

An Empirical Analysis of Funding Costs Spillovers in the EURO-zone with Application to Systemic Risk *

Pietro Bonaldi[†] Ali Hortaçsu[‡] Jakub Kastl[§]

First version: March 2013

This version: June 2015

Abstract

We propose a framework for estimation of spillovers between funding costs of individual banks. The estimation proceeds in three steps: First, using data from liquidity auctions of the European Central Bank, we estimate the funding costs in a given week for each individual bank. In the second step, we apply the adaptive elastic net (a LASSO type estimator) to this panel to estimate the financial network. Finally, using the estimated network we propose new measures of the systemicness and vulnerability of each bank. Our measure of systemicness has quite a natural interpretation, since it can roughly be viewed as the total externality a bank would impose on the funding costs of all other banks in the system. We estimate that most of the banks have fairly weak links and, therefore, if one were to suffer an adverse shock there would likely be a rather limited effect on the other ones. On the other hand, there are a few banks that are quite central: an increase in their funding costs would result in a very significant increase (up to 95 basis points per 100 basis points shock) in the funding costs of the other banks. Our vulnerability scores estimated using data from 2007-2008 are positively correlated with the probability of a bank being bailed out later.

Keywords: systemic risk, financial crisis, multiunit auctions, liquidity

*The paper was previously circulated under the title “An Empirical Analysis of Systemic Risk in the EURO-zone.” We thank Aureo de Paula, Vasco Carvalho, Thomas Phillipon, Marzena Rostek, José-Luis Peydró, David Skeie, Charles M. Kahn, Christian B. Hansen, Azeem Shaikh, Fernando Alvarez and seminar participants at Berkeley, EUI, Minneapolis Fed, NYU, Princeton, Stanford, UCLA, UCL, Wisconsin, 2013 GSE Summer Forum, 2013 Wharton Conference on Liquidity and Financial Crises, 2013 Asian Meeting of the Econometric Society, 2014 NBER Summer Institute, 50th Annual Conference on Bank Structure and Competition, Bank of England’s Conference on Systemic Risk and Macro-Prudential Regulation and 3rd European Meeting on Networks for helpful comments. We are also grateful to Winnie van Dijk for valuable research assistance. Hortaçsu acknowledges financial support from the NSF (SES-1124073 and ICES-1216083). Kastl acknowledges financial support from the NSF (SES-1123314 and SES-1352305) and the Sloan Foundation. The views expressed in this paper are our own and do not necessarily reflect the view of the European Central Bank. All remaining errors are ours.

[†]Department of Economics, University of Chicago

[‡]Department of Economics, University of Chicago and NBER

[§]Department of Economics, Princeton University and NBER

1 Introduction

Starting with Rochet and Tirole (1996) and Allen and Gale (2000), the importance of the interconnectivity structure of financial institutions for the stability of the financial system has been seen as a topic of first-order importance. Additional theoretical research by Freixas, Parigi and Rochet (2000), Eisenberg and Noe (2001), Elliot, Golub and Jackson (2014), Glasserman and Young (2015) and Açemoglu, Ozdaglar and Tahbaz-Salehi (2015) among others, has shown that certain network structures can lead to the presence of “systemic” institutions, i.e. institutions which, due to their connectedness, can transmit adverse shocks to the entire system. The empirical characterization of the actual interconnectivity structure of banks, however, is relatively scarce. This delay is undoubtedly related to the lack of direct data on interbank credit relations, save for information available through intraday payment and settlement systems (see, for example, Boss, Elsinger, Summer and Thurner (2004), Bech and Atalay (2008), Iori, Masi, Precup, Gabbi and Caldarelli (2008), Wetherilt, Zimmerman and Soramäki (2010) and Ashcraft, McAndrews and Skeie (2011)).

In this paper, we attempt to provide empirical tools that would allow Central Banks or potentially other regulators to infer the structure of the network of spillover effects based on comovements of banks’ short-term funding costs. Such comovements can be generated as the result of the propagation of solvency shocks as modeled in Eisenberg and Noe (2001), and Açemoglu et al. (2015). Once we estimate the financial network, we use centrality measures typically used in network analysis (e.g., Ballester, Calvó-Armengol and Zenou (2006)) to estimate the extent to which a particular bank is crucial for the working of the whole system, i.e., to determine the set of systemic banks. Our method provides a measure of such “systemicness” as expressed by that bank’s impact on the future cost of funding of other banks. Similarly, we measure to what extent each bank is affected by shocks to the funding costs of other banks in the system, in other words, how vulnerable to contagion each bank is.

While different proxies for funding costs of banks may be utilized (CDS spreads, LIBOR/EURIBOR quotes, imputed interest rates from overnight markets etc.), we focus on a particular measure that

is (i) *revealed by actions that are directly affecting payoffs* as opposed to stated willingness-to-pay for credit as elicited in a survey, and (ii) available for a large set of banks at a fairly high frequency (weekly). In particular, following our prior work (Cassola, Hortaçsu and Kastl (2013)) we use data on European banks' bids in the ECB's short-term refinancing auctions and infer banks' revealed funding costs using a structural econometric model of strategic bidding in these auctions. In Cassola et al. (2013), we looked at the evolution of banks' behavior before and during the early part of the financial crisis and argued that the bidding data from the main refinancing operations of the ECB is predictive of the financial woes of system banks. We also argued that interpreting the data through a structural model is important, since changes in the bidding behavior itself involve also strategic adjustment to the changes in other bidders' behavior (strategies), which in case of a financial turmoil may be quite important. In other words, a bank i may be bidding higher not because its underlying opportunity cost of funding increased, for example due to unfavorable development of other banks' perceptions of i 's default risk, but because other banks increased their bids and i thus also increased its bid so as to optimally resolve the trade-off between the surplus on the marginal amount of the loan and the probability of winning this marginal amount. That prior paper also showed that the estimates of the changes in banks' funding costs coming from the model are consistent with the ex-post observed changes in accounting measures of banks' performance such as cost-to-income ratio or return on equity.

Given our panel of banks' revealed short-term funding costs, we characterize the network structure by looking at the covariation of a given bank's funding cost with other system banks' (lagged) funding costs, controlling for various sources of comovement due to common asset exposures, such as holdings of sovereign bonds (we conduct numerous robustness exercises for our specification in the paper and in the Appendix, including an instrumental variables strategy). Implementing this through traditional linear regression, however, yields a dimensionality problem: the number of covariates in our regressions potentially exceeds the number of data points we have for each bank in our data set. To combat this problem, which is commonly encountered in Genetics and Machine learning, we utilize the *adaptive elastic net* method of Zou and Zhang (2009) and Zou and Hastie (2005). As will be discussed in Section 4.3, the elastic net is a modification of the Least Absolute

Shrinkage and Selection Operator (LASSO) method of Tibshirani (1996), which has been shown to have several desirable properties in recovering sparse relationship patterns in other contexts. In particular, as opposed to other methods for model selection, like the standard LASSO, the elastic net performs well under the presence of highly correlated covariates, as measured by prediction accuracy, which is a feature that likely applies in the case of financial networks.

Given our estimates of interbank interactions, we build the network of spillover effects as a directed graph with banks for nodes. A link from i to j represents that an increase in the cost of funding of bank i is associated with a positive change in the cost of funding of bank j , one period ahead. Moreover, we assign a weight to this link equal to the predicted response in j 's cost of funding corresponding to a 100bp positive shock to i 's cost of funding. We use the weighted network to calculate each bank's systemicness, using centrality measures developed in the network literature. We chose a generalized version of Katz centrality, traditionally used in social network analysis to determine the influence of a node within the network. Such a measure takes into account the direct effect of a bank on its immediate neighbors (those banks that are directly affected by it) but also on the neighbors of its neighbors, as well as on their neighbors, and so on and so forth. Our results suggest that most of the banks are only weakly connected and, therefore, if one were to suffer an adverse shock there would likely be a rather limited effect on the other ones. On the other hand, there are a few banks that are quite central: an increase in their funding costs would result in a very significant increase (up to 95 basis points per 100 basis points shock) in the funding costs of the other banks.

A vulnerability measure, similar to the systemicness measure, can also be derived from the connectivity matrix. Intuitively, this is a measure of how shocks to banks other than bank i will affect bank i 's funding costs – i.e. which banks are most exposed to shocks affecting other banks in the system. We develop this measure in our setting, which allows us to rank banks based on their vulnerability. Interestingly, we find that banks' vulnerability scores (as measured from our data in 2007 and 2008) contribute significantly to predict the probability that a bank would be subsequently bailed out. We present the statistical evidence for this result in section 5.3.

Several other empirical measures of network linkages among financial institutions have been

developed in the literature, utilizing different sources of data and methods. Battiston, Puliga, Kaushik, Tasca and Caldarelli (2012), analyse the network of cross ownership relations (equity investment) among a small sample of banks that were among the top recipients of aid from the FED through its emergency loans program, during 2008-2010. They use a feedback centrality measure to rank banks according to their systemic impact. Greenwood, Landier and Thesmar (2015) rank banks according to how vulnerable they are to shock propagation. However, they focus exclusively on shocks to equity that force assets liquidation in order to meet a target leverage ratio. In their model, fire sales of assets decrease other banks' equity, providing a channel for propagation. In contrast, our approach could, in principle, capture other sources of contagion, such as direct contractual exposures among banks. Most akin to our measure, Billio, Getmansky, Lo and Pelizzon (2012) have utilized lag relationships between stock returns (in particular bivariate Granger causality tests) to characterize the network structure of the largest financial institutions in the US. Their exercise relies on the assumption of informationally inefficient securities markets, so that predictive correlation patterns between publicly traded securities are possible. Our data on banks' revealed funding costs, however, is not public information, and is observed by the ECB only. The banks only observe their own bids, their final allocation, the market clearing interest rate and the average interest paid in the auction.¹ We discuss the connection between our network estimates and those in Billio et al. (2012) in Section 5.6.1. Diebold and Yilmaz (2011) also estimate a network of connections among financial institutions. As in our case, their reduced form model is a vector autoregression, but their variable of interest are stock return volatilities. Their network is defined as the (h-step) forecast error covariance matrix. From such network, they derive measures of *total connectedness* to, and from other banks, analogous to our systemicness and vulnerability measures, respectively. One advantage of our measure is that the units in which it is reported can be directly interpreted as changes in funding costs. Empirical measures of systemic risk that directly estimate the overall impact of individual institutions on the whole system, have also been proposed

¹While banks potentially may also reveal information through their LIBOR/EURIBOR quotes, the validity of this "stated" information has largely been challenged due to admitted "strategic disclosure" by some large panel banks. In fact, the U.S. Department of Justice and Barclays (2012) have acknowledged manipulation of Barclay's LIBOR and EURIBOR submissions. Barclays is far from an exception, several panel banks have paid billions of dollars in fines and settlements to regulators in the US and Europe, for attempting manipulation of these rates, by submitting quotes that do not accurately reflect their borrowing costs.

in the financial economics literature. Acharya, Pedersen, Philippon and Richardson (2010) provide a simple theoretical framework, in which they propose to measure systemic risk by an institution’s marginal contribution to the shortfall of capital in the financial system that can be expected in a crisis. Instead, Brownlees and Engle (2010) measure it by the expected shortage of capital of an institution given its degree of leverage. Adrian and Brunnermeier (2011) measure the marginal contribution of each institution to the “value at risk” of the whole system, conditional on that institution being under distress relative to its median state (hence the name CoVaR). In section 5.6.2, we compare our estimates to those obtained from a simple specification of CoVaR, for a small set of large publicly traded banks.

The remainder of the paper proceeds as follows. We begin by describing the model of systemic risk in section 2. In section 3 we provide details on the several data sets that we merge together. We continue in section 4 with the description of how we use the data to estimate the parameters of our model. In section 5 we discuss the results of our analysis, and conduct several robustness checks (several additional robustness checks are reported in the Appendix). We conclude in section 6.

2 Model

We model the European banking system, consisting of N risk-neutral banks, as a sequence of periodic markets (e.g., quarterly or monthly markets) indexed by τ . The European Main Refinancing Operations are run every week, and we assume that the periodic markets have lower frequency than that for reasons that will be made clear below. In the beginning of each period τ , banks collect observable information about all rivals (default probability, liquidity ratios, accounting performance, past loan performance etc). This observed information implies a partition of banks into M classes, where $M < N$. Each class m is a sufficient statistic for (ex-ante) characteristic of a bank in that period τ . Based on this partition, banks create new or modify existing links to other banks to form a network, which is then fixed for the next period τ . We do not take a stand on whether the links are based on equity stakes (as in Elliot et al. (2014)) or on debt contracts (as in Açemoglu et al. (2015)) or some combination of these and other exposures. The class m_τ that a bank i belongs to during period τ together with its position in the network determines the distribution,

$F_i(\theta|m_\tau)$, from which i draws its funding costs. We further assume that each τ is partitioned into sub-periods (weeks), indexed by t , which correspond to the frequency with which banks secure their funding (from counterparties or from the central bank) to finance their investment opportunities and everyday operations. The crucial distinction between τ and t is that at each τ , banks have an opportunity to revise their terms with their counterparties by incorporating information revealed over the past periods into their beliefs about each counterparty’s default probability, profitability, network position etc. A bank i ’s class m , therefore, may vary across τ as unexpected shocks to i arrive and also as shocks arrive to banks, to which i is directly or even indirectly connected.

Appendix A.1 describes a more formal network model, which is a straightforward extension of Açemoglu et al. (2015) that incorporates the heterogeneity of banks due to the M different classes. In this model banks extend loans to each other in order to finance each others’ investments. We further assume that the (quantity-weighted average) cost of such loans determines the mean of a distribution of funding costs, from which banks get an independent draw each week within each periodic market. For example, if classes m_1, \dots, m_M were (stochastically) increasing in a bank’s default probability, it would be natural to expect that the distributions of funding costs would also be stochastically increasing in m . At the end of period τ , shocks are realized and payments are settled. The average funding costs of each of their rivals and the outcome of the payments become common knowledge. Based on the realizations of those, banks get again partitioned among the M classes and the next period begins, i.e., some network links might get dissolved and some new ones may be formed. For simplicity, we assume that banks within a class m are ex-ante symmetric, $F_i(\theta|m_\tau) = F(\theta|m_\tau) \forall i \in m_\tau$.

Timing within a period τ :

$t = 0$: Information about past funding costs, default probability and balance sheets becomes public. Banks enter into new or adjust their existing bilateral funding relationships and are subsequently partitioned into M “quality” classes.

$t = 1, \dots, T - 1$: Bank i repays its short-term financing loans and draws a signal θ_{it} , from a distribution $F(\theta|m_\tau)$ where m_τ is the class that i belongs to. θ determines i ’s funding costs in week t , $v(q, \theta)$, that i faces when obtaining short-term financing - either from the central bank or

from a counterparty.

$\tau = T$: Bilateral payments are settled (a payment equilibrium is played) and banks learn each others' realized mean funding costs, $\bar{v}_{i\tau}$.

Our goal is to utilize the dynamics of the cost of funding to recover the network. We will thus be estimating the “average” network during our sample.

Using a model of a discriminatory share auction described in Cassola et al. (2013), we can recover the funding costs corresponding to a bid b_k for a loan of size q_k using the necessary condition for equilibrium bidding:

$$v(q_k, \theta_i) = b_k + \frac{\Pr(b_{k+1} \geq P^c | m)}{\Pr(b_k > P^c > b_{k+1} | m)} (b_k - b_{k+1}) \quad (1)$$

where P^c is the market clearing price, which is random from perspective of each bank. The uncertainty about its realization is what creates a wedge between the bid and the true “willingness-to-pay” or, in our interpretation, the opportunity costs of obtaining funding from a private counterparty. The intuition for our approach to network estimation comes from comparative statics of the payment equilibrium at the end of period τ : Suppose bank i poses a greater risk for the financial system than bank j , but both are important large banks belonging to the same class m . Further suppose that bank i suffers an adverse shock to its balance sheet during period τ , i.e., its average funding cost, $\bar{v}_{i\tau} = \frac{\sum_{t=1}^T q_{it} v(q_{it}, \theta_{it})}{\sum_{t=1}^T q_{it}}$ is high. Ceteris paribus, the new distribution $F(\theta | m_{\tau+1})$ dominates in the first order stochastic sense the distribution $F(\theta | m_\tau)$ since i now has to pay some risk premium on its loans due to a worse balance sheet. This in turn implies an increase in the willingness-to-pay for liquidity obtained in the main refinancing operations from the ECB for bank i . Since i is assumed to pose a greater risk for the financial system than bank j , we should also observe a larger impact on the funding costs of other banks. In particular, banks that have some exposure to i will also now have to pay a higher risk premium when getting a loan as there is a higher risk that a loan to them will get unpaid since i might fail on its obligations to them. Ideally, we would also like to use some information on exposure of individual banks to risk associated with bank i . Unfortunately, such data is currently not available to us. Instead, we will use regression

analysis together with model selection techniques in order to try to infer the overall impact of a shock to bank i 's funding costs on the other banks in the system.

Each bank with particular weekly needs of financing should, on average, be able to get a loan on the interbank market at the (average) rate $\bar{v}_{i\tau}$. There will be week-to-week fluctuations in the actual cost of funding, which may be related to various unexpected shocks, positive or negative, that this bank (or others) may be exposed to.² As described above, every week each bank obtains a (possibly multidimensional) signal, θ_i which affects its cost-of-funding (or, alternatively, of its willingness-to-pay for a loan of size q in the central bank's auctions) for that particular week, $v(q, \theta_i)$, from a distribution $F(\theta_i | m, \tau, X_t)$, where X_t is all public information, such as prevailing overnight unsecured interest rate EONIA, that is available that week. To recover the funding costs from the bids submitted in the discriminatory auctions we assume as is standard in the auction literature that banks play a Bayesian Nash Equilibrium.

Since the main goal of this article is not to provide tools and methodology for estimating the auction part of our model, we refer the reader to our earlier work for more detailed discussion and analysis. We summarize the approach only briefly. Our assumptions discussed above ensure that banks' willingness-to-pay is private and conditionally independent, i.e., that we are in an environment with conditionally independent private values. This means that conditional on all public information, X , available before a given auction, each bank's willingness-to-pay (or opportunity costs of funding) is a function only of its own information and is independent of any rivals' information. While both independence and private values are clearly restrictive assumptions, we showed in our previous work that the estimates produced by such model pass several ex-post tests. We therefore view these estimates as a useful source of information about the situation of individual banks during the crisis.³

Our model of a discriminatory auction is based on the classic Wilson's (1979) paper on share auctions. There is a unit perfectly divisible good to be sold and bidders submit bids for shares

²For example, Freixas, Martin and Skeie (2011) argue that the interbank market is absolutely crucial for a bank with uncertain future liquidity needs.

³As an example of the ex-post tests that we have in mind here, our estimates of the changes in a bank's willingness-to-pay are correlated in the expected direction with changes in several accounting measures of performance between 2006 and 2007, i.e., at the onset of the financial crisis, such as return on equity or cost-to-income ratio (Table V in Cassola et al. (2013)).

of this good. The number of potential bidders (banks), N , is commonly known. The number of bidders in each class m in period τ , $N_{m\tau}$ is also commonly known.

In practice, we use equation (1) to identify the marginal values at all but the last step, and we use equation (A-4) at the last step.⁴ Note that as $K \rightarrow \infty$, (1) and (A-4) coincide in the limit. Proposition 2 in Kastl (2012) provides an equilibrium existence result for our environment. It is also important to point out that as the number of players in each group m grows large, the best responses essentially become symmetric: two banks with the same cost of funding from two different groups would submit the same bids. This is a consequence of an argument based on the Law of Large Numbers, which has been used in Swinkels (2001).

3 Data

Our primary data set consists of all bids in the main refinancing operations of the ECB during 2007 and 2008. In total, there were 91 auctions before the switch to the full allotment took place in mid-October. The data from the MRO auctions is described in detail in Cassola et al. (2013). Using these data we obtain estimates of the banks' funding costs using the estimate willingness-to-pay for different amounts of funding from the ECB at each auction. Hence, we have time series of the quantity weighted funding costs of each bank. As further explained in Section 4.5, we express the funding costs as spreads over the overnight unsecured interest rate, EONIA, which is the interest rate the ECB targets.

For the estimation of the network of spillover effects, we restrict the sample of all banks that participate at the MRO auctions to only 32 large banks in the EURIBOR Panel. These banks belong to the largest and most important financial institutions in Europe. They participate in the MROs virtually every week. While in principle our estimation method could handle using all banks to map out the financial network, since we have only 21 months of data, we prefer to restrict our attention to the financial network only among these 32 largest banks. Besides, for some of our estimations we also need stock and CDS prices, which are not available for most of the smaller banks. We obtain information on the balance sheets of these banks from Bankscope. For all banks

⁴Using (A-4) at all steps leads to qualitatively very similar results, but the estimates turn out to be less precise due to the necessity to numerically estimate the derivative of the distribution of the market clearing price.

Table 1: Balance Sheet Variables of the Banks in the Sample

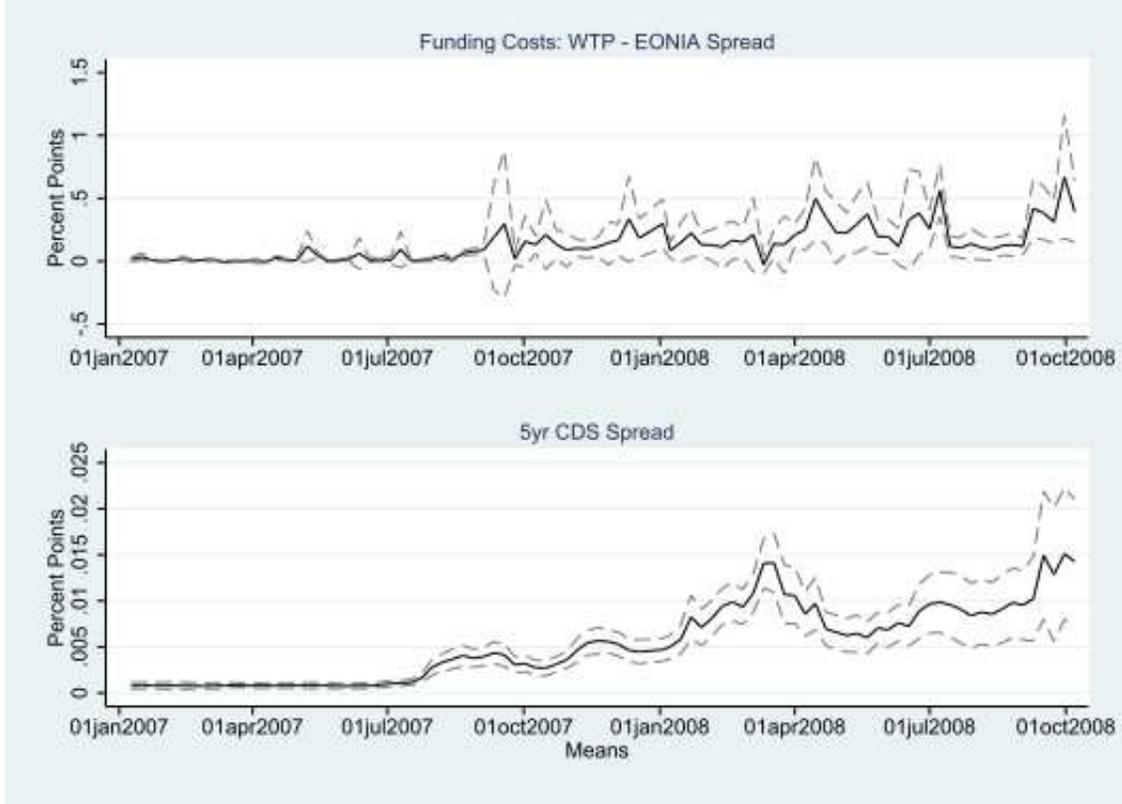
32 Banks in the EURIBOR Panel	2007		2008	
	Mean	Std Dev	Mean	Std Dev
Total Assets (US\$ Billion)	630.162	473.363	691.255	570.602
Return on Assets	0.56%	0.51%	0.11%	0.62%
Return on Equity	11.88%	11.03%	-1.51%	18.45%
Debt to Equity	0.30	0.16	0.41	0.32
Tier 1 Ratio	7.72	1.40	8.20	1.74

in our restricted sample, we gathered additional data on total assets, return on assets, return on equity, debt to equity ratio and Tier 1 Capital ratio. Table 1 shows descriptive statistics of these variables, providing a broad picture of the size and performance of the banks in our final sample.

As additional measures of credit risk, we collected data from Markit on spreads of five-year senior Credit Default Swaps (CDS) referenced to banks' bonds. Also, for a subset of these 32 banks that are publicly traded, we have data on market capitalization from Bloomberg. In Figure 1, we plot the average willingness-to-pay - EONIA spread, and the average CDS spread for the banks in the sample. We include one standard deviation bands to indicate the dispersion of these measures across banks.

We also use reports of the European Commission on government interventions in individual banks (EC 2011). For 629 banks that appear in the liquidity auctions from 2007 and 2008, we identified 20 banks that received targeted government support at least once. Eight of them belong to the group of 32 banks in the restricted sample. On top of these targeted ones, there were also numerous non-targeted bailouts in which various governments helped multiple or all of the banks in their respective countries. Table 2 shows that 50% of these banks in fact received help in multiple rounds. The most notorious recipients of government funds were the Anglo Irish Bank Corporation, which was eventually nationalized, and the IKB Deutsche Industriebank, which after a generous government injection was almost fully privatized in August 2008. From the other 18 banks, 1 more is from Ireland, 8 are German, 2 are from Austria, Belgium and Netherlands, and 1 bank comes from France, Greece and Slovenia. The bailouts can be categorized into several waves. The first wave, which included 5 bailouts, occurred before the collapse of Lehman Brothers between February 2008 and September 2008. After the subsequent change of the ECB's liquidity providing

Figure 1: Average Willingness-to-Pay and CDS Spreads for the Banks in the Sample (one standard deviation bands included)



mechanism from auctions to fixed rate tenders the second wave followed with 14 bailouts before May 2009. Later, there were additional 17 bailouts before July 2011.

In order to control for common sources of variation in the bank's funding costs, we also gathered data from Thomson Reuters' Datastream on government bonds yields, at different maturities, for eleven European countries. In total, we have 30 series of the yield to maturity (YTM) of government bonds for the relevant time range (January 2007 - October 2008). Table 3 lists all these bonds by issuer and maturity.

We now move on to discuss how to link our model and these data to get at a measure of systemic risk.

Table 2: Data Summary: Bailouts

Summary Statistics of Bailouts		
# of Bailouts	# of Banks	Average Size (in mil €)
0	609	
1	10	3,327
2	7	18,397
3	1	4,425
4	1	5,825
5	0	
6	1	7,359

Table 3: List of All Bonds Included as Controls in the Estimation

Issuer\Maturity	2YR	5YR	10YR
Austria	X	X	X
Belgium	X	X	X
Denmark	X	X	X
Spain	X	X	X
France	X	X	X
Finland	X	X	X
Germany	X	X	X
Greece	X	X	X
Ireland			X
Italy	X	X	X
United Kingdom	X		X

We obtained data on bonds yield-to-maturity from Thomson Reuters' Datastream.

4 Estimation

We proceed in two steps. First, we use the equilibrium model of bidding in a discriminatory auction described above to link the bids in the liquidity auctions to the implied willingness-to-pay for repo loans or, equivalently, to the opportunity funding costs. As we described in Cassola et al. (2013) this measure provides us with information on prices (interest rates) that a given bank would have to pay to secure liquidity from other sources on the interbank market. Since there likely are various important factors which affect the evolution of banks' funding costs from week to week, and which are unobserved to the econometrician, we will use an estimation method which uses data only from one auction at a time. In the second step we use the estimated funding costs to estimate the

network structure that is consistent with the correlation patterns. We then propose our measure of systemic risk, which is based on this estimated network structure.

4.1 Estimation of the Cost of Funding

Since this method is discussed in detail in Cassola et al. (2013), we describe it here only briefly. Equation (1) provides us with the link between the observable data (bids) and the variables of interest: banks' funding costs. This inversion of bids is a common approach in the empirical auction literature at least since Guerre, Perrigne and Vuong (2000). In order to invert bids using equation (1) we need to estimate the distribution of the market clearing price, P^c , which is bidder-specific, because it depends on the submitted bid.

To do that we employ the resampling method introduced in Hortaçsu and McAdams (2010) and further developed in Kastl (2011) and Hortaçsu and Kastl (2012). In order to perform this step we impose the assumption of independence and ex-ante (within group) symmetry among banks (within one of M groups). The resampling procedure allows us to simulate the distribution of market clearing price using bids submitted only within one particular auction. By repeatedly drawing with replacement $N_m - 1$ and N_{-m} bids from the observed sample, we can simulate a state of the world, a particular realization of the residual supply from the perspective of a bank from group m . This simulated residual supply intersected with the submitted bid delivers a particular realization of the market clearing price potentially faced by this bank. Repeating this procedure for each bank in each auction yields an empirical distribution of the market clearing prices and thus allows us to evaluate the probabilities in equation (1). In Cassola et al. (2013) we showed that with uncertainty about the available supply our estimator is consistent as the number of bidders within an auction goes to infinity.

We use this method to estimate the willingness-to-pay of each bank that participates in a given auction. Since most banks participate quite frequently (and the major ones participate virtually always), we thus obtain a time series of (opportunity) funding costs for every bank, which rationalizes its bids in the MRO auctions.

4.2 Estimating the Network of Spillovers using the Adaptive Elastic Net

In order to explicitly define what we mean by systemic risk, consider the following simple model of the dynamics of the cost of funding:

$$v_{i,t} = \beta_{0,i} + \sum_j \beta_{i,j} v_{j,t-1} + X_{t-1}^T \gamma_i + \epsilon_{i,t} \quad (2)$$

where $v_{i,t}$ is the cost of funding of bank i in period t , and X_t is a vector of controls intended to reflect aggregate “common” shocks (in practice, we use sovereign bond yields at different maturities). The cross-effect, $\beta_{i,j}$, captures the impact of a shock to j ’s cost of funding, such as a negative shock to its portfolio, on i ’s funding costs. A possible channel through which this effect might operate is the exposure of other banks to j ’s bonds, which are likely to lose value following a negative shock due to j ’s increased probability of default. Another possible mechanism linking banks’ funding costs are asset fire-sales, where the sale of a large quantity of assets by one bank in order to meet capital requirements leads to a devaluation of other banks’ asset portfolios, as in Greenwood et al. (2015). Assuming X_t adequately controls for common shocks, the $\beta_{i,j}$ ’s could be interpreted as the effect due to direct debt contracts between banks or fire-sale spillovers. In Section 5.4-5.6 and the Appendix A.2-A.3, we conduct several robustness exercises, including an instrumental variables strategy.

4.3 Estimation Procedure: The Adaptive Elastic Net

A first important challenge in interpreting a specification like (2) from a causal perspective is the presence of omitted sources of correlation between v_{t-1} and v_t (the vector of all banks’ funding costs at $t-1$ and t , respectively). We attempt to address this issue by adding a number of controls in X_t to account for common shocks. Coupled with the large number of banks in the system (who can all, in principle, be on the right hand side of the regression) this creates yet another challenge, that there are more parameters to estimate than the number of time series observations per bank. For the time being we will assume the vector X_t is given and will now discuss some technical aspects of the estimation procedure.

To deal with the “too many covariates” problem, we use the *elastic net* method of Zou and Hastie (2005), which is a mixture of a Ridge Regression with the Least Absolute Shrinkage and Selection Operator (LASSO) (Tibshirani 1996). For simplicity, consider the classical linear regression model $y = Z\theta + u$, where the sizes of y and Z are $S \times 1$ and $S \times N$, respectively. The adaptive elastic net estimator of this model can be defined as:

$$\hat{\theta}_{enet} = \left(1 + \frac{\lambda}{2}(1 - \alpha_e)\right) \left(\underset{\theta \in \Theta}{\operatorname{argmin}} \|y - Z\theta\|_2^2 + \lambda \left(\frac{(1 - \alpha_e)}{2} \|\theta\|_2^2 + \alpha_e \|\theta\|_1\right)\right) \quad (3)$$

where $\|\cdot\|_1$ and $\|\cdot\|_2$ denote the L^1 and L^2 norms, respectively. The term $(1 + \frac{\lambda}{2}(1 - \alpha_e))$ is a bias correction factor added by Zou and Hastie (2005) to lessen the downward bias due to double penalization.

For a given sample, the estimates $\hat{\theta}_{enet}$ are a function of the penalty parameters λ and α_e . The two extreme cases are LASSO ($\alpha_e = 1$) and Ridge regression ($\alpha_e = 0$). When the L^1 -penalty (LASSO) is sufficiently large, some of the coefficients will be exactly zero, selecting a subset of all possible regressors. This property makes the LASSO particularly suitable for model selection in applications where the researcher has strong reasons to believe that some of the variables in the regression do not belong to the true model generating the data. In our case, this is equivalent to assuming that the network of interbank links is sparse. Moreover, the LASSO can be modified to select the right model consistently, as shown by (Zou 2006). This can be achieved by employing the “adaptive LASSO” which includes adaptive weights to penalize each coefficient differently in the L^1 penalty. The adaptive LASSO also has the oracle property, i.e., it performs as well as an estimator that only includes the variables in the true underlying model, in terms of the convergence rate of the non-zero coefficients.

However, consistent model selection is not our only concern. It is a well known fact that the LASSO does not perform well when (some of) the RHS variables are highly correlated. This problem is of particular relevance to our application, since it is almost by definition of the financial network that the funding costs will be contemporaneously correlated. Zou and Hastie (2005) proposed the elastic net estimator in (3) precisely to deal with this issue.⁵ Unlike the basic LASSO, the elastic

⁵Zou and Hastie (2005) includes a Monte Carlo exercise illustrating the superiority of the elastic net estimator in

net can select groups of strongly correlated predictors, improving the prediction accuracy of the LASSO. Moreover, the elastic net can also be slightly modified to guarantee that the oracle property holds. In fact, the *adaptive elastic net* of Zou and Zhang (2009) combines the weighted L^1 penalty of the adaptive LASSO with the L^2 penalty included in elastic net and it inherits the advantages of both methods over the LASSO. It has the oracle property, as adaptive LASSO does, and handles highly correlated covariates, as elastic net does. The adaptive weights in the L^1 penalty are given by $\hat{w}_j = \left(\left| \hat{\theta}_{enet,j} \right| \right)^{-\delta}$, where δ is a positive constant. Given those weights, the adaptive elastic net estimates are:

$$\hat{\theta}_{enet} = \left(1 + \frac{\lambda}{2} (1 - \alpha_e) \right) \left(\underset{\theta \in \Theta}{\operatorname{argmin}} \|y - Z\theta\|_2^2 + \lambda \left(\frac{(1 - \alpha_e)}{2} \|\theta\|_2^2 + \alpha_e \sum_{j=1}^N \hat{w}_j |\theta_j| \right) \right) \quad (4)$$

where an infinite weight ($\hat{w}_j = \infty$) implies that the j -th variable is not selected ($\hat{\theta}_j = 0$). We choose the adaptive elastic net as our estimation procedure, because it has the desired properties just mentioned, but also because it allows for consistent estimation of the distributions of the corresponding estimators by means of the bootstrap, as shown by Chatterjee and Lahiri (2013).⁶

There are two penalty parameters in the adaptive elastic net: α_e and λ . To perform the estimation, we fix $\alpha_e = 0.5$ and then we use cross validation to find the optimal λ . We use the procedure just described to independently estimate (2) for each bank. Our main focus of interest is the resulting estimate of B , the matrix of direct spillover effects. Our measures of systemicness and vulnerability rely on such estimation. In Appendix A.2 we perform a robustness check to assess the sensitivity of the results to changes in α_e . For $\alpha_e \in \{0.1, 0.5, 0.9\}$ the results are all qualitatively very similar. Besides the sparse “prior” implicit when choosing a LASSO-type estimator, we restrict the estimates $\hat{\theta}_{enet}$ to be non-negative by letting $\Theta = \mathbb{R}_+^N$ in (3) and (4), since our focus are the negative externalities that result from an increase in the funding costs of individual banks. We have estimated the model without imposing the non-negativity constraints and only a few parameters take negative values.

the presence of correlated RHS variables.

⁶As a robustness check on the sensitivity of the results to variation in the individual weights in the adaptive Lasso penalty, we applied the elastic net to our model and compared the estimates to the ones obtained with the adaptive elastic net. The results regarding systemicness and vulnerability are reasonably close, although some differences do arise.

There is very active current literature in statistics and econometrics dealing with inference for these type of estimators. To obtain confidence intervals for our estimates, we apply a refinement of bootstrap proposed by Chatterjee and Lahiri (2013) described in section A.5 in the appendix.

4.4 Systemic Risk and Network Centrality

Equation (2) can be rewritten in matrix form as

$$v_t = \beta_0 + B^T v_{t-1} + \eta_t, \quad \eta_t = \Gamma X_t + \epsilon_t \quad (5)$$

By assumption, B represents the network of one period-ahead spillover effects between banks, as captured by the lagged correlations between their cost of funding. Borrowing some terminology from Network Analysis, B is (the matrix representation of) a weighted directed graph (a network), $B_{i,j}$ is the weight of the link from bank i to bank j , and $B_{i,j} = 0$ indicates that there is no such link. Our measures of systemicness and vulnerability are entirely based on the properties of the matrix B . Since they are closely related to centrality measures from the networks literature, we will briefly review the latter before introducing these two concepts.

For a given node i in a network \mathcal{B} , its ‘‘Katz centrality’’ is defined as:

$$k_i^c(a) = \sum_{s=1}^{\infty} \sum_{j=1}^N a^s (\mathcal{B}^s)_{i,j} \quad (6)$$

where the parameter a is positive and lower than the reciprocal of the largest eigenvalue of B , in absolute value, to guarantee convergence. Notice further that $k_i^c(a)$ is the i -th component of the vector:

$$k^c(a) = \left((I - a\mathcal{B})^{-1} - I \right) \mathbf{1} \quad (7)$$

where $\mathbf{1}$ is a vector of ones. Katz centrality calculates the relative importance of a node within the network by including the sum of its links (the number of direct neighbors in a binary graph) as well as the weighted sum of the links of its indirect neighbors (those to which the node is connected only through links with other nodes). If $a < 1$, the weight assigned to the indirect neighbors decreases the more distant they are from the node.

A clear benefit from our reduced form specification (2) is that it provides a straightforward interpretation of Katz centrality in terms of total effects of idiosyncratic shocks. Assuming stationarity, (5) can be iterated forward to obtain:

$$v_{t+s} = \sum_{\tau=1}^s (B^T)^\tau \beta_0 + (B^T)^\tau v_{t-1} + \sum_{\tau=0}^s (B^T)^{s-\tau} \eta_{t+\tau} \quad (8)$$

Let $\epsilon_{i,t} = 1$ and $\epsilon_{i,t+1+s} = \epsilon_{j,t+s} = 0$ for all j and all $s \geq 0$. Conditioning on the realized values of X_{t+s} , the response in the WTPs of all banks to this shock is given by:

$$\Delta v_{t+s} = \sum_{\tau=0}^s (B^T)^{s-\tau} \epsilon_{t+\tau} = (B^T)^s \epsilon_t \quad (9)$$

Now, if $K = \sum_{s=1}^{\infty} B^s = (I - B)^{-1} - I$, then $K_{i,j}$ can be interpreted as the long run effect of a permanent shock of 100bp to i 's funding costs on j 's funding costs (or $MA(\infty)$ representation of the specified AR(1) process). Moreover, the parameter a in (7) can be used to assign lower weights to more distant responses in the funding costs. We use $a = 0.9$ since at our estimated network matrix, the shocks then completely dissipate after two years, which seems like a reasonably conservative upper bound. We therefore use the Katz centrality measure to define the following notions of systemicness (as the average total effect of a shock to bank i 's funding costs on the funding costs of a bank in the network):

$$k_i^s(a) = \frac{1}{N} \sum_{s=1}^{\infty} \sum_{j=1}^N a^s (B^s)_{i,j} \quad (10)$$

and vulnerability (as the average total effect on bank i 's funding cost of a shock to the funding costs of a bank in the network):

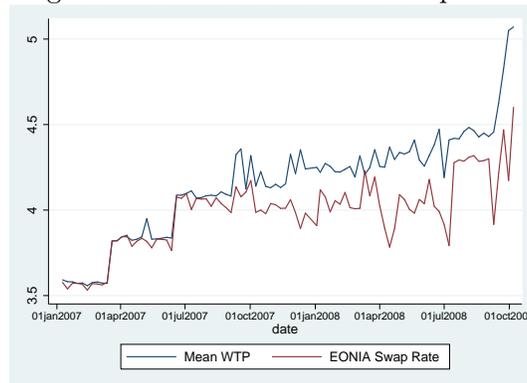
$$k_i^v(a) = \frac{1}{N} \sum_{s=1}^{\infty} \sum_{j=1}^N a^s (B^s)_{j,i} \quad (11)$$

where the difference between the two measures is whether we use the matrix B directly or its transpose.

4.5 Controlling for common shocks

As already mentioned, one important concern with the reduced form specification (5) is the existence of exogenous common components that might be driving the correlation between v_{t-1} and v_t . For instance, there could be a common trend in all the funding costs in the vector v_t , and then the recovered correlations between $v_{i,t-1}$ and $v_{j,t}$ would overestimate the actual spillover effect. To control for the expectations of the market rates, our main variable of interest is the spread between the funding costs and the EONIA swap rate. In fact, adjusting for EONIA accounts for a large share of the common trend in the $v_{i,t}$'s, specially before August 2007, as shown in figure 2. Without adding any other controls, a positive estimate of $\beta_{i,j}$ in (2) implies that a positive shock to bank j 's funding costs spread over EONIA is associated with an increase in the funding costs spread for bank i , one period ahead.

Figure 2: WTP and EONIA Swap Rate



As further controls for omitted common shocks, we also gathered data from Thomson Reuters' Datastream on government bonds yields at different maturities for ten European countries. In total, we have 30 series of the yield to maturity (YTM) of government bonds for the relevant time range (January 2007 - October 2008). Let $X_{c,m,t} = \log(R_{c,m,t-1}) - \log(R_{c,m,t-2})$, where $R_{c,m,t}$ is the (annualized) YTM of a bond from government c with maturity m at time t . The vector X_t in (2) is the collection of all 30 $X_{c,m,t}$'s. We include these lagged differences in government bond yields as regressors to control for possible sources of covariation of the funding costs.⁷ First, it

⁷The bond yields behave as random walks. In fact, the presence of a unit root is not rejected by a Dickey-Fuller test. Hence, we use the lagged differences as controls to be consistent with the assumption that (the EONIA adjusted) funding costs are stationary.

is highly likely that most banks in our sample hold portfolios of sovereign bonds, so they should be commonly exposed to variation in their prices, albeit to different degree. Second, the bond yields could be interpreted as proxies for general macroeconomic conditions in the country issuing the bond. In fact, such yields include risk premia that reflect market expectations on the state of the respective economies. Banks in the same country may be affected by country-level shocks that could be at least partially captured by the variation in bond yields. An important caveat is that prices might decrease due to firesales that result from liquidity or credit tightening conditions (as in Greenwood et al. (2015)). Moreover, banks might be using the bonds as collateral in the repo market, hence price volatility would directly imply an increase in their cost of funding as the quality of their collateral would deteriorate. This may trigger bond liquidations and deleveraging that would increase the price volatility and deteriorate the collateral even further. Thus controlling for bond yields, which are inversely related to their prices, could prevent such firesales from being captured, as a propagation channel, in the estimated residual comovements of the funding costs.

5 Results

5.1 Systemic Risk

Figure 3 illustrates the results of one specification of our model of the financial network. The i -th column corresponds to the regression of $v_{i,t}$ on the vector v_{t-1} . The lighter squares correspond to the variables that are selected by the LASSO component of the elastic net estimation of such regression. The lighter the color, the stronger the effect. More generally, the light squares can be interpreted as positive effects of one-period lagged funding costs of the row (y-axis) bank on the column (x-axis) bank's funding costs. It is evident that shocks to some banks affect a large number of their rivals. The number of lighter squares in a row corresponds to the number of (direct) links starting from that bank in the financial network. This number, the *degree* of the corresponding node in the binary counterpart of this matrix, is one measure of network centrality, but it ignores information on the strength of these links (their weight). The definition of degree can be easily generalized to include such information. We could compute the degree of a node as the sum of the weights of all the links that start from it. We prefer to use the Katz centrality measure that

not only takes the strength of the direct links into account, but also uses the information on the (indirect) paths captured by the powers of the matrix representing the directed weighted network.

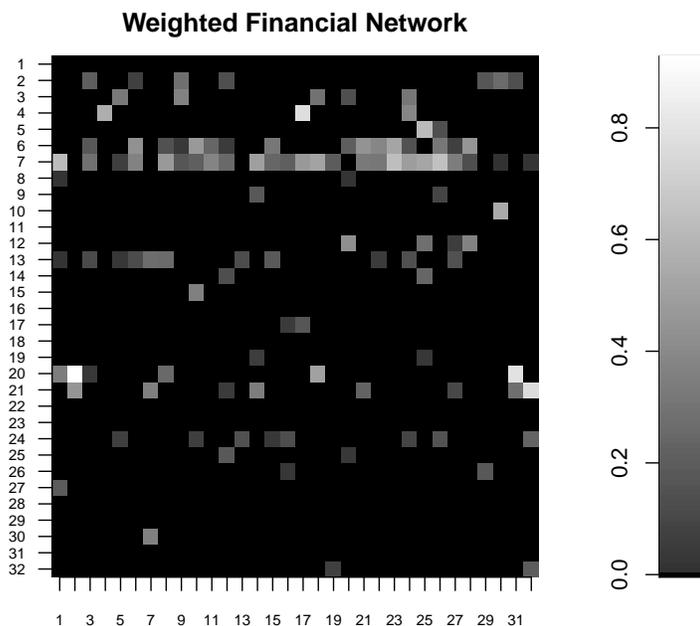


Figure 3: Elastic Net Coefficients: (1st order) Effect of a Shock to Row Bank’s (lagged) Funding Costs on Column Bank’s Funding Costs.

In Figure 4 we plot the estimated financial network. For confidentiality reasons, we cannot report the names of the banks. Each depicted node corresponds to a bank that is a member of the EURIBOR panel. The arrows between them correspond to positive coefficients in the respective adaptive elastic net regression. More precisely, there is an arrow from bank i to bank j if bank i ’s lagged cost of funding is correlated with bank j ’s cost of funding. Different colors correspond to different countries. Finally, the number in each node corresponds to the systemic rank of that node according to the Katz centrality measure (for $a = 0.9$). Even though the strength of the links is an important component in the centrality ranking, it is also the case that the top ranked banks have many links connecting them directly to other banks, as it is clearly the case for the two highest ranked banks.

Tables 4 and 5 offer a quantitative assessment of our estimates (see section A.5 for confidence intervals for these point estimates). In Table 4 we report the results from the adaptive elastic

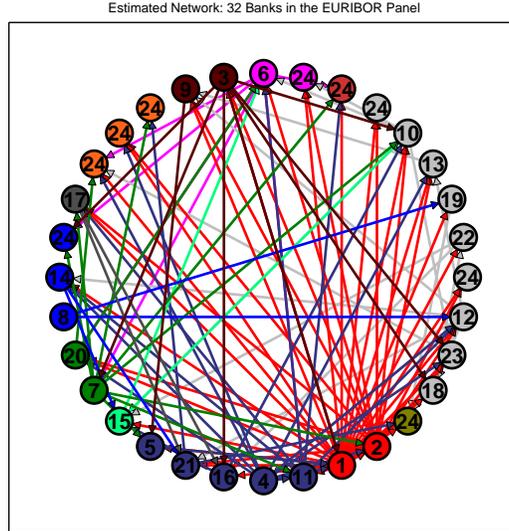


Figure 4: Financial Network

Table 4: First order effect of a 100bp shock to each of the top 5 systemic banks on all banks (Country means reported)

	AT	BE	DE	DK	ES	FR	GR	IE	IT	JP	NL	UK
Spanish Bank (ES)	0.0650	0.5135	0.3367	0.2185	0.1847	0.2886	0.5260	0.2342	0.0236	0.1723	0.3130	0.2730
Spanish Bank (ES)	0.3309	0.0000	0.1527	0.0000	0.2164	0.0989	0.0000	0.0751	0.0000	0.0397	0.1295	0.4565
British Bank (UK)	0.0000	0.0000	0.1025	0.0000	0.1754	0.0692	0.0000	0.2264	0.0921	0.0000	0.0000	0.1161
French Bank (FR)	0.0000	0.0000	0.0530	0.0000	0.1958	0.0469	0.0000	0.1323	0.0104	0.0000	0.0248	0.0000
French Bank (FR)	0.0000	0.0000	0.0000	0.0000	0.1788	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

net regression of the funding costs of the five most systemic banks, i.e., the coefficients of B from equation (2), where we aggregated these effects to the country level. While these coefficients still need to be interpreted with caution (due to omitted variables and other potential confounding factors), we believe that they offer an interesting perspective. For example, all five top ranked banks seem to impact the funding costs of Spanish banks, and one Spanish bank has a large effect on the funding costs of Belgian and Greek banks. If all top 5 banks suffered an increase in their funding costs of 100 basis points, the average funding costs of Spanish banks would rise right away by more than 95 basis points. By contrast, Italian banks do not seem to react right after the shock, the direct impacts are all below 10bp. However, this does not mean that they are not exposed to contagion, as shown in Table 5.

Table 5: Total effect of a 100bp shock to each of the top 5 systemic banks on all banks (Country means reported)

	AT	BE	DE	DK	ES	FR	GR	IE	IT	JP	NL	UK
ES1	0.8937	1.3736	0.9709	0.5674	0.8817	0.9537	1.5968	1.1110	0.4809	0.8688	0.8830	1.1407
ES2	1.1032	1.0339	0.8771	0.3494	1.0523	0.8919	1.0288	1.1240	0.5164	0.8336	0.7824	1.4665
UK	0.5315	0.7345	0.6390	0.2723	0.6793	0.6047	0.9229	0.8541	0.4016	0.6100	0.4762	0.7018
FR	0.4207	0.5801	0.4669	0.2368	0.6058	0.4688	0.6491	0.6228	0.2421	0.4080	0.4028	0.5633
FR	0.2874	0.4417	0.3122	0.1825	0.4443	0.3067	0.5135	0.3573	0.1547	0.2794	0.2840	0.3668

^a Averages across banks within a country reported.

^b The total effect is computed using the Katz centrality measure with $\alpha = 0.9$.

Table 5 reports the total effect, i.e., not only the first order effect reported in Table 4, but also the feedback from indirect effects and persistence captured by all the powers of B in the Katz centrality measure. Comparing the magnitudes across the two tables, the amplification of the shocks due to feedback seems quite high. The total effect of a 100bp shock to the funding costs of any of these top five systemic banks on banks in other countries is, in many cases, larger in magnitude than the original shock. This holds, in particular, for the two most systemic Spanish banks. Notice further than when the total effects are accounted for, the Italian banks are similarly exposed to contagion as banks in other regions. Still, Italian and Danish banks seem to be, on average, the less vulnerable to contagion from the most systemic banks.

We investigate the relative importance of intra-country effects further by highlighting the links in the network that connect two banks in the same country. Heuristically, by means of such comparison we want to test the hypothesis that a large share of systemic risk is due to the effect that the increase in a bank's cost of funding has on the costs of funding of other banks in its own country of origin. Figure 5 displays the binary directed graph (ignoring the weights), highlighting the intra-country links. Clearly, most of the links in the network connect banks in different countries. Moreover, on average the within country links account for no more than 9% of the average direct (first order) effect of a 100bp shock to the cost of funding of the banks in our sample. ⁸

In Figure 6 (left panel) we plot the histogram of our systemicness measure: the average total effect on a rival when a bank suffers a 100 basis point increase in its cost of funding. Our results

⁸In Appendix ?? we illustrate the implications of our estimated network for the financial linkages among countries. To do this, we use data from the Bank of International Settlements introduced in Glover and Richards-Shubik (2014). Unfortunately, these data only contain quite aggregated information so the exercise is not very informative.

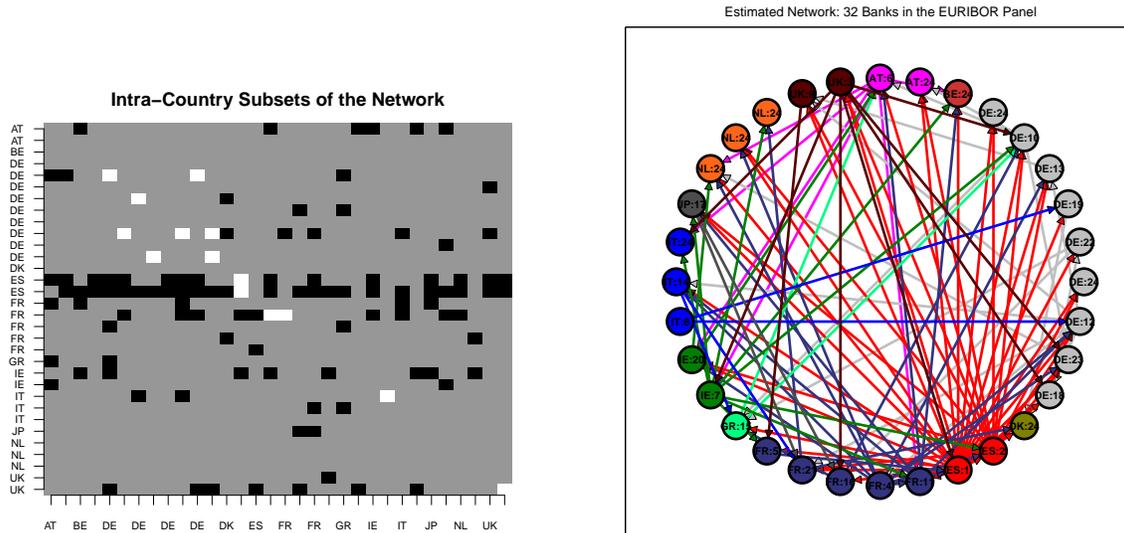


Figure 5: Relative importance of intra-country effects: In the left panel the links between two banks from the same country are highlighted in white. In the right panel, each number is the systemicness ranking of the respective node.

suggest a fair amount of heterogeneity of the impact of such a shock to a bank’s funding costs. While a shock to many banks in the EURO zone would not imply any substantial change in the funding costs of its rivals, there are a few important exceptions. In particular, there are two banks in the EURIBOR panel such that if they were to face a permanent 100 basis point increase in its funding costs, on average the funding costs of the other banks could increase by as much as approximately 90 basis points. We also plot, in the right panel, the histogram of the average total effect on each bank when a rival suffers a 100 basis point increase in its cost of funding. Clearly, the distribution is much less skewed, with most of the banks experiencing an average increase in their funding cost between 15 and 20 basis points.

5.2 Network Structure

If the weights on the links are ignored the network of spillover effects can be represented as a directed graph. That is, a link or edge from bank i to bank j represents a positive correlation between bank i ’s lagged funding costs and bank j ’s funding costs. Moreover, if the direction of the edges is also ignored, the network is simply an undirected graph, where a link between two

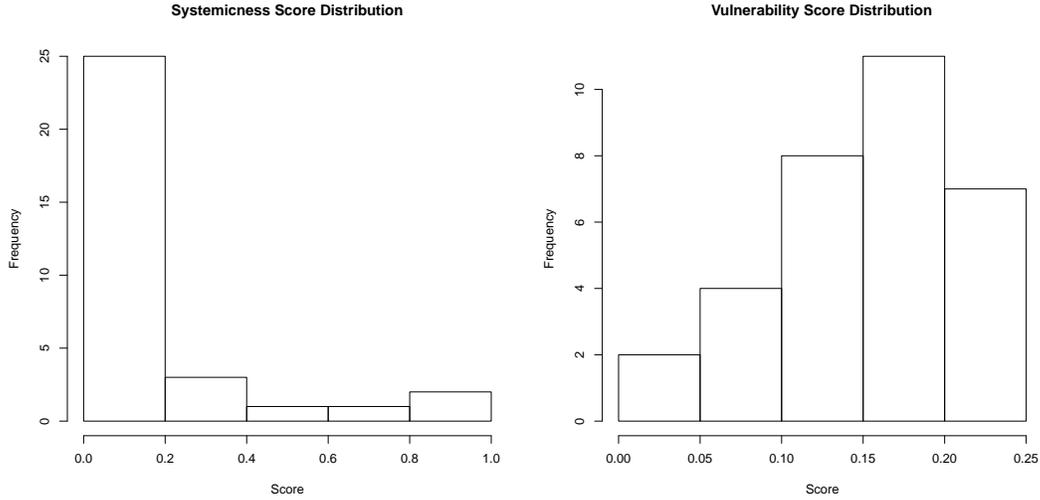
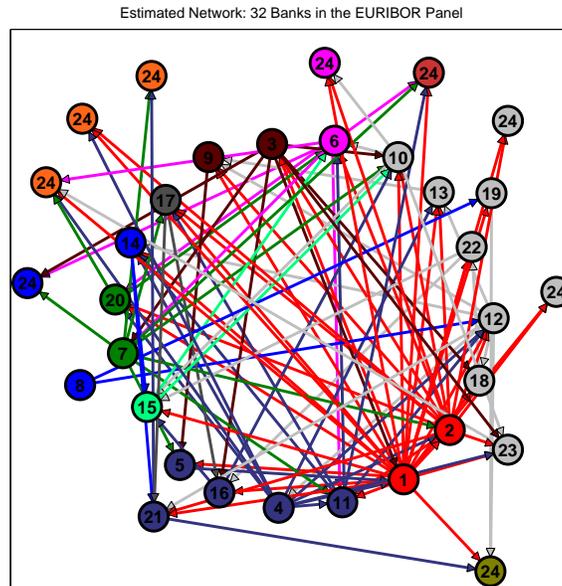


Figure 6: Katz Centrality Distribution. Left: Average total effect on a rival of a 100bp shock to the funding costs of each bank in the sample. Right: Average effect on the funding cost of each bank of a 100bp shock to the funding cost of a rival.

banks represents that either of them is connected to the other one. As can be seen directly by inspecting Figure 7, the undirected graph has only one (connected) component, meaning that there is no subset of banks completely isolated from the rest of the network. Figure 7 also illustrates other salient features of the network. The outer circle, formed by all banks labeled ‘24’ (the lowest position in the systemicness ranking), contains all the banks that have no outgoing edges, that is, there is no edge from these banks to any other bank in the network. However, there are several edges from other banks to all of the banks in the outer circle. Therefore, such banks are exposed to contagion if another bank in the network suffers a shock to its funding costs, but they do not impose any significant risk to any other bank (they do not relay the shocks further). If the main driver of correlation among bank’s funding costs were interbank lending, then such network structure would suggest that around a quarter of the banks in our sample behave as net lenders within the network, which exposes them to counterparty risk. In contrast, banks in the inner circle form a strongly connected component, which means that there is a path (sequence of outgoing edges) from each to any other of them. Furthermore, there is a path from all banks in the inner circle to all other banks in the network, except for one Italian bank. Thus, almost all banks are exposed to contagion

from this subset of roughly two thirds of the banks in our sample.

Figure 7: Network Structure. All banks are labeled according to their systemicness ranking and their color correspond to their country of origin as in figure 5. The outer circle (all banks labeled ‘24’) is formed by banks that have no outgoing edges. The inner circle is a strongly connected component, that is, there is a path (a sequence of outgoing edges) from all these banks to each other.



Similar results have been found both in the theoretical and the empirical literature. Farboodi (2014) develops a model of the interbank lending market with endogenous network formation. Her model predicts a ‘core-periphery’ network structure⁹, where a subset of the banks (the periphery) provides funds to another set of highly interconnected banks (the core) that act as intermediaries and/or invest those funds. Even though we are focusing on a sample of substantially large banks, all of which presumably have profitable investment opportunities, it is worth noting that a similar structure arises.

Bech and Atalay (2008) analyse the federal funds market as a network, where an edge from bank i to bank j means that j borrows fed funds overnight from i . They find that many of the largest banks in the US (in total assets) form a strongly connected component (around 10% of the

⁹For the reader familiar with graph theory, the ‘core-periphery’ terminology might be misleading. Neither term is being used here with its usual meaning in the context of graph theory. Presumably, these words were chosen in some of the papers mentioned in this section to classify banks in mutually exclusive groups according to their level of interconnectedness.

almost 500 banks in their sample). Also, most of the small banks (roughly 58% of all the sample in 2006) are net lenders to banks in the first group. Even though, according to our estimates around 60% of the banks in the Euribor panel belong to the strongly connected component, there is some resemblance between our results and those in Bech and Atalay (2008). Again, we are only considering a subset of the largest banks bidding for funds at the ECB (all of whom belong to the Euribor panel) so it is consistent with what these authors find that most banks in our sample are exposed to each other. What is surprising is that even among these large banks, we find some that are exposed to, but do not expose others to contagion.

Wetherilt et al. (2010) study the network structure of the unsecured overnight loan market in the UK during 2006-2008. Noteworthy, there are only a few (12 to 13) banks trading through CHAPS (the British Clearing House Automated Payment System) and many of them are also included in our sample. These authors also describe what they find as a ‘core-periphery’ network structure. Roughly speaking, they too define the core as a subset of highly interconnected banks. Moreover, they observe that the size of the core increased as the crisis unfolded. However, in contrast with our results and those in the literature previously mentioned, they find that banks in the periphery both borrow from, and lend to, banks in the core. Nevertheless, their definitions of ‘core’ and ‘periphery’ are based on the likelihood that there is an edge (an overnight loan) connecting two banks in each group and not on a partition of the set of nodes induced by strongly connected components. The takeaway is that, despite the differences in how these sets are defined, banks seem to naturally divide into two groups, depending on how tightly or loosely connected they are to each other, and such result seems to hold even when just a small number of large banks is considered.

Based on simulations, Elliot et al. (2014) make an interesting prediction about contagion effects in a network with a related ‘core-periphery’ structure. If the level of integration (roughly, how much of a bank’s assets is owed to its counterparties) among core banks is low, then the failure of any bank is not likely to generate wide-spread contagion. As integration rises to middle levels, the number of institutions that fail as a consequence of the initial failure also rises. Finally, for high levels of integration among core banks, the results depend on whether the bank that suffers the initial shock is in the core or in the periphery. If it is a core bank, the percentage of banks that fail as

a result is weakly higher than for middle-sized integration. However, if it belongs to the periphery, the effect is non-monotonic and the number of failing banks decreases. It is tempting to apply such results to our estimated network to try to predict how exposed to contagion the whole system is. However, there are some prominent differences between the network structures that prevents us from doing so. First, we identify a set of banks that have no outgoing edges (the outer circle in figure 7). Thus, if any of them were to fail (in our model that would be a shock to its funding cost beyond some healthy threshold), our estimates would predict no further failures, regardless of the level of integration. Second, in Elliot et al. (2014) the core of the simulated network is a clique, that is, a set of banks that are completely connected (no missing edges between them). In contrast, we estimate quite a sparse network, and even within the strongly connected component (the inner circle in figure 7) there are a lot of missing edges (as a matter of fact, only 20.5% of all possible edges in such set belong to the network).

5.3 Vulnerability and Bailouts

Eight out of 32 banks in our sample received targeted government support after August 2007 (most of them after October 2008, that is, out of sample). We run probit regressions of the probability of receiving such support (a bailout) on vulnerability (*Vuln_score*), Return on Assets (*ROA2008*), Return on Equity (*ROE2008*), Tier 1 Capital Ratio, a measure of leverage (*Debt to equity 2008*) and Average CDS spread of Dec. 2008 (*cds_dic*). The results are reported in Table 6. Notice that the coefficient on Vulnerability Score has the expected sign and is significant in all regressions. ROA and ROE are also significant, but *Tier_1_Ratio*, *Debt_to_Equity2008*, and CDS spread are not. Therefore, our measure of vulnerability does a better job at predicting bailouts than CDS spreads, Tier 1 Capital and leverage. Our findings regarding Tier 1 Ratio, ROA and ROE are consistent with analogous results recently reported by the European Systemic Risk Board in Alves, Ferrari, Franchini, Heam, Jurca, Langfield, Laviola, Liedorp, Snchez, Tavoraro and Vuillemeier (2013).

We compare our vulnerability score with analogous measures of exposure to contagion based on Euribor quotes, CDS spreads, market capitalization and CoVaR (Adrian and Brunnermeier (2011)), in terms of their ability to predict bailouts.¹⁰ As a benchmark, we use model (4) in Table 6, because

¹⁰In the rest of this Section we describe how these alternative measures are estimated.

	$\mathbb{I}(\text{Bailout})$				
	(1)	(2)	(3)	(4)	(5)
Vuln_score	0.487*	0.457*	0.522*	0.560*	0.541*
	(2.23)	(2.22)	(2.00)	(2.12)	(2.12)
ROA2008	-1.626*		-1.745*		
	(-2.68)		(-2.56)		
Debt_to_Equity2008	-0.402	-0.0904	-0.496	-0.312	
	(-0.41)	(-0.09)	(-0.50)	(-0.30)	
ROE2008		-0.0460*		-0.0540*	-0.0522*
		(-2.71)		(-2.71)	(-2.76)
Tier_1_Ratio			0.0926	0.132	0.127
			(0.48)	(0.70)	(0.67)
CDS			-19.78	-41.47	-37.61
			(-0.35)	(-0.76)	(-0.70)
Const.	-2.979*	-3.126*	-3.588*	-4.059*	-4.102*
	(-2.54)	(-2.67)	(-1.98)	(-2.18)	(-2.19)
N	32	32	32	32	32

t statistics in parentheses

+ $p < 0.10$, * $p < 0.05$

Table 6: Probit Models for the Probability of a Bailout. For model (4), the average marginal effect of a 1bp increase in the vulnerability score on the probability of a bailout is (an increase of) 3.7 percentage points, and the marginal effect for the average bank is 4.4 percentage points.

that is the specification of the probit that includes the most relevant set of controls. We report the results in Table 7, where each column corresponds to a different vulnerability measure included as a covariate (all the other controls are the same as in model (4) in Table 6). As opposed to our vulnerability score, we find that none of these alternative measures significantly predict bailouts at the 5% level. However, the measures based on market capitalization and CoVaR are still significant at the 10% level. In this last two cases, though, the regressions only include the subset of banks in our sample that are publicly traded.

Vulnerability	I(Bailout)			
	Euribor	CDS	Market Cap.	Covar
	0.462	0.121	0.958 ⁺	79.91 ⁺
	(0.474)	(0.724)	(0.507)	(48.33)
<i>N</i>	32	32	23	23

Standard errors in parentheses

⁺ $p < 0.10$, * $p < 0.05$

Table 7: Probit Models for the Probability of a Bailout Using Alternative Proxies for Funding Costs. All the columns correspond to model (4) in Table 6, when the vulnerability score is replaced by alternative measures of exposure to contagion that are not based in our estimates of Banks' funding costs.

5.4 Controlling for Common Shocks

In general terms, we found evidence that the scores and rankings are robust to the inclusion of changes in sovereign bond yields as controls in the estimation. Figure 8 shows the corresponding systemicness and vulnerability scores and rankings. Clearly, most of the changes in the estimates due to controlling for bond yields are minor. This is a direct consequence of the fact that only a few bonds, or in some cases none of them, are selected by the elastic net in the regression for each bank's funding costs. That is, our estimates suggest that the lagged funding costs of the counterparties are a stronger determinant of the funding costs of a bank than the returns on any portfolio of sovereign bonds. Still, the score of the most systemic bank is reduced by 21 basis points when the controls are included, which is not a negligible difference. That is, not including the proper controls would, at the very least, result in an overestimation of the risk that the most systemic bank represents to the system as a whole.

5.5 Alternative Measures of the Cost of Funding: Euribor Quotes and Bids

All the banks in our sample belong to the EURIBOR panel and thus they submit daily quotes of the rate at which they expect a prime bank to lend to another prime bank, through interbank term deposits, within the euro zone. They also bid in the MROs of the ECB for short-term maturity funds. Hence, both the Euribor quotes and the quantity weighted bids could in principle be used to estimate the network of spillover effects between banks, and the corresponding systemicness and

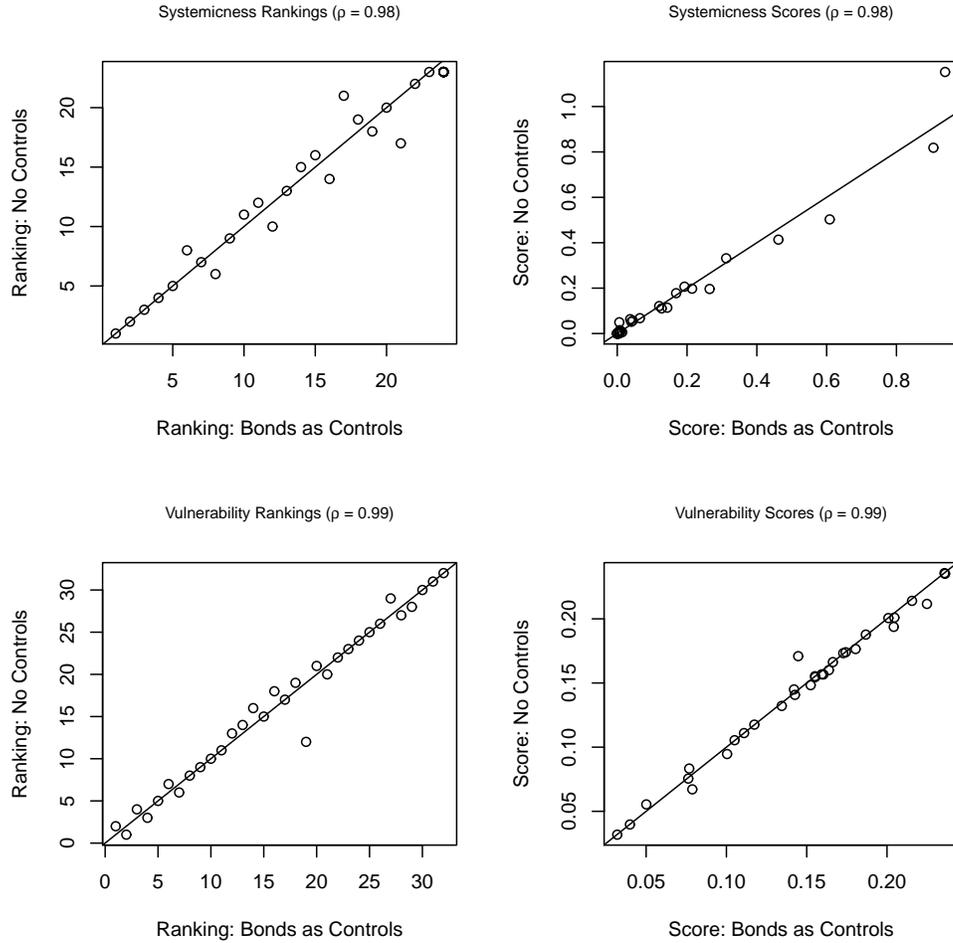


Figure 8: Effect on systemicness and vulnerability measures of controlling for government bonds (each point represents a bank in our sample).

vulnerability measures. As mentioned earlier, there are some shortcomings of using such alternative proxies for the cost of funding faced by banks in the interbank market. Regarding Euribor quotes, they are self-reported and, as recent scandals unveiled, are likely subject to strategic manipulation (see, for example, Barclays (2012)). Bids, as we argued earlier, also contain a strategic component. To assess the effect of this strategic deviations from truthfully revealing their cost of funding, we compare our systemicness and vulnerability measures (based on the inverted bids from the MROs) with the ones resulting from estimating network B in equation (5) using monthly averages of the Euribor quotes and the quantity average bids in the period from January, 2007 to October, 2008.

Figures 9 and 10 summarize the results of these comparisons. As expected, the Euribor quotes provide no information regarding the level of risk that individual banks impose to the system or about the degree to which they are exposed to contagion, as shown in Figure 9. The correlations between the rankings and scores based on the funding costs estimates and those based on Euribor quotes are effectively zero. Only one bank seems to be systemic according to its Euribor quotes, but it is actually one of the least systemic banks when measured by funding costs. Moreover, the matrix B in equation (5) is not stable, when estimated using Euribor quotes instead of funding costs estimated via bid inversion, i.e., it has eigenvalues outside the unit circle, violating the stationarity assumption.¹¹ Additionally, we note that within an auction, the Euribor quotes are correlated neither with the funding costs based on bids inversion, nor with the bids, suggesting that banks with higher costs of funding generally do not submit higher quotes. Not surprisingly, the vulnerability score computed using Euribor quotes instead of funding costs does not predict subsequent bailouts, as shown in Section 5.3. Bids seem to do a better job than Euribor quotes at predicting vulnerability, when compared to our measures, as shown in Figure 10. Both the vulnerability scores and rankings are positively correlated, and the correlation coefficients are both significant.¹² Three of the six most vulnerable banks according to both measures coincide. However, the systemicness scores and rankings are not significantly correlated.

5.6 Alternative Measures of Systemic Risk

As mentioned in the introduction, different measures of systemic risk have already been proposed in the literature. In this section, we compare our results to those obtained from two other measures that can be computed with our data set. These are the Granger causality test of Billio et al. (2012), and the Conditional Value at Risk (CoVaR) of Adrian and Brunnermeier (2011).

¹¹If the matrix B representing the network is not stable, the Katz centrality measures for systemicness and vulnerability are computed using an a lower than the reciprocal of the largest eigenvalue in absolute value, in order to guarantee convergence in (7).

¹²To determine if the correlation coefficient between two rankings is significant, we use the Spearman's rank correlation test.

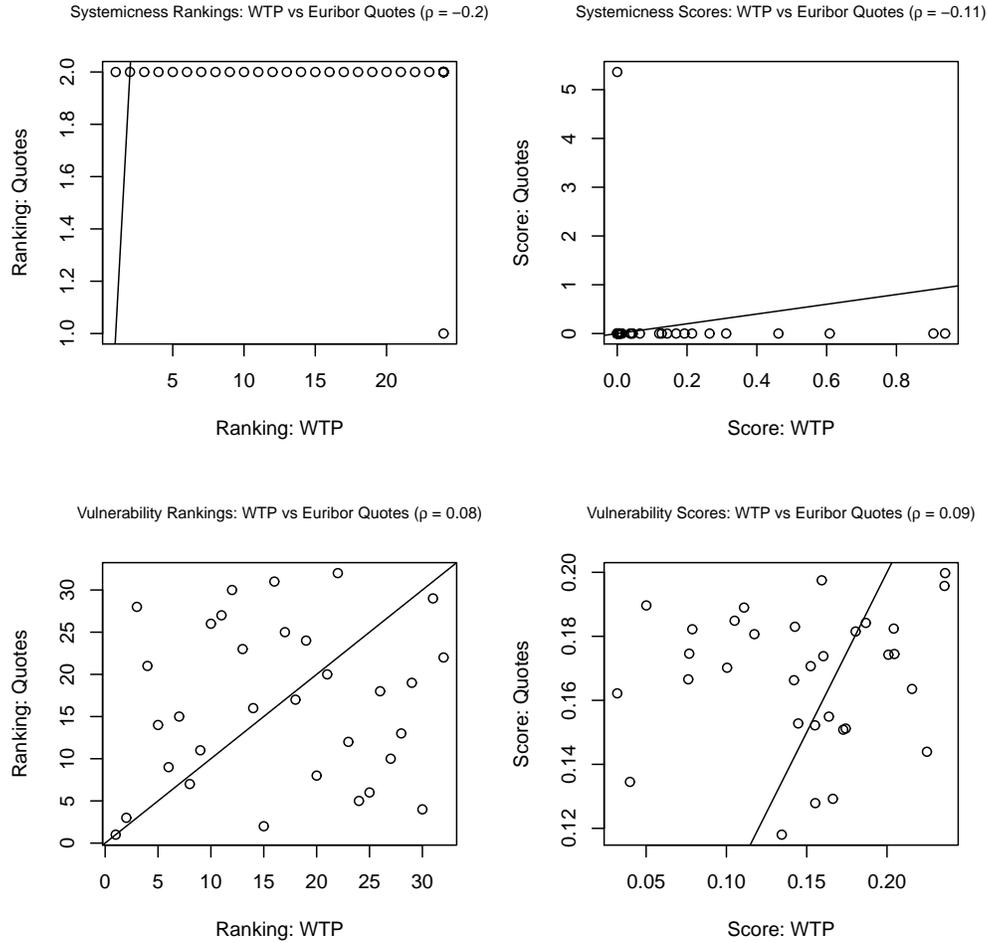


Figure 9: Systemicness and vulnerability rankings and scores: WTP vs Euribor Quotes (correlation coefficients in parenthesis)

5.6.1 Granger Causality: Stock Prices, WTP and Market Efficiency

An advantage of using the liquidity auctions data to recover the network of interbank spillover effects is that the (lagged) dependencies between the estimated funding costs are not ruled out by the efficiency of the market. On the contrary, the systemic risk measures that rely on the lagged correlation of market prices imply certain degree of market inefficiency. Billio et al. (2012) use Granger causality tests to determine if lagged asset returns on stock prices of a given financial institution j , R_{t-1}^j , significantly predict the returns of another institution, after controlling for autocorrelation and heteroskedasticity. More precisely, their linear causality tests are based on the

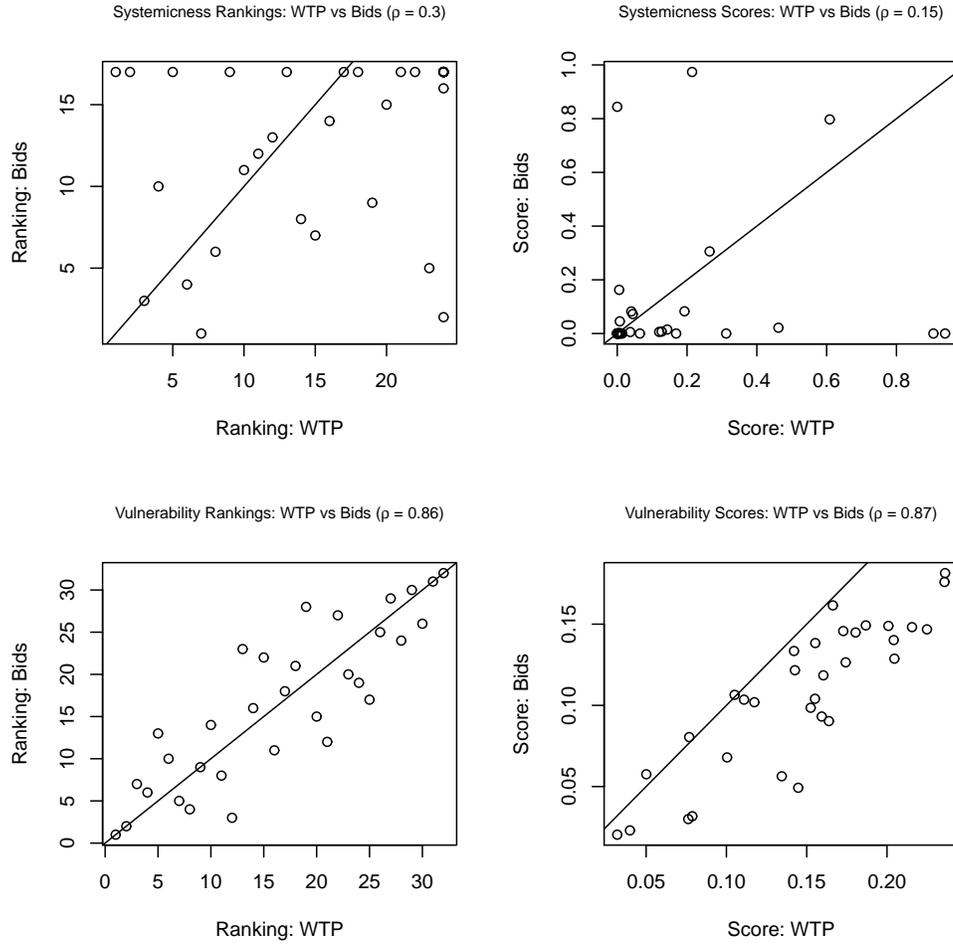


Figure 10: Systemicness and vulnerability rankings and scores: WTP vs Bids (correlation coefficients in parenthesis)

model:

$$R_t^i = a_{ij}R_{t-1}^i + b_{ij}R_{t-1}^j + e_{ij} \quad (12)$$

They say institution j 's returns Granger cause i 's returns if the estimate of b_{ij} is statistically different from zero at the 5% level of significance. To control for heteroskedasticity, they standardize the returns using the maximum likelihood estimates of the time-dependent variance implied by a GARCH(1,1) model. Clearly, market efficiency would imply that $b_{ij} = 0$. That is, if the market can accurately identify the spillover effects, then b_{ij} would be (very close to) zero, even if an

increase in i 's probability of default actually rises j 's future risk, because the current stock prices of both institutions would already reflect such effect. In other words, market efficiency could shade the connections between banks and the resulting estimated network would seem sparser than it actually is. To test this hypothesis, we compare the sparsity of the network of interbank connections when estimated using stock prices, credit default swaps, and Cassola et al. (2013) estimates of the willingness to pay. As reported below, we find that using stock returns seems to result in an underestimation of the number of connections within the network.

For a sample of 23 banks in the EURIBOR panel that are publicly traded, we define R_t^i as the monthly percent change in market capitalization. We use data on market capitalization from January 2007 to September 2008 to estimate the model described in (12) after controlling for heteroskedasticity. We arrange the resulting estimates in a matrix G whose element in the j -th row and i -th column, G_{ji} , is $\hat{b}_{ij} \neq 0$ if j Granger causes i , and 0 otherwise. Alternatively, we consider a VAR(1) process for the vector of returns R_t and we estimate each of the corresponding equations:

$$R_t^i = \sum_j \gamma_{ij} R_{t-1}^j + \eta_{ij} \quad (13)$$

using the same GARCH(1,1) model in Billio et al. (2012) to control for heteroskedasticity. Otherwise, the estimation method is the same described in section 4. Moreover, the term η_{ij} includes the same vector of controls included there. The resulting estimates $\hat{\gamma}_{ij}$ can be arranged in a matrix Γ , defined by $\Gamma_{ji} = \hat{\gamma}_{ij}$. Figure 11 displays the sign patterns of three matrices G , Γ and B , containing the estimated coefficients for the Granger causality test, the VAR(1) on returns and the VAR(1) on funding costs. As a measure of their sparsity, we report the number of interbank links in each estimated network as a fraction of the maximum possible number of links $N(N-1)$, where N is the number of banks in the sample (we do not count the diagonal elements). The number of nonzero connections (NNZ) are 12.5%, 10.1% and 15.4% of all possible off-diagonal connections for the Granger causality tests, the VAR(1) on Returns and the VAR(1) on funding costs, respectively. Using the Granger causality test, Billio et al. (2012) find that around 10% of all possible connections are contained in a network of the 25 largest banks in the US by market capitalization, using data from January 2006 to December 2008. Our estimates based on funding costs suggest a higher

degree of connectivity among large banks in the Euro Zone. In Section A.4 we conduct the same exercise using Credit Default Swaps rather than our estimated funding costs.

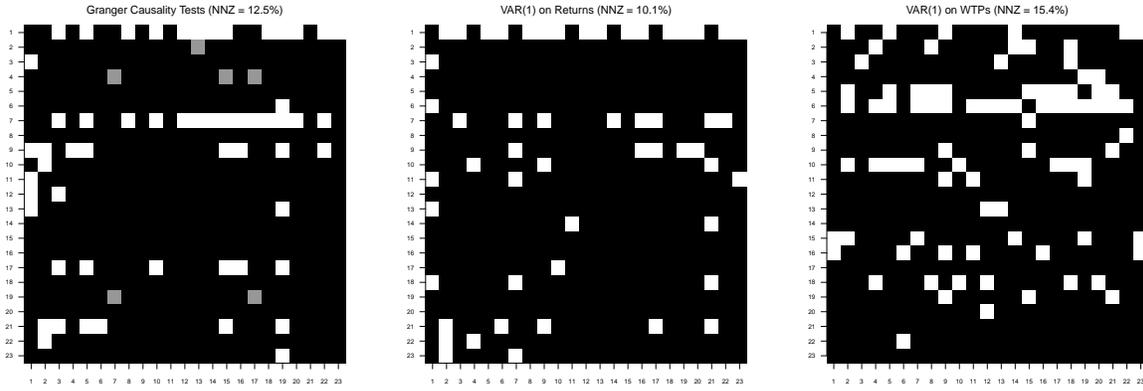


Figure 11: Sparsity of the network of interbank links estimated using data on market capitalization and costs of funding. A dot in the i -th row and j -th column represents an estimated connection from bank i to bank j . White and gray dots represent positive and negative coefficients, respectively. The number of nonzero entries (NNZ) is expressed as a percent of all possible off diagonal links

5.6.2 CoVaR

Adrian and Brunnermeier (2011) also provide a method for estimating the contribution to systemic risk of individual institutions within the system. Their measure is based on how the Value at Risk (VaR) of the whole system changes when conditioning on a specific institution being in distress compared to its value when that same institution is at its median state. Formally, let Z_i be a variable used to measure institution i 's financial soundness. In particular, Adrian and Brunnermeier (2011) use the percentage change in market-valued total assets. The value at risk VaR_i^q is the q -th quantile of the distribution of Z_i . Similarly, $CoVaR_q^{S|C(Z_i)}$ is the value at risk of the system conditional on some event $C(Z_i)$ relative to the level of distress of institution i . Adrian and Brunnermeier (2011) propose to measure the contribution of institution i to systemic risk using

$$\Delta CoVaR_q^{S|i} = CoVaR_q^{S|Z_i=VaR_i^q} - CoVaR_q^{S|Z_i=Median^i} \quad (14)$$

There is more than one way to compute $CoVaR_q^{S|C(Z_i)}$ described in their paper. For our comparison we used their time-invariant CoVaR because it does not require additional data. The methodology

is simple, $CoVaR_q^{S|C(Z_i)}$ is computed as a predicted value from a quantile regression of Z_S (the percentage change in market-valued total assets of the whole system) on Z_i . Correspondingly, $\Delta CoVaR_q^{S|i}$ is just the difference between two predicted values, calculated at different quantiles of the independent variable Z_i . $CoVaR_q^{i|C(Z_S)}$ and $\Delta CoVaR_q^{i|S}$ are similarly defined, to estimate i 's individual exposure to contagion from the system. Using this method, we obtain CoVaR for 23 banks in our sample that are publicly traded. We rank them according to their contribution to systemic risk and their exposure to contagion, implied by CoVaR. In order to compare our own estimates of systemic risk to CoVaR, we alternatively consider estimating the effect of an increase in a bank's lagged funding costs on the average funding costs of all banks in the sample and also, the effect of an increase in the lagged average funding costs of all banks on each bank's funding costs. We do this by estimating the equations

$$v_{i,t} = \bar{b}_{0,i} + \delta_i v_{i,t-1} + \bar{\beta}_i^v \bar{v}_{t-1} + X_t \bar{\gamma}_i + \bar{\epsilon}_{i,t} \quad (15)$$

where $\bar{v}_{t-1} = \frac{1}{N} \sum_j v_{j,t-1}$ is the lagged average WTP, and

$$\bar{v}_t = \bar{b}_0 + \sum_j \bar{\beta}_j^s v_{j,t-1} + X_t \bar{\gamma}_i + \bar{\epsilon}_{i,t} \quad (16)$$

(We use the same symbol for the coefficients on the X 's and the error term in both equations to save on notation). Correspondingly, we use $\bar{\beta}_i^s$ in equation 16 and $\bar{\beta}_i^v$ in equation 15 as alternative measures of bank i 's systemicness and vulnerability, respectively. Moreover, we also rank the banks according to these alternative scores and we compare the scores and the rankings with the ones implied by CoVaR.¹³ We find that the vulnerability rankings are weakly correlated ($\rho = 0.4$ with a p-value of 0.058), as shown in figure 12. However, we do not find any significant correlation between the respective systemicness measures. One possible reason for such discrepancy is that our measures and CoVaR are based on entirely different information about the banks. It is still comforting that both measures provide similar estimates of which banks are more exposed to contagion than others.

¹³We also compare the rankings based on the average funding costs with the ones implied by the Katz centrality measures based on the whole network of spillover effects. All the scores and rankings are positively and significantly correlated.

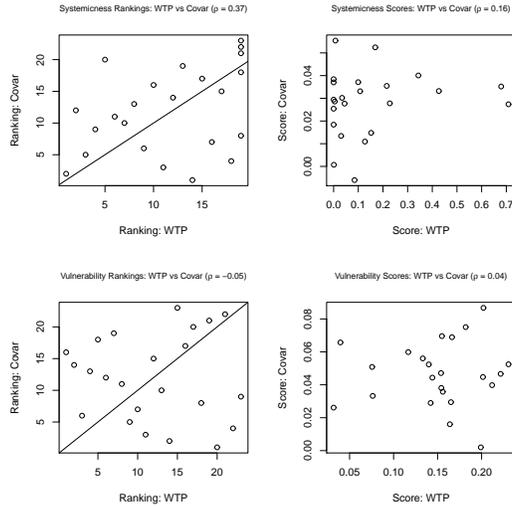


Figure 12: Vulnerability Measures Based on Average Funding Costs and CoVaR

6 Conclusion

In this paper we propose a novel measure of the systemic risk. We use the bidding data from the (weekly) liquidity auctions of the ECB and an equilibrium model of bidding to estimate banks' funding costs. Based on these funding costs we propose to use an elastic net estimator, which is a variant of a LASSO, to estimate the financial network. Our measure of systemicness, based on Katz centrality measure, is directly related to the importance of a bank in this network. Perhaps not surprisingly, our results suggest that there is large heterogeneity among banks. Most of the banks have fairly weak links and, therefore, if one were to suffer an adverse shock there would likely be a rather limited effect on the other ones. On the other hand, there are a few banks that are quite central: an increase in their funding costs would result in a very significant increase in the funding costs of other banks.

References

Açemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "Systemic Risk and Stability in Financial Networks," *American Economic Review*, 2015, 105 (2).

- Acharya, Viral, Lasse Pedersen, Thomas Philippon, and Matt Richardson**, “Measuring Systemic Risk,” 2010. working paper.
- Adrian, Tobias and Markus Brunnermeier**, “CoVaR,” 2011. working paper.
- Afonso, Gara and Ricardo Lagos**, “Trade Dynamics in the Market for Federal Funds,” February 2012. working paper.
- Allen, F. and D. Gale**, “Financial Contagion,” *Journal of Political Economy*, 2000, 108 (1), pp. 1–33.
- Alves, Ivan, Stijn Ferrari, Pietro Franchini, Jean-Cyprien Heam, Pavol Jurca, Sam Langfield, Sebastiano Laviola, Franka Liedorp, Antonio Snchez, Santiago Tavoraro, and Guillaume Vuillemeiy**, “The Structure and Resilience of the European Interbank Market,” 2013. Occasional Paper No. 3.
- Ashcraft, A., J. McAndrews, and D. Skeie**, “Precautionary Reserves and the Interbank Market,” *Journal of Money, Credit and Banking*, 2011, 43 (7), pp. 311–48.
- Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou**, “Who’s Who in Networks. Wanted: The Key Player,” *Econometrica*, 2006, 74 (5), pp. 1403–1417.
- Barclays**, “Statement of facts from the Justice Department,” <http://www.scribd.com/doc/98416026/Barclays-statement-of-facts-from-the-Justice-Department> 2012.
- Battiston, Stephano, M. Puliga, R. Kaushik, P. Tasca, and G. Caldarelli**, “DebtRank: Too Central to Fail? Financial Networks, the FED and Systemic Risk,” *Scientific Reports*, Aug 2012, 2.
- Bech, Morten L. and Enghin Atalay**, “The topology of the federal funds market,” 2008. Federal Reserve Bank of NY Staff Report number 354.

- Belloni, Alexandre, D. Chen, V. Chernozhukov, and C. Hansen**, “Sparse Models and Methods for Optimal Instruments With an Application to Eminent Domain,” *Econometrica*, 2012, *80* (6), 2369–2429.
- Billio, Monica, Mila Getmansky, Andrew W. Lo, and Loriana Pelizzon**, “Econometric measures of connectedness and systemic risk in the finance and insurance sectors,” *Journal of Financial Economics*, 2012, *104* (3), pp. 535–559.
- Boss, Michael, Helmut Elsinger, Martin Summer, and Stefan Thurner**, “Network topology of the interbank market,” *Quantitative Finance*, 2004, *4* (6), 677–684.
- Brownlees, Christian and Robert Engle**, “Volatility, Correlation and Tails for Systemic Risk Measurement,” 2010. working paper.
- Cassola, Nuno, Ali Hortaçsu, and Jakub Kastl**, “The 2007 Subprime Market Crisis in the EURO Area Through the Lens of ECB Repo Auctions,” *Econometrica*, 2013, *81* (4), pp. 1309–1345.
- Chatterjee, A. and S. N. Lahiri**, “Bootstrapping Lasso Estimators,” *Journal of the American Statistical Association*, 2011, *106* (494), 608–625.
- and — , “Rates of convergence of the Adaptive LASSO estimators to the Oracle distribution and higher order refinements by the bootstrap,” *The Annals of Statistics*, 6 2013, *41* (3), 1232–1259.
- Diamond, Peter A.**, “Aggregate demand management in search equilibrium,” *Journal of Political Economy*, 1982, *90*, pp.881–894.
- Diebold, Francis X. and Kamil Yilmaz**, “On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms,” Ko University-TUSIAD Economic Research Forum Working Papers 1124, Koc University-TUSIAD Economic Research Forum October 2011.
- EC**, “State aid: Overview of decisions and on-going in-depth investigations in the context of the financial crisis,”

<http://europa.eu/rapid/pressReleasesAction.do?reference=MEMO/11/616&format=HTML&aged=0&lang=en>
9 2011. European Commission MEMO/11/616.

Eisenberg, Larry and Thomas H. Noe, “Systemic Risk in Financial Systems,” *Management Science*, 2001, *47*, pp.236–249.

Elliot, Matt, Ben Golub, and Matthew Jackson, “Financial Networks and Contagion,” *American Economic Review*, 2014, *104* (10), pp.3115–53.

Farboodi, Maryam, “Intermediation and Voluntary Exposure to Counterparty Risk,” March 2014. working paper.

Freixas, X., Antoine Martin, and D. Skeie, “Bank Liquidity, Interbank Markets and Monetary Policy,” *Review of Financial Studies*, 2011, *24* (8), pp. 2656–92.

Freixas, Xavier, Bruno M. Parigi, and Jean-Charles Rochet, “Systemic Risk, Interbank Relations, and Liquidity Provision by the Central Bank,” *Journal of Money, Credit and Banking*, 2000, *32* (3), pp. 611–638.

Glasserman, Paul and H. Peyton Young, “How Likely Is Contagion in Financial Networks?,” *Journal of Banking and Finance*, January 2015, *50*, pp.383–399.

Glover, Brent and Seth Richards-Shubik, “Contagion in the European Sovereign Debt Crisis,” September 2014. working paper.

Greenwood, Robin, Augustin Landier, and David Thesmar, “Vulnerable Banks,” *Journal of Financial Economics*, 2015, *115* (3), 471 – 485.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong, “Optimal Nonparametric Estimation of First-Price Auctions,” *Econometrica*, 2000, *68* (3), pp. 525–574.

Hortaçsu, Ali and David McAdams, “Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market,” *Journal of Political Economy*, 2010, *118* (5), pp. 833–865.

- and **Jakub Kastl**, “Valuing Dealers’ Informational Advantage: A Study of Canadian Treasury Auctions,” *Econometrica*, 2012, 80 (6), pp.2511–2542.
- Iori, Giulia, Giulia De Masi, Ovidiu Vasile Precup, Giampaolo Gabbi, and Guido Caldarelli**, “A network analysis of the Italian overnight money market,” *Journal of Economic Dynamics and Control*, 2008, 32 (1), 259 – 278. Applications of statistical physics in economics and finance.
- Kastl, Jakub**, “Discrete Bids and Empirical Inference in Divisible Good Auctions,” *Review of Economic Studies*, 2011, 78, pp. 978–1014.
- , “On the Properties of Equilibria in Private Value Divisible Good Auctions with Constrained Bidding,” *Journal of Mathematical Economics*, 2012, 48 (6), pp. 339–352.
- Rochet, Jean-Charles and Jean Tirole**, “Interbank Lending and Systemic Risk,” *Journal of Money, Credit and Banking*, 1996, 28 (4).
- Swinkels, Jeroen**, “Efficiency of Large Private Value Auctions,” *Econometrica*, 2001, 69 (1), pp.37–68.
- Tibshirani, Robert**, “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 1996, 58 (1), pp.267–288.
- Wetherilt, A., P. Zimmerman, and K. Soramäki**, “The Sterling Unsecured Loan Market during 2006-2008: Insights from Network Theory,” 2010. Bank of England Working Paper 398.
- Wilson, Robert**, “Auctions of Shares,” *The Quarterly Journal of Economics*, 1979, 93 (4), pp. 675–689.
- Zou, H.**, “The Adaptive Lasso And Its Oracle Properties,” *Journal of the American Statistical Association*, 2006, 101 (476), pp.1418–1429.
- Zou, Hui and Hao Helen Zhang**, “On the adaptive elastic-net with a diverging number of parameters,” *The Annals of Statistics*, 8 2009, 37 (4), 1733–1751.

— and Trevor Hastie, “Regularization and Variable Selection via the Elastic Net,” *Journal of Royal Statistical Society B*, 2005, 67, pp.301–320.

A Appendix

A.1 Network Model

The setup of each market τ (monthly in our application) is as in Açemoglu et al. (2015), but with the additional bank heterogeneity due to the M classes. Apart from the N risk-neutral banks (index set \mathbb{I}), there is a representative outside financier indexed by 0, which could be the central bank (or lender of last resort). The heterogeneity of banks is captured by the joint distribution, $G_m(K, Z)$, of own capital, K_{it} , and short-term return shocks, Z_{it} , which will be described below. Each market (i.e., quarter) has three dates, $d \in \{0, 1, 2\}$. Date 0 corresponds to the beginning of a quarter, and dates 1 and 2 to the end. The major distinction between dates 1 and 2 is that debt contracts can only be collateralized up to returns realized at date 1 and part of the return on the initial investment (that cannot be pledged) only realizes at $d = 2$. We assume that all financial contracts are settled at $d = 1$. In $d = 0$ each bank is endowed with k_{it} units of own capital, and borrows funds to finance a portfolio (a project) that yields uncertain returns in the subsequent dates. We assume that k_{it} is publicly observed (this may be a reasonable assumption, since balance sheets are reported quarterly). As in Diamond (1982), a bank cannot use its own funds, but must borrow from other banks or the outside financier. Let b_{ijt} denote the amount lent by j to i in period t .

Once enough funds are raised, a bank invests and its portfolio generates a random short-term return of z_{it} in $d = 1$. If the portfolio is not liquidated, it also generates a known, but non-pledgable return A_m in period $d = 2$, where A_m is common to banks within a partition cell. Since this return is non-pledgeable, all debt has to be settled with funds available at $d = 1$, which makes the banks effectively risk-averse. Once the portfolio is invested into, the bank is assumed to also have to cover (potentially random) outside obligations of magnitude $v_t > 0$ at $d = 1$, which are senior to the debt to other banks or the outside financier.

Let R_{ijt} denote the interest on loan from the lender, bank j , to the borrower, bank i in period t .¹⁴ The corresponding bond has a face value $y_{ijt} = R_{ijt}b_{ijt}$. Total debt of bank i at $t = 1$ is therefore $\sum_{j \neq i} y_{ij} + v$. Let x_{ji} denote the amount of the loan that j repays to bank i . The probability of bank i defaulting is therefore equal to the probability that $k + z_i < \sum_{j \neq i, j \in \mathbb{I}} (y_{ij} - x_{ji}) + v$. In case the bank cannot meet its obligations, the project is liquidated and for simplicity we assume as in Açemoglu et al. (2015) that the liquidation destroys all of the non-pledgable return A_m . The endogenous financial network arises in this economy since the debt contracts define a weighted directed graph with n vertices (and size of the loans and rates can be endogenized with some more work, for example, using a search model).

We define the cost of funding of bank i at the beginning of the quarter as the quantity weighted interest that i pays on its loans: $c_{it} \equiv \frac{\sum_j R_{ijt}b_{ijt}}{\sum_j b_{ijt}}$. Part of the steps developed in the subsequent section will describe how this object can be recovered from observables. In particular, we will assume that c_{it} is the mean of a distribution, from which each bank will be drawing its weekly funding costs. An important assumption we will make is that conditional on all public information these (weekly) draws from that distribution are independent of one another. Note, however, that no assumptions are made on the variance or the support of this distribution and hence the independent draws assumption is not as restrictive as one might think.

A.1.1 Payment Equilibrium

We will use the notion of a payment equilibrium (Eisenberg and Noe (2001)). A vector of repayments constitutes a payment equilibrium if and only if given the initial loans the repayments are consistent with the realized shocks and with each other. Formally, in period t , bank i should pay to bank j the face value of any bonds previously sold to it maturing in t , in our notation: $y_{ij,t-1}$. If i 's cash on hand is insufficient, its payment to j , denoted x_{ijt} is strictly smaller and hence $x_{ijt} \in [0, y_{ij,t-1})$. The cash on hand of bank i in period t depends on the realization of the shocks, the repayments of debts owed to i by the other banks and potentially other dividends (if k_t follows some random process). Denoting this by α_{it} , the cash on hand satisfies: $\alpha_{it} = k_t + z_{it} - \sum_{j \neq i} (b_{jit} - x_{ijt})$,

¹⁴We take these rates as exogenous for the sake of simplicity, but it could be endogenized, for example, by employing a search model as in Afonso and Lagos (2012).

where $k - \sum_{j \neq i} b_{ijt-1}$ is the cash left after investing in other banks' portfolios, and $z_{it} + \sum_{j \neq i} x_{ijt}$ is the intermediate return and realization of the repayments of the loans previously issued. Assuming proportional repayments if the full liabilities cannot be met, the repayment by bank i to bank j can be written as:

$$x_{ijt} = \max \left\{ \frac{y_{ijt}}{\sum_j y_{ijt}} \left[\min \left\{ \sum_j y_{ijt}, \alpha_{it} - v_t \right\} \right], 0 \right\} \quad (\text{A-1})$$

Equation (A-1) states that if $\alpha_{it} - v_t > \sum_j y_{ijt}$, the bank i has enough cash left after paying off the outside creditors, v_t , to settle all of the remaining debt - and hence its payment to j is simply y_{ijt} . On the other hand, if the inequality is reversed, then whatever is left after satisfying the outside creditors, if anything, is allocated among the remaining creditors proportionally to the size of their loan to i . Since (A-1) can be viewed as a continuous mapping from a compact and convex subset of a Euclidian space to itself, it is straightforward to argue that it has a fixed point. Moreover, Açemoglu et al. (2015) prove the following nice result:

Proposition 1 (*Acemoglu et al., 2013*) *For any given financial network and any realization of the shocks, a payment equilibrium always exists and is generically unique.*

From an ex-ante perspective (before the beginning of the period $t = 0$ analyzed above) the stochastic process of the shocks (and possibly of own capital k_{it}) results in a distribution over possible payment equilibria. As pointed out above, the ability of bank i to repay its loan depends not only on its own shocks, but also on the realization of the shocks of all banks to whom i extended credit. Assume that the partition of banks into M reliability classes is known.¹⁵ The realizations of (K_{it}, Z_{it}) then govern the transitions of banks between the M groups according to some classification function $\Phi : \mathbb{I} \rightarrow \mathbb{M}$. In our application, the classification function Φ will be a set of $M - 1$ cutoffs specifying changes in the mean cost-of-funding during each quarter, which we will estimate using data from ECB's liquidity auctions as described in Cassola et al. (2013).

To empirically implement the observed heterogeneity, we split the banks into quartiles based on their Tier 1 ratio: the ratio of bank's equity to its risk-weighted assets. Our choice is motivated by the fact that the M groups should capture some (observable) ordering over the perceived riskiness

¹⁵It is sufficient that the number of banks in each class is commonly known, not necessarily their identities.

of each bank.

A.1.2 Auction Model

Q denotes the (random) amount of liquidity offered for sale by the central bank. We assume that the distribution of Q is common knowledge among the bidders. Each bidder receives a private signal, θ_i . The joint distribution of the signals will be denoted by $F(\theta_1, \dots, \theta_N)$. We assume (conditionally) independent private values (CIPV) paradigm. The marginal valuation function has the form $v_i(q, \theta_i)$.

Assumption 1 Bidder i 's signal θ_i is drawn from a common support $[0, 1]^M$, according to an atomless marginal d.f. $F_i(\theta_i|X)$, where X is the vector of available public signals.

Assumption 2 $v_i(q, \theta_i)$ is measurable and bounded, strictly increasing in (each component of) θ_i $\forall q$ and weakly decreasing in q $\forall \theta_i$.

$V_i(q, \theta_i)$ denotes the gross utility: $V_i(q, \theta_i) = \int_0^q v_i(u, \theta_i) du$. Bidders' pure strategies are mappings from private signals to bid functions: $\sigma_i : \Theta_i \rightarrow \mathcal{Y}$. Since in most divisible good auctions in practice, including the liquidity auctions of the ECB, the bidders' choice of bidding strategies is restricted to non-increasing step functions with an upper bound on the number of steps, $K = 10$, we impose the following assumption:

Assumption 3 Each player $i = 1, \dots, N$ has an action set:

$$A_i = \left\{ \begin{array}{l} (\vec{b}, \vec{q}, K_i) : \dim(\vec{b}) = \dim(\vec{q}) = K_i \in \{0, \dots, 10\}, \\ b_{ik} \in B = [0, \infty), q_{ik} \in [0, Q], b_{ik} > b_{ik+1}, q_{ik} < q_{ik+1} \end{array} \right\}$$

Therefore the set \mathcal{Y} includes all non-decreasing step functions with at most 10 steps, $y : \mathbb{R}_+ \rightarrow [0, Q]$, where $y_i(p) = \sum_{k=1}^K q_{ik} I(p \in (b_{ik+1}, b_{ik}])$ where I is an indicator function. A bid function for type θ_i specifies for each price p , how big a share $y_i(p|\theta_i)$ of the securities offered in the auction (type θ_i of) bidder i demands.

If multiple prices were candidates for a market clearing price, we assume consistently with our application that the auctioneer selects the highest price. In case of excess demand at the market clearing price, we assume that demands are rationed pro-rata on-the-margin.

The natural solution concept in this framework with private information is the Bayesian Nash Equilibrium. The expected utility of type θ_i of bidder i who employs a strategy $y_i(\cdot|\theta_i)$ in a discriminatory auction given that other bidders are using $\{y_j(\cdot|\cdot)\}_{j \neq i}$ can be written as:

$$EU_i(\theta_i) = E_{Q, \theta_{-i}|\theta_i} \left[\begin{array}{c} \int_0^{q_i^c(Q, \theta, \mathbf{y}(\cdot|\theta))} v_i(u, \theta_i) du \\ - \sum_{k=1}^K \mathbf{1}(q_i^c(Q, \theta, \mathbf{y}(\cdot|\theta)) > q_k) (q_k - q_{k-1}) b_k \\ - \sum_{k=1}^K \mathbf{1}(q_k \geq q_i^c(Q, \theta, \mathbf{y}(\cdot|\theta)) > q_{k-1}) (q_i^c(Q, \theta, \mathbf{y}(\cdot|\theta)) - q_{k-1}) b_k \end{array} \right] \quad (\text{A-2})$$

where $q_i^c(Q, \theta, \mathbf{y}(\cdot|\theta))$ is the (market clearing) quantity bidder i obtains if the state (bidders' private information and the supply quantity) is (θ, Q) and bidders bid according to strategies specified in the vector $\mathbf{y}(\cdot|\theta) = [y_1(\cdot|\theta_1), \dots, y_N(\cdot|\theta_N)]$, and similarly $P^c(Q, \mathbf{s}, \mathbf{y}(\cdot|\theta))$ will denote the market clearing price associated with state (θ, Q) , which turns out to be the random variable that is most crucial to the analysis. The first term in (A-2) is the gross utility the type θ_i enjoys from his allocation, the second term is the total payment for all units allocated at steps at which the type θ_i was not rationed and the final term is the payment for units allocated during rationing. A Bayesian Nash Equilibrium in this setting is thus a collection of functions such that almost every type θ_i of bidder i is choosing his bid function so as to maximize his expected utility: $y_i(\cdot|\theta_i) \in \arg \max EU_i(\theta_i)$ for a.e. θ_i and all bidders i . Part (i) of the following proposition proved in Kastl (2012) provides necessary conditions characterizing the equilibrium in discriminatory auctions with private values when marginal valuation function is continuous in q . Since continuity of the marginal valuation function might be questionable at the last step (in particular for bidders who submit just one step), we make use of the necessary conditions for optimality with respect to the bid (part (ii)).

Proposition 2 *Under assumptions 1-3 in any Bayesian Nash Equilibrium of a Discriminatory Auction, for almost all θ_i , with a bidder of type θ_i from group m submitting $K_i(\theta_i) \leq 10$ steps, every step k in the equilibrium bid function $y_i(\cdot|\theta_i)$ has to satisfy:*

(i) $\forall k < K_i(\theta_i)$ such that $v(q, \theta_i)$ is continuous in a neighborhood of q_k for a.e. θ_i :

$$v(q_k, \theta_i) = b_k + \frac{\Pr(b_{k+1} \geq P^c | m)}{\Pr(b_k > P^c > b_{k+1} | m)} (b_k - b_{k+1}) \quad (\text{A-3})$$

and at $K_i(\theta_i)$:

$$b_K = v(\bar{q}, \theta_i)$$

where $\bar{q} = \sup_{(Q, \theta_{-i})} q_i^c(Q, \theta, \mathbf{y}(\cdot | \theta))$, i.e., the largest quantity allocated to bidder i in equilibrium.

(ii) $\forall k \leq K_i(\theta_i)$ such that $v(q, \theta_i)$ is a step function in q at step k such that $v(q, \theta_i) = v_k \forall q \in (q_{k-1}, q_k]$ for a.e. θ_i and signals are independently distributed:

$$v_k = b_k + \frac{\Pr(b_k > P^c | m)}{\frac{\partial \Pr(b_k > P^c | m)}{\partial b_k}} \quad (\text{A-4})$$

The last result provides the basis for our first step: the estimation of the panel of funding costs.

A.2 Robustness

Recall that we opted for using the elastic net estimation method, because we suspect that using standard LASSO in our application would not perform as desired due to the likely correlation between our right-hand side variables. Since the elastic net has additional parameter, α_e , we have experimented with different values to verify sensitivity of our results with respect to its choice. Our results remain very stable. In Figure 13 we plot the ranks and scores, respectively, of each bank for two cases: $\alpha_e = 0.5$ (i.e. LASSO and Ridge penalties equally weighted) and $\alpha_e = 0.9$ (higher weight on the LASSO penalty). The correlation between the two sets of estimates is very high - above 0.95. Moreover, even the scores take very close values, except for one of the two most systemic banks for which there is 15bp change in the score when estimated using these two values of α_e . A salient effect of changing the value of α_e is that it affects the sparsity of the estimated matrix \hat{B} , given the tendency of the elastic net to select groups of highly correlated covariates. The lower the value of α_e , the lower the sparsity of the resulting matrix, since higher weight is assigned to the Ridge component of the elastic net penalty. Nevertheless, If banks are ordered according to their contribution to systemic risk, the resulting rankings are very similar for different values

of α_e , as shown in Figures 13 and 14. However, lower sparsity results in higher estimates of the vulnerability and the systemicness scores, for very low values of α_e . Interestingly, the change in the scores is negligible when α_e is increased from 0.5 to 0.9. Consequently, we interpret our results as moderately conservative estimates of the effects on the cost of funding of banks, of shocks to the cost of funding of their counterparties.

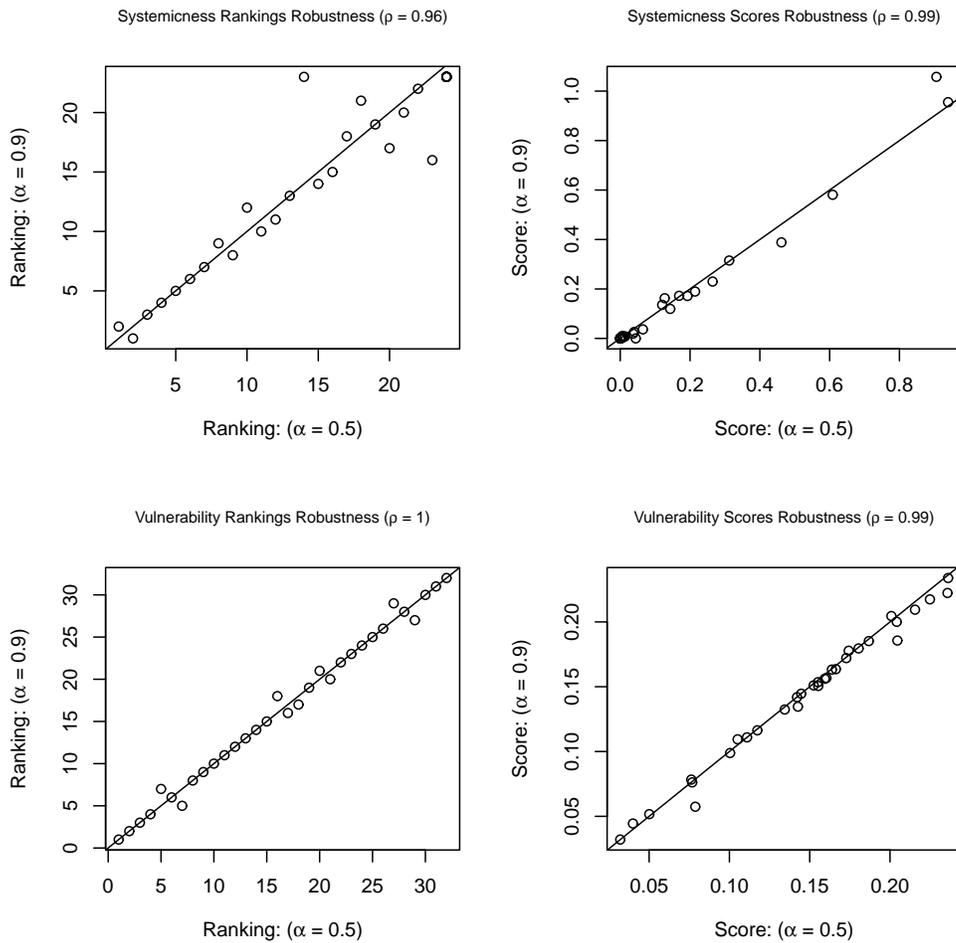


Figure 13: Robustness of systemicness and vulnerability rankings and scores to changes in α_e

A.3 Controlling for common shocks

The estimated matrix B of lag-correlations among banks' funding costs is meant to capture direct spillover effects between banks. However, as discussed in section 4.5, some of these correlations

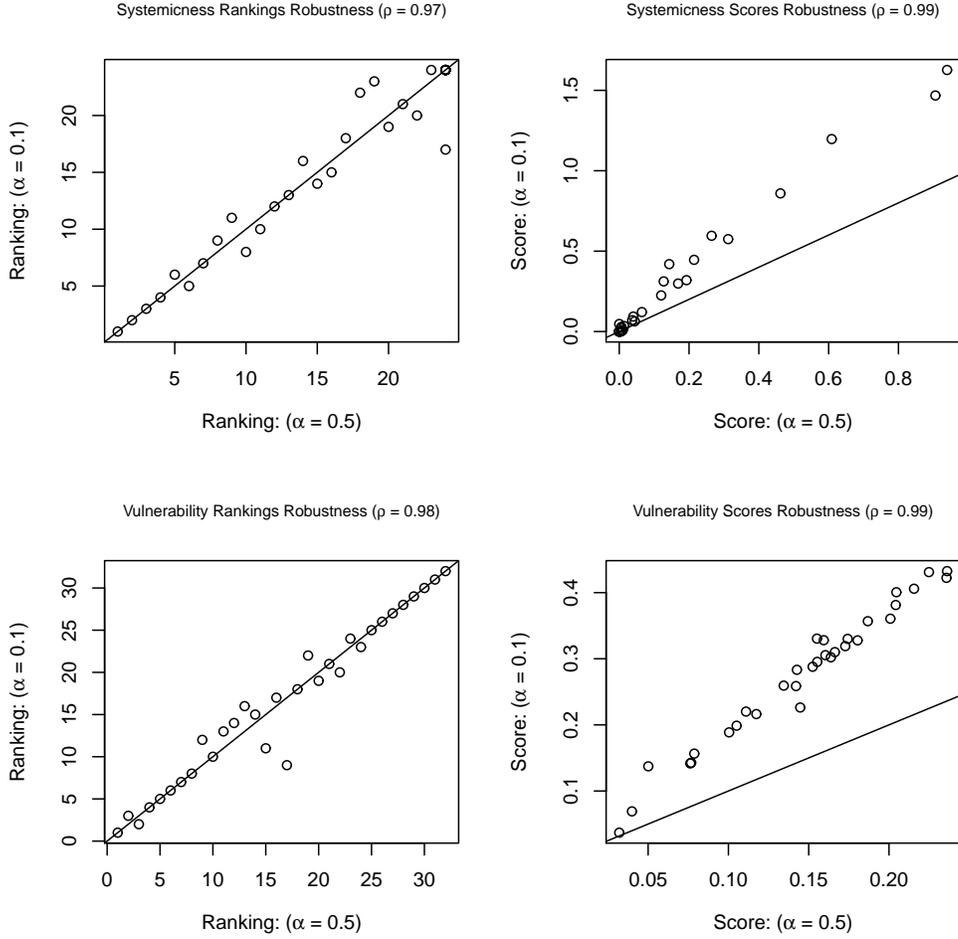


Figure 14: Robustness of systemicness and vulnerability rankings and scores to changes in α_e

might also be reflecting common shocks to the corresponding banks' funding costs.

Formally, we are concerned about the possibility that the error term in equation (2), $\epsilon_{i,t}$, is not orthogonal to some regressor $v_{j,t-1}$, mainly due to omitted variables that behave as unobserved common shocks to $v_{i,t}$ and $v_{j,t-1}$. Suppose z_t is the vector of these omitted variables. Assume for the sake of illustration that there is only one such variable. In order to model endogeneity, let us assume that the following is the true DGP. Suppose, for all i ,

$$v_{i,t} = \beta_{0,i} + \sum_j \beta_{i,j} v_{j,t-1} + \alpha_i z_t + \eta_{i,t}, \quad E[\eta_{i,t} | v_{j,t-1}, X_{t-1}, z_t] = 0$$

With, $\alpha_i \neq 0$. Notice that the source of endogeneity bias when running regression (2) would be serial correlation in z_t . That is, if $Cov(z_{t-1}, z_t) \neq 0$, since $Cov(z_{t-1}, v_{j,t-1}) \neq 0$ and $Cov(z_t, v_{i,t}) \neq 0$ (both by assumption), then, most likely, $Cov(z_{t-1}, v_{i,t}) \neq 0$. In other words, if the shock z_{t-1} is persistent, it could directly affect $v_{j,t-1}$ and $v_{i,t}$.

In order to address the endogeneity bias, one would need to isolate many instruments: exogenous shocks to each single bank's funding costs. Since there are very few bailouts within our sample period (and a bailout might be a valid instrument if one is willing to assume that it was unanticipated), this is a difficult task. One solution, albeit far from perfect, is to use the vector of sovereign bond returns X_{t-1} as instruments to control for common shocks. Intuitively, banks hold sovereign bonds in their portfolios, hence shocks to bond returns should translate into movements in their funding costs. For instance, negative returns could increase funding costs, either because a bank holding such assets might be perceived as more risky, might be willing to pay more for liquidity, or could be subject to higher haircuts to its collateral. Clearly, that is a testable assumption, since we have estimates of both funding costs and bond returns. Thus, let us assume for now that for all banks j there is at least some bond k such that $Cov(v_{j,t-1}, x_{k,t-1}) \neq 0$. Moreover, if we assume further than bond returns are unpredictable from previous returns, then $Cov(x_{l,t-1}, x_{k,t}) = 0$ for all l and k (where $x_{l,t}$ is the return of bond l at time t). Therefore, for all k , $x_{k,t-1}$ would be a valid instrument if $Cov(x_{k,t-1}, z_t) = 0$. Notice that for this condition to fail, z_t must be, to some degree, predictable using information available at $t - 1$, and also have an effect on $v_{i,t}$ that is not already reflected in v_{t-1} . Nevertheless, we acknowledge that we cannot rule out the possibility that some of the information reflected in government bonds might have persistent direct effects on banks' funding costs.

Another potential reason for correlation between $x_{k,t-1}$ and z_t would be the fact that a bank's exposure to sovereign bonds is an endogenous decision, that could be affected by the bank's funding costs and also by its exposure to other banks. If tighter conditions in the interbank market trigger liquidations that impact bond prices, then movements in bond returns would be exposed to the same common shocks that we are trying to control for. In this case, bond returns would also be endogenous, since $Cov(x_{k,t-1}, z_{t-1})$ would not be zero. Still, if other sources of correlation could

be ruled out, a sufficient (though, quite strong) condition for the instrument to be valid would be that banks behave as price takers in the market for sovereign bonds. It is worth noticing that each potentially endogenous regressor $v_{j,t-1}$ is being instrumented with a whole portfolio of bonds (a linear combination of all the bonds in the vector X_{t-1}). Moreover, we use a benchmark index for each country - maturity pair, rather than returns on specific securities. It is therefore more likely that banks' individual asset allocation decisions would have no significant effects on linear combinations of several of such indexes.

A.3.1 IV Estimation

With all the above caveats in place, we now proceed with this IV strategy. In order for the rank condition to be satisfied, we need at least as many instruments as potentially endogenous regressors. Therefore, we extended our original dataset to include 134 sovereign bond price benchmark indexes, provided by Thomson Reuters Datastream. To control for common shocks, we do a two stage IV estimation. In the first stage we project each bank's funding costs, $v_{j,t-1}$, on the vector of (a selected subset of) bond returns. In the second stage, we regress each bank's funding costs, $v_{i,t}$ on all the predicted values from the first stage.

More precisely, the first stage follows directly the procedure described by Belloni, Chen, Chernozhukov and Hansen (2012) to estimate optimal instruments in linear IV with many instruments (potentially more than observations). The conditional expectation of the endogenous regressors, given the vector of instruments, is estimated by running a LASSO regression of each $v_{j,t-1}$ on X_{t-1} to select the optimal instruments, and then computing the predicted values from an OLS regression on the selected instruments. We use data-driven penalty weights for each coefficient in the LASSO regression, as recommended by Belloni et al. (2012). The only modification that we introduce to the procedure described there is that instead of using their suggested value for the penalty parameter we use cross-validation to determine it. In the second stage, we regress each funding cost estimate at t on the vector of predicted values, using adaptive elastic net (as described in Section 4.3).

The results from the first stage suggest that the instruments are relevant. The rank condition always holds, and the average R^2 , across all 32 OLS regressions of each $v_{j,t-1}$ on the selected

instruments, is 0.55.

From the second stage, we obtain estimates of the network of spillover effects between banks, as well as the derived centrality measures and rankings. The vulnerability score and the corresponding ranking are close to those from our original estimation as shown in Figure 15, both correlation coefficients are above 0.85, but the scores are 6bp larger, on average, when no instruments are included. However, the systemicness scores and rankings obtained from both estimations are not significantly correlated. We conclude from this estimation that at least our results concerning vulnerability do not seem to be biased by the omission of possible sources of common variation in a bank’s funding costs and its counterparties lagged funding costs.

A.4 Granger Causality Test Using CDS

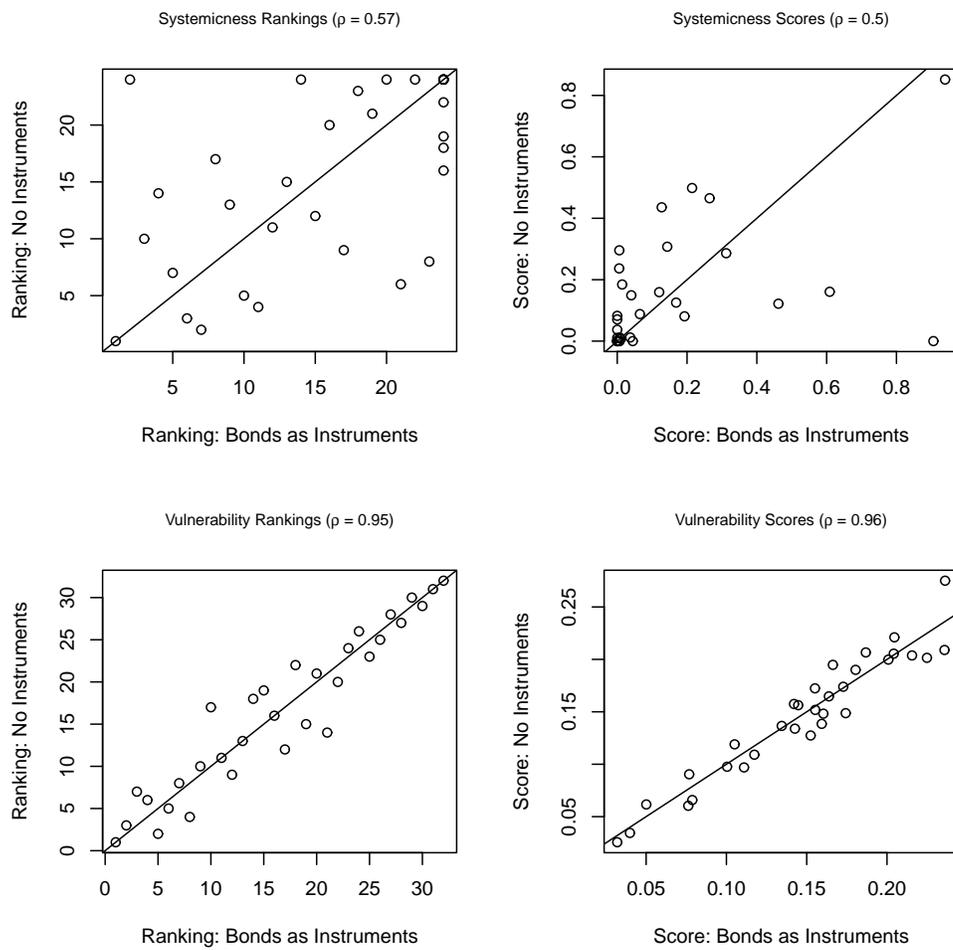
Instead of using market capitalization, as in Section 5.6.1, we repeat the procedure using Credit Default Swaps for all the banks in our sample. We define R_t^i in equations (12) and (13) as the monthly percent change in bank i ’s CDS spread, and we estimate these two models using the procedures described in Section 5.6.1. The number of connections in the estimated networks are 26%, 4.7% and 11% of all possible connections, for the Granger causality tests, the VAR(1) on CDS spread changes, and the VAR(1) on WTPs, respectively. Figure 16 shows the sign patterns of the matrices containing the estimated coefficients.

A.5 Inference for Adaptive Elastic-Net Estimates

Inference for Lasso-type estimators is an area of active research. In particular, the asymptotic distributions of the estimators have zero variance for the zero estimates and, thus, provide poor approximations to the finite-sample variances of such estimates.¹⁶ Very recently, Chatterjee and Lahiri (2011) have shown that the residual bootstrap consistently estimates the distribution of the Adaptive Lasso estimators, even when the number of variables included in the regression grows with the number of observations. Moreover, Chatterjee and Lahiri (2013) compute rates of convergence for the corresponding bootstrap estimates, and show that even for the non-zero estimates

¹⁶Zou and Zhang (2009) derive the asymptotic distribution of the adaptive Elastic-Net, for the case where both the number of observations and the number of variables in the regression go to infinity.

Figure 15: Effect on systemicness and vulnerability measures (Katz centrality) of including bonds as instruments (each point represents a bank in our sample).



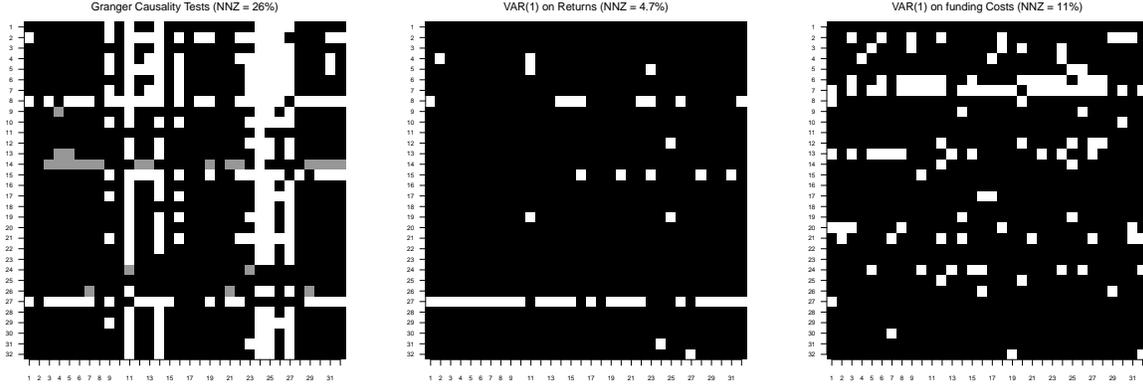


Figure 16: Sparsity of the network of interbank links estimated using data on CDS spreads percent changes and costs of funding. A dot in the i -th row and j -th column represents a estimated connection from bank i to bank j . White and gray dots represent positive and negative coefficients, respectively.

the bootstrap provides a more accurate approximation to the finite-sample distributions than the asymptotic distribution. Given that lack of consensus in the current statistics literature on the best method for post-selection inference, we use the residual bootstrap to obtain confidence intervals for our estimates, based on the results by Chatterjee and Lahiri (2013).

Recall, from equation (5), that $v_t = \theta_0 + B^T v_{t-1} + \Gamma X_{t-1} + \epsilon_t$. Using the procedure described in Section 4, we obtain adaptive Elastic-Net estimates of B , Γ and the covariance matrix of ϵ_t , denoted \hat{B}_{AE} , $\hat{\Gamma}_{AE}$ and $\hat{\Sigma}$, respectively. To perform the residual bootstrap, we take $v_0, X_0, X_1, \dots, X_{S-1}$ as given (from the data), and make 10,000 independent draws of ϵ_t from $N(0, \hat{\Sigma})$, with the k -th draw denoted by ϵ_t^k . Then, we plug in the initial data and each draw of ϵ_t into equation (5) to generate 10,000 samples, each consisting of a sequence (v_0, v_1, \dots, v_S) . Let $(v_0^k, v_1^k, \dots, v_S^k)$ denote the k -th generated sample and $(\theta_0^k, B^k, \Gamma^k)$ the corresponding vector of estimated coefficients obtained by applying adaptive Elastic-Net to the model:

$$v_t^k = \theta_0 + B v_{t-1}^k + \Gamma X_{t-1} + \epsilon_t^k \quad (\text{A-5})$$

Therefore, $(\theta_0^k, B^k, \Gamma^k)$ is an independent draw from the residual bootstrap distribution of the elastic net estimator of (θ_0, B, Γ) . In other words, each $(\theta_0^k, B^k, \Gamma^k)$ is a realization of the random variable, denoted below by $(\theta_0^*, B^*, \Gamma^*)$, which is the bootstrap version of $(\hat{\theta}_{0,AE}, \hat{B}_{AE}, \hat{\Gamma}_{AE})$.

Results for adaptive Elastic-Net

For simplicity, let θ be a vector containing all the parameters in (β_0, B, Γ) , and $\hat{\theta}_{AE}$ be the corresponding estimates obtained using the adaptive Elastic-Net. Also, let θ^* be the bootstrap version of $\hat{\theta}_{AE}$. Correspondingly, let

$$T_S = \sqrt{S} D_S \left(\hat{\theta}_{AE} - \theta \right) \quad (\text{A-6})$$

and

$$T_S^* = \sqrt{S} D_S \left(\theta^* - \hat{\theta}_{AE} \right) \quad (\text{A-7})$$

for a given matrix D_S . (Basically, the $q \times p$ matrix D_S is used to show that the consistency results also holds for vectors of q linear combinations of the parameters)

Theorem 4.1 in Chatterjee and Lahiri (2013) states that, under certain conditions

$$\sup_{B \in \mathcal{C}_q} |P_*(T_S^* \in B) - P(T_S \in B)| = O_p \left(S^{-\frac{1}{2}} \right) \quad (\text{A-8})$$

where \mathcal{C}_q is the collection of all convex measurable subsets of \mathbb{R}^q . A corollary of Theorem 4.1 is that the bootstrap consistently estimates confidence intervals for θ of level α .

We are not particularly interested in doing inference on θ . Instead, we would like to provide some measure of the degree of uncertainty for our estimates of the vulnerability, k^v , and systemicness, k^s . For each $i = 1, \dots, N$ (bank), we want to estimate the asymptotic distribution of the roots

$$R_{i,S}^v = \sqrt{S} \left(\hat{k}_{AE}^v - k_i^v \right) \quad (\text{A-9})$$

and

$$R_{i,S}^s = \sqrt{S} \left(\hat{k}_{AE}^s - k_i^s \right). \quad (\text{A-10})$$

We approximate such distributions with the bootstrap distribution of the roots

$$R_{i,S}^{*v} = \sqrt{S} \left(k_i^{v*} - \hat{k}_{AE}^v \right) \quad (\text{A-11})$$

and

$$R_{i,S}^{*s} = \sqrt{S} \left(k_i^{s*} - \hat{k}_{i,AE}^s \right). \quad (\text{A-12})$$

We use such approximations to build two-sided 90% confidence intervals for each k_i^{v*} and k_i^{s*} (See Table 8 for the results). The confidence intervals are built as follows. For a given level α (0.10), let $q_{\alpha/2}$ and $q_{1-\alpha/2}$ be the $\alpha/2$ and $1 - \alpha/2$ quantiles of the bootstrap distribution of $R_{i,S}^{*v}$. Let

$$CI_i^v = \left\{ k \in [0, \infty) : \hat{k}_{i,AE}^v - S^{-1/2} q_{1-\alpha/2} \leq k \leq \hat{k}_{i,AE}^v - S^{-1/2} q_{\alpha/2} \right\} \quad (\text{A-13})$$

then

$$P(k_i^v \in CI_i^v) \approx 1 - \alpha \quad (\text{A-14})$$

Similarly, we also built confidence intervals for the coefficient and the corresponding p -value on the vulnerability score in the bailout probits. Of course, the p -value is already a measure of uncertainty, but it does not take into account the variance in the estimation of the vulnerability score. It is based on the conditional distribution of the probit estimates, holding the values of the regressors constant. The results are reported in Table 9 below. Finally, Tables 10 and 11 show the 90% confidence intervals for the effects, aggregated at the country level, of a 100bp shock to the five most systemic banks, where Table 10 reports the 1st order effect and Tables 11 the one corresponding to the Katz centrality measure.

Table 8: 90% Confidence Intervals for the Vulnerability and Systemicness Scores

Vulnerability			Systemicness		
Lower	Point	Upper	Lower	Point	Upper
0.00	0.24	0.28	0.87	0.94	1.45
0.04	0.24	0.35	1.18	0.91	1.78
0.07	0.22	0.26	0.58	0.61	1.16
0.09	0.22	0.24	0.23	0.46	0.83
0.09	0.20	0.28	0.13	0.31	0.60
0.09	0.20	0.24	0.00	0.26	0.40
0.00	0.20	0.24	0.09	0.21	0.38
0.09	0.19	0.21	0.00	0.19	0.27
0.08	0.18	0.21	0.00	0.17	0.31
0.04	0.17	0.21	0.00	0.14	0.27
0.06	0.17	0.20	0.00	0.13	0.23
0.02	0.17	0.21	0.00	0.12	0.23
0.04	0.16	0.20	0.00	0.06	0.11
0.05	0.16	0.18	0.00	0.04	0.05
0.07	0.16	0.19	0.00	0.04	0.07
0.06	0.16	0.19	0.00	0.04	0.07
0.00	0.16	0.22	0.00	0.01	0.03
0.06	0.15	0.18	0.00	0.01	0.02
0.00	0.14	0.17	0.00	0.01	0.01
0.05	0.14	0.16	0.00	0.01	0.01
0.01	0.14	0.16	0.00	0.01	0.01
0.05	0.13	0.15	0.00	0.01	0.00
0.02	0.12	0.13	0.00	0.01	0.00
0.00	0.11	0.13	0.00	0.00	0.00
0.00	0.11	0.12	0.00	0.00	0.00
0.01	0.10	0.12	0.00	0.00	0.00
0.02	0.08	0.10	0.00	0.00	0.00
0.00	0.08	0.09	0.00	0.00	0.00
0.00	0.08	0.09	0.00	0.00	0.00
0.00	0.05	0.08	0.00	0.00	0.00
0.00	0.04	0.08	0.00	0.00	0.00
0.00	0.03	0.05	0.00	0.00	0.00

Table 9: 90% Confidence Intervals for the coefficient and the p-value of the vulnerability score in the bailout probit

	lower	point	upper
Coefficient	0.59	0.56	1.05
p_value	0	0.038	0.058

Table 10: 90% Confidence intervals for the first order effect of a 100bp shock to each of the top 5 systemic banks on all banks

	AT	BE	DE	DK	ES	FR	GR	IE	IT	JP	NL	UK
ES1	0.0650	0.5135	0.3367	0.2185	0.1847	0.2886	0.5260	0.2342	0.0236	0.1723	0.3130	0.2730
ES2	0.3309	0.0000	0.1527	0.0000	0.2164	0.0989	0.0000	0.0751	0.0000	0.0397	0.1295	0.4565
UK	0.0000	0.0000	0.1025	0.0000	0.1754	0.0692	0.0000	0.2264	0.0921	0.0000	0.0000	0.1161
FR1	0.0000	0.0000	0.0530	0.0000	0.1958	0.0469	0.0000	0.1323	0.0104	0.0000	0.0248	0.0000
FR2	0.0000	0.0000	0.0000	0.0000	0.1788	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

	AT	BE	DE	DK	ES	FR	GR	IE	IT	JP	NL	UK
ES1	[0,0.13]	[0.31,1.03]	[0.32,0.55]	[0.06,0.44]	[0.1,0.37]	[0.23,0.45]	[0.12,1.05]	[0.01,0.47]	[0,0.05]	[0,0.34]	[0.25,0.5]	[0.17,0.55]
ES2	[0.27,0.66]	[0,0]	[0.16,0.31]	[0,0]	[0.21,0.43]	[0.06,0.2]	[0,0]	[0,0.15]	[0,0]	[0,0.08]	[0.09,0.26]	[0.44,0.91]
UK	[0,0]	[0,0]	[0.04,0.2]	[0,0]	[0.08,0.35]	[0,0.14]	[0,0]	[0.07,0.45]	[0,0.18]	[0,0]	[0,0]	[0.04,0.23]
FR1	[0,0]	[0,0]	[0.0,11]	[0,0]	[0.07,0.39]	[0,0.09]	[0,0]	[0,0.26]	[0,0.02]	[0,0]	[0,0.05]	[0,0]
FR2	[0,0]	[0,0]	[0,0]	[0,0]	[0.07,0.36]	[0,0]	[0,0]	[0,0]	[0,0]	[0,0]	[0,0]	[0,0]

Table 11: 90% Confidence intervals for the total effect of a 100bp shock to each of the top 5 systemic banks on all banks

	AT	BE	DE	DK	ES	FR	GR	IE	IT	JP	NL	UK
ES1	0.8937	1.3736	0.9709	0.5674	0.8817	0.9537	1.5968	1.1110	0.4809	0.8688	0.8830	1.1407
ES2	1.1032	1.0339	0.8771	0.3494	1.0523	0.8919	1.0288	1.1240	0.5164	0.8336	0.7824	1.4665
UK	0.5315	0.7345	0.6390	0.2723	0.6793	0.6047	0.9229	0.8541	0.4016	0.6100	0.4762	0.7018
FR	0.4207	0.5801	0.4669	0.2368	0.6058	0.4688	0.6491	0.6228	0.2421	0.4080	0.4028	0.5633
FR	0.2874	0.4417	0.3122	0.1825	0.4443	0.3067	0.5135	0.3573	0.1547	0.2794	0.2840	0.3668

	AT	BE	DE	DK	ES	FR	GR	IE	IT	JP	NL	UK
ES1	[0.82,1.47]	[1.25,2.25]	[0.82,1.5]	[0.3,0.98]	[0.78,1.42]	[0.86,1.47]	[1.27,2.9]	[0.77,1.85]	[0.25,0.75]	[0.73,1.45]	[0.85,1.37]	[1.04,1.92]
ES2	[1.44,2.18]	[1.29,2.05]	[1.09,1.72]	[0.28,0.69]	[1.41,2.08]	[1.17,1.76]	[1.05,2.05]	[1.36,2.22]	[0.56,1.02]	[1.08,1.65]	[1.04,1.55]	[1.97,2.91]
UK	[0.49,1.02]	[0.66,1.42]	[0.58,1.22]	[0.09,0.52]	[0.6,1.32]	[0.56,1.15]	[0.81,1.82]	[0.72,1.65]	[0.32,0.77]	[0.6,1.18]	[0.44,0.92]	[0.69,1.36]
FR1	[0.2,0.77]	[0.28,1.08]	[0.19,0.84]	[0,0.43]	[0.36,1.14]	[0.26,0.85]	[0.19,1.24]	[0.21,1.15]	[0,0.43]	[0.18,0.75]	[0.23,0.74]	[0.31,1.04]
FR2	[0.12,0.56]	[0.22,0.86]	[0.1,0.6]	[0,0.36]	[0.24,0.87]	[0.13,0.59]	[0.19,1.02]	[0.03,0.69]	[0,0.3]	[0.09,0.54]	[0.15,0.55]	[0.16,0.71]