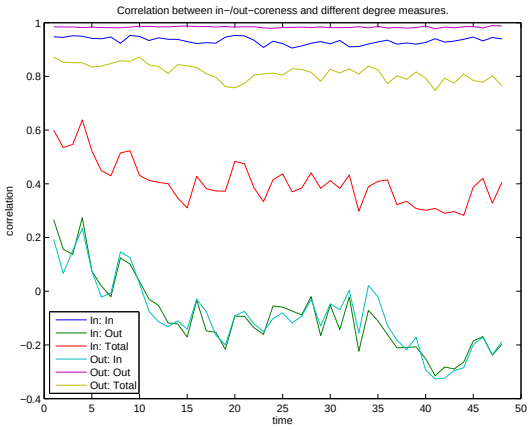

Fig. 5.23: Time-varying correlation between in-/out-coreness and in-degree/out-degree/total degree. Total degree is the degree we would obtain from transforming the directed network into an undirected network.

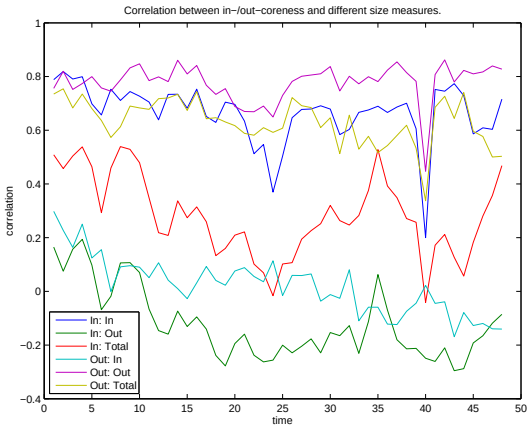

Fig. 5.24: Time-varying correlation between in-/out-coreness and in-size/out-size/total size as defined in the text.

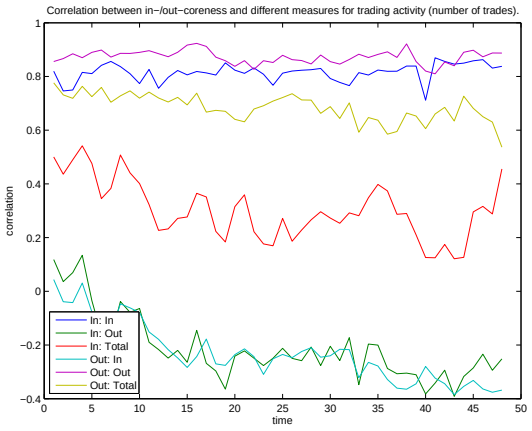

Fig. 5.25: Time-varying correlation between in-/out-coreness and number of in- and out transactions. Total number of transactions is the sum of the two.

5.8.6 *Changing the Aggregation Period*

Here we briefly discuss the results for other than quarterly aggregation periods, by focusing on the results from monthly and annual networks. To be precise, the typical element in $\mathbf{A}$ is $a_{ij} = 1$ if there was at least one transaction from bank i to j during the respective month/year.[86] In both cases, we expect comparable results as for the quarterly networks, however, with certain differences. For example, given the fact that the activity structure of banks is less stable for monthly compared to quarterly networks, we expect the coreness vectors to be less stable over time. Furthermore, due to the aggregation, the total number of active banks in one quarter will be at least as big as the number of active banks in any of the 3 months in this particular quarter. Therefore, we expect the relative size of the core to be somewhat smaller for monthly data. For annual data, it is a priori not clear what might happen to the coreness vectors. While we expect a larger core as compared to the quarterly case (due to the reduction of noise and the relatively large number of active banks per year), the volatility in the coreness vectors (over time) could in fact be higher, given that annual aggregation makes it more likely that two banks being active in different quarters are being put together in a network where they in fact never could have interacted. Again, the anonymity of the data set makes it impossible for us to disentangle these effects.

[86] We do not expect significant changes, if we used a higher threshold here, e.g. banks might have to trade at least twice within a particular period to establish a link.

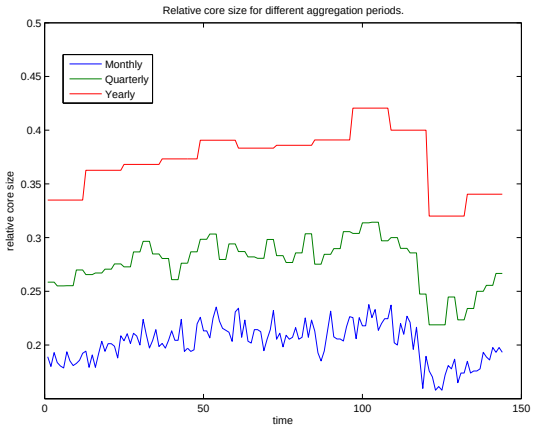

Fig. 5.26: Comparison of the time-varying relative core sizes for different aggregation periods.

Figure 5.26 shows the time-varying relative core sizes (discrete model) of different aggregation periods for the entire sample period of 144 months. As expected, the relative size of the core depends positively on the length of the aggregation period, so the core is largest for yearly networks, consisting of roughly 36% of sample banks before the GFC and close to 30% afterwards. The structural break due to the GFC is clearly visible for yearly and quarterly data. This is not so much true for monthly data, where pre- and post-GFC average core sizes are 20% and 17% respectively. The average core size appears to be relatively stable for monthly data, however with wilder fluctuations as compared to longer frequencies. Thus, there appears to be substantial noise in the monthly networks, backing up our use of quarterly data in the baseline scenario. Similar remarks apply to the error scores in Figure 5.27. As expected, the high level of noise at higher frequencies deteriorates the model fit: the fit is worst for monthly data, with the highest error scores. Interestingly, at the monthly level, the GFC seems less like a big shock as compared to quarterly and yearly data, since the error score increases already before the GFC. In contrast, the fit of the longer aggregation periods drops with the GFC, which is most clearly visible for the yearly data. In any case, it seems that the Italian interbank market shows a core-periphery structure at all frequencies under study, even though for monthly networks the noise level is rather high so that it is much harder to identify the core at such high frequencies.

Concerning the correlation between individual coreness vectors from different models, we just note here that the correlation between the discrete

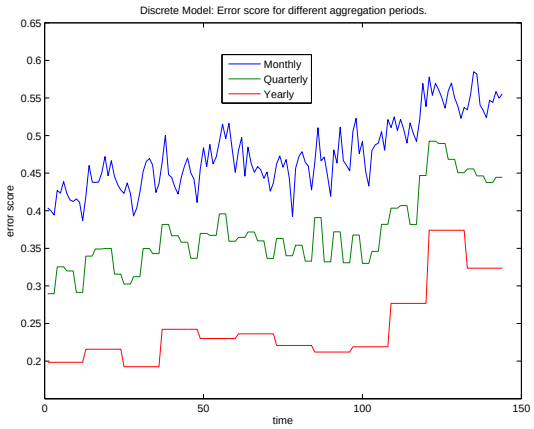

Fig. 5.27: Comparison of the error-scores for different aggregation periods.

and tiering model is roughly .8750 for monthly and .9918 for yearly networks. Thus, the results from the discrete and tiering model are very similar in general, but for monthly data there may be some differences. While the correlation between out-coreness and the discrete model is comparable for all frequencies (around .73), the correlation with in-coreness depends positively on the aggregation period.[87] Thus, it seems that network becomes more symmetric for longer aggregation periods. This is confirmed by the fact that the correlation between in- and out-coreness is -.2126 for monthly data and .1383 for yearly data. Thus, at the yearly level, in- and out-coreness appear to be positively related. In any case, the correlation between the two vectors is rather small in absolute terms, so the asymmetric MINRES/SVD approach captures the inherent asymmetry of the network at all frequencies.

5.8.7 Including Foreign Banks

Analyzing the network formed by foreign banks only, we see average error scores around 45% before the GFC and values close to 90% afterwards. Given that foreign banks' activity is rather unstable, due to the simple fact that they have access to other sources of funding, it comes as no surprise that foreign banks form a structurally different separate subnetwork.

We also analyzed the complete network, with Italian and foreign banks. Interestingly, most of our findings from the baseline scenario with Italian banks only, remain unaffected. However, there is a clear upward trend in the

[87] For monthly data the correlation between the discrete and in-coreness vector is roughly .1578, while the value is .3861 for yearly data.

error score over time which is driven by the rather unstable nature of foreign banks' activity. Thus it seems justified to exclude foreign banks from the analysis.

5.8.8 Continuous Model Using the Number of Transactions

As another robustness check, we ran the continuous model using the quarterly matrices containing the number of (directed) transactions between individual banks $\mathbf{T}$, instead of the total transaction volumes $\mathbf{D}$.[88] We find very similar coreness vectors in both settings. For the SC model the correlation between the vectors is .9618. For the AC model, the correlations are .9594 and .9844 for in- and out-coreness, respectively.[89] Quite interestingly, it seems that the fit (in terms of PRE) of the MINRES model in this case is even worse compared to the values presented in the main text, while the fit of the MINRES/SVD model is slightly better than before. It not not quite clear, why this is the case, but we should stress that all of the results here indicate that the model is quite robust. Most importantly, it seems that the conclusions from above also hold for alternative valued matrices, that measure the intensity of bilateral relations in a meaningful way.

5.8.9 Further Robustness Checks

We performed additional robustness checks for the discrete and continuous models.

- We estimated the discrete models using the correlation-based approach of Stolzenburg and Lux (2011); there we used the correlation between the observed data matrix and the pattern matrix constructed from the coreness vectors as the objective function. Note that the correlation-based approach can only be used for the discrete model in the case of arbitrary off-diagonal blocks. We found very similar results compared to the baseline scenario. In this regard, we also ran the discrete model with valued networks (transaction volumes and number of transactions, both log-transformed) and found very similar results compared to the baseline scenario.

- We used the binary networks, rather than the valued ones, as input matrices in the continuous models. The coreness vectors are very similar

[88] Again we log-transform the data matrix to reduce the level of skewness, cf. section 5.4.2.

[89] We also find similar values for monthly and annual data.

to the baseline case (correlation of above .93), however, with constantly smaller PREs.

- We estimated the continuous models without log-transforming the data. Due to the high level of skewness in the data (driven by many zeros in the matrices), it is not possible to identify a sensible core for the transaction volumes. The coreness vectors are hardly comparable, with a correlation of .4928 between the out- and .3829 between the in-coreness vectors compared to the baseline scenario, respectively. Furthermore the PREs are highly volatile over time. In contrast, using the number of transactions the correlations are .8510 and .7668 for the out- and in-coreness vectors compared to the baseline scenario, respectively. Also the PREs are comparable to the baseline scenario, however, with extreme values around quarter 10 and 39, i.e. the two candidates for structural breaks.

- We also estimated the continuous models using the correlation-based approach. We find identical results in both cases, however, the approach presented in the paper is preferable, since the computation is much faster. We should note that the correlation between the pattern and the observed matrices is always above .70 in the AC model, while it may be as low as .37 for the SC model after the GFC.

6. TRADING STRATEGIES IN THE OVERNIGHT MONEY MARKET: CORRELATIONS AND CLUSTERING ON THE E-MID TRADING PLATFORM

6.1 Introduction and Existing Literature

The global financial crisis (GFC) has shown that the stability of the interbank network is of utmost importance for the macroeconomy. Linkages arising from bilateral lending relationships may lead to systemic risk on the macro-level, see Haldane (2009). Therefore, regulators and policymakers need to understand the structure and functioning of the interbank network in more detail.

Approaches from the natural sciences have been used to investigate the topology of networks formed by interbank liabilities. Examples for the Italian e-MID (electronic market for interbank deposit) include De Masi *et al.* (2006), Iori *et al.* (2008), and Fricke and Lux (2012). Several 'stylized facts' of interbank networks have been identified, among them the finding of disassortative mixing patterns, i.e. high-degree nodes tend to trade with low-degree nodes, and vice versa.[1]. Quite recently, Fricke and Lux (2012) argued that this fact may be a key driver for the observation of a hierarchical core-periphery (CP) structure in interbank networks, see also Craig and von Peter (2010). Thus, preferential lending relationships between individual institutions may lead to community structure at the macro-level. While many authors also note some other form of community structure in the interbank network they analyze (see e.g. Boss *et al.* (2004)), Fricke and Lux (2012) stress that apart from the CP structure, there is no community structure whatsoever in the Italian e-MID network of interbank claims.

Based on these findings, we aim at analyzing the trading strategies of individual institutions in more detail. For this purpose we use a detailed dataset containing all overnight interbank transactions on the e-MID trading platform from January 1999 to December 2010. Splitting the sample into half-yearly subsamples, we define the (intra-) daily net trading volumes of the individual institutions as the trading strategies and analyze the correlations between them.[2] We find evidence for significant and persistent bilateral correlations between institutions' trading strategies. In most semesters we find two anti-correlated clusters of trading strategies, indicating that banks tend to trade preferentially with banks from the other cluster. This finding is both in line with the existence of preferential relationships, and with herding phenomena in banks' trading strategies. Interestingly, the information whether a bank is a continuous net buyer or net seller appears to be the only defining characteristic for its group membership in this anonymous dataset.

[1] See Cocco *et al.* (2009).

[2] See Kyriakopoulos *et al.* (2009) for a similar approach. Comparable studies also exist for the patterns of trading members in stock exchanges, see e.g. Zovko and Farmer (2007) and Tumminello *et al.* (2012).

This is in contrast to the findings of Iori, Renò, Masi, and Caldarelli (2007), who performed a very similar analysis based on the e-MID data during the sample period 1999-2002, and found two clusters containing mostly large and small banks, respectively. We also highlight some problems related to the definition of trading strategies. Overall, our findings are quite persistent over time and appear to be largely unaffected by the GFC. We add further evidence on the fact that preferential lending relationships on the micro-level lead to community structure on the macro-level.

The remainder of the paper is structured as follows: section 6.2 briefly introduces the Italian e-MID interbank data, section 6.3 contains the empirical analysis and section 6.4 concludes.

6.2 Dataset

The Italian electronic market for interbank deposits (e-MID) is a screen-based platform for trading unsecured money-market deposits in Euros, US-Dollars, Pound Sterling, and Zloty operating in Milan through e-MID SpA.[3] The market is fully centralized and very liquid; in 2006 e-MID accounted for 17% of total turnover in the unsecured money market in the Euro area and covers essentially the entire domestic overnight deposit market in Italy. Average daily trading volumes were 24.2 bn Euro in 2006, 22.4 bn Euro in 2007 and only 14 bn Euro in 2008.

Available maturities range from overnight up to one year. Most of the transactions ($> 80\%$) are overnight, which is not surprising given that the loans are unsecured. In August 2011, e-MID had 192 members from EU countries and the US, of which 101 were domestic (Italian) and 61 international banks.[4] This composition is, however, not stable over time as can be seen from Figure 6.1: The left panel shows a clear downward trend in the number of active Italian banks over time, whereas the additional large drop after the onset of the GFC (semester 18) is mainly due to the exit of foreign banks. The right panel shows that the decline of the number of active Italian banks went along with a relatively constant trading volume in this segment until 2008.[5] The overall upward trend of trading volumes was mainly due to the increasing activity of foreign banks until 2008, which virtually faded away after the onset of the crisis. Due to the cyclicality of foreign banks' activity, we will focus on the Italian banks only. Additionally, for reasons

[3] More details can be found on the e-MID website, see http://www.e-mid.it/.

[4] Additionally, 29 central banks and 1 ministry of finance acted as market observers.

[5] Thus, the decline of active Italian banks appears to be driven by domestic mergers and acquisitions.

explained below, we will only consider the relatively active banks within the respective semesters.

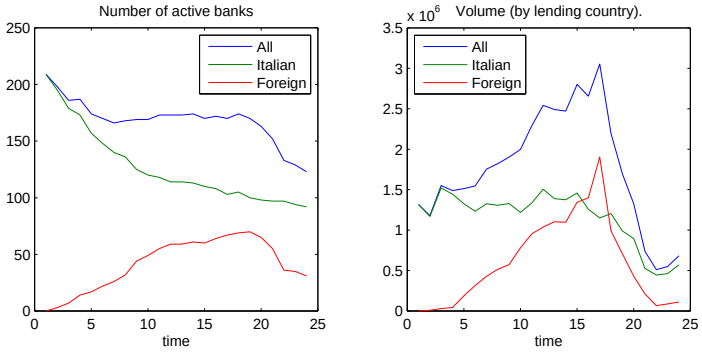

Fig. 6.1: Number of active banks (left) and traded volume (right) over time. We also split the traded volume into money going out from Italian and foreign banks, respectively.

The trading mechanism follows a quote-driven market and is similar to a limit-order-book in a stock market, but without consolidation. The market is transparent in the sense that the quoting banks' IDs are visible to all other banks. Quotes contain the market side (buy or sell money), the volume, the interest rate and the maturity. Trades are registered when a bank (aggressor) actively chooses a quoted order. The minimum quote size is 1.5 million Euros, whereas the minimum trade size is only 50,000 Euros. Thus, aggressors do not have to trade the entire amount quoted. The platform allows for credit line checking before a transaction will be carried out, so trades have to be confirmed by both counterparties. The market also allows direct bilateral trades between counterparties. Contracts are automatically settled through the TARGET2 system.

We have access to all registered trades in Euros in the period from January 1999 to December 2010.[6] For each trade we know the two banks' ID numbers (anonymous to us), their relative position (aggressor and quoter), the maturity and the transaction type (buy or sell). We will focus on all overnight trades conducted on the platform, leaving a total number of 1,317,679 trades. The large sample size allows us to analyze the system's evolution over time. Here we focus on 24 independent semi-annual subsamples.

In the following, we analyze the correlation matrices of the banks' trading strategies. The finding of a core-periphery structure in the interbank network

[6] A detailed description of the dataset can be found in Finger *et al.* (2012) and Fricke and Lux (2012).

indicates the existence of a hierarchical structure. Interestingly, we did not find any further community structure in the network using usual community detecting algorithms. In this paper, we identify clusters of trading strategies in a different type of network based on the (intra-) daily trading patterns of individual banks.

6.3 Empirical Analysis of Trading Strategies

6.3.1 Measuring Correlations Between Strategies

Even though the dataset is anonymous, each bank can be identified by a unique ID-number, so we can observe the behavior of all sample banks over the entire sample period. In the following, we will divide the dataset into 24 half-yearly intervals (semesters) which are treated as independent subsamples.[7] In order to define the trading strategies we further divide each semester into different trading sessions. For example, Iori *et al.* (2007) use a Fourier method to use continuous trade data, while Kyriakopoulos *et al.* (2009) use daily data. Here we find that each day in the sample period can be naturally split into two trading sessions, during which a similar fraction of trades occur on average.[8] Figure 6.2 shows the fraction of trades by the time of the day.[9] For example, we see that roughly 25% of the trades occur between 9 and 10am. This is in line with the findings in Iori *et al.* (2008), and we checked that this observation is in fact stable over the entire sample period.[10] Based on these numbers, we split each day into a morning session (8am-12am) and a afternoon session (12pm-6pm). In total, our sample contains 3,073 days, so we end up with 6,146 trading periods. Due to the heterogeneity of the banks' trading activity, we restrict ourselves to the institutions which were active in a fraction of at least $\theta = 50\%$ of all trading sessions in a particular semester.[11] In this way, we focus on those institutions which were relatively active over a significant part of the sample period. Given the negative trend in the number of active Italian banks, it is unavoidable that the number of banks ending up in the sample is also decreasing over time, see Figure 6.3.

[7] Later on we will see that the subsamples are in fact not independent.

[8] Note that we also checked that the results are not affected by splitting each day into 3 or 4 trading sessions. However, it should be clear that the number of periods without any activity is positively related to the number of subintervals.

[9] The same remarks hold for the transacted volumes, yielding very similar results.

[10] See Iori *et al.* (2007) for an identification of the events that contribute to the intra-day patterns.

[11] We checked the results are qualitatively very similar for other values of θ.

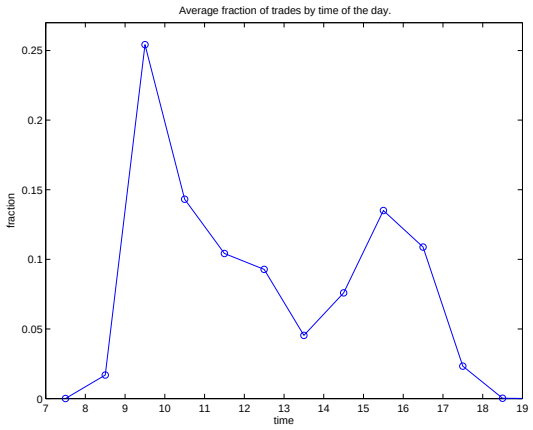

Fig. 6.2: Activity by time of the day, indicated by the fraction of trades occurring at a certain time of the day (hourly intervals).

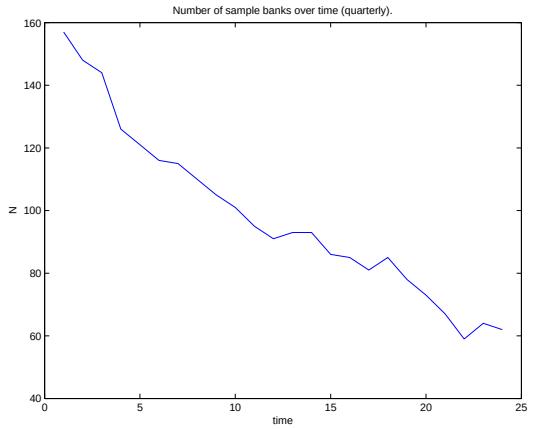

Fig. 6.3: Number of sample banks over time. For each semester, a bank is in the sample if it was active in at least θT trading sessions, with T as the total number of trading sessions in the respective semester.

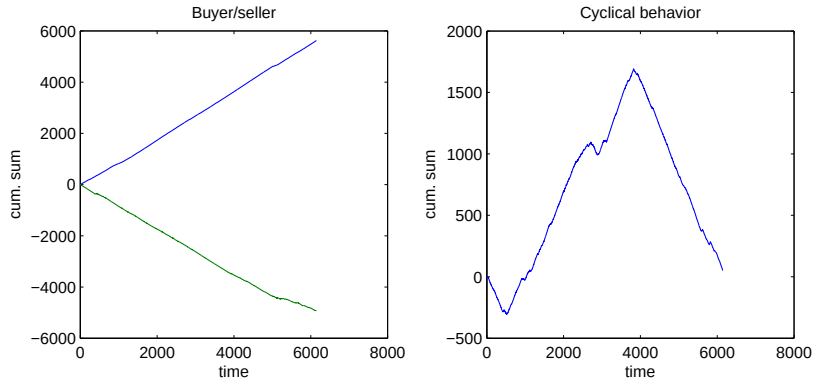

Fig. 6.4: Examples of banks' trading strategies in the Italian interbank market. Here we plot cumulative sums of particular signed strategies (1,-1,0) over the entire sample period. The examples were not chosen randomly, but to illustrate different trading styles. The left panel shows a constant buyer/seller, and the right panel shows a highly cyclical trading strategy.

For each trading period and for each institution, we define the trading strategy as the net traded volume in Euros. Net volume is the total buy volume (money borrowed from other banks) minus total sell volume (money lent to other banks). Thus, a bank with a positive (negative) net volume is a net buyer (seller) of money. Banks with a zero strategy were either inactive in the trading session or ended up with a flat position after buying and selling equal amounts.[12] Figure 6.4 shows examples of cumulative (signed) trading strategies for three institutions which were active during most of the sample period. We see a substantial level of heterogeneity in the employed strategies, with banks acting as continuous buyers/sellers (left panel) and more complicated, possibly cyclical, patterns (right panel).

In the end, we obtain a strategy matrix $\mathbf{S}$ of dimension $N \times T$, where T denotes the number of trading sessions in a particular semester and N the number of sample banks in this semester. Note that in each semester the strategy matrix only includes transactions between the sample banks, thus ig-

[12] We also defined signed strategies (similar to the approach in Zovko and Farmer (2007)), where we assign a +1, -1 or a 0 describing the strategies in a simpler form. However, given that our network is closed, we work with the valued version since it contains more information. A technical reason is that it may happen that a bank has the same strategy in each trading session. In such a case, we cannot compute the correlation of this strategy with other strategies.

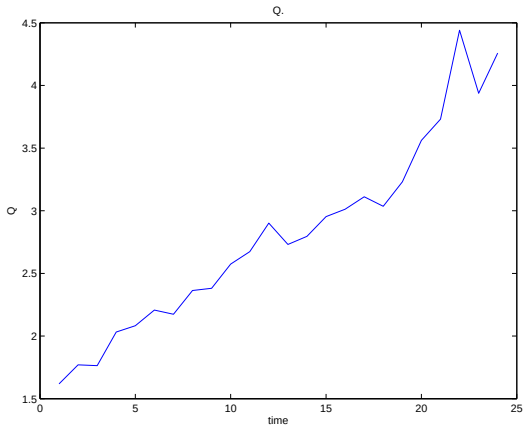

Fig. 6.5: Time evolution of $Q = T/N$.

noring transactions with other banks that might act relatively infrequently.[13] In everything that follows, we will work with the normalized bank-specific time-series (mean equal to 0 and variance equal to 1). For our sample period, $T \sim 240$, while N varies over time as shown above. In total, $N_T = 218$ different institutions are in the final sample, but obviously not each bank is active in each semester. Therefore, the ratio $Q = T/N$ is not constant over time as can be seen from Figure 6.5. We see that this ratio exceeds 1 for the complete sample period, which is an indispensable condition to have meaningful correlation estimates.[14] However, the values are still quite low so we can expect a non-negligible amount of measurement noise in the correlation matrices.

From the strategy matrix $\mathbf{S}_t$ we can construct the $N \times N$ semi-annual correlation matrices $\mathbf{C}_t$ between the institutions' strategies. Due to the normalization of the strategies the correlation matrix can be simply calculated based on the Pearson estimator

$$\mathbf{C}_t = \frac{1}{T-1}\mathbf{S}_t\mathbf{S}_t^T, \tag{6.1}$$

with $()^T$ denoting the transpose of a matrix, see Laloux *et al.* (1999).

A color example of such a correlation matrix for the first half of 2006 can be seen in Figure 6.6. Dark colors represent high absolute correlations, with

[13] In this sense, we work with a closed network of banks.

[14] For a set of N time series of length T, the correlation matrix contains $\frac{N(N-1)}{2}$ entries which have to be determined by NT observations. If T is not very large compared to N, the empirical correlation matrix is to a large extent noisy. We should also note that many results from random matrix theory require $Q \geq 1$, see below.

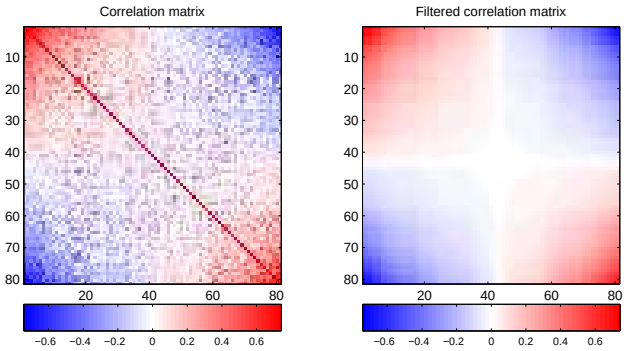

Fig. 6.6: Example of a correlation matrix after applying the clustering approach for the first half of 2006. Left: raw correlation matrix. Right: filtered matrix. Warm colors represent positive correlations, cold colors negative ones. White indicates uncorrelated strategies.

warm and cold colors indicating positive and negative correlations, respectively. As will be explained in more detail below, the left panel shows the raw correlation matrix after applying the clustering approach and the right panel shows the filtered correlation matrix with identical ordering. In both cases, the two (anti-correlated) clusters are readily visible, with the level of noisy correlations significantly reduced in the filtered matrix. We will show below that similar observations hold for most semesters.

As a first rough formal test for the significance of the computed correlations, we use a t-test assuming normally distributed disturbances. Later on, we will provide more thorough significance tests, but even this simple approach is already quite suggestive. Figure 6.7 shows the fraction of significant correlations at the 5% level, among all correlations, over time. Obviously, these values are quite large, with an average value of close to 50% which clearly exceeds the 5% we would expect randomly with a 5% acceptance level of the test. In the following, we will provide a more detailed analysis of the correlation matrices, mostly in terms of their eigenvalues.

6.3.2 Structure in the Correlation Matrices

In this section we investigate the eigenvalue spectrum of the observed correlation matrices. Note that since the correlation matrices are symmetric, all eigenvalues are positive real numbers, see Bouchaud and Potters (2003). We denote v_i as the (normalized) eigenvector corresponding to the ith eigen-

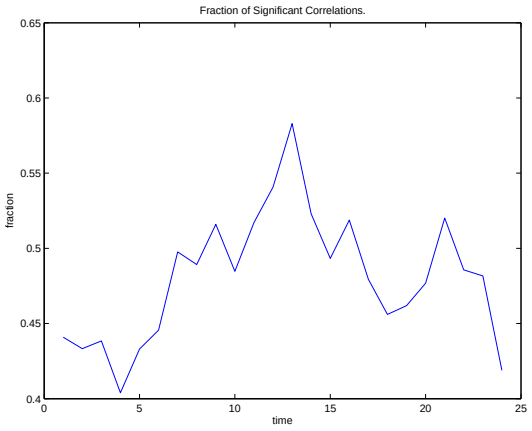

Fig. 6.7: Fraction of significant correlations over time, based on a standard t-test testing the significance of the correlation coefficient at the 5% level.

value λ_i,[15] where the eigenvalues are sorted from largest to smallest, such that $\lambda_1 \geq \lambda_2 \cdots \geq \lambda_N$. By definition, an eigenvector is a linear combination of the different components $j = 1, \cdots, N$ such that $\mathbf{C}v_i = \lambda_i v_i$. Thus, λ_i is the variance of the correlation matrix captured by the weighting based on the corresponding eigenvector.

Time-Persistence of Correlations

As a first step, we investigate whether the observed correlations are stable over time. This is helpful for the dynamic analysis, since we expect trading strategies, and thus the corresponding clusters, to be quite persistent. In this way, the time-persistence of individual correlation pairs offers a first test for spuriousness, since spurious correlations are not likely to persist in time. In contrast, if the correlations are persistent, then the clusters of institutions are also likely to persist in time.

Since we can track individual institutions over time, we can test whether the (time-varying) correlations between institutions' trading strategies are significantly different from zero. In total, we have $N_T = 218$ institutions in our sample. Therefore, we have to evaluate the significance of $(N_T^2 - N_T)/2 = 23,653$ relationships. Given that two sample banks may be active in different semesters, several of these correlations cannot be evaluated. Nevertheless, for

[15] We drop the time-indices in most of the following to save notation. It should be clear that we have a correlation matrix for each semester, and thus time varying eigenvalues and eigenvectors.

all institutions jointly belonging to the sample in at least two semesters, we can check whether the observed correlations are significantly different from zero using a simple t-test assuming Gaussian residuals.[16]

Doing this, we find 3,203 significant combinations at the 5% significance level, i.e. a fraction of roughly 10% of all correlation pairs is significantly different from zero. This value is way above the 5% level we would expect to occur in a purely random fashion.[17] Hence, a large number of the correlations between individual banks' trading strategies are significantly different from zero over the sample period. We can therefore expect any identified clusters to persist over time.

Time Evolution of the Largest Eigenvalues

For a set of infinite length uncorrelated time series, $\mathbf{C}$ would correspond to the identity matrix with all eigenvalues equal to 1. If the time series are of finite length, however, the eigenvalues will never exactly equal 1, see Zovko and Farmer (2007). In fact, the distribution of the eigenvalues in the case of iid Gaussian time-series is known. Below, we will use this information in order to detect significant clustering by finding eigenvalues significantly larger than the predicted values. As a first step, however, we focus on the absolute and relative size of the eigenvalues over time. Since $\sum_i \lambda_i = N$, we can standardize the observed eigenvalues to make the values comparable over time. Figure 6.8 shows the dynamics of the five largest eigenvalues, both in absolute terms (left panel) and standardized (right panel).

Note that the dynamics of the largest eigenvalues (absolute value) are, in qualitative terms, quite similar to the number of sample banks in Figure 6.3. Except for the beginning and the end of the sample period, the relative importance of the leading eigenvalues is rather strong with values around 15-20%.[18] Thus, we might expect at least the largest eigenvalue to carry important information on the structure of the correlation matrix.[19] Given

[16] Note that this is a joint test also on the autocorrelation of observed correlations, since highly volatile values are unlikely to be significant.

[17] Here we disregard those correlation pairs between banks that were never jointly active. Taking this into account, increases the fraction of significant correlations to roughly 20%.

[18] However, we should stress that these values are surprisingly small compared to the findings of Iori *et al.* (2007), where values around 45% were found.

[19] Note that the first eigenvalue monitors information about the primary cluster in the correlation matrix, both in terms of size and its average correlation. The subsequent eigenvalues carry information on subsequent clusters, see Friedman and Weisberg (1981). Given that the sum of the subsequent eigenvalues is of comparable size as the leading eigenvalue, we also checked whether we can extract additional clusters from these. However, we found no evidence in this regard, so in most of the following we will focus on the

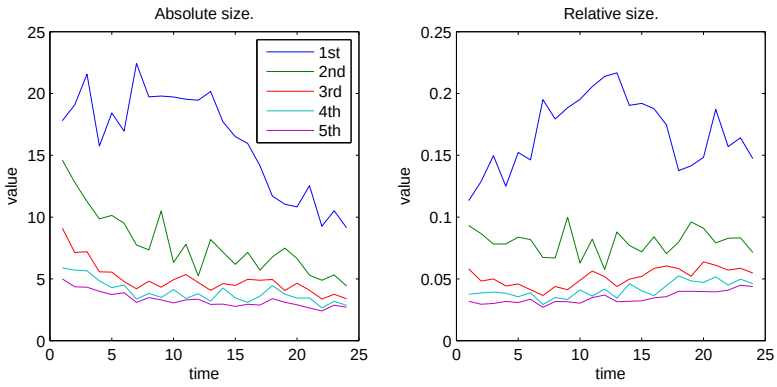

Fig. 6.8: Absolute (left) and relative (right) size of the first 5 Eigenvalues of $\mathbf{C}$ over time. The relative size normalizes the observed eigenvalues by their total sum N.

that the following eigenvalues are quite stable (both in absolute and relative values), most of the variation of the absolute values of the leading eigenvalues is due to changes in N. However, we also see that the relative size of the leading eigenvalue is not stable over time. A large part of this variation is likely to be driven by the instability of the trading strategies (and the trading clusters) at the beginning and the end of the sample period.

While the eigenvalues give us a macroscopic description of our network in terms of the variance, investigating the leading eigenvectors allows us to understand how much a particular principal component affects the original variables. Below, we will use information from the leading eigenvector to define the trading clusters.[20]

Random Matrix Theory and Eigenvalue Density

Here we test the significance of the observed eigenvalues based on results from random matrix theory (RMT). It is well documented that the eigenvalues of a correlation matrix can be used to separate true information from noise. The null hypothesis of uncorrelated strategies translates itself to the assumption of iid random elements in the strategy matrices, which gives the

leading eigenvalues only.

[20] We also performed several bootstrap-tests on the significance of the largest eigenvalues. For example, we randomly reshuffled buy and sell periods for the institutions, constructed the correlation matrix and compared the resulting eigenvalues with those from the observed correlation matrices. The results are comparable with those based on RMT as detailed below, and are available upon request.

so-called random Wishart matrices or Laguerre ensemble of RMT, see Laloux *et al.* (1999).

The basic idea is to compare the observed eigenvalue frequency distribution with a hypothetical one obtained from the correlation matrix for N iid Gaussian random variables of length T. More formally, for N random uncorrelated variables of length T, in the (thermodynamic) limit $T, N \longrightarrow \infty$, with $Q = T/N \geq 1$, the density of the eigenvalues $p(\lambda)$ is given by the functional form

$$p(\lambda) = \frac{Q}{2\pi\sigma^2} \frac{\sqrt{(\lambda^{\max} - \lambda)(\lambda - \lambda_{\min})}}{\lambda}, \tag{6.2}$$

with

$$\lambda_{\min}^{\max} = \sigma^2(1 + 1/Q \pm 2\sqrt{1/Q}), \tag{6.3}$$

where σ^2 being the variance of the time series, which is equal to 1 due to the standardization, and $\lambda \in [\lambda_{\min}, \lambda^{\max}]$. Equations (6.2) and (6.3) define the well-known Marčenko-Pastur (MP) distribution of eigenvalues from iid Gaussian random variables.[21]. Empirical values falling outside the interval defined by Eq. (6.3) are likely to carry important information about the system. Note that the closed-form solution is only a limiting result and, as shown above, the parameter Q changes over time. Consequently, the predicted eigenvalue density under the null changes from semester to semester. Therefore, we run separate tests for each semester, where we compare the observed eigenvalue densities to the range of the eigenvalues under the null, which can be calculated by using the observed Q in Eq. (6.3).

Figure 6.9 shows an illustrative example of the results of our analysis. There we compare the actual eigenvalues (black bars) with the range predicted by RMT (red area). We see that the bulk of the eigenvalues lie within the range of the MP distribution. However, in this particular semester, 5 eigenvalues are outside the upper limit of the MP distribution, with one very large eigenvalue, one of medium size and the remaining eigenvalues of comparable size. The results are comparable for most periods, usually containing 5 or 6 eigenvalues significantly larger than those predicted under the null.[22] This is further evidence for the non-random structure of the correlations between trading strategies, even though we have to bear in mind that both N and T are quite small, such that there may be small-sample biases.[23] Nevertheless, the largest eigenvalue is usually substantially larger than the

[21] See Marčenko and Pastur (1967).

[22] Note that there are often also values smaller than the predicted $\lambda_{\min}$, which we will not comment on in the following.

[23] We found comparable results for other aggregation periods, e.g. yearly data, and other selected samples, e.g. other values for θ. Thus, we are quite confident that the results do not depend crucially on these issues.

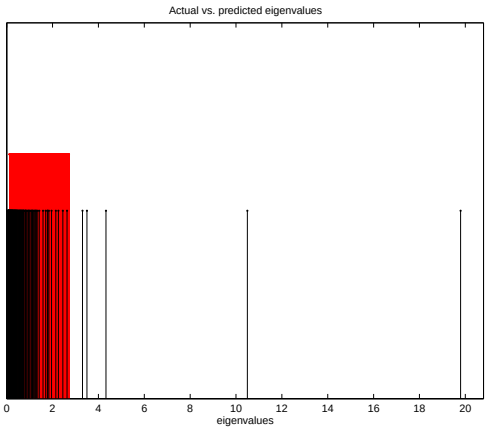

Fig. 6.9: Comparison of observed and predicted eigenvalues from the MP distribution, for semester 9. Black bars indicate observed eigenvalues, the red block shows the range of the predicted distribution, as defined in Eq. 6.3.

theoretical prediction, cf. Figure 6.10. Thus, in most of what follows, we will focus on the largest eigenvalue and the corresponding leading eigenvector.

In the next section we filter out the noisy component of the observed correlation matrices. This will allow us to identify the clusters of trading strategies more convincingly.

Filtering the Correlation Matrices

The presence of a well-defined bulk of eigenvalues in agreement with the MP distribution suggests that the contents of $\mathbf{C}$ are highly affected by noise. In other words, those eigenvalues outside the predicted range are likely to contain substantial information about the underlying community structure, i.e. those components we are actually interested in.[24] In the following, we will therefore filter out the noisy component from the correlation matrices, using the approach of Kim and Jeong (2005). Similar approaches for correlation matrices between stock returns can be found in Laloux *et al.* (1999), Plerou *et al.* (2002), and Livan *et al.* (2011).

With the complete set of eigenvalues and eigenvectors, we can decompose

[24] In the future, it would be very interesting to use the approach of Livan *et al.* (2011) and set up a factor model that might be able to reproduce the observed eigenvalue spectra.

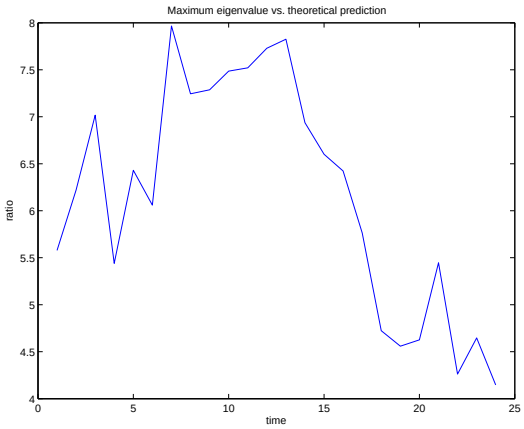

Fig. 6.10: Relative size of the largest eigenvalue compared to $\lambda_{\max}$, as defined in Eq. 6.3.

the correlation matrix $\mathbf{C}$ as

$$\mathbf{C} = \sum_{i=1}^{N} \lambda_i v_i v_i', \tag{6.4}$$

where v_i and λ_i defined as above. Because only the eigenvectors corresponding to the few largest eigenvalues contain the information on significantly non-random structure, we can identify a filtered correlation matrix for these groups by choosing a partial sum of Eq. (6.4), denoted as $\mathbf{C}^g$. Thus, we posit that the eigenvalue spectrum of the correlation matrix is organized into a group part defined by the largest eigenvalues, and a random part containing the bulk of small eigenvalues consistent with the MP distribution. In this regard, we should stress that this is slightly different for the correlations of stock returns, where the first eigenvalue is usually an order of magnitude larger than the rest. The corresponding eigenvector is simply the 'market' itself, containing roughly equal components on all stocks with most components having the same sign.[25] In our case, we did not find evidence for such a 'market' vector, since the distribution of eigenvector components contains a roughly equal number of positive and negative eigenvalues. In this regard, we calculated the inverse participation ratio (IPR) defined as

$$I^k = \sum_{j=1}^{N} [v_j^k]^4, \tag{6.5}$$

[25] See for example Plerou *et al.* (2002).

with v_j^k being the jth component of the kth eigenvector. Here, the IPR serves as a measure for the distribution of eigenvector components. For the case of (1) identical components in eigenvector k, I^k will take the value $1/N$, while for the case of (2) one component equal to 1 and the remainder equal to zero, I^k takes the value of 1.[26] Additionally, it is often useful to compare the results to the IPRs of (Gaussian) random matrices. Figure 6.11 shows the results: The blue (green) line shows the (log) of the IPR for the eigenvector corresponding to the largest (smallest) eigenvalue, while the red line shows the value for the eigenvectors from a Monte-Carlo simulation using random correlation matrices with the same dimensions.[27] For the sake of completeness, we also plot the benchmark of identical components in the eigenvectors (turquoise line). Quite surprisingly, we see that the leading eigenvectors are actually very close to the random case, whereas this is not true for the smallest eigenvectors which consist of significantly less equal components.[28] If anything, it seems that the leading eigenvectors are somewhat closer to the benchmark case of identical components as compared to the random case.[29]

We can write the correlation matrix as

$$\mathbf{C} = \mathbf{C}^g + \mathbf{C}^r = \sum_{i=1}^{N_g} \lambda_i v_i v_i' + \sum_{i=N_g+1}^{N} \lambda_i v_i v_i'. \tag{6.6}$$

Note that the calculation of $\mathbf{C}^g$ involves a choice on the number of leading eigenvalues relevant for its calculation. Here we simply use the leading eigenvector only ($N_g = 1$), but the results do not depend on this assumption.[30] In the following, we will always show the results for the raw and the filtered correlation matrix.

[26] See e.g. Plerou *et al.* (2002).

[27] The results are average values from a Monte-Carlo experiment based on 100 simulations.

[28] This result shows that eigenvalues smaller than the theoretical lower limit might contain information not captured by the largest eigenvalues only. The investigation of these eigenvalues would be an interesting avenue of future research.

[29] In the end, it appears that the eigenvectors, even those corresponding to those eigenvalues exceeding the theoretical threshold, show IPRs which are similar to those of random correlation matrices. This result, surprising as it may seem, should be interpreted with care, since we will see below that these seemingly random eigenvectors display a substantial level of autocorrelation in the signs of their components.

[30] Note that we use the filtering technique to illustrate the clusters graphically (as in Figure 6.6. Since we identify clusters using only the leading eigenvector, the results will not be affected by taking other eigenvalues into account as well.

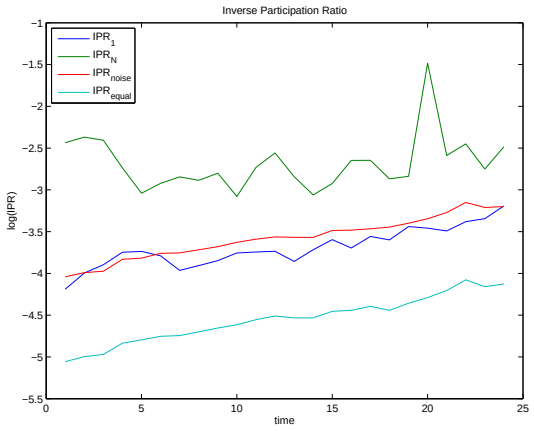

Fig. 6.11: Inverse participation ratio (logarithm) for the first and last eigenvector over time. The plot also shows the results for random matrices and the case of equal components.

6.3.3 Clustering of Trading Behavior

We have both shown that the correlations between banks' trading strategies are quite persistent and that a number of eigenvalues of the correlation matrices significantly exceed those predicted in case of completely uncorrelated strategies. Thus, there is some evidence in favor of clustering, i.e. non-random structure, in the correlation matrices. However, we have also seen that there is a substantial level of noise in the correlation matrices, which is why we constructed filtered correlation matrices, based on the significant eigenvalues and their corresponding eigenvectors only.

In this section we will show that the components of the leading eigenvector allow us to identify clusters of trading behavior.[31] We usually find two clusters of banks, with highly correlated strategies within the clusters, but anti-correlated strategies between the clusters. This finding suggests that the banks in one cluster tend to trade in similar directions and their counterparties are usually from the other cluster. At first sight, this result may seem trivial, because in the money market there is always one buyer and one seller, so some banks will necessarily have correlated strategies. But what is surprising is that the clusters are rather large (compared to the number of sample banks) and persist over time. Thus, banks appear to have preferred counterparties in the market, which they turn to repeatedly when they are

[31] We also use hierarchical clustering techniques, which yielded very similar results. One important disadvantage of these techniques, however, is that they remain silent about the optimal number of clusters.

willing to trade. We will discuss the meaning of the clusters in more detail below.

Clustering Based on the Leading Eigenvectors

Our clustering approach consists of dividing the leading eigenvector into positive and negative components. Components with equal sign end up in the same cluster,[32] i.e. the indicator for the group membership of bank j (g_j) can be defined as

$$g_j = \begin{cases} 1 & \text{if } v_1^j > 0, \\ -1 & \text{if } v_1^j < 0, \\ 0 & \text{if } j \text{ was inactive.} \end{cases} \tag{6.7}$$

We do this for each semester individually, and the results for one particular semester were already shown above in Figure 6.6. For most periods we observe similar patterns, with two relatively large groups with anti-correlated strategies, and a third group containing the remaining banks which are mostly uncorrelated with the rest and therefore have no clear strategy in this particular semester. For some semesters, the usefulness of the filtering approach is more obvious than for the example above, cf. Figure 6.12 for the first half of 1999. In this period, it is hard to identify any structure in the raw correlation matrix, even after reordering the indices according to the leading eigenvector. In contrast, the filtered correlation matrix shows a similar structure as the example above, however, with one of the groups being rather small.

6.3.4 Time-Persistence of Clusters

So far we have only looked at the structure of the correlation matrices for each semester individually, finding that there is in fact evidence for non-random structure in the matrices. Here we aim at assessing the persistence of the identified clusters over the sample period. We already discussed the time-persistence of individual correlation pairs in the trading strategies. Here we check whether two sample banks tend to appear in the same clusters over time. If so, this would be evidence for the non-spuriousness of the identified clusters and therefore suggest that the identified clusters are meaningful. In order to shed light on this issue, we construct a frequency matrix $\mathbf{F}$, where

[32] In principle, we could divide the clusters even further by using the next eigenvectors. However, we find that the usage of the leading eigenvector is already sufficient to identify the large clusters.

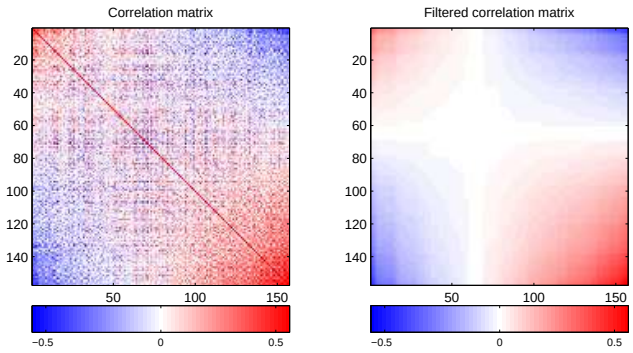

Fig. 6.12: Example of a correlation matrix after applying the clustering approach for the first half of 1999. Left: raw correlation matrix. Right: filtered matrix. Warm colors represent positive correlations, cold colors negative ones. White indicates uncorrelated strategies.

each element $f_{ij} \in [0, 1]$ contains the number of semesters that two banks were in the same cluster minus the number of semesters that they were in different clusters, relative to the total number of semesters that the two banks were jointly active.

Figure 6.13 shows the results of this exercise. The left panel shows the raw matrix, the right panel shows the same matrix, but after reordering the indices according to the size of the components of the leading eigenvector of this matrix. The results are striking: the two clusters appear even for the complete sample period since there are many large entries (in absolute terms) in the frequency matrix. Hence, there is not only stability in the individual correlations, but also in the group memberships of individual banks. Note that this is strong evidence for the non-randomness of the half-yearly matrices, so the identified clusters are indeed significant. If group membership was completely random, it would be impossible to observe this phenomenon.

Another way to investigate the persistence of the cluster memberships is to construct a transition matrix of going from the one to the other cluster. We find that the two clusters contain those banks which mainly bought and sold money (see below), respectively, so we label the clusters B and S. Taking the possibility of exit into account (E), Table 6.1 shows for example, that there is a probability of 68.32% that a bank which was in the buy cluster in the last period (B_{t-1}) will remain in this cluster in the next period. This is way above 50%, the level we would expect if cluster membership was completely random (ignoring the possibility of exit). Note that the diagonal elements in this matrix are largest, so there is a significant level of persis-

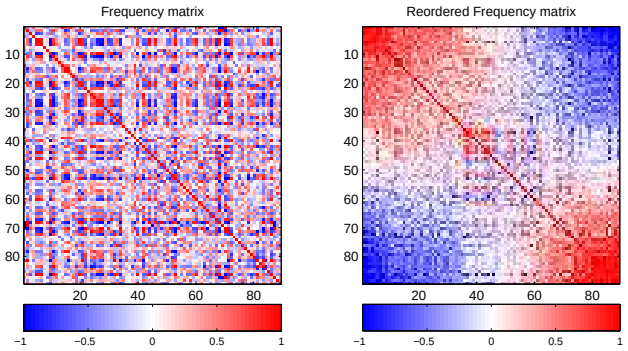

Fig. 6.13: Frequency matrix **F**, as defined in the text. Left: raw matrix. Right: after reordering.

	B_t	S_t	E_t
B_{t-1}	.6832	.2118	.1050
S_{t-1}	.1807	.6996	.1197
E_{t-1}	.0272	.0331	.9397

Tab. 6.1: Transition matrix: B, S and E stand for buyer, seller, and exit, respectively.

tence in the cluster memberships. This is also in line with the findings of Fricke and Lux (2012), who show that there is a substantial level of persistence in the banks' strategies (in terms of being intermediaries, lenders, or borrowers) based on quarterly aggregates. We should also stress that, since our clustering approach is only based on the sign of the components of the leading eigenvectors, there are a number of elements in these which are close to zero i.e. insignificant. Therefore, it comes as no surprise that the probability of switching from the buy to the sell cluster (or vice versa) is also quite high, with a value of 21.18%. This is driven by banks with less clear trading strategies, which might put them randomly in one or the other cluster. On average, this effect should vanish since the switching probabilities are of comparable magnitude.

We checked that these results are very stable over time, cf. Figure 6.14. Interestingly, similar to the remarks in Fricke and Lux (2012), there appears to be a substantial reversal in the trading strategies in the second half of 2001 (semester 5), which Fricke and Lux (2012) identified as a significant structural break in many time-series of the e-MID dataset. Interestingly, the GFC (starting around semester 18) is hardly visible, which is somewhat

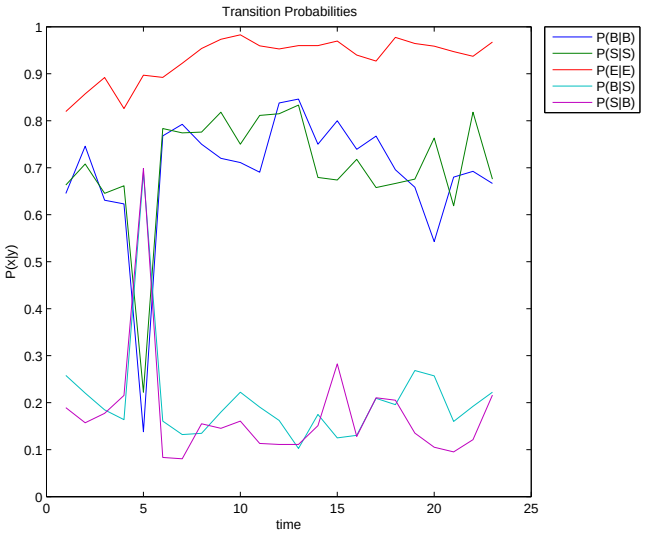

Fig. 6.14: Transition probabilities over time. $P(x|y)$ shows the probability of going from status y to x. B, S and E stand for buyer, seller, and exit, respectively.

surprising given that core banks tended to reverse their strategies during this time.

6.3.5 A Closer Look at the Clusters

In this section we want to take a closer look at the individual clusters. The usual interpretation of the finding of two anti-correlated clusters is that banks in the same cluster have similar strategies (i.e. tend to trade in the same direction), trading with banks in the other cluster. For example, Iori *et al.* (2007) state that the two clusters contain large and small banks, respectively, which tend to trade with each other. At first sight, this interpretation makes intuitive sense, since two strategies can only be positively correlated, if the two banks tend to trade in the same directions during the same trading sessions. Note that the statement *during the same trading sessions* is important, a simple example illustrates this point: suppose bank A is a continuous buyer during a single semester. With equal probability it buys 1 Euro on the market or is inactive in a particular trading session. The cumulative trading volume of bank A will therefore be a monotonically increasing function of time. Now suppose that bank B is also a continuous buyer of money, always buying 1 Euro when bank A is inactive. Again bank B's cumulative trading volume will be a monotonically increasing function of time. What does

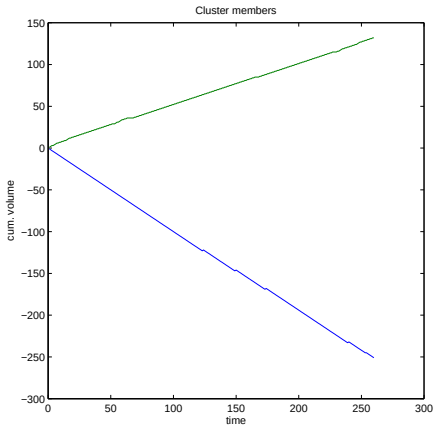

Fig. 6.15: Most anti-correlated (signed) trading strategies in semester 3.

the correlation of the banks' trading strategies tell us about their relationship? By definition, the correlation between the trading strategies of the two banks will be -1. Observed negative correlations may therefore be spurious, in the sense that two strategies may turn out to be identical in terms of the cumulative trading volumes, but not in terms of the individual trading sessions. Quite interestingly, this fact has not been mentioned before and we know of no paper that investigated the trading strategies in this detail after identifying the clusters.[33]

To illustrate these points, consider Figures 6.15 and 6.16, where we show the cumulative trading volumes (signed strategies) of the two most anti-correlated banks in semester 3 and 10, respectively. Figure 6.16 shows the usual pattern, namely one bank continuously buys money in the market, whereas the other bank continuously sells money in the market, and the banks tend to trade in the same trading sessions. In contrast, Figure 6.15 shows the case where banks with, in terms of the sign of the cumulative trading volume, very similar strategies, end up in different clusters simply because they tend to trade asynchronously.

All in all, this indicates that we might not really measure what we would like to measure using our definition of trading strategies. Since we focus only on the relatively active banks, the effect is likely to be small for relatively large values of θ. In order to investigate this in more detail, we checked that one cluster indeed contains banks that tend to buy most of the time, while banks in the other cluster tend to sell most of the time. As another

[33] Note that the continuous approach of Iori *et al.* (2007) also suffers from this problem.

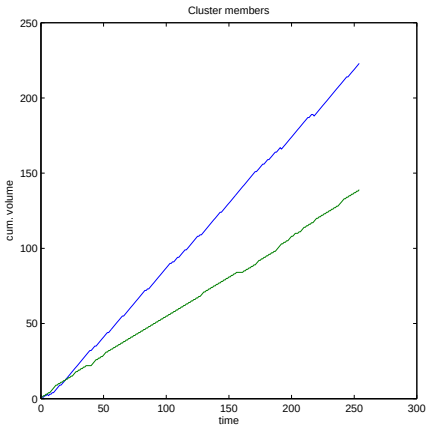

Fig. 6.16: Most anti-correlated (signed) trading strategies in semester 10.

robustness check, we calculated the correlation between trading strategies, using only those trading periods where both banks were jointly active. While this affects particularly those cases where banks tend to trade asynchronously, the general patterns observed in the baseline correlation matrices remain unaffected. We also checked that the total trading volumes in the two clusters are roughly equal, so the clusters indeed tend to trade with each other. Nevertheless, the examples show, that our approach may at times not be able to identify meaningful clusters. Summing up, the story is much simpler than expected: the two clusters contain those banks with clear trading strategies over time, i.e. mostly buying or selling money during a particular semester.

We also checked the results with respect to the CP model in Fricke and Lux (2012), since it might be that one cluster contains the core banks (which are similar across many dimensions and possibly even in terms of their trading strategies), and the other cluster contains the periphery banks. However, we do not find any evidence in this regard.[34] Figure 6.17 shows an example to highlight that there is apparently no relationship between the coreness of banks in one or the other cluster. There we see the same (filtered and reordered) correlation matrix as in Figure 6.6, but with core banks explicitly marked by a black dot on the main diagonal. There seems to be no relationship between the core-periphery model and the trading clusters, since the core banks are spread somewhat evenly over the network. However, given

[34] Note that this is in contrast to the findings of Iori *et al.* (2007), since the coreness of individual banks was found to be highly correlated with transaction volumes (as proxy for bank size) of individual banks.

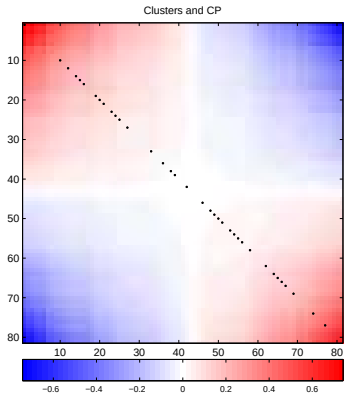

Fig. 6.17: Relation between clusters and core-periphery model. Filtered correlation matrix for the first half of 2006, with core banks highlighted by black dots on the main diagonal.

that core banks tend to act on both market sides, it should be clear that the core banks usually do not appear among the most anti-correlated (buy and sell) strategies. Similar results hold for other sample periods. Thus, the structure of the correlation matrices is not driven by the core membership of individual banks. This is not surprising, given that most core banks act as intermediaries and tend to distribute money across the complete system, rather than building up large positions in a continuous way. Therefore, core banks are spread evenly across the correlation network.

It would, of course, be very interesting to relate our findings to external information about the sample banks, e.g. using balance sheet data. However, the anonymity of our dataset makes such an exercise impossible. It would be interesting to carry out such an analysis in the future based on different non-anonymous datasets. In this way, the aim would be to model the behavior of individual banks based on observable characteristics in order to come up with realistic agent-based models of the interbank market.

6.4 Conclusions

In this paper, we analyzed the correlations in patterns of trading for members of the Italian interbank trading platform e-MID. We showed that there are significant and persistent bilateral correlations between institutions' trading strategies, and in most semesters we find evidence for the existence of two anti-correlated clusters. The two clusters mostly contain continuous net buyers and net sellers of money, respectively. The clusters, in terms

of the individual cluster memberships, are highly persistent. However, we have also seen that there are certain problems related to our definition of trading strategies. Additionally, we highlight some problems related to our definition of trading strategies, since the observed negative correlations may be spurious. The reason lies in the fact that two strategies may be identical in terms of the cumulative trading volumes, but not in terms of the trading time.

Our findings add further evidence on the fact that preferential lending relationships on the micro-level lead to community structure on the macro-level. Furthermore, despite finding no evidence of community structure in the usual network of interbank liabilities apart from the CP structure, see Fricke and Lux (2012), trading on the e-MID platform appears to be a relatively structured process in terms of trading strategies. Given that each trade involves two counterparties, i.e. a buying and a selling side, it may appear trivial that we identify these clusters. However, the high level of persistence shows that most banks tend to have rather stable trading strategies over time, since they tend to appear in the same cluster over time, trading with the same counterparties.

In the future, we need to explore the trading behavior of individual banks in more detail, to understand the evolution of the interbank network in more detail. The main aim is to build an artificial banking system, that allows to test the effects of different regulatory measures in a laboratory setting. We are still at the beginning of understanding the complexity of actual banking networks.

Bibliography

Affinito, M. (2011): "Do Interbank Customer Relationships Exist? And How Did They Function in the Crisis? Learning from Italy," Temi di discussione (Economic working papers) 826, Bank of Italy, Economic Research Department.

Ahn, H.-J., C. Q. Cao, and H. Choe (1996): "Tick Size, Spread, and Volume," *Journal of Financial Intermediation*, 5(1), 2–22.

Albert, R., H. Jeong, and A.-L. Barabasi (2000): "Error and Attack Tolerance of Complex Networks," *Nature*, (406), 378–482.

Alderson, D. L., and L. Li (2007): "Diversity of Graphs with Highly Variable Connectivity," *Physical Review E*, 75, 046102.

Allen, F., and D. Gale (2000): "Financial Contagion," *Journal of Political Economy*, 108(1), 1–33.

Allen, H., and M. P. Taylor (1990): "Charts, Noise and Fundamentals in the London Foreign Exchange Market," *Economic Journal*, 100(400), 49–59.

Anderson, C. W. (1970): "Extreme Value Theory for a Class of Discrete Distributions with Applications to Some Stochastic Processes," *Journal of Applied Probability*, 7(1), 99–113.

Anderson, P. W. (1972): "More Is Different," *Science*, 177(4047), 393–396.

Anufriev, M., and G. Bottazzi (2004): "Asset Pricing Model with Heterogeneous Investment Horizons," LEM Papers Series 2004/22, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy.

Avnir, D., O. Biham, D. Lidar, and O. Malcai (1998): "Is the Geometry of Nature Fractal?," *Science*, 279(5347), 39–40.

Axtell, R. L. (2001): "Zipf Distribution of U.S. Firm Sizes," *Science*, 293(5536), 1818–1820.

Bae, K.-H., G. A. Karolyi, and R. M. Stulz (2003): "A New Approach to Measuring Financial Contagion," *Review of Financial Studies*, 16(3), 717–763.

Bak, P., M. Paczuski, and M. Shubik (1997): "Price Variations in a Stock Market with many Agents," *Physica A: Statistical and Theoretical Physics*, 246(3-4), 430 – 453.

Baltagi, B., D. Li, and Q. Li (2006): "Transaction Tax and Stock Market Behavior: Evidence from an Emerging Market," *Empirical Economics*, 31(2), 393–408.

Bank for International Settlements (2011): "OTC Derivatives Market Activity in the Second Half of 2011 (Press Release)," *http://www.bis.org/press/p120509.htm*.

Barabasi, A.-L., and R. Albert (1999): "Emergence of Scaling in Random Networks," *Science*, 286(5439), 509–512.

Bech, M., and E. Atalay (2010): "The Topology of the Federal Funds Market," *Physica A*, 389(22), 5223–5246.

Beja, A., and M. B. Goldman (1980): "On the Dynamic Behavior of Prices in Disequilibrium," *Journal of Finance*, 35(2), 235–48.

Biais, B., P. Hillion, and C. Spatt (1995): "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse," *Journal of Finance*, 50(5), 1655–89.

Bordo, M., B. Eichengreen, D. Klingebiel, and M. S. Martinez-Peria (2001): "Is the Crisis Problem Growing More Severe?," *Economic Policy*, 16(32), 51–82.

Borgatti, S. P., and M. G. Everett (2000): "Models of Core/Periphery Structures," *Social Networks*, 21(4), 375 – 395.

Borland, L., and J.-P. Bouchaud (2005): "On a Multi-Timescale Statistical Feedback Model for Volatility Fluctuations," Science & Finance (CFM) working paper archive 500059, Science & Finance, Capital Fund Management.

Boss, M., H. Elsinger, M. Summer, and S. Thurner (2004): "Network Topology of the Interbank Market," *Quantitative Finance*, 4(6), 677–684.

Bottazzi, G., G. Dosi, and I. Rebesco (2005): "Institutional Architectures and Behavioral Ecologies in the Dynamics of Financial Markets," *Journal of Mathematical Economics*, 41(1-2), 197–228.

Bouchaud, J.-P. (2008): "Economics Needs a Scientific Revolution," *Nature*, 455, 1181.

Bouchaud, J.-P., J. D. Farmer, and F. Lillo (2008): *How markets slowly digest changes in supply and demand, In: Handbook of Financial Markets (Eds.: T. Hens and K. R. Schenk-Hoppe), Ch. 2*. North Holland.

Bouchaud, J.-P., Y. Gefen, M. Potters, and M. Wyart (2003): "Fluctuations and Response in Financial Markets: The Subtle Nature of 'Random' Price Changes," Science & Finance (CFM) working paper archive 0307332, Science & Finance, Capital Fund Management.

Bouchaud, J.-P., J. Kockelkoren, and M. Potters (2004): "Random Walks, Liquidity Molasses and Critical Response in Financial Markets," Science & Finance (CFM) working paper archive 500063, Science & Finance, Capital Fund Management.

Bouchaud, J.-P., A. Matacz, and M. Potters (2001): "The Leverage Effect in Financial Markets: Retarded Volatility and Market Panic," Science & Finance (CFM) working paper archive 0101120, Science & Finance, Capital Fund Management.

Bouchaud, J.-P., M. Mézard, and M. Potters (2002): "Statistical Properties of Stock Order Books: Empirical Results and Models," *Quantitative Finance*, 2(4), 251–256.

Bouchaud, J.-P., and M. Potters (2003): *Theory of Financial Risk and Derivative Pricing - From Statistical Physics to Risk Management (2nd edition)*. Cambrigde University Press.

Boyd, J. P., W. J. Fitzgerald, and R. J. Beck (2006): "Computing Core/Periphery Structures and Permutation Tests for Social Relations Data," *Social Networks*, 28, 165–178.

Boyd, J. P., W. J. Fitzgerald, M. C. Mahutga, and D. A. Smith (2010): "Computing Continuous Core/Periphery Structures for Social Relations Data with MINRES/SVD," *Social Networks*, 32(2), 125 – 137.

Braeuning, F. (2011): "Relationship Lending and Peer Monitoring: Evidence From Interbank Payment Data," *Tinbergen Institute, mimeo*.

Brandes, U., and D. Wagner (2004): "visone-Analysis and Visualization of Social Networks," in *Graph Drawing Software*, ed. by M. Jünger, and P. Mutzel, pp. 321–340. Springer-Verlag.

Brede, M., and B. J. de Vries (2009): "Networks that Optimize a Trade-off Between Efficiency and Dynamical Resilience," *Physics Letters A*, 373(43), 3910 – 3914.

Brogaard, J. (2010): "High Frequency Trading and its Impact on Market Quality," *Social Science Research Network Working Paper Series*.

Brunnermeier, M. K. (2008): "Deciphering the Liquidity and Credit Crunch 2007-08," Working Paper 14612, National Bureau of Economic Research.

Brusco, M. (2011): "An Exact Algorithm for a Core/Periphery Bipartitioning Problem," *Social Networks*, 33(1), 12 – 19.

Butts, C., and K. Carley (2001): "Multivariate Methods for Interstructural Analysis," CASOS Working Paper, Carnegie Mellon University.

Caccioli, F., T. A. Catanach, and J. Doyne Farmer (2011): "Heterogeneity, correlations and financial contagion," *ArXiv e-prints*.

Caldarelli, G. (2007): *Scale-Free Networks*. Oxford University Press.

Cao, C., O. Hansch, and X. Wang (2008): "Order Placement Strategies In A Pure Limit Order Book Market," *Journal of Financial Research*, 31(2), 113–140.

Challet, D., and R. Stinchcombe (2001): "Analyzing and Modelling 1+1d Markets," *Physica A*, 300(1-2), 285–299.

Challet, D., and Y. C. Zhang (1997): "Emergence of Cooperation and Organization in an Evolutionary Game," *Physica A: Statistical and Theoretical Physics*, 246(3-4), 407 – 418.

Chiarella, C., and G. Iori (2002): "A Simulation Analysis of the Microstructure of Double Auction Markets," *Quantitative Finance*, 2(5), 346–353.

Chiarella, C., G. Iori, and J. Perello (2009): "The Impact of Heterogeneous Trading Rules on the Limit Order Book and Order Flows," *Journal of Economic Dynamics and Control*, 33(3), 525–537.

Christley, R. M., G. L. Pinchbeck, R. G. Bowers, D. Clancy, N. P. French, R. Bennett, and J. Turner (2005): "Infection in Social Networks: Using Network Analysis to Identify High-Risk Individuals," *American Journal of Epidemiology*, 162(10), 1024–1031.

Clauset, A., C. Rohilla Shalizi, and M. E. J. Newman (2009): "Power-law Distributions in Empirical Data," *SIAM Review*, 51, 661–703.

Clauset, A., M. Young, and K. S. Gleditsch (2007): "On the Frequency of Severe Terrorist Events," *Journal of Conflict Resolution*, 51(1), 58–87.

Cliff, D., and J. Bruten (1997): "Zero is Not Enough: On the Lower Limit of Agent Intelligence for Continuous Double Auction Markets," Technical Report HPL-97-141, Hewlett-Packard Laboratories.

Cocco, J. F., F. J. Gomes, and N. C. Martins (2009): "Lending Relationships in the Interbank Market," *Journal of Financial Intermediation*, 18(1), 24–48.

Colander, D., M. Goldberg, A. Haas, K. Juselius, A. Kirman, T. Lux, and B. Sloth (2009): "The Financial Crisis and the Systemic Failure of Academic Economics," *Critical Review*, 21(2-3), 249–267.

Craig, B., and G. von Peter (2010): "Interbank Tiering and Money Center Banks," Discussion Paper, Series 2: Banking and Financial Studies 12/2010, Deutsche Bundesbank.

Daniels, M. G., J. Doyne Farmer, L. Gillemot, G. Iori, and E. Smith (2001): "A Quantitative Model of Trading and Price Formation in Financial Markets," *ArXiv Condensed Matter e-prints*.

Danielsson, J., L. de Haan, L. Peng, and C. de Vries (2001): "Using a Bootstrap Method to Choose the Sample Fraction in Tail Index Estimation," *Journal of Multivariate Analysis*, 76(2), 226 – 248.

de Bandt, O., and P. Hartmann (2000): "Systemic Risk: a Survey," Working Paper Series 35, European Central Bank.

De Masi, G., G. Iori, and G. Caldarelli (2006): "Fitness Model for the Italian Interbank Money Market," *Phys. Rev. E*, 74(6), 66112.

Debreu, G. (1974): "Excess Demand Functions," *Journal of Mathematical Economics*, 1(1), 15–21.

DEMARY, M. (2010): "Transaction Taxes and Traders with Heterogeneous Investment Horizons in an Agent-Based Financial Market Model," *Economics: The Open-Access, Open-Assessment E-Journal*, 4(2010-8).

DORE, R. (2008): "Financialization of the Global Economy," *Industrial and Corporate Change*, 17(6), 1097–1112.

DOREIAN, P., V. BATAGELJ, AND A. FERLIGOJ (2005): *Generalized Blockmodelling*. Cambridge University Press.

DREES, H., AND E. KAUFMANN (1998): "Selecting the Optimal Sample Fraction in Univariate Extreme Value Estimation," *Stochastic Processes and their Applications*, 75(2), 149 – 172.

EHRENSTEIN, G. (2002): "Cont-Bouchaud Percolation Model Including Tobin tax," *International Journal of Modern Physics*, 13, 1323–1331.

ERDÖS, P., AND A. RENYI (1959): "On Random Graphs," *Publicationes Mathematicae*, 6, 290–297.

EUROPEAN CENTRAL BANK (2007): "Euro Money Market Study 2006," Discussion paper, ECB.

EVERETT, M. G., AND S. P. BORGATTI (2000): "Peripheries of Cohesive Subsets," *Social Networks*, 21(4), 397 – 407.

FAMA, E. F. (1965): "The Behavior of Stock-Market Prices," *The Journal of Business*, 38(1), 34–105.

——— (1970): "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance*, 25(2), 383–417.

FARMER, J. D., AND D. FOLEY (2009): "The Economy Needs Agent-Based Modelling," *Nature*, 460, 685–686.

FARMER, J. D., L. GILLEMOT, F. LILLO, S. MIKE, AND A. SEN (2004): "What Really Causes Large Price Changes?," *Quantitative Finance*, 4(4), 383–397.

FARMER, J. D., AND F. LILLO (2004): "On the Origin of Power Law Tails in Price Fluctuations," *Quantitative Finance*, 4(1), 7–11.

FINGER, K., D. FRICKE, AND T. LUX (2012): "Network Analysis of the e-MID Overnight Money Market: The Informational Value of Different Aggregation Levels for Intrinsic Dynamic Processes," Kiel Working Paper 1782, Kiel Institute for the World Economy.

FINGER, K., AND T. LUX (2013): "The Evolution of the Banking Network: An Actor-oriented Approach," unpublished manuscript.

FOUCAULT, T., O. KADAN, AND E. KANDEL (2005): "Limit Order Book as a Market for Liquidity," *Review of Financial Studies*, 18(4), 1171–1217.

FRANKE, R. (2008): "On the Interpretation of Price Adjustments and Demand in Asset Pricing Models with Mean-Variance Optimization," Economics Working Papers 2008,13, Christian-Albrechts-University of Kiel, Department of Economics.

FRANKE, R., AND T. ASADA (2008): "Incorporating Positions Into Asset Pricing Models With Order-Based Strategies," *Journal of Economic Interaction and Coordination*, 3(2), 201–227.

FRANKE, R., AND F. WESTERHOFF (2012): "Structural Stochastic Volatility in Asset Pricing Dynamics: Estimation and Model Contest," *Journal of Economic Dynamics and Control*, 36(8), 1193 – 1211.

FREIXAS, X., B. M. PARIGI, AND J.-C. ROCHET (2000): "Systemic Risk, Interbank Relations and Liquidity Provision by the Central Bank," *Journal of Money, Credit and Banking*, 32(3), 611–638.

FRICKE, D. (2012): "Trading Strategies in the Overnight Money Market: Correlations and Clustering on the e-MID Trading Platform," *Physica A*, (391), 6528–6542.

FRICKE, D., AND T. LUX (2012): "Core-Periphery Structure in the Overnight Money Market: Evidence from the e-MID Trading Platform," Kiel working paper 1759, Kiel Institute for the World Economy.

——— (2013): "On the Distribution of Links in the Interbank Network: Evidence from the e-MID Overnight Money Market," unpublished manuscript.

FRIEDMAN, S., AND H. F. WEISBERG (1981): "Interpreting the First Eigenvalue of a Correlation Matrix," *Educational and Psychological Measurement*, 41(1), pp. 11–21.

GABAIX, X. (1999): "Zipf's Law for Cities: An Explanation," *The Quarterly Journal of Economics*, 114(3), 739–767.

GABAIX, X., AND R. IBRAGIMOV (2011): "Rank -1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents," *Journal of Business & Economic Statistics*, 29(1), 24–39.

Gai, P., A. Haldane, and S. Kapadia (2011): "Complexity, Concentration and Contagion," *Journal of Monetary Economics*, 58(5), 453–470.

Giardina, I., and J.-P. Bouchaud (2003a): "Bubbles, Crashes and Intermittency in Agent Based Market Models," *The European Physical Journal B - Condensed Matter and Complex Systems*, 31, 421–437.

Giardina, I., and J.-P. Bouchaud (2003b): "Volatility Clustering in Agent Based Market Models," *Physica A: Statistical Mechanics and its Applications*, 324(1-2), 6 – 16.

Gigerenzer, G. (2008): *Gut Feelings: The Intelligence of the Unconscious*. Penguin (Non-Classics).

Gode, D. K., and S. Sunder (1993): "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, 101(1), 119–37.

Goh, K.-I., B. Kahng, and D. Kim (2001): "Universal Behavior of Load Distribution in Scale-Free Networks," *Phys. Rev. Lett.*, 87, 278701.

Goldstein, M. L., S. A. Morris, and G. G. Yen (2004): "Problems with Fitting to the Power-law Distribution," *The European Physical Journal B - Condensed Matter and Complex Systems*, 41, 255–258, 10.1140/epjb/e2004-00316-5.

Haldane, A. (2009): "Rethinking the Financial Network," Speech delivered in April 2009 at the Financial Student Association, Amsterdam. http://www.bankofengland.co.uk/publications/speeches/2009/speech386.pdf.

Haldane, A. G., and R. M. May (2011): "Systemic Risk in Banking Ecosystems," *Nature*, 469(7330), 351–355.

Hanke, M., J. Huber, M. Kirchler, and M. Sutter (2010): "The Economic Consequences of a Tobin Tax - An Experimental Analysis," *Journal of Economic Behavior and Organization*, 74, 58–71.

Harris, L. (2003): *Trading and Exchanges - Market Microstructure for Practitioners*. Oxford University Press.

Harris, L., and J. Hasbrouck (1996): "Market vs. Limit Orders: The SuperDOT Evidence on Order Submission Strategy," *Journal of Financial and Quantitative Analysis*, 31(02), 213–231.

Hau, H. (2006): "The Role of Transaction Costs for Financial Volatility: Evidence from the Paris Bourse," *Journal of the European Economic Association*, 4(4), 862–890.

Helbing, D. (2012): *Social Self-Organization: Agent-Based Simulations and Experiments to Study Emergent Social Behavior (Understanding Complex Systems)*. Springer.

Herrera, F., M. Lozano, and J. L. Verdegay (1998): "Tackling Real-Coded Genetic Algorithms: Operators and Tools for Behavioural Analysis," *Artificial Intelligence Review*, 12, 265–319.

Hill, B. M. (1975): "A Simple General Approach to Inference About the Tail of a Distribution," *The Annals of Statistics*, 3(5), 1163–1174.

Ho, T. S. Y., and A. Saunders (1985): "A Micro Model of the Federal Funds Market," *Journal of Finance*, 40(3), 977–88.

Holland, J. H. (1975): *Adaptation in natural and artificial systems*. MIT Press, Cambridge, MA, USA.

Inaoka, H., T. Ninomyia, K. Taniguchi, T. Shimizu, and H. Takayasu (2004): "Fractal Network Derived from Banking Transaction – An Analysis of Network Structures Formed by Financial Institutions," *Bank of Japan Working Papers*, pages(04-E-04), 1–22.

Iori, G., G. De Masi, O. V. Precup, G. Gabbi, and G. Caldarelli (2008): "A Network Analysis of the Italian Overnight Money Market," *Journal of Economic Dynamics and Control*, 32(1), 259–278.

Iori, G., S. Jafarey, and F. G. Padilla (2006): "Systemic Risk in the Interbank Market," *Journal of Economic Behavior & Organization*, 61(4), 525–542.

Iori, G., R. Renò, G. D. Masi, and G. Caldarelli (2007): "Trading Strategies in the Italian Interbank Market," *Physica A: Statistical Mechanics and its Applications*, 376, 467 – 479.

Jackson, P., and A. O'Donnell (1985): "The Effects of Stamp Duty on Equity Transactions and Prices in the UK Stock Exchange," Discussion Paper 25, Bank of England.

Jansen, D. W., and C. G. de Vries (1991): "On the Frequency of Large Stock Returns: Putting Booms and Busts into Perspective," *The Review of Economics and Statistics*, 73(1), 18–24.

Jones, C. M., and P. J. Seguin (1997): "Transaction Costs and Price Volatility: Evidence from Commission Deregulation," *American Economic Review*, 87(4), 728–37.

Kahneman, D., and A. Tversky (1973): "On the Psychology of Prediction," *Psychological Review*, 80(4), 1124–1131.

Kernighan, B. W., and S. Lin (1970): "An Efficient Heuristic Procedure for Partitioning Graphs," *The Bell system technical journal*, 49(1), 291–307.

Keynes, J. M. (1936): *The General Theory of Employment, Interest and Money*. Macmillan, London.

Kim, D.-H., and H. Jeong (2005): "Systematic Analysis of Group Identification in Stock Markets," *Phys. Rev. E*, 72(4), 046133.

Kirman, A. (1991): *Epidemics of Opinion and Speculative Bubbles in Financial Markets, In: Money and Financial Markets (Ed.: M. Taylor)*. Macmillan.

——— (1992): "Whom or What Does the Representatitve Individual Represent?," *Journal of Economic Perspectives*, 6(2), 117–136.

——— (1993): "Ants, Rationality, and Recruitment," *The Quarterly Journal of Economics*, 108(1), 137–56.

Kubelec, C., and F. Sa (2010): "The Geographical Composition of National External Balance Sheets: 1980-2005," Bank of England working papers 384, Bank of England.

Kyle, A. S. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53(6), 1315–35.

Kyriakopoulos, F., S. Thurner, C. Puhr, and S. W. Schmitz (2009): "Network and Eigenvalue Analysis of Financial Transaction Networks," *The European Physical Journal B - Condensed Matter and Complex Systems*, 71(4), 523–531.

La Spada, G., J. Doyne Farmer, and F. Lillo (2011): "Tick Size and Price Diffusion," in *Econophysics of Order-driven Markets*, ed. by F. Abergel, B. Chakrabarti, A. Chakraborti, and M. Mitra, New Economic Windows, pp. 173–187. Springer Milan.

Lachapelle, D. M. D., and D. Challet (2010): "Turnover, Account Value and Diversification of Real Traders: Evidence of Collective Portfolio Optimizing Behavior," *New Journal of Physics*, 12.

Laloux, L., P. Cizeau, J.-P. Bouchaud, and M. Potters (1999): "Noise Dressing of Financial Correlation Matrices," *Phys. Rev. Lett.*, 83(7), 1467–1470.

Lau, A. H.-L., H.-S. Lau, and J. R. Wingender (1990): "The Distribution of Stock Returns: New Evidence against the Stable Model," *Journal of Business & Economic Statistics*, 8(2), 217–23.

Leadbetter, M. (1983): "Extremes and Local Dependence in Stationary Sequences," *Zeitschrift fuer Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 65, 291–306.

LeBaron, B., W. B. Arthur, and R. Palmer (1999): "Time Series Properties of an Artificial Stock Market," *Journal of Economic Dynamics and Control*, 23(9-10), 1487–1516.

LeBaron, B., and L. Tesfatsion (2008): "Modeling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents," *American Economic Review*, 98(2), 246–50.

Lillo, F. (2007): "Limit order Placement as an Utility Maximization Problem and the Origin of Power Law Distribution of Limit Order Prices," *The European Physical Journal B*, 55(4), 453–459.

Livan, G., S. Alfarano, and E. Scalas (2011): "Fine Structure of Spectral Properties for Random Correlation Matrices: An Application to Financial Markets," *Phys. Rev. E*, 84, 016113.

Luo, F., B. Li, X.-F. Wan, and R. Scheuermann (2009): "Core and Periphery Structures in Protein Interaction Networks," *BMC Bioinformatics*, 10(Suppl 4), S8.

Lux, T. (1995): "Herd Behaviour, Bubbles and Crashes," *Economic Journal*, 105(431), 881–96.

——— (2008): *Stochastic Behavioral Asset Pricing Models and the Stylized Facts, In: Handbook of Financial Markets (Eds.: T. Hens and K. R. Schenk-Hoppe), Ch. 3.* North Holland.

——— (2011): "Network Theory is Sorely Required," *Nature*, (469), 303.

Lux, T., and M. Marchesi (1999): "Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market," *Nature*, (397), 498–500.

——— (2000): "Volatility Clustering in Financial Markets: a Microsimulation of Interacting Agents," *International Journal of Theoretical and Applied Finance*, 3(4), 169–196.

Lux, T., and S. Schornstein (2005): "Genetic Learning as an Explanation of Stylized Facts of Foreign Exchange Markets," *Journal of Mathematical Economics*, 41(1-2), 169–196.

Lux, T., and F. Westerhoff (2009): "Economics Crisis," *Nature Physics*, 5, 2–3.

Lynch, P., and G. Zumbach (2003): "Market Heterogeneities and the Causal Structure of Volatility," *Quantitative Finance*, 3(4), 320–331.

Mandelbrot, B. (1963): "The Variation of Certain Speculative Prices," *The Journal of Business*, 36, 394.

Mannaro, K., M. Marchesi, and A. Setzu (2008): "Using an Artificial Financial Market for Assessing the Impact of Tobin-like Transaction Taxes," *Journal of Economic Behavior & Organization*, 67(2), 445–462.

Mantel, R. R. (1974): "On the Characterization of Aggregate Excess Demand," *Journal of Economic Theory*, 7(3), 348–353.

Markose, S., S. Giansante, M. Gatkowski, and A. R. Shaghaghi (2010): "Too Interconnected To Fail: Financial Contagion and Systemic Risk in Network Model of CDS and Other Credit Enhancement Obligations of US Banks," Economics Discussion Papers 683, University of Essex, Department of Economics.

Martinez-Jaramillo, S. (2007): "Artificial Financial Markets: an Agent Based Approach to Reproduce Stylized Facts and to Study the Red Queen Effect," *PhD Thesis, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex.*

Marčenko, V., and L. A. Pastur (1967): "Distribution of Eigenvalues for Some Sets of Random Matrices," *Mat. Sb. (N.S.)*, 72(4), 507–536.

Maslov, S. (2000): "Simple Model of a Limit Order-Driven Market," *Physica A: Statistical Mechanics and its Applications*, 278(3-4), 571 – 578.

MENKHOFF, L. (1998): "The Noise Trading Approach – Questionnaire Evidence From Foreign Exchange," *Journal of International Money and Finance*, 17(3), 547–564.

MENKHOFF, L., AND U. SCHMIDT (2005): "The Use of Trading Strategies by Fund Managers: Some First Survey Evidence," *Applied Economics*, 37(15), 1719–1730.

MIKE, S., AND J. D. FARMER (2008): "An Empirical Behavioral Model of Liquidity and Volatility," *Journal of Economic Dynamics and Control*, 32(1), 200–234.

MILGRAM, S. (1967): "The Small-World Problem," *Psychol. Today*, 2, 60–67.

MISTRULLI, P. E. (2007): "Assessing Financial Contagion in the Interbank Market: Maximum Entropy versus Observed Interbank Lending Patterns," Temi di discussione (Economic working papers) 641, Bank of Italy, Economic Research Department.

MULLINS, N. C., L. L. HARGENS, P. K. HECHT, AND E. L. KICK (1977): "The Group Structure of Cocitation Clusters: A Comparative Study," *American Sociological Review*, 42(4), 552–562.

MUNIZ, A. S. G., AND C. R. CARVAJAL (2006): "Core/Periphery Structure Models: An Alternative Methodological Proposal," *Social Networks*, 28(4), 442 – 448.

MUNIZ, A. S. G., A. M. RAYA, AND C. R. CARVAJAL (2011): "Core Periphery Valued Models in Input-Output Field: A Scope From Network Theory," *Papers in Regional Science*, 90(1), 111–121.

NETOTEA, S., AND S. PONGOR (2006): "Evolution of Robust and Efficient System Topologies," *Cellular Immunology*, 244(2), 80 – 83.

NEWMAN, M. E. J. (2002): "Assortative Mixing in Networks," *Phys. Rev. Lett.*, 89, 208701.

NIER, E., J. YANG, T. YORULMAZER, AND A. ALENTORN (2007): "Network Models and Financial Stability," *Journal of Economic Dynamics and Control*, 31(6), 2033–2060.

PELLIZZARI, P., AND F. WESTERHOFF (2009): "Some Effects of Transaction Taxes Under Different Microstructures," *Journal of Economic Behavior & Organization*, 72(3), 850–863.

Persitz, D. (2009): "Power in the Heterogeneous Connections Model: The Emergence of Core-Periphery Networks," Working Papers 2009.42, Fondazione Eni Enrico Mattei.

Plerou, V., P. Gopikrishnan, B. Rosenow, L. A. N. Amaral, T. Guhr, and H. E. Stanley (2002): "Random Matrix Approach to Cross Correlations in Financial Data," *Phys. Rev. E*, 65, 066126.

Pollin, R., D. Baker, and M. Schaberg (2003): "Securities Transaction Taxes for U.S. Financial Markets," *Eastern Economic Journal*, 29(4), 527–558.

Potters, M., and J.-P. Bouchaud (2003): "More Statistical Properties of Order Books and Price Impact," *Physica A: Statistical Mechanics and its Applications*, 324(1-2), 133 – 140, Proceedings of the International Econophysics Conference.

Raberto, M., S. Cincotti, S. Focardi, and M. Marchesi (2003): "Traders' Long-Run Wealth in an Artificial Financial Market," *Computational Economics*, 22(2), 255–272.

Reiss, R.-D., and M. Thomas (2007): *Statistical Analysis of Extreme Values: with Applications to Insurance, Finance, Hydrology and Other Fields (3rd Edition)*. Birkhäuser Verlag.

Roll, R. (1989): "Price Volatility, International Market Links, and Their Implications for Regulatory Policies," *Journal of Financial Service Research*, 3, 211–246.

Schelling, T. C. (1969): "Models of Segregation," *American Economic Review*, 59(2), 488–93.

Sheldon, G., and M. Maurer (1998): "Interbank Lending and Systemic Risk: An Empirical Analysis for Switzerland," *Swiss Journal of Economics and Statistics (SJES)*, 134(IV), 685–704.

Shiller, R. J. (1981): "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *American Economic Review*, 71(3), 421–36.

Shiller, R. J. (2003): "From Efficient Markets Theory to Behavioral Finance," *Journal of Economic Perspectives*, 17(1), 83–104.

Silverberg, G., and B. Verspagen (2007): "The Size Distribution of Innovations Revisited: An Application of Extreme Value Statistics to Citation and Value Measures of Patent Significance," *Journal of Econometrics*, 139(2), 318–339.

Snijders, T. A. B. (1996): "Stochastic Actor-Oriented Dynamic Network Analysis," *Journal of Mathematical Sociology*, 21, 149–172.

——— (2001): "The Statistical Evaluation of Social Network Dynamics," *Sociological Methodology*, 31(1), 361–395.

Snijders, T. A. B., G. G. van de Bunt, and C. E. G. Steglich (2009): "Introduction to stochastic actor-based models for network dynamics," *Social Networks*, 32(1), 44–60.

Sonnenschein, H. (1972): "Market Excess Demand Functions," *Econometrica*, 40(3), 549–63.

Soramaki, K., M. L. Bech, J. Arnold, R. J. Glass, and W. Beyeler (2007): "The Topology of Interbank Payment Flows," *Physica A*, (379), 317–333.

Sornette, D., and S. V. der Becke (2009): "Crashes and High Frequency Trading," Discussion Paper No. 11-63, Swiss Finance Institute Zurich.

Stewart, G. W. (1993): "On the Early History of the Singular Value Decomposition," *SIAM Review*, 35(4), pp. 551–566.

Stolzenburg, U., and T. Lux (2011): "Identification of a Core-Periphery Structure Among Participants of a Business Climate Survey," *Eur. Phys. J. B*.

Stumpf, M. P. H., and P. J. Ingram (2005): "Probability Models for Degree Distributions of Protein Interaction Networks," *Europhysics Letters*, 71(1), 152–158.

Stumpf, M. P. H., P. J. Ingram, I. Nouvel, and C. Wiuf (2005): "Statistical Model Selection Applied to Biological Network Data," *Proceedings in Computational Systems Biology*, (3), 65–73.

Stumpf, M. P. H., and M. A. Porter (2012): "Critical Truths About Power Laws," *Science*, 335(6069), 665–666.

Stumpf, M. P. H., C. Wiuf, and R. M. May (2005): "Subnets of Scale-free Networks are not Scale-free: Sampling Properties of Networks," *Proceedings of the National Academy of Sciences of the United States of America*, 102(12), 4221–4224.

Thurner, S., J. D. Farmer, and J. Geanakoplos (2010): "Leverage Causes Fat Tails and Clustered Volatility," Cowles Foundation Discussion Papers 1745, Cowles Foundation, Yale University.

Tobin, J. (1978): "A Proposal for International Monetary Reform," *Eastern Economic Journal*, 4(3-4), 153–159.

Toth, B., Z. Eisler, F. Lillo, J. Bouchaud, J. Kockelkoren, and J. Doyne Farmer (2011): "How Does the Market React to Your Order Flow?," *ArXiv e-prints*.

Trichet, J. (2011): "Intellectual Challenges to Financial Stability Analysis in the Era of Macroprudential Oversight," *Financial Stability Review*, (15), 139–149.

Tumminello, M., F. Lillo, J. Piilo, and R. N. Mantegna (2012): "Identification of Clusters of Investors From Their Real Trading Activity in a Financial Market," *New Journal of Physics*, 14(1), 013041.

Tversky, A., and D. Kahneman (1974): "Judgment under Uncertainty: Heuristics and Biases," *Science (New Series)*, 185(4157), 237–251.

Upper, C., and A. Worms (2004): "Estimating Bilateral Exposures in the German Interbank Market: Is there a Danger of Contagion? Cross-Border Bank Contagion in Europe," *European Economic Review*, 48(4), 827–849.

Vriend, N. J. (2000): "An Illustration of the Essential Difference Between Individual and Social Learning, and its Consequences for Computational Analyses," *Journal of Economic Dynamics and Control*, 24(1), 1–19.

Watts, D. J., and S. H. Strogatz (1998): "Collective Dynamics of Small-World Networks," *Nature*, 393(6684), 440–442.

Westerhoff, F. (2003): "Heterogeneous Traders and The Tobin Tax," *Journal of Evolutionary Economics*, 13(1), 53–70.

——— (2004a): "Speculative Dynamics, Feedback Traders and Transaction Taxes: a Note," *Jahrbuch für Wirtschaftswissenschaften/Review of Economics*, 55, 190–195.

——— (2004b): "The Effectiveness of Keynes-Tobin Transaction Taxes When Heterogeneous Agents Can Trade in Different Markets: A Behavioral Finance Approach," Computing in Economics and Finance 2004 14, Society for Computational Economics.

Youssefmir, M., B. A. Huberman, and T. Hogg (1998): "Bubbles and Market Crashes," *Computational Economics*, 12(2), 97–114.

Zhou, H. (2002): "Scaling Exponents and Clustering Coefficients of a Growing Random Network," *Phys. Rev. E*, 66, 016125.

Zovko, I., and D. Farmer (2002): "The Power of Patience: a Behavioural Regularity in Limit-Order Placement," *Quantitative Finance*, 2(5), 387–392.

Zovko, I. I., and J. D. Farmer (2007): "Correlations and clustering in the trading of members of the London Stock Exchange," Quantitative Finance Papers 0709.3261, arXiv.org.

Zumbach, G., and P. Lynch (2001): "Heterogeneous Volatility Cascade in Financial Markets," *Physica A: Statistical Mechanics and its Applications*, 298(3-4), 521 – 529.

Eidesstattliche Erklärung

Ich erkläre hiermit, an Eides Statt, dass ich meine Doktorarbeit 'Coping with the Complexity of Financial Markets' selbständig und ohne fremde Hilfe angefertigt habe und dass ich alle von anderen Autoren wörtlich übernommenen Stellen, wie auch die sich an die Gedanken anderer Autoren eng anlehnenden Ausführungen meiner Arbeit, besonders gekennzeichnet und die Quellen nach den mir angegebenen Richtlinien zitiert habe.

Kiel, Januar 2013

Daniel Fricke

LaTeX