

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

Measuring Interconnectedness between Financial Institutions with Bayesian Time-Varying Vector Autoregressions

Geraci, Marco Valerio; Gnabo, Jean Yves

Published in:

Journal of Financial and Quantitative Analysis

DOI:

[10.1017/s0022109018000108](https://doi.org/10.1017/s0022109018000108)

Publication date:

2018

Document Version

Early version, also known as pre-print

[Link to publication](#)

Citation for published version (HARVARD):

Geraci, MV & Gnabo, JY 2018, 'Measuring Interconnectedness between Financial Institutions with Bayesian Time-Varying Vector Autoregressions', *Journal of Financial and Quantitative Analysis*, vol. 53, no. 3, pp. 1371-1390. <https://doi.org/10.1017/s0022109018000108>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Measuring interconnectedness between financial institutions with Bayesian time-varying vector autoregressions

Marco Valerio Geraci^{†,‡} and Jean-Yves Gnabo^{*‡}

[†]*ECARES, Université libre de Bruxelles, CP139 1050, Belgium,*

Email: mgeraci@ulb.ac.be

[‡]*CeReFiM, Université de Namur, Rempart de la Vierge, 5000, Belgium,*

Email: jygnabo@fundp.ac.be

December 23, 2015

Abstract

We propose a market-based framework that exploits time-varying parameter vector autoregressions to estimate the dynamic network of financial spillovers. In a series of simulation exercises, we show that our framework performs better than the classical approach of Granger causality testing over rolling windows. We apply it to all financials listed in the Standard & Poors 500 index and uncover interesting features at the individual, sectorial and system wide level. This includes a gradual decrease in interconnectedness after the crisis, not observable using the classical rolling window approach, and more stable interconnectedness-based rankings that can be used for monitoring purposes.

Keywords: financial interconnectedness, time-varying parameter, systemic risk

JEL Classification: G01, G18, C32, C51, C63

*We acknowledge support of the Communauté française de Belgique through the Projet d'Actions de Recherche Concertées, 13/17-055 grant.

We are grateful to Matteo Barigozzi, Sophie Béreau, Oscar Bernal, Christiane Baumeister, Monica Billio, Michele Lenza, Matteo Luciani, Nicolas Scholtes, David Veredas and Kamil Yilmaz for helpful comments and suggestions. We also thank participants at the Conference on Financial Risk & Network Theory in Cambridge, the seminar participants at ECARES, naXys and the University of Fordham for their feedback and valuable discussion.

I. Introduction

Since the outbreak of the financial crisis, academics and regulators have developed statistical measures to monitor financial interconnectedness and systemic risk (see among others ?????????). According to the Financial Stability Board (FSB) and the Basel Committee for Banking Supervision (BCBS), financial interconnectedness is a determinant of systemic risk and is defined as the network of contractual obligations which can potentially channel financial distress (?).¹ By their very nature, these contractual obligations continually change over time and are often contingent on the prevailing economic environment. Therefore, statistical measures of financial interconnectedness should take into account the inherently dynamic nature of such an environment.

In this paper we develop a framework, based on Bayesian estimation of time-varying parameter vector autoregressions (TVP-VAR), that models the dynamic nature of connections between financial institutions. The framework allows connections to evolve gradually through time as opposed to the classical approach, which favours sudden, often unjustified, changes in interconnectedness. Paired with graph theory, our framework allows us to reconstruct a continuously evolving network of directed spillover effects. We use our framework to study the evolution of interconnectedness between publicly listed US financial institutions over the past two decades.

Various studies have combined statistical measures of association (e.g. correlation, Granger causality, tail dependence) with network techniques, in order to map and analyse financial interconnectedness (???). However, these standard statistical measures presuppose that the inferred relationships are time-invariant over the sample used for the estimation. To retrieve a dynamic measure of interconnectedness, the usual approach has been to divide the original sample period into multiple subsamples and calculate these statistical measures over rolling windows of data.

We argue that this is potentially unsuitable if the system studied is deemed to be time-varying. By relying on short subsamples, rolling windows lower the power of inference and induce dimensionality problems. Moreover, the rolling window approach is known to be

¹Other determinants that were identified are cross-jurisdictional activity, size, substitutability and financial institution infrastructure and complexity.

susceptible to outliers because, in small subsamples, these have a larger impact on estimates (?). On the other hand, choosing longer windows will lead to estimates that are less reactive to change, biasing results towards time-invariant connections. Thus, the rolling window approach requires the researcher to choose the window size, which involves a trade-off between precision and flexibility (?). The choice of window size is critical and can lead to different results regarding interconnectedness.

Our framework intuitively parallels the Granger causality approach for estimating interconnectedness. We first estimate TVP-VARs (as by ? and ?) on our set of stock returns. Precisely, these can be either a sequence of bivariate TVP-VARs, if we are interested in pairwise unconditional connections, or, if we are interested in conditional connections, a single TVP-VAR containing the stock returns of all firms. We then adopt the methodology of ? to test, in every time period, whether two financial institutions are related. The test is based on the whole sample of observed data and therefore infers connections ex-post.

The major novelty of our framework is that we recover a network of financial spillovers that is entirely dynamic. To do so, we make the modeling assumption that the connection between any two institutions evolves smoothly through time. We consider this assumption reasonable for three main reasons. First, since connections are the result of many financial contracts, it seems natural that they evolve smoothly rather than abruptly. Second, our assumption implies that the best forecast of a connection in the future is the state of that connection today. This is consistent with the notion of forward-looking prices. Third, our assumption allows for high flexibility and for the data to speak for itself.

In a series of simulation exercises, we show that our framework performs well, compared to the classical rolling windows approach, in two respects. First, with respect to the testing efficacy, i.e., the efficacy in determining the existence (or non-existence) of a connection between two institutions. Second, with respect to the precision in estimating the parameters that are common to both approaches, i.e., the strength of connections.

We apply our framework to estimate the US network of financial spillovers. To do this, we use monthly stock price data for all financial institutions listed from 1990 to 2014 on the Standard & Poor's 500 index, including firms that have since gone defunct. We uncover four main results.

First, the empirical application reveals the defiances of the rolling window approach and the gains of adopting the time-varying parameter framework in the network context. Measures of connectivity and centrality computed under the rolling window approach are more volatile because of the approach's high sensitivity to extreme observations. This gives the impression that interconnectedness rose after the crisis, whereas the time-varying parameter framework reveals a decrease, concurring with the credit freeze that occurred.

Second, at the individual institution level, the empirical application identifies some interesting cases. American International Group (AIG) was found to be the largest propagator of financial spillovers, highlighting the potential widespread influence its default could have caused. This result backs the Federal Reserve Bank's decision to rescue AIG. Quite to the contrary, Bear Stearns did not play a major role in the propagation of spillovers but rather was very receptive of incoming spillovers. This could explain why its collapse did not represent a systemic event.

Third, we show that interconnectedness-based rankings, computed using the rolling window approach, are extremely volatile and unlikely to be useful for policy decisions. On the other hand, the time-varying parameter framework produces more stable rankings that could be appropriate for timely monitoring.

Fourth, we examine interconnectedness between the four financial sectors included in our study and find that banks and insurance companies were the largest contributors of financial spillovers. The real estate sector, composed primarily of real estate investment trusts, was the most influenced by these spillovers even though it was found to be also contributing in propagating spillovers at moderately high levels. The combination of these two factors could have enhanced the real estate sector's role in distributing spillovers.

In the next section, we will briefly review the recent literature on financial interconnectedness measures and how this relates to our modelling choices. The remainder of the paper is then structured as follows. In Section ??, we go through the estimation framework, partly reviewing the above-mentioned Granger causality approach. In Section ??, we illustrate the results of the simulation exercises and in Section ?? we discuss the empirical application. We draw our conclusions in Section ??.

II. Literature Review

The greatest difficulty faced by research into financial interconnectedness and stability has been associated with data availability. Little data regarding bank cross-exposures is available, primarily because of banks' confidentiality concerns. For this reason, part of the literature has focused on simulating the data using network models of contagion (e.g., [Battiston et al., 2012](#)) and agent-based models (e.g., [Cacciari and Vespignani, 2011](#)).

Other studies have attempted to recover this data in various ways. [Battiston et al. \(2012\)](#) for example, had to impose the assumption that interbank loans and deposits are equally spread over banks to recover the network of German interbank exposures using balance sheet information. [Battiston et al. \(2012\)](#) resort to a unique dataset of the Bank of International Settlements on cross-border financial flows intermediated by national banking systems. This data is country level, thus cannot allow bank-specific monitoring, and as it concerns flows rather than exposures, it provides more information on liquidity than actual contagion of losses.

Even if data on cross-exposures were available, it might not be sufficient to fully capture interconnectedness between banks. This is because banks can also be connected by off-balance sheet items, such as derivatives, or indirectly by common portfolio holdings ([Battiston et al., 2012](#)). These instruments can potentially have contagion effects far more severe than defaults on bilateral credit exposures ([Battiston et al., 2012](#)). Moreover, concentrating solely on interbank exposures confines the analysis to depository institutions, while neglecting important players in the financial system, such as brokers, insurers and real estate companies who are not part of this market.

In order to overcome such data limitations, a growing area of research with pioneering contributions from [Battiston et al. \(2012\)](#), [Battiston et al. \(2012\)](#) and [Battiston et al. \(2012\)](#), has turned to market data to recover measures of interconnectedness.² These measures are based on the notion of forward-looking markets and the observed phenomenon of comovement in stock prices during times of crisis. The idea is that stock prices include all available information up to the most recent point in time. Then, given the network of contractual obligations, financial distress travelling from one financial institution to another should be reflected by comovement in stock prices.

The first studies to examine comovement in stock prices under a network perspective

²[Battiston et al. \(2012\)](#) analyse different market-based measures of interconnectedness and systemic risk.

come from the field of econophysics. These studies showed that, by interpreting simple linear correlation as a Euclidean distance, it is possible to retrieve a network with a meaningful economic taxonomy (?). That is to say, graphs, evidencing economically founded connections between companies, can be recovered using solely stock price data (?). These studies highlight the incredibly rich informational content of market data and support later studies on market-based measures of interconnectedness.

A subgroup of studies extends the above-mentioned literature on correlation networks. These studies investigate how to avoid the possibility of indirect connections appearing in the graph. The stock prices of two financial institutions may be correlated primarily because, for example, both institutions are influenced by a third. Some studies have chosen to use partial correlation measures to define the network of relations between financial institutions (e.g. ?). This approach can lead to dimensionality problems, which have been addressed in a variety of ways, among which LASSO-based techniques (???).

Both correlation and partial correlation are linear operators measuring the average relationship between two variables. Therefore, comovements that might exist in the extreme tails of stock prices, i.e., during exceptional circumstances, will only be partially captured by correlation. In order to measure connectedness during times of financial distress, when stock prices move far away from their mean, some studies have proposed to examine extreme tail dependence networks. ?, ? and ?, for example, show that interconnectedness measures based on tail dependence networks are good predictors of financial distress.

Alternative non-linear measures have been proposed by ? and ?. The latter use correlations between volatility shocks to measure interconnectedness in the U.S. economy, whereas the former examine interconnectedness between four U.S. financial sectors (commercial banks, investment banks, hedge funds and insurance companies) by proposing a state-dependent sensitivity Value-at-Risk (SDSVaR) model. The SDSVaR model is comparable to our proposed framework as it allows for three states of connections between sectors, according to whether financial markets are in a volatile, normal or tranquil condition. However, our TVP-VAR is more flexible as it allows connections vary freely through time.

Correlation, partial correlation, and extreme dependence are contemporaneous measures of comovement and as such are bidirectional. As a result, it is not possible to distinguish

the *influencing* financial institution from the *influenced* when using these measures (?). To address this issue, studies have turned to Granger causality, a directed measure of interconnectedness.

Granger causality, also known as Wiener-Granger causality after ? and ?, is a well known statistical technique used to infer causality between variables, based on the concepts of *predictability* and *precedence*. It has been adopted in a wide variety of fields ranging from macroeconomics (??) to biology (?) and the neurosciences (?). Especially within the natural sciences, the notion has been paired with graph theory to recover the causal network structure among a set of variables (?).

Recently, these techniques have also been applied in finance to estimate the network of *spillovers* between financial companies using stock market data.³ ?, for example, analysed the interconnectedness of banks, brokers, hedge funds and insurers to find that interconnectedness had grown strongly since the 1990s and peaked during the financial crisis. Moreover, they found that banks and insurers have a more central role than hedge funds and brokers in the network. They ascribe this to the rise of a “shadow hedge-fund system” with banks and insurers taking on risks more appropriate for hedge funds.

As was explained in the introduction, a major issue with all the aforementioned studies is that the underlying connections that obtain among financial institutions plausibly vary over time, whereas the statistical measures adopted (whether correlation, tail-dependence or Granger causality) are designed for time-invariant connections. These studies examined the evolution of interconnectedness through time by applying the given statistical measures to subsequent windows of observations. As will become clear in the next section, this results in a trade-off whereby larger windows lead to greater precision but less flexibility.

In order to allow time-variation while avoiding the classical rolling windows approach, the evolution of connections over time must be modelled. We contribute to the existing literature by proposing an original empirical framework that performs well in addressing this critical issue.

Finally, it is worth mentioning a study put forward ? who estimate the time-varying

³These studies have implicitly assumed that information about distress connecting two financial institutions is impounded in prices with some delay (consistent with the lag structure of the Granger causal relationship).

dependence between European sovereign credit spreads with a lagged spacial model. In their case, the (spacial) network used is time-invariant and exogenously given by the BIS database. On the other hand, we propose a statistical framework for retrieving the underlying time-varying network from market data.

III. Methodology

As mentioned above, we base our framework for measuring interconnectedness on Granger causality testing as in ?. We first go briefly through the classical approach to estimating networks by Granger causality testing. We then present our framework for estimating networks with time-varying connections.

A. Estimating networks by canonical Granger causality

We define a static network graph by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} represents the set of nodes (vertices) and \mathcal{E} is the set of edges (connections). We denote the total number of nodes in the network by $N = |\mathcal{V}|$.

We allow any node $i \in \mathcal{V}$, to be endowed with a measurable and stationary time series of some observable attribute, $x_i = \{x_{i,t}\}_{t=-p}^T$. In our study, nodes represent financial institutions, while the attribute associated with each node will be the stock return of the financial institution.⁴

For any two nodes $i, j \in \mathcal{V}^2$, we draw a directional edge $i \rightarrow j$ if x_i causes x_j in the sense of ?. That is, x_i Granger causes x_j if the past of x_i can improve the in-sample forecast of x_j above and beyond the in-sample forecast based on the past of x_j alone.⁵

?? operationalises the definition very succinctly. We can estimate the network by *conditional* Granger causality testing. This takes into account all variables in a vector autore-

⁴We control for return autocorrelation and heteroskedasticity in our empirical application.

⁵Similarly, if x_j Granger causes x_i we draw an edge in the opposite direction, $j \rightarrow i$. If x_j Granger causes x_i and x_i Granger causes x_j then there is a bi-directional edge, denoted $i \leftrightarrow j$. If no Granger causality is found in neither direction, there is no edge between i and j . Self edges are not allowed.

gression (VAR) system:

$$(1) \quad x_t = c + \sum_{s=1}^p B_s x_{t-s} + u_t,$$

where B_s is $N \times N$ matrix of coefficients and $x_t = [x_{1t}, \dots, x_{Nt}]'$. We assume the zero mean errors in u_t to be uncorrelated with x_{t-s} , $s = 1 \dots p$, and serially uncorrelated between themselves although they can be contemporaneously correlated.

By denoting the (j, i) element of B_s by $B_s^{(j,i)}$, the test for Granger causality from x_i to x_j is based on the null hypothesis that:

$$(2) \quad H_0 : B_1^{(j,i)} = B_2^{(j,i)} = \dots = B_p^{(j,i)} = 0.$$

This can be done using a Wald test by stacking all the coefficients c, B_1, \dots, B_s in the vector B and rewriting equation (??) as

$$(3) \quad x_t = X_t' B + u_t,$$

where $x_t = [x_{1,t}, \dots, x_{N,t}]'$ and $X_t' = I_N \otimes [1, x'_{t-1}, \dots, x'_{t-p}]$.

Then, the null hypothesis defined in (??) reduces to

$$(4) \quad H_0 : \tilde{A} B = \mathbf{0}_{p \times 1},$$

where, \tilde{A} is a $p \times N(1 + Np)$ matrix of zeros and ones placed in accordance to (??). By rearranging the elements of \tilde{A} , we can test in the same way the opposite relation, $j \rightarrow i$.

Conditional Granger causality can become unfeasible if the number of parameters to be estimated, $N(1 + Np)$, exceeds the number of observations, T . This is even more of a problem when using rolling windows to capture time-varying relationships. In fact, by relying on shorter subsamples, the rolling window approach makes it more difficult for this condition to be met. Our framework is less susceptible to this dimensionality issue, because we exploit the whole length of observations.

In order to deal with the dimensionality issue, several studies have relied on *pairwise*

Granger causality testing (e.g. ?). Pairwise Granger causality testing between two variable x_i and x_j does not condition on other variables in the system. Then, rather than conducting the test on the full VAR given in (??), the network is estimated by recursively testing Granger causality, for all pairs, (i, j) , on many bivariate VARs. This of course, is susceptible to omitted variable problems if there are indirect effects between variables in the system, e.g., if one variable is acting as a common cofactor.

B. Estimating networks with time-varying connections

The hypothesis given by equation (??) is conditional on all observations collected from the start of the sample up to period T . If the direction of causality were to change within this time period, the test inference might be affected.

One possibility to account for such changes would be by performing the Granger causality test on a rolling window basis (as, for example, is done by ?, or, with different techniques, by ? and ?). This solution, however, introduces several other issues.

First, by relying on shorter samples, the rolling window approach reduces the power of statistical inference. This means that the Granger causality test will be less reliable in small samples. Moreover, rolling windows can be problematic when testing conditional Granger causality with a large VAR. Effectively, if the number of parameters exceeds the number of observations we fall in the so-called curse of dimensionality. Lastly, the shorter the windows, the more important outliers will be within the subsample. This leads to volatile parameter estimates and, ultimately, imprecise inference on connections (?).

On the other hand, larger windows are associated with more precise inference at the cost of restricting the variability of estimates. Larger subsamples allows greater statistical confidence for inference but at the same time causes estimates to digest change only slowly, as new observations enter the window. This imposes a trade-off when selecting the appropriate window size (?).

Another issue is that Granger causality detected using rolling windows would have to be associated with a given point in time. It is not clear if this point in time should be the end of the window, the middle, or the beginning, as different conventions can be used (?). Generally, rolling window estimates are one-sided backward looking, that is, they are

associated with the end of the window. We keep this convention throughout the rest of the paper.

With a slight abuse of notation, we parallel the Granger causality test in a TVP-VAR setting. In order to do so, we rewrite the unrestricted regression, given in equation (??), as

$$(5) \quad x_t = X_t' B_t + u_t \quad u_t, \sim \mathcal{N}(0, R),$$

where u_t is now assumed to be normally distributed with zero mean and variance R .

Now the coefficients relating any two nodes can vary through time, meaning that the strength of the connection linking any two nodes can also vary. Moreover, time-varying parameters can allow for the connection to even disappear completely and then reappear.

The hypothesis of a connection, from nodes j to node i , existing at time t is then,

$$(6) \quad H_{0,t} : \tilde{A} B_t = \mathbf{0}_{p \times 1}.$$

where \tilde{A} is defined as for (??).

In order to test this hypothesis, we adopt Bayesian techniques for estimating (??) using the Kalman filter and smoother.⁶

Equation (??) can be interpreted as the measurement equation of a state space model.

As is done in the macroeconomic literature (e.g., ???), we let parameters evolve according to a driftless random walk. The state equation of the model is thus,

$$(7) \quad B_{t+1} = B_t + v_{t+1} \quad v_t, \sim \mathcal{N}(0, Q),$$

where we assume that ϵ_t and v_s are independent at all t and s .

It is worth mentioning the general implications of the above assumption for our framework of assessing time-varying connections in a network.

First, equation (??) entails that changes in the strength and direction of causality occur

⁶Non-Bayesian techniques, such as those developed by ?, can also be adopted but can lead to identifiability issues in large systems, with a large number of parameters to estimate.

smoothly. It seems natural that connections between financial institutions evolve gradually rather than in sharp breaks. As explained in the introduction, these connections are given by several layers of contractual obligations between companies. These range from cross-exposures to common portfolio holding and shared derivative contracts. It makes sense that the status of these connections emerge gradually from the behaviour of agents, rather than bursting suddenly.

Second, the transition equation assumes that all shifts in parameters are permanent. That is, once change has occurred it is not reversible to the old state, and the new state is expected to hold forever. This is in stark contrast to alternative models of Granger causality that assume Markov-switching regimes (e.g., ??). Here, states transition back and forth between two outcomes (unless one of the states is absorbing). If the transition probability is constant, the effect that is forecast will be a weighted average of the two states.

The random walk specification offers greater flexibility because it does not limit changes to a restricted number of states. It is well known that Markov-switching models do particularly well if the underlying process has sharp breaks, but not so well if this is not the case. ? show, by simulations, that a time-varying parameters model with random walk dynamics does fairly well even if the true underlying process is Markov-switching.

Finally, some studies from the macro literature choose to impose a stability condition so as to exclude explosive paths for B_t . This is done by assuming that the probability density of B_t takes a value of zero when the roots of the TVP-VAR polynomial are inside the unit circle. Others, such as ?, do not include this condition, because they assume that the model holds for a finite period of time and not forever.⁷ Given that we expect a VAR model on stock returns to have small coefficients (in absolute terms), we follow ? and do not impose a stability condition.

C. Inference

The model is estimated by simulating the posterior distribution of the parameters as in the Bayesian tradition. Given conditionally conjugate priors, the posterior conditional distribution of the states is normal. We retrieve a sample of the joint posterior probability

⁷A random walk process hits any upper or lower bound with probability one.

of the parameters by using the Markov chain Monte Carlo (MCMC) algorithm proposed by ?. An overview of the prior specification and the sampling algorithm are given, respectively, in Appendix A and Appendix B.

The sampling algorithm relies on the Kalman filter and smoother to update the conditional means and variances of the states, meaning that the whole sample of data is used for estimation. In turn, this means that the time dependent hypothesis given in equation (??) is tested conditionally on all the observations observed up to period T . We find that this is an informationally more efficient approach than is recursive testing by rolling windows.

In the Bayesian tradition, hypothesis testing is done by Bayes factor. This gives the odds in favor of the restricted model, implied by the null hypothesis given in equation (??), against the unrestricted model, given by equations (??) and (??). From the Bayes factor we can then retrieve the implied probability, for every point in time t , that equation (??) holds (see ?).

To do so, we first collect the history of the state parameter B_t up to period T in a vector B^T . Denote by Ψ all parameters except the states B^T that is, R , Q and given hyperparameters governing the priors. Then Bayes factor can be written as:

$$K = \frac{p(x^T | M_0)}{p(x^T | M_1)} = \frac{\int \int p(B^T, \Psi | M_0) p(x^T | B^T, \Psi, M_0) dB^T d\Psi}{\int \int p(B^T, \Psi | M_1) p(x^T | B^T, \Psi, M_1) dB^T d\Psi}$$

where M_0 refers to the restricted model imposed under the null hypothesis given in equation (??) and M_1 refers to the unrestricted alternative.

We follow ? who show a convenient way of calculating Bayes factor by the Savage-Dickey density ratio (SDDR).⁸

Let A be a $p \times TN(1 + Np)$ matrix of zeros with ones placed appropriately such that $AB^T = \mathbf{0}_{p \times 1}$ corresponds to the restrictions specified in equation (??). That is, AB^T sets cross-parameters, relating x_i with x_j at time t , to zero and leaves all remaining parameters unrestricted.⁹

⁸The proof for the Savage-Dickey density ratio is given in ?.

⁹If we are conducting the pairwise recursive testing with a bivariate TVP-VAR, A would be a $p \times 2T(1+2p)$ matrix.

Then under the assumption that

$$(8) \quad p(\Psi \mid AB^T = \mathbf{0}_{p \times 1}) = p_0(\Psi)$$

Bayes factor will be given by the SDDR,

$$(9) \quad K = \frac{p(AB^T = \mathbf{0}_{p \times 1} \mid x^T)}{p(AB^T = \mathbf{0}_{p \times 1})}.$$

The assumption given by equation (??) requires the prior for Ψ in the restricted model, $p_0(\Psi)$, to be the same as the prior in the unrestricted model evaluated at the point where the restriction holds, $p(AB^T = \mathbf{0}_{p \times 1} \mid x^T)$. This is amply satisfied if the same prior is used in the restricted and unrestricted model for the parameters that are common to both models.

As explained by ?, an estimate of the nominator in equation (??) can be calculated using the simulations from the conditional posterior $p(B^T \mid x^T, \Psi)$. Given a conjugate Normal conditional prior for B_t , the conditional posterior is known to have a Normal distribution. In turn, this implies that $p(AB^T = \mathbf{0}_{p \times 1} \mid x^T, \Psi)$ is also Normal. By simulating from $p(AB^T = \mathbf{0}_{p \times 1} \mid x^T, \Psi)$ using a Gibbs sampler and averaging across draws, we obtain an estimate of the posterior probability that the null hypothesis holds, $\hat{p}(AB^T = \mathbf{0}_{p \times 1} \mid x^T)$. Similarly, the denominator can be simulated by using a sequential sampler on the conditional priors $p(AB^T \mid \Psi)$, and calculating the average across all draws, $\hat{p}(AB^T = \mathbf{0}_{p \times 1})$.

The estimated Bayes factor \hat{K} gives us the odds that the null hypothesis of no connection at time t holds. The implied probability is then just $\hat{K}/(1 + \hat{K})$. This can be subjected to a threshold to make the decision whether to accept or reject the null hypothesis.

Unlike classical frequentist testing, Bayes factor weighs evidence in favor of M_0 and M_1 equally. Effectively, the threshold is a filtering mechanism and a higher threshold leads to a more dense network with more links. For our simulation exercises and for the empirical application, we make use of a neutral threshold of 50%.

In additional results not shown here, we have repeated our analysis using different thresholds. We have also looked at the possibility of using $1 - \hat{K}/(1 + \hat{K})$ as the weight associated with the given directed link. This is the probability of having a link at period t . The conclusions from these additional results are in line with those shown in the rest of the paper

and are available upon request to the authors.

IV. Simulations

In a series of simulation exercises, we assessed the ability of our time-varying framework to infer the small causal network given in Figure ???. A similar exercise was conducted by ? using the same network. We chose this particular network because it is sparse and sparsity is an observed attribute of financial networks.¹⁰

The network’s underlying system is given by

$$\begin{aligned}
 x_{1,t} &= \alpha_{1,t} + \phi_{1,t} x_{1,t-1} + \epsilon_{1,t} \\
 x_{2,t} &= \alpha_{2,t} + \phi_{2,t} x_{2,t-1} + \beta_{2,1,t} x_{1,t-1} + \epsilon_{2,t} \\
 x_{3,t} &= \alpha_{3,t} + \phi_{3,t} x_{3,t-1} + \beta_{3,1,t} x_{1,t-1} + \epsilon_{3,t} \\
 x_{4,t} &= \alpha_{4,t} + \phi_{4,t} x_{4,t-1} + \beta_{4,1,t} x_{1,t-1} + \beta_{4,5,t} x_{5,t-1} + \epsilon_{4,t} \\
 x_{5,t} &= \alpha_{5,t} + \phi_{5,t} x_{5,t-1} + \beta_{5,4,t} x_{4,t-1} + \epsilon_{5,t}
 \end{aligned}$$

where, $[\epsilon_{1,t} \dots \epsilon_{5,t}]' = \epsilon_t \sim \mathcal{N}(\mathbf{0}, \Omega)$ and $\Omega = \tau I_5$ where τ was set to 0.01. We chose to limit the autoregressive component of the process to one lag, as is done by ?, so as to keep the simulation exercises computationally manageable.

[FIGURE ?? ABOUT HERE]

We performed three different experiments in which the model parameters were allowed to vary according to the following processes:

1. Deterministic fixed constants drawn, at the beginning of each simulation, from a standard uniform distribution.
2. Markov switching between 0 and a random constant drawn, at the beginning of each simulation, from a standard uniform distribution.
3. Smoothly time-varying, according to a unit root process.

¹⁰The interested reader may refer to ? for an in-depth discussion on sparse networks in finance.

For each experiment, we ran 100 simulations each of which involved $T = 300$ time periods after “burning” the first 1000.¹¹

We ran our framework so as to infer all possible connections between variables. Paralleling pairwise and conditional Granger causality, this was done in two alternative ways: 1) by recursively using a bivariate TVP-VAR between every pair of variables, and 2) by running the full TVP-VAR and testing connections conditional on all variables of the system. For means of comparison, we also carried out the same simulation exercises, using the classical approach of Granger causality testing (by pairwise and conditional VARs) over rolling windows. For this, we set the level of significance of the tests to 5%.

We assessed the performance of our framework with respect to three standard measures: the mean-squared error (MSE) of the regression parameter estimates, the receiver-operator characteristic (ROC) curve and the precision-recall (PR) curve. We outline how each measure is computed below.

The MSE of the regression parameter estimates is found by taking the sum, across all time periods, of the squared difference between the estimated regression parameters and the true parameters. This sum is then averaged across all simulations. The formula for the MSE of the cross-parameters $\beta_{i,j,t}$ is given by

$$MSE_C^{TVP} = \sum_{(i,j) \in C} \sum_{t=1}^T (\hat{\beta}_{i,j,t}^{TVP} - \beta_{i,j,t})^2 / T,$$

where $C = \{(2, 1), (3, 4), (3, 5), (4, 1), (4, 5), (5, 4)\}$.

For the classical Granger causality approach, parameters are estimated by ordinary least squares (OLS) over rolling windows of size $w(s)$. Then the MSE is calculated as

$$MSE_{C,w(s)}^{RW} = \sum_{(i,j) \in C} \sum_{t=w(s)+1}^T (\hat{\beta}_{i,j,t}^{RW} - \beta_{i,j,t})^2 / (T - w(s)),$$

where $w = [20, 30, \dots, 200]'$ is the vector of rolling window sizes used to compute estimates. The step size for the rolling window calculation is set to 1.

¹¹A burn-in was not used for the third experiment because such a long time period (1300 observations) would not withstand the stability condition easily.

To allow a fairer comparison between MSE^{RW} and the MSE_C^{TVP} across the same time periods, we look at

$$MSE_{C,w(s)}^{TVP} = \sum_{(i,j) \in C} \sum_{t=w(s)+1}^T (\hat{\beta}_{i,j,t}^{TVP} - \beta_{i,j,t})^2 / (T - w(s)),$$

The cross-parameters can be interpreted as the strength of the connection between two nodes. The larger is $\beta_{i,j,t}$, the stronger is the influence of node j on node i in period t . Having a precise estimate of the cross-parameter is not only important to determine if a link exists or not. It is also important if we are interested in determining the strength of the financial spillover between institutions.¹²

We also compared the performance of our time-varying parameter framework with that of the classical Granger causality approach, by means of the (ROC) and (PR) curves.

The ROC curve plots the true positive rate (TPR) against the false positive rate (FPR). In our case, a positive refers to the existence of a connection between the two nodes in question. Then the TPR is the ratio of the number of correctly estimated connections to the number of existing connections. On the other hand, the FPR is the ratio of incorrectly estimated connections to the number of non-existing connections. A high performing test would combine low FPR with high TPR and therefore have a ROC curve in the upper-left corner of the chart.

For time-varying parameter estimation, the ROC curve was calculated using the implied probability from the estimated Bayes factor (as was explained in Section ??), whereas for the classical Granger causality approach, the p-value was used. All possible connections were tested and results were aggregated over all time periods and across all simulations.

The PR curve plots the precision, also known as the positive predictive value, against the recall, i.e., the TPR. The precision is the fraction of correctly classified positives, i.e., the ratio of connections correctly inferred to the total number of connections inferred. There exists a one-to-one relationship between the ROC and precision-recall curve. If for a given experiment a curve dominates in ROC space, then it will also dominate in precision-recall

¹²Moreover, $\beta_{i,j,t}$ can be interpreted as a weight on the link and used to construct an alternative weighted network, one based on connection strength rather than, as mentioned in the preceding section, the implied probability of a connection existing. In this sense, a low MSE would imply a more precise estimated network.

space (?). However, looking at the PR curve can provide additional insight in situations like ours, where the number of negatives exceeds by far the number of positives.¹³ A high performing test would combine high precision with high TPR and therefore have a PR curve in the upper-right corner of the chart.

A. Experiment 1

The case in which parameters are constant through time corresponds to time-invariant connections. The causal network then corresponds exactly to Figure ?? for all periods $t \in (1, 300)$. The direction and strength of the relationship between any two nodes does not change throughout the 300 time periods simulated.¹⁴

For the first experiment, we fix all regression parameters to constants drawn at the beginning of each simulation.

$$\alpha_{i,t} = a_i, \quad \phi_{i,t} = f_i, \quad \beta_{i,j,t} = b_{i,j}, \quad \forall t \in [0, T],$$

where we draw parameters from a standard uniform distribution at the beginning of each simulation, $a_i, f_i, b_{i,j} \sim \mathcal{U}(0, 1)^3$ for $i = 1, \dots, 5$ and $(i, j) \in C$.

Before launching the burn-in, we checked the stability of the system by making sure the largest eigenvalue of the system was within the unit circle. If this was not the case, we re-drew all parameters simultaneously until an appropriate draw was found. This ensures stationarity of the variables and avoids explosive processes.

[FIGURE ?? ABOUT HERE]

The left panel of Figure ?? shows $MSE_{w(s)}^{RW}$ (light dashed) and $MSE_{w(s)}^{TVP}$ (bold solid). Notice that $MSE_{w(s)}^{RW}$ is downward sloping in window size. This is expected because larger windows lead to more precise estimates at the expense of less variability. Since the underlying parameters are constant, $MSE_{w(s)}^{RW}$ decreases quickly with the window size. $MSE_{w(s)}^{TVP}$ does not possess this downward sloping property because the time-varying parameter estimation is

¹³Even in the constant case, when all connections in our toy-network are present, we only have six connections out of a total of $N^2 - N = 5^2 - 5 = 20$ possible connections.

¹⁴Note however, that the strength of the relationship varies between simulations.

an ex-post procedure. This means that the whole length of the sample is used for estimation, unlike the rolling window approach.

Results show that the time-varying framework performs better than the classical rolling window approach, whether estimation is pairwise (top-left chart) or conditional on the other variables of the system (bottom-left chart). The time-varying parameter framework does well because the Kalman filter and smoother, used for the sampling algorithms, find the best fit with the minimum predictive variance. Even when large rolling windows are used (above 100 observations) the time-varying parameter framework performs comparably well to the classical approach.

We report the performance of our time-varying parameter framework in terms of ROC and PR curves, respectively given in the middle and right panels of Figure ?? (bold solid). We also show the ROC and PR curves associated with the classical Granger causality approach (light dashed) estimated by rolling windows of size 200.¹⁵ This corresponds to two-thirds of the observations in each simulation. It was also one of the best performing window sizes across all three experiments.

The ROC curve for pairwise estimation (top-middle chart) shows that time-varying parameter testing performs comparably well compared to the classical approach with rolling windows. In particular, it does slightly better than the classical approach at low combinations of FPR and TPR, whereas it performs slightly worse at higher combinations of the two. On the other hand, the classical approach with rolling windows appears to perform consistently better than the time-varying parameter approach when testing conditional relationships (bottom-middle chart).

In terms of the PR curve, pairwise time-varying testing does well at combinations with high precision and low recall (upper-right chart). Here the curve associated with time-varying parameters (bold solid) is above that associated with the classical rolling windows approach (light dashed). However, at higher combinations of precision and recall, the two approaches perform similarly, with the PR curve for the classical approach slightly above the time-varying counterpart. As was found for the ROC curve, the PR curve also shows

¹⁵ROC and PR curves calculated at other window sizes have been omitted for space concerns but are available from the authors upon request.

that the time-varying parameter approach performs almost uniformly worse in conditional testing (lower-right chart).

B. Experiment 2

For the second experiment, the cross coefficients, $\beta_{i,j,t}$ with subscripts $(i, j) \in C$, of the system outlined above were assumed to follow a switching process defined as

$$\beta_{i,j,t} = \begin{cases} 0 & s_t^{ij} = 0 \\ b_{i,j} & s_t^{ij} = 1 \end{cases}$$

where $b_{i,j}$ is drawn at the start of the simulation from a standard uniform distribution.

As in the first experiment, the intercept terms $\alpha_{i,t}$ and autoregressive coefficients $\phi_{i,t}$ were drawn from a standard uniform distribution at the beginning of each simulation and were assumed to be constant through time.

Let s_t^{ij} follow a first order Markov chain with the following transition matrix:

$$\mathbf{P} = \begin{bmatrix} \mathbb{P}(s_t^{ij} = 0 \mid s_{t-1}^{ij} = 0) & \mathbb{P}(s_t^{ij} = 1 \mid s_{t-1}^{ij} = 0) \\ \mathbb{P}(s_t^{ij} = 0 \mid s_{t-1}^{ij} = 1) & \mathbb{P}(s_t^{ij} = 1 \mid s_{t-1}^{ij} = 1) \end{bmatrix} = \begin{bmatrix} p_{00} & p_{10} \\ p_{01} & p_{11} \end{bmatrix}$$

where we set $p_{00} = 0.95$ and $p_{11} = 0.90$.

Effectively, the transition matrix holds the probabilities of a link appearing and disappearing between any two nodes $i, j \in C$. In the matrix, p_{00} is the probability of no link occurring between two nodes at time t , given that the two nodes were disconnected at time $t - 1$. Similarly, p_{11} represents the probability of there being a link between two nodes at t , given that these two nodes were already connected at $t - 1$.¹⁶

[FIGURE ?? ABOUT HERE]

The top-left and bottom-left charts of Figure ?? show the MSE for estimates found by, respectively, pairwise estimation and full conditional estimation. Results show that

¹⁶Again, when simulating the parameters, the stability condition was checked such that the system would not allow for explosive processes.

estimation using the time-varying parameter framework is more precise, especially when compared to Granger causality testing carried out with small sized windows. This confirms the Monte Carlo simulation results of ?.

The ROC curves in the top-middle and bottom-middle charts of Figure ?? highlight the gain obtained by using the time-varying parameter framework for detecting connections. For pairwise testing the ROC curve lies completely above the corresponding curve for classical Granger causality testing with rolling windows of size 200. Although less pronounced, one can notice a similar improvement in detecting connections even when estimating conditional connections using the complete full TVP-VAR with all five variables included.

Similarly, the PR curves (top-right and bottom-right charts of Figure ??) show substantial improvements using the time-varying parameter framework when testing connections pairwise. In the case of conditional testing, the PR curve associated with the time-varying parameter framework appears above the corresponding curve for classical testing in areas of the chart with low recall, while it lies slightly below for areas with higher recall. This indicates that, in this case, our framework performs better when the network is sparse.

C. Experiment 3

For the third experiment, the parameters of the system were allowed to evolve according to the following random walk process,

$$\begin{aligned}\alpha_{i,t+1} &= \alpha_{i,t} + v_{i,t+1}^\alpha \\ \phi_{i,t+1} &= \phi_{i,t} + v_{i,t+1}^\phi \\ \beta_{i,j,t+1} &= \beta_{i,j,t} + v_{i,t+1}^\beta,\end{aligned}$$

where, $v_{i,t} = [v_{i,t}^\alpha, v_{i,t}^\phi, v_{i,t}^\beta]'$ and $v_{i,t} \sim \mathcal{N}(\mathbf{0}, \Gamma)$. In turn, the variance of the parameters was set to

$$\Gamma = q^2 \times \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix}$$

where, $q^2 = 0.0002$.

The variance of the parameters was set such that cross coefficients would be more variable than the autoregressive parameters and, in turn, the autoregressive parameters would be more variable than the intercept terms.¹⁷

[FIGURE ?? ABOUT HERE]

As can be noticed from the top-left and bottom-left charts in Figure ??, the time-varying parameters approach results in more precise parameter estimates, in terms of MSE, than the classical approach using rolling windows. Notice that, quite contrary to what was found in Experiment 1, the precision of the classical approach does not increase with window size. Rather, gains from using larger windows reverse for windows of more than 60 observations, showing that precision actually worsens when windows are too large. This highlights the trade-off between higher confidence but less flexibility given by larger windows. The effect is stronger when the network is estimated by recursive pairwise testing using bivariate TVP-VARs, but is still present even when the full TVP-VAR is used for conditional testing.

In terms of ROC curves, shown in the top-middle and bottom-middle charts, the curve referring to the time-varying parameters approach lies completely above the one referring to classical Granger causality testing by rolling windows of size 200.¹⁸

The same result was obtained for the PR curves shown in the top-right and bottom-right charts. Here we notice that there is not much gain from using pairwise testing with recursive bivariate TVP-VARs rather than conditional testing with the full TVP-VAR. On the other hand, in the classical approach, we see that using conditional testing leads to a higher PR curve. However, this curve continues to remain below the curve associated with the time-varying parameter framework meaning that our framework is better at detecting connections across all combinations of precision and recall.

These results, together with those from Experiment 1 and 2, show that our proposed framework provides better estimates and better inference on connections when the true underlying process is changing through time. Our framework performs well whether these changes are abrupt or smooth. On the other hand, when the underlying process is constant,

¹⁷As in the preceding two experiments, when simulating the parameters of the system, unstable simulations were discarded.

¹⁸The results continue to hold for other window sizes.

our framework performs slightly worse (but in some cases just as well) as the classical approach.

V. Empirical application

For the empirical application we are interested in estimating the time-varying network of spillovers between all financial institutions listed on the Standard & Poor's 500 index (S&P 500). We then proceed to analyse interconnectedness at the individual institution level, at the sectorial level and at the overall system level. Such exercise can provide a clear view on the spread of spillovers and the dynamics of interconnectedness during the crisis, which can be particularly important for risk management and supervisory purposes.

As mentioned previously, we adopt a market-based approach by using the stock prices of the sampled institutions. Several studies have performed similar exercises by adopting a rolling window approach (e.g. ???). On the other hand, we apply TVP-VARs to retrieve a completely dynamic network, exploiting the whole sample of data. Moreover, as will be explained further below, we propose a solution for dealing with appearing and disappearing nodes, which represent firms that cease to exist because of a merger or because of bankruptcy.

A. Data and Method

We selected financial institutions with Standard Industrial Classification (SIC) codes from 6000 to 6799 that were components of the S&P 500 from January 1990 to December 2014. For these companies we collected the stock price at monthly close from Thomson Reuters Eikon for the same time period. Initially the sample contained 182 firms but was reduced to 155 after constraining our analysis to stocks with at least 36 monthly observations. Table ?? shows the financial institutions included in the final sample.

The sample can be further subdivided into four sectors based on the SIC code of the companies. These were identified as banks (SIC codes 6000 to 6199), broker/dealers (SIC codes 6200 to 6299), insurers (SIC codes 6300 to 6499) and real estate companies (SIC codes 6500 and 6799). The final sample included 71 banks, 21 brokers/dealers, 40 insurers and 23

real estate companies.

We searched through each company’s history using Factiva to retrieve dates of important events such as acquisitions, mergers or bankruptcy filings. This allowed us to understand why some stock price series were shorter than the whole sample, whether this was because of missing values or of some inherent event. We noticed some companies would inherit the stock price history of an older company as result of a merger. In these cases, in order to avoid multicollinearity problems, we dropped the stock price that the newer company had in common with the older company.

We define monthly stock returns for company i at month t as

$$r_{i,t} = \log p_{i,t} - \log p_{i,t-1},$$

where $p_{i,t}$ is the stock price of company i at the end of month t . The monthly frequency makes it possible to reduce the amount of noise in the data. Data at the intra-daily or even daily frequency reveals a higher number of linkages, because stocks are more susceptible to market shocks that lead to a higher degree of co-movement.

In a second step we detrended stock returns, as is done by [?](#) , in order to control for heteroskedasticity. The standardised series used is then denoted by

$$\tilde{r}_{i,t} = r_{i,t} / \hat{\sigma}_{i,t},$$

where $\hat{\sigma}_{i,t}$ is the estimated according to GARCH(1,1) for every stock in our sample.

Our sample is unbalanced, with several stock time series substantially shorter than the complete sample period. This is due to mergers and bankruptcies. Effectively, when these events occur, it is as if the concerned nodes disappear. In order to deal with such issues, we adopted an approach that can account for time-varying nodes as well as time-varying connections.

We did this by making use of pairwise time-varying connectivity. For each pair, we estimated a bivariate TVP-VAR with one lag, taking into account the longest common time period of available data between any given pair.¹⁹ Using the posterior simulation algorithm

¹⁹According to the BIC criterion, one lag was found to be the most appropriate lag specification for most

described in Appendix A.2 (with 7000 draws, discarding the first 1000 and thinning every 5 draws), we estimate Bayes factor relating to the null hypothesis of no connection between the pair of financial institutions at a given time period (as explained in Section ??). We then use a cut off threshold of 50% to make the decision of whether a connection exists at a given point in time between the pair.

We compare the dynamic network estimated with the TVP-VARs to the network obtained by using Classical Granger causality over rolling windows of data. For this, we take ? as the baseline using rolling windows of 36 months and a critical value of 5% for the Granger causality tests. As for the TVP-VARs, we apply the rolling window approach to the longest common sample per pair.

B. Network measures of connectivity, centrality and stability

In order to analyse the estimated time-varying network we study a set of centrality and connectivity measures.

To study the importance of individual firms in terms of interconnectedness, we computed their degree centrality measures. These give you an indication of importance of a given institution within the network. Since the network is directed, we can measure both the in-degree centrality as well as the out-degree centrality, respectively given by:

$$\text{In-Degree}_{i,t} = \frac{1}{(N_t - 1)} \sum_{j \neq i} (j \rightarrow i), \quad \text{Out-Degree}_{i,t} = \frac{1}{(N_t - 1)} \sum_{j \neq i} (i \rightarrow j)$$

In-degree measures the number of incoming connections of firm i as a ratio of all possible incoming connections. Effectively, it is the number of firms affecting firm i 's stock price. Therefore, it is a measure of firm vulnerability to stock spillovers.

Out-degree measures the number of outgoing links of firm i as a ratio of all possible outgoing links. Thus, a firm with many outgoing connections is influencing many of its neighbours, making it a propagator of spillovers. As suggested by ?, such financial institutions should be monitored closely because they are highly interconnected and potentially drivers of systemic risk.

stock returns.

We can rank institutions according to their in-degree and out-degree centrality, in order to get a sense of which institutions are most important in terms of incoming and outgoing linkages. Let $Z_{i,t}^{in}$ be the ordinal ranking of institution i at time t in terms of in-degree. Similarly, let $Z_{i,t}^{out}$ be its position in the out-degree ranking.

These rankings can potentially be used for monitoring purposes and to take preventive policy actions. For example, the Financial Stability Board (FSB) ranks financial institutions according to their systemic importance and uses this ranking to determine their additional loss absorbency requirements.²⁰ The rankings drawn by the FSB are not only based on interconnectedness but on other determinants of systemic risk such as size and leverage. Nonetheless, it is of interest to study the ranking recovered from our measure of interconnectedness based on degree-centrality.

In particular, ? and ? argue that the usefulness of rankings is severely limited if these are prone to frequent, drastic changes that lead to unmotivated excessive alarm. In an attempt to assess this for the rankings produced by the time-varying framework and the classical rolling window approach, we developed a series of stability measures.

We computed the *quadratic* ranking stability measure as

$$SI_Q^{in} = \sqrt{\sum_{t=2}^T \sum_{i=1}^{N_t} \frac{(Z_{i,t}^{in} - Z_{i,t-1}^{in})^2}{N_t(T-1)}}, \quad SI_Q^{out} = \sqrt{\sum_{t=2}^T \sum_{i=1}^{N_t} \frac{(Z_{i,t}^{out} - Z_{i,t-1}^{out})^2}{N_t(T-1)}}.$$

Similarly, we construct the *absolute* stability measure as

$$SI_A^{in} = \sum_{t=2}^T \sum_{i=1}^{N_t} \frac{|Z_{i,t}^{in} - Z_{i,t-1}^{in}|}{N_t(T-1)}, \quad SI_A^{out} = \sum_{t=2}^T \sum_{i=1}^{N_t} \frac{|Z_{i,t}^{out} - Z_{i,t-1}^{out}|}{N_t(T-1)}.$$

The above-mentioned stability measures are calculated to grasp the suitability, from a policy perspective, of rankings based on degree-centrality. As additional measures of stability, we also computed the average percentage of firms that kept their position in the ranking between adjacent time periods and average changes in the *composition* of the top 10 and top 20 firms in rankings.

²⁰Additional loss absorbency requirements will phase in starting in January 2016 with full implementation by January 2019 (see ?).

In terms of overall connectivity of the network, we computed the network density measure. This is the number of connections inferred at time t as a ratio of all possible connections and is given by

$$(10) \quad \text{Density}_t = \frac{1}{N_t(N_t - 1)} \sum_{i=1}^{N_t} \sum_{j \neq i} (i \rightarrow j),$$

where N_t is the number of nodes present in the network at time t .²¹ By measuring the total number of connections between all firms, this measure captures the level of integration of the financial system.

In order to study connectivity at the sectorial level, we computed inter-sectorial degree measures between every pair of sectors $(m, n) \in \{\text{Banks, Brokers, Insurance, Real Estate}\}$:

$$\text{Inter-Sector-Degree}_{m \rightarrow n, t} = \frac{1}{N_{m,t}(N_{n,t})} \sum_{i=1}^{N_{m,t}} \sum_{i \neq j} \left((i | m) \rightarrow (j | n) \right),$$

where $(i | m)$ denotes node i belonging to sector m and $N_{m,t}$ denotes the number of nodes belonging to sector m in period t . The measure conveys the level of integration between two sectors and in particular, the proportion of financial spillovers from one sector to another.

In a similar manner we computed the intra-sectorial degree as proportion of connections between all firms of a given sector:

$$\text{Intra-Sector-Degree}_{m,t} = \frac{1}{N_{m,t}(N_{m,t} - 1)} \sum_{i=1}^{N_{m,t}} \sum_{i \neq j} \left((i | m) \rightarrow (j | m) \right).$$

The measure quantifies the extent to which firms are financially integrated with the other firms in their given sector.

C. Interconnectedness of US financial institutions

We first compare the degree centrality measures found using our time-varying parameter framework to those found using the classical Granger causality approach with rolling windows. We then proceed to analyse the degree centrality for a set of selected companies.

²¹? refer to this measure as the Granger causal density.

Table ?? shows several summary statistics regarding the in- and out-degree centrality measures, computed under the two approaches, across the firms of the four financial sectors examined. The first noticeable feature is that the centrality measures computed using the time-varying parameter framework are less volatile than the centrality measures computed using the rolling window approach. The reason for this will become clearer further down, when we examine the firms' centrality measures individually.

[TABLE ?? ABOUT HERE]

Table ?? also highlights some interesting patterns regarding the financial sectors analysed in our sample. We can notice that the most interconnected sector in terms of out-degree was the banking sector. This sector is composed of several large commercial banks as well as depository institutions. The most interconnected sector in terms of in-degree was the real estate sector. This sector primarily groups real estate investment trusts. The same sector also has a high out-degree, which highlights the importance of these institutions in terms of propagating financial spillovers.

We selected individual companies from the FSB's list of systemically important financial institutions and systemically important insurers in order to examine the evolution of their degree centrality through time. Figures ?? to Figure ?? depict, respectively, the in- and out-degree of companies in our sample that were identified by the FSB as systemically important financial institutions and as systemically important insurers.

The measures computed from our time-varying parameter framework are shown in bold solid; whereas the light dashed lines depict those computed from the rolling window approach.

[FIGURE ?? TO ?? ABOUT HERE]

By visually inspecting each sub-figure, we can immediately notice that the two methodologies examined, the time-varying parameter framework and rolling window approach, produce very different pictures. In particular, the degree centrality measures computed from the rolling windows approach appear more volatile, corroborating the results of Table ?. Moreover, in several subfigures, such as those relating to the out-degree of Goldman Sachs and Morgan Stanley (Figure ??) and of American International Group (AIG, Figure ??)

or the in-degree of Prudential Financial (Figure ??) show a drop in centrality occurring simultaneously in October 2011.

This drop in the degree centrality is artificially created by the rolling window approach. It does not appear when our time-varying parameter framework is used to calculate the centrality measures. The drop occurs exactly 36 months after the after September 2008, the peak of the financial crisis. In October 2011, when the observations relating to the financial crisis exit the rolling window, the Granger causality approach picks up a strong decrease in the interconnectedness of these financial institutions. By shortening the window size or increasing it, the drop will occur sooner or after. This is one of the limitations of the rolling window approach mentioned earlier.

For AIG, for example, out-degree rose dramatically for during the financial crisis when it was essentially bailed out by the Federal Reserve Bank to avoid its collapse. However, the increase in out-degree began long before the financial crisis, as we can notice it had started rising substantially by end of 2005. On the other hand, the out-degree calculated using the classical rolling window approach gives a measure that jumps suddenly in September 2008 and falls drastically exactly 36 months later.

In addition to highlighting the limitations of the rolling window approach, the degree centrality figures reveal some interesting patterns in the interconnectedness of financial institutions identified as systemically important by the FSB. We can identify the role of these institutions, in terms of interconnectedness, prior and after the crisis.

Figure ?? shows that Bank of New York Mellon and JP Morgan Chase were large absorbers of financial shocks during and after the crisis. On the other hand, Figure ?? shows Goldman Sachs, Morgan Stanley and State Street had a very high outgoing connectivity especially around the crisis. This concurs with the findings of ? and ? who also identified these banks as very central.

The most influential financial institution in terms of out-degree centrality was found to be AIG (shown in Figure ??). This was true for both our time-varying parameter framework and the rolling window approach. During the peak of the crisis, AIG was connected to over 60% of the financial institutions in our sample, therefore playing a central role as a propagator of financial spillovers. This finding seems to support the Federal Reserve Bank's decision to

bailout AIG on 16 September 2008. Under a no-bailout scenario, the high interconnectedness of AIG could have resulted in an even more systemic crisis.

The other insurance companies in our sample deemed to be systemically important by the FSB, were Metlife and Prudential Financial. Quite to the contrary to AIG, these two appear to have a high in-degree (Figure ??) and therefore acted mainly as receivers of spillovers during the crisis. In particular, Metlife seems to have been influenced most by its neighbours just after September 2008.

We analysed four banks that were not included in the list of systemically important financial institutions by the FSB because these banks became defunct at some point during the financial crisis. These were Bear Stearns, Lehman Brothers, Merrill Lynch and Wachovia. Their in and out-degree are shown, respectively, in Figures ?? and ?? (below).

[FIGURES ?? AND ?? ABOUT HERE]

According to Figure ??, the in-degree of Bear Stearns gradually increased through time up until JP Morgan Chase acquired it in March 2008. This means that it was receiving financial spillovers from a growing number of firms, effectively increasing its fragility.

Bear Stearns was the first large bank to collapse as a result of the subprime mortgage crisis of 2007. The fact that it had many incoming connections and few outgoing connections (shown in Figure ??) may explain why the collapse of Bear Stearns was, at the time, not warranted as a systemic event even though it marked the prelude to the financial crisis.

Another interesting aspect that is worth mentioning regarding the interconnectedness of Bear Stearns is that the data involved in calculating its in-degree did not contain the financial meltdown of September 2008. This is because the sample time-series of Bears Stearns stocks ran only until March 2008. Therefore, the high in-degree cannot be attributable to any noise caused by the many events that occurred during September-October 2008.

Figure ?? shows the out-degree for the four defunct financial institutions. We can notice Wachovia's out-degree increased through time until its government-induced forced sale to Wells Fargo, completed in December 2008. Again, we can notice how the out-degree calculated using the classical rolling window approach gives a measure that is more volatile and sporadic. It jumps suddenly in September 2008, while the corresponding measure calculated

with our time-varying parameters framework shows an increase that began long before this event.

Whereas AIG was bailed out, Lehman Brothers was considered by government and Fed officials as not systemic and was therefore allowed to go bankrupt (?). Both our time-varying parameter framework and the rolling window approach show that Lehman Brothers did not have a particularly high out-degree, thus both measures tend to agree with the policy decisions taken at the time (Figure ??). However, whether Lehman Brothers had a role in propagating financial spillovers is still a debated question.

Unfortunately, we do not have the counterfactual event to answer this question but several considerations can be made.

According to an asset and liability exposure analysis conducted by ?, Lehman's bankruptcy was not particularly destabilizing for its direct counterparties. However, without doubt Lehman's bankruptcy triggered a sentiment of fear and uncertainty in markets. Fear was primarily due to the realization that government rescue of large financial institutions was no longer guaranteed whereas uncertainty surrounded the extent of losses incurred by other institutions due to Lehman's default. This led to a freeze in the short-term funding market and effectively a liquidity crisis that affected financial institutions indiscriminately of their connections with Lehman Brothers.

On the other hand, Lehman's reliance on short-term rather than long-term funding made it vulnerable to external shocks (?). This is detectable in the growing in-degree of Lehman, which according to our time-varying framework had begun since 2006 (Figure ??).

D. Stability of centrality-based rankings

Using the approach developed by the Basel Committee on Banking Supervision (BCBS), the FSB ranks financial institutions according to their systemic importance and uses this ranking to determine their additional loss absorbency requirements. In a similar spirit, we rank the institutions in our sample according to their in- and out-degree. These two rankings will give only a partial view on systemic risk, one based only on interconnectedness. Nonetheless, these rankings can help us identify, at every point in time, the most exposed institutions to financial spillovers and the most important propagators of financial spillovers.

The usefulness of rankings for policy makers is severely limited if these are prone to frequent, drastic changes that lead to unmotivated excessive alarm (??). For this reason we developed a series of measures, explained in Section ??, to quantify the stability of our rankings. We compare this stability to the stability of rankings based on the degree centrality measures computed using the classical rolling window approach.

Table ?? (below) shows the results of these measures.

[TABLE ?? ABOUT HERE]

The first four columns of the tables display the *quadratic* and *absolute* stability measures, respectively, for the in- and out-degree centrality rankings. As explained in Section ??, *quadratic* stability indicators, SI_Q^{in} and SI_Q^{out} , measure the average change in the ranking between adjacent time periods. The quadratic term used in the calculation causes large deviations in the rankings to have a much higher weight in the stability indicator compared to smaller deviations. For the *absolute* stability indicator, SI_A^{in} and SI_A^{out} , the weight increases only linearly with the distance between ranking positions. So, for example, if a financial institution ranked first for in-degree were to be ranked 50th in the successive month, this position change would have a greater impact on the quadratic stability indicator than on the absolute stability indicator.

Both measures, for both rankings (in- and out-degree), show that the rankings based on our time-varying parameter framework are far more stable than the rankings based on the classical Granger causality approach over rolling windows. In fact, according to both stability indicators, the rolling window approach appears to produce rankings that are more than twice as unstable as those produced by the time-varying parameter framework.

The successive two columns headed “% Invariance”, denote the average percentage of firms that kept their position in the ranking between adjacent time periods. On any given month, on average, about 20% of firms kept the same position held in the previous month in the in- and out-degree rankings calculated with time-varying parameter framework. Quite to the contrary, only about 10% kept the same position for centrality rankings calculated using the rolling window approach. This confirms the higher stability of rankings found under the time-varying parameter framework.

The columns headed “ Δ Top 10” and “ Δ Top 20” show the average changes (in percentages) in the composition of, respectively, the top 10 and the top 20 financial institutions in the rankings. For example, for the rankings computed using the rolling window approach, we can expect 2 firms to change in the top 10 (19.9% for in-degree, 21.2% for out-degree) every month. On the other hand, for the rankings computed using the time-varying parameter framework, only about 7% of the top ten changed on average, so less than one firm per month. Similar magnitudes of results were obtained for the stability of the top 20 firms in the rankings.

All measures used to quantify stability indicate that the time-varying parameter framework provides more stable rankings compared to the rolling window approach. The high instability of the rankings found using the rolling window approach would make these rankings difficult to use for policy purposes. It would be hard to justify policy decisions based on a ranking that changes, on average, two components in its top 10 most interconnected institutions every month. On the other hand, the time-varying parameter framework offers a generally stable ranking while allowing some degree of flexibility that can be useful to motivate policy intervention.

The reasons for the higher stability offered by the time-varying parameter framework are to be found in the transition law imposed for time-varying connections. By allowing some degree of inertia between successive time periods, large exceptional observations have less influence on the estimated path of connections. On the other hand, with the rolling window approach, these observations have a larger weight in the estimation of connections. Even after degarching, extreme observations entering and exiting the rolling windows can lead to large sudden changes in the degree centrality of financial institutions and therefore in the positions that these financial institutions occupy in the rankings.

E. Density and sectorial interconnectedness

As an indicator of the general connectivity of the financial system, we computed the density of the network estimated using the proposed time-varying parameter framework. The density is the number of connections present in the network at a given moment in time as a proportion of the total possible connections. It gives us a measure of how integrated

the financial system is at a given moment in time. For comparison we also computed the density of the network estimated using the rolling windows of 36 months (as was used above and in ?) and 24 months.

The evolution of the density is given in Figure ?? together with some significant events for the US economy. The figure shows the density under the time-varying parameter framework in bold solid. The density calculated using rolling windows of 36 months is depicted by the light dashed line, whereas that calculated using rolling windows of 24 months is depicted by triangle markers.

[FIGURE ?? ABOUT HERE]

The time-varying parameter framework density emphasizes the strong growth in interconnectedness that occurred prior to the outbreak of the financial crisis. Network density grew gradually from the beginning of 2005 and culminated in October 2008. In between these two periods, the number of connections increased by 46.4%.

Looking at the density measures computed from the rolling window estimates, we notice that these are substantially more volatile than the time-varying parameter counterpart, exhibiting sudden short-lived peaks around the dates identified above. Moreover, there are several discrepancies in the density reported by the two rolling window approaches.

Notice that the peak occurring close to the US downgrade is very high, even higher than that exhibited during September-October 2008. Therefore, according to the rolling window approach, the US financial system was more interconnected during this period than during the financial crisis. However, a closer look at this last peak shows that there is a large discrepancy between the density computed using 36 months rolling windows and that found using 24 rolling windows. We notice that this discrepancy is not as large around the other peaks.

We believe that this peak is generated by the crisis observations, relating to October 2008, exiting the rolling windows and creating an artificial jump in interconnectedness, as was witnessed for the degree centrality measures in Section ?. Notice that for the density computed using 24 months long rolling windows the peak occurs exactly after 24 months, in October 2010. Similarly, for the density found using 36 months long rolling windows, the

peak occurs after 36 months in October 2011, actually two months after the US downgrade.

After the financial crisis, the network density, calculated from the network estimated by the time-varying parameter framework, gradually decreases to below pre-crisis levels. This is a symptom that the policies introduced to counter the financial crisis, such as the Troubled Asset Relief Program and Dodd-Frank act were effective in reducing the perceived interconnectedness of the system. ? found a similar decrease using a realized volatility-based measure of systemic risk.

In order to study the evolution of interconnectedness between and within sectors, we made use of the measures developed in Section ??, Inter-Sector-Degree $_{m \rightarrow n, t}$ and Intra-Sector-Degree $_{m, t}$ for all four sectors analysed in our study. The off-diagonal plots in Figure ?? shows the time-varying interconnectedness between every sector pair, where the sender sector of financial spillovers is represented by the row of the plots and the receiving sector by the column. Along the diagonal, the plots show interconnectedness between institutions of the same sector that is, the intra-sectorial degree of a given sector.

[FIGURE ?? ABOUT HERE]

The first aspect noticeable in Figure ?? is that some sectors appear to have had a stronger increase in interconnectedness compared to others, during the September 2008 financial crisis (denoted by the dashed vertical line in each plot). In particular, banks and insurance companies seem to have had the largest increase in outgoing connections during this period. Between these two sectors, it seems that banks were influencing insurers to a greater extent than insurers were influencing banks. Both sectors appeared to have increased their intra-sectorial connections during the crisis.

Regarding incoming connections, Figure ?? shows that the sector receiving the most spillovers during the financial crisis was the real estate sector, as was also noticeable from Table ?. Specifically, almost 30% of the possible connections between institutions in the real estate sector and institutions in the banking, broker/dealers and insurance sectors were active.

A closer look at the fourth row of the figure shows that the real estate sector also had a moderately high level of outgoing connections. Therefore, the real estate sector was being

influenced by larger and more central financial institutions in other sectors but contemporaneously it played a role in propagating spillovers back to these other sectors. Due to this dual role, the real estate sector can be seen as a “risk distributor” of the type described by ?.

VI. Concluding remarks

By modelling the evolution of connections between financial institutions, we have provided a framework for estimating financial interconnectedness that accounts for the dynamic nature of connections. We have built our framework in a time-varying parameter vector autoregression setting, paralleling studies that infer network connections using Granger causality. In particular, we have adopted a Bayesian test that yields, at every moment in time, the posterior probability of a connection existing between any two financial institutions.

The framework infers connections that are gradually evolving through time, as opposed to the classical approach of sequentially running Granger causality tests over rolling windows of data. The rolling window approach has several limitations that our framework surpasses. First, it relies on shorter samples, which decrease inferential power and favour dimensionality problems. Furthermore, the rolling window approach is highly susceptible to outliers because, in small samples, these have a larger impact on estimates. On the other hand, estimates based on longer windows are less reactive to changes in the underlying relationships and thus have a bias towards static connections. These limitations ultimately mean that the user faces a trade-off when selecting the appropriate window size. We show through our simulation exercises and our empirical application that the choice of window size is thus critical and can lead to different results regarding interconnectedness.

In the simulation exercises we evaluated our framework against the rolling window approach when estimating constant and dynamic networks. Results show that the time-varying framework performs well, both in terms of precision (measured by the MSE of the estimated parameters) and in terms of the efficacy of correctly identifying connections (measured by the ROC and precision-recall curves). Compared to the classical rolling window approach, our proposed time-varying framework performs equally well when the network is constant

and it outperforms when links are changing through time. The reason for this is that the time-varying parameter framework exploits the whole length of data providing a more precise estimate of the evolution of the network.

In the empirical application, we applied the proposed framework to real data so to estimate the network of spillovers for all US financials listed in the S&P 500 between 1990 and 2014. Using the retrieved network, we estimated several connectivity and centrality measures at both the sectorial and company level. We repeated the exercise using the rolling window approach in order to compare the two techniques. Four main results emerged.

First, the empirical application revealed the deficiencies of the rolling window approach and the gains of adopting the time-varying parameter framework in the network context. In particular, degree centrality measures computed under the rolling window approach were more volatile than those computed using the time-varying parameter framework. An indepth analysis of the evolution of degree centrality for individual financial institutions exposed the reasons for this.

It was found that degree centrality measures computed under the rolling window approach were susceptible to sudden unjustified jumps. The jumps actually occurred when extreme observations exited the rolling window. The same phenomenon caused the density of the whole financial system to jump oddly around October 2011, when the observations relating to the peak of the financial crisis exited the 36 month window used for the exercise. Quite to the contrary, the time-varying parameter framework showed interconnectedness gradually decreasing after the crisis, symptom that the policies that had been carried out were effective to this end.

Second, further analysis of the degree centrality results revealed that, during the crisis, American International Group (AIG) was the most interconnected financial institution in terms of propagating spillovers to the rest of the system. At its peak, AIG was connected to over 60% of the financial institutions in our sample, endorsing the Federal Reserve Bank's decision to rescue it. At the same time, Bear Stearns was found to have a high in-degree, making it prone to receiving spillovers from other institutions, but it also had a low out-degree, which could explain why its collapse did not trigger a systemic event.

Third, rankings based on degree centrality calculated according to the rolling window ap-

proach were more unstable than those computed using the time-varying parameter approach. The Financial Stability Board (FSB) produces similar rankings of systemically important financial institutions using five indicators of systemic risk, among which interconnectedness. In order for these rankings to be useful for policy makers, such rankings must be stable through time while allowing for sufficient flexibility to signal important changes in interconnectedness.

Due to the susceptibility of the rolling window approach to extreme observations, the rankings produced by it are too unstable to be used for policy decisions. On the other hand, the time-varying approach produces less variable rankings that appear to combine an appropriate level of flexibility and stability.

Fourth, at the sectorial level, banks and insurance companies have had the largest increase in outgoing connections during the crisis whereas the real estate sector, composed primarily of real estate investment trusts, had the largest increase in incoming connections. The latter also had moderately high levels of outgoing spillovers. The combination of high incoming and outgoing connections renders the real estate sector a “risk distributor”, transferring financial spillovers across the system. When regulating these institutions, policy makers should take this into account.

The major limitation of our framework rests on its computational burden. This is due to the Bayesian estimation technique adopted, which allows greater flexibility but leads to time-costly sampling algorithms. Nonetheless, computational timing should not be an issue if one would like to compute the time-varying network in a recursive manner, as new data enters the information set. For example, the network could be estimated every month for monitoring purposes. For our empirical application, it took about 24 hours to compute the network using a parallel computing cluster. For monthly or weekly monitoring, which is the frequency at which we deem connectivity evolves, recursive updating is surely feasible.

Other limitations of our framework have to do with the actual assumptions in the model. In particular, the assumption of normal errors is often criticized for financial returns. Nonetheless, the assumption is not as critical in the case of degarched log-returns that are used in the empirical application. Moreover, the framework can be extended to account for fat-tailed errors (?).

On the other hand, there are at least two avenues for further research into the time-varying parameter framework for measuring interconnectedness.

First, it is possible to adopt the time-varying parameter framework to forecast the path of the network in the future. This can be done by simulating from the conditional posterior density of the future parameters (see ?). Once this is obtained, we can compute the posterior probability that the future parameters will be zero and thus the probability of observing a link in the future between any two institutions. Thus, rather than yielding an *ex-post* picture of the evolution of the network, the time-varying framework can be used to obtain an *ex-ante* perspective of its future path.

Second, as mentioned in the introduction, interconnectedness is only one aspect of systemic risk. The FSB additionally identifies cross-jurisdictional activity, size, substitutability and financial institution infrastructure and complexity as equally important aspects of systemic risk. In order to obtain a comprehensive measure of systemic risk, future research should concentrate on seeking appropriate ways for combining time-varying measures of interconnectedness with these other aspects characterising financial institutions.

Appendix A

A.1. Priors

The various blocks of parameters are assumed to be independent,

$$p(B_0, R, Q) = p(B_0)p(R)p(Q).$$

The prior for the initial states of the time-varying coefficients, $f(B_0)$, is a Gaussian density,

$$B_0 \sim \mathcal{N}(\bar{B}, \bar{P}),$$

where \bar{B} corresponds to the OLS point estimates of a training sample and \bar{P} to four times the covariance matrix $\hat{V}(\bar{B})$. For both the simulation exercises and the empirical application, we use a training sample of 38 observations. For testing, \bar{B} was estimated by restricting the parameters referring to the connections being tested to zero. On the other hand, the unrestricted estimates were used for $\hat{V}(\bar{B})$.

The prior for Q is inverse-Wishart,

$$Q \sim \mathcal{IW}(\bar{Q}^{-1}, T_0).$$

where T_0 is the prior degrees of freedom. Following ?, we set $T_0 = \dim(B_t) + 1$ so that the prior is only weakly informative and the posterior for Q puts most weight on the data. Through a series of simulation exercises, ? show that this configuration performs well in detecting different levels of time variation.

Following ?, we set $\bar{Q} = (0.01)^2 \times \hat{V}(\bar{B})$ multiplied by the number of observations in the training sample. Although moderately conservative, this choice is not expected to influence the results since the prior is dominated by the sample information.

Priors for R are set as

$$R \sim \mathcal{IW}(I_N, N + 1),$$

where I_N is the identity matrix of size N . In the case of pairwise recursive testing with a

bivariate TVP-VAR, a two by two identity matrix would be used and the degrees of freedom would be set to two instead of N .

A.2. Posterior distribution simulation

We simulate the joint posterior distribution by sequentially drawing from the conditional posterior of the three blocks of parameters: the coefficients B^T and the variance-covariance matrices Q and R .

A.2.1. Parameter states, B^T

The conditional distribution of the TVP-VAR parameters, B^T , can be expressed as:

$$(11) \quad p(B^T | x^T, Q, R) = p(B_T | x^T, Q, R) \prod_{t=1}^{T-1} p(B_t | B_{t+1}, x^T, Q, R)$$

Given the prior assumptions above and the state-space model, the conditional densities are normal and can be simulated using the algorithm proposed by ?.

Precisely, we can compute their means and variances through the forward and backward recursions of the Kalman filter and smoother. The last iteration of the filter provides the mean and variance for the first term in (??),

$$p(B_T | x^T, Q, R) = \mathcal{N}(B_{T|T}, P_{T|T})$$

A draw from the distribution is used in the backward recursions to simulate the remaining terms in (??). Conditional on the information in B_{t+1} , B_t is conditionally normal with mean and variance given respectively by,

$$\begin{aligned} B_{t|t+1} &= B_{t|t} + P_{t|t} P_{t+1|t}^{-1} (B_{t+1} - B_{t|t}), \\ P_{t|t+1} &= P_{t|t} - P_{t|t} P_{t+1|t}^{-1} P_{t|t} \end{aligned}$$

The backward recursions draw sequentially $B_{T-1}, B_{T-2}, \dots, B_1$ from the conditional distri-

bution

$$(12) \quad p(B_t | x^T, Q, R) = \mathcal{N}(B_{t|t+1}, P_{t|t+1}),$$

in order to generate a random trajectory B^T .

A.2.2. Innovation variance-covariance matrix, Q

Conditional on a realization of B^T , the TVP-VAR parameter innovations, v_t , are observable. Under the linear transition law, v_t is i.i.d. normal. Given the natural conjugate prior specified above, the posterior is inverse-Wishart,

$$p(Q | x^T, B^T) = \mathcal{IW}(Q_1^{-1}, T_1),$$

with scale and degree-of-freedom parameters,

$$Q_1 = \bar{Q} + \sum_{t=1}^T v_t v_t' \quad T_1 = T_0 + T.$$

A.2.3. Residual variance-covariance matrix, R

Conditional on a realization of B^T , the TVP-VAR residuals, u_t , are observable. Given the conjugate prior assumption given above, the conditional posterior density of R is given by

$$p(R | x^T, B^T) = \mathcal{IW}(R_1^{-1}, N_1),$$

where,

$$R_1 = I_N + \sum_{t=1}^T u_t u_t' \quad N_1 = N + 1 + T.$$

For each estimation, we perform 6000 iterations of the Gibbs sampler and discard the first 1000 draws. We then keep only the 5th of every draw in order to mitigate autocorrelation

among draws. The remaining sequence of 1000 draws forms a sample of the joint posterior distribution $p(B^T, Q, R | x^T)$. We use this to estimate Bayes factor and test the time-varying hypothesis of no connection between two nodes of the system.

Tables and Figures

Banks		Insurers	Real Estate
AHMANSON (H F) & CO	NORTHERN TRUST CORP	ACE LTD	AMERICAN CAPITAL LTD
AMERICAN EXPRESS CO	OLD KENT FINANCIAL CORP	AETNA INC	AMERICAN TOWER CORP
AMSOUTH BANCORPORATION	PEOPLE'S UNITED FINL INC	AFLAC INC	APARTMENT INVST & MGMT CO
ASSOCIATES FIRST CAP -CL A	PNC FINANCIAL SVCS GROUP INC	ALEXANDER & ALEXANDER	AVALONBAY COMMUNITIES INC
BANK OF AMERICA CORP	PROVIDIAN FINANCIAL CORP	ALLSTATE CORP	BOSTON PROPERTIES INC
BANK OF NEW YORK MELLON CORP	REGIONS FINANCIAL CORP	AMERICAN INTERNATIONAL GROUP	CBRE GROUP INC
BANK ONE CORP	SHAWMUT NATIONAL CORP	AON PLC	CROWN CASTLE INTL CORP
BANKBOSTON CORP	SLM CORP	ASSURANT INC	DDR CORP
BARNETT BANKS INC	STATE STREET CORP	CHUBB CORP	EQUITY OFFICE PROPERTIES TR
BB&T CORP	SUMMIT BANCORP	CIGNA CORP	EQUITY RESIDENTIAL
BENEFICIAL CORP	SUNTRUST BANKS INC	CINCINNATI FINANCIAL CORP	GENERAL GROWTH PPTYS INC
BOATMENS BANCSHARES INC	SYNOVUS FINANCIAL CORP	CNA FINANCIAL CORP	HCP INC
CAPITAL ONE FINANCIAL CORP	U S BANCORP	CNO FINANCIAL GROUP INC	HEALTH CARE REIT INC
CHASE MANHATTAN CORP -OLD	UNION PLANTERS CORP	CONTINENTAL CORP	HFS INC
CIT GROUP INC	VISA INC	COVENTRY HEALTH CARE INC	HOST HOTELS & RESORTS INC
CIT GROUP INC-OLD	WACHOVIA CORP	EXPRESS SCRIPTS HOLDING CO	KIMCO REALTY CORP
CITICORP	WACHOVIA CORP-OLD	GENWORTH FINANCIAL INC	MACERICH CO
CITIGROUP INC	WASHINGTON MUTUAL INC	HANCOCK JOHN FINL SVCS INC	PROLOGIS INC
COMERICA INC	WELLS FARGO & CO	HARTFORD FINANCIAL SERVICES	PUBLIC STORAGE
COMMERCE BANCORP INC/NJ	WELLS FARGO & CO -OLD	HUMANA INC	SIMON PROPERTY GROUP INC
COMPASS BANCSHARES INC	WESTERN UNION CO	JEFFERSON-PILOT CORP	VENTAS INC
CONCORD EFS INC	ZIONS BANCORPORATION	LINCOLN NATIONAL CORP	VORNADO REALTY TRUST
CORESTATES FINANCIAL CORP		LOEWS CORP	WYNDHAM WORLDWIDE CORP
COUNTRYWIDE FINANCIAL CORP		MCLENNAN COS	
DISCOVER FINANCIAL SVCS INC	Brokers	MBIA INC	
FANNIE MAE	AMERIPRISE FINANCIAL INC	METLIFE INC	
FEDERAL HOME LOAN MORTG CORP	BEAR STEARNS COMPANIES INC	MGIC INVESTMENT CORP/WI	
FIFTH THIRD BANCORP	BLACKROCK INC	PROGRESSIVE CORP-OHIO	
FIRST CHICAGO NBD CORP	CME GROUP INC	PROVIDENT COS INC	
FIRST FIDELITY BANCORP	E TRADE FINANCIAL CORP	PRUDENTIAL FINANCIAL INC	
FIRST HORIZON NATIONAL CORP	FEDERATED INVESTORS INC	SAFECO CORP	
FIRST INTERSTATE BNCP	FRANKLIN RESOURCES INC	TORCHMARK CORP	
FLEETBOSTON FINANCIAL CORP	GOLDMAN SACHS GROUP INC	TRAVELERS COS INC	
GOLDEN WEST FINANCIAL CORP	INTERCONTINENTALEXCHANGE GRP	U S HEALTHCARE INC	
GREAT WESTERN FINANCIAL	INVESCO LTD	UNITEDHEALTH GROUP INC	
HSBC FINANCE CORP	JANUS CAPITAL GROUP INC	UNUM GROUP	
HUDSON CITY BANCORP INC	LEGG MASON INC	USF&G CORP	
HUNTINGTON BANCSHARES	LEHMAN BROTHERS HOLDINGS INC	USLIFE CORP	
JPMORGAN CHASE & CO	MERRILL LYNCH & CO INC	WELLPOINT HEALTH NETWRKS INC	
KEYCORP	MORGAN STANLEY	XL GROUP PLC	
M & T BANK CORP	NASDAQ OMX GROUP INC		
MARSHALL & ILSLEY CORP	NYSE EURONEXT		
MASTERCARD INC	PAINE WEBBER GROUP		
MBNA CORP	PRICE (T. ROWE) GROUP		
MELLON FINANCIAL CORP	PRINCIPAL FINANCIAL GRP INC		
MERCANTILE BANCORPORATION	SCHWAB (CHARLES) CORP		
MONEYGRAM INTERNATIONAL INC			
NATIONAL CITY CORP			
NORTH FORK BANCORPORATION			

Table 1: List of financial institutions used for the empirical application

In-Degree					
Rolling window	Mean (%)	Med. (%)	St. Dev (%)	Min. (%)	Max. (%)
<i>Banks</i>	6.4	5.1	5.1	0.3	24.3
<i>Brokers</i>	7.9	5.9	6.3	0.5	28.0
<i>Insurers</i>	6.5	5.1	5.2	3.4	25.0
<i>Real Estate</i>	10.0	6.1	9.4	0.3	41.9
Time-varying parameter	Mean (%)	Med. (%)	St. Dev (%)	Min. (%)	Max. (%)
<i>Banks</i>	12.0	11.8	3.1	0	18.4
<i>Brokers</i>	13.0	13.1	3.5	0	20.0
<i>Insurers</i>	12.0	11.9	3.1	0	18.5
<i>Real Estate</i>	17.1	17.3	4.9	0	26.9
Out-Degree					
Rolling window	Mean (%)	Med. (%)	St. Dev (%)	Min. (%)	Max. (%)
<i>Banks</i>	7.6	6.3	5.6	0.6	28.0
<i>Brokers</i>	7.0	5.0	6.1	0.1	28.2
<i>Insurers</i>	7.6	5.5	6.9	0.1	31.1
<i>Real Estate</i>	7.8	5.0	7.8	0.1	36.0
Time-varying parameter	Mean (%)	Med. (%)	St. Dev (%)	Min. (%)	Max. (%)
<i>Banks</i>	14.1	14.0	3.4	0	20.8
<i>Brokers</i>	13.0	12.8	3.2	0	20.0
<i>Insurers</i>	13.5	13.5	3.9	0	21.3
<i>Real Estate</i>	13.0	13.0	3.4	0	20.0

Table 2: Average degree centrality statistics across financial institutions from different sectors.

	Stability Indicators				% Invariance		Δ Top 10		Δ Top 20	
	<i>quadratic</i>		<i>absolute</i>		<i>in</i>	<i>out</i>	<i>in</i>	<i>out</i>	<i>in</i>	<i>out</i>
	SI_Q^{in}	SI_Q^{out}	SI_A^{in}	SI_A^{out}						
Rolling window	13.8	14.1	8.7	9.0	11.0	10.6	19.9	21.2	18.3	19.1
Time-varying parameter	6.8	6.6	4.1	3.9	19.4	19.9	7.4	7.1	6.9	6.5

Table 3: The table shows the Stability Indicator computed for the centrality-based rankings estimated by the time-varying parameter framework and by the classical Granger causality approach over rolling windows. The Stability Indicator expresses the extent of monthly changes in the rankings based on degree centrality. % Invariance measures the proportion of firms that held the same position in the ranking between adjacent time periods. Δ Top 10 and Δ Top 20 measures the extent of monthly changes in, respectively, the top 10 and top 20 institutions of the rankings.




Figure 1: The Causal Network




Figure 2: Results for Experiment 1 in which underlying parameters are constant. The light dashed line relates to results obtained using classical Granger causality testing over rolling windows, while the bold solid line relates to results obtained using the proposed time-varying parameter framework. The upper panel shows results for recursive estimation using pairwise testing. The lower panel shows results using conditional testing. The left column figures show the mean squared error of cross-parameter estimates. The middle and right column figures show the ROC curves and the PR curves for inference of the underlying network. For inference, Granger causality testing was conducted at 5% significance level, whereas Bayes factor was set to 50% for Bayesian inference of the time-varying hypothesis given by equation (??).

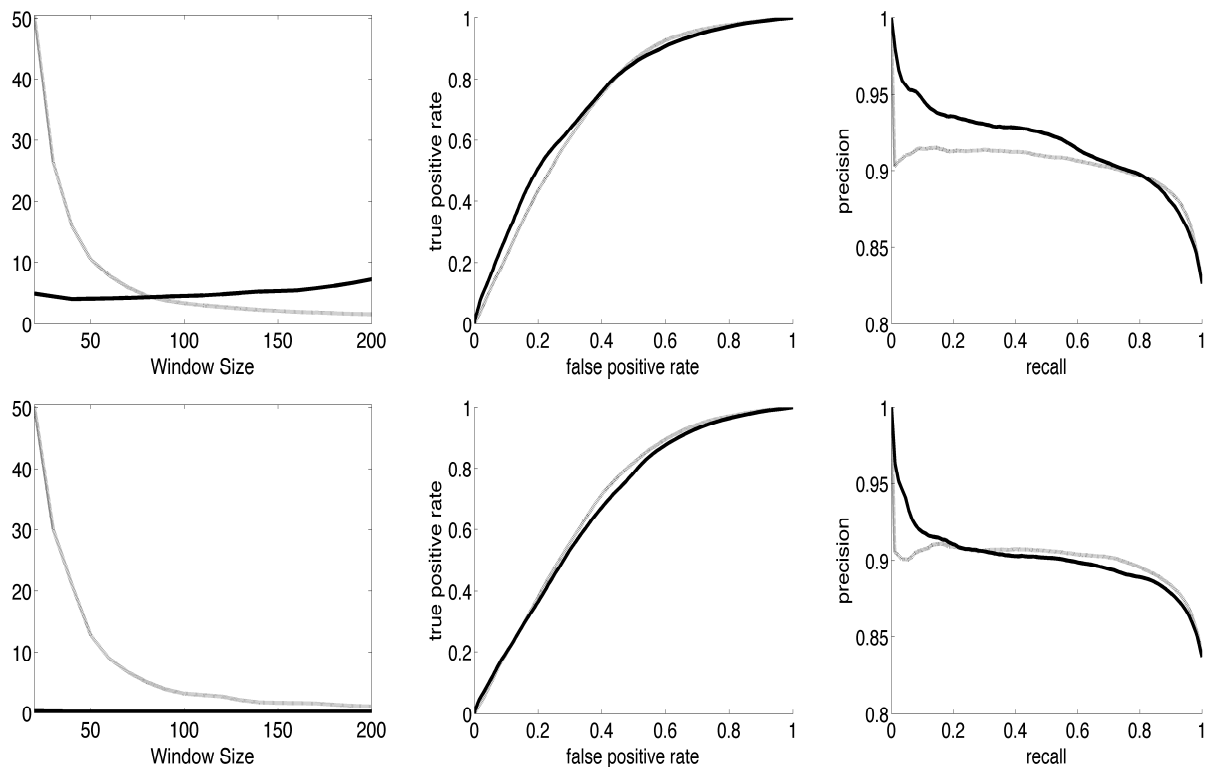


Figure 3: Results for Experiment 2 in which underlying parameters follow a regime switching process. The light dashed line to results obtained using classical Granger causality testing over rolling windows, while the bold solid line relates to results obtained using the proposed time-varying parameter framework. The upper panel shows results for recursive estimation using pairwise testing. The lower panel shows results using conditional testing. The left column figures show the mean squared error of cross-parameter estimates. The middle and right column figures show the ROC curves and the PR curves for inference of the underlying network. For inference, Granger causality testing was conducted at 5% significance level, whereas Bayes factor was set to 50% for Bayesian inference of the time-varying hypothesis given by equation (??).

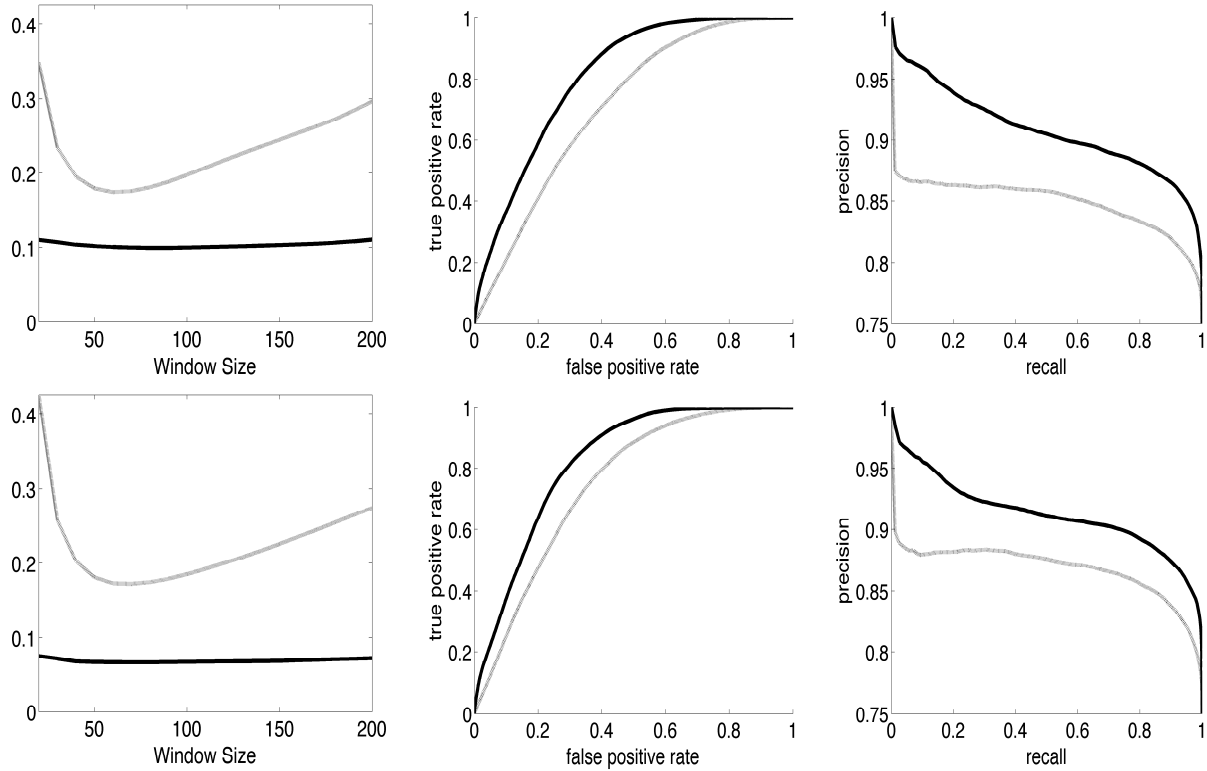


Figure 4: Results for Experiment 3 in which underlying parameters follow a random walk process. The light dashed line relates to results obtained using classical Granger causality testing over rolling windows, while bold solid line relates to results obtained using the proposed time-varying parameter framework. The upper panel shows results for recursive estimation using pairwise testing. The lower panel shows results using conditional testing. The left column figures show the mean squared error of cross-parameter estimates. The middle and right column figures show the ROC curves and the PR curves for inference of the underlying network. For inference, Granger causality testing was conducted at 5% significance level, whereas Bayes factor was set to 50% for Bayesian inference of the time-varying hypothesis given by equation (??).

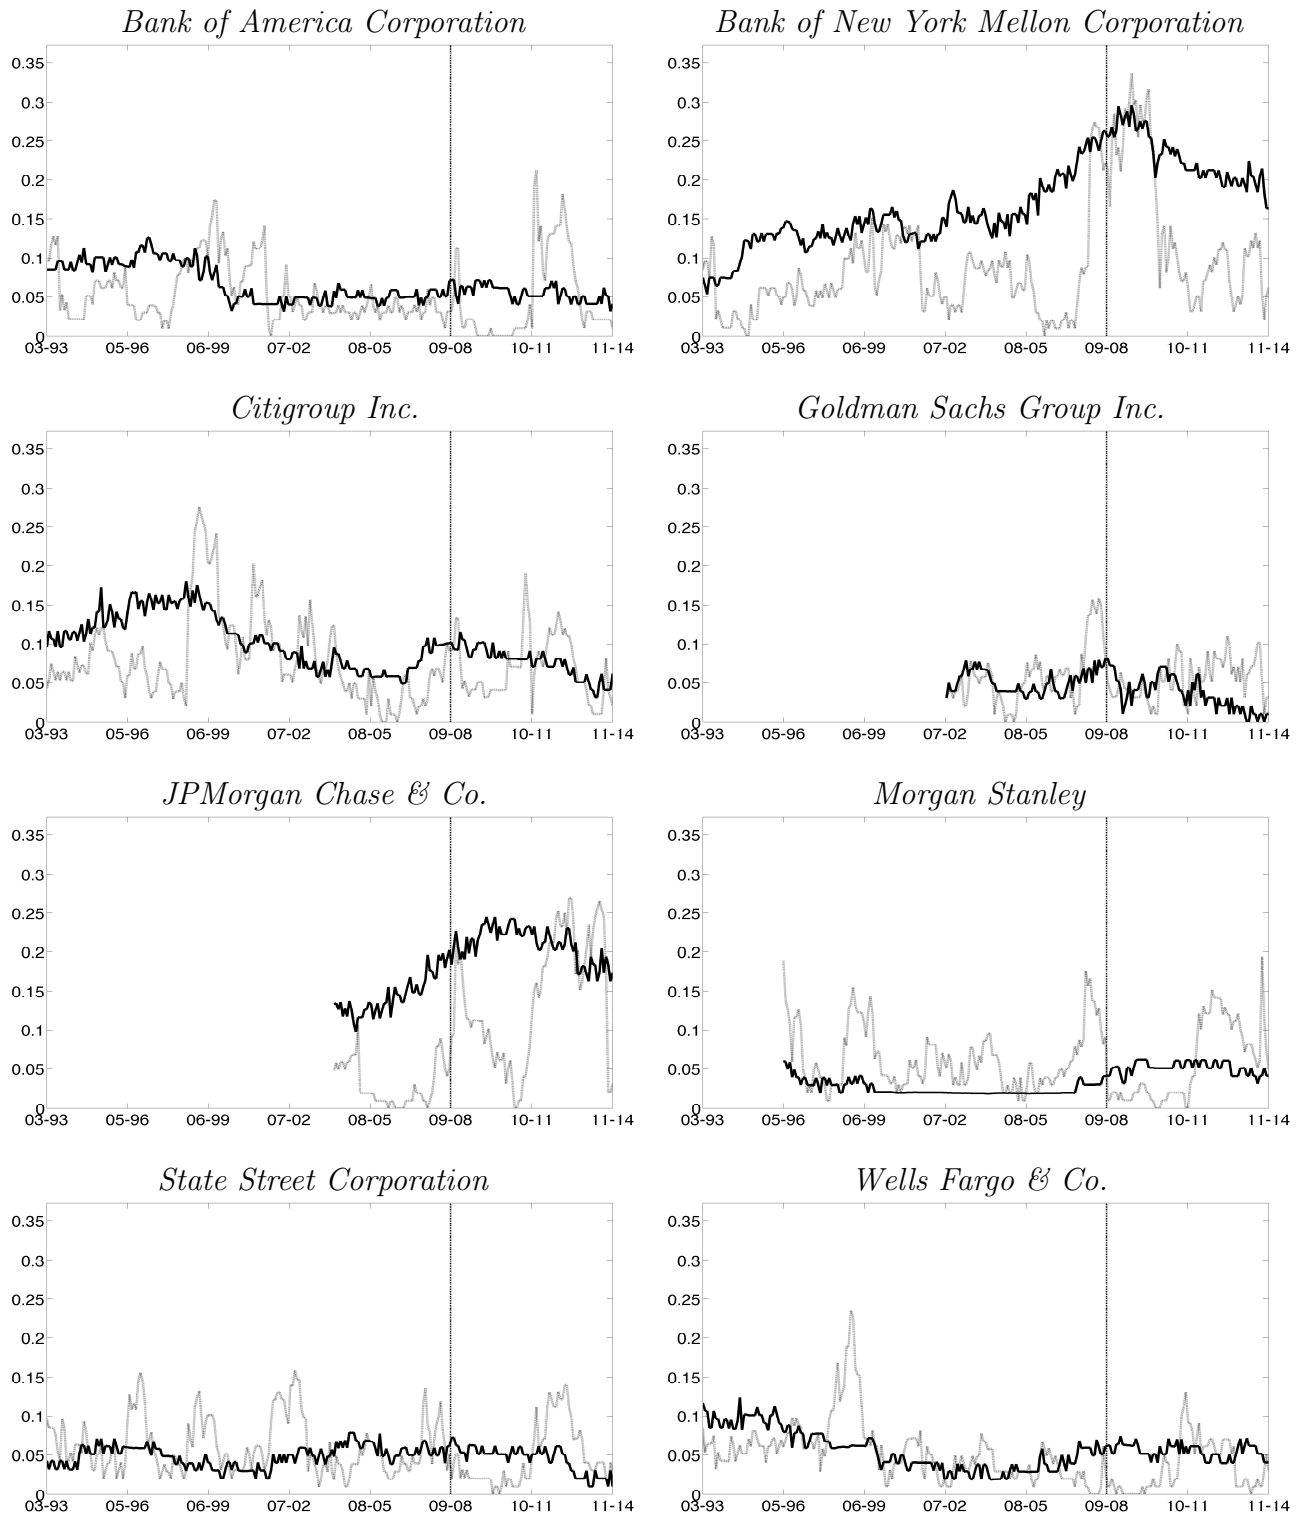


Figure 5: In-degree for US banks identified as global SIFIs by the FSB. The bold solid lines indicate in-degree found from the network estimated using the time-varying parameter framework. The lighter dashed lines indicate in-degree found from the network estimated using the classical approach of Granger causality testing over rolling windows of 38 months. The dashed vertical lines indicate the day of Lehman Brothers bankruptcy, 15 September 2008.

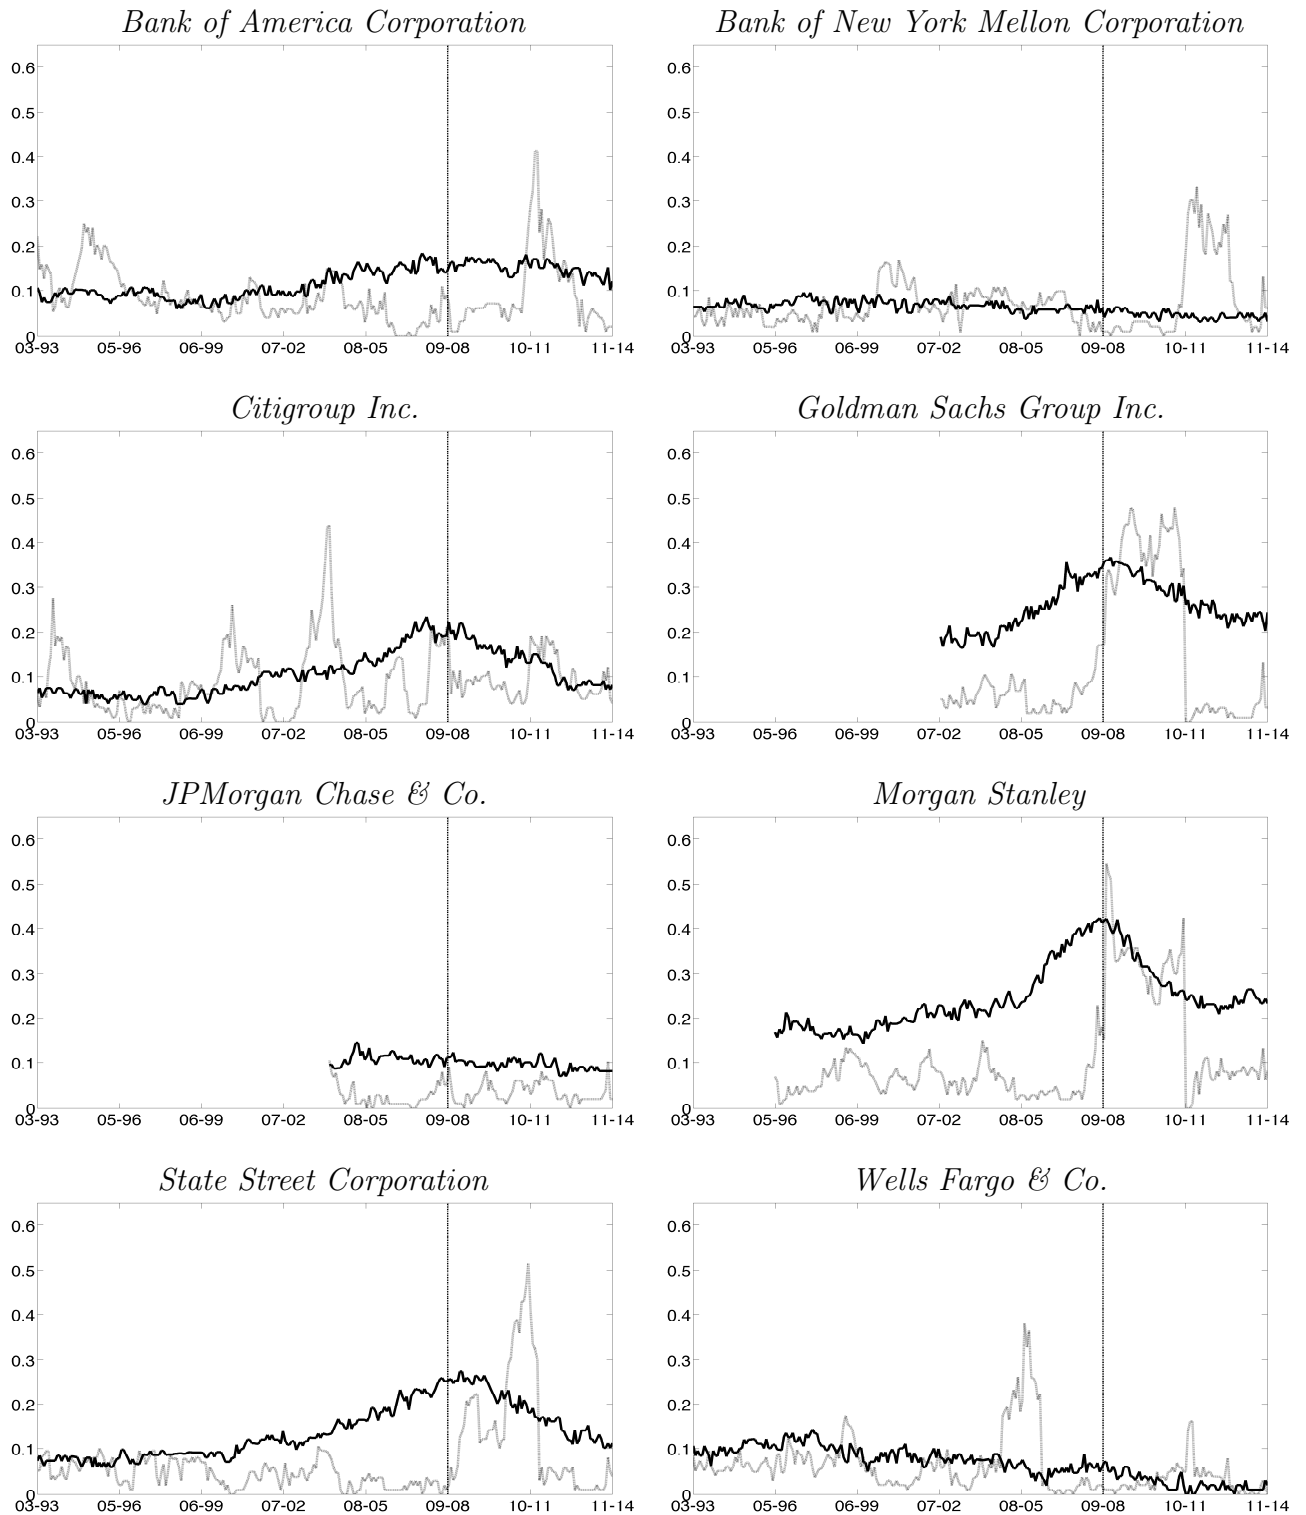


Figure 6: Out-degree for US banks identified as global SIFIs by the FSB. The bold solid lines indicate out-degree found from the network estimated using the time-varying parameter framework. The lighter dashed lines indicate out-degree found from the network estimated using the classical approach of Granger causality testing over rolling windows of 38 months. The dashed vertical lines indicate the day of Lehman Brothers bankruptcy, 15 September 2008.

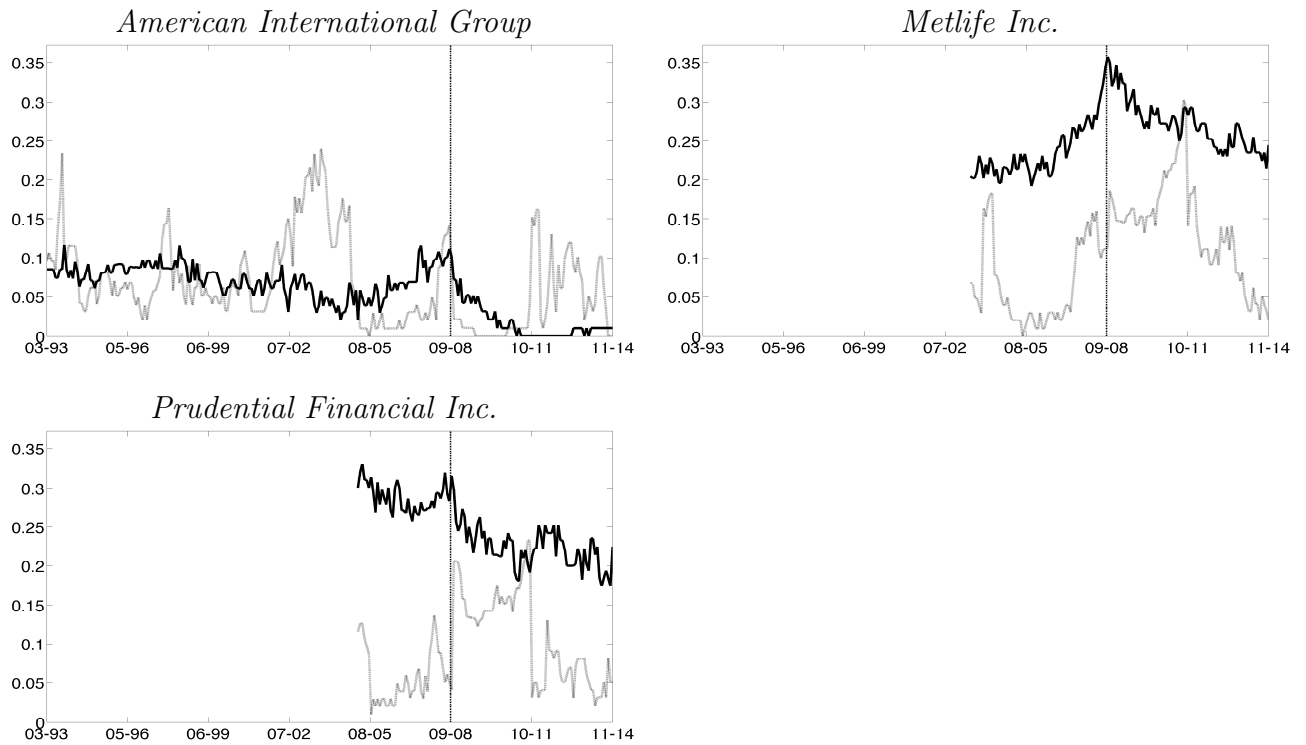


Figure 7: In-degree for US insurance companies identified as global systemically important insurers by the FSB. The bold solid lines indicate in-degree found from the network estimated using the time-varying parameter framework. The lighter dashed lines indicate in-degree found from the network estimated using the classical approach of Granger causality testing over rolling windows of 38 months. The dashed vertical lines indicate the day of Lehman Brothers bankruptcy, 15 September 2008.

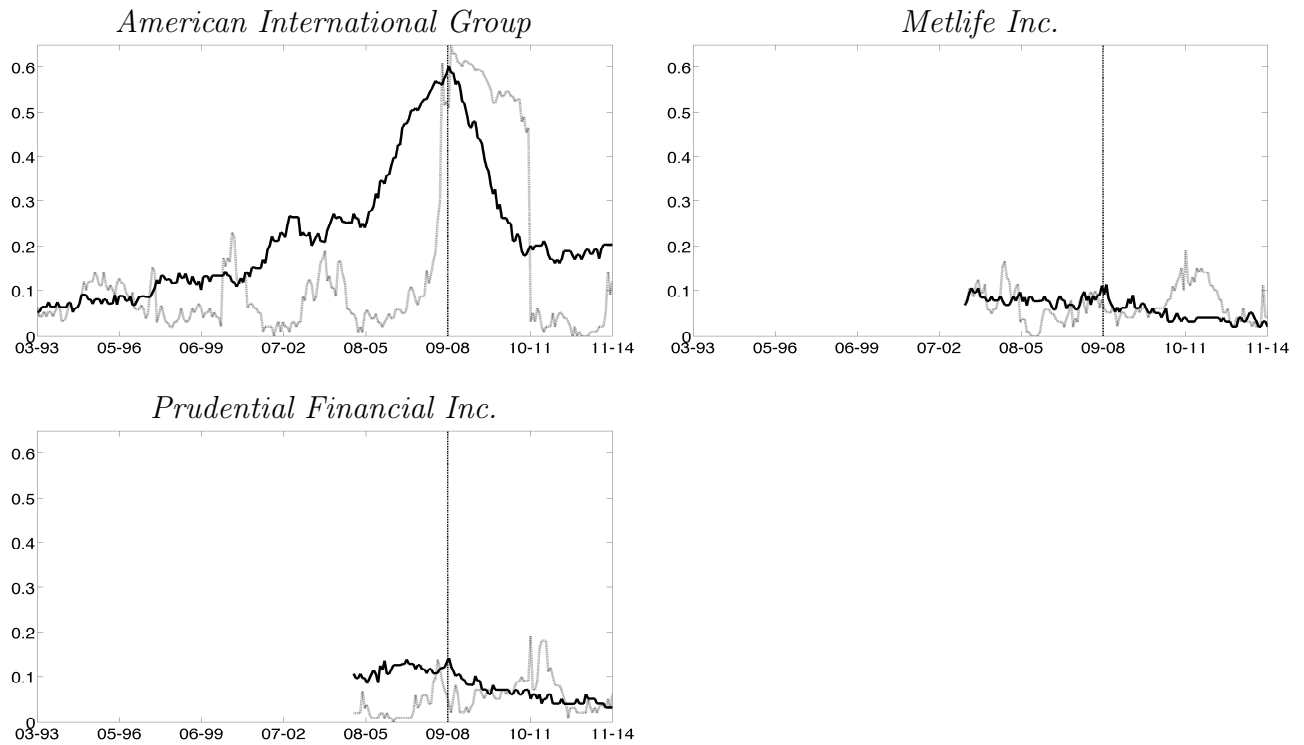


Figure 8: Out-degree for US insurance companies identified as global systemically important insurers by the FSB. The bold solid lines indicate out-degree found from the network estimated using the time-varying parameter framework. The lighter dashed lines indicate out-degree found from the network estimated using the classical approach of Granger causality testing over rolling windows of 38 months. The dashed vertical lines indicate the day of Lehman Brothers bankruptcy, 15 September 2008.

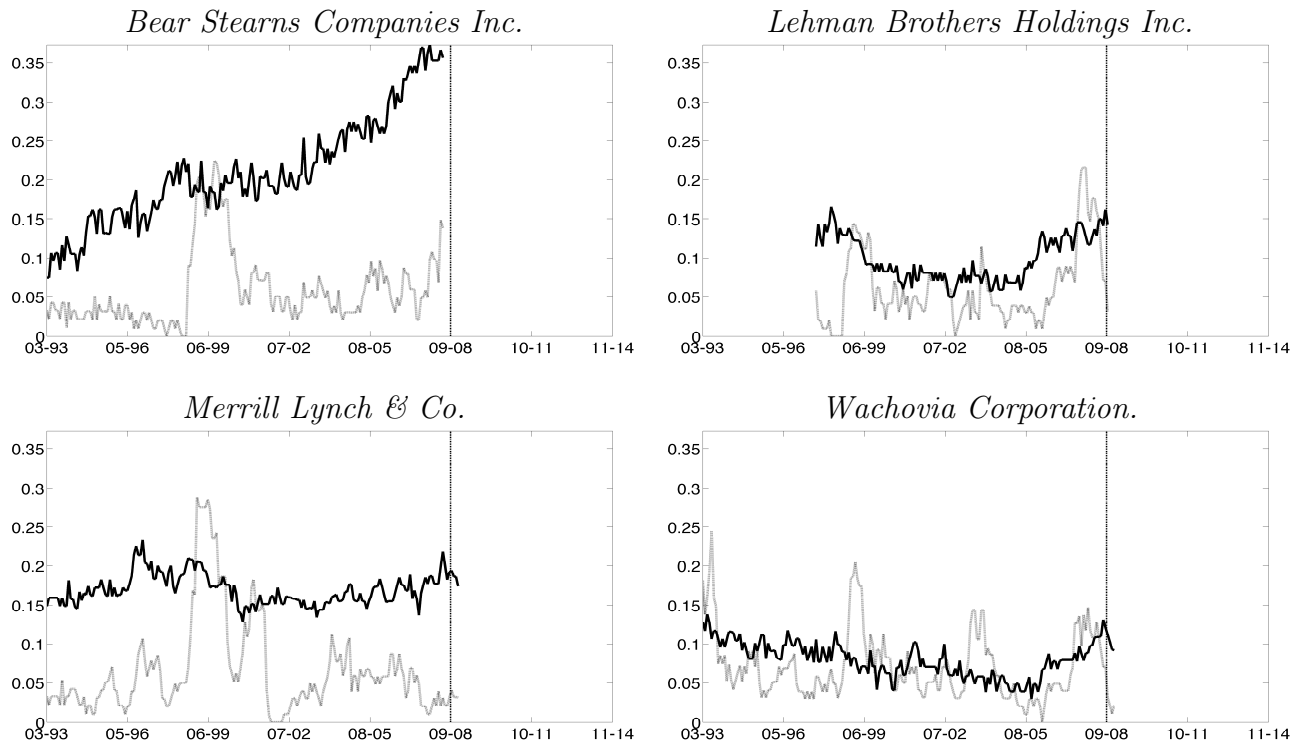


Figure 9: In-degree for four US financial institutions that suffered financial distress during the September 2008 financial crisis. The bold solid lines indicate in-degree found from the network estimated using the time-varying parameter framework. The lighter dashed lines indicate in-degree found from the network estimated using the classical approach of Granger causality testing over rolling windows of 38 months. The dashed vertical lines indicate the day of Lehman Brothers bankruptcy, 15 September 2008.

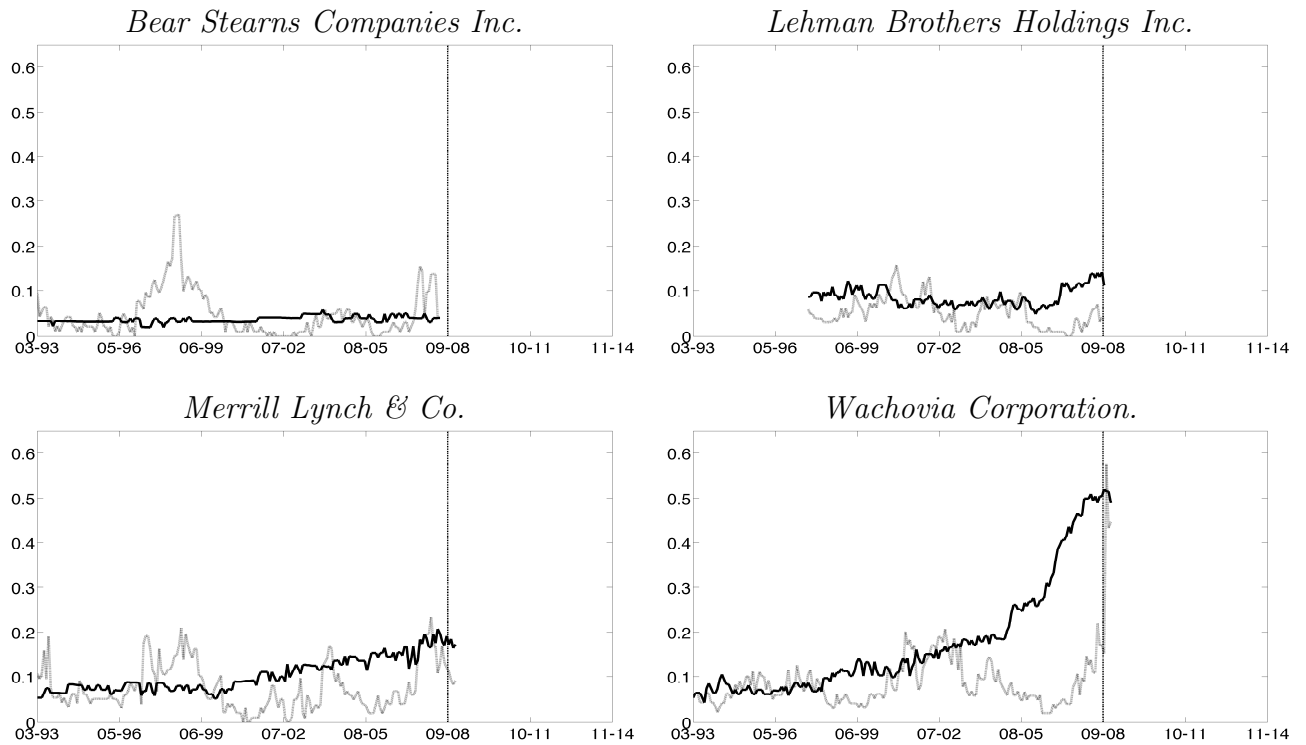


Figure 10: Out-degree for four US financial institutions that suffered financial distress during the September 2008 financial crisis. The bold solid lines indicate out-degree found from the network estimated using the time-varying parameter framework. The lighter dashed lines indicate out-degree found from the network estimated using the classical approach of Granger causality testing over rolling windows of 38 months. The dashed vertical lines indicate the day of Lehman Brothers bankruptcy, 15 September 2008.

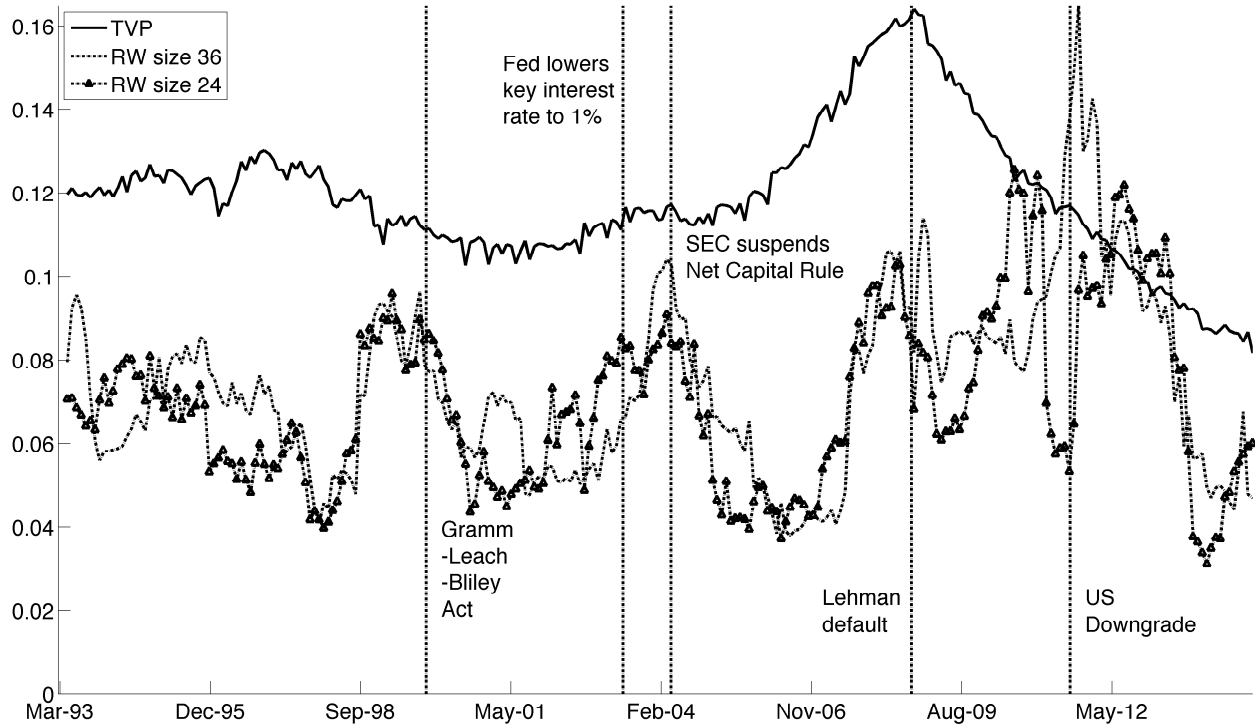


Figure 11: Network density estimated by time-varying parameters (bold solid) and by Granger causality testing over rolling windows of 38 observations (light dashed) and of 24 observations (triangles). Significant events are indicated by the dashed vertical lines.

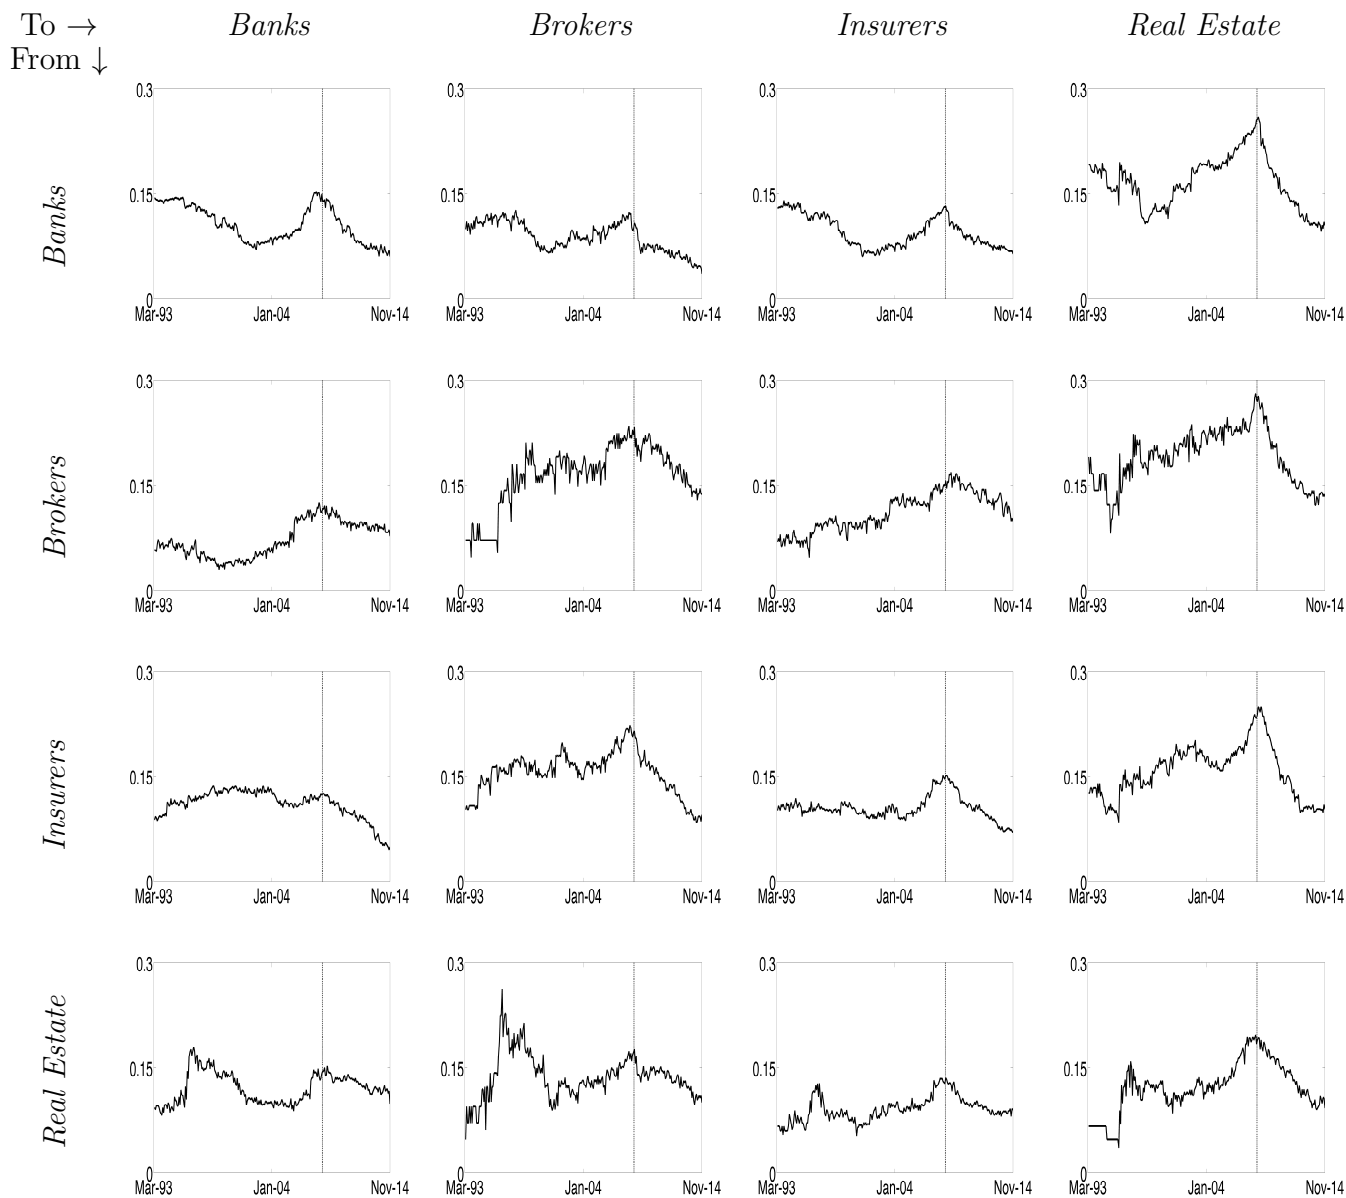


Figure 12: Out-degree at the sectorial level calculated for the network estimated by time-varying parameters. The measure shows the proportion of outgoing connections of all firms in a given sector (indicated by the rows of the chart) reaching financial institutions of another given sector (indicated by the columns of the chart). Plots along the diagonal of the figure show the inter-sectorial interconnectedness. The dashed vertical lines indicate the day of Lehman Brothers bankruptcy, 15 September 2008.

References

- Adams, Z.; R. Füss; and R. Gropp. “Spillover effects among financial institutions: A state-dependent sensitivity value-at-risk approach.” *Journal of Financial and Quantitative Analysis*, 49 (2014), 575–598.
- Alj, A.; K. Jónasson; and G. Mélard. “The exact Gaussian likelihood estimation of time-dependent VARMA models.” In *Computational Statistics and Data Analysis*.
- Allen, F., and D. Gale. “Financial Contagion.” *Journal of Political Economy*, 1 (2000), 1–33.
- Balla, E.; I. Ergen; and M. Migueis. “Tail dependence and indicators of systemic risk for large US depositories.” *Journal of Financial Stability*, 15 (2014), 195–209.
- Barigozzi, M., and C. Brownlees. “NETS: Network Estimation for Time Series.” (2014). Working Paper.
- Barigozzi, M., and M. Hallin. “Networks, Dynamic Factors, and the Volatility Analysis of High-Dimensional Financial Series.” (2015). ECARES Working Paper 2015-34.
- Basel Committee on Banking Supervision. “Global systemically important banks: updated assessment methodology and the higher loss absorbency requirement.” Tech. rep., Bank for International Settlements (2013).
- Baumeister, C., and G. Peersman. “Time-Varying Effects of Oil Supply Shocks on the U.S. Economy.” *American Economic Journal: Macroeconomics*, 5 (2013), 1–29.
- Benoit, S.; G. Colletaz; C. Hurlin; and C. Pérignon. “A theoretical and empirical comparison of systemic risk measures.” (2013). Working paper.
- Billio, M.; M. Getmansky; A. W. Lo; and L. Pelizzon. “Econometric Measures of Connectedness and Systemic Risk in the Finance and Insurance Sectors.” *Journal of Financial Economics*, 104 (2012), 535–559.
- Blasques, F.; S. J. Koopman; A. Lucas; and J. Schaumburg. “Spillover Dynamics for Systemic Risk Measurement Using Spatial Financial Time Series Models.” (2014). Tinbergen Institute Discussion Paper TI2015-107/III.
- Carter, C. K., and R. Kohn. “On Gibbs Sampling for State Space Models.” *Biometrika*, 81 (1994), 541–553.
- Clark, T. E., and M. W. McCracken. “Improving forecast accuracy by combining recursive and rolling forecasts.” *International Economic Review*, 50 (2009), 363–395.
- Cogley, T. “How fast can the new economy grow? A Bayesian analysis of the evolution of trend growth.” *Journal of Macroeconomics*, 27 (2005), 179–207.
- Cogley, T., and T. Sargent. “Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US.” *Review of Economic Dynamics*, 8 (2005), 262–302.

- Committee on Oversight and Government Reform. “Causes and Effects of the Lehman Brothers Bankruptcy.” Hearing (2008). Opening statement of Chairman Henry Waxman.
- Croux, C., and P. Reusens. “Do stock prices contain predictive power for the future economic activity? A Granger causality analysis in the frequency domain.” *Journal of Macroeconomics*, 35 (2013), 93–103.
- D’Agostino, A.; L. Gambetti; and D. Giannone. “Macroeconomic forecasting and structural change.” *Journal of Applied Econometrics*, 28 (2013), 82–101.
- Daniélsson, J.; K. James; M. Valenzuela; and I. Zer. “Model risk of risk models.” (2015). Working Paper.
- Davis, J., and M. Goadrich. “The relationship between precision-recall and roc curves.” In *23rd international conference on Machine learning*. 233–240.
- Demirer, M.; F. X. Diebold; L. Liu; and K. Yılmaz. “Estimating global bank network connectedness.” (2015). Koç University Working Paper Series.
- Diebold, F. X., and K. Yılmaz. “Measuring financial asset return and volatility spillovers, with application to the global equity markets.” *The Economic Journal*, 119 (2009), 158–171.
- Diebold, F. X., and K. Yılmaz. “On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms.” *Journal of Econometrics*, 182 (2014), 119–134.
- Dungey, M.; M. Luciani; and D. Veredas. “Googling SIFIs.” (2013). ECARES Working Paper.
- Fiedor, P. “Granger-causal nonlinear financial networks.” *Journal of Network Theory in Finance*, 1 (2015), 53–82.
- Financial Stability Board. “Policy Measures to Address Systemically Important Financial Institutions.” Tech. rep. (2011).
- Fujita, A.; P. Severino; J. Sato; and S. Miyano. *Advances in Bioinformatics and Computational Biology*, vol. 6268 of *Lecture Notes in Computer Science*, chap. Granger Causality in Systems Biology: Modeling Gene Networks in Time Series Microarray Data Using Vector Autoregressive Models. Springer (2010). 13–24.
- Gai, P., and S. Kapadia. “Contagion in financial networks.” (2010). Bank of England Working Paper No. 383.
- Georg, C.-P. “The effect of the interbank network structure on contagion and common shocks.” *Journal of Banking and Finance*, 37 (2013), 2216–2228.
- Geweke, J. “Measurement of Linear Dependence and Feedback Between Multiple Time Series.” *Journal of the American Statistical Association*, 77 (1982), 304–313.

- Geweke, J. “Measurement of Conditional Linear Dependence and Feedback Between Time Series.” *Journal of the American Statistical Association*, 79 (1984), 907–915.
- Giansante, S.; C. Chiarella; S. Sordi; and A. Vercelli. “Structural contagion and vulnerability to unexpected liquidity shortfalls.” *Journal of Economic Behavior & Organization*, 83 (2012), 558–569.
- Giglio, S.; B. Kell; and S. Pruitt. “Systemic risk and the macroeconomy: An empirical evaluation.” (2013). Chicago Booth Paper No. 12-49.
- Granger, C. W. “Investigating Causal Relations by Econometric Models and Cross-Spectral Methods.” *Econometrica*, 37 (1969), 424–438.
- Hałaj, G.; U. Kochańska; and C. Kok. “Emergence of the EU corporate lending network.” *Journal of Network Theory in Finance*, 1 (2015), 33–76.
- Hautsch, N.; J. Schaumburg; and M. Schienle. “Forecasting systemic impact in financial networks.” *International Journal of Forecasting*, 30 (2014), 781–794.
- Hautsch, N.; J. Schaumburg; and M. Schienle. “Financial network systemic risk contributions.” *Review of Finance*, 19 (2015), 685–738.
- Jacquier, E.; N. G. Polson; and P. E. Rossi. “Bayesian analysis of stochastic volatility models with fat-tails and correlated errors.” *Journal of Econometrics*, 122 (2004), 185–212.
- Kass, R. E., and A. E. Raftery. “Bayes Factor.” *Journal of American Statistical Association*, 90 (1995), 773–795.
- Kenett, D. Y.; M. Tumminello; A. Madi; G. Gur-Gershgoren; R. N. Mantegna; and E. Ben-Jacob. “Dominating clasp of the financial sector revealed by partial correlation analysis of the stock market.” (2010).
- Koop, G.; R. Leon-Gonzalez; and R. W. Strachan. “Dynamic Probabilities of Restrictions in State Space Models: An Application to the Phillips Curve.” *Journal of Business and Economic Statistics*, 28 (2010), 370–379.
- Langfield, S., and K. Soramäki. “Interbank exposure networks.” *Computational Economics* (2014), 1–15.
- Lo, M., and J. Piger. “Is the response of output to monetary policy asymmetric? evidence from a regime-switching coefficient model.” *Journal of Money, Credit, and Banking*, 37 (2005), 865–886.
- Mantegna, R. N. “Hierarchical structure in financial markets.” *The European Physical Journal B*, 11 (1998), 193–197.
- Minoiu, C.; C. Kang; V. Subrahmanian; and A. Berea. “Does financial connectedness predict crises?” In *IMF Working Paper*.

- Montagna, M., and C. Kok. “Multi-layered interbank model for assessing systemic risk.” (2013). Kiel Working Paper No. 1873.
- Nier, E.; J. Yang; T. Yorulmazer; and A. Alentorn. “Network models and financial stability.” *Journal of Economic Dynamics & Control*, 31 (2007), 2033–2060. Tenth Workshop on Economic Heterogeneous Interacting Agents WEHIA 2005.
- Nucera, F.; B. Schwaab; S. J. Koopman; and A. Lucas. “The information in systemic risk rankings.” (2015). Tinbergen Institute Discussion Paper 15-70.
- Peltonen, T. A.; A. Piloju; and P. Sarlin. “Network linkages to predict bank distress.” (2015). Working paper.
- Primiceri, G. “Time Varying Structural Vector Autoregressions and Monetary Policy.” *Review of Economic Studies*, 72 (2005), 821–852.
- Psaradakis, Z.; M. Ravn; and M. Sola. “Markov switching and the money-output relationship.” *Journal of Applied Econometrics*, 20 (2005), 665–683.
- Reusens, P., and C. Croux. “Detecting time variation in the price puzzle: An improved prior choice for time varying parameter var models.” (2014). FEB Research Report KBI 1427.
- Roukney, T.; H. Bersini; H. Pirotte; G. Caldarelli; and S. Battiston. “Default cascades in complex networks: Topology and systemic risk.” *Nature Scientific Reports*, 3 (2013), 1–8.
- Scott, H. S. “Interconnectedness and contagion.” (2012). American Enterprise Institute.
- Seth, A. K. “A matlab toolbox for granger causal connectivity analysis.” *Journal of Neuroscience Methods*, 186 (2010), 262–273.
- Sims, C. “Money, income and causality.” *The American Economic Review*, 62 (1972), 540–552.
- Tam, G.; C. Chang; and Y. Hung. “Gene regulatory network discovery using pairwise granger causality.” *Systems Biology, IET*, 7 (2013), 195–204.
- Tumminello, M.; F. Lillo; and R. N. Mantegna. “Correlation, hierarchies, and networks in financial markets.” *Journal of Economic Behavior & Organization*, 75 (2010), 40–58.
- Upper, C., and A. Worms. “Estimarisk bilateral exposures in the German interbank market: Is there a danger of contagion?” *European Economic Review*, 48 (2004), 827–849.
- Valdes-Sosa, P. A.; J. M. Sanchez-Bornot; A. Lage-Castellanos; M. Vega-Hernandez; J. Bosch-Bayard; L. Melie-Garcia; and E. Canales-Rodriguez. “Estimating Brain Functional Connectivity with Sparse Multivariate Autoregression.” *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360 (2005), 969–981.
- Verdinelli, I., and L. Wasserman. “Computing Bayes Factor Using a Generalization of the Savage-Dickey Density Ratio.” *Journal of the American Statistical Association*, 90 (1995), 614–618.

Wiener, N. *Modern mathematics for engineers*, chap. The theory of prediction. McGraw Hill (1956). 165–90.

Zivot, E., and J. Wang. *Modelling Financial Time Series with S-PLUS*, chap. 9. Springer (2006). 313–360.