

# McMC Estimation of Multiscale Stochastic Volatility Models

**German Molina**

Statistical and Applied Mathematical Sciences Institute. NC 27709. USA  
(german@alumni.duke.edu)

**Chuan-Hsiang Han**

Department of Quantitative Finance, National Tsing-Hua University,  
Hsinchu, Taiwan 30013, ROC. (chhan@mx.nthu.edu.tw)

**Jean-Pierre Fouque**

Department of Statistics and Applied Probability, University of California, Santa Barbara,  
CA 93106-3110, USA. (fouque@pstat.ucsb.edu)

In this paper we propose to use Monte Carlo Markov Chain methods to estimate the parameters of Stochastic Volatility Models with several factors varying at different time scales. The originality of our approach, in contrast with classical factor models is the identification of two factors driving univariate series at well-separated time scales. This is tested with simulated data as well as foreign exchange data.

KEY WORDS: Time scales in volatility, Bayesian estimation, Markov Chain Monte Carlo, Foreign exchange volatility model. Multifactor model.

## 1. INTRODUCTION

Volatility modeling and estimation from historical data has been extensively studied during the last two decades. Constant volatility models have been shown to provide a poor fit of financial time series, while dynamic structures provide a more realistic approach to volatility modeling.

One of the papers that first tackled this problem was the original ARCH work by Engle (1982), followed by its generalization (GARCH) by Bollerslev (1986). In the last decade the GARCH family has expanded enormously, with papers of great relevance, like the proposed EGARCH by Nelson (1991), and other adaptations/generalizations, like threshold, power, (fractionally) integrated as well as multivariate versions.

A second approach to modeling time-varying volatility comes from the stochastic volatility models. GARCH and stochastic volatility models differ in the structure of the volatility process. The former assumes dependence on the previous realized shocks in the series, while the latter assumes that it follows its own stochastic process, with potential dependence coming through the correlation structure between the series and the volatility processes.

Estimation of stochastic volatility models has been an important issue in the literature. Bates (1996) gives a review of the different approaches to fitting these models.

Generalizations of these models have been explored, including jumps in returns and in volatilities and/or fat tails in distributions, as in Jacquier *et al.*, (1998), Chib *et al.*, (2002) or ter Horst (2003), or non-gaussian errors, as in Fridman and Harris (1998).

Other extensions include multivariate stochastic volatility models, as proposed by Harvey et al (1994). Multifactor models for multivariate series have been developed in the last decade, where several return series are being driven by a common (smaller) set of underlying factors, as in Lopes (1999), Aguilar and West (2000) and Chib *et al.*, (2002).

Another approach to multifactor models came with the idea of multiscale modeling. Univariate series can be driven by several factors, with those factors varying at different time scales, as in Alizadeh *et al.*, (2002). In fact, two-factor stochastic volatility models have been shown to produce the kurtosis, fat-tailed return distribution and long memory effect observable in many financial time series. This is well-documented in LeBaron, (2001).

Chernov *et al.*, (2003) used efficient method of moments to calibrate this model. One of the main results they found is that two factors are necessary for log-linear models. In this paper we propose the use of Bayesian methods for the parameter estimation in this class of stochastic volatility models. Bayesian methods have become more popular for estimation of complicated models with the introduction of Markov chain Monte Carlo (McMC) techniques. A review of these methods can be found in Robert and Casella (1998). Apart from its computational advantages, the possibility of incorporating prior knowledge of the parameters in the estimation provides a useful tool to practitioners. Key papers in Bayesian estimation of stochastic volatility models are Jacquier *et al.*, (1994) and Kim *et al.*, (1998) among many others.

Volatility is one of the most sensitive factors in derivative pricing. Black and Scholes introduced a closed form solution to derivative prices under constant volatility. Their work has been extended to incorporating either constant volatility plus jumps or stochastic volatility with/without jumps to provide a better fit of implied volatility surfaces. A recent alternative to these approaches is to consider volatility driven by short and long time scale factors. Under this framework, Fouque *et al.*, (2003) use a combination of singular and regular perturbations to derive a closed form formula to approximate European option prices. They use the term structure of implied volatilities to calibrate some aggregate parameters in their formulas and they find two factor with well-separated time scales provide significantly better fit in comparison with one factor model.

As a consequence of two factor stochastic volatility models, Monte Carlo simulations to evaluate option prices can be computed efficiently, key issue from a practitioner's point of view. Fouque and Han, (2004) propose the use of variance reduction techniques such as importance sampling and control variate methods to evaluate no-arbitrage prices of derivatives.

In this paper we consider a class of stochastic volatility models where the volatility is the exponential of a sum of two mean reverting diffusion processes (more precisely Ornstein-Uhlenbeck

processes). These mean reverting processes vary on well-separated time scales. We then rewrite these models in discrete time and use McMC techniques to estimate their parameters (Section 2). Our study is illustrated on simulated data (Section 3) where we show that we can effectively estimate several factors at well-separated time scales. We also present (Section 4) an application of our technique to a set of foreign exchange data.

This paper focuses on the development of an McMC estimation algorithm for single time series driven by several volatility factors at different time scales. We informally address model selection issues through the diagnostics of the McMC results, but do not enter in proper computation of Bayes Factors or reversible jump algorithms. We instead restrict ourselves to propose a Bayesian alternative to the estimation problems when more than one stochastic volatility factor drive univariate series, and address identifiability problems inherent to such estimation. Our interest resides in applications in the Foreign Exchange market, this being one of the markets more intensively addressed in the statistical literature. The importance of the identification of a second factor, and its relevance for a more accurate estimation of the first factor, is outlined in the paper.

## 2. MULTISCALE MODELING AND McMC ESTIMATION

### 2.1. Continuous Time Model

For purposes of general notation, let  $K$  be the number of volatility factors driving a univariate time series, where  $K$  can be either one or two. We first introduce our  $K$  factor volatility model in continuous time. In this class of models the volatility is simply the exponential of the sum of  $K$  mean reverting Ornstein-Uhlenbeck (O-U) processes:

$$\begin{aligned} dS_t &= \kappa S_t dt + \sigma_t S_t dW_t^{(0)} \\ \log(\sigma_t^2) &= \sum_{j=1}^K F_t^{(j)} \\ dF_t^{(j)} &= \alpha_j (\mu_j - F_t^{(j)}) dt + \beta_j dW_t^{(j)}, \quad j = 1, \dots, K, \end{aligned} \tag{2.1}$$

where  $S$  is the asset price,  $\kappa$  is the constant rate of return and  $W^{(*)}$  are (possibly correlated) Brownian motions. Each of the factors  $F$  is a mean reverting O-U process, where  $\mu_j$  denotes the long-run mean,  $\alpha_j$  is the rate of mean reversion, and  $\beta_j$  is the “volatility of volatility” of the  $F^{(j)}$  factor. The typical time scale of this factor is given by  $1/\alpha_j$ . In our model these time scales are distant in magnitude, which we define as well-separated. They are also ordered, so that, in the two factor case,  $0 < \alpha_1 < \alpha_2$ , and, therefore, the second factor corresponds to the shortest time scale and the first factor corresponds to the longest time scale. We also assume that the variances of the long-run distributions of these factors, given by  $\beta_j^2/2\alpha_j$ , are of the same order. In particular this implies that the  $\beta_j$ ’s are also correspondingly ordered  $\beta_1 < \beta_2 < \dots < \beta_K$ .

## 2.2. Discretization of the Model

The goal of this section is to present a discretized version of the stochastic volatility model given in (2.1). Given a fixed time increment  $\Delta > 0$ , we use the following Euler discretization of the spot price process  $(S_t)$  and volatilities  $(F_t^{(j)})$ ,  $j = 1, \dots, K$  at the times  $t_k = k\Delta$ :

$$\begin{aligned} S_{t_{k+1}} - S_{t_k} &= \kappa S_{t_k} \Delta + \sigma_{t_k} S_{t_k} \sqrt{\Delta} \epsilon_{t_k} \\ \log(\sigma_{t_k}^2) &= \sum_{j=1}^K F_{t_k}^{(j)} \\ F_{t_k}^{(j)} - F_{t_{k-1}}^{(j)} &= \alpha_j (\mu_j - F_{t_{k-1}}^{(j)}) \Delta + \beta_j \sqrt{\Delta} v_{t_k}^{(j)}, \end{aligned}$$

where  $\epsilon_{t_k}$  and  $v_{t_k}^{(j)}$  are (possibly correlated) sequences of i.i.d. standard normal random variables. In this paper we restrict ourselves to the case where the sequence of  $\epsilon$ 's is independent of the  $v^{(j)}$ 's, which, themselves can be correlated. This is usually the case in the Foreign Exchange market, where market participants operate at different time scales (intraday market makers, hedge funds or pension funds operate with quite different schedules, and based on different information sets), but not independently, and where the leverage effect is more rare than in equity markets.

We regroup terms within the equations above such that

$$\begin{aligned} \frac{1}{\sqrt{\Delta}} \left( \frac{S_{t_{k+1}} - S_{t_k}}{S_{t_k}} - \kappa \Delta \right) &= \sigma_{t_k} \epsilon_{t_k} \\ \log(\sigma_{t_k}^2) &= \sum_{j=1}^K F_{t_k}^{(j)} \\ F_{t_k}^{(j)} - \mu_j &= (1 - \alpha_j \Delta) (F_{t_{k-1}}^{(j)} - \mu_j) + \beta_j \sqrt{\Delta} v_{t_k}^{(j)}, \end{aligned}$$

is obtained.

To study the hidden volatility process, we introduce the returns

$$y_{t_k} = \frac{1}{\sqrt{\Delta}} \left( \frac{S_{t_{k+1}} - S_{t_k}}{S_{t_k}} - \kappa \Delta \right),$$

and the discrete driving (vector of) volatilities

$$\mathbf{h}_{t_k} = \mathbf{F}_{t_k} - \boldsymbol{\mu},$$

where  $\mathbf{F}_{t_k}$  (resp  $\boldsymbol{\mu}$ ) denotes the vector  $(F_{t_k}^{(1)}, \dots, F_{t_k}^{(K)})'$  (resp  $(\mu_1, \dots, \mu_K)'$ ). The autoregressive parameters will be denoted by

$$\phi_j = 1 - \alpha_j \Delta, \quad j = 1, \dots, K,$$

and the standard deviation parameters will be defined by

$$\sigma_v^{(j)} = \beta_j \sqrt{\Delta}, \quad j = 1, \dots, K.$$

With an abuse of notation, and since the observations are equally spaced, we will use  $t$  instead of  $t_k/\Delta$  for the discrete time indices, so that  $t \in \{0, 1, 2, \dots, T\}$ . To summarize, our model is a  $K$ -dimensional AR(1) for the log-volatilities, and can be written in state space form as:

$$y_t = e^{(\mathbf{1}'\mathbf{h}_t + \mathbf{1}'\boldsymbol{\mu})/2} \epsilon_t \quad (2.2)$$

$$\mathbf{h}_t = \boldsymbol{\Phi}\mathbf{h}_{t-1} + \boldsymbol{\Sigma}^{1/2}\mathbf{v}_t \quad \text{for } t = 1, \dots, T, \quad (2.3)$$

where  $\mathbf{v}_t$  is the vector  $(v_t^{(1)}, \dots, v_t^{(K)})$  of standard random Gaussian variates,  $\mathbf{1}$  is a ( $K$ -dimensional) vector of ones, the  $K$ -dimensional autoregressive (diagonal) matrix with elements  $\phi_j$  is denoted by  $\boldsymbol{\Phi}$  and the covariance matrix is denoted by  $\boldsymbol{\Sigma}$ . In the following sections we assume that the covariance matrix  $\boldsymbol{\Sigma}$  is diagonal, since the correlation parameter between the factors is indeed not identifiable. For a proof of equivalence between the diagonal and non-diagonal case, see Appendix A.

In the next section we make precise the parametrization of our discrete model, including the initial state and random variable notation.

### 2.3. Discrete time model specification

The simplest one-factor stochastic volatility model can be defined in state-space form as in equation (2.4). We follow the notation in Kim *et al.* (1998), and denote  $h_t$  as the log-volatility at time  $t$ , which follows an AR(1) process reverting to its long-term mean  $\mu$  according to the mean-reversion parameter  $\phi$  and with a volatility of volatility  $\sigma_v$ . The initial state  $h_0$  is assumed to come from its invariant (implied marginal) distribution so that:

$$\begin{aligned} y_t &= e^{(h_t + \mu)/2} \epsilon_t & (2.4) \\ h_t &= \phi h_{t-1} + \sigma_v v_t \\ h_0 &\sim N_1(h_0 \mid 0, \sigma_v^2 / (1 - \phi^2)) \\ \begin{bmatrix} \epsilon_t \\ v_t \end{bmatrix} &\sim N_2\left(\begin{bmatrix} \epsilon_t \\ v_t \end{bmatrix} \mid 0, \mathbf{I}_2\right) \\ (\mu, \phi, \sigma_v) &\sim \pi(\mu, \phi, \sigma_v). \end{aligned}$$

To complete the joint distribution a prior density  $\pi(\mu, \phi, \sigma_v)$  on the remaining parameters must also be specified; this will be more explicitly done in the next section.

For simplicity, the mean return is assumed equal to zero. A trivial modification of the procedure presented here would allow us to consider a non-zero mean in the returns.

Our contribution in this paper does not focus on modeling extreme events (fat tails in the distributions and/or jumps), which is still a possible addition to the model here presented, but instead generalizes the assumption that a single volatility process drives the series.

We will assume that the volatility processes are independent ( $\Sigma$  in (2.5) being diagonal). We show in appendix A that, when the factors are positively correlated, this is equivalent to having independent factors both for estimation and for prediction. Indeed the covariance is not identifiable, since the model is overparametrized.

The one-factor model can be extended to the general case as

$$\begin{aligned}
 y_t &= e^{(\mathbf{1}'\mathbf{h}_t + \mu)/2} \epsilon_t & (2.5) \\
 \mathbf{h}_t &= \Phi \mathbf{h}_{t-1} + \Sigma^{1/2} \mathbf{v}_t \\
 \mathbf{h}_0 &\sim N_K(\mathbf{h}_0 | \mathbf{0}, \Omega) \\
 \begin{bmatrix} \epsilon_t \\ \mathbf{v}_t \end{bmatrix} &\sim N_{K+1} \left( \begin{bmatrix} \epsilon_t \\ \mathbf{v}_t \end{bmatrix} \middle| \mathbf{0}, \mathbf{I}_{K+1} \right) \\
 (\boldsymbol{\mu}, \Sigma, \Phi) &\sim \pi(\boldsymbol{\mu}, \Sigma, \Phi),
 \end{aligned}$$

where

- $\Sigma$  is the ( $K \times K$ ) diagonal matrix of conditional variances  $\{\sigma_k^2\}_{k=1}^K$  of the log-volatility processes;
- $\Omega = \Phi \Omega \Phi + \Sigma$  is the implied (stationary) marginal covariance matrix for the initial state log-volatility vector  $\mathbf{h}_0$ ;
- $\Phi$  is a  $K \times K$  diagonal matrix with mean reversion parameters

Since  $\Phi$  and  $\Sigma$  are diagonal matrices,  $\Omega$  is a diagonal matrix with entries  $\omega_i = \sigma_i^2 / (1 - \phi_i^2)$ .

Notice that we do not have a vector  $\boldsymbol{\mu}$  in the observation equation (2.5), but instead a single parameter  $\mu$ . This is because only the linear combination (sum) is identifiable. Therefore we define the mean log-volatility of the series  $\mu = \mathbf{1}'\boldsymbol{\mu}$ .

One interesting feature of this model is its similarity to models that introduce jumps in volatility. If we set  $\phi_i = 0$  for some  $i$ , the implied model is equivalent to one that introduces a permanent (log-normal) source of independent jumps in the volatility series, as the  $h_{i,t}$  would be (temporally independent) normals added to the (log)variance process driving the series.

## 2.4. Prior specification

We specify independent priors for each of the (sets of) parameters, which is the product measure  $\pi(\mu, \Phi, \Sigma) = \pi_1(\Phi)\pi_2(\mu)\pi_3(\Sigma)$ . We set a uniform prior between -1 and 1 for each of the  $K$  mean reverting parameters. In order to guarantee identifiability of the different series, we specify a prior ordering of the parameters, namely  $\pi_1(\Phi) \propto \mathbb{1}_{1 > \phi_1 > \phi_2 > -1}$ .

Usually, with financial data, there is substantial prior information about  $\mu$  and  $\Sigma$ . We could, in principle, use an informative normal prior for  $\mu$ ,  $\pi_2(\mu) = N_1(\mu|m_0, v_0)$  with fixed hyperparameters  $m_0$  and  $v_0$ . Here, however, a vague prior is used by choosing a large value of  $v_0$ . A similar prior is chosen for the variance hyperparameters:  $\pi_3(\Sigma) \propto \prod_{k=1}^K IG(\sigma_k^2|a_k, b_k)$  for  $k = 1 \dots K$  with fixed  $a_k, b_k$ . We specify values for these hyperparameters such that the prior is flat. The result do not seem to be sensitive to which large values of  $a_k$  and  $b_k$  are chosen.

## 2.5. Estimation

Since the volatility parameters  $\mathbf{h}_t$  are nonlinear in the likelihood, a mixture approximation to the likelihood as detailed in Kim *et al.* (1998), has become a commonly used linear approximation technique. Taking the logarithm of the squares of the mean corrected returns,  $y_t$ , transforms (2.5) into

$$\log(y_t^2) = \mathbf{1}'\mathbf{h}_t + \mu + \underbrace{\log(\epsilon_t^2)}_{\eta_t}, \quad (2.6)$$

and  $\eta_t \sim \log -\chi_1^2$  is then well-approximated as a mixture of seven normals with (known) means  $m_i$ , variances  $v_i$  and weights  $q_i$ . We introduce vectors of indicator variables  $\boldsymbol{\delta}_t = \{\delta_{1t}, \dots, \delta_{7t}\}$ , for the particular component of the mixture at each time point  $t$ , in the same fashion as Kim *et al.* (1998), so that, consequently, we have

$$[\eta_t | \delta_{it} = 1] \sim N_1(\eta_t | m_i - 1.2704, v_i).$$

For simplicity of notation, let  $\mathbf{y}^* = \{\log(y_1^2), \dots, \log(y_T^2)\}$  and  $\mathbf{h} = \{\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_T\}$ .

The new state-space representation can be written as

$$\begin{aligned} [y_t^* | \{\delta_{it}\}_{i=1}^7, \mathbf{h}_t, \mu] &\sim N_1 \left( y_t^* \left| \mathbf{1}'\mathbf{h}_t + \mu - 1.2704 + \sum_{i=1}^7 \delta_{it}m_i, \sum_{i=1}^7 \delta_{it}v_i \right. \right) \\ [\mathbf{h}_t | \mathbf{h}_{t-1}, \Phi, \Sigma] &\sim N_K(\mathbf{h}_t | \Phi\mathbf{h}_{t-1}, \Sigma) \\ [\mathbf{h}_0 | \Phi, \Sigma] &\sim N_K(\mathbf{h}_0 | \mathbf{0}, \Omega) \\ [\mu, \Phi, \Sigma] &\sim \pi(\mu, \Sigma, \Phi) \\ [\boldsymbol{\delta}] &\sim \prod_{t=1}^T MN(\boldsymbol{\delta}_t | 1, q_1, \dots, q_7). \end{aligned} \quad (2.7)$$

A Gibbs sampling scheme can now be employed, and consists of the following steps

1. Initialize  $\boldsymbol{\delta}, \mu, \boldsymbol{\Phi}, \boldsymbol{\Sigma}$
2. Sample  $\mathbf{h} \sim [\mathbf{h}|\mathbf{y}^*, \mu, \boldsymbol{\Sigma}, \boldsymbol{\delta}, \boldsymbol{\Phi}]$
3. Sample  $\boldsymbol{\delta} \sim [\boldsymbol{\delta}|\mathbf{y}^*, \mu, \mathbf{h}]$
4. Sample  $\mu \sim [\mu|\mathbf{y}^*, \mathbf{h}, \boldsymbol{\delta}]$
5. Sample  $\boldsymbol{\Phi} \sim [\boldsymbol{\Phi}|\boldsymbol{\Sigma}, \mathbf{h}]$
6. Sample  $\boldsymbol{\Sigma} \sim [\boldsymbol{\Sigma}|\mathbf{h}, \boldsymbol{\Phi}]$

Details of the full conditionals are given in the Appendix B. In principle one could obtain improved sampling versions by integrating out some parameters from the full conditionals, as in Kim *et al.* (1998). The simpler version given here is efficient enough for our purposes, however.

### 3. SIMULATION STUDY

In this section we study the behavior of the algorithm when the estimated model is the one generating the data, and also when we try to estimate using the wrong model (wrong number of factors). This will be investigated for different sample sizes and parameter settings.

For all experiments, we set the true value of  $\mu = -1.5$ , which is a value similar to the posterior mean of  $\mu$  we obtain using real data. We will denote by  $K^*$  the true number of volatility factors that generated the data, while  $K$  is the value with which we do the estimation.

- We generate the data under three possible parameter settings:

Setting 1:  $K^*=1, \phi_1 = 0.95, \sigma_1 = 0.26$

Setting 2:  $K^*=2, \phi_1 = 0.98, \sigma_1 = 0.20, \phi_2 = 0.60, \sigma_2 = 0.80$

Setting 3:  $K^*=2, \phi_1 = 0.98, \sigma_1 = 0.20, \phi_2 = 0.30, \sigma_2 = 0.80$

- We consider for estimation both the 1 and 2 factor stochastic volatility models:  $K=1,2$ .
- Two sample sizes are considered:  $T=1500$  or  $T=3000$ .

Any combination of the above forms an experiment. As an example, a possible experiment comprises generating 3000 data points with the first parameter setting (one volatility factor) and estimating it with  $K=2$  factors. This would correspond to the last column of Table 2.

We considered  $K^*=3$  as well, but the algorithm failed to identify three factors for the experiments considered, so we restricted ourselves to the  $K^*=2$  factor case of interest.



The results in Tables 1 and 2 are based on 100 Monte Carlo experiments for each of the settings and for each of the experiments. The results reported are Monte Carlo means of the posterior medians (bold), Monte Carlo means of the posterior standard deviations (in parenthesis to the right). Below those numbers are the standard deviations of the posterior medians and in parenthesis the standard deviation of the posterior standard deviations.

Table 1: Estimates of one factor Stochastic Volatility Model

	True value	Estimation with K=1				Estimation with K=2			
		T=1500		T=3000		T=1500		T=3000	
$\phi_1$	0.95	0.95	(0.02)	0.95	(0.01)	0.96	(0.02)	0.96	(0.01)
		0.02	(0.00)	0.01	(0.00)	0.01	(0.01)	0.01	(0.00)
$\phi_2$	—	—	—	—	—	0.66	(0.29)	0.61	(0.29)
		—	—	—	—	0.36	(0.17)	0.38	(0.16)
$\mu$	-1.5	-1.50	(0.14)	-1.50	(0.10)	-1.54	(0.15)	-1.50	(0.10)
		0.14	(0.03)	0.10	(0.01)	0.14	(0.04)	0.01	(0.02)
$\sigma_1$	0.26	0.27	(0.04)	0.27	(0.03)	0.12	(0.06)	0.15	(0.05)
		0.04	(0.01)	0.03	(0.00)	0.09	(0.03)	0.09	(0.03)
$\sigma_2$	—	—	—	—	—	0.29	(0.05)	0.28	(0.04)
		—	—	—	—	0.04	(0.01)	0.02	(0.01)

Parameter setting 1: Results from 100 Monte Carlo simulations generated with  $K^* = 1$  factor. Columns represent data lengths (T=1500,3000) and factors used for the estimation (K=1,2). Top lines are means of the Monte Carlo posterior medians (means of the Monte Carlo standard deviations). Bottom lines are standard deviations of the posterior medians (standard deviations of the posterior standard deviations).

Three main conclusions can be drawn from these tables. First, we can accurately estimate at least two factors with a relatively small data length if the time scales are well separated. Second, overparametrization tends to induce lack of convergence (high Monte Carlo means of the posterior standard deviations of the redundant parameters). Third, underparametrization tends to induce apparent convergence around values of the parameters situated in between the true values, and either visualization of the chain or more proper model selection procedures are necessary.

Figure 1 shows the posterior distribution for  $\phi_k$ , when the data was generated with one volatility scale ( $K^* = 1$ ). After a burn-in of 50,000 iterations, we record, with a thinning of 25, the following 50,000 iterations. Two typical runs are given. Estimation of  $\phi_1$ , under K=1, is very stable around the true value (0.95). On the other hand, estimation of  $\phi_2$  is very unstable, since the data was generated with only one factor. This graph shows the pattern one should

Table 2: Estimates of two factor Stochastic Volatility Models

	True value	Estimation with K=1				Estimation with K=2			
		T=1500		T=3000		T=1500		T=3000	
$\phi_1$	0.98	0.86	(0.03)	0.87	(0.02)	0.97	(0.02)	0.98	(0.01)
		0.05	(0.01)	0.03	(0.00)	0.02	(0.01)	0.01	(0.00)
$\phi_2$	0.60	—	—	—	—	0.53	(0.14)	0.55	(0.08)
		—	—	—	—	0.15	(0.06)	0.07	(0.02)
$\mu$	-1.5	-1.44	(0.14)	-1.46	(0.10)	-1.52	(0.32)	-1.49	(0.21)
		0.26	(0.03)	0.22	(0.02)	0.24	(0.18)	0.17	(0.07)
$\sigma_1$	0.20	0.69	(0.06)	0.67	(0.05)	0.24	(0.08)	0.22	(0.04)
		0.09	(0.01)	0.06	(0.00)	0.07	(0.03)	0.05	(0.01)
$\sigma_2$	0.80	—	—	—	—	0.78	(0.08)	0.82	(0.05)
		—	—	—	—	0.09	(0.01)	0.05	(0.00)

	True value	Estimation with K=1				Estimation with K=2			
		T=1500		T=3000		T=1500		T=3000	
$\phi_1$	0.98	0.89	(0.03)	0.91	(0.02)	0.97	(0.01)	0.98	(0.01)
		0.05	(0.01)	0.04	(0.01)	0.01	(0.01)	0.01	(0.00)
$\phi_2$	0.30	—	—	—	—	0.23	(0.16)	0.28	(0.10)
		—	—	—	—	0.15	(0.04)	0.09	(0.02)
$\mu$	-1.5	-1.42	(0.16)	-1.41	(0.11)	-1.52	(0.28)	-1.50	(0.21)
		0.26	(0.04)	0.21	(0.02)	0.25	(0.10)	0.18	(0.06)
$\sigma_1$	0.20	0.52	(0.07)	0.51	(0.05)	0.22	(0.05)	0.21	(0.03)
		0.12	(0.01)	0.08	(0.01)	0.04	(0.02)	0.03	(0.01)
$\sigma_2$	0.80	—	—	—	—	0.78	(0.08)	0.80	(0.05)
		—	—	—	—	0.08	(0.01)	0.05	(0.00)

Parameter settings 2 and 3: Results from 100 Monte Carlo simulations generated with  $K^* = 2$  factors. Columns represent data lengths (T=1500,3000) and factors used for the estimation (K=1,2). Top lines are means of the Monte Carlo posterior medians (means of the Monte Carlo standard deviations). Bottom lines are standard deviations of the posterior medians (standard deviations of the posterior standard deviations).

expect when the model is overparametrized. Either the estimates of  $\phi_1$  and  $\phi_2$  are confounded (the volatility series splits into two almost equivalent parts) or  $\phi_2$  never stabilizes.

Figure 2 shows the posterior distribution for  $\phi_k$ ,  $k = 1, \dots, K$  when the data was generated with two factors ( $K^*=2$ ,  $\phi_1 = 0.98$  and  $\phi_2 = 0.6$ ). The model with K=1 is underparametrized, leading to an apparent convergence, but the value to which it converges is between the two true




Figure 1: Two typical runs (horizontal) where the estimation of  $\phi_k$  was done using  $K=1$  (1st column) and  $K=2$  factors (2nd and 3rd columns), but the data was generated with  $K^*=1$  ( $\phi_1 = 0.95$ ), after burn-in and thinning of 1 iteration every 25.

values that generated the data. The convergence is fine for the case when we estimate with  $K=2$ .

#### 4. EMPIRICAL APPLICATION: FX DATA

We study the volatility in a foreign exchange time series. Since most time series could be argued to have a positive trend, we construct the detrended series as in Kim *et al.* (1998). In principle we could easily introduce mean processes (constant or time-dependent), but we refrain from doing so to allow comparisons with previous literature and to focus our efforts mainly on the volatility modeling.

We study volatility in the daily exchange rate between the U.K. Sterling and the U.S. Dollar from 10/09/1986 to 08/09/1996. This data is shown in Figure 3.

We transform the spot prices into detrended returns and perform the analysis by running the algorithm with one and two. The results are shown in Table 3. All results are posterior medians (posterior standard deviations).

From the left plot in Figure 4 one can argue that, for the GBP/USD series, using one factor gives reasonably stable results. There is a (very) heavy left tail in the posterior distribution of  $\phi_1$  under  $K=1$ , though. This could be an indication of the existence of a second volatility




Figure 2: Two typical runs (horizontal) where the estimation of  $\phi_k$  was done using  $K=1$  (1st column) and  $K=2$  (2nd and 3rd columns) factors, but the data was generated with  $K^*=2$  ( $\phi_1 = 0.98, \phi_2 = 0.6$ ), after burn-in and thinning of 1 iteration every 25.




Figure 3: Daily GBP/USD spot (top) and returns (bottom) for the period 10/09/1986-08/09/1996.

Table 3: GBP/USD results: Parameter posterior medians (standard deviations), estimated under different number of factors  $K=1,2$ .

	$\phi_1$		$\phi_2$		$\mu$		$\sigma_1^2$		$\sigma_2^2$	
K=1	0.916	(0.057)	—	(—)	-1.220	(0.106)	0.134	(0.130)	—	(—)
K=2	0.988	(0.006)	0.149	(0.090)	-1.470	(0.255)	0.013	(0.007)	0.947	(0.132)

factor running at a faster scale. If that was the case, the single volatility factor model would be picking up a mixture of both the slow and fast mean reverting processes, achieving an intermediate result in terms of its estimation, and showing a left-tail of the distribution heavier than one would expect.

If we run the estimation with  $K=2$  we can see (Figure 4, right plots) that there seems to be (at least) a second volatility factor running at a very fast mean reverting rate. The algorithm remains highly stable around a low value for the mean reverting parameter, and a high value for its volatility. Furthermore, the stability of the parameter estimates for the first factor also improves. Indeed, the inclusion of this second factor in the model allows us to identify that the first factor was moving at a lower frequency than previously thought (the posterior median of the mean reverting parameter for  $K=2$  for the slower scale is 0.988 under  $K=2$ , but it was 0.916 under  $K=1$ ). This is a key result of the introduction of a second factor: a significant change in our estimate of the first factor, whose mean reversion speed, if such second factor exist, is usually inflated. From a practitioner’s point of view, accurate estimation of volatility persistence is a major factor, especially in derivative pricing.

Figure 6 shows the posterior distribution of the mean log-volatilities for each time point  $t$  and for each  $k$  when  $K=2$ , together with their sum and the absolute value of the returns.

Given the limitations of the algorithm to identify more than two factors for the data lengths considered, we do not claim the lack of a third factor in this dataset, but the impossibility to identify it.

Notice in Figure 4 that the short time scale process captures the extreme events, while the long scale factor captures the trends and is much less influenced by extreme events. There is a peak in the returns between the 1900th and the 2000th observation (that is in August 1994). This peak occurs during a period of relative stability (the absolute value of the returns are clearly more stable than in previous periods). The slow mean reverting volatility process remains low at this time and is not affected by this peak, while the fast mean reverting volatility process exhibits a peak coinciding with that day. We do not address the question of whether this should be considered a jump in volatility or the effect of a second volatility process. However, the tight concentration of the posterior around positive values for the second factor, as shown in Table 3, seem to point towards the existence of a clear second volatility factor at a (much) faster scale.




Figure 4: Traceplots of  $\phi_k, k = 1, \dots, K$  for the GBP/USD series estimated with  $K=1$  and  $K=2$  factors, after burn-in and thinning of 1 iteration every 25.




Figure 5: Histograms of  $\sigma_1, \sigma_2$  and  $\mu$  for the GBP/USD series estimated with  $K=2$  factors, after burn-in and thinning of 1 iteration every 25.

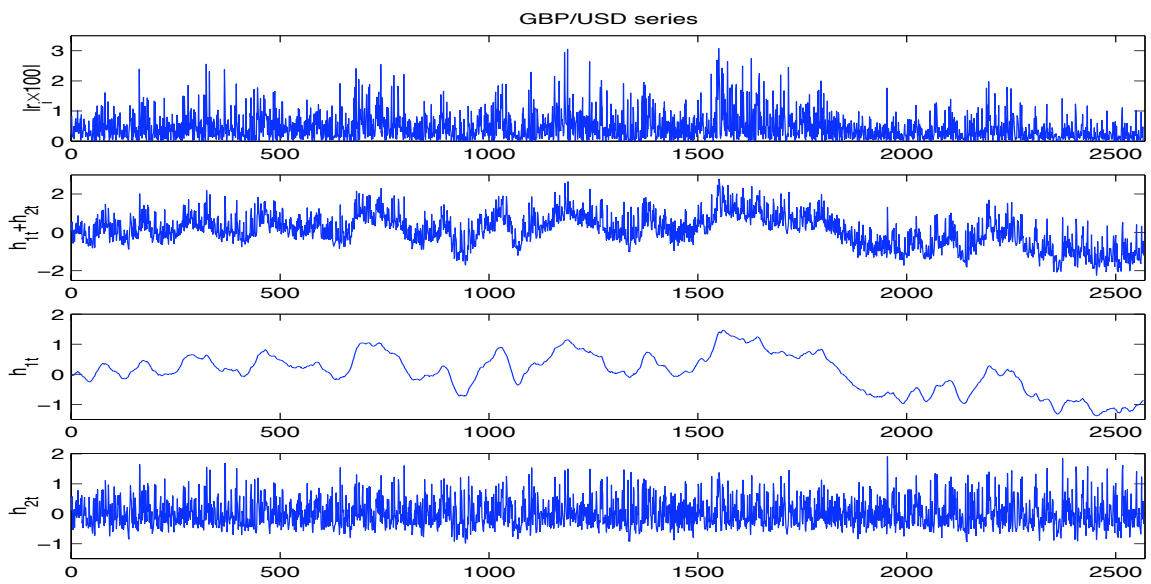


Figure 6: GBP/USD absolute value of the return series (first plot). Posterior means of the sum of the two factors (second plot). Posterior mean of the slow mean reverting factor (third plot) and posterior mean of the fast mean reverting factor (fourth plot), estimated using  $K=2$  factors, after burn-in and thinning of 1 iteration every 25 iterations, for a chain run for 100,000 iterations.

One can also construct easily any functional of the parameters in the model and do inference on it. Indeed, an interesting quantity to look at is the variance of the long-run distributions of the factors. For this dataset these long-run variances (as defined in Section 2.1) are of the same order when  $K=2$  (posterior means are 0.5597 for the slow mean-reverting process and 0.5589 for the fast mean reverting one), while, when estimating them with  $K=1$ , we get a long-run variance of 0.8215. Given that the factors are quite separated, this result gives a very intuitive way to construct joint priors on the mean reverting parameters and the variance parameters.

## 5. IMPLICATION ON DERIVATIVES PRICING AND HEDGING

In previous sections the existence of short and long time scales in driving volatility processes has been justified through the method of McMC. This finding is consistent to those in Alizadeh *et al.* (2002) and Chernov *et al.* (2002) based on data defined under the historical or physical probability space. Due to the fact that these time scales are preserved under change of probability measures, one may find empirical evidence of the existence of multiple time scales in volatility from market prices of derivatives. By means of singular and regular perturbation methods, higher order approximations of option prices are derived by taking into account the effect of stochastic volatility. Fast and slow time scales are justified to provide a good fit to term structure of implied volatility, interest rates, or credit spread yields, etc. These can be found in a series of papers by Fouque *et al.* (2003, 2006), Cotton *et al.* (2004), Papageorgiou and Sircar (2007), and references therein.

Given a fully specified multiscale stochastic volatility model, Fouque and Han (2007, 2008) develop a martingale control variate method to deal with option pricing problems by means of Monte Carlo simulations. The martingale control represents the value of portfolio process consisting of some imperfect hedging strategy. As a result the variance reduced by this method indicates the mean squared error of mis-replication by associated hedging strategy. Han and Lai (2007) incorporate a *randomized Quasi Monte Carlo* and show a significant improvement on the convergence of Monte Carlo simulations. Han and Vestal (2008) estimate the joint default probability of many names under structural models by an importance sampling method. This result is helpful to deal with high-dimensional problems for the use of, for example, credit portfolio (coherent) risk management and the evaluation of Collateralized Debt Obligations and Basket Default Swaps.



## 6. CONCLUSIONS

This paper addressed identifiability and estimation problems for stochastic volatility models with two volatility factors given that those can come at different, and well-separated, time scales.

We show that the number of factors can be visually identified from a simple analysis of the Markov Chains under different models, since overparametrized models will bring clear signs of lack of convergence, while underparametrized models will bring an incorrect appearance of convergence. This is of special relevance to industry practitioners, who might find model selection tools unfeasible under time constraints. This algorithm relies heavily on the factors being well-separated, and the degree of separation, as well as the length of the dataset, does have an influence of the ability to estimate second factors in a given series. We would, therefore, recommend running the algorithm with at least 2 factors as a first approach, as it seems to be quite robust to the length of the datasets and the parameter settings. If there is lack of convergence, reduce to  $K=1$ .

One could argue that, if the second factor is small ( $\phi_2$  close to zero), there is no practical point and no gain in introducing more than one factor. However, as shown in the simulated data and in the application, a (much) more accurate identification of the slow mean-reverting volatility factor is obtained. This will usually come in the form of a more persistent slow mean-reverting factor ( $\phi_1$  closer to one), and, therefore, have a huge impact on volatility predictions.

## ACKNOWLEDGEMENTS

This material was based upon work supported by the National Science Foundation under Agreements No. DMS-0112069 and DMS-0071744. We would like to thank Prof. Mike West for the foreign exchange data used in this paper, as well as the members of the Stochastic Computation research group at the Statistical and Applied Mathematical Sciences Institute (SAMSI), and the SAMSI Workshops participants for their valuable comments. We also thank Prof. Jim Berger for his help and support in the development of this paper.

## APPENDIX A: PROOF OF INDEPENDENT FACTOR EQUIVALENCE

In the case where factors are positively correlated, the inferencial problem simplifies to one of independent volatility processes.

*Proof for the 2-factor case.* Let  $y_t \sim N_1\left(y_t \mid 0, e^{\mathbf{1}'\mathbf{h}_t + \mu}\right)$  be a univariate time series driven by two volatility factors represented in the vector  $\mathbf{h}_t$ . Define an i.i.d. sequence of stochastic (degenerate) vectors  $\{\boldsymbol{\lambda}_t \sim N_2(\boldsymbol{\lambda}_t \mid \mathbf{0}, \mathbf{R})\}_{t=1}^T$ , where  $\mathbf{R}$  is a positive semi-definite 2 by 2 covariance matrix such that  $0 < r_{11} = r_{22} = -r_{12}$ . We can see that  $\mathbf{1}'\boldsymbol{\lambda}_t \equiv 0 \forall t$ . Then we can define the “equivalent” uncorrelated volatility vector  $\mathbf{h}_t^* = \mathbf{h}_t - \boldsymbol{\lambda}_t$  with conditional diagonal covariance matrix  $\mathbf{D}$ . The observation equation remains unchanged, namely

$$y_t \sim N_1\left(y_t \mid \mathbf{0}, e^{\mathbf{1}'\mathbf{h}_t + \mu}\right) = N_1\left(y_t \mid \mathbf{0}, e^{\mathbf{1}'(\mathbf{h}_t^* + \boldsymbol{\lambda}_t) + \mu}\right) = N_1\left(y_t \mid \mathbf{0}, e^{\mathbf{1}'\mathbf{h}_t^* + \mu}\right) \quad (6.8)$$

while we can see in (6.9) that the evolution equations are now independent:

$$\begin{aligned} \mathbf{h}_t &\sim N_1(\mathbf{h}_t \mid \Phi \mathbf{h}_{t-1}, \Sigma) \\ \mathbf{h}_t^* + \boldsymbol{\lambda}_t &\sim N_2(\mathbf{h}_t^* + \boldsymbol{\lambda}_t \mid \Phi(\mathbf{h}_{t-1}^* + \boldsymbol{\lambda}_{t-1}), \Sigma) \\ \mathbf{h}_t^* &\sim N_2\left(\mathbf{h}_t^* \mid \Phi \mathbf{h}_{t-1}^*, \underbrace{\mathbf{R} + \Phi \mathbf{R} \Phi + \Sigma}_{\mathbf{D}}\right). \end{aligned} \quad (6.9)$$

If  $\Sigma_{12} > 0, \phi_1, \phi_2 \neq 0, \exists \mathbf{R} : \mathbf{R} + \Phi \mathbf{R} \Phi + \Sigma = \mathbf{D}$  for some diagonal positive definite matrix  $\mathbf{D}$  with elements  $d_i^2$ . Matching elements gives us that  $r_{11} = r_{22} = \Sigma_{12}/(1 + \phi_1\phi_2) = -r_{12}$ . The variances of these independent volatility processes will be one-to-one versions of the original variances, since  $d_i^2 = \Sigma_{12}(1 + \phi_i^2)/(1 + \phi_1\phi_2) + \Sigma_{ii}$ . Therefore any two positively correlated volatility processes are equivalent to two independent processes, making a parameter like  $\Sigma_{12}$  nonidentifiable.

The implied invariant distribution (marginal) for the initial state would just be a normal centered at zero and with diagonal covariance matrix  $\Omega = \Phi \Omega \Phi + \mathbf{D}$ .

*Proof for the K-factor case.* The proof can be extended to the general case. If our univariate time series was driven by an arbitrary number K of positively correlated volatility factors, represented in the vector  $\mathbf{h}_t$ , we can define a sequence of stochastic vectors  $\{\boldsymbol{\lambda}_t \sim N_K(\boldsymbol{\lambda}_t \mid \mathbf{0}, \mathbf{R})\}_{t=1}^T$  that comply with the diagonality restriction  $\mathbf{R} + \Phi \mathbf{R} \Phi + \Sigma = \mathbf{D}$  and with the equivalence restriction  $\mathbf{1}^T \mathbf{R} \mathbf{1} = 0$ .

Solving  $r_{ij} = -\Sigma_{ij}/(1 + \phi_i\phi_j)$  for  $i \neq j$ , guarantees the diagonality of the transformed factors. One possible solution that guarantees the uniqueness of the equivalence condition

can be found by imposing an auxiliary restriction, namely  $\mathbf{1}^T \mathbf{R} = \mathbf{0}$ , which turns out to satisfy  $r_{ii} = -\sum_{j \neq i} r_{ij}$ .

## APPENDIX B: FULL CONDITIONALS

*Sampling  $\mathbf{h}$*  After approximating the likelihood through a mixture of normals and, conditionally on the mixture component, we can (jointly) sample the vector of volatilities, including the initial state, using its multivariate DLM structure by running the Forward Filtering/Backward Smoothing algorithm as detailed in West and Harrison, (1999).

*Sampling  $\delta$*  Each of the T indicator vectors,  $\delta_t$ , comes from a seven-component multinomial distribution with 1 outcome and with probabilities proportional to  $q_i^* = q_i f_N(\log(y_t^2) | m_i + \mu + \mathbf{1}' \mathbf{h}_t - 1.2704, v_i)$ , where  $m_i, v_i, q_i, i = 1, \dots, 7$  are the prior means, variances and weights of the mixture distributions and  $f_N$  denotes the normal density function. Therefore  $\delta_t \sim MN(\delta_t | 1, q_1^*, \dots, q_7^*)$ .

*Sampling  $\mu$*  Under a normal prior  $\pi(\mu) \sim N(m_0, v_0)$  the full conditional for  $\mu$  is also normal,  $[\mu | \dots] \sim N_1(\mu | m_0^*, v_0^*)$ , where  $m_0^* = (1/v_0^*) \times \sum_{t=1}^T (\log(y_t^2) - \mathbf{1}' \mathbf{h}_t - a_t + 1.2704)/b_t$  and  $v_0^* = (v_0^{-1} + \sum_{t=1}^T b_t^{-1})^{-1}$  and  $a_t, b_t$  are the mixture component means and variances,  $a_t = \sum_{i=1}^7 \delta_{it} m_i$  and  $b_t = \sum_{i=1}^7 \delta_{it} v_i$ .

*Sampling  $\Phi$*  Imposing stationarity and ordering a priori can be done with a uniform prior on the domain  $\phi_1 > \phi_2 > \dots > \phi_K$ . Then we can sample  $\phi_k^* \sim N_1(\phi_k^* | m_k, v_k)$ , where  $m_k = (1/v) \sum_{t=2}^T h_{k,t} h_{k,t-1}$  and  $v_k = \sum_{t=1}^{T-1} h_{k,t}^2$  until the proposed set of values follows the ordering conditions. Then, if we define  $\Omega^* = \Phi \Omega^* \Phi + \Sigma$ , we will set  $\Phi = \Phi^*$  with probability equal to  $\min(1, \exp\{\log[f_N(\mathbf{h}_0 | 0, \Omega^*)] - \log[f_N(\mathbf{h}_0 | 0, \Omega)]\})$ .

*Sampling  $\Sigma$*  We can update the components of  $\Sigma$  either all at once or one at a time. Under a prior of the form  $\pi(\Sigma) \propto \prod_{k=1}^K IG(\sigma_k^2 | a_k, b_k)$  ( $a_k, b_k$  can denote previous information), we can sample  $\sigma_k^2 \sim IG(\sigma_k^2 | \alpha_k/2, \beta_k/2)$ , where  $\alpha_k = T + 1 + 2a_k$  and  $\beta_k = \sum_{t=1}^T [h_{k,t}^2 - \phi_k h_{k,t-1}]^2 + h_{k,0}^2 (1 - \phi_k^2) + 2b_k$ .

## REFERENCES

### References

- [1] Aguilar Aguilar O. and West, M.,(2000) ‘‘Bayesian Dynamic Factor models and variance matrix discounting for portfolio allocation’’, *Journal of Business and Economic Statistics*, 18, pp. 338-357.

- [2] Alizadeh, S., Brandt M. and Diebold F. (2002) "Range-based estimation of stochastic volatility models." *Journal of Finance*, 57, pp 1047-1091.
- [3] Bates D. (1996) "Testing option pricing models, in Statistical Methods in Finance," Vol. 14 of Handbook of Statistics, eds. G.Maddala and C. Rao, North Holland, Amsterdam (1996) Chap. 20, 567-611.
- [4] Bollerslev, T. (1986) "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics* 31, 307-327.
- [5] Chernov M., Gallant R., Ghysels E., and Tauchen G. (2003) "Alternative models for stock price dynamics." *Journal of Econometrics*, Vol. 116, 1-2, 225-257.
- [6] Chib, S., Nardari, F., and Shephard, N. (2002), "Markov Chain Monte Carlo methods for stochastic volatility models." *Journal of Econometrics*. 108, 281-316.
- [7] Cotton, P., Fouque, J.-P., Papanicolaou, G., and Sircar, R. (2004) "Stochastic Volatility Corrections for Interest Rate Derivatives," *Mathematical Finance*, Vol. 14, No.2, 173-200.
- [8] Engle, R.F. (1982) "Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation", *Econometrica*, 50, 987-1008.
- [9] Fouque, J.-P. and Han, C.-H. (2004) "Variance Reduction for Monte Carlo Methods to Evaluate Option Prices under Multi-factor Stochastic Volatility Models", *Quantitative Finance*, Volume 4(5), 597-606.
- [10] Fouque, J.-P. and Han, C.-H. (2007) "A Martingale Control Variate Method for Option Pricing with Stochastic Volatility", *ESAIM Probability & Statistics* 11, 40-54.
- [11] Fouque, J.-P. and Han, C.-H. (2008) "Asymmetric Variance Reduction for Pricing American Options", to appear in *Mathematical Modeling and Numerical Methods in Finance*, Edited by A. Bensoussan and Q. Zhang.
- [12] Fouque, Jean-Pierre, Papanicolaou, G., Sircar, R., and Sølna, K. (2003) "Multiscale Stochastic Volatility Asymptotics," *SIAM Journal on Multiscale Modeling and Simulation*, 2(1), 22-42.
- [13] Fouque, Jean-Pierre, Sircar, R., and Sølna, K. (2006) "Stochastic Volatility Effects on Defaultable Bonds", *Applied Mathematical Finance*, 13(3), 215-244.
- [14] Fridman, M. and Harris, L. (1998) "A maximum likelihood approach for non-Gaussian stochastic volatility models," *Journal of Business and Economic Statistics*, 16, pp. 284-291.

- [15] Han, C.-H. and Lai, Y. (2007) “Generalized Control Variate Methods for Pricing Asian Options”, submitted.
- [16] Han, C.-H. and Vestal, D. (2008) “Importance Sampling for Estimating Joint Default Probability under Structural Models,” working paper.
- [17] Harvey, A., Ruiz, E. and Shephard N. (1994) “Multivariate stochastic variance models,” *Review of Economic Studies*, 61, pp 247-264.
- [18] ter Horst, E. (2003), “A Levy generalization of compound Poisson Processes in Finance: Theory and applications”, Ph.D. thesis, Institute of Statistics and Decision Sciences.
- [19] Jacquier, E., Polson, N. G. and Rossi, P. E. (1994) “Bayesian analysis of stochastic volatility models,” *Journal of Business and Economic Statistics*, 12, pp 413-417.
- [20] Jacquier, E., Polson, N.G. and Rossi, P.E., (1998) “Stochastic Volatility: Univariate and Multivariate extensions”, *CIRANO* working paper.
- [21] Kim, S., Shephard, N. and Chib, S. (1998) “Stochastic Volatility: Likelihood Inference and Comparison with ARCH models”, *The Review of Economic Studies*, 65, pp 361-393.
- [22] LeBaron, B. (2001) “Stochastic volatility as a simple generator of apparent financial power laws and long memory”. *Quantitative Finance*, Vol 1, No 6: 621-31.
- [23] Lopes, H. (1999) Ph.D. thesis, Institute of Statistics and Decision Sciences.
- [24] Melino A. and Turnbull S. (1990) “Pricing foreign currency options with stochastic volatility,” *Journal of Econometrics*, 45, pp 239-265.
- [25] Nelson, D. (1991) “Conditional heteroskedasticity in asset returns: A new approach,” *Econometrica*, 59(2), pp 347-370.
- [26] Papageorgiou, E. and Sircar, R. (2007) “Multiscale Intensity Models and Name Grouping for Valuation of Multi-name Credit Derivatives”, submitted.
- [27] Robert, C., and Casella, G. (1999), ”Monte Carlo Statistical Methods”, Springer Texts in Statistics.
- [28] West, M., and Harrison, J. (1999), “Bayesian Forecasting and Dynamic Models”, Springer-Verlag.