Bank Credit Risk Networks: Evidence from the Eurozone[☆]

Christian Brownlees[a,b,*], Christina Hans[a], Eulalia Nualart[a,b]

[a]*Department of Economics and Business, Universitat Pompeu Fabra*
[b]*Barcelona GSE*

Abstract

This work proposes a credit risk model for large panels of financial institutions in which default intensity interdependence is induced by exposure to common factors as well as dependence between entity specific idiosyncratic shocks. In particular, the idiosyncratic shocks have a sparse partial correlation structure that we call the bank credit risk network. A LASSO estimation procedure is introduced to recover the network from CDS data. The methodology is used to study credit risk interdependence among European financial institutions. The analysis shows that the network captures a substantial amount of interconnectedness in addition to what is explained by common factors.

Keywords: Credit Risk, Networks, CDS, LASSO

JEL: C33, C55, E44, F36, G12, G13, G15, G18

1. Introduction

One of the lessons learnt in Europe in recent years is the systemic relevance of the financial sector and the potential risks of excessive interconnectedness. In Germany, several banks suffered a sharp increase in their CDS prices in January 2009 following financial turmoil at Commerzbank surrounding their take over of Dresdner Bank. In Italy, a scandal about secret derivative trading to conceal losses conducted by Monte dei Paschi di Siena lead to a surge in CDS prices of the entire Italian banking sector in January 2013. In both cases, cross-border linkages with banks in the Eurozone propagated the distress throughout Europe. As a result, several European countries introduced costly rescue packages for banks perceived as "too big to fail" or "too interconnected to fail" in order to mitigate the crises in their banking sectors. In response to these events, current bank regulation focuses on systemically relevant institutions. However, the detection of the network of credit risk interconnections between financial institutions and the identification of highly interconnected firms is still to this date an empirically challenging task.

The finance literature identifies two broad channels that induce dependence in the default risk of financial institutions: common exposure to a systematic shock and dependence between the idiosyncratic shocks of individual banks. As explained in Ang and Longstaff (2013), the systematic channel is associated with both macroeconomic or financial shocks. The effect of macroeconomic or financial shocks on the financial system has been the scope of extensive research, such as Calomiris and Mason (2003), Kritzman *et al.* (2011) and Stein (2012). At the same time, dependence can arise among the idiosyncratic shocks to banks, both through direct and indirect connections. Direct counterparty exposures between banks stem from the interbank market or obligations such as syndication and have been studied in Allen and Gale (2000), Mistrulli (2011), Tian *et al.* (2013) or Hale *et al.* (2016). Additionally, banks can be linked indirectly when holding similar portfolios, as shown in Gai *et al.* (2011), and Caballero and Simsek (2013).

Ang and Longstaff (2013) introduce a credit risk model that focuses on systematic shocks.

The authors build upon the standard reduced form models for pricing credit derivatives used in the finance literature (e.g. Duffie and Singleton, 1999) and propose a multifactor affine model in which defaults of individual financial institutions can be triggered by systematic or idiosyncratic shocks. Their empirical analysis forcefully shows the relevance of systematic factors.

Despite the undisputed relevance of the systematic channel, in a study of forty three financial crises Alfaro and Drehmann (2009) find that only half of them occurred before the macroeconomy experienced adverse economic shocks. This motivates us to focus on studying the dependence among the default intensities of financial entities conditional on the systematic factor. In particular, we assume that, conditional on the systematic factors, the idiosyncratic shocks of the financial entities have a sparse partial correlation structure. The set of non-zero partial correlation relations among financial entities is then used to define a network that we name the bank credit risk network. We assume that the econometrician does not know which financial institutions are partially correlated among each other and the key problem of interest of this work consists in recovering such network from the data.

We develop an estimation strategy to recover the bank credit risk network of large panels of financial entities from CDS data. The global systematic default intensity of each entity is identified as the one of the German sovereign, whereas the country-specific systematic default intensity is identified as the one of the respective sovereign. Standard pricing formulas for single-name Credit Default Swap (CDS) contracts derived in Ang and Longstaff (2013) are applied to bootstrap risk neutral idiosyncratic default intensities from CDS data. Next we apply a LASSO procedure to determine which bank idiosyncratic default intensities are partially correlated with each other. In particular, in our framework we have that network detection can be cast as the problem of estimating a covariance matrix subject to sparsity constraints on its inverse and in this work we rely on the Adaptive Graphical LASSO (Adaptive GLASSO) algorithm to carry out this task (Yuan and Lin, 2007, Friedman *et al.*, 2011, Banerjee and Ghaoui, 2008, Fan *et al.*, 2009).

Our modeling approach has a number of highlights with respect to the literature on factor and network (credit) risk models. In relation to factor models such as Ang and Longstaff (2013) and Oh and Patton (2018) we remark that while there is evidence that these models are able to capture a large amount of co-variability in the data, they typically do not capture *all* interdependence. In this strand of the literature, it is also sometimes the case that the analysis is silent about the degree of interdependence in the data not explained by the factors. This may be a concern as ignoring residual idiosyncratic dependence may lead to substantial underestimation of the probability of joint defaults. In relation to network models based on the Connectedness Tables a là Demirer *et al.* (2017) we point out that our bank credit risk network can be naturally embedded in standard reduced form credit risk models, whereas the Connectedness Tables is based on a Gaussian VAR that is harder to justify for this specific type of data. We also point out that unlike Demirer *et al.* (2017), our model controls for the variability explained by the factors, thus it provides an assessments on the degree of interconnectedness conditional on systematic factors. Another highlight of the approach put forward here in relation to reduced form approaches proposed in the literature is that through bootstrapping intensities rather than working with CDS spreads directly we can make use of the entire term structure of CDS contracts, and interpret obtained partial correlations as interconnections between the default probabilities of two entities rather than their CDS prices.[1]

We apply this methodology to study a sample of 71 top financial institutions from 10 selected Eurozone countries in between 2006 and 2013. The sample includes two dramatic periods for Eurozone financial institutions: the financial crisis of 2007–2009 and the European sovereign debt crisis of 2010–2012. A number of empirical findings emerge from our analysis.

First of all, we find that the network channel captures a substantial amount of (positive) cross sectional dependence. The important implication of this is that the probability of joint

---

[1]It is also important to emphasize that the bootstrapped intensity adjusts for the yield curve, which indeed changed dramatically throughout the sample period of our analysis.

defaults can be severely underestimated when one does not take into account the network channel. Estimation results reveal that the channel is more relevant for core countries. We interpret this as a consequence of the sovereign debt crisis. As the crisis widens and credit risk increases, banks in Greece, Italy, Ireland, Portugal and Spain become more tightly interrelated with their respective sovereigns.

As far as the structure of the network is concerned, we find evidence of both intra- and inter-country linkages between banks in Europe. The network reveals that the most central banks are typically large financial institutions from the largest European economies. The network is also fairly concentrated, with the top 10 most interconnected entities accounting for roughly 40% of total interconnectedness of the system, signaling the presence small world effects in credit risk interdependence.

A rolling-window analysis shows that during crisis periods, heavily affected financial institutions become hubs in the center of the bank credit risk network. This is relevant from a contagion perspective, since otherwise healthy institutions in core countries can be affected by idiosyncratic shocks to troubled banks in the periphery. In crisis periods, these hubs can quickly spread adverse shocks and lead to major downturns, such that their identification and monitoring is crucial for the health of the financial system.

Finally, an out-of-sample validation exercise is used to backtest our methodology. We use our model to construct forecasts of the covariance matrix of the idiosyncratic default intensities as well as forecasts of the degree of interdependence of the financial entities in the panel. Results show that our network based estimator provides accurate forecasts that perform favourably relative to a number of alternative benchmarks.

This research is related to a number of contributions in the literature. First, this work is related to the literature on credit risk through CDS in finance and financial econometrics which includes, among others, the work of Duffie and Singleton (1999), Lando (1998), Longstaff *et al.* (2005) and Oh and Patton (2018). Second, our paper is related to the literature on network estimation, in particular using LASSO based techniques. The list of

5

contributions in this area includes, among others, the work of Diebold and Yılmaz (2014), Demirer *et al.* (2017), Hautsch *et al.* (2014), Brownlees *et al.* (2018), Barigozzi and Brownlees (2018). In the finance literature, LASSO techniques have successfully been applied, for instance, in Rapach *et al.* (2013). Last, our paper relates to the literature on market based measures of risk for financial institutions, see Adrian and Brunnermeier (2016) and Brownlees and Engle (2017),

The rest of the paper is structured as follows. Section 2 introduces the model and the estimation procedure. Section 3 contains the empirical analysis of the paper. Concluding remarks follow in Section 4.

2. Methodology

We introduce a reduced form credit risk model in which default dependence among financial entities arises through three channels: a global factor, a country factor and a bank network channel. The model is a variant of standard affine multifactor models used in the credit risk literature.

2.1. Credit Risk Model

Credit events are modelled as jumps of a Poisson process with stochastic intensity. The global shock is modelled as the jump of a Poisson process $M_G(t)$ with intensity parameter λ_G that follows a standard square root process,

$$d\lambda_G(t) = a_G(m_G - \lambda_G(t))dt + b_G\sqrt{\lambda_G(t)}dW_G(t)\,,$$

where a_G, m_G and b_G are positive with $2a_G m_G > b_G^2$, and $W_G(t)$ denotes a Brownian motion.

Next, we consider a set of s different sovereigns. The default of sovereign ℓ can be triggered by two different types of credit events: The first type is a systematic global shock, affecting all sovereigns simultaneously. Conditional on each global shock, the probability that sovereign ℓ defaults is denoted as $\gamma_{\ell,G} \in (0,1)$. The second type is a country-specific

6

shock that triggers default of the respective sovereign with certainty. It is modelled as the jump of a Poisson process $M_\ell(t)$ with intensity parameter λ_ℓ that follows a standard square root process,

$$d\lambda_\ell(t) = a_\ell(m_\ell - \lambda_\ell(t))dt + b_\ell\sqrt{\lambda_\ell(t)}dW_\ell(t)\,,$$

where a_ℓ, m_ℓ and b_ℓ are positive with $2a_\ell m_\ell > b_\ell^2$, and $W_\ell(t)$ denotes a Brownian motion independent of $W_G(t)$.

Last, we consider a panel of m financial entities each belonging to one of the s sovereigns, and observed over a period of n days. The default of institution i can now be triggered by three different types of credit events: a systematic global shock, a systematic sovereign shock and an entity-specific idiosyncratic shock. The probability that entity i defaults following a systematic global shock is denoted as $\gamma_{i,G} \in (0,1)$, while the probability of default following a systematic sovereign shock is denoted as $\gamma_{i,\ell} \in (0,1)$. The idiosyncratic shock of firm i is modelled as the first jump of a Poisson process $N_i(t)$ with intensity parameter ξ_i that follows a standard square root process,

$$d\xi_i(t) = \alpha_i(\mu_i - \xi_i(t))dt + \sqrt{\xi_i(t)}dB_i(t)\,,$$

where a_i, m_i and b_i are positive with $2a_i m_i > b_i^2$, and $B_i(t)$ denotes an entity specific Brownian motion independent of $W_G(t)$ and $W_\ell(t)$. Following an idiosyncratic shock, firm i defaults with certainty. Finally, we denote by $\mathcal{F}_t$ the natural σ-algebra generated by the Brownian motions $B_i(t)$, $W_G(t)$ and $W_\ell(t)$.

The Brownian motion vector driving the idiosyncratic default intensities has a positive definite instantaneous covariance matrix $\Sigma = [\sigma_{ij}]$, that is $B(t) = (B_1(t), \ldots, B_m(t))' \sim \mathcal{N}(0, \Sigma t)$. Rather than charactering the interdependence structure of the Brownian motion $B(t)$ using the instantaneous covariance Σ in this work we rely on partial correlations. The partial correlation between $B_i(t)$ and $B_j(t)$ measures the linear dependence between the two variables after partialling out the influence of the remaining variables in the panel. We

7

formally define it as $\rho^{ij} = \text{Corr}(e_i(t), e_j(t))$, where $e_i(t)$ and $e_j(t)$ are the prediction errors of the best linear predictors of $B_i(t)$ and $B_j(t)$, respectively, based on $\{B_s(t) : 1 \leq s \neq i, j \leq m\}$. It is well known that the linear partial dependence structure of the system is embedded in the concentration matrix (Dempster, 1972). In fact, we have that the instantaneous concentration matrix $K = \Sigma^{-1}$ is related to the partial correlations through the identity

$$\rho^{ij} = -\frac{k_{ij}}{\sqrt{k_{ii}k_{jj}}}, \quad i \neq j , \tag{1}$$

where k_{ij} denotes the i, j-th entry of the instantaneous concentration matrix K. A key assumption of this work is that the partial correlation structure of the innovation process driving the idiosyncratic intensities is sparse. In other words, we assume that the Brownian motion driving the idiosyncratic default intensity of a given financial entity is not partially correlated with *all* the Brownian motions of the other financial entities but only a subset. The set of nonzero partial correlation relations among the financial entities can be used to define a network. Notice that it follows from the identity in (1) that our sparsity assumption on the partial correlations is equivalent to assuming that the concentration matrix K is sparse. As we shall detail precisely in what follows, this property turns out to be extremely useful for estimation.

It is straightforward to check that the idiosyncratic intensity vector $\xi(t) = (\xi_1(t), \ldots, \xi_m(t))'$ has the same network conditional dependence structure of the Brownian motion vector $B(t)$. The instantaneous covariance matrix $\Sigma_\xi(t)$ of the intensity vector $\xi(t)$ is given by

$$(\Sigma_\xi(t))_{ij} = \sqrt{\xi_i(t)\xi_j(t)} \, \sigma_{ij}.$$

The integrated covariance matrix Σ_ξ^* of the intensity vector $\xi(t)$ over the time interval $[0, n - 1]$, which measures the total covariation over the entire sample, is given by

$$\left(\Sigma_\xi^*\right)_{ij} = \sigma_{ij} \int_0^{n-1} \sqrt{\xi_i(t)\xi_j(t)} \, dt,$$

8

and the corresponding integrated concentration matrix is $K_\xi^* = (\Sigma_\xi^*)^{-1}$. It can be shown that if K is sparse then K_ξ^* is also sparse. Indeed, since the trajectories of the intensities ξ_i are a.s. continuous, we can apply the mean value theorem for integrals, and we get the existence of a point $s \in (0, n-1)$ such that

$$\int_0^{n-1} \sqrt{\xi_i(t)\xi_j(t)}\, dt = (n-1)\sqrt{\xi_i(s)\xi_j(s)} \quad \text{a.s.}$$

Therefore, we have that

$$\Sigma_\xi^* = (n-1)D_s \Sigma D_s,$$

where D_s is the diagonal matrix with the vector $(\sqrt{\xi_1(s)}, \ldots, \sqrt{\xi_n(s)})$ on the diagonal. In particular, the integrated concentration matrix can be written as

$$K_\xi^* = \frac{1}{n-1} D_s^{-1} K D_s^{-1},$$

which proves that if K is sparse so will be K_ξ^* (almost surely). Observe that in our argument we have also used the fact that $\xi_i(s) > 0$ almost surely.

We define the bank credit risk network as the sparse partial correlation network implied by the integrated concentration matrix K_ξ^* of the idiosyncratic intensity vector $\xi(t)$, and we use the squared partial correlations implied by the integrated concentration matrix K_ξ^* as a measure of the strength of the relationship between the idiosyncratic intensities. The partial correlation between entity i and j is defined as

$$\rho^{ij} = \frac{-k_{ij}^*}{\sqrt{k_{ii}^* k_{jj}^*}},$$

and measures the correlation between the idiosyncratic intensities of bank i and j obtained after netting out the influence of the other intensities. The squared partial correlation coefficient can then be interpreted as the proportion of the variance of the idiosyncratic intensity of bank i explained by bank j after netting out from i and j the influence of the

9

remaining intensities, and is a natural measure of the strength of a link in this context.[2]

In order to bring the model to the data, in this work we make a number of identification assumptions that are similar to the ones made in Ang and Longstaff (2013). In particular, in our application to a panel of Eurozone institutions we assume that the global intensity is identified with the intensity of the German sovereign whereas the sovereign intensities are identified with the intensity of the corresponding sovereigns of each entity. We also point out here that in our framework systematic and sovereign defaults can cause bank default, but not the other way round. While this makes the estimation of the model fairly straightforward it is important to highlight that such an assumption may be quite restrictive in practice from a structural perspective. Overall, it is important to emphasize that our model estimates reduced form interdependence among the financial entities in the panel that a structural interpretation is more challenging.

2.2. CDS Pricing

One of the key quantities of interest for credit derivative pricing is the $\mathcal{F}_t$-conditional probability of survival to a future time T for a financial entity. In our framework, the probability that entity i has not defaulted by time T equals the probability that no idiosyncratic shock occurs until time T times the probability that the entity does not default following any of potentially many global shocks (with probability $1 - \gamma_{i,G}$ each) and systematic shocks

---

[2]We point out that while it is natural to use the squared partial correlation coefficient as a measure of the strength of a link, other choices are also possible. For example, one could opt for the absolute partial correlation. This measure will also typically exhibit a lower degree of cross-sectional concentration in comparison to the squared partial correlation.

(with probability $1 - \gamma_{i,\ell}$ each). In terms of conditional probabilities, this writes as

$$
\begin{aligned}
p_i&(t,T) \\
&= \mathsf{P}\left(N_i(T) - N_i(t) = 0 \,|\, \mathcal{F}_t\right) \mathsf{E}\left((1 - \gamma_{i,G})^{M_G(T) - M_G(t)} \,|\, \mathcal{F}_t\right) \mathsf{E}\left((1 - \gamma_{i,\ell})^{M_\ell(T) - M_\ell(t)} \,|\, \mathcal{F}_t\right) \\
&= \mathsf{E}\left(\exp\left(-\int_t^T \xi_i(s)\,ds\right) \times \sum_{g=0}^\infty \frac{1}{g!} \exp\left(-\int_t^T \lambda_G(s)\,ds\right) \left((1 - \gamma_{i,G})\int_t^T \lambda_G(s)\,ds\right)^g \right. \\
&\qquad \left. \times \sum_{j=0}^\infty \frac{1}{j!} \exp\left(-\int_t^T \lambda_\ell(s)\,ds\right) \left((1 - \gamma_{i,\ell})\int_t^T \lambda_\ell(s)\,ds\right)^j \,|\, \mathcal{F}_t\right) \\
&= \mathsf{E}\left(\exp\left(-\int_t^T (\gamma_{i,G}\lambda_G(s) + \gamma_{i,\ell}\lambda_\ell(s) + \xi_i(s))\,ds\right) \,|\, \mathcal{F}_t\right),
\end{aligned}
$$

where we are assuming that all computations are done under a risk-neutral probability measure. It follows from the last equation that the standard reduced form framework can be applied for valuing credit derivatives by setting the instantaneous probability of default for entity i proportional to $\gamma_{i,G}\lambda_G(s) + \gamma_{i,\ell}\lambda_\ell(s) + \xi_i(s)$. Also notice that in this modelling framework the instantaneous firm default intensity has a factor type representation.

A Credit Default Swap (CDS) is a financial swap agreement through which the CDS seller compensates the CDS buyer in case of a credit event (e.g. default). We denote by s_{it}^k the CDS spread of entity $i = 1, \ldots, m$ on day t with maturity k equal to $1, \ldots, 5$ corresponding to the maturities of 2, 3, 5, 7 and 10 years. We assume the spread to be paid continuously. Next to the CDS spread, we assume that there exists a risk-free asset, and we denote the associated (continuously compounded) risk-free rate by r_t and the price at time t of the zero-coupon bond with maturity T by $D(t,T)$, so that $D(t,T) = \mathsf{E}\left[\exp\left(-\int_t^T r(s)ds\right) \,|\, \mathcal{F}_t\right]$. We assume that the risk-less rate is independent of all intensity processes.

The CDS contract consists in two legs, the spread leg and the protection leg. The value of the CDS spread leg at time t of entity i is given by

$$
s_{it}^k \int_t^T D(t,s)\mathsf{E}\left[\exp\left(-\int_t^s (\gamma_{i,G}\lambda_G(u) + \gamma_{i,\ell}\lambda_\ell(u) + \xi_i(u))du\right) \,|\, \mathcal{F}_t\right] ds.
$$

11

with T equal to $t+2$, $t+3$, $t+5$, $t+7$ and $t+10$ for k equal respectively to $1,\ldots,5$. The value at time t if the CDS protection leg of entity i is given by

$$
\begin{aligned}
\text{CDS(protection leg)}_t = \omega \int_t^T D(t,s)\mathsf{E}\Big[& (\gamma_{i,G}\lambda_G(s) + \gamma_{i,\ell}\lambda_\ell(s) + \xi_i(s)) \\
& \times \exp\Big(-\int_t^s (\gamma_{i,G}\lambda_G(u) + \gamma_{i,\ell}\lambda_\ell(u) + \xi_i(u))du \Big) \mid \mathcal{F}_t \Big] ds,
\end{aligned}
\tag{2}
$$

where $1-w$ is the recovery rate. To meet the no arbitrage condition, the protection leg and the premium leg of a CDS contract must be equal, and we can get out the value of the premium payments,

$$
s_{it}^k = \frac{\text{CDS(protection leg)}_t}{\int_t^T D(t,s)\mathsf{E}\left[\exp\left(-\int_t^s (\gamma_{i,G}\lambda_G(u) + \gamma_{i,\ell}\lambda_\ell(u) + \xi_i(u))du \right) \mid \mathcal{F}_t \right] ds}.
$$

As shown in the Online Appendix, we can rewrite the CDS spread for each entity i as

$$
s_{it}^k = \frac{\text{CDS(protection leg)}_t}{\int_t^T D(t,s)F_{s,t}^{i,G}(\lambda_G)F_{s,t}^{i,\ell}(\lambda_\ell)G_{s,t}(\xi_i)ds},
\tag{3}
$$

with

$$
\begin{aligned}
\text{CDS(protection leg)}_t = \omega \int_t^T D(t,s)\Big(& \big(\gamma_{i,G}I_{s,t}^{i,G}(\lambda_G)F_{s,t}^{i,\ell}(\lambda_\ell) \\
& + \gamma_{i,\ell}I_{s,t}^{i,\ell}(\lambda_\ell)F_{s,t}^{i,G}(\lambda_G)\big)G_{s,t}(\xi_i) + F_{s,t}^{i,G}(\lambda_G)F_{s,t}^{i,\ell}(\lambda_\ell)H_{s,t}(\xi_i) \Big) ds,
\end{aligned}
$$

where $\lambda_G = \lambda_G(t)$, $\lambda_\ell = \lambda_\ell(t)$ and $\xi_i = \xi_i(t)$, and the functions F, G, I, H are standard and defined in the Online Appendix. It is important to stress that despite the idiosyncratic shocks network dependence assumption, the pricing of single–name CDS carries through unaltered.

175 *2.3. Estimation*

We carry out inference on the bank credit risk network by combining standard estimation techniques based on CDS prices (Duffie and Singleton, 1999, Ang and Longstaff, 2013)

12

together with LASSO-type estimation (Tibshirani, 1996, Friedman *et al.*, 2011).

First, we estimate the bank credit risk model parameters $\alpha_i, \mu_i, \sigma_i$ for each bank in the panel by minimizing the squared pricing errors between the model implied CDS prices $\hat{s}_{it}^k$ and the observed CDS price s_{it}^k, that is

$$\hat{\theta}_i = \arg\min_{\theta_i} \sum_{t=1}^{n} \sum_{k=1}^{5} (s_{it}^k - \hat{s}_{it}^k)^2 \ ,$$

where $\theta_i = (\alpha_i, \mu_i, \sigma_i, \gamma_i)'$. We assume a constant recovery rate and fix the value of the recovery fraction at 50%, which is in line with Ang and Longstaff (2013). Since we apply the recovery rate symmetrically to both the protection leg and the premium leg of the CDS contract, its magnitude does not play a role for estimation results. Note that the evaluation of this objective function requires performing a series of intermediate optimizations. For a given value of θ_i we "bootstrap" the corresponding idiosyncratic intensity ξ_i for each day. That is, for each day we find the ξ_i which minimizes the squared CDS pricing error across all maturities keeping the value of the parameters fixed to θ_i. This is not available in closed form but can be easily computed by nonlinear least squares. Then, the value of the objective function is computed as the sum the squared CDS pricing errors corresponding to θ_i and the sequence of bootstrapped idiosyncratic intensities ξ_i. The CDS mispricing error objective function is then minimized using a gradient based algorithm. As Ang and Longstaff (2013) point out, note that ξ captures the level of the CDS term structure while the $\alpha_i, \mu_i, \sigma_i$ capture its shape. In order to estimate the bank credit risk model we need to identify the global and sovereign intensities λ_G and λ_ℓ. The global intensity is bootstrapped from the factor credit risk model using German sovereign CDS prices following the same estimation strategy used for the bank credit risk model. Analogously, the sovereign intensity is bootstrapped from the factor credit risk model using sovereign CDS prices.

Next, using the bootstrapped idiosyncratic default intensity $\hat{\xi}(t)$ obtained in the previous

13

step, we estimate the integrated covariance matrix Σ_ξ^* as

$$\widehat{\Sigma}_\xi^* = \sum_{t=2}^{n} (\hat{\xi}(t) - \hat{\xi}(t-1))(\hat{\xi}(t) - \hat{\xi}(t-1))' \, .$$

Realized covariance estimators have a long tradition in finance (cf Merton, 1980) in the estimation of equity volatility, and have recently been rediscovered in the financial econometrics literature (cf Andersen *et al.*, 2003) which has thoroughly analysed the properties of these type of estimators.

Finally, we use the Adaptive Graphical LASSO procedure (Adaptive GLASSO) to estimate the integrated concentration matrix K_ξ^* and the bank credit risk network. The estimator is defined as

$$\widehat{K}_\xi^* = \arg \min_{K \in \mathcal{S}^n} \left\{ \mathrm{tr}(\widehat{\Sigma}_\xi K) - \log \det(K) + \kappa w_{ij} \sum_{i \neq j} |k_{ij}| \right\} , \tag{4}$$

where $\kappa \geq 0$ is the LASSO tuning parameter, $w_{ij} \geq 0$ is an adaptive penalty weight which is proportional to the reciprocal of the (i,j)-the entry of a pilot estimate of the concentration matrix K_ξ^*, and $\mathcal{S}^n$ is the set of $n \times n$ symmetric positive definite matrices. We denote by $\widehat{k}_{ij}$ the entries of the realized network estimator $\widehat{K}_\xi^*$. The bank credit risk network estimator is a shrinkage type estimator. If we set $\kappa = 0$ in (4), the estimator is equal to the inverse realized covariance estimator $(\widehat{\Sigma}_\xi^*)^{-1}$. If κ is positive, (4) becomes a penalized objective function with penalty equal to the sum of the absolute values of the non-diagonal entries in the estimator. The important feature of the absolute value penalty is that for $\kappa > 0$ some of the entries of the realized network estimator are going to be set to exact zeros. Thus, the highlight of LASSO estimation is that it simultaneously estimates and selects the nonzero entries of K_ξ^*. The effect of the adaptive weights w_{ij} is to penalise more severely coefficients that are close to zero while penalising less large coefficients. It is well know that the standard (G)LASSO (i.e. the estimator corresponding to $w_{ij} = 1$ for each i, j) generate biases in large coefficients. The Adaptive and SCAD GLASSO algorithms are alternative procedures that

14

have been introduced in the literature to overcome the limitation of the standard estimator. We refer to Fan *et al.* (2009) for more details on such procedures. Numerically, Friedman *et al.* (2011) and Fan *et al.* (2009) show that minimizing the objective function in (4) is equivalent to carrying out a series of LASSO regression. An appealing feature of this LASSO estimator is that it is guaranteed to provide a sparse positive definite matrix estimate of the concentration matrix. Moreover, the algorithm is suitable for the analysis of sparse large dimensional systems containing, say, hundreds of series. Note that the estimator depends on the choice of the tuning parameter κ which determines the sparsity of the K_ξ^*. It is chosen in a data driven way using the BIC model selection criterion, which is widely used in the literature (cf Yuan and Lin, 2007, Peng *et al.*, 2009). We point out that the theoretical properties of this procedure have been studied in an analogous setting in Brownlees *et al.* (2018). The implementation details of the algorithm are detailed in the Online Appendix.

3. Empirical Analysis

We use the methodology introduced in Section 2 to study the bank credit risk network of our sample of Eurozone financial institutions. We carry out both static and dynamic network analysis. The static analysis consists in estimating the network over the full sample period while in the dynamic analysis we use a rolling window estimation scheme to obtain a time series of networks. Last, we carry out a forecasting exercise to assess if our network methodology is able to produce accurate predictions of the future degree of interdependence among the institutions in the panel.

3.1. Data

We consider a sample of 71 large Eurozone financial institutions from January 1st, 2006 until December 31st, 2013. Note that, for simplicity, in what follows we refer to all of our financial institutions as banks. We focus on financial firms headquartered in Austria (AT), Belgium (BE), Germany (DE), Spain (ES), France (FR), Greece (GR), Ireland (IE), Italy (IT), Netherlands (NL) and Portugal (PT). For each of these countries we select all financial

institutions for which CDS data is available for the entire sample period. The complete list of banks included in the sample is reported in Table 1. For each bank in the sample we obtain daily mid-market spreads for one-year, two-year, three-year, five-year, seven-year and ten-year CDS contracts. Additionally, we include spreads for sovereign CDS contracts of the same maturity for all ten countries. The data used in this study are obtained from Markit, who collects CDS quotes from more than thirty market participants on a daily basis, and provides a composite spread only if on a given date observations from at least two different participants are available. As outlined in Fontana and Scheicher (2016), while sovereign CDSs in the euro area can be denominated both in Euro and in USD, USD-denominated CDSs might be traded with a premium to hedge for a depreciation risk arising in case of a credit event, which might be accompanied by a depreciation of the bond's currency. Furthermore, as outlined in for example De Santis (2015), the size of quanto CDS spreads, the spread between euro and USD-denominated CDS, varies over time and according to different sovereigns, which could introduce potential distortions in our model. To avoid any issues arising from this, we focus the analysis on CDS contracts for which the notional is denominated in Euro whenever available, and enhance our sample with Markit composite spreads for the notional denominated in US dollars otherwise. Since the CDS spreads themselves are denominated in basis points, we do not face the challenge of currency conversion and whenever both series are available, we can see that their correlation, both in levels and in first differences, is close to one. In order to calculate the values for zero-coupon bonds in the CDS pricing formulas, we refer to the Nelson-Siegel-Svensson curves estimated by Deutsche Bundesbank with daily frequency. Descriptive statistics on the CDS series are reported in the Online Appendix (Tables A-1 and A-2).

In the presentation of the empirical results we shall often consider four sub-periods capturing different phases of the recent history of the Eurozone financial system. The first period runs from January 1st, 2006 until August 1st, 2008 and is what we refer to as the pre financial crisis period preceding the bankruptcy filing of Lehman Brothers on September

16

15th, 2008. We let the pre financial crisis period end some weeks before the actual filing for Chapter 11 protection to avoid including a period of anticipation in the first subsample. The second period runs from August 1st, 2008 until April 1st, 2010 and is what we refer to as the financial crisis period. April 2010 is chosen as a breaking point, since it coincides with the official filing for financial help by Greece on April 23rd. We take this as a starting point for our third subsample, which we refer to as the sovereign debt crisis. This third period finishes on September 1st 2012, reflecting the initiation of the legal framework for Outright Monetary Transactions (OTM) by the ECB to face the European debt crisis. Our fourth subsample accordingly runs from September 2012 until the end of our sample period on December 31st, 2013.

3.2. Default Intensities

We begin by estimating the credit risk model introduced in Section 2 and bootstrapping the risk-neutral default intensities of each sovereign and bank in the panel. We report the full set of estimates in the Online Appendix (Tables A-3 and A-4). Figure 1 plots the bootstrapped sovereign/systematic risk–neutral default intensities, divided into core (Austria, Belgium, France, Germany, Netherlands) and periphery (Spain, Greece, Ireland, Italy, Portugal) countries. The scale of the plot is such that an intensity level of one corresponds approximately to a 1% probability of default over the next year. The time series profiles of the sovereign default intensities are similar but there are clear differences in the levels of the series for core and periphery countries. A principal component analysis on the first differences of the sovereign intensities shows that the amount of variability explained by the first principal component is 40%. The mean default intensity for core and periphery countries amounts to, respectively, 71 and 424 basis points. Note that through the height of the sovereign debt crisis, CDS spreads for Greek sovereign debt increased up to more than 23'000 basis points (or 230 percentage points) for 5-year CDS contracts, implying an instantaneous default probability higher than 100% for the period from September 21, 2011

17

to June 6, 2013.[3] Inspection of the bank idiosyncratic default intensity (reported in the On-line Appendix in Figure A-1) shows that there is a moderate degree of heterogeneity in the dynamics of idiosyncratic intensities. A principal component analysis on the first differences of the bank intensities shows that the amount of variability explained by the first principal component is less than 20%.

<div align="center">FIGURE 1 ABOUT HERE</div>

3.3. Static Analysis

We estimate the bank credit risk network over the full sample. We use the BIC to select the optimal level of shrinkage κ, which delivers a degree of sparsity of roughly 25%.[4] We point out that in order to get insights on the network estimation sensitivity, in the Online Appendix we report the so called trace plot in Figure A-2, that is the plot showing the estimated partial correlations as a function of the amount of shrinkage used in the estimation. Albeit being sparse, the estimated network is too interconnected for visualization purposes. To this extent we construct a network plot using a shrinkage parameter κ that delivers a degree of sparsity equal to 10%. We plot this network in Figure 2. In the plot, financial entities correspond to vertices and edges correspond to non-zero partial correlation relations. The vertex size is proportional to the weighted degree of each financial entity, which is defined as the sum of the square partial correlations relative to that entity, and the vertex color is set according to the bank's country of origin. The edge width is proportional to its weight, which is defined as the square of the corresponding partial correlation. The network layout algorithm[5] chosen to create the plot is such that the most interconnected banks in the network correspond to the most central vertices in the plot.

<div align="center">FIGURE 2 ABOUT HERE</div>

---

[3]In our analysis we truncate the default intensities at 100%.

[4]That is, the number of non-zero partial correlation is 25% of the total.

[5]We use the Fruchterman & Reingold (1991) force-direct graph drawing algorithm.

The analysis of the network provides interesting insights on the interdependence structure of Eurozone financial institutions. First, we use the partial correlation and weighted degree distributions, which are reported in Figure 3, to summarize the global properties of the network. The partial correlation distribution shows that the vast majority of the dependence in the network is positive. This facilitates the interpretation of the graph in Figure 2 in that the presence of an edge almost always signals that the idiosyncratic default intensity of a bank is positively related to the one of its neighbours. Turning to the weigheted degree distribution, we note that there is strong heterogeneity in the degree of the banks in the network. In particular, we have that the weighted degree of the top 10 most interconnected institutions in the panel accounts for roughly 40% of the total. The network literature typically emphasies that networks with higher concentration exhibit "small world effects", that is the distance (i.e. the smallest number of connecting edges) between any two nodes is proportional to the log of the total number of vertices. Small world effects imply that even if the network is large and sparse, all banks in the system are strongly interrelated. More specifically in our context small world effects imply a non-negligible probability of joint distress of a substantial number of institutions following the idiosyncratic shock to an individual entity.

FIGURE 3 ABOUT HERE

A number of interesting country clustering patterns also emerge from the network. To investigate these in more detail, in Table 2 we report the sum of edge weights among Eurozone countries in the network. The table shows that there is a high proportion of within-country linkages: after controlling for the global and country factor, banks belonging to the same country still exhibit a high degree of interdependence. This phenomenon is referred to as national fragmentation and has been documented by, among others, Betz *et al.* (2016). As far as between-country linkages are concerned, we observe that banks headquartered in France, Germany, Italy and Spain have the highest number of connections. It is interesting to note that the number of cross-border linkages is correlated with banks' international exposures. In

19

the Online Appendix in Table A-5 we report the ranking of countries by number of linkages and by the BIS estimate of the total foreign claims of each country computed as of the end of the sample. We note that the two rankings are similar, especially for top ranked institutions.

TABLES 2 ABOUT HERE

Last, we study which banks are most central in the network using the page-rank algorithm and report rankings of the top ten banks in Table 3. Central banks in the network can be interpreted as yellow canaries of distress in the system, that is highly interdependent institutions whose distress correlates with distress in a large fraction of the entire system. We analyse centrality using two different versions of the eigenvector centrality algorithm: one for unweighted networks and the other for weighted networks. The unweighted eigenvector centrality algorithm computes centrality on the basis of whether entities are connected or not, whereas the weighted algorithm also takes into account the edge weight, which in our case is the square partial correlation coefficient. Inspection of the rankings reveals that size is an important determinant of interconnectedness: the most central banks in the network correspond to the largest banks in the sample. Different pictures however arise depending the choice of the algorithm. When considering the unweighted network, the most interconnected banks are typically headquartered in core countries, especially in France and Germany. On the other hand, when considering weighted rankings, we see that the top ten includes a large number of troubled entities in periphery countries (Spain, Portugal and Italy) and in particular the Spanish Banco Santander (SAN) turns out to be the most interconnected institution in the Eurozone system.

TABLE 3 ABOUT HERE

The results above raise the question about which underlying channels of bank interconnectedness could be driving the observed patterns. Among others, potential candidates for actual channels are direct channels such as interbank lending and loan syndication as well

20

as indirect channels such as asset commonality. While we are unable to perform a com-
parison of observed correlations with different channels of interconnectedness for the entire
sample of banks due to unavailability of data, two of the authors have performed an anal-
ysis to answer the question for a subsample of German banks. The results of this analysis
have been published in the paper Abbassi *et al.* (2017). In that paper we identify direct
channels of interconnectedness through bilateral exposures in the wholesale funding market.
Indirect channels are measured through asset allocations, decomposed into banks' securities
investments and loans granted to the real economy. We show (for the restricted sample
of banks considered) that observed correlations strongly reflect both direct and indirect
channels. On the funding side, we find that bank pairs show higher observed correlations
whenever both counterparties have higher Tier 1 capital-weighted interbank exposures. For
indirect channels, we document that both banks' exposure to the real economy and their
securities investments are linked to higher observed correlations. Bank pairs show up as
more interconnected whenever their lending practices to the real economy are more similar
or whenever both counterparties have higher exposures to risky securities. Lastly, we show
that the relation between observed correlations and actual channels of interconnectedness
varies over time and in the cross-section. While interbank lending is a relevant driver of
observed correlations mainly during crisis times as other sources of financing become hard
to obtain, banks' securities investments show asymmetric effects in the cross-section. Bank
pairs with higher exposures to troubled security classes exhibit higher observed correlations,
while commonality in securities investments related to security classes which were unaffected
by the crisis does not induce higher dependence.

3.4. Time-Varying Analysis

In this section we carry out a time-varying analysis to study the evolution of the bank
credit risk network throughout the sample. The network time series is obtained by estimating
the model at the end of each month from January 2008 until the end of the sample using
the last two years of observations available. The BIC is used to select the optimal amount

21

of shrinkage.

Figure 4 shows the time series plot of the average weighted degree in the panel together with the time series plot of the 5% an 95% quantiles of the weighted degree distribution. We observe that the amount of interconnectedness exhibits substantial time series variation throughout the sample. The degree of interconnectedness builds up during the great financial crisis and it peaks in April 2010 as the first signs of the Europe sovereign crisis become apparent. Interestingly, we see a sharp decline in the degree of interconnectedness in the correspondence of Mario Draghi's announcements about the Outright Monetary Transaction program of the ECB.

FIGURE 4 ABOUT HERE

In order to quantify the amount of dependence captured by the factor and network components of the model, we define R^2 type goodness of fit indices that we call factor and network R^2. The factor R^2 of a bank is defined as the R^2 of the regression of its log-CDS spread difference on the log-CDS spread difference of its respective sovereign. The network R^2 of a bank is defined as the additional R^2 obtained by adding as explanatory variables of the previous regression all the log-CDS spread differences of all the neighbouring banks detected in the bank credit risk network. The factor and network R^2 are computed on the basis of rolling estimates in order to have a time series of values. Figure 5 shows the plot of the factor and network R^2 averages by country. The left and right panels show respectively the plots for the core and periphery countries. For core countries we note that the network channel is more relevant than the factor one and that the time series profile of the R^2s is roughly stable over time. Periphery countries exhibit a rather different behaviour. At the beginning of the sample, the network channel dominates the factor. However, as the sovereign debt crisis unwinds, banks progressively become more dependent on their respective sovereign and the relevance of the interconnectedness with other banks declines. This trend stops in the second half of 2012 in correspondence with the beginning of the Outright Monetary Transactions program of the ECB. Overall, the plots convey that both the factor and the

22

network channel explain a significant amount of comovement and exhibit different time series evolution throughout the sample.

Last, we summarise the degree of interconnectedness of two banks in our panel, Commerzbank and Monte dei Paschi di Siena, by displaying their weighted degree as well as the number of non-zero partial correlation through the sample. We point out that despite the fact that we use two fairly different measures to summarise the degree of interconnectedness of these two entities, the overall time series pattern that emerge are fairly similar. Commerzbank was facing major difficulties in the times surrounding the acquisition of Dresdner Bank starting from December 2008. In May 2009, Commerzbank received a liquidity injection by the government effectively constituting a partial nationalization of the bank. These events are reflected in the number of linkages as we see a sharp fall in mid 2009 following the nationalization. For Monte dei Paschi di Siena, in January 2013 news on the scandal surrounding derivative deals to conceal previous losses materialized, leading to a large drop in stock prices. Accordingly, the plot shows that the number of connections of this entity peaks around the same time.

3.5. Predictive Analysis

We carry out a predictive analysis to assess if the bank credit risk network methodology provides advantages for forecasting. In particular, we carry out two different evaluation exercises. In the first exercise we evaluate the performance of our network estimator in terms of covariance forecasting whereas in the second we evaluate the estimator in terms of providing useful rankings of the most interconnected institutions in the panel.

From an estimation perspective, the network methodology we propose can be interpreted as a regularization procedure of a moderately large dimensional covariance matrix. As forcefully put forward, among others, by Ledoit and Wolf (2004), precise estimation of a covariance

matrix is challenging when the number of series considered is large, and in these cases co-variance regularization can provide substantial gains. Accordingly, we design our predictive

evaluation exercise as follows. On each day from the 1st of January 2008 to the end of the sample, we compute the bank credit risk network concentration matrix estimator and then use the corresponding covariance estimator as a forecast of the covariance of the bootstrap idiosyncratic intensities changes of the following day. We estimate the network using a two year rolling window and we use the BIC to choose the shrinkage parameter κ. We use the

multivariate QLIKE loss function of Patton (2011) (see also Patton and Sheppard, 2009) using the first differences of the bootstrap idiosyncratic default intensities for the next day to measure forecast accuracy. The network estimator is compared against two alternatives: the in-sample realized covariance estimator based on the ξ differences; and the in-sample realized covariance estimator with all its off-diagonal entries truncated to zero.[6]

The second predictive exercise consists in assessing whether our methodology is successful in identifying highly interconnected institutions. We design this second predictive evaluation exercise as follows. On the last day of each month from January 2008 to the end of the sample, we compute the bank credit risk network concentration matrix estimator as well as the out-of-sample realized covariance of the idiosyncratic intensities over the following 12 months.

We estimate the network using a two year rolling window and we use the BIC to choose the shrinkage parameter κ. Next, we compute the weighted degree (sum of squared partial correlations) of each entity based on the in-sample network estimator and out-of-sample realized covariance matrix. Finally, we compute Spearman's rank correlation between the in-sample and out-of-sample degree. As an alternative to the network estimator we consider

the degrees computed on the basis of the in-sample realized covariance.

We report the results of the two predictive analysis in Table 4 where we show the average prediction losses and the average rank correlation over the full out-of-sample period as well

---

[6]Additionally in the Online Appendix we compare the performance of the rolling network estimation strategy considered here against a simple time-varying covariance model based on exponential smoothing.

as different subsamples. The ranking between the different estimators is clear. The realized network estimator produces more accurate forecasts over the entire sample as well as all the different subsamples. In particular we note that the average QLIKE loss is higher in periods of larger distress and that in particularly it drops substantially in the last subsample, the one that corresponds with the initiation of the OMT program of the ECB. In the covariance forecasting exercise we point out that a Diebold-Mariano predictive ability tests conveys that the realized network out-performs the diagonal covariance estimator at the 1% significance level. On the other hand, there is no evidence of out-performance of the unconstrained estimator. We remark that further refinements on the predictive ability of these different methodologies may be obtained by using a proper time-varying parameter model and this may be explored in future research.

<div align="center">TABLE 4 ABOUT HERE</div>

Overall, results convey that the bank credit risk network methodology is not only useful to represent the dependence structure of idiosyncratic shocks but it also provides more precise estimates of the covariance of the idiosyncratic shocks when the conditional dependence structure of these shocks is sufficiently sparse.

4. Conclusions

The recent financial crisis in Europe has forcefully shown the potential impact of high levels of interconnectedness in the financial system and brought forward a renewed interest in monitoring the current state of interconnections in the financial sector as well as identifying its most central institutions. However, in the absence of regulatory data, this remains a challenging task empirically.

In this work we introduce a new approach to estimating interconnectedness through extending the standard reduced form credit risk model commonly used in finance. In our proposed model, interdependence in credit risk stems from two components: common exposure to a systematic factor and pairwise dependence among idiosyncratic shocks. We then

use this methodology to study credit risk interdependence in a sample of financial institutions located in ten selected Eurozone countries over an eight-year period from 2006 to 2013.

We find that the network channel captures a substantial amount of interdependence on top of what is explained by systematic factors. A cross-sectional analysis shows that the large financial institutions from the largest Eurozone economies are the most interconnected institutions in the panel. We find evidence of linkages both intra- and inter-country even after controlling for each respective sovereign. Furthermore, the structure of the credit risk network is fairly concentrated, implying that the correlation between the idiosyncratic shocks of any two institutions may be non-negligible and adverse shocks may spread rapidly through the network.

A time-varying analysis reveals that distressed financial institutions become hubs in the credit risk network during crisis times. These hub institutions can spread negative shocks within the network through an increased number and strength of their linkages, making their monitoring an important task for financial stability.

Our results have important implications from a systemic risk perspective. First, we find that frequently used models which do not account for idiosyncratic dependence can severely underestimate the joint default probability of two institutions. Second, our time-varying analysis shows that the network potential for spreading contagion of single institutions can vary over time, creating the need for constant monitoring of systemic importance which is not just defined by more stable characteristics such as bank size.

Abbassi, P., Brownlees, C., Hans, C., and Podlich, N. (2017). Credit Risk Interconnectedness: What Does the Market Really Know? *Journal of Financial Stability*, **29**, 1–12.

Adrian, T. and Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, **106**, 1705–1741.

Alfaro, R. and Drehmann, T. (2009). Macro stress tests and crises: what can we learn? *BIS Quarterly Review, December 2009*, pages 29–41.

Allen, F. and Gale, D. (2000). Bubbles and Crises. *Economic Journal*, **110**, 236–255.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003). Modeling and Forecasting Realized Volatility. *Econometrica*, **71**(2), 579–625.

Ang, A. and Longstaff, F. (2013). Systemic Sovereign Credit Risk: Lessons from the U.S. and Europe. *Journal of Monetary Economics*, **60**(5), 493–510.

Banerjee, O. and Ghaoui, L. E. (2008). Model Selection Through Sparse Maximum Likelihood Estimation for Multivariate Gaussian or Binary Data. *Journal of Machine Learning Research*, **9**, 485–516.

Barigozzi, M. and Brownlees, C. (2018). NETS: Network Estimation for Time Series. *Journal of Applied Econometrics*, **34**(3), 347–364.

Bauwens, L., Laurent, S., and Rombouts, J. (2006). Multivariate garch models: a survey. *Journal of Applied Econometrics*, **21**, 79–109.

Betz, F., Hautsch, N., Peltonen, T. A., and Schienle, M. (2016). Systemic Risk Spillovers in the European Banking and Sovereign Network. *Journal of Financial Stability*, **25**(C), 206–224.

Brownlees, C. and Engle, R. (2017). SRISK: A Conditional Capital Shortfall Measure of Systemic Risk. *Review of Financial Studies*, **30**(1), 48–79.

27

Brownlees, C., Nualart, E., and Sun, Y. (2018). Realized Networks. *Journal of Applied Econometrics*, **33**(7), 986–1006.

Caballero, R. and Simsek, A. (2013). Fire Sales in a Model of Complexity. *Journal of Finance*, **68**(6), 2549–2587.

Calomiris, C. and Mason, J. (2003). Fundamentals, Panics, and Bank Distress During the Depression. *American Economic Review*, **93**, 1615–1647.

De Santis, R. (2015). A measure of redenomination risk. *ECB Working Paper 1875 / April 2015*.

Demirer, M., Diebold, F. X., Liu, L., and Yilmaz, K. (2017). Estimating global bank network connectedness. *Journal of Applied Econometrics*, **33**(1), 1–15.

Dempster, A. (1972). Covariance Selection. *Biometrics*, **28**(1), 157–175.

Diebold, F. X. and Yılmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, **182**, 119–134.

Duffie, D. and Singleton, K. J. (1999). Modeling Term Structures of Defaultable Bond Yields. *Review of Financial Studies*, **12**(4), 687–720.

Fan, J., Feng, Y., and Wu, Y. (2009). Network Exploration via the adaptive and SCAD penalties. *The Annals of Applied Statistics*, **3**(2), 521–541.

Fontana, A. and Scheicher, M. (2016). An analysis of euro area sovereign CDS and their relation with government bonds. *Journal of Banking and Finance*, **62**(C), 126–140.

Friedman, J., Hastie, T., and Tibshirani, R. (2011). Sparse Inverse Covariance Estimation with the Graphical Lasso. *Biostatistics*, **9**(3), 432–441.

Gai, P., Haldane, A., and Kapadia, S. (2011). Complexity, Concentration and Contagion. *Journal of Monetary Economics*, **58**(5), 453–470.

Hale, G., Kapan, T., and Minoiu, C. (2016). Crisis Transmission in the Global Banking Network. Technical report, IMF Working Paper 16/91.

Hautsch, N., Schaumburg, J., and Schienle, M. (2014). Financial Network Systemic Risk Contributions. *Review of Finance*, **19**(2), 685–738.

Kritzman, M., Yuanzhen, L., Page, S., and Rigobon, R. (2011). Principal Components as a Measure of Systemic Risk. *The Journal of Portfolio Management*, **37**(4), 112–126.

Lando, D. (1998). On cox processes and credit risky securities. *Review of Derivatives Research*, **2**, 99–120.

Ledoit, O. and Wolf, M. (2004). A Well-Conditioned Estimator for Large–Dimensional Covariance Matrices. *Journal of Multivariate Analysis*, **88**, 365–411.

Longstaff, F., Mithal, S., and Neis, E. (2005). Corporate yield spreads: Default risk or liquidity? new evidence from the credit default swap market. *The Journal of Finance*, **60**(5), 2213–2253.

Merton, R. C. (1980). On estimating the expected return on the market : An exploratory investigation. *Journal of Financial Economics*, **8**(4), 323–361.

Mistrulli, P. (2011). Assessing financial contagion in the interbank market: Maximum entropy versus observed interbank lending patterns. *Journal of Banking and Finance*, **35**(5), 1114–1127.

Oh, D. H. and Patton, A. (2018). Time-Varying Systemic Risk: Evidence From a Dynamic Copula Model of CDS Spreads. *Journal of Business and Economics Statistics*, **36**(2), 181–195.

Patton, A. J. (2011). Data-based ranking of realised volatility estimators. *Journal of Econometrics*, **161**(2), 284–303.

Patton, A. J. and Sheppard, K. (2009). Evaluating volatility and correlation forecasts. In T. G. Andersen, R. A. Davis, J.-P. Kreiß, and T. Mikosch, editors, *Handbook of Financial Time Series*. Springer, Berlin.

Peng, J., Wang, P., Zhou, N., and Zhu, J. (2009). Partial Correlation Estimation by Joint Sparse Regression Models. *Journal of the American Statistical Association*, **104**, 735–746.

Rapach, D. E., Strauss, J. K., and Zhou, G. (2013). International Stock Return Predictability: What Is the Role of the United States? *Journal of Finance*, **68**(486), 1633–1662.

Silvennoinen, A. and Tersvirta, T. (2008). Multivariate garch models. SSE/EFI Working Paper Series in Economics and Finance 669, Stockholm School of Economics.

Stein, J. (2012). Monetary Policy as Financial Stability Regulation. *Quarterly Journal of Economics*, **127**(1), 57–95.

Tian, S., Yang, Y., and Zhang, G. (2013). Bank Capital, Interbank Contagion and Bailout Policy. *Journal of Banking and Finance*, **37**(8), 2765–2778.

Tibshirani, R. (1996). Regression shrinkage and selection via the LASSO. *Journal of Royal Statistical Society B.*, **58**(1), 267–288.

Yuan, M. and Lin, Y. (2007). Model selection and estimation in the Gaussian graphical model. *Biometrika*, **94**(1), 19–35.

Table 1: EUROPEAN FINANCIAL INSTITUTIONS

Abbr.	Name of the institution	Abbr.	Name of the institution
	Austria		**Italy**
WAG	Austria	ASG	Assicurazioni Generali
ERS	Erste Bank Group	LAV	Banca Nazionale de Lavoro
RAI	Raiffeisen Bank International	LEA	Banca Italease S.p.A.
	Belgium	MED	Mediobanca S.p.A.
FOR	Fortis N.V. / Ageas Holding N.V.	INT	Intesa Sanpaolo S.p.A
KBC	KBC BANK	MIL	Banca Popolare di Milano
DEX	Dexia Crédit Local	MPS	Banca Monte dei Paschi di Siena S.p.A
	Germany	UDP	Unione di Banche Italiane SCPA
ALL	Allianz AG	POP	Banco Popolare S.C.
DBA	Deutsche Bank AG	UNI	UniCredit S.p.A.
COM	Commerzbank AG		**Netherlands**
DBZ	DZ Bank	SAN	Espirito Santo Financial Group, SA
MRV	Münchner Rückversicherung	AEG	Aegon NNV
NLB	Norddeutsche Landesbank	NIB	NIBC Bank NV
HSH	HSH Nordbank	RAB	Rabobank
LBW	Landesbank Baden-Württemberg	SNS	SNS Bank NV
BLB	Bayerische Landesbank	SNS	SNS Bank NV
LBB	Landesbank Berlin	VAN	F. van Lanschot Bankiers NV
LHT	Landesbank Hessen - Thüringen	ING	Ing Bank NV
HRE	Hypo Real Estate Holding AG	ABN	Abn Amro Bank NV
WLB	West LB / Portigon AG	ROD	Rodamco Europe NV
	Spain	ACH	Achmea Holding NV
SAN	Banco Santander S.A.		**Portugal**
BBV	Banco Bilbao Vizcaya Argentaria S.A.	BCO	Banco Comercial Portugues, SA
CAB	Caja de Ahorros y Pensiones de Barcelona	CAI	Caixa Geral de Depositos, SA
MPM	Caja de Ahorros y Monde de Piedad de Madrid	BPI	Banco Portugues de Investimiento
POP	Banco Popular Espanol, S.A.		

Abbr.	Name of the institution
	Spain (cont'd)
PAS	Banco Pastor, S.A. (PAS)
SAB	Banco de Sabadell, S.A.
SAB	Banco de Sabadell, S.A.
MED	Caja de Ajorros del Mediterraneo
CAV	Caixa de Ahorros de Valencia, Castelln y Alicante / Bancaja
INT	Bankinter, S.A.
SOP	Sophia GE
	France
AXA	AXA France
BQE	Banque Fédérative du Crédit Mutuel
PEU	Banque PSA Finance
BNP	BNP Paribas
AGR	Crédit Agricole
CIC	Crédit Industriel et Comercial
LYO	Crédit Lyonnais
WEN	Wendel
SOC	Société Générale
	Greece
NBG	National Bank of Greece
EFG	EFG Eurobank Ergasias S.A.
GEC	Gecine
	Ireland
ANG	Anglo Irish Bank
DEP	Depfa PLC
GOV	Governor and Company of the Bank of Ireland
NAT	Irish Nationwide Bank
ILP	Irish Life and Permanent

The table reports the list of financial institutions included in the sample with their country of origin and abbreviation.

Table 2: NETWORK SUMMARY

	Austria	Belgium	Germany	Spain	France	Greece	Ireland	Italy	Netherlands	Portugal	Links
Austria	70.4	2.0	11.2	6.6	0.9	0.7	1.5	1.0	5.0	0.6	24.95
Belgium	12.8	17.4	7.2	2.0	31.8	0.1	2.7	6.8	15.9	3.5	3.92
Germany	3.7	0.4	59.3	2.2	12.2	0.0	0.4	6.4	12.3	3.1	76.44
Spain	1.8	0.1	1.9	83.4	2.3	0.0	0.7	4.1	1.1	4.6	89.78
France	0.2	1.3	9.5	2.1	69.6	0.1	1.2	3.8	11.5	0.8	98.80
Greece	17.0	0.2	0.7	1.9	8.4	42.6	0.3	14.6	12.0	2.3	1.10
Ireland	2.9	0.8	2.5	4.7	9.5	0.0	65.5	2.5	4.4	7.0	12.87
Italy	0.4	0.4	6.5	4.9	5.0	0.2	0.4	76.2	3.4	2.6	74.76
Netherlands	2.6	1.3	19.7	2.1	23.8	0.3	1.2	5.4	41.5	2.1	47.68
Portugal	0.3	0.3	5.5	9.5	1.7	0.1	2.1	4.4	2.3	73.7	43.23

The (i, j) entry of the table reports the sum of the edge weights between country i and j relative to the total sum of weights of country i. The last column of the table reports the sum of edge weights of each country. The edge weights are defined as the squared partial correlation coefficients implied by the network estimator introduce in Section 2 estimated using the Adaptive GLASSO algorithm over the entire sample.

Table 3: CENTRALITY

	Weighted	Unweighted
1	SAN (es)	AGR (fr)
2	BBV (es)	COM (de)
3	SAN (pt)	SOC (fr)
4	BCO (pt)	ING (nl)
5	AGR (fr)	SAN (es)
6	BNP (fr)	DBA (de)
7	LYO (fr)	LYO (fr)
8	SOC (fr)	BBV (es)
9	MRV (de)	MRV (de)
10	MPS (it)	ALL (de)

The table shows the ranking of the top 10 most central banks by eigenvector centrality. The eigenvector centrality algorithm is based on the adjacency matrix implied by the bank credit risk network estimator estimated over the entire sample period using the Adaptive GLASSO procedure presented in Section 2. Two different versions of the adjacency matrix are used to construct the rankings. The weighted rankings are based the weighted version of the adjacency matrix of the bank credit risk network where the squared partial correlation coefficients are used as link weights. The unweighted rankings are based on the unweighted version of adjacency matrix of the bank credit risk network.

Table 4: OUT-OF-SAMPLE EVALUATION

	QLIKE			Rank Corr	
Sample	Network	Sample Cov	Diag.	Network	Sample Cov
2008-01-01 − 2008-08-01	12.147	12.504	39.604	0.251	0.200
2008-08-01 − 2010-04-01	51.674	52.779	139.408	0.392	0.392
2010-04-01 − 2012-09-01	9.486	11.539	95.626	0.278	0.205
2012-09-01 − 2013-12-31	-20.058	-16.095	49.046	0.012	0.012
2008-01-01 − 2013-12-31	14.197696	15.822	78.852	0.295	0.254

The left panel of the table shows the out-of-sample QLIKE of the network covariance estimator, realized covariance and diagonal realized covariance estimator over different subsamples. The right panel of the table shows Spearman's rank correlation between the in-sample and out-of sample weighted degree of the financial entities in the panel over different subperiods. The in-sample weighted degree distribution is estimated using the network covariance estimator and the sample covariance estimator.

Figure 1: Systematic Default Intensity

Systematic Default Intensity, Core

Systematic Default Intensity, Periphery
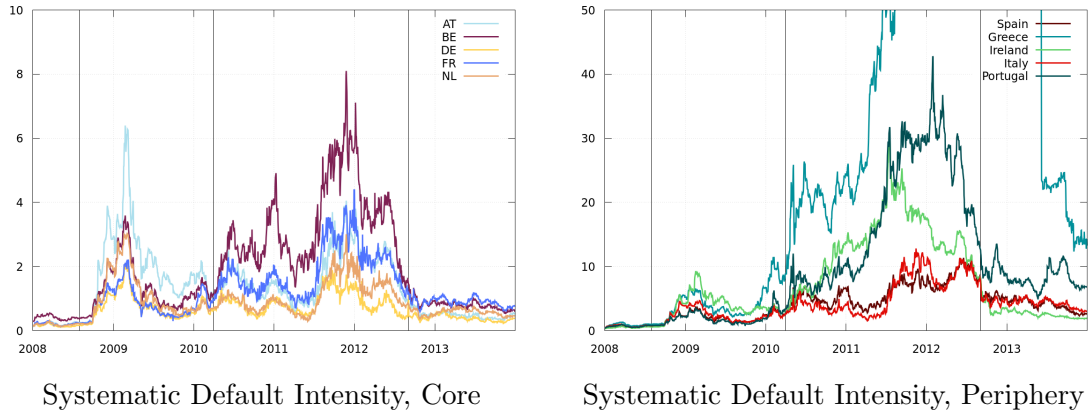
The figure shows the time series of sovereign risk-neutral default intensities for core countries bootstrapped from CDS prices of 1-year, 3-year, 5-year, 7-year and 10-year maturity on the basis of of the credit risk model introduced in Section 2 estimated over the full sample. The intensities are rescaled such that the level of one corresponds to a 1% probability of default over the next year. The figure displays the risk-neutral intesities for core counties in the left panel (Austria, Belgium, France, Germany, Netherlands) and for periphery countries in the right panel (Spain, Greece, Ireland, Italy, Portugal) countries.

Figure 2: Bank Credit Risk Network
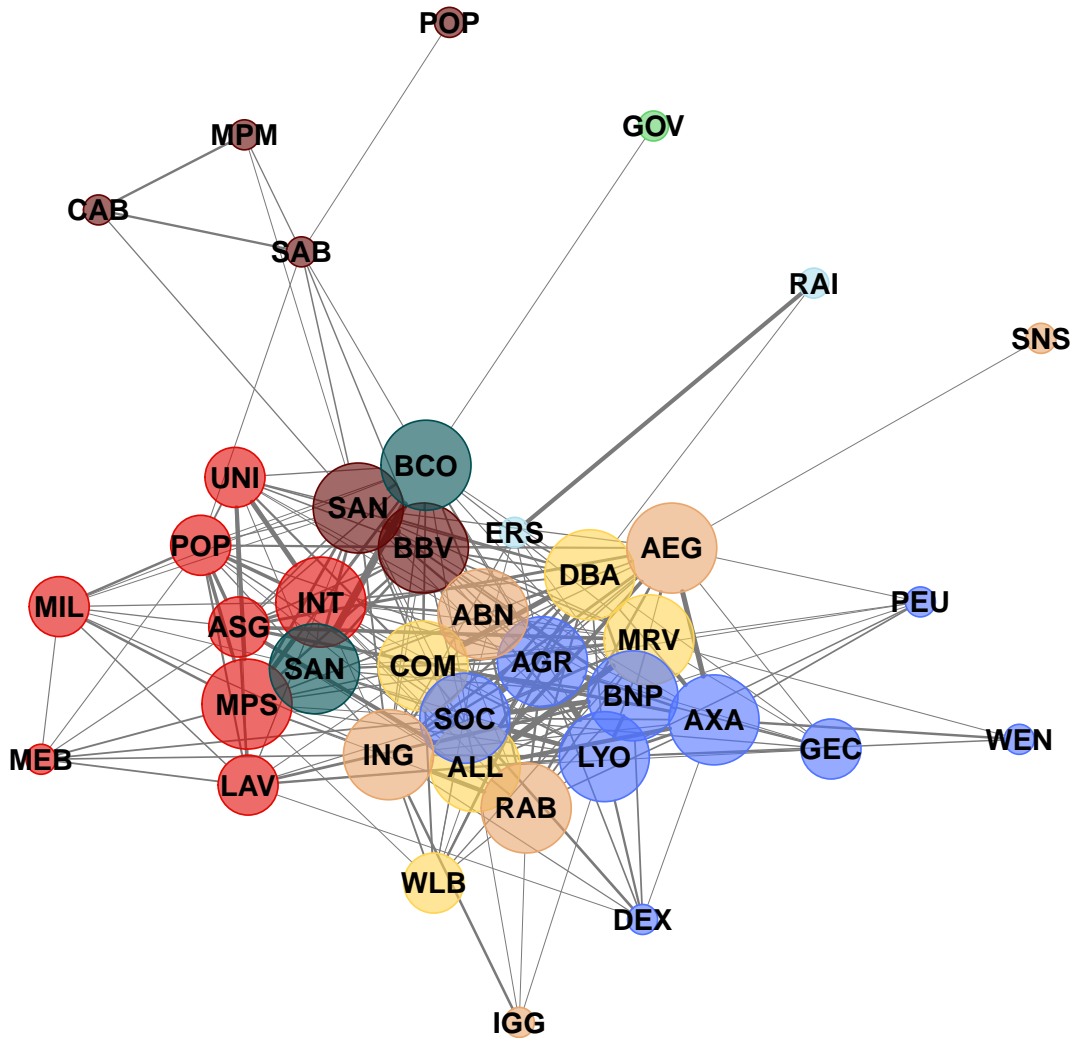
The figure shows the bank credit risk network. The bank credit risk network is defined as the set of non-zero partial correlation relations implied by the concentration matrix of the idiosyncratic defaul intensities. This is estimated over the entire sample period using the Adaptive GLASSO procedure presented in Section 2. The vertex size is proportional to the degree of each financial entity the vertex color is set according to the bank's country of origin. The edge width is proportional to the square of the corresponding partial correlation.

Figure 3: DEGREE AND PARTIAL CORRELATION DISTRIBUTION

Weighted Degree Distribution   Partial Correlation Distribution
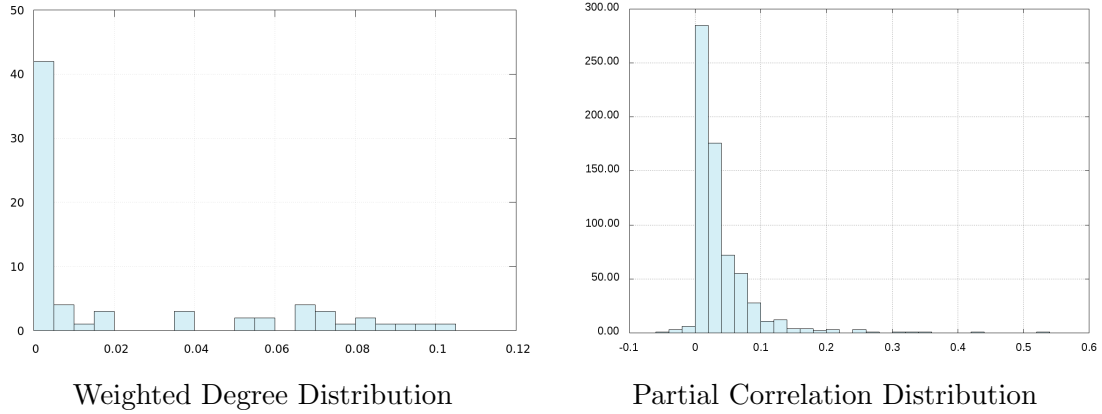
The figure shows histograms of the weighted degrees (left) and non-zero partial correlations (right). The weighted degree of a financial entity in the panel is defined as the sum of the square partial correlations relative to that entity. The partial correlations are the ones implied by the concentration matrix of the idiosyncratic default intensities estimated over the entire sample period using the Adaptive GLASSO procedure presented in Section 2.

Figure 4: Network Density

The figure shows the monthly time-series of the average weighted degree together with the 5% and 95% quantiles of the weighted degree distribution obtained from the rolling estimates of the bank credit risk network from January 2008 to December 2013. The partial correlations are the ones implied by the concentration matrix of the idiosyncratic default intensities estimated by the Adaptive GLASSO procedure presented in Section 2 using a two-year rolling window.

Figure 5: FACTOR AND NETWORK R SQUARES

Factor R^2, Core

Factor R^2, Periphery

Network R^2, Core

Network R^2, Periphery
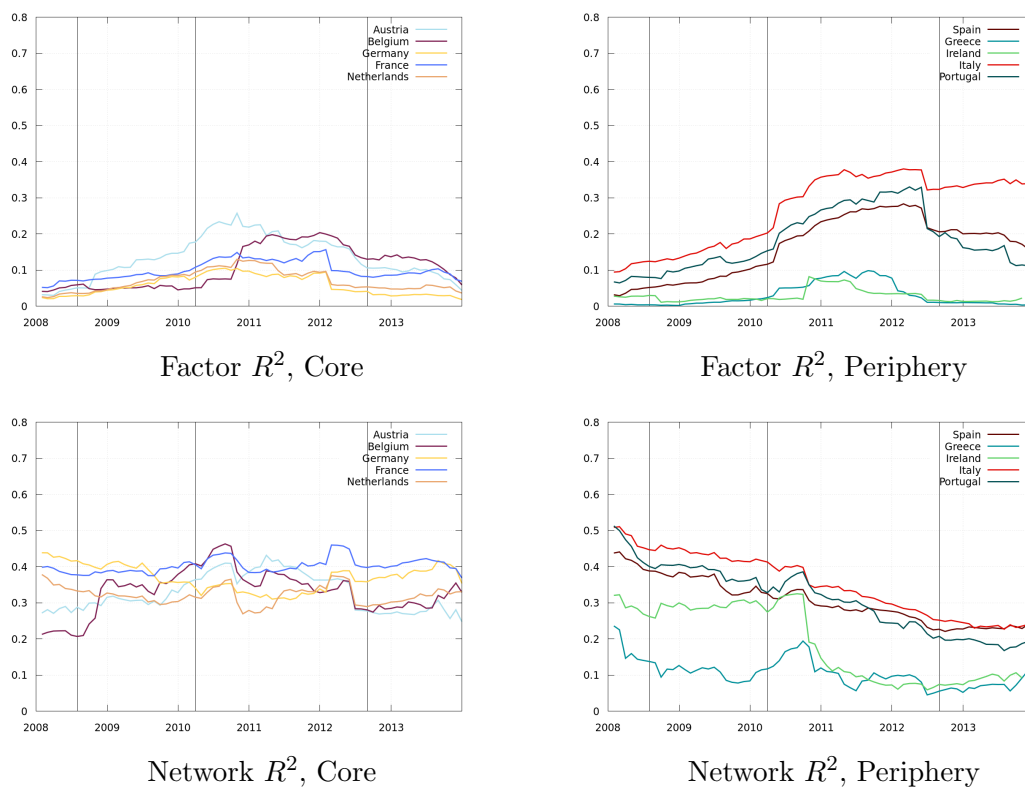
This figure shows the time series of factor and network R squares for core and periphery countries. The factor R^2 of a bank is defined as the R^2 of the regression of its log-CDS spread difference on the log-CDS spread difference of its respective sovereign. The network R^2 of a bank is defined as the additional R^2 obtained by adding as explanatory variables of the previous regression all the log-CDS spread differences of all the neighbouring banks detected in the bank credit risk network. The factor and network R^2 are computed on the basis of two-year rolling estimates. The factor and network R^2 are computed on the basis on rolling estimates in order to have a time series of values.

Figure 6: DEGREE FOR INDIVIDUAL BANKS

COM   MPS

The figure shows the monthly time-series of the weighted degree (thick line/left y-scale) as well as the total number of non-zero partial correlations coefficients (thin line/right y-scale) obtained by the rolling estimates of the bank credit risk network from January 2018 to December 2013 for Commerzbank (COM) and Monte Dei Paschi (MPS). The partial correlations are the ones implied by the concentration matrix of the idiosyncratic default intensities estimated by the Adaptive GLASSO procedure presented in Section 2 using a two-year rolling window.