An Empirical Analysis of Systemic Risk in the EURO-zone *

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

First version: March 2013

This version: October 2014

Abstract

We propose a new measure of systemic risk, which is based on estimating 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 elastic net (a LASSO type estimator) to this panel to estimate the financial network. Finally, using the estimated network we propose a new measure of 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 140 basis points per 100 basis points shock) in the funding costs of the other banks.

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

JEL Classification: D44, E58, G01

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

*We thank Aureo de Paula, Thomas Phillipon, Marzena Rostek, José-Luis Peydró, David Skeie and seminar participants at Berkeley, EUI, NYU, Princeton, Stanford, UCLA, UCL, 2013 GSE Summer Forum, 2013 Wharton Conference on Liquidity and Financial Crises and 2014 NBER Summer Institute for helpful comments. Hortaçsu acknowledges financial support from the NSF (SES-1124073). Kastl acknowledges financial support from the NSF (SES-1123314) 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

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) and Açemoglu, Ozdaglar and Tahbaz-Salehi (forthcoming) 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. The chief culprit underlying this delay is the lack of direct data on interbank credit relations, save for information available through intraday payment and settlement systems (Bech and Atalay (2008) or 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 credit network 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. (forthcoming). Once we characterize the credit network, we use centrality measures typically used in network analysis to measure the extent to which a particular bank is crucial for the working of the whole system, i.e., for determining the set of systemic banks. Our method provides a measure of this "systemicness" as expressed by that bank's impact on the future cost of funding of other banks.

While different proxies for funding costs of banks may be utilized (CDS spreads, LIBOR/EURIBOR quotes), we focus on a particular measure that is (i) a *revealed by actions that are directly affecting payoff* as opposed to stated willingness-to-pay for credit as elicited in a survey, and (ii) available for a large set of banks. 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 willingness-to-pay 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 willingness-to-pay 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. The paper showed that the estimates of the changes in banks' willingness-to-pay 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 essentially running regressions of a given bank's funding cost on other system banks' (lagged) funding costs, controlling for various sources of comovement due to common asset exposures, such as holdings of sovereign bonds. This, 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 *elastic net* method of 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.

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 the Katz centrality measure, 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 140 basis points per 100 basis points shock) in the funding costs of the other banks.

As discussed by Greenwood, Landier and Thesmar (2011), a "vulnerability" measure similar to the "systemicity" measure can 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) and whether they were subsequently bailed out are (albeit weakly) correlated.

Several other empirical measures of systemicity (and vulnerability) 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. 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 Billio

[1]While 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. Moreover, the "statement of facts" makes explicit reference to the belief held by some Barclay's employees that other Contributor Panel banks where also submitting LIBOR quotes below their borrowing costs.

et al.'s estimates in Section 5.3.1.

Several other empirical measures of systemic risk have been proposed in the financial economics literature. Acharya, Pedersen, Philippon and Richardson (2010) provide a simple theoretical framework, in which they propose to measure the systemic risk by an institution's marginal contribution to the shortfall of capital in the financial system that can be expected in a crisis. Brownlees and Engle (2010) instead propose to measure the systemic risk by the expected shortage of capital of an institution given its degree of leverage. Adrian and Brunnermeier (2011) propose a way to empirically assess the degree of systemic risk associated with a bank, which relies on marginal contribution of each bank to the "value at risk" of the whole system conditional on that bank being under distress relative to its median state (hence the name CoVAR).[2] Aït-Sahalia, Cacho-Diaz and Laeven (2011) offer a model of financial contagion and propose an estimation method.[3]

Our approach is in some ways very related to this prior literature, especially Adrian and Brunnermeier (2011), in the following sense: Since there are a lot of players in the financial system, estimating bilateral relationships between them, controlling for developments in other players' financial "states" , is an estimation problem with too many covariates, and not enough data on the dependent variables. Given this, the prior literature has aggregated the independent variables in this "regression" framework (either by a weighted mean, or by principal components) and has focused on the comovement between individual banks' financial "state" and the "market state." Here, we utilize shrinkage methods to try to overcome this dimensionality problem, which helps us shed more light into the bilateral connectivity structure of banks.[4] The remainder of the paper proceeds as follows. We begin by laying out a formal 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 we conclude in section 6.

[2]Schwarcz (2011) offers an interesting perspective on systemic risk and its measurement from the point of the law.

[3]Since our approach exploits the covariances between the funding costs, it is related to the empirical literature on contagion (Dungey, Fry, González-Hermosillo and Martin (2006) and many others).

[4]Indeed, the ridge regression method, which is a component of the elastic net method we use, can be interpreted as over-weighting the most important principal components of the design matrix in a regression.

2 Model

We model the European banking system (consisting of n risk-neutral banks) as a sequence of quarterly markets.[5] We describe the market in detail below, but first we begin by a quick overview. When the market opens (every quarter) the banks are partitioned into M classes that are determined by their balance sheet structure (access to deposits and other liquid assets, availability of high quality collateral etc.) and availability of investment projects (their return and risk). This partition is common knowledge among the market participants. The banks interact among themselves and the relationships they enter into (loans extended etc.) give rise to a financial network. At the end of the quarter, the agreements (loans) get settled and the market re-opens. We will describe these agreements in detail below. In between the two quarterly markets, the banks finance their operations (and/or investment objectives) using liquidity that is either obtained by issuing debt in the quarterly market or obtained when needed on the interbank market or by utilizing one of the channels offered by the central bank, such as the repo-auctions.

The setup of each quarterly market is as in Açemoglu et al. (forthcoming), 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.

[5] Balance sheets are published quarterly for most financial institutions.

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.[6] 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. (forthcoming) 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.

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.

2.1 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

[6] We 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).

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_{ijt-1}. If i's cash on hand is insufficient, its payment to j, denoted x_{ijt} is strictly smaller and hence $x_{ijt} \in [0, y_{ijt-1})$. 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: $\alpha_{it} = k_t + z_{it} - \sum_{j \neq i} (b_{jit} - x_{ijt})$, 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:

$$x_{ijt} = \max \left\{ \frac{y_{ijt}}{\sum_j y_{ijt}} \left[\min \left\{ \sum_j y_{ijt}, \alpha_{it} - v_t \right\} \right], 0 \right\} \tag{1}$$

Equation (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 (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. (forthcoming) 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.[7] The realizations of (K_{it}, Z_{it}) then govern the transitions of banks between the M groups according to

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

some classification function $\Phi : \mathbb{I} \to \mathbb{M}$. 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 this 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 riskeness of each bank.

2.2 Empirical implications of payment equilibrium

Our approach is based on the dynamics of the cost of funding. In particular, we are interested in recovering the correlation patterns both within the cross-section and over time. The intuition for our approach comes from comparative statics of the payment equilibrium: Suppose bank i poses a greater risk for the financial system than bank j, but both are important large banks (i.e., indistinguishable based on observables). Further suppose that bank i suffers an adverse shock to its balance sheet in period t (the new distribution $\hat{H}(K) \equiv \int_{infty}^{K} \int_{-\infty}^{\infty} g(u,v)\, dvdu$ is dominated in the first order stochastic sense by the original $G(K,Z)$). This on average translates into higher cost of funding on the interbank market. 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. Suppose a similar scenario occurs with bank j, but in period t'. Since i' s contribution to systemic risk is higher by assumption, we should expect other banks' willingness-to-pay for liquidity in the MROs to increase shortly after the adverse shock to i's balance sheet. It should increase more than following the shock to j's balance sheet.

Obviously, an important caveat is that for banks that have healthy balance sheets and rely mostly on loans secured by high quality collateral, the willingness-to-pay for liquidity does not necessarily have to increase after such shocks. To address this caveat, we make use of the simple model of the composition of the willingness-to-pay as in Cassola et al. (2013). In particular, for banks submitting multiple steps in an auction, which large banks typically do, our method allows us to obtain an estimate of the spread between funding costs secured via unsecured and secured loans,

$u_i - s_i$, which captures to a large extent the default risk of bank i. Our goal is therefore to identify banks whose significant changes in the quantity weighted willingness-to-pay and in the spread $u_i - s_i$ propagate further into the system. In other words, whether such changes are followed by similar changes in the spreads for other banks. 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 panel data techniques in order to try to infer the overall impact of a shock to bank i's willingness-to-pay on the other banks in the system.

2.3 Auction Equilibrium

The payment equilibrium characterized in the previous section gives rise to the opportunity cost of obtaining liquidity from the central bank. Each bank with particular weekly needs of financing should, on average, be able to get a loan on the interbank market at the rate $\bar{R}_i$, which is the mean of the rates it pays on its loans. 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.[8] We assume that it is common knowledge among the market participants, to which of the M groups each bank belongs. Every week, each bank then 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 $F_m(\theta_i|X)$, where X is all public information, such as prevailing overnight unsecured interest rate EONIA, that is available that week. To recover the willingness-to-pay (i.e., the cost of funding) from the bids submitted in the discriminatory auctions we assume that banks play a Bayesian Nash Equilibrium, which we will describe below.

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

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

public information, X, available before a given auction, each bank's willingness-to-pay 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. As 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.

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 of this good. The number of potential bidders (banks), N, is commonly known. The number of bidders in each class m, N_m is also commonly known. Q denotes the amount of liquidity offered for sale by the central bank, i.e., the good to be divided between the bidders, and we assume Q is perfectly divisible. Q is itself a random variable since the ECB augments the supply based on current situation after the bids have been submitted. We assume that the distribution of Q is common knowledge among the bidders. Each bidder receives a private (possibly multidimensional) signal, θ_i, which is the only private information about the underlying value of the auctioned goods. The joint distribution of the signals will be denoted by $F(\theta_1, ..., \theta_N)$. As mentioned above, we assume (conditionally) independent private values (CIPV) paradigm.[9] In this case the θ_i's are distributed independently across bidders, and bidders' values do not depend on private information of other bidders, i.e., 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 $\forall \theta_i$.*

[9] Bindseil, Nyborg and Strebulaev (2009) provide some econometric evidence that private values might be appropriate in case of repo auctions.

$V_i(q,\theta_i)$ 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: $\sigma_i : \Theta_i \to \mathcal{Y}$. 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 $i = 1, ..., N$ has an action set:*

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

Therefore the set $\mathcal{Y}$ includes all non-decreasing step functions with at most 10 steps, $y : \mathbb{R}_+ \to [0,Q]$, 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.

Finally, since bidders use step functions, a situation may arise in which multiple prices would clear the market. If that is the case, we assume consistently with our application that the auctioneer selects the most favorable price from his perspective, i.e., the highest price. In case of excess demand at the market clearing price, we assume consistenly with our application 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:

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

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 $\mathbf{y}(\cdot|\theta) = [y_1(\cdot|\theta_1), ..., y_N(\cdot|\theta_N)]$, 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 (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:

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

and at $K_i(\theta_i)$:

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

where $\overline{q} = \sup_{(Q,\theta_{-i})} q_i^c(Q, \theta, \mathbf{y}(\cdot|\theta))$, 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:

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

In practice, we use equation (3) to identify the marginal values at all but the last step, and we use equation (4) at the last step.[10] Note that as $K \to \infty$, (3) and (4) coincide in the limit. Kastl

[10] Using (4) at all steps leads to qualitatively very similar results, but the estimates turn out to be less precise due

(2012) further proves that there exists an equilibrium of a discriminatory auction in distributional strategies in this constrained game whenever signals are not too dependent. It is also important to point out that as the number of players in each group m grows large, the best responses 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 main data set consists of all bids in the main refinancing operations and long term refinancing operations of the ECB during 2007 and 2008. In total there were 91 auctions before the switch to the full allottment took place in mid-October. The data from the MRO auctions is summarized in Table 1.

To obtain information on balance sheets of individual banks we also linked our data with Bankscope. We successfully linked the data for 390 European banks. We used these data to obtain information on 40 balance sheet variables.

Finally, we obtained data from the 2011 stress tests conducted on the 90 (presumably largest or most important) banks selected by the ECB or the national central banks. These data contain responses of many important financial indicators to various simulated scenarios: a negative outlook for the Euro-zone economy (-3% deviation of the GDP) and a shock to sovereign yields that should be similar in magnitude to the Greek crisis.

We also use reports of European Commission on government interventions in individual banks (EC 2011). For 629 banks that appear in our data from liquidity auctions from 2007 and 2008, we identified 20 banks that received targeted government support at least once. 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 recepients of government funds were the Anglo Irish Bank Corporation, which was eventually nationalized, and the IKB Deutsche Indus-

to the necessity to numerically estimate the derivative of the distribution of the market clearing price.

Table 1: Data Summary: Before and After May 2008

Summary Statistics				
	Mean		Std Dev	
	Before	After	Before	After
Bidders	284.2	414	28.92	58.46
Submitted steps	2.10	2.66	1.44	1.98
Price bid	4.18	4.37	0.07	0.21
Price bid spread[a]	0.17	0.15	0.07	0.22
Quantity bid[b]	0.005	0.003	0.01	0.01
Issued Amount (billion €)	173.94	147.48	2.49	2.60

[a] Spread against EONIA rate.
[b] Bid expressed as a fraction of the issued amount.

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

triebank, 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, occured before the collapse of Lehman Brothers between February 2008 and September 2008. After the subsequent change of the ECB's liquidity providing 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.

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

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. 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' willingness-to-pay for liquidity 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 willingness-to-pay 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 (3) provides us with the link between the observable data (bids) and the variables of interest: banks' willingness-to-pay. 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 (3) 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 (within group) independence and ex-ante (within group) symmetry among banks (within one of M groups), where we allow four groups of banks based on the quartile of the Tier 1 ratio. 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 $N_g - 1$ and N_{-g} 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 g. This simulated residual supply interestected with the submitted bid delivers a particular realization of

the market clearing price potential 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 (3). 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 one participate virtually always), we thus obtain a time series of willingness-to-pay for every bank, which rationalizes its bids in the MRO auctions.

4.2 Aggregate and Systemic Risk

If a bank's willingness-to-pay (or even the bid) exceeds the unsecured rate (EURIBOR, for Euro denominated loans), it means that that bank did not believe it could get an unsecured loan at that rate during that particular week. Solving for α_{it} from the equation $(1 - \alpha_{it}) * Euribor + \alpha_{it} * Eurepo = v_{it}$, we obtain a parameter summarizing roughly bank i's financial situation. If EURIBOR and Eurepo were the true unsecured and risk-free rates that bank i would face, i's funding costs would have to be bounded between these two interest rates. In such a case, the convex weight α would summarize whether i has sufficient high quality collateral and can thus obtain funding at a rate close to the risk-free one ($\alpha = 1$) or if it has no good collateral, in which case it has to rely on unsecured loans ($\alpha = 0$). If i would have to pay an unsecured rate that is higher than EURIBOR, α would be negative. We can therefore use our estimates of willingness-to-pay to find a share of banks that are likely unable to obtain an unsecured loan at EURIBOR, i.e., for whom $\alpha < 0$. Figure 1 plots this share (with the y-axis on the left) and also the unsecured rate (y-axis on the right). It clearly shows that there are large movements in this share, which suggests presence of systemic risk factors: not surprisingly, banks' financial situation tends to be correlated.

Let us now use our data to investigate this systemic risk and its distribution across countries and/or banking groups in more detail. In order to explicitly define what we mean by systemic risk,

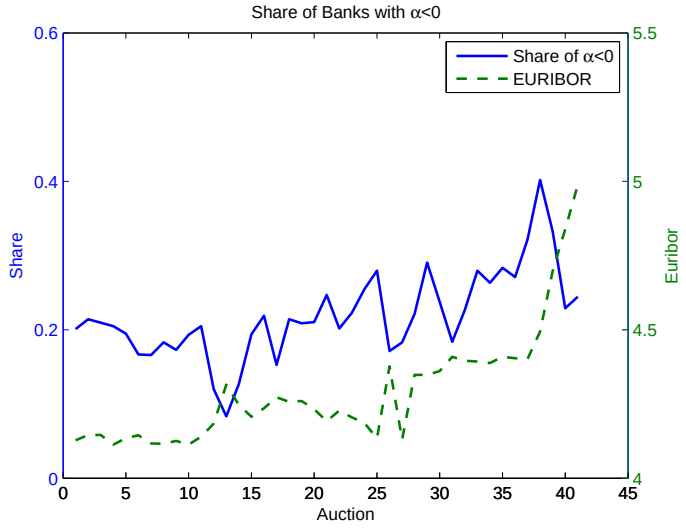


Figure 1: Share of Banks unable to borrow at EURIBOR

consider the following simple model of the dynamics of the cost of funding:

$$v_{i,t} = \beta_{0,i} + \sum_j \beta_{i,j} v_{j,t-1} + X_{t-1}^T \gamma_i + \epsilon_{i,t} \tag{5}$$

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 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 part of 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. Assuming X_t controls for common shocks, the $\beta_{i,j}$'s could be interpreted as causal effects due to direct debt contracts between banks or fire-sale spillovers.

4.3 Estimation Procedure: The Elastic Net

A first important challenge in interpreting a specification like (5) 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 WTP 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 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 the sizes of y and Z are $N \times 1$ and $N \times P$, respectively. The elastic net estimator of this model can be defined as:

$$\hat{\theta}_{enet} = \left(1 + \frac{\lambda}{2}(1-\alpha_e)\right)\left(\underset{\theta\in\Theta}{\operatorname{argmin}} \|y - Z\theta\|_2^2 + \lambda\left(\frac{(1-\alpha_e)}{2}\|\theta\|_2^2 + \alpha_e\|\theta\|_1\right)\right) \tag{6}$$

where $\|\cdot\|_1$ and $\|\cdot\|_2$ denote the L^1 and L^2 norms, respectively. The term $\left(1 + \frac{\lambda}{2}(1-\alpha_e)\right)$ is a bias correction factor added by Zou and Hastie (2005) to lessen the downward biased 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 consistently select the right model, 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 (or their lags) will be correlated. Zou and Hastie (2005) proposed the "elastic net" estimator in (6) to deal with this issue. Unlike the basic LASSO, the elastic net makes strongly correlated predictors be in or out of the model together, improving the prediction accuracy of the LASSO. Moreover, the elastic net can also be slightly modified to guarantee that the oracle property holds [11] .

There are two tuning parameters in the elastic net: α_e and λ. To perform the estimation, we fix α_e and then we use cross validation to find the optimal λ. We use the procedure just described to independently estimate (5) for each bank. We are particularly interested in obtaining an estimate of B, the matrix of direct spillover effects. Our measures of systemicness and vulnerability rely on such estimation. In Appendix A we perform a robustness check to assess the sensitivity of the results to changes in α_e. A non-negligible 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. Nevertheless, If banks are ordered according to their contribution to systemic risk, the resulting rankings are very similar for different values of α_e. Besides the sparsity "prior" implicit when choosing a LASSO-type estimator, we restrict the estimates $\hat{\theta}_{enet}$ to be non-negative

[11] 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:

$$\hat{\theta}_{anet} = \left(1 + \frac{\lambda_R}{N}\right)\left(\underset{\theta\in\Theta}{\operatorname{argmin}} \|y - Z\theta\|_2^2 + \lambda_R \|\theta\|_2^2 + \lambda_L^* \sum_{j=1}^{P} \hat{w}_j |\theta_j|\right) \tag{7}$$

where an infinite weight ($\hat{w}_j = \infty$) implies that the $j-th$ variable is not selected ($\hat{\theta}_j = 0$). As a robustness check on the choice of the estimation procedure, we applied the adaptive elastic net to our model and compared the estimates to the ones obtained with the elastic net. The results regarding systemicness and vulnerability are very similar for a wide range of parameter values.

by leting $\Theta = \mathbb{R}_+^P$ in (6) since our main interest 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 very small number of parameters take negative values.

4.4 Systemic Risk and Network Centrality

Equation (5) can be rewritten as

$$v_t = \beta_0 + B^T v_{t-1} + \eta_t,\ \eta_t = X_t^T \gamma + \epsilon_t \tag{8}$$

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 B, its "Katz centrality" is defined as:

$$k_i^c(a) = \sum_{s=1}^{\infty} \sum_{j=1}^{n} a^s \left(B^s\right)_{i,j} \tag{9}$$

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:

$$k^c(a) = \left((I - aB)^{-1} - I\right)\mathbf{1} \tag{10}$$

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.

4.5 Systemic Risk

A clear benefit from our reduced form specification (5) is that it provides a straightforward interpretation of Katz centrality in terms of total effects of idiosyncratic shocks. Assuming stationarity, (8) can be iterated forward to obtain:

$$v_{t+s} = \sum_{\tau=1}^{s} \left(B^T\right)^{\tau} \beta_0 + \left(B^T\right)^{\tau} v_{t-1} + \sum_{\tau=0}^{s} \left(B^T\right)^{s-\tau} \eta_{t+\tau} \tag{11}$$

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:

$$\Delta v_{t+s} = \sum_{\tau=0}^{s} \left(B^T\right)^{s-\tau} \epsilon_{t+\tau} = \left(B^T\right)^{s} \epsilon_t \tag{12}$$

Now, if $K = \sum_{s=1}^{\infty} B^s = (I - B)^{-1} - I$, then $K_{i,j}$ can be interpreted as the sum across s of the responses in $j's$ WTP at $t+s$ to a 100bp temporary idiosyncratic shock to $i's$ WTP at t (assuming v_t is measured in percentage points). Similarly, it can also be interpreted as the long run effect on $j's$ WTP of a permanent shock of 100bp to $i's$ WTP. Moreover, the parameter a in (10) can be used to assign lower weights to more distant responses in the WTP. Even more generally, instead of using a single parameter, we allow α to be a function, denoted by $\alpha = (\alpha_0, \alpha_1, \alpha_2, ...)$, and correspondingly we define the **systemicness** of bank i as

$$k_i^s(\alpha) = \frac{1}{P} \sum_{s=1}^{\infty} \sum_{j=1}^{n} \alpha_s \left(B^s\right)_{i,j} \tag{13}$$

where P is the number of banks in the network. That is, k_i^s is the average, across all banks, of the (weighted) sum of the subsequent effects on their WTP of an exogenous temporary shock of 100bp to i's WTP. Equivalently, k_i^s is the average total effect of a permanent shock of 100bp to i's WTP on all other banks' WTP. Similarly, we define the **vulnerability** of bank i as

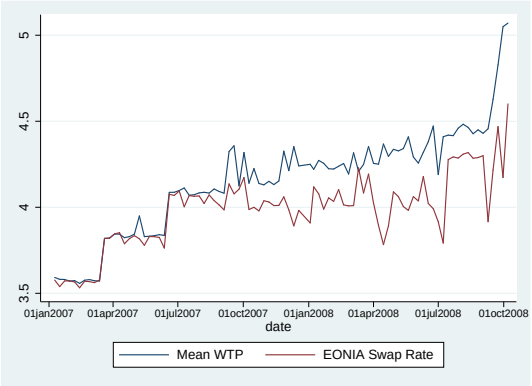
$$k_i^v(\alpha) = \frac{1}{P} \sum_{s=1}^{\infty} \sum_{j=1}^{n} \alpha_s \left(B^s\right)_{j,i} \tag{14}$$

i.e., as the total effect on i's WTP of an exogenous aggregate shock to all bank's WTP, normalized by the number of banks. It is the average total effect on i's WTP of a 100bp shock to any other bank, since the shocks are additive due to the assumed linear structure. The corresponding vectors $k^s(\alpha)$ and $k^v(\alpha)$ can be used to rank banks in terms of their systemicness and their vulnerability.

4.6 Controlling for common shocks

As already mentioned, a main concern with the reduced form specification (8) 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 WTPs 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. Cassola et al. (2013) show that the rate for a secured collateralized loan issued in the interbank market (EUREPO) is a lower bound for the WTP since counterparties in secondary markets require higher quality collateral than the ECB. Moreover, the EONIA swap rate can be used as an indicator of the expectations of the market rates. As in Cassola et al. (2013), our main variable of interest is the spread between the WTP and the EOINA 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 (5) implies that a positive deviation on bank j's WTP from EONIA predicts and increase in the WTP - EONIA 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 $X_{c,m,t} = R_{c,m,t-1} - R_{c,m,t-2}$, 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 (5) 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 WTPs.[12] First, it 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. 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. (2011)). 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

We begin by illustrating what was going on with the aggregate willingness-to-pay for liquidity. Figures 3 and 4 depict aggregate bids and values before and after the collapse of Lehman Brothers. The willingness-to-pay for liquidity obtained in the MROs of the ECB increased substantially during 2007-08. In fact, the marginal value for what amounts to 50% of the supply already exceeded the rate necessary to obtain a loan at the discount window at the very end of the auctions. This

[12]The 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 WTPs are stationary.

eventually led to the ECB's decision to abandon the auctions altogether and implement a full allotment mechanism instead. We discuss this development and the eventual demise of auctions in detail in Cassola et al. (2013).

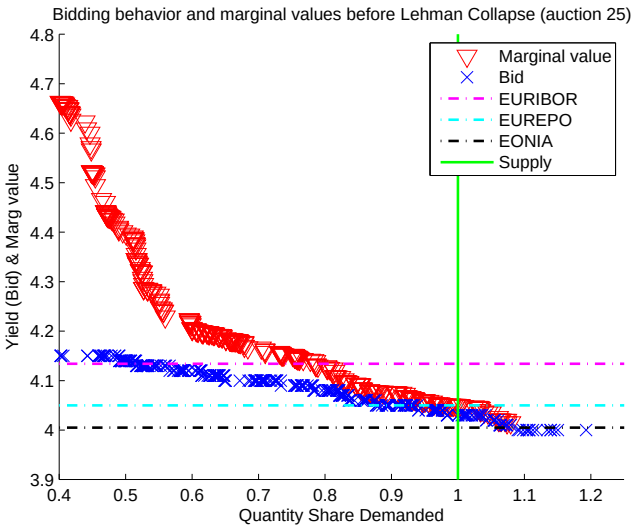


Figure 3: Aggregate Bid and Marginal Values Before Lehman Collapse

Figure 5 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 is a potential measure of centrality as well. Indeed, it is defined as the *degree* of the corresponding node in the binary counterpart of this matrix. It would, however, ignore 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

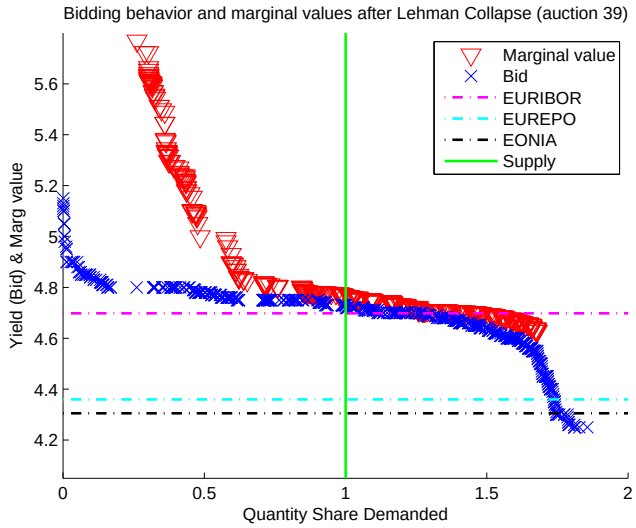


Figure 4: Aggregate Bid and Marginal Values After Lehman Collapse

uses the information on the (indirect) paths captured by the powers of the matrix representing the directed weighted network.

5.1 The parameter α in the Katz centrality measure

In principle, a link in the estimated network ($B_{i,j} > 0$) may represent both a direct spillover effect between two banks, for instance, due to a bilateral contractual obligation or and indirect effect through a third party. The vector of parameters α in (13) can be used to evaluate how persistent are the shocks to individual banks, that propagate to the whole system through the spillover effects. There are several reasons why a transitory shock to a bank may take time to vanish. Generally speaking, such persistence may reflect the average ability of banks to adjust to sudden changes in their credit conditions. Other possible source of persistence is the existence of indirect connections between banks that are not captured in the matrix of first order effects. To simplify the argument imagine that B reflects only direct effects due to, lets say, bilateral contractual obligations. That is, $B_{i,j} > 0$ implies that bank j holds a bond from bank i. Now let $B_{j,l} > 0$ as well, meaning that bank l is lending to bank j. If $v_{i,t}$ receives a positive shock, bank l's cost of funding would eventually increase, but not necessarily at $t + 1$ (it could be the case that bank i has no direct

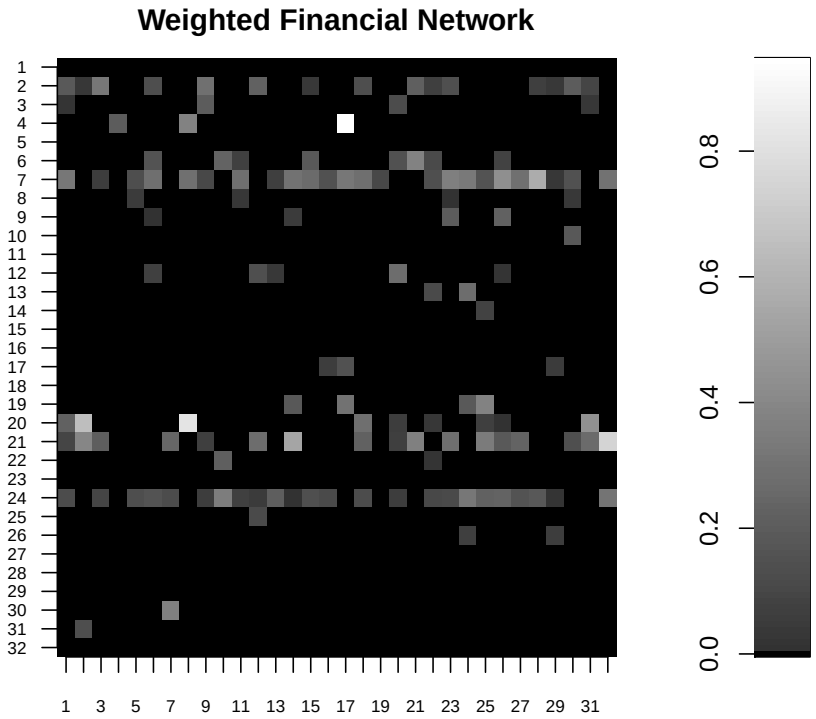


Figure 5: Elastic Net Coefficients: (1st order) Effect of a Shock to Row Bank's (lagged) Funding Costs on Column Bank's Funding Costs.

contractual obligation with l). Nevertheless, such indirect effect would be reflected in the matrix B^2, in fact $B_{i,j} > 0$ and $B_{j,l} > 0$ imply $\left(B^2\right)_{i,l} > 0$. By comparing $k^s(\alpha)$ for different choices of α, we can investigate how much information is added to our systemicness measure as higher powers of B are included. This, in turn, gives us an idea of how persistent are the shocks.

Let α^∞ be defined by $\alpha_t^\infty = 1$ for all $t \geq 1$. We use $k^s(\alpha^\infty)$ as a benchmark for our robustness check on α, since it contains all the information regarding direct and indirect spillover effects that could be extracted from B. Similarly, let $\alpha_t^\tau = 1$ for all $t \leq \tau$ and $\alpha_t^\tau = 0$, otherwise. We use α^τ to build a vector of systemicness scores that contains only the first τ powers of B. Clearly, $k^s(\alpha^\tau) = \frac{1}{P}\left(\sum_{t=1}^{\tau} \alpha_t B^t\right)\mathbf{1}$. We use the L^∞-norm to compare $k^s(\alpha^\tau)$ with $k^s(\alpha^\infty)$, that is we define

$$D_\tau = \left\| k^s(\alpha^\infty) - k^s(\alpha^\tau) \right\| = \max_i \left\{ k_i^s(\alpha^\infty) - k_i^s(\alpha^\tau) \right\} \tag{15}$$

Table 3 contains the values of D_τ for some values of τ from 1 to 32. The shocks appear to be

quite persistent. Two years after the shock ($\tau = 24$) the maximum difference between the average total effect of a 100bp to one bank is still around 8bp. That is, even very high powers of B add significant information on the systemicness of each bank. This justifies the use of Katz centrality as our measure of systemicness as opposed to other alternative measures of centrality that rely only on the first order effects.

Table 3: The Spillover Effects

τ	1	8	16	24	32
D_τ	1.2600	0.5470	0.2114	0.0817	0.0316

The table shows the maximum difference between the benchmark systemicness scores and the scores resulting from including only the first τ powers of B in the computation of the Katz centrality measure.

In Figure 6 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 a significant coefficient in the respective 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 rank of that node according to the Katz centrality measure $k^s(\alpha^\infty)$, where all the powers of the matrix B are included and equally weighted. 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.

Tables 4 and 5 offer a quantitative assessment of our estimates. In Table 4 we report the results from the elastic net regression of the funding costs of the 5 most central banks, i.e., the coefficients of B from equation (5), where we aggregated these effects to 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 ranking banks seem to impact the funding costs of Belgian banks, even though for one of the Spanish banks the effect is less than 1bp. If all top 5 banks suffered an increase in their funding costs of 100 basis points, the average funding costs of Belgian banks would rise right away by almost 80 basis

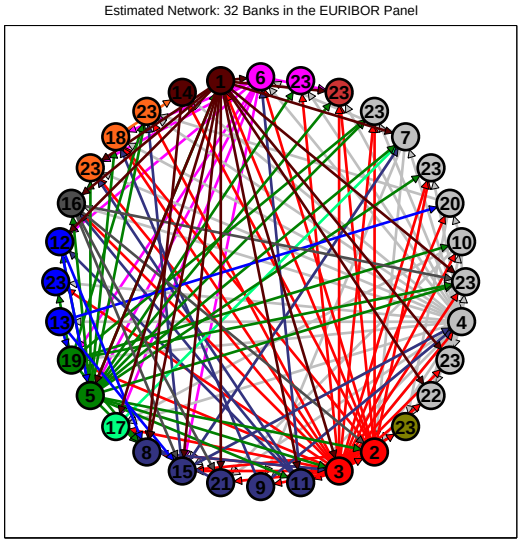


Figure 6: Financial Network

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	GB
British Bank (GB)	0.0362	0.2248	0.1721	0.0000	0.1250	0.2159	0.3379	0.1964	0.0880	0.0692	0.0323	0.1840
Spanish Bank (ES)	0.0753	0.0090	0.0304	0.0000	0.0818	0.0171	0.0000	0.0000	0.0000	0.0000	0.0383	0.3077
Spanish Bank (ES)	0.2837	0.2833	0.2526	0.1554	0.1431	0.2040	0.1636	0.1474	0.0449	0.1031	0.1675	0.0000
German Bank (DE)	0.1206	0.1219	0.1324	0.1176	0.1454	0.1123	0.2292	0.0006	0.0468	0.0621	0.0841	0.1762
Irish Bank (IE)	0.0337	0.1361	0.0461	0.0000	0.0693	0.1046	0.0000	0.0172	0.0332	0.2982	0.1021	0.1117

points. The funding costs of Greek banks seem to be quite sensitive to British banks. By contrast, Italian banks do not seem to react right after the shock, the direct impacts are all well below 10bp. However, this does not mean that they are not exposed to contagion, as shown in Table 5.

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 captures 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 dramatic. 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 is particularly true of the top most systemic British bank. In fact, this is the only bank in the sample for whom the systemicness score is, in fact, larger than

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	GB	GR	IE	IT	JP	NL
GB	1.3204	2.0684	1.4247	0.4150	1.4459	1.6106	2.0234	2.0656	0.7486	1.5471	0.9873	1.6121
ES1	0.8093	1.1707	0.7957	0.2256	0.8873	0.8937	1.0687	1.2358	0.4283	0.8574	0.6066	1.2535
ES2	0.9121	1.1176	0.8950	0.4387	0.8339	0.8667	1.0620	0.9352	0.3526	0.6991	0.6552	0.8519
DE	0.7844	0.9994	0.7854	0.3949	0.8295	0.7967	1.0759	0.8475	0.3659	0.6867	0.5905	1.0024
IE	0.6963	1.0673	0.6788	0.1939	0.7664	0.8235	0.8462	1.0003	0.3762	1.0523	0.5732	0.8820

[a] Averages across banks within a country reported.
[b] The total effect is computed using the Katz centrality measure with $\alpha = 1$.

one. Notice further than when the total effects are accounted for, the Italian banks are similarly exposed to contagion as banks in other regions. By contrast, Danish banks seem to be, on average, the less vulnerable.

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 7 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 11.9% of the average direct (first order) effect of a 100bp shock to the cost of funding of the top five most systemic banks.

In figure 8 we plot the histogram of the average total effect on a rival when a bank suffers a 100 basis point increase in its cost of funding. Our results 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 the other banks, there are important exceptions. In particular, there is one bank in the EURIBOR panel such that if it 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 140 basis points. Again, the shock is amplified by feedback due to indirect effects and high persistence.

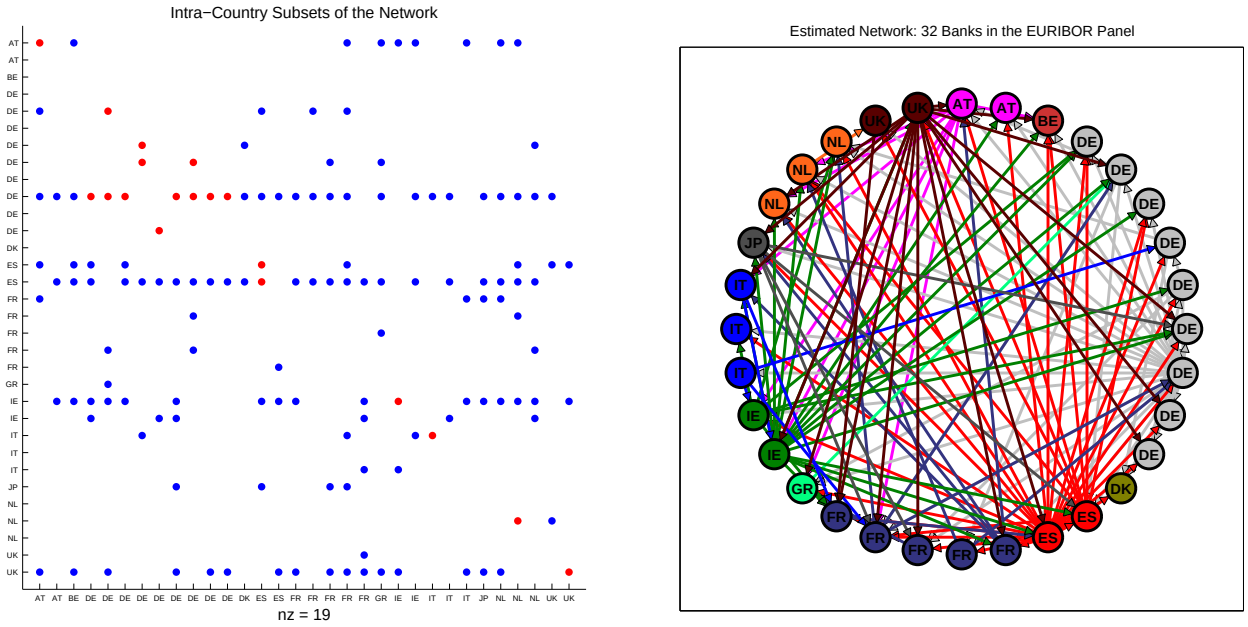


Figure 7: Relative importance of intra-country effects: In the left panel the links between two banks from the same country are highlighted in red.

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 WTP and bank j's WTP. Moreover, if the direction of the edges is also ignored, the network is simply an undirected graph, where a link between two banks represents that either of them is connected to the other one. As it can be seen directly by inspecting figure 9, 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 9 also illustrates other salient features of the network. The outer circle, formed by all banks labeled '23' (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 propagate the shocks). If the main driver of correlation among bank's

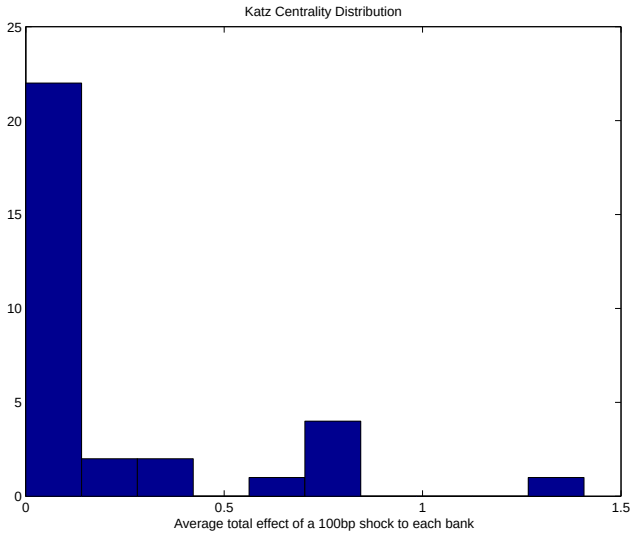


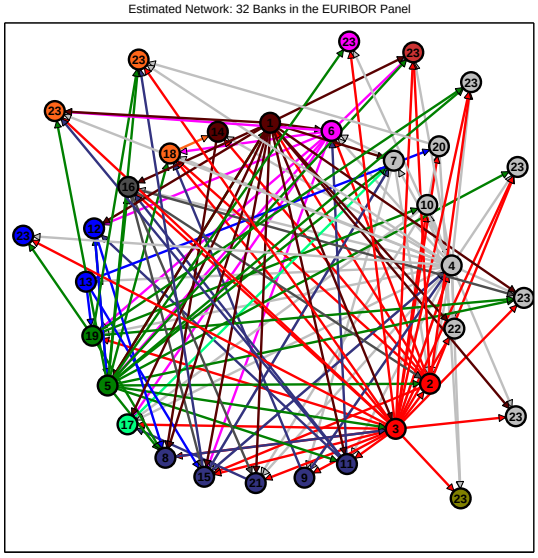
Figure 8: Katz Centrality Distribution: Average total effect on the funding costs of 32 banks in the EURIBOR panel of a 100bp shock to each of these banks.

funding costs was interbank lending then such network structure would suggest that almost one third 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. Thus, all banks are exposed to contagion from this subset of roughly two thirds of the banks in our sample.

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 structure[13], 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.

[13]For the reader familiarized 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 where chosen in some of the papers mentioned in this section to classify banks in mutually exclusive groups according to their level of interconnectedness.

Figure 9: Network Structure. All banks are labeled according to their systemicness rankings and their color correspond to their country of origin as in figure 7. The outer circle (all banks labeled '23') 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.



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 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 66% 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 remarkable is that even among these large banks, we find some that are exposed to, but do not expose others to contagion.

Wetherilt, Zimmerman and Soramäki (2010) study the network structure of the unsecured overnight loan market in the UK during 2006-2008. Noteworthily, 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 own 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 are 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 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 9). Thus, if any of them where 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 a quite sparse network, and even within the strongly connected component

(the inner circle in figure 9) there are a lot of missing edges (as a matter of fact, only 21.2% of all possible edges in such set belong to the network). Moreover, there are no cliques with more than two banks in our estimated network.

Vulnerability Predicts 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 probits of the probability of receiving such support (a bailout) on vulnerability (vuln_score), Return on Assets (ROA2008), Return on Equity (ROE2008), Tier1ratio and Average CDS spread of Dec. 2008 (cds_dic). The results are reported in table 5.1. Notice that the coefficient on vulnerability is significant in all regressions and it has the expected sign. ROA and ROE are also significant, but tier1ratio and CDS spread are not. Therefore, our measure does a better job at predicting bailouts than CDS spreads and tier1ratio. Our findings regarding tier1ratio, ROA and ROE are consistent with analogous results recently reported by the European Systemic Risk Board (The Structure and Resilience of the European Interbank Market 2013)

Average Willingness to Pay

Instead of using the full estimated network of direct spillover effects to measure how systemic or vulnerable a bank is, we alternatively consider estimating the effect of an increase in one bank's lagged WTP on the average WTP of all banks in the sample and also, the effect of an increase in the lagged average WTP of all banks on each bank's WTP. We do this by estimating the equations

$$v_{i,t} = \bar{b}_{0,i} + \delta_i v_{i,t-1} + \bar{\beta}_i^v \bar{v}_{t-1} + X_t \bar{\gamma}_i + \bar{\epsilon}_{i,t} \tag{16}$$

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

$$\bar{v}_t = \bar{b}_0 + \sum_j \bar{\beta}_j^s v_{j,t-1} + X_t \bar{\gamma}_i + \bar{\epsilon}_{i,t} \tag{17}$$

bailout	(1)	(2)	(3)	(4)	(5)	(6)
vuln_score	0.299*	0.309*	0.292*	0.300*	0.305*	0.333*
	(2.53)	(2.60)	(2.45)	(2.53)	(2.26)	(2.41)
ROA2008	-1.542*		-1.561*		-1.558*	
	(-2.58)		(-2.56)		(-2.48)	
ROE2008		-0.0461*		-0.0476*		-0.0477*
		(-2.75)		(-2.70)		(-2.71)
tier1ratio			0.0615	0.0837		
			(0.30)	(0.42)		
cds_dic					-5.070	-19.28
					(-0.09)	(-0.37)
_cons	-2.913*	-3.137*	-3.355^{+}	-3.753*	-2.885*	-3.045*
	(-3.02)	(-3.20)	(-1.88)	(-2.08)	(-2.86)	(-3.00)
N	32	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

(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 17 and $\bar{\beta}_i^v$ in equation 16 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 based on the full network of pairwise effects. Figure 10 shows the results of such comparison. All the scores and rankings based on the average WTP are positively correlated with their counterparts based on the full network. Four banks appear in the top five positions of both systemicness rankings and three other banks are among the top six most vulnerable according to both measures. However, the fifth most vulnerable bank, as measured by the full network, occupies the 25th position of the vulnerability ranking computed using the average WTP. Except for this anomaly, we do not find this results surprising. The positive correlation between the rankings follows from the way they are constructed. For instance, regarding the systemicness measure, if a bank's lagged WTP has an strong effect on many other's WTPs then it should also have an effect on the mean, and the

converse should also hold. Nevertheless, it is worth emphasising that the estimation of equations 17 and 16 only captures first-order effects. That is, the effects directly reflected in the matrix of coefficients B representing the network, ignoring all the higher order powers of B that are, by contrast, also included in the Katz centrality measures of the full network and that carry substantial information, as previously argued. Therefore, although positive, the correlations are far from one.

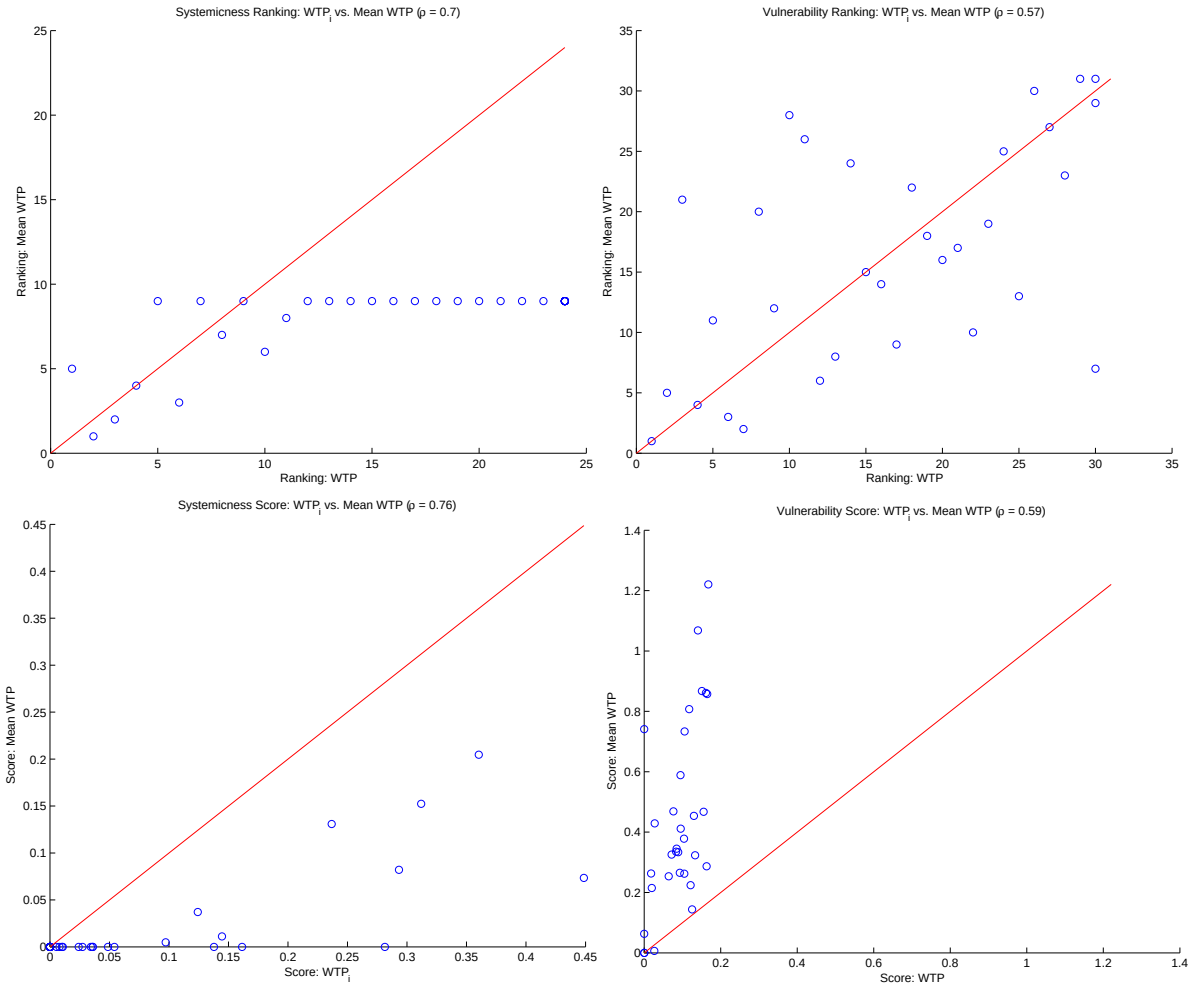


Figure 10: Systemicness and vulnerability rankings and scores: WTP vs. Mean WTP (correlation coefficients in parenthesis)

Controlling for Common Shocks

In general terms, we found evidence that the rankings are robust to the inclusion of bonds as controls in the estimation. However, the scores show some sensitivity to such an inclusion, specially at the

top positions of the corresponding ranking. To begin with, only a few bonds, or in some cases none of them, are selected by the elastic net in the regression for each bank's WTP. Both the systemicness and the vulnerability scores are in general, very similar, however, the vulnerability score of the second most vulnerable bank is reduced by 15 basis points after removing the controls. Similarly, the systemicness score of the most systemic score is reduced by approximately five basis points. Figure 11 shows the correlation between the systemicness and vulnerability scores and rankings.

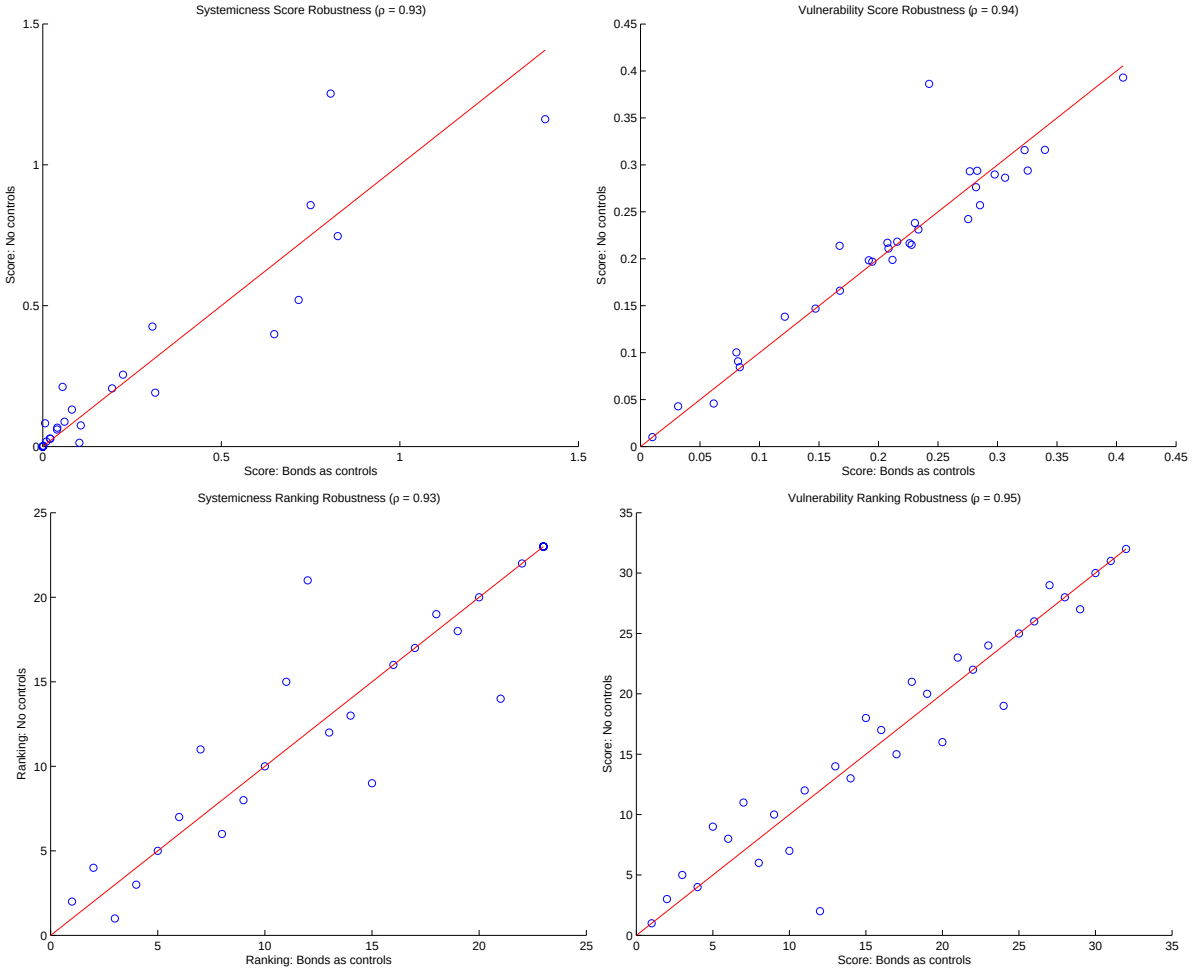


Figure 11: Effect on systemicness and vulnerability measures of controlling for government bonds (each point represents a bank in our sample).

Additionally, including the bonds yields controls seems to improve the accuracy with which our vulnerability score predicts observed bailouts, as measured by the difference between the average

Table 7: Vulnerability Rankings and Scores of the Bailed out Banks.

	No Controls	Bonds as Controls
	3 (32)	2 (34)
	11 (26)	7 (29)
	26 (14)	26 (12)
	12 (24)	11 (28)
	8 (29)	6 (30)
	25 (15)	25 (15)
	5 (29)	3 (33)
	1 (39)	1 (40)
Mean Scores		
Bailed out	26	27
Non-bailed out	20	19

[a] A comparison of the mean vulnerability scores of the bailed out vs. the non-bailed out banks is included in the bottom.
[b] The scores (in parenthesis) are measured in basis points.

scores of the bailed-out and the non bailed-out banks. Eight banks in the sample received targeted government support after August 2007. Table 7 reports the position of these banks in the vulnerability rankings. Under both specifications, the mean vulnerability score is higher for the subset of bailed out banks that for the banks that were not intervened. However, such difference is larger when bonds yields are included as controls.

5.2 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. Obviously, they also bid in the MRO 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 spill over effects between banks, and the corresponding systemicness and 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 might be subject to strategic manipulation for two main reasons. First, some of the bank's assets returns might be indexed to the EURIBOR.

Second, they are publicly known, and hence may be use by banks to signal their soundness. Bids, in the other hand, also have a strategic component. For instance, a bank could bid higher than in previous auctions, not because its underlying willingness to pay have increased, but just because it expects other banks to bid more aggressively. 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 WTP) with the ones resulting from estimating network B in equation (8) using monthly averages of the Euribor quotes and the quantity average bids in the period from January, 2007 to October, 2008. Figures 12 and 13 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 (relative to each other), as shown in Figure 12. The correlations between the rankings and scores based on WTPs, in one hand, and Euribor quotes, in the other one, are effectively zero or even negative. None of the top 5 most systemic banks according to each measure coincide, and the top 1 bank in each ranking is at the bottom of the other ranking. Moreover, the matrix B in equation (8) is not stable, when estimated using Euribor quotes instead of WTPs, i.e., it has eigenvalues outside the unit circle, violating the stationarity assumption.[14] Additionally, we notice that within an auction, the Euribor quotes are not correlated with the WTPs, nor with the bids, suggesting that banks with higher costs of funding generally do not submit higher quotes. Bids seem to do a better job than Euribor quotes at predicting systemicness and vulnerability, when compared to our measures, as shown in Figure 13. Both the systemicness and the vulnerability rankings are positively correlated, and the correlation coefficients are both significant.[15] Four of the six most vulnerable banks according to each measure coincide, and five of the top six most systemic in each ranking coincide as well. However, the second most vulnerable bank according to its WTP is ranked in the 21st position when the bids are used instead. Also, for all but two banks and for all but one bank their systemicness and ulnerability scores, respectively, are way lower when

[14] If the matrix B representing the Network is not stable, the Katz centrality measures for systemicness and vulnerability are computed using an α lower than the reciprocal of the largest eigenvalue in absolute value in order to guarantee convergence in (10).

[15] To determine if the correlation coefficient between two rankings of length N is significant, we compare the estimated value with the distribution of the Spearman's rank correlation coefficient between two randomly selected permutations of the set $1, ..., N$.

measured using bids. For completeness, we include in Appendix ??? the corresponding comparison between the rankings and scores obtained after estimating the model using bids and Euribor quotes, respectively.

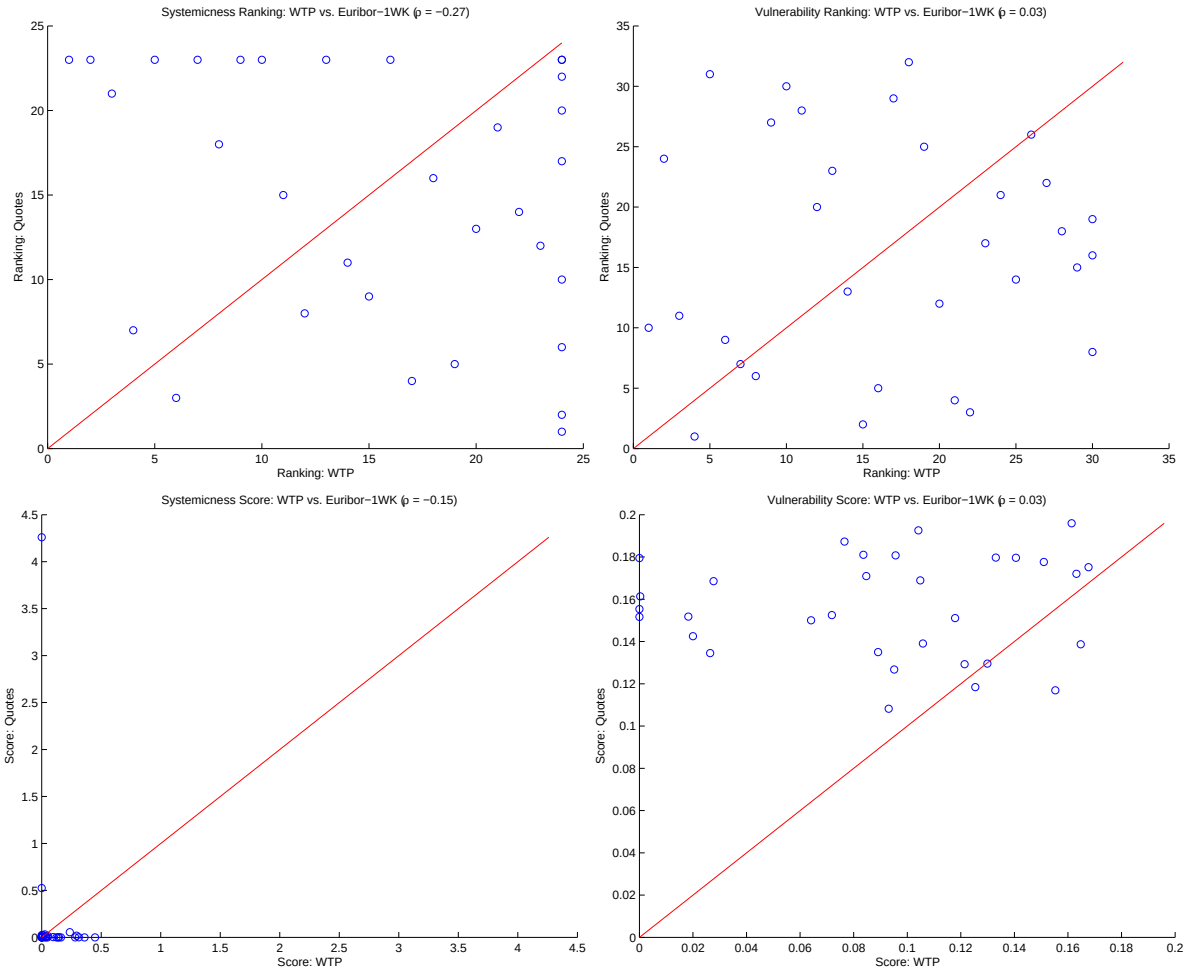


Figure 12: Systemicness and vulnerability rankings and scores: WTP vs Euribor Quotes (correlation coefficients in parenthesis)

5.3 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).

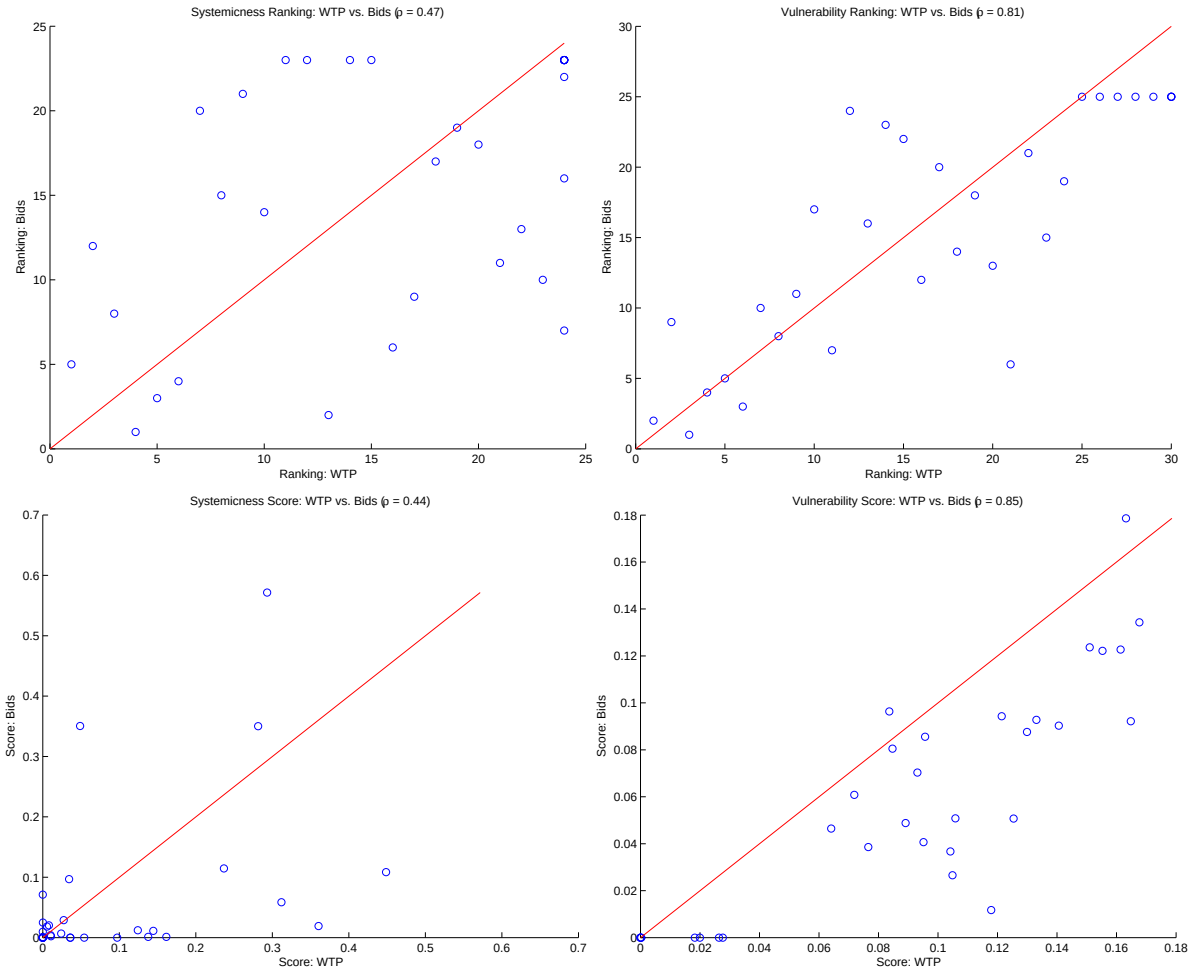


Figure 13: Systemicness and vulnerability rankings and scores: WTP vs Bids (correlation coefficients in parenthesis)

5.3.1 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 willingness to pay of different banks for a loan from the ECB 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 determined if lagged asset returns on stock prices of a given financial institution j, R^j_{t-1}, significantly predict the returns of another institution, after controlling for autocorrelation and heteroskedasticity. More precisely, their linear causality tests are based on the model:

$$R_t^i = a_{ij} R_{t-1}^i + b_{ij} R_{t-1}^j + e_{ij} \tag{18}$$

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, we define R_t^i as the monthly percent change in market capitalization. We use data on market capitalization from January 2007 to December 2008 to estimate the model described in (18) 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 $\hat{b}_{ij}$ 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:

$$R_t^i = \sum_j \gamma_{ij} R_{t-1}^j + \eta_{ij} \tag{19}$$

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 14 displays the sign patterns of three matrices Γ, G and B, containing the estimated coefficients for the VAR(1) on returns, the Granger causality test and the VAR(1) on

WTPs. 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 $P(P-1)$, where P is the number of banks in the sample (we do not count the diagonal elements). The number of (positive) connections are 13.6%, 8.7% and 17.8% of all possible off-diagonal connections for the VAR(1) on Returns, the Granger causality tests and the VAR(1) on WTPs, respectively. Using the Granger causality test, Billio et al. (2012) find that 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.

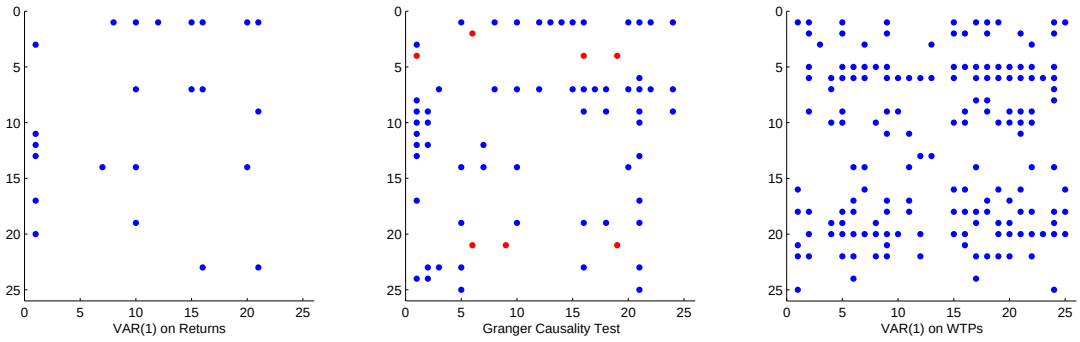


Figure 14: Sparsity of the network of interbank links estimated using data on market capitalization and costs of funding. A dot in the i-th row and j-th column represents a estimated connection from bank i to bank j. Blue and red dots represent positive and negative coefficients, respectively.

5.3.2 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

$$\Delta CoVaR_q^{S|i} = CoVaR_q^{S|Z_i=VaR_i^q} - CoVaR_q^{S|Z_i=Median^i} \tag{20}$$

There is more than one way to compute $CoVaR_q^{S|C(Z_i)}$ described in Adrian and Brunnermeier (2011). 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 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. 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 implied by CoVaR. In figure 15 we compare the results with the systemicness scores and ranking based on Katz centrality for the same subsample of banks. The correlation between the ranking implied by the two measures is significant, in the sense explain in footnote 15. However, there are some obvious discrepancies.

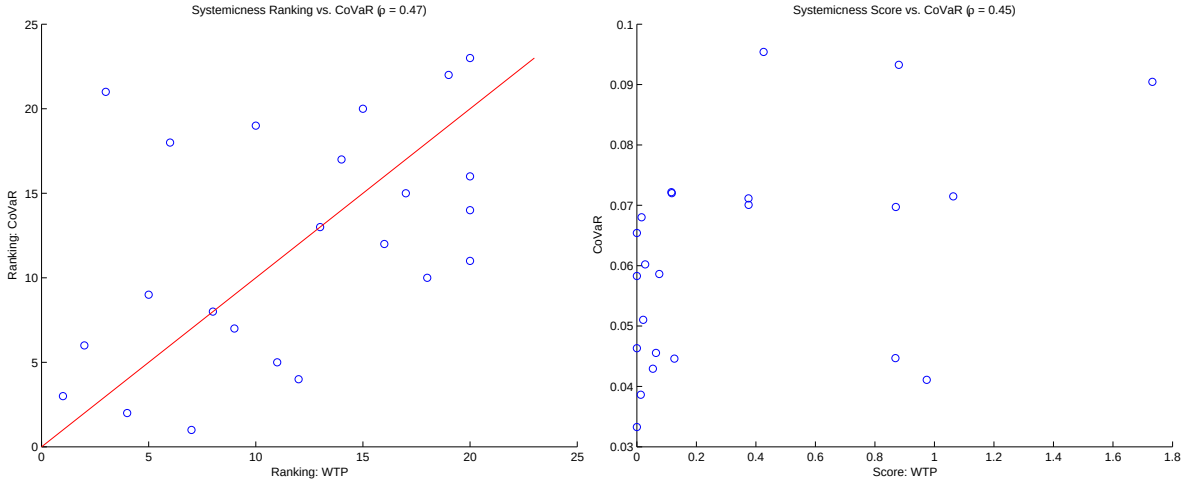


Figure 15: Contribution to systemic risk measured with Katz centrality and CoVaR

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 a variant of the LASSO known as elastic net to estimate the financial network. Our measure of systemicness, which is a version of Katz centrality, is directly related to the importance of a bank in this network. Perhaps not surprisingly, our results suggest that there is a large heterogeneity among the 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 (up to 140 basis points per 100 basis points shock) in the funding costs of the other banks.

In our future work we hope to be address the issue of causality, which we sidestep up until now. In particular, we hope to use the bailouts reported in section 3 as instruments in our estimation. This might allow us to obtain consistent estimates of some of the structural parameters.

References

Açemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "Systemic Risk and Stability in Financial Networks," *American Economic Review*, forthcoming.

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.

Aït-Sahalia, Yacine, Julio Cacho-Diaz, and Roger Laeven, "Modeling Financial Contagion Using Mutually Exciting Jump Processes," November 2011. working paper.

Allen, F. and D. Gale, "Financial Contagion," *Journal of Political Economy*, 2000, *108* (1), pp. 1–33.

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.

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.

Bindseil, Ulrich, Kjell Nyborg, and Ilya Strebulaev, "Bidding and Performance in Repo Auctions: Evidence from ECB Open Market Operations," *Journal of Money, Credit and Banking*, 2009, *41* (7), pp.1391–1421.

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.

Diamond, Peter A., "Aggregate demand management in search equilibrium," *Journal of Political Economy*, 1982, *90*, pp.881–894.

Dungey, Mardi, Renée Fry, Brenda González-Hermosillo, and Vance Martin, "Contagion in international bond markets during the Russian and the LTCM crises," *Journal of Financial Stability*, 2006, *2* (1), pp.1–27.

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\&language=EN\&guiLanguage=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.

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.

Greenwood, Robin, Augustin Landier, and David Thesmar, "Vulnerable Banks," November 2011. NBER working paper 18537.

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.

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).

Schwarcz, Steven, "Identifying and Managing Systemic Risk: An Assessment of Our Progress," *Harvard Business Law Review Online*, 2011, pp. 94–104.

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 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 Robustness

Recall that we opted for using the elastic net estimation method, because we are worried about using "pure" LASSO in our application due to the likely correlation between our right-hand side variables. We have, therefore, experimented with different values of α_e and fortunately the results remain qualitatively very stable. In figure 16 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 $\alpha_e = 1$ (pure LASSO). The correlation between the two sets of estimates is very high - above 0.9. However, all the LASSO estimates of the scores are lower in magnitude that the elastic net ($\alpha_e = 0.5$) estimates.

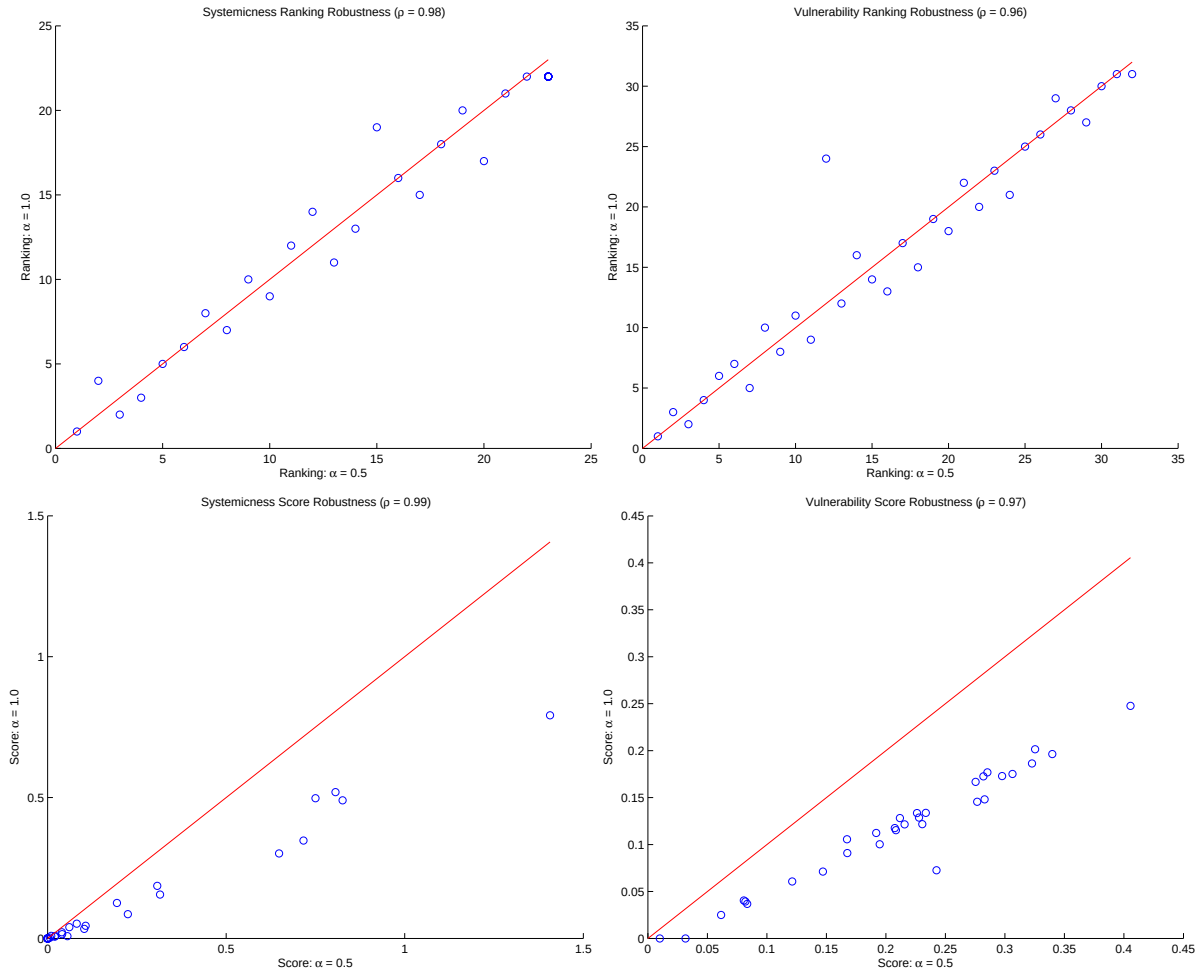


Figure 16: Robustness of systemicness and vulnerability rankings and scores to changes in α_e

A.2 Controlling for common shocks

The estimated matrix B of lag-correlations among banks' WTP is meant to capture direct spill over effects between banks. However, as discussed in section 4.6, some of these correlations might also be reflecting common shocks to the corresponding banks' funding costs.

Formally, we are concerned with the possibility that the error term in equation 5, $\epsilon_{i,t}$, 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. (Actually, lets assume for the sake of conciseness that there is only one of such variables). In order to model endogeneity, let us assume that the following is the true causal model generating the data. Suppose, for all i,

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

With, $\alpha_i \neq 0$. Notice that the source of endogeneity bias when running regression 5 would be serial correlation in z_t. That is, if $Cov\left(z_{t-1}, z_t\right) \neq 0$, since $Cov\left(z_{t-1}, v_{j,t-1}\right) \neq 0$ and $Cov\left(z_t, v_{i,t}\right) \neq 0$ (both by assumption), then, most likely, $Cov\left(z_{t-1}, v_{i,t}\right) \neq 0$. (For example, it could be the case that $z_t = \rho z_{t-1} + \xi_t$, and the argument above follows).

In order to address the endogeneity bias, the idea 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 both WTP and bond returns are observable. Thus, let us assume for now that for all bank j there is at least some bond k such that $Cov\left(v_{j,t-1}, x_{k,t-1}\right) \neq 0$. Moreover, if we assume further than bond returns are unpredictable from previous returns, then $Cov\left(x_{l,t-1}, x_{k,t}\right) = 0$ 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 $Cov\left(x_{k,t-1}, z_t\right) = 0$. Notice that for this condition to fail, z_t must be, to some degree, predictable using information available at $t-1$, and yet it has an effect on $v_{i,t}$ that is not already reflected in v_{t-1}.

A 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 bonds returns would be exposed to the same common shocks that we are trying to control for. In a few words, bond returns would also be endogenous, since $Cov\left(x_{k,t-1}, z_{t-1}\right)$ would not be zero, and thus they would not serve as valid instruments. 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[16].

A.2.1 IV Estimation

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 WTP, $v_{j,t-1}$, on the vector of (a selected subset of) bond returns. In the second stage, we regress each WTP, $v_{i,t}$ 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

[16] An alternative could be to use credit default swaps on sovereign bonds as instruments. The rationale for this is that default probabilities are correlated to bonds prices, but liquidations or firesales of bonds shouldn't increase the probability of a credit event and thus, neither the corresponding CDS spreads.

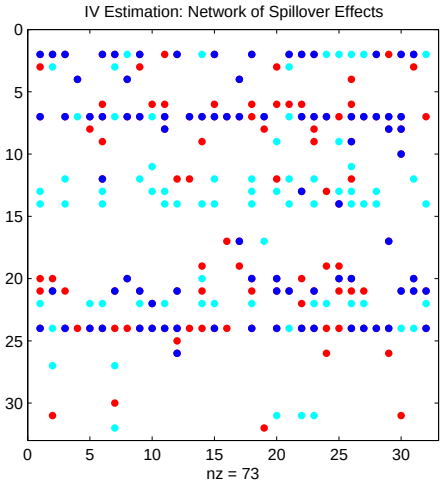


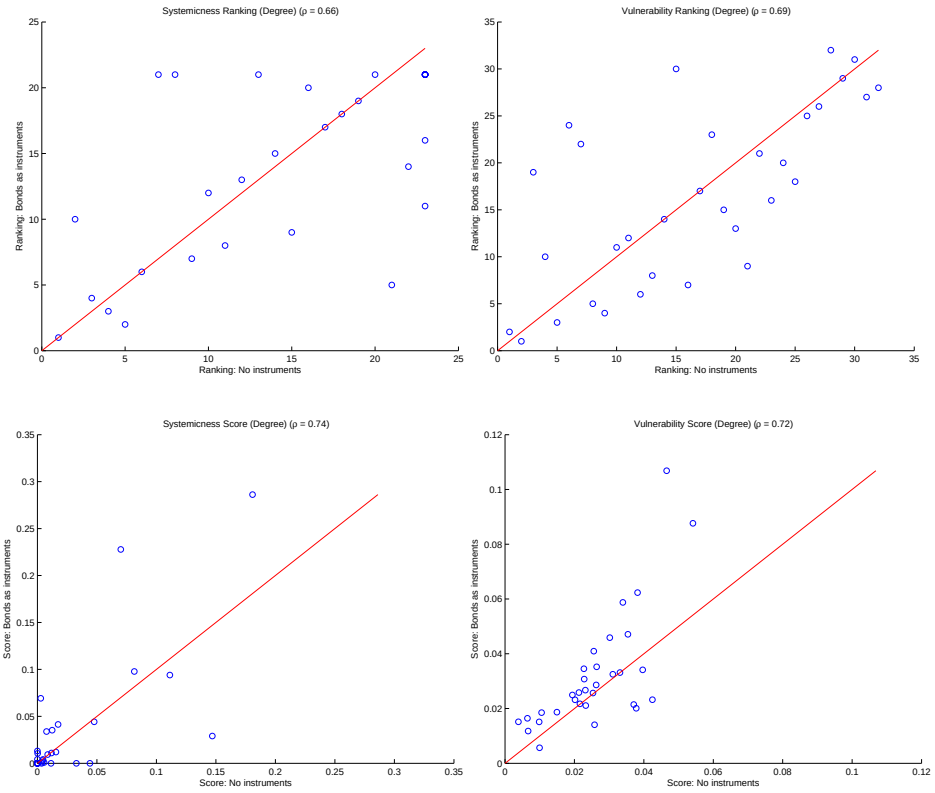
Figure 17: Effect on the network Structure Estimation of using bonds returns as instruments. Blue dots correspond to positive estimates under both estimations (with and without IVs). Red dots represent links that belong to the non-IV graph but not to the IV graph. Cyan dots are links that belong to the later but not to the former.

regression, as recommended by Belloni et al. (2012). In the second stage, we regress each WTP at t on the vector of predicted values, using Elastic net (as described in section 4.3).

The results from the first stage confirm 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 0.52. However, the number of instruments selected is close to the number of observations and roughly 47% of the OLS coefficients on selected bond returns are positive, which cast some doubt on the precision of the conditional expectation estimates.

From the second stage, we obtain estimates of the network of spillover effects between banks, as well as the derived centrality measures and rankings. If weights are ignored, the directed graph representing spillover effects is close to the one presented in section 5 (when bonds are included as controls, but not used as instruments). Figure 17 compares the two graphs, the blue dots correspond to links than belong two both, the red dots represent links that belong to the graph in section 5 but not to the one resulting from the IV estimation and, finally, the cyan dots are links that belong to the later but not to the former. Including the zeros, out of 1024 estimated coefficients, 890 (87%) coincide.

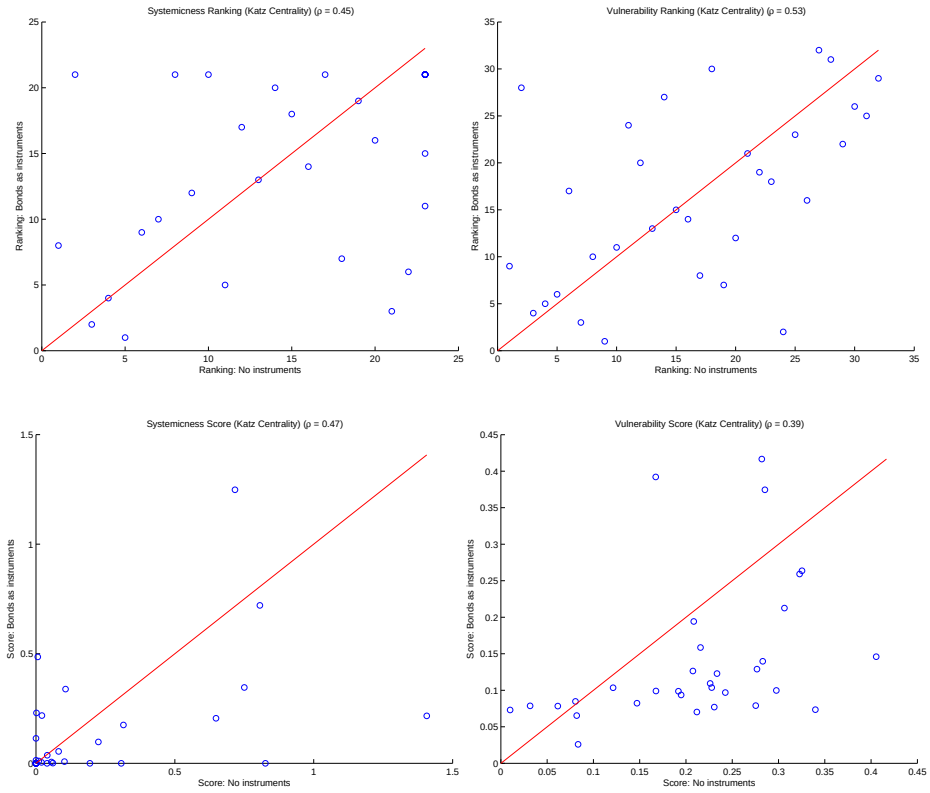
Figure 18: Effect on systemicness and vulnerability degrees of including bonds as instruments (each point represents a bank in our sample).



The rankings of banks according to how systemic and vulnerable they are, are also close to those from our original estimation, specially when only the first degree (direct) effects are considered. In figure 18, we compare the corresponding rankings and scores based on the degree of each node in the network, that is, based on the average first order responses to a 100bp shock to each bank's WTP. The correlations between the systemicness and vulnerability rankings are 0.67 and 0.7 respectively. The scores have similar correlations but they differ in magnitude, with the average IV scores being 24.4% higher.

When higher degree effects are consider, that is, when Katz centrality is used to score and rank banks, the differences between the two sets of estimates being concidered are larger. Figure 19, compares the corresponding rankings and scores based on the Katz centrality of each node in the network. The correlations are lower than the ones reported in figure 18 (between 0.4 and 0.5), and

Figure 19: Effect on systemicness and vulnerability measures (Katz centrality) of including bonds as instruments (each point represents a bank in our sample).



the average IV scores are 33% lower.

A.3 Granger Causality Test Using CDS

We have data on Credit Default Swaps for all the banks in our sample. We now define R_t^i in equations (18) and (19) as the monthly return on bank i's CDS and we estimate them using the same procedure just described. The number of connections in the estimated networks are 5.65%, 21.2% and 13.8% of all possible connections for the VAR(1) on CDS returns, the Granger causality tests and the VAR(1) on WTPs, respectively. Figure 20 shows the sign patterns of the matrices containing the estimated coefficients.

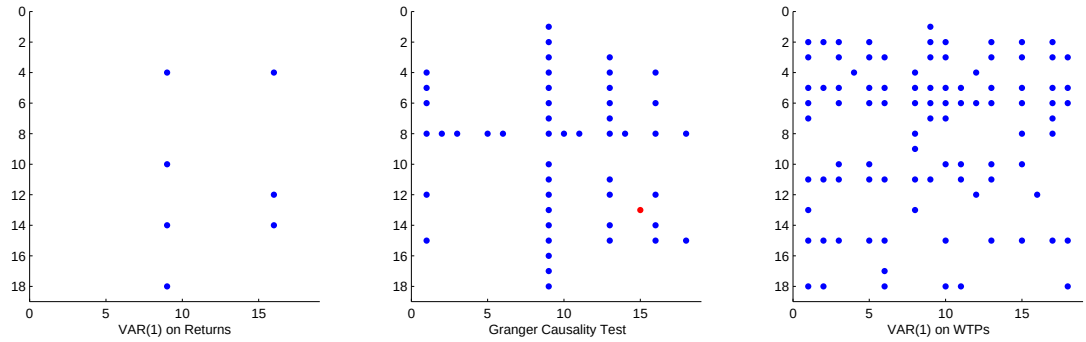


Figure 20: Sparsity of the network of interbank links estimated using data on CDS returns and costs of funding. A dot in the i-th row and j-th column represents a estimated connection from bank i to bank j. Blue and red dots represent positive and negative coefficients, respectively.

A.4 Bids vs Euribor Quotes

We estimated the model using both bids and Euribor Quotes, instead of WTP. Below we present figures showing how the resulting rankings and scores compare to each other when using these two alternative measures of the cost of funding. Essentially there is no correlation. This confirms our suspicion that the Euribor quotes are quite noisy.

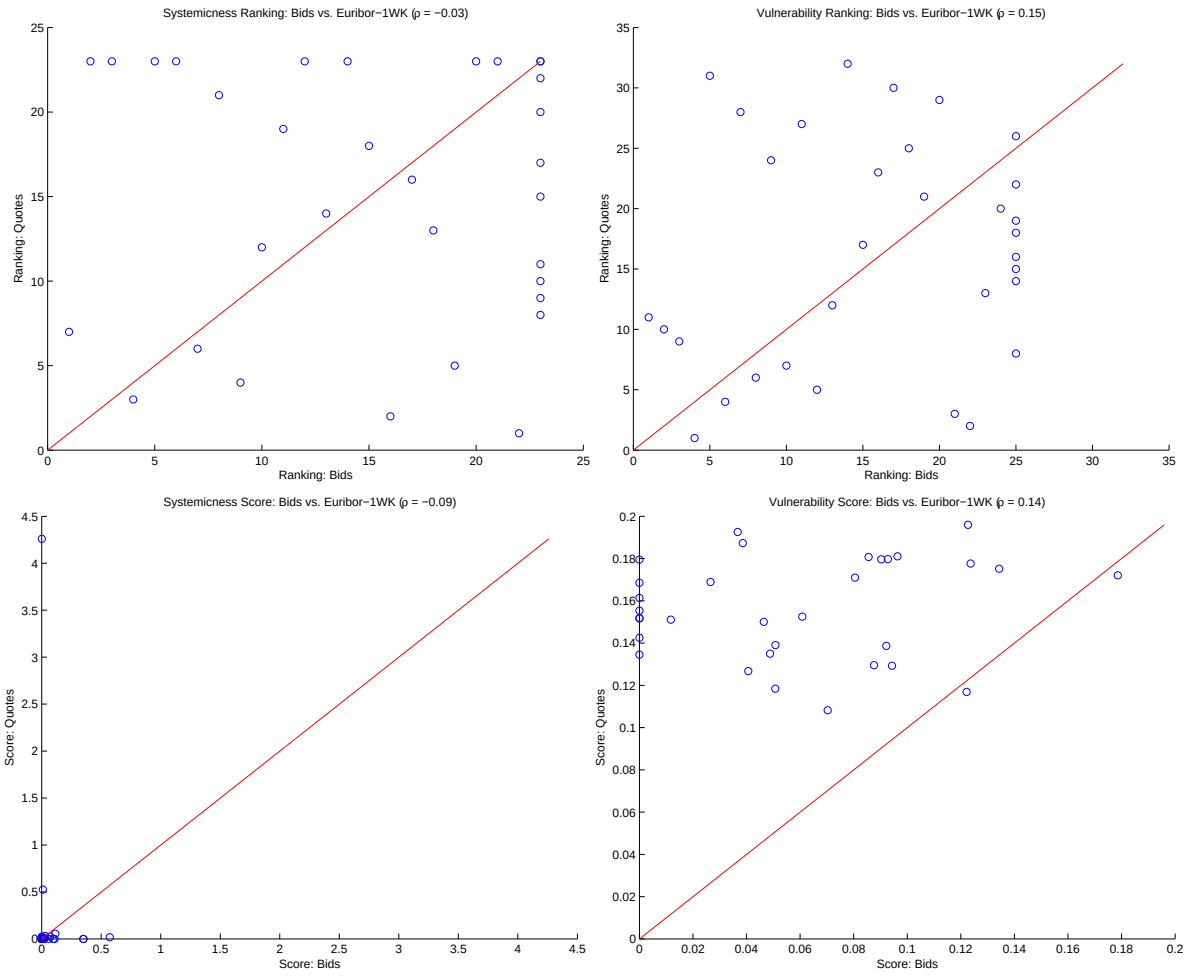


Figure 21: Systemicness and vulnerability rankings and scores: Bids vs Euribor Quotes (correlation coefficients in parenthesis)