Synchronizing Multivariate Financial Time Series

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

Revised version: September 2003

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 given by the fact that information continues to flow for closed markets while others are still open. We propose a synchronization technique, including further modelling of synchronized returns simultaneously: the overall method can then be described in terms of a new model for asynchronous returns.

Besides the nice interpretation of synchronization, we found empirically that the method potentially increases the predictive performance of many reasonable models for a seven-dimensional time series of daily equity index returns and is more appropriate for the calculation of portfolio risk measures.

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, varies because not all markets have the same trading hours. For example, between the US and Japan, there are no common opening hours, and between the US and Europe, there is only partial overlap. As a consequence, correlations across the assets are often too small[1] when using such asynchronous data. Therefore, the value of real global portfolios constructed on daily data across different markets is never known at a fixed point in time and 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) may lead to inaccurate and misleading results.

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 asset values may change before the market re-opens. Synchronizing data involves estimates of asset values at a specified (synchronization) time point for every day[2]. The estimated asset values at the same synchronization time across markets are then called *synchronized*. Unlike Burns et al. (1998) who do not test the forecasting power of the obtained synchronized data, we propose

[1]See, for example, Burns et al. (1998).

[2]For instance, we always use the closing time of the New York stock exchange, i.e. 4 pm local New York time, as synchronization time point.

to use any reasonable multivariate model for the constructed synchronized data and to test their out-of-sample predictive performance for applications in risk management. In particular, to model the dynamics of the synchronized data, we consider the CCC-GARCH(1,1) model introduced by Bollerslev (1990) allowing for time varying conditional variances and covariances but imposing constant conditional correlations. The CCC-GARCH(1,1) model for synchronized data represents a different and new model in terms of the original asynchronous data. This new model is called *synchronous CCC-GARCH(1,1)* and allows an estimation of synchronization and model parameters in a simultaneous way.

In our empirical investigations, we compare the performance results obtained using synchronized and asynchronous data for a real global portfolio with daily log-returns of seven equity indices around the world. When using synchronized data the resulting gains are sometimes considerable, depending on how we measure performance. Thus, this finding supports the usefulness of synchronizing data in a first step. Our model for synchronized returns also yields some improvements over the synchronization method from Burns et al. (1998).

Moreover, our multivariate modelling approach, taking into account the synchronization of the data, is superior to univariate models for a portfolio index which are exposed to an information loss by averaging previous individual prices[3]. In contrast, we show that no gain over a simple univariate GARCH(1,1) model for the portfolio returns can be achieved using a multivariate GARCH model without synchronization, confirming the recent results of Berkowitz and O'Brien (2002). We also find that univariate modelling yields too conservative risk estimates (validated by using back-tests). For example, the capital needed to cover possible portfolio losses is usually overestimated by univariate modelling.

The plan of the paper is as follows. Section 2 presents our synchronization model and the corresponding estimation procedure. The empirical goodness of fit results for a real global portfolio of seven equity indices around the world are summarized in Section 3. Results are computed using our model in comparison with other standard approaches for synchronized and asynchronous data. In Section 4, we discuss the impact of synchronization on the calculation of risk measures such as the VaR and expected shortfall for the same seven-dimensional real data example. Section 5 includes a summary and presents our conclusion.

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

2.1 Synchronization of the data

Our synchronization technique follows very closely that proposed by Burns et al. (1998). As an illustrative example, consider a global portfolio, including stocks traded in New York and London. At the closing time of trade in New York, the value of the portfolio should be measured with an estimate of the value of the London stocks (at the New York closing time). For example, to take the closing prices of the London stocks at a day when the US market goes down 1 percent (after London has closed) is inappropriate for pricing the portfolio at New York closing time[4]. We associate with synchronization some estimates of the prices of the share traded in London at the closing time in New York[5].

[3]Moreover, note that in the more realistic case where the portfolio weights changes over time, stationarity of all individual prices does not imply stationarity of the portfolio prices which makes the direct univariate portfolio modelling difficult.

[4]It will follow that the US shares in the portfolio decline today, while London's will decline tomorrow, mainly due to the asynchronous trade at different places.

[5]Clearly, from the viewpoint of a British investor, the data could also be synchronized at the closing time in London. The choice of the synchronization time point is arbitrary.

We denote by $S_{t,j},\ j = 1, \cdots M$ 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 on 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 4 hours before New York at 5:00 pm, the observed closing price of asset 2 in London on the first day would be denoted by $S_{0.83,2}$. This is also illustrated by the following Figure taken from Burns et al. (1998).

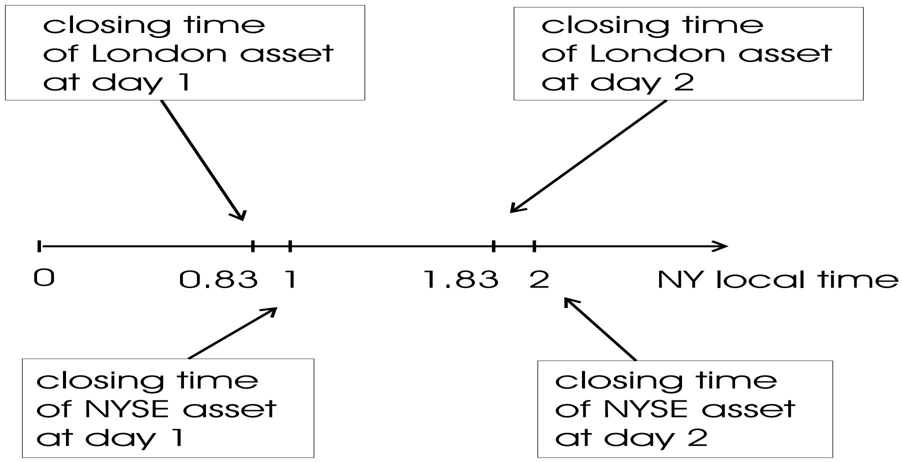

Generally, the observed data is taken at closing times of different markets. It has the structure

$$S_{t_j,j}\ (j = 1, \cdots, M),\ \text{where}$$
$$t_j = t_1 - c_j\ (0 \le c_j < 1),\ j = 1, \ldots, M.$$

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

Let us define the *synchronized prices* $S^s_{t,j}$ by

$$\log\left(S^s_{t,j}\right) = \mathbb{E}\Big[\log\left(S_{t,j}\right) \mid \mathcal{F}_t\Big],\ \text{where}\ \mathcal{F}_t = \Big\{S_{t_j,j}\ ;\ t_j \le t,\ j = 1, \ldots, M\Big\}, \tag{2.1}$$

where the logarithms are used to be consistent with continuously compounded returns. Hence, the synchronized log-prices are defined as the best predicted log-prices at t given the complete information $\mathcal{F}_t$ of all *recorded* prices up to time t. Note that $\mathcal{F}_t$ contains only the prices $S_{t_j,j}$ with closing times $t_j \le t$, and often with a strict relation $t_j < t$ if the trading place for equity j has a closing time other than the first equity in New York.

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 from predictions at time t to predicted prices at $t_j + 1$ are unpredictable:

$$\log\left(S^s_{t,j}\right) = \mathbb{E}\Big[\log\left(S_{t,j}\right) \mid \mathcal{F}_t\Big] = \mathbb{E}\Big[\log\left(S_{t_j+1,j}\right) \mid \mathcal{F}_t\Big],\quad t_j \le t < t_j + 1\ (t \in \mathbb{N}). \tag{2.2}$$

The first equality holds by the definition in (2.1). As we will see, the approximation (2.2) allows us to derive the main synchronization formula (2.7).

We denote the vector of negative log-returns (in percentages)[6], in different markets and at various time points on day t by $\mathbf{X_t}$,

$$\mathbf{X_t} = -100 \cdot \begin{pmatrix} \log\left(\frac{S_{t_1,1}}{S_{t_1-1,1}}\right) \\ \vdots \\ \log\left(\frac{S_{t_M,M}}{S_{t_M-1,M}}\right) \end{pmatrix} = -100 \cdot \Big(\log\left(\mathbf{S_t}\right) - \log\left(\mathbf{S_{t-1}}\right)\Big), \tag{2.3}$$

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

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

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

The synchronized returns depend 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

$$\mathbf{X_t} = A \cdot \mathbf{X_{t-1}} + \epsilon_t \ , \tag{2.5}$$

with errors ϵ_t such that $\mathbb{E}[\epsilon_t \mid \mathcal{F}_{t-1}] \equiv 0$, and A an $M \times M$ matrix. Contrary to the approach of Burns et al. (1998), which proposes a first order vector moving average, we choose a first order vector autoregressive synchronization that considerably simplifies the analysis, since $\mathbb{E}[\mathbf{X_t}|\mathcal{F}_{t-1}]$ in (2.5) depends only on the previous $\mathbf{X_{t-1}}$ (due to the Markovian structure of an autoregressive model). As we will show in (2.7), 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 $\mathbf{X_{t-1}}, \mathbf{X_{t-2}}, \ldots$ as in Burns et al. (1998). Moreover, we will also provide empirical evidence that our synchronization model in connection with (2.5) is better (in terms of predictive performance) than using a first order vector moving average.

Substituting (2.2) into (2.4) yields the synchronized returns

$$\mathbf{X}_t^s = -100 \cdot \Big(\log\left(\mathbf{S}_t^s\right) - \log\left(\mathbf{S}_{t-1}^s\right)\Big) = -100 \cdot \Big(\mathbb{E}\left[\log\left(\mathbf{S_{t+1}}\right) \mid \mathcal{F}_t\right] - \mathbb{E}\left[\log\left(\mathbf{S_t}\right) \mid \mathcal{F}_{t-1}\right]\Big). \tag{2.6}$$

Now, we want the synchronized returns to be written only in terms of asynchronous returns. For this purpose, we add and subtract on the right hand side of (2.6) the terms $\mathbb{E}\left[\log\left(\mathbf{S_t}\right) \mid \mathcal{F}_t\right] = \log\left(\mathbf{S_t}\right)$ for t and $t-1$. Substituting (2.3) and (2.5) into (2.6) then yields

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

and thus

$$\mathbf{X}_t^s = \mathbf{X_t} + A \cdot \Big(\mathbf{X_t} - \mathbf{X_{t-1}}\Big) \ . \tag{2.7}$$

[6]Analogously to McNeil and Frey (2000), we use the negative returns, i.e. the losses, because they are of major interest in risk analysis.

The synchronized returns equal the asynchronous returns plus a correction, which consists of linear combinations of increments from time point $t-1$ to t, representing some aspects of the dynamics of the multivariate return process. 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 (2.7) boils down to estimation of A in model (2.5) or a more specific version. An estimation procedure for A 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), allowing for time varying conditional variances and covariances but imposing constant conditional correlations:

$$\begin{aligned} \mathbf{X}_t^s &= \mu_t^s + \epsilon_t^s = \mu_t^s + \Sigma_t^s\ \mathbf{Z}_t \quad (t \in \mathbb{Z})\ , \\ \mathbf{X}_t^s &= \mathbf{X_t} + A\ \big(\mathbf{X_t} - \mathbf{X_{t-1}}\big) = (I_M + A)\ \mathbf{X_t} - A\ \mathbf{X_{t-1}} \quad \text{as in (2.7)}, \end{aligned} \tag{2.8}$$

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 $\mathrm{Cov}(\mathbf{Z}_t) = I_M$ 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 $i, j = 1, \ldots, M$;

(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 $\alpha_0^{(i)}, \alpha_1^{(i)}$, $\beta^{(i)} > 0,\ \ \alpha_1^{(i)} + \beta^{(i)} < 1$ for $i = 1, \ldots, M$;

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

Proposition 1.
Assume that the matrix $(I_M + A)$ is invertible. Then, the synchronous CCC-GARCH(1,1) model (2.8) can be represented with asynchronous returns $\mathbf{X_t}$ as

$$\mathbf{X_t} = A \cdot \mathbf{X_{t-1}} + (I_M + A)^{-1}\ \Sigma_t^s\ \mathbf{Z_t}\ , \tag{2.9}$$

where the matrix Σ_t^s has the same CCC-GARCH(1,1) structure already defined in (2.8).

Proposition 1 implies that the synchronous CCC-GARCH(1,1) model is still a constant conditional correlation model in terms of asynchronous 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 (2.7), the fact that $\mathbf{X_{t-1}} \in \mathcal{F}_{t-1}$ and (2.5), we calculate the conditional mean of the synchronized returns as

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

It follows that (2.8) is equivalent to

$$\mathbf{X}_t^s = A^2 \cdot \mathbf{X_{t-1}} + \Sigma_t^s \ \mathbf{Z_t}.$$

Using (2.7), we obtain the assertion:

$$\begin{aligned} & \mathbf{X}_t^s = \big(I_M + A\big) \cdot \mathbf{X_t} - A \cdot \mathbf{X_{t-1}} = A^2 \cdot \mathbf{X_{t-1}} + \Sigma_t^s \ \mathbf{Z_t} \\ \Longleftrightarrow \quad & \mathbf{X_t} = \big(I_M + A\big)^{-1} \big(A + A^2\big) \cdot \mathbf{X_{t-1}} + \big(I_M + A\big)^{-1} \ \Sigma_t^s \ \mathbf{Z_t} \\ \Longleftrightarrow \quad & \mathbf{X_t} = A \cdot \mathbf{X_{t-1}} + \big(I_M + A\big)^{-1} \ \Sigma_t^s \ \mathbf{Z_t} \ . \end{aligned}$$

□

2.3 Estimating the model

Model structure. The synchronous CCC-GARCH(1,1) model involves the matrix A: we insist on sparseness by setting the matrix elements to zero if they are found to be statistically insignificant. This is important to reduce the number of parameters in the case of high-dimensional portfolios with 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 model (2.8).

Step 1. Find the estimates for the M^2 parameters of the matrix A and for the matrix Σ using the Yule-Walker estimator[7]. The Yule-Walkers covariance relations for a multivariate AR(1) model are given by

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

Calculate some model-based standard errors of the estimated elements of A using a bootstrap strategy as follows. Compute residuals

$$\widehat{\mathbf{Z}}_t = (\widehat{\Sigma}_t^s)^{-1} (\widehat{\mathbf{X}}_t^s - \widehat{\mu}_t^s),$$

where $\widehat{\mathbf{X}}_t^s$ and $\widehat{\mu}_t^s$ involve the Yule-Walker estimate $\widehat{A}$ and $\widehat{\Sigma}_t^s$ involves the estimates $\hat{\alpha}_0^{(i)}, \hat{\alpha}_1^{(i)}, \hat{\beta}^{(i)}$ $(i = 1, \dots M)$ and $\hat{\rho}_{ij}$ $(i, j = 1, \dots, M)$. Now do an i.i.d. resampling from the empirical distribution of the residuals $\widehat{\mathbf{Z}}_t$ to obtain

$$\mathbf{Z}_1^*, \mathbf{Z}_2^*, \dots \mathbf{Z}_n^*,$$

and generate recursively the bootstrap sample

$$\begin{aligned} & \mathbf{X}_t^* = (I_M + \widehat{A})^{-1} (\mathbf{X}_t^{s*} + \widehat{A} \mathbf{X}_{t-1}^*), \\ & \mathbf{X}_t^{s*} = \widehat{\mu}_t^{s*} + \Sigma_t^{s*} \mathbf{Z}_t^*, \\ & \widehat{\Sigma}_t^{s*} \text{ satisfies (A2) and (A3) from Section 2.2 with estimated parameters} \\ & \qquad \text{and the lagged } X_{t-1,i}^{s*} \text{ and } \sigma_{t-1,i}^{s*}, \\ & \widehat{\mu}_t^{s*} \text{ satisfies (A4) from Section 2.2 with estimated } \widehat{A} \text{ and lagged } \mathbf{X}_{t-1}^*. \end{aligned}$$

[7] For more details on the Yule-Walker estimator, see Brockwell and Davis (1991) or Reinsel (1991).

This is a semi-parametric model-based bootstrap, related to an early idea in Freedman (1984). Now, calculate standard errors

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

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-statistics

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

and $A_{ij} = 0$ for all j, and i corresponding to the New York stocks (in our case $i = 1$).

Parameter estimation. The parameters A, $\alpha_0^{(j)}$, $\alpha_1^{(j)}$, $\beta^{(j)}$, ρ^s_{ij} $(j = 1, \cdots, M)$ in the synchronous CCC-GARCH(1,1) model (2.9) can be estimated with the maximum likelihood method. We 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 number of degrees of freedom ν also has to be estimated, i.e. $\mathbf{Z}_t \sim t_\nu(0, I_M)$. The negative log-likelihood is then given by

$$-l(\theta; \mathbf{X}_2^n) = \frac{TM}{2}\log(\pi\nu) - T\log\Big(\frac{\Gamma(\frac{M+\nu}{2})}{\Gamma(\frac{\nu}{2})}\Big) + \frac{T}{2}\log|\ R^s\ | + \sum_{t=1}^{T}\Big(\log|\ \sqrt{(\nu-2)/\nu}D_t^s\ |\Big) +$$
$$+\frac{M+\nu}{2}\sum_{t=2}^{T}\left(\log\Big(1 + \frac{(\epsilon_t^s)'(R^s)^{-1}\epsilon_t^s}{\nu}\Big)\right) - T\log\big(\ |\ (I_M + A)\ |\ \big), \tag{2.10}$$

where, from the CCC-construction, $H_t^s = D_t^s R^s D_t^s$ with D_t^s a diagonal $M \times M$ matrix with diagonal-elements $\sigma^s_{t,1}, \ldots, \sigma^s_{t,M}$, $R^s = \left[\rho^s_{ij}\right]_{1\leq i,j\leq M}$ and $\epsilon_t^s = (\sqrt{(\nu-2)/\nu}D_t^s)^{-1}(\mathbf{X}_t^s - \mu_t^s)$. 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 a preliminary correlation matrix estimate $\hat{R}^s = I_M$, we estimate the remaining parameters of the matrix A and the parameters $\alpha_0^{(j)}, \alpha_1^{(j)}, \beta^{(j)}$ $(j = 1, \ldots, M)$ by minimizing the negative log-likelihood in (2.10). This yields estimates $\hat{\mu}_t^s = (\hat{\mu}^s_{t,1}, \ldots, \hat{\mu}^s_{t,M})$ and $\hat{\sigma}^s_{t,j}$, $j = 1, \ldots, M$. We then construct the estimate for the correlation matrix R^s as follows. We build the residuals

$$\hat{\varepsilon}_{t,j} = (\hat{X}^s_{t,j} - \hat{\mu}^s_{t,j})/\hat{\sigma}^s_{t,j},\ t = 1, \ldots, T$$

and define

$$\hat{R}^s = T^{-1}\sum_{t=1}^{T}\hat{\varepsilon}_t\hat{\varepsilon}_t^T,\ \hat{\varepsilon}_t = (\hat{\varepsilon}_{t,1}, \ldots, \hat{\varepsilon}_{t,M})^T. \tag{2.11}$$

We can then iterate (once) by minimizing the negative log-likelihood in (2.10) using $\hat{R}^s$ from (2.11).

3 Numerical results

We consider a global portfolio of seven equity indices: the US Dow Jones Industrial Average (DJIA), the French CAC40 Index, the German Deutsche Aktien (DAX), the Italian BCI General Index, the Dutch CBS All-Share, the UK FT-SE-A All-Share Index (FTAS) and the Japanese NIKKEI 225 Average (NIK). The daily data is taken from the time period between January 17, 1990 and June 22, 1994, corresponding to 1000 days without holidays in the different countries (i.e. a holiday in one country led to that day being left out in all the components of the whole multivariate series). The closing times of the seven market indices are given in Table 1.

TABLE 1 ABOUT HERE.

We use here (negative) relative difference 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 as used in Section 3.3. Nevertheless, we still synchronize such relative difference returns as in (2.7).

The aim of this Section is to support empirically the effect of synchronization. We compare our synchronous CCC-GARCH(1,1) model in (2.9) with the asynchronous classical CCC-GARCH(1,1) model and the synchronous approach of Burns et al. (1998). Note that all goodness-of-fit measures and out-of-sample tests of this Section are computed for the usual asynchronous returns. This allows us to compare our results with those from other approaches. For numerical optimization of log-likelihoods, we use a quasi-Newton method.

3.1 Estimate of A and synchronization

We first examine the effect of synchronization from a descriptive point of view. The parsimoniously estimated matrix A is obtained using the procedure for structure determination described in Section 2.3 and from maximum likelihood in (2.10):

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

where the variables are ordered as DJIA, CAC40, DAX, BCI, CBS, FTAS, NIK. The column with the highest (in magnitude) coefficients corresponds to the DJIA: there is substantial predictability of all other markets from the DJIA the day before. Besides a major determining effect of the US market for the financial world, the observed pattern is natural, since the exchange in New York is the last to close. There also seems to be predictability of all other markets from the Japanese returns (NIK), although in this case, the coefficients are small and negative. The negative sign, acting 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. In addition, the German DAX and the British FTAS seem to be autocorrelated. The coefficients for the three markets that close simultaneously (French, Italian and Dutch) are all equal to zero, except for two which are still close to zero but have t-statistics below 2.3. All the other coefficients have t-statistics greater than 3 except the three coefficients in Germany and the -0.0164 and -0.0401 in Japan.

Using $\widehat{A}$ from (3.1) and the synchronization formula (2.7), we obtain the synchronized returns $\widehat{\mathbf{X}}_t^s$. The effect of synchronization in terms of empirical correlations is described in Table 2: synchronized data often exhibit larger instantaneous correlations between different returns from indices from the same day.

TABLE 2 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 would always yield higher correlations. This result is consistent with and similar to the analysis in Burns et al. (1998).

However, it is important to remark here that empirical (unconditional) correlations have no direct relation to empirical conditional quantities such as volatility or risk measures like value at risk (VaR). In the following Sections, we will demonstrate empirically that our synchronous CCC-GARCH(1,1) model has lower mean volatility, when averaged over time and multivariate components, and that it leads to lower, less conservative risk estimates than the asynchronous CCC-GARCH(1,1) model.

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-likelihood $+$ 2 # parameters

the outsample $-$ log-likelihood: $-$ log-likelihood$\Big(\widetilde{\mathbf{X}}_1^T; \widehat{A}, \widehat{\nu}, \widehat{R}^s, \{\widehat{\alpha}_0^{(j)}, \widehat{\alpha}_1^{(j)}, \widehat{\beta}^{(j)};\ j = 1, \cdots, M\}\Big)$,

where $\widetilde{\mathbf{X}}_1^T = \widetilde{\mathbf{X}}_1, \ldots, \widetilde{\mathbf{X}}_T$ are new test data and the parameter estimates, equipped with hats, are from the training sample $\mathbf{X}_1^n = \mathbf{X}_1, \ldots \mathbf{X}_n$. The likelihood itself is given in (2.10). The two statistics are measures for out-of-sample performance. A low value for the statistics indicates that the model is better. In our analysis, we take $n = 1000$ and the test set values $\widetilde{\mathbf{X}}_1^T = \mathbf{X}_{n+1}^{n+500}$ 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 equity indices are believed to be stationary (at least approximately).

The resulting values for the AIC statistic and the out-of-sample negative log-likelihood (OS-neg.LL) statistic are:

synchronous CCC-GARCH(1,1):	19058.3 (AIC)	4235.9 (OS-neg.LL),
classical CCC-GARCH(1,1):	19137.9 (AIC)	4285.9 (OS-neg.LL).

The synchronous CCC-GARCH(1,1) model is better than the asynchronous CCC-GARCH(1,1) model with respect to both statistics, although the difference is small (in order of 1 percent). Moreover, the synchronous CCC-GARCH(1,1) model also yields some improvements over the model for synchronized returns proposed by Burns et al. (1998) whose resulting values are:

synchronous model from Burns et al. (1998): 19064.9 (AIC) 4247.2 (OS-neg.LL).

We also report here the mean of absolute empirical correlations between actual outsample values $\tilde{X}_{t,i}\tilde{X}_{t,j}$ $(t = 1, \ldots, T)$ and one-step ahead predicted values of the conditional covariance

$\text{Cov}(\tilde{X}_{t,i}, \tilde{X}_{t,j} | \widetilde{\mathbf{X}}_{t-1}, \widetilde{\mathbf{X}}_{t-2}, \ldots)$ $(t = 1, \ldots, T)$, averaged over all possible components $1 \leq i \leq j \leq M$:

synchronous CCC-GARCH(1,1):	0.1325
synchronous model from Burns et al. (1998):	0.1233
classical CCC-GARCH(1,1):	0.1139

As before, the differences among the models seem to be small, although we see some improvement using our methodology. Such small differences could be obscured by low signal to noise ratio when replacing the unobservable conditional covariances by their corresponding actual return values, which are noisy estimates. 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. Testing our synchronization method on the seven-dimensional asynchronous equity index returns against both the asynchronous CCC-GARCH(1,1) model and the synchronization method from Burns et al. (1998) we find evidence of statistical significance, implying a preference of our synchronous model over both alternatives. The exact construction of the tests, as well as the description of the results, is deferred to Appendix A.

3.3 Estimating the performance at the portfolio level

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

$$P_t = \sum_{j=1}^{7} \alpha_j \; S_{t,j} \quad t = 1, \ldots, m. \tag{3.2}$$

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

$$\begin{aligned} \Delta_t &= -100 \cdot \big(\frac{P_t - P_{t-1}}{P_{t-1}}\big) = -100 \cdot \frac{\sum_{j=1}^{7} \alpha_j \; (S_{t,j} - S_{t-1,j})}{P_{t-1}} \\ &= \sum_{j=1}^{7} \Big(\frac{\alpha_j \; S_{t-1,j}}{P_{t-1}} \; \big(-100 \cdot \frac{S_{t,j} - S_{t-1,j}}{S_{t-1,j}}\big)\Big) = \beta_{t-1}' \mathbf{X_t}, \end{aligned}$$

where

$$\beta_{t-1,j} = \alpha_j \; \frac{S_{t-1,j}}{P_{t-1}}, \quad j = 1, \ldots, 7. \tag{3.3}$$

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 the following four models. The multivariate synchronous CCC-GARCH(1,1) model (2.9), the synchronous model introduced by Burns et al. (1998), the asynchronous CCC-GARCH(1,1) model and a classical GARCH(1,1) univariate analysis (and extensions thereof) for the portfolio returns Δ_t. Clearly, the different approaches give rise to different $\mu_{t,P}$ and $\sigma_{t,P}$. Note that in the more realistic case, the weights $\alpha_i = \alpha_{t,i}$ depend on t. As a consequence, the univariate analysis of the returns Δ_t will be inappropriate, because the returns of portfolio prices would typically be far from stationarity.

For univariate analysis based on returns Δ_t, we always assume the model

$$\Delta_t = \mu_{t,P} + \epsilon_t = \phi\Delta_{t-1} + \sigma_{t,P}\ Z_t, \tag{3.4}$$

where $\sigma_{t,P}$ is a measurable function of previous returns $\Delta_{t-1}, \Delta_{t-2}, \ldots$ 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 $\mathrm{Var}(Z_t) = 1$. The univariate GARCH(1,1) specification is

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

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

$$\begin{aligned} -\ell(\alpha_0, \alpha_1, \beta, \phi, \nu; \Delta_2^n) &= -\sum_{t=2}^{n} \log\left(c(\nu)^{-1}\sigma_{t,P}^{-1} f_{t_\nu}\left(\frac{\Delta_t - \phi\Delta_{t-1}}{c(\nu)\sigma_{t,P}}\right)\right), \\ c(\nu) &= ((\nu-2)/\nu)^{1/2}, \end{aligned} \tag{3.5}$$

where f_{t_ν} denotes the density of the univariate t_ν distribution. Minimizing the negative log-likelihood yields estimates $\widehat{\mu}_{t,P} = \widehat{\phi}\Delta_{t-1}$ and $\widehat{\sigma}_{t,P}^2 = \widehat{\alpha}_0 + \widehat{\alpha}_1\epsilon_{t-1}^2 + \widehat{\beta}\sigma_{t-1,P}^2$. In particular, we find the parameter estimates $\widehat{\alpha}_0 = 0.020, \widehat{\alpha}_1 = 0.049, \widehat{\beta} = 0.917$, $\widehat{\phi} = 0.154$ and $\widehat{\nu} = 5.659$.

When using the multivariate synchronous CCC-GARCH(1,1) model, we calculate estimates of the portfolio conditional means $\widehat{\mu}_{t,P}$ and variances $\widehat{\sigma}_{t,P}^2$, $t = 1, \ldots, n$ as follows. We always take the innovations $\mathbf{Z_t}$ of the model (2.9) to be multivariate t_ν distributed (ν unknown) with zero mean and covariance matrix $\mathrm{Cov}(\mathbf{Z}) = I_M$. 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

$$\mathbf{X_t} \mid \mathcal{F}_{t-1} \sim t_\nu\big(A\ \mathbf{X_{t-1}}\ ,\ (I_M + A)^{-1}\ \Sigma_t^s\ ((I_M + A)^{-1}\ \Sigma_t^s)'\big).$$

Exploiting a nice property of elliptical distributions[8] we see 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:

$$\Delta_t \mid \mathcal{F}_{t-1} \sim t_\nu\big(\beta_{t-1}'\ A\ \mathbf{X_{t-1}}\ ,\ \beta_{t-1}'\ (I_M + A)^{-1}\ \Sigma_t^s\ (\beta_{t-1}'\ (I_M + A)^{-1}\ \Sigma_t^s)'\big),$$

where the vector of coefficients β_{t-1} is given in (3.3). Thus, we compute

$$\begin{aligned} \widehat{\mu}_{t,P} &= \beta_{t-1}'\ \widehat{A}\ \mathbf{X_{t-1}} \qquad \text{and} \\ \widehat{\sigma}_{t,P}^2 &= \beta_{t-1}'\ (I_M + \widehat{A})^{-1}\ \widehat{\Sigma}_t^s\ (\beta_{t-1}'\ (I_M + \widehat{A})^{-1}\ \widehat{\Sigma}_t^s)'\ , \end{aligned}$$

where $\widehat{A}$ and $\widehat{\Sigma}_t^s$ are the maximum likelihood estimates in the model (2.9). The estimates from the classical asynchronous CCC-GARCH(1,1) model are of the same form, but with $\widehat{A} = 0$ and $\widehat{\Sigma}_t^{\text{asynch.}}$.

The predicted portfolio conditional variance using the synchronous CCC-GARCH(1,1) model, the asynchronous CCC-GARCH(1,1) model and an univariate GARCH(1,1) model for the portfolio returns are plotted in Figure 1.

[8]See, for example, Fang et al. (1990).

FIGURE 1 ABOUT HERE.

The predicted (squared) volatilities are generally larger when using the univariate approach. This seems to be due to the fact that univariate approaches are exposed to an information loss by averaging individual prices. It may cause a bias that result in higher risk estimates for the data set considered here, see Section 4. On the contrary, multivariate methods do not estimate directly the portfolio dynamics but are based on a multivariate analysis of all individual return series before constructing the portfolio volatility predictions. The differences between the two multivariate methods are (visibly) much smaller.

Now, we test the goodness of the residuals

$$\widehat{Z}_t = \frac{\widetilde{\Delta}_t - \widehat{\mu}_{t,P}}{\widehat{\sigma}_{t,P}}, \; t = 1, \ldots, T, \tag{3.6}$$

in the different approaches. Here $\widetilde{\Delta}_t$ is from new test set data $\Delta_{n+1}, \ldots, \Delta_{n+500}$ over the next 500 days; $\hat{\mu}_{t;P}$ and $\hat{\sigma}_{t;P}$ are from the different models, estimated with the training data $\Delta_1, \ldots, \Delta_n$, naturally 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 (3.4). Under the null hypothesis and assuming Gaussian innovations, the statistic $\sqrt{T}\overline{Z}$ is approximately standard normally distributed. The observed values, including the corresponding P-values in parentheses, of the test for the synchronous CCC-GARCH(1,1), the asynchronous CCC-GARCH(1,1) and the univariate model are -1.570 (0.058), -1.898 (0.029) and -1.527 (0.063), respectively. Thus, only the asynchronous CCC-GARCH(1,1) model seems incompatible with the data.

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

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

where, as in (3.6), $\widetilde{\Delta}_1^T = \Delta_{n+1}^{n+500}$ are new test data and $\widehat{\mu}_{t,P}, \widehat{\sigma}_{t,P}$ are estimated from the training data. The exact form of the log-likelihood is given in (3.5). In addition, we also consider the following out-of-sample prediction loss statistics:

$$\text{OS-PL}_i = \sum_{t=1}^{T} \mid \widehat{\sigma}_{t,P}^2 - \left(\widetilde{\Delta}_t - \widehat{\mu}_{t,P}\right)^2 \mid^i, \; i = 1, 2.$$

The OS-PL statistics and the portfolio out-of-sample log-likelihood are, as before, measures for predictive performance. A low value for the statistics indicates that the model is better.

The values of these statistics are summarized in Table 3. Results are computed using the multivariate synchronous CCC-GARCH(1,1) model, the synchronous model introduced by Burns et al. (1998) (denoted by BEM, according to the names of the authors), the asynchronous CCC-GARCH(1,1) model and a standard univariate GARCH(1,1) analysis of the negative portfolio returns.

TABLE 3 ABOUT HERE.

Analogously to the results obtained in Section 3.2, we also find that the synchronous CCC-GARCH(1,1) model is better at the portfolio level than the asynchronous CCC-GARCH(1,1) model or the univariate approach with respect to all goodness of fit measures. Table 3 also shows that, in this case, improvements are more relevant than at the multivariate level. In particular with respect to the OS-PL performances the values decrease by 3-6%. The synchronous

CCC-GARCH(1,1) model also yields some improvements over the synchronous BEM model. In contrast, the univariate GARCH(1,1) analysis exhibits slight advantages over multivariate modelling without using synchronization.

As already mentioned in Section 3.2, more impressive gains may be masked by a low signal to noise ratio. The t- and sign-type tests described in the Appendix A can be used to better compare the accuracy of the estimates from the different approaches, also at the portfolio level. The results are summarized in Table 4.

TABLE 4 ABOUT HERE.

The t-type tests yield significant differences only in the comparison between the synchronous CCC-GARCH(1,1) model and the univariate approach, preferring the former over the latter. This may be just a fact of low power due to non-Gaussian observations. In contrast, the sign-type tests, which are robust against deviations from Gaussianity, yield significant results in most cases. The synchronous CCC-GARCH(1,1) model is better than both alternatives, whereas the univariate approach is about as good (maybe slightly better) than the asynchronous CCC-GARCH(1,1) model. It seems that classical multivariate modelling without synchronization does not yield any particular gain over a simple univariate analysis of the portfolio returns, confirming the results recently found by Berkowitz and O'Brien (2002).[9]

4 Estimating risk measures

We will now test the effect of synchronization on the computation of conditional (dynamical) risk measures for negative portfolio returns Δ_t following (3.3)-(3.4) given the information $\mathcal{F}_{t-1}$ from previous prices. The most popular risk measure, which has also been adopted for regulatory purposes[10], is *Value at Risk* (VaR). A one-day VaR is given by

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

where $F_{\Delta_t|\mathcal{F}_{t-1}}(\cdot)$ denotes the cumulative distribution function of Δ_t given $\mathcal{F}_{t-1}$ and q is the confidence level at which we want to compute the VaR. This is the quantile of the predictive distribution of the negative portfolio return over the next day.

A second, widely used risk measure is the so-called *expected shortfall or conditional VaR*, which is defined as the expected loss, on condition that the loss has exceeded VaR. Thus, a one-day expected shortfall is

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

where, as before, q is the confidence level. We typically choose $q \in \{0.90, 0.95, 0.99\}$ (note that we consider negative returns). The expected shortfall is a coherent measure of risk.[11] The

[9]We also analyzed, whether the improvements with the synchronous CCC-GARCH(1,1) model could 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 than the classical CCC-GARCH(1,1) (see Proposition 1), but motivated from the view of synchronization with a simple linear transform. We analyzed different extensions of the asynchronous classical GARCH(1,1) model (for example including a conditional mean term and more complex, potentially high dimensional parameterizations for approximating more general volatility functions). Considering the same goodness of fit measures as used above, we found that more sophisticated models (not being of synchronization type) did not show worthwhile improvements. Hence, synchronization seems to have a more substantial effect than trying to improve the model-dynamics. This finding confirms the results of Hansen and Lunde (2002) for univariate series.

[10]More specifically, the 1996 Market Risk Amendment to the Basel Accord stipulates that the minimum capital requirement for market risk should be based on a 10-day VaR at a 99% confidence level. The amendment allows 10-day VaR to be measured as a multiple of one-day VaR.

[11]See Artzner et al. (1999), and Rockafeller and Uryasev (2002) for more details.

empirical investigations of the next Sections are performed for the same real global portfolio consisting of the seven equity indices listed in Section 3.3.

4.1 The estimates

We assume that the dynamics of the negative asynchronous portfolio returns Δ_t $(t \in \mathbb{Z})$ are given by (3.4). Since

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

the risk measures can then be written as

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

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

Estimates for the VaR and for the expected shortfall are constructed using the assumption of scaled t_ν distributed innovations Z_t in (3.4), 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 for the expected shortfall is given by

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

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 $\widehat{\nu}$ is the maximum likelihood estimate from the multivariate or univariate models, as before. Clearly, the t_ν assumption made for the distribution of the innovations is not restrictive. Alternatively, we can use, for example, extreme value theory and the peaks over the threshold method[12] to model the tails of $F_Z(\cdot)$.

For illustrative purposes, in Figure 2 we show estimates of the conditional expected shortfall. The estimates are constructed using the multivariate synchronous CCC-GARCH(1,1) model and the asynchronous CCC-GARCH(1,1) model on the seven-dimensional series of negative equity index returns, and a standard univariate GARCH(1,1) model on the portfolio negative returns.

FIGURE 2 ABOUT HERE.

Figure 2 shows that expected shortfall estimates constructed using multivariate models are in general lower and change more slowly than those from an univariate approach. Moreover, estimates from the synchronous CCC-GARCH(1,1) model wiggle more than those from the asynchronous CCC-GARCH(1,1) model and exhibit small scale fluctuations. Performance results from Section 3 suggest that these small scale movements are a good feature.

[12] See Embrechts et al. (1997) or McNeil and Frey (2000) for a detailed description of the method.

4.2 Backtesting

Backtesting expected shortfall estimates can be very difficult since a tail phenomenon is estimated. As a descriptive tool, rather than a formal test, in Figure 3 we show boxplots of residuals

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

FIGURE 3 ABOUT HERE.

Under model assumptions (3.4) and ignoring estimation effects, we can easily show that the theoretical residuals R_t are an i.i.d. sequence with expected value zero. Figure 3 yields additional evidence that expected shortfall estimates from a classical univariate GARCH(1,1) analysis of the negative portfolio returns are too conservative (too low values of residuals). However, standard tests for unconditional coverage examining whether losses exceed VaR estimates at the $q \cdot 100\%$ confidence level more frequently than q percent of the time, never rejected the null hypothesis of unconditional unbiasedness of VaR estimates[13]. Therefore, all the models seem to be compatible with the data.

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. Apart from the neat interpretative structure of synchronization, we found empirically that it improved the predictive performance of the CCC-GARCH(1,1) model for a seven-dimensional time series of daily equity index returns. The predictive gain with synchronization often seems to be bigger than when trying to extend the CCC-GARCH(1,1) model to a more complex model for approximating more general volatility functions and allowing for leverage effects (Audrino (2002), Section 3.5). The question of whether more sophisticated volatility models are able to outperform the simple GARCH(1,1) model has also been recently investigated by Hansen and Lunde (2002) at the univariate level. In their empirical analysis, they found that it is very difficult to beat the simple GARCH(1,1) model. Our empirical comparisons confirm that in our particular example, synchronization yields more prominent improvements than extending the multivariate GARCH to a more complex model.

Backtesting VaR and expected shortfall estimates computed using our synchronous CCC-GARCH(1,1) model on a real global portfolio consisting of seven equity indices around the world, we provide empirical evidence of the power of our synchronization methodology. Backtest results support the evidence that synchronization leads to better risk estimates when compared to those from a classical CCC-GARCH(1,1) model without synchronization and from a direct univariate GARCH(1,1) analysis of the portfolio returns. Risk estimates from our model change more slowly and exhibit more small scale fluctuations than those from alternative approaches. We also found that the univariate analysis of the portfolio returns seems to be biased for the calculation of risk measures and turned out to be too conservative. In constrast, we showed that no improvement (in terms of predictive performance and measurement of risk) over a simple univariate GARCH(1,1) model for the portfolio returns could be achieved using a multivariate GARCH model without synchronization. This also emphasizes the importance of synchronization in multivariate approaches.

[13] Note that this can be due to the very low power of such tests in detecting errors in VaR estimates; see, for example, Kupiec (1995).

Acknowledgements: We thank Michel Dacorogna for some interesting remarks. The data was provided by Olsen & Associates, Zürich. We also thank an anonymous referee and the editor Philippe Jorion for constructive comments.

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. (2002). Statistical Methods for High-Multivariate Financial Time Series. Dissertation No. **14565**, ETH Zürich.

[3] Berkowitz, J. and O'Brien J. (2002). How Accurate are Value-at-Risk Models at Commercial Banks. *Journal of Finance* **57**, Issue 3, 1093-1111.

[4] 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.

[5] Brockwell, P.J. and Davis, R.A. (1991). *Time Series: Theory and Methods.* Springer, New York.

[6] Burns, P., Engle, R.F. and Mezrich, J. (1998). Correlations and volatilities of asynchronous data. Journal of Derivatives, Summer, 1–12.

[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] Freedman, D. (1984). On bootstrapping two-stage least-squares estimates in stationary linear models. Annals of Statistics **12**, 827–842.

[10] Hansen, P.R. and Lunde, A. (2002). A Forecast Comparison of Volatility Models: Does Anything Beat a GARCH(1,1)? Manuscript, Brown University.

[11] Kupiec, P. (1995). Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives* **3** (Winter), 73-84.

[12] McCullagh, P. and Nelder, J.A. (1989). *Generalized linear Models.* Chapman & Hall, London.

[13] McNeil, A.J. and Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance* **7**, 271-300.

[14] Reinsel, G.C. (1991). *Elements of Multivariate Time Series Analysis.* Springer, New York.

[15] Rockafellar, R.T. and Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance* **26**, 1443-1471.

A Appendix

A.1 t-type and sign-type tests

Consider differences $\widehat{D}_t,\ t = 1, \dots, T$, of performance terms

$$\widehat{D}_t = \widetilde{U}_{t;\text{model}_1} - \widetilde{U}_{t;\text{model}_2},\ t = 1, \cdots, T\ ,$$

where the sum $\sum_{t=1}^{T} \widetilde{U}_{t;\text{model}}$ equals the total out-of-sample performance of a model[14]. We are now testing the null hypothesis that the differences $\widehat{D}_t$ have mean zero against the alternative of mean less than zero, i.e. the estimates from model$_1$ are better than the ones from model$_2$. 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

$$\sqrt{T}\ \frac{\overline{D}}{\widehat{\sigma}_{D;\infty}}\ , \quad \text{where } \overline{D} = \frac{1}{T} \sum_{t=1}^{T} \widehat{D}_t\ . \tag{A.1}$$

In (A.1), $\widehat{\sigma}^2_{D;\infty} = (2\pi)\widehat{f}_{\widehat{D}}(0)$, where $\widehat{f}_{\widehat{D}}(0)$ is a smoothed periodogram estimate at frequency zero, based on $\widehat{D}_1, \dots, \widehat{D}_T$; see for example Brockwell and Davis (1991). The motivation for this estimate is based on the assumption that $\{\widehat{D}_t\}_t$ is stationary (conditional on the training data) and satisfies suitable dependence conditions, e.g. mixing. Then, conditional on the training data,

$$\begin{aligned} &\sqrt{T}\big(\overline{D} - E[\widehat{D}_t]\big) \Longrightarrow \mathcal{N}(0, \sigma^2_{D;\infty}) \quad (T \to \infty)\ , \\ &\sigma^2_{D;\infty} = \sum_{k=-\infty}^{+\infty} \text{Cov}\big[\widehat{D}_0, \widehat{D}_k\big] = (2\pi) f_{\widehat{D}}(0), \end{aligned} \tag{A.2}$$

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

Thus, using (A.2) for the test statistic in (A.1), and conditional on the training data,

$$\sqrt{T}\frac{\overline{D}}{\widehat{\sigma}_{D;\infty}} \Longrightarrow \mathcal{N}(0, 1) \quad (T \to \infty) \tag{A.3}$$

under the null hypothesis.

Analogously, the version of the sign test in the case of dependent observations is based on the number of negative differences

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

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

$$\sqrt{T}\ \frac{\overline{W} - \frac{1}{2}}{\widehat{\sigma}_{W;\infty}}, \quad \text{where } \overline{W} = \frac{1}{T} \sum_{t=1}^{T} \widehat{W}_t \tag{A.4}$$

[14]In case of the negative log-likelihood loss, $\widehat{D}_t$ is the difference between out-of-sample deviance residuals, up to a change of signs; see McCullagh and Nelder (1989).

and $\widehat{\sigma}^2_{W;\infty}$ as in (A.1) but based on $\widehat{W}_1, \ldots, \widehat{W}_T$. As in the derivation of the t-type test above, we have, conditional on the training data,

$$\sqrt{T}\,\frac{\overline{W} - \frac{1}{2}}{\widehat{\sigma}_{W;\infty}} \Longrightarrow \mathcal{N}(0,1) \quad (T \to \infty) \tag{A.5}$$

under the null hypothesis.

A.2 Multivariate negative out-of-sample log-likelihood results

We report here the results of the t-type and sign-type tests for the multivariate negative out-of-sample log-likelihood statistic introduced in Section 3.2. The observed values for the t-type test statistic (A.1) equal -1.7822 and -1.6923 with corresponding P-values of 0.037 and 0.045 when testing our synchronization method against the asynchronous CCC-GARCH(1,1) model and against the synchronization method from Burns et al. (1998), respectively. These results imply that our synchronized model is significantly better than both alternative approaches. Analogously, we find 308 and 254 negative differences (among a total of $T = 500$) and the observed values of the sign-type test statistic (A.4) are 1.708 and 0.186 with corresponding P-values of 0.044 and 0.426, respectively. Therefore, these tests also lead to one rejection of the null hypothesis (i.e. equal performance of models) at the 5% significance level, implying a preference of our synchronous CCC-GARCH(1,1) model over the asynchronous CCC-GARCH(1,1) model.

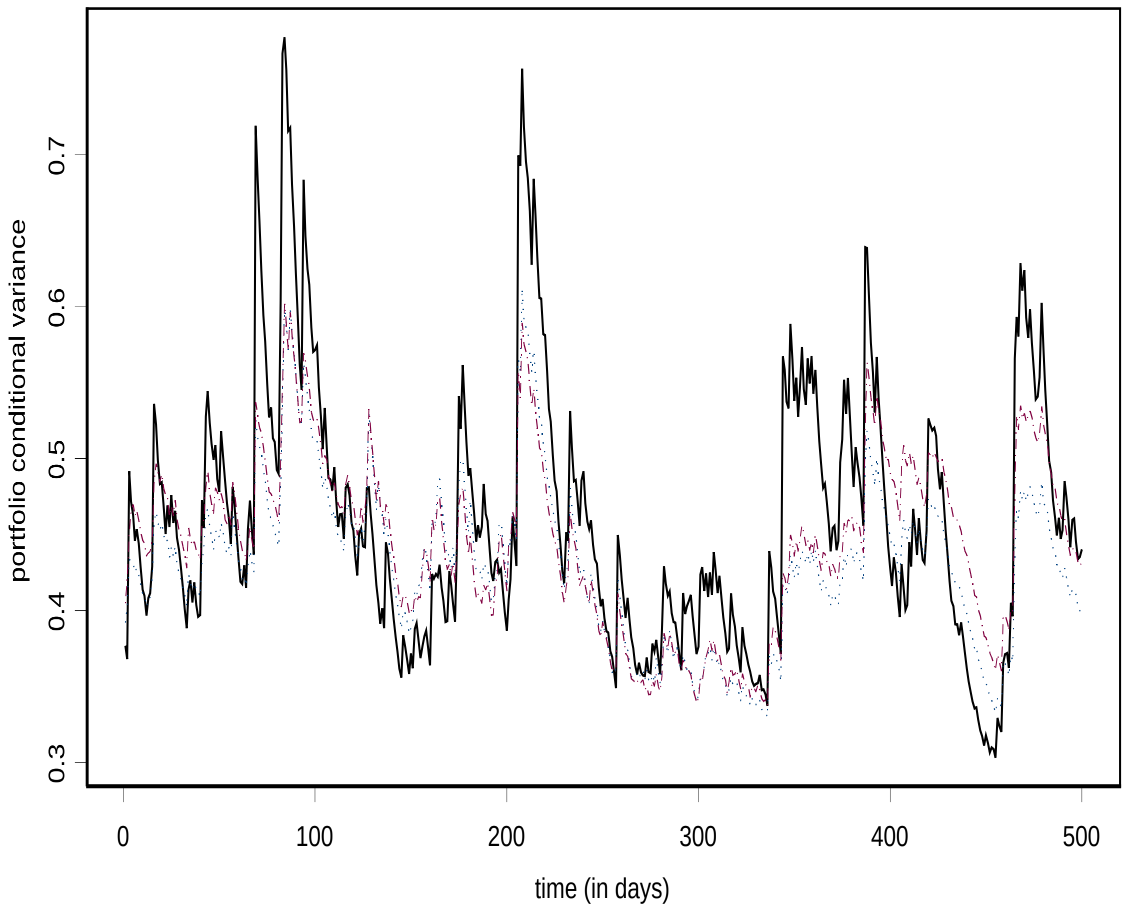

Figure 1: Predicted portfolio conditional variances $\widehat{\sigma}^2_{t,P}$ for the backtesting period between June 23, 1994 and September 9, 1996 (500 trading days). Predictions are constructed using a standard univariate GARCH(1,1) model on the (negative) portfolio returns (solid line), and both the synchronous CCC-GARCH(1,1) model (dotted line) and the standard CCC-GARCH(1,1) model without synchronization (dashed line) on the seven-dimensional series of (negative) equity index returns.

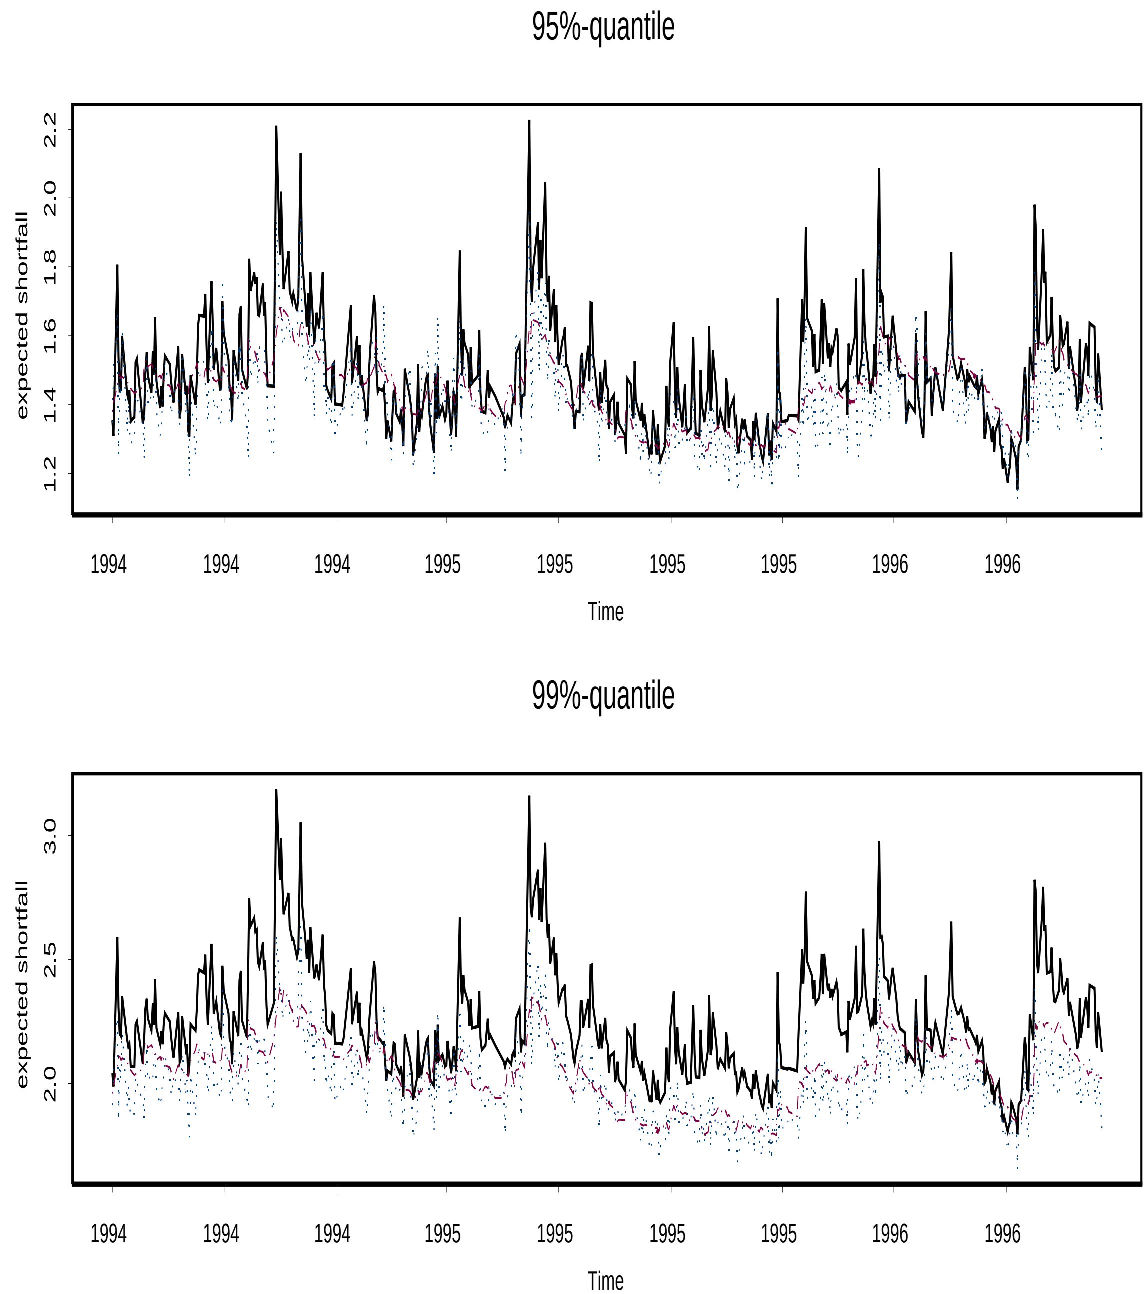

Figure 2: Expected shortfall estimates $\widehat{S}_t^q$ for the (negative) portfolio returns $\widetilde{\Delta}_t$ during the back-testing period beginning June 23, 1994 and ending September 9, 1996 (for a total of 500 trading days). The estimates $\widehat{S}_t^q$ for $q = 0.95$ (top) and $q = 0.99$ (bottom) are obtained using a standard univariate GARCH(1,1) model on the (negative) portfolio returns (solid line), and both the synchronous CCC-GARCH(1,1) model (dotted line) and the standard CCC-GARCH(1,1) model without synchronization (dashed line) on the seven-dimensional series of (negative) equity index returns.

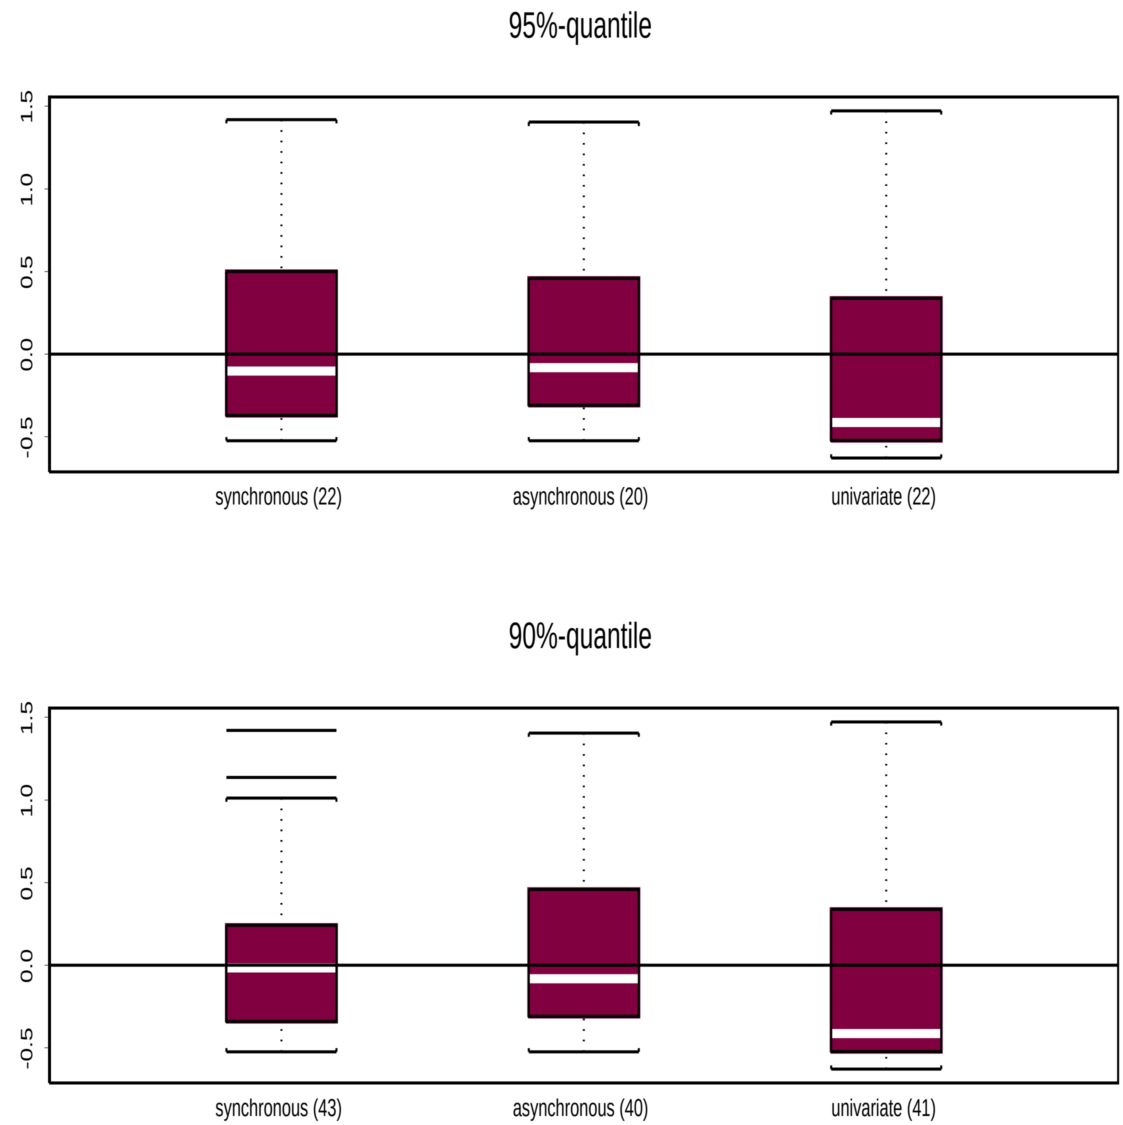

Figure 3: Boxplots of residuals $\widehat{R}_t$ from the expected shortfall estimates computed during the back-testing period beginning June 23, 1994 and ending September 9, 1996 (for a total of 500 trading days). Expected shortfall estimates are constructed at the 95% (top) and at the 90% (bottom) confidence levels using the synchronous CCC-GARCH(1,1) model (left) and the standard CCC-GARCH(1,1) model without synchronization (center) on the seven-dimensional series of (negative) equity index returns, and a standard univariate GARCH(1,1) model (right) on the (negative) portfolio returns. The number of violations is given between parentheses. If the model is correct, the expected numbers of violations are 25 (top) and 50 (bottom), respectively. The horizontal lines indicate the expected value of the residuals equalling zero if the model is correct.

Closing times

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

Table 1: Closing times for a seven-dimensional real data example consisting of (negative) equity index return series of developed capital markets around the world.

Raw empirical correlations

	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

Synchronous empirical correlations

	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 2: Top: instantaneous empirical correlations between components of usual asynchronous returns $\mathbf{X_t}$ for a seven-dimensional series of (negative) equity index returns. Bottom: instantaneous empirical correlations between components of estimated synchronized returns $\widehat{\mathbf{X}}_t^s$ using the synchronous CCC-GARCH(1,1) model.

Portfolio performance results

	synchronous CCC-GARCH(1,1)	synchronous BEM model	CCC-GARCH(1,1) without synchronization	Univariate approach
OS-PL_1	210.0009	211.6876	218.7766	216.5164
OS-PL_2	227.6059	230.9892	241.4637	234.3302
Port. out. log-lik.	459.5995	461.8611	469.9058	464.3049

Table 3: Values of different portfolio out-of-sample goodness of fit statistics. Results are computed using the synchronous CCC-GARCH(1,1) model, the synchronous BEM model of Burns et al. (1998) and the CCC-GARCH(1,1) model without synchronization on the seven-dimensional series of (negative) equity index returns, and a standard AR(1)-GARCH(1,1) model on the univariate series of (negative) portfolio returns.

t-type tests

Model 1	Model 2	Performance measure	
		Portfolio out. log-likelihood	OS-PL_2
Synchronous CCC-GARCH(1,1)	Asynchronous CCC-GARCH(1,1)	-1.3944 (0.082)	-1.3738 (0.085)
Synchronous CCC-GARCH(1,1)	Univariate approach	-1.7781 (0.038)	-2.1519 (0.016)
Asynchronous CCC-GARCH(1,1)	Univariate approach	0.6822 (0.248)	0.7238 (0.235)

Sign-type tests

Model 1	Model 2	Performance measure	
		Portfolio out. log-likelihood	OS-PL_2
Synchronous CCC-GARCH(1,1)	Asynchronous CCC-GARCH(1,1)	2.0393 (0.021)	2.1820 (0.015)
Synchronous CCC-GARCH(1,1)	Univariate approach	-0.5808 (0.281)	1.7631 (0.039)
Asynchronous CCC-GARCH(1,1)	Univariate approach	-1.7456 (0.040)	0.0486 (0.481)

Table 4: Testing differences of performance terms between Model 1 and Model 2 at the portfolio level. We consider three different models: the synchronous CCC-GARCH(1,1) model and the classical CCC-GARCH(1,1) model without synchronization on the seven-dimensional series of (negative) equity index returns, and a standard AR(1)-GARCH(1,1) model on the univariate series of (negative) portfolio returns. The values of t-type and sign-type test statistics adapted to the case of dependent observations are summarized. The corresponding P-values are given between parentheses.