

Synchronizing Multivariate Financial Time Series

Francesco Audrino*
University of Southern Switzerland
and
Peter Bühlmann
ETH Zürich, Switzerland

May 2002
(Revised version)

Abstract

Prices or returns of financial assets are most often collected in local times of the trading markets. The need to synchronize multivariate time series of financial prices or returns is motivated by the fact that information continues to flow for closed markets while others are still open. We propose here a synchronization technique which takes this into account.

Besides the nice interpretation of synchronization, the method potentially increases the predictive performance of any reasonable model and is more appropriate for the calculation of portfolio risk measures such as for example the expected shortfall. We found empirically that this was the case for the CCC-GARCH(1,1) model for a 7-dimensional time series of daily exchange rate returns. Since multivariate analysis is generally important for analyzing time-changing portfolios and for better portfolio predictions (even when portfolio weights are time-constant), synchronization is a valuable technique for a variety of problems with multivariate financial data.

Keywords. CCC-GARCH model, Expected shortfall, Multivariate time series, Likelihood estimation, Value at Risk.

1 Introduction

The time of measurement of daily financial data, typically the closing time, is often different in markets which do not have the same trading hours. For example, between US and Japan there are no common open hours and between US and Europe there is only partial overlap. Therefore, the value of real global portfolios constructed on daily data across different markets is never known at a fixed point in time and consequently the calculation of risk measures such as the Value at Risk (quantile of the Profit-and-Loss distribution of a given portfolio over a prescribed holding period) and the conditional Value at Risk or expected shortfall (the expected loss given that the loss exceeds VaR) could give misleading results. A consequence of using such asynchronous data is the fact that correlations across the assets are often small, see Burns et al. (1998). This may lead to inaccurate (estimated) risk calculations such as VaR or expected shortfall.

We propose here a synchronization of daily data in real global portfolios. Proceeding as in Burns et al. (1998), our general approach recognizes that even when markets are closed, the

*Corresponding author: Institute of Finance, University of Southern Switzerland, Via Buffi 13, 6900 Lugano.

asset values may change before the market reopens. Synchronizing data involves estimates of asset values at a specified (synchronization) time point in every day; we always use the closing time of the New York stock exchange, i.e. 4 pm local New York time, as the synchronization time point. The estimated asset values at the same synchronization time across markets are called *synchronized*. Having constructed synchronized data, we then propose to use any reasonable multivariate model for the synchronized data. This is different from Burns et al. (1998) which do not model the obtained synchronized data any further. We consider here the CCC-GARCH(1,1) model (Bollerslev, 1990) with time varying conditional variances and covariances but constant conditional correlations for synchronized data which represents a different and new model for the original asynchronous data. This new model is called *synchronous CCC-GARCH(1,1)* and is presented in Section 2.

In Section 3, we compare empirically the performance of the results obtained with synchronized and asynchronous data, using the same multivariate CCC-GARCH(1,1) model. The analysis is for a real global portfolio with negative daily log-returns (in percentages) $X_{t,i} = -100 \cdot \log \frac{S_{t,i}}{S_{t-1,i}}$ of seven indices all over the world, where $S_{t,i}$ denotes the price of the asset i at day t (time synchronized or asynchronous, respectively). In Section 5, we discuss the influence of synchronization on the calculation of risk measures such as VaR and expected shortfall.

The resulting gains when using synchronous data are sometimes considerable, depending on how we measure performance. We argue in Section 4 that these relevant gains are not due to strong model-misspecification of the volatilities in the CCC-GARCH(1,1) model with asynchronous data. More sophisticated threshold models with asynchronous data like the multivariate extension of the univariate tree-structured GARCH model (Audrino and Bühlmann, 2001), which constructs a potentially high dimensional approximation of a general non-parametric CCC model, yields only marginal improvements over the standard CCC-GARCH(1,1) model. Thus, it emphasizes the power of synchronizing data in a first step.

Our empirical comparisons in Sections 3 and 5 also include univariate approaches for a portfolio index price

$$P_t = \sum_i \alpha_i S_{t,i} \tag{1.1}$$

modelled by previous P_{t-1}, P_{t-2}, \dots . Note that in the more realistic case where the weights $\alpha_i = \alpha_{t,i}$ are depending (deterministically) on t , stationarity of all $\{S_{t,i}\}_t$ is not implying stationarity of the portfolio prices $\{P_t\}_t$; but a multivariate analysis of $\{\{S_{t,i}\}_t; i = 1, \dots, M\}$ still yields the (conditional) distribution of P_t (given the past). It turns out that even when focusing on a time-constant portfolio index as in (1.1), multivariate approaches are better than sophisticated univariate modelling of P_t based on P_{t-1}, P_{t-2}, \dots , which is exposed to an information loss by averaging previous individual prices.

Summarizing, we collect strong empirical evidence that:

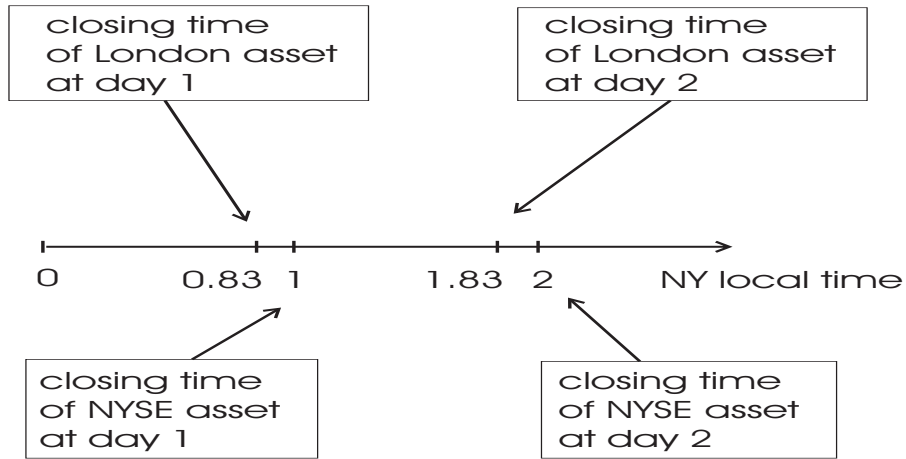
1. the multivariate modelling approach is superior over univariate models for a portfolio index as in (1.1). In particular, we find that univariate modelling yields estimates which are too conservative.
2. synchronizing data, proposed here as a novelty, leads to further improvements in multivariate modelling and in calculating risk measures.

2 The synchronous CCC-GARCH(1,1) model

2.1 Synchronization of the data

Consider a global portfolio including stocks traded in New York and London. At the closing time of the trading in New York, the value of the portfolio should be measured with an estimate of the value of the London stocks at the closing time in New York. For example, to take the closing prices of the London stocks at one day when the US market goes down 1 percent (after London closes) for pricing the portfolio at New York closing time is highly unrealistic. It will follow that the US share of the portfolio declines today while the London share declines tomorrow. We associate with synchronization some estimates of the prices of the share traded in London at the closing time in New York (from the viewpoint of a British investor, the data could also be synchronized at the closing time in London).

We denote by $S_{t,j}$, $j = 1, \dots, M$ the continuous time price of an asset j . The time t is here always measured as New York local time (in units of days) and $t \in \mathbb{N}$ corresponds to 4:00 pm New York local time on day t . For example, $S_{1,1}$ denotes the price of an asset of the NYSE at 4:00 pm New York local time on the first day. Since 4:00 pm corresponds to 9:00 pm in London and since London closes at 5:00 pm, 4 hours before New York closes, the observed closing price of an asset in London on the first day would be denoted by $S_{0.83,2}$. This is also illustrated by the following Figure.



Generally, the observed data is taken at closing times of different markets. It has the structure

$$S_{t_j,j} \quad (j = 1, \dots, M), \text{ where} \\ t_j = t_1 - c_j \quad (0 \leq c_j \leq 1), \quad j = 2, \dots, M.$$

We always synchronize to the closing time t_1 (in New York) of asset $j = 1$, where $t_1 \in \{1, 2, \dots, T\}$. The goal is to construct synchronized prices $S_{t,j}^s$ with $t \in \{1, 2, \dots, T\}$ for all j . These prices, or returns thereof, are more appropriate for many multivariate discrete time series models.

We define the *synchronized prices* $S_{t,j}^s$ by

$$\log(S_{t,j}^s) = \mathbb{E} \left[\log(S_{t,j}) \mid \mathcal{F}_t \right], \text{ where } \mathcal{F}_t = \left\{ S_{t_j,j} ; t_j \leq t, j = 1, \dots, M \right\}. \quad (2.1)$$

The logarithms are used to be consistent with continuously compounded returns, and \mathcal{F}_t is the complete information of all recorded prices up to time t . We usually only need $S_{t,j}^s$ for discrete $t \in \mathbb{N} = \{1, 2, \dots\}$.

Clearly, if the closing price S is observed at time $t \in \mathbb{N}$, then its conditional expectation given \mathcal{F}_t is the observed price. This is the case for the stocks from New York. If the market closes before t , then its past prices and all the other markets are potentially useful in predicting S at time t .

As a simplifying but reasonable approximation, we assume that, given the information \mathcal{F}_t , the best predicted log-prices at t and at the nearest succeeding closing time $t_j + 1$ remain the same, saying that future changes up to $t_j + 1$ are unpredictable

$$\log(S_{t,j}^s) = \mathbb{E}[\log(S_{t,j}) | \mathcal{F}_t] = \mathbb{E}[\log(S_{t_j+1,j}) | \mathcal{F}_t], \quad t_j \leq t < t_j + 1 \quad (t \in \mathbb{N}). \quad (2.2)$$

The first equality holds by definition in (2.1). Thus, given the information at time t , the next predicted future values are given by the log-transformed synchronized prices.

Now, denote the vector of negative log-returns (in percentages), in different markets and at various time points on day t as \mathbf{X}_t ,

$$\mathbf{X}_t = -100 \cdot \begin{pmatrix} \log\left(\frac{S_{t_1,1}}{S_{t_1-1,1}}\right) \\ \vdots \\ \log\left(\frac{S_{t_M,M}}{S_{t_M-1,M}}\right) \end{pmatrix} = -100 \cdot \left(\log(\mathbf{S}_t) - \log(\mathbf{S}_{t-1}) \right), \quad (2.3)$$

where $\mathbf{t} = (t_1, t_2, \dots, t_M)$ is a multi-index.

We define the *synchronized returns* as the change in the logarithms of the synchronized prices

$$\mathbf{X}_t^s = -100 \cdot \begin{pmatrix} \log\left(\frac{S_{t,1}^s}{S_{t-1,1}^s}\right) \\ \vdots \\ \log\left(\frac{S_{t,M}^s}{S_{t-1,M}^s}\right) \end{pmatrix} = -100 \cdot \left(\log(\mathbf{S}_t^s) - \log(\mathbf{S}_{t-1}^s) \right), \quad t \in \mathbb{N}. \quad (2.4)$$

The synchronized returns are depending on unknown conditional expectations and have to be modelled (and estimated). We assume a simple ‘‘auxiliary’’ multivariate AR(1) model for the synchronization, given by

$$\mathbf{X}_t = A \cdot \mathbf{X}_{t-1} + \epsilon_t, \quad (2.5)$$

with errors ϵ_t such that $\mathbb{E}[\epsilon_t | \mathcal{F}_{t-1}] \equiv 0$, and A a $M \times M$ matrix. Contrary to the approach of Burns et al. (1998) which propose a first order vector moving average, we choose a synchronization which is somewhat simpler since $\mathbb{E}[\mathbf{X}_t | \mathcal{F}_{t-1}]$ in (2.5) depends only on the previous \mathbf{X}_{t-1} (as in a Markovian model). As we will show in (2.6), the synchronized returns with (2.5) are then functions of \mathbf{X}_t and \mathbf{X}_{t-1} only and not of unobservable innovations ϵ_t (or infinitely many lagged variables $\mathbf{X}_{t-1}, \mathbf{X}_{t-2}, \dots$ as in Burns et al. (1998)

Substituting (2.2), (2.3) and (2.5) into (2.4) gives the synchronized returns as

$$\begin{aligned} \mathbf{X}_t^s &= -100 \cdot \left(\log(\mathbf{S}_t^s) - \log(\mathbf{S}_{t-1}^s) \right) = -100 \cdot \left(\mathbb{E}[\log(\mathbf{S}_{t+1}) | \mathcal{F}_t] - \mathbb{E}[\log(\mathbf{S}_t) | \mathcal{F}_{t-1}] \right) \\ &= -100 \cdot \left(\mathbb{E}[\log(\mathbf{S}_{t+1}) - \log(\mathbf{S}_t) | \mathcal{F}_t] - \mathbb{E}[\log(\mathbf{S}_t) - \log(\mathbf{S}_{t-1}) | \mathcal{F}_{t-1}] + \log\left(\frac{\mathbf{S}_t}{\mathbf{S}_{t-1}}\right) \right) \\ &= \mathbb{E}[\mathbf{X}_{t+1} | \mathcal{F}_t] - \mathbb{E}[\mathbf{X}_t | \mathcal{F}_{t-1}] + \mathbf{X}_t = \mathbf{X}_t + A \cdot \mathbf{X}_t - A \cdot \mathbf{X}_{t-1}, \end{aligned}$$

and thus

$$\mathbf{X}_t^s = \mathbf{X}_t + A \cdot (\mathbf{X}_t - \mathbf{X}_{t-1}) . \quad (2.6)$$

Clearly, if A is the zero matrix, $\mathbf{X}_t^s = \mathbf{X}_t$ and the data are already synchronized. Since the New York market data are already synchronized, the row of A corresponding to the New York stocks is a zero row.

Computing synchronized returns from (2.6) boils down to estimation of A in model (2.5) or a more specific version. The estimation procedure is described in Section 2.3.

2.2 The model

For the synchronized returns, we consider the standard CCC-GARCH(1,1) model, introduced by Bollerslev (1990):

$$\begin{aligned} \mathbf{X}_t^s &= \mu_t^s + \epsilon_t^s = \mu_t^s + \Sigma_t^s \mathbf{Z}_t \quad (t \in \mathbb{Z}) , \\ \mathbf{X}_t^s &= \mathbf{X}_t + A (\mathbf{X}_t - \mathbf{X}_{t-1}) = (I_M + A) \mathbf{X}_t - A \mathbf{X}_{t-1}, \end{aligned} \quad (2.7)$$

where we make the following assumptions:

- (A1) $(\mathbf{Z}_t)_{t \in \mathbb{Z}}$ is a sequence of i.i.d. multivariate innovation variables with spherical distribution (e.g. the multivariate normal or the multivariate t distribution) with zero mean, covariance matrix $\text{Cov}(\mathbf{Z}_t) = I_M$ and \mathbf{Z}_t independent from $\{\mathbf{X}_k^s, k < t\}$;
- (A2) (CCC construction) $\Sigma_t^s (\Sigma_t^s)' = H_t^s$ is almost surely positive definite for all t , where the typical element of H_t^s is $h_{ij,t}^s = \rho_{ij}^s (h_{ii,t}^s h_{jj,t}^s)^{\frac{1}{2}}$, for $i, j = 1, \dots, M$;
- (A3) (GARCH(1,1) part) $h_{ii,t}^s = (\sigma_{t,i}^s)^2 = \alpha_0^{(i)} + \alpha_1^{(i)} (X_{t-1,i}^s)^2 + \beta^{(i)} (\sigma_{t-1,i}^s)^2$, with $\alpha_0^{(i)}, \alpha_1^{(i)}, \beta^{(i)} > 0$ for $i = 1, \dots, M$;
- (A4) $\mu_t^s = \mathbb{E}[\mathbf{X}_t^s \mid \mathcal{F}_{t-1}] = (I_M + A) \mathbb{E}[\mathbf{X}_t \mid \mathcal{F}_{t-1}] - A \mathbf{X}_{t-1}$, $\mathbb{E}[\mathbf{X}_t \mid \mathcal{F}_{t-1}] = A \mathbf{X}_{t-1}$ (as in (2.5)).

We call this the *synchronous CCC-GARCH(1,1) model*. Note that ρ_{ij}^s in (A2) equals the constant conditional correlation $\text{Corr}(X_{t,i}^s, X_{t,j}^s \mid \mathcal{F}_{t-1})$.

Proposition 1.

Assume that the matrix $(I_M + A)$ is invertible. Then, the synchronous CCC-GARCH(1,1) model (2.7) can be represented with asynchronized returns \mathbf{X}_t :

$$\mathbf{X}_t = A \cdot \mathbf{X}_{t-1} + (I_M + A)^{-1} \Sigma_t^s \mathbf{Z}_t , \quad (2.8)$$

where the matrix Σ_t^s has the same CCC-GARCH(1,1) structure as defined in (2.7).

Proposition 1 implies that the synchronous CCC-GARCH(1,1) model is still a constant conditional correlation model (in terms of asynchronized data). Moreover, we should view it as a super-model of the classical CCC-GARCH(1,1): setting $A = 0$ yields the classical sub-model. Generally A is a sparse parameter matrix whose structure will be estimated from data, see Section 2.3.

Proof of Proposition 1. Using (2.6) and the fact that $\mathbf{X}_{t-1} \in \mathcal{F}_{t-1}$, we calculate the conditional mean of the synchronized returns as

$$\begin{aligned} \mu_t^s &= \mathbb{E}[\mathbf{X}_t^s \mid \mathcal{F}_{t-1}] = \mathbb{E}[(I_M + A) \cdot \mathbf{X}_t \mid \mathcal{F}_{t-1}] - \mathbb{E}[A \cdot \mathbf{X}_{t-1} \mid \mathcal{F}_{t-1}] = \\ &= (I_M + A) \cdot \mathbb{E}[\mathbf{X}_t \mid \mathcal{F}_{t-1}] - A \cdot \mathbf{X}_{t-1} = (I_M + A) \cdot \mu_t - A \cdot \mathbf{X}_{t-1} = \\ &= (I_M + A) \cdot A \cdot \mathbf{X}_{t-1} - A \cdot \mathbf{X}_{t-1} = A^2 \cdot \mathbf{X}_{t-1} . \end{aligned}$$

It follows that (2.7) is equivalent to $\mathbf{X}_t^s = A^2 \cdot \mathbf{X}_{t-1} + \Sigma_t^s \mathbf{Z}_t$. By (2.6), we obtain the assertion:

$$\begin{aligned} \mathbf{X}_t^s &= (I_M + A) \cdot \mathbf{X}_t - A \cdot \mathbf{X}_{t-1} = A^2 \cdot \mathbf{X}_{t-1} + \Sigma_t^s \mathbf{Z}_t \\ \iff \mathbf{X}_t &= (I_M + A)^{-1} (A + A^2) \cdot \mathbf{X}_{t-1} + (I_M + A)^{-1} \Sigma_t^s \mathbf{Z}_t \\ \iff \mathbf{X}_t &= A \cdot \mathbf{X}_{t-1} + (I_M + A)^{-1} \Sigma_t^s \mathbf{Z}_t . \end{aligned}$$

□

2.3 Estimating the model

Model structure. The synchronous CCC-GARCH(1,1) involves the matrix A : we insist on sparseness by setting some elements (to be selected from the data) to zero. This is important to reduce the number of parameters in the case of high-dimensional portfolios with dozens up to hundreds of assets. We proceed with a computationally fast and feasible procedure for estimating the structure (the non-zero elements) of the matrix A ; the actual values of A will then be estimated by maximum likelihood in the working model in (2.7).

Step 1. Find the estimates for the M^2 parameters of the matrix A and for the matrix Σ using the Yule-Walker estimator. The Yule-Walkers covariance relations for a multivariate AR(1) model are given by

$$\begin{aligned} R(0) &= R(-1) \cdot A' + \Sigma = R(1)' \cdot A' + \Sigma \\ R(1) &= R(0) \cdot A' , \quad \text{where } R(k) = \mathbb{E}[\mathbf{X}_{t-k} \cdot \mathbf{X}_t'] . \end{aligned}$$

For more details about the Yule-Walker estimator, see Brockwell and Davis (1991) or Reinsel (1991).

Calculate some model-based standard errors of the estimated elements of A using a bootstrap strategy, i.e. an i.i.d. resampling of the residuals estimated from our working model (2.7) with full matrix A and recursively generating a time series using the fitted parameters in model (2.7); this is called a semiparametric model-based bootstrap, for more details see Efron and Tibshirani (1993)

$$\text{s.e.}(\widehat{A}_{ij}) = \sqrt{\widehat{\text{Var}}(\widehat{A}_{ij})} = \sqrt{\frac{1}{B-1} \sum_{b=1}^B (\widehat{A}_{ij}^*{}_{(b)} - \overline{\widehat{A}^*}_{ij})^2} ,$$

where $\overline{\widehat{A}^*}_{ij} = \frac{1}{B} \sum_{b=1}^B \widehat{A}_{ij}^*{}_{(b)}$, $\widehat{A}_{ij}^*{}_{(b)}$ is the estimate of the ij -th element of the matrix A in the b -th bootstrap iteration and B is the number of bootstrap iterations.

Step 2. Set $A_{ij} = 0$ if the t-ratio

$$t_{ij} = \left| \frac{\widehat{A}_{ij}}{\text{s.e.}(\widehat{A}_{ij})} \right| \leq 1.96 \quad (5\% \text{ significance level})$$

and $A_{1j} = 0$ for all j corresponding to the New York stocks.

Parameter estimation. The parameters A , $\alpha_0^{(j)}$, $\alpha_1^{(j)}$, $\beta^{(j)}$, ρ_{ij}^s ($j = 1, \dots, M$) in the synchronous CCC-GARCH(1,1) model can be estimated with the maximum likelihood method. We usually assume the innovations \mathbf{Z}_t to be multivariate t_ν distributed with zero mean and

covariance matrix $\text{Cov}(\mathbf{Z}_t) = I_M$, where the degrees of freedom ν have to be estimated as well, i.e. $\mathbf{Z}_t \sim t_\nu(0, I_M)$. The negative log-likelihood is then given by

$$-l(\theta; \mathbf{X}_2^n) = \frac{TM}{2} \log(\pi\nu) - T \log \left(\frac{\Gamma(\frac{M+\nu}{2})}{\Gamma(\frac{\nu}{2})} \right) + \frac{T}{2} \log |R^s| + \sum_{t=1}^T \left(\log |D_t^s| \right) + \frac{M+\nu}{2} \sum_{t=2}^T \left(\log \left(1 + \frac{(\epsilon_t^s)'(R^s)^{-1}\epsilon_t^s}{\nu} \right) \right) - T \log (| (I_M + A) |), \quad (2.9)$$

where $H_t^s = D_t^s R^s D_t^s$ and from the CCC-construction, D_t^s a diagonal $M \times M$ matrix with diagonal-elements $\sigma_{t,1}^s, \dots, \sigma_{t,M}^s$, $R^s = [\rho_{ij}^s]_{1 \leq i, j \leq M}$ and $\epsilon_t^s = (D_t^s)^{-1}(\mathbf{X}_t^s - \mu_t^s)$; θ denotes the vector of all parameters involved and $\mathbf{X}_t^s = \mathbf{X}_t + A(\mathbf{X}_t - \mathbf{X}_{t-1})$ as before. Also, we use the sparse structure of the matrix A as described above.

For preliminary $R^s = I_M$, we estimate the remaining parameters $A, \alpha_0^{(j)}, \alpha_1^{(j)}, \beta^{(j)}$ ($j = 1, \dots, M$) by minimizing the negative log-likelihood in (2.9). This yields estimates $\hat{\mu}_t^s = (\hat{\mu}_{t,1}^s, \dots, \hat{\mu}_{t,M}^s)$ and $\hat{\sigma}_{t,j}^s$, $j = 1, \dots, M$. Then, we construct the estimate for the correlation matrix R^s as follows: build the residuals

$$\hat{\epsilon}_{t,j} = (\hat{X}_{t,j}^s - \hat{\mu}_{t,j}^s) / \hat{\sigma}_{t,j}^s, \quad t = 1, \dots, T$$

and define

$$\hat{R}^s = T^{-1} \sum_{t=1}^T \hat{\epsilon}_t \hat{\epsilon}_t^T, \quad \hat{\epsilon}_t = (\hat{\epsilon}_{t,1}, \dots, \hat{\epsilon}_{t,M})^T. \quad (2.10)$$

We can then iterate (once) by minimizing the negative log-likelihood in (2.9) using \hat{R}^s from (2.10).

3 Numerical results

We consider a real global portfolio of seven market indices: US Dow Jones Industrial Average (DJIA), French CAC40 Index, German Deutsche Aktien (DAX), Italian BCI General Index, Dutch CBS All-Share, UK FT-SE-A All-Share Index (FTAS) and Japanese NIKKEI 225 Average (NIK). The daily data is from the time period between January 17, 1990 and June 22, 1994, corresponding to 1000 days without holidays in the different countries. The closing times of the seven market indices are given in the following Table. We use here (negative) relative difference

Index	Closing local time	Closing NY local time
NIKKEI	3:00 PM	2:00 AM
CBS	5:30 PM	11:30 AM
BCI	5:30 PM	11:30 AM
CAC40	5:30 PM	11:30 AM
FTAS	5:00 PM	12:00 PM
DAX	8:00 PM	2:00 PM
DJIA	4:00 PM	4:00 PM

returns (in percentages) $X_{t_j,j} = -100 \cdot \frac{S_{t_j,j} - S_{t_{j-1},j}}{S_{t_{j-1},j}}$, where $S_{t_j,j}$ denotes the price of the asset

j at the local closing time t_j of the day t , because they are close approximations of the log-returns and because they allow for much simpler portfolio and risk computations; see Section 3.3 for more details on the construction of the negative portfolio returns. Nevertheless, we still synchronize such relative difference returns as in (2.6).

The aim is to support empirically the effect of synchronization and to compare the synchronous CCC-GARCH(1,1) model in (2.7) with the asynchronous classical CCC-GARCH(1,1) model. For the numerical optimization, we use a quasi-Newton method.

3.1 Estimate of A and synchronization

We examine here first the effect of synchronization from a descriptive point of view. The parsimoniously estimated matrix A is obtained from the procedure for structure determination illustrated in Section 2.3 and from maximum likelihood in (2.9):

$$\hat{A} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.2223 & 0 & 0.0189 & 0.0212 & 0 & 0 & -0.0663 \\ 0.3012 & 0.0873 & -0.0086 & 0 & 0 & 0 & -0.0916 \\ 0.2883 & 0 & 0 & 0 & -0.0107 & 0.0970 & -0.0164 \\ 0.2493 & 0 & -0.0033 & 0 & 0 & 0 & -0.0401 \\ 0.1749 & 0 & 0 & 0.0073 & 0 & 0.0412 & -0.0507 \\ 0.3168 & 0.0510 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad (3.1)$$

where the variables are ordered as DJIA, CAC40, DAX, BCI, CBS, FTAS, NIK. The column with the largest coefficients corresponds to the DJIA: there is substantial predictability of all other markets from the DJIA at the previous day. Besides a major determining effect of the US market for the financial world, the observed pattern is natural since the exchange in New York closes last. There seems to be also predictability of all other markets from the Japanese returns (NIK), although in this case, the coefficients are small and negative. The negative sign, which seems to work as a kind of correction impulse for the European indices, could be explained by some joint effect from the DJIA and the NIK index and could be a consequence of the big impact of the US on the Japanese market. The German DAX and the British FTAS seem also to be autocorrelated. The coefficients which correspond to the three markets which close simultaneously (French, Italian and Dutch) are all equal to zero except for two which are still close to zero but have t-ratios (from Section 2.1) smaller than 2.3. All the other coefficients have t-ratios greater than 3 except the three negative coefficients in Germany and the -0.0164 and -0.0401 in Japan.

Using \hat{A} from (3.1) and the synchronization formula (2.6), we obtain the synchronized returns $\hat{\mathbf{X}}_t^s$. The effect of synchronization in terms of empirical correlations is described in Table 3.1: synchronized data often exhibit larger instantaneous correlations between different returns from indices at the same day.

TABLE 3.1 ABOUT HERE.

The empirical correlations are typically too small for highly asynchronous markets. This is the case for example of the US and the Japanese markets: the empirical correlation between DJIA and NIK is much bigger when synchronizing (0.328 vs. 0.189). Of course, there is no reason to believe that synchronization generally yields higher correlations. This result is consistent and similar to the analysis in Burns et al. (1998).

3.2 Estimates for the synchronous CCC-GARCH(1,1) model and its performance

The parameters are estimated by maximum likelihood as in Section 2.3. For quantifying the goodness of fit of the models, we consider the following statistics:

the AIC statistic : $-2 \log\text{-likelihood} + 2 \# \text{ parameters}$

the outsample $-\log\text{-likelihood}$: $-\log\text{-likelihood}\left(\tilde{\mathbf{X}}_1^T; \hat{A}, \hat{\nu}, \hat{R}^s, \{\hat{\alpha}_0^{(j)}, \hat{\alpha}_1^{(j)}, \hat{\beta}^{(j)}; j = 1, \dots, M\}\right)$,

where $\tilde{\mathbf{X}}_1^T = \tilde{\mathbf{X}}_1, \dots, \tilde{\mathbf{X}}_T$ are new test data and the parameter estimates, equipped with hats, are from the training sample $\mathbf{X}_1^n = \mathbf{X}_1, \dots, \mathbf{X}_n$. The likelihood itself is given in (2.9). Both quantities are measures for out-sample performance: their values (low is better) indicate closeness to the true data generating model with respect to the Kullback-Leibler divergence. In our analysis, we take $n = 1000$ and the test set values $\tilde{\mathbf{X}}_1^T = \mathbf{X}_{n+1}^{n+500}$ are the next 500 consecutive observations (days between June 23, 1994 and September 9, 1996). We take $T = 500$ (little more than two years), because it seems a reasonable time period where the multivariate return series of the seven indices are believed to be (at least approximately) stationary.

The resulting values for the AIC and the out-of-sample negative log-likelihood statistic are 19058.6 and 4235.798 respectively using the synchronous CCC-GARCH(1,1) model and are 19136.796 and 4284.736 respectively using the classical CCC-GARCH(1,1) model without synchronization. The synchronous CCC-GARCH(1,1) model is better than the asynchronous CCC-GARCH(1,1) model with respect to both goodness of fit statistics, although the difference is small (in order of 1 percent). Such small differences could be obscured by low signal to noise ratio. It is often useful to consider differences of performance terms and use the concept of hypothesis testing, rather than quantifying differences in terms of percentages.

We consider the differences of each term in the out-of-sample negative log-likelihood,

$$\hat{D}_t = \tilde{U}_{t;\text{sync}} - \tilde{U}_{t;\text{async}}, \quad t = 1, \dots, T,$$

where

$$\sum_{t=1}^T \tilde{U}_{t;\text{model}} = \text{outsample} - \log\text{-likelihood}.$$

Note that, up to a change of signs, \hat{D}_t is the difference between deviance residuals, see McCullagh and Nelder (1989). We are now testing the null hypothesis that the differences \hat{D}_t have mean zero against the alternative of mean less than zero, i.e. the estimates from the synchronous CCC-GARCH(1,1) are better than the ones from the asynchronous classical CCC-GARCH(1,1). For this purpose, we use versions of the t-test and sign-test, adapted to the case of dependent observations. The t-type test statistic is

$$\sqrt{T} \frac{\bar{D}}{\hat{\sigma}_{D;\infty}}, \quad \text{where } \bar{D} = \frac{1}{T} \sum_{t=1}^T \hat{D}_t. \quad (3.2)$$

In (3.2), $\hat{\sigma}_{D;\infty}^2 = (2\pi)\hat{f}_{\hat{D}}(0)$, where $\hat{f}_{\hat{D}}(0)$ is a smoothed periodogram estimate at frequency zero, based on $\hat{D}_1, \dots, \hat{D}_T$; see for example Brockwell and Davis (1991). The motivation for this estimate is based on the assumption that $\{\hat{D}_t\}_t$ is stationary and satisfies suitable dependence

conditions, e.g. mixing. Then

$$\begin{aligned} \sqrt{T}(\bar{D} - E[\widehat{D}_t]) &\implies \mathcal{N}(0, \sigma_{D;\infty}^2) \quad (T \rightarrow \infty), \\ \sigma_{D;\infty}^2 &= \sum_{k=-\infty}^{+\infty} \text{Cov}[\widehat{D}_0, \widehat{D}_k] = (2\pi)f_{\widehat{D}}(0), \end{aligned} \quad (3.3)$$

where $f_{\widehat{D}}(0)$ is the spectral density at zero of $\{\widehat{D}_t\}_t$.

Thus, using (3.3) for the test statistic in (3.2)

$$\sqrt{T} \frac{\bar{D}}{\widehat{\sigma}_{D;\infty}} \implies \mathcal{N}(0, 1) \quad (T \rightarrow \infty) \quad (3.4)$$

under the null hypothesis. The observed value for the test statistic (3.2) equals -1.7077 with a corresponding P -value of 0.044 indicating at the 5% significance level that the synchronous model is better.

The version of the sign test is based on the number of negative differences

$$\widehat{W}_t = I_{\{\widehat{D}_t \leq 0\}}, \quad t = 1, \dots, T,$$

for the null hypothesis that the negative differences \widehat{W}_t have mean $\frac{1}{2}$ against the alternative of mean bigger than $\frac{1}{2}$. The test statistic is given by

$$\sqrt{T} \frac{\bar{W} - \frac{1}{2}}{\widehat{\sigma}_{W;\infty}}, \quad \text{where } \bar{W} = \frac{1}{T} \sum_{t=1}^T \widehat{W}_t \quad (3.5)$$

and $\widehat{\sigma}_{W;\infty}^2$ as in (3.2) but based on $\widehat{W}_1, \dots, \widehat{W}_T$. As in the derivation of the t-type test above, we have

$$\sqrt{T} \frac{\bar{W} - \frac{1}{2}}{\widehat{\sigma}_{W;\infty}} \implies \mathcal{N}(0, 1) \quad (T \rightarrow \infty) \quad (3.6)$$

under the null hypothesis. There are 307 negative differences (of total $T = 500$) and the observed value of the test statistic is 1.661 with a corresponding P -value of 0.048. Therefore, this test leads to the rejection of the null hypothesis at the 5% significance level, implying a preference of the synchronous over the asynchronous CCC-GARCH(1,1) model.

Finally, we illustrate the resulting estimates of the synchronized returns obtained from the synchronous CCC-GARCH(1,1) model.

FIGURE 3.1 ABOUT HERE.

Figure 3.1 indicates that synchronization has primarily an effect on returns of small or medium absolute size. In spite of this, if we consider a sample of the synchronized returns on a day (for example May 8, 1990) with an extreme negative return we see that synchronization increases the value (in absolute terms) of the negative asynchronized returns (6.614 vs. 6.143 for the CAC40 index, 7.111 vs. 6.515 for the DAX index, 3.833 vs. 3.240 for the NIKKEI index and so on) implying that the losses will be underestimated when using the asynchronized returns for the calculation of the portfolio value at New York closing time.

3.3 Estimating the performance at the portfolio level

We now examine the effect of synchronization for the estimation of volatility in a portfolio. Denote by P_t the price of a portfolio at day t

$$P_t = \sum_{i=1}^7 \alpha_i S_{t,i} \quad t = 1, \dots, m. \quad (3.7)$$

This portfolio employs a constant asset division. For illustrative purposes, we use the data $S_{t,i}$ from before and choose $\alpha_1 = 0.4$, $\alpha_2 = \dots = \alpha_6 = 0.08$ and $\alpha_7 = 0.2$, corresponding to the volumes of the different stock exchanges. We also translate all the prices to US dollars, using daily currencies. It is known that the (negative) portfolio returns Δ_t at day t become then a linear combination of the individual (negative) asset returns \mathbf{X}_t

$$\begin{aligned} \Delta_t &= -100 \cdot \left(\frac{P_t - P_{t-1}}{P_{t-1}} \right) = \beta'_{t-1} \mathbf{X}_t, \\ \beta_{t-1,i} &= \alpha_i \frac{S_{t-1,i}}{P_{t-1}}, \quad i = 1, \dots, 7. \end{aligned} \quad (3.8)$$

Our general model for Δ_t is

$$\Delta_t = \mu_{t,P} + \epsilon_t = \mu_{t,P} + \sigma_{t,P} Z_t,$$

where $\mu_{t,P} \in \mathbb{R}$ and $\sigma_{t,P} \in \mathbb{R}^+$ are measurable functions of \mathcal{F}_{t-1} (see (2.1)).

We compare portfolio volatility estimates from multivariate synchronous and asynchronous CCC-GARCH(1,1) models with the ones from a classical GARCH(1,1) univariate analysis (and extensions thereof) for the portfolio returns Δ_t : all give rise to different $\mu_{t,P}$ and $\sigma_{t,P}$ above. Note that in the more realistic case, the weights $\alpha_i = \alpha_{t,i}$ are depending on t . As a consequence, the univariate analysis of the returns Δ_t is inappropriate, because the returns of portfolio prices would typically be far from stationary. For univariate analyses based on returns Δ_t , we always assume the model

$$\Delta_t = \mu_{t,P} + \epsilon_t = \phi \Delta_{t-1} + \sigma_{t,P} Z_t, \quad (3.9)$$

where $\sigma_{t,P}$ is a measurable function of previous returns $\Delta_{t-1}, \Delta_{t-2}, \dots$ and i.i.d. innovations $Z_t \sim \sqrt{(\nu-2)/\nu} t_\nu$. The scaling factor $\sqrt{(\nu-2)/\nu}$ is used so that $\text{Var}(Z_t) = 1$. The univariate GARCH(1,1) specification is

$$\sigma_{t,P}^2 = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta \sigma_{t-1,P}^2, \quad \epsilon_{t-1} = \Delta_{t-1} - \phi \Delta_{t-2},$$

where $\alpha_0, \alpha_1, \beta > 0$. The negative log-likelihood, conditioned on the first observation Δ_1 and some starting value $\sigma_{1,P}$, (e.g. the square root of the sample variance) is then

$$\begin{aligned} -\ell(\phi, \alpha_0, \alpha_1, \beta; \Delta_2^n) &= - \sum_{t=2}^n \log \left(c(\nu)^{-1} \sigma_{t,P}^{-1} f_{t_\nu} \left(\frac{\Delta_t - \phi \Delta_{t-1}}{c(\nu) \sigma_{t,P}} \right) \right), \\ c(\nu) &= ((\nu-2)/\nu)^{1/2}, \end{aligned} \quad (3.10)$$

where f_{t_ν} denotes the density of the univariate t_ν distribution. Minimizing the negative log-likelihood yields estimates $\hat{\mu}_{t,P} = \hat{\phi} \Delta_{t-1}$ and $\hat{\sigma}_{t,P}^2 = \hat{\alpha}_0 + \hat{\alpha}_1 \epsilon_{t-1}^2 + \hat{\beta} \sigma_{t-1,P}^2$.

With the multivariate synchronized CCC-GARCH(1,1) model, we calculate estimates of the portfolio conditional means $\hat{\mu}_{t,P}$ and variances $\hat{\sigma}_{t,P}^2$, $t = 1, \dots, n$ as follows. We always take the innovations \mathbf{Z}_t of the model in (2.7) or (2.8) to be multivariate t_ν distributed (ν unknown)

with zero mean and covariance matrix $\text{Cov}(\mathbf{Z}) = I_M$. Using the representation in Proposition 1, assuming that $(I_M + A)^{-1}$ exists, it follows that the asynchronous returns \mathbf{X}_t given the information up to time $t - 1$ are multivariate t_ν distributed

$$\mathbf{X}_t | \mathcal{F}_{t-1} \sim t_\nu(A \mathbf{X}_{t-1}, (I_M + A)^{-1} \Sigma_t^s ((I_M + A)^{-1} \Sigma_t^s)'),$$

Now, using a nice property of elliptical distributions (see Fang et al., 1990) we find that the portfolio return Δ_t given the information up to time $t - 1$ is univariate t_ν distributed with the following mean and variance:

$$\Delta_t | \mathcal{F}_{t-1} \sim t_\nu(\beta'_{t-1} A \mathbf{X}_{t-1}, \beta'_{t-1} (I_M + A)^{-1} \Sigma_t^s (\beta'_{t-1} (I_M + A)^{-1} \Sigma_t^s)'),$$

where the vector of coefficients β_{t-1} is given in (3.8). Thus, we calculate

$$\begin{aligned} \hat{\mu}_{t,P} &= \beta'_{t-1} \hat{A} \mathbf{X}_{t-1} \quad \text{and} \\ \hat{\sigma}_{t,P}^2 &= \beta'_{t-1} (I_M + \hat{A})^{-1} \hat{\Sigma}_t^s (\beta'_{t-1} (I_M + \hat{A})^{-1} \hat{\Sigma}_t^s)', \end{aligned}$$

where \hat{A} and $\hat{\Sigma}_t^s$ are the maximum likelihood estimates in the model (2.8). The estimates from the classical asynchronous CCC-GARCH(1,1) model are of the same form, but with $\hat{A} = 0$ and $\hat{\Sigma}_t^{\text{asynch}}$.

The predicted portfolio conditional mean and variance estimates with the synchronous and the asynchronous CCC-GARCH(1,1) and the univariate GARCH(1,1) model are plotted in Figure 3.2.

FIGURE 3.2 ABOUT HERE.

The predicted volatilities are substantially larger when using the univariate approach. The differences between the two multivariate methods are (visually) much smaller. The conditional mean in the synchronized model is small (relative to the magnitude of the square root of conditional variance).

Now, we test the goodness of the residuals

$$\hat{Z}_t = \frac{\tilde{\Delta}_t - \hat{\mu}_{t,P}}{\hat{\sigma}_{t,P}}, \quad t = 1, \dots, T, \quad (3.11)$$

in different models. Here $\tilde{\Delta}_t$ is from new test set data $\Delta_{n+1}, \dots, \Delta_{n+500}$ over the next 500 days; $\hat{\mu}_{t,P}$ and $\hat{\sigma}_{t,P}$ are from the different models, estimated with the training data $\Delta_1, \dots, \Delta_n$ but of course evaluated using the immediate lagged values in the test set.

We are particularly interested in the null hypothesis that the dynamics of the (negative) portfolio returns follow model (3.9) against the alternative of a misspecified model. Under the null hypothesis, the statistic $\sqrt{T}\bar{Z}$ is standard normally distributed. The observed values, and the corresponding P -values (given in parenthesis), of the test for the synchronous CCC-GARCH(1,1), the asynchronous CCC-GARCH(1,1) and the univariate model are -1.859 (0.063), -2.201 (0.0277) and -1.534 (0.125), respectively. It shows that the model assumed for the portfolio returns is rejected only by the asynchronous CCC-GARCH(1,1) model.

For quantifying the goodness of fit of the models, we consider again the outsample log-likelihood performance

$$\text{portfolio outsample negative log-likelihood: } -\log\text{-likelihood}\left(\tilde{\Delta}_1^T; \hat{\mu}_{t,p}, \hat{\sigma}_{t,P}, \hat{\nu}\right),$$

where, as in (3.11), $\tilde{\Delta}_1^T = \Delta_{n+1}^{n+500}$ are new test set data and $\hat{\mu}_{t,P}, \hat{\sigma}_{t,P}$ are estimated from the training data. The log-likelihood is from (3.10). In addition, we also consider the following statistics:

$$\begin{aligned} \text{IS-PL}_i &= \sum_{t=1}^n |\hat{\sigma}_{t,P}^2 - (\Delta_t - \hat{\mu}_{t,P})^2|^i, \quad i = 1, 2, \quad (\text{in-sample prediction loss}), \\ \text{OS-PL}_i &= \sum_{t=1}^T |\hat{\sigma}_{t,P}^2 - (\tilde{\Delta}_t - \hat{\mu}_{t,P})^2|^i, \quad i = 1, 2, \quad (\text{out-sample prediction loss}). \end{aligned}$$

The OS-PL statistics and the portfolio outsample log-likelihood are measures for predictive performance. For this reason, we consider them as more important than the in-sample IS-PL statistics.

These statistics are given in Table 3.2 for the multivariate synchronous and asynchronous CCC-GARCH(1,1) model and for the univariate GARCH(1,1) model.

TABLE 3.2 ABOUT HERE.

Analogously to the results obtained in the Section 3.2, we also find at the portfolio level that the synchronous CCC-GARCH(1,1) model is better (with respect to all goodness of fit measures) than the asynchronous classical CCC-GARCH(1,1) model and than the univariate approach. We can also see that the improvement in this case is more relevant than at the multivariate level, in particular with respect to the OS-PL performances (5-10% less). Also, multivariate modelling shows substantial advantages over the univariate GARCH(1,1) analysis, even without using synchronization.

As mentioned in Section 3.2, more impressive gains may be masked by a low signal to noise ratio. Similar t- and sign-type tests to the ones in (3.2)-(3.4) and (3.5)-(3.6) can be done for the goodness of fit statistics at the portfolio level. The results are given in Table 3.3.

TABLE 3.3 ABOUT HERE.

As we can see from Table 3.3, the t-type test never yields significant differences between the models. This may be just a fact of low power due to non-Gaussian observations. On the other hand, the sign-type test which is robust against deviations from Gaussianity yields very significant results: the synchronous is better than the asynchronous CCC-GARCH(1,1) model, and multivariate is better than univariate modelling.

3.4 More sophisticated asynchronous models as another source for improvements?

We question here, whether the improvements with the synchronous CCC-GARCH(1,1) model could also be achieved or even surpassed by more sophisticated models for volatilities or conditional means. Of course, the synchronous CCC-GARCH(1,1) model is also a more complex model for asynchronous data (see Proposition 1), but motivated from the view of synchronization with a simple linear transform.

The first extension of the asynchronous classical GARCH(1,1) model is to include a conditional mean term

$$\begin{aligned} \mathbf{X}_t &= \mu_t + \Sigma_t \mathbf{Z}_t \quad (t \in \mathbb{Z}), \\ \mu_t &= B \cdot X_{t-1}, \quad B = \text{diag}(b_1, \dots, b_M), \end{aligned} \tag{3.12}$$

where \mathbf{Z}_t and Σ_t are exactly as in the model (2.7).

The model in (3.12) can then be further extended to more complex, potentially high dimensional parameterizations for approximating more general volatility functions, such as for example piecewise linear functions. Fitting “best” piecewise linear functions is computationally prohibitive. A tree-structured (sub-optimal) technique for univariate series has been proposed by Audrino and Bühlmann (2001). The same idea can also be used for every individual series in the multivariate set-up.

Considering the same goodness of fit measures already introduced in this paper for the model in (3.12) and its extensions, we see that the main implication is that more sophisticated models (not being of synchronization type) do not show worthwhile improvements (for example for the model (3.12) we find a value of 472.4394 for the portfolio out-of-sample negative log-likelihood statistic and a value of 243.4545 for the OS-PL₂ statistic). Hence, synchronization seems to play a more substantial effect than trying to improve the model-dynamics as in (3.12) or its extensions.

4 Estimating risk measures

We consider here the effect of volatility estimates for conditional (dynamical) risk measures for negative returns Δ_t from a portfolio as in (3.8)-(3.9) given the information \mathcal{F}_{t-1} from previous prices. The most popular risk measures are the *Value at Risk* (VaR) given \mathcal{F}_{t-1}

$$\delta_q^t = \inf\{\delta \in \mathbb{R} : F_{\Delta_t|\mathcal{F}_{t-1}}(\delta) \geq q\}, \quad 0 < q < 1,$$

where $F_{\Delta_t|\mathcal{F}_{t-1}}(\cdot)$ denotes the cumulative distribution function of Δ_t given \mathcal{F}_{t-1} . This is the quantile of the predictive distribution of the negative portfolio return over the next day. Another measure of risk is the *expected shortfall* given \mathcal{F}_{t-1}

$$S_q^t = E[\Delta_t \mid \Delta_t > \delta_q^t, \mathcal{F}_{t-1}], \quad 0 < q < 1.$$

We typically choose $q \in \{0.90, 0.95, 0.99\}$ (note that we consider negative returns). Unlike VaR, the expected shortfall is coherent and satisfies a sub-additivity requirement, see Artzner et al. (1999). The illustrations are for the real global portfolio already introduced in Section 3.3.

4.1 The estimates

We assume that the dynamics of the negative asynchronized portfolio returns Δ_t ($t \in \mathbb{Z}$) are given by (3.9). Since

$$F_{\Delta_t|\mathcal{F}_{t-1}}(\delta) = P[\mu_{t,P} + \sigma_{t,P} Z_t \leq \delta \mid \mathcal{F}_{t-1}] = F_Z\left(\frac{\delta - \mu_{t,P}}{\sigma_{t,P}}\right),$$

the risk measures can then be written as

$$\begin{aligned} \delta_q^t &= \mu_{t,P} + \sigma_{t,P} z_q, \quad 0 < q < 1 \text{ and} \\ S_q^t &= \mu_{t,P} + \sigma_{t,P} E[Z \mid Z > z_q], \quad 0 < q < 1, \end{aligned}$$

where z_q is the q -th quantile of $F_Z(\cdot)$ which by assumption does not depend on time t .

Now, estimates for the VaR and for the expected shortfall are constructed using the assumption of scaled t_ν distributed innovations Z_t in (3.9), i.e. $Z_t \sim \sqrt{(\nu-2)/\nu} t_\nu$. Thus, an estimate for the VaR is given by

$$\widehat{\delta}_q^t = \widehat{\mu}_{t,P} + \widehat{\sigma}_{t,P} \sqrt{\frac{\widehat{\nu}-2}{\widehat{\nu}}} \widetilde{z}_q,$$

and an estimate of the expected shortfall is given by

$$\widehat{S}_q^t = \widehat{\mu}_{t,P} + \widehat{\sigma}_{t,P} \sqrt{\frac{\widehat{\nu} - 2}{\widehat{\nu}}} \left(\frac{1}{1 - q} c \frac{\widehat{\nu}}{\widehat{\nu} - 1} \left(1 + \frac{(\widetilde{z}_q)^2}{\widehat{\nu}} \right)^{\frac{1 - \widehat{\nu}}{2}} \right),$$

where the constant c equals $\Gamma(\frac{1}{2}(\widehat{\nu} + 1))/\Gamma(\frac{\widehat{\nu}}{2})(\widehat{\nu}\pi)^{-1/2}$, \widetilde{z}_q is the q -th quantile of a standard t_ν distributed random variable and $\widehat{\nu}$ the maximum likelihood estimate from the multivariate or univariate models as before. As an alternative method, one could use extreme value theory to model the tails of $F_Z(\cdot)$; for a detailed description, see Embrechts et al. (1997).

We only show here estimates of the conditional expected shortfall, for illustrative purposes. The results, using the multivariate synchronous and asynchronous CCC-GARCH(1,1) and the univariate GARCH(1,1) model are shown in Figure 4.1 (note that we consider negative returns).

FIGURE 4.1 ABOUT HERE.

The multivariate models produce more progressive estimates; and the synchronous model exhibits more small scale fluctuations. Also, the results from Section 3 support that the multivariate models are better. Putting together we gain evidence that the more progressive behavior is appropriate. The better performance of the synchronized model in Section 3 also suggests that the small scale movements of the corresponding estimate of the expected shortfall, which wiggles more, is a good feature.

4.2 Backtesting

Backtesting the results for the expected shortfall is generally very difficult since a tail phenomenon is involved. As a descriptive tool, rather than a formal test, we show boxplots of residuals

$$\widehat{R}_t = \frac{\Delta_t - \widehat{S}_q^t}{\widehat{\sigma}_{t,P}} I_{\{\Delta_t > \delta_q^t\}}. \quad (4.1)$$

in Figure 4.2.

FIGURE 4.2 ABOUT HERE.

Under our model assumptions (3.9) and ignoring estimation effects, we can easily show that these (theoretical) residuals are an i.i.d. sequence with expected value zero. Figure 4.2 yields then additional evidence that the estimate from the univariate model is too conservative (too low values of residuals) and it favors the more progressive behavior of the multivariate models.

5 Conclusions

The need to synchronize multivariate financial time series is strongly motivated by the fact that information continues to flow for closed markets while others are still open. Besides the nice interpretative structure of synchronization, we found empirically that it improved the predictive performance of the CCC-GARCH(1,1) model for a 7-dimensional time series of daily exchange rate returns. The predictive gain with synchronization seems also much more prominent than extending the GARCH to a more complex model for approximating more general functions of volatility.

For analyzing returns of a univariate portfolio, multivariate analysis of the individual returns still plays a key role. If the portfolio weights change over time, which is most often the case in practice, only the multivariate framework allows to treat such a time-changing portfolio,

assuming that the individual returns are stationary. But even for a time-constant portfolio, there is a substantial gain when using a multivariate CCC-GARCH(1,1) model for individual returns instead of a univariate GARCH(1,1) model for portfolio returns. Moreover, the univariate analysis of the portfolio returns seems to be too conservative for the calculation of risk measures (such as for example the expected shortfall) if confronted with the risk estimates from our multivariate synchronous CCC-GARCH(1,1) model. It emphasizes the importance of the multivariate approach in general and the wide range of applications where synchronization can be very valuable.

Acknowledgements: We thank Michel Dacorogna for some interesting remarks. The data was provided by Olsen & Associates, Zürich.

References

- [1] Artzner, P., Delbaen, F., Eber, J.-M. and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance* **3**, 203–228.
- [2] Audrino, F. and Bühlmann, P. (2001). Tree-structured GARCH models. *Journal of the Royal Statistical Society, Series B*, **63**, Part 4, 727–744.
- [3] Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. *The Review of Economics and Statistics* **72**, 498–505.
- [4] Brockwell, P.J. and Davis, R.A. (1991). *Time Series: Theory and Methods*. Springer, New York.
- [5] Burns, P., Engle, R.F. and Mezrich, J. (1998). Correlations and volatilities of asynchronous data. *Journal of Derivatives*, Summer, 1–12.
- [6] Efron, B. and Tibshirani, R.J. (1993). *An Introduction to the Bootstrap*. Chapman & Hall, London.
- [7] Embrechts, P., Klüppelberg, C. and Mikosch, T. (1997). *Modelling extremal events for insurance and finance*. Springer, Berlin.
- [8] Fang, K.T., Kotz, S. and Ng, K.W. (1990). *Symmetric multivariate and related Distributions*. Chapman & Hall, London.
- [9] McCullagh, P. and Nelder, J.A. (1989). *Generalized linear Models*. Chapman & Hall, London.
- [10] Reinsel, G.C. (1991). *Elements of Multivariate Time Series Analysis*. Springer, New York.




Figure 3.1: Asynchronized returns $X_{t_i,i}$ against estimated synchronized returns $\widehat{X}_{t_i,i}^s$ from the synchronous CCC-GARCH(1,1) model (using \widehat{A} in (3.1)) for the French CAC40, the German DAX, the Italian BCI, the Dutch CBS, the British FTAS and the Japanese NIKKEI index.




Figure 3.2: Predicted portfolio conditional variances $\hat{\sigma}_{t,P}^2$ (top) and means $\hat{\mu}_{t,P}$ (center and bottom) for the test data days between June 23, 1994 and September 9, 1996 (500 days) using the univariate GARCH(1,1) model (3.9) for the negative portfolio returns (solid line), the synchronous CCC-GARCH(1,1) model (dotted line) and the standard CCC-GARCH(1,1) model without synchronization (dashed line).

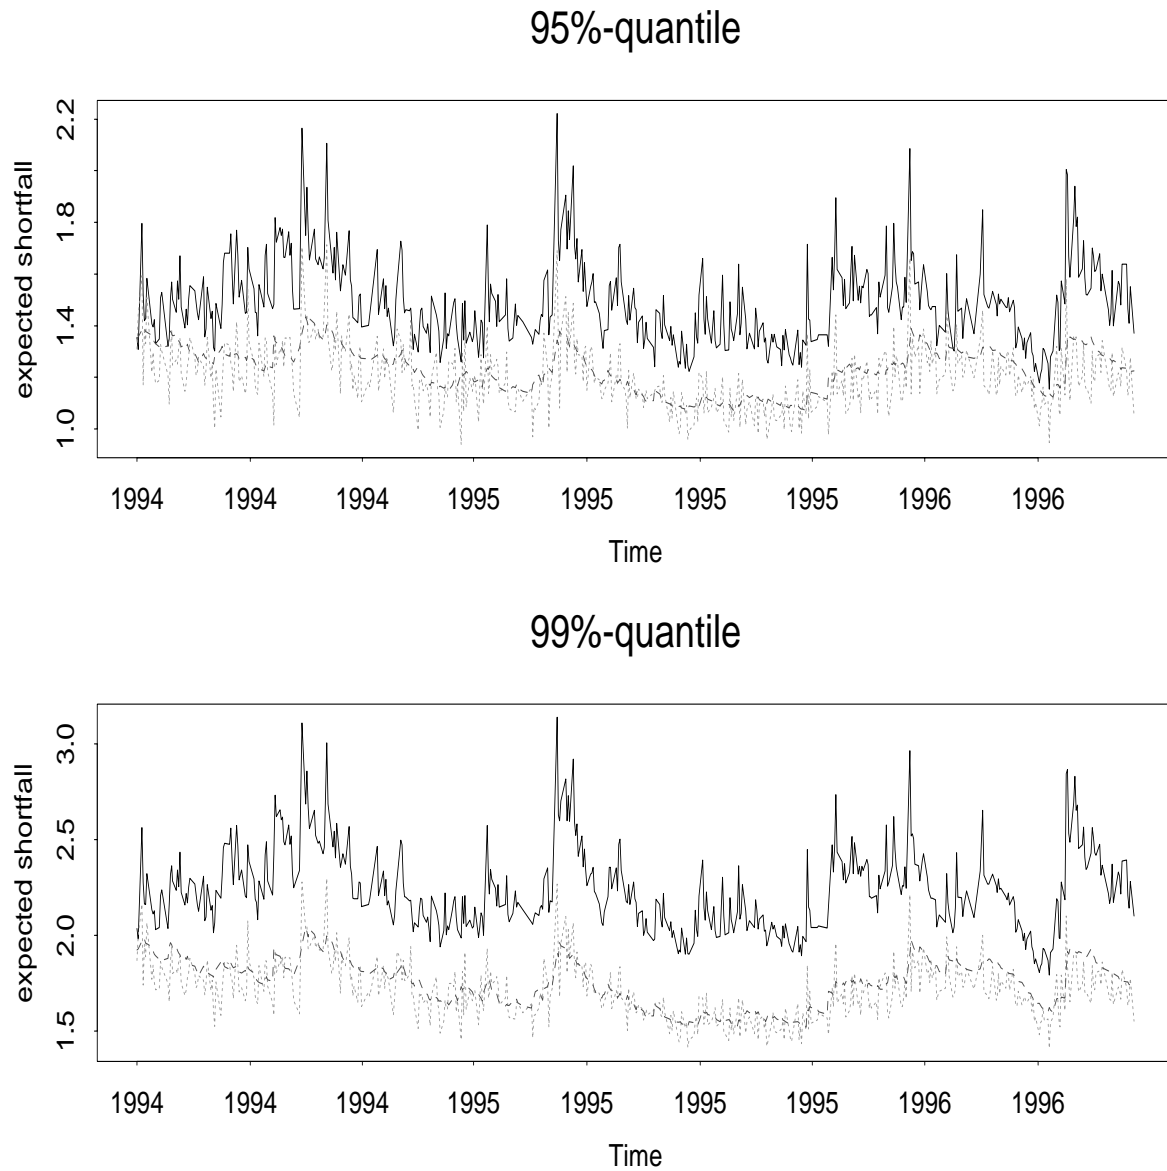


Figure 4.1: The expected shortfall for the portfolio returns $\tilde{\Delta}_t$ for the test data days beginning on June 23, 1994 and ending on September 9, 1996. The estimates \tilde{S}_t^q for $q = 0.95$ (top) and $q = 0.99$ (bottom) are obtained using the univariate GARCH(1,1) model (solid line), the synchronous CCC-GARCH(1,1) model (dotted line) and the standard CCC-GARCH(1,1) model without synchronization (dashed line).

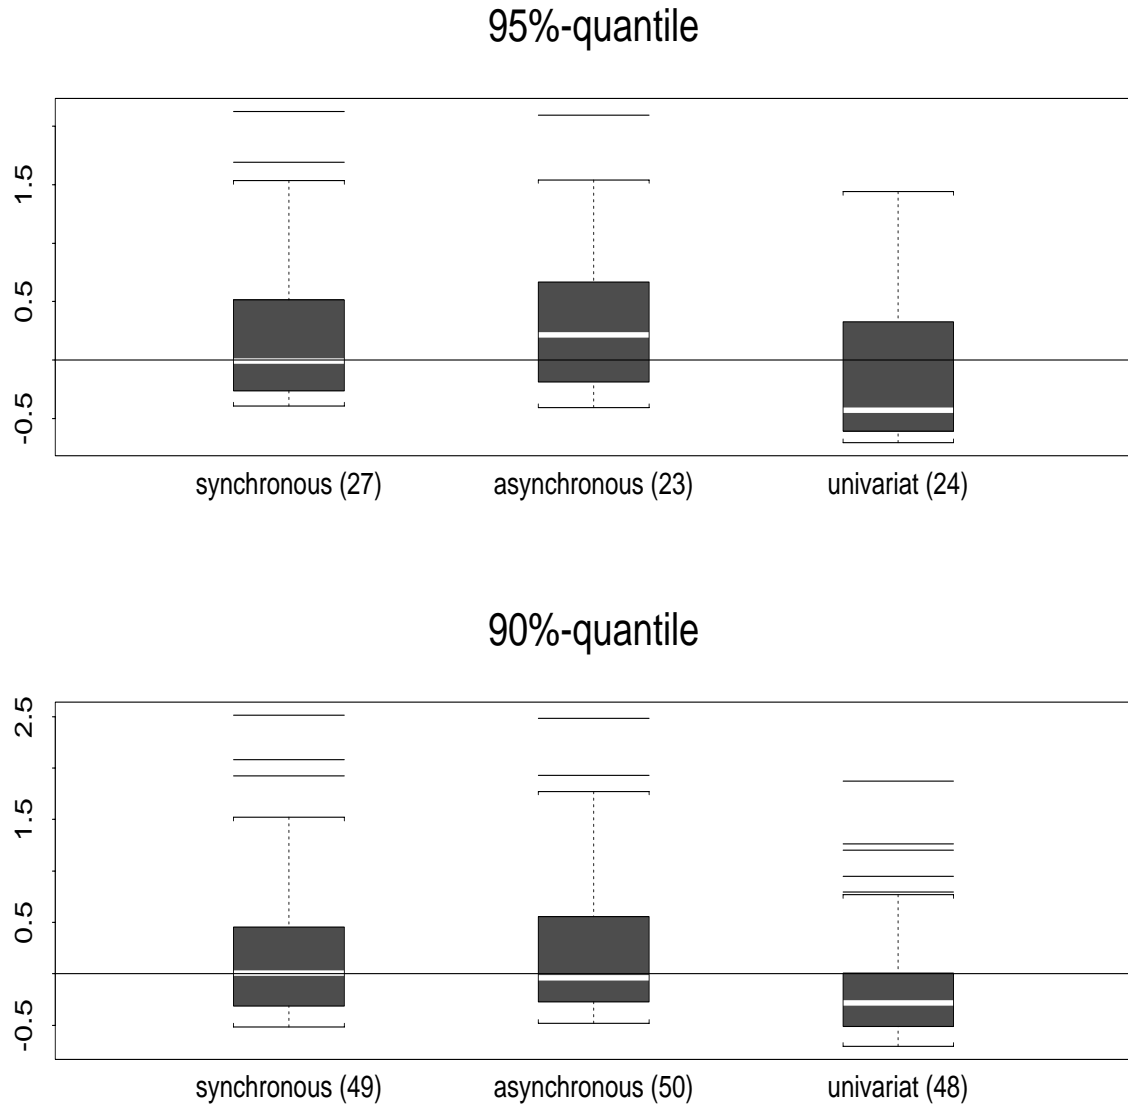


Figure 4.2: Boxplots of residuals \hat{R}_t in (4.1) for test data with $q = 0.95$ (top) and $q = 0.90$ (bottom) from the synchronous CCC-GARCH(1,1) model (left), the standard CCC-GARCH(1,1) model without synchronization (center) and the univariate GARCH(1,1) model (right). The number of violations $\Delta_t > \hat{\delta}_q^t$ is given in parenthesis.

	DJIA	CAC40	DAX	BCI	CBS	FTAS	NIK
DJIA	1	0.26569	0.23928	0.11794	0.13863	0.21754	0.18972
CAC40	0.26569	1	0.75794	0.51368	0.76965	0.70239	0.31629
DAX	0.23928	0.75794	1	0.55304	0.76055	0.59974	0.32281
BCI	0.11794	0.51368	0.55304	1	0.56079	0.47805	0.25427
CBS	0.13863	0.76965	0.76055	0.56079	1	0.73521	0.29812
FTAS	0.21754	0.70239	0.59974	0.47805	0.73521	1	0.30296
NIK	0.18972	0.31629	0.32281	0.25427	0.29812	0.30296	1

	DJIA	CAC40	DAX	BCI	CBS	FTAS	NIK
DJIA	1	0.36616	0.38101	0.27009	0.29308	0.31922	0.32770
CAC40	0.36616	1	0.80167	0.56694	0.78070	0.70587	0.31191
DAX	0.38101	0.80167	1	0.57977	0.77477	0.62009	0.28068
BCI	0.27009	0.56694	0.57977	1	0.58814	0.53707	0.28324
CBS	0.29308	0.78070	0.77477	0.58814	1	0.73965	0.28781
FTAS	0.31922	0.70587	0.62009	0.53707	0.73965	1	0.28548
NIK	0.32770	0.31191	0.28068	0.28324	0.28781	0.28548	1

Table 3.1: Top: instantaneous empirical correlations between components of asynchronous returns \mathbf{X}_t . Bottom: instantaneous empirical correlations between components of estimated synchronized returns $\widehat{\mathbf{X}}_t^s$ using the synchronous CCC-GARCH(1,1) model with $\widehat{A}_{\text{synchron.}}$ in (3.1).

	synchronous CCC-GARCH(1,1)	CCC-GARCH(1,1) without synchronization	Univariate approach
IS-PL ₁	534.0402	573.9434	546.91
IS-PL ₂	1424.506	1732.441	1606.679
OS-PL ₁	185.2044	195.5282	185.9926
OS-PL ₂	227.3589	240.7812	244.8054
Portfolio out. log-likelihood	460.8094	470.5776	474.4924

Table 3.2: Goodness of fit for the synchronous CCC-GARCH(1,1) model, the CCC-GARCH(1,1) model without synchronization and the univariate model defined in (3.9) for the portfolio introduced in Section 3.

Model 1	Model 2	Performance measure	
		Portfolio out. log-likelihood	OS-PL ₂
Synchronous CCC-GARCH(1,1)	Asynchronous CCC-GARCH(1,1)	-1.3177 (0.094)	-1.2956 (0.098)
Synchronous CCC-GARCH(1,1)	Univariate approach	-0.7659 (0.222)	-1.4904 (0.068)
Asynchronous CCC-GARCH(1,1)	Univariate approach	0.5207 (0.301)	0.6143 (0.270)

Model 1	Model 2	Performance measure	
		Portfolio out. log-likelihood	OS-PL ₂
Synchronous CCC-GARCH(1,1)	Asynchronous CCC-GARCH(1,1)	4.386 ($6 \cdot 10^{-6}$)	5.043 ($2 \cdot 10^{-7}$)
Synchronous CCC-GARCH(1,1)	Univariate approach	6.835 ($4 \cdot 10^{-12}$)	10.95 ($3 \cdot 10^{-28}$)
Asynchronous CCC-GARCH(1,1)	Univariate approach	4.054 ($3 \cdot 10^{-5}$)	13.54 ($5 \cdot 10^{-42}$)

Table 3.3: Testing differences of performance on the portfolio level. Given are the values of test statistics and corresponding P -values (in parenthesis). Top: t-type test (analogously to (3.2)-(3.4)). Bottom: sign-type test (analogously to (3.5) and (3.6)).