Centrality-Based Capital Allocations*

Adrian Alter,[a] Ben R. Craig,[b,c] and Peter Raupach[b]
[a]International Monetary Fund
[b]Deutsche Bundesbank
[c]Federal Reserve Bank of Cleveland

We look at the effect of capital rules on a banking system that is connected through correlated credit exposures and interbank lending. Keeping total capital in the system constant, the reallocation rules, which combine individual bank characteristics and interconnectivity measures of interbank lending, are to minimize a measure of system-wide losses. Using the detailed German credit register for estimation, we find that capital rules based on eigenvectors dominate any other centrality measure, saving about 15 percent in expected bankruptcy costs.

JEL Codes: G21, G28, C15, C81.

1. Introduction

"The difficult task before market participants, policymakers, and regulators with systemic risk responsibilities such as the

*The views expressed in this paper are those of the authors and do not necessarily reflect those of the Deutsche Bundesbank, the Eurosystem, the Federal Reserve Bank of Cleveland, the International Monetary Fund (IMF), its Executive Board, or IMF policies. We are very thankful for comments from Günter Franke, Andrew Haldane, Moritz Heimes, Christoph Memmel, Camelia Minoiu, Rafael Repullo, Almuth Scholl, Vasja Sivec, Martin Summer, Alireza Tahbaz-Salehi, participants at the Annual International Journal of Central Banking Research Conference hosted by the Federal Reserve Bank of Philadelphia, the Final Conference of the Macro-prudential Research Network (MaRs) hosted by the ECB, the EUI Conference on Macroeconomic Stability, Banking Supervision and Financial Regulation, and seminar participants at the Bundesbank and the IMF. Author contact: Alter: International Monetary Fund, 700 19th St. NW, Washington DC, 20431, USA; e-mail: aalter@imf.org. Craig: Deutsche Bundesbank, Research Centre, Wilhelm-Epstein-Str. 14, 60431 Frankfurt-am-Main, Germany and Federal Reserve Bank of Cleveland, Cleveland, OH, USA; e-mail: ben.r.craig@clev.frb.org. Raupach: Deutsche Bundesbank, Research Centre, Wilhelm-Epstein-Str. 14, 60431 Frankfurt-am-Main, Germany; e-mail: peter.raupach@bundesbank.de.

> Federal Reserve is to find ways to preserve the benefits of interconnectedness in financial markets while managing the potentially harmful side effects." Yellen (2013)

This paper examines capital requirements that mitigate the harmful side effects of interconnectedness in the context of a model of interbank contagion. Although the model is fairly classical in the way it handles contagion, it uses a very rich data set of credit exposures of a large domestic banking system, the fifth largest system in the world. What we find is that the same nationwide amount of required capital, when distributed in part based on the interconnectedness of the system, performs better in terms of total losses to the system than the same amount of capital when allocated on the basis of banks' individual risk-weighted assets alone. Indeed, a capital allocation based upon our best network centrality measure saves some 15 percent of expected bankruptcy costs, which is our preferred measure of total system losses.

The idea of tying capital charges to interbank exposures and interconnectedness in order to improve the stability of the banking system—i.e., to minimize expected social costs (arising from bailouts, growth effects, or unemployment, for example)—is in the spirit of the regulatory assessment methodology for systemically important financial institutions (SIFIs) proposed by the Basel Committee on Banking Supervision (2011). In contrast to that methodology, however, our study determines an optimal rule for capital charges that is based on interconnectedness measures as well as on the portfolio risk of bank assets, and we then compare the results under the different capital allocations.

We focus on two main sources of systemic risk: correlated credit exposures and interbank connectivity. First, banks' balance sheets can be simultaneously affected by macro or industry shocks, since the credit risk of their non-bank borrowers is correlated. If losses are large, capital of the entire system is eroded, making the system less stable. Second, these shocks can trigger the default of certain financial institutions and, again, erode bank capital further. The second effect is modeled in the interbank market. Since banks are highly connected through interbank exposures, we focus on those negative tail events in which correlated losses of their portfolios trigger contagion in the interbank market.

Our model comes close to the framework proposed by Elsinger, Lehar, and Summer (2006) and Gauthier, Lehar, and Souissi (2012), which combines common credit losses with interbank network effects and externalities in the form of asset fire sales. The aim of our paper is different. We propose a tractable framework to reallocate capital for large financial systems in order to minimize contagion effects and the possible costs of a public bailout. We contrast two different capital allocations: the benchmark case, in which we allocate capital based on the risks in individual banks' portfolios, and a comparison case, where we allocate capital based partly on some interbank network metrics (such as *degree*, *eigenvector*, or *betweenness*) that capture the potential contagion risk of individual banks. Literature has shown that the network structure matters. For instance, Sachs (2014) randomly generates interbank networks and investigates contagion effects in different setups. She finds that the distribution of interbank exposures plays a crucial role for system stability and confirms the "knife-edge" or tipping-point feature (as mentioned by Haldane 2009) of highly interconnected networks.

We compare different capital structures in which the total capital requirement to the entire banking system is constant but the capital charge varies across banks, based on the network metric chosen and the weight we put on it. Both the choice of a metric and the weight are optimized. Among various sensible target functions, to minimize, we select *total expected bankruptcy costs* of defaulted banks. This measure of system losses is especially interesting, as it represents a deadweight social loss and is independent of distributional considerations which would arise if we focused on the losses incurred by a certain group of bank claimants such as depositors.

We use the credit register for the German banking system. It records every bilateral lending relationship in excess of €1.5 million, including interbank lending. The richness of our data set allows us to do two things. First, we can compute centrality measures accurately. Second, we achieve a comparably high precision in exploring the implications of both the joint credit risk and the interconnected direct claims in the banking system. Using a state-of-the-art credit portfolio model, we can derive the joint distribution functions of the shocks to the banks within the system and feed the shocks into the interbank lending network, so that we can simulate how they work their way through the system.

To model the credit risk arising from exposures to the real economy, which we call *fundamental* credit risk, we generate correlated credit losses by means of the risk engine CreditMetrics, which is often used in bank risk management (Bluhm, Overbeck, and Wagner 2003, ch. 2). Based on a multi-factor credit risk model, the engine helps us to deal with risk concentration caused by large exposures to a single sector or highly correlated sectors. Even explicit common credit exposures, caused by firms borrowing from multiple banks, are precisely addressed. CreditMetrics assigns realistic, well-founded probabilities to those scenarios that have particularly large losses across the entire banking system. These bad scenarios are our main focus, since capital across financial institutions is eroded simultaneously and the banking system becomes more prone to interbank contagion.

Moreover, we model interbank contagion as in Rogers and Veraart (2013), which extends Eisenberg and Noe (2001) to include bankruptcy costs. This allows us to measure expected contagion losses and to observe the propagation process. To empirically implement our framework, we use several sources of information: the German central credit register (covering large loans), aggregated credit exposures (small loans), balance sheet data (e.g., total assets), market data (e.g., to compute sector correlations in the real economy or credit spreads), and data on rating transitions (to calibrate the CreditMetrics model). The approach can be applied in any country or group of countries where this type of information is available.

A major advantage of our framework is that policymakers can deal with large banking systems, making the regulation of systemic risk more tractable: while Gauthier, Lehar, and Souissi (2012) state that their model requires substantial numeric effort even with the six Canadian banks considered in their paper, German regulators have to deal with more than 1,500 banking groups, which is possible when taking our approach.

This study is related to several strands of the literature including applications of network theory to economics, macroprudential regulations, and interbank contagion. Cont, Moussa, and Santos (2013) find that not only banks' capitalization and interconnectedness are important for spreading contagion but also the vulnerability of the neighbors of poorly capitalized banks. Gauthier, Lehar, and Souissi (2012) use different holdings-based systemic risk measures (e.g., marginal expected shortfall, ΔCoVaR, Shapley value) to reallocate capital in the banking system and to determine macroprudential capital requirements. Using the Canadian credit register data for a system of six banks, they rely on an "Eisenberg-Noe"-type clearing mechanism extended to incorporate asset fire-sale externalities. In contrast to their paper, we reallocate capital based on centrality measures extracted directly from the network topology of the interbank market. Webber and Willison (2011) assign systemic capital requirements, optimizing over the aggregated capital of the system. They find that systemic capital requirements are directly related to bank size and interbank liabilities. Tarashev, Borio, and Tsatsaronis (2010) claim that systemic importance is mainly driven by size and exposure to common risk factors. In order to determine risk contributions, they utilize the Shapley value. In the context of network analysis, Battiston, Puliga et al. (2012) propose a measure closely related to eigenvector centrality to assign the systemic relevance of financial institutions based on their centrality in a financial network. Similarly, Soramäki and Cook (2012) try to identify systemically important financial institutions in payment systems by implementing an algorithm based on Markov chains. Employing simulation techniques, they show that the proposed centrality measure, SinkRank, highly correlates with the disruption of the entire system. In accordance with the latter two studies, we also find measures that focus on banks "being central," especially eigenvector centrality, to dominate size as a measure of systemic importance.

As the subprime crisis has shown, banks do not have to be large to contribute to systemic risk, especially where banks are exposed to correlated risks (e.g., credit, liquidity, or funding risk) via portfolios and interbank interconnectedness. Assigning risks to individual banks might be misleading. Some banks might appear healthy when viewed as single entities, but they could threaten financial stability when considered jointly. Gai and Kapadia (2010) find that greater complexity and concentration in the network of bank connections can amplify systemic fragility. Anand et al. (2013) extend their model to include asset fire-sale externalities and macroeconomic feedback on top of network structures, in order to stress-test financial systems. These studies illustrate the tipping point at which the financial system breaks down based on the severity of macroeconomic shocks that affect probabilities of corporate default or asset liquidity. Battiston, Gatti et al. (2012) show that interbank connectivity increases systemic risk, mainly due to a higher contagion risk. Furthermore, Acemoglu, Ozdaglar, and Tahbaz-Salehi (2013) claim that financial network externalities cannot be internalized and thus, in equilibrium, financial networks are inefficient. This creates incentives for regulators to improve welfare by bailing out SIFIs.

In our analysis we keep the total amount of capital in the system constant; otherwise, optimization would be simple but silly: more capital for all, ideally 100 percent equity funding for banks. As a consequence, when we require some banks to hold more capital, we are willing to accept that others may hold less capital, as in the benchmark case. Taken literally, there would be no lower limit to capital except zero. However, we also believe that there should be some minimum capital requirement that applies to all banks for reasons of political feasibility, irrespective of their role in the financial network. Implementing a uniform maximum default probability for all banks, as we actually do in our reallocation mechanism, might be one choice.

Finally, we realize that this is just the first step in calculating optimal capital requirements from network positions to prevent systemic risk. Clearly our results are subject to the standard critique that banks will adjust their network position in response to their new capital requirements. This is not a paper of endogenous network formation, but rather a first step in describing how the system could improve its capital allocation with a given network structure. Given the current German structure, we find that those network measures most influential in reducing total system losses are based on eigenvectors of the adjacency matrix, or closeness measures, and to a lesser extent on the number of lenders a bank has, and a measure that combines this number with the indebtedness of the bank to the rest of the system. These measures will all be described in detail in section 2.2. At its best combination with the benchmark capital requirement, the eigenvector measure can reduce the expected systemic losses by about 15 percent. It works by focusing the capital requirements on a few important banks.

The rest of this paper is structured as follows. In section 2 we describe our risk engine that generates common credit losses to banks' portfolios and our interconnectedness measures. In section 3 we describe our data sources and the network topology of the

German interbank market. Section 4 gives an overview of the contagion algorithm and section 5 describes how capital is optimized. In section 6 we present our main results, and we make some final remarks in section 7.

2. Methodology

Our procedure can be summarized in two stages, along with our initial condition. In the initial state, we use each bank's measured portfolio, which is composed of large and small credit exposures (e.g., loans, credit lines, derivatives) to real-economy and interbank (IB) borrowers. On the liability side, banks hold capital, interbank debt, and deposits. Depositors and other creditors are senior to interbank creditors.

Capital is either set to a benchmark case that is based solely on the loss distributions of their portfolios or according to other capital allocations that partly rely on network measures. Details of how portfolio risk is mixed with network measures are explained in section 5.

In the first stage, we simulate correlated exogenous shocks to all banks' portfolios that take the form of returns on individual large loans (where loans that are shared among multiple lenders are accounted for) and aggregated small loans. Due to changes in value of borrowers' assets, their credit ratings migrate (or they default), and banks make profits or losses on their investments in the real-economy sectors. At the end of this stage, in the case of portfolio losses, capital deteriorates and some banks experience negative capital and default. Thus, we are able to generate correlated losses that affect the capital of each bank simultaneously.

In the second stage, we model interbank contagion. To each simulation round of the first stage we apply an extended version (Rogers and Veraart 2013) of the fictitious contagion algorithm as introduced by Eisenberg and Noe (2001), augmented with bankruptcy costs and a macroeconomic proxy for fire sales. Fundamental bank defaults generate losses to other interbank creditors and trigger some new defaults. Hence, bank defaults can induce domino effects in the interbank market. We refer to new bank failures from this stage as *contagious defaults*.

Finally, we repeat the previous stages for different capital allocations. We discuss the optimization procedure in section 5. Moreover, section 2.2 offers an overview of the interconnectedness measures calculated with the help of network analysis and utilized in the optimization process.

2.1 Credit Risk Model

Our credit risk engine is essential to our study for two reasons. First, it leads to our initial set of bank defaults and helps us determine capital with which our banks face the contagion event in a second stage. Just as important to our model is that the risk engine establishes our benchmark capital allocation as described above. Our results turn out to be sensitive to our choice of benchmark capital, so that it is important to get the credit risk engine close to some realistic risk process. Second, to the best of our knowledge, this is the first paper to incorporate correlated losses and defaults in the first stage for a large banking system.[1] As such, it is important to work with our rich data (with whatever limitations it might have) using a risk model consistent with models that could be used by actual banking risk officers. However, we also have to make a few concessions to the data supplied by the Deutsche Bundesbank. In particular, pricing data of the loans is not available, and our study relies only on the loan portion of the bank portfolios and credit exposures arising from derivatives.

In order to model credit risk, we utilize lending information from two data sources at different levels of aggregation: large loans and small loans. These loans are given to the "real economy." Since borrowers of large loans are explicitly known, along with various parameters such as the loan volume, probability of default, and sector, we can model their credit risk with high precision. When simulating defaults and migrations of individual borrowers, we can even account for the fact that loans given by different banks to the same borrower should migrate or default synchronously.

[1]Elsinger, Lehar, and Summer (2006) and Gauthier, Lehar, and Souissi (2012) do model correlated portfolio losses. However, both the Austrian and the Canadian banking systems consist of much fewer banks than the German one.

We cannot keep this level of precision for small loans because we only know their exposures as a lump sum to each sector. Accordingly, we simulate their credit risk on portfolio level.

2.1.1 Large Loans

In modeling credit portfolio risk, we closely follow the ideas of CreditMetrics (Gupton, Finger, and Bhatia 1997; Bluhm, Overbeck, and Wagner 2003). In the form we use, it is a one-period model, and all parameters are calibrated to a one-year time span. We start with a vector $Y \sim N(0, \Sigma)$ of systematic latent factors. Each component of Y corresponds to the systematic part of credit risk in one of the *risk-modeling* (*RM*) sectors (see section 3.1 for details). The random vector is normalized such that the covariance matrix Σ is actually a correlation matrix. In line with industry practice, we estimate correlations from co-movements of stock indices. For each borrower k in RM sector j, the systematic factor Y_j assigned to the sector is coupled with an independent idiosyncratic factor $Z_{j,k} \sim N(0, 1)$. Thus, the "asset return" of borrower (j, k) can be written as

$$X_{j,k} = \sqrt{\rho} Y_j + \sqrt{1 - \rho} Z_{j,k}. \tag{1}$$

The so-called intrasector asset correlation ρ is common to all sectors.[2] The term "asset return" should not be taken literally; i.e., the link between asset returns and loan losses is not established by the contingent-claims analysis of a structural credit model. Rather, the latent factor $X_{j,k}$ is mapped into rating migrations via a threshold model, and it is a rating migration matrix that the model is calibrated to. If a loan does not default, a loss on it may arise from the fact that the credit spread used in the loan-pricing formula is sensitive to the credit rating.

We use sixteen Standard & Poor's rating classes, including notches AAA, AA+, AA, ..., B–, plus the aggregated "junk" class CCC–C. Moreover, we treat the default state as a further rating (D) and relabel ratings as numbers from 1 (AAA) to 18 (default). Let

[2] This assumption could be relaxed but would require the inclusion of other data sources. In the simulations we use a value of 0.20, which is very close to a value reported by Zeng and Zhang (2001). It is the average over their sub-sample of firms with the lowest number of missing observations.

R_0 denote the initial rating of a borrower and R_1 the rating one year later. A borrower migrates from R_0 to rating state R_1 whenever

$$X \in [\theta(R_0, R_1), \theta(R_0, R_1 - 1)],$$

where θ is a matrix of thresholds associated with migrations between any two ratings. For one-year migration probabilities $p(R_0, R_1)$ from R_0 to R_1, which are given by empirical estimates,[3] the thresholds are chosen such that

$$\mathbf{P}\left(\theta(R_0, R_1) < X_{j,k} \leq \theta(R_0, R_1 - 1)\right) = p(R_0, R_1),$$

which is achieved by formally setting $\theta(R_0, 18) = -\infty$, $\theta(R_0, 0) = +\infty$ and calculating

$$\theta\left(R_0, R_1\right) = \Phi^{-1}\left(\sum_{R>R_1} p\left(R_0, R\right)\right), \quad 1 \leq R_0, R_1 \leq 17.$$

The present value of each non-defaulted loan depends on notional value, rating, loan rate, and time to maturity. In this section we ignore the notional value and focus on D, the *discount factor*. A loan is assumed to pay an annual loan rate C until maturity T, at which all principal is due. We set T equal to a uniform value of 4 years, which is the digit closest to the mean maturity of 3.66 estimated from Bundesbank's borrower statistics (BS).[4] Payments are discounted at a continuous rate $r_f + s(R)$, where r_f is the default-free interest rate and $s(R)$ are rating-specific credit spreads. The term structure of spreads is flat. We ignore the risk related to the default-free interest rate and set $r_f = 2\%$ throughout. The discount factor for a non-defaulted, R-rated loan at time t is

$$D\left(C, R, t, T\right) \equiv \sum_{u=t+1}^{T} \left(C + I_{\{u=T\}}\right) e^{-(r_f + s(R))(u-t)}. \tag{2}$$

[3] We use the 1981–2010 average one-year transition matrix for a global set of corporates (Standard & Poor's 2011).

[4] The borrower statistics report exposures in three maturity buckets. Exposure-weighted averages of maturities indicate only small maturity differences between BS sectors. By setting the maturity to four years, we simplify loan pricing substantially, mainly since the calculation of sub-annual migration probabilities is avoided.

If the loan is not in default at time 1, it is assumed to have just paid a coupon C. The remaining future cash flows are priced according to equation (2), depending on the rating at time $t = 1$, so that the loan is worth $C + D(C, R_1, 1, T)$. If the loan has defaulted at time 1, it is worth $(1+C)(1-LGD)$, where LGD is an independent random variable drawn from a beta distribution with expectation 0.39 and standard deviation 0.34.[5] This means the same relative loss is incurred on loan rates and principal. The spreads are set such that each loan is priced at par at time 0:

$$C\left(R_0\right) \equiv e^{r_f + s(R_0)} - 1, \quad D\left(C\left(R_0\right), R_0, 0, T\right) = 1.$$

Each loan generates a return equal to

$$ret\left(R_0, R_1\right) = -1 + \begin{cases} D\left(C\left(R_0\right), R_1, 1, T\right) + C\left(R_0\right) & \text{if } R_1 < 18 \\ \left(1 + C\left(R_0\right)\right)\left(1 - LGD\right) & \text{if } R_1 = 18. \end{cases}$$

Besides secure interest, the expected value of $ret(R_0, R_1)$ incorporates credit risk premia that markets require in excess of the compensation for expected losses. We assume that the same premia are required by banks and calibrate them to market spreads, followed by minor manipulations to achieve monotonicity in the ratings.[6]

Having specified migrations and the revaluation on a single-loan basis, we return to the portfolio perspective. Assuming that the loan index k in (j, k) runs through all sector-j loans of all banks, we

[5] We have chosen values reported by Davydenko and Franks (2008), who investigate LGDs of loans to German corporates, similar to Grunert and Weber (2009), who find a very similar standard deviation of 0.36 and a somewhat lower mean of 0.275.

[6] Market spreads are derived from a daily time series of Merrill Lynch euro corporate spreads covering all maturities, from April 1999 to June 2011. The codes are ER10, ER20, ER30, ER40, HE10, HE20, and HE30. Spreads should rise monotonically for deteriorating credit. We observe that the premium does rise in general but has some humps and troughs between BB and CCC. We smooth these irregularities out, as they might have substantial impact on bank profitability but lack economic reason. To do so, we fit $\mathbf{E}return(R_0)$ by a parabola, which turns out to be monotonous, and calibrate spreads afterwards to make the expected returns fit the parabola perfectly. Spread adjustments have a magnitude of 7 basis points for A– and better, and 57 basis points for BBB+ and worse. Ultimate credit spreads for ratings without notches are AAA: 0.47 percent; AA: 0.66 percent; A: 1.22 percent; BBB: 2.2476 percent; BB: 4.10 percent; B: 8.35 percent; and CCC–C: 16.40 percent.

denote by $R_1^{j,k}$ the rating of loan (j,k), which is the image of asset return $X_{j,k}$ at time 1. If bank i has given a (large) loan to borrower (j,k), the variable $LL_{i,j,k}$ denotes the notional exposure; otherwise, it is zero. Then, the euro return on the large loans of bank i is

$$ret_{\text{large},i} = \sum_{j,k} LL_{i,j,k} ret\left(R_0^{j,k}, R_1^{j,k}\right).$$

This model accounts for common exposures of banks not only to the same sector but also to individual borrowers. If several banks lend to the same borrower, they are synchronously hit by its default or rating migration.

2.1.2 Small Loans

As previously described, for each bank we have further information on the exposure to loans that fall short of the €1.5 million reporting threshold of the credit register, which is the database for the large loans. However, we know the exposures to small loans only as a sum for each RM sector, so we are forced to model its risk portfolio wise. As each individual loan exposure in a sub-portfolio of small loans is limited by definition, a relationship between portfolio size and the degree of diversification is likely to exist. We account for this relationship by making the idiosyncratic part of portfolio risk dependent on portfolio size.

We sketch the setup only; further details are available from the authors on request. Let us consider all small loans in a bank's portfolio that belong to the same sector j; these are the loans that are too small to be covered by the credit register. They are commonly driven by the sector's systematic factor Y_j and idiosyncratic risk, as in equation (1). If we knew all individual exposures and all initial ratings, we could just run the same risk model as for the large loans. It is central to notice that the returns on individual loans in portfolio j would be independent, *conditional* on Y_j. Hence, if the exposures were extremely granular, the corresponding returns would get very close to a deterministic function of Y_j, as a consequence of the conditional law of large numbers.[7]

[7] This idea is the basis of asymptotic credit risk models. The model behind Basel II is an example of this model class.

We do not go that far, since small portfolios will not be very granular; instead, we utilize the central limit theorem for conditional measures, which allows us to preserve an appropriate level of idiosyncratic risk. Once Y_j is known, the total of losses on an increasing number of loans converges to a (conditionally!) normal random variable. This conditional randomness accounts for the presence of idiosyncratic risk in the portfolio.

Our simulation of losses for small loans has two steps. First, we draw the systematic factor Y_j. Second, we draw a normal random variable, where the mean and variance are functions of Y_j that match the moments of the exact Y_j-conditional distribution. The Y_j-dependency of the moments is crucial to preserve important features of the exact portfolio distribution, especially its skewness. That dependency also assures that two banks suffer correlated losses if both have lent to sector j. An exact fit of moments is not achievable for us, as it would require knowledge about individual exposures and ratings of the small loans, but an approximate fit can be achieved based on the portfolio's Herfindahl-Hirschman Index (HHI) of individual exposures. As the HHI is also unknown, we employ an additional large sample of small loans provided by a German commercial bank to estimate the relationship between portfolio size and HHI. The sector-specific estimate provides us with a *forecast* of the actual HHI, depending on the portfolio's size and the sector. The HHI forecast is the second input (besides Y_j) to the function that gives us Y_j-conditional variances of the (conditionally normal) portfolio losses. A detailed analytical calculation of the conditional moments and a description of the calibration process are available from the authors on request.

This modeling step ends up with a (euro) return on each bank's small loans, denoted by $ret_{\text{small},i}$.

2.2 Centrality Measures

In order to assign the interconnectedness relevance/importance to each bank of the system, we rely on several centrality characteristics. The descriptive statistics of our centrality measures are summarized later in table 2. The information content of an interbank network is best summarized by a matrix X in which each cell x_{ij} corresponds to the liability amount of bank i to bank j. As each positive entry represents an edge in the graph of interbank lending, an edge goes from the borrowing to the lending node. Furthermore, the adjacency matrix (A) is just a mapping of matrix X, in which $a_{ij} = 1$ if $x_{ij} > 0$, and $a_{ij} = 0$ otherwise. In our case, the network is directed, and by X we use the full information regarding an interbank relationship, not only its existence. We do not net bilateral exposures. Our network has a *density* of only 0.7 percent given that it includes 1,764 nodes and 22,752 links.[8] This sparsity is typical for interbank networks (see, for example, Soramäki et al. 2007).

As outlined by Newman (2010), the notion of centrality is associated with several metrics. In economics the most-used measures are *out degree* (the number of links that originate from each node) and *in degree* (the number of links that end at each node), *strength* (the aggregated sum of interbank exposures), *betweenness centrality* (based on the number of shortest paths that pass through a certain node), *eigenvector centrality* (centrality of a node given by the importance of its neighbors), or *clustering coefficient* (how tightly a node is connected to its neighbors).[9]

The out degree, one of the basic indicators, is defined as the total number of direct interbank creditors that a bank borrows from:

$$k_i = \sum_{j}^{N} a_{ij}. \tag{3}$$

In economic terms—for example, in the case of a bank default—the *out degree* defines the number of banks that will suffer losses in the interbank market, assuming equal seniority.

Similarly, we can count the number of banks that a bank lends to (*in degree*). *Degree* is the sum of *out degree* and *in degree.*

We furthermore compute each node's *strength*, that is, its total amount borrowed from other banks:

$$s_i = \sum_{j}^{N} x_{ij}. \tag{4}$$

[8]The *density* of a network is the ratio of the number of existing connections divided by the total number of possible links. In our case of a directed network, the total number of possible links is 1,764 × 1,763 = 3,109,932.

[9]For a detailed description of centrality measures related to interbank markets, see Gabrieli (2011) and Minoiu and Reyes (2013).

In other words, the *strength* of a node is simply a bank's total of interbank liabilities. Similarly, we calculate each bank's interbank assets, which would be labeled the strength of inbound edges in network terminology.

The empirical distribution of degrees shows a tiered interbank structure. A few nodes are connected to many banks. For example, 20 banks (around 1 percent) lend to more than 100 banks each. On the borrowing side, 30 banks have a liability to at least 100 banks. These banks are part of the *core* of the network as defined by Craig and von Peter (2014). In terms of *strength*, 158 banks have a total IB borrowed amount in excess of €1 billion, while only 27 banks have total interbank liabilities in excess of €10 billion. On the assets side, 103 banks lend more than €1 billion and 25 banks have interbank assets in excess of €10 billion.

Opsahl, Agneessens, and Skvoretz (2010) introduce a novel centrality measure that we label *Opsahl centrality*. This measure combines the out degree (equation (3)) with the borrowing strength (total IB liabilities, equation (4)) of each node, using a tuning parameter φ:[10]

$$OC_i = k_i^{(1-\varphi)} \times s_i^{\varphi}.$$

The intuition of Opsahl centrality is that, in the event of default, a node with a high value is able to infect many other banks with high severity. This ability is expected to translate into a higher probability of contagion (conditional on the node's default), compared with other nodes.

Before we define *closeness*, let us define a *path* from node A to B to be a consecutive sequence of edges starting from A and ending in B. Its *length* is the number of edges involved. The (directed) *distance* between A and B is the minimum length of all paths between them. For the definition of *closeness* centrality, we follow Dangalchev (2006):

$$C_i = \sum_{j:j\neq i} 2^{-d_{ij}},$$

[10] In our analysis we set $\varphi = 0.5$, leading to the geometric mean between strength and degree.

where d_{ij} is the distance from i to j, which is set to infinity if there is no path from i to j. This formula has a very nice intuition. If "farness" measures the sum of the distances (in a network sense) of the shortest paths from a node to all of the other nodes, then closeness is the reciprocal of the farness.

Bonacich (1987) proposes an *eigenvalue centrality*, based on the adjacency matrix A. If κ_1 is the largest eigenvalue of A, then eigenvector centrality is given by the corresponding normalized eigenvector v so that $Av = \kappa_1 v$. Eigenvector centralities of all nodes are non-negative.

The *weighted eigenvector centrality* is defined by the eigenvector belonging to the largest eigenvalue of the liabilities matrix X.

As a third version of eigenvectors, we calculate a *weighted normalized* eigenvector based on a modification of X where each row is normalized to sum up to 1 (if it contains a non-zero entry). This normalization ignores the *amount* that a bank borrows from others, once it borrows at all, but accounts for the relative size of IB borrowings.

The *global clustering coefficient,* as in Watts and Strogatz (1998), refers to the property of the overall network, while *local clustering coefficients* refer to individual nodes. This property is related to the mathematical concept of *transitivity.*

$$Cl_i = \frac{\text{number of pairs of neighbors of } i \text{ that are connected}}{\text{number of pairs of neighbors of } i}$$

Here, a neighbor of i is defined as any bank that is connected to it either by lending to or borrowing from it. The local clustering coefficient can be interpreted as the "probability" that a pair of i's neighbors is connected as well. The local clustering coefficient of a node with the degree 0 or 1 is equal to zero. Note that this clustering coefficient refers to undirected graphs where a neighbor to a node is any other node connected to it, and where any link in either direction between two nodes means that they are connected.

The *betweenness centrality* relies on the concept of *geodesics.* A path between two nodes is called geodesic if there is no other path of shorter length. Betweenness centrality simply answers the following question: Of all the geodesics, how many of them go through a given node? More formally, if we let g_{ij} be the number of possible geodesic paths from i to j (there might be more than a single shortest path) and n_{ij}^q be the number of geodesic paths from i to j that pass through node q, then the betweenness centrality of node q is defined as

$$B_q = \sum_{\substack{i,j \\ j \neq i}} \frac{n_{ij}^q}{g_{ij}},$$

where by convention $\frac{n_{ij}^q}{g_{ij}} = 0$ in the case where g_{ij} or n_{ij}^q are zero. The intuition here is that a node of high *betweenness* is more likely to lie on a shortest route that is likely to be taken between two nodes.

We also use *total assets* as a "centrality measure" on which to base the capital allocation. Finally, in addition, we measure the effect of capital allocations based on centrality measures that are summaries of all the other centrality measures: we take the first and second principal component of all of our (normalized) centrality measures.

3. Data Sources

Our model builds on several data sources. In order to construct the interbank network, we rely on the Large-Exposures Database (LED) of the Deutsche Bundesbank. Furthermore, we infer from the LED the portfolios of credit exposures (including loans, bond holdings, and derivatives) to the real economy of each bank domiciled in Germany. This data set is not enough to get the entire picture, since especially the smaller German banks hold plenty of assets falling short of the reporting threshold of €1.5 million for the LED. We therefore use balance sheet data and the so-called borrower statistics.

When calibrating the credit risk model, we rely on stock market indices to construct a sector correlation matrix and utilize a migration matrix for credit ratings from Standard & Poor's. Rating-dependent spreads are taken from the Merrill Lynch corporate spread indices.

3.1 Large-Exposures Database (LED)

The Large-Exposures Database represents the German central credit register.[11] Banks report exposures to a single borrower or a borrower unit (e.g., a banking group) which have a notional exceeding a threshold of €1.5 million. The definition of an exposure includes bonds, loans, or the market value of derivatives and off-balance-sheet items.[12] In this paper, we use the information available at the end of 2011:Q1. The interbank market consists of 1,764 active lenders.[13] Including exposures to the real economy, they have in total around 400,000 credit exposures to more than 163,000 borrower units.

Borrowers in the LED are assigned to 100 fine-grained sectors according to the Bundesbank's customer classification. In order to calibrate our credit risk model, we aggregate these sectors to sectors that are more common in risk management. In our credit risk model, we use EURO STOXX's nineteen industry sectors (and later its corresponding equity indices). Table 1 lists risk-management sectors and the distribution characteristics of the probabilities of default (PDs) assigned to them. These twenty-one sectors represent the *risk-model* (*RM*) *sectors* of our model.

There are two additional sectors (households, including nongovernmental organizations, and the public sector) that are not

[11]The Bundesbank labels this database as *Gross- und Millionenkreditstatistik*. A detailed description of the database is given by Schmieder (2006).

[12]Loan exposures also have to be reported if they are larger than 10 percent of a bank's total regulatory capital. If such an exposure falls short of €1.5 million, it is not contained in our data set of large exposures. Such loans represent a very small amount compared with the exposures that have to be reported when exceeding €1.5 million; they are captured in the borrower statistics, though, and hence part of the "small loans"; see section 3.2.

It is also important to notice that, while the data are quarterly, the loan volume trigger is not strictly related to an effective date. Rather, a loan enters the database once its actual volume has met the criterion at some time throughout the quarter. Furthermore, the *definition of credit* triggering the obligation to report large loans is broad: besides *on-balance-sheet* loans, the database conveys bond holdings as well as *off-balance-sheet* debt that may arise from open trading positions, for instance. We use *total exposure* of one entity to another. Master data of borrowers contains its nationality as well as assignments to *borrower units*, when applicable, which is a proxy for the joint liability of borrowers. We have no information regarding collateral in this data set.

[13]Each lender is considered at an aggregated level (i.e., as "Konzern"). At the single-entity level, there are more than 4,000 different lending entities that report data.

Table 1. Risk-Model (RM) Sectors

No.	Risk Model Sector	No. of Borrowers	Volume Weight	Default Probabilities (Percent)					
				5%	25%	50%	75%	95%	Mean
1	Chemicals	3,200	0.9%	0.008	0.19	0.59	1.54	6.66	1.66
2	Basic Materials	14,419	1.5%	0	0.27	0.85	2.05	10.00	2.20
3	Construction and Materials	17,776	1.3%	0	0.12	0.66	1.85	7.95	1.99
4	Industrial Goods and Services	73,548	15.1%	0	0.23	0.77	2.07	13.00	2.57
5	Automobiles and Parts	1,721	0.7%	0.001	0.31	1.05	3.00	14.98	2.91
6	Food and Beverage	13,682	0.8%	0.001	0.27	0.82	1.85	8.00	1.92
7	Personal and Household Goods	21,256	1.3%	0	0.17	0.74	1.99	14.90	2.75
8	Health Care	16,460	1.0%	0	0.03	0.12	0.86	3.84	0.98
9	Retail	25,052	1.6%	0	0.17	0.79	2.67	11.40	2.37
10	Media	2,534	0.2%	0	0.17	0.45	1.71	7.90	1.81
11	Travel and Leisure	8,660	0.7%	0	0.36	1.17	2.92	20.00	3.16
12	Telecommunications	299	0.8%	0	0.12	0.36	2.32	6.32	1.82
13	Utilities	15,679	3.2%	0	0.09	0.39	1.26	6.77	1.62
14	Insurance	1,392	4.1%	0.029	0.03	0.09	0.66	4.82	1.15
15	Financial Services	23,634	22.5%	0.021	0.03	0.05	0.57	4.82	1.07
16	Technology	2,249	0.2%	0	0.20	0.50	1.60	4.68	1.31
17	Foreign Banks	3,134	22.1%	0.003	0.03	0.09	0.88	7.95	1.40
18	Real Estate	56,451	11.4%	0	0.10	0.51	1.68	8.31	1.82
19	Oil and Gas	320	0.5%	0.038	0.22	0.81	2.96	12.86	3.03
20	Households (incl. NGOs)	79,913	1.3%	0	0.05	0.35	1.24	6.00	1.34
21	Public Sector	1,948	9.2%	0	0	0	0	0	0
	Total	388,327	100.0%						1.5

Notes: Volume weight refers to credit exposure.

linked to equity indices. We consider exposures to the public sector to be risk free (and hence exclude them from our risk engine) since the federal government ultimately guarantees all public bodies in Germany. Households include credit risk according to the PDs in the data set. Similar to other sectors, their intrasector correlation ρ equals 20 percent such that all banks' losses in the households sector are correlated through the same systematic factor. However, this factor is uncorrelated to the ones of other sectors.

Information regarding borrowers' default probabilities is included as well in the LED. We report several quantiles and the mean of sector-specific PD distributions in table 1. Since only internal-ratings-based (IRB) banks and savings and loan (S&L) banks report this kind of information, borrowers without a reported PD are assigned random PDs drawn from a sector-specific empirical distribution.

3.2 Borrower and Balance Sheet Statistics

While the LED is a unique database, the threshold of €1.5 million of notional is still a substantial restriction. Although large loans build the majority of money lent by German banks, the portfolios of most German banks would not be well represented by them. That does not come as a surprise, as the German banking system is dominated (in numbers) by rather small S&L and cooperative banks. Many banks hold few loans large enough to enter the LED, though these banks are, of course, much better diversified. For two-thirds of banks, the LED covers less than 54 percent of the total exposures. We need to augment the LED with information on smaller loans.

Bundesbank's borrower statistics (BS) data set reports lending to German borrowers by each bank on a quarterly basis. Focusing on the calculation of money supply, it reports only those loans made by banks and branches situated in Germany; e.g., a loan originated in the London office of a German bank would not enter the BS, even if the borrower is German. Corporate lending is structured in eight main industries, of which two are further split up.[14]

[14] The main sectors are agriculture, basic resources and utilities, manufacturing, construction, wholesale and retail trade, transportation, financial intermediation and insurance, and services.

Loans to households and non-profit organizations are also reported.[15]

While lending is disaggregated into various sectors, the level of aggregation is higher than in the LED, and sectors are different from the sectors in the risk model. We treat this mismatch by a linear mapping of exposures from BS to RM sectors. Detailed information on the mapping is available on request.

In addition to borrower statistics, we use figures from the monthly balance sheet statistics, which is also run by the Bundesbank. These sheets contain lending to domestic insurances, households, non-profit organizations, social security funds, and so-called other financial services companies. Lending to foreign entities is measured by a total figure that covers all lending to non-bank companies and households. The same applies to domestic and foreign bond holdings which, if large enough, are also included in the LED.

3.3 Market Data

Market credit spreads are derived from a daily time series of Merrill Lynch option-adjusted euro spreads covering all maturities, from April 1999 to June 2011. The codes are ER10, ER20, ER30, ER40, HE10, HE20, and HE30.

Asset correlations used in the credit portfolio model are computed from EURO STOXX weekly returns of the European sector indices for the period April 2006–March 2011, covering most of the financial crisis. The European focus of the time series is a compromise between a sufficiently large number of index constituents and the actual exposure of the banks in our sample, which is concentrated on German borrowers but also partly European wide.

The credit ratings migration matrix is provided by Standard & Poor's (2011).

[15] A financial institution has to submit BS forms if it is a monetary financial institution (MFI), which does not necessarily coincide with being obliged to report to the LED. There is one state-owned bank with substantial lending that is exempt from reporting BS data by German law. As it is backed by a government guarantee, we consider this bank neutral to interbank contagion.

3.4 The German Interbank Market

In this section we explore the German market for interbank lending in more detail. In the existing literature, Craig and von Peter (2014) use the German credit register to analyze the German interbank market. They find that a *core-periphery* model can be well fitted to the German interbank system: core banks build a complete sub-network (i.e., there exist direct links between any two members of the subset), while periphery banks are less connected by lending. The core-periphery structure turns out to be very stable through time. Roukny, Georg, and Battiston (2014) use the same data source, spanning the period 2002–12. Providing a thorough analysis of how German interbank lending develops over time, they find most of the characteristics to be very stable, including the distributions of various centrality measures utilized in our paper. Because of these findings, complemented by our own analysis for the period 2005–11, we neglect the time dimension completely and focus on a single point of time.[16]

At the end of 2011:Q1, 1,921 MFIs were registered in Germany, holding a total balance sheet of €8,233 billion.[17] The German banking system is composed of three major types of MFIs: 282 commercial banks (including 4 big banks and 110 branches[18] of foreign banks) that hold approximately 36 percent of total assets, 439 saving banks (including 10 *Landesbanken)* that hold roughly 30 percent of the system's assets, and 1,140 credit cooperatives (including 2 regional institutions) that hold around 12 percent of market share. Other banks (i.e., mortgage banks, building and loan associations, and special-purpose vehicles) are in total sixty MFIs and represent approximately 21 percent of the system's balance sheet.

Our interbank (IB) network consists of 1,764 active banks (i.e., aggregated banking groups). These banks are actively lending and/or borrowing in the interbank market. They hold total assets worth €7,791 billion, from which 77 percent represent large loans and 23 percent small loans.

[16]Tables with these measures are available on request.

[17]Source: Deutsche Bundesbank's Monthly Report, March 2011.

[18]If a foreign bank runs a branch in Germany, the branch has to report large loans to the LED. Subsidiaries of foreign banks are treated as German banks.

Table 2 presents the descriptive statistics of the main characteristics and network measures of the German banks utilized in our analysis. The average size of a bank-individual IB exposure is around €1 billion. As figures show, there are few very large total IB exposures, since the mean is between the 90th and 95th percentile, making the distribution highly skewed. Similar properties are observed for total assets, the total of large loans and out degrees, supporting the idea of a tiered system with few large banks that act as interbank broker-dealers connecting other financial institutions (see, e.g., Craig and von Peter 2014).[19]

4. Modeling Contagion

As introduced in section 2, we differentiate between *fundamental defaults* and *contagious defaults* (see also Elsinger, Lehar, and Summer 2006 or Cont, Moussa, and Santos 2013, for instance). Fundamental defaults are related to losses from the credit risk of "real-economy" exposures, while contagious defaults are related to the interbank credit portfolio (German only).[20]

Moreover, we construct an *interbank clearing mechanism* based on the standard assumptions of interbank contagion (see, e.g., Upper 2011): First, banks have *limited liability*. Next, interbank liabilities are senior to equity but *junior* to non-bank liabilities (e.g., deposits).[21] Losses related to bank defaults are *proportionally* shared

[19] One aspect that needs to be mentioned here is that the observed IB network is not the complete picture, since interbank liabilities of German banks raised outside Germany are not reported to the LED. For example, the LED captures a loan made by Goldman Sachs to Deutsche Bank only if it is made by Goldman Sachs' German subsidiary. This aspect might bias downwards centrality measures of big German banks that might borrow outside Germany.

[20] Loans to foreign banks are included in sector 17 of the "real-economy" portfolio. Their losses are correlated through a common systematic factor, as for loans to the real economy, and their PDs are taken from the LED. However, if a foreign bank has a branch registered in Germany (or multiple of them) and if that branch appears as a lender in the LED, a loan to that foreign bank (or the branch) does not belong to sector 17 but is part of our interbank network. This is so because both the foreign parent bank and the German branch are matched to the same borrower unit, as explained in section 3.1. Subsidiaries of foreign banks simply appear as German banks.

[21] As Upper (2011) points out, deposits by non-banks are senior to some interbank liabilities in the German banking system, but not to all of them. Our

Table 2. Interbank (IB) Market and Network Properties

	Quantiles								
	5%	**10%**	**25%**	**50%**	**75%**	**90%**	**95%**	**Mean**	**Std. Dev.**
Total IB Assets[a]	7,591	13,089	35,553	100,450	310,624	868,299	1,647,666	990,433	7,906,565
Total IB Liabilities[a]	2,640	6,053	19,679	61,450	180,771	527,460	1,212,811	990,433	7,782,309
Total Assets[a]	37,798	63,613	160,741	450,200	1,290,698	3,412,953	7,211,608	4,416,920	37,938,474
Total Large Loans[a]	8,719	17,293	60,692	208,475	675,199	2,054,840	4,208,323	3,424,445	30,221,363
Total Small Loans[a]	8,906	34,932	85,741	217,516	550,855	1,287,390	2,253,149	992,475	8,748,757
Out Degree	1	1	1	2	4	9.1	16	13	82
In Degree	1	1	4	9	14	19	25	13	37
Total Degree	2	3	5	11	18	27	38	26	111
Opsahl Centrality	51.5	80.6	165.8	345.9	791.6	2,071.2	4,090	3,342	24,020
Eigenvector Centrality	0.000003	0.000019	0.000078	0.000264	0.001092	0.004027	0.009742	0.003923	0.023491
Weighted Betweenness	0	0	0	0	0	401.5	6,225.8	10,491	82,995
Weighted Eigenvector	0.000004	0.000012	0.000042	0.000142	0.000611	0.002217	0.004839	0.003148	0.023607
Closeness Centrality	253.2	328.8	347.8	391.8	393.6	411.5	427.4	371.22	69.13
Clustering Coefficient	0	0	0	0.00937	0.04166	0.12328	0.16667	0.0379	0.0677
Number of banks[b]	1,764								
Number of links	22,752								

Notes: Data point: 2011:Q1.
[a]In thousands euro.
[b]Number of banks active in the interbank market.

among interbank creditors, based on the share of their exposure to total interbank liabilities of the defaulted bank. In other words, its interbank creditors suffer the same loss given default.[22] Finally, non-bank assets of a defaulted bank are liquidated at a certain discount. This extra loss is referred to as *fire sales* and is captured by *bankruptcy costs,* defined in section 4.1. The clearing mechanism follows the idea of Eisenberg and Noe (2001) which, however, is not designed to include bankruptcy costs. To account for these costs, we follow Rogers and Veraart (2013), who propose a simple algorithm which converges to the fixed point with minimum losses.

When a group of banks default, they trigger losses in the interbank market. If interbank losses (plus losses on loans to the real economy) exceed the remaining capital of the banks that lent to the defaulted group, this can develop into a domino cascade. At every simulation when interbank contagion arises, we follow Rogers and Veraart (2013) to compute losses that take into account the above assumptions.

More formally, first recall that each bank makes a (euro) return on its large and small loans, defined in section 2.1. We switch the sign and define *fundamental losses* as

$$L_i^{\text{fund}} \equiv -\left(ret_{\text{large},i} + ret_{\text{small},i}\right).$$

Each bank incurs total portfolio losses L_i equal to its fundamental losses and losses on its interbank loans:

$$L_i = L_i^{\text{fund}} + L_i^{\text{IB}}, \tag{5}$$

where L_i^{IB} is yet to be determined. A bank defaults if its capital K_i cannot absorb the portfolio losses. We define the default indicator as

$$D_i = \begin{cases} 1 & \text{if} K_i < L_i, \\ 0 & \text{otherwise.} \end{cases}$$

attempts to find out the exact relationship failed. Thus, we consider the most pessimistic approach from the interbank creditors' point of view. If this assumption were relaxed and, for instance, replaced by proportional loss sharing among *all* creditors of a bank, the effects of interbank contagion would probably weaken.

[22]We do not have any precise information related to collateral or the seniority of claims among interbank loans.

Below, bankruptcy costs are modeled so that their (potential) extent BC_i is known before contagion; i.e., they are just a parameter of the contagion mechanism, but only accrue when $D_i = 1$.

Total portfolio losses and bankruptcy costs are now distributed to the bank's claimants. If capital is exhausted, further losses are primarily borne by interbank creditors, since their claim is junior to other debt. Bank i causes its interbank creditors an aggregate loss of

$$\Lambda_i^{\mathrm{IB}} = \min\left(l_i, \max\left(0, L_i + BC_i D_i - K_i\right)\right),$$

which is zero if the bank does not default. The Greek letter signals that Λ_i^{IB} is a loss *on the liability side* of bank i, which *causes* a loss on the assets of its creditors.

The term x_{ij} denotes interbank liabilities of bank i against bank j, and the row sum $l_i = \sum_{j=i}^{N} x_{ij}$ defines total interbank liabilities of bank i. This gives us a *proportionality matrix* π to allocate losses, given by

$$\pi_{ij} = \begin{cases} \frac{x_{ij}}{l_i} & \text{if } l_i > 0; \\ 0 & \text{otherwise.} \end{cases}$$

If the loss amount Λ_i^{IB} is proportionally shared among the creditors, bank j incurs a loss of $\pi_{ij}\Lambda_i^{\mathrm{IB}}$ because of the default of i. Also, the bank i we have started with may have incurred interbank losses; they amount to

$$L_i^{\mathrm{IB}} = \sum_{k=1}^{N} \pi_{ji}\Lambda_j^{\mathrm{IB}}, \tag{6}$$

which provides the missing definition in equation (5). This completes the equation system (5)–(6), which defines our allocation of contagion losses.

Let us write the system in vector form, for which we introduce $l = (\sum_{j=i}^{N} x_{ij})_i$ as the vector of interbank liabilities, $\wedge$ as the element-wise minimum operator, and $\tilde{B}$ as a diagonal matrix with bankruptcy costs BC_i on the diagonal. Default dummies are subsumed in the vector D so that the product $\tilde{B}D$ defines *actual* bankruptcy costs, i.e., the ones that become real. Losses in the IB market *caused* by bank i and incurred by the other banks can then be written as

$$\Lambda^{IB} = l \wedge \left[L + \tilde{B}D - K\right]^{+}.$$

With Π being the proportionality matrix, interbank losses—now on the asset side—amount to

$$L^{\mathrm{IB}} = \Pi^{\top} \Lambda^{\mathrm{IB}}.$$

We consider the total portfolio losses $L = L^{\mathrm{fund}} + L^{\mathrm{IB}}$ (not containing bankruptcy costs) as a solution.[23] Altogether, we have to solve the equation

$$L = \Phi\left(L\right) \equiv L^{\mathrm{fund}} + \Pi^{\top}\left(l \wedge \left[L + \tilde{B}I_{[L>K]} - K\right]^{+}\right),$$

where the indicator function and the relational operator are defined element by element. According to our definition of default, the operator Φ is left-continuous on $\mathbb{R}^N$.

Rogers and Veraart (2013) show that a simple repeated application of the monotonic operator Φ to L^{fund} generates a sequence of losses that must converge to a unique loss vector L^{∞}, simply because the sequence is monotonic in each dimension and the solution space is compact. Since the operator Φ is left-continuous in our setup, L^{∞} is also a fixed point. As shown by Rogers and Veraart (2013), it has minimum losses among all fixed points.

4.1 Bankruptcy Costs

In our analysis, we are particularly interested in bankruptcy costs (henceforth BCs), since they represent a deadweight loss to the economy. We model them as the sum of two parts. The first one is a function of a bank's total assets, because there is empirical evidence for a positive relationship between size and BCs of financial institutions; see Altman (1984). The second part incorporates fire sales and their effect on the value of the defaulted bank's assets. For their

[23] We could also search for a solution for $L^{all} \equiv L + \tilde{B}D$, which turns out to be equivalent but more complicated.

definition, recall that we want to model BCs such that their extent is known before contagion, which is why we make them exclusively dependent on the *fundamental* portfolio losses L_i^{fund}. If that loss of bank i exceeds capital K_i, the bank's creditors suffer a loss equal to $\max(0, L_i^{\text{fund}} - K_i)$. In the whole economy, fundamental losses add up to

$$\overline{L}_{\text{fund}} \equiv \sum_i \max\left(0, L_i^{\text{fund}} - K_i\right).$$

It is this total fundamental loss in the system by which we want to proxy lump-sum effects of fire sales. The larger $\overline{L}_{\text{fund}}$, the more assets will the creditors of defaulted banks try to sell quickly, which puts asset prices under pressure. We proxy this effect by defining a system-wide relative loss ratio λ that is a monotonic function of $\overline{L}_{\text{fund}}$. In total, if bank i defaults, we define BCs as the sum of two parts related to total assets and fire sales:

$$BC_i \equiv \phi\left(TotalAssets_i - L_i^{\text{fund}}\right) + \lambda\left(\overline{L}_{\text{fund}}\right)\max\left(0, L_i^{\text{fund}}\right). \quad (7)$$

We consider ϕ to be the proportion of assets lost due to litigation and other legal costs. In our analysis we set $\phi = 5\%$.[24] It is for convenience rather than for economic reasons that we set the monotonic function λ equal to the cumulative distribution function of $\overline{L}_{\text{fund}}$. Given this choice, the larger total fundamental losses in the system are, the closer λ gets to 1.[25]

In the optimization process, we minimize a measure of system losses (i.e., the target function). The mechanism of contagion proposed by us has several sets of agents, each of whom suffers separate kinds of losses. There are many conflicting arguments regarding

[24]Our results remain robust also for other values $\phi \in \{1\%, 3\%, 10\%\}$. Alessandri et al. (2009) and Webber and Willison (2011) use contagious BCs as a function of total assets, and set ϕ to 10 percent. Given the second term of our BC function that incorporates fire-sales effects, we reach a stochastic function with values between 5 percent and 15 percent of total assets.

[25]We acknowledge that real-world BCs would probably be sensitive to the amount of interbank credit losses, which we ignore. This simplification, however, allows us to calculate potential BCs before we know which bank will default through contagion, such that we do not have to update BCs in the contagion algorithm. If we did, it would be difficult to preserve proportional loss sharing in the Eisenberg-Noe allocation.

which agent's losses the regulator should particularly be interested in—for instance, those of depositors (as a proxy for "the public" that is likely to be the party that ultimately bails banks out) or even those of bank equityholders (who are at risk while offering a valuable service to the real economy). While all of them may be relevant, our primary target function is the *expected BCs*, which is just the sum of the expected BCs of defaulted banks:

$$EBC = \mathbf{E}\sum_i BC_i D_i, \tag{8}$$

where D_i is the default indicator of bank i. BCs are a total social deadweight loss that does not include the initial portfolio loss due to the initial shock. While this is a compelling measure of social loss, there are distributional reasons that it might not be the only measure of interest.

Losses of equation (8) are free of distributional assumptions. However, some amount of BCs can be acceptable as a side effect of otherwise desirable phenomena (e.g., the plain existence of bank business), such that minimizing expected BCs does not necessarily lead to a "better" system in a broader sense. For example, minimum expected BCs might entail an undersupply of credit more harmful to the real economy than the benefits from low BCs.

The *expected total loss to equityholders* and the *expected total loss to non-bank debtholders* are therefore at least of interest, if not justifiable alternative target functions. Also, their sum appears as a natural choice for a target function, this way treating the interests of bank owners and non-bank debtholders as equally important. However, it can easily be shown that this sum is equivalent to our target function, the expected BCs.

It is important to note that we consider banks as institutions only, meaning that any loss hitting a bank must ultimately hit one of its *non-bank* claimants. In our model, these are simply non-bank debtors and bank equityholders, as we split bank debt into non-bank and interbank debt only. Counting total losses on the asset side of banks' balance sheets (or just their credit losses) makes little sense anyway, as they involve interbank losses and therefore a danger of double-counting.

We considered all of the above losses in our simulations, but we report only those related to equation (8) with the purpose of exposition in this paper.

5. Optimization

In this section, we define the way we reallocate capital. Several points should be made about it.

First, the rules themselves are subject to a variety of restrictions. These include the fact that the rules must be simple and easily computed from observable characteristics, and they should preferably be smooth to avoid cliff effects. Simplicity is important not just because of computational concerns. Simple formal rules are necessary to limit discretion on the ultimate outcome. Too many model and estimation parameters set strong incentives for banks to lobby for a design in their particular interest. While this is not special to potential systemic risk charges, it is clear that those banks that will most likely be confronted with increased capital requirements are the ones with the most influence on politics. Vice versa, simplicity can also help to avoid arbitrary punitive restrictions imposed upon individual banks. In this sense, the paper cannot offer deliberately fancy first-best solutions for capital requirements.

Second, as noted in the Introduction, the rules must keep the total capital requirement the same so that we do not mix the effects of capital reallocation with the effect of increasing the amount of capital in the entire system.

Finally, for reasons of exposition, we focus on capital reallocations based on a single centrality measure. While we did explore more complicated reallocation rules that were optimized over combinations of centrality measures, these reallocations gave only marginal improvements and so are not reported.

We introduce a range of simple capital rules over which our chosen loss measure, the expected BCs, is minimized. We minimize over two dimensions, the first being the choice of a centrality measure and the second being represented by a gradual deviation from a VaR-based benchmark allocation towards an allocation based on the chosen centrality measure.

In our benchmark case, capital requirements focus on a bank's individual portfolio risk (and not on network structure). For our analysis we require banks to hold capital equal to its portfolio VaR on a high security level of $\alpha = 99.9\%$, which is in line with the level used in Basel II rules for the banking book.[26] There is one specialty of this VaR, however. In line with Basel II again, the benchmark capital requirement treats interbank loans just as other loans. For the determination of bank i's benchmark capital $K_{\alpha,i}$ (and *only* for this exercise), each bank's German interbank loans (on the asset side of the balance sheet) are merged with loans to foreign banks into portfolio sector 17, where they contribute to losses just as other loans.[27]

In the whole system, total required capital adds up to $TK_\alpha \equiv \sum_i K_{\alpha,i}$. To establish a "level playing field" for the capital allocation rules, TK_a is to be the same for the various allocations tested.

The basic idea is that we give banks a—hypothetical—proportional capital relief from their benchmark capital, which we then redistribute according to a rule in which capital is scaled up or down by a centrality measure.

Given C_i to be one of the centrality measures introduced in section 2.2, we subtract a fraction β from each bank's benchmark capital $K_{\alpha,i}$ for redistribution. Some required capital is added back, again as a fraction of $K_{\alpha,i}$, which is proportional to the centrality measure:

$$K_{\text{simple},i}\left(\beta\right) \equiv K_{\alpha,i}\left(1-\beta+\beta a C_i\right). \tag{9}$$

The parameter a is chosen such that total system capital remains the same as in the benchmark case, which immediately leads to

$$a = \frac{\sum_j K_{\alpha,j}}{\sum_j K_{\alpha,j} C_j}, \quad \text{for all } \beta.$$

[26]To check the numerical stability of our results, we reran several times the computation of VaR measures at quantile $\alpha = 99.9\%$. For this computation we employed a new set of one million simulations and kept the same PDs for loans where unreported values had been reported by random choices. Results are very similar, with an average variance of under 2 percent. VaR measures at 99 percent have a variance of under 0.5 percent.

[27]Default probabilities for these loans are taken from the Large-Exposures Database in the same way as for loans to the real economy.

This simple capital rule is not yet our final allocation because it has a flaw, although it is almost exactly the one we use in the end. Among those banks for which capital decreases in β (these are the banks with the lowest C_i), some banks' capital buffer may get so small that its probability of default (PD) rises to an unacceptable level. We want to limit the PDs (in an imperfect way) to 1 percent, which we assume to be politically acceptable. To this end, we require each bank to hold at least capital equal to bank i's VaR at $\alpha = 99\%$, which we denote by $K_{\min,i}$. In other words, we set a floor on $K_{\text{simple},i}(\beta)$. The VaR is again obtained from the model used for benchmark capital, which treats IB loans as ordinary loans and hence makes the upper PD limit imperfect.[28]

If we plainly applied the floor to $K_{\text{simple},i}(\beta)$ (and if the floor was binding somewhere), we would require more capital to be held in the system than in the benchmark case. We therefore introduce a tuning factor $\tau(\beta)$ to reestablish TK_α. The final capital rule is

$$K_{\text{centr},i}\left(\beta\right) \equiv \max\left(K_{\min,i}, K_{\alpha,i}\left[1-\beta+\beta\,\tau\left(\beta\right)\,a\,C_i\right]\right). \qquad (10)$$

For a given β, the tuning factor $\tau(\beta)$ is set numerically by root finding. To anticipate our results, tuning is virtually obsolete. Even if 30 percent of benchmark capital is redistributed—which is by far more than in the optimum, as will turn out—we find $0.999 < \tau(\beta) \leq 1$ throughout, and there is a maximum of only fourteen banks for which the floor becomes binding.

In general, the approach is not limited to a single centrality measure. As regards the formal approach, we could easily replace C_i in equation (10) with a linear combination of centrality measures $\sum_k a_k C_i^k$ and optimize the a_k along with β; the degree of freedom would be equal to the number of centrality measures included. However, optimization quickly becomes prohibitively expensive, as every step requires its own extensive simulation. We carry out a couple of

[28]If banks held exactly $K_{\min,i}$ as capital, actual bank PDs after contagion could be below or above 1 percent, depending on whether quantiles of portfolio losses in a risk model where interbank loans are directly driven by systematic factors are larger than in the presence of contagion (but without direct impact of systematic factors). However, the probability of bank defaults through *fundamental* losses cannot exceed 1 percent, given that asset correlations in the factor model are positive.

bivariate optimizations, focusing on the centrality measures found most powerful in one-dimensional optimizations. Improvements over the one-dimensional case are negligible and therefore not reported.

In this section we have dealt only with capital *requirements*. To assess their consequences on actual systemic risk, we also have to specify in which way banks intend to obey these requirements. In practice, banks hold a buffer on top of required capital in order to avoid regulatory distress. However, as our model runs over one period only, we do not lose much generality when we abstract from additional buffers and assume banks to hold exactly the amount of capital they are required to hold according to the capital rule in force; data on banks' actual capital holdings are not directly involved.

Depending on the capital regime supposed, varying amounts of capital must be reconciled with the other parts of the balance sheet. Our central assumption is that all positions given by data remain the same throughout this paper. There is especially no reaction of the interbank lending network to capital; endogenous interbank lending or borrowing is beyond the scope of our paper.

For implementation, assume a bank holds (interbank and non-bank) loans of size $Loans$ on the asset side, while it has interbank debt liabilities $Liab^{IB}$. These positions are given by data and kept fixed. In addition, the bank holds a varying, rules-dependent amount of capital K. If the bank does not excessively fund itself by interbank borrowing, $Loans$ will exceed $Liab^{IB} + K$. In that case, the liabilities gap $Loans - Liab^{IB} - K$ is assumed to be funded by non-bank debt. If, in contrast, the gap is on the asset side, the bank has no additional liabilities and the gap is closed by an investment in riskless assets such as cash. In other words, funding by non-bank creditors or cash are the only adaptive positions; the balance sheet is complete if we assume

$$\begin{aligned} Liab^{NB} &= \max\left(0, Loans - Liab^{IB} - K\right) \\ Cash &= \max\left(0, Liab^{IB} + K - Loans\right). \end{aligned}$$

If K moves up or down, there is no effect on lending to the real economy (and the corresponding credit losses) or interbank lending.

Figure 1. Comparison of Centrality-Based Capital Allocations

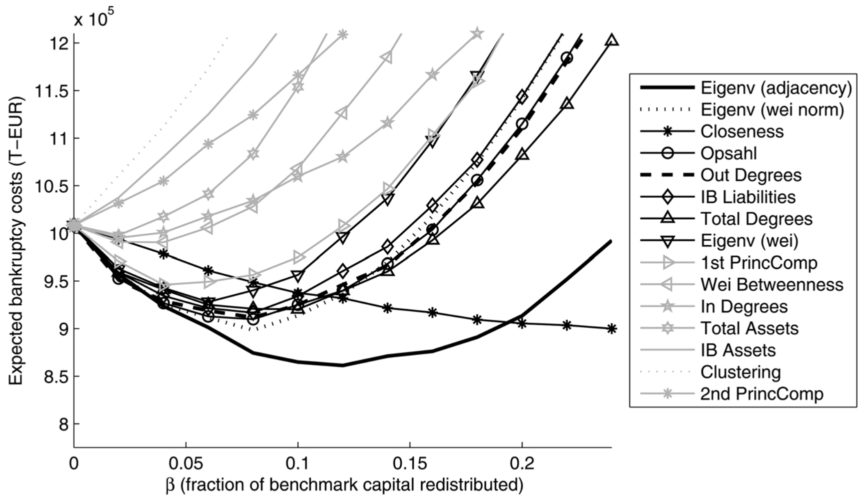

Note also that there is no dynamic aspect of these adjustments in our analysis, as we simply compare alternative static choices.

6. Results

Our main results from the reallocation of capital are depicted in figure 1. Table 3 presents the corresponding optimum values in numbers. As defined in (8), the target function that we compare across centrality measures is the total expected bankruptcy costs. Our benchmark allocation (based only on VaR($\alpha = 99.9\%$)) is represented by the point where $\beta = 0$. As β increases, more of the capital is allocated to an allocation rule based upon the centrality measure as described in equation (10). The extreme right end of the diagram still allocates 75 percent of the capital based on the original VaR($\alpha = 99.9\%$). To allocate less would probably be difficult to support politically and, as can be seen in the figure, covers the area for the minimum losses in all cases but one. Even in this case, the minimum for the closeness centrality is very close to the minimum in the plotted range.

Table 3. Total Expected Bankruptcy Costs after Contagion, Applying Optimized Centrality-Based Capital Allocation

Centrality Measure	Expected BCs (T-EUR)	Saving (Percent)	Optimal Beta (Percent)
Eigenvector (Adjacency)	861,000	14.6	12
Eigenvector (Weighted Normalized)	898,000	10.9	8
Closeness	900,000	10.7	24
Opsahl Centrality	910,000	9.8	8
Out Degree	912,000	9.6	8
IB Liabilities	917,000	9.1	8
Degree	920,000	8.7	10
Eigenvector (Weighted)	928,000	8.0	6
First Principal Component	946,000	6.2	4
Weighted Betweenness	991,000	1.8	4
In Degree	995,000	1.3	2
Total Assets	998,000	1.0	2
IB Assets	1,008,000	0.0	0
Clustering	1,008,000	0.0	0
Second Principal Component	1,008,000	0.0	0
(Benchmark, Purely VaR Based)	1,008,000	—	—

Notes: Expected BCs include fundamental and contagious defaults. Saving in percent refers to the benchmark case. β has been optimized in bins of 2 percent size.

Several patterns are clearly evident in the diagram and supported by findings in table 3. The first observation in figure 1 is that some centrality measures are completely dominated by others. Over the region of $0 \leq \beta \leq 24\%$, capital reallocations based on three of the centrality measures were destructive to the system over the entire region: clustering, IB assets, and the second principal component.

Capital allocations based on all other centrality measures help to improve the stability of the banking system in terms of expected total bankruptcy costs over at least some of the region of β.

One measure stands out among all of them: the adjacency eigenvector. A reallocation of about 12 percent of baseline capital to one based on the adjacency eigenvector saves the system 14.6 percent in system losses as measured by expected BCs. Radically giving 24 percent of the allocation to one based on closeness saves the system less (10.7 percent), but these two capital allocations behave differently from the rest in two senses. First, they are more effective at reducing losses, meaning that the next most effective group of centrality measures, Opsahl, weighted eigenvector, out degree, total degree, and weighted normalized eigenvector, stabilize the system so that it reduces losses between 9 percent and 10 percent, in contrast to 14.6 percent. Second, they all do so at a level of capital reallocation that involves only 8 percent being reallocated from the benchmark to the centrality-based allocation. The two most successful reallocations redistribute more. Finally, the rest of the centrality measures generate even smaller savings at even smaller amounts redistributed. Because the allocations based on closeness and the adjacency eigenvector perform best, we focus on them in the discussion that follows (although the weighted normalized eigenvector actually performs second best, we skip it from further analyses because of its similarity to the adjacency eigenvector). To a lesser extent, we also discuss the effect of the next-best-performing allocations, Opsahl centrality and out degree.

Figure 2 presents expected BCs under the best-performing centrality measures for both pre- and post-contagion losses. The case of closeness stands out from the other three. For closeness, the fundamental defaults mimic the baseline case throughout the range of capital reallocation. They barely increase despite changes in the distribution of bank defaults. As β increases, the decline in losses from contagion is not matched by an increase in pre-contagion losses, and the total losses decrease. The best-performing allocation, based on the adjacency eigenvector, for the smaller range of β, does not increase the fundamental losses much over the benchmark and gives striking savings in post-contagion losses compared with all other measures. It does exactly what one would expect from a capital reallocation based on a wider set of information: strongly reduce the cascading losses at a small expense of fundamental ones. Once the fundamental losses start increasing steeply, at an inflection point of about $\beta = 15\%$, the total costs rise rapidly as well. Below this value, the increase in fundamental costs is relatively smooth, just as the fall-off in benefits from the post-contagion savings are smooth. This reflects the large range of β where the capital reallocation is very

Figure 2. Expected Bankruptcy Costs: All Defaults, Fundamental Defaults, and Contagious Defaults

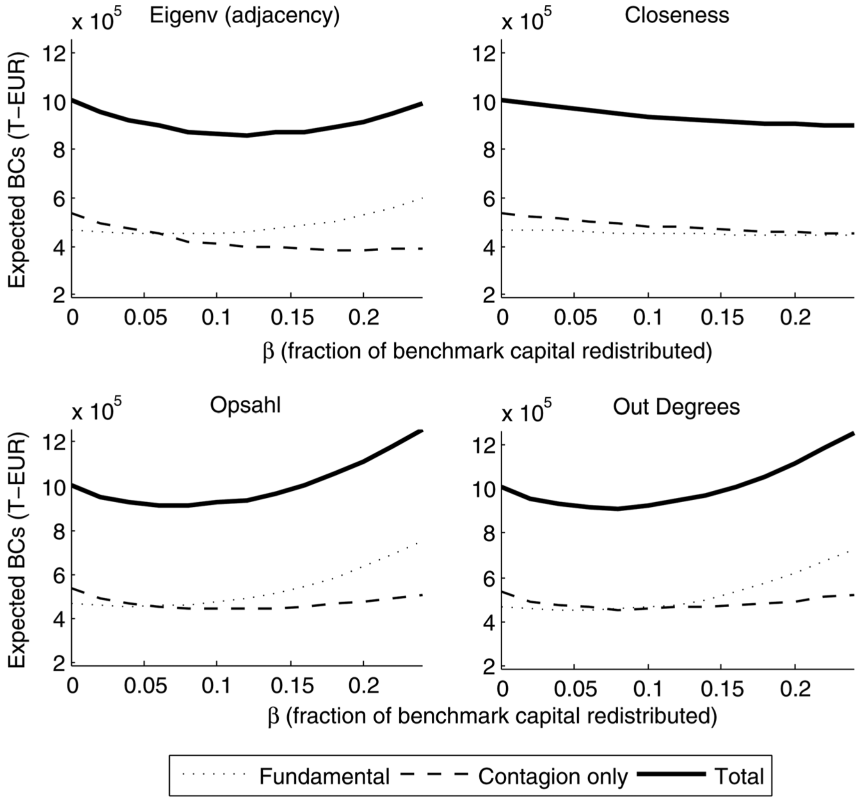

Notes: The y-axes represent expected bankruptcy costs (as measured by equation (7)) from fundamental defaults, contagious defaults, and all defaults, under different capital allocations. On the x-axis, β represents the redistributed fraction of benchmark capital, VaR($\alpha = 99.9\%$).

effective. The other two measures are similar to the behavior of the allocation based on the adjacency eigenvector. As before, the inflection points of fundamental losses lead to a minimum, but featuring some differences. First, the post-contagion losses stop decreasing at around the same value of β; second, the inflection points are sharper; and finally, all of this happens at smaller values of β, and with smaller declines in post-contagion bankruptcy costs. These are less effective measures upon which to base capital reallocations.

Figure 3. Frequency Distributions of Individual Bank PDs

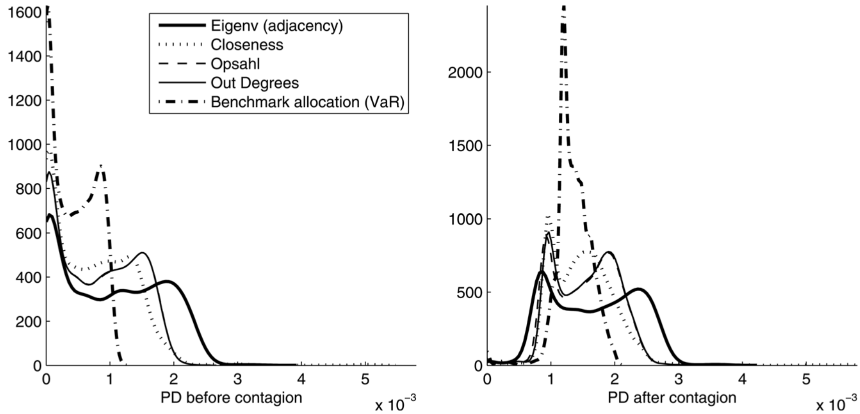

Notes: On the y-axis is presented the estimated density of the distribution of bank PDs. On the x-axis are represented PDs (per bank). Results were obtained with 500,000 simulations.

Figure 3 displays frequency distributions of all banks' default probabilities before and after contagion in the benchmark case (VaR-based capital)[29] and compares them with distributions generated by the best centrality-based allocations. These four capital reallocations are taken at their optimal β, according to the values presented in table 3. All densities show a similar picture for the fundamental defaults ("PD before contagion"). The reallocations all spread the distribution to the right on fundamental defaults. Indeed, fundamental defaults for the best basis, the adjacency eigenvector, spread the fundamental defaults most to the right. We know from the costs of the fundamental defaults that total bankruptcy costs remain the same. They were redistributed such that more (presumably smaller) banks defaulted to save a few whose default triggers potentially larger bankruptcy costs. This is most pronounced for the adjacency-eigenvector-based capital allocation, but it is true for all the other three measures as well: pre-contagion default frequencies shift to the right with a longer and fatter tail than in the benchmark

[29]Recall that the VaR used for capital is not identical with the actual loss quantile in the model that includes contagion, be it before or after contagion; see section 5.

case. This is not surprising, as the PDs before contagion are limited to 0.1 percent by construction (cf. footnote 28); any exceedance is caused by simulation noise. The best allocations based on the other three centrality measures lie somewhere in between the fat-tailed adjacency-eigenvector-based distribution and the thin-tailed benchmark. Still, the lesson from the left-hand side of figure 3 is that the capital allocations based on each of our centrality measures in the initial set of defaults sacrifice a large number of less relevant banks by shifting capital to the relevant banks. There is also, of course, a relationship to size, but we cannot present related graphs for confidentiality reasons.

Post-contagion probabilities of default on the right-hand side of figure 3 show further interesting patterns with respect to the benchmark case. The capital reallocation based on each of our centrality measures continues to sacrifice more of the (presumably) smaller banks in order to reduce the probability of default, post-contagion, for the larger banks. The end result of this is a distribution that is considerably wider in its default probability both pre- and post-contagion. The cost in terms of the banks that are more likely to fail is made up through a few banks (otherwise suffering—or elsewhere causing—large losses, in expectation) that are less likely to fail during the contagion phase of the default cycle. It is remarkable that each of our reallocations behave in this way. The extra capital gained from the reduction in the benchmark capital rule which focuses on pre-contagion default has not reduced post-contagion defaults across the board for any tested reallocation. Instead, in contrast to benchmark allocation, our centrality-based capital rules perform less consistently (at least in this sense), sacrificing many defaults to save those which matter in terms of our loss function.

Figure 3 can also be used as a validation of the "traditional" way of measuring the risk of interbank loans, that is, by treating them as in the benchmark case; cf. section 5. In the latter case, they are part of an ordinary industry sector and driven by a common systematic factor. This treatment is very much in line with the approach taken in the Basel III framework. The bank PDs sampled in figure 3 are default probabilities from the model including contagion. If the bank-individual loss distributions generated under the "traditional" treatment were a perfect proxy of the ones after contagion, we should observe the histogram of PDs after contagion in

Figure 4. Relationship between Default Probabilities of Individual Banks Before and After Contagion

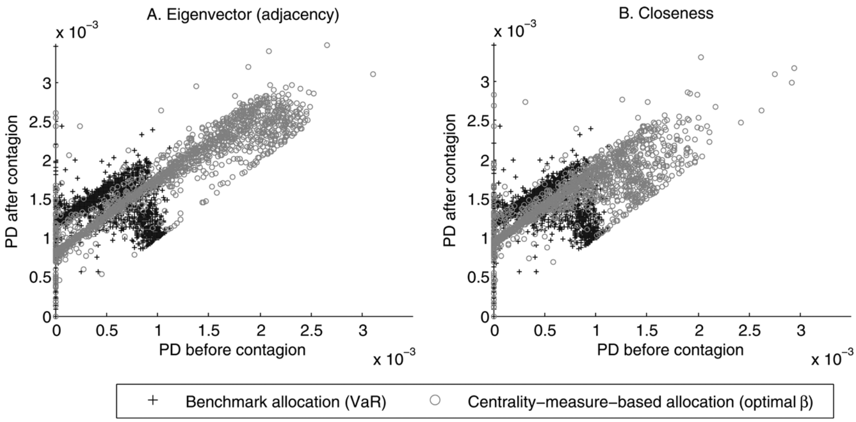

Notes: PDs are not to be interpreted as real; e.g., actual capital holdings of German banks have not found entry into their calculation.

the VaR-based benchmark case to be strongly concentrated around 0.1 percent. Instead, the PDs are widely distributed and even on average 35 percent higher than that possibly concluded from the label "99.9 percent VaR." While the link to the Basel framework is rather of a methodological nature, our analysis clearly documents that interbank loans are special and that correlating their defaults by Gaussian common factors may easily fail to capture the true risk. As there are, however, also good reasons to remain with rather simple "traditional" models, such as those behind the Basel rules, our modeling framework lends itself to a validation of the capital rules for interbank credit. This exercise is beyond the scope of this paper.

In figure 4, each observation represents a bank for which we calculate the probability of fundamental defaults (x-axis) and the PD including both fundamental and contagious defaults (y-axis). (Obviously, there is nothing lower than the 45-degree line in these diagrams.) In the case of the VaR-based benchmark allocation (black markers), the relationship is clearly non-linear. Most banks, although their fundamental PD is effectively limited to 0.1 percent, experience much higher rates of default due to contagion. Suggested by the graph, there seems to be a set of around thirty events where the system fully breaks down, leading to the default of most banks, irrespectively of their default propensity for fundamental reasons. The benchmark capital requirement does a very good job in limiting the probability of fundamental defaults to 0.1 percent, which is what causes the relationship between fundamental and total default probabilities to be non-linear.

In contrast to the benchmark capital, the optimized capital allocation based on our best-performing measure, the adjacency eigenvector, has a stronger linear relationship between fundamental and total PD, although with heteroskedastic variance that increases with fundamental default probability. Indeed, it looks as if the probability of total default could be approximated by the probability of fundamental default plus a constant. This leads us to a conclusion that imposing a capital requirement based on the adjacency eigenvector causes many banks to have a higher probability of default than in the baseline case, but because these tend to be smaller banks, they impose smaller bankruptcy costs on the system as a whole. The reallocation based on closeness does this less. There are more banks with a lower default probability in the post-contagion world than with the adjacency-eigenvector-based allocation. The closeness allocation induces less of a linear relationship between fundamental and total default probabilities. The lesson here is much the lesson of figure 3: the sacrifice of a lot of small banks to reduce the system losses caused by default of the few banks which are most costly. The additional information of these diagrams is that the same banks sacrificed in the pre-contagion phase are those banks sacrificed in the contagion phase. This relationship is strongest for our best-performing allocation basis, the adjacency eigenvector.

In order to get the full picture, we now focus on the distribution of bankruptcy costs. We provide in figure 5 the tail distribution of BCs for benchmark capital and for capital based on the adjacency eigenvector and closeness, plotted for BCs exceeding €100 billion (note that the lines are ordinary densities plotted in the tail; they are not tail conditional). One is cautioned not to place too much emphasis on the BCs being multi-modal. The default of an important bank might trigger the default of smaller banks which are heavily exposed to it, creating a cluster.

Several things are apparent from figure 5. Both closeness and the adjacency eigenvector perform better than the benchmark in

Figure 5. Kernel Density of BCs Greater than €100 Billion

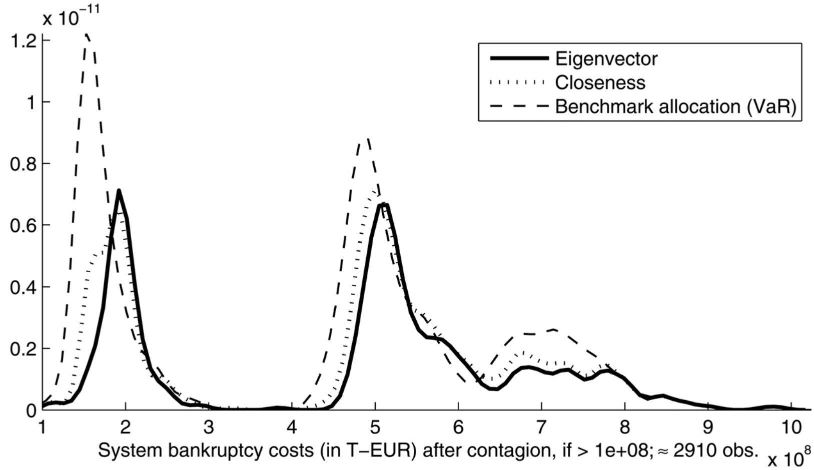

Notes: This figure zooms into a detail of unconditional densities. Densities for BCs smaller than €100 billion are almost identical.

precisely those catastrophic meltdowns where we would hope they would help. The benchmark is dominated by both reallocations at each of the three modes exhibited by the benchmark allocation. Second, the optimal adjacency eigenvector reallocation dominates closeness, although to a lesser extent than the savings from the benchmark allocation. Our optimal capital rule, based on the adjacency eigenvector, shows the best performance in the area where the banking authority would like capital requirements to have teeth. Finally, all rules perform equally poorly for extreme BCs of greater than €780 billion. When the catastrophe is total, allocating an amount of capital that is too small does not make much difference. Almost the entire capital is used and plenty of banks fail.

At this point the question arises whether a capital reallocation combining two centrality measures might perform better than one based on a single measure. Indeed, it could not perform worse, by definition. However, figure 5 shows one of the problems with this logic. Although the two centrality measures perform differently in the end, and despite their low correlation compared with other combinations of centrality measures, the way in which they reshape the loss tail in

Figure 6. Relative Changes of Required Capital for Different Capital Allocations

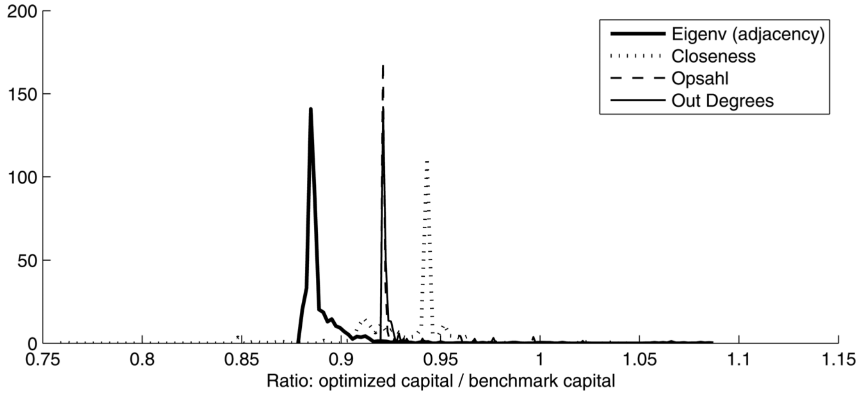

Notes: The horizontal axis is capital under the optimal capital allocations (on the various centrality measures) divided by capital in the VaR-based benchmark case. The maximum is truncated at the 99.5 percent quantiles.

figure 5 is rather similar; in a sense, there is a lack of "orthogonality" in their impact on the loss distribution. As a consequence, the resulting optimum from the combined allocation rule leads only to a 0.2 percent improvement in expected BCs. We also combined other centrality measures and never observed an improvement; all optimal choices were boundary solutions that put full weight on only one of the two centrality measures.

Finally, figure 6 shows, under various capital allocations, bank-individual ratios of centrality-based capital divided by VaR-based benchmark capital. The first observation for all of the measures is that ratios involve a cut to the capital requirement for the vast majority of banks, and the cut is highly concentrated at the same amount within the centrality measure. Each centrality measure cuts a different amount for the majority of the banks, except that Opsahl and out degree have very similar effects in the distribution of their capital requirements. The adjacency eigenvector cuts most drastically, by 12 percent of benchmark capital, and closeness is least drastic, at a 5 percent decrease for the vast majority of banks.

To conclude, we infer for each allocation that the capital reallocation is a comparably "mild" modification of the benchmark case for almost all of the banks because the extra capital allocated to some banks is not excessive. Another intuitive inference is that the other centrality measures might mis-reallocate capital. For example, let's say that benchmark capital for a bank represents 8 percent of the total assets. Forcing this bank to hold thirteen times more capital under the new allocation rule would mean it would have to hold more capital than total assets. Due to this misallocation likelihood, we infer that more constraints should be set in order to achieve better-performing allocations using centrality measures. We leave this extra mile for future research. Clearly, our general approach to capital reallocation is focused on a few particular banks for each of the centrality measures, simply because the measures focus on them. Beyond that, we cannot say much without revealing confidential information about individual banks.

For most of our centrality-based capital reallocations, there is a reduction of the expected BCs in the whole system, at least over a range of capital redistributed. In the allocations we examined most closely, these new requirements function on the same principle of shifting capital away from those banks that were presumably small and had little effect on contagion towards a few banks that had a large effect. The shift has the outcome one might expect: more banks are likely to fail in the fundamental defaults phase of our experiment, even in those instances where total costs due to bankruptcy do not increase during this phase. The less expected outcome is that even in the post-contagion phase, more banks fail under the capital reallocation. However, failed banks have a lower systemic cost in terms of bankruptcy, on average, than those banks that are saved through the reallocation. Under some circumstances, this is true even in the pre-contagion phase of defaults. Our two most effective reallocations, based on the adjacency eigenvector and closeness, give considerable savings in terms of total bankruptcy costs: 14.6 percent and 10.7 percent, respectively. The reallocations based solely on these measures perform just about as well as when the capital requirement is based on a combination of both.

7. Conclusion

In this paper we present a tractable framework that allows us to analyze the impact of different capital allocations on the financial stability of large banking systems. Furthermore, we attempt to provide some empirical evidence of the usefulness of network-based centrality measures. Combining simulation techniques with confidential bilateral lending data, we test our framework for different capital reallocations. Our aim is twofold: first, to provide regulators and policymakers with a stylized framework to assess capital for systemically important financial institutions; second, to give a new direction to future research in the field of financial stability using network analysis.

Our main results show that there are certain capital allocations that improve financial stability, as defined in this paper. Focusing on the system as a whole and defining capital allocations based on network metrics produces results that outperform the benchmark capital allocation which is based solely on the portfolio risk of individual banks. Our findings come as no surprise when considering a stylized contagion algorithm. The improvement in our capital allocations comes from the fact that they take into account the "big picture" of the entire system where interconnectedness and centrality play a major role in triggering and amplifying contagious defaults. As for the best network measure, we find that the capital allocation based on the adjacency eigenvector dominates any other centrality measure tested. These results strengthen the claim that systemic capital requirements should depend also on interconnectedness measures that take into account not only individual bank centrality but also the importance of their neighbors. In the optimal case, by reallocating 12 percent of capital based on eigenvector centrality, expected system losses (measured by expected bankruptcy costs) decrease by almost 15 percent from the baseline case.

As shown by Löffler and Raupach (2013), market-based systemic risk measures can be unreliable when they assign capital surcharges for systemically important institutions (as can other alternatives such as the systemic risk tax proposed by Acharya et al. 2012). What we propose in this paper is a novel tractable framework to improve system stability based on network and balance sheet measures. Our study complements the methodology proposed by Gauthier, Lehar, and Souissi (2012). As in that paper, we take into account the fact that market data for all financial intermediaries does not exist when dealing with large financial systems. Instead we propose a method that relies mainly on the information extracted from a central credit register.

We are not providing further details on how capital reallocation in the system could be implemented by policymakers. This aspect is complex, and the practical application involves legal and political issues. Further, we assume the network structure to be exogenous. In reality, banks will surely react to any capital requirement that is based on measurable characteristics by adjusting the characteristics. Such considerations are beyond the scope of the paper. Future research is needed to measure the determinants of endogenous network formation in order to make a more precise statement about the effects of a capital regulation as suggested in this paper.

Further future research could go in several directions. For example, bailout rules used for decisions on aid for banks that have suffered critical losses could be based on information gained from centrality measures. How this bailout mechanism would be funded and the insurance premium assessed to each bank are central to this line of research. Finally, the method can be extended by including information on other systemically important institutions (e.g., insurance companies or shadow banking institutions), although further reporting requirements for those institutions would be necessary.

References

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi. 2013. "Systemic Risk and Stability in Financial Networks." NBER Working Paper No. 18727.

Acharya, V. V., L. H. Pedersen, T. Philippon, and M. Richardson. 2012. "Measuring Systemic Risk." CEPR Discussion Paper No. 8824.

Alessandri, P., P. Gai, S. Kapadia, N. Mora, and C. Puhr. 2009. "Towards a Framework for Quantifying Systemic Stability." *International Journal of Central Banking* 5 (3): 47–81.

Altman, E. 1984. "A Further Empirical Investigation of the Bankruptcy Cost Question." *Journal of Finance* 39 (4): 1067–89.

Anand, K., P. Gai, S. Kapadia, S. Brennan, and M. Willison. 2013. "A Network Model of Financial System Resilience." *Journal of Economic Behavior and Organization* 85: 219–35.

Basel Committee on Banking Supervision. 2011. "Global Systemically Important Banks: Assessment Methodology and the Additional Loss Absorbency Requirement." Bank for International Settlements.

Battiston, S., D. D. Gatti, M. Gallegati, B. Greenwald, and J. E. Stiglitz. 2012. "Liaisons Dangereuses: Increasing Connectivity, Risk Sharing, and Systemic Risk." *Journal of Economic Dynamics and Control* 36 (8): 1121–41.

Battiston, S., M. Puliga, R. Kaushik, P. Tasca, and G. Caldarelli. 2012. "Debtrank: Too Central to Fail? Financial Networks, the Fed and Systemic Risk." *Scientific Reports* 2 (Article 541): 1–6.

Bluhm, C., L. Overbeck, and C. Wagner. 2003. *An Introduction to Credit Risk Modeling*. Chapman & Hall.

Bonacich, P. 1987. "Power and Centrality: A Family of Measures." *American Journal of Sociology* 92 (5): 1170–82.

Cont, R., A. Moussa, and E. B. Santos. 2013. "Network Structure and Systemic Risk in Banking Systems." In *Handbook on Systemic Risk*, ed. J.-P. Fouque and J. A. Langsam, 327–68 (chapter 13). Cambridge University Press.

Craig, B., and G. von Peter. 2014. "Interbank Tiering and Money Center Banks." *Journal of Financial Intermediation* 23 (3): 322–47.

Dangalchev, C. 2006. "Residual Closeness in Networks." *Physica A* 365 (2): 556–64.

Davydenko, S. A., and J. R. Franks. 2008. "Do Bankruptcy Codes Matter? A Study of Defaults in France, Germany, and the U.K." *Journal of Finance* 63 (2): 565–608.

Eisenberg, L., and T. H. Noe. 2001. "Systemic Risk in Financial Systems." *Management Science* 47 (2): 236–49.

Elsinger, H., A. Lehar, and M. Summer. 2006. "Risk Assessment for Banking Systems." *Management Science* 52 (9): 1301–14.

Gabrieli, S. 2011. "Too-Interconnected versus Too-Big-to-Fail: Banks Network Centrality and Overnight Interest Rates." Available at SSRN.

Gai, P., and S. Kapadia. 2010. "Contagion in Financial Networks." *Proceedings of the Royal Society, Series A: Mathematical, Physical and Engineering Sciences* 466 (2120): 2401–23.

Gauthier, C., A. Lehar, and M. Souissi. 2012. "Macroprudential Capital Requirements and Systemic Risk." *Journal of Financial Intermediation* 21 (4): 594–618.

Grunert, J., and M. Weber. 2009. "Recovery Rates of Commercial Lending: Empirical Evidence for German Companies." *Journal of Banking and Finance* 33 (3): 505–13.

Gupton, G., C. Finger, and M. Bhatia. 1997. "CreditMetrics – Technical Document." JP Morgan & Co.

Haldane, A. 2009. "Rethinking the Financial Network." Speech delivered at the Financial Student Association, Amsterdam, April 28.

Löffler, G., and P. Raupach. 2013. "Robustness and Informativeness of Systemic Risk Measures." Discussion Paper No. 04/2013, Deutsche Bundesbank.

Minoiu, C., and J. A. Reyes. 2013. "A Network Analysis of Global Banking: 1978–2009." *Journal of Financial Stability* 9 (2): 168–84.

Newman, M. 2010. *Networks: An Introduction.* Oxford, UK: Oxford University Press.

Opsahl, T., F. Agneessens, and J. Skvoretz. 2010. "Node Centrality in Weighted Networks: Generalizing Degree and Shortest Paths." *Social Networks* 32 (3): 245–51.

Rogers, L., and L. A. Veraart. 2013. "Failure and Rescue in an Interbank Network." *Management Science* 59 (4): 882–98.

Roukny, T., C.-P. Georg, and S. Battiston. 2014. "A Network Analysis of the Evolution of the German Interbank Market." Discussion Paper No. 22/2014, Deutsche Bundesbank.

Sachs, A. 2014. "Completeness, Interconnectedness and Distribution of Interbank Exposures—A Parameterized Analysis of the Stability of Financial Networks." *Quantitative Finance* 14 (9): 1677–92.

Schmieder, C. 2006. "European Data Watch: The Deutsche Bundesbank's Large Credit Database (BAKIS-M and MiMiK)." *Schmollers Jahrbuch* 126 (4): 653–63.

Soramäki, K., K. M. Bech, J. Arnolda, R. J. Glass, and W. E. Beyeler. 2007. "The Topology of Interbank Payment Flows." *Physica A* 379 (1): 317–33.

Soramäki, K., and S. Cook. 2012. "Algorithm for Identifying Systemically Important Banks in Payment Systems." Discussion Paper No. 2012-43.

Standard and Poor's. 2011. "2010 Annual Global Corporate Default Study and Rating Transitions."

Tarashev, N., C. Borio, and K. Tsatsaronis. 2010. "Attributing Systemic Risk to Individual Institutions." BIS Working Paper No. 308.

Upper, C. 2011. "Simulation Methods to Assess the Danger of Contagion in Interbank Markets." *Journal of Financial Stability* 7 (3): 111–25.

Watts, D. J., and S. Strogatz. 1998. "Collective Dynamics of 'Small-World' Networks." *Nature* 393 (6684): 440–42.

Webber, L., and M. Willison. 2011. "Systemic Capital Requirements." Working Paper No. 436, Bank of England.

Yellen, J. L. 2013. "Interconnectedness and Systemic Risk: Lessons from the Financial Crisis and Policy Implications." Speech at the American Economic Association/American Finance Association Joint Luncheon, San Diego, California, January 4.

Zeng, B., and J. Zhang. 2001. "An Empirical Assessment of Asset Correlation Models." Moody's Investors Service.