Modeling Financial Sector Joint Tail Risk in the Euro Area

Lucas, A.; Schwaab, B.; Zhang, X.

MODELING FINANCIAL SECTOR JOINT TAIL RISK IN THE EURO AREA

ANDRÉ LUCAS,[a] BERND SCHWAAB[b]* AND XIN ZHANG[c]

[a] *Vrije Universiteit Amsterdam and Tinbergen Institute, Netherlands*
[b] *Financial Research Department, European Central Bank, Frankfurt, Germany*
[c] *Research Division, Sveriges Riksbank, Stockholm, Sweden*

We develop a novel high-dimensional non-Gaussian modeling framework to infer measures of conditional and joint default risk for numerous financial sector firms. The model is based on a dynamic generalized hyperbolic skewed-*t* block equicorrelation copula with time-varying volatility and dependence parameters that naturally accommodates asymmetries and heavy tails, as well as nonlinear and time-varying default dependence. We apply a conditional law of large numbers in this setting to define joint and conditional risk measures that can be evaluated quickly and reliably. We apply the modeling framework to assess the joint risk from multiple defaults in the euro area during the 2008–2012 financial and sovereign debt crisis. We document unprecedented tail risks between 2011 and 2012, as well as their steep decline following subsequent policy actions. Copyright © 2016 John Wiley & Sons, Ltd.

Supporting information may be found in the online version of this article.

1. INTRODUCTION

In this paper we develop a novel high-dimensional non-Gaussian modeling framework to infer conditional and joint risk measures for many financial sector firms. The model is based on a dynamic generalized hyperbolic skewed-*t* copula with time-varying volatility and dependence parameters. Such a framework naturally accommodates asymmetries and heavy tails, as well as nonlinear and time-varying default dependence. To balance the need for parsimony as well as flexibility in a high-dimensional cross-section, we endow the dynamic model with a score-driven block equicorrelation structure. Furthermore, we demonstrate that a conditional law of large numbers applies in our setting, which allows us to define risk measures that can be evaluated semi-analytically and within seconds. The modeling framework allows us to assess the joint risk from multiple financial firm defaults in the euro area during the financial and sovereign debt crisis. We document unprecedented tail risks between 2011 and 2012, as well as a sharp decline in joint (but not conditional) tail risk probabilities following a sequence of announcements by the European Central Bank (ECB) that introduced its Outright Monetary Transactions (OMT) program.[1]

Since the onset of the financial crisis in 2007, financial stability monitoring has become a key priority for many central banks, which added to their respective monetary policy mandates; see, for example, Acharya *et al.* (2012) and Adrian *et al.* (2013). The new prudential responsibilities involve monitoring financial risks to and from a large system of interconnected financial intermediaries. The cross-sectional dimensions of these systems are typically high, even if attention is restricted to only

---

* Correspondence to: Bernd Schwaab, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt, Germany. E-mail: bernd.schwaab@ecb.int

[1] See ECB (2012) and Coeuré (2013). The OMT is a non-standard monetary policy measure within which the ECB could, under certain conditions, make purchases in secondary markets of bonds issued by euro area member states.

large and systemically important institutions. Our modeling framework is directly relevant for such monitoring tasks. In addition, our framework is interesting for financial institutions and clearing houses that are required to actively set risk limits (often in real time) and maintain economic capital buffers to withstand bad risk outcomes due to exposures to a large number of credit-risky counterparties. Finally, with the benefit of hindsight, evaluating the time variation in conditional and joint risks allows us to assess the impact of non-standard policy measures that central banks (or other actors) take on the risk of a simultaneous and widespread failure of financial intermediaries.

Our starting point for modeling time-varying joint and conditional risks is a dynamic copula framework, also considered by Segoviano and Goodhart (2009), Christoffersen *et al.* (2012), Oh and Patton (2014) and Lucas *et al.* (2014). In each case, a collection of firms is seen as a portfolio of firms whose multivariate dependence structure is inferred from equity returns or CDS data. Our current framework extends the Lucas *et al.* (2014) approach in two ways. First, by considering grouped equicorrelation structures, our current framework allows us to fit a cross-sectional dimension much larger than, say, 15 firms, while retaining the ability to capture the salient data features such as skewness, fat tails and time-varying correlations. We thus fix the drawback that parameter estimation breaks down when considering many firms in this class of models due to a well-known curse of dimensionality; see Engle and Kelly (2012) for a discussion. Second, we show how to evaluate joint and conditional risk measures within seconds in the current framework. Estimating time-varying parameters and portfolio risk measures is computationally relatively inexpensive because explicit expressions are available for the likelihood and the joint and conditional portfolio credit risk measures. Simulation-based methods are not required, but are used in our empirical study to provide points of comparison from a robustness perspective.

We apply our general framework by studying joint and conditional default probabilities for financial sector firms in the euro area, based on weekly data from January 1999 to September 2013. We consider two applications. First, using a limited sample of $N = 10$ firms, we verify that our dynamic correlations based on the block equicorrelation assumption closely track the average correlations from a full-correlation-matrix model analysis. We demonstrate that the loss of cross-sectional dispersion in correlations hardly matters when evaluating joint credit risk measures, at least for our sample of firms. Finally, we verify that our semi-analytical approximations to compute joint and conditional risk measures already work well even if the cross-sectional dimension is as low as 10 firms.

We then turn to our final high-dimensional study of $N = 73$ firms. We document unprecedented joint tail risks for a set of large euro area financial sector firms during the financial and sovereign debt crises, and also establish a clear peak of financial sector joint default risk in the summer of 2012. Based on time variation in our joint risk measures we argue that three events collectively ended the most acute phase of extreme financial sector tail risks in the euro area. These events are a speech by the ECB President in London to do '*whatever it takes*' to save the euro on 26 July 2012, the announcement of the ECB's OMT program on 2 August 2012 and the disclosure of the OMT details on 6 September 2012. This is a startling finding, as the ECB's OMT program provides conditional and partial insurance to governments and not to financial sector firms. Conditional tail probabilities did not decline as much, indicating that risk spillovers might have remained a concern. Based on the OMT's strong impact on financial sector joint risks, we conclude that the design and implementation of unconventional monetary policies and financial stability (tail) outcomes are strongly related. This finding suggests substantial scope for the coordination of monetary, macro-prudential and bank supervision policies. This is relevant as both monetary policy as well as banking supervision has been carried out jointly by the ECB since November 2014.

Our study relates to several directions of current research. First, we draw from the growing literature on non-Gaussian dependence modeling as well as the literature on credit risk measurement and portfolio loss asymptotics. Time-varying parameter models for volatility and dependence were considered, among others, by Engle (2002), Creal *et al.* (2011), Christoffersen *et al.* (2012), Engle and

Kelly (2012) and Creal and Tsay (2015). Similarly, credit risk models and portfolio tail risk measures were studied, for example, by Vasicek (1987), Lucas *et al.* (2001, 2003), Gordy (2000, 2003), Koopman *et al.* (2011, 2012) and Giesecke *et al.* (2015). Combining results from both strands of literature allows us to obtain joint credit risk measures at a relatively high frequency, such as daily or weekly. Second, to balance the need for parsimony and flexibility, we consider a variant of the block equicorrelation structure for the covariance matrix in Engle and Kelly (2012); see also Christoffersen *et al.* (2014a) and Creal and Tsay (2015) for applications with DECO and grouped dependence structures, respectively. The version of the block equicorrelation model developed in this paper still allows us to draw on the machinery developed in the credit portfolio tail risk literature mentioned earlier. Third, an additional strand of literature investigates joint and conditional default dependence from a financial stability perspective; see, for example, Hartmann *et al.* (2007), Acharya *et al.* (2012) and Suh (2012). Finally, to introduce time variation into our econometric model specification, we endow our model with observation-driven dynamics based on the score of the conditional predictive log-density. Score-driven time-varying parameter models are actively researched; see for example, Creal *et al.* (2011, 2013, 2014), Harvey (2013) and Oh and Patton (2014).[2]

The two papers that are most closely related to ours are Oh and Patton (2014) and Christoffersen *et al.* (2014b). Oh and Patton (2014) propose a class of dynamic copula factor models for high dimensions, which facilitates the estimation of a wide variety of systemic risk measures. Firms in their framework load on a common factor that has the skewed Student's *t* density of Hansen (1994). Additional idiosyncratic shocks are modeled by a symmetric-*t* distribution. They use their model for studying the risk from a large number of US corporates. Time-varying dependence is modeled in a score-driven way, as in this paper. Our study differs in that we use the generalized hyperbolic skewed-*t* (GHST) distribution for both the marginal and the copula modeling; this means that our factor copula has a nonlinear two-factor rather than a linear one-factor structure. Moreover, our application focuses on the euro area financial sector during the financial and sovereign debt crisis, and on the impact evaluation of central bank unconventional monetary policy measures during that time. Finally, we introduce and discuss the efficient evaluation of semi-analytic risk measures. Christoffersen *et al.* (2014b) study diversification benefits among US corporates, and use Hansen's (1994) *t*-distribution for the innovations in univariate generalized autoregressive conditional heteroskedasticity (GARCH) models. Their GHST copula is the same as used in our paper. Our paper differs in that we provide a score-driven approach to the modeling of dynamic dependence, which is particularly attractive in a non-Gaussian context when data are fat-tailed and skewed; see Zhang *et al.* (2011) and Blasques *et al.* (2015). We do not rely exclusively on composite likelihood methods, but employ them only to provide alternative points of comparison from a robustness perspective. Composite likelihood techniques are feasible in high dimensions, but also statistically inefficient. Instead, we proceed by proposing block equicorrelation models, at least for settings where prior information is available on which subsets of firms are likely to move together as a group, and risks need to be evaluated in a computationally efficient way. We explicitly highlight what information is lost when moving from a full to an equicorrelation copula structure. Finally, we introduce semi-analytic risk measures and focus our application on the euro area sovereign debt crisis and non-standard monetary policy measures.

Section 2 introduces our statistical framework, the dynamic GHST block equicorrelation model, and discusses parameter estimation. Section 3 demonstrates how a conditional law of large numbers can be applied to reliably and quickly compute portfolio risk measures in a GHST factor copula setting. Section 4 applies the modeling framework to the euro area financial sector during the financial and sovereign debt crisis. Section 5 concludes. A supplementary web Appendix is provided as supporting information, which presents proofs and additional results.

---

[2] We refer to http://www.gasmodel.com for an extensive enumeration of recent work in this area. Computer code for this paper is available from this source.

2. STATISTICAL MODEL

2.1. The Dynamic Generalized Hyperbolic Skewed-t Copula Model

Following Lucas *et al.* (2014), we consider a copula model based on the GHST distribution. Let

$$y_{it} = (\varsigma_t - \mu_\varsigma)\gamma_i + \sqrt{\varsigma_t}\, \tilde{\Sigma}_{it}^{1/2} \epsilon_t, \quad i = 1, \ldots, N \tag{1}$$

where $y_t = (y_{1t}, \ldots, y_{Nt})'$ is a vector of firm-specific log asset values, $\varsigma_t \in \mathbb{R}^+$ is an inverse-Gamma distributed common risk factor that affects all firms simultaneously, $\varsigma_t \sim \text{IG}\left(\frac{\nu}{2}, \frac{\nu}{2}\right)$, $\gamma \in \mathbb{R}^N$ is a vector controlling the skewness of the copula, $\tilde{\Sigma}_{it}$ is the ith row of the GHST copula scale matrix $\tilde{\Sigma}_t \in \mathbb{R}^{N \times N}$, and $\epsilon_t \in \mathbb{R}^N$ is a vector of standard normally distributed risk factors. We assume that the two random variables ς_t and ϵ_t are independent and set $\mu_\varsigma = \text{E}[\varsigma_t] = \nu/(\nu-2)$, such that y_{it} has zero mean if $\nu > 2$. Matrix $\tilde{\Sigma}_t^{1/2}$ is the Choleski matrix square root of the scale matrix $\tilde{\Sigma}_t$.

In our copula framework, a firm defaults if its log asset value y_{it} falls below its default threshold y_{it}^*. Below, we approximate the copula of the (unobserved) log asset values with the copula estimated from equity returns. Implicitly, this assumes that default is sensitive to changes in the market value of equity, while the market value of debt remains approximately constant over a relevant horizon. The cross-sectional dependence in defaults captured by equation (1) thus stems from two sources: common exposures to the normally distributed risk factors ϵ_t as captured by the time-varying matrix $\tilde{\Sigma}_t$; and an additional common exposure to the scalar risk factor ς_t. The former captures connectedness through correlations, while the latter captures such effects through the tail dependence of the copula. To see this, note that if ς_t is non-random the first term in equation (1) drops out of the equation and there is zero tail dependence. Conversely, if ς_t is large, all asset values are affected at the same time, making joint defaults of two or more firms more likely.

Earlier applications of the GHST distribution to financial and economic data include Mencía and Sentana (2005), and Aas and Haff (2006). Alternative skewed t-distributions have been proposed as well, such as Branco and Dey (2001), Gupta (2003), Azzalini and Capitanio (2003) and Bauwens and Laurent (2005); see also the overview of Aas and Haff (2006). The GHST distribution closely connects to a continuous-time finance literature utilizing Lévy processes for stock price processes and firm asset values; see Bibby and Sørensen (2003) for a survey.

We denote the 1-year-ahead default probability for firm i at time t as p_{it}, such that

$$p_{it} = \Pr\left[y_{it} < y_{it}^*\right] = F_{it}(y_{it}^*) \Leftrightarrow y_{it}^* = F_{it}^{-1}(p_{it}) \tag{2}$$

where F_{it} is the univariate GHST cumulative distribution function (cdf) of y_{it}. In our application, we assume that we observe p_{it} as the expected default frequency of firm i reported at time t by Moody's Analytics (formerly Moody's KMV). Instead of focusing on the individual default probabilities p_{it}, our focus is on the time-varying *joint* probabilities, such as $\Pr\left[y_{it} < y_{it}^*, y_{jt} < y_{jt}^*\right]$, and on conditional probabilities such as $\Pr\left[y_{it} < y_{it}^* | y_{jt} < y_{jt}^*\right]$, for firms $i \neq j$. Below, we first develop a dynamic version of the GHST copula model. Then we consider a dynamic (block) equicorrelation version of the model in the spirit of Engle and Kelly (2012), which turns out to be particularly useful to study joint and conditional default probabilities in a parsimonious way for large-dimensional systems.

2.2. The Dynamic GHST Model

To describe the dynamics of the scale parameter $\tilde{\Sigma}_t$ in the GHST model (1), we use the generalized autoregressive score (GAS) dynamics as proposed in Creal *et al.* (2011, 2013); see also Harvey (2013). These dynamics easily adapt to the skewed and fat-tailed nature of the GHST density and improve the

stability of dynamic volatility and correlation estimates; see Blasques *et al.* (2015). Our version of the model is different from that of Lucas *et al.* (2014) owing to a different parametrization. We consider the scale matrix (rather than the covariance matrix) in order to fully employ the block equicorrelation structure later on for our conditional law of large numbers result.

To derive the score dynamics for the GHST model, we need the conditional density of y_t, which we parametrize as

$$p(y_t; \tilde{\Sigma}_t, \gamma, \nu) = \frac{2 (\nu/2)^{\nu/2}}{\Gamma(\nu/2) |2\pi \tilde{\Sigma}_t|^{1/2}} \cdot \frac{K_{(\nu+N)/2}\left(\sqrt{d(y_t) \cdot d(\gamma)}\right) e^{\gamma' \tilde{\Sigma}_t^{-1}(y_t - \tilde{\mu})}}{(d(y_t)/d(\gamma))^{(\nu+N)/4}} \tag{3}$$

$$d(y_t) = \nu + (y_t - \tilde{\mu})' \tilde{\Sigma}_t^{-1} (y_t - \tilde{\mu}) \tag{4}$$

$$d(\gamma) = \gamma' \tilde{\Sigma}_t^{-1} \gamma, \quad \tilde{\mu} = -\frac{\nu}{\nu - 2} \gamma \tag{5}$$

where $K_a(b)$ is the modified Bessel function of the second kind; see Bibby and Sørensen (2003). The parameters $\gamma = (\gamma_1, \ldots, \gamma_N)' \in \mathbb{R}^N$ and $\nu \in \mathbb{R}^+$ are the skewness and degrees of freedom (or kurtosis) parameter, respectively, while $\tilde{\mu}$ and $\tilde{\Sigma}_t$ denote the location vector and scale matrix, respectively. Note that if y_t has a multivariate GHST distribution with parameters $\tilde{\mu}$, $\tilde{\Sigma}$, γ and ν as given in equation (3), then $Ay_t + b$ for some matrix A and vector b also has a GHST distribution, with parameters $A\tilde{\mu} + b$, $A\tilde{\Sigma}A'$, $A\gamma$ and ν. In particular, the marginal distributions of y_{it} also have a GHST distribution. The GHST density (3) nests the symmetric-t ($\gamma = 0$) and multivariate normal ($\gamma = 0$ and $\nu \to \infty$) distributions as special cases.

We parametrize the time-varying matrix $\tilde{\Sigma}_t$ as in Engle (2002), i.e.

$$\tilde{\Sigma}_t = D(f_t)\tilde{R}(f_t)D(f_t) \tag{6}$$

where f_t is a vector of time-varying parameters, $D(f_t)$ is a diagonal matrix holding the scale parameters of y_{it} and $\tilde{R}(f_t)$ captures the dependence parameters. In our current copula setup, we use univariate models for $D(f_t)$, and the multivariate model for $\tilde{R}(f_t)$. The web Appendix provides further details on our univariate volatility modeling approach for $D(f_t)$. In the remainder of this section, we concentrate on the matrix $\tilde{R}(f_t)$.

Following Creal *et al.* (2011, 2013), we endow f_t with score-driven dynamics using the derivative of the log conditional observation density (3). The transition dynamics for f_t are given by

$$f_{t+1} = \tilde{\omega} + \sum_{i=0}^{p-1} A_i s_{t-i} + \sum_{j=0}^{q-1} B_j f_{t-j} \tag{7}$$

$$s_t = \mathcal{S}_t \nabla_t, \quad \nabla_t = \partial \ln p(y_t | \mathcal{F}_{t-1}; f_t, \theta)/\partial f_t \tag{8}$$

where $\tilde{\omega} = \tilde{\omega}(\theta)$ is a vector of fixed intercepts, $A_i = A_i(\theta)$ and $B_j = B_i(\theta)$ are fixed parameter matrices that depend on the vector θ containing all time-invariant parameters in the model, and $\mathcal{S}_t$ a scaling function.

The key element in equation (7) is the scaled score s_t. If y_t has a zero mean GHST distribution $p(y_t; \tilde{\Sigma}_t, \gamma, \nu)$ and the time-varying scale matrix is driven by equations (7)–(8), then the score is given by

$$\nabla_t = \Psi_t' H_t' \text{vec}\left(w_t \cdot (y_t - \tilde{\mu})(y_t - \tilde{\mu})' - 0.5\tilde{\Sigma}_t - \gamma(y_t - \tilde{\mu})' - \check{w}_t \cdot \gamma\gamma'\right) \tag{9}$$

where

$$w_t = \frac{\nu + N}{4 d(y_t)} - \frac{k'_{0.5(\nu+N)}\left(\sqrt{d(y_t)d(\gamma)}\right)}{2\sqrt{d(y_t)/d(\gamma)}},$$

$$\check{w}_t = \frac{\nu + N}{4 d(\gamma)} + \frac{k'_{0.5(\nu+N)}\left(\sqrt{d(y_t)d(\gamma)}\right)}{2\sqrt{d(\gamma)/d(y_t)}},$$

$$H_t = \tilde{\Sigma}_t^{-1} \otimes \tilde{\Sigma}_t^{-1}, \quad \Psi_t = \frac{\partial \text{vec}(\tilde{\Sigma}_t)'}{\partial f_t}$$

and where $k_\nu(\cdot) = \ln K_\nu(\cdot)$ with first derivative $k'_\nu(\cdot)$. The matrices Ψ_t and H_t are time-varying and parametrization-specific; both matrices depend on f_t but not on the data y_t. We refer to the web Appendix for a derivation of equation (9).

Equation (8) reveals the key feature of the score-driven specification. In essence, the score-driven mechanism takes a Gauss–Newton improvement step for the scale matrix to better fit the most recent observation. Equation (8) shows that f_t reacts to deviations between $\tilde{\Sigma}_t$ and the observed $(y_t - \tilde{\mu})(y_t - \tilde{\mu})'$. The reaction is asymmetric if $\gamma \neq 0$, in which case there is also a reaction to the level $(y_t - \tilde{\mu})$ itself. The reaction to $(y_t - \tilde{\mu})(y_t - \tilde{\mu})'$ is modified by the weight w_t. If $\nu < \infty$, the GHST distribution is fat-tailed and the weight decreases in the Mahalanobis distance $d(y_t)$; compare the discussion for the symmetric Student's t case in Creal *et al.* (2011). This feature gives the model a robustness flavor in that incidental large values of y_t have a limited impact on future volatilities and correlations. The remaining expressions for H_t and Ψ_t only serve to transform the dynamics of the covariance matrix in equation (6) into the dynamics of the unobserved factor f_t.

To scale the score in equation (8) we set the scaling matrix $\mathcal{S}_t$ equal to the inverse conditional Fisher information matrix of the symmetric Student's t-distribution; see the web Appendix for details. Zhang *et al.* (2011) demonstrate that this choice of scaling matrix results in a stable model that outperforms alternative models if the data are fat-tailed and skewed. We do not use the inverse conditional Fisher information matrix of the multivariate GHST distribution because it is not available in closed form.

2.3. Dynamic Block Equicorrelation Structure

As we want to use our model in the context of a large cross-sectional dimension to describe the joint (tail) risk dynamics in a large system of financial institutions, we refrain from modeling all dependence parameters in $\tilde{R}(f_t)$ individually. Instead, we adopt the approach of Engle and Kelly (2012) and impose a block equicorrelation structure on the matrix $\tilde{R}(f_t)$. By limiting the number of free parameters, we facilitate the estimation process to a large extent while retaining the ability to capture dynamic patterns in the dependence structure among financial firms, particularly in times of stress.

For the single block equicorrelation model, we assume that $\tilde{\Sigma}_t$ takes the form

$$\tilde{\Sigma}_t = (1 - \rho_t^2)\text{I}_N + \rho_t^2 \ell_N \ell_N' \tag{10}$$

where $\rho_t = (1 + \exp(-f_t))^{-1} \in (0, 1)$, and ℓ_N is a $N \times 1$ vector of ones. In this case, the expressions for Ψ_t and H_t simplify. In particular,

$$\Psi_t = \frac{\partial \text{vec}(\tilde{\Sigma}_t)'}{\partial f_t} = (\ell_{N^2} - \text{vec}(\text{I}_N)) \frac{2\exp(-f_t)}{(1 + \exp(-f_t))^3} \tag{11}$$

which implies that the score ∇_t reduces to a scalar process over time. We refer to the web Appendix for a derivation. We can easily generate the equicorrelation structure (10) from model (1) by specifying

$$y_{it} = \left(\varsigma_t - \mu_\varsigma\right)\gamma_i + \sqrt{\varsigma_t}\left(\rho_t\,\kappa_t + \sqrt{1 - \rho_t^2}\,u_{it}\right) \tag{12}$$

where κ_t and u_{it} are two independent standard normal random variables. Note that equation (12) is a special case of equation (1). The logistic parametrization $\rho_t = (1 + \exp(f_t))^{-1}$ forces the correlation parameter to be in the unit interval, irrespective of the value of $f_t \in \mathbb{R}$. While the original equicorrelation specification of Engle and Kelly (2012) also allows for (slightly) negative equicorrelations, such values are typically unrealistic in the type of applications we consider later on. The parametrization with equicorrelation parameter $\rho_t^2 > 0$ therefore suffices for our current purposes.

A two-block equicorrelation model may be considered in settings for which the restriction of a single correlation parameter characterizing the entire scaling matrix $\tilde{\Sigma}_t$ might be too restrictive empirically. For example, we might want to allow for different dependence between financial firms in stressed and non-stressed countries in the context of the euro area sovereign debt crisis. For two blocks containing N_1 and N_2 firms, respectively, the two-block equicorrelation specification is given by

$$\tilde{\Sigma}_t = \begin{bmatrix} \left(1 - \rho_{1,t}^2\right)\mathrm{I}_{N_1} & 0 \\ 0 & \left(1 - \rho_{2,t}^2\right)\mathrm{I}_{N_2} \end{bmatrix} + \begin{pmatrix} \rho_{1,t}\ell_{N_1} \\ \rho_{2,t}\ell_{N_2} \end{pmatrix} \cdot \begin{pmatrix} \rho_{1,t}\ell'_{N_1} & \rho_{2,t}\ell'_{N_2} \end{pmatrix} \tag{13}$$

The two-block equicorrelation structure (13) differs from the setup in Engle and Kelly (2012) in that there is a direct relation between the equicorrelation in the off-diagonal blocks and the diagonal blocks. The main advantage of this specification is that, conditional on ς_t, equation (13) preserves the Vasicek (1987) single-factor credit risk structure of y_t. In particular, specification (12) depends on two common factors only, namely κ_t and ς_t, and conditionally on ς_t the model is linear. We use this feature extensively to compute joint and conditional risk measures fast and reliably in Section 3 in settings where standard simulation methods quickly become inefficient.[3] If $\tilde{\Sigma}_t$ is given by equation (13) with $\rho_{j,t} = (1 + \exp(-f_{j,t}))^{-1}$ for $j = 1, 2$, then the time-varying factor $f_t = (f_{1,t}, f_{2,t})' \in \mathbb{R}^{2\times 1}$ follows equation (9), with

$$
\begin{aligned}
\Psi_t &= \frac{\partial \mathrm{vec}(\tilde{\Sigma}_t)'}{\partial f_t} = \frac{\partial \mathrm{vec}(\tilde{\Sigma}_t)'}{\partial \rho_t}\frac{\mathrm{d}\rho_t'}{\mathrm{d}f_t}, \\
\frac{\mathrm{d}\rho_t'}{\mathrm{d}f_t} &= \begin{pmatrix} \frac{\exp(-f_{1,t})}{(1+\exp(-f_{1,t}))^2} & 0 \\ 0 & \frac{\exp(-f_{2,t})}{(1+\exp(-f_{2,t}))^2} \end{pmatrix}, \\
\frac{\partial \mathrm{vec}(\tilde{\Sigma}_t)'}{\partial \rho_t} &= \begin{pmatrix} \mathrm{vec}\begin{pmatrix} \mathrm{I}_{N_1} & 0 \\ 0 & 0 \end{pmatrix}, & \mathrm{vec}\begin{pmatrix} 0 & 0 \\ 0 & \mathrm{I}_{N_2} \end{pmatrix} \end{pmatrix} \cdot \begin{pmatrix} -2\rho_{1,t} & 0 \\ 0 & -2\rho_{2,t} \end{pmatrix} \\
&\quad + \left(\left(\begin{pmatrix} \rho_{1,t}\ell_{N_1} \\ \rho_{2,t}\ell_{N_2} \end{pmatrix} \otimes \mathrm{I}_N + \mathrm{I}_N \otimes \begin{pmatrix} \rho_{1,t}\ell_{N_1} \\ \rho_{2,t}\ell_{N_2} \end{pmatrix}\right) \cdot \left(\begin{pmatrix} \ell_{N_1} \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ \ell_{N_2} \end{pmatrix}\right)\right)
\end{aligned} \tag{14}
$$

where $\rho_t = (\rho_{1,t}, \rho_{2,t})'$ and $N = N_1 + N_2$.[4]

The m-block equicorrelation structure is a straightforward generalization of the two-block case (14). Rather than providing the (lengthy) expressions in the main text, we refer to our web Appendix for

---

[3] We demonstrate in Section 4 that the tail risk measurements obtained from applying equation (10) or (13) are close to those based on a full correlation matrix analysis.

[4] We can introduce further flexibility to the model by extending the support of $\rho_{1,t}$ and $\rho_{2,t}$ from $(0, 1)$ to $(-1, 1)$ by defining $\rho_{j,t} = (\exp(f_{j,t}) - 1)/(\exp(f_{j,t}) + 1)$ for $j = 1, 2$. The advantage of this extension is that the correlations between blocks can now become negative, whereas the within-block correlation remains positive. This extension is not needed for our empirical application in Section 4, however.

the precise formulations. These formulations are used in our empirical analysis in Section 4, where we also consider a three-block equicorrelation specification.

2.4. Parameter Estimation

We can estimate the static parameters θ of the dynamic GHST model through standard maximum likelihood procedures. Parameter estimation is straightforward as the likelihood function is known in closed form using a standard prediction error decomposition. Deriving the asymptotic behavior for time-varying parameter models with GAS dynamics is non-trivial. We refer to Blasques *et al.* (2012, 2014) for details.

We split the estimation problem into two parts by adopting a copula perspective. As a result, the number of parameters that need to be estimated in each step is reduced substantially. In addition, the copula perspective has the advantage that we can add more flexibility to modeling the marginal distributions. For example, when working with a multivariate GHST density, all marginal distributions must have the same kurtosis parameter ν. By adopting a copula perspective, we can relax this restriction considerably.

The two stages of the estimation process can be summarized as follows. In a first step, we estimate univariate dynamic GHST models using the equity returns for each firm i; see the web Appendix for details. Based on the estimated univariate models with parameters $\tilde{\mu}_{it}$, $\tilde{\sigma}_{it}$, γ_i, and ν_i, we transform the observations into their probability integral transforms $u_{it} \in [0, 1]$. In the second step, we estimate the matrix $\tilde{\Sigma}_t = \tilde{R}_t$ as parametrized in Section 2.3, using the probability integral transforms u_{it} constructed in the first step. For the single-block equicorrelation model, the GHST copula parameters are 0, $\tilde{R}_t$, $\gamma \cdot (1, \ldots, 1)'$ for $\gamma \in \mathbb{R}$ and ν, respectively. The respective static parameter vector θ includes γ, ν and $\tilde{\omega}$, A_j and B_j of the dynamic equation (7). For equicorrelation models with multiple blocks, we consider block-specific skewness parameters γ_j.

3. JOINT AND CONDITIONAL RISK MEASURES

This section defines joint and conditional risk measures and demonstrates how to compute these efficiently and reliably based on the application of a conditional law of large numbers (cLLN). Using a cLLN in the credit risk context was popularized by Vasicek (1987) and studied further, for example, in Gordy (2000, 2003) and Lucas *et al.* (2001, 2003).

Conceptually, the simplest way to compute joint and conditional default probabilities is based on Monte Carlo simulations of firms' asset values. For example, one can generate many paths for the joint evolution of $(y_{1t}, \ldots, y_{Nt})$ and check how many simulations lie in a joint distress region of the type $\{y_t \mid y_{jt} < y_{jt}^* \ \forall j \in J\}$ for some set of firms $J \subset \{1, 2, \ldots, N\}$, where y_{jt}^* denotes the default threshold for firm j at time t. Such a simulation-based approach quickly becomes inefficient if the cross-sectional dimension of the data and the number of firms considered in the set J become large: because marginal default probabilities are typically small, we need a large number of simulations to obtain a sufficient number of realizations of joint defaults, particularly if three, four or even more joint defaults are considered. A partial remedy could be to use simulations based on importance sampling methods as in Glassermann and Li (2005), but the computational burden would remain high compared to the simple semi-analytic approach proposed in this section.

The semi-analytic cLLN approximation we propose is based on the observation that equation (12) is a nonlinear two-factor model. More specifically, conditional on the common factor ς_t model (12) simplifies to a (heterogeneous) Gaussian one-factor model as in Vasicek (1987). This holds even if we replace ρ_t in equation (12) by ρ_{it} using the block equicorrelation structure from Section 2. We exploit this feature to obtain reliable alternative risk measures that can be evaluated semi-analytically.

We define our joint tail risk measure (JRM) as the time-varying probability that a certain fraction of firms defaults over a pre-specified period. Let $D_{N,t}$ denote the fraction of firms defaulting over period t, e.g. $D_{N,t} = 5\%$, with

$$D_{N,t} = \frac{1}{N} \sum_{i=1}^{N} 1\{y_{it} < y_{it}^*\} \tag{15}$$

Since the indicators $1\{y_{it} < y_{it}^*\}$ are conditionally independent given κ_t and ς_t, we can apply a conditional law of large numbers to obtain

$$D_{N,t} \approx \frac{1}{N} \sum_{i=1}^{N} \mathrm{E}[1\{y_{it} < y_{it}^*\}|\kappa_t, \varsigma_t] = \frac{1}{N} \sum_{i=1}^{N} \mathrm{P}\left[y_{it} < y_{it}^*|\kappa_t, \varsigma_t\right] \equiv C_{N,t} \tag{16}$$

for large N, as under standard regularity conditions $|D_{N,t} - C_{N,t}| \overset{a.s.}{\to} 0$ for $N \to \infty$; see, for example, Vasicek (1987), Gordy (2000, 2003) and Lucas *et al.* (2001). Note that

$$\mathrm{P}[y_{it} < y_{it}^*|\kappa_t, \varsigma_t] = \Phi\left(\frac{y_{it}^* - (\varsigma_t - \mu_\varsigma)\gamma_i - \sqrt{\varsigma_t}\,\rho_t \kappa_t}{\sqrt{\varsigma_t\,(1 - \rho_t^2)}}\right) \tag{17}$$

where $\Phi(\cdot)$ denotes the cumulative standard normal distribution. Also note that $C_{N,t}$ is a function of the random variables κ_t and ς_t only, and not of u_{it} in (12). We now define the joint tail risk measure as

$$\mathrm{P}(D_{N,t} > \bar{c}) \approx \mathrm{P}(C_{N,t} > \bar{c}) = \mathrm{P}(C_{N,t}(\kappa_t, \varsigma_t) > \bar{c}) \equiv p_t \tag{18}$$

i.e. we approximate the probability that the fraction of credit portfolio defaults $D_{N,t}$ exceeds the threshold $\bar{c} \in [0, 1]$ by the quantity p_t. Following Vasicek (1987) and using the one-factor structure of equation (12) for given ς_t, we note that $C_{N,t}$ is monotonically decreasing in κ_t for $\rho_t > 0$. This is intuitive: an increase in κ_t (for instance, due to improved business cycle conditions) implies fewer defaults in the portfolio and vice versa. We exploit this to efficiently compute unique threshold levels $\kappa_{N,t}^*(\bar{c}, \varsigma^{(g)})$ for a number of grid points $\varsigma^{(g)}$, $g = 1, \ldots, G$. This can be done by solving the equations $C_{N,t}(\kappa_{N,t}^*(\bar{c}, \varsigma^{(g)}), \varsigma^{(g)}) = \bar{c}$ numerically for the threshold values $\kappa_{N,t}^*(\bar{c}, \varsigma^{(g)})$ for each grid point $\varsigma^{(g)}$ and time t. Given a grid of threshold values, we can then use standard numerical integration techniques to efficiently compute the joint default probability

$$p_t = \mathrm{P}(C_{N,t} > \bar{c}) = \int \mathrm{P}(\kappa_t < \kappa_{t,N}^*(\bar{c}, \varsigma_t)) p(\varsigma_t) \mathrm{d}\varsigma_t \tag{19}$$

Our second measure is a conditional tail risk measure (CRM). Let

$$D_{N-1,t}^{(-i)} = \frac{1}{N-1} \sum_{j \neq i} 1\left[y_{jt} < y_{jt}^*|\kappa_t, \varsigma_t\right] \approx \frac{1}{N-1} \sum_{j \neq i} \mathrm{P}[y_{jt} < y_{jt}^*|\kappa_t, \varsigma_t] \equiv C_{N-1,t}^{(-i)}$$

where $D_{N-1,t}^{(-i)}$ is the fraction of defaulted companies excluding firm i, which we approximate using the cLLN by $C_{N-1,t}^{(-i)}$. We define the CRM as the probability of $C_{N-1,t}^{(-i)}$ exceeding $\bar{c}^{(-i)}$ conditional on the default of firm i, i.e.

$$
\begin{aligned}
\mathrm{P}\left(C_{N-1,t}^{(-i)} > \bar{c}^{(-i)}|y_{it} < y_{it}^*\right) &= p_{it}^{-1} \cdot \mathrm{P}\left(C_{N-1,t}^{(-i)} > \bar{c}^{(-i)}, \; y_{it} < y_{it}^*\right) \\
&= p_{it}^{-1} \cdot \int \mathrm{P}\left(\kappa_t < \kappa_{N-1,t}^*\left(\bar{c}^{(-i)}, \varsigma_t\right), y_{it} < y_{it}^* \Big| \varsigma_t\right) p(\varsigma_t) \mathrm{d}\varsigma_t \\
&= p_{it}^{-1} \cdot \int \Phi_2\left(\kappa_{N-1,t}^*\left(\bar{c}^{(-i)}, \varsigma_t\right), y_{it}^{**}(\varsigma_t); \rho_t\right) p(\varsigma_t) \mathrm{d}\varsigma_t
\end{aligned}
\tag{20}
$$

where $y_{it}^{**}(\varsigma_t) = (y_{it}^* - (\varsigma_t - \mu_\varsigma)\gamma_i)/\sqrt{\varsigma_t}$, and $\Phi_2(\cdot, \cdot; \rho_t)$ denotes the cumulative distribution function of the bivariate normal with standard normal marginals and correlation parameter ρ_t. To obtain the last equality in equation (20), note that the GHST distribution becomes Gaussian conditional on ς_t. The conditional probability (20) is a time-varying higher-frequency extension of the multivariate extreme spillovers measure of Hartmann *et al.* (2004, 2007).

Both the joint probability (19) and the conditional probability (20) can be computed quickly, using simple one-dimensional numerical integration techniques.[5] In addition, the model is easily extended to fit the m-block equicorrelation structure explained in Section 2.3. The fact that the $\rho_{i,t}$ parameters are different between blocks does not distort the one-factor structure à la Vasicek (1987). The above derivations remain valid if ρ_t is replaced by $\rho_{i,t}$, particularly in equations (17) and (20), and all computations remain of similar structure and speed.

4. EURO AREA FINANCIAL SECTOR JOINT TAIL RISK

We apply our model to a high-dimensional dataset of $N = 73$ euro area financial firms. We first present a preliminary analysis for a subsample of $N = 10$ large firms headquartered in different euro area countries. We then present the results for the entire dataset. This allows us to benchmark the results for the dynamic GHST block equicorrelation model against a model with a fully specified time-varying correlation matrix and to investigate the sensitivity of joint and conditional tail risk measures to the (block) equicorrelation assumption.

4.1. Equity and EDF Data

Our equity data come from Bloomberg. We use 73 listed financial firms that are located in 11 euro area countries: Austria (AT), Belgium (BE), Germany (DE), Spain (ES), Finland (FI), France (FR), Greece (GR), Ireland (IE), Italy (IT), the Netherlands (NL) and Portugal (PT). Firms were selected if (i) they were financial firms headquartered in the euro area, and (ii) were listed as of 2011:Q1 as a component of the STOXX Europe 600 index. For each firm, we construct recursively demeaned weekly equity returns. The sample comprises large commercial banks as well as large financial non-banks such as insurers and investment companies. The total panel covers 762 weeks from January 1999 to August 2013. The panel is unbalanced in that some data are missing in the first part of the sample. We assume that this is not related to the volatility dynamics or the credit risk mechanism in the data. The scores in equation (9) automatically correct for the unbalancedness of the data.

For the marginal default probabilities p_{it}, we use 1-year-ahead expected default frequencies (EDF) obtained from Moody's Analytics. EDFs are widely used measures of time-varying 1-year-ahead marginal default probabilities (see Duffie *et al.*, 2007). We do not use the EDFs to estimate the dependence structure. Rather, we only use the EDFs to calibrate the model at any time to current market perceptions of 1-year-ahead marginal default risk conditions for each firm in the sample. The copula as estimated from the equity data subsequently takes care of the dependence structure when computing the joint and conditional tail risk measures.

4.2. Small-Sample Study of 10 Banking Groups

Before presenting the results for the full sample of $N = 73$ institutions, we first study a geographically diversified subsample of 10 large financial firms from 10 euro area countries. We do this to study

---

[5] It is straightforward to add exposure weights e_i to the definition of $D_{N,t}$ in equation (15). The computations in that case remain equally efficient. If the exposures are very unevenly distributed, however, the approximation error of the cLLN in equation (16) might increase. To mitigate such an effect, one could try to implement a second-order expansion using a conditional central limit theorem rather than a cLLN only.

how the equicorrelation assumption may affect our joint and conditional risk measures compared to a model with an unrestricted correlation structure. The subsample contains no missing observations and consists of Erste Bank Group (AT), Dexia (BE), Deutsche Bank (DE), Santander (ES), BNP Paribas (FR), National Bank of Greece (GR), Bank of Ireland (IE), UniCredito (IT), ING (NL) and Banco Comercial Portugues (PT). We distinguish firms across different countries given the interdependence of bank risk and sovereign risk as an important feature of the euro area sovereign debt crisis; see, for example, ECB (2012, 2014).

4.2.1. Dependence Modeling

This section compares correlation estimates across four different dependence models. Since we use a copula approach, the models share the same structure for the univariate volatilities. The descriptive statistics in the web Appendix reveal that the return data are significantly negatively skewed and fat-tailed. We therefore use the dynamic GHST model for the marginals.

We consider three score-driven equicorrelation specifications with one, two and three blocks, respectively. The two-block model distinguishes firms in countries that experienced pronounced stress during the sovereign debt crisis (Greece, Ireland, Italy, Portugal, Spain) and firms headquartered in non-stressed countries; see Eser and Schwaab (2015) for a similar grouping of countries. The three-block model further distinguishes firms in smaller stressed countries that entered the sovereign debt crisis earlier (Greece, Portugal, Ireland), and larger stressed countries that entered the sovereign debt crisis at a relatively later stage (Spain and Italy); see ECB (2014). As a benchmark, we consider a model with a full correlation matrix with DCC dynamics as in Christoffersen *et al.* (2014b), which we estimate through composite likelihood methods. As in the block equicorrelation specifications, the DCC full correlation matrix model uses a common scalar skewness parameter γ in the GHST copula. For $N = 10$ firms the model with full correlation matrix contains 45 pairwise correlation coefficients, and thus 45 dynamic factors. Correlation targeting is used to estimate the intercepts in the transition equation for the correlations.

The first columns in Table I report parameter estimates and model selection criteria for the one-, two- and three-block equicorrelation models for $N = 10$. The correlation dynamics are highly persistent in all specifications given the high values of B, or of $A + B$ in the case of the DCC specification. The unconditional correlation levels are monotonically increasing in the factor intercept parameters $\tilde{\omega}_j$ (see equation (7)), and are thus highest for firms in non-stressed countries (block 1). For $N = 10$, both the degrees of freedom parameter and the common skewness parameter of the GHST copula have the same sign and similar magnitudes for the block equicorrelation models (BEq[m]) and the DCC specification.

Figure 1 plots the estimated dynamic correlations. The top left panel plots the single block equicorrelation along with the average pairwise correlations of each firm with the other (nine) firms. The latter correlation estimates are based on the full correlation model with 45 time-varying parameters. The correlation estimates reveal a pronounced commonality in the correlation dynamics. This is intuitive, as we model firms from the same industry which operate in a single-currency area and are subject to similar regulatory requirements.[6] All correlations tend to increase over the sample period, possibly reflecting gradual financial integration and economic convergence in the euro area following the inception of the euro in 1999. All correlations remain elevated during the global financial crisis from 2008 to 2010, and peak at a time when Greece, Ireland and Portugal needed the assistance of third parties, such as the EU and the IMF in mid 2010 (see ECB, 2014). Such shared correlation dynamics can be captured simply and conveniently by block equicorrelation structures.

We can capture a larger share of the cross-sectional dispersion in correlations when we allow for multiple blocks. The top right panel of Figure 1 plots the dynamic correlation estimates for two groups.

---

[6] A principal components analysis of the 45 correlation pairs from the full correlation model suggests that the first three components explain 55.6%, 16% and 8.5% of the total variation, respectively. The first two factors therefore explain approximately 72% of the total variation in correlations.

Table I. Parameter estimates for the copula models

	10 firms				73 firms		
	GAS-BEq [1]	GAS-BEq [2]	GAS-BEq [3]	DCC-CL	GAS-BEq [1]	GAS-BEq [1]-t	GAS-BEq [3]
A	0.207	0.109	0.093	0.023	0.502	0.252	0.152
	(0.059)	(0.022)	(0.015)	(0.002)	(0.054)	(0.039)	(0.018)
B	0.983	0.992	0.993	0.963	0.990	0.986	0.996
	(0.014)	(0.007)	(0.005)	(0.003)	(0.009)	(0.013)	(0.004)
ω_1	0.581	0.778	0.745		0.089	0.276	0.006
	(0.214)	(0.272)	(0.273)		(0.336)	(0.288)	(0.311)
ω_2		0.231	0.977				0.339
		(0.305)	(0.415)				(0.302)
ω_3			-0.301				-0.542
			(0.422)				(0.448)
ν	15.497	16.022	16.540	13.176	28.838	30.135	27.755
	(1.156)	(2.074)	(1.566)	(1.212)	(2.579)	(1.959)	(1.538)
γ_1	-0.156	-0.118	-0.100	-0.221	-0.452		-0.475
	(0.073)	(0.095)	(0.084)	(0.034)	(0.070)		(0.062)
γ_2		-0.219	-0.126				-0.318
		(0.092)	(0.108)				(0.069)
γ_3			-0.261				-0.451
			(0.097)				(0.078)
Log-lik.	1509.2	1549.2	1644.9	1736.3	10945.4	10911.0	11047.6
AIC	-3008.4	-3084.5	-3271.7	-3464.7	-21880.7	-21812.0	-22077.2
BIC	-2976.1	-3039.1	-3213.4	-3438.8	-21848.4	-21779.6	-22019.0

Note: Parameter estimates for multivariate GAS-GHST models for financial firms' equity returns. The left-hand and right-hand blocks refer to the copula models for $N = 10$ and $N = 73$ firms, respectively. Univariate GAS-GHST models are used to model the marginal volatilities. The two-block model distinguishes between firms from non-stressed countries in the euro area, i.e. Austria, Belgium, France, Germany and the Netherlands (block 1), and firms from the remaining stressed countries, i.e. Greece, Portugal, Spain, Italy, Ireland (block 2). The three-block model further distinguishes between financial firms from non-stressed countries (block 1), larger stressed countries (Spain, Italy; block 2), and smaller stressed countries (Greece, Portugal, Ireland; block 3). Standard errors are in parentheses. Column DCC-CL reports parameter estimates obtained by maximizing a composite likelihood. The reported DCC log-likelihood is the full (non-composite) log-likelihood evaluated at the estimates obtained from maximizing the composite likelihood. We report the unconditional factor mean $\omega_j = \tilde{\omega}_j/(1 - B)$; see equation (7). Standard errors for the time-invariant parameters are constructed from the numerical second derivatives of the log-likelihood function.

The first group contains the Bank of Ireland, Banco Comercial Portugues, Santander, National Bank of Greece and UniCredito. The second group includes BNP Paribas, Deutsche Bank, Dexia, Erste Bank Group and ING. The overall dependence dynamics are similar.[7] The bottom left panel plots the correlation estimates for the three-block model, which allows for further cross-sectional dispersion in correlation estimates. Finally, the bottom right panel compares the average correlations across models. In addition, we provide average correlations estimated from a 52-week rolling window. The average correlations are similar across all models. Only relatively minor deviations are observed. We conclude that, despite its restrictive nature, the equicorrelation model reliably estimates the salient trends in average correlation dynamics.

Although the block equicorrelation models work well in capturing the average correlations, a substantial share of the cross-sectional dispersion in correlations may be lost when using block equicorrelation structures; see the top left panel of Figure 1. This may or may not matter when evaluating joint and conditional credit risk measures. To investigate how much cross-sectional dispersion in correlations is lost, the upper panel of Figure 2 plots R^2 statistics that correspond to repeated *cross-sectional* regressions of 45 correlation pairs from the GHST full-correlation-matrix model on a constant and the corresponding correlation estimates from a two-block and three-block equicorrelation model, respectively. By construction, the one-block equicorrelation model is unable to account for any

---

[7] Perhaps surprisingly, the correlation between firms from non-stressed countries lies above that of firms in stressed countries, also before the financial and sovereign debt crisis starting 2008. It is therefore probably related to different degrees of financial integration, rather than to shared exposure to heightened market turmoil in a crisis.
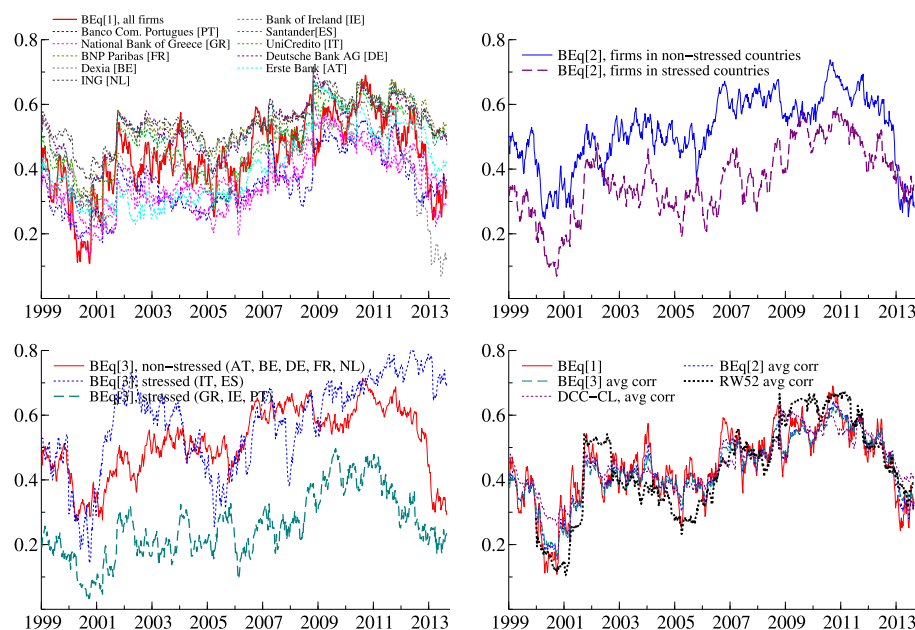
Figure 1. Filtered correlation estimates for $N = 10$ firms. The top left panel plots the correlation estimates based on the BEq[1] model, along with the average pairwise correlation of each firm with the other nine firms. The latter are based on the specification with full correlation matrix with 45 time-varying parameters and DCC dynamics. The top right and bottom left panels plot the block equicorrelation estimates based on the two-block (BEq[2]) and three-block (BEq[3]) model specification, respectively. The bottom right panel compares the average correlation estimates from the one-block, two-block and three-block models to the average (over 45 pairs) correlation estimates from a GHST full correlation DCC model, as well as the average correlation based on a 52-week rolling window. This figure is available in color online at wileyonlinelibrary.com/journal/jae

cross-sectional variation in correlations, as it collapses the latter to a single number (i.e. a regression on a constant). For the current sample, the two-block and three-block equicorrelation models are able to account for approximately 20% and 50% of the cross-sectional dispersion in correlations.

The lower panel of Figure 2 plots kernel density estimates and a histogram of *time series R^2* statistics. These correspond to a histogram of 45 time series regressions, with $T = 762$ weekly observations each, of GHST full-correlation-matrix model correlations on the corresponding estimates from a one-block, two-block and three-block equicorrelation specification, respectively. The one-block, two-block and three-block equicorrelation models are able to account for approximately 30–50%, 40–60% and 50–70% of the time series variation in the full correlation estimates, respectively. Both findings confirm that the simplified equicorrelation assumption can capture a large part of the variation in correlations. In the next section we highlight that the percentage of variation captured by the BEq[m] models appears adequate for reliable estimates of our systemic risk measures.

4.2.2. Joint and Conditional Probabilities

This section compares joint and conditional default probabilities as defined in Section 3 and as implied by different copula model specifications. We argue that the cLLN works well even when the data dimension is very small. This is because it eliminates a source of risk (idiosyncratic risk) that does not matter in the tail of the credit loss distribution. Only κ_t and ς_t are common to all firms and drive extreme default clustering. In addition, block equicorrelation and full correlation models lead to approximately similar patterns for joint and (to a lesser degree) also for conditional tail probability estimates.
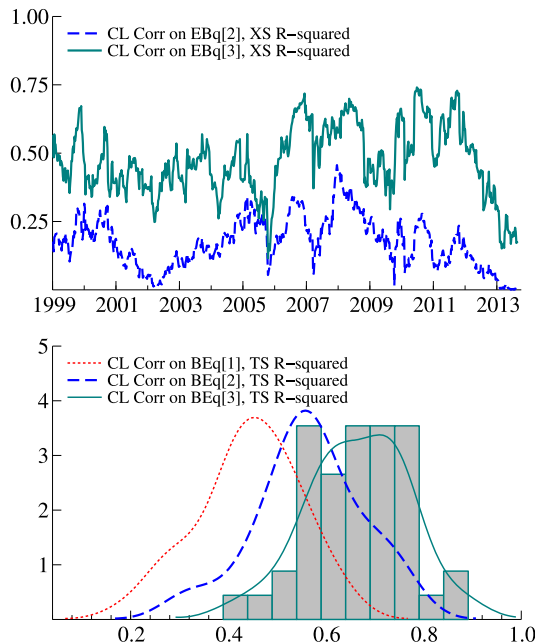
Figure 2. R^2 statistics. The upper panel plots R^2 statistics corresponding to repeated (over time) *cross-sectional* regressions of 45 correlation pairs from the $N = 10$ GHST full-correlation-matrix model on a constant and the corresponding correlation estimates from a two-block and three-block equicorrelation model. The lower panel plots kernel density estimates and a histogram of 45 R^2 statistics corresponding to *time series* regressions with $T = 762$ observations each of correlations from a GHST full-correlation-matrix model on the corresponding correlation estimates from a one-block, two-block and three-block equicorrelation model, respectively. The histogram refers to R^2's from the three-block equicorrelation model. This figure is available in color online at wileyonlinelibrary.com/journal/jae

We compare the full correlation specification with a one-block and three-block copula model. In addition, we assess the adequacy of the cLLN approximation by comparing it to a simulation-based approach to compute joint and conditional default probabilities. The default thresholds are obtained by inverting the GHST distribution function at the observed EDF levels. For the simulation-based computations, we use 500,000 simulations, each at time t, where we save the number and identities of the defaulted firms in each simulation.

The top left and right panels in Figure 3 refer to the probability of observing three or more defaults (out of 10 firms) over a 1-year-ahead horizon. The left-hand panel compares the one-block specification with the full-correlation-matrix outcome; the right-hand panel considers the three-block specification. Risk measures are either simulated or computed semi-analytically based on equation (19). As an important finding, the joint default probabilities are similar in each of the six cases. The losses from moving from a full correlation matrix to a much simpler equicorrelation structure are generally small. Also the losses from applying the cLLN approximation *vis-á-vis* the simulation approach are small in these cases, even though $N = 10$ is far from infinity.

The bottom panels in Figure 3 plot the conditional probability of three or more (out of nine possible) credit events given a credit event for a specific financial firm, averaged over all 10 firms. Again, the left- and right-hand panels refer to the one-block and three-block model, respectively. In the one-block case, the loss in fit from moving from a full-correlation-matrix model to the equicorrelation model is more pronounced. This is in line with the R^2 results shown in Figure 2. The cLLN approximation, however, works relatively well. In the three-block case, the respective loss in fit is less pronounced. Here, however, the cLLN approximation appears not to work as well. This reveals an interesting trade-off when
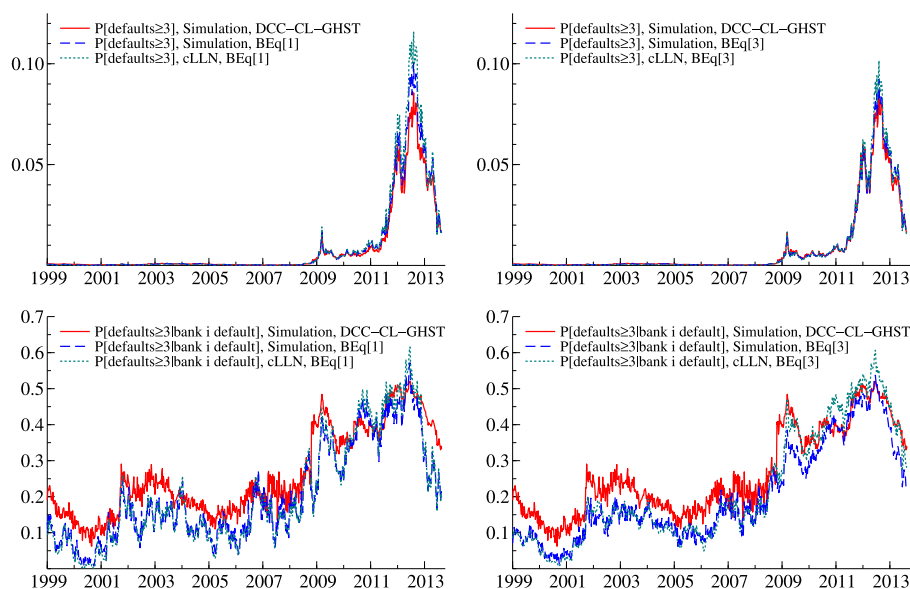
Figure 3. Joint and conditional probabilities. The two top panels plot the joint default probability (19) of three or more financial firm defaults out of 10. The two bottom panels plot the (average) conditional probability (20), which is the probability of three or more (out of nine possible) defaults given a default of a specific financial firm, averaged over all 10 firms. The probabilities are either computed using 500,000 simulations at each point in time t, or alternatively using the cLLN approximation as discussed in Section 3 for block equicorrelation models. The left-hand (top and bottom) panels are based on the estimated one-block equicorrelation model, while the right-hand (top and bottom) panels are based on the estimated three-block equicorrelation model. This figure is available in color online at wileyonlinelibrary.com/journal/jae

choosing the number of blocks. On the one hand, increasing the number of blocks can help in capturing more of the cross-sectional dispersion in correlations and increases the fit in the time dimension (see Figure 2). On the other hand, increasing the number of blocks also means fewer firms in each block, which in turn implies that the cLLN approximation does not work as well within each block.[8] Finally, the simulation-based approach can suffer from sizable simulation noise in non-crisis periods (when marginal default probabilities are low, see bottom panels before 2007). No such problems are encountered for the cLLN based approximation.

4.3. All 73 Euro Area Financial Firms

Given the encouraging preliminary results for the subsample of $N = 10$ institutions, we now turn to the joint and conditional tail risk measures based on the full panel of 73 large financial-sector firms. The sample contains commercial banks as well as financial non-banks such as insurers and investment companies. All are listed at a stock exchange.[9] Based on descriptive statistics as presented in the web Appendix, univariate GHST models seem appropriate given the skewness and kurtosis features of the equity return data.

---

[8] Recall that the three-block model contains five, three and two firms in each block. The cLLN hardly applies in the latter two cases, particularly given the low loading to the common κ_t for the stressed small countries shown in Figure 1.

[9] Freezing the set of firms as the constituents of a broad-based equity index in 2011:Q1 means that we may underestimate total euro area financial sector risk prior to this date (due to sample selection; weaker firms may have dropped out of the index by then). While this concern should be kept in mind, it is unlikely to be large, as most financial institutions under stress during the financial crisis continued to operate, also due to substantial government aid and the extension of public sector guarantees. Sample selection is no issue after 2011:Q1.

4.3.1. Dependence Modeling

Figure 4 plots the correlation estimates. For parameter estimates, we refer to Table I. We use the same model specifications as described in Section 4.2, but now estimated on the full sample of $N = 73$ firms. The equicorrelation estimates (top left) range from low values of approximately 0.1 in 2000 to values as high as 0.6 towards the start of the sovereign debt crisis in 2010. Using the two block structure (top right), there appear to be only small differences between financial firms from stressed versus non-stressed countries, except during the peak of the sovereign debt crisis between 2011 and 2012. The bottom left panel of Figure 4 introduces further heterogeneity by modeling the dynamic correlation for firms from Greece, Ireland and Portugal on the one hand, and Italy and Spain on the other. Again, the rationale for this grouping is that Ireland, Greece and Portugal entered the euro area sovereign debt crisis earlier, and were relatively more stressed, compared with Spain and Italy. This distinction may matter for our inference on financial-sector tail risks around 2012. We obtain a lower correlation level among financial institutions from the first group of smaller stressed countries. The correlation rises after the financial crisis up to the start of 2010, and then decreases to pre-crisis levels around 2012. By contrast, the correlation for financial firms from Italy and Spain starts to rise earlier, peaks higher and remains high until the end of the sample. If we consider the average correlations across all $73(73 − 1)/2 = 2628$ pairs in the bottom right panel, the picture emerging from all three model specifications is similar.

Our equity panel dataset is characterized by a substantial amount of common variation across correlations. A principal components analysis (PCA) of 52-week rolling window correlations suggests that the first three components explain approximately 43%, 22% and 12% of the total variation in pair-wise correlations, respectively. This suggests that a three-block model specification—which allows for six different correlations at each point in time—should capture a substantial share (approximately 60–75%) of the common time series variation of all 2628 full correlations. This is the case. Figure 5
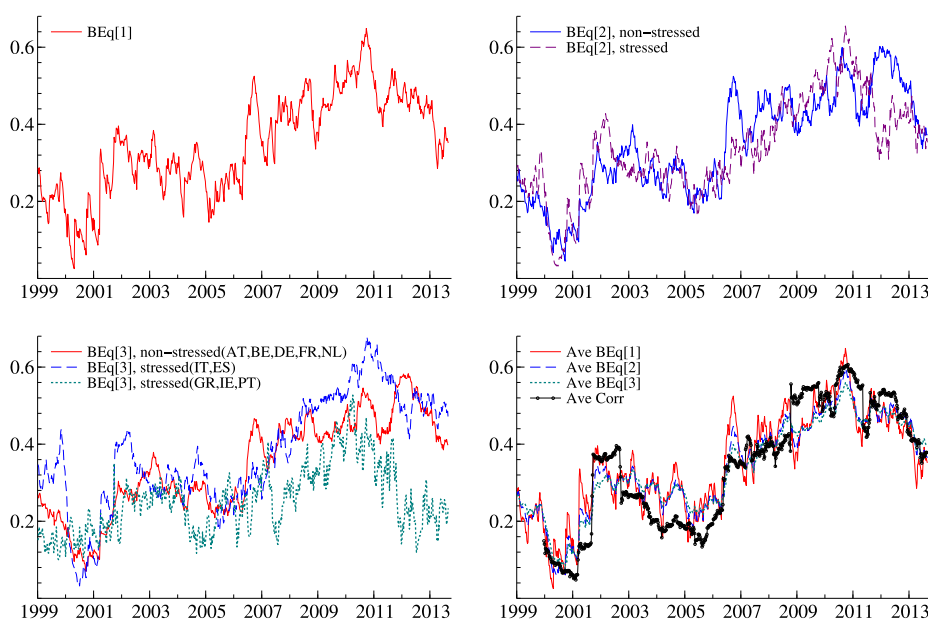

Figure 4. Filtered correlation estimates for all $N = 73$ financial firms. The panels plot GHST block equicorrelation estimates from different block equicorrelation models: one-block (top left), two-block (top right) and three-block (bottom left). The fourth panel benchmarks the block equicorrelation estimates to the average correlation taken across $73(73 − 1)/2 = 2628$ pairwise estimates based on a 52-week rolling window. This figure is available in color online at wileyonlinelibrary.com/journal/jae
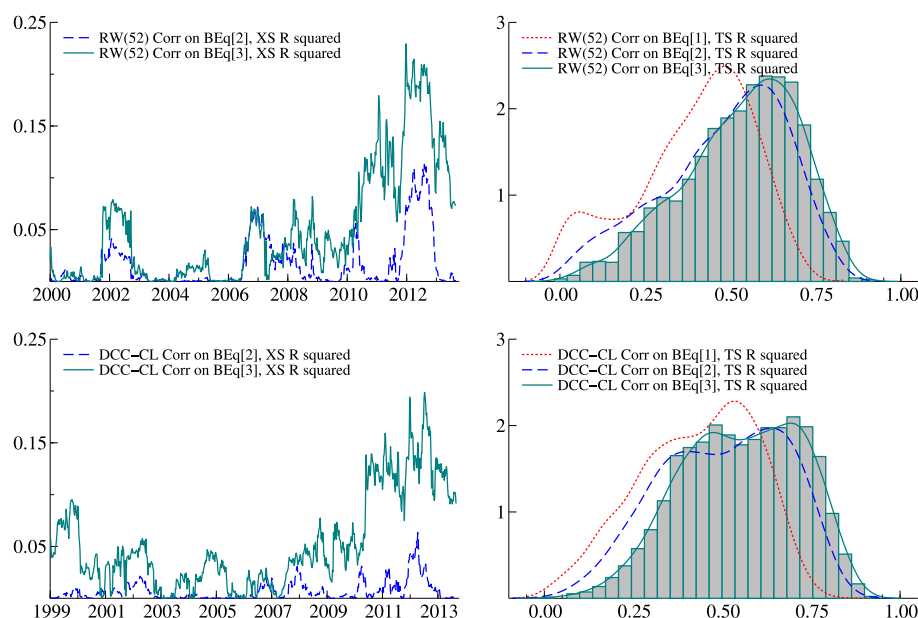
Figure 5. R^2 statistics. The upper left panel plots R^2 statistics corresponding to repeated (over time) *cross-sectional* regressions of 2628 correlation pairs from the $N = 73$ rolling window correlation (window size 52) model on a constant and the corresponding correlation estimates from a two-block and three-block equicorrelation model. The upper right panel plots kernel density estimates and a histogram of 2628 R^2 statistics corresponding to *time series* regressions with $T = 762$ observations each of rolling window correlation (window size 52) on the corresponding correlation estimates from a one-block, two-block and three-block equicorrelation model, respectively. The histogram pertains to R^2's from the three-block equicorrelation model. The bottom two panels are equivalent to the respective above panels, but use the 2628 correlation pairs from the $N = 73$ DCC-CL-GHST model instead of rolling window correlations as the left-hand-side variable. This figure is available in color online at wileyonlinelibrary.com/journal/jae

reports cross-sectional (left) and time series (right) R^2 statistics for our large-dimensional $N = 73$ case. Full correlation estimates are obtained from either rolling window correlations, or are based on the maximization of a composite likelihood as in Christoffersen *et al.* (2012).[10]

4.3.2. Joint and Conditional Tail Risk

This section presents our joint and conditional tail risk estimates for financial sector firms in the euro area. For the joint default probability, we consider the probability that more than $\bar{c} = 7/73 \approx 10\%$ of currently active financial institutions experience a credit event over the next 12 months. The probability of such widespread and massive failure should typically be very small during non-crisis times. We plot the result in the upper panel of Figure 6.

Different copula specifications yield strikingly similar results. As the joint probability moves relatively little before 2008, we only plot it over the period 2006–2013. For a comparison of joint and conditional risk outcomes for a variety of GHST and symmetric-t copula specifications we refer to the web Appendix.[11] The joint probability moves sharply upwards after the bankruptcy of Lehman Brothers in September 2008, and reaches a first peak during the Irish sovereign debt crisis in the Spring of

---

[10] Alternatively, we could get full correlation estimates by running bivariate DCC models per pair rather than going though the composite likelihood approach. We expect differences to be minor.

[11] The likelihood of the single block equicorrelation model drops by approximately 34.4 points if the restriction of symmetry is imposed; see Table I. This pronounced difference in log-likelihood, in line with Akaike information criterion (AIC) and Bayesian information criterion (BIC), strongly suggests a better fit of the asymmetric model.
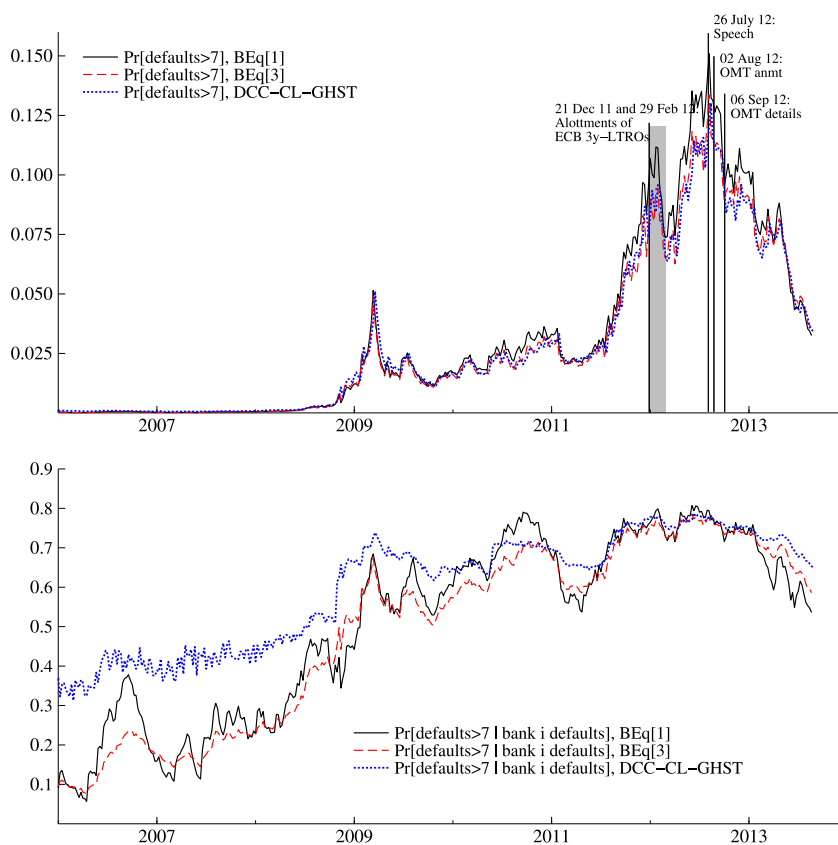
Figure 6. Joint and average conditional tail probabilities. The top panel plots the joint tail probability that more than seven financial firms ($\bar{c} = 7/73 \approx 10\%$) experience a credit event over a 12-month-ahead horizon, at any time t. The bottom panel plots the average conditional risk measure, the average (over all 73 financial firms) conditional probability that more than $\bar{c}^{(-i)} = 7/72 \approx 10\%$ financial firms default given a default for a given firm i. In each case, computations are based on a GAS-BEq[1], GAS-BEq[3] and DCC-CL full correlation model, respectively. Computations rely on the conditional law of large numbers approximation for the GHST BEq[m] models, and on simulation for the DCC composite likelihood model. The vertical lines in the top panel indicate the following events: (a) the allotments of two 3-year long-term refinancing operations by the ECB on 21 December 2011 and 29 February 2012; (b) a speech by the ECB President ('whatever it takes') on 26 July 2012; (c) the OMT announcement on 2 August 2012; and (d) the announcement of the OMT technical details on 6 September 2012. This figure is available in color online at wileyonlinelibrary.com/journal/jae

2009. It remains approximately constant thereafter until the onset of the sovereign debt crisis in early 2010. It reaches a first peak in late 2011, followed by a second and final peak mid 2012.

The vertical lines in the top panel of Figure 6 indicate a number of relevant policy announcements. In late 2011, the announcement and implementation of two 3-year long-term refinancing operations (3y-LTROs, allotted in December 2011 and February 2012) had a visible but temporary impact on financial sector tail risk.[12] In particular, the announcement and allotment of two 3y-LTROs in December 2011 and February 2012 appear to have lowered financial-sector joint tail risk, temporarily, for a few months. In the first half of 2012, financial-sector joint tail risk picked up again and rose to unprecedented levels until July 2012. The three vertical lines from July to September 2012, together with the time variation in the joint default probability, strongly suggest that three events collectively ended the

---

[12] See ECB (2011) for the official press release and monetary policy objectives.

J. Appl. Econ. **32**: 171–191 (2017)
                                                                              DOI: 10.1002/jae

most acute phase of extreme financial-sector joint tail risks. These are a speech by the ECB President in London to do '*whatever it takes*' to save the euro on 26 July 2012, the announcement of the ECB's Outright Monetary Transactions program on 2 August 2012, and especially the disclosure of the OMT details on 6 September 2012. The joint default probability decreases sharply, all the way until the end of the sample. The OMT is a program within which the ECB can purchase ('outright transactions') bonds issued by euro area member states in secondary sovereign bond markets, under certain conditions. We refer to ECB (2012) and Coeuré (2013) for details. No purchases have yet been undertaken by the ECB within the OMT. Instead, the mere announcement of the program was sufficient to trigger the substantial decline in financial sector joint tail risk.

The bottom panel of Figure 6 presents the average (across institutions) conditional tail probability for the 73 firms in our sample. Again, different copula model specifications yield similar results, although there are some discrepancies in the first third of the sample leading up to the financial crisis. Conditional probability estimates are also more sensitive to the choice of skewed or non-skewed copula; see the web Appendix for details. Also, the DCC-based conditional measures appear to be much less responsive to the different historical events. This is in line with the well-known phenomenon that the dynamics of DCC models become relatively flat in high dimensions. The equicorrelation model is much less susceptible to this bias. Though the composite likelihood approach for estimation partly corrects the potential bias in the DCC estimates, Figure 6 illustrates that the equicorrelation model picks up more of the dynamics in the series.

At the start of the sample, there is little evidence of systemic clustering on average with low levels of the CRM between 10% and 40%. The conditional probability rises following the Lehman Brothers bankruptcy to levels of approximately 60%, and then to approximately 80% around the peak of the sovereign debt crisis. Such high levels of conditional tail probabilities signal strong interconnectedness among euro area financial institutions. Interestingly, the conditional tail probability is still quite high towards the end of the sample, despite the collapse in joint risk as shown in the top panel. The OMT announcements apparently did not lower market perceptions of conditional (or contagion) risks in the euro area financial system as a whole to a comparable extent.

As a final result, unconventional monetary policy measures, such as the 3y-LTROs and the OMT announcements, and financial stability tail risk outcomes appear strongly related. This suggests substantial scope for the coordination of monetary, macro-prudential and bank supervision policies. This is relevant as both monetary policy and banking supervision have been carried out jointly by the ECB since November 2014.

5. CONCLUSION

We developed a novel modeling framework for estimating joint and conditional tail risk probabilities over time in a financial system that consists of numerous financial sector firms. For this purpose, we used a copula approach based on the generalized hyperbolic skewed Student's *t* (GHST) distribution, endowed with score-driven dynamics. Parsimony and flexibility were traded off by using a dynamic block equicorrelation structure. Using this structure, we were able to implement efficient approximations based on a conditional law of large numbers to compute joint and conditional tail probabilities of multiple defaults for a large set of firms. An application to euro area financial firms from 1999 to 2013 revealed unprecedented joint default risks between 2011 and 2012. We also document the collapse of these joint default risks (but not conditional risks) after a sequence of announcements pertaining to the ECB's Outright Monetary Transactions program in August and September 2012.

REFERENCES

Aas K, Haff I. 2006. The generalized hyperbolic skew student's *t* distribution. *Journal of Financial Econometrics* **4**(2): 275–309.

Acharya V, Engle R, Richardson M. 2012. Capital shortfall: a new approach to ranking and regulating systemic risks. *American Economic Review* **102**(3): 59–64.

Adrian T, Covitz D, Liang NJ. 2013. Financial stability monitoring. *Fed Staff Reports* **601**: 347–370.

Azzalini A, Capitanio A. 2003. Distributions generated by perturbation of symmetry with emphasis on a multivariate skew *t* distribution. *Journal of the Royal Statistical Society B* **65**: 367–389.

Bauwens L, Laurent S. 2005. A new class of multivariate skew densities, with application to generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics* **23**(3): 346–354.

Bibby B, Sørensen M. 2003. Hyperbolic processes in finance. In *Handbook of Heavy Tailed Distributions in Finance*, Rachev ST (ed). Elsevier Science: Amsterdam; 211–248.

Blasques F, Koopman SJ, Lucas A. 2012. Stationarity and ergodicity of univariate generalized autoregressive score processes. *Technical Report Tinbergen Institute Discussion Paper 12–059*, Tinbergen Institute, Amsterdam.

Blasques F, Koopman SJ, Lucas A. 2014. Maximum likelihood estimation for generalized autoregressive score models. *Technical Report Tinbergen Institute Discussion Paper 14–029*, Tinbergen Institute, Amsterdam.

Blasques F, Koopman SJ, Lucas A. 2015. Information theoretic optimality of observation driven time series models for continuous responses. *Biometrika* **102**(2): 325–343.

Branco M, Dey D. 2001. A general class of multivariate skew-elliptical distributions. *Journal of Multivariate Analysis* **79**: 99–113.

Christoffersen P, Errunza V, Langlois H, Jacobs K. 2012. Is the potential for international diversification disappearing? *Review of Financial Studies* **25**: 3711–3751.

Christoffersen P, Errunza V, Jacobs K, Jin X. 2014a. Correlation dynamics and international diversification benefits. *International Journal of Forecasting* **30**(3): 807–824.

Christoffersen P, Jacobs K, Jin X, Langlois H. 2014b. Dynamic dependence in corporate credit. *Technical Report Working paper*, Stern NYU, New York. Available: http://www.stern.nyu.edu/sites/default/files/assets/documents/con_041385.pdf [13 March 2016].

Coeuré B. 2013. Outright monetary transactions, one year on, *Speech at the Conference The ECB and its OMT Programme*: Berlin. 2 September 2013.

Creal DD, Tsay RS. 2015. High dimensional dynamic stochastic copula models. *Journal of Econometrics* **189**(2): 335–345.

Creal D, Koopman S, Lucas A. 2011. A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations. *Journal of Business and Economic Statistics* **29**(4): 552–563.

Creal D, Koopman S, Lucas A. 2013. Generalized autoregressive score models with applications. *Journal of Applied Econometrics* **28**(5): 777–795.

Creal D, Schwaab B, Koopman SJ, Lucas A. 2014. An observation driven mixed measurement dynamic factor model with application to credit risk. *Review of Economics and Statistics* **96**(5): 898–915.

Duffie D, Saita L, Wang K. 2007. Multi-period corporate default prediction with stochastic covariates. *Journal of Financial Economics* **83**(3): 635–665.

ECB. 2011. ECB announces measures to support bank lending and money market activity. ECB press release, 8 December 2011.

ECB. 2012. Technical features of Outright Monetary Transactions. ECB press release, 6 September 2012.

ECB. 2014. The determinants of euro area sovereign bond yield spreads during the crisis. *ECB Monthly Bulletin* May 2014.

Engle R. 2002. Dynamic conditional correlation. *Journal of Business and Economic Statistics* **20**(3): 339–350.

Engle R, Kelly B. 2012. Dynamic equicorrelation. *Journal of Business and Economic Statistics* **30**(2): 212–228.

Eser F, Schwaab B. 2015. Evaluating the impact of unconventional monetary policy measures: empirical evidence from the ECBs Securities Markets Programme. *Journal of Financial Economics* **119**(1): 147–167.

Giesecke K, Spiliopoulos K, Sowers RB, Sirignano JA. 2015. Large portfolio asymptotics for loss from default. *Mathematical Finance* **25**(1): 77–114.

Glassermann P, Li J. 2005. Importance sampling for portfolio credit risk. *Management Science* **51**: 1643–1656.

Gordy M. 2000. A comparative anatomy of credit risk models. *Journal of Banking and Finance* **24**: 119–149.

Gordy MB. 2003. A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation* **12**(3): 199–232.

Gupta A. 2003. Multivariate skew-*t* distribution. *Statistics* **37**(4): 359–363.

Hansen BE. 1994. Autoregressive conditional density estimation. *International Economic Review* **35**(3): 705–730.

Hartmann P, Straetmans S, de Vries C. 2004. Asset market linkages in crisis periods. *Review of Economics and Statistics* **86**(1): 313–326.

Hartmann P, Straetmans S, de Vries C. 2007. Banking system stability: a cross-Atlantic perspective. In *The Risks of Financial Institutions*, Carey M, Stulz RM (eds). NBER and University of Chicago Press: Chicago, IL; 1–61.

Harvey AC. 2013. *Dynamic Models for Volatility and Heavy Tails*. Cambridge University Press: Cambridge, UK.

Koopman SJ, Lucas A, Schwaab B. 2011. Modeling frailty correlated defaults using many macroeconomic covariates. *Journal of Econometrics* **162**(2): 312–325.

Koopman SJ, Lucas A, Schwaab B. 2012. Dynamic factor models with macro, frailty, and industry effects for U.S. default counts: the credit crisis of 2008. *Journal of Business and Economic Statistics* **30**(4): 521–532.

Lucas A, Klaassen P, Spreij P, Straetmans S. 2001. An analytic approach to credit risk of large corporate bond and loan portfolios. *Journal of Banking and Finance* **25**(9): 1635–1664.

Lucas A, Klaassen P, Spreij P, Straetmans S. 2003. Tail behaviour of credit loss distributions for general latent factor models. *Applied Mathematical Finance* **10**(4): 337–357.

Lucas A, Schwaab B, Zhang X. 2014. Conditional euro area sovereign default risk. *Journal of Business and Economic Statistics* **32**(2): 271–284.

Mencía J, Sentana E. 2005. Estimation and testing of dynamic models with generalized hyperbolic innovations. *Technical Report CEPR Discussion Paper No. 5177*, CEPR, London.

Oh DH, Patton A. 2014. Time-varying systemic risk: evidence from a dynamic copula model of cds spreads. *Technical Report Working paper*, Department of Economics, Duke University, Durham.

Segoviano MA, Goodhart C. 2009. Banking stability measures. IMF working paper.

Suh S. 2012. Measuring systemic risk: a factor-augmented correlated default approach. *Journal of Financial Intermediation* **21**(2): 341–358.

Vasicek O. 1987. Probability of loss on loan portfolio. *Technical Report Working Paper,* Moody's Analytics Corporation, San Francisco.

Zhang X, Creal D, Koopman S, Lucas A. 2011. Modeling dynamic volatilities and correlations under skewness and fat tails. *Discussion Paper 11-078/2/DSF22*, Tinbergen Institute, Amsterdam.