

Conditional probabilities and contagion measures for Euro Area sovereign default risk*

Xin Zhang,^(a) Bernd Schwaab,^(b) André Lucas^(a,c)

^(a) VU University Amsterdam, and Tinbergen Institute

^(b) European Central Bank, Financial Research

^(c) Duisenberg school of finance

March 27, 2012

*Author information: Xin Zhang, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, Email: x.zhang@vu.nl. Bernd Schwaab, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt, Germany, Email: bernd.schwaab@ecb.int. Andre Lucas, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, Email: a.lucas@vu.nl. We thank Markus Brunnermeier and Laurent Clerc for early comments, and seminar participants at Bundesbank, ECB, HEC Lausanne, Riksbank, and the Royal Economics Society 2012 meeting at Cambridge. Andre Lucas thanks the Dutch National Science Foundation (NWO) for financial support. The views expressed in this paper are those of the authors and they do not necessarily reflect the views or policies of the European Central Bank or the European System of Central Banks.

Conditional probabilities and contagion measures for Euro Area sovereign default risk

Abstract

The Eurozone debt crisis raises the issue of measuring and monitoring interconnected sovereign credit risk. We propose a novel empirical framework to assess the likelihood of joint and conditional failure for Euro Area sovereigns. Our model captures all the salient features of the data, including skewed and heavy-tailed changes in the price of CDS protection against sovereign default, as well as dynamic volatilities and correlations to ensure that failure dependence can increase in times of stress. We apply the model to Euro Area sovereign CDS spreads from 2008 to mid-2011. Our results reveal significant time-variation in risk dependence and considerable spill-over effects in the likelihood of sovereign failures. We further investigate distress dependence around a key policy announcement on 10 May 2010, and demonstrate the importance of capturing higher-order time-varying moments during times of crisis for the assessment of interactional risks.

Keywords: sovereign credit risk; higher order moments; time-varying parameters; financial stability surveillance.

JEL classification: C32, G32.

1 Introduction

In this paper we construct a novel empirical framework to assess the likelihood of joint and conditional failure for Euro Area sovereigns. The new framework allows us to estimate marginal, joint, and conditional probabilities of sovereign default from observed prices for credit default swaps (CDS) on sovereign debt. We define failure as any credit event that would trigger a sovereign CDS contract. Examples of such failure are the non-payment of principal or interest when it is due, a forced exchange of debt into claims of lower value, also a moratorium or official repudiation of the debt. Our methodology is novel in that our risk measures are derived from a multivariate framework based on a dynamic Generalized Hyperbolic (GH) skewed- t density that naturally accommodates all relevant empirical data features, such as skewed and heavy-tailed changes in individual country CDS spread changes, as well as time variation in their volatilities and dependence. Moreover, the model can easily be calibrated to match current market expectations regarding the marginal probabilities of default, similar to for example Segoviano and Goodhart (2009) and Huang, Zhou, and Zhu (2009).

We make three main contributions to the literature on risk assessment. First, we provide estimates of the time variation in Euro Area joint and conditional sovereign default risk using CDS data from January 2008 to June 2011. For example, the conditional probability of a default on Portuguese debt given a Greek failure is estimated to be around 30% at the end of our sample. Similar conditional probabilities for other countries are also reported. At the same time, we may infer which countries are more exposed than others to a certain credit event. Second, we analyze the extent to which parametric assumptions matter for such joint and conditional risk assessments. Perhaps surprisingly, and despite the widespread use of joint risk measures to guide policy decisions, we are not aware of a detailed investigation of how different parametric assumptions matter for joint and conditional risk assessments. We therefore report results based on a dynamic multivariate Gaussian and symmetric- t

density in addition to a GH skewed- t (GHST) specification. The distributional assumptions matter most for our conditional assessments, whereas simpler joint failure estimates are less sensitive to the assumed dependence structure. Third, our modeling framework allows us to investigate the presence and severity of market implied spill-overs in the likelihood of sovereign failure. Specifically, we document spill-overs from the possibility of a Greek failure to the perceived riskiness of other Euro Area countries. For example, at the end of our sample we find a difference of about 30% between the one-year conditional probability of a Portuguese default given that Greece does versus that Greece does not default. This suggests that the cost of debt refinancing in some European countries depends to some extent on developments in other countries. We conclude our empirical analysis by investigating the impact on sovereign joint and conditional risks of a key policy announcement on 09 May 2010 by Euro Area heads of state.

Policy makers routinely track joint and conditional probabilities of failure for financial sector surveillance purposes. Model based estimates of stress and spill-overs can be useful as quick gauges of market expectations. An additional reason for their popularity in policy circles may be that they are less expensive in terms of manpower than detailed studies based on micro data. We give two examples for joint failure measures that are used in practice. First, the European Central Bank's semi-annual Financial Stability Review contains a plot of the estimated probabilities of two or more failures in a portfolio of about twenty large and complex European financial firms, see ECB (2011, p. 95). The estimate relies on CDS data to infer marginal risks, and equity return correlations to infer the dependence structure. The setup is based on a multi-factor credit risk model and Gaussian dependence, see Avesani, Pascual, and Li (2006) and Hull and White (2004). As a second example, the International Monetary Fund's Global Financial Stability Review (2009, p. 132) features an estimate of the time varying probability that all banks in a certain portfolio become distressed at the same time. This joint probability estimate is based on Segoviano and Goodhart (2009).

Their CIMDO framework is based on a multivariate prior distribution, usually Gaussian or symmetric- t , which can be calibrated to match marginal risks as implied by the CDS market. The multivariate density becomes discontinuous at so-called threshold levels: some parts of the density are shifted up, others are shifted down, while the parametric tail and extreme dependence implied by the prior remains intact at all times.

From a risk perspective, bad outcomes are much worse if they occur in clusters. What seems manageable in isolation may not be so if the rest of the system is also under stress, see Acharya, Pedersen, Philippon, and Richardson (2010). Though this argument is mainly applied when looking at the risks of banks or other financial institutions, the same intuition applies, possibly to an even greater extent, to sovereign default risk in the Euro Area. While adverse developments in one country's public finances could perhaps still be handled with the support of a healthy remainder of the union, the situation may quickly become untenable if two, three, four, five, or more, of its members are also distressed at the same time. Multivariate sovereign risk dependence is a natural consequence of a large degree of legal, economic, and financial interconnectedness between member states. Such dependence can also be expected to be time varying, as market participants adapt to changes in policy and in the institutional setting.

Figure 1 illustrates the particular situation in the Euro Area from January 2008 to June 2011. In addition, the figure reveals some of the statistical difficulties that need to be addressed. The top panel in Figure 1 shows the price of CDS protection against sovereign failure for ten Euro Area countries. The Euro Area sovereign debt crisis has progressively spread across various member countries. After the intensification of tensions in the Greek government bond market in Spring 2010, Ireland, Portugal and eventually also Spain and Italy became increasingly engulfed in the sovereign crisis. French and German sovereign credit default swaps (CDS) have also been affected since then. The bottom panel of Figure 1 plots the daily changes in the CDS data. Strong volatility clustering is present in all spread

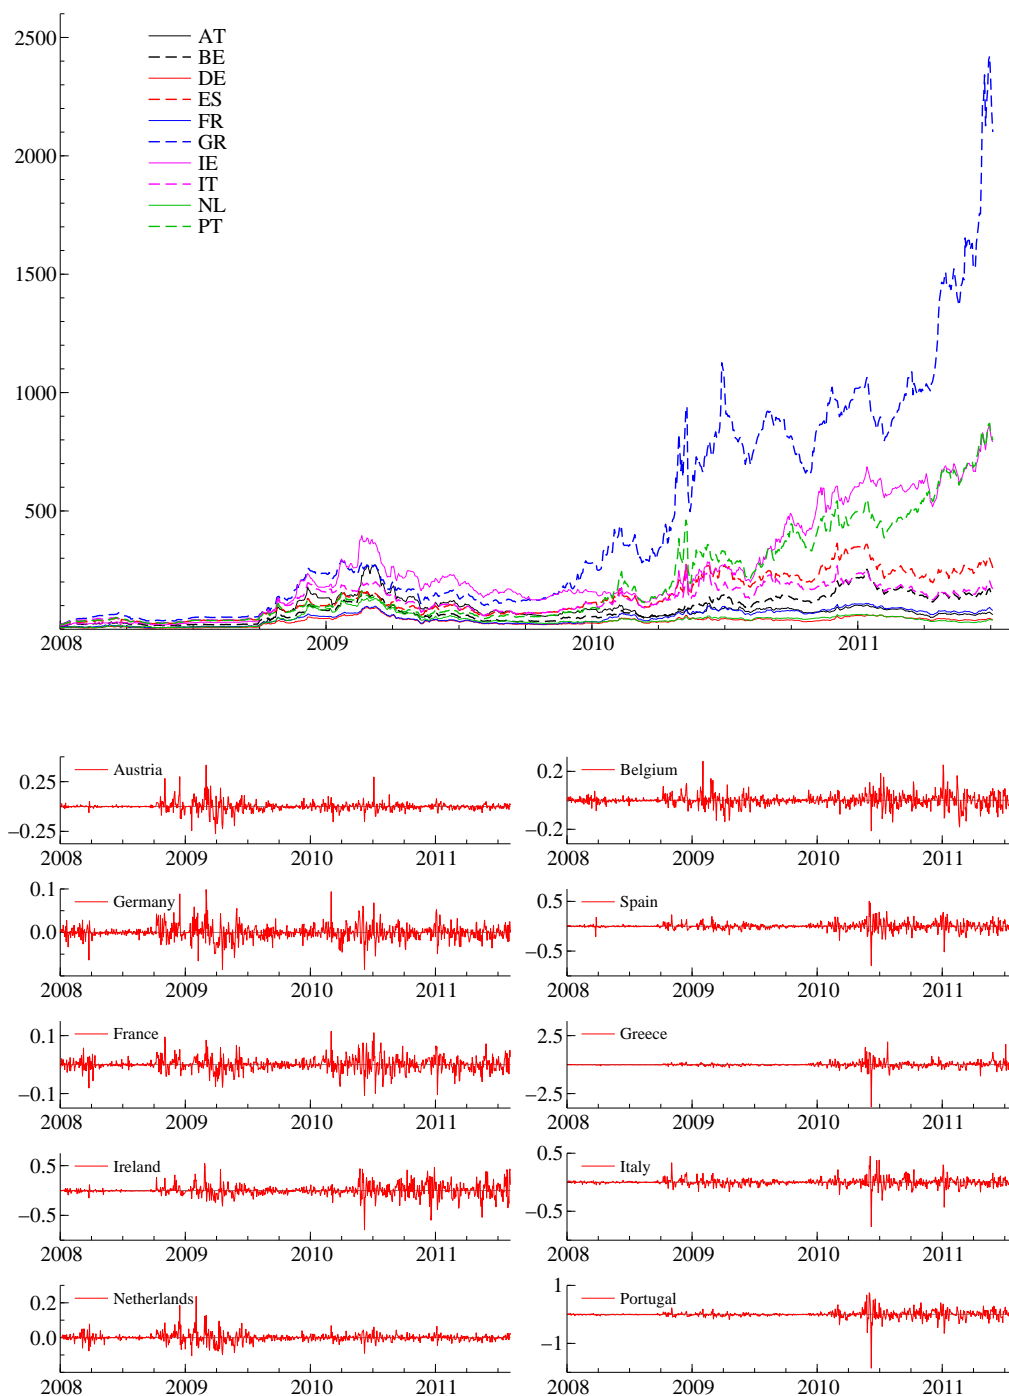
changes. In addition, skewness and fat tails are important features in each time series.

Our framework for joint sovereign risk assessments has four attractive features that each match empirical stylized facts: volatility clustering, dynamic correlations, non-trivial tail dependence, and the ability to handle a cross sectional dimension of intermediate size. As seen from Figure 1, any adequate model needs to accommodate the time-variation in volatilities and the skewness and fat-tailedness of the data before estimating the correlation structure. Second, the correlation structure between individual countries is likely to be time-varying. Correlations tend to increase during times of stress, see for example Forbes and Rigobon (2002). Policy decisions such as the introduction of the European Financial Stability Facility in May 2010, as well as direct bond purchases by the ECB starting around the same time, may also have had a direct impact on the dependence structure. We therefore propose a model with a dynamic correlation matrix, treating each correlation pair as a latent process. Third, we want to allow for multivariate non-Gaussian features such as extreme tail dependence. Clearly, if an important aspect of overall system risk is that of simultaneous failures, then the multivariate distribution should not rule out extreme dependence a priori, e.g., by assuming a Gaussian copula. Fourth, the model needs to be flexible enough to be calibrated repeatedly to current market conditions, such as current CDS spread levels. In particular, there is a preference in policy circles to impose the constraint that the marginal probabilities of default obtained from the multivariate model should be equal to the market implied probabilities of default as obtained from current CDS spread levels. Finally, and most obviously, an application to Euro Area sovereign risk requires us to handle dimensions larger than what is usual in the non-Gaussian copula literature (less than five). Our results are therefore also be interesting from this more technical perspective.

The remainder of the paper is set up as follows. Section 2 introduces a conceptual framework for joint and conditional risk measures. Section 3 introduces the multivariate model for failure dependence. The empirical results are discussed in Section 4. Section 5

Figure 1: CDS levels and spread changes for ten Euro Area sovereigns

The top panel reports the price of CDS protection against sovereign default for ten Euro Area countries as the CDS spread (in basis points). The sample is 01 Jan 2008 to 30 Jun 2011. The bottom panels plot changes in CDS spreads with vertical axes denoted in points (so 0.2 denoting 20 percentage points).



concludes.

2 The conceptual framework

In a corporate credit risk setting, the probability of failure is often modeled as the probability that the value of a firm's assets falls below the value of its debt at (or before) the time when its debt matures, see Merton (1974) and Black and Cox (1976). To allow for default clustering, default times of the individual firms are linked together using a copula function such as the Gaussian or Student t copula, see for example McNeil, Frey, and Embrechts (2005). In a sovereign credit risk setting, a similar modeling approach can be adopted, though the interpretation has to be altered given the different nature of a sovereign compared to a corporate. Rather than to consider asset levels falling below debt values, it is more convenient for sovereign credit risk to compare costs and benefits of default, see Eaton and Gersovitz (1981) and Calvo (1988). Costs from defaulting may arise from losing credit market access for at least some time, obstacles to conducting international trade, difficulties borrowing in the domestic market, etc., while benefits from default include immediate debt relief.

In our model below, failure is triggered by a variable v_{it} that measures the time-varying changes in the difference between the perceived benefits and cost of default for country i at time t . If v_{it} rises above a fixed threshold value, default is triggered. Since a cost, or penalty, can always be recast in terms of a benefit, we incur no loss of generality if we focus on a model with time-varying benefits of default and fixed costs, or vice versa, see Calvo (1988).

To model the dependence structure between sovereign defaults, we take a slightly different approach than for corporates defaults. As the cross-sectional dimension for corporate credit risk studies is typically large, the common approach there is to impose a simple factor structure on the copula function in order to capture default dependence, see for example McNeil et al. (2005, Chapter 8) for a textbook treatment. For sovereign credit risk, the cross-sectional dimension is typically smaller and contagion concerns are much more pronounced.

Our model is set in discrete time, similar to the familiar models for corporate credit risk such as CreditMetrics (2007). The model describes v_{it} , i.e., the changes in the differences between perceived benefits and costs from default, as

$$v_{it} = (\varsigma_t - \mu_\varsigma) \tilde{L}_{it} \gamma + \sqrt{\varsigma_t} \tilde{L}_{it} \epsilon_t, \quad i = 1, \dots, n, \quad (1)$$

where $\epsilon_t \in \mathbb{R}^n$ is a vector of standard normally distributed risk factors such as exposure to the business cycle or to monetary policy, \tilde{L}_t is a $n \times n$ matrix with i th row \tilde{L}_{it} holding the sensitivities of v_{it} to the n risk factors, $\gamma \in \mathbb{R}^n$ is a vector of sovereign-specific constants that controls the skewness of v_{it} , ς_t is a positively valued scalar random risk factor common to all sovereigns and independent of ϵ_t , and $\mu_\varsigma = \mathbb{E}[\varsigma_t]$ assuming that the expectation exists. Note that $\mathbb{E}[v_{it}] = 0$ and that if ς_t is non-random, the first term in (1) drops out.

The basic structure of (1) is very similar to the standard Gaussian or Student t copula model. The main two differences lie in the use of the additional risk factor ς_t and the skewness parameters γ . The factor ς_t is a scalar risk factor, unlike the normally distributed risk factors ϵ_t . If ς_t is large, all sovereigns are affected at the same time, making joint defaults of two or more sovereigns more likely.

In this paper, we assume that ς_t has an inverse-Gamma distribution, thus obtaining a generalized hyperbolic skewed t (GHST) distribution for v_{it} . For $\gamma = 0$, we obtain the familiar Student t copula model. The GHST model could be further generalized to the GH model by assuming a generalized inverse Gaussian distribution for ς_t , see McNeil et al. (2005). The current simpler GHST model, however, already accounts for all the empirical features in the CDS data at hand, including skewness and fat tails. As ς_t affects both the variance and kurtosis of v_{it} and (via γ) also its mean and skewness, model (1) can capture cross-sectional default spill-overs and contagion concerns through correlations (z_t) as well as through tail dependence (ς_t).

We assume that country i fails if v_{it} exceeds a predefined sovereign-specific threshold c_{it} .

The probability of default p_{it} is then given by

$$p_{it} = \Pr[v_{it} > c_{it}] = 1 - F_i(c_{it}) \quad \Leftrightarrow \quad c_{it} = F_i^{-1}(1 - p_{it}), \quad (2)$$

where $F_i(\cdot)$ is the cumulative distribution function of v_{it} . The thresholds may be imputed from market-implied estimates \hat{p}_{it} of the probability of default based on, for example, CDS data. It is easily seen from the mean-variance mixture (1) construction that if (v_{1t}, \dots, v_{nt}) are jointly GHST distributed, then each v_{it} is marginally GHST distributed for $i = 1, \dots, n$.

Given γ_i , \tilde{L}_t , and the remaining model parameters, we can use (1) to obtain *joint* and *conditional* probabilities of sovereign failure from CDS data in four main steps.

As a first step, we operationalize joint sovereign risk as the probability of joint credit events in a subset of $n^* \leq n$ countries. The dependence structure is taken from observed CDS data. In particular, we calibrate the copula structure in (1) on the copula of observed changes in the daily CDS spreads of the different sovereigns. This produces both the multivariate dependence structure and the implied marginal densities.

In a second step, we obtain the thresholds c_{it} that determine individual country default risk. These can be obtained by direct simulation or by a numerical inversion of the GHST distribution $F_i(\cdot)$ as explained in (2). Both approaches require a CDS market-implied estimate of the probability of default. To obtain such an estimate, we make a number of simplifying assumptions. First, we fix the recovery rate at $rec_i = 50\%$ for all countries. This number is roughly in the middle of the 13% to 73% range for sovereign haircuts reported in Sturzenegger and Zettelmeyer (2008), and close to the number discussed for Greek debt at the end of 2011. In addition, assuming a slightly higher value for expected recoveries is conservative, since expected recovery and implied risk neutral pd's are positively related given the CDS spread, see below. Second, we assume a risk free rate of $r_t = 2\%$ and a flat term structure for both interest rates and default intensities until the maturity of the CDS contract. Finally, we assume that the premium payments are paid out continuously. Alternative specifications are clearly possible and can be adopted as well. Using these as-

sumptions, the standard CDS pricing formula of Hull and White (2000) simplifies and can be inverted to extract the market-implied risk neutral probability of default p_{it} , see for example Brigo and Mercurio (2006, Chapter 21). This probability is a direct function of the observed CDS spread s_{it} ,

$$p_{it} = \frac{s_{it} \times (1 + r_t)}{1 - rec_i}. \quad (3)$$

In a third step, we use the time-varying estimates of the multivariate dependence structure between CDS spreads, obtained using the techniques discussed in the next section, Section 3. This dependence structure allows the correlations to vary over time and is based on pre-filtered and standardized changes in CDS spreads.

In a fourth step, we obtain measures of marginal, joint, and conditional failure of one, two, or more sovereigns by means of simulation. In the simulation setup, we draw repeatedly from the multivariate distribution, say, 10,000 times at each time point t . In each simulation, country i fails if and only if the draw for country i exceeds the default threshold c_{it} computed in the second step. In this way, we control the marginal probabilities of default for country i at time t to closely match the CDS implied probabilities. Joint and conditional failure probabilities are obtained similarly by counting the joint exceedances and dividing by the number of simulations.

3 Statistical model

3.1 Generalized Autoregressive Score dynamics

As mentioned in Section 2, we use sovereign CDS data to estimate the model's time-varying dependence structure and to calibrate the model's marginal default probabilities to current market data. Our econometric specification closely follows the economic model structure of the previous section. We observe a vector $y_t \in \mathbb{R}^n$ of changes in sovereign CDS spreads and

assume that

$$y_t = \mu + L_t e_t, \quad (4)$$

with $\mu \in \mathbb{R}^n$ a vector of fixed unknown means, and e_t a GHST distributed random variable with mean zero and covariance matrix I . To ease the notation, we set $\mu = 0$ in the remaining exposition. For $\mu \neq 0$, all derivations go through if y_t is replaced by $y_t - \mu$. We concentrate on the case where e_t follows a GHST distribution, such that y_t has the density

$$p(y_t; \tilde{\Sigma}_t, \gamma, \nu) = \frac{\nu^{\frac{\nu}{2}} 2^{1-\frac{\nu+n}{2}}}{\Gamma(\frac{\nu}{2}) \pi^{\frac{n}{2}} |\tilde{\Sigma}_t|^{\frac{1}{2}}} \cdot \frac{K_{\frac{\nu+n}{2}} \left(\sqrt{d(y_t) \cdot (\gamma' \gamma)} \right) e^{\gamma' \tilde{L}_t^{-1} (y_t - \tilde{\mu}_t)}}{d(y_t)^{\frac{\nu+n}{4}} \cdot (\gamma' \gamma)^{-\frac{\nu+n}{4}}}, \quad (5)$$

$$d(y_t) = \nu + (y_t - \tilde{\mu}_t)' \tilde{\Sigma}_t^{-1} (y_t - \tilde{\mu}_t), \quad (6)$$

$$\tilde{\mu}_t = -\frac{\nu}{\nu-2} \tilde{L}_t \gamma, \quad (7)$$

where $\nu > 4$ is the degrees of freedom parameter, $\tilde{\mu}_t$ is the location vector and $\tilde{\Sigma}_t = \tilde{L}_t \tilde{L}_t'$ is the scale matrix,

$$\tilde{L}_t = L_t T, \quad (8)$$

$$(T' T)^{-1} = \frac{\nu}{\nu-2} I + \frac{2\nu^2}{(\nu-2)^2(\nu-4)} \gamma \gamma', \quad (9)$$

and $K_a(b)$ is the modified Bessel function of the second kind. The matrix L_t characterizes the time-varying covariance matrix $\Sigma_t = L_t L_t'$. We consider the standard decomposition

$$\Sigma_t = L_t L_t' = D_t R_t D_t, \quad (10)$$

where D_t is a diagonal matrix containing the time-varying volatilities of y_t , and R_t is the time-varying correlation matrix.

The fat-tailedness and skewness of the data y_t creates challenges for standard dynamic specifications for volatilities and correlations, such as standard GARCH or DCC type dynamics, see Engle (2002). In the presence of fat tails, large absolute observations y_{it} occur regularly even if volatility is not changing rapidly. If not properly accounted for, such observations lead to biased estimates of the dynamic behavior of volatilities and correlations. The

Generalized Autoregressive Score (GAS) framework of Creal, Koopman, and Lucas (2012) as applied in Zhang, Creal, Koopman, and Lucas (2011) to the case of GHST distributions provides a coherent approach to deal with such settings. The GAS model creates an explicit link between the distribution of y_t and the dynamic behavior of Σ_t , L_t , D_t , and R_t . In particular, if y_t is fat-tailed, observations that lie far outside the center automatically have less impact on future values of the time-varying parameters in Σ_t . The same holds for observations in the left-hand tail if y_t is left-skewed. The intuition for this is that the score dynamics attribute the effect of a large observation y_t partly to the distribution properties of y_t and partly to a local increase of volatilities and/or correlations. The estimates of dynamic volatilities and correlations thus become more robust to incidental influential observations. This is important for the CDS data displayed in Figure 1. We refer to Creal, Koopman, and Lucas (2011) and Zhang, Creal, Koopman, and Lucas (2011) for more details.

We assume that the time-varying covariance matrix Σ_t is driven by a number of unobserved dynamic factors f_t , or $\Sigma_t = \Sigma(f_t) = L(f_t)L(f_t)'$. The dynamics of f_t are specified using the GAS framework for GHST distributed random variables and are given by

$$f_{t+1} = \omega + \sum_{i=0}^{p-1} A_i s_{t-i} + \sum_{j=0}^{q-1} B_j f_{t-j}; \quad (11)$$

$$s_t = \mathcal{S}_t \nabla_t, \quad (12)$$

$$\nabla_t = \partial \ln p(y_t; \tilde{\Sigma}(f_t), \gamma, \nu) / \partial f_t, \quad (13)$$

where ∇_t is the score of the GHST density with respect to f_t , $\tilde{\Sigma}(f_t) = L(f_t)TT'L(f_t)'$, ω is a vector of fixed intercepts, A_i and B_j are appropriately sized fixed parameter matrices, \mathcal{S}_t is a scaling matrix for the score ∇_t , and $\omega = \omega(\theta)$, $A_i = A_i(\theta)$, and $B_j = B_j(\theta)$ all depend on a static parameter vector θ . Typical choices for the scaling matrix \mathcal{S}_t are the unit matrix or inverse (powers) of the Fisher information matrix \mathcal{I}_{t-1} , where

$$\mathcal{I}_{t-1} = \mathbb{E}[\nabla_t \nabla_t' | y_{t-1}, y_{t-2}, \dots].$$

For example, $\mathcal{S}_t = \mathcal{I}_{t-1}^{-1}$ accounts for the curvature in the score ∇_t .

For appropriate choices of the distribution, the parameterization, and the scaling matrix, the GAS model (11)–(13) encompasses a wide range of current familiar models such as the (multivariate) GARCH model. Further details on the parameterization $\Sigma_t = \Sigma(f_t)$, $D_t = D(f_t)$, and $R_t = R(f_t)$, as well as the choice for the scaling matrix \mathcal{S}_t used in this paper are presented in the appendix. In particular, f_t contains the log-volatilities of y_t and a spherical representation of the correlation matrix R_t . The latter ensures that $R_t = R(f_t)$ is a correlation matrix by construction, irrespective of the value of f_t .

Using the GHST specification of (5), we obtain that

$$\nabla_t = \Psi_t' H_t' \text{vec} \left(w_t \cdot y_t y_t' - \tilde{\Sigma}_t - \left(1 - \frac{\nu}{\nu-2} w_t \right) \tilde{L}_t \gamma y_t' \right), \quad (14)$$

where w_t is a scalar weight function defined in the appendix that depends on the distance $d(y_t)$ from equation (6), and where Ψ_t and H_t are time-varying matrices that depend on f_t , but not on the data. See the appendix for further details.

Due to the presence of w_t in (14), observations that are far in the tails receive a smaller weight and therefore have a smaller impact on future values of f_t . This robustness feature is directly linked to the fat-tailed nature of the GHST distribution and allows for smoother correlation and volatility dynamics in the presence of heavy-tailed observations.

For skewed distributions ($\gamma \neq 0$), the score in (14) shows that positive CDS changes have a different impact on correlation and volatility dynamics than negative ones. As explained earlier, this aligns with the intuition that CDS changes from for example the left tail are less informative about changes in volatilities and correlations if the (conditional) observation density is itself left-skewed. For the symmetric Student's t case, we have $\gamma = 0$ and the asymmetry term in (14) drops out. If furthermore the fat-tailedness is ruled out by considering $\nu \rightarrow \infty$, one can show that the weights w_t tend to 1 and that ∇_t collapses to the intuitive form for a multivariate GARCH model, $\nabla_t = \Psi_t' H_t' \text{vec}(y_t y_t' - \Sigma_t)$.

3.2 Parameter estimation

The parameters of the dynamic GHST model can be estimated by standard maximum likelihood procedures as the likelihood function is known in closed form using a standard prediction error decomposition. The joint estimation of all parameters in the model, however, is rather cumbersome. Therefore, we split the estimation in two steps relating to (i) the marginal behavior of the coordinates y_{it} and (ii) the joint dependence structure of the vector of standardized residuals $D_t^{-1}y_t$. Similar two-step procedures can be found in Engle (2002), Hu (2005), and other studies that are based on a multivariate GARCH framework.

In the first step, we estimate a dynamic GHST model for each series y_{it} separately using the GAS dynamic specification and taking our time-varying parameter f_t as the log-volatility $\log(\sigma_{it})$. The skewness parameter γ_i is also estimated for each series separately, while the degrees of freedom parameter ν is fixed at a pre-determined value. This restriction ensures that the univariate GHST distributions are the marginal distributions from the multivariate GHST distribution and that the model is therefore internally consistent.

In the second step, we consider the standardized data $z_{it} = y_{it}/\hat{\sigma}_{it}$, where $\hat{\sigma}_{it}$ are obtained from the first step. Using $z_t = (z_{1t}, \dots, z_{nt})'$, we estimate a multivariate dynamic GHST model using the GAS dynamic specification. The GHST distribution in this second step has mean zero, skewness parameters $\hat{\gamma}_i$ as estimated in the first step, the same pre-determined value for ν as for the marginal models, and covariance matrix $\text{cov}(z_t) = R_t = R(f_t)$, where f_t contains the spherical coordinates of the choleski decomposition of the correlation matrix R_t , see the Appendix for further details.

The advantages of the two-step procedure for estimation efficiency are substantial, particularly if the number n of time series considered in y_t is large. The univariate models of the first step can be estimated at low computational cost. Using these estimates, the univariate dynamic GHST models are used as a filter to standardize the individual CDS spread changes. In the second step, only the parameters that determine the dynamic correlations

remain to be estimated.

4 Empirical application: Euro Area sovereign risk

4.1 CDS data

We compute joint and conditional probabilities of failure for a set of ten countries in the Euro Area. We focus on sovereigns that have a CDS contract traded on respective reference bonds since the beginning of our sample in January 2008. We select ten countries: Austria (AT), Belgium (BE), Germany (DE), Spain (ES), France (FR), Greece (GR), Ireland (IE), Italy (IT), the Netherlands (NL) and Portugal (PT). CDS spreads are available for these countries at a daily frequency from 01 January 2008 to 30 June 2011, yielding $T = 913$ observations. The CDS contracts are denominated in U.S. dollars and therefore do not depend on foreign exchange risk concerns should a European credit event materialize. The contracts have a five year maturity and are traded in fairly liquid over-the-counter markets. All time series data is obtained from Bloomberg. Table 1 provides summary statistics for daily de-meaned changes in these CDS spreads. All time series have significant non-Gaussian features under standard tests and significance levels. All series are covariance stationary according to standard unit root (ADF) tests.

4.2 Marginal and joint risk

We model the CDS spread changes with the framework explained in Section 3 based on the dynamic GHST distribution with $p = q = 1$ in (11), which we label a GAS(1,1) specification. We consider three different choices for the parameters, corresponding to a Gaussian, a Student- t , and a GHST distribution, respectively. We fix the degrees of freedom parameter for the fat-tailed specifications to $\nu = 5$. Although this may seem high at first sight given some of the large spread changes in Figure 1, the value is small enough to result in a substantial robustification of the results, both in terms of likelihood evaluation as well as in terms of the volatility and correlation dynamics.

Table 1: Data descriptive statistics

The summary statistics correspond to daily changes in observed sovereign CDS spreads for ten Euro Area countries from January 2008 to June 2011. Mean, Median, Standard Deviation, Minimum and Maximum are multiplied by 100. Almost all skewness and excess kurtosis statistics have p -values below 10^{-4} , except the skewness parameters of France and Ireland.

	Mean	Median	Std.Dev.	Skewness	Kurtosis	Minimum	Maximum
Austria	0.00	0.00	0.05	1.07	18.74	-0.27	0.42
Belgium	0.00	0.00	0.04	0.33	8.29	-0.21	0.27
Germany	0.00	0.00	0.02	0.41	7.98	-0.09	0.10
Spain	0.00	0.00	0.08	-0.71	18.47	-0.79	0.50
France	0.00	0.00	0.02	0.14	6.38	-0.11	0.11
Greece	0.00	-0.02	0.30	-0.31	46.81	-3.64	2.91
Ireland	0.00	-0.01	0.12	0.02	9.13	-0.79	0.55
Italy	0.00	0.00	0.07	-0.82	25.54	-0.77	0.45
Netherlands	0.00	0.00	0.02	1.62	19.59	-0.10	0.24
Portugal	0.00	-0.01	0.13	-2.60	51.49	-1.85	0.74

The assumed statistical model directly influences the volatility estimates. Figure 2 plots estimated volatility levels for the three different models along with the squared CDS changes. The volatilities from the univariate Gaussian models repeatedly seem to be too high. The thin tails of the Gaussian imply that volatility needs to increase sharply in response to an extreme jump in the CDS spread, see for example the Spanish CDS spread around April 2008, and many countries around Spring 2010. In particular, the magnitude appears too large when compared to the subsequent squared CDS spread changes. The volatility estimates based on the Student- t and GHST distribution change less abruptly after incidental large changes than the Gaussian ones. The results for the two fat-tailed distribution are mutually similar and in line with the subsequent squared changes in CDS spreads. Some differences are visible for the series that exhibit significant skewness, such as the time series for Greece, Spain, and Portugal.

Table 2 reports the parameter estimates for the ten univariate country-specific models. In all cases, volatility is highly persistent, i.e., B is close to one. Note that the parameterization of our score driven model is different than that of a standard GARCH model. In particular,

Figure 2: Estimated time varying volatilities for changes in CDS for EA countries

We report three different estimates of time-varying volatility that pertain to changes in CDS spreads on sovereign debt for 10 countries. The volatility estimates are based on different parametric assumptions regarding the univariate distribution of sovereign CDS spread changes: Gaussian, symmetric t , and GHST. For comparison, the squared CDS spread changes are plotted as well.


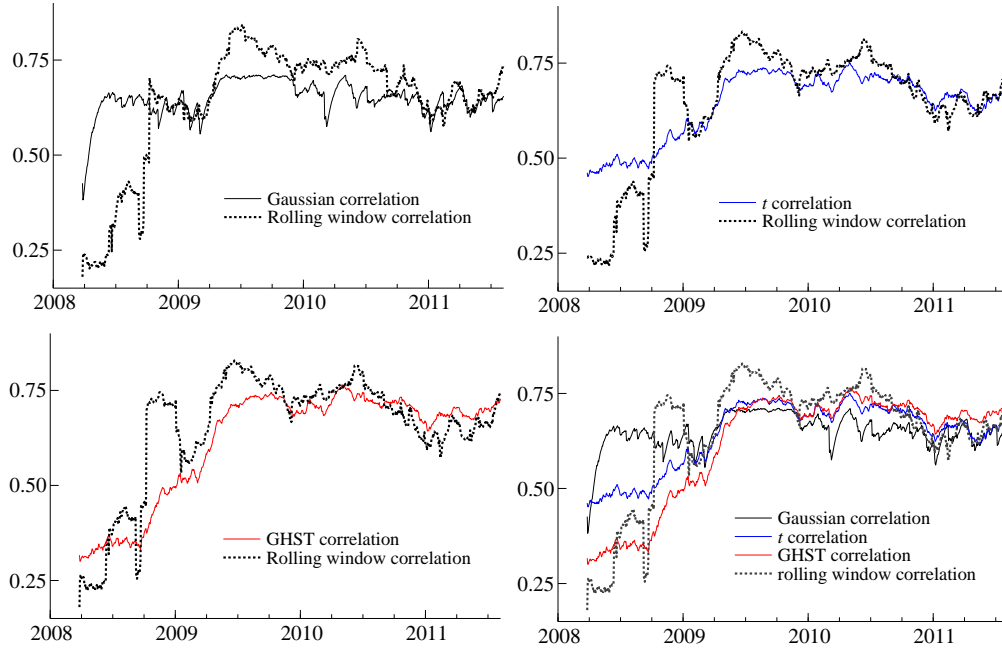


Figure 3: Average correlation over time

We plot the estimated average correlation over time, where averaging takes place over 45 estimated correlation coefficients. The correlations are estimated based on different parametric assumptions: Gaussian, symmetric t , and GH Skewed- t (GHST). The time axis runs from March 2008 to June 2011. The corresponding rolling window correlations are each estimated using a window of sixty business days of pre-filtered CDS changes. The bottom-right panel collects the six series for comparison.



the persistence is completely captured by B rather than by $A + B$ as in the GARCH case. Also note that ω sometimes takes on negative values. This is natural as we define f_t to be the log-volatility rather than the volatility itself.

Next, we estimate the dynamic correlation coefficients from the standardized CDS spread changes. Given $n = 10$, there are 45 different elements in the correlation matrix. Figure 3 plots the average correlation, averaged across 45 time varying bivariate pairs, for each model specification. As a robustness check, we benchmark each multivariate model-based estimate with the average over 45 correlation pairs obtained from a 60 business days rolling window. Over each window we use the pre-filtered marginal data as for the multivariate model estimates.

Table 2: Model parameter estimates

The table reports parameter estimates that pertain to three different model specifications. The sample consists of daily changes from January 2008 to June 2011. The degree of freedom parameter ν is set to five for the t distributions. Parameters in γ are estimated in the marginal distributions. Almost all parameters are statistically significant at the 5% level. Some parameters for the volatility factor means in the t -model are insignificant.

	AT	BE	DE	ES	FR	GR	IE	IT	NL	PT	Correlation
Gaussian											
A	0.06 (0.00)	0.10 (0.01)	0.08 (0.01)	0.15 (0.02)	0.11 (0.01)	0.12 (0.01)	0.08 (0.01)	0.11 (0.01)	0.08 (0.01)	0.16 (0.02)	0.02 (0.00)
B	0.99 (0.00)	0.98 (0.01)	0.97 (0.01)	0.94 (0.01)	0.97 (0.01)	0.99 (0.00)	0.96 (0.01)	0.99 (0.00)	0.97 (0.01)	0.99 (0.00)	0.96 (0.01)
ω	-0.03 (0.01)	-0.07 (0.02)	-0.14 (0.03)	-0.18 (0.02)	-0.12 (0.03)	0.00 (0.00)	-0.09 (0.01)	0.00 (0.00)	-0.11 (0.03)	0.00 (0.00)	1.00 (0.00)
t											
A	0.28 (0.07)	0.30 (0.31)	0.35 (0.31)	0.39 (0.18)	0.40 (0.68)	0.42 (0.00)	0.30 (0.22)	0.34 (0.17)	0.26 (0.04)	0.36 (0.04)	0.01 (0.00)
B	0.99 (0.00)	0.98 (0.00)	0.95 (0.00)	0.98 (0.00)	0.96 (0.00)	0.98 (0.00)	0.99 (0.00)	0.98 (0.00)	0.97 (0.00)	0.99 (0.00)	0.99 (0.00)
ω	0.07 (0.38)	0.05 (1.62)	-0.07 (2.03)	0.09 (0.79)	0.00 (4.14)	0.14 (0.00)	0.09 (0.84)	0.08 (0.82)	-0.02 (0.28)	0.11 (0.15)	1.01 (0.01)
ν	5 -	5 -	5 -	5 -	5 -	5 -	5 -	5 -	5 -	5 -	5 -
GHST											
A	0.13 (0.02)	0.15 (0.02)	0.21 (0.03)	0.16 (0.02)	0.22 (0.02)	0.17 (0.02)	0.14 (0.02)	0.16 (0.02)	0.16 (0.02)	0.15 (0.02)	0.01 (0.00)
B	0.99 (0.01)	0.98 (0.01)	0.93 (0.02)	0.98 (0.01)	0.95 (0.01)	0.97 (0.01)	0.98 (0.01)	0.98 (0.01)	0.96 (0.01)	0.98 (0.01)	0.99 (0.00)
ω	-0.04 (0.02)	-0.08 (0.03)	-0.29 (0.07)	-0.05 (0.03)	-0.18 (0.05)	-0.05 (0.02)	-0.05 (0.02)	-0.05 (0.03)	-0.18 (0.05)	-0.05 (0.02)	1.05 (0.01)
ν	5 -	5 -	5 -	5 -	5 -	5 -	5 -	5 -	5 -	5 -	5 -
γ	0.11 (0.04)	0.17 (0.04)	0.04 (0.04)	0.12 (0.04)	0.12 (0.04)	0.35 (0.04)	0.22 (0.04)	0.10 (0.04)	0.06 (0.04)	0.29 (0.04)	- -

If we compare the correlation estimates across the different specifications, the GHST model matches the rolling window estimates most closely. Rolling window and GHST correlations are low in the beginning of the sample at around 0.3 and increase to around 0.75 during 2010 and 2011. In the beginning of the sample the GHST-based average correlation is lower than that implied by the two alternative specifications. The pattern reverses in the second half of the sample. This result is in line with correlations that tend to increase during times of stress.

The correlation estimates across all model specifications vary considerably over time. Estimated dependence across Euro Area sovereign risk increases sharply for the first time around 15 September 2008, on the day of the Lehman failure, and around 30 September 2008, when the Irish government issued a blanket guarantee for all deposits and borrowings of six large financial institutions. Average GHST correlations remain high afterwards, around 0.75, until around 10 May 2010. At this time, Euro Area heads of state introduced a rescue package that contained government bond purchases by the ECB under the so-called Securities Markets Program, and the European Financial Stability Facility, a fund designed to provide financial assistance to Euro Area states in economic difficulties. After an eventual decline to around 0.6 towards the end of 2010, average correlations increase again towards the end of the sample.


The parameter estimates for volatility and correlations are shown in Table 2. Unlike the raw sample skewness, the estimated skewness parameters are all positive, indicating a fatter right tail of the distribution. The negative raw skewness may be the result of several influential outliers. These are accommodated in a model specification with fat-tails.

4.3 Measures of Eurozone financial stress

This section reports marginal and joint risk estimates that pertain to Euro Area sovereign default. First, Figure 4 plots estimates of CDS-implied probabilities of default (pd) over a one year horizon based on (3). These are directly inferred from CDS spreads, and do not

Figure 4: Implied marginal failure probabilities from CDS markets

We plot risk neutral marginal probabilities of failure for ten Euro Area countries extracted from CDS markets. The time axis is from January 2008 to June 2011.



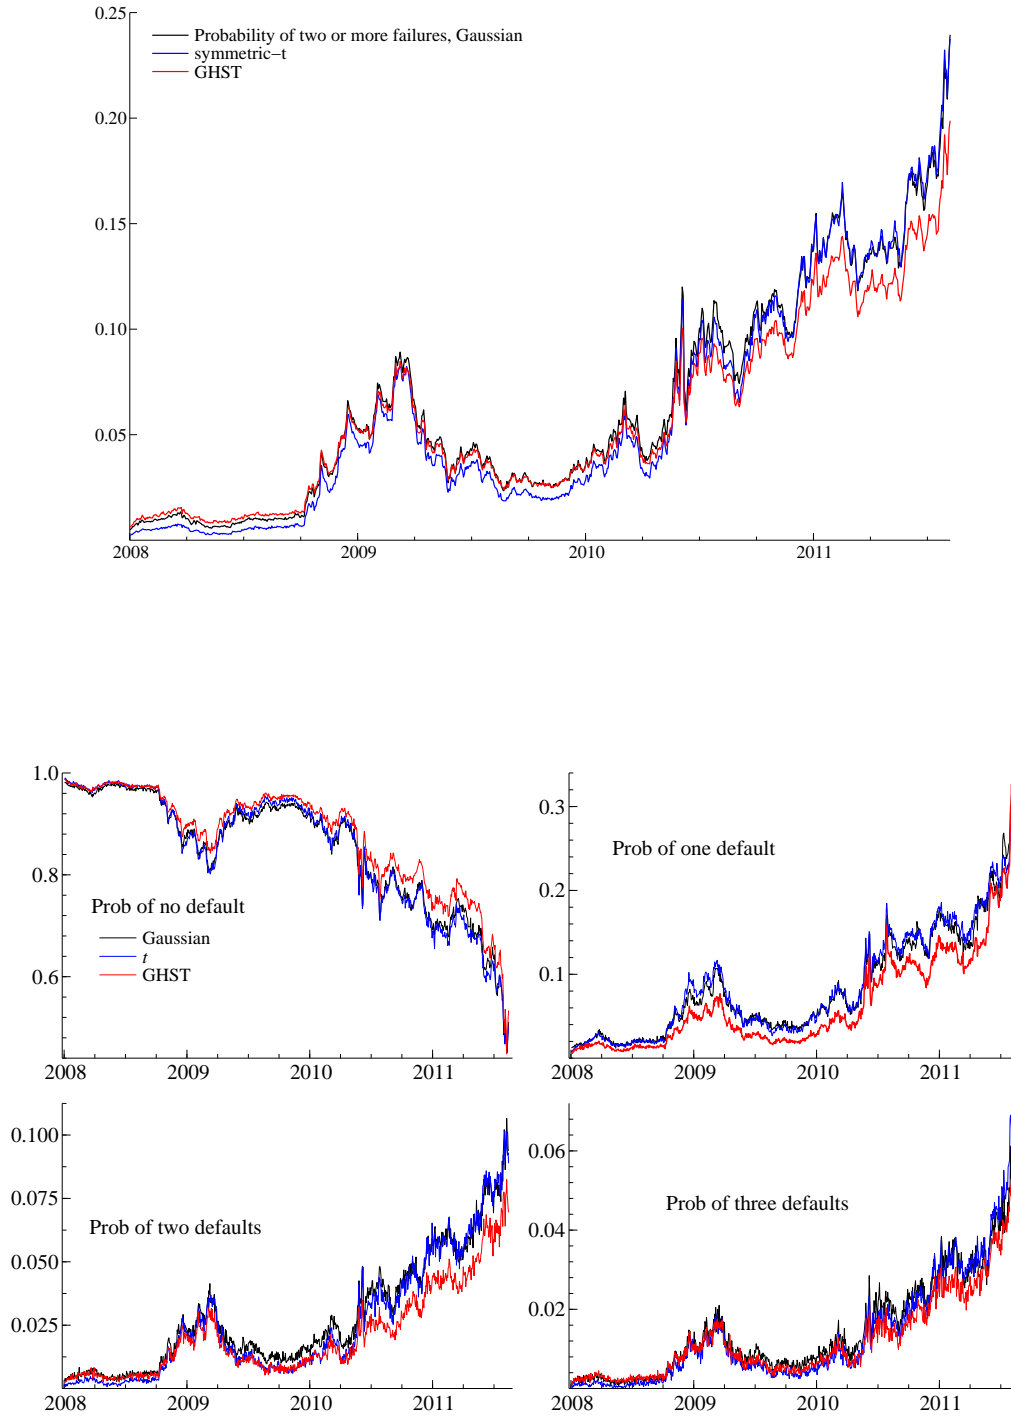
depend on parametric assumptions regarding their joint distribution. Market-implied pd's range from around 1% for Germany and the Netherlands to above 10% for Greece, Portugal, and Ireland.

The top panel of Figure 5 tracks the market-implied probability of two or more failures among the ten Euro Area sovereigns in the portfolio over a one year horizon. The joint failure risk estimate is calculated by simulation, using 10,000 draws at each time t . This simple estimate combines all marginal and joint failure information into a single time series plot and reflects the deterioration of debt conditions since the beginning of the Eurozone crisis. The overall dynamics are roughly similar across the different distributional specifications.

The risk of two or more failures over a one year horizon, as reported in Figure 5, starts to pick up in the weeks after the Lehman failure and the Irish blanket guarantee in September 2008. The joint risk estimate peaks in the first quarter of 2009, at the height of the Irish debt

Figure 5: Probability of two or more failures

The top panel plots the time-varying probability of two or more failures (out of ten) over a one-year horizon. Estimates are based on different distributional assumptions regarding marginal risks and multivariate dependence: Gaussian, symmetric- t , and GH skewed- t (GHST). The bottom panel plots model-implied probabilities for n^* sovereign failures over a one year horizon, for $n^* = 0, 1, \dots, 4$.



crisis, then decreases until the third quarter of 2009. It is increasing since then until the end of the sample. The joint probability decreases sharply, but only temporarily, around the 10 May 2010 announcement of the the European Financial Stability Facility and the European Central Bank's intervention in government debt markets starting at around the same time.

In the beginning of our sample, the joint failure probability from the GHST model is higher than the ones from the Gaussian and symmetric- t model. This pattern reverses towards the end of the sample, when the Gaussian and symmetric- t estimates are slightly higher than the GHST estimate. Towards the end of the sample, the joint probability measure is heavily influenced by the possibility of a credit event in Greece and Portugal. The CDS changes for each of these countries are positively skewed, i.e., have a longer right tail. As the crisis worsens, we observe more frequent positive and extreme changes, which increase the volatility in the symmetric models more than in the skewed setting. Higher volatility translates into higher marginal risk, or lower estimated default thresholds. This explains the (slightly) different patterns in the estimated probabilities of joint failures.

The bottom panel in Figure 5 plots the probability of a pre-specified number of failures. The lower level of our GHST joint failure probability in the top panel of Figure 5 towards the end of the sample is due to the higher probability of no defaults in that case. Altogether, the level and dynamics in the estimated measures of joint failure from this section do not appear to be very sensitive to the precise model specification.

4.4 Spillover measures: What if ... failed?

This section investigates conditional probabilities of failure. Such conditional probabilities relate to questions of the "what if?" type and reveal which countries may be most vulnerable to the failure of a given other country. We condition on a credit event in Greece to illustrate our general methodology. We pick this case since it has by far the highest market-implied probability of failure at the end of our sample period. To our knowledge, this is the first attempt in the literature on evaluating the spill-over effects and conditional probability of

sovereign failures.

Figure 6 plots the conditional probability of default for nine Euro Area countries if Greece defaults. We distinguish four cases, i.e., Gaussian dependence, symmetric- t , GHST, and GHST with zero correlations. The last experiment is included to disentangle the effect of correlations and tail dependence, see our discussion below equation (1). Regardless of the parametric specification, Ireland and Portugal seem to be most affected by a Greek failure, with conditional probabilities of failure of around 30%. Other countries may be perceived as more ‘ring-fenced’ as of June 2011, with conditional failure probabilities below 20%. The level and dynamics of the conditional estimates are sensitive to the parametric assumptions. The conditional pd estimates are highest in the GHST case. The symmetric- t estimates in turn are higher than those obtained under the Gaussian assumption. The bottom right panel of Figure 6 demonstrates that even if the correlations are put to zero, the GHST still shows extreme dependence due to the mixing variable ς_t in (1). The correlations and mixing construction thus operate together to capture the dependence in the data.

Figure 7 plots the pairwise correlation estimates for Greece with each of the remaining nine Euro Area countries. The estimated correlations for the GHST model are higher than for the other two models in the second half of the sample. This is consistent with the higher level of conditional pd’s in the GHST case compared to the other distribution assumptions, as discussed above for Figure 5. Interestingly, the dynamic correlation estimates of Euro Area countries with Greece increased most sharply in the first half of 2009. These are the months before the media attention focused on the Greek debt crisis, which was more towards the end of 2009 up to Spring 2010.

Figure 8 plots the difference between the conditional probability of failure of a given country given that Greece fails and the respective conditional probability of failure given that Greece does not fail. We refer to this difference as a spillover component or contagion effect as the differences relate to the question whether CDS markets perceive any spillovers

Figure 6: Conditional probabilities of failure given that Greece fails

We plot annual conditional failure probabilities for nine Euro Area countries given a Greek failure. We distinguish estimates based on a Gaussian dependence structure, symmetric- t , GH skewed- t (GHST), and a GHST with zero correlations.

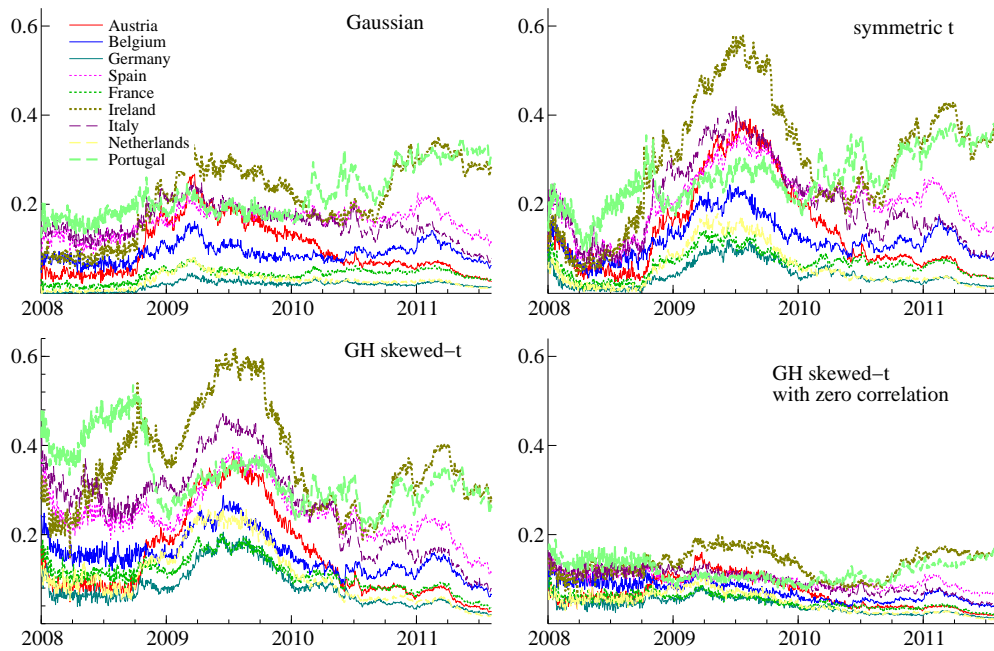


Figure 7: Dynamic correlation of Euro Area countries with Greece

We plot time-varying bivariate correlation pairs for nine Euro Area countries and Greece. The correlation estimates are obtained from the ten-dimensional multivariate model with a Gaussian, symmetric- t , and GH skewed- t (GHST) dependence structure, respectively.

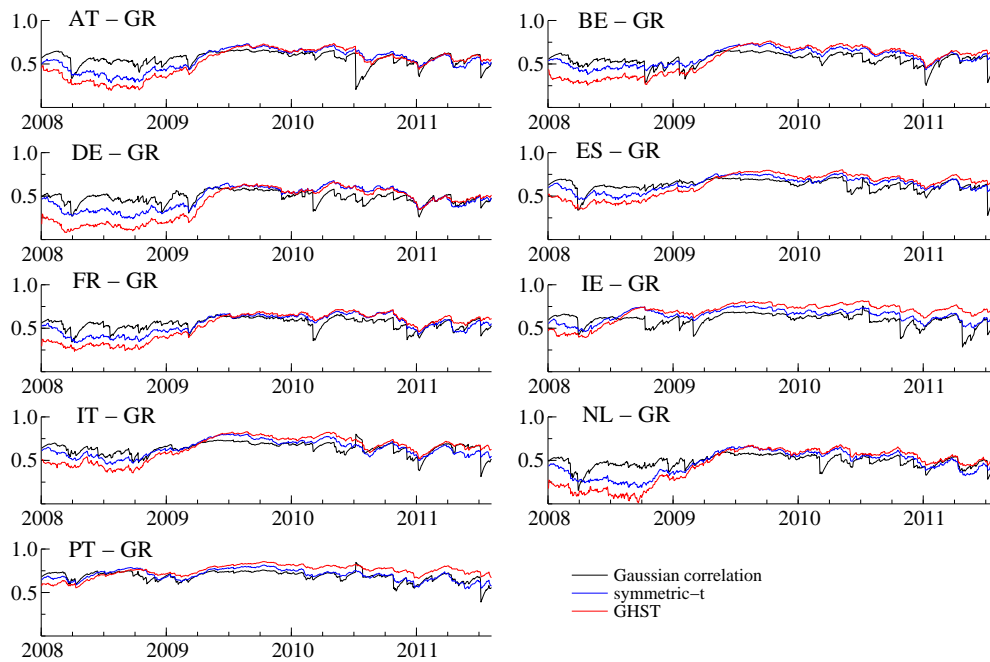
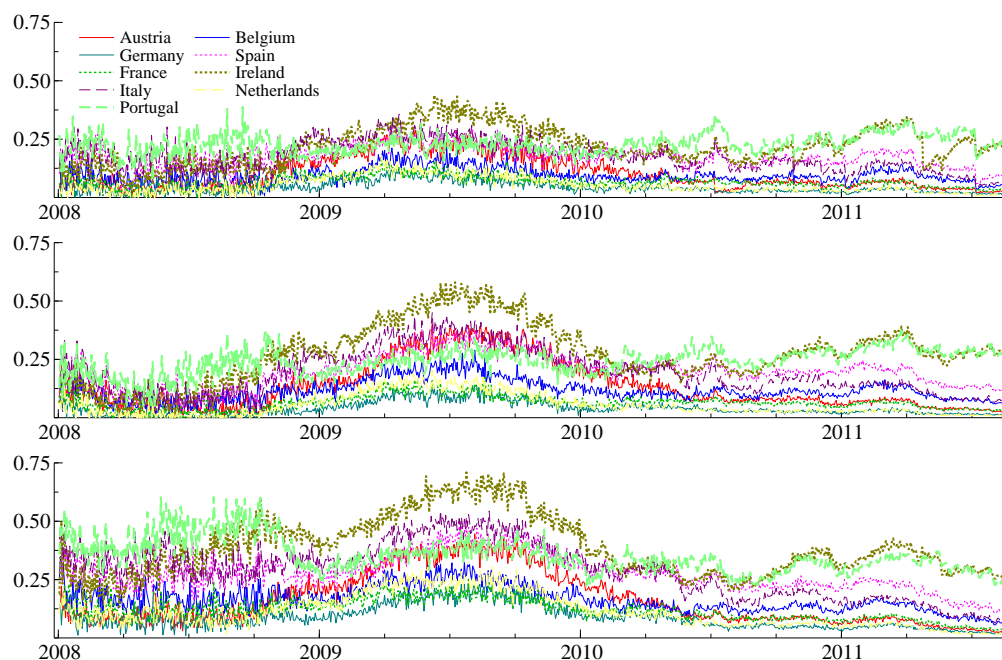


Figure 8: Risk spillover components

We plot the difference between the (simulated) probability of failure of i given that Greece fails and the probability of failure of i given that Greece does not fail. The underlying distributions are multivariate Gaussian, symmetric- t , and GH skewed- t (GHST), respectively.

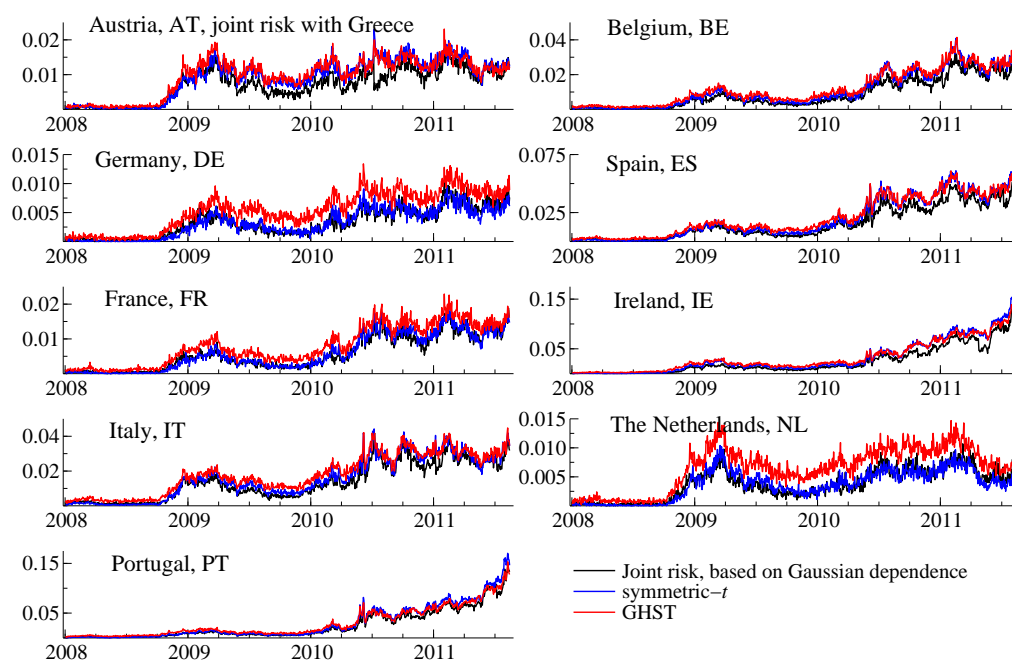


from a potential Greek default to the likelihood of other Euro Area countries failing. The level of estimated spillovers are substantial. For example, the difference in the conditional probability of a Portuguese failure given that Greece does or does not fail, is about 30%. The spillover estimates do not appear to be very sensitive to the different parametric assumptions. In all cases, Portugal and Ireland appear the most vulnerable to a Greek default since around mid-2010.

The conditional probabilities can be scaled by the time-varying marginal probability of a Greek failure to obtain pairwise joint failure risks. These joint risks are increasing towards the end of the sample and are higher in 2011 than in the second half of 2009. Annual joint probabilities for nine countries are plotted in Figure 9. For example, the risk of a joint failure over a one year horizon of both Portugal and Greece, as implied by CDS markets, is about

Figure 9: Joint default risk with Greece

We plot the time-varying probability of two simultaneous credit events in Greece and a given other Euro Area country. The estimates are obtained from a multivariate model based on a Gaussian, symmetric- t , and GH skewed- t (GHST) density, respectively.



16% at the end of our sample.

4.5 Event study: the 09 May 2010 rescue package and risk dependence

During a weekend meeting on 08-09 May, 2010, Euro Area heads of state ratified a comprehensive rescue package to mitigate sovereign risk conditions and perceived risk contagion in the Eurozone. This section analyses the impact of the resulting simultaneous announcement of the European Financial Stability Facility (EFSF) and the ECB's Securities Markets Program (SMP) on Euro Area joint risk and conditional risk as implied by our empirical model. We do so by comparing CDS-implied risk conditions closely before and after the 9th May 2010 announcement.

The agreed upon rescue fund, the European Financial Stability Facility (EFSF), is a

limited liability company with an objective to preserve financial stability of the Euro Area by providing temporary financial assistance to Euro Area member states in economic difficulties. Initially committed funds were 440bn Euro (which were increased to 780bn Euro in June 2011). The announcement made clear that EFSF funds can be combined with funds raised by the European Commission of up to 60bn Euro, and funds from the International Monetary Fund of up to 250bn Euro, for a total safety net up to 750bn Euro.

A second key component of the 09 May 2010 package consisted of the ECB's government bond buying program, the SMP. Specifically, the ECB announced that it would start to intervene in secondary government bond markets to ensure depth and liquidity in those market segments that are qualified as being dysfunctional. These purchases were meant to restore an appropriate transmission of monetary policy actions targeted towards price stability in the medium term. The SMP interventions were almost always sterilized through additional liquidity-absorbing operations.

The joint impact of the 09 May 2010 announcement of the EFSF and SMP as well as of the initial bond purchases on joint risk estimates can be seen in the top panel of Figure 5. The figure suggests that the probability of two or more credit events in our sample of ten countries decreases from about 12% to approximately 6% before and after the 09 May 2010 announcement. Figure 1 and 4 indicate that marginal risks decreased considerably as well. The graphs also suggest that these decreases were temporary. The average correlation plots in Figure 3 do not suggest a wide-spread and prolonged decrease in dependence. Instead, there seems to be an up-tick in average correlations. Overall, the evidence so far suggest that the announcement of the policy measures and initial bond purchases may have substantially lowered joint risks, but not necessarily through a decrease in joint dependence.

To further investigate the impact on joint and conditional sovereign risk from actions communicated on 09 May 2010 and implemented shortly afterwards, Table 3 reports model-based estimates of joint and conditional risk. We report our risk estimates for two dates,

Thursday 06 May 2010 and Tuesday 11 May 2011, i.e., two days before and after the announced change in policy. The top panel of Table 3 confirms that the joint probability of a credit event in, say, both Portugal and Greece, or Ireland and Greece, declines from 7.3% to 3.1% and 4.8% to 2.6%, respectively. These are large decreases in joint risk. For any country in the sample, the probability of that country failing simultaneously with Greece or Portugal over a one year horizon is substantially lower after the 09 May 2010 policy announcement than before.

The bottom panel of Table 3, however, indicates that the decrease in joint risk is generally not due to a decline in failure dependence, ‘interconnectedness’, or ‘contagion’. Instead, the conditional probabilities of a credit event in for example Greece or Ireland given a credit event in Portugal increases from 78% to 84% and from 43% to 49%, respectively. Similarly, the conditional probability of a credit event in Belgium or Ireland given a credit event in Greece increases from 10% to 15% and from 24% to 25%, respectively.

As a bottom line, based on the initial impact of the two policy measures on CDS risk pricing, our analysis suggests that the two policies may have been perceived to be less of a ‘firewall’ or ‘ringfence’ measure, i.e., intended to lower the impact and spread of an adverse development should it actually occur. Markets perceived the measures much more as a means to affect the probability of individual adverse outcomes downwards, but without decreasing dependence.

5 Conclusion

We proposed a novel empirical framework to assess the likelihood of joint and conditional failure for Euro Area sovereigns. Our methodology is novel in that our joint risk measures are derived from a multivariate framework based on a dynamic Generalized Hyperbolic skewed- t (GHST) density that naturally accommodates skewed and heavy-tailed changes in marginal risks as well as time variation in volatility and multivariate dependence. When applying the

Table 3: Joint and conditional failure probabilities

The top and bottom panels report model-implied joint and conditional probabilities of a credit event for a subset of countries, respectively. For the conditional probabilities $\Pr(i \text{ failing} \mid j \text{ failed})$, the conditioning events j are in the columns (PT, GR, DE), while the events i are in the rows (AT, BE, ..., PT). Avg contains the averages for each column.

Joint risk, $\Pr(i \text{ and } j \text{ failing})$						
	Thu 06 May 2010			Tue 11 May 2010		
	PT	GR	DE	PT	GR	DE
AT	1.54%	1.64%	0.86%	0.92%	1.04%	0.58%
BE	1.84%	1.97%	0.90%	1.34%	1.57%	0.65%
DE	1.43%	1.45%		0.83%	1.01%	
ES	4.67%	5.01%	1.22%	1.91%	2.41%	0.71%
FR	1.45%	1.44%	0.89%	1.06%	1.22%	0.67%
GR	7.28%		1.45%	3.06%		1.01%
IR	4.00%	4.75%	1.16%	1.79%	2.57%	0.77%
IT	4.06%	4.34%	1.23%	1.89%	2.29%	0.72%
NL	1.29%	1.29%	0.72%	0.92%	1.10%	0.65%
PT		7.28%	1.43%		3.06%	0.83%
Avg	3.06%	3.24%	1.10%	1.52%	1.81%	0.73%

Conditional risk, $\Pr(i \text{ failing} \mid j \text{ failed})$						
	Thu 06 May 2010			Tue 11 May 2010		
	PT	GR	DE	PT	GR	DE
AT	17%	8%	56%	25%	10%	49%
BE	20%	10%	58%	37%	15%	55%
DE	15%	7%		23%	10%	
ES	50%	25%	79%	52%	23%	60%
FR	16%	7%	58%	29%	12%	56%
GR	78%		94%	84%		85%
IR	43%	24%	75%	49%	25%	65%
IT	44%	22%	80%	52%	22%	61%
NL	14%	7%	47%	25%	11%	55%
PT		37%	93%		30%	70%
Avg	33%	16%	71%	42%	17%	62%

model to Euro Area sovereign CDS data from January 2008 to June 2011, we find significant time variation in risk dependence, as well as considerable spillover effects in the likelihood of sovereign failures. We also documented how parametric assumptions, including assumptions about higher order moments, matter for joint and conditional risk assessments. Using the 09 May 2010 new policy measures of the European heads of state, we illustrated how the model contributes to our understanding of market perceptions about specific policy measures.

References

- Acharya, V. V., L. H. Pedersen, T. Philippon, and M. Richardson (2010). Measuring systemic risk. *NYU working paper*.
- Avesani, R. G., A. G. Pascual, and J. Li (2006). A new risk indicator and stress testing tool: A multifactor nth-to-default cds basket.
- Black, F. and J. C. Cox (1976). Valuing corporate securities: Some effects of bond indenture provisions. *The Journal of Finance* 31(2), 351–367.
- Brigo, D. and F. Mercurio (2006). *Interest Rate Models: Theory and Practice*. Wiley Finance.
- Calvo, G. A. (1988). Servicing the public debt: The role of expectations. *American Economic Review* 78(4), 647–661.
- Creal, D., S. J. Koopman, and A. Lucas (2011). A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations. *Journal of Economic and Business Statistics* 29(4), 552–563.
- Creal, D., S. J. Koopman, and A. Lucas (2012). Generalized Autoregressive Score Models with Applications. *Journal of Applied Econometrics*, forthcoming.
- CreditMetrics (2007). CreditMetrics (TM) - Technical Document, RiskMetrics Group. www.riskmetrics.com/pdf/dnldtechdoc/CMTD1.pdf.
- Eaton, J. and M. Gersovitz (1981). Debt with potential repudiation: Theory and estimation. *Review of Economic Studies* 48(April), 289–309.
- ECB (2011). *European Central Bank, Financial Stability Review, June 2011*. Frankfurt: www.ecb.int/pub/pdf/other/financialstabilityreview201106en.pdf.
- Engle, R. (2002). Dynamic conditional correlation. *Journal of Business and Economic Statistics* 20(3), 339–350.
- Forbes, K. and R. Rigobon (2002). No contagion, only interdependence: measuring stock market comovements. *The Journal of Finance* 57(5), 2223–2261.
- Hu, W. (2005). *Calibration Of Multivariate Generalized Hyperbolic Distributions Using The EM Algorithm, With Applications In Risk Management, Portfolio Optimization And Portfolio Credit Risk*. Ph. D. thesis.
- Huang, X., H. Zhou, and H. Zhu (2009). A framework for assessing the systemic risk of major financial institutions. *Journal of Banking and Finance* 33, 2036–2049.
- Hull, J. and A. White (2004). Valuation of a cdo and an nth-to-default cds without monte carlo simulation. *Journal of Derivatives* 12(2).
- Hull, J. C. and A. White (2000). Valuing credit default swaps i: No counterparty default risk. *The Journal of Derivatives* 8, 29–40.
- IMF (2009). *International Monetary Fund, Global Financial Stability Review, April 2009*. Washington D.C.: <http://www.imf.org/External/Pubs/FT/GFSR/2009/01/pdf/text.pdf>.
- McNeil, A. J., R. Frey, and P. Embrechts (2005). *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press.
- Merton, R. (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance* 29(2), 449–470.
- Segoviano, M. A. and C. Goodhart (2009). Banking stability measures. *IMF Working Paper*.

- Sturzenegger, F. and J. Zettelmeyer (2008). Haircuts: Estimating investor losses in sovereign debt restructurings, 1998-2005. *Journal of International Money and Finance* 27, 780–802.
- Zhang, X., D. Creal, S. Koopman, and A. Lucas (2011). Modeling dynamic volatilities and correlations under skewness and fat tails. *TI-DSF Discussion paper 11-078/DSF22*.

A Appendix: the dynamic GH skewed- t (GHST) model

The generalized autoregressive score model for the GH skewed- t (GHST) density (5) adjusts the time-varying parameter f_t at every step using the scaled score of the density at time t . This can be regarded as a steepest ascent improvement of the parameter using the local (at time t) likelihood fit of the model. Under the correct specification of the model, the scores form a martingale difference sequence.

We partition f_t as $f_t = (f_t^v, f_t^c)$ for the (diagonal) matrix $D_t^2 = D(f_t^v)^2$ of variances and correlation matrix $R_t = R(f_t^c)$, respectively, where $\Sigma_t = D_t R_t D_t = \Sigma(f_t)$. We set $f_t^v = \ln(\text{diag}(D_t^2))$, which ensures that variances are always positive, irrespective of the value of f_t^v . For the correlation matrix, we use the hypersphere transformation also used in Zhang et al. (2011). This ensures that R_t is always a correlation matrix, i.e., positive semi-definite with ones on the diagonal. We set $R_t = R(f_t^c) = X_t X_t'$, with f_t^c as a vector containing $n(n-1)/2$ time-varying angles $\phi_{ijt} \in [0, \pi]$ for $i > j$, and

$$X_t = \begin{pmatrix} 1 & c_{12t} & c_{13t} & \cdots & c_{1nt} \\ 0 & s_{12t} & c_{23t}s_{13t} & \cdots & c_{2nt}s_{1nt} \\ 0 & 0 & s_{23t}s_{13t} & \cdots & c_{3nt}s_{2nt}s_{1nt} \\ 0 & 0 & 0 & \cdots & c_{4nt}s_{3nt}s_{2nt}s_{1nt} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & c_{n-1,nt} \prod_{\ell=1}^{n-2} s_{\ell nt} \\ 0 & 0 & 0 & \cdots & \prod_{\ell=1}^{n-1} s_{\ell nt} \end{pmatrix}, \quad (\text{A1})$$

where $c_{ijt} = \cos(\phi_{ijt})$ and $s_{ijt} = \sin(\phi_{ijt})$. The dimension of f_t^c thus equals the number of correlation pairs.

As implied by equation (13), we take the derivative of the log-density with respect to f_t ,

and obtain

$$\nabla_t = \frac{\partial \text{vech}(\Sigma_t)' \partial \text{vech}(L_t)' \partial \text{vec}(\tilde{L}_t)' \partial \ln p^{GH}(y_t|f_t)}{\partial f_t \partial \text{vech}(\Sigma_t) \partial \text{vech}(L_t) \partial \text{vec}(\tilde{L}_t)} \quad (\text{A2})$$

$$= \Psi_t' H_t' \left(w_t (y_t \otimes y_t) - \text{vec}(\tilde{\Sigma}_t) - \left(1 - \frac{\nu}{\nu-2} w_t\right) (y_t \otimes \tilde{L}_t \gamma) \right) \quad (\text{A3})$$

$$= \Psi_t' H_t' \text{vec} \left(w_t y_t y_t' - \tilde{\Sigma}_t - \left(1 - \frac{\nu}{\nu-2} w_t\right) \tilde{L}_t \gamma y_t' \right), \quad (\text{A4})$$

$$\Psi_t = \partial \text{vech}(\Sigma_t) / \partial f_t', \quad (\text{A5})$$

$$H_t = (\tilde{\Sigma}_t^{-1} \otimes \tilde{\Sigma}_t^{-1}) (\tilde{L}_t \otimes \text{I}) \left((T' \otimes \text{I}_n) \mathcal{D}_n^0 \right) (\mathcal{B}_n (\text{I}_{n^2} + \mathcal{C}_n) (L_t \otimes \text{I}_n) \mathcal{D}_n^0)^{-1}, \quad (\text{A6})$$

$$w_t = \frac{\nu + n}{2 \cdot d(y_t)} - \frac{k'_{(\nu+n)/2} \left(\sqrt{d(y_t)} \cdot \gamma' \gamma \right)}{\sqrt{d(y_t) / \gamma' \gamma}}, \quad (\text{A7})$$

where $k'_a(b) = \partial \ln K_a(b) / \partial b$ is the derivative of the log modified Bessel function of the second kind, \mathcal{D}_n^0 is the duplication matrix $\text{vec}(L) = \mathcal{D}_n^0 \text{vech}(L)$ for a lower triangular matrix L , \mathcal{D}_n is the standard duplication matrix for a symmetric matrix S $\text{vec}(S) = \mathcal{D}_n \text{vech}(S)$, $\mathcal{B}_n = (\mathcal{D}_n' \mathcal{D}_n)^{-1} \mathcal{D}_n'$, and \mathcal{C}_n is the commutation matrix, $\text{vec}(S') = \mathcal{C}_n \text{vec}(S)$ for an arbitrary matrix S . For completeness, we mention that $\tilde{L}_t = L_t T$, $\tilde{\Sigma}_t = \tilde{L}_t \tilde{L}_t'$, and

$$(T' T)^{-1} = \frac{\nu}{\nu-2} \text{I} + \frac{2\nu^2}{(\nu-2)^2(\nu-4)} \gamma \gamma'.$$

To scale the score ∇_t , Creal, Koopman, and Lucas (2012) propose the use of powers of the inverse information matrix. The information matrix for the GHST distribution, however, does not have a tractable form. Therefore, we scale by the information matrix of the symmetric Student's t distribution,

$$\mathcal{S}_t = \left\{ \Psi' (\text{I} \otimes \tilde{L}_t^{-1})' [gG - \text{vec}(\text{I}) \text{vec}(\text{I})'] (\text{I} \otimes \tilde{L}_t^{-1}) \Psi \right\}^{-1}, \quad (\text{A8})$$

where $g = (\nu + n)(\nu + 2 + n)$, and $G = \text{E}[x_t x_t' \otimes x_t x_t']$ for $x_t \sim \text{N}(0, \text{I}_n)$. Zhang et al. (2011) demonstrate that this results in a stable model that outperforms alternatives such as the DCC if the data is fat-tailed and skewed.

Using the dynamic GH model for the individual CDS series, we first estimate the parameters for the f_t^v process. Apply the equations (A4) to (A7) in the univariate setting, we

compute the f_t^v s and use them to filter the data. The time varying factor for country i 's volatility follows

$$f_{i,t+1}^v = (1 - a_i^v - b_i^v)\omega_i^v + a_i^v s_{i,t}^v + b_i^v f_{i,t}^v, \quad (\text{A9})$$

with a_i^v and b_i^v scalar parameters corresponding to the i th series.

Next, we estimate the parameters for the f_t^c process using the filtered data. Assuming the variances are constant ($D_t = I_n$), the covariance matrix Σ_t is equivalent to R_t . The matrix Ψ_t should only contain the derivative with respect to R_t . The dynamic model can be estimated directly as explained above. For parsimony, we follow a similar parameterization of the dynamic evolution of f_t^c as in the DCC model and assume

$$f_{t+1}^c = (1 - A^c - B^c)\omega^c + A^c s_t^c + B^c f_t^c, \quad (\text{A10})$$

where $A^c, B^c \in \mathbb{R}$ are scalars, and ω^c is an $n(n-1)/2$ vector. To reduce the number of parameters in the maximization, we obtain ω^c from transformed correlation matrix. All parameters are estimated by maximum likelihood. Inference is carried out by taking the negative inverse Hessian of the log likelihood at the optimum as the covariance matrix for the estimator.