

Conditional Volatility and Correlations of Weekly Returns and the VaR Analysis of 2008 Stock Market Crash*

Bahram Pesaran
Wadhwani Asset Management, LLP

M. Hashem Pesaran
Cambridge University and USC

April 2010

Abstract

Modelling of conditional volatilities and correlations across asset returns is an integral part of portfolio decision making and risk management. Over the past three decades there has been a trend towards increased asset return correlations across markets, a trend which has been accentuated during the recent financial crisis. We shall examine the nature of asset return correlations using weekly returns on futures markets and investigate the extent to which multivariate volatility models proposed in the literature can be used to formally characterize and quantify market risk. In particular, we ask how adequate these models are for modelling market risk at times of financial crisis. In doing so we consider a multivariate t version of the Gaussian dynamic conditional correlation (DCC) model proposed by Engle (2002), and show that the t -DCC model passes the usual diagnostic tests based on probability integral transforms, but fails the value at risk (VaR) based diagnostics when applied to the post 2007 period that includes the recent financial crisis.

JEL Classifications: C51, C52, G11

Key Words: Volatilities and Correlations, Weekly Returns, Multivariate t , Financial Interdependence, VaR diagnostics, 2008 Stock Market Crash.

*This is a substantially revised and updated version of a paper previously distributed under the title "Modelling Volatilities and Conditional Correlations in Futures Markets with a Multivariate t Distribution", 2007, *IZA Discussion Papers*, No. 2906. We are grateful to Enrique Sentana and Elisa Tosetti for useful discussions and comments.

1 Introduction

Modelling of conditional volatilities and correlations across asset returns is an integral part of portfolio decision making and risk management. In risk management the value at risk (VaR) of a given portfolio can be computed using univariate volatility models, but a multivariate model is needed for portfolio decisions. Even in risk management the use of a multivariate model would be desirable when a number of alternative portfolios of the same universe of m assets are under consideration. By using the same multivariate volatility model marginal contributions of different assets towards the overall portfolio risk can be computed in a consistent manner. Multivariate volatility models are also needed for determination of hedge ratios and leverage factors.

The literature on multivariate volatility modelling is large and expanding. Bauwens, Laurent, and Rombouts (2006) provide a recent review. A general class of such models is the multivariate generalized autoregressive conditional heteroscedastic (MGARCH) specification. (Engle and Kroner (1995)). However, the number of unknown parameters of the unrestricted MGARCH model rises exponentially with m and its estimation will not be possible even for a modest number of assets. The diagonal-VEC version of the MGARCH model is more parsimonious, but still contains too many parameters in most applications. To deal with the curse of dimensionality the dynamic conditional correlations (DCC) model is proposed by Engle (2002) which generalizes an earlier specification by Bollerslev (1990) by allowing for time variations in the correlation matrix. This is achieved parsimoniously by separating the specification of the conditional volatilities from that of the conditional correlations. The latter are then modelled in terms of a small number of unknown parameters, which avoids the curse of the dimensionality. With Gaussian standardized innovations Engle (2002) shows that the log-likelihood function of the DCC model can be maximized using a two step procedure. In the first step, m univariate GARCH models are estimated separately. In the second step using standardized residuals, computed from the estimated volatilities from the first stage, the parameters of the conditional correlations are then estimated. The two step procedure can then be iterated if desired for full maximum likelihood estimation.

DCC is an attractive estimation procedure which is reasonably flexible in modeling individual volatilities and can be applied to portfolios with a large number of assets. However, in most applications in finance the Gaussian assumption that underlies the two step procedure is likely to be violated. To capture the fat-tailed nature of the distribution of asset returns, it is more appropriate if the DCC model is combined with a multivariate t distribution, particularly for risk analysis where the tail properties of return distributions are of primary concern. But Engle's two-step procedure will no longer be applicable to such a t -DCC specification and a simultaneous approach to the estimation of the parameters of the model, including the degree-of-freedom parameter of the multivariate t distribution would be needed. This paper develops such an estimation procedure and proposes the use of devolatilized returns computed as returns standardized by realized volatilities rather than by GARCH type volatil-

ity estimates. Devolatilized returns are likely to be approximately Gaussian although the same cannot be said about the standardized returns. (Andersen, Bollerslev, Diebold, and Ebens (2001), and Andersen, Bollerslev, Diebold and Labys (2001)).

The t -DCC estimation procedure is applied to a portfolio composed of 6 currencies, four 10 year government bonds, and seven equity index futures over the period May 27, 1994 to October 30, 2009; split into an estimation sample (1994 to 2007) and an evaluation sample (2008 to 2009). To avoid the non-synchronization of daily returns across markets in different time zones we estimate the volatility models using weekly rather than daily returns.

Main features of the empirical results are as follows:

- The estimation results strongly reject the normal-DCC model in favour of a t -DCC specification.
- The t -DCC specification passes the non-parametric Kolmogorov-Smirnov tests, but fails the VaR test due to the extreme events in September and October of 2008.
- Important changes to asset return volatilities have taken place which are shared across assets and markets.
- The 2008 financial crisis resulted in the reversal of the trend volatilities from its low levels during 2003-2007 to unprecedented heights in 2008.
- Asset return correlations have been rising historically. Recent crisis has accentuated this trend rather than leading to it.
- The rise in asset return correlations seems to be more reflective of underlying trends - globalization and integration of financial markets, and cannot be attributed to the recent financial crisis. More research on this topic is clearly needed.

The plan of the paper is follows. Section 2 introduces the t -DCC model and discusses the devolatilized returns and the rational behind their construction. Section 3 considers recursive relations for real time analysis. The maximum likelihood estimation of the t -DCC model is set out in Section 4, followed by a review of diagnostic tests in Section 5. The empirical application to weekly returns is discussed in Sections 6 and 7. The evolution of asset return volatilities and correlations is discussed in Section 8, followed by some concluding remarks in Section 9.

2 Modelling Conditional Correlation Matrix of Asset Returns

Let \mathbf{r}_t be an $m \times 1$ vector of asset returns at close day t assumed to have a conditional multivariate t distribution with means, $\boldsymbol{\mu}_{t-1}$, and the non-singular

variance-covariance matrix Σ_{t-1} , and $v_{t-1} > 2$ degrees of freedom. Here we are not concerned with how mean returns are predicted and take μ_{t-1} as given.¹ For specification of Σ_{t-1} we follow Bollerslev (1990) and Engle (2002) consider the decomposition

$$\Sigma_{t-1} = \mathbf{D}_{t-1} \mathbf{R}_{t-1} \mathbf{D}_{t-1}, \quad (1)$$

where

$$\mathbf{D}_{t-1} = \begin{pmatrix} \sigma_{1,t-1} & & & & \\ & \sigma_{2,t-1} & & & \\ & & 0 & & \\ & & & \ddots & \\ & & & & \sigma_{m,t-1} \end{pmatrix},$$

$$\mathbf{R}_{t-1} = \begin{pmatrix} 1 & \rho_{12,t-1} & \rho_{13,t-1} & \cdots & \rho_{1m,t-1} \\ \rho_{21,t-1} & 1 & \rho_{23,t-1} & \cdots & \rho_{2m,t-1} \\ \vdots & & \ddots & & \vdots \\ \vdots & & & \ddots & \vdots \\ \rho_{m1,t-1} & \cdots & \cdots & \rho_{m,m-1,t-1} & 1 \end{pmatrix},$$

$\mathbf{R}_{t-1} = (\rho_{ij,t-1}) = (\rho_{ji,t-1})$ is the symmetric $m \times m$ correlation matrix, and \mathbf{D}_{t-1} is the $m \times m$ diagonal matrix with $\sigma_{i,t-1}, i = 1, 2, \dots, m$ denoting the conditional volatility of the i -th asset return. More specifically

$$\sigma_{i,t-1}^2 = V(r_{it} | \Omega_{t-1}),$$

and $\rho_{ij,t-1}$ are conditional pair-wise return correlations defined by

$$\rho_{ij,t-1} = \frac{\text{Cov}(r_{it}, r_{jt} | \Omega_{t-1})}{\sigma_{i,t-1} \sigma_{j,t-1}},$$

where Ω_{t-1} is the information set available at close of day $t-1$. Clearly, $\rho_{ij,t-1} = 1$, for $i = j$.

Bollerslev (1990) considers (1) with a constant correlation matrix $\mathbf{R}_{t-1} = \mathbf{R}$. Engle (2002) allows for \mathbf{R}_{t-1} to be time-varying and proposes a class of multivariate GARCH models labeled as dynamic conditional correlation (DCC) models. An alternative approach would be to use the conditionally heteroskedastic factor model discussed, for example, in Sentana (2000) where the vector of unobserved common factors are assumed to be conditionally heteroskedastic. Parsimony is achieved by assuming that the number of the common factors is much less than the number of assets under considerations.

The decomposition of Σ_{t-1} in (1) allows separate specification of the conditional volatilities and conditional cross-asset returns correlations. For example, one can utilize the GARCH (1,1) model for $\sigma_{i,t-1}^2$, namely

$$V(r_{it} | \Omega_{t-1}) = \sigma_{i,t-1}^2 = \bar{\sigma}_i^2 (1 - \lambda_{1i} - \lambda_{2i}) + \lambda_{1i} \sigma_{i,t-2}^2 + \lambda_{2i} r_{i,t-1}^2, \quad (2)$$

¹ Although, the estimation of μ_{t-1} and Σ_{t-1} are inter-related, in practice mean returns are predicted by least squares techniques (such as recursive estimation or recursive modelling) which do not take account of the conditional volatility. This might involve some loss in efficiency of estimating μ_{t-1} , but considerably simplifies the estimation of the return distribution needed in portfolio decisions and risk management.

where $\bar{\sigma}_i^2$ is the unconditional variance of the i -th asset return. Under the restriction $\lambda_{1i} + \lambda_{2i} = 1$, the unconditional variance does not exist and we have the integrated GARCH (IGARCH) model used extensively in the professional financial community, which is mathematically equivalent to the “exponential smoother” applied to the r_{it}^2 's²

$$\sigma_{i,t-1}^2(\lambda_i) = (1 - \lambda_i) \sum_{s=1}^{\infty} \lambda_i^{s-1} r_{i,t-s}^2 \quad 0 < \lambda_i < 1, \quad (3)$$

or written recursively

$$\sigma_{i,t-1}^2(\lambda_i) = \lambda_i \sigma_{i,t-2}^2 + (1 - \lambda_i) r_{i,t-1}^2. \quad (4)$$

For cross-asset correlations Engle proposes the use of the following exponential smoother applied to the “standardized returns”

$$\hat{\rho}_{ij,t-1}(\phi) = \frac{\sum_{s=1}^{\infty} \phi^{s-1} z_{i,t-s} z_{j,t-s}}{\sqrt{\sum_{s=1}^{\infty} \phi^{s-1} z_{i,t-s}^2} \sqrt{\sum_{s=1}^{\infty} \phi^{s-1} z_{j,t-s}^2}}, \quad (5)$$

where the standardized returns are defined by

$$z_{it} = \frac{r_{it}}{\sigma_{i,t-1}(\lambda_i)}. \quad (6)$$

For estimation of the unknown parameters, $\lambda_1, \lambda_2, \dots, \lambda_m$, and ϕ , Engle (2002) proposes a two-step procedure whereby in the first step individual GARCH(1,1) models are fitted to the m asset returns separately, and then the coefficient of the conditional correlations, ϕ , is estimated by the Maximum Likelihood method assuming that asset returns are conditionally Gaussian. This procedure has two main drawbacks. First, the Gaussianity assumption does not hold for daily returns and its use can under-estimate the portfolio risk. Second, the two-stage approach is likely to be inefficient even under Gaussianity.

2.1 Pair-wise correlations based on realized volatilities

In this paper we consider an alternative formulation of $\rho_{ij,t-1}$ that makes use of realized volatilities, or their approximations based on daily or weekly observations when realized volatility measures are not available. In a series of papers Andersen, Bollerslev and Diebold show that daily returns on foreign exchange and stock returns standardized by realized volatility are approximately Gaussian. See, for example, Andersen, Bollerslev, Diebold, and Ebens (2001), and Andersen, Bollerslev, Diebold and Labys (2001). The transformation of returns to Gaussianity is important since as recently shown by Embrechts et al. (2003), the use of correlation as a measure of dependence can be misleading

²See, for example, Litterman and Winkelmann (1998).

in the case of (conditionally) non-Gaussian returns. In contrast, estimation of correlations based on devolatilized returns that are nearly Gaussian is likely to be more generally meaningful. Denote the realized volatility of i^{th} return in day (week) t by $\sigma_{it}^{realized}$ and standardize the returns by the realized volatilities to obtain

$$\tilde{r}_{it} = \frac{r_{it}}{\sigma_{it}^{realized}}. \quad (7)$$

To avoid confusions we refer to \tilde{r}_{it} as the “devolatilized returns”, and refer to z_{it} defined by (6) as the standardized returns. The conditional pair-wise return correlations based on \tilde{r}_{it} are now given by

$$\tilde{\rho}_{ij,t-1}(\phi) = \frac{\sum_{s=1}^{\infty} \phi^{s-1} \tilde{r}_{i,t-s} \tilde{r}_{j,t-s}}{\sqrt{\sum_{s=1}^{\infty} \phi^{s-1} \tilde{r}_{i,t-s}^2} \sqrt{\sum_{s=1}^{\infty} \phi^{s-1} \tilde{r}_{j,t-s}^2}}, \quad (8)$$

where $-1 < \tilde{\rho}_{ij,t-1}(\phi) < 1$ for all values of $|\phi| < 1$.

As compared to z_{it} , the use of \tilde{r}_{it} is more data intensive and requires intradaily observations. Although, intradaily observations are becoming increasingly available across a large number of assets, it would still be desirable to work with a version of \tilde{r}_{it} that does not require intradaily observations, but is nevertheless capable of rendering the devolatilized returns approximately Gaussian. One of the main reasons for the non-Gaussian behavior daily returns is presence of jumps in the return process as documented for a number of markets in the literature (see, for example, Barndorff-Nielsen and Shephard (2002)). The standardized return, z_{it} , used by Engle does not deal with such jumps, since the jump process that affects the numerator of z_{it} in day t does not enter the denominator of z_{it} which is based on past returns and exclude the current return, r_t . The problem is accentuated due to the facts that jumps are typically independently distributed over time. The use of realized volatility ensures that the numerator and the denominator of the devolatilized returns, \tilde{r}_{it} , are both affected by the same jumps in day t .

In the absence of intradaily observations the following simple estimate of σ_{it} based on daily or weekly returns, *inclusive* of the contemporaneous value of r_{it} , seem to work well in practice

$$\tilde{\sigma}_{it}^2(p) = \frac{\sum_{s=0}^{p-1} r_{i,t-s}^2}{p}. \quad (9)$$

The lag-order, p , needs to be chosen carefully. We have found that for weekly returns a value of $p = 13$ tends to render the devolatilized returns, $\tilde{r}_{it} \approx r_{it}/\tilde{\sigma}_{it}(p)$, nearly Gaussian, with approximately unit variances, for all asset classes foreign exchange, equities, bonds or commodities.³ Note that $\tilde{\sigma}_{it}^2(p)$ is not the same of the rolling historical estimate of σ_{it} defined by

$$\hat{\sigma}_{it}^2(p) = \frac{\sum_{s=1}^p r_{i,t-s}^2}{p}.$$

³Also see Section 6. For daily observations $p = 20$ seems to work well. See the earlier version of this paper, Pesaran and Pesaran (2007), for further details.

Specifically

$$\tilde{\sigma}_{it}^2(p) - \hat{\sigma}_{it}^2(p) = \frac{r_{it}^2 - r_{i,t-p}^2}{p}.$$

It is the inclusion of the current squared returns, r_{it}^2 , in the estimation of $\tilde{\sigma}_{it}^2$ that seems to be critical in transformation of r_{it} (which is non-Gaussian) into \tilde{r}_{it} which seems to be approximately Gaussian.

3 Real Time Risk Analysis and Updates

In financial analysis estimation and evaluation are in general recursive and the unknown parameters need to be updated over time.⁴ The frequency by which parameters are updated depends on the processing costs and the expected benefit from the updates. When processing costs are negligible parameter updates are carried out on the arrival of new data or shortly thereafter. For daily or weekly observations (the focus of the present paper) monthly or quarterly updates are recommended. Daily or weekly updates can be quite time consuming for large portfolios, and the expected benefit of the more frequent updates unclear. For model evaluation, however, a weekly frequency seems desirable. Clearly, model evaluation need not be carried out at the same frequency with which parameters are updated.

The implementation of the real time analysis is very much facilitated using recursive formulae in the estimation and the evaluation process. For computational of $\rho_{ij,t-1}$, given by (5) and (8), as noted by Engle (2002) we have

$$\tilde{\rho}_{ij,t-1}(\phi) = \frac{q_{ij,t-1}}{\sqrt{q_{ii,t-1}q_{jj,t-1}}} \quad (10)$$

where

$$q_{ij,t-1} = \phi q_{ij,t-2} + (1 - \phi) \tilde{r}_{i,t-1} \tilde{r}_{j,t-1}. \quad (11)$$

The recursive expression for $\hat{\rho}_{ij,t-1}(\phi)$ is identical except that instead of devolatilized returns the standardized returns, z_{it} , given by (6) are used.

The above models for $\rho_{ij,t-1}$ are non-mean reverting. A more general mean-reverting specification is given by

$$q_{ij,t-1} = \bar{\rho}_{ij}(1 - \phi_1 - \phi_2) + \phi_1 q_{ij,t-2} + \phi_2 \tilde{r}_{i,t-1} \tilde{r}_{j,t-1}, \quad (12)$$

where $\bar{\rho}_{ij}$ is the unconditional correlation of r_{it} and r_{jt} and $\phi_1 + \phi_2 < 1$. One would expect $\phi_1 + \phi_2$ to be close to unity. The non-mean reverting case can be obtained as a special case by setting $\phi_1 + \phi_2 = 1$. In practice it is impossible to be sure if $\phi_1 + \phi_2 < 1$ or not. The unconditional correlations, $\bar{\rho}_{ij}$, can be estimated using an expanding window. In the empirical applications we shall consider the mean reverting as well as the non-mean reverting specifications, and experiment with the two specifications of the conditional correlations that are based on standardized and devolatilized returns.

⁴A general discussion of real time econometric analysis is provided in Pesaran and Timmermann (2005).

3.1 Initialization, Estimation and Evaluation Samples

Suppose daily or weekly observations are available on m returns in the $m \times 1$ vector \mathbf{r}_t over the period $t = 1, 2, \dots, T, T+1, \dots, T+N$. The first T_0 observations are used for computation of (9), the initialization of the recursions (12), and the estimation of sample variances and correlations, namely $\bar{\sigma}_i^2$ and $\bar{\rho}_{ij}$, used in (2) and (12), respectively. Let s denote the starting point of the most recent sample of observations to be used in estimation. Clearly, we must have $T > s > T_0 > p$. The size of the estimation window will then be given by $T_e = T - s + 1$. The remaining N observations can then be used for evaluation purposes. More specifically, the initialization sample will be given by $\mathcal{S}_0 = \{\mathbf{r}_t, t = 1, 2, \dots, T_0\}$, the estimation sample by $\mathcal{S}_e = \{\mathbf{r}_t, t = s, s+1, \dots, T\}$, and the evaluation sample, $\mathcal{S}_{\text{eval}} = \{\mathbf{r}_t, t = T+1, T+2, \dots, T+N\}$. This decomposition allows us to vary the size of the estimation window ($T_e = T - s + 1$) by moving the index s along the time axis in order to accommodate estimation of the unknown parameters using expanding or rolling observation windows, with different estimation update frequencies. For example, for an expanding estimation window we set $s = T_0 + 1$. For a rolling window of size W we need to set $s = T + 1 - W$. The whole estimation process can then be rolled into the future with an update frequency of h by carrying the estimations at $T + h, T + 2h, \dots$, using either expanding or rolling estimation samples from $t = s$. Note that model (risk) evaluation can be carried out using observations $t = T+1, T+2, \dots$, irrespective of the update frequency parameter h .

3.2 Mean Reverting Conditional Correlations

In the mean reverting case we also need the estimates of the unconditional volatilities and the correlation coefficients. These can be estimated by

$$\bar{\sigma}_{i,t}^2 = \frac{\sum_{\tau=1}^t r_{i\tau}^2}{t}, \quad (13)$$

$$\bar{\rho}_{ij,t} = \frac{\sum_{\tau=1}^t r_{i\tau} r_{j\tau}}{\sqrt{\sum_{\tau=1}^t r_{i\tau}^2} \sqrt{\sum_{\tau=1}^t r_{j\tau}^2}}. \quad (14)$$

The index t refers to the end of the available estimation sample which in real time will be recursively rolling or expanding, namely $t = T, T + h, T + 2h, \dots$

4 Maximum Likelihood Estimation of the t -DCC Model

In its most general formulation (the non-mean reverting specifications given by (2) and (12)) the DCC(1,1) model contains $2m+3$ unknown parameters; $2m$ coefficients $\boldsymbol{\lambda}_1 = (\lambda_{11}, \lambda_{12}, \dots, \lambda_{1m})'$ and $\boldsymbol{\lambda}_2 = (\lambda_{21}, \lambda_{22}, \dots, \lambda_{2m})'$ that enter the individual asset returns volatilities, the 2 coefficients ϕ_1 and ϕ_2 that enter the

conditional correlations, and the degrees of freedom of the multivariate t distribution, v . The parameters $\bar{\sigma}_i^2$ and $\bar{\rho}_{ij}$ in (2) and (12) refer to the unconditional volatilities and return correlations and can be estimated using the estimation sample or the estimation plus initialization sample. See (13) and (14). In the non-mean reverting case these intercept coefficients disappear, but for the initialization of the recursive relations (2) and (12) it is still advisable to use unconditional estimates of the correlation matrix and asset returns volatilities.

Denote the unknown coefficients by

$$\boldsymbol{\theta} = (\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \phi_1, \phi_2, v)'.$$

Then based on a sample of observations on returns, $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_t$, available at time t , the time t log-likelihood function based on the decomposition (1) is given by

$$l_t(\boldsymbol{\theta}) = \sum_{\tau=s}^t f_{\tau}(\boldsymbol{\theta}), \quad (15)$$

where $s < t$ is the start date of the estimation window (see above). Under t -DCC specification $f_{\tau}(\boldsymbol{\theta})$ refers to the density of the multivariate distribution with v degrees of freedom which can be written in terms of the $\boldsymbol{\Sigma}_{t-1} = \mathbf{D}_{t-1} \mathbf{R}_{t-1} \mathbf{D}_{t-1}$ as⁵

$$\begin{aligned} f_{\tau}(\boldsymbol{\theta}) = & -\frac{m}{2} \ln(\pi) - \frac{1}{2} \ln |R_{\tau-1}(\boldsymbol{\theta})| - \ln |D_{\tau-1}(\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2)| \\ & + \ln \left[\Gamma\left(\frac{m+v}{2}\right) / \Gamma\left(\frac{v}{2}\right) \right] - \frac{m}{2} \ln(v-2) \\ & - \left(\frac{m+v}{2}\right) \ln \left[1 + \frac{\mathbf{e}_{\tau}' D_{\tau-1}^{-1}(\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2) R_{\tau-1}^{-1}(\boldsymbol{\theta}) D_{\tau-1}^{-1}(\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2) \mathbf{e}_{\tau}}{v-2} \right], \end{aligned} \quad (16)$$

where

$$\mathbf{e}_{\tau} = \mathbf{r}_{\tau} - \boldsymbol{\mu}_{\tau-1},$$

and

$$\ln |D_{\tau-1}(\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2)| = \sum_{i=1}^m \ln [\sigma_{i,\tau-1}(\lambda_{1i}, \lambda_{2i})]. \quad (17)$$

It is worth noting that under Engle's specification \mathbf{R}_{t-1} depends on $\boldsymbol{\lambda}_1$ and $\boldsymbol{\lambda}_2$ as well as on ϕ_1 and ϕ_2 . Under the alternative specification advanced here (based on devolatilized returns) \mathbf{R}_{t-1} does not depend on $\boldsymbol{\lambda}_1$ and $\boldsymbol{\lambda}_2$, but depends on ϕ_1 and ϕ_2 , and p , the lag order used in the devolatilization process.

The ML estimate of $\boldsymbol{\theta}$ based on the sample observations, $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_T$, can now be computed by maximization of $l_t(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$, which we denote by $\hat{\boldsymbol{\theta}}_t$. More specifically

$$\hat{\boldsymbol{\theta}}_t = \underset{\boldsymbol{\theta}}{\text{Arg max}} \{l_t(\boldsymbol{\theta})\}, \text{ for } t = T, T+h, T+2h, \dots, T+N, \quad (18)$$

⁵Typically the multivariate t density is written in terms of a scale matrix. But assuming $v > 2$ ensures that $\boldsymbol{\Sigma}_{t-1}$ exists and therefore the scale matrix of the multivariate t distribution can be written in terms of $\boldsymbol{\Sigma}_{t-1}$, which is more convenient for the analysis of multivariate volatility models. See, for example, Bauwens and Laurent (2005).

where h is the (estimation) update frequency, and as before N refers to the length of the evaluation sample. The standard errors of the ML estimates are computed using the asymptotic formulae⁶

$$\widehat{Cov}(\hat{\theta}_t) = \left\{ \sum_{\tau=s}^t \left[\frac{-\partial^2 f_{\tau}(\theta)}{\partial \theta \partial \theta'} \right]_{\theta=\hat{\theta}_t} \right\}^{-1}.$$

In comparison with general specifications of multivariate GARCH model, the model set out in this paper is quite parsimonious. The number of unknown coefficients of the general MGARCH model rises as a quadratic function of m , while the parameters of the DCC model rises linearly with m . Nevertheless, in practice the simultaneous estimation of all the parameters of the DCC model could be problematic, namely can encounter convergence problems, or could lead to a local maxima of the likelihood function. When the returns are conditionally Gaussian one could simplify (at the expense of some loss of estimation efficiency) the computations by adopting Engle's two-stage estimation procedure. But for our preferred distributional assumption the use of such a two-stage procedure does not seem possible and can lead to contradictions. For example, estimation of separate $t - GARCH(1,1)$ models for individual asset returns can lead to different estimates of v , while the multi-variate t distribution requires v to be the same across all assets.⁷

5 Simple Diagnostic Tests of the t -DCC Model

Consider a portfolio based on the m assets with the return vector \mathbf{r}_t using the $m \times 1$ vector of pre-determined weights, \mathbf{w}_{t-1} . The return on this portfolio is given by

$$\rho_t = \mathbf{w}_{t-1}' \mathbf{r}_t. \quad (19)$$

Suppose that we are interested in computing the capital Value at Risk (VaR) of this portfolio expected at the close of business on day $t - 1$ with probability $1 - \alpha$, which we denote by $VaR(\mathbf{w}_{t-1}, \alpha)$. For this purpose we require that

$$\Pr [\mathbf{w}_{t-1}' \mathbf{r}_t < -VaR(\mathbf{w}_{t-1}, \alpha) | \Omega_{t-1}] \leq \alpha.$$

Under our assumptions, conditional on Ω_{t-1} , $\mathbf{w}_{t-1}' \mathbf{r}_t$ has a Student t distribution with mean $\mathbf{w}_{t-1}' \boldsymbol{\mu}_{t-1}$, the variance $\mathbf{w}_{t-1}' \boldsymbol{\Sigma}_{t-1} \mathbf{w}_{t-1}$, and the degrees of freedom v . Hence

$$z_t = \sqrt{\frac{v}{v-2}} \left(\frac{\mathbf{w}_{t-1}' \mathbf{r}_t - \mathbf{w}_{t-1}' \boldsymbol{\mu}_{t-1}}{\sqrt{\mathbf{w}_{t-1}' \boldsymbol{\Sigma}_{t-1} \mathbf{w}_{t-1}}} \right),$$

⁶ An analytical expression for the information matrix for the multivariate t -GARCH model is provided by Florentini, Sentana, and Calzolari (2003). But in the applications considered in this paper we did not encounter any problems using numerical derivatives to compute the information matrix.

⁷ Marginal distributions associated with a multi-variate t -distribution with v degrees of freedom are also t -distributed with the same degrees of freedom.

conditional on Ω_{t-1} will also have a t distribution with v degrees of freedom. It is easily verified that $E(z_t|\Omega_{t-1}) = 0$, and $V(z_t|\Omega_{t-1}) = v/(v-2)$. Denoting the cumulative distribution function of a Student t with v degrees of freedom by $F_v(z)$, $VaR(\mathbf{w}_{t-1}, \alpha)$ will be given as the solution to

$$F_v \left(\frac{-VaR(\mathbf{w}_{t-1}, \alpha) - \mathbf{w}'_{t-1} \boldsymbol{\mu}_{t-1}}{\sqrt{\frac{v-2}{v}} (\mathbf{w}'_{t-1} \boldsymbol{\Sigma}_{t-1} \mathbf{w}_{t-1})} \right) \leq \alpha.$$

But since $F_v(z)$ is a continuous and monotonic function of z we have

$$\frac{-VaR(\mathbf{w}_{t-1}, \alpha) - \mathbf{w}'_{t-1} \boldsymbol{\mu}_{t-1}}{\sqrt{\frac{v-2}{v}} (\mathbf{w}'_{t-1} \boldsymbol{\Sigma}_{t-1} \mathbf{w}_{t-1})} = F_v^{-1}(\alpha) = -c_\alpha,$$

where c_α is the $\alpha\%$ critical value of a Student t distribution with v degrees of freedom. Therefore,

$$VaR(\mathbf{w}_{t-1}, \alpha) = \tilde{c}_\alpha \sqrt{(\mathbf{w}'_{t-1} \boldsymbol{\Sigma}_{t-1} \mathbf{w}_{t-1})} - \mathbf{w}'_{t-1} \boldsymbol{\mu}_{t-1}, \quad (20)$$

where $\tilde{c}_\alpha = c_\alpha \sqrt{\frac{v-2}{v}}$.

Following Christoffersen (1998) and Engle and Manganelli (2004), a simple test of the validity of t -DCC model can be computed recursively using the VaR indicators

$$d_t = I(\mathbf{w}'_{t-1} \mathbf{r}_t + VaR(\mathbf{w}_{t-1}, \alpha)) \quad (21)$$

where $I(A)$ is an indicator function which is equal to unity if $A > 0$ and zero otherwise. These indicator statistics can be computed in-sample or preferably can be based on recursive out-of-sample one-step ahead forecast of $\boldsymbol{\Sigma}_{t-1}$ and $\boldsymbol{\mu}_{t-1}$, for a given (pre-determined set of portfolio weights, \mathbf{w}_{t-1}). In such an out-of-sample exercise the parameters of the mean returns and the volatility variables ($\boldsymbol{\beta}$ and $\boldsymbol{\theta}$, respectively) could be either kept fixed at the start of the evaluation sample or changed with an update frequency of h periods (for example with $h = 5$ for weekly updates, or $h = 20$ for monthly updates). For the evaluation sample, $\mathcal{S}_{\text{eval}} = \{\mathbf{r}_t, t = T+1, T+2, \dots, T+N\}$, the mean hit rate is given by

$$\hat{\pi}_N = \frac{1}{N} \sum_{t=T+1}^{T+N} d_t. \quad (22)$$

Under the t -DCC specification, $\hat{\pi}_N$ will have mean $1 - \alpha$ and variance $\alpha(1 - \alpha)/N$. The standardized statistic,

$$z_\pi = \frac{\sqrt{N} [\hat{\pi}_N - (1 - \alpha)]}{\sqrt{\alpha(1 - \alpha)}}, \quad (23)$$

will have a standard normal distribution for a sufficiently large evaluation sample size, N . This result holds irrespective of whether the unknown parameters are

estimated recursively or fixed at the start of the evaluation sample. In the case of the latter the validity of the test procedure requires that $N/T \rightarrow 0$ as $(N, T) \rightarrow \infty$. For a proof see Pesaran, Schleicher and Zaffaroni (2009).

The z_π statistic provides evidence on the performance of Σ_{t-1} and μ_{t-1} in an average (unconditional) sense. (Lopez (1999)). An alternative conditional evaluation procedure, proposed by Berkowitz (2001), can be based on probability integral transforms⁸

$$\hat{U}_t = F_v \left(\frac{\mathbf{w}'_{t-1} \mathbf{r}_t - \mathbf{w}'_{t-1} \hat{\mu}_{t-1}}{\sqrt{\frac{v-2}{v} \mathbf{w}'_{t-1} \hat{\Sigma}_{t-1} \mathbf{w}_{t-1}}} \right), \quad t = T+1, T+2, \dots, T+N. \quad (24)$$

Under the null hypothesis of correct specification of the t -DCC model, the probability transform estimates, \hat{U}_t , are serially uncorrelated and uniformly distributed over the range $(0, 1)$. Both of these properties can be readily tested. The serial correlation property of \hat{U}_t can be tested by Lagrange multiplier tests using OLS regressions of \hat{U}_t on an intercept and the lagged values $\hat{U}_{t-1}, \hat{U}_{t-2}, \dots, \hat{U}_{t-s}$. The maximum lag length, s , can be selected by the application of the AIC criteria, for example. The uniformity of the distribution of \hat{U}_t over t can be tested using the Kolmogorov-Smirnov statistic defined by, $KS_N = \sup_x |F_{\hat{U}}(x) - U(x)|$, where $F_{\hat{U}}(x)$ is the empirical cumulative distribution function (CDF) of the \hat{U}_t , for $t = T+1, T+2, \dots, T+N$, and $U(x) = x$ is the CDF of iid $U[0, 1]$. Large values of the Kolmogorov-Smirnov statistic, KS_N , indicate that the sample CDF is not similar to the hypothesized uniform CDF.⁹

6 Volatilities and Conditional Correlations in Weekly Returns

We estimated alternative versions of the t -DCC model for a portfolio composed of weekly returns on

- 6 currencies: British pound (GBP), euro (EU), Japanese yen (JPY), Swiss franc (CHF), Canadian dollar (CAD), and Australian dollar (AD).
- 4 government bonds: US T-Note 10Y (BU), Europe Euro Bund 10Y (BE), Japan Government Bond 10Y (BJ), and, UK Long Gilts 8.75-13Y (BG).
- 7 equity index futures S&P 500 (SP), FTSE 100 (FTSE), German DAX (DAX), French CAC40 (CAC), Swiss Market Index (SM), Australia SPI200 (AUS), Nikkei 225 (NK).

The weekly returns are computed from daily prices obtained from Datastream and cover the period from 07-Jan-94 to 30-Oct-2009.

⁸See also Christoffersen (1998) for a related test that applied to the VaR indicators, d_t , defined by (21).

⁹For details of the Kolmogorov-Smirnov test and its critical values see, for example, Massey (1951), and Neave and Worthington (1992, pp.89-93).

6.1 Devolitized returns and their properties

Table 1 provides summary statistics for the weekly returns (r_{it} , in percent) and the devolitized weekly returns $\tilde{r}_{it} = r_{it}/\tilde{\sigma}_{it}(p)$, where in the absence of intradaily observations $\tilde{\sigma}_{it}^2(p)$ is defined by (9), with $p = 13$ weeks. The choice of $p = 13$ was guided by some experimentation with the aim of transforming r_{it} into an approximately Gaussian process. A choice of p well above 13 does now allow the (possible) jumps in r_{it} to become adequately reflected in $\tilde{\sigma}_{it}(p)$, and a value of p well below 13 transforms r_{it} to an indicator looking function. In the extreme case where $p = 1$ we have $\tilde{r}_{it} = 1$, if $r_{it} > 0$, and $\tilde{r}_{it} = -1$, if $r_{it} < 0$, and $\tilde{r}_{it} = 0$, if $r_{it} = 0$. We did not experiment with other values of p for the sample under consideration and set $p = 13$ for all the 17 assets.

For the non-devolitized returns the results are as to be expected from previous studies. The returns seem to be symmetrically distributed with kurtosis in some cases well in excess of 3 (the value for the Gaussian distribution). The excess kurtosis is particularly large for equities, mostly around 5 or more. For currencies the kurtosis coefficient is particularly large for yen, British pound, and Singapore dollar. In comparison the weekly returns on government bonds are less fat-tailed with kurtosis coefficients only marginally above 3. In contrast, none of the 17 devolitized returns show any evidence of excess kurtosis. For example, for equities the excess kurtosis of weekly returns on SP, FTSE and Nikkie fall from 8.01, 10.40, 9.65 to -0.124, -0.0132 and -0.147, respectively after the returns are devolitized. For currencies, the excess kurtosis of the weekly returns on AD, BP, and JY fall from 7.89, 5.35, and 9.46 to -0.112, -0.020, and 0.139, respectively. Out of the four ten year government bonds only the weekly returns on Japanese government bond show some degree of excess kurtosis which is eliminated once the returns are devolitized. It is also interesting to note that the standard deviations of the devolitized returns are now all very close to unity, that allows a more direct comparison of the devolitized returns across assets.

The extent to which the devolatilization has been effective in transforming the returns into Gaussian variates can be seen in Figures 1-17. The top panel of each figure gives the histograms, a kernel density fitted to the returns together with the normal density and the normal QQ-plots. These plots graphically compare the distribution of returns to the normal distribution (represented by a straight line in the case of the QQ-plots). The figures on the bottom panel display the same graphs for the devolitized returns. These figures clearly show that devolatilization has been quite effective in achieving Gaussianity to a high degree of approximation. This can be seen particularly if one compares QQ-plots of returns and their devolitized counterparts. For the devolitized returns the QQ-plots generally lie on the straight-line with a few exceptions. But for the raw returns there are important departures from normality, particularly in tails of the return distributions.

7 ML estimates of the DCC models

It is well established that daily or weekly returns are approximately mean zero serially uncorrelated processes and for the purpose of risk analysis it is reasonable to assume that $\boldsymbol{\mu}_{t-1} = \mathbf{0}$. Using the ML procedure described above, initially we estimate a number of DCC models on the 17 weekly returns over the period 27-May-94 to 28-Dec-2007 (710 observations). We then use the post estimation sample observations from January 4, 2008 to October 30, 2009 for the evaluation of the estimated volatility models using the VaR and distribution free diagnostics.¹⁰ We also provide *separate* t -DCC models for currencies, bonds and equities for purposes of comparisons.

We begin with the unrestricted version of the DCC(1,1) model with asset-specific volatility parameters $\boldsymbol{\lambda}_1 = (\lambda_{11}, \lambda_{12}, \dots, \lambda_{1m})'$, $\boldsymbol{\lambda}_2 = (\lambda_{21}, \lambda_{22}, \dots, \lambda_{2m})'$, and common conditional correlation parameters, ϕ_1 and ϕ_2 , and the degrees-of-freedom parameter, v , under conditionally t distributed returns. (note that $m = 17$). We did not encounter any convergence problems, and obtained the same ML estimates when starting from different initial parameter values. But to achieve convergence in some applications we had to experiment with different initial values. In particular we found the initial values $\lambda_{1i} = 0.95$, $\lambda_{2i} = 0.05$, $\phi_1 = 0.96$, $\phi_2 = 0.03$ and $v = 12$ to work relatively well. Also the sum of unrestricted estimates of λ_1 and λ_2 for the Canadian dollar exceeded 1, and to ensure a non-explosive outcome we estimated its volatility equation subject to the restriction $\lambda_{1,CD} + \lambda_{2,CD} = 1$.

To evaluate the statistical significance of the multivariate t distribution for the analysis of return volatilities, in Table 2 we first provide the maximized log-likelihood values under multivariate normal and t distributions for currencies, bonds and equities separately, as well as for all the 17 assets jointly. We report these results both for standardized and devolatilized returns. It is firstly clear from these results that the normal-DCC specifications are strongly rejected relative to the t -DCC models for all asset categories. The maximized log-likelihood values for the t -DCC models are significantly larger than the ones for the normal-DCC models. The estimated degrees of freedom of the multivariate t -distribution for different asset classes are quite close and range from 8 (for equities) to 11 (for bonds), all well below the values of 30 and above that one would expect for a multivariate normal distribution. For the full set of 17 assets the estimate of v is closer to 12. There seems to be a tendency for the estimate of v to rise as more assets are included in the t -DCC model.

The above conclusions are robust to the way returns are scaled for computation of cross asset return correlations. The maximized log-likelihoods for the standardized and devolatilized returns are very close, although due to the non-nested nature of the two return transformations no definite conclusions can be reached as to their relative merits. The specifications where the returns are standardized by the conditional volatilities tend to fit better (give higher log-likelihood values). But this is to be expected since the maximization of the

¹⁰The ML estimation and the computation of the diagnostic statistics are carried out using Microfit 5. See Pesaran and Pesaran (2009).

log-likelihood function in this case is carried out with respect to the parameters of the scaling factor, unlike the case where scaling is carried out with respect to the realized volatilities which do not depend on the unknown parameters of the likelihood function. In what follows we base our correlation analysis on the devolatilized returns on the grounds of their approximate Gaussianity, as argued above.

7.1 Asset specific estimates

Table 3 presents the ML estimates of the t -DCC model including all the 17 assets computed over the period 27-May-94 to 28-Dec-07 (710 weekly returns). The asset-specific estimates of the volatility decay parameters are all highly significant, with the estimates of λ_{1i} , $i = 1, 2, \dots, 17$ falling in the range of 0.818 (for Japanese government bond) to 0.986 (for Canadian dollar).¹¹ The average estimate of λ_1 across assets is 0.924 which is somewhat smaller than the values in the range of 0.95 to 0.97 recommended by Riskmetrics for computation of *daily* volatilities using their exponential smoothing procedure. This is not surprising, since one would expect the exponential smoothing parameter for computing the volatility of weekly returns to be smaller than the one used for computing the volatility of daily returns.

There are, however, notable differences across asset groups with λ_{i1} estimated to be larger for currencies as compared to the estimates for equities and bonds. The average estimate of λ_1 across currencies is 0.95 as compared to 0.93 for equities and 0.88 for bonds. The correlation parameters, ϕ_1 and ϕ_2 are very precisely estimated and $\hat{\phi}_1 + \hat{\phi}_2 = 0.9846(0.0028)$, and suggest very slow but statistically significant mean reverting conditional correlations.

The sum of the estimates of λ_{1i} and λ_{2i} are very close to unity, but the hypothesis that $\lambda_{1i} + \lambda_{2i} = 1$ (the integrated GARCH hypothesis) against the one-sided alternative $\lambda_{1i} + \lambda_{2i} < 1$ is rejected for 10 out of the 17 assets at the 5% significance level; the exceptions being British pound, Swiss franc, Nikkei, S&P 500, and Australian SPI200. For Canadian dollar to ensure a non-explosive outcome, as noted earlier estimation is carried out subject to the restriction $\lambda_{1,CD} + \lambda_{2,CD} = 1$. If the test is carried out at the 1% significance level, the integrated GARCH hypothesis is rejected only in the case of the JGB (Japanese Government Bond).

The integrated GARCH (IGARCH) hypothesis is implicit in the approach advocated by Riskmetrics, but as shown by Zaffaroni (2008) can lead to inconsistent estimates. However, in the present applications the unrestricted parameter estimates and those obtained under IGARCH are very close and one can view the restrictions $\lambda_{1i} + \lambda_{2i} = 1$ as a first order approximation that avoids explosive outcomes. We also note that the diagnostic test results, to be reported below, are not qualitatively affected by the imposition of the restrictions, $\lambda_{1i} + \lambda_{2i} = 1$.

Finally, it is worth noting that there is statistically significant evidence of

¹¹ Recall that for Canadian dollar the volatility model is estimated subject to the restriction $\lambda_{1,CD} + \lambda_{2,CD} = 1$.

parameter heterogeneity across assets, which could lead to misleading inference if these differences are ignored.

7.2 Post estimation evaluation of the t -DCC model

The evaluation sample, 04-Jan-08 to 30-Oct-09, covers the recent periods of financial crisis and include 96 weeks of post estimation sample of portfolio returns. The parameter values are estimated using the sample 27-May-04 to 28-Dec-07 and then fixed throughout the evaluation sample. To evaluate the t -DCC model we first consider the tests based on probability integral transforms (PIT), \hat{U}_t , defined by (24). We have already seen that under the null hypothesis that the t -DCC model is correctly specified, \hat{U}_t , are serially uncorrelated and uniformly distributed over the range $(0, 1)$. To compute \hat{U}_t we consider an equal-weighted portfolio, with all elements of \mathbf{w} in (19) set to $1/17$, and use the risk tolerance probability of $\alpha = 1\%$, which is the value typically assumed in practice. We considered two versions of the t -DCC model: a version with no restrictions on λ_{1i} and λ_{2i} (except for $i = CD$), and an integrated version where $\lambda_{1i} + \lambda_{2i} = 1$, for all i .

Using the Lagrange multiplier statistic to test the null hypothesis that \hat{U}_t 's are serially uncorrelated we obtained the values of $\chi^2_{12} = 4.74$ and $\chi^2_{12} = 5.31$ for the unrestricted and the restricted t -DCC specifications. These statistics are computed assuming a maximum lag order of 12, and are asymptotically distributed as chi-squared variates with 12 degrees of freedom. It is clear that both specifications of the t -DCC model pass this test.

Next we applied the Kolmogorov-Smirnov statistic to \hat{U}_t 's to test the null hypothesis that the PIT values are draws from a uniform distribution. The KS statistics for the unrestricted and the restricted versions amounted to 0.0646 and 0.0454, respectively. Both these statistics are well below the KS critical value of 0.1388 (at the 5% level).¹² Therefore, the null hypothesis that the sample CDF of \hat{U}_t 's is similar to the hypothesized uniform CDF cannot be rejected. Figure 18 provides a graphical display of the KS test for the unrestricted version of the model, and confirms the close match between the empirical CDF of the \hat{U}_t 's and the CDF of the uniform distribution (45 degree line)

It is interesting that neither of the tests based on \hat{U}_t 's are capable of detecting the effects of the financial turmoils that took place in 2008. A test based on the violations of the VaR constraint is likely to be more discriminating, since it focusses on the tail properties of the return distributions. For a tolerance probability of $\alpha = 0.01$, we would expect only one violation of the VaR constraint in 100 observations (our evaluation sample contains 96 observations). The unrestricted specification results in three violations of the VaR constraint, and the restricted specification in four violations. Both specifications violate the VaR constrain in the weeks starting on 5-Sep-08, 3-Oct-08 and 10-Oct-08. The restricted version also violates the VaR in the week starting in 18-Jan-08. The test statistics associated with these violations are -2.09 and -3.12 which are

¹²See Table 1 in Massey (1951).

normally distributed. Thus both specifications are rejected by the VaR violation test.¹³ Not surprisingly, the rejection of the test is due to the unprecedented market volatility during the weeks in September and October of 2008. This period covers the Federal take over of Fannie Mae and Freddie Mac, the collapse of Lehman Brothers, and the downgrading of the AIG's credit rating. In fact during the two weeks starting on 3-Oct-08, the S&P 500 dropped by 29.92%, which is larger than the 20% market decline experienced during the October Crash of 1987.

7.3 Recursive estimates and the VaR diagnostics

We now consider if the excess VaR violations documented above could have been avoided if the parameter estimates of the t -DCC model were updated at regular intervals. To simplify the computations we focussed on the IGARCH version of the model and re-estimated all its parameters (including the degree-of-freedom parameter v) every 13 weeks (or four times in a year). Using the recursive estimates of the PIT, U_t , and the VaR indicator d_t we obtained similar results for the post 2007 period. The KS statistic for the recursive estimates is 0.0518 as compared to the 5% critical value of 0.1381 and does not reject the null hypothesis that the recursive PIT values are draws from a uniform distribution. We also could not find any evidence of serial correlation in the PIT values. But as before, the violations of the VaR constraint were statistically significant with $z_\pi = -3.09$. The violations occur exactly on the same dates as when the parameters were fixed at the end of 2007. Updating of the estimates of the t -DCC model seem to have little impact on the diagnostic test outcomes.

8 Changing Volatilities and Correlations

The time series plots of volatilities are displayed in Figures 19-21 for returns on currencies, bonds and equities, respectively. Conditional correlations of Euro with other currencies, US 10 year bond futures with other bond futures, and S&P futures with other equity future indices are shown in Figures 22 to 24, respectively. To reduce the impact of the initialization on the plots of volatilities and conditional correlations initial estimates for 1994 are not shown. These figures clearly show the declining trends in volatilities over the 2003-2006 period just before the financial crisis which led to unprecedented rise in volatilities, particularly in the currency and equity markets. It is, however, interesting to note that return correlations have been rising historically and seem to be only marginally accentuated by the recent crisis. These trend could reflect the advent of Euro and a closer integration of the world economy, particularly in the euro

¹³We also carried out the VaR diagnostic test for the higher risk tolerance value of $\alpha = 5\%$, but did not find a statistically significant evidence against the t -DCC specifications. For both versions of the model the VaR constraint was violated 8 times, 3 more than one would have expected, giving $\hat{\pi} = 0.9167$ and $z_\pi = -1.50$ which is not significant at the 5% level. It is, however, interesting that all the eight violations occurred in 2008 with five of them occurring over the crisis months of 5-Sep-08 to 21-Nov-08.

area. Return correlations across asset types have also been rising, although to a lesser extent. An overall measure of the extent of the correlations across all the 17 assets under consideration is given by the maximum eigenvalue of the 17 by 17 matrix of asset return correlations. Figure 25 displays the conditional estimates of this eigenvalue over time and clearly show the sharp rise in asset return correlations particularly over the past two years.

9 Concluding Remarks

This paper applies the t -DCC model to the analysis of asset returns as a way of dealing with the fat-tailed nature of the underlying distributions. It is shown that the t -DCC model captures some of the main features of weekly asset returns. It fits the data reasonably well and seems to be computationally stable even for a moderate number of returns (17 in our application). Also when tested out of sample, it passes the serial correlation and Kolmogorov-Smirnov tests applied to probability integral transforms even over the highly turbulent weeks of 2008-2009 period. However, the model fails the VaR diagnostic test and the weekly returns on an equal weighted portfolio violates the VaR constraint 3 times over the six weeks from 5-Sep-08 to end 10-Oct-08. Two of these violations occur in two successive weeks. Conditional on the t -DCC model being valid, such an event could be expected to occur every 192 years! Of course, it could be argued that it is the inadequacy of the t -DCC model that has given rise to such an outcome, and a better model could have done better and such events are not as rare as suggested by the application of the t -DCC to the post 2007 observations. This is an important open question and its resolution is beyond the scope of the present paper. But it seems doubtful if modifications of the t -DCC suggested in the literature, such as allowing for asymmetry or leverage effects, could resolve the DCC's poor performance during crisis periods. The use of more fat-tailed distributions, such as mixtures of multivariate normal distributions as considered in Pesaran, Schleicher, and Zaffaroni (2009) is likely to be more effective. But the problem of matching volatility models to the data in normal as well as in crisis times would be a real challenge. A fat-tailed distribution suited to the crisis period might yield outcomes that are too conservative in normal times, whilst a model with satisfactory performance in normal times generally performs poorly during a crisis period. Developing a model that switches between the two states seems a sensible strategy, but it requires a reliable early warning system that is capable of accurately identifying periods of crisis *ex ante*, a goal which might not be attainable.

Our analysis also shows falling conditional volatilities and rising correlations during the 2003-2007, before the emergence of the financial crisis in 2008. These trends seem to have been important contributory factors to the emergence of the crisis. Low levels of volatilities might have tempted many investors and traders to take more risks, at times when asset return correlations had been rising. The crisis led to a reversal of the trend in volatilities and accentuated the rising correlations, particularly across the equity returns. Although volatilities have

fallen substantially from their heights in 2008, they are still high by historical standards. Return correlations continue to be high and in some cases are rising even. Further research is clearly needed for a better understanding of asset return correlations and their evolution over time.

Table 1: Summary Statistics for Raw Weekly Returns and Devolatilized Weekly Returns - 01-Apr-94 to 30-Oct-09

Asset	Returns				Devolatilized Returns			
	Mean	S.D.	Skewness	Ex. Kurtosis	Mean	S.D.	Skewness	Ex. Kurtosis
Currencies								
Australian dollar	0.044	1.690	-1.163	7.886	0.059	1.005	-0.214	-0.112
British pound	0.019	1.297	-0.831	5.348	0.037	1.013	-0.148	-0.197
Canadian dollar	0.035	1.136	-0.739	7.443	0.031	1.023	-0.040	-0.266
Swiss franc	0.053	1.517	0.210	1.071	0.044	0.994	0.146	-0.299
Euro	0.039	1.381	-0.043	1.424	0.044	1.012	-0.008	-0.281
Yen	0.031	1.669	1.326	9.462	-0.009	1.016	0.328	0.139
Bonds								
Euro Bunds	0.070	0.755	-0.378	0.910	0.123	1.000	-0.210	-0.205
UK Gilt	0.051	0.893	-0.013	1.744	0.068	1.008	-0.015	-0.290
Japan JGB	0.072	0.578	-0.436	2.323	0.152	1.007	-0.364	0.022
US TNote	0.077	0.894	-0.359	0.954	0.084	1.004	-0.243	-0.188
Equities								
S&P 500	0.094	2.575	-0.749	8.018	0.054	1.011	-0.314	-0.124
Nikkei	-0.017	3.175	-0.979	9.645	-0.005	0.996	-0.235	-0.147
FTSE	0.060	2.535	-0.858	10.399	0.042	1.002	-0.264	-0.132
CAC	0.107	3.116	-0.656	5.473	0.043	1.003	-0.216	-0.478
DAX	0.113	3.398	-0.559	5.673	0.055	1.008	-0.312	-0.220
SM	0.137	2.819	-0.734	10.174	0.077	1.005	-0.349	0.077
AUS	0.083	2.118	-0.670	4.698	0.066	1.001	-0.224	-0.253

Table 2: Maximized log-likelihood Values of DCC Models Estimated with Weekly Returns over 27-May-94 to 28-Dec-07

Assets	Standardized Returns			Devolatized returns		
	Normal	t-distribution	D.F.	Normal	t-distribution.	D.F.
Currencies (6)	-5783.7	-5689.8	9.62 (1.098)	-5790.6	-5694.1	9.24 (0.94)
Bonds (4)	-2268.5	-2243.5	11.28 (2.00)	-2270.7	-2246.9	11.35 (5.53)
Equities (7)	-9500.1	-9380.7	7.96 (0.74)	-9504.4	-9383.2	7.79 (0.72)
All 17	-17509.2	-17244.8	11.84 (0.90)	-17510.4	-17250.4	12.11 (0.92)

Note: D.F. is the estimated degrees of the freedom of the multivariate t-distribution.
Standard errors of the estimates are given in round brackets.

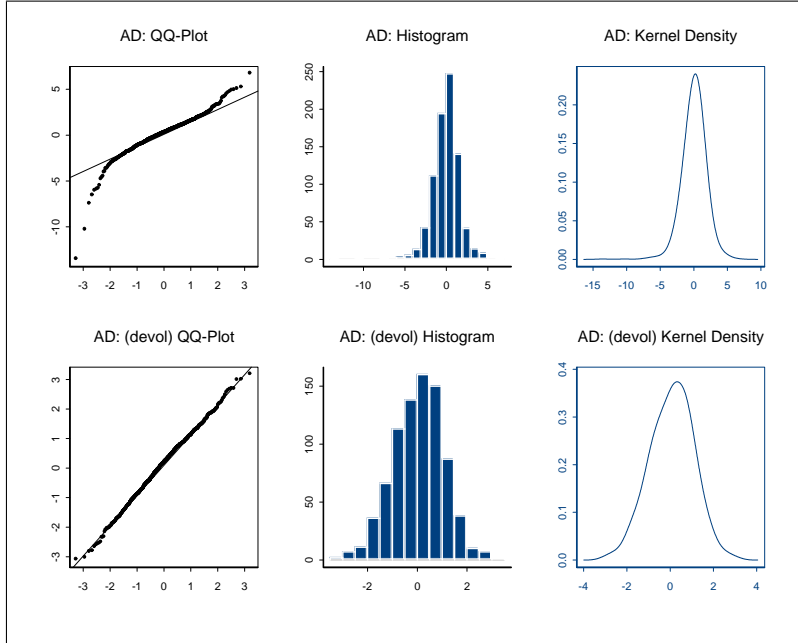


Figure 1: Australian dollar weekly returns (simple and devolatized) 01-Apr-1994 to 30-Oct-2009

Table 3: ML Estimates of t-DCC Model Estimated with Weekly Returns over the Period 27-May-94 to 28-Dec-07

Asset	ML Estimates		
	$\hat{\lambda}_1$	$\hat{\lambda}_2$	$1 - \hat{\lambda}_1 - \hat{\lambda}_2$
Currencies			
Australian dollar	0.9437 (0.0201)	0.0361 (0.0097)	0.0201 (0.0140)[1.44]
British pound	0.9862 (0.0110)	0.0124 (0.0056)	0.0014 (0.0081)[0.18]
Canadian dollar	0.9651 (0.0102)	0.0349 (0.0102)	0 (N/A)[N/A]
Swiss franc	0.9365 (0.0517)	0.0303 (0.0157)	0.0332 (0.0378)[0.88]
Euro	0.9222 (0.0264)	0.0487 (0.0133)	0.0291 (0.0154)[1.89]
Yen	0.9215 (0.0235)	0.0586 (0.0151)	0.01992 (0.0107)[1.86]
Bonds			
Euro Bunds	0.9031 (0.0237)	0.0703 (0.0149)	0.0266 (0.0118)[2.26]
UK Gilt	0.9062 (0.0304)	0.0774 (0.0224)	0.0164 (0.0091)[1.80]
Japan JGB	0.8179 (0.0369)	0.1444 (0.0268)	0.0377 (0.0141)[2.74]
US TNote	0.9072 (0.0249)	0.0714 (0.0165)	0.0216 (0.0115)[1.87]
Equities			
CAC	0.9252 (0.0118)	0.0674 (0.0099)	0.0074 (0.0033)[2.23]
DAX	0.9267 (0.0117)	0.0653 (0.0095)	0.0080 (0.0039)[2.03]
Nikkei	0.9552 (0.0305)	0.0402 (0.0210)	0.0046 (0.0109)[0.42]
S&P 500	0.9326 (0.0194)	0.0582 (0.0150)	0.0091 (0.0060)[1.53]
FTSE	0.9298 (0.0144)	0.0589 (0.0109)	0.0112 (0.0052)[2.16]
SM	0.9066 (0.0225)	0.0774 (0.0165)	0.0160(0.0076)[2.11]
AUS	0.9393 (0.0295)	0.0370 (0.0128)	0.0237(0.0194)[1.22]
$\hat{v} = 12.11 (0.9233)$, $\hat{\phi}_1 = 0.9673 (0.0037)$, $\hat{\phi}_2 = 0.0172 (0.0012)$ [5.49]			
Note: Standard errors of the estimates are given in round brackets, t -statistics are given in square brackets. λ_{1i} and λ_{2i} are the asset-specific volatility parameters. ϕ_1 and ϕ_2 are the common conditional correlation parameters.			

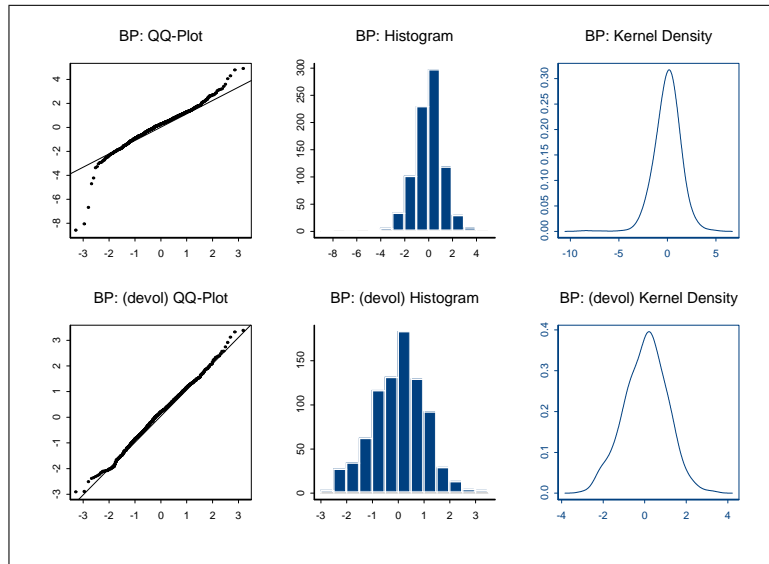


Figure 2: British pound weekly returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

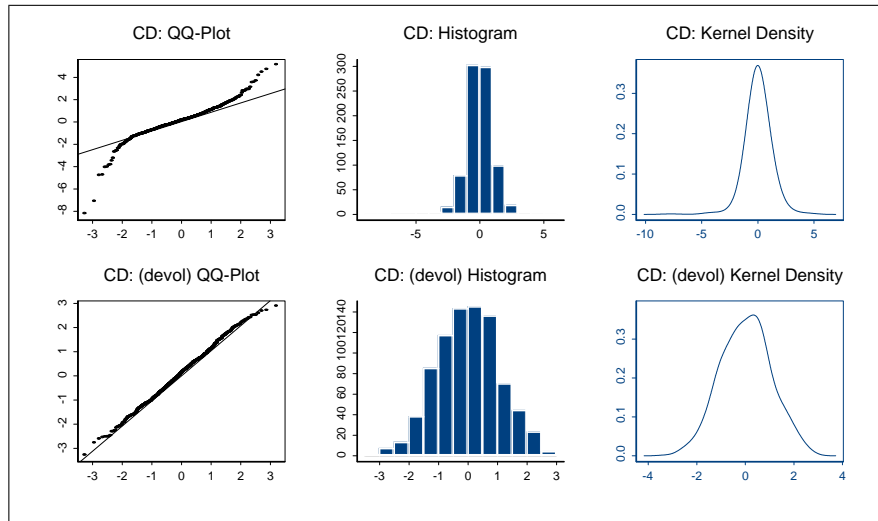


Figure 3: Canadian dollar weekly returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

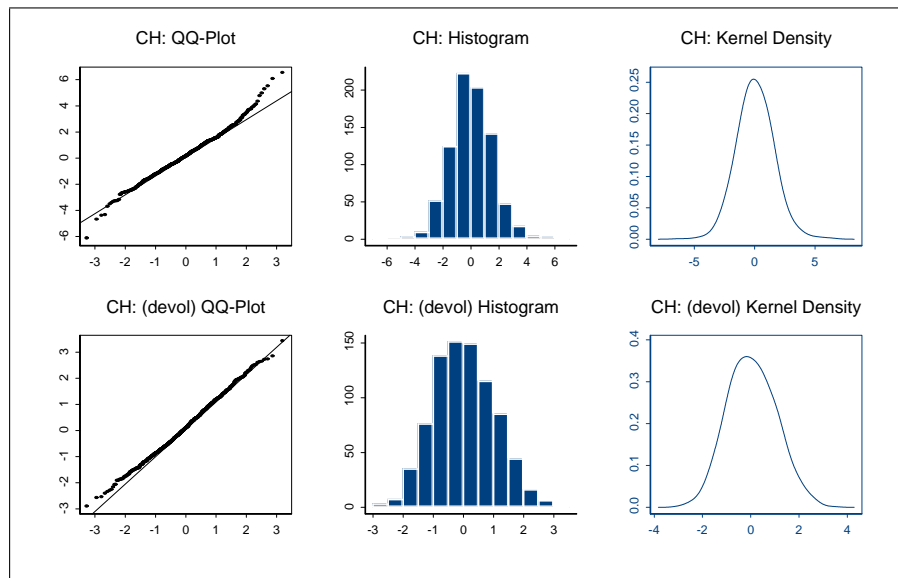


Figure 4: Swiss franc weekly returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

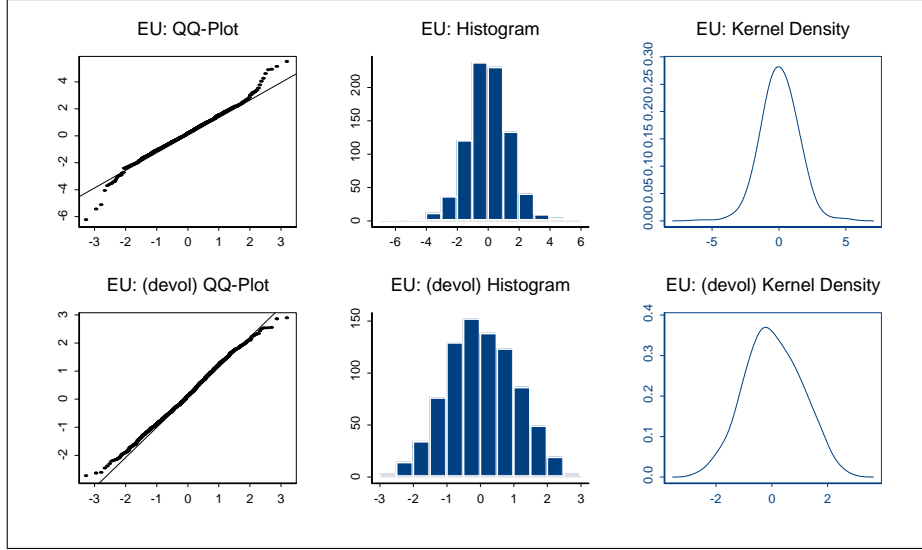


Figure 5: Euro weekly returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

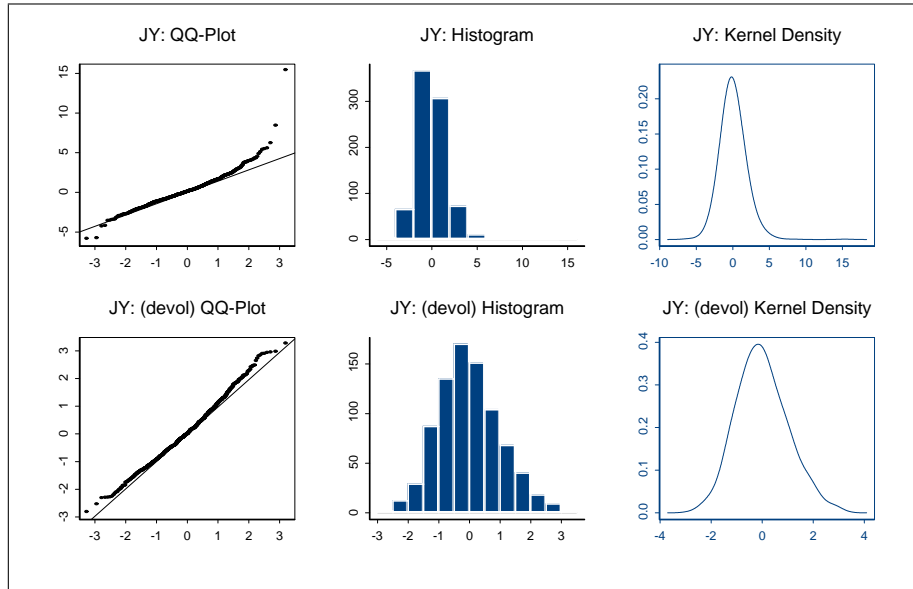


Figure 6: Yen weekly returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

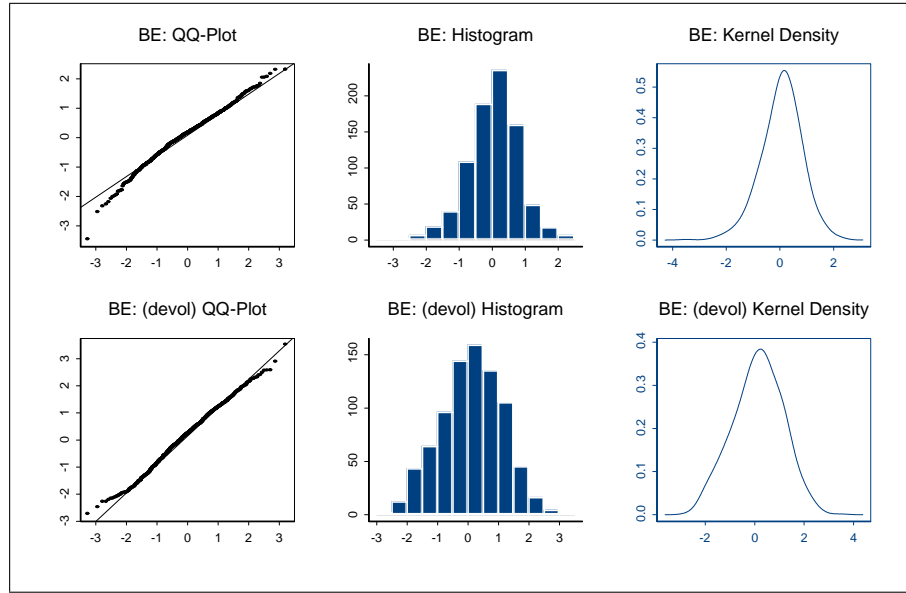


Figure 7: Euro Bunds weekly returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

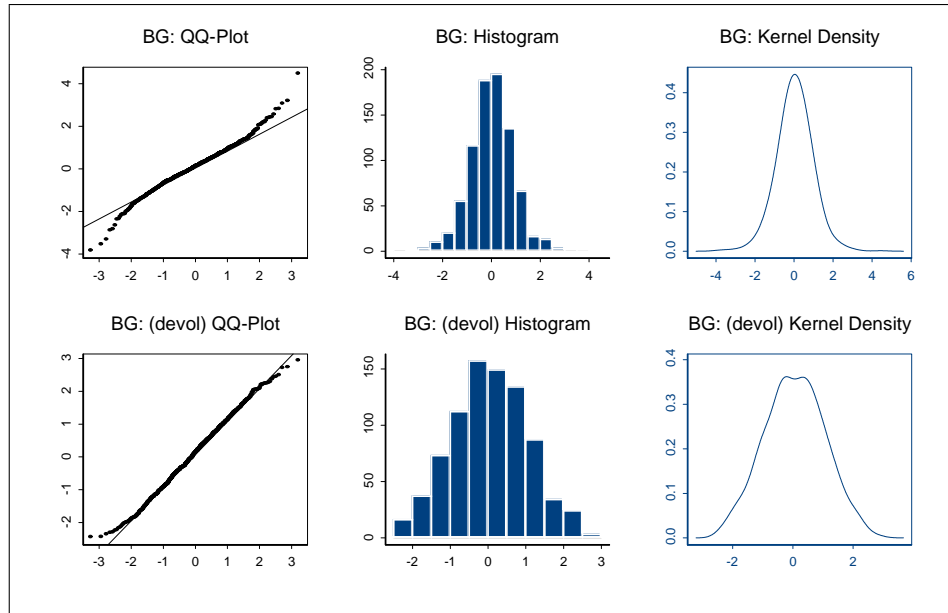


Figure 8: UK Gilt weekly returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

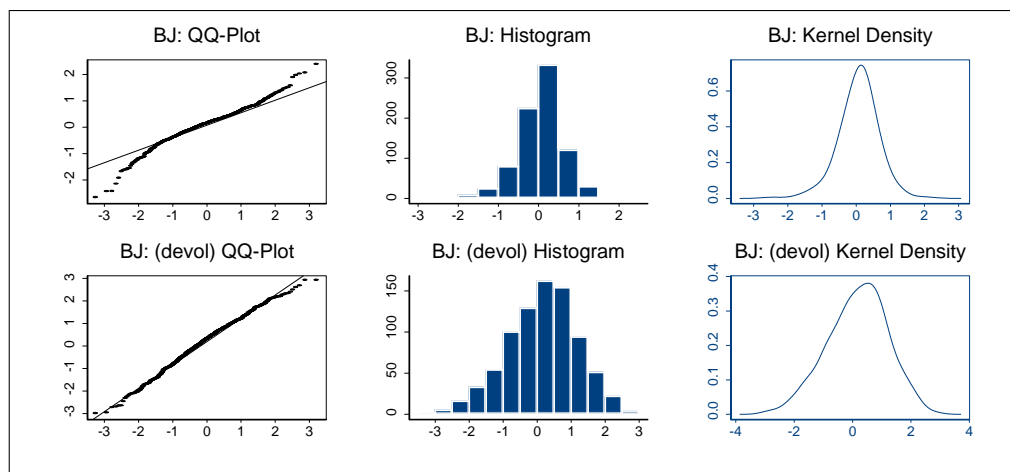


Figure 9: Japanese Government bond weekly returns (simple and devolatilized)
01-Apr-1994 to 30-Oct-2009

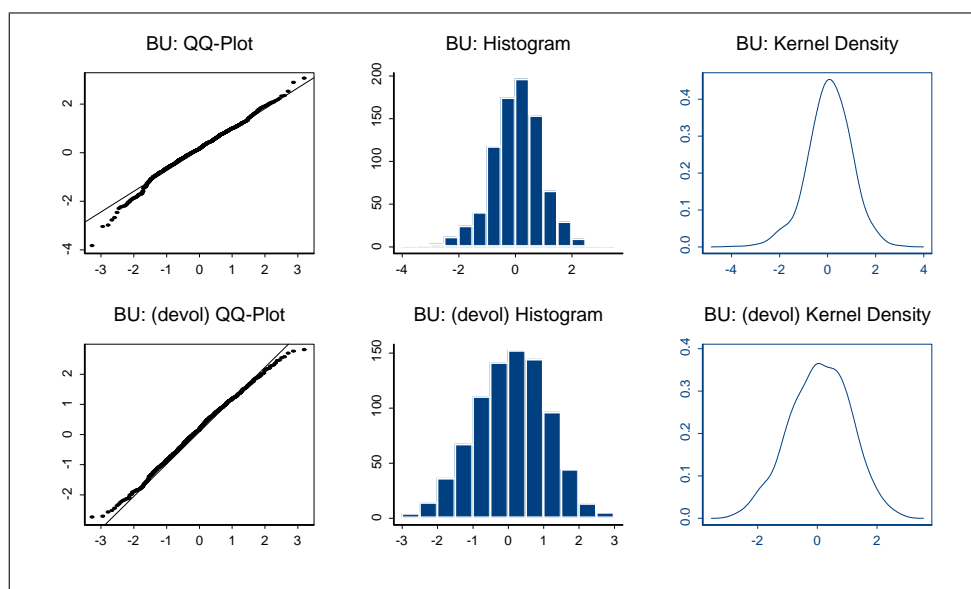


Figure 10: US TNote weekly returns (simple and devolatilized) 01-Apr-1994 to
30-Oct-2009

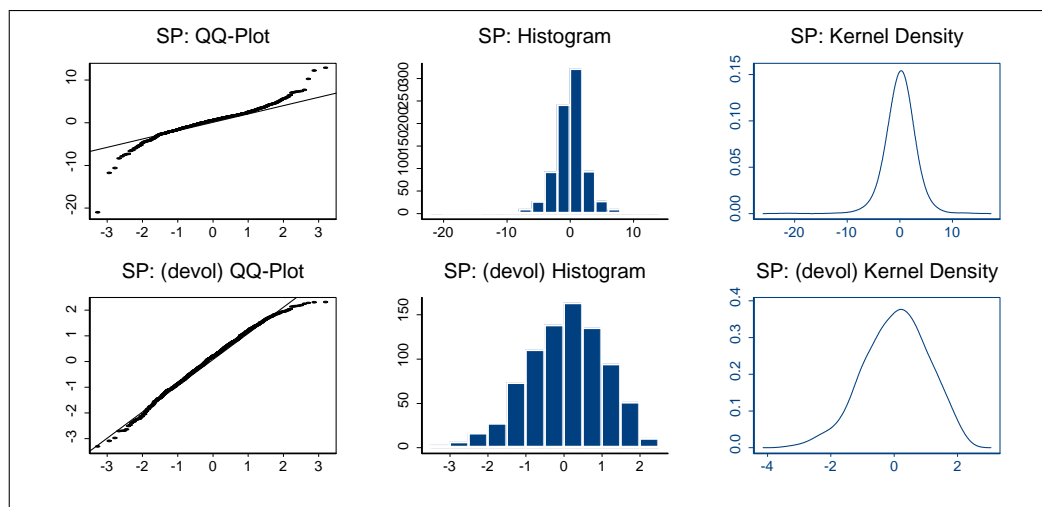


Figure 11: S&P 500 weekly returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

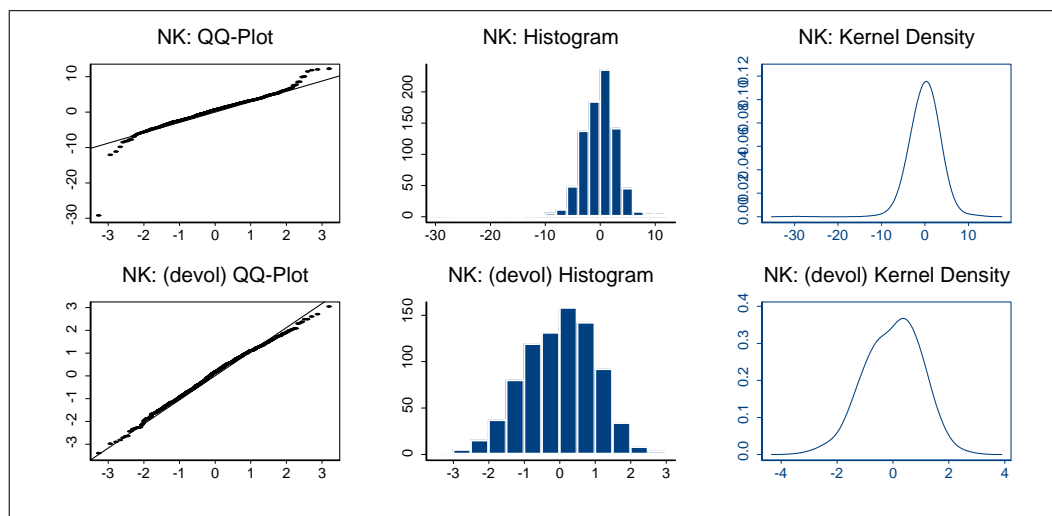


Figure 12: Nikkei weekly returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

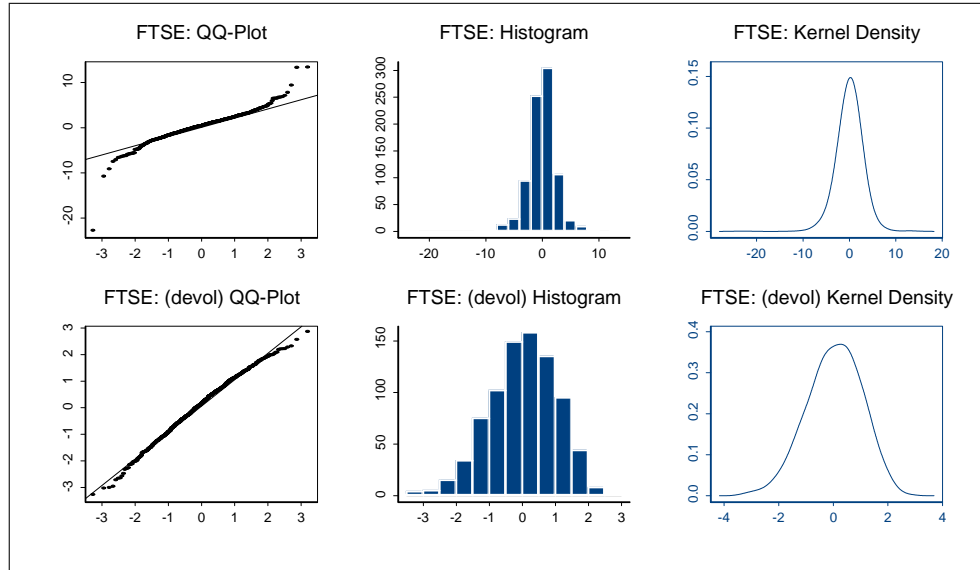


Figure 13: FTSE weekly returns (simple and devolitized) 01-Apr-1994 to 30-Oct-2009

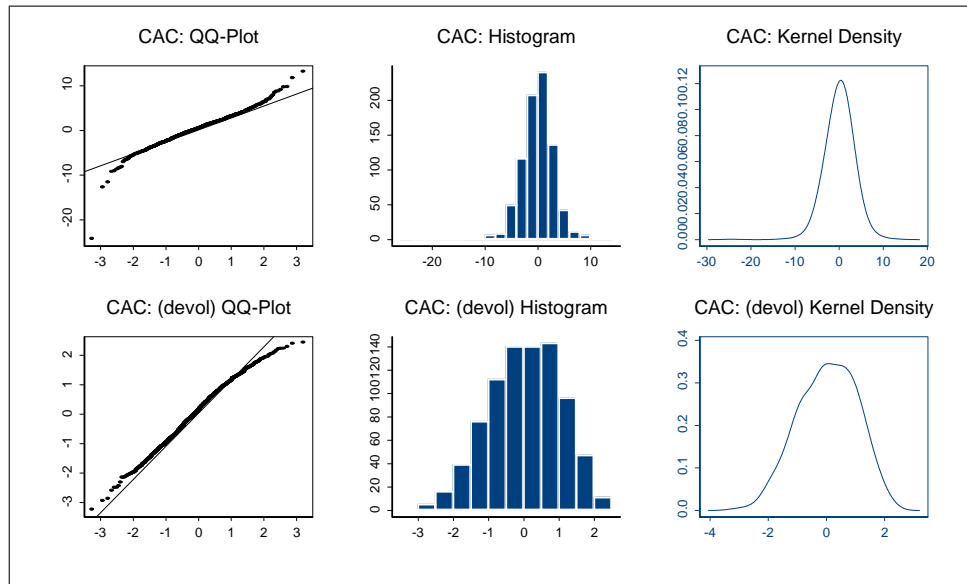


Figure 14: CAC weekly returns (simple and devolitized) 01-Apr-1994 to 30-Oct-2009

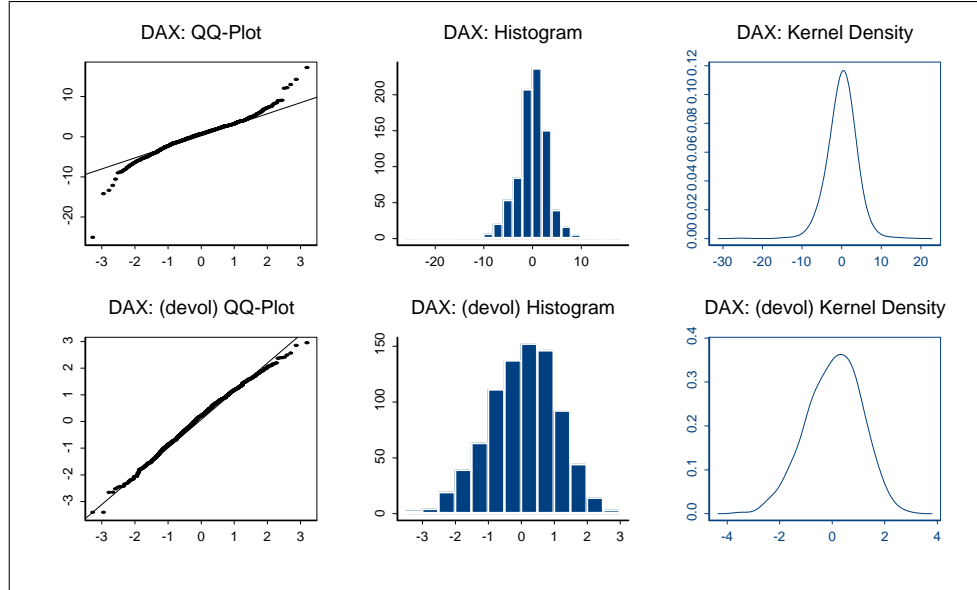


Figure 15: DAX weekly returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

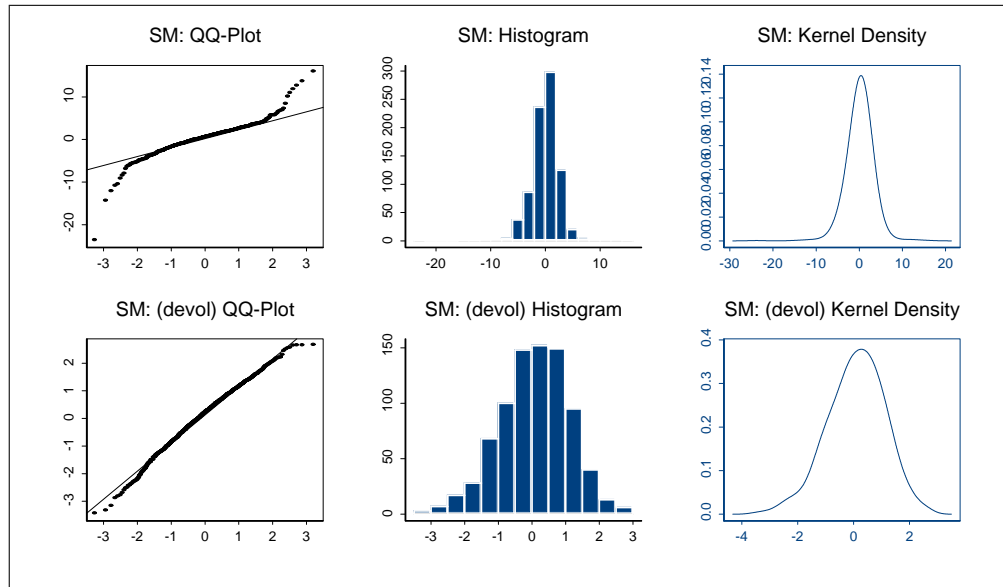


Figure 16: Swiss weekly equity returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

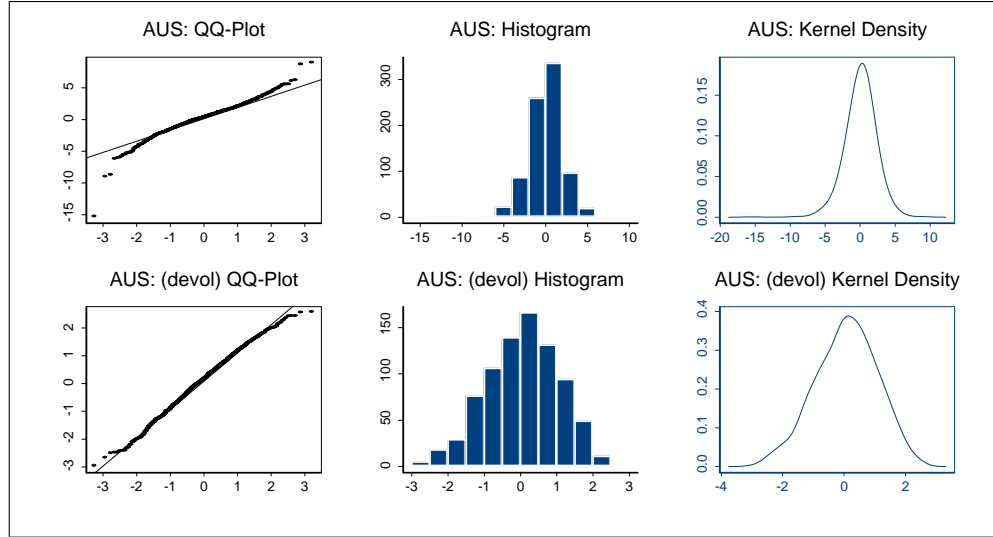


Figure 17: Australian weekly equity returns (simple and devolatilized) 01-Apr-1994 to 30-Oct-2009

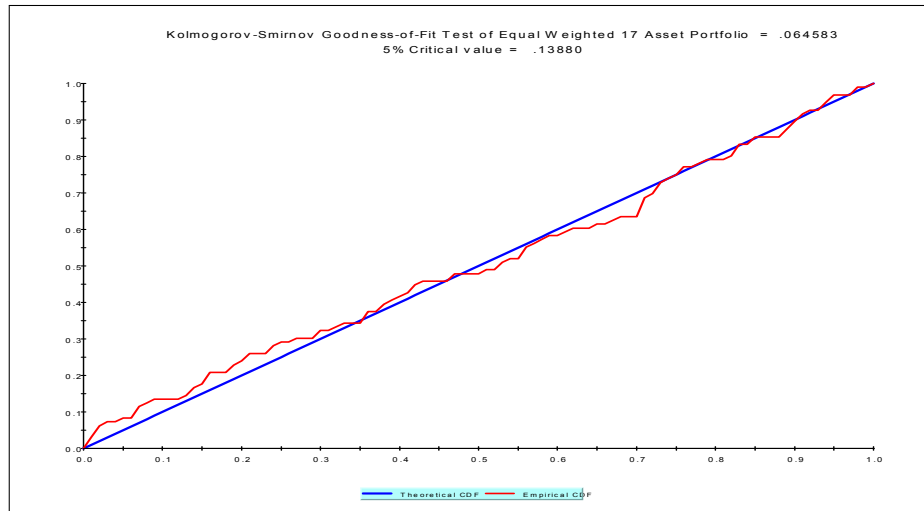


Figure 18: Kolmogorov-Smirnov goodness-of-fit test for the full t-DCC model over the evaluation sample 4-Jan-08 to 30-Oct-09




Figure 19: Currency Returns Volatilities




Figure 20: Bond Returns Volatilities




Figure 21: Equity Return Volatilities

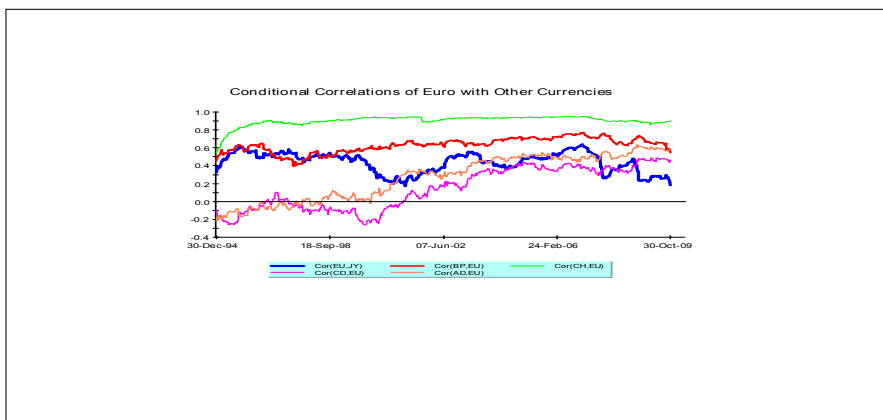


Figure 22: Conditional Correlation of Euro




Figure 23: Conditional Correlations of US 10 Year Bond

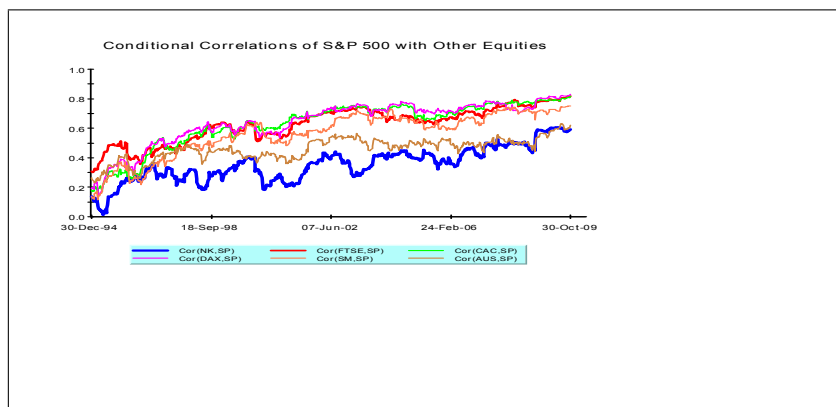


Figure 24: Conditional Correlations of S&P 500




Figure 25: Maximum Eigenvalue of Conditional Correlation Matrix

References

- [1] Andersen, T., Bollerslev, T., Diebold, F.X. and Ebens, H. (2001), "The Distribution of Realized Stock Return Volatility", *Journal of Financial Economics*, 61, 43-76.
- [2] Andersen, T. Bollerslev, T., Diebold, F.X. and Labys, P. (2001), "The Distribution of Realized Exchange Rate Volatility", *Journal of the American Statistical Association*, 96, 42-55.
- [3] Barndorff-Nielsen, O. E. and N. Shephard N. (2002), "Econometric Analysis of Realised Volatility and its use in Estimating Stochastic Volatility Models", *Journal of the Royal Statistical Society, Series B*, 64, 253-280.
- [4] Bauwens, L., Laurent, S., Rombouts, J.V.K., (2006), "Multivariate GARCH models: A survey", *Journal of Applied Econometrics*, 21, 79-109.
- [5] Berkowitz, J., (2001), "Testing Density Forecasts with Applications to Risk Management", *Journal of Business & Economic Statistics*, 19, 465-474.
- [6] Bollerslev, T. (1990), "Modelling the Coherence in Short Run Nominal Exchange Rates: a Multivariate Generalized ARCH Model" *Review of Economics and Statistics*, 72, 498-505.
- [7] Christoffersen, P.F. (1998), "Evaluating Interval Forecasts", *International Economic Review*, 39, 841-862.
- [8] Embrechts, P. A. Hoing, and A. Juri (2003) "Using Copulas to Bound VaR for Functions of Dependent Risks", *Finance and Stochastics*, 7, 145-167.
- [9] Engle, R. F., and K. Kroner, (1995) "Multivariate Simultaneous GARCH", *Econometric Theory*, 11, 122-150.
- [10] Engle, R.F. (2002), "Dynamic Conditional Correlation - A Simple Class of Multivariate GARCH Models", *Journal of Business Economics & Statistics*, 20, 339-350.
- [11] Engle, R.F., and S. Manganelli, (2002), "CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles", *Journal of Business & Economic Statistics*, 22, 367-381.
- [12] Fiorentini, G., E. Sentana, and G. Calzolari, (2003), "Maximum Likelihood Estimation and Inference in Multivariate Conditionally Heteroscedastic Dynamic Regression Models With Student t Innovations", *Journal of Business & Economic Statistics*, 21, pp. 532-546.
- [13] Litterman, R., and K. Winkelmann (1998), "Estimating Covariance Matrices", *Risk Management Series*, Goldman Sachs.
- [14] Lopez, J.A. (1998), "Methods for Evaluating Value-at-Risk Estimates", *Economic Policy Review, Federal Reserve Bank of New York*, 119-124.

- [15] Massey, F. J. (1951), "The Kolmogorov-Smirnov Test for Goodness of Fit", *Journal of the American Statistical Association*, 46, 68-78.
- [16] Neave, H.R. and P.L. Worthington (1992), *Distribution-free Tests*, Routledge, London.
- [17] Pesaran B., and M.H. Pesaran (2007), "Modelling Volatilities and Conditional Correlations in Futures Markets with a Multivariate t Distribution", *Cambridge Working Papers in Economics*, No. 0734, and *IZA Discussion Papers* No. 2906, Institute for the Study of Labor.
- [18] Pesaran B., and M.H. Pesaran (2009), *Time Series Econometrics using Microfit 5.0*, Oxford University Press, Oxford.
- [19] Pesaran M.H., and A. Timmermann (2005), "Real Time Econometrics", *Econometric Theory*, 21, 212-231.
- [20] Pesaran, M.H., C. Schleicher, and P. Zaffaroni (2009), "Model Averaging in Risk Management with an Application to Futures Markets", *Journal of Empirical Finance*, 16, 280-305.
- [21] Sentana, E. (2000), "The Likelihood Function of Conditionally Heteroskedastic Factor Models", *Annales d'Economie et de Statistique* 58, 1-19.
- [22] Zaffaroni, P. (2008), "Estimating and forecasting volatility with large scale models: theoretical appraisal of professionals' practice", *Journal of Time Series Analysis*, 581-599, 2008.