

Conditional probabilities for euro area sovereign default risk*

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

^(a) VU University Amsterdam, and Tinbergen Institute

^(b) Duisenberg school of finance

^(c) European Central Bank, Financial Research

November 12, 2012

*Author information: André Lucas, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, Email: a.lucas@vu.nl. Bernd Schwaab, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt, Germany, Email: bernd.schwaab@ecb.int. Xin Zhang, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, Email: x.zhang@vu.nl. We thank Markus Brunnermeier and Laurent Clerc for early comments, and seminar participants at Bundesbank, ECB, HEC Lausanne, Riksbank, the Royal Economics Society 2012 meeting at Cambridge, the ILMC risk management conference in Rome, and the Society for Financial Econometrics 2012 meeting in Oxford. André 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 for euro area sovereign default risk

Abstract

We propose an empirical framework to infer the likelihood of joint and conditional sovereign defaults from observed CDS prices. Our model is based on a dynamic skewed-t copula which 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 risk dependence can increase in times of stress. We apply the framework to euro area sovereign CDS spreads from 2008 to mid-2011. Our results reveal significant time-variation in distress dependence and considerable spill-over effects in the likelihood of sovereign default. We also investigate distress dependence around a policy announcement on May 9, 2010, and demonstrate the importance of capturing higher-order time-varying moments during times of crisis for the correct assessment of interacting risks.

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

JEL classifications: C32, G32.

Non-technical summary

The issue of measuring and monitoring interconnected sovereign default risk has received a lot of interest in the wake of the sovereign debt crisis in the euro area. This paper introduces a novel empirical framework to estimate marginal, joint, and conditional probabilities of sovereign default from observed prices for credit default swaps (CDS) on sovereign debt. Such joint and conditional probabilities of default can be informative for risk management and financial sector surveillance purposes. Furthermore, conditional probabilities shed light on the extent to which euro area capital markets are affected by contagion concerns, and to what extent priced credit risks are affected by policy measures. Our methodology is novel in that our joint risk measures are derived from a multivariate framework that is based on a flexible multivariate density that naturally accommodates skewed and heavy tailed changes in marginal risks as well as time variation in uncertainty and multivariate dependence.

This paper makes four main contributions to the literature on risk assessment. First, we provide time series estimates of the variation in euro area joint and conditional sovereign default risk based on our modeling framework and CDS data from January 2008 to June 2011. For example, we compute the conditional probability of, say, a default on Portuguese debt given a credit event in Greece, as expected by the market. In addition, we may infer which countries appear more exposed to certain credit events than others. As the model-implied probabilities are inferred solely from market prices, such estimates reflect what is expected by market participants at a given point in time, and are uninformative about any policy stance. As a second contribution, we analyze the extent to which parametric assumptions matter for joint and conditional sovereign risk assessments. We demonstrate that 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 spill-overs in the likelihood of sovereign failure. In this regard we document spill-overs from the possibility

of a credit event in one euro area country to the perceived riskiness of other countries. These spill-overs effects are pronounced and suggest that the cost of debt refinancing in some European countries depends on developments in other countries. Fourth, we provide an in-depth analysis of the impact on sovereign joint and conditional risks of a key policy announcement on May 9, 2010. This event study shows how our model can be used to disentangle market assessments of joint and conditional probabilities.

Our modeling framework has several attractive features that each match empirical stylized facts regarding changes in euro area sovereign CDS prices, i.e., volatility clustering, dynamic correlation, non-trivial tail dependence, and the ability to handle a cross sectional dimension of intermediate size. Univariate models first accommodate the pronounced volatility clustering in the marginal risks — in addition to skewness and substantial fat tails. Second, the correlation structure between individual countries' risk is likely to be time-varying, as correlations tend to increase during times of stress, and policy decisions may have a direct impact on the implied dependence structure. We consequently propose a model with a dynamic correlation matrix from the outset. Third, we would like to allow for non-Gaussian features such as multivariate extreme tail dependence when modeling interconnected joint risks. Fourth, the model needs to be flexible enough to be calibrated repeatedly to current market conditions, such as current CDS spread levels. Finally, and most obviously, an application to euro area sovereign risk requires us to handle dimensions larger than usual in the non-Gaussian dependence literature (less than five). These considerations directly favor the use of the novel modeling framework employed in this paper.

1 Introduction

The Eurozone debt crisis raises the issue of measuring and monitoring interconnected sovereign credit risk. In this paper we construct a novel empirical framework to assess the likelihood of joint and conditional default for euro area sovereigns. This 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 default, or failure, as any credit event that would trigger a sovereign CDS contract. Examples of such events are the non-payment of principal or interest when it is due, a forced exchange of debt into claims of lower value, or a moratorium or official repudiation of the debt. Unlike marginal probabilities, conditional probabilities of sovereign default cannot be obtained from raw market data alone, but instead require a proper joint modeling framework. Our methodology is novel in that our probability assessments are derived from a multivariate framework based on a dynamic Generalized Hyperbolic (GH) skewed- t density that naturally accommodates all relevant empirical features of the data, such as skewed and heavy-tailed changes in individual country CDS spreads, 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 four main contributions. First, we provide estimates of the time variation in euro area joint and conditional sovereign default risk using a new model and a 10-dimensional data set of sovereign CDS spreads from January 2008 to June 2011. For example, we estimate the conditional probability of a default on Portuguese debt given a Greek failure to be around 30% at the end of our sample in mid-June 2011. We report similar conditional probabilities for other countries. At the same time, we infer which countries are more exposed than others to certain credit events.

Second, we analyze the extent to which parametric modeling assumptions matter for such

joint and conditional risk assessments. Perhaps surprisingly, and despite the potential 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, symmetric- t , and GH skewed- t (GHST) specification. The distributional assumptions turn out to be most important for our conditional assessments, whereas simpler joint failure probability estimates are less sensitive to the assumed dependence structure. In particular, and much in line with Forbes and Rigobon (2002), we show that it is important to account for the different salient features of the data, such as non-zero tail dependence and skewness, when interpreting estimates of time-varying volatilities and increases in correlation in times of stress.

Third, our modeling framework allows us to investigate the presence and severity of market implied spill-overs in the likelihood of sovereign default. Specifically, we document spill-overs from the possibility of a Greek default to the perceived riskiness of other euro area countries. For example, at the end of our sample we find a difference of about 25% between the one-year conditional probability of a Portuguese default given that Greece does versus that Greece does not default. This indicates that the cost of debt refinancing in some European countries may depend to a considerable extent on developments in other countries.

Fourth, we provide an in-depth analysis of the impact on sovereign joint and conditional risks of a key policy announcement on May 9, 2010. On this day, euro area heads of state announced a comprehensive rescue package to mitigate sovereign risk conditions and perceived risk contagion in the Eurozone. The rescue package contained the European Financial Stability Facility (EFSF), a rescue fund, and the ECB's Securities Markets Program (SMP), under which the central bank can purchase government bonds in secondary markets. This event study shows how our model can be used to disentangle market assessments of joint and conditional probabilities. In particular, for May 9, 2010 we find that market perceptions of *joint* sovereign default risk have decreased, while market perceptions of *conditional*

sovereign default risk have not decreased at the same time. From a risk perspective, our joint approach is in line with for example Acharya, Pedersen, Philippon, and Richardson (2010) who focus on financial institutions: 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. While adverse developments in one country's public finances could perhaps still be handled with the support of the remaining healthy countries in the Eurozone, the situation may become more and more problematic if two, three, or more countries are in distress. Relevant questions regarding joint and conditional sovereign default risks would be hard if not impossible to answer without an empirical model such as the one proposed in this paper.

The use of CDS data to estimate market implied failure probabilities means that our probability estimates combine physical failure probabilities with the price of sovereign default risk. As a result, our risk measures constitute an upper bound for an investor worried about losing money due to a joint sovereign failure. This has to be kept in mind when interpreting the empirical results later on. Estimating failure probabilities directly from observed defaults, however, is impossible in our context, as OECD defaults are not observed over our sample period. Even if such defaults would have been observed, they would not have allowed us to perform the detailed empirical analysis on the dynamics of joint and conditional failure probabilities.

The literature on sovereign credit risk has expanded rapidly and branched off into different fields. Part of the literature focuses on the theoretical development of sovereign default risk and strategic default decisions; see for example Guembel and Sussman (2009) or Yue (2010). Another part of the literature tries to disentangle the different priced components of sovereign credit risk using asset pricing methodology, including the determination of common risk factors across countries; see for example Pan and Singleton (2008), Longstaff, Pan, Pedersen, and Singleton (2011), or Ang and Longstaff (2011). Finally, there is a line of

literature that investigates the link between sovereign credit risk, country ratings, and macro fundamentals; see for example Haugh, Ollivaud, and Turner (2009), Hilscher and Nosbusch (2010), or DeGrauwe and Ji (2012).

Our paper primarily relates to the empirical literature on sovereign credit risk as proxied by sovereign CDS spreads and focuses on spill-over risk as perceived by financial markets. We take a pure time-series perspective instead of assuming a specific pricing model as in Longstaff, Pan, Pedersen, and Singleton (2011) or Ang and Longstaff (2011). The advantage of such an approach is that we are much more flexible in accommodating all the relevant empirical features of CDS changes given that we are not bound by the analytical (in)tractability of a particular pricing model. This appears particularly important for the data at hand. In particular, our paper relates closely to the statistical literature for multiple defaults, such as for example Li (2001), Hull and White (2004) or Avesani, Pascual, and Li (2006). These papers, however, typically build on a Gaussian or sometimes symmetric Student t dependence structure, whereas we impose a dependence structure that allows for non-zero tail dependence, skewness, and time variation in both volatilities and correlations. Our approach therefore also relates to an important strand of literature on modeling dependence in high dimensions, see for example Demarta and McNeil (2005), Christoffersen, Errunza, Jacobs, and Langlois (2011), Oh and Patton (2012), and Engle and Kelly (2009), as well as to a growing literature on observation-driven time varying parameter models, such as for example Patton (2006), Harvey (2010), and Creal, Koopman and Lucas (2011, 2012). Finally, we relate to the CIMDO framework of Segoviano and Goodhart (2009). This is based on a multivariate prior distribution, usually Gaussian or symmetric- t , that can be calibrated to match marginal risks as implied by the CDS market. Their multivariate density becomes discontinuous at so-called threshold levels: some parts of the density are shifted up, others are shifted down, while the parametric tails and extreme dependence implied by the prior remain intact at all times. Our model does not have similar discontinuities, while it allows

for a similar calibration of default probabilities to current CDS spread levels as Segoviano and Goodhart (2009).

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

2 Conceptual framework

In a corporate credit risk setting, the probability of default 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 the debt matures, see Merton (1974) and Black and Cox (1976). To allow for default clustering, the default processes of individual firms can be linked together using a copula function, see for example McNeil, Frey, and Embrechts (2005). In a sovereign credit risk setting, a similar approach can be adopted, though the interpretation has to be slightly altered given the different nature of a sovereign compared to a corporate default. 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 for example Calvo (1988). Default costs may arise from losing credit market access for some time, obstacles to conducting international trade, difficulties in borrowing in the domestic market, etc., while default benefits include immediate debt relief.

To accommodate this interpretation, we introduce a variable v_{it} that triggers default if v_{it} exceeds a threshold value c_{it} . The variable v_{it} captures the time-varying changes in the difference between the perceived benefits and cost of default for sovereign i at time t . 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). The v_{it} s, $i = 1, \dots, n$, are linked together via a Generalized Hyperbolic

Skewed Student's t (GHST) copula,

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

where $\epsilon_t \in \mathbb{R}^n$ is a vector of standard normally distributed risk factors, \tilde{L}_t is an $n \times n$ matrix of risk factor sensitivities, and $\gamma \in \mathbb{R}^n$ is a vector controlling the skewness of the copula. The random scalar $\varsigma_t \in \mathbb{R}^+$ is assumed to be an inverse-Gamma distributed risk factor that affects all sovereigns simultaneously, where ς_t and ϵ_t are independent, and $\mu_\varsigma = \mathbb{E}[\varsigma_t]$. The GHST model can 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.

Default dependence in model (1) stems from two sources: common exposures to the normally distributed risk factors ϵ_t as captured by the time-varying matrix \tilde{L}_t ; and a common exposure to the scalar risk factor ς_t . The former captures spillover effects through the 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 (1) drops out of the equation and there is zero tail dependence. Conversely, if ς_t is large, all sovereigns are affected at the same time, making joint defaults of two or more sovereigns more likely.

The probability of default p_{it} of sovereign i at time t is 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} . In our case, $F_i(\cdot)$ is the univariate GHST distribution, which follows directly from the mean-variance mixture construction in equation (1). Our main interest, however, is not in the marginal default probability p_{it} , but rather in the joint default probability $\Pr[v_{it} > c_{it}, v_{jt} > c_{jt}]$ or the conditional default probability $\Pr[v_{it} > c_{it} \mid v_{jt} > c_{jt}]$, for $i \neq j$. The (market implied) marginal default probabilities are typically estimated directly from CDS market data under a number of

simplifying assumptions. We follow this practice. First, we fix the recovery rate at a stressed level of $rec_i = 25\%$ for all countries and use the 6 months LIBOR rate as the discount rate r_t . We assume that the premium payments occur continuously, such that the standard CDS pricing formula as in for example Hull and White (2000) simplifies and can be inverted to extract the market-implied probability of default p_{it} . The relation is given by

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

where s_{it} is the CDS spread for sovereign i at time t , and r_t is our discount rate; see also Brigo and Mercurio (2006, Chapter 21) and Segoviano and Goodhart (2009).

Given our market implied estimates of the default probabilities, we can make use of our multivariate model in (1) to infer the magnitude and time-variation in joint and conditional default probabilities. To do this, we proceed in two simple steps. In the first step, we estimate the dependence structure in (1) from observed CDS data as explained in Section 3, and we infer the threshold values c_{it} by inverting the univariate GHST distributions using our market implied estimates of the default probabilities. In the second step, we then use the calibrated thresholds c_{it} and the estimated dependence structure of the v_{it} s to simulate joint and conditional default probabilities. We show in Section 4 how the combination of marginal default probabilities calibrated to current CDS spread levels with the time-varying copula structure in (1) can lead to new insights into sovereign credit spread spillovers.

3 Statistical model

3.1 Generalized Autoregressive Score dynamics

As mentioned in Section 2, we use sovereign CDS spreads to estimate the time-varying dependence structure in (1) and to calibrate the model's marginal default probabilities through equation (3). The statistical model, therefore, closely follows the set-up of the previous section while allowing for time variation in the parameters using the Generalized Autoregressive

Score dynamics of Creal, Koopman, and Lucas (2012).

We assume that we observe a vector $y_t \in \mathbb{R}^n$, $t = 1, \dots, T$, of changes in sovereign CDS spreads for sovereign $i = 1, \dots, n$, where

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

with $\mu \in \mathbb{R}^n$ a vector of fixed unknown means, and e_t a GHST distributed random variable with zero mean, ν degrees of freedom, skewness parameter γ , and covariance matrix I . Equation (4) is the empirical counterpart to (1). 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 $y_t - \mu$. The density of y_t is denoted by

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

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

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

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 CDS data y_t creates challenges for standard dynamic specifications of 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 distributional 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, which are prevalent in the CDS data used in our empirical analysis. 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 number of factors coincides with the number of free elements in Σ_t in our empirical application later on, but may also be smaller. 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

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

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 familiar models such as the (multivariate) GARCH model, the autoregressive conditional duration (ACD) model, and the multiplicative error model (MEM); see Creal, Koopman, and Lucas (2012) for more examples. Details on the parameterization $\Sigma_t = \Sigma(f_t)$, $D_t = D(f_t)$, and $R_t = R(f_t)$, and the scaling matrix \mathcal{S}_t used in our empirical application can be found in the appendix.

Using the GHST specification in equation (5), the appendix shows 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 that decreases in the Mahalanobis distance of y_t from its center $\tilde{\mu}_t$ as defined in (6). The matrices Ψ_t and H_t are time-varying, parameterization specific and depend on f_t , but not on the data. Due to the presence of w_t in (14), observations that are far out 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 (i.e., $\nu < \infty$).

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 a GAS(1,1) dynamic specification with $p = q = 1$ 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 again a GAS(1,1) dynamic specification. The GHST distribution in this second step has mean zero, skewness parameters $\hat{\gamma}_i$, $i = 1, \dots, n$, as estimated in the first step, the same pre-determined value for ν , 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 computational 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 their 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 January 1, 2008 to June 30, 2011, yielding $T = 913$ observations. The CDS contracts have a five year maturity. They are denominated in U.S. dollars and therefore do not depend on foreign exchange risk concerns should a European credit event materialize. Such contracts are also far more liquidly traded than their Euro denominated counterparts. All time series data are obtained from Bloomberg. We prefer CDS spreads to bond yield spreads as a measure of sovereign default risk since the former are less affected by liquidity and flight-to-safety issues, see for example Pan and Singleton (2008) and Ang and Longstaff (2011). In addition, our CDS series are likely to be less affected than bond yields by the outright government bond purchases that might have taken place under the Securities Markets Program during the second half of our sample, see Section 4.5 below.

Table 1 provides summary statistics for daily de-meaned changes in these ten CDS

spreads. All time series have significant non-Gaussian features under standard tests and significance levels. In particular, we note the non-zero skewness and large values of kurtosis for almost all time series in the sample. 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 specification (11). We consider three different choices for the parameters, corresponding to a Gaussian, a Student- t , and a GHST distribution, respectively. We treat the degrees of freedom parameter ν as a robustness parameter; compare Franses and Lucas (1998). This implies we fix the degrees of freedom at $\nu = 5$ rather than estimating it. The advantage of such an approach is that it further simplifies the estimation process, while retaining many of the robustness features of model (11). In particular, fixing ν at $\nu = 5$ may seem high at first sight given some of the high kurtosis values in Table 1. The value is small enough, however, to result in a substantial robustification of the results via the weights w_t in (14), both in terms of likelihood evaluation as well as in terms of the volatility and correlation dynamics.

Figure 1 plots estimated volatility levels for the three different models along with the squared CDS changes. The assumed statistical model (Gaussian, Student- t , GHST) directly influences the volatility estimates. The volatilities from the univariate Gaussian models repeatedly seem to be too high. The thin tails of the Gaussian distribution imply that volatility increases sharply in response to a 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 of the increase in volatility 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 due to the weighting mechanism in (14). The results for the Student- t and GHST

are very 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, 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 for the standardized CDS spread changes. Given $n = 10$, there are 45 different elements in the correlation matrix. Figure 2 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 to the average over 45 correlation pairs obtained from a 60 business days rolling window. Over each window we use the same pre-filtered marginal data as for the multivariate model estimates.

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 vary considerably over time across all model specifications considered. Estimated dependence across euro area sovereign risk increases sharply for the first time around September 15, 2008, on the day of the Lehman failure, and around September




Figure 1: 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. As a direct benchmark, the squared CDS spread changes are plotted as well.

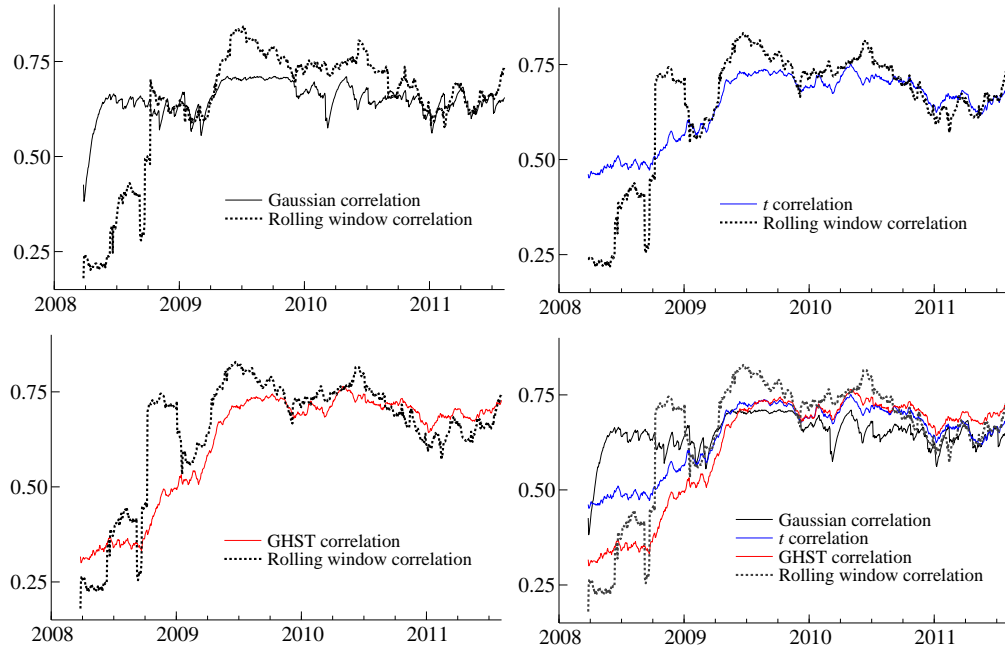


Figure 2: Average correlation over time

Plots of 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 four series for comparison.

30, 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 May 10, 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 of CDS changes. The negative raw skewness may be the result of several influential outliers. These are accommodated in a model specification with fat-tails.

4.3 Joint probabilities of Eurozone financial stress

This section reports marginal and joint risk estimates that pertain to euro area sovereign default. First, Figure 3 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 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 at the end of our sample.


The top panel of Figure 4 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 probability is calculated by simulation, using 50,000 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 probability of two or more failures over a one year horizon, as reported in Figure

Table 1: CDS 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

**Figure 3: Implied marginal failure probabilities from CDS markets**

The 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.

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.

	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)
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)
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)	- -

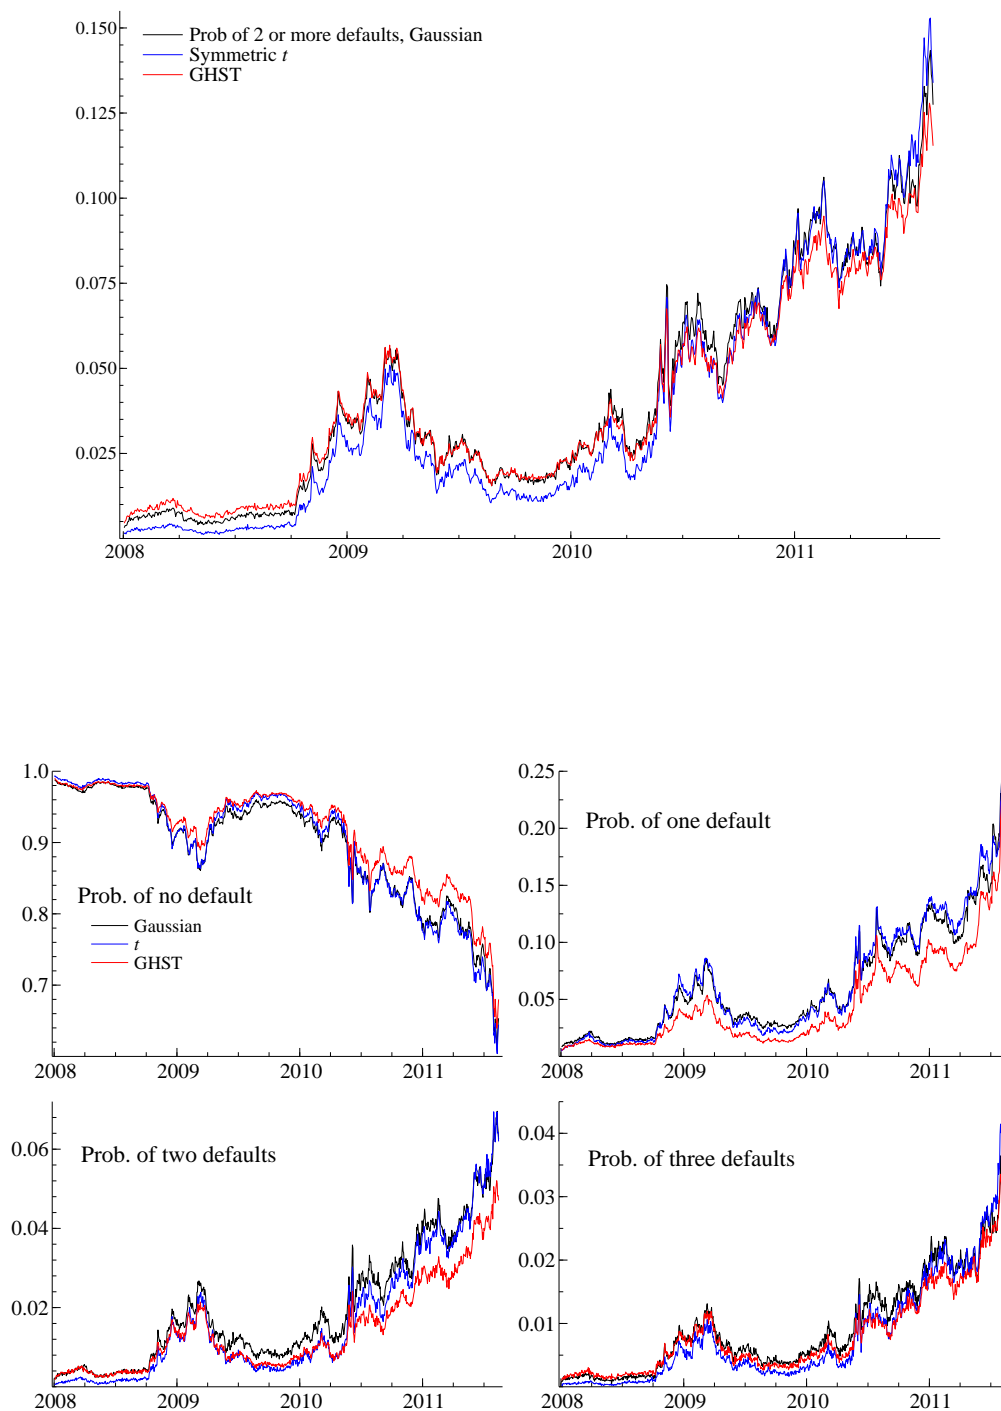


Figure 4: Probability of two or more failures

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 $n^* = 0, 1, 2, 3$.

4, starts to pick up in the weeks after the Lehman failure and the Irish blanket guarantee in September 2008. The joint probability estimate peaks in the first quarter of 2009, at the height of the Irish debt 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 May 10, 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. We come back to this later.

In the beginning of our sample, the joint failure probability from the GHST model is higher than that 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 4 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 4 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. Clearly, conditioning on a credit event is different from conditioning on incremental changes in other countries' risks, see Caceres, Guzzo, and Segoviano (2010) and Caporin, Pelizzon, Ravazzolo, and Rigobon (2012).

Figure 5 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 default probability 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 5 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 6 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 probabilities of default in the GHST case compared to the other distributional assumptions, as discussed above for Figure 4. 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




Figure 5: Conditional probabilities of failure given that Greece fails

Plots of 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.

crisis, which was more towards the end of 2009 up to Spring 2010.

Figure 7 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 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 25%. 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 8. 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 10% at the end of our sample.

4.5 Event study: the May 9, 2010 rescue package and risk dependence

During a weekend meeting on May 8–9, 2010, euro area heads of state agreed on 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 announcement of May 9, 2010.

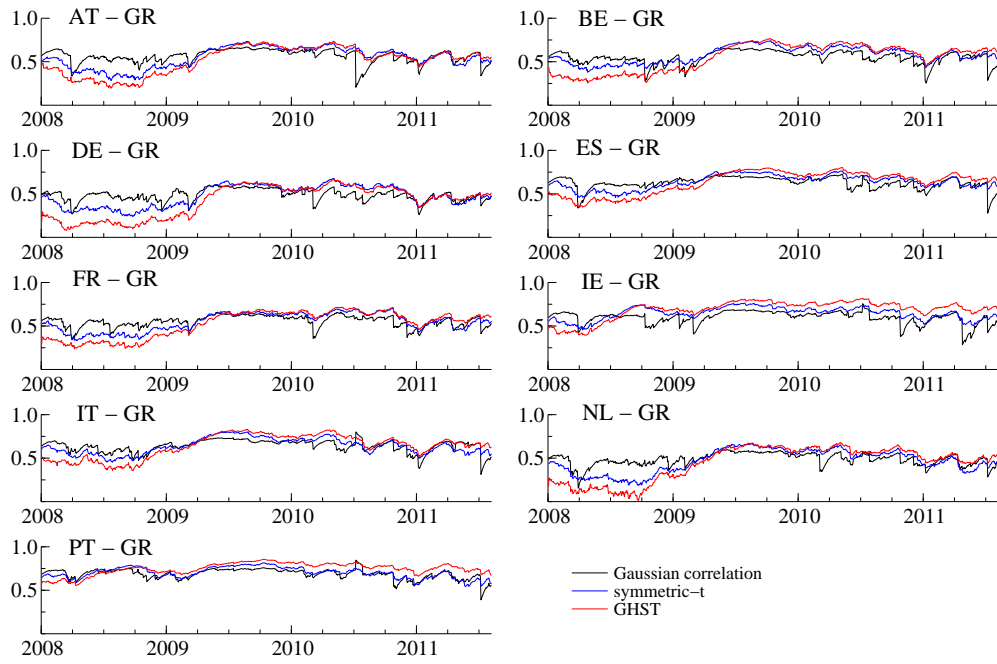


Figure 6: Dynamic correlation of euro area countries with Greece

The 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 7: Risk spillover components

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.

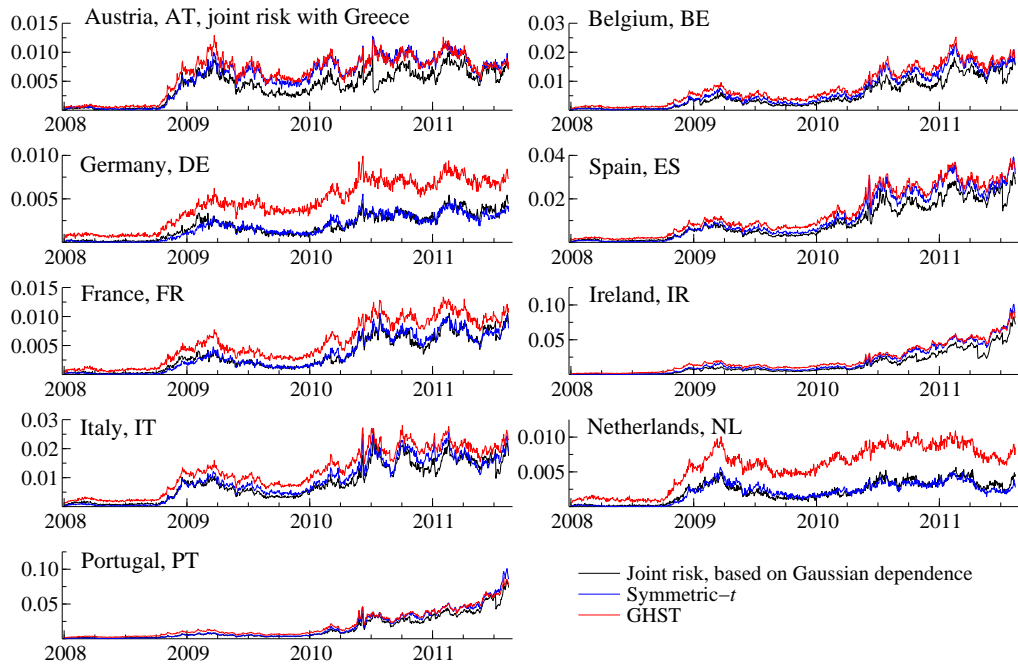


Figure 8: Joint default risk with Greece

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.

The agreed upon rescue fund, the European Financial Stability Facility (EFSF), is a limited liability facility 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. 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 May 9, 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 May 9, 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 4. The figure suggests that the probability of two or more credit events in our sample of ten countries decreases from about 7% to approximately 3% before and after the May 9, 2010 announcement. Figure 3 indicates that marginal risks decreased considerably as well. The graphs also suggest that these decreases were temporary. The average correlation plots in Figure 2 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 May 9, 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 May 6, 2010 and Tuesday May 11, 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 4.8% to 2.3% and 3.1% to 1.8%, 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 May 9, 2010 policy announcement than before.

The bottom panel of Table 3, however, indicates that the decrease in joint failure probabilities is generally not due to a decline in failure dependence. Instead, the conditional probabilities of a credit event in for example Greece or Ireland given a credit event in Portugal increases from 78% to 80% 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 23% to 26%, respectively.

As a bottom line, based on the initial impact of the two policy measures on CDS prices, 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. Rather, markets perceived the measures much more as a means to affect the probability of individual adverse outcomes downwards, but without decreasing dependence. These findings are robust to, for example, alternative choices for the degrees of freedom parameter ν in the copula, and different choices for the expected recovery rate in case of defaults.

5 Conclusion

We have 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

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 May 6, 2010			Tue May 11, 2010		
	PT	GR	DE	PT	GR	DE
AT	1.1%	1.1%	0.6%	0.6%	0.7%	0.4%
BE	1.2%	1.4%	0.7%	0.9%	1.0%	0.6%
DE	1.0%	1.1%		0.8%	0.8%	
ES	3.0%	3.3%	0.9%	1.5%	1.6%	0.6%
FR	1.0%	1.0%	0.6%	0.8%	0.9%	0.6%
GR	4.8%		1.1%	2.3%		0.8%
IR	2.6%	3.1%	0.8%	1.4%	1.8%	0.6%
IT	2.8%	2.9%	0.9%	1.4%	1.5%	0.6%
NL	0.9%	0.9%	0.5%	0.6%	0.7%	0.5%
PT		4.8%	1.0%		2.3%	0.8%
Avg	2.0%	2.2%	0.8%	1.1%	1.2%	0.6%

Conditional risk, $\Pr(i \text{ failing} \mid j \text{ failed})$						
	Thu May 6, 2010			Tue May 11, 2010		
	PT	GR	DE	PT	GR	DE
AT	17%	8%	52%	22%	10%	46%
BE	20%	10%	60%	32%	15%	61%
DE	16%	8%		26%	12%	
ES	49%	25%	78%	50%	23%	63%
FR	16%	8%	58%	28%	12%	62%
GR	78%		99%	80%		86%
IR	43%	23%	75%	49%	26%	68%
IT	45%	22%	77%	49%	21%	64%
NL	14%	7%	49%	21%	10%	50%
PT		36%	91%		33%	81%
Avg	33%	17%	71%	40%	18%	64%

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 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 May 9, 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*.
- Ang, A. and F. Longstaff (2011). Systemic sovereign credit risk: Lessons from the US and Europe. *NBER discussion paper 16983*.
- 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. *IMF Working Paper, WP/06/105*.
- 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.
- Caceres, C., V. Guzzo, and M. Segoviano (2010). Sovereign Spreads: Global Risk Aversion, Contagion or Fundamentals? *IMF working paper WP/10/120*.
- Calvo, G. A. (1988). Servicing the public debt: The role of expectations. *American Economic Review* 78(4), 647–661.
- Caporin, M., L. Pelizzon, F. Ravazzolo, and R. Rigobon (2012). Measuring sovereign contagion in Europe. *mimeo*.
- Christoffersen, P., V. Errunza, K. Jacobs, and H. Langlois (2011). Is the Potential for International Diversification Disappearing? *working paper*.
- 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.
- DeGrauwe, P. and Y. Ji (2012). Mispricing of sovereign risk and multiple equilibria in the Eurozone. *working paper*.
- Demarta, S. and A. J. McNeil (2005). The t copula and related copulas. *International Statistical Review* 73, 111–129.
- Engle, R. (2002). Dynamic conditional correlation. *Journal of Business and Economic Statistics* 20(3), 339–350.
- Engle, R. F. and B. T. Kelly (2009). Dynamic Equicorrelation. *SSRN eLibrary*.
- Forbes, K. and R. Rigobon (2002). No contagion, only interdependence: measuring stock market comovements. *The Journal of Finance* 57(5), 2223–2261.
- Franses, P. and A. Lucas (1998). Outlier detection in cointegration analysis. *Journal of Business & Economic Statistics*, 459–468.
- Guembel, A. and O. Sussman (2009). Sovereign debt without default penalties. *Review of Economic Studies* 76(4), 1297–1320.
- Harvey, A. (2010). Exponential Conditional Volatility Models. *working paper*.
- Haugh, D., P. Ollivaud, and D. Turner (2009, July). What drives sovereign risk premiums? an analysis of recent evidence from the euro area. OECD Economics Department Working Papers 718, OECD Publishing.
- Hilscher, J. and Y. Nosbusch (2010). Determinants of sovereign risk: Macroeconomic fundamentals and the pricing of sovereign debt. *Review of Finance* 14(2), 235–262.

- 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. C. and A. White (2000). Valuing credit default swaps i: No counterparty default risk. *The Journal of Derivatives* 8, 29–40.
- Hull, J. C. and A. White (2004). Valuation of a cdo and an nth-to-default cds without monte carlo simulation. *Journal of Derivatives* 12(2).
- Li, D. (2001). On default correlation: a copula function approach. *Journal of Fixed Income* 9, 43–54.
- Longstaff, F., J. Pan, L. Pedersen, and K. Singleton (2011). How sovereign is sovereign credit risk? *American Economic Journal: Macroeconomics* 3(2), 75–103.
- 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.
- Oh, D. H. and A. J. Patton (2012). Modelling dependence in high dimensions with factor copulas. *working paper*.
- Pan, J. and K. Singleton (2008). Default and recovery implicit in the term structure of sovereign CDS spreads. *The Journal of Finance* 63(5), 2345–84.
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review* 47 (2), 527–556.
- Segoviano, M. A. and C. Goodhart (2009). Banking stability measures. *IMF Working Paper*.
- Yue, V. (2010). Sovereign default and debt renegotiation. *Journal of International Economics* 80(2), 176–187.
- 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*.

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

The Generalized Autoregressive Score model of Creal et al. (2011, 2012) 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 Creal et al. (2011) and 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

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

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

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

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

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

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

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 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,

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

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

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

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 $y_{it}/\exp(f_{it}^v/2)$. Assuming the variances are constant ($D_t = \mathbf{I}_n$), 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

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

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 the hypersphere transformation of the unconditional correlation matrix of the transformed data. All remaining 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.