Dynamic copula models for multivariate high-frequency data in finance

Alexandra Dias[a], Paul Embrechts[b,*]

[a] *Department of Mathematics, ETH Zurich and Fac. Ciências e Tecnologia, UN Lisboa*
[b] *Department of Mathematics, ETH Zurich and London School of Economics*

Abstract

The stylized facts of univariate high-frequency data in finance are well known; see Dacorogna et al. (2001). In this paper, we concentrate on multivariate high-frequency data, and analyse the dependence structure, through the notion of copula. More in particular, we analyse the conditional copula for two dimensional high-frequency data. We find evidence of non-constancy over time of the conditional copula and investigate this fact at six different time aggregation frequencies. Further, we investigate the existence of change-points in the conditional copula. The data analysed are high-frequency foreign exchange spot rates for US Dollar quoted against German Mark and quoted against Japanese Yen, previously analysed in a non-dynamic setting in Breymann et al. (2003).

JEL classification: C32; C51; F31

Keywords: Multivariate time series; Conditional copula; Dynamic copula; Change-points; High-frequency foreign exchange data

*Corresponding author. Department of Mathematics, ETH Zurich, CH – 8092 Zürich, Switzerland. Tel.: +41 44 632 34 19. Fax.: +41 44 632 15 23. E-mail: embrechts@math.ethz.ch

1 Introduction

On dealing with multivariate risks, the dynamics of the dependence structure is of importance. For example, there is considerable interest in the dynamic behaviour of correlation between different risks as a function of time; see for instance Boyer et al. (1999), Longin and Solnik (2001) and Loretan and Phillips (1994). Because of the fundamental importance of the notion of linear correlation in finance and insurance, such dynamics may have a non-trivial impact on the pricing and hedging of underlying instruments, or on the risk measurement of such positions. As a consequence, a more systematic modelling for the dynamic behaviour of the dependence structure underlying multivariate risks is called for.

In Breymann et al. (2003) we investigated the stylized facts of the dependence structure in a set of high-frequency data, namely tick-by-tick observations of foreign exchange (FX) spot rates for US Dollar quoted against German Mark (USD/DEM) and quoted against Japanese Yen (USD/JPY). After an initial data deseasonalization, bivariate log-return time series for six different frequencies were considered, from one hour up to one day. At each frequency we evaluated the dependence structure, fitting copula-based models by the pseudo log-likelihood method introduced by Genest et al. (1995); see also Chen and Fan (2002). It is important to stress that the above analysis assumed the vectors to be independent and identically distributed (iid). We know however that this assumption is violated in practice due for instance to volatility effects; papers like Fortin and Kuzmics (2002), Patton (2002), Rockinger and Jondeau (2001) discuss this issue.

In the present paper, starting from the deseasonalized FX data, we first model the volatility dynamics with GARCH type models. The dependence between the resulting residuals is then analysed through dynamic copula models. Of main concern will be the detection of change-points.

There is an enormous econometric literature on the use of regime changes for describing non-stationary economic data. In the context of time series analysis, see for instance Hamilton (1990). In the field of monetary policy, regime switches may come about as a consequence of new policy implementation, see Francis and Owyang (2004). In the latter context, tools have been developed to measure the influence of both regime changes as well as policy shocks. A further series of publications concentrates on the relationship between long-memory and the existence of structural breaks, so-called "spurious long memory process"; see Choi and Zivot (2003). In many of the above publications, the issue of changes in correlations is important. One of the main reasons for writing our RiskLab report, Embrechts et al. (1999, 2002), was to encourage a more structured dialogue between practicioners and academics on the issue of dependence. Through many examples, we now know that "there is more to dependence as can be measured through (linear) correlation". It turns out that the notion of copula as discussed in Embrechts et al. (2002) yields an excellent tool for the modelling of nonlinear dependences. By now, these techniques have achieved the text book level, see for instance Cherubini et al. (2004). The present paper contributes to our knowledge of dependence modelling beyond linear correlation by adding an analysis on conditional dependence across sampling frequencies in two-dimensional FX data, and also explores the existence of structural changes in the dependence. Concretely, we introduce a new methodology for testing for change-points in conditional copulae. The main aim of the paper is to give, through an economically relevant example, the financial experts a new tool that we hope will have potential for a more careful analysis of financial data. As such, our paper is very much in line with the goals set out in Andersen et al. (2004).

After a preliminary data analysis in Section 2, in Section 3.1 we first filter the data through univariate GARCH models and analyze, in Section 3.2, the copula function of the residuals. In Section 3.3 we propose a multivariate dynamic model where the copula is time-varying. The model presented allows the use of any copula to link the univariate innovations. In Section 3.3.1 we explain how to estimate the parameters of the time-varying copula-based model. We apply this model to the several frequencies of FX returns on USD/DEM and USD/JPY in Section 3.3.2. Based on the models from Section 3.2, Section 4.2 is devoted to testing for the existence of change points in the dependence parameters; we estimate both the size and time of the changes.

2 Preliminary data analysis

The data considered are bivariate log-returns of FX spot rates USD/DEM and USD/JPY after being deseasonalized, as in Breymann et al. (2003). The observations cover the period from April 27, 1986 until October 25, 1998. The six different frequencies considered are one, two, four, eight, twelve hours and one day periods. In order to make the paper more self-contained, in the next section we summarise the deseasonalization procedure. For further details we refer the reader to Breymann et al. (2003) and the references therein.

2.1 Deseasonalization of the data

After being collected by Olsen Data, the observations (tick-by-tick FX quotes) are corrected for transmission errors, fake quotes originated by transmission tests, etc. A description of this filtering process is to be found in Dacorogna et al. (2001).

Originally the number of quotes emitted by the FX market is very high (around ten million for USD/DEM in the given period) and they are irregularly spaced in time. Regular time series are obtained through a reduction of the observations to a time step δ of five minutes. Missing observations are obtained by linear interpolation. For a given currency, a single quote at time t consists of a bid price, $p_{t,\mathrm{bid}}$ and an ask price, $p_{t,\mathrm{ask}}$. As we are not interested in the effects related to the bid-ask spread, we consider logarithmic middle prices, defined at time t as:

$$\bar{p}_t = \frac{1}{2} \left(\log\, p_{t,\mathrm{bid}} + \log\, p_{t,\mathrm{ask}} \right).$$

These are displayed in Figure 1 for the USD/DEM and USD/JPY spot rates. The logarithmic return at t with respect to the time horizon (or frequency) Δt, $\tilde{x}_{t,\Delta t}$ is given as

$$\tilde{x}_{t,\Delta t} = \bar{p}_t - \bar{p}_{t-\Delta t}. \tag{1}$$

We may occasionally drop the frequency subscript Δt in (1).

The changing market activity induces a cyclic behaviour on the variability of the returns. In other words, the market volatility possesses a seasonal component, the so-called seasonal volatility pattern. This is very pronounced in high-frequency data.

Specific approaches to the deseasonalization of high-frequency data can be found in the literature; at a text book level, see Dacorogna et al. (2001). We use a deseasonalization method referred to as volatility weighting, which consists of standardising the time series of the returns by a volatility, estimated conditionally on the time of the week. Formally, denote by $\widetilde{X}_t$ the log-return at time t for the time frequency of five minutes. Suppose that the return time series data is a realisation of the process

$$\widetilde{X}_t = \mu + v_t\, X_t, \tag{2}$$

where μ is the mean log-return, v_t is the five minutes expected volatility at t and X_t is a random component corresponding to a deseasonalized log-return at moment t for the five minutes frequency. We assume in (2) that the trend component of the process is constant, the expected volatility however depends strongly on market activity, a_t. For our data we have to consider the activity of the European, the American and East Asian markets and hence split a_t in three components:

$$a_t = a_{t,\text{America}} + a_{t,\text{East Asia}} + a_{t,\text{Europe}}.$$

For a normal day, the market activity a_t is taken to be one. If a given time t_0 falls within a public holiday, in Europe say, then $a_{t_0,\text{Europe}} = 0$ and a_{t_0} is less than one. Moreover, the expected volatility is subject to the shift of daylight saving time (DST) periods. Hence, we have to compute separately the expected volatility for winter and for summer DST periods. We have to estimate the volatility also conditional on the moment of the week because there is a weekly and intra-day seasonality. We denote $t = t_i + \tau\delta$, where t_i is the beginning of week i (Sunday 00:00:00 GMT) and $\tau \in \{0, 1, \ldots, T\}$ with T the number of five minutes periods δ in a seven days week. Taking the above facts into consideration, v_t is modelled as

$$v_t^2 = a_t \left(v_{\tau\delta}^{(d)}\right)^2, \tag{3}$$

where $v_{\tau\delta}^{(d)}$ is the expected volatility at moment $\tau\delta$ of the week in the DST period d, winter or summer. Then, $v_{\tau\delta}^{(d)}$ is estimated as

$$\left(\bar{v}_{\tau\delta}^{(d)}\right)^2 = \frac{1}{N_d} \sum_{i=1}^{N_d} \left(\tilde{x}_{t_i+\tau\delta}\right)^2,$$

where N_d is the number of weeks in the DST period d.

Using the volatility patterns estimated by (3), we can compute the five minutes deseasonalized log-returns

$$x_t = \frac{\tilde{x}_t - \bar{x}}{\bar{v}_t}, \tag{4}$$

where $\bar{x}$ is the sample mean of the log-returns $\tilde{x}_t$ which were computed from (1) with $\Delta t = \delta$. For simplicity, and without loss of generality, assume in the following that the returns $\widetilde{X}$ are already centred by their mean μ. Now we need to know how to obtain the deseasonalized returns for any time frequency Δt rather than for five minutes. Consider the simple generalisation of (2) from a five minutes to a Δt time frequency process:

$$\widetilde{X}_{t,\Delta t} = v_{t,\Delta t}\, X_{t,\Delta t}. \tag{5}$$

The expected volatility at time t for an arbitrary time frequency Δt is defined in terms of the five minutes volatility as

$$v_{t,\Delta t} = \left(\sum_{i=0}^{n-1} v^2_{t-i\delta} \right)^{1/2},$$

where $n\delta = \Delta t$. Denote by $\overline{P}_t$ the logarithmic middle price at time t. Then, (5) can be rewritten in the form

$$\overline{P}_t - \overline{P}_{t-\Delta t} = \left(\sum_{i=0}^{n-1} v^2_{t-i\delta} \right)^{1/2} X_{t,\Delta t}, \tag{6}$$

with $n\delta = \Delta t$. Using (6), we can compute the deseasonalized log-returns for any time frequency Δt through

$$x_{t,\Delta t} = \frac{\bar{p}_t - \bar{p}_{t-\Delta t}}{\left(\bar{v}^2_t + \bar{v}^2_{t-\delta} + \ldots + \bar{v}^2_{t-(n-1)\delta} \right)^{1/2}}; \tag{7}$$

this is a function of the logarithmic middle prices and the five minutes volatility pattern estimated from the data.

The mean activity of the FX market during weekends is very low. The usual approach to handle this cyclic behaviour consists of dropping the weekends from the middle price series. However, big jumps may appear between closing Friday and opening Monday prices, so that dropping the weekends without a special treatment could produce wrong large return values. This would then induce spurious seasonality. Therefore, a weekend volatility has to be calculated in order to avoid this. Define the beginning and the end of a weekend respectively as t_{w_0} = Friday, 21:00:00 GMT and t_{w_1} = Sunday, 21:00:00 GMT. Let $\Delta t_w = t_{w_0} - t_{w_1}$ be the weekend length. The expected volatility for weekends is estimated as

$$\bar{v}_t = \left| \tilde{x}_{t_{w_1},\Delta t_w} \right| \left(\frac{\delta}{\Delta t_w} \right)^{1/2}, \tag{8}$$

if $t \in (t_{w_0}, t_{w_1}]$. After the volatility pattern estimation is done, the weekends can be dropped in the middle price series and in the volatility. Because of (8), the first volatility estimate of the week is in tune with a possible price jump during the weekend and no spurious seasonality remains in the deseasonalized returns.

With the volatility patterns estimated, we computed the series of deseasonalized log-returns for six frequencies: one, two, four, eight, twelve hours and one day.

2.2 Analysis of the margins

As mentioned before, the FX data set available covers the period from April 1986 until October 1998. Tables 1 and 2 contain some summary statistics of the deseasonalized log-returns. From those tables we observe that neither exchange rate return shows a significant trend. The skewness of the USD/JPY deseasonalized returns is negative for all frequencies. For the USD/DEM returns, the skewness is closer to zero and even positive for the two and eight hours frequencies. At all frequencies, both series show excess kurtosis.

The univariate normality of the deseasonalized returns can be formally tested using the Jarque–Bera test from Jarque and Bera (1987). We ran this test and normality is rejected throughout; see Table 3.

Furthermore, Figures 2 and 3 contain normal QQ-plots for the empirical quantiles of the deseasonalized returns. As it is usually found in the econometric literature, from those two figures we can see that the univariate distributions are clearly heavy tailed for short time horizons and become more thin tailed as the frequency decreases.

In order to test for conditional heteroscedasticity, we consider the ARCH(p) process

$(X_t)_{t\in\mathbb{Z}}$ defined as

$$\begin{aligned} X_t &= \mu + \epsilon_t \\ \epsilon_t &= \sigma_t Z_t \\ \sigma_t^2 &= \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \ldots + \alpha_p \epsilon_{t-p}^2, \end{aligned} \tag{9}$$

where $(Z_t)_{t\in\mathbb{Z}}$ is a sequence of iid random variables with zero mean and unit variance. Moreover, $\alpha_0 > 0$, $\alpha_i \geq 0$ for $i = 1, 2, \ldots, p$ and Z_t is independent of $(X_s)_{s\leq t}$ for all t; see Engle (1982). At all the frequencies, the test for absence of ARCH effects is rejected; see Table 4.

3 Conditional copula modelling across frequencies

Our main goal is to model the conditional dependence underlying the bivariate USD/DEM, USD/JPY spot rate returns across the different frequencies. By now, a standard approach is based on the notion of conditional copula, as discussed in Patton (2002). In this two-stage procedure, one first models the marginal dynamics.

3.1 Modelling the marginal dynamics

In the marginal tests performed in Section 2.2, the deseasonalized returns reveal the presence of time-varying variance and heavy taileness. In our discrete-time setting, we model stochastic volatility effects by GARCH type models. In particular, we fit univariate ARMA-GARCH models to each of the marginal series with innovations assumed to come from a t-distribution.

Formally, consider the sequence of iid random variables with zero mean and unit variance $(Z_t)_{t\in\mathbb{Z}}$. The process $(X_t)_{t\in\mathbb{Z}}$ is an ARMA(p_1, q_1)-GARCH(p_2, q_2) if it satisfies the equations

$$
\begin{aligned}
X_t &= \mu_t + \epsilon_t \\
\mu_t &= \mu + \sum_{i=1}^{p_1} \phi_i \left(X_{t-i} - \mu\right) + \sum_{j=1}^{q_1} \theta_j \epsilon_{t-j} \\
\epsilon_t &= \sigma_t Z_t \\
\sigma_t^2 &= \alpha_0 + \sum_{i=1}^{p_2} \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^{q_2} \beta_j \sigma_{t-j}^2
\end{aligned}
\tag{10}
$$

where $\alpha_0 > 0$, $\alpha_i \geq 0$ for $i = 1, 2, \ldots, p_2$, $\beta_j \geq 0$ for $j = 1, 2, \ldots, q_2$ and Z_t is independent of $(X_s)_{s \leq t}$. The polynomials $\phi(z) = 1 - \phi_1 z - \ldots - \phi_{p_1} z^{p_1}$ and $\theta(z) = 1 - \theta_1 z - \ldots - \theta_{p_1} z^{q_1}$ have no common roots and no roots inside the unit circle. See Brockwell and Davis (1991) and Zivot and Wang (2003) for more details.

We fitted univariate ARMA-GARCH models by maximum likelihood to each of the marginal series assuming that the innovations Z_t come from a t-distribution with ν degrees of freedom. Table 5 gives the order of the models fitted and the estimates of ν for the t-innovations; we refrain from listing the other parameter values as, for our analysis, these are less important. Note that, the t-distributions fitted at the one, two and four hours frequencies have infinite kurtosis ($\hat{\nu} < 4$) and so the fourth moment does not exist.

In each univariate model we considered an extra parameter γ in the GARCH dynamics. The parameter γ attempts to take into account that innovations of different signs may have asymmetric impacts on the future variance; see for example Bollerslev et al. (1992) and references therein. This improvement is also possible in the model specified in (10); see Ding et al. (1993) and Zivot and Wang (2003) where the GARCH component of model (10) is treated as a special case of a power GARCH model. The last equation in (10) then becomes

$$
\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p_2} \alpha_i (|\epsilon_{t-i}| + \gamma_i \epsilon_{t-i})^2 + \sum_{j=1}^{q_2} \beta_j \sigma_{t-j}^2 .
$$

We use the usual t-statistic $\hat{\alpha}/\hat{\sigma}_{\hat{\alpha}}$ to test whether the general model parameter α is zero. For the USD/DEM returns we cannot reject the null hypothesis of $\gamma_k = 0$ for the estimated γ parameter, and this at all frequencies. In the case of USD/JPY, the situation is almost the complete reverse. In this case, we reject the null hypothesis for all frequencies lower than (and including) four hours. In Table 5 we put a superscript * in the ARCH order of such frequencies. Rejecting the null hypothesis for the USD/JPY model parameter $\gamma_k = 0$, and using that the estimated values $\hat{\gamma}$ are negative, we have that negative shocks on the USD/JPY rate have a larger impact on volatility than positive shocks.

From the fitted ARMA-GARCH model parameters we recover the residuals or filtered returns $\hat{z}_t$ for each univariate time series $(x_1, x_2, \dots, x_n)$:

$$\hat{z}_t = (x_t - \hat{\mu}_t)/\hat{\sigma}_t, \qquad t = 1, 2, \dots, n. \tag{11}$$

Once the univariate models are selected and fitted, the dynamics as well as the goodness-of-fit of the t-density, must be checked. We use the Ljung–Box (L–B) test for testing for serial correlation and the Anderson–Darling (A–D) test for the goodness-of-fit of the t-density. For the USD/DEM residuals, in Table 6 we report the p-values of the L–B test for the residuals and for the absolute values of the residuals, and the p-values of the goodness-of-fit test. The corresponding results for the USD/JPY residuals are in Table 7.

The L–B test for the residuals gives indication of no autocorrelation on the lower frequencies of one day, twelve and eight hours. Also the one hour frequency of USD/JPY residuals do not reveal serial correlation. The remaining frequencies fail the L–B test. For the absolute values of the residuals we have that only the one hour frequency fails the autocorrelation test for a significance level of 5%. The t goodness-of-fit of the marginal densities, according to the A–D test, is not rejected at all frequencies except for the one and two hours residuals. We also tried to include crossed lagged returns in these models but without success. Hence, the non-inclusion of a model taking into account the crossed lagged dependence is not responsible for the bad fit at the higher frequencies. On the other hand, there is a significant contemporaneous linear correlation between the two residuals time series, see Table 8, and this at all frequencies. The estimated linear correlation turns out to be of the same order as the one obtained for the deseasonalized returns, listed in Table 2. It is this dependence that we want to model using copulae. In the next section we perform a copula analysis of the bivariate residuals or filtered returns $(\hat{\mathbf{z}}_t)$ in (11).

3.2 A static copula model

Though the main aim of this paper concerns a dynamic model for the residual dependence, we first perform a static copula analysis in order to restrict the class of parametric copula families we want to concentrate on.

Figure 4 shows the scatter-plots of residuals $(\hat{\mathbf{z}}_t)$ obtained through the GARCH modelling in Section 3.1. Assuming at first stationarity, suppose that the USD/DEM residuals are represented by the random variable Z_1 and the USD/JPY by the random variable Z_2. Assume that (Z_1, Z_2) has multivariate distribution function F and continuous univariate marginal distribution functions F_1 and F_2. In order to investigate the residual dependence, we fit copula-based models of the type

$$F(z_1, z_2; \boldsymbol{\theta}) = C(F_1(z_1), F_2(z_2); \boldsymbol{\theta}), \tag{12}$$

where C is a copula function, which we know to exist uniquely by Sklar's Theorem (Sklar (1959)), parameterized by the vector $\boldsymbol{\theta} \in \mathbb{R}^q$ with $q \in \mathbb{N}$. The corresponding model density is the product of the copula density c and the marginal densities f_1 and f_2:

$$f(z_1, z_2; \boldsymbol{\theta}) = c(F_1(z_1), F_2(z_2); \boldsymbol{\theta}) \prod_{i=1}^{2} f_i(z_i),$$

where c is the copula density of model (12) and is given by

$$c(u_1, u_2; \boldsymbol{\theta}) = \frac{\partial^2 C(u_1, u_2; \boldsymbol{\theta})}{\partial u_1 \, \partial u_2}, \quad (u_1, u_2) \in [0, 1]^2.$$

Denote by $\mathbf{Z} = \{(Z_{1i}, Z_{2i}) : i = 1, 2, \dots, n\}$ a general random sample of n bivariate observations. The dependence parameter $\boldsymbol{\theta}$ of C is estimated by the pseudo log-likelihood estimator introduced by Genest et al. (1995) where the marginal distribution functions F_i, $i = 1, 2$, are estimated by the rescaled empirical distribution functions $F_{in}(z) = \frac{1}{n+1} \sum_{j=1}^{n} \mathbb{I}_{\{y \in \mathbb{R} : y \leq z\}}(Z_{ij})$. As usual $\mathbb{I}_A$ denotes the indicator function of the set A. After the marginal transformations to so-called pseudo observations $(F_{1n}(Z_{1i}), F_{2n}(Z_{2i}))$ for $i = 1, 2, \dots, n$, the copula family C is fitted. Suppose that its density exists, we then maximize the pseudo log-likelihood function

$$L(\boldsymbol{\theta}; \mathbf{z}) = \sum_{i=1}^{n} \log c(F_{1n}(z_{1i}), F_{2n}(z_{2i}); \boldsymbol{\theta}). \tag{13}$$

The pseudo log-likelihood estimator $\hat{\boldsymbol{\theta}}$ that maximizes (13) is known to be consistent and asymptotically normally distributed when the data are known to be iid; see Genest et al. (1995). Note however that the preliminary ARMA-GARCH filtering may have increased the variance of the estimates $\hat{\boldsymbol{\theta}}$.

The copula families used are: t, Frank, Plackett, Gaussian, Gumbel, Clayton and the mixtures Gumbel with survival Gumbel, Clayton with survival Clayton, Gumbel with Clayton and survival Gumbel with survival Clayton; for details on these classes see Embrechts et al. (2002), Joe (1997) and Nelsen (1999). Denoting the copula family A with parameter $\boldsymbol{\theta}$ by $C^A(\cdot, \cdot; \boldsymbol{\theta})$, the fitted mixtures have distribution functions of the form

$$C(u_1, u_2; \boldsymbol{\theta}) = \theta_3 \, C^A(u_1, u_2; \theta_1) + (1 - \theta_3) \, C^B(u_1, u_2; \theta_2). \tag{14}$$

The above choice of time invariant-copula models is partly based on previous analyses, on tractability and flexibility and also to allow for a fairly broad class with respect to extremal clustering and possible asymmetry. The Gaussian copula is included mainly for comparison.

The models were ranked using Akaike's information criterion

$$\mathrm{AIC} = -2L(\hat{\boldsymbol{\theta}}; \mathbf{x}) + 2q$$

where q is the number of parameters of the family fitted. Parameter estimates and standard errors (s.e.) for all fitted models are listed in Tables 9 and 10.

Overall, the t-copula comes out as the best. The exception is for the daily observations where a Gumbel mixture performs slightly better.

Based on this analysis, in Section 3.3 we will analyse our data with a dynamic t-copula model. Further support for the t-based models is to be found in Breymann et al. (2003), Demarta and McNeil (2004), Daul et al. (2003), Rosenberg and Schuermann (2004) and Pesaran et al. (2004).

This may seem in contrast to Patton (2002) where for the daily data, an asymmetric copula model was used. We performed a likelihood ratio (LR) test for possible asymmetry in a Gumbel mixture model as in (14); we tested for the null hypothesis

$$H_0 : \theta_1 = \theta_2 \qquad \text{and} \qquad \theta_3 = 0.5$$

versus the alternative

$$H_A : \theta_1 \neq \theta_2 \qquad \text{or} \qquad \theta_3 \neq 0.5.$$

A low p-value indicates that a three parameter asymmetric Gumbel mixture model is significantly better than the one model with one parameter symmetric one. This turns out to be the case for the frequencies other than daily and four hours; see Table 11. As the t-copula came out best based on the AIC, we decided to continue our analysis with this copula model.

3.3 A dynamic copula model

In this section we start again from the deseasonalized FX data (x_t) in (7). A dynamic copula model will be achieved through a combination of two univariate ARMA-GARCH models with a time-varying copula. With such a procedure we investigate the constancy of the conditional dependence structure allowing for time-varying dependence parameters and assuming a fixed copula family. As before, we look at several frequencies for the FX data considered. This is a multivariate GARCH-type model; for a survey on multivariate GARCH models, see Bauwens et al. (2003).

Let $(\mathbf{X})_{t\in\mathbb{Z}}$ be a sequence of observable d-dimensional random vectors. Consider the process description given by

$$\begin{aligned} \mathbf{X}_t &= \mathbf{c} + \boldsymbol{\epsilon}_t \\ \boldsymbol{\epsilon}_t &= \boldsymbol{\sigma}_t \mathbf{Z}_t \\ \boldsymbol{\sigma}_t^2 &= A_0 + \sum_{i=1}^{p} A_i \otimes (\boldsymbol{\epsilon}_{t-i}\boldsymbol{\epsilon}_{t-i}^t) + \sum_{j=1}^{q} B_j \otimes \boldsymbol{\sigma}_{t-j}^2 \end{aligned} \tag{15}$$

where A_i for $i = 0, 1, \ldots, p$ and B_j with $j = 1, 2, \ldots, q$ are diagonal $d \times d$ matrices, $\mathbf{c}$ is a vector in $\mathbb{R}^d$ and p and q are positive integers. In (15), $\otimes$ stands for the Hadamard product, the element by element multiplication. Moreover, $(Z_{i,t})_{t\in\mathbb{Z}}$ for $i = 1, 2, \ldots, d$ are assumed to be univariate, strict white noise processes with zero mean and unit variance. The set of equations (15) defines each marginal process as a univariate GARCH. Now we couple the d processes (15) assuming a copula family C for the multivariate distribution of $\mathbf{Z}_t$ with time dependent parameter vector $\boldsymbol{\theta}_t$. The search for suitable dynamics for $\boldsymbol{\theta}_t$ depends very much on the interpretation that a specific dependence parameter may have; see Section 3.3.2. Patton (2002) and Rockinger and Jondeau (2001) contain specific examples for the Gaussian, the symmetrised Joe-Clayton and Plackett copula. A further interesting paper is Cappiello et al. (2003).

3.3.1 Model estimation: general theory

The natural estimation method for (15) is (conditional) maximum likelihood. Furthermore, the definition of the model suggests a two step estimation procedure. In fact, this is used in similar situations in Engle and Sheppard (2001), Patton (2002) and Rockinger and Jondeau (2001).

For a random sample $(\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_n)$, assume that the conditional distribution of $\mathbf{X}_t$ can be written as

$$F(\mathbf{x}; \boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_d, \boldsymbol{\theta}_t) = C(F_1(x_1; \boldsymbol{\alpha}_1), \ldots, F_d(x_d; \boldsymbol{\alpha}_d); \boldsymbol{\theta}_t)$$

where we assume that the F_i's are absolutely continuous with density f_i; the vectors $\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_d$ parameterize the marginal distribution functions; the time-varying parameter $\boldsymbol{\theta}_t$ parameterizes the copula family. Assuming that C has density c, then the conditional density of $\mathbf{X}_t$ is given by:

$$\begin{aligned} f(\mathbf{x}; \boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_d, \boldsymbol{\theta}_t) &= \\ &= c(F_1(x_1; \boldsymbol{\alpha}_1), \ldots, F_d(x_d; \boldsymbol{\alpha}_d); \boldsymbol{\theta}_t) \prod_{i=1}^{d} f_i(x_i; \boldsymbol{\alpha}_i). \end{aligned}$$

The conditional log-likelihood function of the model therefore is

$$\sum_{t=m+1}^{n} \left(\log c(F_1(x_{1,t}; \boldsymbol{\alpha}_1), \ldots, F_d(x_{d,t}; \boldsymbol{\alpha}_d); \boldsymbol{\theta}_t) + \sum_{i=1}^{d} \log f_i(x_{i,t}; \boldsymbol{\alpha}_i) \right), \qquad (16)$$

where $m = \max(p, r)$. Numerical maximization of (16) gives the maximum likelihood estimates of the model. However, the optimization of the likelihood function with possibly many parameters is numerically difficult and time consuming. It is more tractable to estimate first the marginal model parameters and then the dependence model parameters using the estimates from the first step. In order to do so, the d marginal likelihood functions

$$\sum_{t=p+1}^{n} \log f_i(x_{i,t}; \boldsymbol{\alpha}_i), \qquad i = 1, 2, \ldots, d, \qquad (17)$$

are independently maximised, yielding the estimates $\hat{\boldsymbol{\alpha}}_1, \ldots, \hat{\boldsymbol{\alpha}}_d$. The final function to maximise then is

$$\sum_{t=m+1}^{n} \log c(F_1(x_{1,t}; \hat{\boldsymbol{\alpha}}_1), \ldots, F_d(x_{d,t}; \hat{\boldsymbol{\alpha}}_d); \boldsymbol{\theta}_t). \tag{18}$$

From this, estimates for the dependence parameter $\boldsymbol{\theta}_t$ are obtained. In Section 3.3.2 below we apply the above procedure to our FX data.

3.3.2 Fitting the time-varying copula model to the FX returns

For the USD/DEM and USD/JPY spot rate returns, in Section 3.2 we found that the t-copula yields the best model for the cross dependence. For the t-copula, $\boldsymbol{\theta}_t = (\nu_t, \rho_t)'$ with possibly time-varying degrees of freedom ν_t and correlation coefficient ρ_t. In a first attempt, we take as a parsimonious model:

$$\begin{aligned} \nu_t &= \nu \qquad \text{for all } t, \\ \rho_t &= h^{-1}(r_0 + r_1 z_{1,t-1} z_{2,t-1} + s_1 h(\rho_{t-1})), \end{aligned} \tag{19}$$

where $h(\cdot)$ is Fisher's transformation

$$h(\rho) = \log\left(\frac{1+\rho}{1-\rho}\right), \qquad -1 < \rho < 1.$$

This choice of model is motivated by the fact that in first instance we want to model the dependence structure of the data and hence concentrate on the dynamics of the correlation function ρ_t; a generalisation of the procedure to time-varying ν_t is definitely possible. The regression-type dynamics in the z_i's in (19) have as a natural consequence that when both innovations have the same sign we have a positive contribution in (19). In case innovations have opposite signs, a negative contribution results.

3.3.3 Dynamic copula estimation: the results

From the marginal ARMA-GARCH models from Section 3.1, we obtain the estimated degrees of freedom $\hat{\nu}$; see Table 5. Denote by $\{(\hat{z}_{1,t}, \hat{z}_{2,t}) : t = 0, \ldots, n\}$ the standardised residual return series obtained from this GARCH filtering. The standardised residuals

$$\left(\sqrt{\frac{\hat{\nu}_i}{\hat{\nu}_i - 2}}\,\hat{z}_{i,t}\right), \qquad i = 1, 2,$$

are mapped into the unit square by the standard t probability-integral transformation, yielding the pseudo-observations:

$$\left\{\left(t_{\hat{\nu}_1}\left(\sqrt{\frac{\hat{\nu}_1}{\hat{\nu}_1 - 2}}\,\hat{z}_{1,t}\right), t_{\hat{\nu}_2}\left(\sqrt{\frac{\hat{\nu}_2}{\hat{\nu}_2 - 2}}\,\hat{z}_{2,t}\right)\right) : t = 1, \dots, n\right\}. \tag{20}$$

These data are plugged into (18) with c being the t-copula density function and using (19) for the dynamics of the dependence parameters. Maximisation yields $\hat{\boldsymbol{\theta}} = (\hat{\nu}, \hat{\rho}_t)'$ For the one and two hours frequencies we used the empirical distribution function instead of the t-distribution. The reason was the poor results obtained for the goodness-of-fit test reported in Tables 6 and 7. The results are reported in Table 12. We added the t-copula parameter estimates from the static copula model.

The AIC of the time-varying copula model is lower than the AIC of the static copula model; see Table 12, hence we have an improvement in the goodness-of-fit. The estimate for r_0 in (19) can be considered zero for eight hours, twelve hours and daily returns. But r_1 and s_1 are definitely different from zero for all frequencies. In other words, the estimated (copula) correlation depends on the marginal returns and on the correlation from the previous moment in time. From the estimated parameters for the correlation dynamics we compute, through the second equation of (19), the time-varying estimated correlation which is plotted in Figures 5 and 6. These also show the estimated constant correlation with a 95% confidence interval (see Section 3.2). Given these results we can infer that the conditional dependence is definitely time-varying and so its dynamics has to be modelled.

The degrees of freedom estimated for the dependence structure is always larger for the dynamic copula model than for the static one; compare the values listed in Table 12. We observe an increasing pattern in the estimated degrees of freedom from higher to lower frequencies.

4 A change-point model

Looking at Figures 5 and 6, the implied correlation function estimated in Section 3.3 shows possible level changes. We analyse this issue though a change-point detection technique for parametric copula models. There are several well known tests on structural breaks in econometric time-series analysis; see for instance Bai (1997), Bai and Perron (1998) and Hansen (2001). See also Polzehl and Spokoiny (2004). Here, we test for changes in the copula parameters, estimate the size of those changes and the corresponding time of occurrence. For related work, see for instance Gombay and Horváth (1999). For a detailed treatment of the change-point theory underlying our approach, see Csörgő and Horváth (1997) and references therein, and Dias and Embrechts (2002).

4.1 Detecting change-points in copula parameters

Let $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_n$ be a sequence of independent random vectors in $[0,1]^d$ with univariate uniformly distributed margins and copulae $C(\mathbf{u}; \boldsymbol{\theta}_1, \boldsymbol{\eta}_1)$, $C(\mathbf{u}; \boldsymbol{\theta}_2, \boldsymbol{\eta}_2), \ldots, C(\mathbf{u}; \boldsymbol{\theta}_n, \boldsymbol{\eta}_n)$ respectively, where $\boldsymbol{\theta}_i$ and $\boldsymbol{\eta}_i$ are the copula parameters such that $\boldsymbol{\theta}_i \in \Theta^{(1)} \subseteq \mathbb{R}^p$ and $\boldsymbol{\eta}_i \in \Theta^{(2)} \subseteq \mathbb{R}^q$. We will consider the $\boldsymbol{\eta}_i$ as nuisance parameters and look for one single change-point in $\boldsymbol{\theta}_i$. Formally, we test the null hypothesis

$$H_0 : \boldsymbol{\theta}_1 = \boldsymbol{\theta}_2 = \ldots = \boldsymbol{\theta}_n \qquad \text{and} \qquad \boldsymbol{\eta}_1 = \boldsymbol{\eta}_2 = \ldots = \boldsymbol{\eta}_n$$

versus the alternative

$$H_A : \boldsymbol{\theta}_1 = \ldots = \boldsymbol{\theta}_{k^*} \neq \boldsymbol{\theta}_{k^*+1} = \ldots = \boldsymbol{\theta}_n \qquad \text{and} \qquad \boldsymbol{\eta}_1 = \boldsymbol{\eta}_2 = \ldots = \boldsymbol{\eta}_n.$$

If we reject the null hypothesis, k^* is the time of the change. All the parameters of the model are supposed to be unknown under both hypotheses. If $k^* = k$ were known, the null hypothesis would be rejected for small values of the likelihood ratio

$$\Lambda_k = \frac{\sup_{(\boldsymbol{\theta},\boldsymbol{\eta})\in\Theta^{(1)}\times\Theta^{(2)}} \prod_{1\le i\le n} c(\mathbf{u}_i;\boldsymbol{\theta},\boldsymbol{\eta})}{\sup_{(\boldsymbol{\theta},\boldsymbol{\theta}',\boldsymbol{\eta})\in\Theta^{(1)}\times\Theta^{(1)}\times\Theta^{(2)}} \prod_{1\le i\le k} c(\mathbf{u}_i;\boldsymbol{\theta},\boldsymbol{\eta}) \prod_{k<i\le n} c(\mathbf{u}_i;\boldsymbol{\theta}',\boldsymbol{\eta})}, \tag{21}$$

where we assume that C has a density c. The estimation of Λ_k is carried out through maximum likelihood and so all the necessary conditions of regularity and efficiency have to be assumed (see Lehmann and Casella (1998)).

Denote

$$L_k(\boldsymbol{\theta},\boldsymbol{\eta}) = \sum_{1\le i\le k} \log c(\mathbf{u}_i;\boldsymbol{\theta},\boldsymbol{\eta})$$

and

$$L_k^*(\boldsymbol{\theta},\boldsymbol{\eta}) = \sum_{k<i\le n} \log c(\mathbf{u}_i;\boldsymbol{\theta},\boldsymbol{\eta}).$$

Then the likelihood ratio equation can be written as

$$-2\log(\Lambda_k) = 2\left(L_k(\hat{\boldsymbol{\theta}}_k,\hat{\boldsymbol{\eta}}_k) + L_k^*(\boldsymbol{\theta}_k^*,\hat{\boldsymbol{\eta}}_k) - L_n(\hat{\boldsymbol{\theta}}_n,\hat{\boldsymbol{\eta}}_n)\right).$$

As k is unknown, H_0 will be rejected for large values of

$$Z_n = \max_{1\le k<n} (-2\log(\Lambda_k)). \tag{22}$$

4.1.1 Asymptotic critical values

The asymptotic distribution of $Z_n^{1/2}$ is known but has a very slow rate of convergence; see Csörgő and Horváth (1997), page 22. In the same reference we can also find an approximation for the distribution of $Z_n^{1/2}$ derived to give better small sample rejection regions. Indeed, for $0 < h < l < 1$, the following approximation holds:

$$P\left(Z_n^{1/2} \ge x\right) \approx \frac{x^p \exp(-x^2/2)}{2^{p/2}\Gamma(p/2)}. \left(\log\frac{(1-h)(1-l)}{hl} - \frac{p}{x^2}\log\frac{(1-h)(1-l)}{hl} + \frac{4}{x^2} + O\left(\frac{1}{x^4}\right)\right), \tag{23}$$

as $x \to \infty$ and where h and l can be taken as $h(n) = l(n) = (\log n)^{3/2}/n$. Note that in (23) p is the number of parameters that may change under the alternative. This result turns out to be very accurate as shown in a simulation study in Dias and Embrechts (2002) where it is applied to the Gumbel copula.

4.1.2 The time of the change

If we assume that there is exactly one change-point, then the maximum likelihood estimator for the time of the change is given by

$$\hat{k}_n = \min\{1 \leq k < n : Z_n = -2\log(\Lambda_k)\}. \tag{24}$$

In the case that there is no change, $\hat{k}_n$ will take a value near the boundaries of the sample. This holds because under the null hypothesis, and given that all the necessary regularity conditions hold, for $n \to \infty$, $\hat{k}_n/n \xrightarrow{d} \xi$, where $P(\xi = 0) = P(\xi = 1) = 1/2$; see Csörgő and Horváth (1997), page 51. This behavior was verified in a simulation study for the Gumbel copula under the no-change hypothesis in Dias and Embrechts (2002).

4.1.3 Multiple Changes

The detection of several change-points in multidimensional processes with unknown parameters can be done using the so called binary segmentation procedure. This method was proposed by Vostrikova (1981) and allows us to simultaneously detect the number and the location of the change-points. The method consists of first applying the likelihood ratio test for one change. If H_0 is rejected then we have the estimate of the time of the change $\hat{k}_n$. Next, we divide the sample in two subsamples $\{\mathbf{u}_i : 1 \leq i \leq \hat{k}_n\}$ and $\{\mathbf{u}_i : \hat{k}_n < i \leq n\}$ and test H_0 for each one of them. If we find a change in any of the sets we continue this segmentation procedure until we don't reject H_0 in any of the subsamples. An alternative method is discussed in Mercurio and Spokoiny (2004).

4.2 An application to the FX data

In this section we test for the occurrence of change-points in the copula parameters of the deseasonalized FX filtered returns. Concretely we use the procedures from the previous section to estimate change-points in the correlation parameter of a constant parameter t-copula fitted to the residuals of daily USD/DEM and USD/JPY returns. A similar analysis can be done at other frequencies. For the change-points found, we estimate the size of those changes and the corresponding time of occurrence. We also look at macro economic reasons possibly triggering the changes.

After filtering the univariate returns using the GARCH type models specified in Table 5 of Section 3.1, we start with the assumption that the residuals do not depend on time. Hence, we can use (21) and (23) for detecting possible change-points in the parameters of the multivariate contemporaneous conditional distribution and in particular in the copula.

For the copula fitting, we use the empirical distribution function to map the residuals into the unit square. In a first step, we assume that the degrees of freedom of the copula are constant over time and hence we only test for change-points in the correlation parameter. We evaluate Λ_k for $k = 1, 2, \ldots, n$ where $n = 3\,259$; see (21). The values obtained are displayed in the top panel of Figure 7. The test statistic (21) takes the value $z_{n\,obs}^{1/2} = 13.26$ and by (23) we have that $P(Z_n^{1/2} > 13.26) \approx 0$. The null hypothesis of no change-point is to be rejected and the estimated time of the change is $\hat{k}_n =$ November 8, 1989; coinciding with the fall of the Berlin wall. In the next step, the sample is divided in two sub-samples, one up to November 8, 1989 and another from the estimated time of change onwards. For each sub-sample Λ_k is computed as well as $Z_n^{1/2}$. The middle panel of Figure 7 plots these estimates and Table 13 has the values for $z_{n\,obs}^{1/2}$ and all the information about the testing procedure. As the obtained p-values are close to zero we reject the null hypothesis of no change for each sub-sample and estimate two more times of change, December 29, 1986 and June 18, 1997. The later date corresponds to the beginning of the Asia crisis starting with the devaluation of the Thai Baht. Each sub-sample is again divided in two and the procedure is repeated yielding the estimates in the bottom panel of Figure 7.

For the results of the analysis shown in the bottom panel of Figure 7 only for the maximum attained at October 23, 1990 the null hypothesis is rejected at a 5% level. So we still have to split this sub-sample further. The first from November 8, 1989 until October 23, 1990 and the second from this date up to June 18, 1997. The $z_{n\,obs}^{1/2}$ obtained in these cases are low, see Table 13, and we do not reject the null hypothesis of no change in both cases. In Table 13 we give the time where the test statistic (22) is attained for each sub-sample. If the null hypothesis is not rejected the date found is put in parentheses. In summary we found four change-points: December 29, 1986, November 8, 1989, October 23, 1990 and June 18, 1997. For the five periods between the times of change we estimated for the copula correlation ($\hat{s.e.}$): $\hat{\rho}_1 = 0.6513$ (0.0384), $\hat{\rho}_2 = 0.8312$ (0.0113), $\hat{\rho}_3 = 0.3099$ (0.0608), $\hat{\rho}_4 = 0.5752$ (0.0149) and $\hat{\rho}_5 = 0.3505$ (0.0460). To visualise these results we redisplay in Figure 8 the time-varying correlation estimated in Section 3.3, first plotted in the bottom panel of Figure 6 (daily data), super-imposed with the estimated change-point cross-correlation for the five periods between the times of change. The change-point analysis seems to detect the main features of the changes in the dynamic correlation curve. No change-points were detected from October 23, 1990 until June 18, 1997 which may seem a rather long period for the correlation to be constant. It may be interesting to note that the former date (October 23, 1990) corresponds to the burst in the Japanese asset price bubble. On October 18, 1990, the USD/JPY ended a fall from about 158 to 125.

We can compare the AIC value for the three models discussed: static copula (Section 3.2), dynamic copula (Section 3.3) and change-points (Section 4). The values obtained are:

	non-dynamic	dynamic	change-points
AIC	-1644.549	-1881.760	-1925.482

The t-copula model with the change-points in the correlation parameter yields a superior fit.

We also tested the existence of change points in the parameters of the distribution of the univariate residuals. If the marginal models are appropriate then the residuals should come from a t-distribution with constant parameters (location, scale and degrees of freedom). The null hypothesis corresponds to no change points in any of the three parameters. For the daily USD/DEM we obtain a test statistic value of $z^{1/2}_{n\,obs} = 3.195$ to which corresponds a p-value of 42% according to (23). In the case of the USD/JPY residuals we have that $z^{1/2}_{n\,obs} = 3.915$ which implies a p-value of 9.2%.

The above analysis were partly repeated at the higher frequencies yielding finer estimates of the possible change points. In particular, we put some data windows around the change point found in the daily data and then proceeded as above. We shall come later to some of these findings in further work. We want to stress at this point that, for instance at the higher frequency of one hour, the computational complexity for the full change point procedure is considerable.

5 Conclusion

The aim of the paper is essentially twofold. First of all, we want to contribute to the ongoing discussion between practitioners and academics in order to advance the methodological basis for risk measurement technology. No doubt, thinking beyond linear correlation, and this through the notion of copula, contributes to this goal. Second, through the example of the high-frequency, two-dimensional FX data, we were able to come up with a parsimonious conditional dependence model taking changes in the dependence structure into account. The change points found relate to specific macro-economic events. We have not made the next step: undertaking the full economic reasons underlying the final model presented. We do however hope that the tools introduced will eventually lead to such, more economic studies.

Acknowledgement

We would like to thank Wolfgang Breymann and Olsen Data for providing the FX data. This paper was written while the second author was Centennial Professor of Finance at the London School of Economics. We also acknowledge useful discussions with Alexander McNeil and La Fischer.

We thank the referees for the detailed reports which permitted to improve this work.

References

Andersen, T. G., Bollerslev, T., Christoffersen, P. F., and Diebold, F. X. (2004). Practical volatility and correlation modeling for financial market risk management. Working paper, `http://www.nber.org/~confer/2004/Riskf04/andersen.pdf`.

Bai, J. (1997). Estimating multiple breaks one at a time. *Econometric Theory*, 13:551–563.

Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 66:47–78.

Bauwens, L., Laurent, S., and Rombouts, J. V. K. (2003). Multivariate garch models: a survey. Discussion paper, CREST, Laboratoire de Finance et Assurance and CORE, Université catholique de Louvain.

Bollerslev, T., Chou, R. Y., and Kroner, K. (1992). ARCH modeling in finance. *J. Econometrics*, 52:5–59.

Boyer, B. H., Gibson, M. S., and Loretan, M. (1999). Pitfalls in tests for changes in correlations. *Board of Governors of the Federal Reserve System, International Finance Discussion Papers*, 597.

Breymann, W., Dias, A., and Embrechts, P. (2003). Dependence structures for multivariate high-frequency data in finance. *Quant. Finance*, 3:1–14.

Brockwell, P. J. and Davis, R. A. (1991). *Time Series: Theory and Methods.* Springer Series in Statistics. Springer-Verlag, New York, second edition.

Cappiello, L., Engle, R., and Sheppard, K. (2003). Asymmetric dynamics in the correlations of global equity and bond returns. Working paper 204, European Central Bank.

Chen, X. and Fan, Y. (2002). Estimation of copula-based semiparametric time series models. Working paper, New York University and Vanderbilt University.

Cherubini, U., Luciano, E., and Vecchiato, W. (2004). *Copula Methods in Finance.* Wiley, Chichester.

Choi, K. and Zivot, E. (2003). Long memory and structural changes in the forward discount: An empirical investigation. Preprint, Ohio University and University of Washington.

Csörgő, M. and Horváth, L. (1997). *Limit Theorems in Change-Point Analysis.* Wiley, Chichester.

Dacorogna, M. M., Gençay, R., Müller, U. A., Olsen, R. B., and Pictet, O. V. (2001). *An Introduction to High-Frequency Finance.* Academic Press, San Diego, CA.

Daul, S., Giorgi, E. D., Lindskog, F., and McNeil, A. (2003). Using the grouped t-copula. *RISK Magazine*, pages 73–76.

Demarta, S. and McNeil, A. J. (2004). The t copula and related copulas. Working paper, ETH Zurich, Department of Mathematics.

Dias, A. and Embrechts, P. (2002). Change-point analysis for dependence structures in finance and insurance. In: *Novos Rumos em Estatística* (Ed. C. Carvalho, F. Brilhante and F. Rosado), Sociedade Portuguesa de Estatística, Lisbon, pages 69–86. Also in: *New Risk Measures in Investment and Regulation* (Ed. G. Szego), John Wiley and Sons, New York (2003).

Ding, Z., Granger, C. W. J., and Engle, R. F. (1993). A long memory property of stock market returns and a new model. *J. Empirical Finance*, 1:83–106.

Embrechts, P., McNeil, A. J., and Straumann, D. (1999). Correlation: pitfalls and alternatives. *RISK, May*, pages 69–71.

Embrechts, P., McNeil, A. J., and Straumann, D. (2002). Correlation and dependence in risk management: Properties and pitfalls. In: *Risk Management: Value at Risk and Beyond* (Ed. M. Dempster), Cambridge University Press, Cambridge, pages 176–223.

Engle, R. and Sheppard, K. (2001). Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. Working paper 2001–15, UCSD.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4):987–1007.

Fortin, I. and Kuzmics, C. (2002). Tail dependence in stock return-pairs: Towards testing ellipticity. Working paper, Institute For Advanced Studies, Vienna and Faculty of Economics and Politics, University of Cambridge.

Francis, N. and Owyang, M. T. (2004). Monetary Policy in a Markov-Switching VECM: Implications for the Cost of Disinflation and the Price Puzzle. Working paper, Federal Reserve Bank of St. Louis.

Genest, C., Ghoudi, K., and Rivest, L.-P. (1995). A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika*, 82(3):543–552.

Gombay, E. and Horváth, L. (1999). Change-points and bootstrap. *Environmetrics*, 10:725–736.

Hamilton, J. D. (1990). Analysis of time series subject to changes in regime. *J. Econometrics*, 45(1-2):39–70.

Hansen, B. E. (2001). The new econometrics of structural change: Dating breaks in U.S. Labor Productivity. *Journal of Economic Perspectives*, 15:117–128.

Jarque, C. M. and Bera, A. K. (1987). A test for normality of observations and regression residuals. *Internat. Statist. Rev.*, 55(2):163–172.

Joe, H. (1997). *Multivariate Models and Dependence Concepts.* Chapman & Hall, London.

Lehmann, E. L. and Casella, G. (1998). *Theory of Point Estimation.* Springer Texts in Statistics. Springer-Verlag, New York, second edition.

Longin, F. and Solnik, B. (2001). Extreme correlation of international equity markets. *J. Finance*, LVI(2):649–676.

Loretan, M. and Phillips, P. C. B. (1994). Testing the covariance stationarity of heavy-tailed time series: An overview of the theory with applications to several financial data sets. *J. Empirical Finance*, 2:211–248.

Mercurio, D. and Spokoiny, V. (2004). Estimation of time dependent volatility via local change point analysis. Preprint, no. 904, Weierstrass Institute for Applied Analysis and Stochastics.

Nelsen, R. B. (1999). *An Introduction to Copulas*, volume 139 of *Lecture Notes in Statistics*. Springer-Verlag, New York.

Patton, A. J. (2002). Modelling time-varying exchange rate dependence using the conditional copula. Working paper, UCSD.

Pesaran, M. H., Schuermann, T., and Weiner, S. M. (2004). Modeling regional interdependencies using a global error-correcting macroeconometric model. *Journal of Business Economics and Statistics*, 22(2):129–162.

Polzehl, J. and Spokoiny, V. (2004). Varying coefficient GARCH versus local constant volatility modeling. Comparison of the predictive power. Preprint, no. 977, Weierstrass Institute for Applied Analysis and Stochastics.

Rockinger, M. and Jondeau, E. (2001). Conditional dependency of financial series: An application of copulas. Working paper, HEC-School of Management, Departement of Finance.

Rosenberg, J. V. and Schuermann, T. (2004). A general approach to integrated risk management with skewed, fat-tailed risks. Staff report, Federal Reserve Bank of New York.

Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris*, 8:229–231.

Vostrikova, L. J. (1981). Detecting "disorder" in multidimensional random processes. *Soviet Mathematics Doklady*, 24(1):55–59.

Zivot, E. and Wang, J. (2003). *Modeling Financial Time Series with S–Plus.* Springer-Verlag, New York.

6 Tables and figures

	USD/DEM		USD/JPY	
Frequency	Mean	Std. dev.	Mean	Std. dev.
1 hour	0.0031	0.9539	0.0001	0.9521
2 hours	0.0030	0.9724	-0.0010	0.9585
4 hours	-0.0005	0.9877	-0.0050	0.9509
8 hours	-0.0033	0.9878	-0.0092	0.9611
12 hours	-0.0039	0.9767	-0.0100	0.9457
1 day	-0.0112	1.0182	-0.0128	1.0026

Table 1: Summary statistics of the USD/DEM and USD/JPY deseasonalized returns at the six frequencies.

	USD/DEM		USD/JPY		Linear
Frequency	Skewness	Kurtosis	Skewness	Kurtosis	correlation
1 hour	-0.0113	5.9903	-0.1831	7.3674	0.5604
2 hours	0.0615	6.9665	-0.1058	7.7071	0.5858
4 hours	-0.0035	5.5195	-0.1855	5.4950	0.5973
8 hours	0.0026	3.9248	-0.2718	5.0043	0.6112
12 hours	-0.1036	2.4768	-0.3448	3.5822	0.6186
1 day	-0.1234	1.6838	-0.3171	2.7735	0.6243

Table 2: Summary statistics of the USD/DEM and USD/JPY deseasonalized returns at the six frequencies.

	USD/DEM		USD/JPY	
Frequency	Test statistic	P-value	Test statistic	p-value
1 hour	114 785.4	0.0	174 000.0	0.0
2 hours	78 207.3	0.0	95 757.0	0.0
4 hours	24 619.7	0.0	24 551.5	0.0
8 hours	6 246.1	0.0	10 276.3	0.0
12 hours	1 668.0	0.0	3 594.3	0.0
1 day	391.3	0.0	1 094.2	0.0

Table 3: Jarque–Bera test statistic values and the p-values for the USD/DEM and USD/JPY deseasonalized returns at the six frequencies.

	USD/DEM		USD/JPY	
Frequency	Test statistic	P-value	Test statistic	p-value
1 hour	2 893.13	0.0	3 083.30	0.0
2 hours	857.24	0.0	1 178.70	0.0
4 hours	578.51	0.0	614.08	0.0
8 hours	190.61	0.0	336.04	0.0
12 hours	117.80	0.0	160.43	0.0
1 day	35.25	0.0	75.99	0.0

Table 4: ARCH effects test statistic values and the p-values for the USD/DEM and USD/JPY deseasonalized returns at six frequencies.

USD/DEM

Frequency	p_1	q_1	p_2	q_2	$\hat{\nu}$ ($\hat{s.e.}$)
1 hour	-	-	1	1	3.693 (0.054)
2 hours	2	2	2	1	3.708 (0.044)
4 hours	-	5	1	1	3.975 (0.105)
8 hours	2	4	1	1	4.679 (0.234)
12 hours	1	-	1	1	5.385 (0.326)
1 day	1	-	1	1	5.797 (0.556)

USD/JPY

Frequency	p_1	q_1	p_2	q_2	$\hat{\nu}$ ($\hat{s.e.}$)
1 hour	-	-	1	1	3.654 (0.052)
2 hours	1	-	2	1	3.759 (0.077)
4 hours	4	4	2^*	1	3.819 (0.109)
8 hours	2	2	1^*	1	4.357 (0.195)
12 hours	1	-	1^*	1	4.574 (0.251)
1 day	10	-	1^*	1	4.889 (0.412)

Table 5: Order of the ARMA(p_1, q_1)-GARCH(p_2, q_2) models fitted to the USD/DEM and USD/JPY returns at the several frequencies. Degrees of freedom estimated for the marginal conditional distribution t of the innovations and corresponding standard errors are also given.

	L–B test		A–D
Frequency	z_t	$\|z_t\|$	test
1 hour	0.0002	0.0000	0.0000
2 hours	0.0000	0.0911	0.0014
4 hours	0.0000	0.0836	0.3137
8 hours	0.4107	0.8764	0.1904
12 hours	0.3917	0.3114	0.5189
1 day	0.1204	0.8456	0.6452

Table 6: Autocorrelation and density goodness-of-fit p-values for the USD/DEM residuals at the six frequencies.

	L–B test		A–D
Frequency	z_t	$\|z_t\|$	test
1 hour	0.0795	0.0000	0.0000
2 hours	0.0000	0.0478	0.0100
4 hours	0.0046	0.3618	0.0832
8 hours	0.1702	0.2291	0.1640
12 hours	0.1423	0.9169	0.0953
1 day	0.5328	0.6803	0.3105

Table 7: Autocorrelation and density goodness-of-fit p-values for the USD/JPY residuals at the six frequencies.

Frequency	Linear correlation
1 hour	0.5441
2 hours	0.5723
4 hours	0.5836
8 hours	0.6018
12 hours	0.6139
1 day	0.6180

Table 8: Linear correlation of the USD/DEM and USD/JPY bivariate residuals at the six frequencies.

Frequency	Copula model	$\hat{\theta}$ (s.e.)	AIC
1 hour	Clayton	0.859 (0.006)	-23401.101
	Frank	3.979 (0.024)	-27032.306
	Gaussian	0.550 (0.002)	-28267.108
	Gumbel	1.562 (0.004)	-28146.727
	Plackett	6.503 (0.061)	-29324.002
2 hour	Clayton	0.913 (0.009)	-12730.906
	Frank	4.200 (0.035)	-14806.051
	Gaussian	0.571 (0.002)	-15483.028
	Gumbel	1.605 (0.006)	-15506.480
	Plackett	7.038 (0.093)	-16020.855
4 hour	Clayton	0.944 (0.013)	-6652.147
	Frank	4.341 (0.050)	-7821.259
	Gaussian	0.584 (0.004)	-8176.122
	Gumbel	1.634 (0.009)	-8251.189
	Plackett	7.361 (0.137)	-8446.380
8 hour	Clayton	0.984 (0.019)	-3536.107
	Frank	4.563 (0.072)	-4260.632
	Gaussian	0.603 (0.005)	-4413.271
	Gumbel	1.669 (0.013)	-4412.292
	Plackett	7.752 (0.201)	-4533.552
12 hour	Clayton	1.025 (0.024)	-2487.922
	Frank	4.659 (0.088)	-2941.280
	Gaussian	0.615 (0.006)	-3092.874
	Gumbel	1.681 (0.016)	-3007.219
	Plackett	7.949 (0.252)	-3113.152
1 day	Clayton	1.034 (0.035)	-1252.289
	Frank	4.599 (0.124)	-1446.464
	Gaussian	0.617 (0.009)	-1552.695
	Gumbel	1.679 (0.023)	-1500.065
	Plackett	7.772 (0.350)	-1526.993

Table 9: Residuals on USD/DEM and USD/JPY log-returns. Estimates and standard errors of dependence parameters in Clayton, Frank, Gaussian, Gumbel and Plackett models. For each model fitted we provide the AIC value. The reading of this table must be complemented with Table 10.

Freq.	Copula model	$\hat{\theta}_1$ (s.e.)	$\hat{\theta}_2$ (s.e.)	$\hat{\theta}_3$ (s.e.)	AIC
	Cl & s. Cl	1.125 (0.025)	1.171 (0.027)	0.516 (0.007)	-29924.80
	Cl & Gumbel	1.568 (0.014)	1.363 (0.066)	0.659 (0.009)	-30642.91
1 hour	s.Cl & s.Gum	1.552 (0.010)	1.510 (0.065)	0.701 (0.008)	-30665.31
	Gum & s.Gum	2.038 (0.030)	1.405 (0.009)	0.421 (0.010)	-31061.35
	t	4.935 (0.108)	0.558 (0.002)	-	-31517.70
	Cl & s. Cl	1.164 (0.033)	1.316 (0.041)	0.517 (0.010)	-16430.36
	Cl & Gumbel	1.674 (0.034)	1.233 (0.114)	0.650 (0.014)	-16803.66
2 hour	s.Cl & s.Gum	1.576 (0.013)	1.723 (0.086)	0.695 (0.011)	-16801.39
	Gum & s.Gum	2.109 (0.039)	1.420 (0.013)	0.441 (0.013)	-17015.86
	t	4.822 (0.147)	0.580 (0.003)	-	-17192.73
	Cl & s. Cl	1.238 (0.048)	1.325 (0.051)	0.499 (0.014)	-8653.704
	Cl & Gumbel	1.682 (0.032)	1.359 (0.128)	0.674 (0.017)	-8863.199
4 hour	s.Cl & s.Gum	1.640 (0.024)	1.535 (0.115)	0.669 (0.016)	-8847.032
	Gum & s.Gum	1.501 (0.028)	1.991 (0.064)	0.545 (0.020)	-8932.445
	t	4.748 (0.201)	0.593 (0.004)	-	-9088.884
	Cl & s. Cl	1.265 (0.060)	1.472 (0.071)	0.502 (0.018)	-4607.028
	Cl & Gumbel	1.771 (0.053)	1.265 (0.170)	0.667 (0.024)	-4722.141
8 hour	s.Cl & s.Gum	1.663 (0.028)	1.710 (0.140)	0.668 (0.021)	-4713.239
	Gum & s.Gum	1.991 (0.072)	1.534 (0.040)	0.496 (0.027)	-4764.398
	t	5.323 (0.343)	0.612 (0.006)	-	-4818.328
	Cl & s. Cl	1.492 (0.095)	1.286 (0.084)	0.503 (0.024)	-3157.357
	Cl & Gumbel	1.653 (0.031)	1.893 (0.184)	0.679 (0.025)	-3242.862
12 hour	s.Cl & s.Gum	1.787 (0.060)	1.307 (0.206)	0.673 (0.030)	-3248.558
	Gum & s.Gum	1.556 (0.046)	2.018 (0.088)	0.511 (0.033)	-3281.252
	t	5.837 (0.505)	0.621 (0.007)	-	-3304.250
	Cl & s. Cl	1.548 (0.120)	1.280 (0.099)	0.494 (0.032)	-1599.798
	Cl & Gumbel	1.665 (0.045)	1.844 (0.249)	0.671 (0.037)	-1629.394
1 day	s.Cl & s.Gum	1.816 (0.071)	1.234 (0.195)	0.656 (0.039)	-1632.435
	Gum & s.Gum	1.588 (0.072)	1.952 (0.117)	0.501 (0.048)	-1642.460
	t	6.012 (0.786)	0.620 (0.010)	-	-1640.061

Table 10: Residuals on USD/DEM and USD/JPY log-returns. Estimates and standard errors of parameters for the t-copula and for the four mixture copulae considered. In case of the mixture copulae, θ_1 and θ_2 are the dependence parameters respectively for the first and second terms of the mixture. θ_3 is the mixture parameter which gives the proportion of the first term. For the t-copula, θ_1 are the degrees of freedom and θ_2 is the correlation. For each model fitted we provide the AIC. The reading of this table must be complemented with Table 9.

Frequency	p-values of the LR test for the Gumbel mixture
1 hour	0.0000
2 hour	0.0000
4 hour	0.1371
8 hour	0.0048
12 hour	0.0164
1 day	0.1842

Table 11: p-values for a likelihood ratio test for the Gumbel and survival Gumbel mixture model with three versus one parameter.

Time frequency	Parameter Estimates ($\hat{s.e.}$) non-dynamic		dynamic	
1 hour	$\hat{\nu}$	4.935 (0.108)	$\hat{\nu}$	6.330 (0.167)
	$\hat{\rho}$	0.558 (0.002)	$\hat{r}_0$	0.0005 (0.0002)
			$\hat{r}_1$	0.0193 (0.0010)
			$\hat{s}_1$	0.9921 (0.0005)
	AIC	-31517.70	AIC	-34488.72
2 hours	$\hat{\nu}$	4.822 (0.147)	$\hat{\nu}$	6.203 (0.230)
	$\hat{\rho}$	0.580 (0.003)	$\hat{r}_0$	-0.0004 (0.0002)
			$\hat{r}_1$	0.0128 (0.0009)
			$\hat{s}_1$	0.9952 (0.0004)
	AIC	-17192.73	AIC	-19349.29
4 hours	$\hat{\nu}$	4.669 (0.195)	$\hat{\nu}$	6.072 (0.313)
	$\hat{\rho}$	0.592 (0.005)	$\hat{r}_0$	-0.0008 (0.0002)
			$\hat{r}_1$	0.0147 (0.0011)
			$\hat{s}_1$	0.9947 (0.0004)
	AIC	-9085.848	AIC	-10262.23
8 hours	$\hat{\nu}$	5.296 (0.339)	$\hat{\nu}$	7.206 (0.584)
	$\hat{\rho}$	0.612 (0.006)	$\hat{r}_0$	0.0005 (0.0005)
			$\hat{r}_1$	0.0173 (0.0014)
			$\hat{s}_1$	0.9927 (0.0006)
	AIC	-4813.6	AIC	-5456.312
12 hours	$\hat{\nu}$	5.830 (0.499)	$\hat{\nu}$	8.053 (0.884)
	$\hat{\rho}$	0.620 (0.008)	$\hat{r}_0$	0.0002 (0.0008)
			$\hat{r}_1$	-0.0249 (0.0023)
			$\hat{s}_1$	0.9901 (0.0010)
	AIC	-3299.16	AIC	-3744.28
1 day	$\hat{\nu}$	5.945 (0.758)	$\hat{\nu}$	8.573 (1.455)
	$\hat{\rho}$	0.619 (0.011)	$\hat{r}_0$	-0.0023 (0.0017)
			$\hat{r}_1$	-0.0343 (0.0041)
			$\hat{s}_1$	0.9846 (0.0021)
	AIC	-1644.549	AIC	-1881.760

Table 12: Parameter estimates, standard errors and AIC values for the two copula models, without and with dynamics in the correlation, fitted to the hourly up to daily returns on USD/DEM and USD/JPY rates.

$z_{n\,obs}^{1/2}$	n	$P\left(Z_n^{1/2} > z_{n\,obs}^{1/2}\right)$	$H_0(0.95)$	Time of change	Event
13.26	3 259	0	reject	8 Nov. 1989	Fall of the Berlin wall
5.96	923	0.0000004	reject	29 Dec. 1986	
5.31	2 336	0.0000143	reject	18 June 1997	Beginning of Asia crisis
2.99	176	0.0689621	not rej.	(23 June 1986)	
3.10	747	0.0709747	not rej.	(31 July 1989)	
5.86	1 985	0.0000007	reject	23 Oct. 1990	Burst in the Japanese asset price bubble
2.36	351	0.3380491	not rej.	(8 Sep. 1998)	
2.78	1 736	0.1873493	not rej.	(21 Oct. 1996)	
2.86	249	0.1061709	not rej.	(21 March 1990)	

Table 13: Change-point analysis for USD/DEM and USD/JPY spot rate residuals. In the last column we refer to economic or political events around the estimated date of the change-point that could have triggered it.

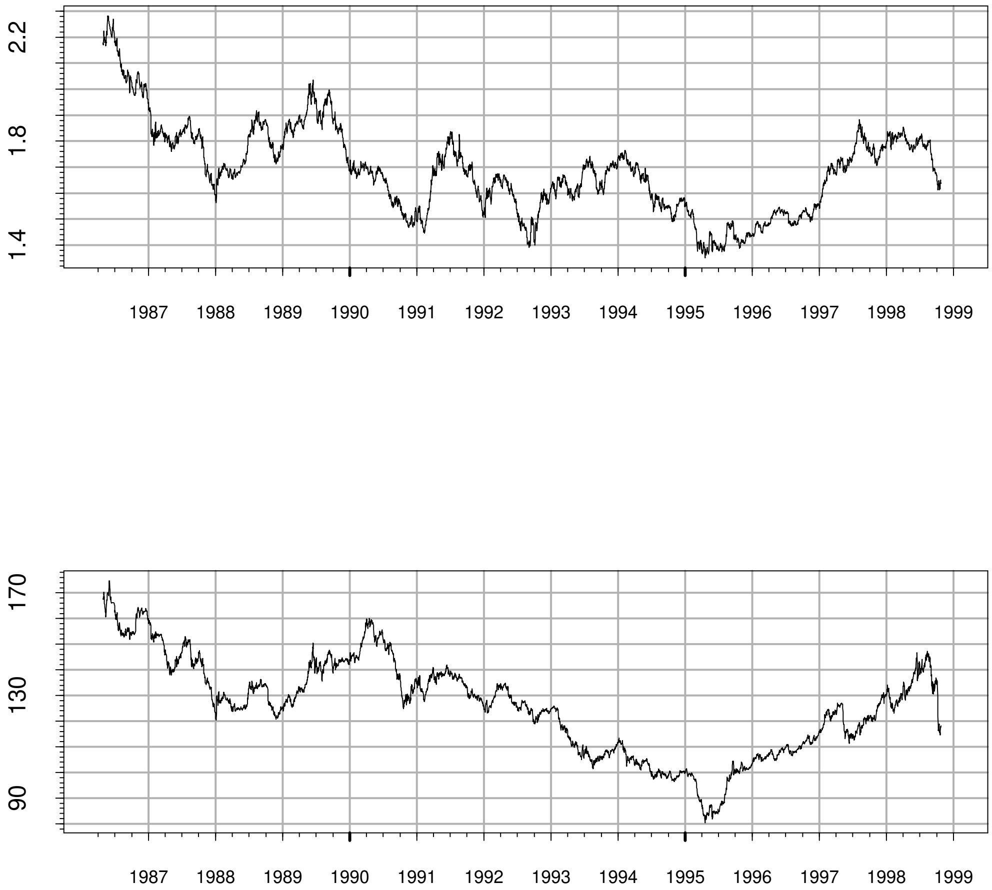

Figure 1: Logarithmic middle prices for USD/DEM (top) and USD/JPY (bottom) spot rates.

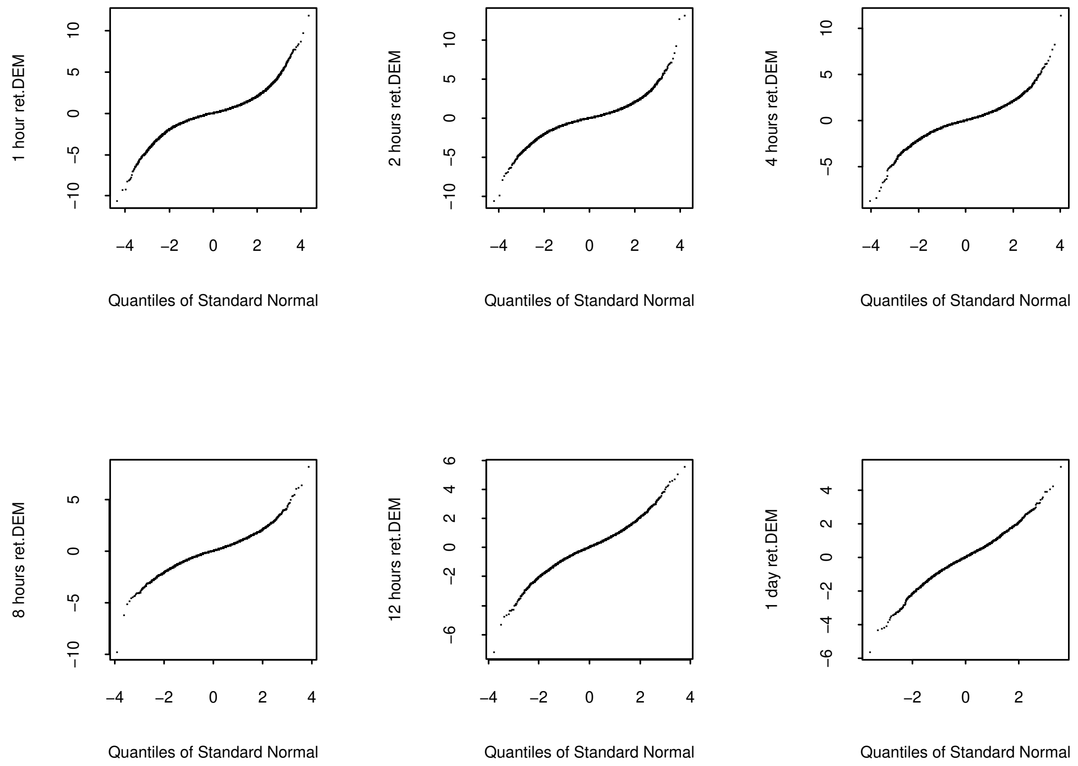

Figure 2: QQ-plots of the normal versus the empirical quantiles of deseasonalized log-returns on USD/DEM spot rate for the six frequencies considered.

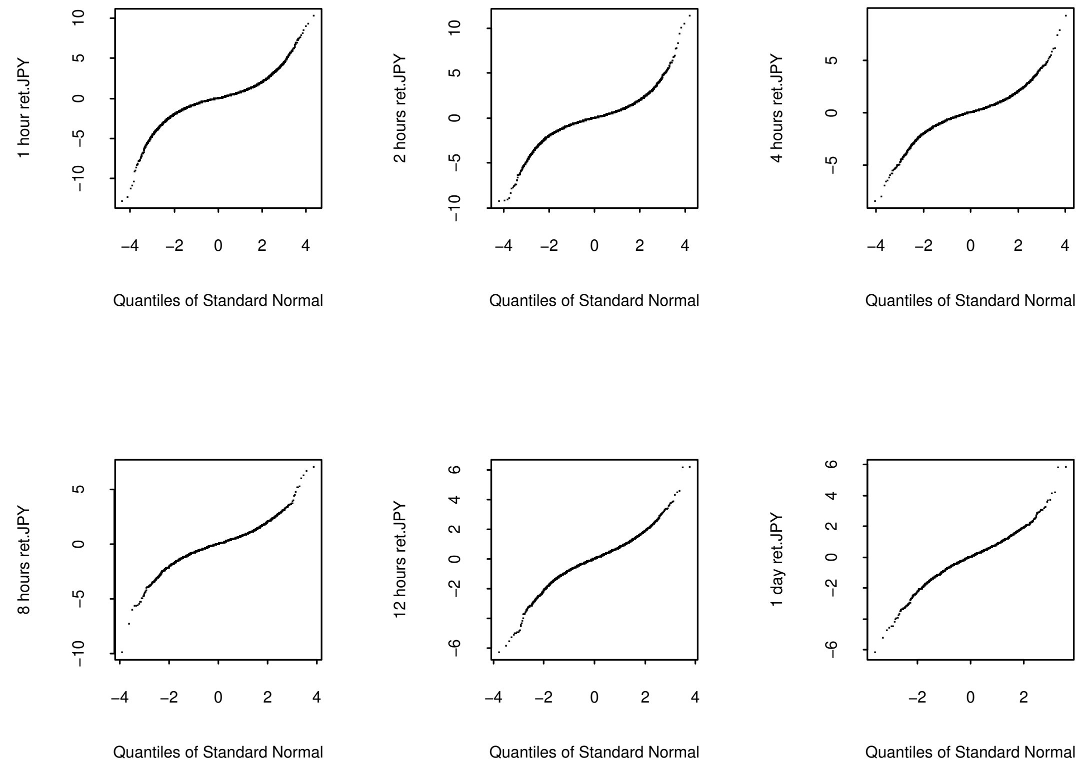

Figure 3: QQ-plots of the normal versus the empirical quantiles of deseasonalized log-returns on JPY/JPY spot rate for the six frequencies considered.

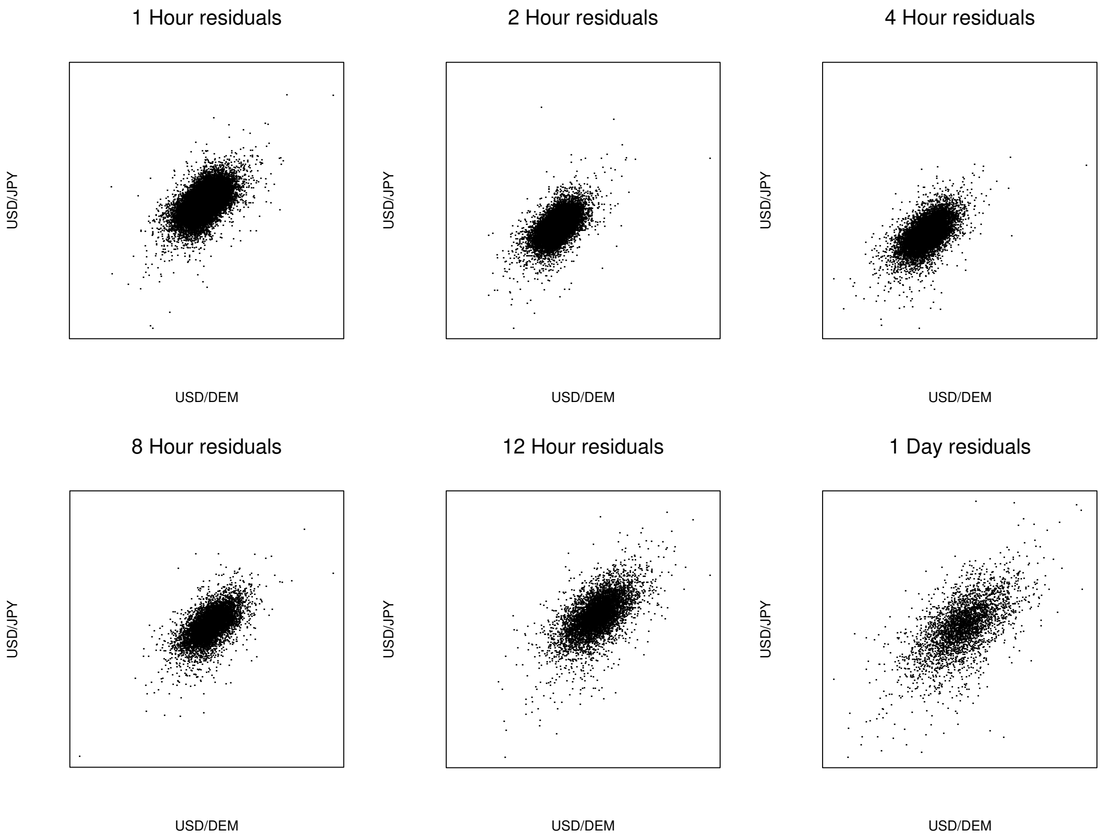

Figure 4: FX spot rates for USD/DEM and USD/JPY. The figure displays the scatter-plots of the filtered returns for the several frequencies.

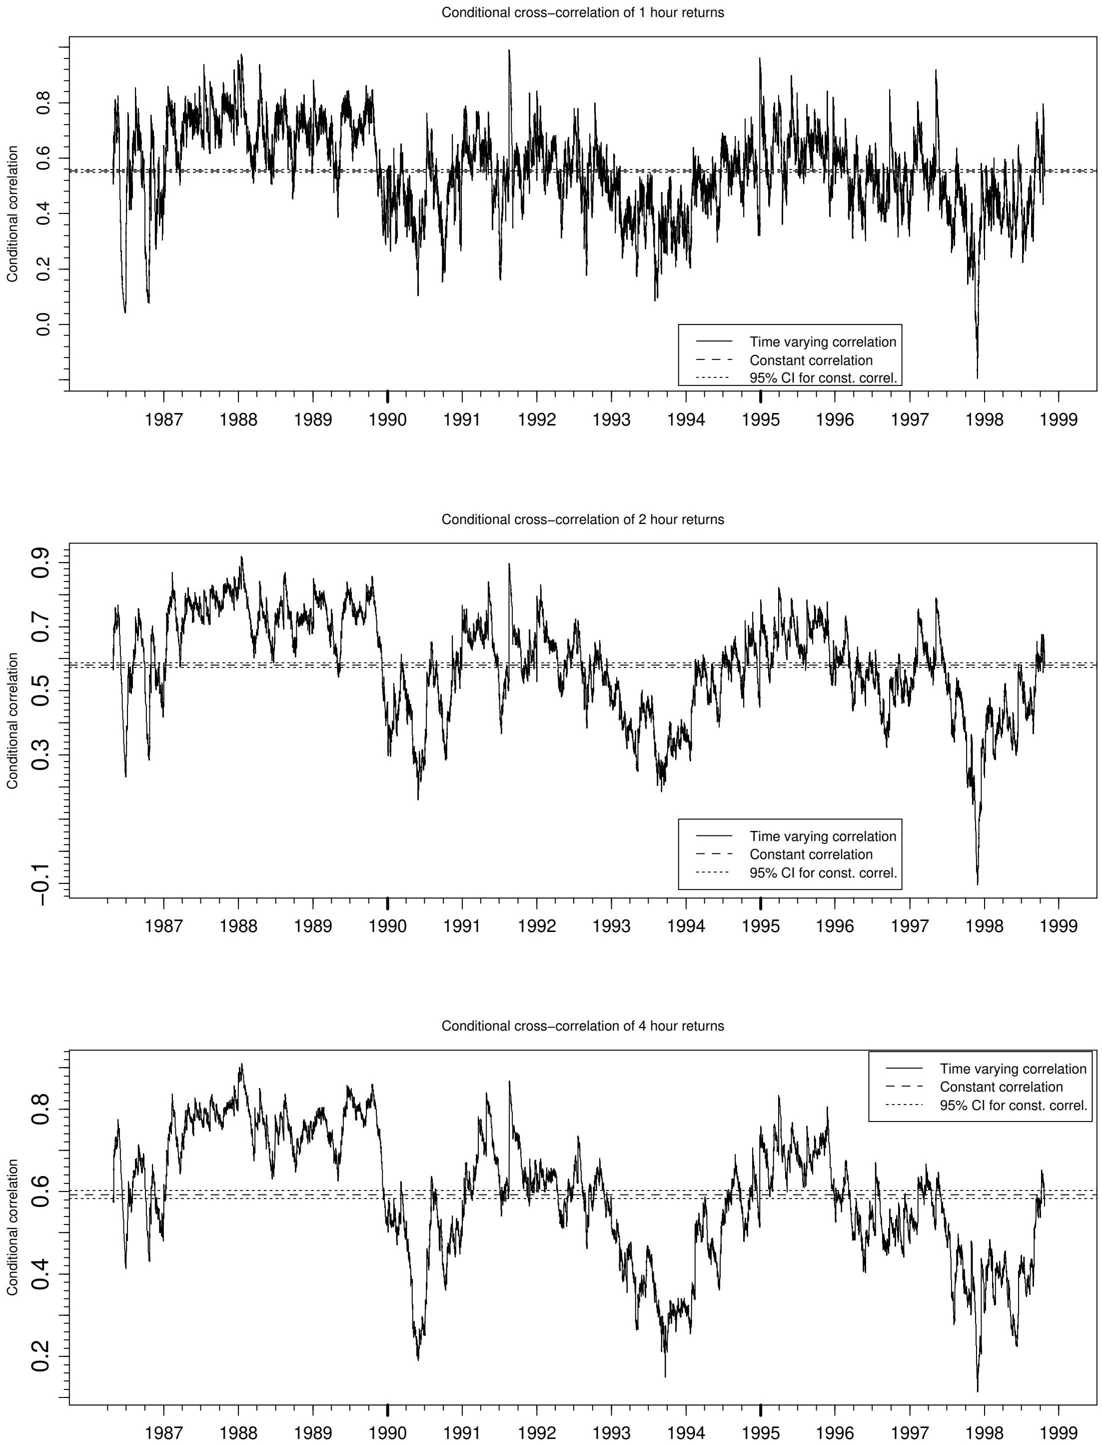

Figure 5: Time-varying cross-correlations estimated by a time-varying copula-based model for the one, two and four hours returns on the FX USD/DEM and USD/JPY spot rates.

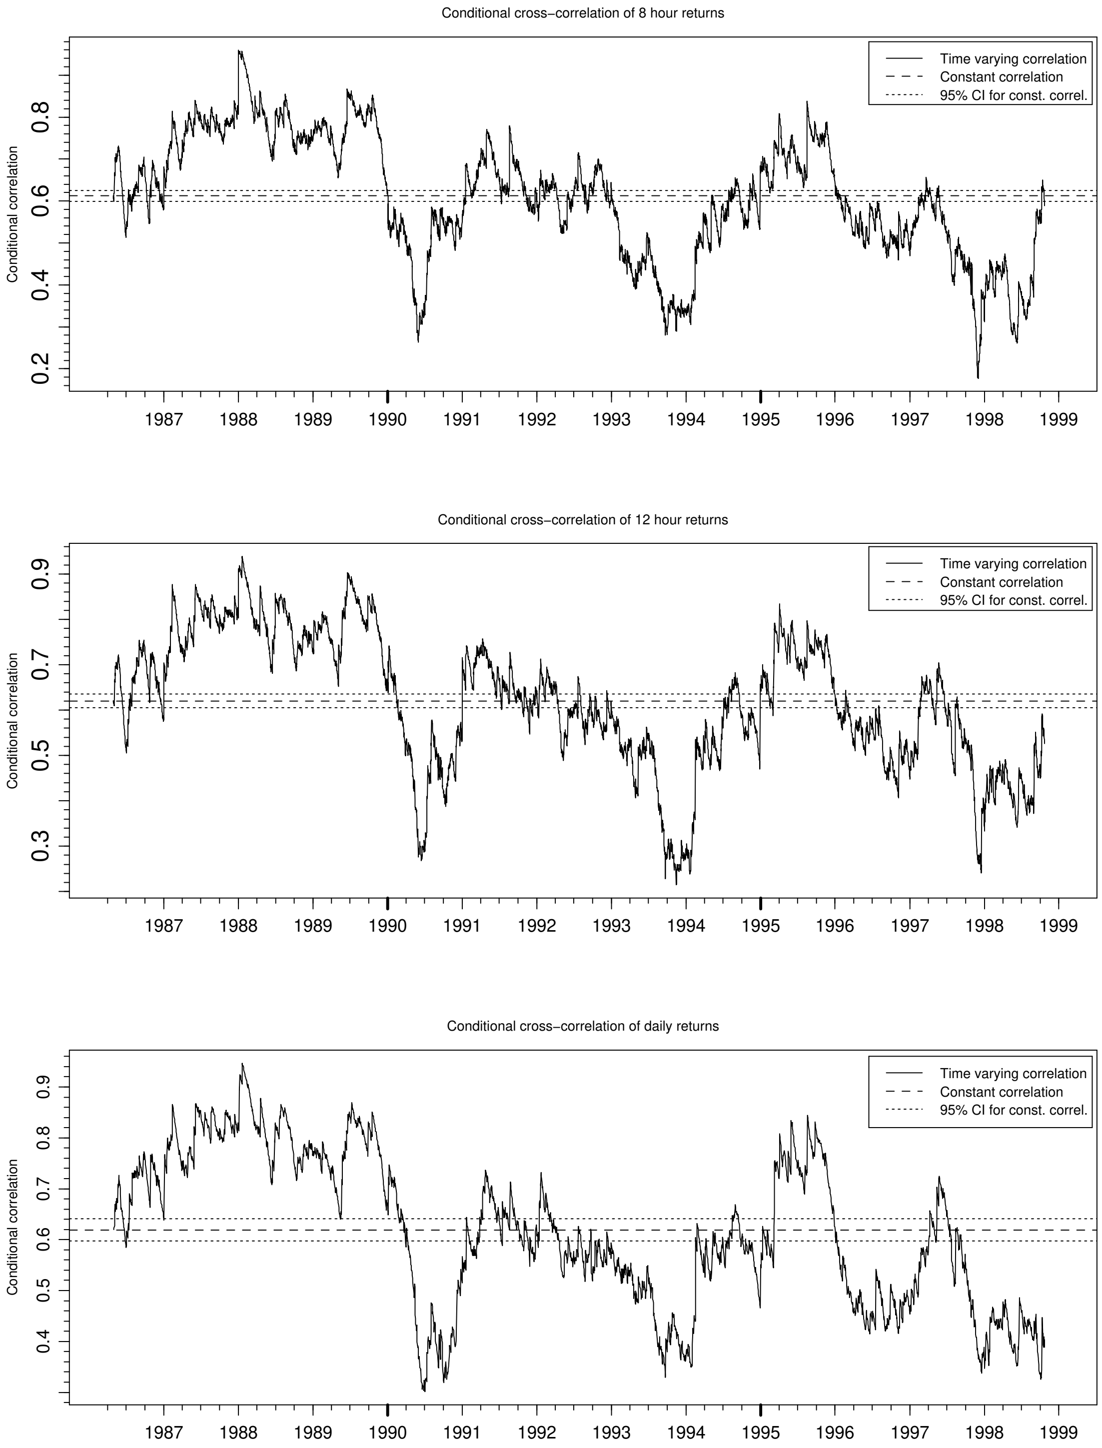

Figure 6: Time-varying cross-correlations estimated by a time-varying copula-based model for the eight hours, twelve hours and daily returns on the FX USD/DEM and USD/JPY spot rates.

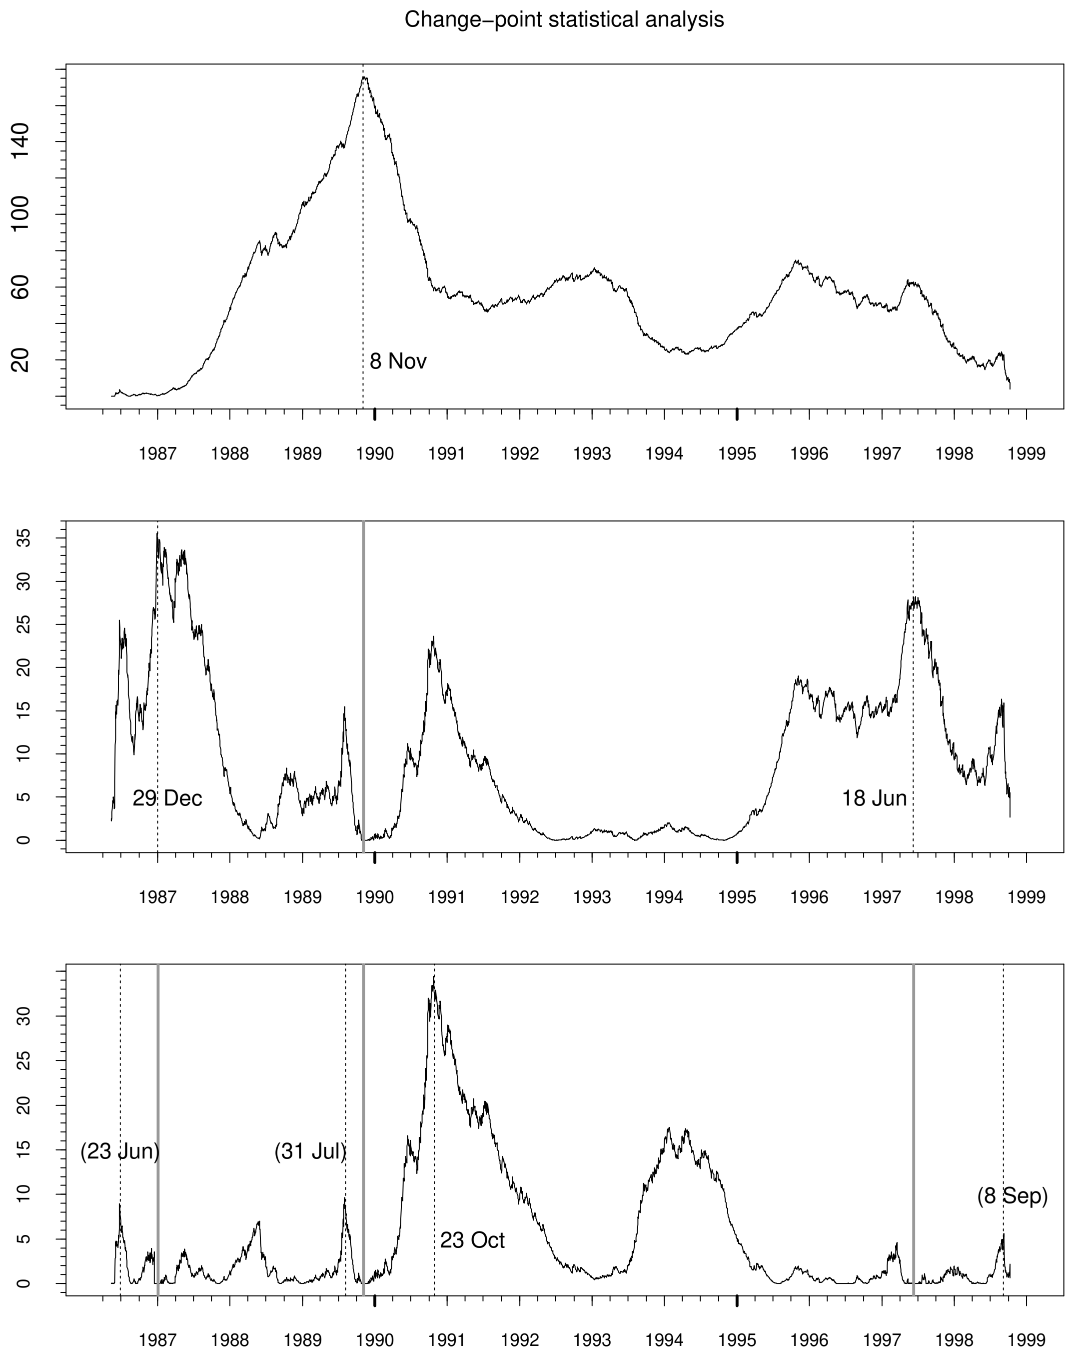

Figure 7: Change-point analysis of daily returns on the FX rates USD/DEM and USD/JPY spot rates. The three panels display three steps of the change-point analysis. Each panel plots the likelihood ratio values $-2\log(\Lambda_k)$ for $k = 1, 2, \ldots, n$. In each sub-sample its maximum, the test statistic Z_n, gives the time of the change in case the no change null hypothesis is rejected. If the null hypothesis is not rejected the moment where Z_n is achieved is put in parentheses.

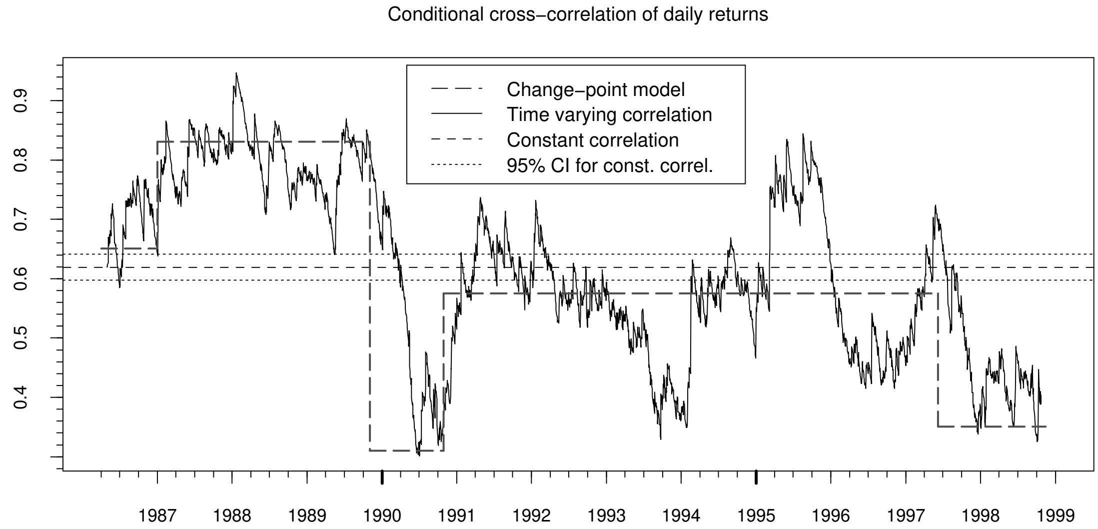

Figure 8: Estimated t-copula correlation paths of daily returns on the FX USD/DEM and USD/JPY spot rates. The long-dashed line is the estimated correlation by the change-point tests. This is super-imposed on the estimated correlation using the time-varying copula model from Chapter 3. The short-dotted line is the time-invariant correlation estimate after marginal GARCH filtering. The change-points model seems to react quicker to important economic events than the time-varying copula model but ignores smaller changes, at least at the daily frequency.