Author/s:
Greenwood-Nimmo, M;Nguyen, VH;Shin, Y

Title:
What is mine is yours: Sovereign risk transmission during the European debt crisis

Date:
2023-04

Citation:
Greenwood-Nimmo, M., Nguyen, V. H. & Shin, Y. (2023). What is mine is yours: Sovereign risk transmission during the European debt crisis. Journal of Financial Stability, 65, https://doi.org/10.1016/j.jfs.2023.101103.

Persistent Link:
https://hdl.handle.net/11343/337197

What is Mine is Yours: Sovereign Risk Transmission during the European Debt Crisis*

Matthew Greenwood-Nimmo
Department of Economics, University of Melbourne
Centre for Applied Macroeconomic Analysis, Australian National University
Email: mgreenwood@unimelb.edu.au Tel: +61 (0)383-445-354

Viet Hoang Nguyen
Melbourne Institute of Applied Economic and Social Research
Email: vietn@unimelb.edu.au Tel: +61 (0)390-353-621

Yongcheol Shin
Department of Economics and Related Studies, University of York
Email: yongcheol.shin@york.ac.uk Tel: +44 (0)1904-323-757

Abstract

We develop an empirical network model to characterize the density of bilateral sovereign credit risk spillovers during the European debt crisis. We show that the spillover density is often asymmetric with heavy tails and that its location and shape vary strongly and systematically in relation to published indicators of systemic stress. Using auxiliary panel data models, we show that the intensity of bilateral spillovers is related to the portfolio investment exposures among country-pairs. Because our spillover statistics can be updated daily, they represent a valuable supplement to existing weekly and monthly measures of systemic stress.

JEL Codes: C58, F45, G15, H63.

Keywords: Sovereign credit risk network, credit default swaps (CDS), contagion and systemic stress, network models and connectedness, spillover density.

*Address for correspondence: M.J. Greenwood-Nimmo, 3.12 Faculty of Business and Economics, 111 Barry Street, University of Melbourne, VIC 3053, Australia. Email: matthew.greenwood@unimelb.edu.au. Tel: +61 (0)3 8344 5354. We thank the Editor, Iftekhar Hasan, and two anonymous referees whose input has greatly enriched our analysis. We are grateful to Kathryn St John and Catherine Tanuwidjaja for conscientious research assistance and for the constructive discussion of In Choi, Juan Carlos Cuestas, Renee Fry-McKibbin, Jongsuk Hahn, Jingong Huang, John Hunter, Minjoo Kim, Faek Menla Ali, Barry Rafferty, Chris Skeels, Tomasz Woźniak and Eliza Wu. This paper has benefited from comments raised by participants at the 2015 Conference of the Paul Woolley Centre for the Study of Capital Market Dysfunctionality (Sydney, December 2015), the Australasian Meeting of the Econometric Society (Sydney, July 2016), the 2016 Asian Meeting of the International Finance and Banking Society (Brunei, August 2016), the 2016 KAEA-KIPF Conference (Sejong City, August 2016), the 2016 University of Melbourne–Toulouse School of Economics Workshop (Melbourne, December 2016), the 2017 Singapore Economic Review Conference (Singapore, August 2017), the 72nd European Meeting of the Econometric Society (Manchester, August 2019) and by seminar participants at Deakin and Sogang Universities and the Universities of Catania and Piraeus. Greenwood-Nimmo acknowledges financial support from the Australian Research Council (Grant Number DE150100708) and the Faculty of Business and Economics at the University of Melbourne. The usual disclaimer applies.

I. Introduction

As concerns over the sustainability of sovereign debt swept across Europe in the wake of the global financial crisis, tackling sovereign credit contagion emerged as a key priority among European policymakers (e.g. Constâncio, 2012). Among the challenges that they faced were an incomplete understanding of the network structure of sovereign credit risk and the factors that influence it, as well as the absence of techniques for monitoring systemic stress in the market for sovereign debt in close-to-real time. In this paper, we make progress on each front. We develop a dynamic network model to characterize the comovement of idiosyncratic sovereign credit risk among a group of 23 European sovereigns over the period January 2$^{\text{nd}}$ 2006 to July 27$^{\text{th}}$ 2015, with excess comovement relative to a benchmark providing evidence of sovereign credit contagion. By scrutinizing the cross-sectional and dynamic properties of the network, we identify phenomena that expose sovereigns to credit contagion and develop new indicators of systemic stress in the market for sovereign debt that are available at higher frequency than existing alternatives.

Sovereign credit risk can be measured using sovereign credit default swap (SCDS) spreads. An SCDS operates like an insurance contract in which a bondholder pays a premium to transfer the default risk of the bond onto the protection seller over a given time-frame. Due to its liquidity and to the engagement of many well-informed institutional investors, the market for credit default swaps is the leading forum for credit risk price discovery (Blanco et al., 2005). Sovereign credit contagion is associated with excess comovement among SCDS spreads relative to the level of comovement observed in normal states of the world. Excess comovement can be driven by changes in the sensitivity of SCDS spreads to common factors or by rising cross-sectional comovement of the idiosyncratic components of SCDS spreads. We focus on the latter, which represents pure contagion unrelated to fundamentals and which corresponds to the notion of residual contagion laid out by Bekaert et al. (2014) in their study of equity market contagion.

We begin by purging the SCDS spreads of the influence of a range of global and domestic factors identified in the literature, including global measures of investor risk appetite and funding liquidity and country-specific measures of economic performance (e.g. Pan and Singleton, 2008; Remolona et al., 2008; Longstaff et al., 2011; Ang and Longstaff, 2013; Montfort and Renne, 2014; Augustin, 2018). The defactored SCDS spreads that we obtain in this way isolate the idiosyncratic component of each sovereign's credit risk. We then

use the vector autoregressive (VAR) approach to network analysis developed by Diebold and Yilmaz (2009; 2014, hereafter collectively DY) to infer the network structure of the defactored SCDS spreads. DY show that the forecast error variance decomposition of a VAR model can be interpreted as a weighted directed network. For an m-variable VAR, the forecast error variance decomposition estimates how much of the future uncertainty associated with variable i can be attributed to shocks affecting variable j, $i, j = 1, \ldots, m$.

By defactoring and estimating over rolling samples, we are able to track the evolution of the credit risk network over time. The established method to measure the intensity of bilateral spillovers in DY networks is via the spillover index, which is proportional to the mean bilateral spillover. Increases in the spillover index are typically interpreted in relation to systemic stress and contagion (e.g. Alter and Beyer, 2014; Claeys and Vašíček, 2014; Ballester et al., 2016; Bostanci and Yilmaz, 2020). However, this approach does not utilize the information in the higher moments of the spillover density (i.e. the density of the bilateral spillover effects). Neglecting this information may be costly, because we show that the spillover density obtained from the DY method is typically asymmetric with heavy tails. In such a setting, it is easy to conceive of scenarios in which meaningful changes in the shape of the spillover density are not reflected in its mean. Such behavior may sound contrived but we show that this exact phenomenon arises during the European debt crisis, as the decoupling of the peripheral eurozone sovereigns from the eurozone core discussed by Antonakakis and Vergos (2013) generates growth in the right tail of the spillover density while its mean is falling. To address this issue, we develop a framework to characterize and track the entire spillover density via kernel density estimation (KDE).

Some degree of interdependence in idiosyncratic sovereign credit risk across countries is to be expected in normal states of the world. To identify episodes of contagion, we must distinguish between network structures that are consistent with normal levels of interdependence and those that are not. We treat the estimated network obtained over a tranquil period prior to the global financial crisis (GFC) as the benchmark network. This is an obvious choice because, as Acharya et al. (2014) point out, there is no evidence that sovereign credit risk was a concern for developed markets at this time. Having defined a suitable benchmark, evidence of idiosyncratic sovereign credit risk spillovers in excess of the benchmark is evidence of residual sovereign credit contagion.

The benchmark spillover density has a peak close to zero and a long right tail, driven

by a handful of stronger spillovers among country-pairs with close geopolitical linkages, such as Hungary and Poland. These features are maintained until the outbreak of the subprime crisis in mid-2007, which instigates a rightward shift in the spillover density coupled with a reduction in its skewness that occurs in several steps that coincide with adverse news regarding Bear Stearns and with the failure of Lehman Brothers. The intensification of idiosyncratic sovereign credit risk spillovers over this time relative to the pre-GFC benchmark constitutes evidence of residual sovereign credit contagion.

The evidence of contagion is sustained from late-2008 until mid-2011, with the location of the spillover density remaining largely unchanged. Claeys and Vašíček (2014) document a similar plateau in spillover intensity at this time. Many key events of the European debt crisis, including the dissolution of the Greek parliament, the revelation of inaccuracies in Greek economic data and a number of sovereign bailouts are scarcely reflected in the spillover index. After the Portuguese request for aid in mid-2011, the spillover density gradually shifts leftward until the end of our sample. Over this entire period, changes in the shape of the spillover density convey strong signals about changes in the transmission of idiosyncratic sovereign risk. For example, the Greek request for aid in 2010 generates a surge in the right tail of the spillover density that reflects an intensification of credit risk spillovers among peripheral sovereigns. Furthermore, even as the mean of the spillover density falls late in our sample, its right tail grows to reflect strengthening spillovers among the GIIPS (Greece, Ireland, Italy, Portugal and Spain). This is a manifestation of the decoupling of the GIIPS from the eurozone core and offers an excellent illustration of the value of characterizing the entire spillover density, because the reduction in the mean of the spillover density masks evidence of contagion in its right tail at this time.

Having documented evidence of residual contagion over our sample period, we turn our attention to the factors that explain the strength of idiosyncratic credit risk transmission. Using an auxiliary panel data model, we show that the intensity of bilateral spillovers is positively related to bilateral portfolio investment exposures and negatively related to bilateral portfolio investment flows. Cross-border portfolio investment positions expose sovereigns to foreign shocks and the retrenchment of international investments offers a channel through which idiosyncratic sovereign risk shocks can propagate internationally. Furthermore, we show that fast-growing countries and those with large cyclically adjusted budget deficits have a higher propensity both to generate and to receive credit risk

spillovers. These findings are consistent with a body of work that emphasizes the role of financial linkages and portfolio investment in episodes of contagion (e.g. Van Rijckeghem and Weder, 2001; Caramazza et al., 2004; Forbes, 2013; Fry-McKibbin et al., 2014) and the destabilizing influence of mutual funds documented by Raddatz and Schmukler (2012).

Our final contribution is to show that changes in the shape of the spillover density are strongly and systematically related to changes in systemic stress, measured by the European Central Bank's (ECB) Composite Indicator of Systemic Stress (CISS) and its sovereign counterpart (SovCISS). Increases in both CISS and SovCISS are associated with statistically significant rightward shifts in the spillover density, reduced skewness and a greater concentration of mass close to the mean. This finding has an important practical implication. The CISS and SovCISS are only reported at weekly and monthly frequency, respectively. By contrast, our model can be used to construct a new spillover density daily, providing a close-to-real time indicator of systemic stress in the market for sovereign debt. By equipping policymakers with timely indications of changes in systemic stress, our framework can make a valuable contribution to the formulation and implementation of stabilization policies during fast-moving crises.

Our work provides several contributes to the literature. First, we add to the literature that has used variants of the DY technique to study sovereign credit risk. The paper that is most closely related to ours is Claeys and Vašíček (2014), which is among the first applications of the DY technique to account for common factors. In their analysis of credit risk spillovers among 16 EU sovereign bond markets up to 2012, the authors document a pronounced increase in bilateral spillovers up to 2008, after which they plateau, in keeping with our results. Other related papers from this literature include Alter and Beyer (2014), Buse and Schienle (2019), Bostanci and Yilmaz (2020) and Ando et al. (2022). Unlike these papers, we exploit the entire spillover density and draw explicit comparisons against a benchmark density to look for evidence of sovereign credit contagion.

Second, our paper sits within the broader literature on sovereign credit risk that has considered additional channels beyond pure residual contagion. Gómez-Puig and Sosvilla-Rivero (2016) use a dynamic Granger causality framework and find evidence of both pure and fundamentals-based contagion during the sovereign debt crisis. Giordano et al. (2013) and Beirne and Fratzscher (2013) both find evidence of wake-up call contagion, whereby a shock prompts investors to pay greater attention to the underlying country-specific

fundamentals that ultimately determine sovereign creditworthiness. By contrast, using quantile regressions, Caporin et al. (2018) argue that changes in the mechanisms by which shocks propagate (so-called shift contagion) did not contribute to the European debt crisis. Meanwhile, Aït-Sahalia et al. (2014) model eurozone SCDS spreads as a mutually-exciting jump process. Their results indicate that both self-excitation and asymmetric cross-excitation explain the evolution of SCDS spreads.

This paper proceeds as follows. In Section II, we introduce our dataset and summarize our analytical framework. In Section III, we present our estimation results and scrutinize the properties of the sovereign credit risk network. We conclude in Section IV. Further details of our dataset and the results of several additional tests and robustness exercises is presented in the appendices.

II. Analytical Framework

A. The Dataset

Let c_{it} denote the SCDS spread for sovereigns $i = 1, 2, \ldots, m$ in periods $t = 1, 2, \ldots, T$, measured in basis points. We work with SCDS spreads obtained from Markit on US dollar denominated SCDS contracts for unsecured sovereign debt with a tenor of five years under a complete restructuring clause. Our selection of these specific SCDS contract parameters accords with the market conventions documented by Bai and Wei (2017).

We assume that c_{it} can be decomposed into components driven by exposure to global and domestic factors and an idiosyncratic component, as follows:

$$\Delta c_{it} = \boldsymbol{\lambda}_i^{f\prime} \Delta \boldsymbol{f}_t + \boldsymbol{\lambda}_i^{d\prime} \Delta \boldsymbol{d}_{it} + u_{it}, \tag{1}$$

where $\boldsymbol{\lambda}_i^f$ and $\boldsymbol{\lambda}_i^d$ are $f \times 1$ and $d \times 1$ vectors of heterogeneous loadings on the global and domestic factors, $\boldsymbol{f}_t$ and $\boldsymbol{d}_{it}$ are $f \times 1$ and $d \times 1$ vectors of global and domestic factors, respectively, and u_{it} is the idiosyncratic innovation specific to sovereign i. In principle, sovereign credit contagion can arise through variations in the factor loadings or via increased comovement of the u_{it}s in the cross-section dimension. We focus on the latter and so our first step is to defactor the SCDS spreads, thereby removing the influence of the common factors. The common factors may be either observed or latent. We limit our

attention to the case of observed factors. There are several reasons to favor observed factor models over latent factor models in general, including their ease of implementation and their straightforward interpretation. Furthermore, in the specific case of SCDS spreads, there is compelling evidence that a range of market-determined variables act as common factors. Longstaff et al. (2011) show that a selection factors related to US stock and high-yield markets explain a considerable proportion of the variation in the cross-section of SCDS spreads. Augustin (2018) further develops this finding, demonstrating that the same set of factors also affects the slope of the term structure of SCDS spreads.

In keeping with the results of Longstaff et al. (2011) and Augustin (2018), we employ the following US macroeconomic and financial indicators as common factors, which proxy for global economic and financial conditions:[1]

(i) *US stock market performance,* q_{0t}. We measure US stock market performance using the value-weighted excess return on CRSP firms incorporated in the US and listed on the NYSE, AMEX or NASDAQ.

(ii) *US Treasury market conditions,* r_t. We use the change in the five-year constant maturity Treasury (CMT) yield to capture expectations regarding US and global macroeconomic conditions.

(iii) *The variance risk premium (VRP),* v_t. Bollerslev et al. (2009) define the VRP as the difference between the one-month-ahead implied variance and a forecast of the realized variance over the same period. We compute the VRP as $VRP_t = VIX_t^2 - E\left[RV_{t+1}^{(22)}\right]$, where VIX_t^2 denotes the de-annualized squared VIX and $RV_t^{(22)}$ denotes the realized variance for the S&P 500 measured over the next 22 trading days as the sum of squared five-minute intraday returns. In light of the model comparison carried out by Bekaert and Hoerova (2014), we generate out-of-sample forecasts of the realized variance over our sample period using a heterogeneous autoregressive model supplemented with the squared VIX (model 8 in Bekaert and Hoerova) estimated using data for the pre-sample period 03-Jan-2000 to 30-Dec-2005.

[1]It is natural to use US data to proxy for global conditions, as the US sovereign is not in our model and there is ample evidence showing that US economic and financial conditions exert a strong global influence (e.g. Chudik and Fratzscher, 2011; Helbling et al., 2011; Longstaff et al., 2011).

(iv) *The Treasury term premium,* p_t. The term premium measures the excess yield required by investors to hold a long-term bond as opposed to a sequence of shorter-term bonds. It conveys valuable information on investors' time preferences and their expectations. We include the 5-year Treasury term premium derived from the five-factor no-arbitrage term structure model of Adrian et al. (2013).

(v) *US investment grade and high yield spreads,* i_t *and* h_t. To capture changes in the required rate of return on investment grade (IG) and high yield (HY) corporate bonds, we include both the IG (BBB-AAA) and HY (BB-BBB) spreads.

(vi) *The TED spread,* l_t. 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.

In addition, we include the following Europe-specific global factor:

(vii) *The Euribor-DeTBill spread,* l_t^e. We measure funding liquidity in Europe using the spread between the 3-month Euribor and the 3-month German Treasury bill yield.

Lastly, we include the following domestic factors for each sovereign:[2]

(viii) *Local stock market returns,* $q_{it},\ i = 1, 2, \ldots, m$. Augustin (2018) shows that domestic stock returns are the most significant domestic determinant of SCDS spreads, so we use local stock returns to capture country-specific financial risk.

(ix) *Bilateral spot exchange rate returns,* $s_{jt},\ j = 1, 2, \ldots, \tilde{m}$.[3] As the US dollar is the contract currency for our CDS data, we control for currency fluctuations against the dollar. To this end, we include the daily log-return on the bilateral spot exchange rate for each unique currency in our sample in units of foreign currency per USD.

[2]Local macroeconomic and political conditions have also been shown to affect CDS pricing (e.g. Remolona et al., 2008). However, these phenomena cannot be reliably measured at daily frequency, so we are unable to include them as factors directly. Nonetheless, their influence should be reflected in the local stock market return and USD bilateral exchange rate, which are included among our observed factors.

[3]All of the Euro area countries in our sample share a common currency over our sample period, so the total number of unique exchange rates against the US dollar is $\tilde{m} < m$.

Details of the data sources for the observed factors may be found in Appendix A. Our selection of global and local factors incorporates the majority of the explanatory variables considered by Longstaff et al. (2011).[4] By regressing the first difference of the ith SCDS spread on the contemporaneous first differences of the global factors and the domestic factors for sovereign i, we are able to isolate the idiosyncratic innovations in the credit risk of the ith sovereign, u_{it}. As expected, unit root testing and autocorrelation analysis indicates that our idiosyncratic credit risk measures are stationary and display little autocorrelation (results are available on request).

B. The Diebold–Yilmaz Framework

We apply the DY framework to characterize the comovement in the cross-section of idiosyncratic sovereign credit risk on a rolling sample basis. Suppose that the length of the full sample is $t = 1, 2, \ldots, T$ and that a window length of $\tau < T$ is selected (the selection of τ is discussed below). Given τ, we estimate VAR models on rolling samples of length τ, each of which is indexed by $r = 1, 2, \ldots, R$, where $R = T - \tau + 1$. The $p^{(r)}$-th order VAR model for $\boldsymbol{u}_t = (u_{1t}, u_{2t}, \ldots, u_{mt})'$ estimated on the r-th rolling sample is given by:

$$\boldsymbol{u}_t = \boldsymbol{\alpha}^{(r)} + \sum_{j=1}^{p^{(r)}} \boldsymbol{A}_j^{(r)} \boldsymbol{u}_{t-1} + \boldsymbol{\varepsilon}_t, \tag{2}$$

for $t = r, \ldots, r + \tau - 1$, where the regression residuals $\boldsymbol{\varepsilon}_t \sim N(0, \boldsymbol{\Omega}^{(r)})$ and the covariance matrix $\boldsymbol{\Omega}^{(r)}$ is positive-definite. We select the lag order for the VAR model in the r-th

[4]Our use of daily data necessitates the omission of some variables used by Longstaff et al. (2011) that are sampled at lower frequency, such as the price-earnings ratio for the S&P 100 index and equity and bond mutual fund flows. In addition, Longstaff et al. (2011) include cross-section averages of the SCDS spreads to proxy for regional and global factors, which can be viewed as a type of latent factor similar to those employed in the common correlated effects estimator developed by Pesaran (2006). Because we focus on observed factors, we do not include these cross-section average terms. Note that our selection of observed factors impounds the information contained in the factors used by other recent studies of sovereign credit risk transmission. For example, Caporin et al. (2018) control for the VRP of the Euro Stoxx 50 (which is correlated with the S&P 500 VRP), the Euribor-EONIA spread (which is correlated with the TED spread and the Euribor-DeTBill spread) and the first difference of the Euribor (which is related to the Euribor-DeTBill spread and is correlated with the change in the US five-year CMT yield).

rolling sample, $p^{(r)}$, by minimization of the Schwarz Information Criterion.

Suppressing the vector of intercepts for simplicity, the Wold representation of (2) is:

$$\boldsymbol{u}_t = \sum_{j=0}^{\infty} \boldsymbol{B}_j^{(r)} \boldsymbol{\varepsilon}_{t-j}^{(r)}, \tag{3}$$

where $\boldsymbol{B}_j^{(r)} = \boldsymbol{A}_1^{(r)} \boldsymbol{B}_{j-1}^{(r)} + \boldsymbol{A}_2^{(r)} \boldsymbol{B}_{j-2}^{(r)} + \ldots$, for $j = 1, 2, \ldots$ with $\boldsymbol{B}_0^{(r)} = \boldsymbol{I}_m$ and $\boldsymbol{B}_j^{(r)} = \boldsymbol{0}_m$ for $j < 0$. Following Pesaran and Shin (1998), the generalized forecast error variance decomposition (GVD) at horizon $h = 0, 1, 2, \ldots, H$ is defined as follows:

$$\vartheta_{i \leftarrow j}^{(r,h)} = \frac{\left(\sigma_{\varepsilon,jj}^{(r)}\right)^{-1} \sum_{\ell=0}^{h} \left(\mathbf{e}_i' \boldsymbol{B}_\ell^{(r)} \boldsymbol{\Omega}^{(r)} \mathbf{e}_j\right)^2}{\sum_{\ell=0}^{h} \mathbf{e}_i' \boldsymbol{B}_\ell^{(r)} \boldsymbol{\Omega}^{(r)} \boldsymbol{B}_\ell^{(r)\prime} \mathbf{e}_i}, \tag{4}$$

for $i, j = 1, \ldots, m$, where $\sigma_{\varepsilon,jj}^{(r)}$ is the (j,j)-th element of $\boldsymbol{\Omega}^{(r)}$ and $\mathbf{e}_i$ is a $m \times 1$ vector, with the i-th element equal to 1 and all others equal to 0. Consequently, $\vartheta_{i \leftarrow j}^{(r,h)}$ expresses the proportion of the h-step-ahead forecast error variance (FEV) of variable i that can be attributed to shocks in the equation for variable j in the r-th rolling sample. Because the GVD is scaled by the residual variance in each rolling sample, our results do not suffer from the volatility bias described by Forbes and Rigobon (2002).

DY were the first to show that GVDs can be used to form a weighted directed network. One may construct the so-called h-step-ahead 'connectedness matrix' for $\boldsymbol{u}_t$ as follows:

$$\boldsymbol{\vartheta}^{(r,h)} = \begin{pmatrix} \vartheta_{1\leftarrow 1}^{(r,h)} & \vartheta_{1\leftarrow 2}^{(r,h)} & \cdots & \vartheta_{1\leftarrow m}^{(r,h)} \\ \vartheta_{2\leftarrow 1}^{(r,h)} & \vartheta_{2\leftarrow 2}^{(r,h)} & \cdots & \vartheta_{2\leftarrow m}^{(r,h)} \\ \vdots & \vdots & \ddots & \vdots \\ \vartheta_{m\leftarrow 1}^{(r,h)} & \vartheta_{m\leftarrow 2}^{(r,h)} & \cdots & \vartheta_{m\leftarrow m}^{(r,h)} \end{pmatrix}. \tag{5}$$

For a given horizon, h, and rolling sample, r, $\underset{(m\times m)}{\boldsymbol{\vartheta}^{(r,h)}}$ characterizes the spillover of idiosyncratic sovereign credit risk among the m sovereigns in our sample.

In a reduced form VAR such as ours, $\boldsymbol{\Omega}^{(r)}$ is non-diagonal and the DY spillover measures do not have a structural interpretation but can be interpreted as a generalized directed type of correlation (Greenwood-Nimmo et al., 2020). The GVD is designed to accommodate this non-diagonality, although the variance shares will sum to more than one in this case. Diebold and Yilmaz (2014) tackle this scaling issue by normalizing the

elements of $\boldsymbol{\vartheta}^{(r,h)}$ as follows:

$$\theta_{i\leftarrow j}^{(r,h)} = \vartheta_{i\leftarrow j}^{(r,h)} / \sum_{j=1}^{m} \vartheta_{i\leftarrow j}^{(r,h)} \quad , \quad i,j = 1,2,\ldots,m, \tag{6}$$

such that we may write the $m \times m$ normalized h-step-ahead connectedness matrix as:

$$\boldsymbol{\theta}^{(r,h)} = \begin{pmatrix} \theta_{1\leftarrow 1}^{(r,h)} & \theta_{1\leftarrow 2}^{(r,h)} & \cdots & \theta_{1\leftarrow m}^{(r,h)} \\ \theta_{2\leftarrow 1}^{(r,h)} & \theta_{2\leftarrow 2}^{(r,h)} & \cdots & \theta_{2\leftarrow m}^{(r,h)} \\ \vdots & \vdots & \ddots & \vdots \\ \theta_{m\leftarrow 1}^{(r,h)} & \theta_{m\leftarrow 2}^{(r,h)} & \cdots & \theta_{m\leftarrow m}^{(r,h)} \end{pmatrix}. \tag{7}$$

C. Accounting for Compositional Changes

Unlike the simple case assumed above, the number of sovereigns in our sample changes through time (Bostanci and Yilmaz, 2020, face the same issue). On the one hand, data for more sovereigns becomes available over our sample period as the SCDS market broadens. Given that important sovereigns including the the UK and Switzerland are among the group without SCDS data at the start of our sample period, it would be desirable to allow our cross-section of sovereigns to grow as more data becomes available. On the other hand, and a more pressing concern for any analysis of the European sovereign debt crisis, there is a large gap in the Greek SCDS data, reporting of which ceases on 08-Mar-2012, when the spread is listed as 37,030.49 basis points.[5] Reporting of Greek SCDS data resumes on 10-Jun-2013. The gap in the Greek SCDS data is a result of the Greek debt crisis, where fears of an imminent credit event prompted a switch away from trading on a running spread in favour of points upfront trading.[6]

Faced with a sample of changing composition, we have two options: (i) to eliminate all those sovereigns for which we do not have data spanning the entire sample period;

[5]Such a high SCDS spread is a reflection of illiquidity, so we trim the Greek data back to 14-Feb-2012, the last observation for which the SCDS spread is less than 10,000 basis points.

[6]Points upfront trading is an alternative to operating with a running spread, in which the protection buyer pays the protection seller in an ongoing manner through the life of the contract. Points upfront trading involves the payment of the present value of the protection contract (or a fraction thereof) at the start of the trade. This eliminates the risk that the protection seller may be required to pay out before having received an income from the contract.

or (ii) to adapt our framework to accommodate compositional changes between rolling samples. We pursue the latter option to avoid excluding important countries from our analysis. Table 1 records how the number of sovereigns in our sample changes over our time, from 02-Jan-2006 to 27-Jul-2015. At the start of our sample, we have data for 20 sovereigns. The UK, Bulgaria and Switzerland enter our sample in March 2006, April 2006 and January 2009, respectively. Consequently, our largest sample includes 23 sovereigns.

— Insert Table 1 here —

In light of the compositional changes documented in Table 1, the dimension of the VAR model in the r-th rolling sample is not simply m but $m^{(r)}$. This changing dimension does not pose any problems for estimation of the rolling VAR models nor for computation of the connectedness matrix following the DY approach within each rolling sample. However, it does complicate comparisons across rolling samples, as easily demonstrated by a simple thought experiment. Suppose that we have data for 20 sovereigns in the r-th rolling sample and for 21 sovereigns in the $(r+1)$-th rolling sample. For simplicity, suppose also that spillover activity is perfectly uniform in both rolling samples — hence, every element of the connectedness matrix for the r-th rolling sample is equal to $(1/20) \times 100\% = 5\%$ while the equivalent figure for the $(r+1)$-th rolling sample is $(1/21) \times 100\% \approx 4.75\%$. In the absence of an appropriate adjustment, it would therefore appear that the average spillover is weaker in samples containing more sovereigns.

To ensure that our spillover measures remain comparable across rolling samples, we multiply every element of $\boldsymbol{\theta}^{(r,h)}$ by a scale factor such that:

$$s_{i\leftarrow j}^{(r,h)} = \frac{m^{(r)}}{\max\limits_{r\in\{1,\ldots,R\}}(m^{(r)})} \times \theta_{i\leftarrow j}^{(r,h)}, \tag{8}$$

This ensures that spillovers in every rolling sample are expressed on the same scale as in the sample with the largest cross-section of sovereigns. Returning to our thought experiment, after applying our adjustment to the r-th rolling sample, we obtain a value of $(20/21) \times (1/20) \times 100\% \approx 4.75\%$ and the connectedness matrices from the two rolling samples are now measured on the same scale. The final re-scaled connectedness matrix

is:

$$\mathbb{S}^{(r,h)} = \begin{pmatrix} s^{(r,h)}_{1\leftarrow 1} & s^{(r,h)}_{1\leftarrow 2} & \cdots & s^{(r,h)}_{1\leftarrow m^{(r)}} \\ s^{(r,h)}_{2\leftarrow 1} & s^{(r,h)}_{2\leftarrow 2} & \cdots & s^{(r,h)}_{2\leftarrow m^{(r)}} \\ \vdots & \vdots & \ddots & \vdots \\ s^{(r,h)}_{m^{(r)}\leftarrow 1} & s^{(r,h)}_{m^{(r)}\leftarrow 2} & \cdots & s^{(r,h)}_{m^{(r)}\leftarrow m^{(r)}} \end{pmatrix}. \tag{9}$$

The (i,i)-th diagonal element of $\mathbb{S}^{(r,h)}$ measures the proportion of the h-step-ahead FEV of idiosyncratic risk in sovereign i in the r-th rolling sample that can be attributed to shocks to affecting sovereign i itself. Meanwhile, the (i,j)-th off-diagonal element of $\mathbb{S}^{(r,h)}$ measures the proportion of the h-step-ahead FEV of idiosyncratic risk in sovereign i in the r-th rolling sample that can be attributed to shocks affecting sovereign j — that is, the spillover of idiosyncratic credit risk from sovereign j to i. It is the density of these $m^{(r)}\left(m^{(r)}-1\right)$ bilateral spillover effects that we now set out to characterize.

D. Summarizing the Bilateral Spillover Density

Given that the number of elements in $\mathbb{S}^{(r,h)}$ is quadratically increasing in $m^{(r)}$, reductive methods are needed to summarize variations in spillover activity across rolling samples. DY propose the aggregate spillover index, $S^{(r,h)}$:

$$S^{(r,h)} = \frac{\boldsymbol{e}'_{m^{(r)}}\mathbb{S}^{(r,h)}\boldsymbol{e}_{m^{(r)}} - \text{trace}\left(\mathbb{S}^{(r,h)}\right)}{\boldsymbol{e}'_{m^{(r)}}\mathbb{S}^{(r,h)}\boldsymbol{e}_{m^{(r)}}} \times 100\%, \tag{10}$$

where, due to the scale factor that we introduce in (8):

$$\boldsymbol{e}'_{m^{(r)}}\mathbb{S}^{(r,h)}\boldsymbol{e}_{m^{(r)}} = \frac{m^{(r)}}{\max\limits_{r\in\{1,\ldots,R\}}\left(m^{(r)}\right)} \times m^{(r)}.$$

Consequently, it follows that the DY spillover index is proportional to the mean bilateral spillover (i.e. the mean of the off-diagonal elements of $\mathbb{S}^{(r,h)}$). To see this, note that:

$$\mathbb{E}\left[s^{(r,h)}_{i\leftarrow j, j\neq i}\right] = \frac{\boldsymbol{e}'_{m^{(r)}}\mathbb{S}^{(r,h)}\boldsymbol{e}_{m^{(r)}} - \text{trace}\left(\mathbb{S}^{(r,h)}\right)}{m^{(r)}\left(m^{(r)}-1\right)} \equiv \frac{S^{(r,h)}}{m^{(r)}-1} \times \frac{m^{(r)}}{\max\limits_{r\in\{1,\ldots,R\}}\left(m^{(r)}\right)}. \tag{11}$$

For a symmetric distribution, the mean may represent an adequate summary statistic. However, reliance on the mean may be insufficient when the spillover density is asymmetric. Given that the GVD in (4) is defined as a ratio of quadratic forms, the off-diagonal

elements of $\mathbb{S}^{(r,h)}$ are likely to follow an asymmetric distribution. In Figure 1, based on simulations from a simple stationary and ergodic VAR(1) data generating process, we verify that this is the case. The spillover density is not only right-skewed but also markedly leptokurtic under a range of parameterizations (see Appendix B for a detailed summary of the simulation results). In light of this asymmetry, the mean bilateral spillover is likely to provide a poor summary of spillover activity. Measures designed to summarize spillover activity should take account of both the location and shape of the spillover density.

— Insert Figure 1 here —

We proceed by estimating the probability density function of the set of bilateral spillovers in rolling sample r, $\left\{s_{i\leftarrow j}^{(r,h)}\right\}_{i,j=1,i\neq j}^{m^{(r)}}$, by KDE. To this end, we define an $\eta \times 1$ vector of grids, $\boldsymbol{\nu} = (\nu_1, \nu_2, \ldots, \nu_\eta)'$. We follow the established practice and set $\eta = 1,024$. We then estimate the spillover density as:

$$\widehat{g}^{(r,h)}(\nu_l) = \frac{1}{b^{(r,h)}} \left(\frac{1}{m^{(r)}(m^{(r)}-1)} \right) \sum_{i,j=1;i\neq j}^{m^{(r)}} K\left(\frac{\nu_l - s_{i\leftarrow j}^{(r,h)}}{b^{(r,h)}} \right), \; l = 1, 2, \ldots, \eta, \tag{12}$$

where K is a chosen kernel and $b^{(r,h)}$ denotes the bandwidth. To ensure that $\widehat{g}^{(r,h)}(\nu_l)$ integrates to unity over the selected range of grid points, we normalize as follows:

$$\widehat{f}^{(r,h)}(\nu_l) = \frac{\widehat{g}^{(r,h)}(\nu_l)}{\text{RSUM}\left(\widehat{\boldsymbol{g}}^{(r,h)}\right)}, \tag{13}$$

where $\text{RSUM}\left(\widehat{\boldsymbol{g}}^{(r,h)}\right)$ is a numerical Riemann sum of $\widehat{\boldsymbol{g}}^{(r,h)} = \left(\widehat{g}^{(r,h)}(\nu_1), \ldots, \widehat{g}^{(r,h)}(\nu_\eta)\right)'$. Proceeding in this way, for a given window length (τ) and forecast horizon (h), we construct a sequence of $r = 1, 2, \ldots, R$ spillover densities, one for each rolling sample.

E. Summarizing the Evolution of Spillover Activity

In our framework, there is evidence of residual contagion in rolling sample r if the bilateral spillover of idiosyncratic sovereign risk exceeds normal levels. Consequently, determining if there is evidence of residual contagion requires: (i) a method to formally compare densities over time that is sufficiently simple to be implemented in a rolling-sample setting; and (ii) the selection of a benchmark density that is representative of normal times.

Our solution is to compare each of the R estimated spillover densities against a pre-GFC benchmark density, f_0, using the following divergence criteria:

$$D_{HN}\left(\widehat{f}^{(r,h)}, f_0\right) = \sup_{\nu}\left|\widehat{f}^{(r,h)}(\nu) - f_0(\nu)\right| \Big/ \sup_{\nu} f_0(\nu) \tag{14}$$

$$D_{KL}\left(\widehat{f}^{(r,h)}, f_0\right) = \int \widehat{f}^{(r,h)}(\nu) \ln\left\{\widehat{f}^{(r,h)}(\nu) \Big/ f_0(\nu)\right\} d\nu, \tag{15}$$

where $\widehat{f}^{(r,h)}$ is the estimated density under evaluation. D_{HN} is the Hilbert norm and D_{KL} is the Kullback Leibler Information Criterion (KLIC). Both measures are strictly non-negative and take the value zero only if $\widehat{f}^{(r,h)} = f_0$. Consequently, large values of the divergence criteria are evidence of residual contagion.

The choice of benchmark density is an important one. In principle, one may define a theoretical benchmark spillover density. This is an intriguing possibility but it is beyond the scope of this paper — we pursue a simpler approach. Given that we are particularly interested in tracing the evolution of the spillover density through time, our suggestion is to treat the estimated density from the first rolling sample as the benchmark — that is, to set $f_0 = \widehat{f}^{(1,h)}$. Our first rolling sample spans a period of relative tranquility in the financial markets during 2006 and so the resulting spillover density can be considered typical for 'normal' times. This proposition is validated in our empirical analysis, where the spillover densities in all rolling samples ending prior to the outbreak of the subprime crisis in mid-2007 are found to be qualitatively and quantitatively similar.

III. Results

To implement our technique, it is necessary to select an appropriate window length (τ), forecast horizon (h) and KDE implementation. Based on extensive sensitivity analysis, the results of which are summarized in Appendix C, we find that the properties of the estimated spillover density are highly robust to choices of $\tau \in \{200, 250, 300\}$ trading days, $h \in \{5, 10, 15\}$ trading days and to the use of several alternative KDEs. We therefore proceed by using the Gaussian kernel with the asymptotically optimal bandwidth and by setting $\tau = 250$ trading days and $h = 10$ trading days, without loss of generality.[7]

[7]In response to comments from an anonymous referee, we also conducted a range of tests to evaluate the robustness of our analysis to the use of alternative factor specifications and to the use of a shrinkage

A. Time-Varying Spillover Density

We begin by providing a high-level overview of the contours of the spillover density on a rolling sample basis in Figure 2. To verify that our KDE approach provides an accurate summary of spillover activity in each rolling sample, the figure provides a pair of contour plots tracking the evolution of the spillover histogram (panel (a)) and the estimated spillover density (panel (b)) over our sample period. The black line in each panel of the figure reports the mean bilateral spillover, which is proportional to the DY spillover index. The red line shows the median bilateral spillover. The difference between the mean and the median offers an intuitive spillover asymmetry measure (SAM), the value of which is shown by the cross-hatched area in each contour plot.[8] A positive (negative) value of the SAM indicates that the spillover density exhibits a right (left) skew. The dark gray shaded region shows the interquartile range of the spillover density, while the pale gray shading shows the region bounded by the 5th and 95th percentiles. Recall that the off-diagonal elements of the connectedness matrix are idiosyncratic risk spillovers, measured in percent. Therefore, the vertical axes in each panel of Figure 2 measure the strength of bilateral spillover effects in percent.[9] The similarity of the two contour plots indicates that our KDE implementation offers a good approximation of the spillover density.

— Insert Figure 2 here —

Figure 2 reveals that both the location and shape of the spillover density change markedly over time. Our benchmark spillover density is strongly asymmetric, with a peak close to zero and a long right tail. These characteristics of the spillover density are maintained until mid-2007, indicating that our benchmark density is representative of spillover activity in the period prior to the subprime crisis. Further direct evidence of the validity of our benchmark density can be obtained from the divergence criteria, which

and selection estimator to counter concerns of over-fitting. We found that our results are very robust to these changes. To conserve space, we do not report details of these exercises here, but they are available from the authors on request.

[8]Of course, there are many possible alternative formulations of the SAM. Note that, despite the similar terminology, our spillover asymmetry measure differs from that of Baruník et al. (2016), which is based on a signed decomposition of the variables entering the VAR model.

[9]The spillover density is bounded to the left by zero because GVDs are non-negative by definition.

take values very close to zero over this period, indicating the stability of the spillover density prior to the subprime crisis (the maximum values recorded by the Hilbert Norm and the KLIC prior to July 2007 are 0.064 and 0.068, respectively).

As the subprime crisis and the GFC unfold, the peak of the spillover density moves rightward and it becomes less skewed. This increase in aggregate spillover activity reflects an intensification of the average bilateral spillover of idiosyncratic credit risk among sovereigns rather than the influence of common sources of variation, for which our model explicitly controls. This is evidence of residual sovereign credit contagion during the GFC, a finding that adds to the evidence of contagion at this time put forth by Arghyrou and Kontonikas (2012) and Kalbaska and Gatkowski (2012), among others.

From late-2008 until mid-2011, the location of the spillover density remains above its benchmark level and is remarkably stable, providing evidence of continued strong spillovers of idiosyncratic credit risk – that is, evidence of continued contagion. A similar plateau in spillover activity starting in 2008 is documented by Claeys and Vašíček (2014) using data on bond yield spreads up to 2012. The stability of the mean bilateral spillover over this time is striking, particularly given that this was a period of profound instability associated with the European sovereign debt crisis. However, while the mean of the spillover density remains largely constant, its shape continues to change, with notable variation in tail mass. It is only after the Portuguese request for aid in April 2011 that the peak of the spillover density starts to gradually move leftward.

To develop richer intuition for the time-variation in the spillover density, Figure 3 presents several slices through the KDE-based contour plot in Figure 2(b), showing both the spillover CDF and PDF in a familiar format. The figure shows the benchmark spillover density in comparison to the spillover density obtained for eight other rolling samples ending on the day of the following events: Bernanke's warning of subprime risks in March 2007; the acquisition of Bear Stearns by JP Morgan in March 2008; the bankruptcy of Lehman Brothers in September 2008; the first Greek bailout as well as the Irish, Portuguese and Spanish bailouts, which occurred between April 2010 and June 2012; and Greece's failure to uphold its payment schedule to the IMF on 30-Jun-2015. Note that, even though our focus is on European sovereigns, we consider a range of global events to place our findings in their temporal context.

— Insert Figure 3 here —

When presented in this form, the extent of the asymmetry in the benchmark spillover density is readily apparent. Figure 3(a) reveals that the bulk of the probability mass is close to zero, with approximately two-thirds of all bilateral spillovers being no larger than 1% and approximately nine out of ten being no larger than 3%. The median bilateral spillover is just 0.51% but the right skew is sufficiently marked that the mean bilateral spillover takes a much higher value of 1.16%. Inspection of the connectedness matrix reveals that the subset of strong spillovers in the right tail of the spillover density in the benchmark sample occur among isolated country-pairs with close geopolitical relations, such as Hungary and Poland. The shape of the spillover density is similar shortly before the eruption of the subprime crisis. As of March 2007, when Bernanke (2007) warned of systemic risks related to government-sponsored mortgage enterprises, the median and mean bilateral spillovers are 0.63% and 1.26%, respectively.

An intensification of spillovers is evident at the time of JP Morgan's acquisition of Bear Stearns in March 2008. Six months later, when Lehman Brothers fails, the spillover density has shifted markedly to the right and the right skew has diminished. At this time, the median and mean bilateral spillovers are 2.74% and 3.22%, respectively, and roughly 1 in 10 bilateral spillovers are no larger than 1%, while a similar proportion are 6.5% or larger. The change in the shape of the spillover density is so profound that the mean bilateral spillover during the GFC is comparable to the upper decile of spillovers observed prior to the GFC. This is strong evidence of residual sovereign credit contagion. The magnitude of the response of the sovereign credit risk network to the GFC can be understood in light of the existing literature, which has demonstrated the nexus between financial crises and sovereign credit risk and which has shown that Europe was particularly vulnerable to financial shocks due to its institutional design (e.g. Reinhart and Rogoff, 2011; Lane, 2012; Acharya et al., 2014).

Spillovers of idiosyncratic sovereign credit risk remain at high levels at the time of the Greek request for financial assistance in April 2010, when the median and mean bilateral spillovers are 2.74% and 3.35%, respectively. This fragile environment is sustained throughout several subsequent sovereign bailouts and indicates a prolonged period of sovereign credit contagion that was of considerable concern to policymakers at the time (e.g. Constâncio, 2012). The mean bilateral spillover remains largely unchanged over the entire period spanning the Greek, Irish and Portuguese bailouts, but the higher-order

moments of the spillover density continue to evolve, with a notable jump in right tail mass at the time of the Portuguese request for assistance.

The spillover density shifts back to the left toward the end of our sample. Nonetheless, as of the end of June 2015, when Greece missed its scheduled IMF payment, both the median and mean bilateral spillovers exceed their benchmark levels, taking values of 1.46% and 2.55%, respectively. Although the falling average spillover intensity may be taken as evidence of the end of the contagion episode, it is important to realize that the strongest spillovers at this time occur among a well-defined cluster of sovereigns (the GIIPS), a setting that differs significantly from the isolated strong pairwise spillover effects observed in the benchmark sample. This is evidence of ongoing residual contagion among the GIIPS and is consistent with the widely discussed decoupling of the GIIPS sovereigns from the core European sovereigns at this time (e.g. Antonakakis and Vergos, 2013).[10]

The fact that the decoupling of the GIIPS primarily manifests in the right tail of the spillover density suggests that the shape of the spillover density may be of at least as much concern to policymakers as its location. Figure 4 offers several insights into the value of characterizing the entire spillover density in addition to its location. The figure records the location of bilateral credit risk spillovers among the GIIPS within the spillover density at the time of each of the events addressed in Figure 3, as well as two additional events for Greece — the agreement over the second Greek bailout on 21-Feb-2012 and the Greek debt-swap agreement on 09-Mar-2012. The figure shows that the bilateral spillovers among the GIIPS have typically remained well inside the upper quartile of the spillover density since the onset of the European debt crisis in late-2009, indicating a worrisome concentration of contagion risk in the European periphery (e.g. Alter and Schüler, 2012; Kalbaska and Gatkowski, 2012; Alter and Beyer, 2014; Aït-Sahalia et al., 2014).

— Insert Figure 4 here —

Let us first focus on an adverse event, Greece's departure from its IMF repayment

[10] A large literature has distinguished between the core and peripheral regions of the EU. A good example is Hale and Obstfeld (2016), who describe a pattern of heavy borrowing, domestic lending booms and rapid asset price inflation that became typical of the peripheral economies in the pre-crisis period, fueled by the low borrowing costs associated with EMU membership. Much of this borrowing was intermediated by financial institutions from the core economies, creating a financial boom in the core at the cost of elevated systemic exposure to peripheral risk.

schedule (labeled 'GRD'). Figure 4 reveals a surge in spillover activity among many of the GIIPS at this time that is not apparent using typical mean-based summary statistics like the DY spillover index. 2015 was a period of political unrest in Greece, with the election of the populist Syriza government raising the possibility of a disorderly Greek default or even the departure of Greece from the eurozone (Hodson, 2015). Either scenario would have had severe consequences for European sovereign bond markets. However, the instability in Greece posed particularly acute risks to peripheral sovereigns, especially those with populist anti-austerity movements, where the risk of political contagion (i.e. the risk of other peripheral states using the Greek precedent in an effort to renegotiate bailout terms subject to the threat that they may otherwise leave the eurozone) added to the threat of sovereign credit risk contagion. In its May 2015 Financial Stability Report, the European Central Bank (2015, pp. 11) notes that "[f]inancial market reactions to the developments in Greece have been muted to date, but in the absence of a quick agreement on structural implementation needs, the risk of an upward adjustment of the risk premia demanded on vulnerable euro area sovereigns could materialise".

Moving on to beneficial events, sovereign bailouts are motivated in large part by the desire to ameliorate spillovers from troubled sovereigns. As such, they may implicitly target the right tail of the spillover density. Comparing adjacent pairs of events in Figure 4 reveals that some bailouts achieved a sustained reduction in outward credit risk spillovers from the affected sovereign. Alter and Beyer (2014) also document reduced sovereign risk spillovers following EU/IMF bailouts, which they interpret as evidence of sovereign risk transmission onto the European Financial Stability Facility. For example, the Portuguese bailout (labeled 'PT') occurs when outward spillovers from Portugal are very strong, lying well inside the upper decile of the spillover density. By the time of the second Greek bailout (labeled 'GR2'), outward spillovers from Portugal to both Italy and Spain are weaker and lie more centrally within the spillover density, indicating a reduction in contagion risk. Similarly, the first Greek bailout appears to have reduced outward spillovers from Greece and the Irish bailout appears to have slightly reduced spillovers from Ireland to both Greece and Italy, if not significantly to the other GIIPS states.[11] Again, these effects are not easily discerned at the mean of the spillover density.

[11]No such inferences can be drawn for the Spanish bailout based on Figure 4, as the gap between the Spanish bailout (event 'ES') and the next event (Greece's missed IMF payment, 'GRD') is too large.

B. Summarizing Changes in the Shape of the Spillover Density

In any setting where a subset of strong bilateral spillovers may be of greater consequence to policymakers than the average bilateral spillover, techniques to monitor not just the location but the shape of the spillover density over time will enhance the utility of empirical network models for policy analysis. We use divergence criteria for this purpose. Figure 5 plots the rolling sample DY spillover index against the Hilbert Norm and the KLIC. To assist the reader, the timing of several major events is marked in the figure, including bank failures, sovereign crises, sovereign bailouts and ECB policy interventions. Furthermore, in Table 2, we provide a simple comparison of the magnitude of changes to the DY spillover index and the divergence criteria to accompany the figure. For each event marked in Figure 5, the table records the proportional change in the DY spillover index, the Hilbert Norm and the KLIC. In addition, in light of the time-varying asymmetry in the spillover density documented above, the table also reports the proportional change in the SAM. The analysis is based on the comparison of three time periods: (i) the average over the five days prior to the event, (ii) the day of the event; and (iii) the average over the five days after the event. To facilitate comparisons, we normalize such that the average value in the five days prior to the event is set to unity.

— Insert Table 2 and Figure 5 here —

Because the DY spillover index is proportional to the mean bilateral spillover reported in Figure 2, it exhibits the same prolonged period of relative stability during the European sovereign debt crisis. The proportional change in the spillover index is close to zero for many of the major events of the European debt crisis, yielding little signal of changes in the credit risk environment at this time. Therefore, in the absence of any information on the shape of the spillover density, one may naïvely conclude that there was no substantial change in spillover activity at this time. This would be an erroneous conclusion.

Because the divergence criteria fully exploit the informational content of the entire spillover density and not just its location, they both provide clear signals of changes in the credit risk environment when mean-based measures do not. For example, the abrupt jump in both divergence criteria after the dissolution of the Greek parliament in late-2009 reflects the increasing concentration of the spillover density around its mean relative to the benchmark density, with a marked contraction in the left tail indicating a notable

strengthening of the weakest spillovers in the system. The DY spillover index does not reflect this development and remains unchanged. In contrast, the SAM, Hilbert Norm and KLIC rise by 11%, 4% and 10% respectively over the following week. Similarly, at the time of the Greek request for aid in April 2010, the proportional change in the spillover index is -2%. By contrast, on the day of the request, the SAM jumps by 27%, indicating a large increase in right tail mass. Because this right-skewed shape is more similar to the benchmark spillover density, the Hilbert Norm falls by 4% and the KLIC falls by 10%.

However, the usefulness of the divergence criteria is best demonstrated toward the end of our sample, once the GIIPS have decoupled from the eurozone core. In May 2015, Janet Yellen's discussion of a possible interest rate rise in the US was met with a 5% proportional reduction in average spillovers coupled with large drops in both divergence criteria, as prospects for a normalization of monetary policy reassured investors. However, less than six weeks later, Greece missed its scheduled repayment to the IMF, instigating a surge in spillovers among the European periphery indicative of contagion among the GIIPS, as discussed above. The 4% proportional increase in the spillover index at this time masks the size of the effect on credit risk spillovers among the GIIPS. By contrast, the 8% reduction in the SAM on impact coupled with jumps of 36% in the Hilbert Norm and 28% in the KLIC, draws immediate attention to the tails of the spillover density.

This exercise has two important implications. First, the stability of the mean bilateral spillover from one period to the next should not be interpreted as evidence that spillover activity has not changed. Second, changes in the mean bilateral spillover are often accompanied by oppositely signed changes in the skewness of the spillover density, indicating that the behavior of the strongest spillovers in the system may be dissimilar to the behavior of the average spillover. This implies that density-based summary statistics such as the SAM, Hilbert Norm and KLIC can identify episodes of contagion even when it is isolated to a subset of sovereigns and is not strongly reflected in measures of the average intensity of bilateral spillovers.

C. What Explains the Shape of the Spillover Density?

So far, we have documented variations in the location and shape of the spillover density that are indicative of residual contagion during and after the GFC. However, the question of how changes in spillover activity are distributed among the edges in the network remains

open. This is an important question, because changes in the shape of the spillover density reflect underlying changes in the intensity of bilateral spillovers among sovereign pairs. Understanding what drives these changes in pairwise spillovers is therefore crucial to understanding the evolution of the spillover density over time.

To develop intuition on this issue, we begin by analyzing and comparing the structure of the sovereign credit risk network in two rolling samples. Figure 6(a) shows the spillover density overlaid on the spillover histogram for the benchmark sample, from 02-Jan-2006 to 15-Dec-2006. Figure 6(b) shows the spillover density and histogram for the rolling sample covering the period 22-Apr-2010 to 06-Apr-2011, the end of which corresponds to the Portuguese request for activation of the aid mechanism in the midst of the sovereign debt crisis.[12] To investigate how the distribution of spillovers among the edges in the network changes between these two samples, Figure 7 plots the corresponding ten-days-ahead connectedness matrices, $S^{(r,10)}$, as heatmaps. To facilitate their interpretation, each heatmap has been clustered using an agglomerative single linkage algorithm.[13] Consequently, the order of the sovereigns differs between panels (a) and (b) of Figure 7. Furthermore, due to changes in the cross-sectional composition of our sample over time, 20 sovereigns are included in Figure 7(a) and 23 in Figure 7(b), although the reported spillover effects are directly comparable due to the scale factor defined in (8).

— Insert Figures 6 & 7 here —

Figure 7(a) reveals that, with the exception of the bidirectional spillovers between Poland and Hungary, the bidirectional spillovers between the Netherlands and Austria and the spillover from Norway to Sweden, the remaining spillovers are mostly negligible in the benchmark sample. Figure 7(b) presents a stark contrast. We have already seen that spillover activity has intensified markedly by the time of the Portuguese request for assistance but Figure 7(b) reveals that increases in spillover intensity are not uniformly

[12]Our choice of the Portuguese bailout sample as a comparison case is essentially arbitrary — we could use any rolling sample during the sovereign debt crisis period and the principal implications that we highlight below would still hold. Detailed results for other samples are available on request.

[13]This is an iterative 'bottom-up' algorithm. To start, every vertex is considered as a cluster in its own right. In each iteration, the pair of clusters that are closest to one-another are merged into a new cluster. The closeness of two clusters is measured by the single strongest connection that exists between members of those clusters. The algorithm terminates when all vertices are included in a single cluster.

distributed over the edges in the network. At one extreme, Switzerland and Iceland are very weakly connected to most other sovereigns. At the other extreme, the European core and peripheral sovereigns are now strongly connected to one-another. The extent of the regional variation in spillover activity revealed by Figure 7 is striking. The clustering algorithm draws attention to distinct groups of sovereigns, which can be seen as blocks along the prime diagonal of the connectedness matrices. Clusters of EU core and peripheral sovereigns can be discerned and a cluster of Eastern sovereigns is particularly well-defined. A natural explanation for this clustering is that the intensity of bilateral credit risk spillovers — and, by extension, the shape of the spillover density — is related to the underlying economic and financial linkages between countries, such as trade flows and financial exposures, which themselves display regional variation. A large literature has documented the relevance of trade linkages, spatial proximity, the level of economic development, fiscal sustainability and European-specific factors in explaining episodes of contagion (e.g. Glick and Rose, 1999; Van Rijckeghem and Weder, 2001; Caramazza et al., 2004; Kali and Reyes, 2010; Forbes, 2013; Bekaert et al., 2014; Fry-McKibbin et al., 2014; Ters and Urban, 2018; Niemann and Pichler, 2020; Hsiao and Morley, 2022).

To investigate this possibility, we estimate an unbalanced panel regression model in which the bilateral credit risk spillover between country-pairs is regressed on a set of indicators of real and financial bilateral linkages, while controlling for several observable characteristics of both countries, as well as country-pair and time fixed effects. Specifically, we model the bilateral credit risk spillover from country i to country j as a function of the following explanatory variables:

(i) *Bilateral portfolio investment exposure.* The exposure of country j to bilateral portfolio investments with respect to country i is captured by the sum of country i's declared holdings of assets from country j and country j's declared holdings of assets from country i, expressed as a percentage of the GDP of country j.

(ii) *Bilateral portfolio investment flow.* The magnitude of the gross bilateral portfolio investment flow between countries i and j from the perspective of country j is captured by the quarterly change in the sum of country i's declared holdings of assets from country j and country j's declared holdings of assets from country i, expressed as a percentage of the GDP of country j.

(iii) *Bilateral trade exposure.* The exposure of country j to bilateral trade with country i is captured by the sum of exports and imports between countries i and j expressed as a percentage of the GDP of country j.

(iv) *Structural budget balance.* The cyclically-adjusted fiscal position of sovereign i (j) is captured by its structural budget balance quoted as a percentage of potential GDP.

(v) *Real GDP growth.* The annual growth rate of real GDP in country i (j) relative to the same quarter of the previous year, expressed in percent.

(vi) *CPI inflation.* The annual rate of CPI inflation in country i (j) relative to the same quarter of the previous year, expressed in percent.

These variables capture a great deal of information about real and financial linkages among sovereigns, as well as their macroeconomic performance. The sampling frequency of these series varies, from monthly (in the case of trade data, for example) to annual (e.g. the structural budget balance). To avoid excessive interpolation of the lower frequency series, we work at quarterly frequency. To this end, we convert our daily spillover measures to quarterly frequency using the period average value. Full details of the construction of the panel dataset may be found in Appendix A.

Our regression results are reported in Table 4. Because our goal is to understand what drives bilateral credit risk spillovers, we are primarily concerned with the coefficients on the three country-pair-specific variables listed at the top of the table.[14] Two key results emerge. First, after controlling for conditions in the source and recipient countries, an increase in the bilateral portfolio exposure between countries i and j is associated with an increase in the intensity of the idiosyncratic sovereign credit risk spillover from country i to country j. Meanwhile, the gross bilateral portfolio investment flow between countries

[14]Even though our primary focus is on the country-pair-specific variables, the estimated parameters on the country-specific controls in Table 4 reveal some interesting findings. For example, sovereigns with larger structural budget deficits (indicating an impaired fiscal position) tend to generate stronger outward credit risk spillovers and are also more vulnerable to inward spillovers. The importance of debt in the study of contagion has been emphasized by Niemann and Pichler (2020) and Hsiao and Morley (2022), for example. Faster growth rates in both the source and recipient countries are also associated with stronger spillovers, at least after controlling for other observable characteristics and for country-pair and period fixed effects.

i and j is negatively related to the spillover intensity. These findings indicate that closer financial linkages provide a channel for contagion. Several studies including Forbes (2013) and Fry-McKibbin et al. (2014) obtain similar results for a variety of crises.

Our results also suggest that retrenchments in bilateral portfolio investment can provide a channel for contagion. In extreme cases, such retrenchments may reflect fire sales, where distressed investors seeking to sell assets are obliged to accept prices below fair value. In this vein, Raddatz and Schmukler (2012) have shown that capital flows from mutual funds are procyclical and serve to expose countries to foreign shocks. Second, we find no significant effect of the bilateral trade exposure on sovereign credit risk spillovers. This leads us to conclude that it is financial exposures among countries rather than bilateral trade relations that contribute to the transmission of idiosyncratic sovereign credit risk across borders, echoing similar findings from Van Rijckeghem and Weder (2001), Caramazza et al. (2004) and Fry-McKibbin et al. (2014) across a range of recent crises.

— Insert Table 4 here —

D. Monitoring Systemic Stress

As a final exercise, we note that the speed with which the location and shape of the spillover density adjust in response to systemic events raises the possibility that the spillover density can be used to monitor systemic stress in the market for sovereign debt.[15] To investigate this possibility, we estimate a set of auxiliary models in which we regress a selection of our network statistics on two popular measures of systemic stress published by the ECB. The first is the Composite Indicator of Systemic Stress (CISS) described by Holló et al. (2012), which is constructed by aggregating 15 different measures of financial stress in a manner that accounts for time-variation in their cross-correlations. The

[15]The use of density information for monitoring purposes is related to the literature on the use of quantile regressions in the analysis of systemic stress (e.g. Caporin et al., 2018; Chavleishvili and Manganelli, 2019; Ando et al., 2022). The main conceptual distinction between our approach and quantile regression methods is that we seek to characterise the distribution of bilateral spillover effects estimated from a conditional mean model and observe how it changes over rolling samples, while quantile regression allows the coefficients of a model to differ across the quantiles of the conditional distribution. Exploring the complementarity between our approach and quantile regression methods represents an interesting avenue for continuing research.

second is its sovereign counterpart, SovCISS, developed by Garcia-de-Andoain and Kremer (2018), which applies a similar methodology to data from European sovereign bond markets. Both CISS and SovCISS are bounded between zero and one, with higher values indicating greater stress. CISS is reported weekly, while SovCISS is reported monthly. Therefore, in regressions involving CISS (SovCISS), we convert our network statistics to weekly (monthly) frequency by taking the period average. The results are presented in Table 3.

— Insert Table 3 here —

First, consider the top two rows of the table, which relate to the rolling sample mean and median of the spillover density. As expected, both the mean and median are positively correlated with CISS (0.50 and 0.51, respectively) and strongly positively correlated with SovCISS (0.71 and 0.72, respectively).[16] However, the median is considerably more responsive to changes in both CISS and SovCISS than the mean. A one unit increase in CISS is associated with a 1.50% increase in the mean and a 2.01% increase in the median. Likewise, a one unit increase in SovCISS is associated with a 2.84% increase in the mean and a 3.32% increase in the median.

The greater sensitivity of the median than the mean to changes in systemic stress can be explained by the asymmetry of the spillover density. The SAM is negatively correlated with both CISS and SovCISS (-0.46 and -0.30, respectively) and exhibits a strong negative contemporaneous association with both — the relevant slope coefficients in Table 3 are -0.55 (CISS) and -0.48 (SovCISS). On average, an increase in systemic stress is associated with a reduction in the right skew of the spillover density, as strong spillover effects that would be confined to the right tail of the spillover density in normal times become more widespread in times of stress. Furthermore, the skewness and kurtosis of the spillover density exhibit a strong negative contemporaneous association with CISS and SovCISS, indicating that the spillover density does not only move right and become more symmetric but that it also becomes more concentrated around its mean in periods of systemic stress.

The systematic responses of both the location and shape of the spillover density to variations in systemic stress indicate that the network statistics that we develop can

[16]The fact that our network statistics are more strongly correlated with SovCISS than CISS is a natural finding given that our model focuses on spillovers of idiosyncratic sovereign credit risk as opposed to idiosyncratic risk spillovers among a broader set of asset markets.

be used to monitor systemic stress. Crucially, our model can be used to update these measures at daily frequency, providing a more timely indication of systemic stress than either the CISS or the SovCISS. Access to such timely indications of systemic stress would be extremely useful for policymakers in the midst of a fast-moving crisis and may contribute to the formulation of better targeted stabilization policies.

IV. Concluding Remarks

In this paper, we develop and implement an innovative extension on the DY framework for network analysis to look for evidence of residual sovereign credit contagion in Europe between 2006 and 2015. By studying the comovement of the cross-section of idiosyncratic sovereign credit risk relative to a pre-crisis benchmark, we find evidence of residual contagion, starting with the outbreak of the GFC and continuing over the period of the European debt crisis. Furthermore, we find evidence that contagion becomes increasingly localized among the GIIPS sovereigns as the European debt crisis progresses, which is consistent with evidence of a decoupling of the GIIPS from the core of the eurozone.

A distinctive feature of our framework is that, unlike prior studies that have used DY networks to study sovereign credit contagion (e.g. Claeys and Vašíček, 2014), it fully exploits information on both the location and the shape of the density of bilateral credit risk spillovers. In our efforts to understand what determines the shape of the spillover density, we show that international portfolio investment positions expose sovereigns to foreign shocks and that the retrenchment of international investments represents a channel by which shocks to credit risk can propagate internationally. This finding is consistent with a body of evidence that emphasizes the role of financial linkages and portfolio investment in episodes of contagion (e.g. Van Rijckeghem and Weder, 2001; Caramazza et al., 2004; Forbes, 2013; Fry-McKibbin et al., 2014) and is consistent with the destabilizing influence of mutual funds documented by Raddatz and Schmukler (2012).

Finally, we demonstrate that the location and shape of the spillover density vary systematically with changes in systemic stress, as measured by the CISS and SovCISS indices published by the ECB. The strength of the contemporaneous association between the moments of the spillover density and the CISS and SovCISS indicates that the spillover density can be used as a tool to monitor systemic stress in the market for sovereign debt

at daily frequency, providing a valuable and timely supplement to the ECB's weekly CISS release and its monthly SovCISS release.

In conclusion, we note three avenues for continuing research. First, the factor dependence of the SCDS spreads could be handled within the model rather than by defactoring prior to estimation. This would provide an avenue to study additional channels of sovereign credit contagion in future work.[17] Second, the development of a theoretically-grounded benchmark spillover density against which the empirical spillover density could be compared would open intriguing possibilities for counterfactual analysis. Finally, our technique has obvious applications to other markets. The foreign exchange market provides a natural example, where the forced unwinding of carry trades during periods of illiquidity described by Brunnermeier et al. (2009) is likely to create a concentration of strong spillovers among a small group of high-yield currencies.

References

Acharya, V. V., I. Drechsler, and P. Schnabl (2014): "A Pyrrhic Victory? Bank Bailouts and Sovereign Credit Risk," *Journal of Finance*, 69, 2689–2739.

Adrian, T., R. K. Crump, and E. Moench (2013): "Pricing the Term Structure with Linear Regressions," *Journal of Financial Economics*, 110, 110–138.

Aït-Sahalia, Y., R. J. Laeven, and L. Pelizzon (2014): "Mutual Excitation in Eurozone Sovereign CDS," *Journal of Econometrics*, 183, 151–167.

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

Alter, A. and Y. S. Schüler (2012): "Credit Spread Interdependencies of European States and Banks During the Financial Crisis," *Journal of Banking and Finance*, 36, 3444–3468.

Ando, T., M. Greenwood-Nimmo, and Y. Shin (2022): "Quantile Connectedness: Modeling Tail Behavior in the Topology of Financial Networks," *Management Science*, 68, 2377–3174.

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

[17] Some indirect evidence on the role of the common factors in the transmission of sovereign credit risk may be found in Appendix D, which contrasts the results of our model using defactored data against an alternative model using raw SCDS spreads, which are subject to the influence of common factors.

Antonakakis, N. and K. Vergos (2013): "Sovereign Bond Yield Spillovers in the Euro Zone during the Financial and Debt Crisis," *Journal of International Financial Markets, Institutions and Money*, 26, 258–272.

Arghyrou, M. G. and A. Kontonikas (2012): "The EMU Sovereign-debt Crisis: Fundamentals, Expectations and Contagion," *Journal of International Financial Markets, Institutions and Money*, 22, 658–677.

Augustin, P. (2018): "The Term Structure of CDS Spreads and Sovereign Credit Risk," *Journal of Monetary Economics*, 96, 53–76.

Bai, J. and S.-J. Wei (2017): "Property Rights and CDS Spreads: When Is There a Strong Transfer Risk from the Sovereigns to the Corporates?" *Quarterly Journal of Finance*, 7, 1–36.

Ballester, L., B. Casu, and A. González-Urteaga (2016): "Bank Fragility and Contagion: Evidence from the Bank CDS Market," *Journal of Empirical Finance*, 38, 394–416.

Baruník, J., E. Kočenda, and L. Vácha (2016): "Asymmetric Connectedness on the U.S. Stock Market: Bad and Good Volatility Spillovers," *Journal of Financial Markets*, 27, 55–78.

Beirne, J. and M. Fratzscher (2013): "The Pricing of Sovereign Risk and Contagion during the European Sovereign Debt Crisis," *Journal of International Money and Finance*, 34, 60–82.

Bekaert, G., M. Ehrmann, M. Fratzscher, and A. Mehl (2014): "The Global Crisis and Equity Market Contagion," *Journal of Finance*, 69, 2597–2649.

Bekaert, G. and M. Hoerova (2014): "The VIX, the Variance Premium and Stock Market Volatility," *Journal of Econometrics*, 183, 181–192.

Bernanke, B. S. (2007): "GSE Portfolios, Systemic Risk, and Affordable Housing," Speech before the Independent Community Bankers of America's Annual Convention and Techworld, Honolulu, Hawaii.

Blanco, R., S. Brennan, and I. Marsh (2005): "An Empirical Analysis of the Dynamic Relationship between Investment-Grade Bonds and Credit Default Swaps," *Journal of Finance*, 60, 2255–2281.

Bollerslev, T., G. Tauchen, and H. Zhou (2009): "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 22, 4463–4492.

Bostanci, G. and K. Yilmaz (2020): "How Connected is the Global Sovereign Credit Risk Network?" *Journal of Banking and Finance*, forthcoming.

Botev, Z., J. Grotowski, and D. Kroese (2010): "Kernel Density Estimation via Diffusion," *Annals of Statistics*, 38, 2916–2957.

Brunnermeier, M. K., S. Nagel, and L. H. Pedersen (2009): "Carry Trades and Currency Crashes," *NBER Macroeconomics Annual 2008*, 23, 313–347.

Buse, R. and M. Schienle (2019): "Measuring Connectedness of Euro Area Sovereign Risk," *International Journal of Forecasting*, 35, 25–44.

Caporin, M., L. Pelizzon, F. Ravazzolo, and R. Rigobon (2018): "Measuring Sovereign Contagion in Europe," *Journal of Financial Stability*, 34, 150–181.

Caramazza, F., L. Ricci, and R. Salgado (2004): "International Financial Contagion in Currency Crises," *Journal of International Money and Finance*, 23, 51–70.

Chavleishvili, S. and S. Manganelli (2019): "Forecasting and Stress Testing with Quantile Vector Autoregression," Working Paper 2330, European Central Bank, Frankfurt.

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

Claeys, P. and B. Vašíček (2014): "Measuring Bilateral Spillover and Testing Contagion on Sovereign Bond Markets in Europe," *Journal of Banking and Finance*, 46, 151–165.

Constâncio, V. (2012): "Contagion and the European Debt Crisis," *Banque de France Financial Stability Review*, 16, 109–121.

Diebold, F. and K. Yilmaz (2009): "Measuring Financial Asset Return and Volatility Spillovers, with Application to Global Equity Markets," *Economic Journal*, 119, 158–171.

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

European Central Bank (2015): "Financial Stability Review, May 2015," European Central Bank, Frankfurt.

Forbes, K. (2013): "The "Big C": Identifying and Mitigating Contagion," in *The Changing Policy Landscape. 2012 Jackson Hole Symposium hosted by the Federal Reserve Bank of Kansas City*, 23–87.

Forbes, K. and R. Rigobon (2002): "No Contagion, Only Interdependence: Measuring Stock Market Comovements," *Journal of Finance*, 57, 2223–2261.

Fry-McKibbin, R., C. Y.-L. Hsiao, and C. Tang (2014): "Contagion and Global Financial Crises: Lessons from Nine Crisis Episodes," *Open Economies Review*, 25, 521–570.

Garcia-de-Andoain, C. and M. Kremer (2018): "Beyond Spreads: Measuring Sovereign Market Stress in the Euro Area," Working Paper 2185, European Central Bank, Frankfurt.

Giordano, R., M. Pericoli, and P. Tommasino (2013): "Pure or Wake-up-call Contagion? Another Look at the EMU Sovereign Debt Crisis," *International Finance*, 16, 131–160.

Glick, R. and A. K. Rose (1999): "Contagion and Trade: Why are Currency Crises Regional?" *Journal of International Money and Finance*, 18, 603–617.

Gómez-Puig, M. and S. Sosvilla-Rivero (2016): "Causes and Hazards of the Euro Area Sovereign Debt Crisis: Pure and Fundamentals-based Contagion," *Economic Modelling*, 56, 133–147.

Greenwood-Nimmo, M. J., V. H. Nguyen, and B. Rafferty (2016): "Risk and Return Spillovers among the G10 Currencies," *Journal of Financial Markets*, 31, 43–62.

Greenwood-Nimmo, M. J., V. H. Nguyen, and Y. Shin (2020): "Measuring the Connectedness of the Global Economy," *International Journal of Forecasting*, in press.

Hale, G. and M. Obstfeld (2016): "The Euro and the Geography of International Debt Flows," *Journal of the European Economic Association*, 14, 115–144.

Helbling, T., R. Huidrom, M. Kose, and C. Otrok (2011): "Do Credit Shocks Matter? A Global Perspective," *European Economic Review*, 55, 340–353.

Hodson, D. (2015): "Eurozone Governance: Deflation, Grexit 2.0 and the Second Coming of Jean-Claude Juncker," *Journal of Common Market Studies*, 53, 144–161.

Holló, D., M. Kremer, and M. Lo Duca (2012): "CISS – A Composite Indicator of Systemic Stress in the Financial System," Working Paper 1426, European Central Bank, Frankfurt.

Hsiao, C. Y.-L. and J. Morley (2022): "Debt and Financial Market Contagion," *Empirical Economics*, 62, 15991648.

Kalbaska, A. and M. Gatkowski (2012): "Eurozone Sovereign Contagion: Evidence from the CDS Market (2005–2010)," *Journal of Economic Behavior and Organisation*, 83, 657–673.

Kali, R. and J. Reyes (2010): "Financial Contagion on the International Trade Network," *Economic Inquiry*, 48, 1072–1101.

Lane, P. R. (2012): "The European Sovereign Debt Crisis," *Journal of Economic Perspectives*, 26, 49–67.

Longstaff, F., J. Pan, L. Pedersen, and K. Singleton (2011): "How Sovereign Is Sovereign Credit Risk?" *American Economic Journal: Macroeconomics*, 3, 75–103.

Montfort, A. and J.-P. Renne (2014): "Decomposing Euro-Area Sovereign Spreads: Credit and Liquidity Risks," *Review of Finance*, 18, 2103–2151.

Niemann, S. and P. Pichler (2020): "Optimal Fiscal Policy and Sovereign Debt Crises," *Review of Economic Dynamics*, 37, 234–254.

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

Pesaran, M. and Y. Shin (1998): "Generalized Impulse Response Analysis in Linear Multivariate Models," *Economics Letters*, 58, 17–29.

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

Raddatz, C. and S. Schmukler (2012): "On the International Transmission of Shocks: Micro-Evidence from Mutual Fund Portfolio," *Journal of International Economics*, 88, 357–374.

Reinhart, C. M. and K. S. Rogoff (2011): "From Financial Crash to Debt Crisis," *American Economic Review*, 101, 1676–1706.

Remolona, E., M. Scatigna, and E. Wu (2008): "The Dynamic Pricing of Sovereign Risk in Emerging Markets: Fundamentals and Risk Aversion," *Journal of Fixed Income*, 17, 57–71.

Silverman, B. W. (1986): *Density Estimation for Statistics and Data Analysis*, London: Chapman and Hall.

Ters, K. and J. Urban (2018): "Intraday Dynamics of Credit Risk Contagion Before and During the Euro Area Sovereign Debt Crisis: Evidence from Central Europe," *International Review of Economics Finance*, 54, 123–142.

Van Rijckeghem, C. and B. Weder (2001): "Sources of Contagion: Is it Finance or Trade?" *Journal of International Economics*, 54, 293–308.

Sovereign	ISO Code	Sample Begins	5-year CDS Spreads			Stock Return			Observed Factors			
			Mean	Min	Max	Mean	Min	Max		Mean	Min	Max
Greece	GR	02-Jan-06	4.65	-2455.84	1674.74	-6.11	-1366.89	1343.11	GBP	0.39	-447.45	391.82
Austria	AT	02-Jan-06	0.01	-27.49	44.28	-1.61	-1025.26	1202.10	BGN	0.24	-461.26	384.84
Belgium	BE	02-Jan-06	0.01	-57.38	35.11	0.17	-831.93	922.13	CZK	-0.05	-521.92	553.98
Czech Rep.	CZ	02-Jan-06	0.02	-42.31	46.03	-0.28	-1548.05	1463.79	DKK	0.24	-462.34	385.56
Denmark	DK	02-Jan-06	0.01	-15.77	17.60	3.84	-1147.58	951.98	EUR	0.26	-373.33	278.48
Finland	FI	02-Jan-06	0.01	-8.48	11.82	1.60	-760.64	946.20	HUF	1.06	-519.98	630.47
France	FR	02-Jan-06	0.01	-29.83	23.02	0.18	-947.15	1059.46	ISK	2.98	-1322.93	1341.53
Germany	DE	02-Jan-06	0.00	-13.36	11.75	2.86	-743.35	1079.75	LTL	0.24	-461.76	383.76
Hungary	HU	02-Jan-06	0.05	-83.68	125.23	-0.06	-1355.28	1475.57	NOK	0.74	-645.81	501.51
Iceland	IS	02-Jan-06	0.05	-126.49	279.82	-5.47	-10960.25	507.83	PLN	0.53	-669.73	569.62
Ireland	IE	02-Jan-06	0.02	-152.45	113.79	-0.61	-1396.36	973.31	RUR	2.91	-1552.30	1426.83
Italy	IT	02-Jan-06	0.04	-74.04	71.80	-1.80	-859.81	1087.69	SEK	0.26	-554.74	354.11
Lithuania	LT	02-Jan-06	0.03	-76.87	108.38	0.41	-1193.78	1100.15	CHF	-1.28	-1141.89	847.48
Norway	NO	02-Jan-06	0.00	-11.88	12.80	1.28	-1127.56	1101.57	Term Prem.	0.01	-33.04	20.31
Poland	PL	02-Jan-06	0.02	-63.29	55.34	1.50	-828.88	608.37	RmRf	-0.01	-1221.00	1314.00
Portugal	PT	02-Jan-06	0.06	-167.25	170.66	-1.67	-1037.92	1019.59	CMT	-0.11	-46.00	34.00
Russia	RU	02-Jan-06	0.11	-168.80	206.86	1.80	-2065.71	2522.61	IG	0.01	-152.00	153.00
Spain	ES	02-Jan-06	0.03	-69.60	59.21	0.15	-958.59	1348.36	HY	0.01	-57.00	50.00
Sweden	SE	02-Jan-06	-0.03	-79.33	65.77	2.00	-751.27	986.50	VRP	0.00	-188.98	174.52
Netherlands	NL	02-Jan-06	0.01	-13.43	25.69	0.39	-959.03	1002.83	TED	-0.01	-80.00	99.63
UK	UK	21-Mar-06	0.01	-16.74	17.48	0.34	-926.56	938.43	De-TBill	0.01	-65.80	73.60
Bulgaria	BG	17-Apr-06	0.06	-86.22	91.81	-2.58	-1136.00	729.24				
Switzerland	CH	19-Jan-09	-0.04	-43.92	50.20	3.09	-907.03	490.27				

NOTE: Currencies are identified as follows: British Pound (GBP); Bulgarian Lev (BGN); Czech Koruna (CZK); Danish Krone (DKK); Euro (EUR); Hungarian Forint (HUF); Lithuanian Lita (LTL); Norwegian Krone (NOK); Polish Złoty (PLN); Russian Rouble (RUR); Swedish Krona (SEK); and Swiss Franc (CHF). The other observed common factors are abbreviated as follows: term premium (Term Pm.); US stock market excess return (Rm-Rf); US 5-year CMT Treasury yield (CMT); investment grade spread (IG); high yield spread (HY); variance risk premium (VRP); TED spread (TED); and the Euribor-DeTBill spread (EU TED). Descriptive statistics for the log-return on the domestic stock market and the log-return on the bilateral spot exchange rate against the US dollar are reported in basis points. All exchange rates are expressed in units of foreign currency per US dollar. Descriptive statistics for the VRP are reported in variance units. Descriptive statistics for the first difference of all other variables are reported in basis points.

Table 1: Cross-Sectional Composition and Descriptive Statistics for the Estimation Sample

Event	DY	SAM	HN	KLIC	Event	DY	SAM	HN	KLIC
5 days prior	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
A	1.01	0.95	0.89	0.95	K	1.01	0.82	1.07	1.09
5 days after	1.00	0.97	0.90	0.95		1.01	0.84	1.07	1.08
	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
B	1.01	0.91	1.09	1.14	L	1.00	0.93	1.02	1.01
	1.28	0.77	1.50	1.90		1.00	1.01	1.02	1.02
	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
C	1.00	0.94	1.03	1.01	M	1.00	1.32	1.00	1.00
	1.01	0.86	1.10	1.21		1.01	1.11	1.02	1.04
	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
D	1.05	0.87	1.12	1.40	N	1.00	1.43	0.94	0.92
	1.04	0.94	1.08	1.31		1.00	1.39	0.90	0.88
	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
E	1.00	1.39	0.97	0.95	O	0.92	1.02	0.63	0.57
	1.02	0.59	1.07	1.08		0.92	0.85	0.74	0.68
	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
F	1.00	0.97	1.00	0.99	P	1.00	1.05	0.96	0.95
	1.00	1.11	1.04	1.10		1.00	1.02	0.99	0.98
	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
G	1.01	0.95	1.05	1.13	Q	0.98	1.21	0.98	0.95
	1.00	1.57	1.04	1.09		0.94	1.13	0.99	0.95
	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
H	0.98	1.27	0.96	0.90	R	1.01	1.04	1.05	1.04
	0.99	1.11	0.92	0.86		1.10	0.97	1.29	1.42
	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
I	1.00	1.04	0.99	0.98	S	0.95	1.12	0.62	0.73
	0.99	0.97	0.98	0.91		0.95	1.11	0.67	0.75
	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
J	1.00	0.98	0.98	0.97	T	1.04	0.92	1.36	1.28
	1.01	1.09	1.00	0.99		1.04	0.93	1.37	1.29

NOTE: The table reports the change in the DY spillover index, the SAM, the Hilbert Norm and the KLIC for 20 events, details of which are provided below. For each event, we focus on three values: the average value of the selected indicator over the 5 trading days prior to the event, the value on the event day itself and the average over the 5 trading days after the event. To facilitate comparisons, we normalize such that the average value in the five days prior to the event is set to unity. The following events are considered: (A) two Bear Stearns hedge funds collapse; (B) Bear Stearns is downgraded by S&P; (C) Bear Stearns is acquired by JP Morgan; (D) Lehman Brothers files for bankruptcy; (E) RBS reports record losses; (F) Greek parliament is dissolved; (G) European Commission releases a report on the falsification of Greek economic data; (H) Greece requests aid; (I) Ireland requests aid; (J) Portugal requests aid; (K) second Greek bailout; (L) Spain requests aid; (M) Cypriot bailout announced; (N) ECB cuts interest rates to a record low of 0.5%; (O) US federal government shutdown; (P) Greece returns to the bond market; (Q) ECB announces negative interest rate policy; (R) the October 2014 *flash-crash*; (S) Yellen discusses the case for a rate rise; (T) Greece misses its IMF payment deadline. For statistics based on the spillover density, computation is based on the Gaussian kernel with the asymptotically optimal bandwidth and the density in the first rolling sample is used as the benchmark density. The window length is set at $\tau = 250$ trading days and the forecast horizon at $h = 10$ trading days.

Table 2: Responsiveness of Selected Indicators to Major Events

Network Statistic	(a) Weekly CISS			(b) Monthly SovCISS		
	Corr.	Slope	S.E.	Corr.	Slope	S.E.
Mean	0.501	1.498	0.247	0.705	2.838	0.747
Median	0.596	2.052	0.265	0.724	3.317	0.676
Variance	-0.117	-1.278	1.217	0.132	1.847	2.833
Skewness	-0.473	-1.653	0.311	-0.675	-3.211	0.895
Kurtosis	-0.392	-10.024	2.474	-0.619	-21.901	7.840
SAM	-0.455	-0.550	0.127	-0.304	-0.478	0.246
Hilbert Norm	0.549	0.519	0.079	0.647	0.811	0.229
KLIC	0.506	1.384	0.224	0.561	2.001	0.629

NOTE: CISS denotes the ECB's Composite Indicator of Systemic Stress, which is available at weekly frequency. SovCISS is the ECB's Composite Indicator of Sovereign Stress, which is available at monthly frequency. To obtain the results under the heading 'CISS', we convert our network statistics from daily frequency to weekly frequency by taking the weekly average. Likewise, results under the heading 'SovCISS' are obtained by converting our network statistics to monthly frequency by taking the monthly average. 'Corr.' denotes the correlation between the named network statistic and either CISS or SovCISS. 'Slope' is the slope coefficient from a regression of the named network statistic on a constant and either CISS or SovCISS, with 'S.E.' denoting the Newey West heteroskedasticity and autocorrelation consistent standard error. CISS and SovCISS are available from the ECB's Statistical Data Warehouse, with series identifiers CISS.D.U2.Z0Z.4F.EC.SS_CI.IDX and CISS.M.U2.Z0Z.4F.EC.SOV_EW.IDX, respectively.

Table 3: Changes in the Shape of the Spillover Density vs. CISS and SovCISS

	Coeff.	SE	95% CI
Bilateral Portfolio Exposure	0.011***	0.004	[0.003 , 0.019]
Bilateral Portfolio Flow	-0.007**	0.003	[-0.014 , 0.000]
Bilateral Trade Exposure	-0.026	0.034	[-0.092 , 0.040]
Source Structural Balance	-0.076***	0.016	[-0.108 , -0.044]
Recipient Structural Balance	-0.030*	0.015	[-0.060 , 0.000]
Source Inflation	0.010	0.014	[-0.017 , 0.036]
Recipient Inflation	0.000	0.014	[-0.029 , 0.028]
Source Real GDP Growth	0.100***	0.008	[0.083 , 0.117]
Recipient Real GDP Growth	0.038***	0.009	[0.019 , 0.056]
Cross Sections	506		
Total observations	16,534		
Within R-squared	0.203		
Between R-squared	0.029		
Overall R-squared	0.092		

NOTE: The table reports estimation results obtained from an unbalanced panel data model where the quarterly average value of the spillover from the source sovereign to the recipient sovereign is regressed on the named explanatory variables, controlling for both country-pair and period fixed effects. Estimated values of the constant and the parameters on the country-pair and period dummies are suppressed to save space. Inference is based on heteroskedasticity and autocorrelation robust standard errors. Significance at the 10%, 5% and 1% levels is denoted by ***, ** and *, respectively. See Appendix A for details of the data sources.

Table 4: The Determinants of Bilateral Credit Risk Spillovers

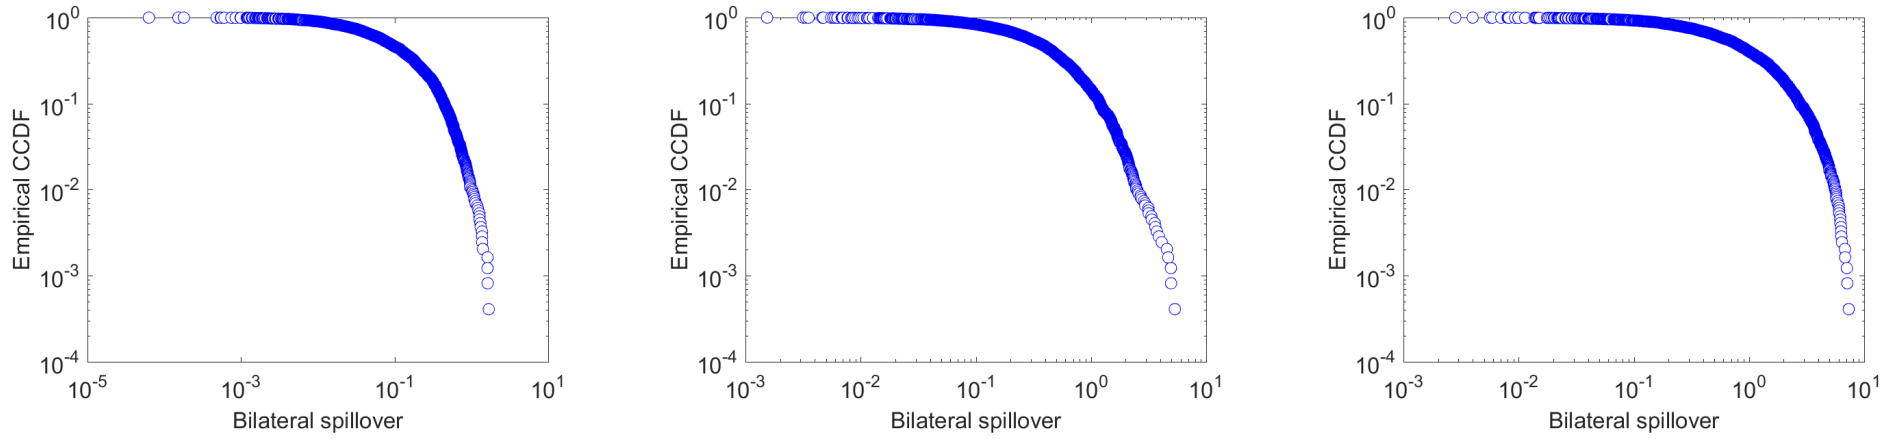

(a) Uncorrelated shocks (b) Weakly correlated shocks (c) Strongly correlated shocks

NOTE: The figure plots the empirical counter-cumulative distribution function (CCDF) for spillovers from a simulated DY network on a log scale under three different assumptions about the correlation of the shocks in the underlying VAR model. Simulations are based on the VAR(1) process $\boldsymbol{A}\boldsymbol{y}_t = \boldsymbol{B}\boldsymbol{y}_{t-1} + \boldsymbol{e_t}$ where $\boldsymbol{e_t} \sim IIN\Big(\boldsymbol{0}, \boldsymbol{\Sigma}\Big)$ assuming a cross-sectional dimension of $N = 50$. To obtain uncorrelated shocks, we set $\boldsymbol{A} = \boldsymbol{I}_N$ and $\boldsymbol{\Sigma} = \boldsymbol{I}_N$ and draw $\boldsymbol{B} \sim MN(\boldsymbol{0}, \boldsymbol{\Omega}_B)$ where $\boldsymbol{\Omega}_B = \tilde{\boldsymbol{B}}\tilde{\boldsymbol{B}}'$ and the (i,j)th element of $\tilde{\boldsymbol{B}}$ is drawn as $\omega_{\tilde{B},ij} \sim U(-1/2N, 1/2N)$ for $i, j \in \{1, 2, \ldots, N\}$. To obtain correlated shocks, we retain the same assumptions except that we now set the i-th diagonal element of $\boldsymbol{A}$ to $a_{ii} = 1\ \forall i$ while we draw the a_{ij}'s from the multivariate normal distribution with mean $\boldsymbol{0}$ and covariance matrix $\boldsymbol{\Omega}_A = \tilde{\boldsymbol{A}}\tilde{\boldsymbol{A}}'$ where the (i,j)th element of $\tilde{\boldsymbol{A}}$ is drawn as $\omega_{\tilde{A},ij} \sim U(-c, c)$. For weakly correlated shocks we set $c = 1/2N$ and for strongly correlated shocks we use $c = 1/N$. In each case, we verify that the resulting VAR model is stationary and ergodic. Results are shown for a single representative draw. To account for the random variation across draws, Appendix Table B.1 provides additional results based on 10,000 draws of $\boldsymbol{A}$ and $\boldsymbol{B}$.

Figure 1: Asymmetry of the Weighted Degree Distribution in the DY Framework

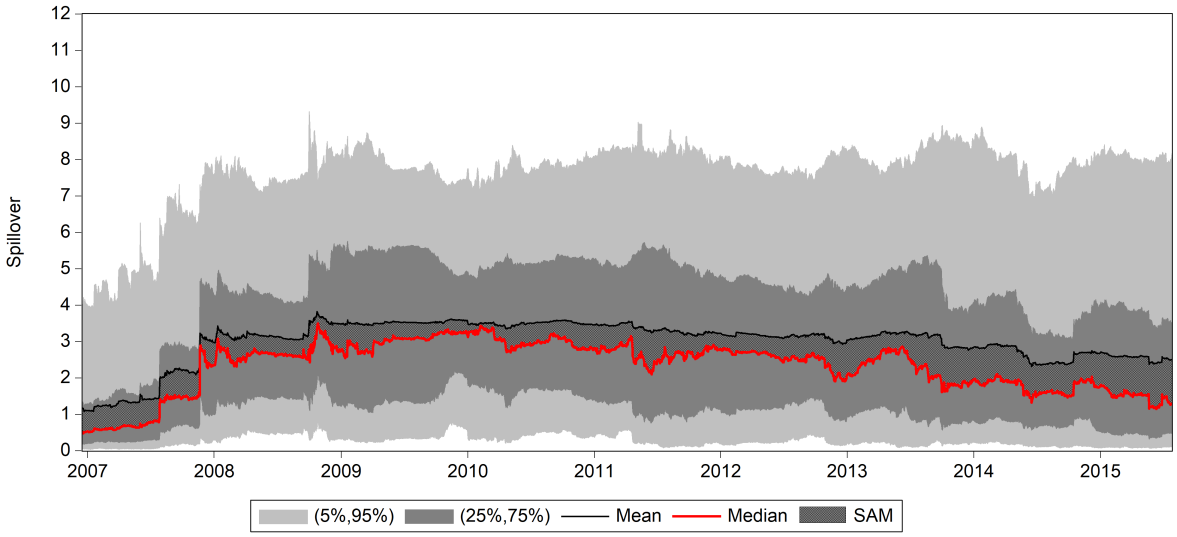

(a) Contour plot of the spillover histograms over rolling samples

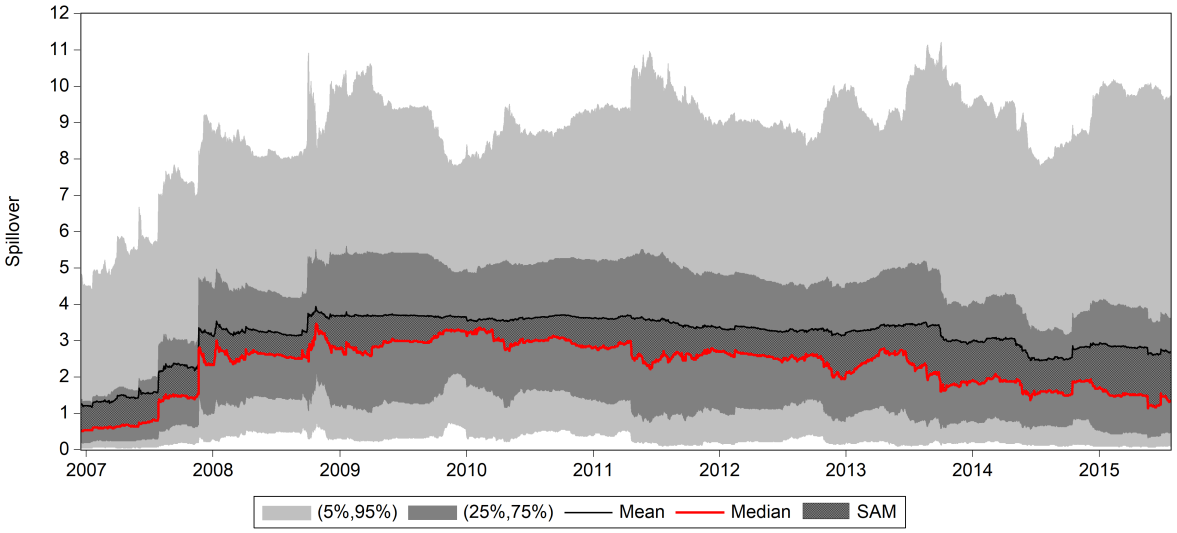

(b) Contour plot of the spillover densities over rolling samples

NOTE: Panels (a) and (b) present contour plots which trace the mean, the median, the spillover asymmetry measure (SAM) and selected percentiles of the spillover histogram and spillover density (respectively) over rolling samples. In each rolling sample, the spillover density is computed using the Gaussian kernel with the asymptotically optimal bandwidth. The window length is set at $\tau = 250$ trading days and the forecast horizon at $h = 10$ trading days. The unit of measurement on the vertical axis is percent. The dates shown on the horizontal axis correspond to the end of each rolling sample.

Figure 2: Comparison of the Spillover Density and the Spillover Histogram

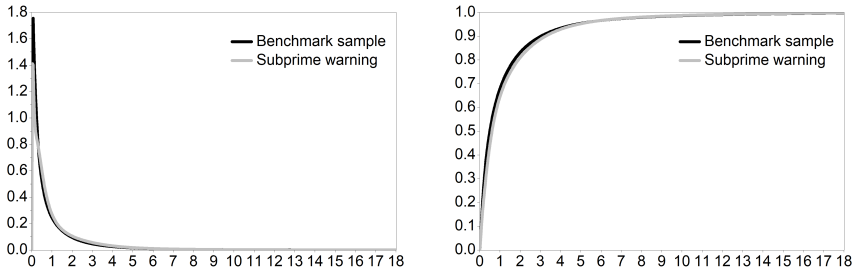

(a) Benchmark Sample and Bernanke's Subprime Warning

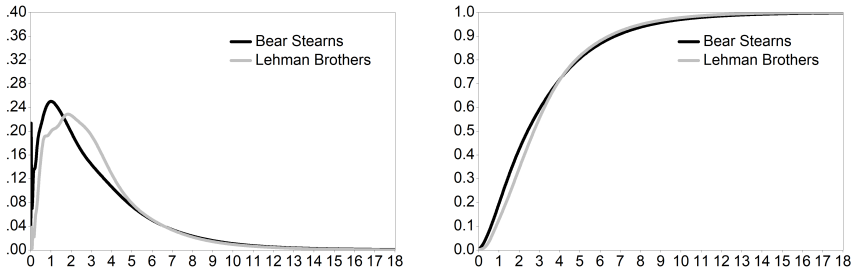

(b) Sale of Bear Stearns and Failure of Lehman Brothers

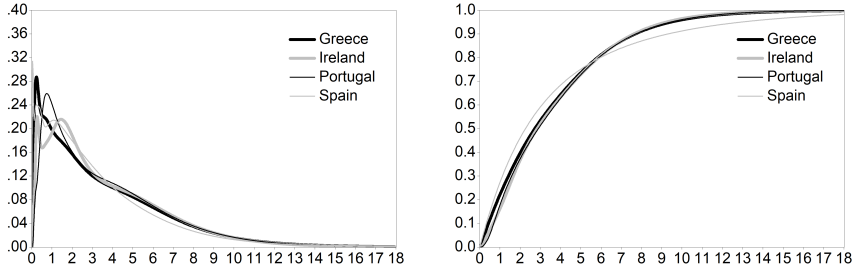

(c) GIIPS Sovereign Bailouts

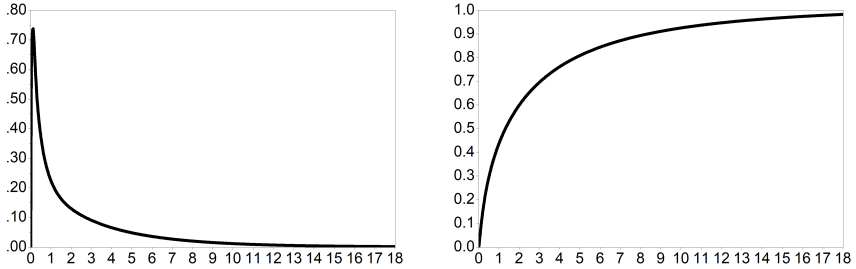

(d) Greek default on IMF repayments

Note: The figure presents PDFs (left side of each panel) and CDFs (right side) of the spillover density for the following events: the benchmark sample ending on 15-Dec-2006; Bernanke's warning of systemic risk in the subprime mortgage market on 06-Mar-2007; the acquisition of Bear Stearns by JP Morgan on 17-Mar-2008; the collapse of Lehman Brothers on 15-Sep-2008; the Greek request for aid on 23-Apr-2010; the Irish request for aid on 22-Nov-2010; the Portuguese request for aid on 06-Apr-2011; the Spanish request for aid on 25-Jun-2012; and the Greek default on its IMF debt repayments on 30-Jun-2015. If an event occurs on a non-trading day, the event date that we report refers to the next available trading day in our sample. The forecast horizon is $h = 10$ trading days in all cases and the unit of measurement on the horizontal axis is percent.

Figure 3: Changes in the Spillover Density, Selected Events

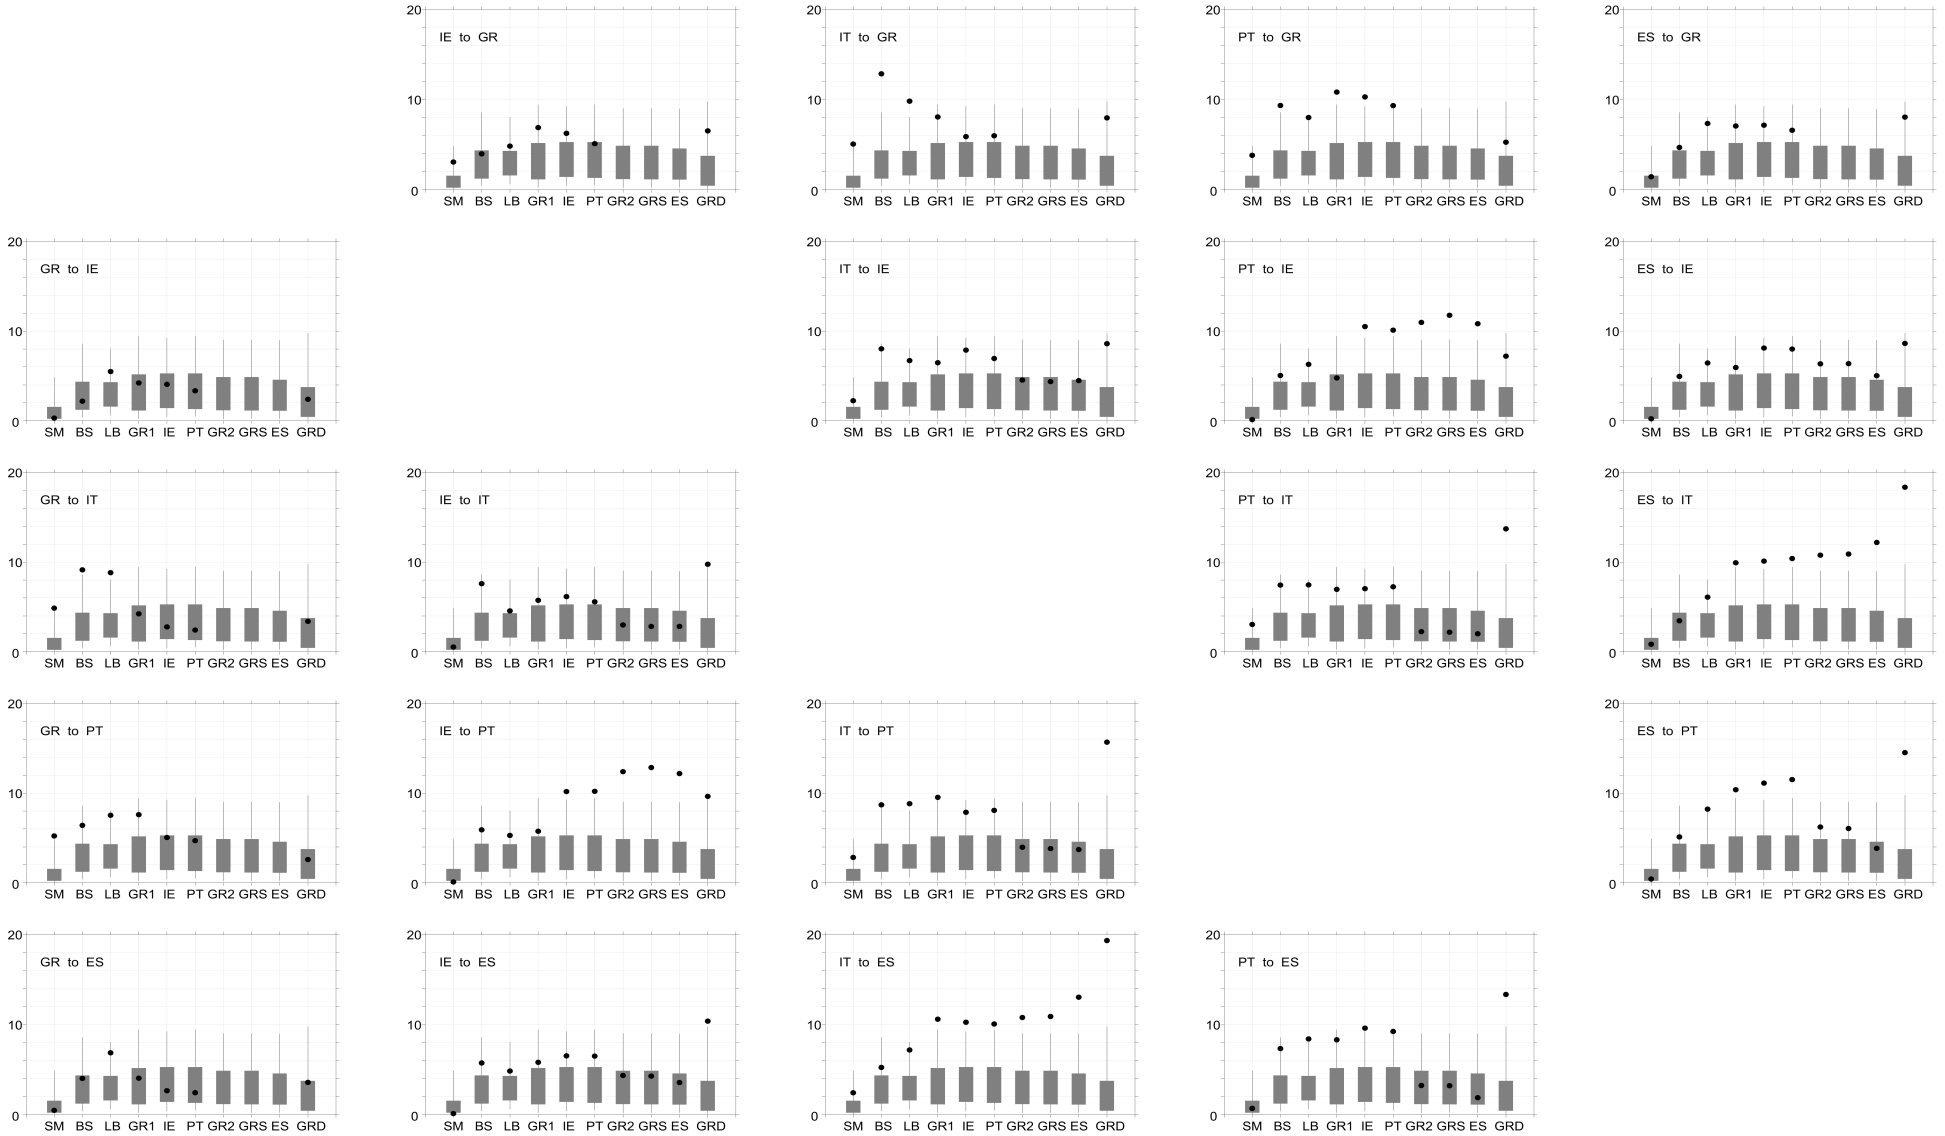

NOTE: The figure reports bilateral spillovers among the GIIPS for the following events: Bernanke's warning of systemic risk in the subprime mortgage market on 06-Mar-2007 (SM); the acquisition of Bear Stearns by JP Morgan on 17-Mar-2008 (BS); the collapse of Lehman Brothers on 15-Sep-2008 (LB); the Greek request for aid on 23-Apr-2010 (GR1); the Irish request for aid on 22-Nov-2010 (IE); the Portuguese request for aid on 06-Apr-2011 (PT); the agreement over the second Greek bailout on 21-Feb-2012 (GR2); the Greek debt-swap agreement on 09-Mar-2012 (GRS); the Spanish request for aid on 25-Jun-2012 (ES); and Greece's failure to meet its IMF payments deadline on 30-Jun-2015 (GRD). If an event occurs on a non-trading day, the event date that we report refers to the next available trading day in our sample. Missing points in the panels relating to Greece arise due to the lack of Greek CDS data between 15-Feb-2012 and 10-Jun-2013. For each event, the interquartile range of the spillover density is shown as a box and the 5-95% percentile range as whiskers. The vertical axis is in percent.

Figure 4: Bilateral Credit Risk Spillovers among the GIIPS

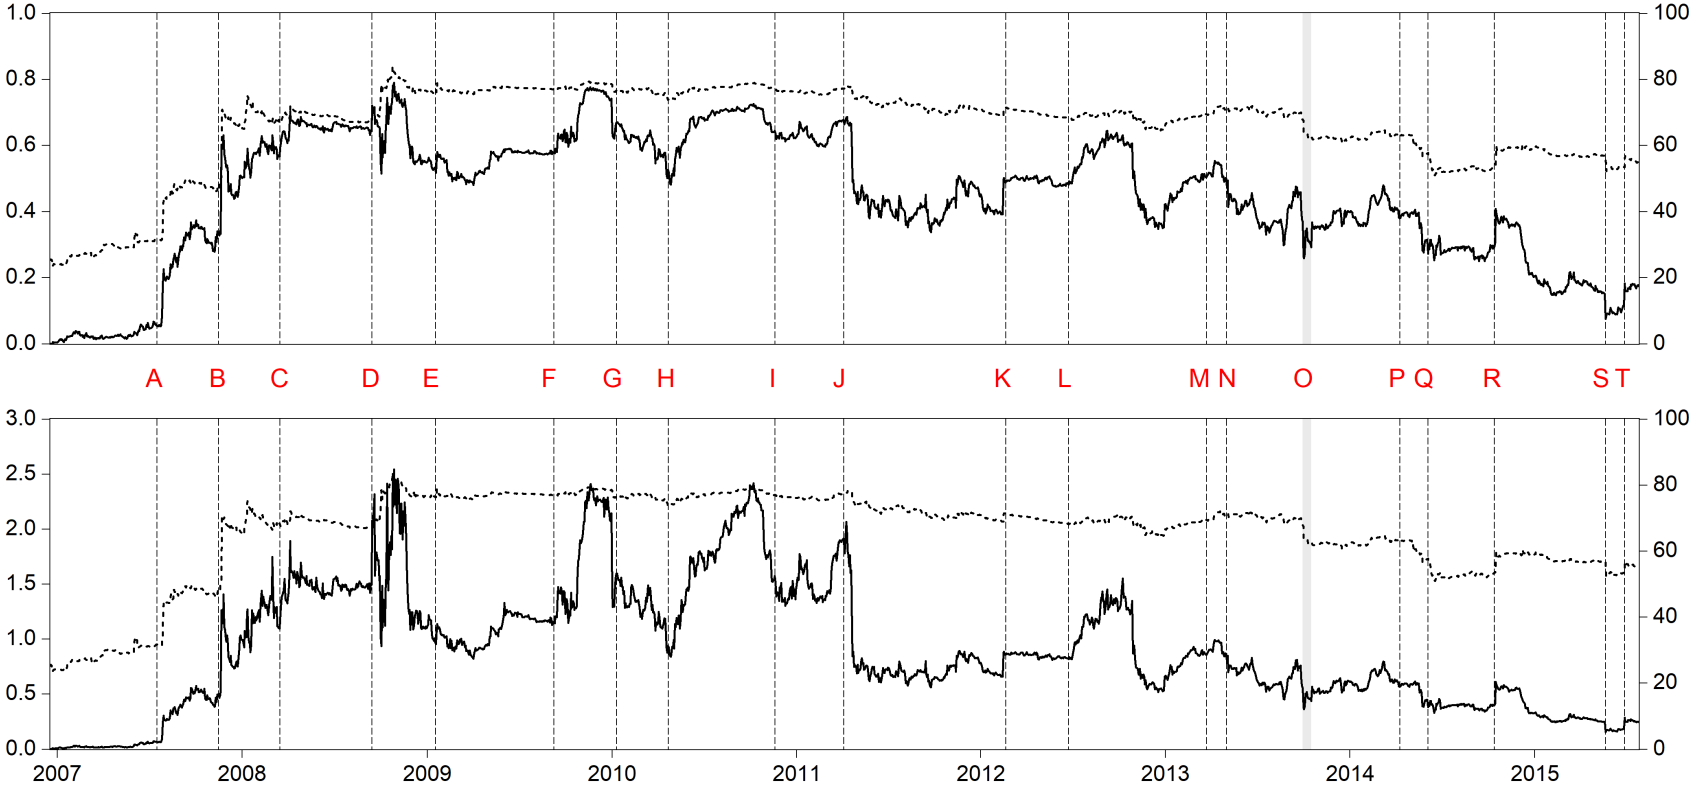

NOTE: The solid black lines represent the Hilbert norm (top panel) and the KLIC (bottom panel) which are plotted on the left axis. The dashed black line in each case is the DY spillover index plotted on the right axis. The spillover density is computed using the Gaussian kernel with the asymptotically optimal bandwidth and the density in the first rolling sample is used as the benchmark density, f_0. The window length is set at $\tau = 250$ trading days and the forecast horizon at $h = 10$ trading days. The dates shown on the horizontal axis correspond to the end of each rolling sample. Vertical lines/shading indicate the following events: (A) two Bear Stearns hedge funds collapse; (B) Bear Stearns is downgraded by S&P; (C) Bear Stearns is acquired by JP Morgan; (D) Lehman Brothers files for bankruptcy; (E) RBS reports record losses; (F) Greek parliament is dissolved; (G) European Commission releases a report on the falsification of Greek economic data; (H) Greece requests aid; (I) Ireland requests aid; (J) Portugal requests aid; (K) second Greek bailout; (L) Spain requests aid; (M) Cypriot bailout announced; (N) ECB cuts interest rates to a record low of 0.5%; (O) US federal government shutdown; (P) Greece returns to the bond market; (Q) ECB announces negative interest rate policy; (R) the October 2014 *flash-crash*; (S) Yellen discusses the case for a rate rise; (T) Greece misses its IMF payment deadline.

Figure 5: Measuring Changes in the Shape of the Spillover Density via Divergence Criteria

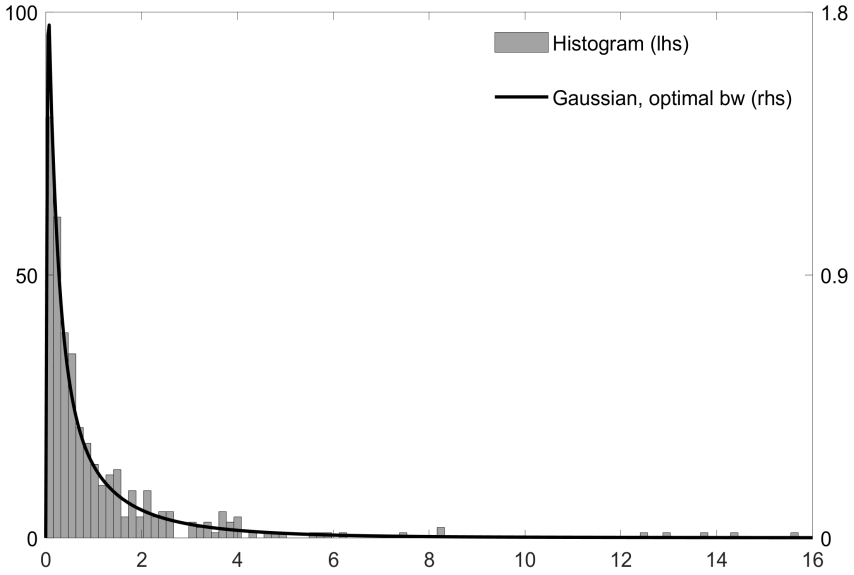

(a) First Rolling Sample (Benchmark)

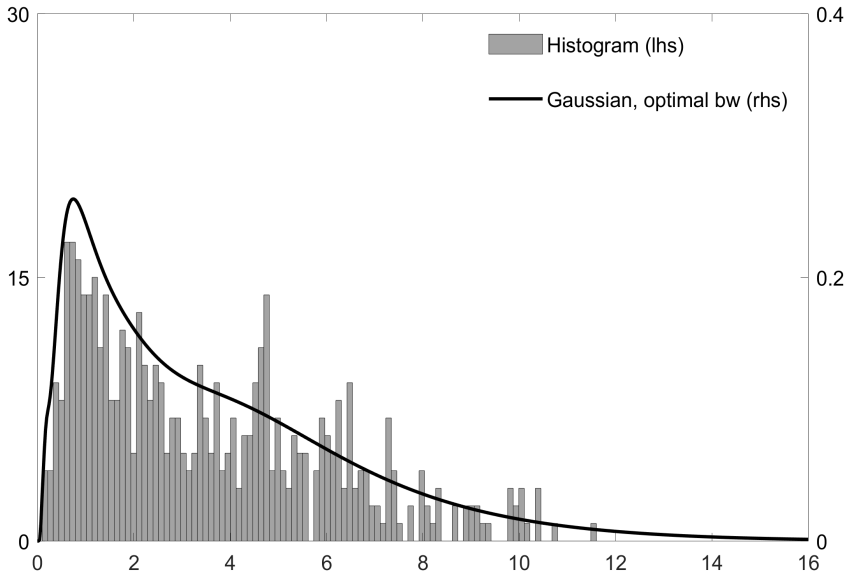

(b) Portuguese Request for Aid (6-Apr-2011)

NOTE: The density plots are computed using the bilateral spillovers (i.e. the off-diagonal elements) of the ten-days-ahead spillover matrix. Panel (a) shows the spillover density in the first rolling sample, 02-Jan-2006 to 15-Dec-2006, which we treat as a benchmark. Panel (b) shows the spillover density for the sample spanning 22-Apr-2010 to 06-Apr-2011, the end of which corresponds to Portugal's request for financial assistance. In both cases, the spillover density is estimated using the Gaussian kernel with the asymptotically optimal bandwidth. The unit of measurement on the horizontal axis is percent. For an analysis of the effect of kernel choice on these plots, see Figures C.1 and C.2 in Appendix C.

Figure 6: Spillover Densities for the Benchmark Sample and after the Portuguese Request for Aid

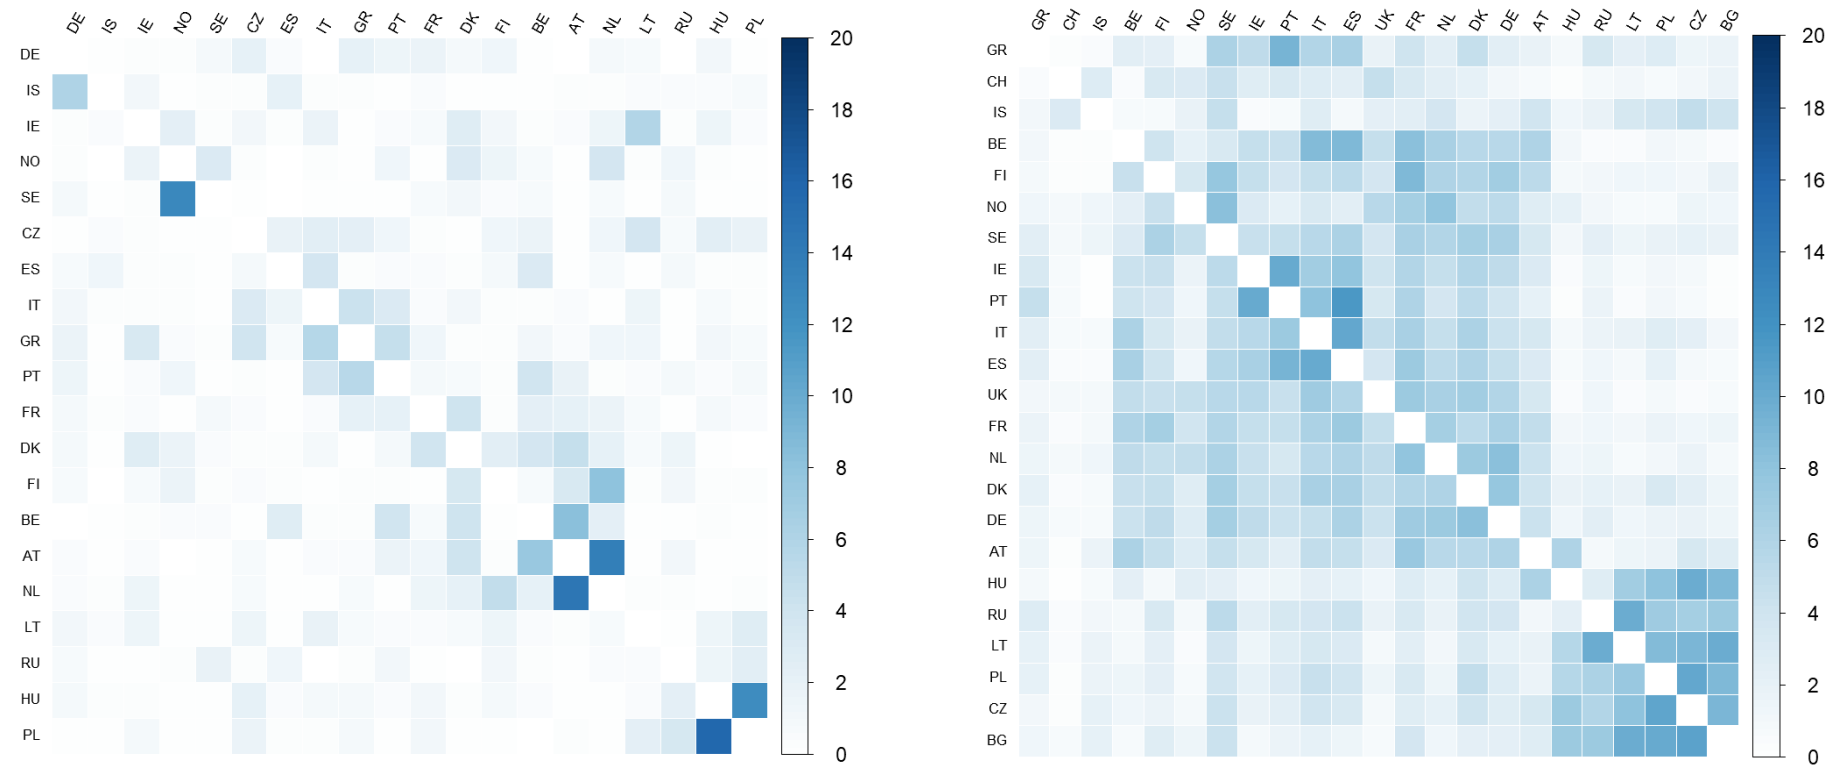

(a) Pre-Crisis Benchmark Sample

(b) Portuguese Request for Aid

NOTE: The heat maps summarize the ten-days-ahead spillover matrix, $S^{(r,10)}$ defined in equation (9) for two rolling samples: (a) the first rolling sample, 02-Jan-2006 to 15-Dec-2006, which we treat as a benchmark; and (b) the rolling sample spanning 22-Apr-2010 to 06-Apr-2011, the end point of which coincides with the Portuguese request for financial assistance. Own-variable effects on the prime diagonal have been suppressed in each case and the heatmaps have been clustered according to a single linkage algorithm for clarity of presentation. The benchmark sample contains 20 sovereigns, so each element of the spillover matrix reported in panel (a) is adjusted by a scale factor of 20/23 according to the procedure described in Section II. The Portuguese bailout sample contains all 23 sovereigns, so the scale factor in panel (b) is equal to unity. The unit of measurement is percent.

Figure 7: Connectedness Heat Maps for the Benchmark Sample and after the Portuguese Request for Aid

Technical Appendices

Intended for Online Distribution

Appendix A: Construction of the Dataset

SCDS Spreads

The SCDS spreads are mid-market rates for sovereign debt where the CDS contract has a tenor of 5 years and a complete restructuring clause. The contract currency is US dollars in all cases. The CDS data are sourced from Markit using the series identifiers listed in Table A.1.

Sovereign	Series ID	Sovereign	Series ID	Sovereign	Series ID
Austria	AUST	Greece	GREECE	Poland	POLAND
Belgium	BELG	Hungary	REPHUN	Portugal	PORTUG
Bulgaria	BGARIA	Iceland	ICELND	Russia	RUSSIA
Czech Rep.	CZECH	Ireland	IRELND	Spain	SPAIN
Denmark	DENK	Italy	ITALY	Sweden	SWED
Finland	FINL	Lithuania	LITHUN	Switzerland	SWISS
France	FRTR	Norway	NORWAY	UK	UKIN
Germany	DBR	Netherlands	NETHRS		

Table A.1: Series Identifiers for the SCDS Spreads

Explanatory Variables used in the Panel Data Model

The explanatory variables used in the estimation of the fixed effects panel data model in Section II.C are constructed as follows.

Bilateral Portfolio Investment Exposure

The exposure of country j to bilateral portfolio investments with respect to country i is captured by the sum of country i's declared holdings of assets from country j and country j's declared holdings of assets from country i, expressed as a percentage of the GDP of country j. The data on portfolio assets is sourced from the IMF's *Coordinated Portfolio Investment Survey* at annual frequency until 2012 and at semiannual frequency thereafter. The data is then converted to quarterly frequency by linear interpolation. Nominal GDP is obtained from the IMF's *International Financial Statistics* in local currency units and is then converted to US dollars using the end-of-period exchange rate, which is also obtained from the IMF's *International Financial Statistics*.

Bilateral Portfolio Investment Flow

The magnitude of the gross bilateral portfolio investment flow between countries i and j from the perspective of country j is captured by the quarterly change in the sum of country i's declared holdings of assets from country j and country j's declared holdings of assets from country i, expressed as a percentage of the GDP of country j. Data sources are as above.

Bilateral Trade Exposure

The exposure of country j to bilateral trade with country i is captured by the sum of exports and imports between countries i and j expressed as a percentage of the GDP of country j. The trade data is sourced from the IMF's *Direction of Trade Statistics* at quarterly frequency. The nominal GDP data is obtained as in the case of bilateral portfolio exposures above.

Structural Budget Balance

The cyclically-adjusted fiscal position of ith (jth) sovereign is captured by its structural budget balance expressed as a percentage of its potential GDP. Our data is sourced from the IMF's *World Economic Outlook* at annual frequency. The data is reported as a percentage of potential GDP. We obtain quarterly series by linear interpolation of the annual data.

Real GDP Growth

The annual growth rate of real GDP in country i (j) relative to the same quarter of the previous year, expressed in percent. Our data is sourced from the IMF's *International Financial Statistics* at quarterly frequency.

CPI Inflation

The annual rate of CPI inflation in country i (j) relative to the same quarter of the previous year, expressed in percent. Our data is constructed as the growth rate of the CPI index relative to the same quarter in the previous year. We obtain CPI indices for each country from the IMF's *International Financial Statistics* at quarterly frequency.

Appendix B: Additional Simulation Results

	Mean	Variance	Skewness	Kurtosis
Uncorrelated shocks				
Lower Decile	0.157	0.042	2.484	11.128
Median	0.165	0.047	2.751	13.820
Upper Decile	0.175	0.053	3.131	18.755
Weakly correlated shocks				
Lower Decile	0.443	0.189	1.825	7.248
Median	0.472	0.221	2.063	9.048
Upper Decile	0.508	0.264	2.402	12.385
Strongly correlated shocks				
Lower Decile	1.068	1.277	1.555	5.526
Median	1.188	1.522	1.867	7.218
Upper Decile	1.347	1.797	2.141	9.180

Note: All simulations are based on the VAR(1) process $\boldsymbol{A}\boldsymbol{y}_t = \boldsymbol{B}\boldsymbol{y}_{t-1} + \boldsymbol{e_t}$ where $\boldsymbol{e_t} \sim IIN\left(\boldsymbol{0}, \boldsymbol{\Sigma}\right)$ assuming a cross-sectional dimension of $N = 50$. To obtain uncorrelated shocks, we set $\boldsymbol{A} = \boldsymbol{I}_N$ and $\boldsymbol{\Sigma} = \boldsymbol{I}_N$ and draw $\boldsymbol{B} \sim MN(\boldsymbol{0}, \boldsymbol{\Omega}_B)$ where $\boldsymbol{\Omega}_B = \tilde{\boldsymbol{B}}\tilde{\boldsymbol{B}}'$ and the (i,j)th element of $\tilde{\boldsymbol{B}}$ is drawn as $\omega_{\tilde{B},ij} \sim U(-1/2N, 1/2N)$ for $i, j \in \{1, 2, \ldots, N\}$. To obtain correlated shocks, we retain the same assumptions except that we now set the i-th diagonal element of $\boldsymbol{A}$ to $a_{ii} = 1\ \forall i$ while we draw the a_{ij}'s from the multivariate normal distribution with mean $\boldsymbol{0}$ and covariance matrix $\boldsymbol{\Omega}_A = \tilde{\boldsymbol{A}}\tilde{\boldsymbol{A}}'$ where the (i,j)th element of $\tilde{\boldsymbol{A}}$ is drawn as $\omega_{\tilde{A},ij} \sim U(-c, c)$. For weakly correlated shocks we set $c = 1/2N$ while for strongly correlated shocks $c = 1/N$. In each case, we generate 10,000 draws of $\boldsymbol{A}$ and $\boldsymbol{B}$ and verify that the resulting VAR model is stationary and ergodic. We then apply the Diebold-Yilmaz technique within each draw to estimate the set of $N(N-1)$ bilateral spillovers between variables. For each draw, we compute the sample mean, variance, skewness and kurtosis of the set of bilateral spillovers. The table reports selected percentiles of the empirical distributions of these four central moments obtained over the 10,000 draws.

Table B.1: Evaluating the Shape of the Spillover Density by Simulation

Appendix C: Sensitivity Tests

Sensitivity to the Window Length and Forecast Horizon

While there is no simple method to select either the window length or the forecast horizon optimally, the existing literature provides a useful guide. Applications of the DY method to daily data have used a range of values. Diebold and Yilmaz (2014) obtain their main results using $\tau = 100$ and $h = 12$ days, Baruník et al. (2016) use $\tau = 200$ and $h = 10$ days and Greenwood-Nimmo et al. (2016) use $\tau = 250$ and $h = 10$ days. Given this uncertainty, the emergent norm in the literature is to evaluate the sensitivity of selected network statistics to several candidate window lengths and forecast horizons. Adopting this approach, we begin by evaluating the sensitivity of the DY spillover index and the KLIC to varying $\tau \in \{200, 250, 300\}$ and $h \in \{5, 10, 15\}$ trading days.[18]

Table C.1(a) reports the common sample correlation coefficients between the DY spillover indices computed under all pairwise combinations of $\tau \in \{200, 250, 300\}$ and $h \in \{5, 10, 15\}$ trading days. The reported correlations are remarkably high, indicating that the behavior of the DY spillover index is not driven either by the choice of window length or horizon. In fact, for any given window length, the correlation across different values of the forecast horizon is close to perfect, which indicates that the GVDs have largely converged to their long-run values within five days. This reflects the low degree of serial correlation in the defactored SCDS spreads and is a common finding in the existing literature on DY networks (e.g. Greenwood-Nimmo et al., 2016).

Table C.1(b) reports the common sample correlations of the KLIC as we vary $h \in \{5, 10, 15\}$ and $\tau \in \{200, 250, 300\}$. For this exercise, the KLIC is computed based on the rolling sample spillover densities obtained using the Gaussian kernel with the asymptotically optimal bandwidth and the spillover density in the first rolling sample is treated as the benchmark density. The reported correlations reveal that the KLIC exhibits slightly more sensitivity to the choice of window and forecast horizon than is observed in Table C.1(a) using the DY spillover index, although the correlation is still strongly positive in general and is never lower than 0.77 to two decimal places.

Based on these findings, we set $\tau = 250$ and $h = 10$ without loss of generality.

Sensitivity to the KDE Implementation

Given that a variety of kernels are in widespread use in economics and finance, we consider four options: (i) the Gaussian kernel with Silverman's (1986) rule-of-thumb bandwidth; (ii) the Gaussian kernel with the asymptotically optimal bandwidth; (iii) the Epanechnikov kernel with the asymptotically optimal bandwidth; and (iv) the adaptive kernel density estimator of Botev et al. (2010, BGK).

Table C.2 compares the four implementations of the KDE that we consider. For a given divergence criterion (either the Hilbert Norm or the KLIC), the table reports the correlation between the results obtained using the four different implementations of the KDE. In all cases, the window length is $\tau = 250$ trading days and the forecast horizon is $h = 10$ trading days. The reported correlation is never lower than 0.97 to two decimal places, indicating that the properties of the spillover density are highly robust to the choice of KDE.

[18] Given the dimension of our VAR model, we consider it imprudent to set $\tau < 200$ trading days.

As a visual illustration of this robustness to the KDE implementation, Figures C.1 and C.2 recreate panels (a) and (b) of Figure 6 using each of the four KDEs discussed above. In Figure C.1, regardless of the kernel that is considered, the estimated density fits the spillover histogram very well. The majority of the probability mass lies close to zero, reflecting our finding that most bilateral spillovers are negligible in normal times. The density exhibits a marked skew, with a long right tail corresponding to the subset of strong spillovers noted in the main text.

It does not follow that the admirable performance of each kernel will necessarily be maintained as the spillover density evolves and takes on different shapes. In Figure C.2, we focus on the rolling sample ending with the Portuguese bailout on 06-Apr-2011, a sample where we know from Figure 2 that the spillover density will have shifted rightward and will exhibit a lower degree of skewness. Some differences in the behavior of the four KDEs are now apparent. The Gaussian kernel using Silverman's rule-of-thumb bandwidth mildly over-estimates the mass in both tails and fails to capture the apparent bimodality of the spillover histogram. When the asymptotically optimal bandwidth is used, both the Gaussian and Epanechnikov kernels perform relatively well. The variable bandwidth of the BGK adaptive KDE helps to capture the bimodality of the estimated spillovers in this case, although the shape of the resulting density is rather complex.

In light of these experiments and in deference to the principle of parsimony, we adopt the simple Gaussian kernel with the optimal bandwidth as our chosen KDE without loss of generality.

		$w = 200$ days			$w = 250$ days			$w = 300$ days		
		$h = 5$	$h = 10$	$h = 15$	$h = 5$	$h = 10$	$h = 15$	$h = 5$	$h = 10$	$h = 15$
	$h = 5$	1.000	1.000	1.000	0.971	0.971	0.971	0.951	0.951	0.951
$w = 200$	$h = 10$		1.000	1.000	0.971	0.971	0.971	0.951	0.951	0.951
	$h = 15$			1.000	0.971	0.971	0.971	0.951	0.951	0.951
	$h = 5$				1.000	1.000	1.000	0.982	0.982	0.982
$w = 250$	$h = 10$					1.000	1.000	0.982	0.982	0.982
	$h = 15$						1.000	0.982	0.982	0.982
	$h = 5$							1.000	1.000	1.000
$w = 300$	$h = 10$								1.000	1.000
	$h = 15$									1.000

(a) DY Spillover Index

		$w = 200$ days			$w = 250$ days			$w = 300$ days		
		$h = 5$	$h = 10$	$h = 15$	$h = 5$	$h = 10$	$h = 15$	$h = 5$	$h = 10$	$h = 15$
	$h = 5$	1.000	1.000	1.000	0.836	0.837	0.837	0.770	0.770	0.770
$w = 200$	$h = 10$		1.000	1.000	0.836	0.837	0.837	0.770	0.771	0.771
	$h = 15$			1.000	0.836	0.837	0.837	0.770	0.771	0.771
	$h = 5$				1.000	1.000	1.000	0.878	0.878	0.878
$w = 250$	$h = 10$					1.000	1.000	0.878	0.878	0.878
	$h = 15$						1.000	0.878	0.878	0.878
	$h = 5$							1.000	1.000	1.000
$w = 300$	$h = 10$								1.000	1.000
	$h = 15$									1.000

(b) Kullback-Leibler Information Criterion

NOTE: The table reports balanced sample correlation coefficients for the DY spillover index and the Kullback Leibler information criterion as the window length is varied over $\tau = \{200, 250, 300\}$ trading days and the forecast horizon is varied over $h = \{5, 10, 15\}$ trading days.

Table C.1: Sensitivity to the Choice of Window Length and Forecast Horizon

	Silverman	Gaussian	Epanechnikov	BGK
Silverman	1.000	0.996	0.996	0.987
Gaussian		1.000	1.000	0.991
Epanechnikov			1.000	0.992
BGK				1.000

(a) Hilbert Norm

	Silverman	Gaussian	Epanechnikov	BGK
Silverman	1.000	0.996	0.996	0.980
Gaussian		1.000	1.000	0.973
Epanechnikov			1.000	0.973
BGK				1.000

(b) Kullback-Leibler Information Criterion

NOTE: The table reports the correlation between the Hilbert Norm and Kullback Leibler Information Criterion over our rolling samples using four different kernels: the Gaussian kernel with Silverman's rule of thumb bandwidth (Silverman); the Gaussian kernel with the asymptotically optimal bandwidth (Gaussian); the Epanechnikov kernel with the asymptotically optimal bandwidth (Epanechnikov); and the BGK kernel with a data-driven bandwidth (BGK). In all cases, the window length is $\tau = 250$ trading days and the forecast horizon is $h = 10$ trading days.

Table C.2: Sensitivity to the Choice of Kernel

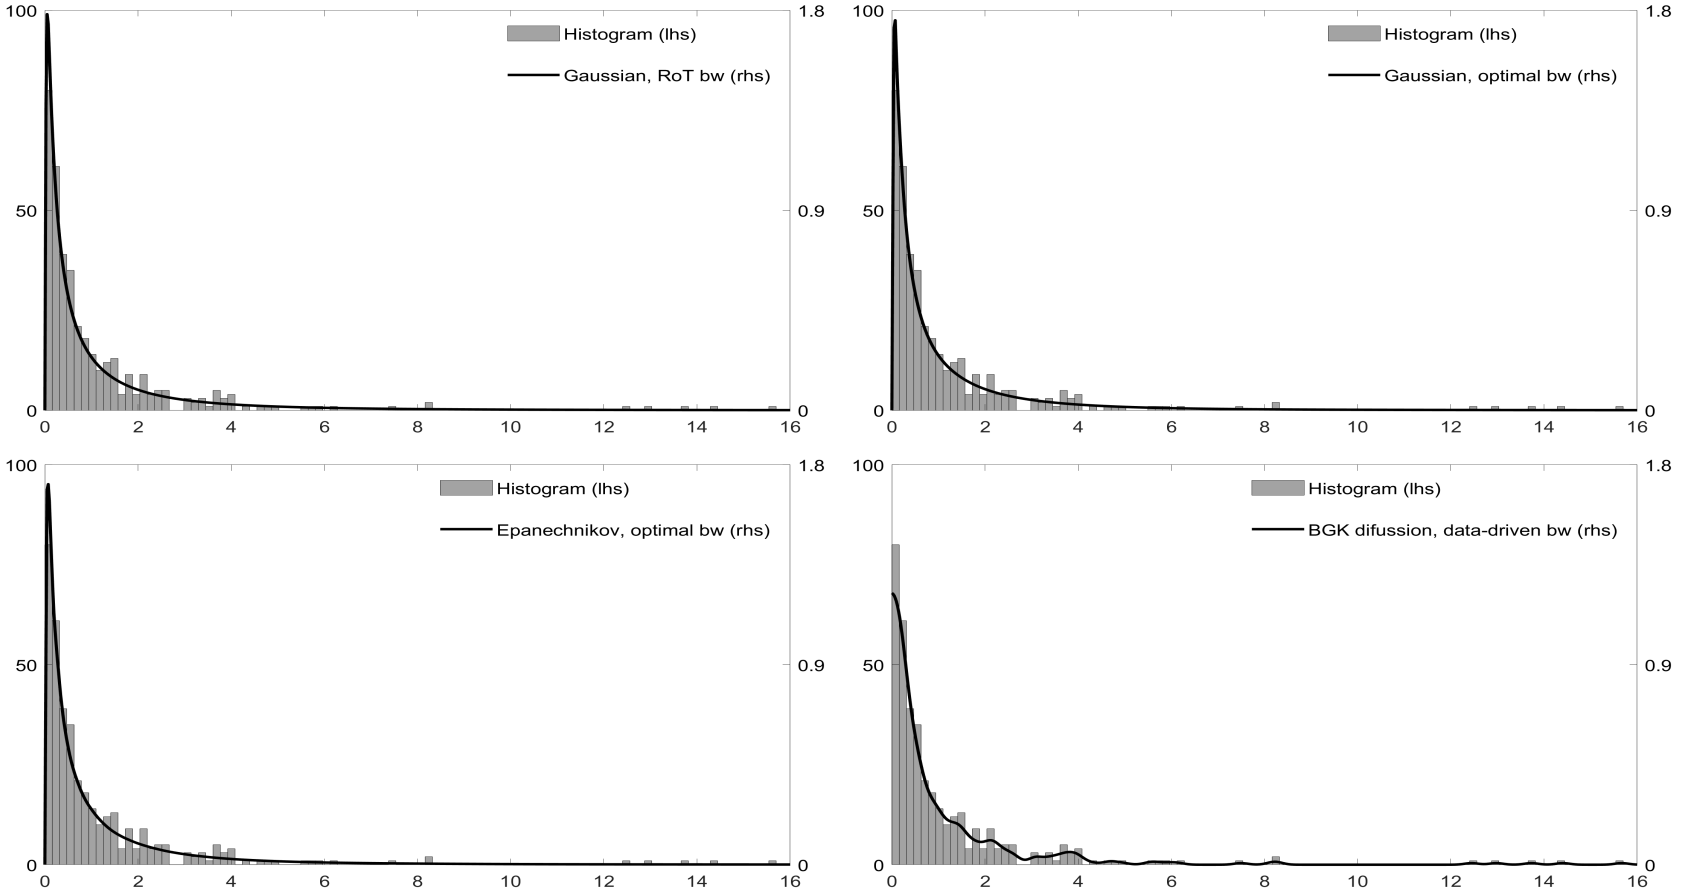

NOTE: The density plots are computed using the bilateral spillovers (i.e. the off-diagonal elements) of the ten-days-ahead spillover matrix for the first rolling sample, 02-Jan-2006 to 15-Dec-2006. The spillover density is estimated using four different implementations of the KDE: (i) the Gaussian kernel using Silverman's (1986) rule-of-thumb bandwidth; (ii) the Gaussian kernel using the asymptotically optimal bandwidth; (iii) the Epanechnikov kernel using the asymptotically optimal bandwidth; and (iv) the BGK adaptive approach using the data-driven bandwidth. The unit of measurement on the horizontal axis is percent.

Figure C.1: Spillover Densities for the First Rolling Sample (Benchmark Case) using Different KDEs

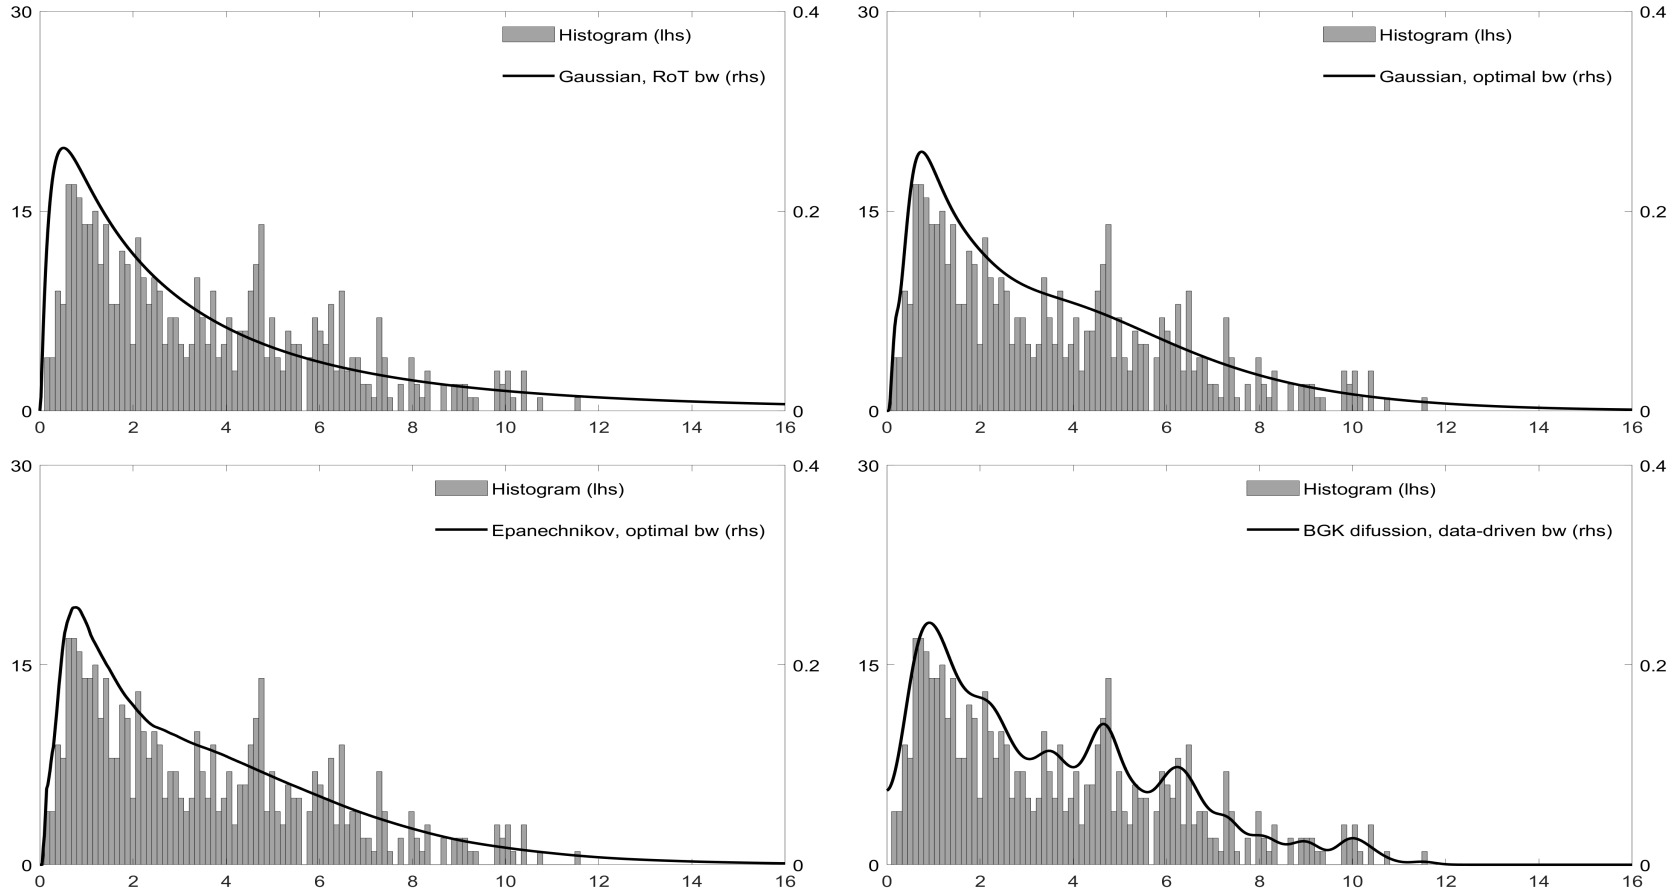

NOTE: The density plots are computed using the bilateral spillovers (i.e. the off-diagonal elements) of the ten-days-ahead spillover matrix for the sample spanning 22-Apr-2010 to 06-Apr-2011. The end point of this sample corresponds to Portugal's request for financial assistance. As in Figure C.1, the spillover density is estimated using four different implementations of the KDE. The unit of measurement on the horizontal axis is percent.

Figure C.2: Spillover Densities after the Portuguese Request for Aid (6-Apr-2011) using Different KDEs

Appendix D: A Comparison using Raw SCDS Spreads

Our focus in this paper is on the spillover of idiosyncratic sovereign credit risk using defactored SCDS spreads. The purpose of this appendix is to compare the evolution of the spillover density obtained under an alternative setting using raw SCDS spreads (i.e. without defactoring). This is an interesting exercise for two reasons. First, it provides information on the role of the common factors in determining the observed comovement of SCDS spreads. Second, it provides a basis for comparison against papers that apply the DY method to the study raw SCDS spreads (e.g. Bostanci and Yilmaz, 2020).

Figure D.1 reports the difference between the results obtained using defactored SCDS spreads and those obtained without de-factoring the SCDS spreads for each of the following indicators: (a) the DY spillover index; (b) the KLIC; (c) the median spillover; and (d) the SAM, which we define as the mean spillover minus the median spillover. For the reader's convenience, the events that are marked in Figure 5 are also shown in Figure D.1.

The discrepancy between the spillover statistics obtained from the de-factored and raw SCDS spreads reflects the comovement in the cross-section of SCDS spreads that is attributable to common sources of variation. To illustrate how the figure is interpreted, consider panel (a). At the start of the sample, the plot records a positive value, indicating that the DY spillover index obtained using defactored SCDS spread is larger than that obtained using the raw spreads. The opposite is true at the time of JP Morgan's acquisition of Bear Stearns (event C). The other panels are interpreted similarly.

Figure D.1 reveals several interesting phenomena. First, for most of the sample, the spillover density obtained using raw SCDS spreads lies to the right of the spillover density obtained using defactored data. This is to be expected, because factor dependence represents an additional source of comovement in the cross-section of SCDS spreads.

Second, the spillover density for the defactored SCDS spreads is typically less right-skewed than the spillover density for the raw SCDS spreads. This effect arises because the peak of the spillover density shifts right under the influence of the common factors, reducing the relative length of the right tail.

Third, the KLIC typically takes larger values in the model estimated using raw SCDS data. This indicates that the evolution of the spillover density is even more marked in the model without factors, presumably because changes in spillover activity in this model can arise through more channels than in the model using defactored SCDS spreads. This observation is borne out by the fact that the discrepancy between the spillover measures obtained from the defactored and raw data is non-random, highly persistent and displays immediate and substantial responses to many major financial market and regulatory events. This implies that variations in the factors and/or the factor loadings provide a vector for shock propagation.

We are grateful to an anonymous referee for suggesting that we investigate the impact of excluding the factor structure from our model and we agree with their assessment that decomposing CDS spread comovements into a part that arises from the spillover of fundamental risk and a part due to the equilibrium pricing mechanism represents an interesting and potentially fruitful avenue for continuing research.

NOTE: This figure reports the difference between selected network statistics obtained from our model estimated using de-factored SCDS spreads and an alternative version estimated using raw SCDS spreads. The following network statistics are considered: (a) the DY spillover index; (b) the KLIC; (c) the median spillover; and (d) the spillover asymmetry measure (SAM), which we define as the mean spillover minus the median spillover. The spillover density is computed using the Gaussian kernel with the asymptotically optimal bandwidth and the density in the first rolling sample is used as the benchmark density, f_0. The window length is set at $\tau = 250$ trading days and the forecast horizon at $h = 10$ trading days. The dates shown on the horizontal axis correspond to the end of each rolling sample. Vertical lines/shading indicate the following events: (A) two Bear Stearns hedge funds collapse; (B) Bear Stearns is downgraded by S&P; (C) Bear Stearns is acquired by JP Morgan; (D) Lehman Brothers files for bankruptcy; (E) RBS reports record losses; (F) Greek parliament is dissolved; (G) European Commission releases a report on the falsification of Greek economic data; (H) Greece requests aid; (I) Ireland requests aid; (J) Portugal requests aid; (K) second Greek bailout; (L) Spain requests aid; (M) Cypriot bailout announced; (N) ECB cuts interest rates to a record low of 0.5%; (O) US federal government shutdown; (P) Greece returns to the bond market; (Q) ECB announces negative interest rate policy; (R) the October 2014 *flash-crash*; (S) Yellen discusses the case for a rate rise; (T) Greece misses its IMF payment deadline.

Figure D.1: Effect of Omitting the Factor Structure on Selected Spillover Measures