Quantile Connectedness: Modelling Tail Behaviour in the Topology of Financial Networks*

Tomohiro Ando Matthew Greenwood-Nimmo Yongcheol Shin

December 30, 2017

Abstract

We develop a new technique to estimate vector autoregressions by quantile regression. A factor structure is used to remove cross-section correlation in the residuals such that the system can be estimated on an equation-by-equation basis using existing quantile regression toolboxes. We use our model to study credit risk spillovers among a panel of 18 sovereigns and their respective financial sectors between January 2006 and February 2012. We show that idiosyncratic credit risk shocks do not propagate strongly at the median but that powerful spillovers occur in both tails. Furthermore, rolling sample analysis reveals marked time-varying tail-dependence. These important features of credit risk transmission are obscured in models estimated using conventional conditional mean estimators.

JEL Codes: C31, C32, C58, E44, G01.

Keywords: Network Modelling; Quantile Vector Autoregression with Common Factors; Quantile Connectedness; Financial–Sovereign Credit Risk Transmission; Tail-Dependence.

*Declarations of interest: None. Role of funding sources in the preparation of this manuscript: None.

1 Introduction

The topology of financial networks is central to the study of systemic risk. An adverse idiosyncratic shock to one part of the financial system poses a threat to systemic stability if there are linkages through which it can propagate to other parts of the system. Measuring the nature and strength of financial market linkages is not only important for risk management strategies but also to inform the policy response to systemic crises. Several techniques for the econometric analysis of financial networks have been proposed in recent years, including those based on Granger-causality and on innovation accounting (e.g. Billio, Getmansky, Lo and Pelizzon, 2012; Diebold and Yilmaz, 2009, 2014; Alter and Beyer, 2014). A common feature of this literature is the reliance on conventional conditional mean estimators such as ordinary least squares. The result is an estimate of the *average* network structure which prevails when an *average* shock affects the system. However, *systemic shocks* are likely to be much larger than average and it need not be the case that large shocks propagate in the same way as smaller shocks. To address this possibility, we develop a new framework which uses regression quantiles to investigate whether the topology of a network changes with the size of the shocks that affect the system.

Our framework builds upon that of Diebold and Yilmaz (2009, 2014), where the $m\,(m-1)$ bilateral interactions among an m-vector of variables, $\boldsymbol{y}_t$, are approximated by the h-step-ahead forecast error variance decomposition (FEVD) of an underlying vector autoregression (VAR) in $\boldsymbol{y}_t$. Consequently, the Diebold–Yilmaz framework answers the question *'how much of the future uncertainty associated with variable i can be attributed to shocks coming from variable j?'*. A major advantage of the Diebold–Yilmaz framework relative to the Granger-causal network analysis of Billio et al. (2012) is that the resulting network is not only directed but also weighted and therefore provides an estimate of the strength of bilateral spillover effects. It has proven to be an influential technique, with applications including the connectedness among equity markets (e.g. Diebold and Yilmaz, 2009), foreign exchange markets (e.g. Baruník, Kočenda and Vácha, 2016; Greenwood-Nimmo, Nguyen and Rafferty, 2016a), systemically important financial institutions (e.g. Demirer, Diebold, Liu and Yilmaz, 2017) and sovereign and corporate credit spreads (e.g. Bostanci and Yilmaz, 2015; Greenwood-Nimmo, Huang and Nguyen, 2017).

Rolling sample analysis is typically used to track the evolution of the network over time, with abrupt increases in connectedness typically being interpreted in relation to systemic shocks. However, there is a tension between this interpretation and the fact that existing applications of the Diebold–Yilmaz framework rely on a range of conventional estimators, including least

squares (e.g. Diebold and Yilmaz, 2009, 2014), the least absolute shrinkage and selection operator, or LASSO (e.g. Greenwood-Nimmo et al., 2017) and elastic net regularisation (e.g. Demirer et al., 2017). Each of these estimators evaluates the relationship between $\boldsymbol{y}_t$ and $\boldsymbol{z}_t = \{\boldsymbol{y}_{t-1}, \boldsymbol{y}_{t-2}, \ldots, \boldsymbol{y}_{t-p}\}$ at the mean of the conditional distribution of $\boldsymbol{y}_t|\boldsymbol{z}_t$. The parameters of a VAR model estimated by least squares, for example, capture the dynamic relationships among the variables in $\boldsymbol{y}_t$ under the assumption that *average* shocks — which are precisely equal to zero by definition — affect the system. The tension arises because systemic shocks are likely to be *larger than average*. Consequently, studies in this literature must implicitly assume that the relationship which prevails at the conditional mean can be generalised to the entire conditional distribution. This is a strong assumption but it is vital if rolling sample analysis based on conditional mean estimators is to provide a valid signal regarding the impact of large systemic shocks.

We relax this assumption by developing a new approach based on the premise that if one wishes to analyse the network structure associated with a large shock — for example a shock in the ninety-fifth percentile of the size distribution of shocks — then one must set aside conditional mean estimators and instead fit the VAR model at the ninety-fifth percentile by quantile regression. Following Koenker and Bassett (1978), quantile regression can be used to evaluate the dependence of $\boldsymbol{y}_t$ on $\boldsymbol{z}_t$ over the entire range of the conditional distribution of $\boldsymbol{y}_t|\boldsymbol{z}_t$. At the time of writing, two approaches to the estimation of quantile VAR models have been proposed by Cecchetti and Li (2008) and Schüler (2014), respectively.[1] Cecchetti and Li develop an equation-by-equation estimation framework for VAR models with cross-sectionally correlated residuals, while Schüler develops a Bayesian framework for the analysis of structural VAR models.

Our framework is distinct from these existing methods by virtue of our treatment of the VAR residuals. We assume that the cross-sectional correlation of the VAR residuals is driven by a finite number of common factors. This assumption serves two purposes. First, by purging the common component from the VAR residuals, we are able to isolate the idiosyncratic shock to each variable in the system. Not only does this align our framework with the large literature on systematic vs. idiosyncratic variations in finance — an important consideration if one wishes to use network models to analyse the propagation of idiosyncratic shocks, for example — but it also

[1] The method of Cecchetti and Li (2008) has subsequently been applied by Linnemann and Winkler (2016) and Zhu, Su, Guo and Ren (2016). A related paper by Chavleishvili and Manganelli (2016) develops a framework for quantile impulse response analysis of a bivariate system with one endogenous and one exogenous variable.

addresses the likelihood that the failure to account for sources of common variation may generate substantial biases in the analysis of networks. Specifically, if an omitted common factor affects all variables in the VAR system, then the proportion of the forecast error variance which should rightly be attributed to that factor must instead be attributed to one or more of the endogenous variables in the system, thereby upwardly biasing the estimated spillover effects.[2] The issue of spurious Granger causality arising due to the omission of sources of common variation is well-known but the impact of omitted common factors on Diebold–Yilmaz networks has not been adequately explored to date (Greenwood-Nimmo, Nguyen and Shin, 2016b). Second, the use of a factor structure simplifies the estimation problem substantially because it renders the system of equations that comprise the VAR model cross-sectionally independent. Given that the errors are uncorrelated across equations, the factor VAR system can be directly estimated on an equation-by-equation basis using the standard quantile regression commands built in to many statistical software packages (such as Roger Koenker's `quantreg` package in R).

Our approach answers a modified version of the question addressed by the Diebold–Yilmaz framework: *'how much of the future uncertainty associated with variable i can be attributed to idiosyncratic shocks coming from variable j as the shock size varies?'*. Consequently, our technique is ideally suited to the study of the propagation of idiosyncratic risk shocks and contagion, the latter of which is often defined in relation to the difference in the way that shocks propagate during rare events relative to normal times (e.g. Caporin, Pelizzon, Ravazzolo and Rigobon, 2013). We introduce the term *quantile connectedness* to distinguish between our quantile-based approach to network analysis and the established mean-based approach.

We apply our technique to study spillovers of idiosyncratic credit risk between sovereigns and national financial sectors over the period January 2006 to February 2012, both within and across borders. The study of credit risk transmission has become an important focus among policy institutions, with a particular concern for the emergence of feedback loops and destabilising spirals in credit markets (e.g. International Monetary Fund, 2013, pp. 65-6). Acharya, Drechsler and Schnabl (2014) provide compelling evidence of just such an adverse feedback effect. They demonstrate that the financial sector bailouts undertaken by many developed countries in 2008 amounted to a substantial transfer of private sector credit risk onto the public sector at a time of rapid public debt growth. This combination ultimately proved untenable in several countries

[2] In this paper, we focus on analysing direct spillovers of idiosyncratic credit risk having controlled for common systematic factors. Although it is not our focus here, one may also be interested in analysing the indirect propagation of shocks via the common factors. The Diebold–Yilmaz framework accounts for both direct and indirect linkages but it does not allow them to be analysed separately.

and led to a resurgence of systemic risk driven by the emergence of an adverse feedback loop between sovereign credit risk and financial sector credit risk. Variants of the Diebold–Yilmaz technique have been applied to the analysis of the sovereign–financial credit risk network by Alter and Beyer (2014) and Greenwood-Nimmo et al. (2017), although our paper represents the first attempt to study quantile-variation in the structure of the credit risk network.

We follow the existing literature and use credit default swap (CDS) spreads to measure credit risk.[3] The existence of a factor structure in the cross-section of CDS spreads has been documented by Pan and Singleton (2008), Longstaff, Pan, Pedersen and Singleton (2011), Fender, Hayo and Neuenkirch (2012) and Ang and Longstaff (2013), among others. To isolate the idiosyncratic variation in the CDS spreads, our model includes the following observed common factors, which draw on those identified by Longstaff et al. (2011): (i) global stock market factors measured via the three Fama-French factors; (ii) US macroeconomic conditions proxied by the five-year US Treasury yield; (iii) funding liquidity measured via the TED spread and the Euribor-DeTBill spread; (iv) investor risk appetite captured by the S&P 500 variance risk premium, the five-year US Treasury term premium and the US corporate investment grade and high yield spreads; (v) a selection of ITRAXX credit default indices which capture pan-European credit risk factors; and (vi) bilateral exchange rate fluctuations measured for each currency against the US Dollar.

Our first finding is that the topology of the credit risk network varies significantly with the shock size. We find that the effects of small idiosyncratic credit risk shocks in the central 70% of the conditional distribution are predominantly localised. Bilateral spillovers account for no more than 9% of the five-days-ahead forecast error variance (FEV) over this range and for just 4.34% at the conditional median. However, this is not true of large shocks. Large adverse shocks in the right tail of the conditional distribution spread forcefully through the financial system, with bilateral spillovers accounting for 22.87% and 83.32% of FEV at the ninety-fifth and ninety-ninth conditional quantiles, respectively. This finding accords with the existing evidence on increased financial market comovements under adverse conditions (e.g. Ang and Bekaert, 2002). Interestingly, we also find evidence that large beneficial shocks propagate strongly, with bilateral spillovers accounting for 21.07% and 76.57% of the FEV at the fifth

[3] A CDS contract operates like an insurance agreement negotiated between two parties, one of whom holds a risky bond and the other of whom agrees to absorb losses arising in the event that the bond issuer defaults. The CDS spread defines the price that the bondholder must pay to the protection seller. Blanco, Brennan and Marsh (2005) and Gyntelberg, Hördahl, Ters and Urban (2013) show that the CDS market is the leading forum for credit risk price discovery, providing more timely signals of changes in the credit risk environment than bond yield spreads.

and first conditional quantiles, respectively. The finding that large shocks in both tails spillover strongly is consistent with the existing literature on good and bad contagion (e.g. Londono, 2016). Crucially, when the model is estimated at the conditional mean by OLS, this quantile variation is averaged out and bilateral spillovers are found to account for 11.49% of the FEV. This result arises by definition because OLS is equivalent to an equally-weighted average of quantile regression estimators over the entire conditional distribution. This cautions against the use of network models estimated using conditional mean estimators to analyse the spillovers associated with extreme events.

Our second result is that the adverse feedback loop documented by Acharya et al. (2014) manifests as a marked intensification of the bidirectional feedback between each sovereign and its domestic financial sector in the presence of large adverse shocks in the right tail of the conditional distribution. However, we once again find that this behaviour is not unique to adverse shocks — there is a similar intensification of bidirectional feedback in the left tail. This leads us to conclude that the feedback loop described by Acharya et al. (2014) is associated with a vicious circle which leads to the amplification of bad news but that the same feedback effect can rapidly reduce credit spreads when beneficial shocks occur.

Our final result is derived from rolling sample estimation of our model. In this way, we demonstrate that the dependence structure that exists among the cross-section of sovereigns and financial institutions displays marked time-variation. The time-variation of bilateral credit risk spillovers has already been demonstrated at the conditional mean (e.g. Greenwood-Nimmo et al., 2017). However, we are the first to demonstrate that the time-variation in the network topology observed in the tails of the conditional distribution does not closely resemble the time-variation observed at either the conditional mean or median. This is an important finding given the relevance of tail-dependence for financial supervision and risk management (e.g. Betz, Hautsch, Peltonen and Schienle, 2016) and it suggests that the implications derived from network models evaluated by conventional conditional mean estimators cannot necessarily be generalised to the tails. We show that major adverse events are associated with an increase in average connectedness but that their effects on the tails differs. Specifically, we find that right-tail-dependence increases while left-tail-dependence decreases. This combination implies that bad news leads to an increase in the propensity for the destabilising propagation of adverse shocks coupled with a reduction in the stabilising propagation of beneficial shocks. By contrast, stabilising policy interventions which reduce average connectedness tend to increase

left-tail-dependence while reducing right-tail-dependence. These findings lead us to develop a new measure of relative tail-dependence which draws attention to the negative correlation of left- and right-tail-dependence. We suggest that this negative correlation may arise from the aggregate behaviour of market participants if the information revealed by a major event in either tail causes a non-trivial proportion of market participants to focus disproportionately on further events occurring in that tail while paying less attention to events in the other tail.

Aside from the work on empirical network modelling, we wish to highlight three strands of literature to which our paper is related. First, our use of quantile regression to study extreme events closely resembles a branch of the systemic risk literature which is well-represented by Caporin et al. (2013) and Betz et al. (2016), both of which use quantile regression to study the propagation of adverse shocks through the financial markets. Second, our concept of quantile connectedness has a natural link to value-at-risk (VaR) and associated concepts such as CoVaR (Adrian and Brunnermeier, 2016). VaR is widely used by investors to measure the potential loss that they may endure on their positions due to adverse shocks over a defined horizon and at a predetermined confidence level. To illustrate the conceptual link, assume that investors maintain sufficient capital reserves to cover the VaR at the 95% confidence level. In this case, losses up to the VaR can be absorbed within the capital buffer. However, the probability of an adverse shock sufficiently severe to generate a loss in excess of the VaR is 5%. Investors may find the losses caused by such large shocks untenable, raising the possibility of default and insolvency. As a result, the transmission of risk among counterparties may be considerably stronger in the case of large shocks than small shocks. This offers a partial explanation of the quantile-variation that we document in the topology of the credit risk network. Finally, the notion that the topology of a network may vary with the size of the shocks affecting the system is related to Acemoglu, Ozdaglar and Tahbaz-Salehi's (2015) insight that a phase change may occur whereby dense financial networks are resilient to small shocks but can be vulnerable to cascading failures in the event of a large adverse shock.

This paper proceeds as follows. In Section 2, we outline the quantile factor VAR model and derive the associated generalised forecast error variance decomposition which is then used to construct network statistics. We provide a detailed discussion of the construction and properties of our dataset in Section 3 before presenting our estimation results in Section 4. We conclude in Section 5.

2 Quantile Connectedness

In this section, we propose a new framework for the graphical analysis of VAR models estimated at a given conditional quantile, $\tau \in (0,1)$. To this end, we develop a framework for the equation-by-equation estimation of a VAR model by quantile regression where a factor structure is used to distinguish between the common and idiosyncratic components of the error process. We then derive the associated h-step-ahead forecast error variance decomposition, which forms the basis for network analysis in the tradition of Diebold and Yilmaz (2009, 2014).

2.1 The Quantile VAR Model

Consider a multi-country environment where countries are indexed by $i = 1, 2, \ldots, N$ and where time periods are indexed by $t = 1, 2, \ldots T$. Let $\boldsymbol{y}_{it} = (\Delta s_{it}, \Delta f_{it})'$, where s_{it} and f_{it} denote the sovereign and financial sector CDS spreads for the ith country, which will be discussed in detail below. Moreover, let the $m \times 1$ vector $\boldsymbol{y}_t = (\boldsymbol{y}'_{1t}, \boldsymbol{y}'_{2t}, \ldots, \boldsymbol{y}'_{Nt})'$ collect the credit spreads for all N countries in the system, where $m = 2N$. Implementation of the Diebold–Yilmaz technique requires us to specify a model capturing the dynamics of $\boldsymbol{y}_t$. To this end, consider the following VAR model which expresses the sovereign and financial sector CDS spreads for the ith country as a function of the lagged sovereign and financial sector CDS spreads for every country in the system:

$$\boldsymbol{y}_{it} = \boldsymbol{\mu}_i + \sum_{j=1}^{p} \boldsymbol{\Phi}_{ij} \boldsymbol{y}_{t-j} + \boldsymbol{e}_{it} \tag{1}$$

where $\underset{2\times 1}{\boldsymbol{\mu}_i}$ is a vector of intercepts, $\underset{2\times m}{\boldsymbol{\Phi}_{ij}}$ is the parameter matrix at the jth lag and the regression residuals $\underset{2\times 1}{\boldsymbol{e}_{it}} \sim (\boldsymbol{0}, \underset{2\times 2}{\boldsymbol{\Sigma_i}})$, where $\underset{2\times 2}{\boldsymbol{\Sigma_i}}$ is a positive definite covariance matrix. By stacking (1) for all countries in the system, we obtain the following VAR system for $\boldsymbol{y}_t$:

$$\boldsymbol{y}_t = \boldsymbol{\mu} + \sum_{j=1}^{p} \boldsymbol{\Phi}_j \boldsymbol{y}_{t-j} + \boldsymbol{e}_t \tag{2}$$

where $\underset{m\times 1}{\boldsymbol{\mu}} = (\boldsymbol{\mu}'_1, \boldsymbol{\mu}'_2, \ldots, \boldsymbol{\mu}'_N)'$ is a vector of intercepts, $\underset{m\times m}{\boldsymbol{\Phi}_j} = \left(\boldsymbol{\Phi}'_{1j}, \boldsymbol{\Phi}'_{2j}, \ldots, \boldsymbol{\Phi}'_{Nj}\right)'$ is the jth autoregressive parameter matrix and the residual process $\underset{m\times 1}{\boldsymbol{e}_t} = (\boldsymbol{e}'_{1t}, \boldsymbol{e}'_{2t}, \ldots, \boldsymbol{e}'_{Nt})' \sim (\boldsymbol{0}, \boldsymbol{\Sigma})$ where $\underset{m\times m}{\boldsymbol{\Sigma}}$ is positive definite.

The order of the VAR model (2), p, is estimated consistently using the Schwarz Information Criterion and should be sufficient to yield serially uncorrelated residuals. Nonetheless, the residuals will typically exhibit contemporaneous cross-section correlation and so $\boldsymbol{\Sigma}$ is likely to

be non-diagonal.[4] Conditional mean estimation of unrestricted VARs of this type is straightforward, and can be achieved on an equation-by-equation basis using ordinary least squares. Conditional quantile estimation is more challenging, however. Given that each equation in (2) shares a common set of right hand side variables, the estimation problem has a seemingly unrelated regressions (SUR) structure which has led several authors to pursue an equation-by-equation quantile regression estimation strategy (e.g. Cecchetti and Li, 2008; Linnemann and Winkler, 2016; Zhu et al., 2016). However, from equation (5) in Zellner and Ando (2010), it is clear that this approach sets the off-diagonal elements of the covariance matrix of $\boldsymbol{e}_t$ to zero, which amounts to the assumption of cross-section independence. The failure to adequately account for the cross-section correlation among the regression residuals is likely to bias the resulting parameter estimates.

We address this issue by modelling the cross-section correlation in the residuals as the result of an $f \times 1$ vector of common factors. In this way, we are able to separate the systematic and idiosyncratic components of the error process, thereby aligning our approach with the large literature on the distinction between systematic and idiosyncratic risks (see Feldhütter and Nielsen, 2012, for a recent example which focuses on credit spreads). Furthermore, where one wishes to analyse spillover effects between variables, it is important to focus on the idiosyncratic variation having purged any systematic variation or else one is likely to obtain a biased estimate of the spillover intensity, a phenomenon that we demonstrate in Section 4.1 below. Formally, we assume that:

$$\boldsymbol{e}_{it} = \boldsymbol{\lambda}_i' \boldsymbol{f}_t + \boldsymbol{v}_{it} \tag{3}$$

and, by extension, that:

$$\boldsymbol{e}_t = \boldsymbol{\Lambda} \boldsymbol{f}_t + \boldsymbol{v}_t \tag{4}$$

where $\underset{f\times 1}{\boldsymbol{f}_t}$ is a vector of common factors, $\underset{m\times f}{\boldsymbol{\Lambda}} = \left(\boldsymbol{\lambda}_1', \boldsymbol{\lambda}_2', \ldots, \boldsymbol{\lambda}_N'\right)'$ is a matrix of heterogeneous factor loadings and where $\underset{m\times 1}{\boldsymbol{v}_t} = (\boldsymbol{v}_{1t}', \boldsymbol{v}_{2t}', \ldots, \boldsymbol{v}_{Nt}')' \sim (\boldsymbol{0}, \boldsymbol{\Omega})$ contains the idiosyncratic components of $\boldsymbol{e}_t$ which are mutually uncorrelated such that $\underset{m\times m}{\boldsymbol{\Omega}}$ is diagonal. Combining (2) and (4)

[4] When using VARs for macroeconomic analysis, it is common to transform the reduced form VAR (2) into a structural counterpart with uncorrelated disturbance terms to which one can attach a structural interpretation. This is typically achieved either by imposing Wold-causality as in Sims (1986), short-run exclusion restrictions as in Blanchard and Watson (1986) or long-run restrictions as in Blanchard and Quah (1989). However, the application of these traditional methods becomes increasingly challenging as the dimension of the model increases.

yields the following factor VAR model:

$$\boldsymbol{y}_t = \boldsymbol{\mu} + \sum_{j=1}^{p} \boldsymbol{\Phi}_j \boldsymbol{y}_{t-j} + \boldsymbol{\Lambda} \boldsymbol{f}_t + \boldsymbol{v}_t \tag{5}$$

The factors, $\boldsymbol{f}_t$, may be either observed or latent — we limit our attention to the case of observed factors. This confers at least two benefits relative to unobserved factor approaches such as principal components analysis (e.g. Bai, 2003, 2009) or the common correlated effects framework of Pesaran (2006). First, it is often easier to achieve an economically meaningful interpretation of observed factors than latent factors. Second, an observed factor structure is parsimonious and can be feasibly implemented in relatively small datasets. By contrast, the reliable estimation of principal components requires a relatively high-dimensional dataset. Likewise, the common correlated effects approach can lead to a proliferation of estimated parameters because the unobserved common factors are approximated by augmenting the model with p lags of the cross-section averages of $\boldsymbol{y}_t$.

Given that the error terms in (5) are uncorrelated by construction, the model can be estimated by equation-by-equation quantile regression without loss of generality. We write the quantile factor VAR (QFVAR) model evaluated at the τth conditional quantile as follows:

$$\boldsymbol{y}_t = \boldsymbol{\mu}_{(\tau)} + \sum_{j=1}^{p} \boldsymbol{\Phi}_{j(\tau)} \boldsymbol{y}_{t-j} + \boldsymbol{\Lambda}_{(\tau)} \boldsymbol{f}_t + \boldsymbol{v}_{t(\tau)} \tag{6}$$

where $\tau \in (0, 1)$ is a given quantile index. Following Koenker and Xiao (2006), we assume that the optimal lag order for the conditional mean model remains valid at every conditional quantile. Assuming that $Q_\tau \left(\boldsymbol{v}_{t(\tau)} | \mathcal{F}_{t-1} \right) = 0$, where $\mathcal{F}_{t-1}$ denotes the information set available at time $t-1$, then:[5]

$$Q_\tau \left(\boldsymbol{y}_t | \mathcal{F}_{t-1} \right) = \boldsymbol{\mu}_{(\tau)} + \sum_{j=1}^{p} \boldsymbol{\Phi}_{j(\tau)} \boldsymbol{y}_{t-j} + \boldsymbol{\Lambda}_{(\tau)} \boldsymbol{f}_t \tag{7}$$

To illustrate the quantile regression procedure, it is useful to first re-write the ith equation of (6) compactly as follows:

$$y_{it} = \boldsymbol{\beta}'_{i(\tau)} \boldsymbol{z}_t + v_{it(\tau)} \tag{8}$$

for $i = 1, 2, \ldots, m$ where $\boldsymbol{z}_t$ is an $(mp + f + 1) \times 1$ vector containing all of the regressors including

[5] Note that the residual covariance matrix, $\boldsymbol{\Omega}$, is the same for all τ. This assumption conforms with the precedent in the frequentist literature (e.g. Cecchetti and Li, 2008), although a Bayesian technique for the estimation of a VAR model with a quantile-dependent covariance matrix has been proposed by Schüler (2014).

the intercept and $\boldsymbol{\beta}_{i(\tau)}$ contains the corresponding regression coefficients evaluated at the τth conditional quantile. We impose the usual assumption that the conditional quantile model is correctly specified:

$$E\left(\psi_\tau\left(v_{it(\tau)}\right)|\boldsymbol{z}_t\right)=0 \tag{9}$$

where $\psi_\tau(z)=\tau-1_{[z\leq 0]}$. This assumption implies that $\int_{-\infty}^{\boldsymbol{\beta}_{i(\tau)}\boldsymbol{z}_t} f_{y_{it}|\boldsymbol{z}_t}(t|\boldsymbol{z}_t)dt=\tau$, where $f_{y_{it}|\boldsymbol{z}_t}(t|\boldsymbol{z}_t)$ is the density of y_{it} conditional on $\boldsymbol{z}_t$. For a fixed value of τ, the single-step quantile regression estimates are obtained as follows:

$$\min_{\boldsymbol{\beta}_{i(\tau)}}\sum_{t=1}^{T}\xi_\tau\left(y_{it}-\boldsymbol{\beta}_{i(\tau)}'\boldsymbol{z}_t\right) \tag{10}$$

where $\xi_\tau(z)$ is the check loss function defined as $\xi_\tau(z)=z(\tau-1_{[z\leq 0]})$ as in Koenker and Hallock (2001). The solution to this minimisation, denoted $\widehat{\boldsymbol{\beta}}_{i(\tau)}$, is consistent and asymptotically normal under the assumption that the quantile specification is correct and subject to a number of mild regularity conditions (see Koenker, 2005, for further details).

2.2 Quantile Forecast Error Variance Decomposition

The functional form of the factor VAR model (5) is identical to a VARX$(p,0)$ model — that is, a VAR model with p lags of a set of endogenous variables augmented with the contemporaneous values of a set of exogenous variables — in the special case that the residual covariance matrix is diagonal.[6] Two approaches to innovation accounting with VARX models have been pursued in the literature: (i) to conduct innovation accounting by conditioning on the values of the exogenous variables; and (ii) to augment the VARX specification with an auxiliary marginal model for the exogenous variables and then to undertake innovation accounting with respect to the augmented system. Given that our interest is in modelling the propagation of the idiosyncratic shocks as opposed to the effects of systemwide shocks, we pursue the former option. In addition, as the computation of the FEVD for VARX models evaluated at the conditional mean is well established (e.g. Garratt, Lee, Pesaran and Shin, 2006), we limit our attention to the QFVAR case.

[6]The model may be easily extended to include lags of the factors — i.e. one may estimate a VARX(p,q) model with $q>0$ — although this would substantially increase the number of parameters to be estimated.

We start by re-writing the QFVAR model (6) as follows:

$$\boldsymbol{y}_t = \sum_{j=1}^{p} \boldsymbol{\Phi}_{j(\tau)} \boldsymbol{y}_{t-j} + \boldsymbol{\Lambda}^*_{(\tau)} \boldsymbol{f}^*_t + \boldsymbol{v}_{t(\tau)} \tag{11}$$

where $\boldsymbol{\Lambda}^*_{(\tau)} = \left(\boldsymbol{\mu}_{(\tau)}, \boldsymbol{\Lambda}_{(\tau)}\right)$, $\boldsymbol{f}^*_t = \left(1, \boldsymbol{f}'_t\right)'$. The Wold representation of (11) can be written as:

$$Q_\tau\left(\boldsymbol{y}_t | \mathcal{F}_{t-1}\right) = \sum_{j=0}^{\infty} \boldsymbol{B}_{j(\tau)} \boldsymbol{v}_{t-j(\tau)} + \sum_{j=0}^{\infty} \boldsymbol{C}_{j(\tau)} \boldsymbol{f}^*_{t-j} \tag{12}$$

where $\boldsymbol{B}_{j(\tau)} = \boldsymbol{\Phi}_{1(\tau)} \boldsymbol{B}_{j-1(\tau)} + \boldsymbol{\Phi}_{2(\tau)} \boldsymbol{B}_{j-2(\tau)} + \ldots$ for $j = 1, 2, \ldots$ with $\boldsymbol{B}_{0(\tau)} = \boldsymbol{I}_m$ and $\boldsymbol{B}_{j(\tau)} = \boldsymbol{0}$ for $\ell < 0$ and where $\boldsymbol{C}_{j(\tau)} = \boldsymbol{B}_{j(\tau)} \boldsymbol{\Lambda}^*_{(\tau)}$.

Following the precedent established by Cecchetti and Li (2008, p. 12) in the context of multivariate forecasting with dynamic quantile regressions, we assume that the quantile index is fixed throughout the forecast horizon.[7] Under this assumption, (12) implies that the vector of forecast errors associated with the prediction of $\boldsymbol{y}_{t+h}$ conditional on the information at time $t-1$ and on the common factors is given by:

$$\boldsymbol{u}_{t+h(\tau)} = \sum_{\ell=0}^{h} \boldsymbol{B}_{\ell(\tau)} \boldsymbol{v}_{t+h-\ell(\tau)}$$

and the total forecast error variance matrix is:

$$Cov\left(\boldsymbol{u}_{t+h(\tau)}\right) = \sum_{\ell=0}^{h} \boldsymbol{B}_{\ell(\tau)} \boldsymbol{\Omega} \boldsymbol{B}'_{\ell(\tau)} \tag{13}$$

Now, consider the covariance matrix of the forecast errors associated with predicting $\boldsymbol{y}_{t+h}$ given values of the shocks to the ith equation, $v_{it(\tau)}, v_{i,t+1(\tau)}, \ldots, v_{i,t+h(\tau)}$:

$$\boldsymbol{u}^{(i)}_{t+h(\tau)} = \sum_{\ell=0}^{h} \boldsymbol{B}_{\ell(\tau)} \left(\boldsymbol{v}_{t+h-\ell(\tau)} - E\left(\boldsymbol{v}_{t+h-\ell(\tau)} | v_{i,t+h-\ell(\tau)}\right)\right) \tag{14}$$

Under the assumption that $\boldsymbol{v}_{t(\tau)} \sim i.i.d.\left(\boldsymbol{0}, \boldsymbol{\Omega}\right)$ with $\boldsymbol{\Omega} = diag\left(\omega_{11}, \omega_{22}, \ldots, \omega_{mm}\right)$, we have:

$$\begin{aligned} E\left(\boldsymbol{v}_{t+h-\ell(\tau)} | v_{i,t+h-\ell(\tau)}\right) &= \left(\omega_{ii}^{-1} \boldsymbol{\Omega} \mathbf{e}_i\right) v_{i,t+h-\ell(\tau)} \\ &= \mathbf{e}_i v_{i,t+h-\ell(\tau)} \end{aligned} \tag{15}$$

[7] In principle, one could allow for the quantile index to vary across the forecast horizon, although it is not clear how one could systematically determine the time-path of the quantile index from one horizon to the next.

where $\mathbf{e}_i$ is an $m \times 1$ selection vector with its ith element set to 1 and all other elements set to zero and where $\omega_{ii}^{-1}\boldsymbol{\Omega}\mathbf{e}_i = \mathbf{e}_i$. Substituting this result into (14), we obtain:

$$\boldsymbol{u}_{t+h(\tau)}^{(i)} = \sum_{\ell=0}^{h} \boldsymbol{B}_{\ell(\tau)} \left(\boldsymbol{v}_{t+h-\ell(\tau)} - \left(\omega_{ii}^{-1}\boldsymbol{\Omega}\mathbf{e}_i \right) v_{i,t+h-\ell(\tau)} \right) \tag{16}$$

and taking the unconditional expectation yields:

$$Cov\left(\boldsymbol{u}_{t+h(\tau)}^{(i)} \right) = \sum_{\ell=0}^{h} \boldsymbol{B}_{\ell(\tau)}\boldsymbol{\Omega}\boldsymbol{B}_{\ell(\tau)}' - \omega_{ii}^{-1} \sum_{\ell=0}^{h} \boldsymbol{B}_{\ell(\tau)}\boldsymbol{\Omega}\mathbf{e}_i\mathbf{e}_i'\boldsymbol{\Omega}\boldsymbol{B}_{\ell(\tau)}' \tag{17}$$

Therefore, the decline in the h-step-ahead forecast error variance of $\boldsymbol{y}_t$ obtained as a result of conditioning on future shocks to the ith equation is given by:

$$\begin{aligned} \boldsymbol{\Delta}_{ih(\tau)} &= Cov\left(\boldsymbol{u}_{t+h(\tau)} - \boldsymbol{u}_{t+h(\tau)}^{(i)} \right) \\ &= \omega_{ii}^{-1} \sum_{\ell=0}^{h} \boldsymbol{B}_{\ell(\tau)}\boldsymbol{\Omega}\mathbf{e}_i\mathbf{e}_i'\boldsymbol{\Omega}\boldsymbol{B}_{\ell(\tau)}' \end{aligned} \tag{18}$$

Scaling the jth diagonal element of $\boldsymbol{\Delta}_{ih(\tau)}$ — that is, $\mathbf{e}_j'\boldsymbol{\Delta}_{ih(\tau)}\mathbf{e}_j$ — by the h-step-ahead forecast error variance of the jth variable in $\boldsymbol{y}_t$, we obtain:

$$\begin{aligned} FEVD\left(y_{jt}; u_{it(\tau)}, h \right) &= \frac{\omega_{ii}^{-1} \sum_{\ell=0}^{h} \mathbf{e}_j' \left(\boldsymbol{B}_{\ell(\tau)}\boldsymbol{\Omega}\mathbf{e}_i\mathbf{e}_i'\boldsymbol{\Omega}\boldsymbol{B}_{\ell(\tau)}' \right) \mathbf{e}_j}{\sum_{\ell=0}^{h} \mathbf{e}_j'\boldsymbol{B}_{\ell(\tau)}\boldsymbol{\Omega}\boldsymbol{B}_{\ell(\tau)}'\mathbf{e}_j} \\ &= \frac{\omega_{ii}^{-1} \sum_{\ell=0}^{h} \left(\mathbf{e}_j'\boldsymbol{B}_{\ell(\tau)}\boldsymbol{\Omega}\mathbf{e}_i \right)^2}{\sum_{\ell=0}^{h} \mathbf{e}_j'\boldsymbol{B}_{\ell(\tau)}\boldsymbol{\Omega}\boldsymbol{B}_{\ell(\tau)}'\mathbf{e}_j} \end{aligned} \tag{19}$$

for $\ell = 0, 1, \ldots, h$ and $i, j = 1, \ldots, m$, where $\mathbf{e}_j$ selects the predicted variable and $\mathbf{e}_i$ selects the source innovation. Consequently, $FEVD\left(y_{jt}; u_{it(\tau)}, h \right)$ measures the proportion of the h-step-ahead forecast error variance of the jth variable in $\boldsymbol{y}_t$ accounted for by the ith idiosyncratic innovation, $v_{it(\tau)}$. Note that any variation in the FEVD over quantiles is due to the quantile-dependence of the parameters of the Wold representation of the QFVAR model, the $\boldsymbol{B}_{\ell(\tau)}$'s. The quantile-variation in the parameters reflects changes in the dynamic relationships among $\boldsymbol{y}_t$ as shocks of different size affect the system. To facilitate the comparison of the FEVDs across quantiles, rather than allowing the size of the shock to differ over quantiles, we scale the ith shock to one standard deviation of the ith regression residual for all τ. By considering an identical shock at every quantile, we are able to focus keenly on the quantile-variation in the

dependence structure captured by the parameters of the QFVAR model.

2.3 The Diebold–Yilmaz Approach to Network Analysis

Based on our definition of the quantile FEVD in (19), it is straightforward to generalise the Diebold–Yilmaz framework for conditional mean network analysis to the conditional quantile setting. The h-step-ahead $m \times m$ spillover matrix for $\boldsymbol{y}_t$ evaluated at the τth conditional quantile may be written as follows:

$$\mathbb{A}^{(h)}_{(\tau)} = \begin{bmatrix} \theta^{(h)}_{1\leftarrow 1,(\tau)} & \theta^{(h)}_{1\leftarrow 2,(\tau)} & \cdots & \theta^{(h)}_{1\leftarrow m,(\tau)} \\ \theta^{(h)}_{2\leftarrow 1,(\tau)} & \theta^{(h)}_{2\leftarrow 2,(\tau)} & \cdots & \theta^{(h)}_{2\leftarrow m,(\tau)} \\ \vdots & \vdots & \ddots & \vdots \\ \theta^{(h)}_{m\leftarrow 1,(\tau)} & \theta^{(h)}_{m\leftarrow 2,(\tau)} & \cdots & \theta^{(h)}_{m\leftarrow m,(\tau)} \end{bmatrix} \tag{20}$$

where we define $\theta^{(h)}_{j\leftarrow i,(\tau)} \equiv FEVD\left(y_{jt}; u_{it(\tau)}, h\right)$ to simplify the notation and where $\theta^{(h)}_{j\leftarrow i,(\tau)}$ measures the spillover of idiosyncratic shocks affecting variable i onto variable j. Note that we need not apply the row sum normalisation used by Diebold and Yilmaz (2014) due to the diagonality of the covariance matrix, Ω, which ensures that $\sum_{i=1}^{m} \theta^{(h)}_{j\leftarrow i,(\tau)} = 1$, $j = 1, 2, \ldots, m$ by construction.

Based on $\mathbb{A}^{(h)}_{(\tau)}$, we may define the following summary measures of the network topology at the τth conditional quantile:

$$\begin{aligned} O^{(h)}_{i\leftarrow i,(\tau)} &= \theta^{(h)}_{i\leftarrow i,(\tau)} \\ F^{(h)}_{i\leftarrow \bullet,(\tau)} &= \sum_{j=1, j\neq i}^{m} \theta^{(h)}_{i\leftarrow j,(\tau)} \\ T^{(h)}_{\bullet\leftarrow i,(\tau)} &= \sum_{j=1, j\neq i}^{m} \theta^{(h)}_{j\leftarrow i,(\tau)} \\ N^{(h)}_{i\leftarrow i,(\tau)} &= T^{(h)}_{\bullet\leftarrow i,(\tau)} - F^{(h)}_{i\leftarrow \bullet,(\tau)} \end{aligned} \tag{21}$$

The proportion of the h-step-ahead FEV of the i-th variable that can be attributed to shocks to variable i itself is known as the *own variance share*, $O^{(h)}_{i\leftarrow i,(\tau)}$. The *from* (or *in*) degree of variable i, $F^{(h)}_{i\leftarrow \bullet,(\tau)}$, measures the total spillover from the system to variable i. As such, it measures the dependence of variable i on external conditions. Likewise, the *to* (or *out*) degree of variable i, $T^{(h)}_{\bullet\leftarrow i,(\tau)}$, captures the total spillover from variable i to the system, which measures the influence

of the ith node in the network. $N^{(h)}_{i\leftarrow i,(\tau)}$ is therefore a natural measure of the net directional connectedness of variable i. Note that $O^{(h)}_{i\leftarrow i,(\tau)} + F^{(h)}_{i\leftarrow\bullet,(\tau)} = 1$, $i = 1, 2, \ldots, m$ by construction but that $T^{(h)}_{\bullet\leftarrow i,(\tau)}$ can be greater than or less than one. Finally, the spillover index evaluated at the τth conditional quantile is given by:

$$S^{(h)}_{(\tau)} = m^{-1} \sum_{i=1}^{m} F^{(h)}_{i\leftarrow\bullet,(\tau)} \tag{22}$$

3 Credit Spread Data

Our model includes the following $N = 18$ countries: Austria†, Australia, Belgium†, China, France†, Germany†, Greece†, Ireland†, Italy†, Japan, the Netherlands†, Norway*, Portugal†, Russia, Spain†, Sweden*, the U.K.* and the U.S.A. Eurozone members are marked with a dagger symbol, while European countries which do not use the Euro are marked with an asterisk. Our dataset is sampled at daily frequency over the period 03-Jan-2006 to 14-Feb-2012, with the end date being determined by data availability. Specifically, the Greek sovereign CDS spread exceeds 10,000bp on 15-Feb-2012, shortly before Greek sovereign CDS contracts started trading points upfront due to the expectation of an imminent credit event. For each country, we include two country-specific credit spreads, one for the sovereign and one for the financial sector. We also include an array of common factors building on those identified by Longstaff et al. (2011).

3.1 Sovereign Credit Risk

We measure the change in sovereign credit risk using the first difference of the five-year sovereign CDS spread, expressed in basis points. Following the CDS market conventions outlined by Bai and Wei (2017), we work with US dollar denominated CDS in all cases except for the US, where we employ Euro denominated CDS. In addition, we use CDS contracts with a complete restructuring clause for every sovereign except Australia, where we use CDS contracts with a modified restructuring clause as the data is more complete.

3.2 Financial Sector Credit Risk

We track changes in financial sector credit risk in the ith country using the first difference of the synthetic sector-wide CDS spreads constructed by Greenwood-Nimmo et al. (2017). Taking inspiration from the approach of ADS, Greenwood-Nimmo et al. (2017) construct a synthetic credit spread for the ith country as an equally-weighted average of the CDS spreads for firms

which satisfy a variety of selection criteria. Among these criteria, firms must: (i) have USD denominated five-year CDS spread data in the Markit database which conforms to the corporate CDS market conventions documented by Bai and Wei (2017); (ii) be classified by Markit as *financials*; (iii) be classified as either banking or insurance firms in Bureau van Dijk's *Osiris* database; (iv) be identified by Markit as operating in the ith country; and (v) hold assets of USD10bn or more. The large majority of the firms included in the indices are publicly traded although there are two notable exceptions: (i) in Austria, Raiffeisen Zentralbank is included in the sample to ensure that the index is not based on data for a single firm; and (ii) in China, data for four large state-sponsored banks is used as there is not enough CDS data for privately held Chinese banks to construct a meaningful index. Similarly, rather than simply dropping failed banks from the sample, Greenwood-Nimmo et al. (2017) include CDS data for several institutions which became state-owned as a result of the crisis, such as the Irish Bank Resolution Corporation.[8]

3.3 Observed Common Factors

As noted by Longstaff et al. (2011), a large number of global variables may exert a common influence on credit spreads. The authors propose a parsimonious approach to the selection of factors, focusing on market-determined variables on the grounds that they should impound a wide array of information relevant to investors. We follow this precedent and include the majority of the observable explanatory variables considered by Longstaff et al. that are reported at daily frequency as well as several variables that they do not consider.[9] Following Longstaff et al. (2011), we start with a selection of US macroeconomic and financial indicators, which proxy for global economic and financial conditions:[10]

(i) *US stock market performance.* To capture the key risk factors affecting US equity markets, we include the Rm-Rf, SMB and HML factors developed by Fama and French (1993). Rm-Rf, accounts for a market factor, while SMB and HML account for risk factors related to firm size and book-to-market equity, respectively. The Fama-French factors are freely

[8]For a detailed discussion of the construction and properties of the financial sector CDS spreads, the reader is referred to Greenwood-Nimmo et al. (2017) and, in particular, to their Data Supplement.

[9]Our use of daily data necessitates the exclusion of several of the explanatory variables used by Longstaff et al. that are sampled at lower frequency, including bond and equity flows, for example.

[10]The use of US data to approximate global factors is supported by the wealth of evidence that US macrofinancial conditions exert a powerful and widespread influence on global economic and financial performance (e.g. Dees, di Mauro, Pesaran and Smith, 2007; Chudik and Fratzscher, 2011; Helbling, Huidrom, Kose and Otrok, 2011; Longstaff et al., 2011; Pesaran and Chudik, 2013).

available from Ken French via http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

(ii) *US Treasury market conditions.* We include the change in the five-year constant maturity Treasury (CMT) yield to capture expectations regarding macroeconomic conditions in the US and, by extension, in the world economy. In addition, given that investors regard US Treasury debt as a safe haven asset, Longstaff et al. note that the CMT yield may incorporate a flight-to-liquidity component. The CMT yield is published by the Federal Reserve in its H.15 Statistical Release.

Next, in light of the evidence that variations in funding liquidity and in the risk appetite of investors played an important role in the propagation of the global financial crisis (GFC) (e.g. Brunnermeier and Pedersen, 2009; Chudik and Fratzscher, 2011; Greenwood-Nimmo et al., 2016a; Pelizzon, Subrahmanyam, Tomio and Uno, 2016), we include the following factors:

(iii) *The TED spread.* The TED spread is the difference between the 3-month USD LIBOR and the 3-month US Treasury bill yield. Variations in the TED spread reflect changes in counterparty risk and liquidity in the US interbank market. Consequently, it is widely used as an indicator of funding liquidity. The TED spread is available from the Federal Reserve Economic Data Service (FRED) via https://fred.stlouisfed.org/.

(iv) *The Euribor-DeTBill spread.* To capture European-specific variations in funding liquidity, we include the spread between the 3-month Euribor and the 3-month German Treasury bill yield. The Euribor and the German yield data are available from Datastream.

(v) *The variance risk premium (VRP).* Bollerslev, Tauchen and Zhou (2009) define the VRP as the difference between the one-month-ahead implied variance and a forecast of the realised variance over the same period. Under this definition, the VRP is typically positive, with higher values indicating a reduced risk appetite. We forecast the realised variance using the same augmented version of Corsi's (2009) heterogeneous autoregressive model used by Bekaert and Hoerova (2014). We compute the VRP as $VRP_t = VIX_t^2 - E\left[RV_{t+1}^{(22)}\right]$, where VIX_t^2 denotes the de-annualised squared VIX and $RV_t^{(22)}$ denotes the realised variance for the S&P 500 measured over the next 22 trading days as the sum of squared five-minute intraday returns. The VIX data is available from FRED, while we obtain the daily realised variance from the Oxford Man Institute's Realized Library (Heber, Lunde, Shephard and

Sheppard, 2009, ver. 0.2).[11]

(vi) *The Treasury term premium.* The term premium measures the excess yield required to induce investors to hold a long-term bond as opposed to a sequence of shorter-term bonds. Consequently, it conveys valuable information on investors' time preferences as well as their expectations regarding the macroeconomic outlook. We include the 5-year Treasury term premium derived from the five-factor no-arbitrage term structure model of Adrian, Crump and Moench (2013) which is freely available from the Federal Reserve Bank of New York via `https://www.newyorkfed.org/research/data_indicators/term_premia.html`.

(vii) *US investment grade and high yield spreads.* To capture changes in the required rate of return on investment grade (IG) and high yield (HY) corporate bonds, we include both the IG and HY spreads. We define the IG spread as the spread between the Bank of America Merrill Lynch US corporate BBB and AAA effective yields and the HY spread as the difference between the Bank of America Merrill Lynch US corporate BB and BBB effective yields. The corporate bond yield data is available from FRED.

Longstaff et al. (2011) use cross-sectional averages of the credit spreads to proxy for regional and global factors in a manner reminiscent of the way that unobserved factors are approximated in the common correlated effects framework of Pesaran (2006). However, given our focus on observable factors, we elect to include an array of tradeable credit spread indices instead:

(viii) *ITRAXX indices to capture pan-European credit risk factors.* To account for European credit risk factors not captured elsewhere in our factor structure, we include five 5-year ITRAXX indices to isolate European credit risk factors. Specifically, we include the ITRAXX Europe index, the ITRAXX High Volatility index, ITRAXX Crossover index, and the ITRAXX Senior and Subordinated Financials indices. ITRAXX data is available via Datastream.[12]

Lastly, we control for currency fluctuations relative to the US dollar:

(ix) *Bilateral spot exchange rate returns.* To capture exchange rate fluctuations, we include the daily log-return on the bilateral spot exchange rate for each currency in our sample in

[11] Longstaff et al. (2011) also include the equity premium approximated at monthly frequency by the price-earnings ratio for the S&P 100 index. We are obliged to exclude the equity premium because earnings per share is unavailable at daily frequency. However, the variance risk premium should be highly correlated with the equity premium and should impound much of its informational content.

[12] We also experimented with the inclusion of North American and emerging markets CDX indices among our factors but we found that they did not add substantially to the information content of our other observed factors.

units of foreign currency per USD. The exchange rate data is obtained from Datastream. By including every exchange rate in each equation of the factor VAR model, we are able to control for portfolio adjustments affecting multiple currencies simultaneously.

3.4 Properties of the Dataset

Table 1 provides elementary summary statistics for the dataset. Preliminary analysis of the autocorrelation structure in the data indicates that each series is stationary with relatively limited serial correlation (results are available on request). The countries in our sample form two natural groups, one composed of the GIIPS (Greece, Ireland, Italy, Portugal and Spain) and Russia, which display high and volatile credit spreads, and the other composed of the remaining countries in our sample, which display considerably lower and less volatile credit spreads. For any given country, the sovereign CDS spread is typically lower and less volatile than the financial sector CDS spread. In general, sovereign credit spreads should act as a lower bound on financial sector credit spreads, not least because of the implicit guarantee that the sovereign extends to the financial sector. However, due to their deep sovereign crises, this is not the case in Greece, Italy, Portugal or Spain.

— Insert Table 1 here —

The credit spreads for all countries display pronounced excess kurtosis. This is also a feature of our common factors and is a natural reflection of the severity of the shocks affecting global financial markets over our sample period. The evidence of heavy tails in the data suggests that estimation by quantile regression may be preferable to the use of simple conditional mean estimators not just because it can illuminate tail relationships but also because it is more robust in the presence of extreme observations.

4 Estimation Results

Before we proceed, we must first determine appropriate values for both the QFVAR lag order and the forecast horizon used in the construction of our connectedness measures. We follow the precedent of Koenker and Xiao (2006) and set the lag order at every conditional quantile equal to the optimal lag order which is selected at the conditional mean by minimisation of the Schwarz Information Criterion. This results in a first-order model which we verify is dynamically stable — the largest eigenvalue of the companion matrix is just 0.39. Unfortunately, there is

no similar rule to select an optimal forecast horizon for connectedness analysis. For this reason, we compare the properties of the adjacency matrix evaluated at the conditional mean for three different forecast horizons, $h \in \{3, 5, 10\}$ trading days. We select relatively short horizons in light of the fact that the FEVDs obtained from a first order VAR model estimated on data with a low degree of persistence are likely to rapidly converge to their long-run values. This observation is borne out by our finding that the network statistics are largely invariant to the choice of horizon within this range, with the spillover indices obtained under $h = 3$, $h = 5$ and $h = 10$ being identical to the first decimal place (a detailed elementwise comparison of the adjacency matrices is available on request). A similar degree of invariance with respect to the forecast horizon has been documented by Greenwood-Nimmo et al. (2016a) in the context of a stationary first order VAR model applied to the analysis of risk spillovers among the G10 currencies. We therefore adopt a forecast horizon of five trading days without loss of generality, although we shall return to the issue of horizon selection in the rolling sample context in Section 4.4 below.

4.1 Performance of the Observed Factors

In practice, a finite number of observed common factors is unlikely to completely eliminate the cross-section correlation among the regression residuals but it should render it sufficiently weak that the off-diagonal elements of the residual covariance matrix can be set to zero without loss of generality. Consequently, the adequacy of our factors is ultimately an empirical question. To evaluate the performance of our factors, we estimate two models at the conditional mean — a simple VAR(1) model with no factors and our factor VAR(1) model — and compare the residual cross-correlations in each case. Figure 1(a) reveals considerable correlation among the residuals of the simple VAR(1) model, with more than half of the pairwise correlation coefficients exceeding 0.2 in absolute value and more than 20% exceeding 0.4. Figure 1(b) shows that the factors remove a great deal of this correlation. In the factor VAR(1) model, more than three-quarters of the pairwise correlations are weaker than 0.2 in absolute value and 91% are weaker than 0.4. This strongly supports the validity of our factors and suggests that it is a reasonable simplification to treat the error terms of the factor VAR(1) model as cross-sectionally uncorrelated.

— Insert Figure 1 here —

Next, we examine how the introduction of common factors affects the network statistics

obtained at the conditional mean. To this end, Figure 2 plots the distribution of bilateral spillover effects — the *spillover density* to adopt the terminology of Greenwood-Nimmo et al. (2016b) — for the models with and without factors. In the factor VAR case, each row of the adjacency matrix defined in (20) sums to unity by virtue of the diagonality of the covariance matrix, $\mathbf{\Omega}$. To obtain comparable values for the simple VAR model where the covariance matrix is non-diagonal, we apply the row-sum normalisation suggested by Diebold and Yilmaz (2014). For convenience, we multiply each element of the adjacency matrix by 100 so that the estimated spillover effects can be interpreted as percentages rather than proportions. In principle, therefore, the spillover density has support $[0, 100]$ although, in practice, the limiting cases of 0 and 100 will only arise from restricted VAR models where the parameter and covariance matrices are sparse. Greenwood-Nimmo et al. (2016b) have shown that the spillover density of a network constructed by the Diebold–Yilmaz method resembles a power law.[13] This is a natural finding given that the elements of the adjacency matrix are constructed from variance decompositions, which are defined as ratios of quadratic forms. Consequently, for ease of interpretation, we follow Acemoglu et al. (2012) and report empirical counter cumulative distribution functions (CCDFs) on a logarithmic scale in Figure 2.

— Insert Figure 2 here —

The CCDF for the model without factors lies considerably to the right of the CCDF for the model with factors, indicating that the omission of common factors leads to stronger estimated spillover effects. The difference between the two CCDFs is profound, with the spillover index defined in (22) obtained from the model without factors taking a value of 72.50% compared to just 11.49% for the model with factors. The mechanism driving this result is straightforward. If a common component which contributes to the h-step-ahead FEV of the system is omitted in estimation, then the share of the FEV that should rightly be attributed to that common factor must be attributed to one or more of the endogenous variables included in the model. The omission of relevant common factors will therefore upwardly bias the estimated bilateral spillover effects. It is important to distinguish between common and idiosyncratic sources of variation if one wishes to measure true bilateral spillovers which are free of common components. For this reason, we henceforth focus exclusively on factor VAR models.

[13] Power law behaviour is pervasive in economic and financial networks — see Acemoglu, Carvalho, Ozdaglar and Tahbaz-Salehi (2012) and Gabaix (2016), for example.

4.2 Conditional Mean vs. Conditional Median

To establish a point of reference for the conditional quantile analysis that follows, we first compare the network structure evaluated at the conditional mean by OLS and at the conditional median by LAD over the full sample. If the results from network models evaluated at the conditional mean are to be generalised across the conditional distribution — as is currently the norm, at least implicitly — then the results obtained under OLS and LAD should be similar. However, unlike OLS, LAD belongs to the class of robust M-estimators and is therefore less susceptible to the influence of outliers. Given the extreme credit spread movements recorded in several countries during our sample — notably among the GIIPS and Russia — it is likely that the OLS and LAD estimates will differ. The comparison of OLS and LAD also provides an initial glimpse of the value of estimation by quantile regression. The OLS estimator is equivalent to an equally weighted average of the quantile regression estimators for $\tau \in (0, 1)$, while the LAD estimator is simply the quantile regression estimator at $\tau = 0.5$. It follows, therefore, that if we observe differences between the network under OLS and LAD, then there must also be differences between the network under LAD and at other conditional quantiles — that is, the network will display quantile variation. The more pronounced the differences between OLS and LAD, the greater the quantile variation is likely to be.

Figures 3(a) and 3(b) plot the spillover density in both the OLS and LAD cases. Several features of the spillover densities under OLS and LAD are noteworthy. First, the right tails of both densities are similar, which indicates that the strongest spillovers in the system are of comparable magnitude at the conditional mean and median. The three strongest bilateral spillovers under OLS are from the Spanish sovereign to the Spanish financial sector (8.55%), the Chinese sovereign to the Russian sovereign (8.12%) and the Russian sovereign to the Russian financial sector (5.03%). By contrast, under LAD, the three strongest pairwise spillovers all arise from the Spanish sovereign and affect the Italian sovereign (5.81%), the Spanish financial sector (5.07%) and the German sovereign (4.51%). Although these values are small compared to the spillover effects reported in much of the existing literature on Diebold–Yilmaz networks, recall that we employ a factor structure to isolate uncorrelated idiosyncratic shocks. In this context, a spillover which accounts for 5% of FEV at the five-days-ahead horizon in a model with thirty-six endogenous variables represents a strong bilateral linkage. The evidence of strong spillovers originating from Spain and affecting the local financial sector and other European sovereigns reflects the integration of financial markets within the EU.

— Insert Figure 3 here —

The similarity in the right tail of the spillover density does not extend to other parts of the density, however. The CCDF is less curved and displays considerably greater left tail mass under LAD than under OLS, implying a higher proportion of weak spillovers at the median than at the mean. The weakest bilateral spillover at the conditional mean takes a value of $3.65 \times 10^{-4}\%$. At the conditional median, there are fifty-seven bilateral spillovers that are weaker than this and the weakest spillover is two orders of magnitude smaller at just $4.37 \times 10^{-6}\%$. To appreciate the difference in the shape of the two spillover densities more clearly, consider an arbitrary threshold — of the 1,260 bilateral spillovers that we study, 551 are weaker than 0.1% at the conditional mean but this value rises to 947 at the conditional median.

The granular differences documented above accumulate substantially under aggregation. Table 2 reports the to (weighted out-degree), from (weighted in-degree) and net statistics for each node in the system. In every case, the OLS estimate is larger than the LAD estimate — often substantially so. If one computes the ratio of the OLS estimate to the LAD estimate for each of the reported weighted in- and out-degrees, the minimum, mean and maximum values are 1.23, 6.14 and 124.15, respectively. In addition to these scale differences in the estimated spillover effects, the ranking of the most influential nodes in the system — measured by the weighted out-degree — also differs across estimators. Under OLS, the three most influential nodes are the Spanish (35.83%), Italian (25.85%) and Austrian (24.00%) sovereigns. Meanwhile, at the conditional median, the most influential nodes are the Spanish sovereign (29.21%), the French financial sector (14.15%) and the Dutch sovereign (9.59%).

— Insert Table 2 here —

To provide an impression of how the differences surveyed above affect the network as a whole, Figure 4 provides a visual comparison of the network topology under OLS and LAD. Sovereigns are represented by white nodes and financial sectors by shaded nodes, while the strength of bilateral spillovers is indicated by the relative thickness of the edges. The layout of both graphs is identical and is determined using the force-directed algorithm of Fruchterman and Reingold (1991) applied at the conditional mean. The networks display some similar features, notably the centrality of the Spanish sovereign and the disposition of many of the strongest bilateral spillovers. However, the excess connectedness of the network evaluated by OLS relative to LAD is easily seen and is clearly reflected in the spillover index, which takes a value of 11.49% at the

conditional mean compared to 4.34% at the conditional median.

— Insert Figure 4 here —

The results obtained under OLS suggest that, with a few notable exceptions, idiosyncratic credit risk shocks do not propagate strongly. The results obtained under LAD indicate that credit risk spillovers are even weaker than under the OLS case. However, the large difference between the spillover intensity under OLS and LAD suggests that stronger spillovers prevail at some non-central location or locations in the conditional distribution. It is to this issue that we now turn.

4.3 Quantile Variation: Network Topology and the Shock Size

Our interpretation of the quantile regression estimates below will be predicated on the distinction between large adverse shocks and large beneficial shocks. To see this, note that in the right tail of the conditional distribution, the observed changes in the vector of credit spreads are large conditional on the values taken by the explanatory variables — that is, credit spreads are increasing sharply due to the effect of large adverse shocks. By contrast, in the left tail, credit spreads are falling sharply conditional on the explanatory variables due to the impact of large beneficial shocks. In light of the evidence that financial market comovements increase under adverse conditions (e.g. Ang and Bekaert, 2002), it is natural to think that strong spillovers should be observed in the right tail of the conditional distribution, where the largest adverse shocks affect the system. In practice, the pattern of quantile variation in the spillover index shown in Figure 5 reveals that strong spillovers occur in both the left and right tails of the conditional distribution, indicating that spillover intensity increases with shock size for both adverse (right tail) and beneficial (left tail) shocks. This is consistent with the literature on good and bad contagion, which emphasises the transmission of unexpected events in both the left and right tail (e.g. Londono, 2016).

— Insert Figure 5 here —

Over the central 70% of the conditional distribution, the spillover index never exceeds 9% and the influence of idiosyncratic credit risk shocks is largely localised. By comparison, when large idiosyncratic shocks affect the system, bilateral spillovers play a profound role in shaping the evolution of sovereign and financial sector credit risk. At the 1st, 5th, 10th, 90th, 95th and

99th percentiles, the spillover index takes values of 76.57%, 21.07%, 10.04%, 12.58%, 22.87% and 83.32%, respectively. It is interesting to note that the pattern of quantile-variation in the spillover index is roughly symmetric. Given that the residual covariance matrix is quantile-invariant, this effect arises because of similarities in the dynamic parameters of the QFVAR model at quantiles $\tau = \alpha$ and $\tau = 1 - \alpha$. In practice, this near-symmetry arises by chance over the full sample and is not a general feature of our results — we will shortly demonstrate that it breaks down in rolling sample analysis.

The increased influence of bilateral credit risk spillovers in both tails can be seen clearly in Figures 3(c)–(f). Note that how the spillover density moves rightward and becomes more peaked as $\tau \to 1$ and $\tau \to 0$. It is also readily apparent in Figure 6, which shows network plots for $\tau = \{0.05, 0.95\}$ drawn on the same scale and with the same layout as Figure 4. The increase in connectedness in the tails is marked, implying that large idiosyncratic credit risk shocks propagate considerably more forcefully than weaker shocks.

— Insert Figure 6 here —

These large tail effects are not simply due to a lack of effective observations in the tails of the conditional distribution. For any given τ, quantile regression makes use of every data point with non-zero weight and our sample of 1,596 trading days is relatively sizable compared to the dimensionality of the QFVAR system. Rather, our results are consistent with the hypothesis that the informational content of large shocks is greater than that of small shocks, a point which is well-established in the volatility literature (e.g. Dendramis, Kapetanios and Tzavalis, 2015). When a large idiosyncratic shock affects a given sovereign or financial sector, investors gain significant news which may lead to a reappraisal of the level of risk associated with other nodes in the system. Consequently, the quantile variation documented in Figure 5 can be interpreted like a regime-switching process where the transition between regimes of strongly beneficial news at one extreme and strongly adverse news at the other occurs smoothly as the shock size varies.

An obvious question to ask as this stage is whether the ranking of nodes by centrality is preserved across quantiles. That is, does the group of nodes that display the strongest outward spillovers vary with the shock size? To this end, in Figure 7, we plot the weighted out-degree rank of each node in the system, with the nodes ranked in decreasing order of influence. The weighted out-degree captures the total strength of all outward spillovers originating from a given node. As such, it represents a natural measure of the influence of a node.

— Insert Figure 7 here —

The figure is organised with the GIIPS on the top row, the other European countries on the next two rows and the non-European countries on the final row. Several features of Figure 7 are noteworthy. First, the weighted out-degree rank displays marked quantile variation in the majority of cases. Consider the GIIPS to begin with. With the exception of Greece, the GIIPS sovereigns are typically toward the top of the weighted out-degree ranking across all quantiles. This is particularly apparent for Spain and Italy, where the combination of a high and rising debt servicing cost with a substantial stock of outstanding debt generated a significant risk of contagion — the scale of the losses generated by a Spanish or Italian default would have posed a grave threat to global financial stability. Greece behaves quite differently than the other GIIPS sovereigns, demonstrating a relatively high weighted out-degree rank in the middle of the conditional distribution but a much lower rank in both tails. This suggests that the propagation of large idiosyncratic Greek sovereign risk shocks was relatively mild. At first sight, this appears to stand at odds with historical experience of the European debt crisis, where Greek contagion was widely discussed.[14] However, several studies have since shown no Greek contagion at this time (e.g. Mink and de Haan, 2013; Pragidis, Aiellia, Chionis and Schizas, 2015). A more nuanced interpretation of our finding is that although direct bilateral spillovers of Greek sovereign risk may be weak, the crisis in Greece may have affected global markets indirectly via its impact on global factors such as investor risk appetite.

A general feature of many of the European sovereigns — GIIPS and non-GIIPS alike — is that the weighted out-degree rank is typically higher in the tails of the conditional distribution than in the central region. The same is true of the of the US. The increased influence of these sovereigns in the tails is a reflection of their centrality in the global financial crisis and the European debt crisis. By contrast, the other non-European sovereigns stand apart. In Australia, the weighted out-degree rank is low and relatively stable across conditional quantiles, reflecting the country's muted experience of both the global financial crisis and the European debt crisis. In China, Japan and Russia, the weighted out-degree rank is considerably lower in the tails than at the median, indicating that these countries were predominantly the recipients of large external shocks during both crises as opposed to the source of large influential shocks.

Another result which emerges from Figure 7 is that, for a given country, the weighted out-degree rank of the sovereign typically exceeds that of the financial sector. This effect can be

[14] A good example may be found in Nouriel Roubini's column in the *Wall Street Journal* on May 6, 2010 (https://www.forbes.com/2010/05/05/greece-bailout-imf-opinions-columnists-nouriel-roubini.html).

understood in relation to the credit risk transfers implicit in financial sector bailouts. Acharya et al. (2014) demonstrate that if a sovereign elects to bail out the domestic financial sector in order to maintain the stability of the financial system, a significant amount of private sector risk is transferred onto the sovereign. Consequently, dysfunction in the ith financial sector is likely to be felt disproportionately by the ith sovereign, not least because the sovereign guarantee acts to partially insulate both domestic and foreign investors from shocks originating in the ith financial sector. By contrast, in an environment where financial institutions hold internationally diversified portfolios of debt instruments, dysfunction in a given sovereign debt market may rapidly propagate to the financial sector both domestically and internationally. This effect is likely to be particularly pronounced among the Eurozone member states, where financial market integration, common monetary policy and a shared currency create an environment where shocks may propagate forcefully and where many of the instruments of stabilisation policy cannot be manipulated on a country-by-country basis.

The relationship between the credit risk of the ith sovereign and the ith financial sector is central to the analysis of Acharya et al. (2014). Specifically, the authors document an adverse feedback loop between sovereign credit risk and financial sector credit risk which emerges after a financial sector bailout. In their analysis, the transfer of credit risk from the financial sector to the sovereign associated with a financial sector bailout coupled with the fiscal burden of the bailout leads to an increase in the sovereign's credit risk. This undermines the value of the sovereign's implicit guarantee of the financial sector moving forward and reduces the value of the sovereign debt portfolios held in the financial sector. This causes financial sector credit risk to rise which, in turn, further exacerbates sovereign risk because it raises the likelihood of further sovereign intervention in the financial sector and so the cycle continues.

In our model, the feedback between the ith sovereign and the ith financial sector is captured by the sum of the bilateral spillovers between the two nodes, $\mathcal{T}^{(5)}_{s_i \leftrightarrow f_i,(\tau)} = T^{(5)}_{s_i \leftarrow f_i,(\tau)} + T^{(5)}_{f_i \leftarrow s_i,(\tau)}$. Figure 8 reports the quantile variation in $\mathcal{T}^{(5)}_{s_i \leftrightarrow f_i,(\tau)}$ for all 18 countries. For most countries, the feedback effect is negligible throughout the centre of the conditional distribution but intensifies markedly in the tails. The mechanism described by Acharya et al. (2014) focuses on the feedback associated with adverse shocks, in the right tail of the conditional distribution. Our results indicate that the same feedback loop acts upon the arrival of large beneficial shocks. A sovereign bailout is a natural example of such a beneficial shock. Suppose that the ith sovereign receives a bailout which reduces its credit risk. This increases the ability of the ith sovereign to stabilise

the ith financial sector while simultaneously lowering the default risk associated with sovereign bonds held by the ith financial sector. This leads to a reduction in financial sector credit risk which lowers probability that the sovereign will be required to intervene in the financial sector, further lowering sovereign risk and so on. This suggests that the same feedback loop that promotes instability in the analysis of Acharya et al. (2014) can act to restore stability if policymakers are able to generate large beneficial shocks.

— Insert Figure 8 here —

4.4 Time-Varying Tail-Dependence

To this point, our analysis has focused exclusively on full-sample statistics. We have demonstrated substantial quantile-variation in the topology of the credit risk network and have shown that bilateral spillovers of idiosyncratic credit risk are an order of magnitude stronger in the tails than they are at the conditional median. This implies that network models estimated at the conditional mean are unlikely to adequately capture the extent of dependence observed when large shocks occur. As noted by Betz et al. (2016), it is tail-dependence that should be of the greatest interest for surveillance and regulatory purposes. Unlike the existing literature on Diebold–Yilmaz networks, by estimating our QFVAR model on a rolling sample basis, we can directly study time-variation not only in average connectedness but also in the extent of left- and right-tail-dependence.

Before we proceed, we must first choose an appropriate window length for our rolling samples. Existing studies in the Diebold–Yilmaz network literature which work with daily data have used a variety of window lengths, including 100 days (e.g. Diebold and Yilmaz, 2014), 150 days (e.g. Demirer et al., 2017), 200 days (e.g. Baruník et al., 2016) and 250 days (e.g. Greenwood-Nimmo et al., 2016a, 2017). In the absence of a firm precedent, we follow Greenwood-Nimmo et al. (2016a) and evaluate the sensitivity of our results to a set of three candidate window lengths, $w \in \{200, 250, 300\}$ trading days — we do not consider shorter windows to ensure that we do not encounter small-sample issues in estimation. In addition, we take this opportunity to further explore the robustness of our estimation results to the choice of forecast horizon, $h \in \{3, 5, 10\}$.

Figure 9 reports the spillover index under the nine possible combinations of window length and forecast horizon at the conditional mean and median as well as the fifth and ninety-fifth conditional quantiles. First, consider panels (a) and (b). At the conditional mean and median, the choice of forecast horizon has little discernible effect on the spillover index, a result which

reinforces our findings over the full sample. The choice of window length affects the level of the spillover index, with shorter windows yielding somewhat higher values. Nonetheless, it is the dynamics of the spillover index which are of primary concern and they are largely unaffected by the choice of window length. The correlation among the spillover indices obtained under the nine different combinations of w and h is close to one in all cases. The same basic features are also apparent at the fifth and ninety-fifth percentiles, although with greater noise. Critically, however, the correlation across different combinations of window length and forecast horizon remains substantial in the tails. We therefore conclude that the choice of window length does not exert an undue influence on our results and we proceed with $w = 250$ and $h = 5$ trading days without loss of generality. This leaves us with 1,347 rolling samples.

— Insert Figure 9 here —

With the forecast horizon and window length set, Figure 10(a) re-plots the spillover index evaluated at the conditional mean and median — both of which are different measures of the average connectedness of the system — on the same axes. To assist the reader, the dates of several important events marked by vertical dashed lines. Figure 10(b) plots the spillover index evaluated at the fifth conditional quantile as a measure of left-tail-dependence. This captures the propensity for beneficial shocks that reduce credit risk to propagate through the system. All else equal, stronger left-tail-dependence is stabilising. Figure 10(c) plots the the spillover index at the ninety-fifth conditional quantile as a measure of right-tail-dependence, the propensity for destabilising adverse shocks to spread. Lastly, Figure 10(d) plots the linear combination $RTD = S^{(5)}_{0.95} - S^{(5)}_{0.05}$. The time-variation in RTD clearly demonstrates that the near-symmetry of Figure 5 in the full-sample setting is not preserved over rolling samples. Of greater interest, however, is the natural interpretation of RTD as a measure of *relative tail-dependence*, with positive (negative) values indicating stronger (weaker) dependence in the right tail than in the left tail. We interpret increases (decreases) in RTD as evidence of rising (falling) financial fragility. The correlations between the the four different spillover indices and our RTD measure are reported in Table 3.

— Insert Figure 10 & Table 3 here —

The spillover indices evaluated at the conditional mean and median display broadly similar behaviour. Both increase abruptly as a result of major adverse shocks such as the freezing of redemptions in selected investment funds by BNP Paribas in August 2007 and the bankruptcy

of Lehman Brothers in September 2008. Similarly, beneficial shocks such as the announcement of the Troubled Asset Relief Program (TARP) and the GIIPS sovereign bailouts cause both indices to fall gradually. However, the spillover index evaluated at the conditional mean falls much more markedly than its counterpart evaluated at the conditional median. As a result, Table 3 reveals that the two are only moderately correlated (0.57).

A natural question to ask at this stage is whether the spillover indices evaluated at the conditional mean and median share common dynamics with the spillover indices evaluated in the tails of the conditional distribution. The correlations in Table 3 provide a striking answer. Both are positively correlated with our measure of right-tail-dependence (0.79 and 0.55, respectively) but are essentially uncorrelated with our measure of left-tail-dependence (-0.11 and 0.07, respectively). Furthermore, our measures of left- and right-tail-dependence are mutually negatively correlated (-0.30). It is the last result which is most interesting as it suggests that changes in right-tail-dependence coincide with oppositely-signed changes in left-tail-dependence. For example, a period of growing fragility associated with an increased propensity for adverse shocks to propagate is also likely to be a period where the propagation of beneficial shocks becomes weaker and vice-versa. This result may arise from the aggregate behaviour of market participants if the information revealed by a major event in either tail causes a non-trivial proportion of market participants to focus disproportionately on further events occurring in that tail while paying less attention to events in the other tail.

The negative association between left- and right-tail-dependence is easily seen in Figures 10(b)-(d). Consider the revelation of major losses at the Bear Stearns High-Grade Structured Credit Fund and the Bear Stearns High-Grade Structured Credit Enhanced Leveraged Fund in July 2007 as an example of a large adverse shock which led market participants to re-evaluate the level of risk associated with mortgage-backed securities. At this time, we observe a prolonged downward drift in left-tail-dependence, indicating a reduced propensity for spillovers of good news. At the same time, there is a sharp and sustained increase in right-tail-dependence, indicating a substantial increase in the sensitivity of market participants to adverse shocks. This combination naturally generates a marked increase in RTD.

Now, consider the announcement of TARP as an example of a major beneficial shock. TARP represented a major government intervention into the financial markets, providing funds for the purchase of toxic assets and equity from troubled financial institutions. The introduction of TARP is associated with a mild reduction in right-tail-dependence but with a strong and sus-

tained increase in left-tail-dependence. Recall that we have controlled for a raft of factors including liquidity conditions in the US so this result is not simply a manifestation of the improvement in the economic outlook and in funding liquidity in particular brought about by TARP. Rather, it suggests that major policy interventions can generate a pronounced intensification of stabilising beneficial spillovers.

As a final exercise, to rule out the possibility that the behaviour of the relative tail-dependence documented above is simply an artefact of our choice to work with the fifth and ninety-fifth conditional quantiles, we plot two alternative measures of relative tail-dependence in Figure 11. Specifically, alongside our benchmark 5% RTD, we plot the 10% and 1% RTDs defined as $RTD_{10\%} = S^{(5)}_{0.90} - S^{(5)}_{0.10}$ and $RTD_{1\%} = S^{(5)}_{0.99} - S^{(5)}_{0.01}$, respectively. The dynamics of the three RTD measures are remarkably similar, indicating that our results are robust to the precise definition of the left- and right-tail-dependence measures.

— Insert Figure 11 here —

5 Concluding Remarks

In this paper, we develop a new technique for the econometric analysis of financial networks where the topology of the network is allowed to vary with the size of the shocks that affect the system. Our approach is based on a novel framework for the estimation of vector autoregressions by quantile regression. We employ a factor structure to isolate the idiosyncratic component of the error process from the systematic component. Not only does this align our model with the long literature on systematic and idiosyncratic risk but it also simplifies the estimation problem as it renders the system of equations cross-sectionally independent. As a result, we are able to estimate the model on an equation-by-equation basis using the standard quantile regression toolboxes built into many statistical software packages. Our approach is therefore considerably easier to implement than the existing frameworks for the estimation of quantile VAR models associated with Cecchetti and Li (2008) and Schüler (2014).

We apply our technique to study the transmission of credit risk among a panel of eighteen sovereigns and their respective financial sectors between January 2006 and February 2012. We document marked quantile variation in the topology of the network. We show that idiosyncratic credit risk shocks do not propagate strongly at the conditional mean or median but that powerful spillovers occur in both tails of the conditional distribution. In addition, by studying the

bidirectional feedback between each sovereign and its domestic financial sector, we find that the adverse feedback loop between sovereign credit risk and financial sector credit risk documented by Acharya et al. (2014) manifests as a marked intensification of feedback in the right tail, where large adverse shocks occur. Interestingly, however, we note a similar intensification in the left tail, which indicates that the same feedback loop can act in a stabilising manner in the presence of large beneficial shocks, such as sovereign bailouts. Finally, we use rolling sample analysis to document time-variation in the degree of tail-dependence. This reveals an interesting phenomenon — the level of left-tail-dependence is negatively correlated with the level of right-tail-dependence. Specifically, our results indicate that major adverse (beneficial) events are associated with a reduction (intensification) in stabilising left-tail spillovers coupled with a simultaneous intensification (reduction) in destabilising right-tail spillovers. Furthermore, although the dynamic evolution of spillover activity in the right tail is broadly similar to that observed at the conditional mean and median, this is not true of spillover activity in the left tail. Consequently, the evolution of relative tail-dependence is obscured when network models are estimated at the conditional mean by OLS, as is common in the literature. We therefore conclude that our framework for the analysis of quantile connectedness represents a valuable addition to the existing literature on empirical network modelling.

References

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "Systemic Risk and Stability in Financial Networks," *American Economic Review*, 2015, *105*, 564–608.

_ **, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, "The Network Origins of Aggregate Fluctuations," *Econometrica*, 2012, *80*, 1977–2016.

Acharya, Viral V., Itamar Drechsler, and Philipp Schnabl, "A Pyrrhic Victory? Bank Bailouts and Sovereign Credit Risk," *Journal of Finance*, 2014, *69*, 2689–2739.

Adrian, Tobias and Markus K. Brunnermeier, "CoVaR," *American Economic Review*, 2016, *106*, 1705–1741.

_ **, Richard K. Crump, and Emanuel Moench**, "Pricing the Term Structure with Linear Regressions," *Journal of Financial Economics*, 2013, *110*, 110–138.

Alter, Adrian and Andreas Beyer, "The Dynamics of Spillover Effects during the European Sovereign Debt Turmoil," *Journal of Banking and Finance*, 2014, *42*, 134–153.

Ang, Andrew and Francis A. Longstaff, "Systemic Sovereign Credit Risk: Lessons from the U.S. and Europe," *Journal of Monetary Economics*, 2013, *60*, 493–510.

_ **and Geert Bekaert**, "International Asset Allocation with Regime Shifts," *Review of Financial Studies*, 2002, *15*, 1137–1187.

Bai, Jennie and Shang-Jin Wei, "Property Rights Gaps and CDS Spreads: When is there a Strong Transfer Risk from the Sovereigns to the Corporates?," *Quarterly Journal of Finance*, 2017, *In press.*

Bai, Jushan, "Inferential Theory for Factor Models of Large Dimensions," *Econometrica*, 2003, *71*, 135–171.

_ , "Panel Data Models with Interactive Fixed Effects," *Econometrica*, 2009, *77*, 1229–1279.

Baruník, Jozef, Evžen Kočenda, and Lukáš Vácha, "Asymmetric Connectedness on the U.S. Stock Market: Bad and Good Volatility Spillovers," *Journal of Financial Markets*, 2016, *27*, 55–78.

Bekaert, Geert and Marie Hoerova, "The VIX, the Variance Premium and Stock Market Volatility," *Journal of Econometrics*, 2014, *183*, 181–192.

Betz, Frank, Nikolaus Hautsch, Tuomas A. Peltonen, and Melanie Schienle, "Systemic Risk Spillovers in the European Banking and Sovereign Network," *Journal of Financial Stability*, 2016, *25*, 206–224.

Billio, Monica, Mila Getmansky, Andrew Lo, and Loriana Pelizzon, "Econometric Measures of Connectedness and Systemic Risk in the Finance and Insurance Sectors," *Journal of Financial Economics*, 2012, *104*, 535–559.

Blanchard, Olivier J. and Danny Quah, "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review*, 1989, *79*, 655–673.

_ **and Mark W. Watson**, "Are all Business Cycles Alike?," in R.J. Gordon, ed., *The American Business Cycle*, Chicago: University of Chicago Press, 1986, pp. 123–160.

Blanco, Roberto, Simon Brennan, and Ian W. Marsh, "An Empirical Analysis of the Dynamic Relationship between Investment-Grade Bonds and Credit Default Swaps," *Journal of Finance*, 2005, *60*, 2255–2281.

Bollerslev, Tim, George Tauchen, and Hao Zhou, "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 2009, *22*, 4463–4492.

Bostanci, Görkem and Kamil Yilmaz, "How Connected is the Global Sovereign Credit Risk Network?," Working Paper 1515, Koç University, Istanbul August 2015.

Brunnermeier, Markus K. and Lasse H. Pedersen, "Market Liquidity and Funding Liquidity," *Review of Financial Studies*, 2009, *22*, 2201–2238.

Caporin, Massimiliano, Loriana Pelizzon, Francesco Ravazzolo, and Roberto Rigobon, "Measuring Sovereign Contagion in Europe," Working Paper 18741, NBER, Cambridge (MA) January 2013.

Cecchetti, Stephen G. and Hong Li, "Measuring the Impact of Asset Price Booms Using Quantile Vector Autoregressions," Mimeo: Brandeis University February 2008.

Chavleishvili, Sulkhan and Simone Manganelli, "Quantile Impulse Response Functions," Mimeo: Goethe University Frankfurt January 2016.

Chudik, Alexander and Marcel Fratzscher, "Identifying the Global Transmission of the 2007–2009 Financial Crisis in a GVAR Model," *European Economic Review*, 2011, *55*, 325–339.

Corsi, Fulvio, "A Simple Approximate Long Memory Model of Realized Volatility," *Journal of Financial Econometrics*, 2009, *7*, 174–196.

Dees, Stephane, Filippo di Mauro, M. Hashem Pesaran, and L. Vanessa Smith, "Exploring The International Linkages Of The Euro Area: A Global VAR Analysis," *Journal of Applied Econometrics*, 2007, *22* (1), 1–38.

Demirer, Mert, Francis X. Diebold, Laura Liu, and Kamil Yilmaz, "Estimating Global Bank Network Connectedness," *Journal of Applied Econometrics*, 2017, *in press.*

Dendramis, Yiannis, George Kapetanios, and Elias Tzavalis, "Shifts in Volatility Driven by Large Stock Market Shocks," *Journal of Economic Dynamics and Control*, 2015, *55*, 130–147.

Diebold, Francis X. and Kamil Yilmaz, "Measuring Financial Asset Return and Volatility Spillovers, with Application to Global Equity Markets," *The Economic Journal*, 2009, *119*, 158–171.

_ and _ , "On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms," *Journal of Econometrics*, 2014, *182*, 119–134.

Fama, Eugene F. and Kenneth R. French, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 1993, *33*, 3–56.

Feldhütter, Peter and Mads Stenbo Nielsen, "Systematic and Idiosyncratic Default Risk in Synthetic Credit Markets," *Journal of Financial Econometrics*, 2012, *10*, 292–324.

Fender, Ingo, Bernd Hayo, and Matthias Neuenkirch, "Daily Pricing of Emerging Market Sovereign CDS Before and During the Global Financial Crisis," *Journal of Banking and Finance*, 2012, *36*, 2786–2794.

Fruchterman, Thomas M.J. and Edward M. Reingold, "Graph Drawing by Force-Directed Placement," *Software – Practice and Experience*, 1991, *21*, 1129–1164.

Gabaix, Xavier, "Power Laws in Economics: An Introduction," *Journal of Economic Perspectives*, 2016, *30*, 185–206.

Garratt, Anthony, Kevin Lee, M. Hashem Pesaran, and Yongcheol Shin, *Global and National Macroeconometric Modelling: A Long-Run Structural Approach*, Oxford: Oxford University Press, 2006.

Greenwood-Nimmo, Matthew J., Jingong Huang, and Viet H. Nguyen, "Financial Sector Bailouts, Sovereign Bailouts and the Transfer of Credit Risk," Unpublished Manuscript March 2017.

_ , Viet H. Nguyen, and Barry Rafferty, "Risk and Return Spillovers among the G10 Currencies," *Journal of Financial Markets*, 2016, *31*, 43–62.

_ , _ , and Yongcheol Shin, "What's Mine is Yours: Sovereign Risk Transmission during the European Debt Crisis," Mimeo, University of Melbourne May 2016.

Gyntelberg, Jacob, Peter Hördahl, Kristyna Ters, and Jörg Urban, "Intraday Dynamics of Euro Area Sovereign CDS and Bonds," Working Paper 423, Bank for International Settlements, Basel September 2013.

Heber, Gerd, Asger Lunde, Neil Shephard, and Kevin Sheppard, "Oxford-Man Institute's Realized Library," Dataset, Oxford-Man Institute, University of Oxford 2009.

Helbling, Thomas, Raju Huidrom, M. Ayhan Kose, and Christopher Otrok, "Do Credit Shocks Matter? A Global Perspective," *European Economic Review*, 2011, *55*, 340–353.

International Monetary Fund, "Old Risks, New Challenges," Global Financial Stability Report, April 2013, IMF, Washington DC April 2013.

Koenker, Roger, *Quantile Regression*, Cambridge: Cambridge University Press, 2005.

_ and Gilbert Bassett, "Regression Quantiles," *Econometrica*, 1978, *46*, 33–50.

_ and Kevin F. Hallock, "Quantile Regression," *Journal of Economic Perspectives*, 2001, *15*, 143–156.

_ and Zhijie Xiao, "Quantile Autoregression," *Journal of the American Statistical Association*, 2006, *101*, 980–90.

Linnemann, Ludger and Roland Winkler, "Estimating Nonlinear Effects of Fiscal Policy using Quantile Regression Methods," *Oxford Economic Papers*, 2016, *68*, 1120–1145.

Londono, Juan M., "Bad Bad Contagion," International Finance Discussion Papers 1178, Board of Governors of the Federal Reserve System, Washington D.C. September 2016.

Longstaff, Francis A., Jun Pan, Lasse H. Pedersen, and Kenneth J. Singleton, "How Sovereign Is Sovereign Credit Risk?," *American Economic Journal: Macroeconomics*, 2011, *3*, 75–103.

Mink, Mark and Jakob de Haan, "Contagion during the Greek Sovereign Debt Crisis," *Journal of International Money and Finance*, 2013, *34*, 102–113.

Pan, Jun and Kenneth J. Singleton, "Default and Recovery Implicit in the Term Structure of Sovereign CDS Spreads," *Journal of Finance*, 2008, *63*, 2345–2384.

Pelizzon, Loriana, Marti G. Subrahmanyam, Davide Tomio, and Jun Uno, "Sovereign Credit Risk, Liquidity, and European Central Bank Intervention: Deus Ex Machina?," *Journal of Financial Economics*, 2016, *122*, 86–115.

Pesaran, M. Hashem, "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure," *Econometrica*, 2006, *74*, 967–1012.

_ **and Alexander Chudik**, "Econometric Analysis of High Dimensional VARs Featuring a Dominant Unit," *Econometric Reviews*, 2013, *32*, 592–649.

Pragidis, Ioannis C., G.P. Aiellia, Dionysios Chionis, and Panagiotis Schizas, "Contagion Effects during Financial Crisis: Evidence from the Greek Sovereign Bonds Market," *Journal of Financial Stability*, 2015, *18*, 127–138.

Schüler, Yves S., "Asymmetric Effects of Uncertainty over the Business Cycle: A Quantile Structural Vector Autoregressive Approach," Working Paper 2014-02, University of Konstanz January 2014.

Sims, Christopher A., "Are Forecasting Models Usable for Policy Analysis?," *Minneapolis Federal Reserve Bank Quarterly Review*, 1986, *Winter*, 2–16.

Zellner, Arnold and Tomohiro Ando, "A Direct Monte Carlo Approach for Bayesian Analysis of the Seemingly Unrelated Regression Model," *Journal of Econometrics*, 2010, *159*, 33–45.

Zhu, Huiming, Xianfang Su, Yawei Guo, and Yinghua Ren, "The Asymmetric Effects of Oil Price Shocks on the Chinese Stock Market: Evidence from a Quantile Impulse Response Perspective," *Sustainability*, 2016, *8*, 1–19.

	Sovereign CDS				Financial CDS					Observed Factors			
	Mean	S.D.	Skew	Kurt	Mean	S.D.	Skew	Kurt		Mean	S.D.	Skew	Kurt
European Countries									*Factors*				
Austria	0.105	4.102	0.877	19.095	0.177	6.694	1.190	21.021	Rm-Rf	2.475	152.971	-0.089	7.223
Belgium	0.141	5.285	-0.676	20.709	0.148	7.481	9.342	287.651	SMB	1.182	65.265	0.034	4.097
France	0.115	3.202	-0.293	17.066	0.141	5.200	-0.204	13.751	HML	-0.484	81.816	0.688	8.164
Germany	0.051	1.751	0.211	12.120	0.118	10.167	-1.806	68.381	US 5y CMT Yield	-0.222	7.006	-0.113	3.169
Greece	5.720	196.789	-0.735	58.131	1.428	32.204	-0.385	16.559	TED Spread	-0.011	7.759	0.547	41.879
Ireland	0.357	12.760	-0.493	28.233	0.565	49.298	-1.137	56.276	Euribor-DeTBill	0.056	7.547	0.471	26.162
Italy	0.245	8.303	0.430	19.354	0.207	6.767	0.169	12.675	Equity VRP	0.008	19.349	-0.225	35.710
Netherlands	0.059	2.114	1.178	23.699	0.137	9.848	1.544	84.734	Term Premium	0.010	4.715	-0.240	3.204
Norway	0.020	1.403	0.496	27.955	0.077	3.187	-0.328	46.948	IG Spread	0.085	6.922	0.444	299.608
Portugal	0.733	16.821	-0.248	30.262	0.664	15.979	0.404	31.613	HY Spread	-0.005	7.642	0.231	8.454
Spain	0.236	8.115	-0.182	13.801	0.271	8.085	-0.665	18.275	ITRAXX Europe	0.061	4.674	0.027	8.000
Sweden	0.035	2.094	0.365	15.740	0.087	3.131	0.626	22.296	ITRAXX High Vol.	0.080	8.318	0.283	13.336
UK	0.047	2.345	-0.258	10.737	0.106	5.547	-0.253	26.297	ITRAXX Crossover	0.194	18.008	-0.060	9.420
									ITRAXX Sen. Fin.	0.130	6.643	-0.162	10.796
									ITRAXX Sub. Fin.	0.217	11.388	0.094	12.391
Non-European Countries									AUD/USD	-2.358	102.044	0.777	9.658
Australia	0.040	2.575	0.502	17.005	0.118	5.886	-4.337	111.245	CNY/USD	-1.552	10.257	0.071	6.969
China	0.065	4.809	0.103	42.408	0.126	16.692	-0.846	199.949	EUR/USD	-0.651	69.107	-0.044	1.841
Japan	0.077	2.557	2.135	38.132	0.084	4.190	1.402	18.931	GBP/USD	0.590	65.966	0.472	2.986
Russia	0.093	16.530	2.225	57.542	0.241	31.436	-0.442	64.587	JPY/USD	-2.552	71.902	0.131	6.189
US	0.025	1.523	0.645	11.995	0.214	24.257	-0.380	148.344	NOK/USD	-1.016	91.735	0.113	2.667
									RUR/USD	0.287	61.365	1.149	13.403
									SEK/USD	-1.086	94.605	0.028	2.449

NOTES: Descriptive statistics are reported over a period of 1,596 trading days from 03-Jan-2006 to 14-Feb-2012. With the exception of the Fama-French factors, the equity VRP and the exchange rate data, every series is expressed as a first difference and is measured in basis points. Descriptive statistics for the Fama-French factors are reported in basis points for the levels of the series. The equity VRP is expressed as a first difference and is measured in squared percentage units. The foreign exchange rates are bilateral spot rates expressed in foreign currency units per US dollar. For each exchange rate, we compute the daily log-return in basis points.

Table 1: Descriptive Statistics

	Sovereign Credit Risk						Financial Sector Credit Risk					
	To		From		Net		To		From		Net	
	OLS	LAD	OLS	LAD	OLS	LAD	OLS	LAD	OLS	LAD	OLS	LAD
Australia	6.90	1.29	19.52	9.06	-12.62	-7.77	8.14	3.25	16.99	5.18	-8.85	-1.92
Austria	24.00	5.21	9.83	6.00	14.18	-0.80	10.47	3.29	13.21	3.32	-2.73	-0.03
Belgium	11.79	3.10	14.52	6.19	-2.72	-3.08	7.80	2.78	4.13	0.77	3.67	2.02
China	19.06	4.77	13.68	7.89	5.38	-3.12	10.30	3.18	11.46	0.45	-1.16	2.73
France	17.04	5.71	9.03	6.85	8.02	-1.14	18.33	14.15	12.89	7.66	5.45	6.49
Germany	11.91	3.48	13.85	10.87	-1.93	-7.39	4.93	1.14	7.14	0.85	-2.21	0.29
Greece	3.82	3.05	7.30	0.06	-3.47	2.99	5.59	1.24	8.20	2.39	-2.61	-1.15
Ireland	10.11	7.16	6.46	3.03	3.65	4.13	2.11	1.10	7.94	0.42	-5.83	0.68
Italy	25.85	7.07	11.54	8.70	14.32	-1.63	18.17	4.32	10.41	5.38	7.76	-1.06
Japan	7.18	2.29	13.85	6.65	-6.67	-4.36	4.77	1.20	11.59	5.23	-6.83	-4.02
Netherlands	20.39	9.59	7.37	5.63	13.02	3.96	3.50	0.48	7.62	0.65	-4.12	-0.17
Norway	7.67	1.32	11.00	2.45	-3.33	-1.13	5.87	0.50	5.77	1.06	0.11	-0.56
Portugal	11.99	9.33	10.27	1.85	1.72	7.48	9.55	3.90	11.71	5.81	-2.16	-1.92
Russia	20.42	7.49	19.68	3.09	0.74	4.40	9.64	1.08	17.18	1.48	-7.54	-0.40
Spain	35.83	29.21	8.90	3.04	26.93	26.17	4.50	0.72	18.39	8.23	-13.89	-7.51
Sweden	10.83	3.17	12.53	3.97	-1.70	-0.80	9.15	1.39	10.52	3.74	-1.37	-2.35
UK	14.47	3.39	12.63	7.20	1.84	-3.81	7.66	2.60	14.07	3.68	-6.41	-1.08
US	10.23	2.40	14.36	7.07	-4.12	-4.67	3.64	0.95	8.14	0.42	-4.50	0.53

NOTES: Values are reported in percent.

Table 2: Topology of the Credit Risk Network at the Conditional Mean and Median

	Mean	Med.	5%	95%	RTD
Mean	1.00	0.56	-0.11	0.79	0.58
Med.	0.56	1.00	0.07	0.55	0.32
5%	-0.11	0.07	1.00	-0.30	-0.78
95%	0.79	0.55	-0.30	1.00	0.83
RTD	0.58	0.32	-0.78	0.83	1.00

Table 3: Correlation among Spillover Indices Evaluated at Selected Quantiles

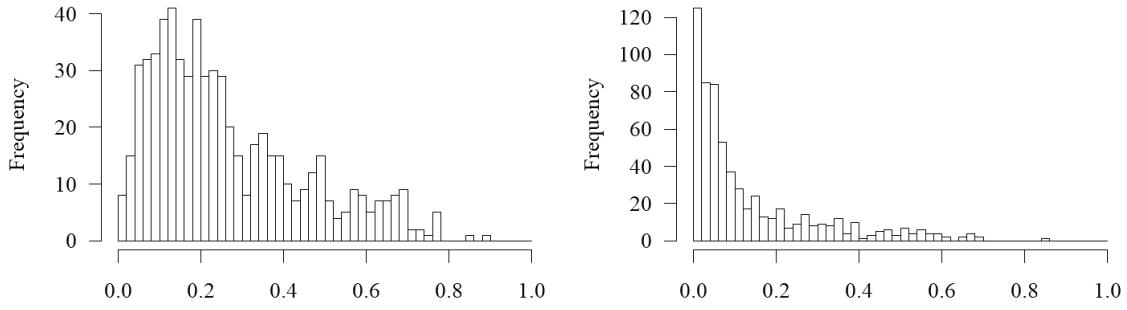

(a) Simple VAR(1) Model

(b) Factor VAR(1) Model

NOTES: The histograms show the distribution of the absolute pairwise correlations between the residuals of the simple VAR(1) model and our factor VAR(1) model evaluated at the conditional mean by OLS.

Figure 1: Comparison of Absolute Residual Correlations, with and without Factors

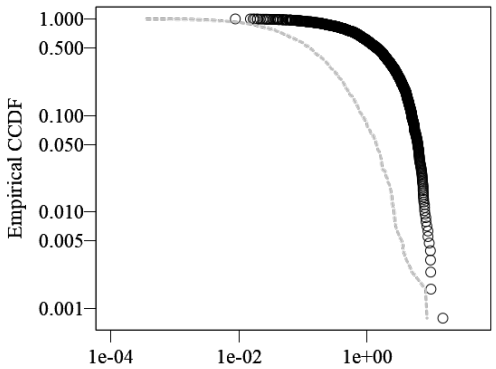

NOTES: The spillover density for the model with factors is shown as a dashed gray line, while the black circles show the spillover density for the model without factors.

Figure 2: Spillover Density at the Conditional Mean, with and without Factors

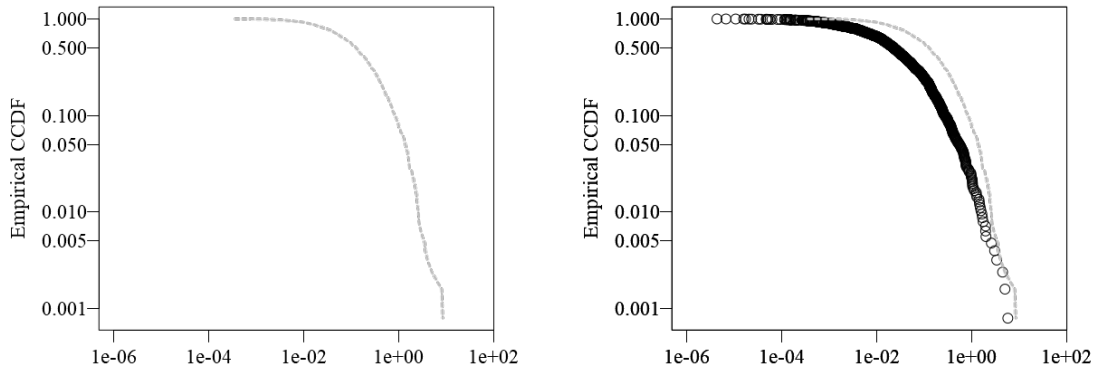

(a) OLS

(b) $\tau = 0.50$

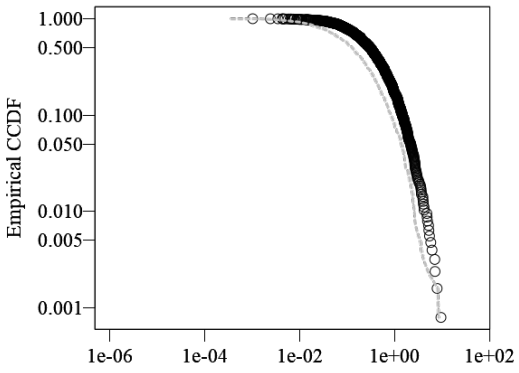

(c) $\tau = 0.05$

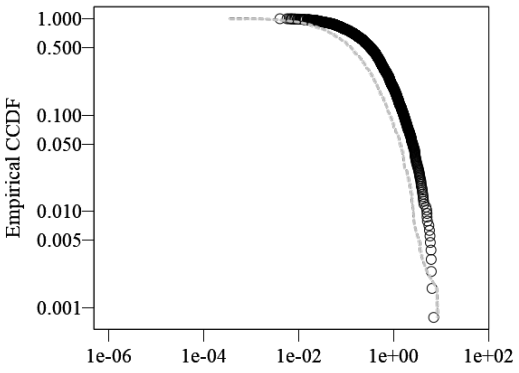

(d) $\tau = 0.95$

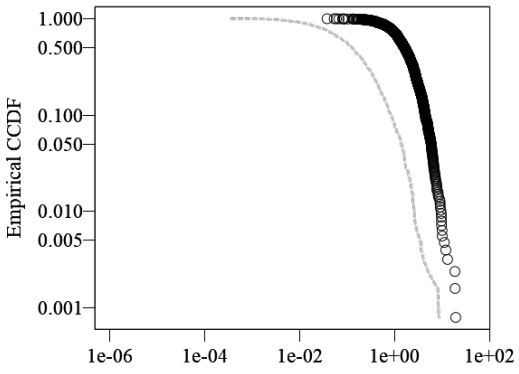

(e) $\tau = 0.01$

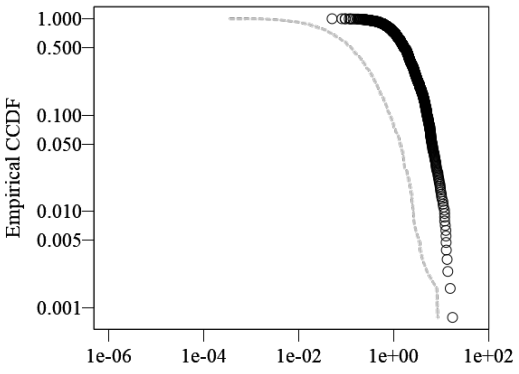

(f) $\tau = 0.99$

NOTES: The figure reports the empirical counter cumulative distribution function (CCDF) of the $m(m-1)$ off-diagonal elements of the adjacency matrix on a logarithmic scale. The CCDF under OLS is shown as a dashed line in every panel for comparison.

Figure 3: Quantile Variation in the Shape of the Spillover Density

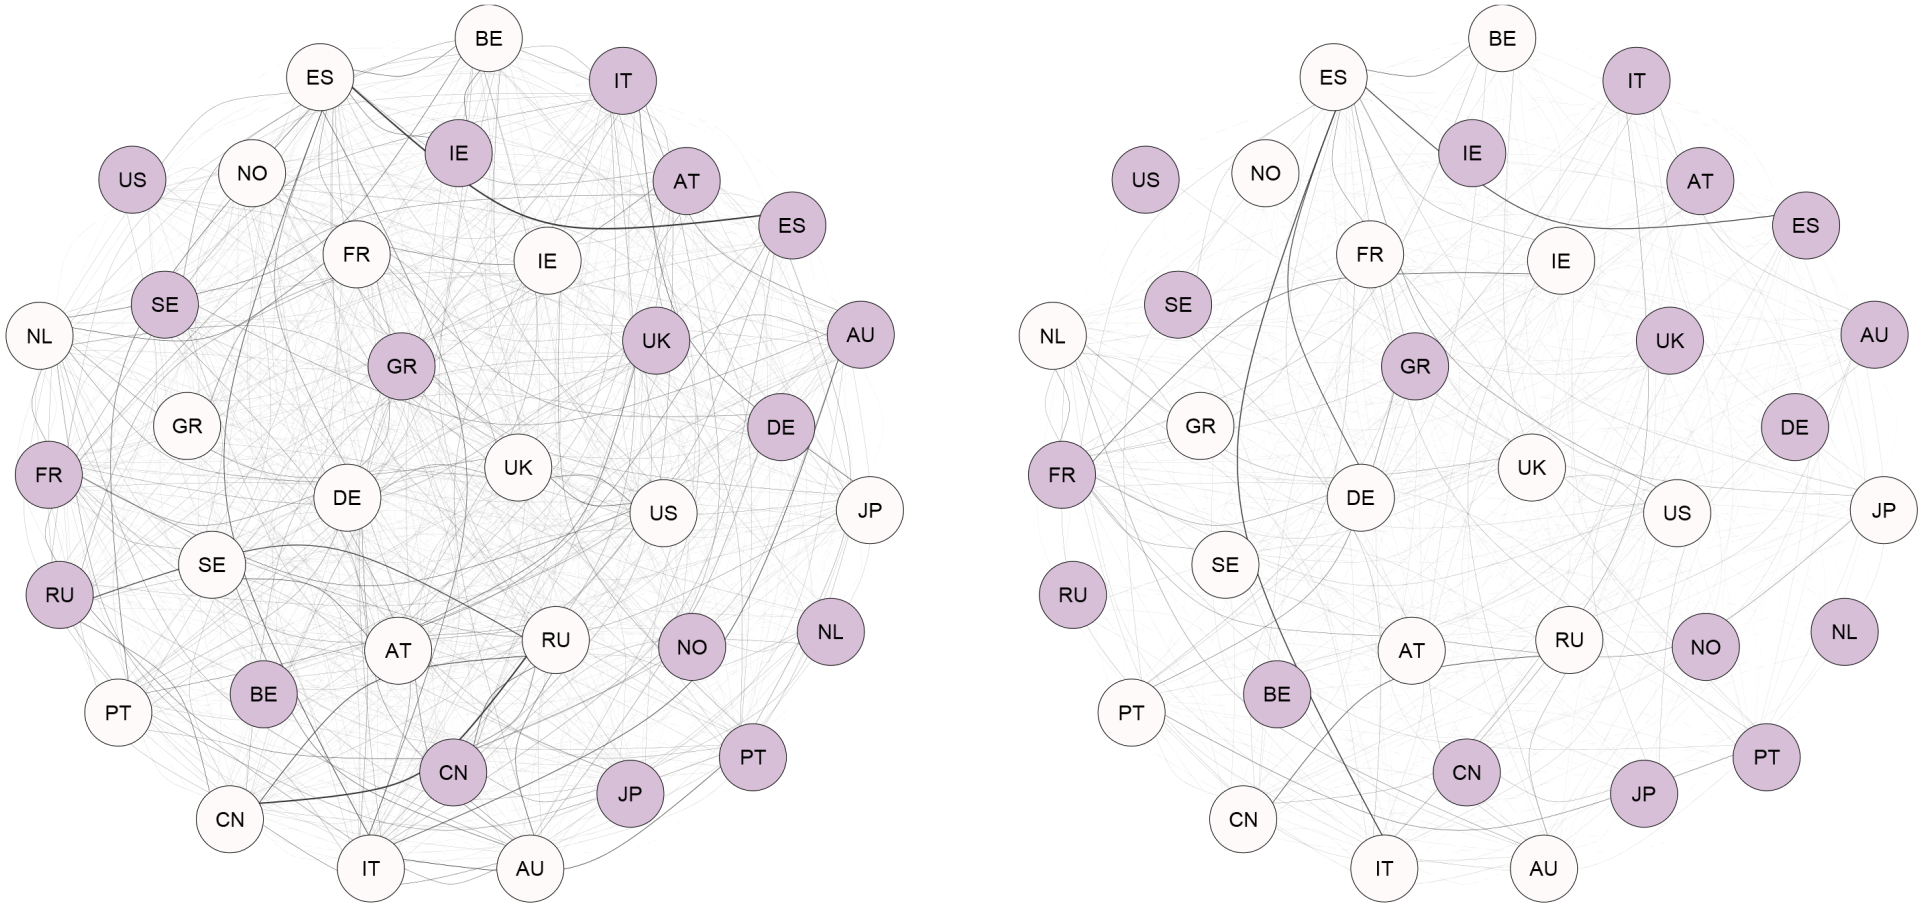

(a) Conditional Mean (Ordinary Least Squares Estimator)

(b) Conditional Median (Least Absolute Deviations Estimator)

NOTES: Countries are identified by their respective alpha-2 ISO codes, with white nodes corresponding to sovereigns and shaded nodes to national financial sectors. Edges are drawn as curves and the direction of pairwise spillovers is anticlockwise (i.e. the strong spillover between the Spanish sovereign and the Spanish financial sector goes from the sovereign to the financial sector). The thickness of an edge is proportional to its weight. The layout of the nodes is determined by applying the force-directed algorithm of Fruchterman and Reingold (1991) to the network evaluated at the conditional mean.

Figure 4: Network Visualisations for the Model Evaluated at the Conditional Mean and at the Conditional Median

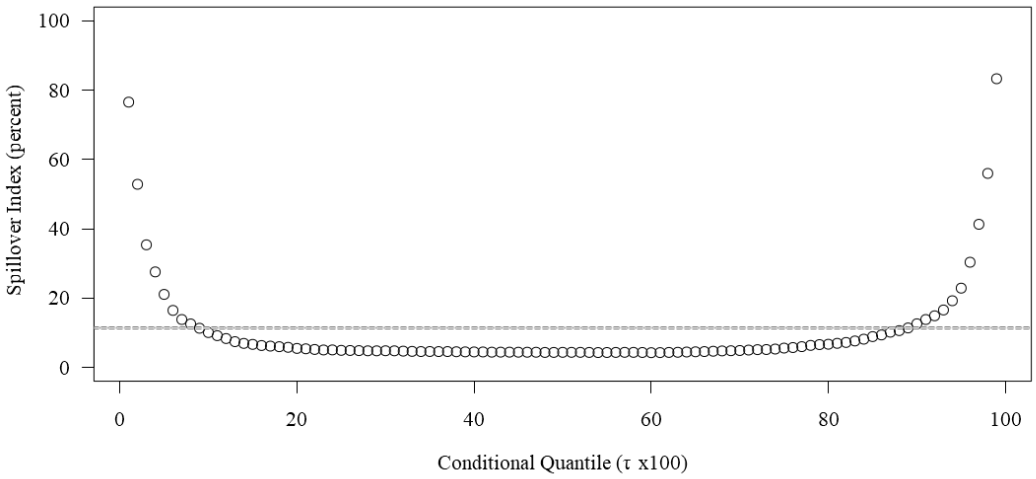

NOTES: The figure reports the value of the spillover index defined in (22) evaluated at the τth conditional quantile (plotted as a circle) relative to the value at the conditional mean (shown by the dashed line).

Figure 5: Variation in the DY Spillover Index over the Conditional Distribution

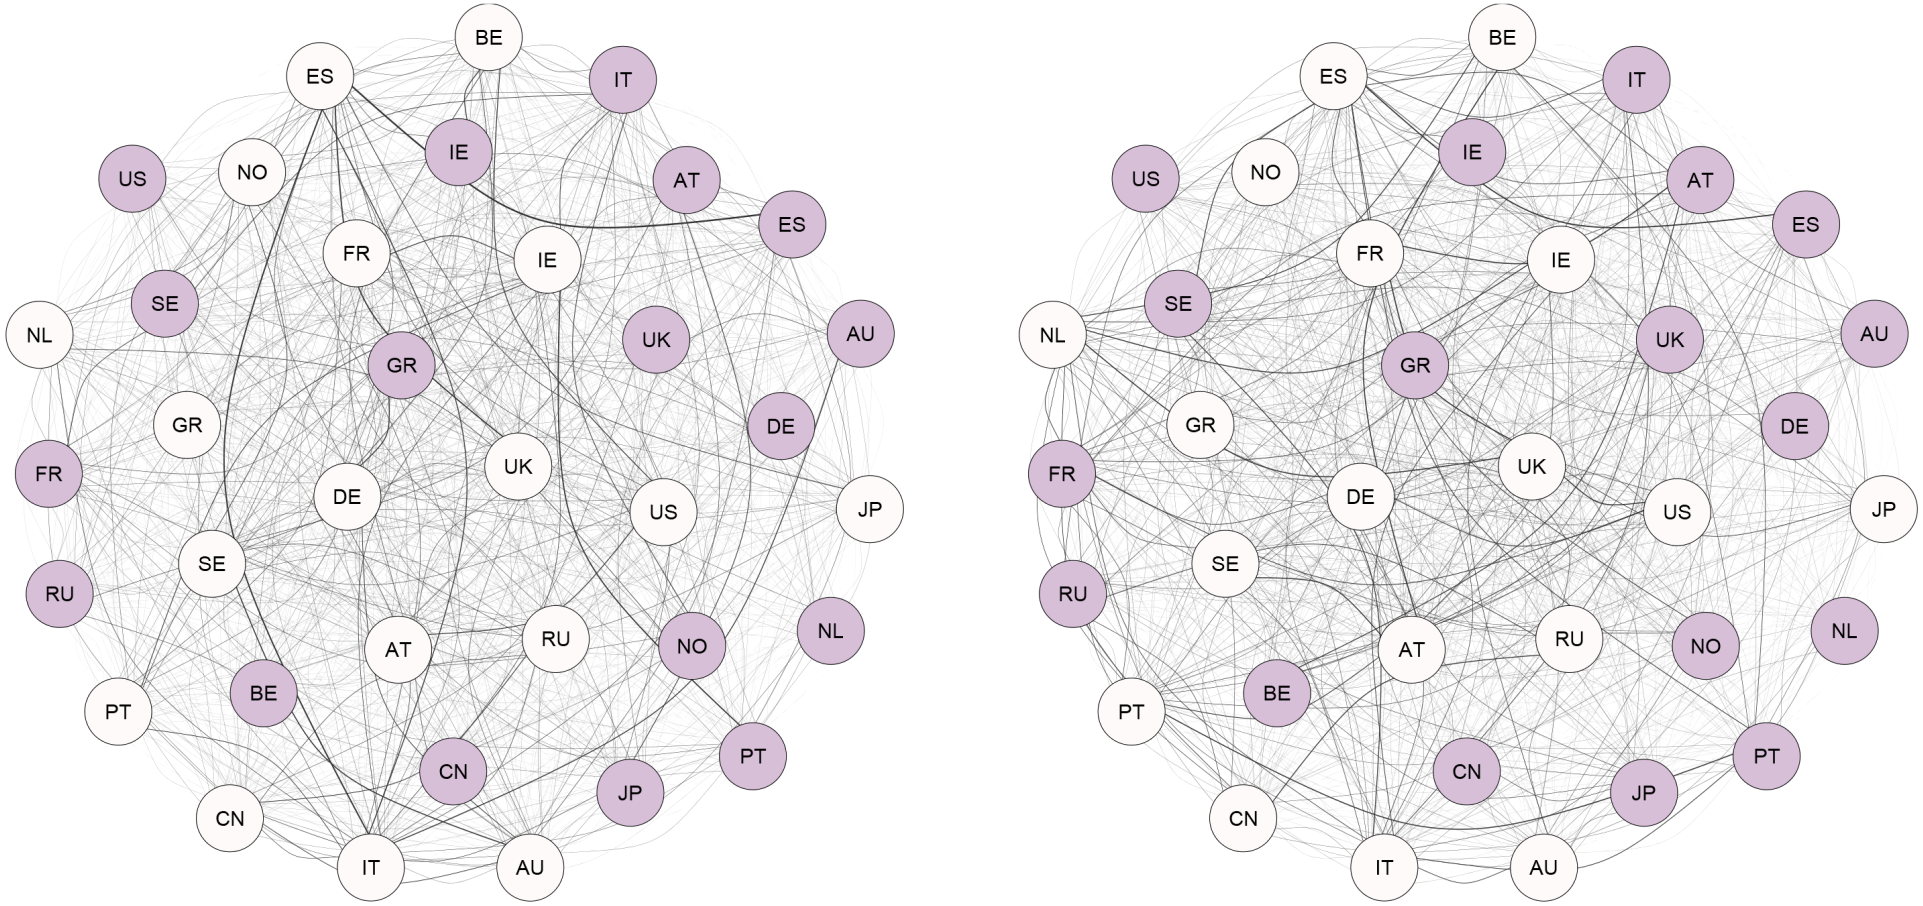

(a) $\tau = 0.05$ (b) $\tau = 0.95$

NOTES: Countries are identified by their respective alpha-2 ISO codes, with white nodes corresponding to sovereigns and shaded nodes to national financial sectors. Edges are drawn as curves and the direction of pairwise spillovers is anticlockwise. The thickness of an edge is proportional to its weight. The layout of the nodes is determined by applying the force-directed algorithm of Fruchterman and Reingold (1991) to the network evaluated at the conditional mean.

Figure 6: Network Visualisations for the Model Evaluated by Quantile Regression at $\tau = 0.05$ and $\tau = 0.95$

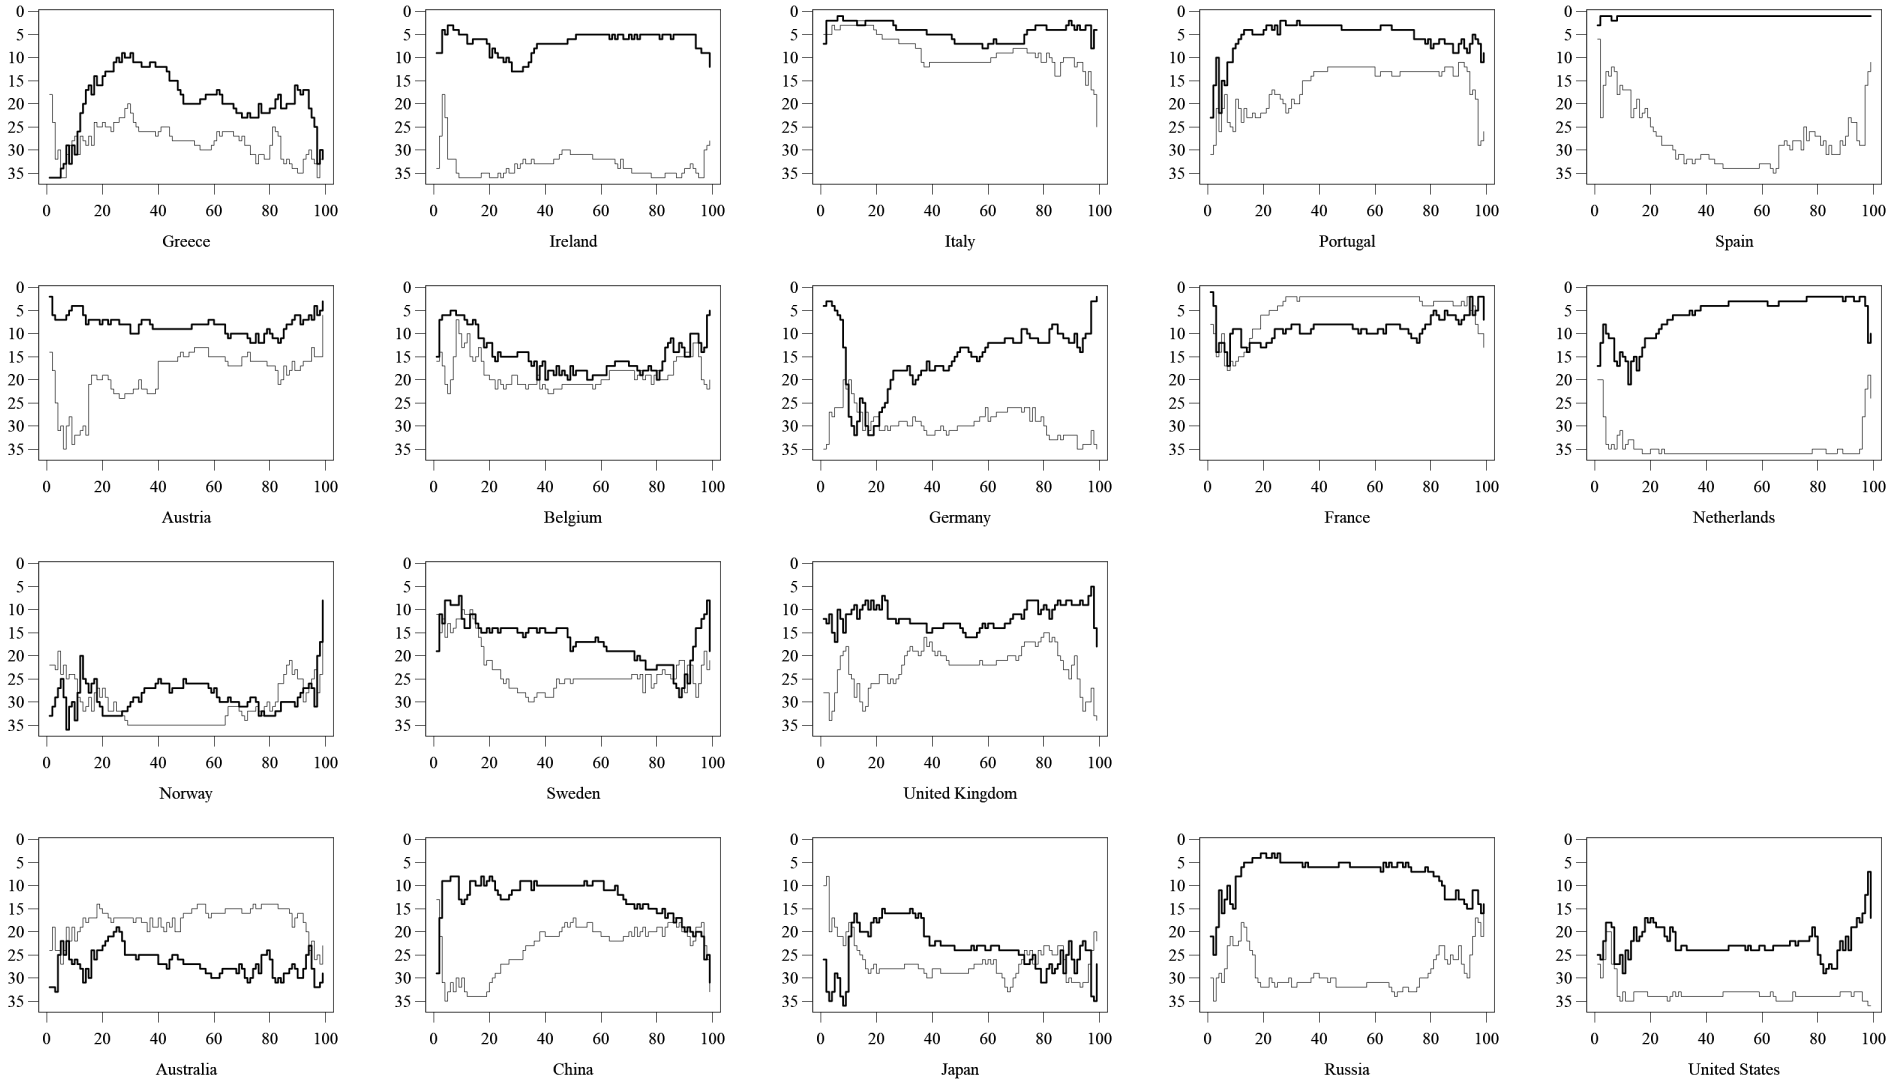

NOTES: The horizontal axis shows the quantile index multiplied by 100 while the vertical axis records the weighted out-degree rank from 1 (strongest) to 36 (weakest).

Figure 7: Weighted Out-Degree Rank of the Financial Sector (Fine Line) and the Sovereign (Heavy Line), by Country

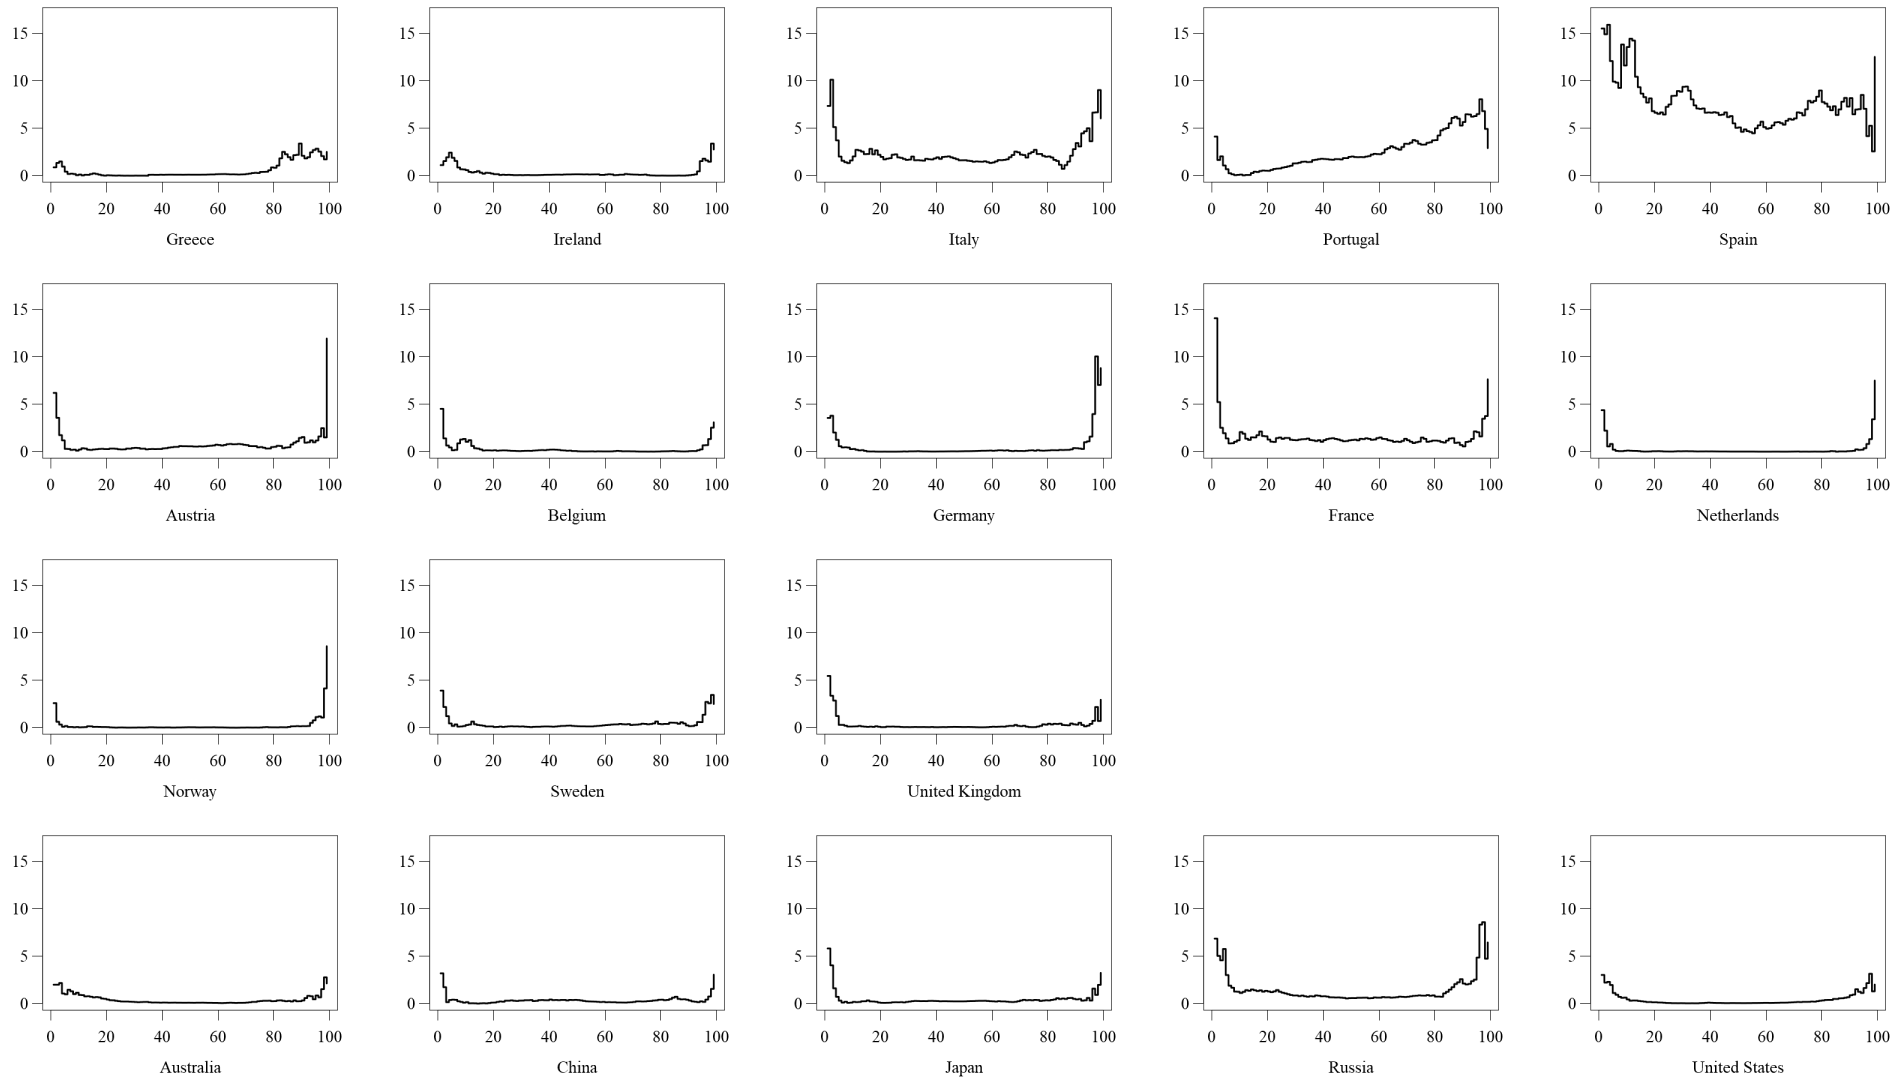

NOTES: The horizontal axis shows the quantile index multiplied by 100 while the vertical axis records the strength of the sovereign–financial feedback effect, $\mathcal{T}^{(5)}_{s_i \leftrightarrow f_i,(\tau)} = T^{(5)}_{s_i \leftarrow f_i,(\tau)} + T^{(5)}_{f_i \leftarrow s_i,(\tau)}$, measured in percent.

Figure 8: Total Feedback between the Financial Sector and the Sovereign by Country

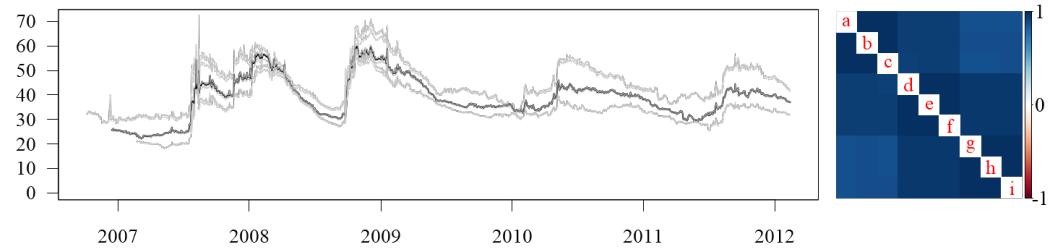

(a) Conditional Mean

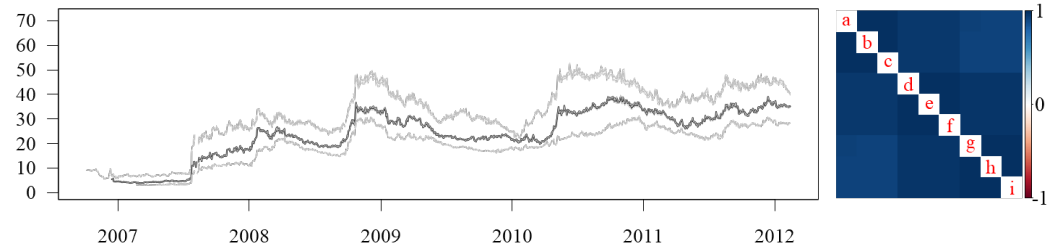

(b) Conditional Median

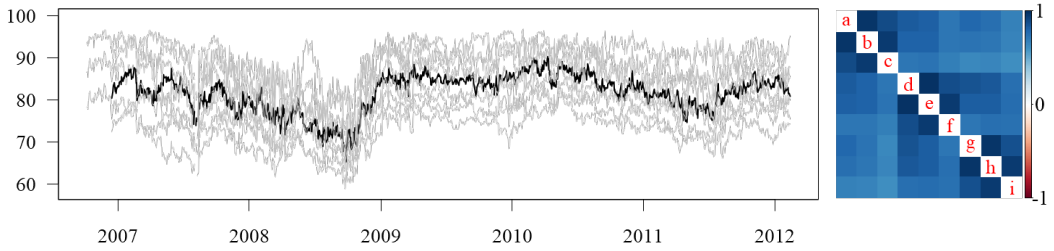

(c) Fifth Conditional Quantile

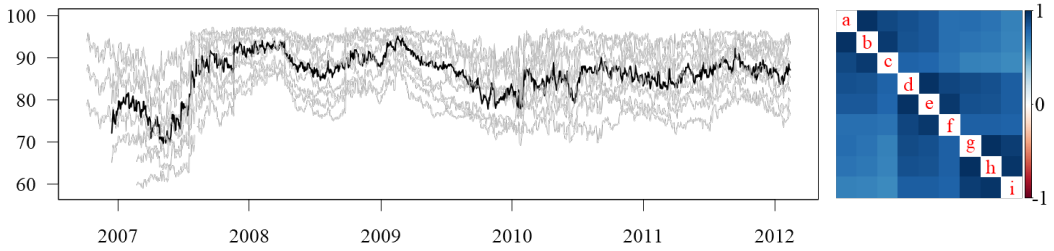

(d) Ninety-Fifth Conditional Quantile

NOTES: In each panel, results for our baseline setting with $w = 250$ and $h = 5$ trading days are shown as a heavy black line. Results for each other combination of $w \in \{200, 250, 300\}$ and $h \in \{3, 5, 10\}$ trading days are shown as fine gray lines. Letters a–i in the common-sample correlation heatmaps refer to the following combinations: (a) $w = 200$, $h = 3$; (b) $w = 200$, $h = 5$; (c) $w = 200$, $h = 10$; (d) $w = 250$, $h = 3$; (e) $w = 250$, $h = 5$; (f) $w = 250$, $h = 10$; (g) $w = 300$, $h = 3$; (h) $w = 300$, $h = 5$; and (i) $w = 300$, $h = 10$.

Figure 9: Robustness to the Choice of Rolling Window and Forecast Horizon

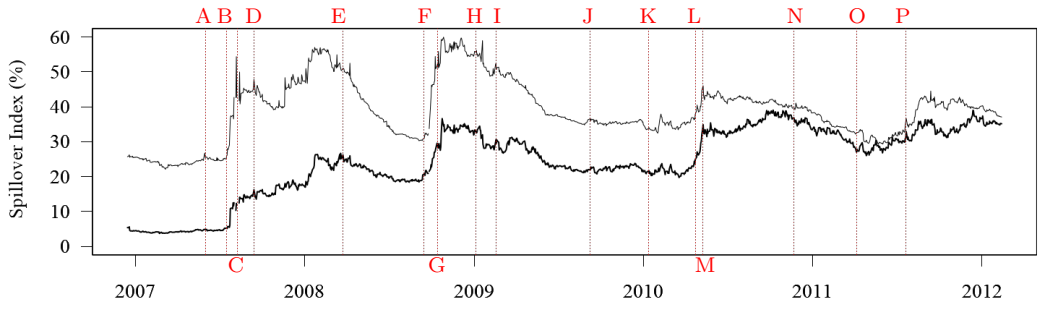

(a) Conditional Mean and Median

(b) Fifth Conditional Quantile

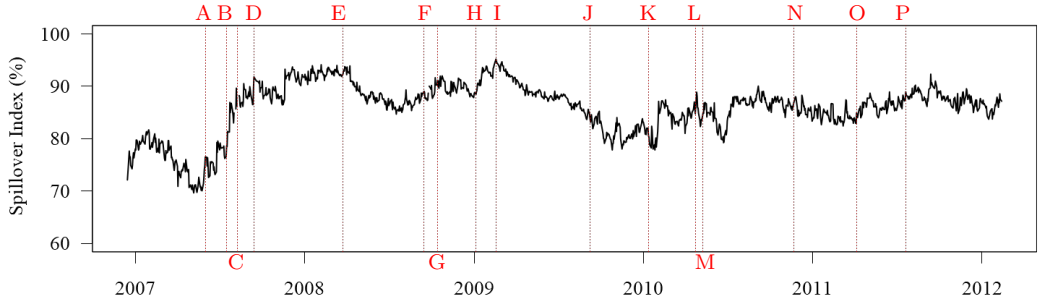

(c) Ninety-Fifth Conditional Quantile

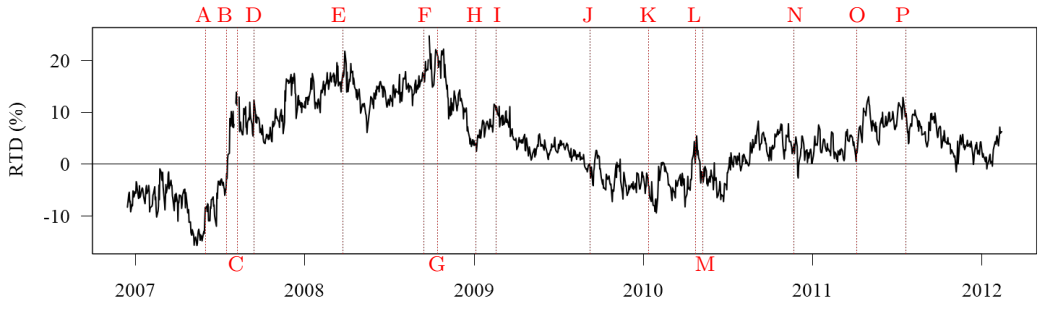

(d) Relative Tail-Dependence

NOTES: The figure reports the value of the spillover index defined in (22) evaluated at the mean (shown as a fine line in panel (a)) and at the 5th, 50th and 95th conditional quantiles (shown as heavy lines in panels (a)–(c)) as well as the difference between the spillover index at the 95th and 5th conditional quantiles (shown as a heavy line in panel (d)). The results are obtained from rolling regressions with a window length of 250 trading days. The date shown corresponds to the last day of each rolling window. We suppress four rolling samples where the system exhibits instability at the conditional mean. The following events are marked: A: S&P and Moody's downgrade bonds backed by subprime loans (01-Jun-07); B: Bear Stearns announces hedge fund losses (17-Jul-07); C: BNP Paribas halts redemptions on three investment funds (09-Aug-07) D: UK Treasury announces liquidity support for Northern Rock (14-Sep-07); E: Bear Stearns is acquired by JP Morgan (24-Mar-08); F: Lehman Brothers files for bankruptcy (15-Sep-08); G: the Troubled Asset Relief Program is announced (14-Oct-08); H: the Fed begins purchasing mortgage-based securities issues by Fannie Mae and Freddie Mac (05-Jan-09); I: signing of the American Recovery and Reinvestment Act (17-Feb-09); J: Greek parliament is dissolved (08-Sep-09); K: report on the falsification of Greek data released (12-Jan-10); L: Greece requests aid (23-Apr-10); M: the European Financial Stability Facility is announced (09-May-10); N: Ireland requests aid (22-Nov-10); O: Portugal requests aid (06-Apr-11); and P: second Greek bailout (22-Jul-11).

Figure 10: Time-Varying Dependence at Selected Quantiles

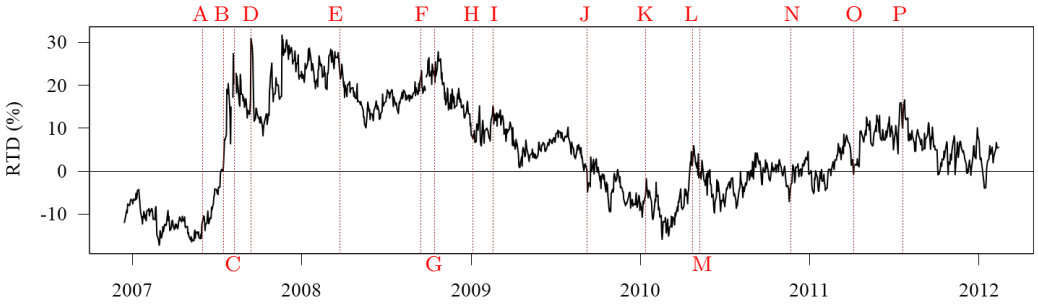

(a) 10% Relative Tail-Dependence, $S^{(5)}_{0.90} - S^{(5)}_{0.10}$

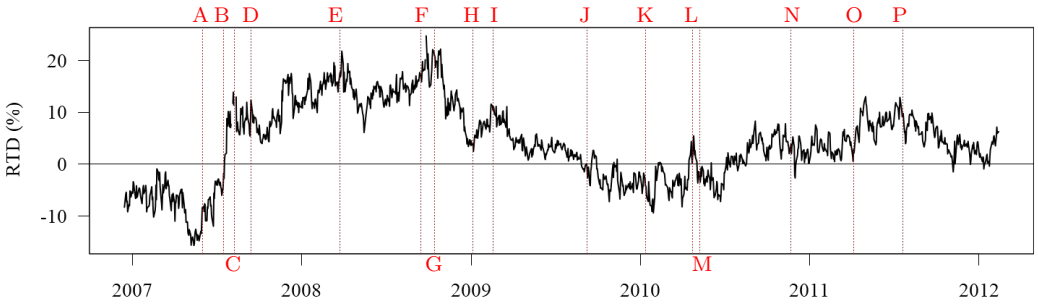

(b) 5% Relative Tail-Dependence, $S^{(5)}_{0.95} - S^{(5)}_{0.05}$

(c) 1% Relative Tail-Dependence, $S^{(5)}_{0.99} - S^{(5)}_{0.01}$

NOTES: The results are obtained from rolling regressions with a window length of 250 trading days. The date shown corresponds to the last day of each rolling window. We suppress four rolling samples where the system exhibits instability at the conditional mean. The following events are marked: A: S&P and Moody's downgrade bonds backed by subprime loans (01-Jun-07); B: Bear Stearns announces hedge fund losses (17-Jul-07); C: BNP Paribas halts redemptions on three investment funds (09-Aug-07) D: UK Treasury announces liquidity support for Northern Rock (14-Sep-07); E: Bear Stearns is acquired by JP Morgan (24-Mar-08); F: Lehman Brothers files for bankruptcy (15-Sep-08); G: the Troubled Asset Relief Program is announced (14-Oct-08); H: the Fed begins purchasing mortgage-based securities issues by Fannie Mae and Freddie Mac (05-Jan-09); I: signing of the American Recovery and Reinvestment Act (17-Feb-09); J: Greek parliament is dissolved (08-Sep-09); K: report on the falsification of Greek data released (12-Jan-10); L: Greece requests aid (23-Apr-10); M: the European Financial Stability Facility is announced (09-May-10); N: Ireland requests aid (22-Nov-10); O: Portugal requests aid (06-Apr-11); and P: second Greek bailout (22-Jul-11).

Figure 11: Alternative Measures of Relative Tail-Dependence