Preprint, April 2010

Market Correlations in the Euro Changeover Period With a View to Portfolio Management

Gernot Müller*

Keywords: European Monetary Union · European Currencies · Markov Chain Monte Carlo · Minimum Variance Portfolio · Stochastic Volatility · Value at Risk · Volatility Correlations

JEL classification: C11 · C32 · F21 · F33 · G15

*Centre of Mathematics, Technische Universität München, 85748 Garching, Germany, email: mueller@ma.tum.de, phone: +49 89 289 17441, fax: +49 89 289 17435

Introduction

Although the idea of a common European currency can be traced back to about 1970, the first major step towards this common currency was made on July 1, 1990, when restrictions on trade and movement of capital were lifted between the members of the European Union. In January 1994, the European Monetary Institute (as a predecessor of the European Central Bank) was founded, and one started to control consequently the budgetary situation of the member countries. In order to join the European Monetary Union, the nations had to meet strict standards on inflation, currency stability and deficit spending. In December 1995 the European Council fixed the name of the new common currency: euro. As a last step, the European parliament formally approved on May 2, 1998 the historic decision to launch the euro with 11 founder nations: Germany, France, Italy, Spain, the Netherlands, Belgium, Finland, Portugal, Austria, Ireland and Luxembourg. On January 1, 1999, when the eleven founding member countries of the European Economic and Monetary Union (EMU) surrendered their right to print money, the exchange rates among the currencies of those nations were fixed both against each other and against the euro. The euro was first introduced as a virtual currency for accounting purposes, but country specific currencies continued to be used for cash payments and investments until the euro was issued physically on January 1, 2002.

With the introduction of the euro in January 1999, exchange rate problems were eliminated and the prices became transparent between the member countries. Immediately the euro was accepted as reliable international currency. Whereas the currency union was originally seen to be irrevocable, now, about ten years later, some experts recommend to take into consideration the exclusion of certain member nations which do not meet the strict stability standards. At the time of writing, Greece, which became a member of the EMU in 2001, comes under pressure due to its high national deficit: Stock markets reacted quite sensitively in December 2009, when a leading rating agency downgraded Greece's credit rating.

Whatever the consequences for Greece will be, it is clear that the discussion on the advantages and disadvantages of the euro will go on, and analyses which can shed light on the past performance of the euro will be of great interest. Our first goal is to study the changes in market linkages which came along with the introduction of the euro. We restrict ourselves to the two major EMU countries Germany and France, but include Switzerland in our analysis, since it is neither a member of the EMU nor a member of the European Union, so that the effects caused by the euro can be separated from general market integration effects in the past 20 years. The analysis is conducted using a trivariate stochastic volatility (SV) model, which is fitted to the data in a Bayesian framework using Markov chain Monte Carlo (MCMC). The second aim of the paper is to exploit the Bayesian approach to illustrate the uncertainty of estimates used for portfolio selection and management. In particular, we calculate, as a common risk measure, VaR and explicitly derive credible intervals for it, and illustrate the uncertainty of the estimated weights of a minimum variance portfolio.

Of course, there is a great variety of publications on the effects of the euro, and we can only mention a few which are close to our analysis. The study of Westermann (2004) investigates daily index data of the three major EMU countries Germany, France and Italy, comparing the dynamic linkages only for the years 1998 and 1999. Moreover, Westermann also included the US in his study, confirming its leading impact on daily prices in the European countries. Moore (2007) investigates dynamic market linkages between the earlier established euro-zone countries and some later-entry countries (the Czech Republic, Hungary, Poland and Slovakia). The analysis of Kim, Moshirian, and Wu (2005) is based on bivariate E-GARCH models to assess integration among equity markets in the European Union over the period January 1989 to May 2003. They find a clear regime shift in stock market comovements within the European Union with the introduction of the euro. The model we use can be considered as a multivariate extension of the SV model in Fleming, Kirby, and Ostdiek (1998). It allows to check for integration of market information and volatility spillovers between the markets.

The paper is organized as follows. After describing the data, we introduce the multivariate SV model and explain the Bayesian setting. Then the results from model-fitting are discussed as well as goodness-of-fit. After that, we illustrate how the Bayesian approach allows to measure uncertainty in portfolio management. In particular, we show how Bayesian value at risk estimates can be derived as well as Bayesian estimates for a minimum variance portfolio. Finally we summarize our results.

Data

Our analysis is based on daily log-returns from the DAX, CAC and SMI[1], for the two periods January 1, 1993 to December 31, 1998 and January 1, 1999 to December 31,

[1] *Deutscher Aktien Index* = German Stock Index, CAC takes its name from the Paris Bourse's early automatization system *Cotation Assistée en Continue* = Continuous Assisted Quotation, and SMI = Swiss Market Index

Figure 1: *Log-returns in USD in the pre-euro period, 1993-1998: DAX, CAC, SMI.*

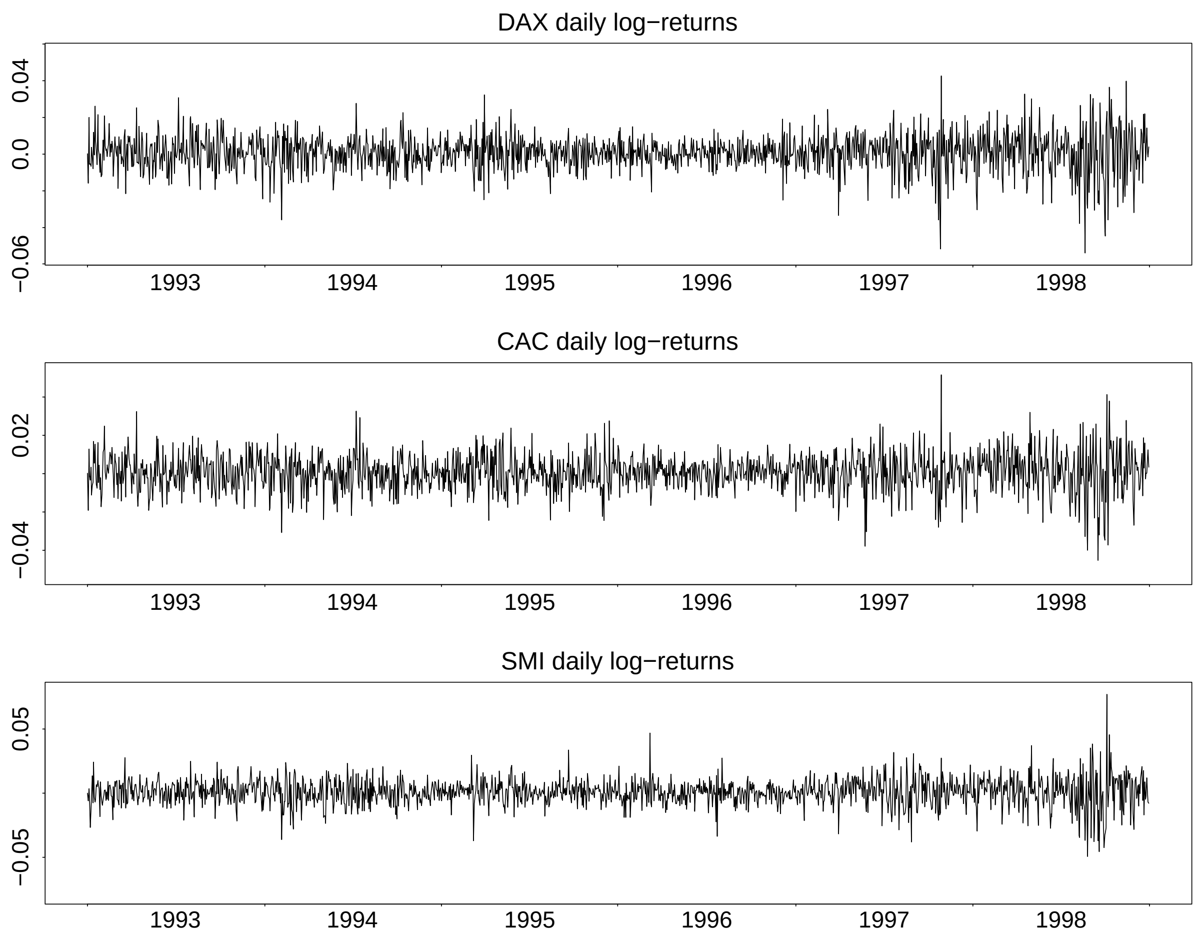

2004, which we call briefly the pre-euro and the post-euro period, respectively, in the following. The pre-euro data consists of 1565 log-returns, and the post-euro data of 1566 log-returns.

To have a common basis for comparison before the introduction of the euro, we convert all returns to USD at the exchange rate at the time of the return. For consistency, we do the same for the post-euro period, so that, as for the portfolio analysis part of the paper, we can think of taking on the position of an American investor.

Figures 1 and 2 show the USD based log-returns for the three indices. The similarity in the volatility patterns is obvious. Indeed, using straightforward sample correlation estimates one can already support the hypothesis that the market integration between

Figure 2: *Log-returns in USD in the post-euro period, 1999-2004: DAX, CAC, SMI.*

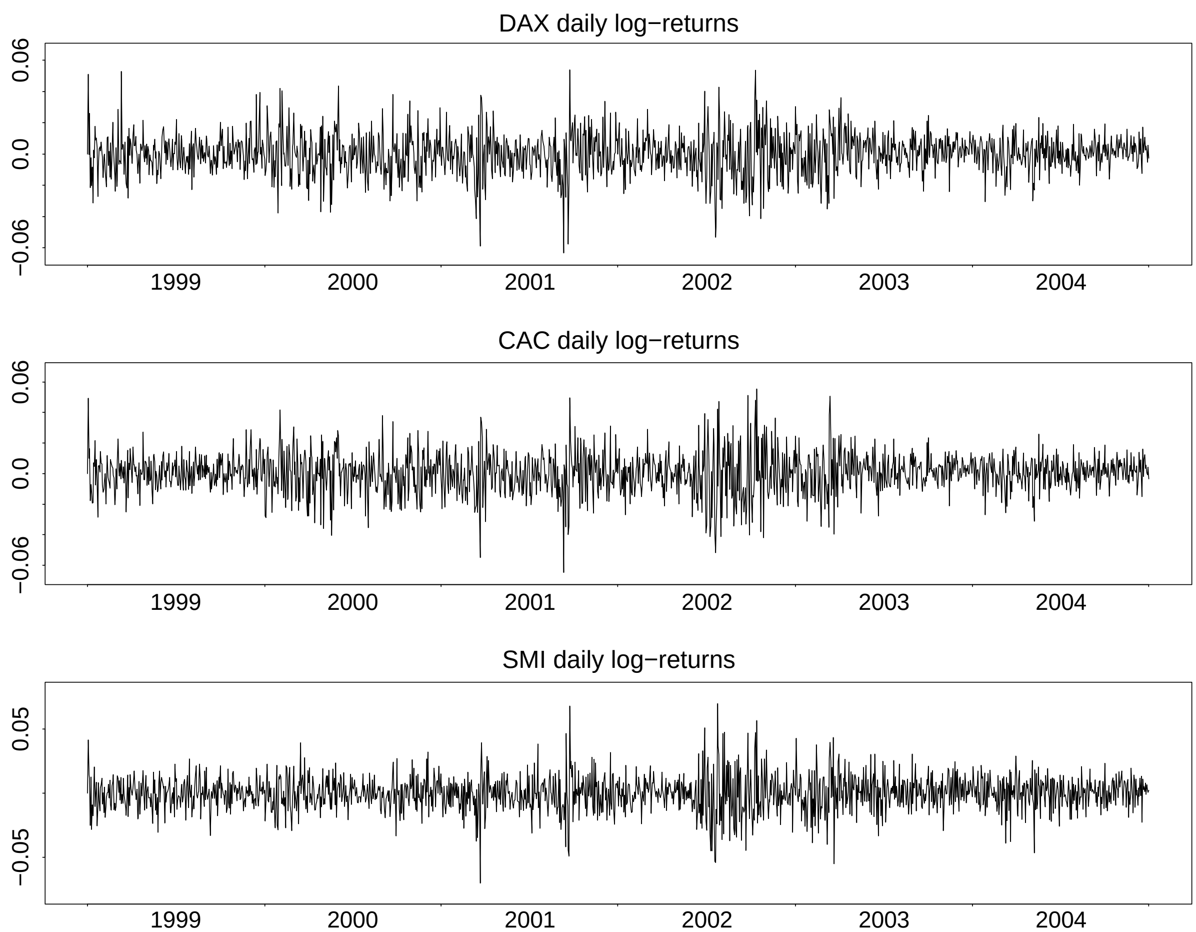

Germany and France increased from the pre- to the post-euro period. But to get a clear answer whether these changes were significant, and to which extent also the correlations of Switzerland to Germany and France were affected by the introduction of the euro, an appropriate model is needed. As mentioned before, we will use a multivariate SV model, which is introduced in the following section.

Model and MCMC method

First we briefly outline the trivariate stochastic volatility model which is used in our analysis. Then we show how residuals in this model can be derived, which is of interest for assessing goodness-of-fit, and summarize the Bayesian MCMC approach used to fit the model.

The multivariate stochastic volatility model

The stochastic volatility model we use is an N-variate extension of the model developed in Taylor (1994) and Ruiz (1994), where the observations, y_{it}, are the mean corrected log returns of the i-th country's stock market index, $i = 1, 2, \ldots, N$, at time t, while the log squared volatilities follow AR(1) processes. In our case, $N = 3$ since we restrict ourselves to the three countries Germany, France and Switzerland, so we have a trivariate stochastic volatility model. The model is specified as

$$\mathbf{y}_t = \mathbf{H}_t^{1/2} \boldsymbol{\varepsilon}_t, \tag{1}$$

where $\mathbf{y}_t = (y_{1t}, y_{2t}, y_{3t})$, and the subscript t indicates daily observations over a time period $t = 1, 2, \ldots, T$. The matrix $\mathbf{H}_t = \mathrm{diag}(e^{h_{1t}}, e^{h_{2t}}, e^{h_{3t}})$ is a 3×3 diagonal matrix whose diagonal log-elements satisfy the recursion

$$\mathbf{h}_t = \boldsymbol{\mu} + \boldsymbol{\Phi}(\mathbf{h}_{t-1} - \boldsymbol{\mu}) + \boldsymbol{\eta}_t, \;\; t = 1, 2, \ldots, T \tag{2}$$

with $\mathbf{h}_t = (h_{1t}, h_{2t}, h_{3t})$ and $\mathbf{h}_0$ a given vector. The matrix $\boldsymbol{\Phi} = (\varphi_{ij})_{i,j \in \{1,2,3\}}$ parameterises a lagged regression of log-volatilities on previous log-volatilities, allowing for a possible cross-country lagged relationship. Long-term mean levels of the log-volatilities are captured by the vector $\boldsymbol{\mu}$.

The vectors $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \varepsilon_{2t}, \varepsilon_{3t})$ and $\boldsymbol{\eta}_t = (\eta_{1t}, \eta_{2t}, \eta_{3t})$ are assumed normally distributed over $t = 1, 2, \ldots, T$, with mean vectors 0 and finite covariance matrices $\boldsymbol{\Sigma}_\varepsilon$ and $\boldsymbol{\Sigma}_\eta$. The collections $\{\boldsymbol{\varepsilon}_t;\ t = 1, 2, \ldots, T\}$ and $\{\boldsymbol{\eta}_t;\ t = 1, 2, \ldots, T\}$ are assumed independent of each other.

The correlation matrix $\boldsymbol{\Sigma}_\varepsilon$ quantifies the correlations of the log-returns, standardized for the dynamic stochastic volatility relationship in Equation (2), across the three countries, while the matrix $\boldsymbol{\Sigma}_\eta$ measures the variances and cross-country covariances in the error terms of the log-volatilities. These matrices are assumed to have the forms:

$$\boldsymbol{\Sigma}_\varepsilon = \begin{pmatrix} 1 & \rho_{\varepsilon 12} & \rho_{\varepsilon 13} \\ \rho_{\varepsilon 21} & 1 & \rho_{\varepsilon 23} \\ \rho_{\varepsilon 31} & \rho_{\varepsilon 32} & 1 \end{pmatrix} \quad \text{and} \quad \boldsymbol{\Sigma}_\eta = \begin{pmatrix} \sigma_1^2 & \sigma_{12} & \sigma_{13} \\ \sigma_{21} & \sigma_2^2 & \sigma_{23} \\ \sigma_{31} & \sigma_{32} & \sigma_3^2 \end{pmatrix}. \tag{3}$$

Let $\rho_{\eta ij} = \sigma_{ij}/(\sigma_i \sigma_j)$ for $i, j = 1, 2, 3$ denote the correlations of the log volatilities. Note that $\mathbf{\Sigma}_\varepsilon$ and $\mathbf{\Sigma}_\eta$ are symmetric matrices whereas the matrix $\mathbf{\Phi}$ is not necessarily symmetric, reflecting possible asymmetric transmission of information between countries.

Residuals

Investigating residuals is a standard method in assessing goodness-of-fit. Residuals can be calculated by solving Equation (1) in the form

$$\boldsymbol{\varepsilon}_t = \mathbf{H}_t^{-1/2}\mathbf{y}_t,$$

and replacing the matrix $\mathbf{H}_t$ by an estimated version. For fixed t we assume the $\boldsymbol{\varepsilon}_t$ to be normally distributed with mean vector 0 and covariance matrix $\mathbf{\Sigma}_\varepsilon$, hence the components of $\boldsymbol{\varepsilon}_t$ are correlated. To investigate whether the dependencies between the different univariate time series are modelled reasonably well, it is, therefore, useful to transform the vector $\boldsymbol{\varepsilon}_t$ into a vector $\boldsymbol{\varepsilon}_t^*$ with uncorrelated components, and to check (after fitting the model) whether the empirical correlations between the components of $\boldsymbol{\varepsilon}_t^*$ are indeed negligible. We deal with this in a standard way by decomposing $\mathbf{\Sigma}_\varepsilon$ as given in Equation (3) into

$$\mathbf{\Sigma}_\varepsilon = \mathbf{P}\mathbf{\Lambda}\mathbf{P}'$$

where $\mathbf{P}$ is orthogonal and $\mathbf{\Lambda}$ is diagonal and contains the (non-negative) eigenvalues of $\mathbf{\Sigma}_\varepsilon$. The existence of such a decomposition is guaranteed (cf. Graybill (1983)), since $\mathbf{\Sigma}_\varepsilon$ is symmetric and non-negative definite. Defining $\mathbf{\Sigma}_\varepsilon^{-1/2} := \mathbf{P}\mathbf{\Lambda}^{-1/2}\mathbf{P}'$, we can transform $\boldsymbol{\varepsilon}_t$ into the random vector

$$\boldsymbol{\varepsilon}_t^* := \mathbf{\Sigma}_\varepsilon^{-1/2}\boldsymbol{\varepsilon}_t = \mathbf{\Sigma}_\varepsilon^{-1/2}\mathbf{H}_t^{-1/2}\mathbf{y}_t \tag{4}$$

which has a standard multivariate normal distribution, i.e., with mean vector 0 and identity covariance matrix. Similarly one can calculate standardized residuals $\boldsymbol{\eta}_t^*$ with uncorrelated components from Equation (2).

Bayesian Inference and the MCMC Method

The latent variable setting of the model is very well suited to the use of an MCMC method (Chib, Nardari, and Shephard, 2002). For inference, we used the WinBUGS 1.4.1 program, a freely available Bayesian software (Spiegelhalter, Thomas, Best, and Gilks, 2003), which performs the Gibbs sampling algorithm (Gilks, Richardson, and Spiegelhalter, 1996; Robert and Casella, 2000) for a user-defined model. Most standard distributions, also multivariate ones, are implemented in WinBUGS so that one has no difficulties in setting up the trivariate SV model specified by Equations (1) and (2).

Of course, in our Bayesian framework we have to specify prior distributions for the parameters $\rho_{\varepsilon ij}$, μ_i, φ_{ij}, σ_i, and $\rho_{\eta ij}$ $(i, j = 1, 2, 3)$. Prior independence of all these parameters is assumed. For the correlations $\rho_{\varepsilon 12}$, $\rho_{\varepsilon 13}$, $\rho_{\varepsilon 23}$, $\rho_{\eta 12}$, $\rho_{\eta 13}$ and $\rho_{\eta 23}$ we used independent uniform priors on $[-1, 1]$, while the components of the mean vector $\boldsymbol{\mu}$ were assumed to have $N(-9, 25)$ priors. This was based on the fact that an initial exploratory analysis of the data indicated average daily log-volatilities of approximately -9. However, in order not to impose too informative priors on the components of $\boldsymbol{\mu}$, we used a relatively large variance of 25. Further, while we might expect the standard deviations σ_i to be rather small, we did not want to impose any further restrictions on their prior distributions. Hence, we simply chose uniform priors on $[0, 1]$ for the σ_i, which subsequently proved to have a sufficiently large support. Note that our choice of the priors for $\rho_{\varepsilon ij}$, $\rho_{\eta ij}$, and σ_i $(i, j = 1, 2, 3, i \neq j)$ guarantees that the covariance matrices $\boldsymbol{\Sigma}_\varepsilon$ and $\boldsymbol{\Sigma}_\eta$ are positive semi-definite. Since we expect the daily volatility for each country to be quite persistent, we took $\text{Beta}(20, 1.5)$ priors for the transformed parameters $(\varphi_{ii} + 1)/2$, $i = 1, 2, 3$. Finally, for the parameters φ_{ij} $(i \neq j)$ uniform priors on $[-1, 1]$ are used. Further general information about prior selection can be found in the Chapter of Robert and Rousseau (2010) in this book.

Before we continue with estimating our data using WinBUGS, we mention that in several simulation studies WinBUGS produced very reliable results for this model and simulated data sets. However, in these simulations WinBUGS was, for the model at hand, very sensitive to the initial values. When the initial values were not very carefully choosen, the program stopped after a few iterations and returned an error message. It turned out that a good strategy to avoid these problems is to run the program for simpler submodels, where a few of the parameters are set to zero. The final state of the sampler can be used as collection of starting values for the next more complicated model (then, of course, initial values for the additional parameters have to be chosen).

The model described above is the main model of interest in this paper; however, some of the parameters, in particular in the matrix $\boldsymbol{\Phi}$, turned out to be not significant in terms of 95% credible intervals. Hence, we reduced the full model step by step, by equating some parameters to zero, until all parameters were significant. The final model derived in this way contains estimates of all elements in $\boldsymbol{\Sigma}_\eta$ and $\boldsymbol{\Sigma}_\varepsilon$, and of course of $\boldsymbol{\mu}$, but only for the diagonal elements of $\boldsymbol{\Phi}$.

Changes in Market Correlations

We now fit the trivariate SV model to the the (centered) returns for all three indices and investigate how the parameters, in particular the correlations $\rho_{\varepsilon ij}$ between the three markets, have changed from the pre- to the post-euro period. After this we discuss further topics related to goodness-of-fit.

Model Estimates

We ran the sampler for a total of 15000 iterations, where the first 5000 iterations were discarded for burn-in. These choices were made after conducting simulation trials with parameters close to those found for the data set at hand. The chains always converged within 3000 to 4000 iterations, and gave very consistent estimates from the iterations after this initial burn-in period. Moreover, to reduce the sample autocorrelations, we used subsampling. In the following, all estimates are based always on the iterations $5005, 5010, \ldots, 15000$, i.e. on 2000 samples.

Table 1 shows posterior mean estimates and 95% credible intervals for all model parameters. All parameters are significantly different from zero, since all credible intervals do not contain zero.

The estimates of μ_1, μ_2 and μ_3 in Table 1 represent the long-term mean levels of the log-volatilities. Since the corresponding credible intervals do not overlap, the increase in this level is significant from the pre- to the post-euro period, for all three countries. Moreover, the levels for the three countries are not significantly different. The parameters σ_1, σ_2 and σ_3 represent the long-term standard deviations of the log-volatilities. However, again judging from the credible intervals, the changes from the pre- to the post-euro period seem not to be very significant. The estimates of φ_{11}, φ_{22} and φ_{33} quantify the dependence of log-volatilities on previous log-volatilities. From the pre- to the post-euro period, they have all increased slightly. Overall, the volatility seems to be stationary, with autocorrelations at most around 0.98.

Our main interest is in the estimates of $\rho_{\varepsilon 12}$, $\rho_{\varepsilon 13}$ and $\rho_{\varepsilon 23}$ which measure the cross-country correlations of returns in our model. Of course, we expect that the introduction of the euro might have increased the connection between the markets in Germany and France, whereas the situation remained quite unchanged for Switzerland. Indeed, our model confirms this conjecture: The correlation $\rho_{\varepsilon 12}$ between Germany and France was about 0.61 before the introduction of the euro, and increased heavily to a level of about

Table 1: *Posterior mean estimates and 95% credible intervals for the parameters in the full model. Subscript 1 denotes Germany, subscript 2 denotes France, and subscript 3 denotes Switzerland. 1565 daily returns in the pre-period and 1566 in the post period.*

	Pre-Period		**Post-Period**	
Parameters	Post. Mean	95% Cred. Int.	Post. Mean	95% Cred. Int.
μ_1	**-9.4281**	[-9.7010,-9.1690]	**-8.6842**	[-8.9310,-8.4270]
μ_2	**-9.4053**	[-9.5770,-9.2230]	**-8.6875**	[-8.9200,-8.4220]
μ_3	**-9.3349**	[-9.5870,-9.0720]	**-8.6986**	[-8.9020,-8.4610]
σ_1	**0.1416**	[0.1176, 0.1714]	**0.1451**	[0.1205, 0.1719]
σ_2	**0.1597**	[0.1349, 0.1871]	**0.1176**	[0.0998, 0.1352]
σ_3	**0.1304**	[0.1084, 0.1600]	**0.1205**	[0.0927, 0.1552]
φ_{11}	**0.9649**	[0.9500, 0.9779]	**0.9812**	[0.9697, 0.9895]
φ_{22}	**0.9591**	[0.9415, 0.9726]	**0.9813**	[0.9705, 0.9892]
φ_{33}	**0.9699**	[0.9526, 0.9823]	**0.9767**	[0.9588, 0.9881]
$\rho_{\varepsilon 12}$	**0.6140**	[0.5793, 0.6443]	**0.8686**	[0.8573, 0.8793]
$\rho_{\varepsilon 13}$	**0.4586**	[0.4140, 0.4980]	**0.5238**	[0.4904, 0.5569]
$\rho_{\varepsilon 23}$	**0.5006**	[0.4641, 0.5355]	**0.5583**	[0.5259, 0.5913]
$\rho_{\eta 12}$	**0.7842**	[0.7045, 0.8459]	**0.9384**	[0.9077, 0.9616]
$\rho_{\eta 13}$	**0.7560**	[0.6298, 0.8541]	**0.8013**	[0.6502, 0.9189]
$\rho_{\eta 23}$	**0.8845**	[0.8002, 0.9505]	**0.8385**	[0.7374, 0.9268]

0.87 after the introduction of the euro. Judging from the credible intervals, this increase is most significant. The correlation between Germany and Switzerland and between France and Switzerland (represented by $\rho_{\varepsilon 13}$ and $\rho_{\varepsilon 23}$, respectively), however, increased only very slightly, and, looking again at the credible intervals, not significantly.

It is interesting that we find similar results also for the parameters $\rho_{\eta 12}$, $\rho_{\eta 13}$ and $\rho_{\eta 23}$ which quantify the volatility linkages. Only $\rho_{\eta 12}$, measuring the linkage between Germany and France, has increased significantly from the pre- to the post-euro period. Again, the change for Germany-Switzerland and France-Switzerland is not significant.

To illustrate the precision of the parameter estimates in Table 1 and highlight the most significant changes associated with the introduction of the euro, Figure 3 contains density plots of the estimated marginal posterior distributions for the most important model parameters. Each of the nine plots shows the estimated densities for both the pre- and

Figure 3: *Estimated marginal posterior density plots, pre- and post-euro. Diagonal: Volatilities* σ_1, σ_2, σ_3. *Upper triangle: Return Correlations* $\rho_{\varepsilon 12}$, $\rho_{\varepsilon 13}$, $\rho_{\varepsilon 23}$. *Lower triangle: Log-Volatility Correlations* $\rho_{\eta 12}$, $\rho_{\eta 13}$, $\rho_{\eta 23}$.

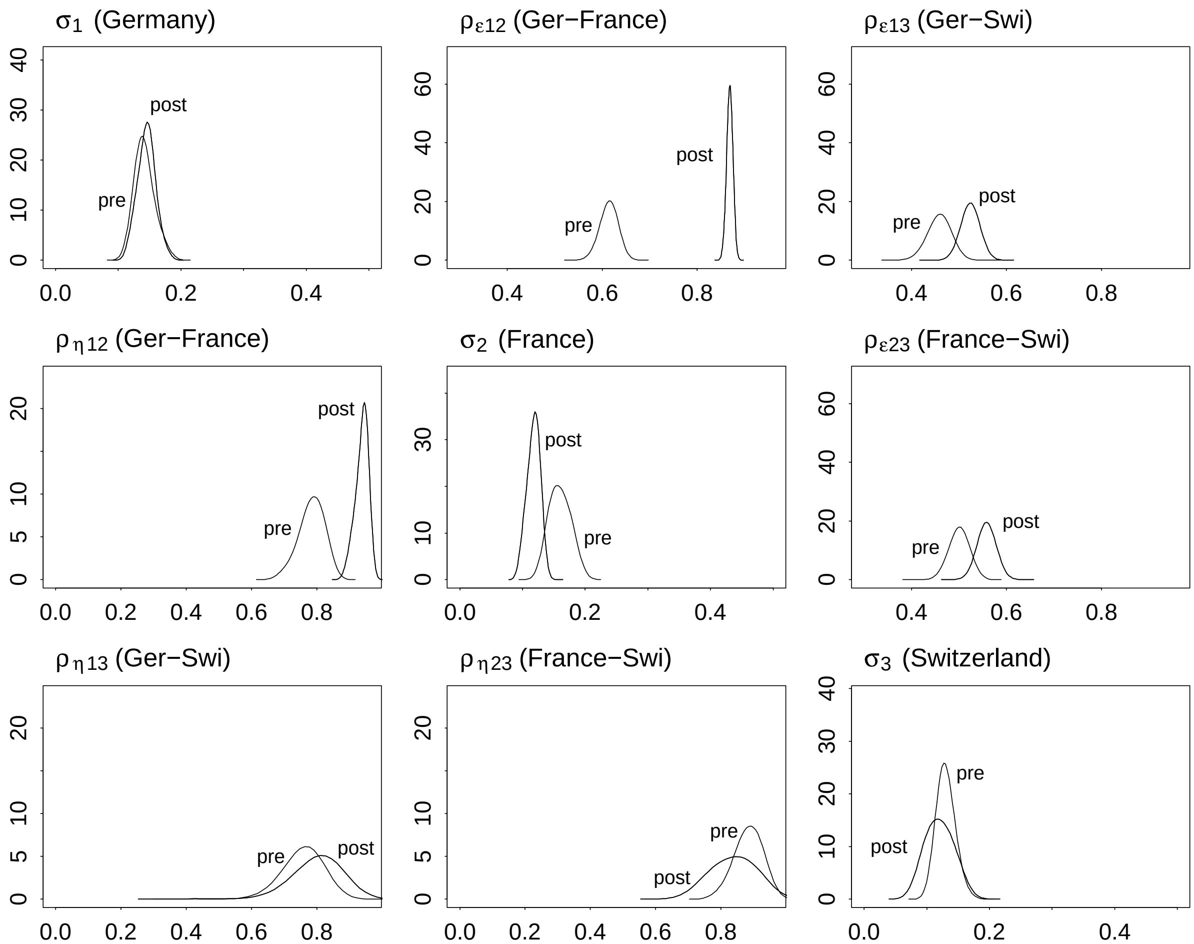

the post-euro period. The diagonal contains the volatilities σ_1, σ_2, and σ_3 in the log-volatility Equation (2) of the trivariate SV model. The upper triangle shows the densities for the return correlations $\rho_{\varepsilon 12}$, $\rho_{\varepsilon 13}$, and $\rho_{\varepsilon 23}$ in Equation (1), the lower triangle the log-volatility correlations $\rho_{\eta 12}$, $\rho_{\eta 13}$, and $\rho_{\eta 23}$. The increase in the correlations between Germany and France is clearly seen, whereas the correlations between Germany and Switzerland and between France and Switzerland remain quite the same.

Goodness-of-Fit

For assessing goodness-of-fit of a model to the observed data within a Bayesian framework, a suitable approach is a *posterior predictive analysis* (PPA). This kind of analysis has been thoroughly studied in the literature; see for example Gelman, Carlin, Stern, and Rubin (1995) or Gelman, Meng, and Stern (1996).

The PPA requires the estimation of the *posterior predictive distribution*. Samples from this distribution can be obtained from the MCMC procedure. We did this for the pre-euro data based on USD, using the iterations numbered $k = 5010, 5020, 5030, \ldots, 15000$ from the WinBUGS MCMC output. Thus the analysis is based on 1000 replicated data sets each of length 1565. The crux of the idea is to check whether the replicated data which is simulated according to the postulated, fitted, model shows the same characteristics as the original data set. We make the test here by overlaying on histograms of the replicated, simulated statistics, the corresponding observed values of the statistics.

The choice of discrepancy statistic is one of convenience. To reflect the major features of the SV model, we use discrepancy functions based on squared returns and correlations between returns. Thus, 3 of our discrepancy statistics are simply the sums of squares of the returns for the 3 countries, and another 3 are the correlations between the standardized sums of squares, for the 3 countries.

The 1000 values of each of the 3 discrepancy functions for the squared returns from the *replicated* data sets are summarized in the histograms in the first row of Figure 4. The corresponding *observed* values for the pre-euro data are added as vertical lines. Since these vertical lines lie well in the middle of the support of the histograms, we conclude that there is no significant evidence against the model.

Similarly, the 1000 values of the correlation discrepancy statistics for the replicated data sets are summarized in histograms in the second row of Figure 4, along with vertical solid lines representing the means of the observed data and vertical dashed lines representing the corresponding 2.5% and 97.5% quantiles. Thus, as in the first row, the histograms display the 1000 discrepancy statistics calculated from the replicated data, whereas the vertical (solid and dashed) lines refer to the same statistics derived from the observed pre-euro data. Again the diagrams indicate that our model describes the correlation structure of the data very well.

We finally note also that from similar plots for the post-euro data we again concluded that there is no significant evidence against the model. Since the plots look very similar to those given in Figure 4, we omit them here.

Figure 4: *Posterior Predictive Analysis. Histograms of discrepancy statistics for replicated data. Vertical (solid and dashed) lines refer to the discrepancy statistics derived for the pre-euro data. The vertical solid lines represent the means of the observed data whereas the vertical dashed lines indicate the 2.5% and 97.5% quantiles. The discrepancy statistics are the sums of squared returns for the 3 countries, Germany, France and Switzerland ($D_{(1)}$, $D_{(2)}$, $D_{(3)}$), and the correlations between returns for the 3 countries ($D_{(1,2)}$, $D_{(1,3)}$, $D_{(2,3)}$).*

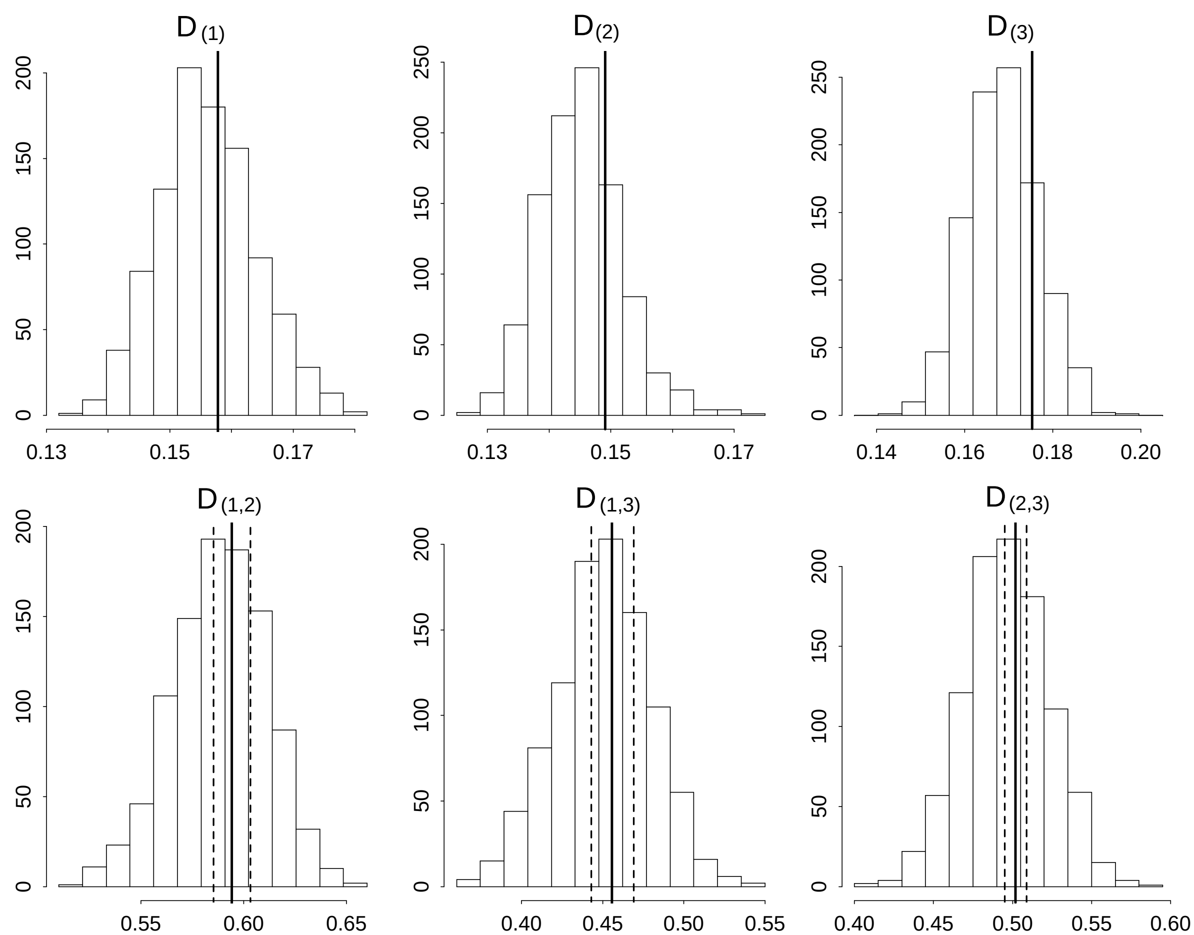

We note that it may be worth to repeat the above analysis with other (also not aggregated) statistics. Besides, it is, of course, possible and useful to assess goodness-of-fit by other methods, from which we just mention a few exemplarily. One of the most common method for checking goodness-of-fit is, without any doubt, to investigate residuals. One can find the residuals according to Equation (4), by replacing all h_{it} (appearing in $\mathbf{H}_t$) and $\rho_{\varepsilon ik}$ (appearing in $\mathbf{\Sigma}_\varepsilon$) by their sample values $\hat{h}_{it}^{(k)}$ and $\hat{\rho}_{\varepsilon ij}^{(k)}$, respectively, from iteration k of the MCMC chain. Hence, for each iteration k of the MCMC sampler one gets a series of (in this case three-dimensional) residuals, which can be analyzed in terms of auto- and crosscorrelations. Furthermore, one can check the normality assumption for the components of the residuals by deriving, for instance, Jarque-Bera statistics for all components. This provides, for each component, a Bayesian posterior distribution for each of the Jarque-Bera statistics.

Another issue is to check the prior sensitivity of the analysis. Due to the relatively large number of observations, one does not expect a too high sensitivity. Indeed, in our study, the selection of other (not too concentrated) priors did not lead to significantly other results.

Portfolio management

In this section we illustrate how the Bayesian approach can help to quantify risks and uncertainty an investor is faced with when selecting and holding a portfolio. We choose two classical approaches, value-at-risk analysis and minimum variance portfolio selection. Of course, all ideas can be easily carried over to other portfolio management strategies.

Value-at-Risk

Here we are interested in estimates for the value-at-risk at a confidence level of 99% over the time horizon of one day resulting from our stochastic volatility model. As an example, we consider the post-euro data at time $t = 1566$. We calculate the VaR(99%) for the three different countries, Germany, France and Switzerland, assuming that we have three portfolios representing the DAX, CAC, and SMI, valued 100 USD on day 1566, each. Hence, the VaR is the amount of loss in USD which is exceeded within one day only with probability 1%. Additionally, our Bayesian approach enables us to derive the uncertainty of the estimated VaRs for the three portfolios. For each of the iterations numbered $k = 5005, 5010, \ldots, 15000$ from the MCMC output, we can calculate the VaRs(99%) using the samples $\hat{h}_{i,1566}^{(k)}$ of $h_{i,1566}$, $j = 1, 2, 3$, by looking at the 1% quantiles of a corresponding distribution. Recall that the variables $\boldsymbol{\varepsilon}_t$ are assumed to be normal and that we use centered data in our analysis. Hence, the VaRs have to be corrected finally by the long-run mean of the returns (which is, of course, quite negligible over the

Figure 5: *Estimated densities for the one-day VaR(99%), for the DAX, CAC and SMI portfolio, valued 100 USD at time* $t = 1566$ *of the post-euro period. Dashed lines indicate posterior mean estimates of the VaRs.*

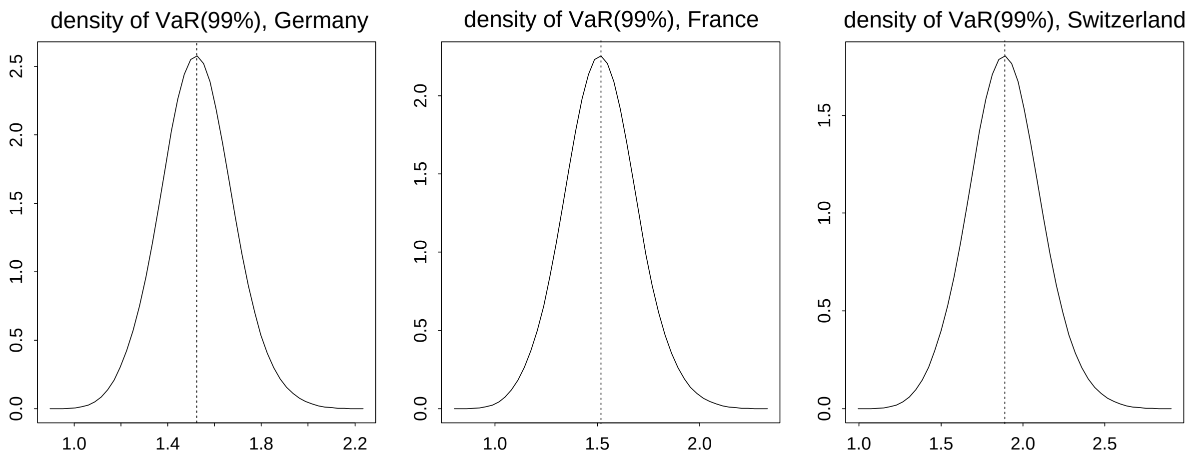

period of one day compared to the standard deviation). Figure 5 shows the estimated densities for the VaR(99%), for the DAX, CAC and SMI portfolio, valued 100 USD at time $t = 1566$ of the post-euro period.

The dashed lines indicate the posterior mean estimates of the VaRs. The posterior mean estimates are 1.53 for the DAX, 1.52 for the CAC, and 1.89 for the SMI. It is interesting to compare these values with quantile estimates derived directly from the data. Using the daily long-run means over the whole post-period as well as estimated daily volatilities from the last 100 days of the post-period, one gets VaR(99%) estimates of 1.54, 1.59, and 2.17, for the DAX, CAC and SMI, respectively. Note, however, that these numbers are quite sensitive to the length of the window which is used for the volatility estimation.

Judging from the posterior mean estimates, the CAC portfolio has the lowest risk, since with probability 99% the loss within one day is not greater than 1.52 USD (i.e. 1.52% of the value of the portfolio). For the DAX portfolio, this loss is not greater than 1.53 USD within one day. For the SMI portfolio, however, there is a chance of 1% that the loss within one day is greater than 1.89 USD. The density curves, however, show that one should treat these numbers with care. A 95% credible interval for the DAX-VaR is $[1.26, 1.80]$, for the CAC-VaR $[1.22, 1.83]$, and for the SMI-VaR $[1.51, 2.28]$, corresponding to the relative large uncertainties of about 35%, 40%, and 41%, respectively. This indicates that the risk of holding a portfolio may be underestimated significantly by relying only on point estimates, such as the posterior mean or the maximum likelihood estimates. Hence, this example clearly shows that the uncertainty about VaR estimates plays an important role in practice. In particular, credible intervals (or other information about the uncertainty, as e.g. marginal density plots) should always be provided when risk figures are reported or used for decision making.

Minimum variance portfolios

Next we illustrate how the uncertainty, an investor is faced with when selecting a portfolio, can be quantified. We consider the classical case that a risk-averse investor likes to invest 1 USD in a minimum variance portfolio comprised of an amount w_1 of the DAX, an amount w_2 of the CAC, and an amount of w_3 of the SMI. At time t the portfolio is, therefore, described by the weights $\mathbf{w}_t = (w_{1t}, w_{2t}, w_{3t})'$, with $w_{1t} + w_{2t} + w_{3t} = 1$, and its variance v_t^2 at time t is

$$v_t^2 = w_{1t}^2 v_{1t}^2 + w_{2t}^2 v_{2t}^2 + w_{3t}^2 v_{3t}^2 + 2w_{1t}w_{2t}\mathrm{cov}_{12t} + 2w_{1t}w_{3t}\mathrm{cov}_{13t} + 2w_{2t}w_{3t}\mathrm{cov}_{23t}, \qquad (5)$$

where v_{1t}^2, v_{2t}^2 and v_{3t}^2 denote the variances of the three assets at time t, respectively, and cov_{12t}, cov_{13t} and cov_{23t} the covariances between these assets at time t. Setting $w_{3t} = 1 - w_{1t} - w_{2t}$, Equation (5) can be considered a function of the two variables w_{1t} and w_{2t}. It achieves its minimum at

$$\begin{aligned} w_{1t} &= \frac{z_{23t}(v_{3t}^2 - \mathrm{cov}_{13t}) - x_{3t}(v_{3t}^2 - \mathrm{cov}_{23t})}{z_{13t}z_{23t} - x_{3t}^2}, \\ w_{2t} &= \frac{z_{13t}(v_{3t}^2 - \mathrm{cov}_{23t}) - x_{3t}(v_{3t}^2 - \mathrm{cov}_{13t})}{z_{13t}z_{23t} - x_{3t}^2}, \end{aligned}$$

where $z_{13t} = \mathrm{var}(\mathrm{asset}_1 - \mathrm{asset}_3) = v_{1t}^2 + v_{3t}^2 - 2\mathrm{cov}_{13t}$, $z_{23t} = \mathrm{var}(\mathrm{asset}_2 - \mathrm{asset}_3) = v_{2t}^2 + v_{3t}^2 - 2\mathrm{cov}_{23t}$, and $x_{3t} = v_{3t}^2 + \mathrm{cov}_{12t} - \mathrm{cov}_{13t} - \mathrm{cov}_{23t}$. Hence, one can easily determine the minimum variance portfolio given estimates of v_{1t}, v_{2t}, v_{3t}, cov_{12t}, cov_{13t} and cov_{23t}.

In our framework, it is straightforward to plug in the corresponding estimates from the MCMC iterations while fitting the trivariate SV model to the data. In particular, at time t we have $v_{it}^2 = e^{h_{it}}$ and $\mathrm{cov}_{ijt} = e^{(h_{it}+h_{jt})/2}\rho_{\varepsilon ij}$, for $i, j = 1, 2, 3$, $i \neq j$.

Based on the data sets from 1993 to 1998 and 1999 to 2004, respectively, we can now calculate estimates of $\mathbf{w}_t$ in each iteration k of the MCMC sampler, using the actual samples $\hat{h}_{it}^{(k)}$ of h_{it} and $\hat{\rho}_{\varepsilon ij}^{(k)}$ of $\rho_{\varepsilon ij}$. This gives us, for each time $t = 1, \ldots, T$, 2000 estimates $\hat{\mathbf{w}}_t^{(k)}$ of $\mathbf{w}_t$, with $k = 5005, 5010, \ldots, 15000$.

Figure 6: *Estimated minimum variance portfolios consisting of DAX, CAC and SMI. The lower curve indicates the posterior mean estimates for the DAX part, the difference between the upper and the lower curve the posterior mean estimates for the CAC part, and the difference between 1 and the upper curve the SMI part in the minimum variance portfolio.*

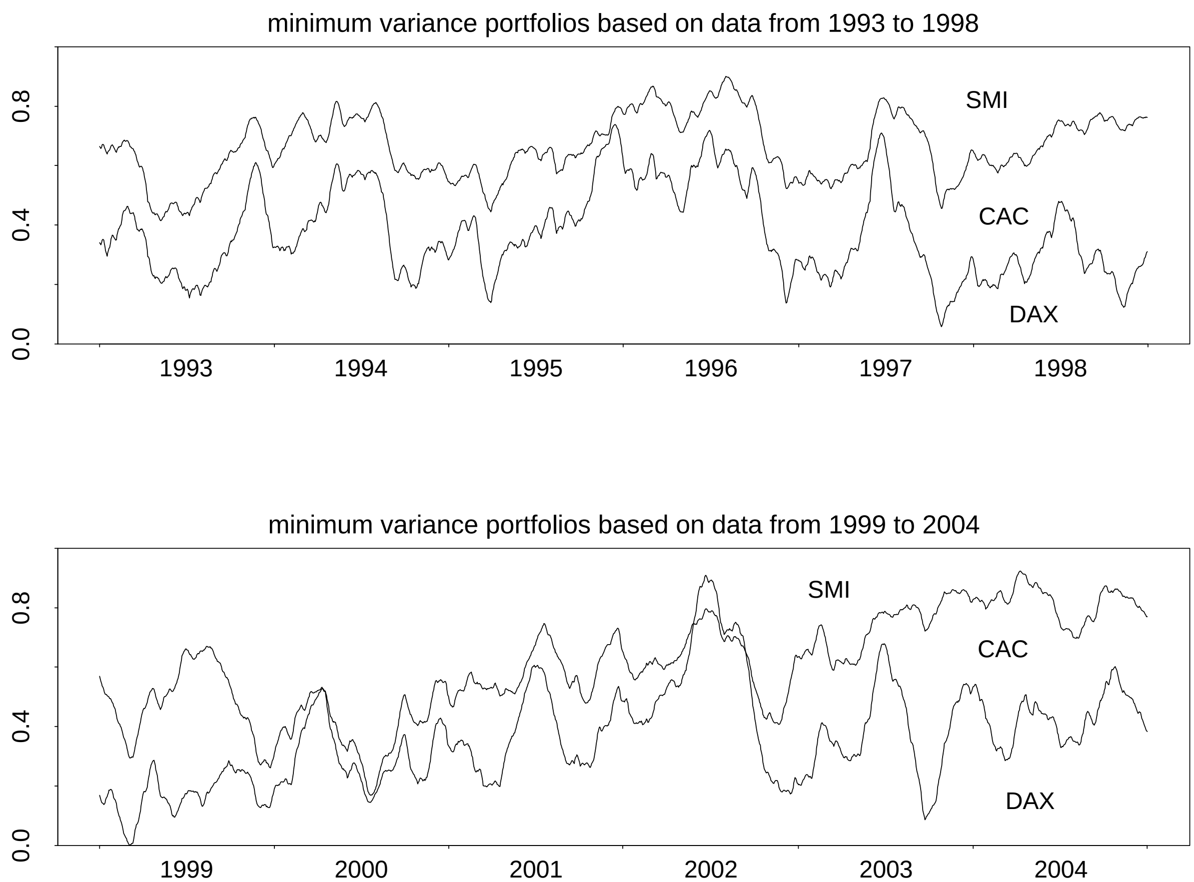

Figure 6 shows, for each $t = 1, \ldots, T$ the corresponding posterior mean estimates of w_{1t} and $w_{1t}+w_{2t}$. Hence, the lower curve indicates the posterior mean estimates for the DAX part, the difference between the upper and the lower curve the posterior mean estimates for the CAC part, and the difference between 1 and the upper curve the SMI part in the minimum variance portfolio. The result is no surprise, the weights vary according to the estimated variances and correlations in the trivariate SV model. Note, however, that all estimates are based on the whole data sets. Of course, for $t < T$ it might be more realistic to estimate the variances and correlations using only the data up to time t, meaning that one has to refit the model 1565+1566 times to get the corresponding plot. To do this *ex post*, however, is computationally quite expensive: Since one model fit would take about 2.5 hours on average, it would take approximately 300 days to create a plot where the portfolio weights are always based only on the data before.

In particular, we get the posterior mean estimates 0.3106 and 0.4514 for $w_{1,1565}$ and $w_{2,1565}$, respectively, in the pre-euro period, and 0.3734 and 0.3662 for $w_{1,1566}$ and $w_{2,1566}$, respectively, in the post-euro period. The investor might use these values to choose the portfolio, however, he/she should be aware that due to the estimation of the variances and correlations, there remains some uncertainty about the optimality of this choice. To illustrate how big this uncertainty is in this case, we investigate the 2000 pairs $(\hat{w}_{1T}^{(k)}, \hat{w}_{2T}^{(k)})$ for $k = 5005, 5010, \ldots, 15000$ (recall again that $T = 1565$ in the pre-euro period and $T = 1566$ in the post-euro period). Since we expect w_{1T} and w_{2T} to show a certain correlation, a scatterplot of the pairs seems to be more informative in our situation than two univariate marginal density plots to assess the uncertainty around the posterior mean estimates. Figure 7 shows these pairs for Dec 31, 1998 and Dec 31, 2004. In both cases, the uncertainty about the (w.r.t. minimum variance) optimal portfolio is quite large: The samples vary heavily around the posterior means (0.3106,0.4514) (for Dec 31, 1998) and (0.3734,0.3662) (for Dec 31, 2004), which are indicated by the dashed lines. Moreover, also negative weights and weights greater than 1 occur, and w_{1T} and w_{2T} are obviously negatively correlated.

Summary

Our analysis confirms that the market linkages between Germany and France were significantly increased by the introduction of the euro. The estimate of the correlation parameter $\rho_{\varepsilon 12}$ increased from 0.61 in the pre-euro period to 0.87 in the post-euro period. By contrast, the correlations between Germany and Switzerland and between France and Switzerland have hardly changed in 1999, and remained at around 0.46 to 0.56. The same is true for the volatility correlation, where we found a significant increase only between Germany and France. The use of Bayesian methods was illustrated on two examples in the area of portfolio management. Our approach allowed to derive Bayesian estimates of value at risk and of portfolio weights for minimum variance portfolios. This strategy may help for a more realistic investigation of the risk and uncertainty coming from selecting and holding asset portfolios. In particular, we showed that typical figures

Figure 7: *Samples from the bivariate posterior density for the minimum variance portfolios on Dec 31, 1998 and Dec 31, 2004, based on the data within 6 years before each. The x-axis shows the estimated proportion for the DAX part, the y-axis the estimated proportion for the CAC part. The dashed lines indicate the posterior mean estimates.*

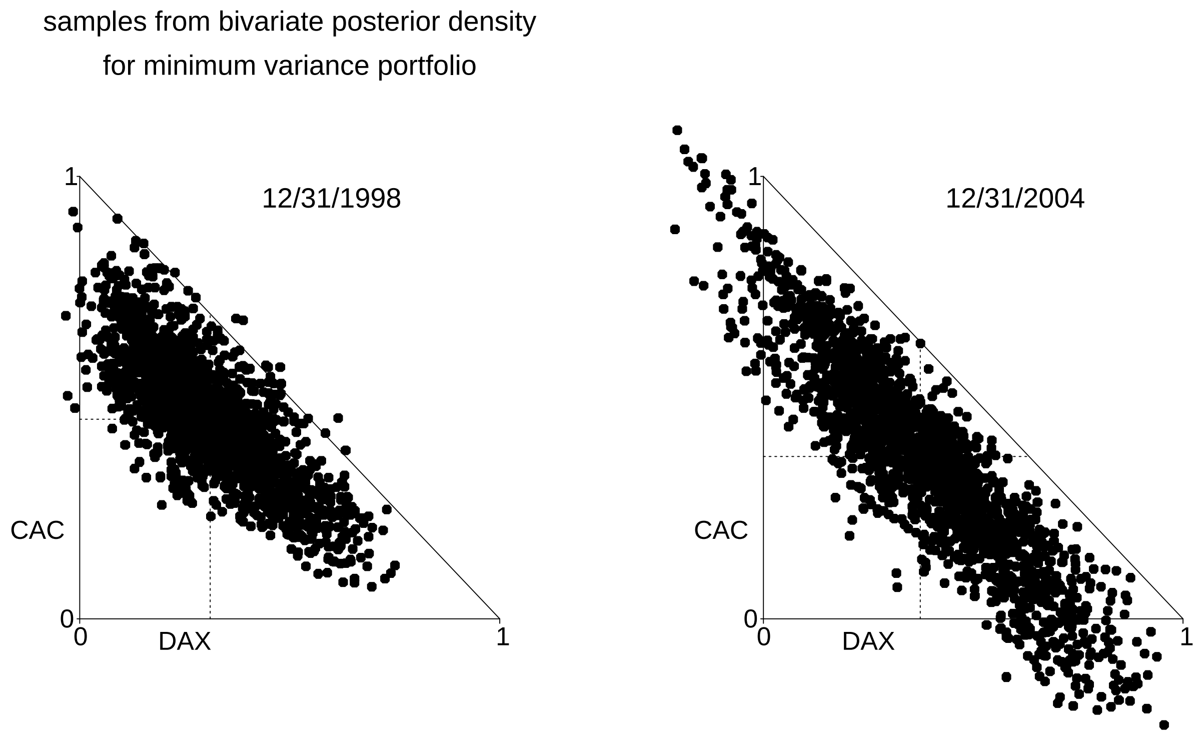

used in risk management practice (such as VaR or portfolio weights) are afflicted with high uncertainties that can hardly be neglected.

References

Chib, S., F. Nardari, and N. Shephard (2002). Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics 108*, 281–316.

Fleming, J., C. Kirby, and B. Ostdiek (1998). Information and Volatility Linkages in the Stock, Bond and Money Markets. *Journal of Financial Economics 49*, 111–137.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin (1995). *Bayesian Data Analysis*. London: Chapman and Hall.

Gelman, A., X.-L. Meng, and H. S. Stern (1996). Posterior Predictive Assessment of Model Fitness via Realized Discrepancies. *Statistica Sinica 6*, 733–807.

Gilks, W. R., S. Richardson, and D. J. Spiegelhalter (1996). *Markov Chain Monte Carlo in Practice*. London: Chapman and Hall.

Graybill, F. A. (1983). *Matrices with Applications in Statistics*. Belmont, CA: Wadsworth.

Kim, S. J., F. Moshirian, and E. Wu (2005). Dynamic stock market integration driven by the european monetary union: An empirical analysis. *Journal of Banking and Finance 29*, 2475–2502.

Moore, T. (2007). Has entry to the european union altered the dynamic links of stock returns for the emerging markets? *Applied Financial Economics 17*, 1431–1446.

Robert, C. P. and G. Casella (2000). *Monte Carlo Statistical Methods*. New York: Springer.

Ruiz, E. (1994). Quasi-Maximum Likelihood Estimation of Stochastic Variance Models. *Journal of Econometrics 63*, 284–306.

Spiegelhalter, D., A. Thomas, N. Best, and W. Gilks (2003). *WinBUGS User Manual (Version 1.4)*. Cambridge, UK: MRC Biostatistics Unit.

Taylor, S. J. (1994). Modelling Stochastic Volatility: A Review and Comparative Study. *Mathematical Finance 4*, 183–204.

Westermann, F. (2004). Does the Euro affect the dynamic interactions of stock markets in Europe? Evidence from France, Germany and Italy. *The European Journal of Finance 10*, 139–148.