Computationally intensive techniques for a fully Bayesian, decision theoretic approach to financial forecasting and portfolio selection.

Andrew Simpson
Andrew.Simpson@ncl.ac.uk

Darren J Wilkinson
d.j.wilkinson@ncl.ac.uk

Department of Statistics, University of Newcastle,
Newcastle upon Tyne, NE1 7RU, U.K.

Abstract

This paper considers the problem of modelling and forecasting for multivariate financial time series. The use of Dynamic Linear State Space models and Stochastic Volatility models with Kalman filtering techniques to address this problem is considered in the context of providing a modular software implementation. The combination of these two approaches is presented with an illustrative example. We also show how a marginal posterior forecast distribution may be used in order to implement a fully Bayesian decision theoretic approach to portfolio selection.

Keywords: Bayesian; DLMs; Finance; Forecasting; Multivariate, Monte Carlo Markov Chain (MCMC), Stochastic Volatility, Portfolio Management.

1 Introduction

The multivariate behaviour of stock and currency prices has long been of interest to statistical researchers. The prices of related stocks often exhibit dependencies between series as well as the usual dependencies over time. This has lead to the proposal of many different classes of model to attempt to explain this behaviour. The multivariate dynamic linear state space models of West and Harrison (1997) are often appropriate for explaining log-price behaviour. The basic model and approaches to inference are laid out in section 2 with particular reference to the development of an object-oriented software library to make inference on these models.

These simple models suffer from a failure to allow the volatility within the model to evolve over time. Various models have been suggested in the literature to explain this temporal evolution, amongst the most popular have been the ARCH Engle (1982) and GARCH Bollersev (1987) family of models. In this family of models the serial correlation of the volatility over time is captured in a deterministic model. Other approaches rely upon Stochastic Volatility models for which new computational techniques are required. Barndorff-Nielsen and Shephard (2001) apply non-Gaussian Ornstein-Uhlenbeck processes to these models, while other authors including Kim, Shephard, and Chib (1998), Jacquier, Polson, and Rossi (1994, 1999) and Aguilar and West (2000) and Aguilar, Huerta, Prado, and West (1999) have applied MCMC techniques to the basic log-Gaussian Stochastic Volatility model. In section 3 we will outline an approach to making inference on the models. In section 4 we will look at multivariate extensions of the basic Stochastic Volatility model.

In section 5 we will consider methods for developing composite models incorporating stochastic volatility within the basic dynamic linear model form. Section 6 considers the application of the results from these inferences to the portfolio selection question. Finally section 7 illustrates the process through use of a simple example.

2 The Dynamic Linear Model (DLM)

The multivariate dynamic linear state space models of West and Harrison (1997) are, as we have already stated, often appropriate for explaining log-price behaviour in financial time series. The basic form of these models is as follows

$$\begin{aligned} &\text{The Observation Equation:} \quad && Y_t = F'\theta_t + \nu_t \quad && \nu_t \sim N(0, V_t) \\ &\text{The System Equation:} \quad && \theta_t = G\theta_{t-1} + \omega_t \quad && \omega_t \sim N(0, W_t). \end{aligned}$$

If the variance is known, or is constant, this is a relatively simple problem, requiring a simulation smoother, or other Kalman filtering techniques. This is extensively covered in the literature, for example see West and Harrison (1997). We have written and developed a sather class to carry out the necessary calculations. Sather is a fast efficient typesafe object oriented programming language which was initially developed at ICSI and is now a GNU project. Sather is particularly well suited to complex problems in scientific computing.

The class KALMAN packages together a set of routines that are required when making inference about data being modelled by a DLM. It includes a Kalman filter and a smoother; the latter being included to allow for inference via an EM algorithm, see Koopman (1993) for an example of this approach. The MCMC approach outlined in West and Harrison (1997) requires a simulation smoother which is also included. Two further routines are also included in the class to simulate forecasts from both the state and observations at future time points. These are of key importance to the advantages of the MCMC approach as they allow simulation from the full joint posterior density of the hidden states and the future observables, thus allowing analyses of forecasting problems which integrate over all uncertainties of the model.

To make inference on this dynamic linear model a block Gibbs sampler may be constructed using the KALMAN class as one major element, sampling from the state, given the data and the model parameters. Samples from the conditional distribution of the variance/covariance matrices can be generated using semi-conjugate inverse Wishart priors.

The algorithm for inference on this model is as follows

1. Sample from the state θ_t by running Kalman filter and simulation smoother, using KALMAN class in sather. This gives posterior samples conditional on the data and all other parameters

2. Update V using semi-conjugate inverse Wishart prior.

3. Update W using semi-conjugate inverse Wishart prior.

4. Repeat steps 1 to 3 until convergence.

It is, however, unlikely that the variance/covariance matrices are either known or fixed and this prompts us to consider alternative approaches.

3 The Univariate Stochastic Volatility Model (ISV)

Let us set to one side the dynamic linear model format for a while and consider approaches to modelling time varying stochastic volatility.

In this case the series of interest y_t could be the one step returns of a univariate time series, or the error term ν_t or ω_t from a univariate DLM. Consider the following model

$$y_t = \sigma_t \varepsilon_t, \qquad \text{where} \qquad \varepsilon_t \sim N(0,1)$$

Here, the series of interest is made up of some white noise multiplied by a standard deviation σ_t. If we write $\alpha_t = \log {\sigma_t}^2$, this becomes

$$y_t = \varepsilon_t \exp\left(\frac{\alpha_t}{2}\right)$$
$$\Rightarrow \log {y_t}^2 = \alpha_t + \log {\varepsilon_t}^2.$$

If we let α_t (the log volatility) evolve through time according to a centred $AR(1)$ process

$$\alpha_t = \mu + \phi(\alpha_{t-1} - \mu) + \eta_t \qquad \text{where} \qquad \eta_t \sim N(0, {\sigma_\eta}^2),$$

we can see similarities to our DLM from section 2. Using the notation of Pitt and Shephard (1999) we can represent this as $y_t \sim ISV_n(\phi; \sigma_\eta; \mu)$ that is a series $y = (y_1, \ldots, y_n)'$ arises from a stochastic volatility model, conditionally independent of any other series.

We obviously want to make inference on this Stochastic Volatility model. The mean and variance of $\log {\varepsilon_t}^2$ are known standard results, being -1.27 and $\pi^2/2$ respectively. Initially we assume that $\log {\varepsilon_t}^2$ is Gaussian with these parameters, as per Harvey, Ruiz, and Shephard (1994) who used this assumption to allow Kalman filtering and smoothing to produce quasi-maximum likelihood estimators of the model parameters. Assuming normality we can sample from the state of our model α_t conditional on the data and all other model parameters using Kalman filtering and simulation smoothing techniques. We have incorporated elements of our KALMAN class into a new sather class ISV which we have written to allow inference on these models.

As with our dynamic linear model we can include the simulation smoothing routine into a block Gibbs sampler, so having sampled from the state α_t given the data and the model parameters, we can the sample from the full-conditional of each of our parameters in turn.

The algorithm for inference on this model is as follows.

1. Sample from the state α using Kalman filtering and simulation smoothing.
2. Sample from the full-conditional of σ_η, given the states and all other model parameters; see section 3.1 for details.
3. Sample from the full-conditional of μ, given the states and all other model parameters; see section 3.2 for details.
4. Sample from the full-conditional of ϕ, given the states and all other model parameters; see section 3.3 for details.
5. Repeat steps 1 to 4 until convergence.

3.1 Update of σ_η

We know that $\eta_t \sim N(0, {\sigma_\eta}^2)$ and hence we can use a standard conjugate Gamma update on the precision. Assuming a prior of ${\sigma_\eta}^{-2} \sim \Gamma(a,b)$

$$\sigma_\eta^{-2}|\cdot \sim \Gamma\left(a + \frac{n}{2}, b + \frac{1}{2}\sum_{t=1}^{n} {\eta_t}^2\right).$$

From this full-conditional we can sample a realisation of σ_η. The necessary summation is provided in the KALMAN class, and replicated within the ISV class.

3.2 Update of μ

Assume a Normal prior on μ i.e. $\mu \sim N(a,b)$ and using standard results

$$\mu|\cdot \sim N(m,v)$$

where

$$m = v\left(\frac{a}{b} + \frac{\sum_{t=1}^{n}(\alpha_t - \phi\alpha_{t-1})(1-\phi)}{{\sigma_\eta}^2}\right) \qquad v = \left(\frac{1}{b} + n\frac{(1-\phi)^2}{\sigma_\eta^2}\right)^{-1}.$$

We can sample μ from it's full-conditional. Again the necessary summations are provided within the ISV class.

3.3 Update of ϕ

We could assume a semi-conjugate normal prior for ϕ allowing a simple Gibbs update. Unfortunately this leads to identifiability and convergence problems, due to the fact that the prior is not restricted in the range $(0,1)$, which would capture the prior belief that the volatility process is stationary with positive dependence.

To overcome this we assume a Beta prior for ϕ, i.e. $\phi \sim \beta(a,b)$ as this is a flexible class of distributions in the range $(0,1)$. This is a non-conjugate prior, however we can easily sample from the posterior

$$\pi(\phi|\cdot) \propto \exp\left(\frac{-1}{2v}(\phi - m)^2\right)\phi^{(a-1)}(1-\phi)^{(b-1)}$$

where

$$m = v\left(\frac{\sum_{t=2}^{n}(\alpha_t - \mu)(\alpha_{t-1} - \mu)}{{\sigma_\eta}^2}\right) \qquad v = \left(\frac{\sum_{t=2}^{n}(\alpha_{t-1} - \mu)^2}{\sigma_\eta^2}\right)^{-1}$$

using a random walk Metropolis-Hastings scheme. This relies upon the symmetric nature of the random walk to reduce the acceptance probability to $\min\{1, \pi(\phi^*|\cdot)/\pi(\phi|\cdot)\}$, where ϕ^* is a new value proposed for ϕ from the random walk and of course ϕ is the current value. The algorithm for the update of ϕ is as follows

1. Simulate a proposed ϕ called ϕ^* from a random walk based on the current value of ϕ and some set small variance.

2. Calculate the acceptance probability using the posterior we have for ϕ and compare with a random uniform and accept the new value if this is smaller than the acceptance probability.

It is however, for stability reasons at the computational level, necessary to work with the log of the acceptance probability, which is as follows, and compare this with a log uniform value.

$$\log(A) = \log(\pi(\phi^*|\cdot)) - \log(\pi(\phi|\cdot))$$

The ISV class contains these calculations and the update can be called from a sather program that has defined an object as class ISV.

3.4 Correcting for the Normal Approximation

The samples of α_t produced by the Kalman filter simulation smoother are dependent on the Gaussian approximation used for $\log \varepsilon_t{}^2$, which is not valid. It is possible to correct for this approximation using a mixture of normals as per Kim, Shephard, and Chib (1998); however we choose to adopt the approach of Pitt and Shephard (1999) incorporating a Metropolis-Hastings step into our algorithm to make this correction.

We know we have a target density, $p(\alpha|\cdot)$ under the true model, so we can propose values for α using the normal approximation and only accept a proportion of these proposed moves using a Metropolis-Hasting acceptance probability. The proposed value is the value of the state α, generated by the Kalman filter/simulation smoother under the current values of each parameter. We will call this proposed value α^* and the likelihood under the normal approximation we will denote by $q(\alpha|\cdot)$. The Metropolis-Hastings acceptance probability is $\min\{1, A\}$where A is

$$A = \frac{p(\alpha^*|\cdot)}{q(\alpha^*|\cdot)} \bigg/ \frac{p(\alpha|\cdot)}{q(\alpha|\cdot)}$$

The algorithm for this corrective step is therefore as follows.

1. Generate a sample from the state, α given the data and all other model parameters.
2. Evaluate the acceptance probability A and compare with a random uniform.
3. If A is greater than the uniform accept the proposed value for α.
4. If A is less than the uniform reject the proposed value for α and keep the current value.

This Metropolis-Hastings step is included in the ISV class with the option of being used or not. Pitt and Shephard (1997) have noted the problems with convergence of this type of algorithm in the stochastic volatility context. The solution they propose and that which is adopted here is to use blocking of the data series and perform the sampling from the state in blocks which are less than the size of the whole data set. Currently the **sather** class ISV does this using equal block sizes.

4 The Multivariate Factor Stochastic Volatility Model (FSV)

In the section 3 we have considered a univariate Stochastic Volatility model which we designated as $ISV_n(\phi;\sigma_\eta;\mu)$. However in most real world applications we really require a multivariate model allowing us to consider dependencies between series. Harvey, Ruiz, and Shephard (1994) demonstrate the limitations of a direct multivariate extension of this model in a Quasi Maximum Likelihood context. These direct multivariate extensions become unwieldy with high dimensionality resulting in computational problems, see Aguilar and West (2000) for further discussion on this. One approach to overcome this is to consider a latent factor model where the multivariate nature of the problem is expressed in a limited number of factors. This was proposed by Harvey, Ruiz, and Shephard (1994) and applications have been described by Pitt and Shephard (1999) and Aguilar and West (2000). We will adopt the approach taken by Pitt and Shephard (1999), which is the following Factor Stochastic Volatility (FSV) model:

$$
\begin{aligned}
&y_t = \beta f_t + \tau_t, && \text{where} \quad t = 1, \ldots, n \\
&\tau_j \sim ISV_n(\phi^{\tau_j}; \sigma_\eta^{\tau_j}; \mu^{\tau_j}) && \text{where} \quad j = 1, \ldots, N \\
&f_i \sim ISV_n(\phi^{f_i}; \sigma_\eta^{f_i}; 0) && \text{where} \quad i = 1, \ldots, K \quad \text{and} \quad K < N.
\end{aligned}
$$

In this model N represents the number of individual series and K represents the number of factors, always less than N. β is an $N \times K$ matrix of factor loadings, whilst f_t is a $K \times 1$ vector of the unobserved factors at time t. Note that the mean of the log-volatility process for the unobserved factors (the f_t's) is constrained to be zero, for reasons of identifiability. The advantage of this model is that the multivariate nature of the model is captured by a limited number of unobserved latent factors, whilst the within series volatility is captured by the idiosyncratic error terms τ_j. The second advantage of this model is that it is constructed from a number of univariate ISV_n series of which we can can already make inference as described in section 3. The only new elements introduced to the model are f_t and β.

We have written and developed a new sather class FSV to group the routines necessary to make inference on these Factor Stochastic Volatility models. This class draws upon ISV to perform updates on the parameters of the individual ISV_n's. Only two new updates are required to make full inference on this model. The algorithm for inference on this model is as follows.

1. Sample from the states α^{f_i} for $i = 1, \cdots, K$, using the ISV class; see section 3 for details.

2. Sample from θ^{f_i} for $i = 1, \cdots, K$ where $\theta^{f_i} = \left(\phi^{f_i}; {\sigma_\eta}^{f_i}\right)$, using the ISV class; see section 3 for details.

3. Sample from the states α^{τ_j} for $j = 1, \cdots, N$, using the ISV class; see section 3 for details.

4. Sample from θ^{τ_j} for $j = 1, \cdots, N$ where $\theta^{\tau_j} = (\phi^{\tau_j}; {\sigma_\eta}^{\tau_j}; \mu^{\tau_j})$, using the ISV class; see section 3 for details.

5. Sample from the conditional posterior for f; see section 4.1 for details.

6. Sample from the conditional posterior for β; see section 4.2 for details.

7. repeat steps 1 to 6 until convergence.

4.1 Update of f_t

The f_t's are conditionally independent, so for each time point we can sample f_t from its full conditional.

$$\pi(f_t|\cdot) \propto N\left(f_t; \beta' T y_t, F + \beta' T \beta\right)$$

where

$$F = diag\left(\exp\left\{-\alpha_t^f\right\}\right) \qquad T = diag\left(\exp\left\{-\alpha_t^{\tau_j}\right\}\right).$$

Here $\mathcal{N}$ denotes the canonical parameterisation of the normal density. The canonical parameterisation is useful as it allows for easy updating of the parameters. $\mathcal{N}(x;h,K)$ denotes the density of a normal with mean $K^{-1}h$ and variance K^{-1}. The advantage of this parameterisation lies in the following property

$$\mathcal{N}(x;h_1,K_1)\,\mathcal{N}(x;h_2,K_2) = \mathcal{N}(x;h_1+h_2,K_1+K_2).$$

For further details on this parameterisation see Wilkinson and Yeung (2002). The FSV class performs this update.

4.2 Update of β

In the K factor model, the elements of β are constrained as follows $\beta_{ii} = 1, i = 1,\ldots,K$ and $\beta_{ij} = 0, \text{for} j > i$, i.e. unit diagonal lower triangular. These constraints are made to ensure identifiability of the model; for further details see Pitt and Shephard (1999) and Kim, Shephard, and Chib (1998). This means that the updating algorithm is concerned only with updating the non-fixed elements.

These restrictions and the block diagonal nature of the covariance matrix allows us to consider only those elements which need to be updated. Writing $\beta^{(i)}$ as the updateable elements of of the i^{th} column of β and $\beta^{(i\backslash)}$ as all other columns of the matrix β. Then placing a normal prior on each column $\beta^{(i)} \sim \mathcal{N}(h,k)$ then $\beta^{(i)}|\beta^{(i\backslash)},y_t,\alpha^\tau,f$ is also normal. Considering one particular column $\beta^{(m)}$ it has a conditional posterior of the form

$$\beta^{(m)}|\cdot \sim \mathcal{N}\left(H^{(m)}+h, K^{(m)}+k\right)$$

where

$$H^{(m)} = \sum_{t=1}^{n}\left[f_{mt}T_t\left(y_t - \sum_{i\neq m} f_{it}\beta^{(i\backslash)}\right)\right] \qquad K^{(m)} = \sum_{t=1}^{n} f_{mt}^2 T_t$$

where $T_t = \text{diag}(\exp\{-\alpha^{\tau_j}\})$ for $j = i,\ldots,N$. Therefore to update the whole matrix one simply iterates through all the columns for $i = 1,\cdots,K$. The FSV class performs this update.

5 The Composite DLM/FSV

The techniques discussed in sections 3 and 4 for making inference on ISV and FSV models have been used by amongst others Aguilar and West (2000), Pitt and Shephard (1997, 1999) and Kim, Shephard, and Chib (1998) to directly model the returns distributions of financial time series; an approach they hold in common with the ARCH/GARCH literature. Simpson and Wilkinson (2000) demonstrate an alternative approach based on directly modelling the natural logarithm of the share price, with a time evolving mean. This can now be extended to include both a time evolving mean and stochastic volatility. For simplicity let us consider the following model

$$\begin{aligned} Y_t &= \theta_t + \nu_t & \nu_t &\sim FSV(\phi;\sigma_\eta;\mu) \\ \theta_t &= \theta_{t-1} + \omega_t & \omega_t &\sim N(0,W) \end{aligned}$$

so that,

$$\nu_t = \beta f_t + \tau_t$$

$$\tau^j \sim ISV_n\left(\phi^{\tau^j}; \sigma_\eta{}^{\tau^j}; \mu^{\tau^j}\right) \qquad f^j \sim ISV_n\left(\phi^{f^i}; \sigma_\eta{}^{f^i}\right).$$

Initially this looks like a complicated model, but we already have all the necessary routines included in the KALMAN and FSV classes in order to be able to make inference on this model. We can sample θ_t, the state of the DLM, conditioned on all other parameters and the data using the KALMAN class. As in section 2 we can update the W matrix using a semi-conjugate inverse Wishart update. Then taking the error terms ν_t as the data for a FSV object, we can sample the state of this object, conditioned on all other parameters using the FSV class, as well as performing all the other updates as covered in section 4. We can then produce a time evolving covariance matrix V_t, which with transformation of the data structure allows us to re-sample from θ_t. This is repeated until convergence. The algorithm for this is as follows.

1. Sample from the state θ_t; see section 2 for details.
2. Update W using semi-conjugate inverse Wishart prior; see section 2 for details.
3. Place the ν's into a FSV object.
4. Update all parameters associated with this FSV object as described in section 4.
5. Set $V_t = diag(\exp\{\alpha_t{}^{\tau_j}\})$.
6. Reset data for DLM as $y_t - \beta f_t$ to conform with new V_t matrices.
7. Repeat steps 1 to 6 until convergence.

Forecasting future states and observables is a natural extension to this, allowing portfolio selection to be considered.

6 Portfolio Selection

Traditional portfolio management, selection and analysis as proposed by Markowitz (1959) and discussed many times since has as a cornerstone the need to provide estimates of the mean and covariance matrix of the returns distribution. The underlying theory is based on a very simple model for price evolution, and uncertainty about the first and second order moments can lead to a sub-optimal portfolio. Traditional approaches have relied on historic point estimates of the mean vector and covariance matrix, ignoring uncertainty. More recently the Bayes-Stein method has been shown to provide better point estimates than conventional sample estimators; Jorion (1986) and Broad and Sutcliffe (1994). Polson and Tew (2000) have demonstrated a multivariate MCMC approach to portfolio selection based on investigation of the returns distribution, while Aguilar and West (2000) consider an approach based on Multivariate Factor Stochastic Volatility models. Their portfolio selection criteria, however, is based on achieving a set return for the minimum variance in the classical Markowitz style. In Simpson and Wilkinson (2000) a fully Bayesian decision theoretic approach to portfolio selection is adopted and this is now extended to a Stochastic Volatility context. The MCMC approach outlined above allows for more realistic modelling of price evolution, and simulation from the full joint posterior probability distribution for the future returns bypasses the need for calculating point estimates by directly integrating over all uncertainties in the model.

As an example we may assume an investor has an exponential utility for money of the form $u(x) = 1 - \exp^{-\lambda x}$, where λ is a measure of the investor's risk preference. This is only one of many potential forms of the investors risk function and is chosen because of its' relationship to the traditional mean/variance approach. If our investor has a portfolio designated by a vector of weights κ and an amount of money M to invest, in shares with returns R_t then it is possible to consider investment strategies. The traditional mean/variance approach would calculate point estimates of the second order moments of the returns information and maximise the portfolio based solely on this. The MCMC approach to modelling allows us to adopt other approaches.

The MCMC scheme gives us a collection of non-independent simulated returns vectors $\left\{R^{(i)}\right\}$, $i \leq i \leq N$. If we assume a particular investment strategy κ it is trivial to compute the the utility for any one particular returns vector $\left\{R^{(i)}\right\}$, i.e.

$$u\left(R^{(i)}\right) = 1 - \exp^{-M\lambda\kappa' R^{(i)}}$$

From utility theory, we know that the utility of a gamble is the expected utility of that gamble, $u(G) = E\left(u(G)\right)$. The expected utility for the investment strategy κ is the expected utility of that strategy. Hence the utility of a particular investment strategy κ is the sample mean of the utility for a large collection of returns vectors. This can be maximised with respect to κ in order to find the utility optimising portfolio, using standard numerical optimisation techniques.

7 Example

To help clarify the inference and portfolio selection process the following simple example of the multivariate factor stochastic model is included. If we consider trying to select a portfolio from four shares, where we are trying to maximise our investors utility for a given level of risk aversion, as discussed in section 6.

We simulated a single realisation consisting of 2000 time points for four series of one step returns from a known multivariate factor stochastic volatility model. We then ran our block sampler as described in section 4 on this data. Vague prior specifications were placed on the μ and β parameters, while more informative priors were placed on the ϕ's and σ_η's, as these are hard to identify from the data alone.

Figure 1 shows the last thousand simulated values from the returns series and the 5,50 and 95 percent predictive intervals generated using the FSV class, based on a two factor model. This shows the models' ability to establish both the local and global volatility of the data.

Figure 2 shows 1 and 50 step ahead portfolios for varying risk aversion parameters, based on the forecast returns values calculated by the FSV class.The left hand plot shows the 1 step ahead portfolio, this shows little variation after $\lambda = 1$. This might suggest that in the short term the investors' risk preference has little effect on their choice of portfolio, except for those individuals who are particularly risk neutral. The right hand plot tells a different story. In the longer term the greater volatility expressed by some shares leads to a rebalancing of the portfolio for all investors. Whilst the portfolios are superficially similar for risk neutral individuals, the more risk averse investors tend to abandon the more volatile shares and adopt a more balanced portfolio.

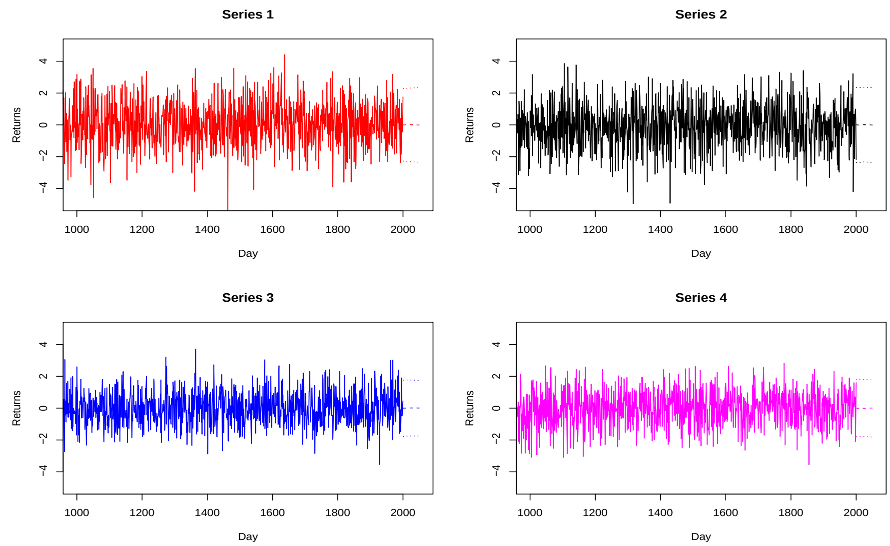

Figure 1: Observed simulated returns and pointwise predictive intervals for four series.

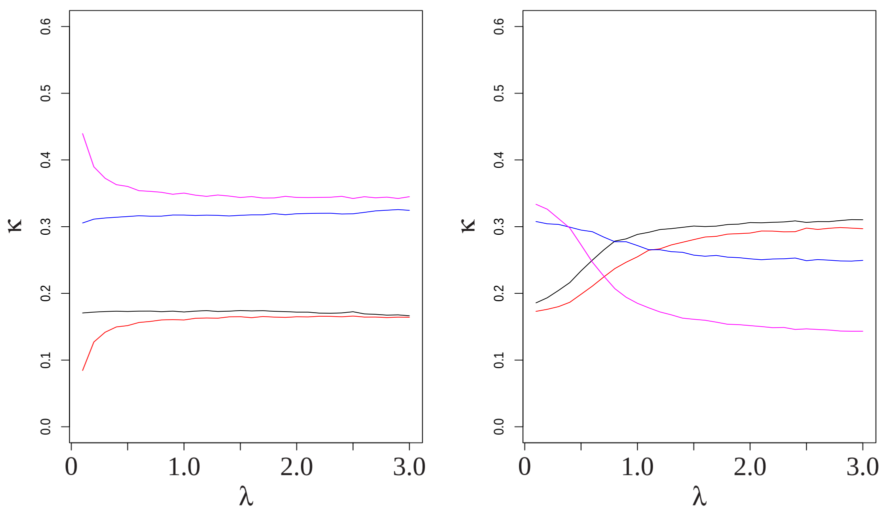

Figure 2: Portfolio weights against risk aversion for one and 50 step ahead returns.

8 Conclusions

In this paper we have outlined the need for multivariate modelling of financial time series. We have presented a number of models that are applicable to this task and written modular software libraries for making Bayesian inference on these models. We have also provided a simple model to illustrate the use of these models in the context of portfolio selection.

References

Aguilar, O., G. Huerta, R. Prado, and M. West (1999). Bayesian inference on latent structure time series. In J. M. B. et al (Ed.), *Bayesian Statistics 6*, Oxford, pp. 3–26. Oxford University Press.

Aguilar, O. and M. West (2000). Bayesian dynamic factor models and portfolio selection. *Journal of Business and Economic Statistics 18*, 338–357.

Barndorff-Nielsen, O. and N. Shephard (2001). Non-Gaussian OU based models and some of their uses in financial economics. *Journal of the Royal Statistical Society, Series B 63*, 167–241.

Bollersev, T. (1987). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics 51*, 307–327.

Broad, J. and C. Sutcliffe (1994). Estimation methods in portfolio selection and effectiveness of short sales restrictions:U.K. evidence. *Management Science 40*, 516–534.

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

Harvey, A., E. Ruiz, and N. Shephard (1994). Multivariate stochastic variance models. *The Review of Economic Studies 61*(2), 247–264.

Jacquier, E., N. Polson, and P. Rossi (1994). Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics 12*(4), 371–417.

Jacquier, E., N. Polson, and P. Rossi (1999). Stochastic volatility: univariate and multivariate extensions. *Unpublished manuscript*.

Jorion, P. (1986). Bayes-Stein estimation for portfolio analysis. *Journal of Financial and Quantitative Analysis 21*, 279–292.

Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility:likelihood inference and comparison with ARCH models. *Review of Economic Studies 65*, 361–393.

Koopman, S. (1993). Disturbance smoother for state space models. *Biometrika 80*, 117–126.

Markowitz, H. (1959). *Portfolio Selection: efficient diversification of investment* (second ed.). Oxford: Blackwell Publishers Ltd.

Pitt, M. and N. Shephard (1997). Likelihood analysis of non-Gaussian measurement time series. *Biometrika 84*(3), 653–667.

Pitt, M. and N. Shephard (1999). Time varying covariances: a factor stochastic volatility approach, (with discussions). In J. M. B. et al (Ed.), *Bayesian Statistics 6*, Oxford, pp. 3–26. Oxford University Press.

Polson, N. and B. Tew (2000). Bayesian portfolio selection: An empirical analysis of the S & P Index 1970-1996. *Journal of Business and Economic Statistics 18*(2), 164–173.

Simpson, A. and D. J. Wilkinson (2000). Multivariate DLMs for financial time series, with application to the management of portfolios. In J. G. Bethlehem and P. G. M. van der Heijden (Eds.), *Proceedings in Computational Statistics 2000*, Heidelberg, pp. 457–462. Physica-Verlag.

West, M. and J. Harrison (1997). *Bayesian forecasting and dynamic models* (second ed.). New York: Springer.

Wilkinson, D. J. and S. K. H. Yeung (2002). Conditional simulation from highly structured Gaussian systems, with application to blocking-MCMC for the Bayesian analysis of very large linear models. *Statistics and Computing 12*(x), xx–xx. To appear.