ELSEVIER

Factor stochastic volatility with time varying loadings and Markov switching regimes

Hedibert Freitas Lopes[a],*, Carlos Marinho Carvalho[b]

[a]*Graduate School of Business, University of Chicago, 5807 South Woodlawn Avenue, Chicago, IL 60637, USA*
[b]*Institute of Statistics and Decision Sciences, Duke University, 214 Old Chemistry Building, Durham, NC 27708-0251, USA*

Abstract

We generalize the factor stochastic volatility (FSV) model of Pitt and Shephard [1999. Time varying covariances: a factor stochastic volatility approach (with discussion). In: Bernardo, J.M., Berger, J.O., Dawid, A.P., Smith, A.F.M. (Eds.), Bayesian Statistics, vol. 6, Oxford University Press, London, pp. 547–570.] and Aguilar and West [2000. Bayesian dynamic factor models and variance matrix discounting for portfolio allocation. J. Business Econom. Statist. 18, 338–357.] in two important directions. First, we make the FSV model more flexible and able to capture more general time-varying variance–covariance structures by letting the matrix of factor loadings to be time dependent. Secondly, we entertain FSV models with jumps in the common factors volatilities through So, Lam and Li's [1998. A stochastic volatility model with Markov switching. J. Business Econom. Statist. 16, 244–253.] Markov switching stochastic volatility model. Novel Markov Chain Monte Carlo algorithms are derived for both classes of models. We apply our methodology to two illustrative situations: daily exchange rate returns [Aguilar, O., West, M., 2000. Bayesian dynamic factor models and variance matrix discounting for portfolio allocation. J. Business Econom. Statist. 18, 338–357.] and Latin American stock returns [Lopes, H.F., Migon, H.S., 2002. Comovements and contagion in emergent markets: stock indexes volatilities. In: Gatsonis, C., Kass, R.E., Carriquiry, A.L., Gelman, A., Verdinelli, I., Pauler, D., Higdon, D. (Eds.), Case Studies in Bayesian Statistics, vol. 6, pp. 287–302].

© 2007 Elsevier B.V. All rights reserved.

1. Introduction

A vast literature on Bayesian analysis of univariate stochastic volatility (SV) processes have appeared after the seminal work of Jacquier et al. (JPR 1994) who perform fully Bayesian inference through a Markov chain Monte Carlo (MCMC) scheme. For instance, Kim et al. (1998) replace JPR's scheme by a forward filtering–backward sampling (FFBS) step (Carter and Kohn, 1994; Frühwirth-Schnatter, 1994) when sampling the log-volatilities. Jensen (2004) develops semiparametric inference for long-memory SV processes, while So et al. (1998) and Carvalho and Lopes (2007) accommodate Markov jumps in the log-volatilities.

---

* Corresponding author. Tel.: +1 773 834 5458; fax: +1 773 702 0458.
  E-mail address: hlopes@ChicagoGSB.edu (H.F. Lopes).

The literature on multivariate SV models is less abundant, with the works of Harvey et al. (1994), Pitt and Shephard (1999), Aguilar and West (2000), Lopes and Migon (2002) and Chib et al. (2006) forming the basis for the developments we consider here. Roughly speaking, they model the levels (or first differences) of a set of (financial) time-series by a standard normal factor model (Lopes and West, 2004) in which both the common factor variances and the specific (or idiosyncratic) time-series variances are modeled as univariate SV processes. The main practical and computational advantage of the factor stochastic volatility (FSV) model is its parsimony, where all the variances and covariances of a vector of time-series are modeled by a low-dimensional SV structure dictated by common factors. It is fairly common to find that, for large vectors of time-series, the number of common factors is usually one or two orders of magnitude smaller, which speeds up computation and estimation considerably.

Our main contribution is twofold. First, we let the factor loading matrix to be time-dependent allowing the FSV model to capture more general time-varying correlation structures. Second, we extend So et al. (1998) and Carvalho and Lopes (2007) Markov jumps to model the common factors' stochastic log-volatilities. We start, in Section 2, with a general introduction to FSV models followed by the description of the proposed extensions in Section 3. Bayesian inference and computation are developed in Section 4 where customized MCMC algorithms are presented. The extensions are illustrated through two financial time-series examples in Section 5 followed by a final section of discussions and thoughts on directions for future investigation.

2. FSV models: a brief review

In this section we briefly review standard factor analysis and FSV models, as well as modeling and identification issues.

2.1. Standard factor analysis

Standard normal factor analysis is the backbone of FSV models. In this context, data on p related variables are considered to arise through random sampling from a zero-mean multivariate normal distribution where $\mathbf{\Omega}$ denotes a $p \times p$ non-singular variance matrix. For any specified positive integer $q \leqslant p$, the standard q-factor model relates each $\mathbf{y}_t$ to an underlying q-vector of random variables $\boldsymbol{f}_t$, the common factors, via

$$\mathbf{y}_t = \boldsymbol{\beta f}_t + \boldsymbol{\epsilon}_t, \tag{1}$$

where (i) the factors $\boldsymbol{f}_t$ are independent with $\boldsymbol{f}_t \sim \mathrm{N}(\mathbf{0}, \mathbf{I}_q)$, (ii) the $\boldsymbol{\epsilon}_t$ are independent normal p-vectors with $\boldsymbol{\epsilon}_t \sim \mathrm{N}(\mathbf{0}, \mathbf{\Sigma})$, and $\mathbf{\Sigma} = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_p^2)$, (iii) $\boldsymbol{\epsilon}_t$ and $\boldsymbol{f}_s$ are independent for all t and s and (iv) $\boldsymbol{\beta}$ is the $p \times q$ factor loadings matrix.

Under this model, the variance–covariance structure of the data distribution is constrained with $\mathbf{\Omega} = V(\mathbf{y}_t | \mathbf{\Omega}) = V(\mathbf{y}_t | \boldsymbol{\beta}, \mathbf{\Sigma}) = \mathbf{\Omega} = \boldsymbol{\beta}\boldsymbol{\beta}' + \mathbf{\Sigma}$. Conditional on the common factors, observable variables are uncorrelated. In other words, the common factors explain all the dependence structures among the p variables. For any element y_{it} and y_{jt} of $\mathbf{y}_t$ and conditionally on $\boldsymbol{\beta}$ and $\mathbf{\Sigma}$, we have the characterizing moments, (i) $\mathrm{var}(y_{it}|f) = \sigma_i^2$, (ii) $\mathrm{cov}(y_{it}, y_{jt}|f) = 0$, (iii) $\mathrm{var}(y_{it}) = \sum_{l=1}^q \beta_{il}^2 + \sigma_i^2$ and (iv) $\mathrm{cov}(y_{it}, y_{jt}) = \sum_{l=1}^q \beta_{il}\beta_{jl}$.

In practical problems, especially with larger values of p, the number of factors q will often be small relative to p, so most of the variance–covariance structure is explained by a small number of common factors. The *uniquenesses*, or *idiosyncratic variances*, σ_i^2 measure the residual variability in each of the data variables once that contributed by the factors is accounted for. Modern MCMC-based posterior inference in standard factor analysis appears in, among others, Geweke and Zhou (1996) and Lopes and West (2004).

2.2. Model structure and identification issues

The q-factor model is invariant under transformations of the form $\boldsymbol{\beta}^* = \boldsymbol{\beta P}'$ and $\boldsymbol{f}_t^* = \boldsymbol{P f}_t$, where $\boldsymbol{P}$ is any orthogonal $q \times q$ matrix. There are many ways of identifying the model by imposing constraints on $\boldsymbol{\beta}$, including constraints to orthogonal $\boldsymbol{\beta}$ matrices, and constraints such that $\boldsymbol{\beta}'\mathbf{\Sigma}^{-1}\boldsymbol{\beta}$ is diagonal. The alternative preferred here is to constrain $\boldsymbol{\beta}$ so that it is a block lower triangular matrix, assumed to be of full rank, with diagonal elements strictly positive. This form is used, for example, in Geweke and Zhou (1996), Aguilar and West (2000) and Lopes and Migon (2002), and provides

both identification and, often, useful interpretation of the factor model. With p non-zero σ_i parameters, the resulting factor form of $\mathbf{\Omega}$ has $p(q+1) - q(q-1)/2$ parameters, compared with the total $p(p+1)/2$ in an unconstrained (or $q = p$) model; leading to the constraint that $p(p+1)/2 - p(q+1) + q(q-1)/2 \geqslant 0$, which provides an upper bound on q. Even for small p, the bound will often not matter as relevant q values will not be so large. In realistic problems, with p in double digits or more, the resulting bound will rarely matter. Finally, note that the number of factors can be increased beyond such bounds by setting one or more of the residual variances σ_i to zero.

Lopes and West (2004) examine the problem of determining the number of factors in a static factor model while showing that the order of the variable is immaterial. The same is not necessarily true when both common and specific factors exhibit temporal dynamics. Vrontos et al. (2003), for instance, examine the importance of the order of the variables in a full-factor GARCH when modeling financial time-series and perform Bayesian model averaging.

2.3. FSV models

Posterior inference for the standard FSV models appears in Pitt and Shephard (1999) and Aguilar and West (2000). In this more general context, Eq. (1) is replaced by

$$(\mathbf{y}_t | f_t, \boldsymbol{\beta}, \boldsymbol{\Sigma}_t) \sim N(\boldsymbol{\beta} f_t; \boldsymbol{\Sigma}_t), \tag{2}$$

where $(f_t | \mathbf{H}_t) \sim N(\mathbf{0}; \mathbf{H}_t)$, with $\mathbf{H}_t = \text{diag}(h_{1t}, \ldots, h_{qt})$ the common factors' variance matrix. The main departure from standard factor analysis lies in the time-varying structure of $\boldsymbol{\Sigma}_t$ and $\mathbf{H}_t$, with $\eta_{it} = \log \boldsymbol{\Sigma}_{ii,t}$ following a standard SV evolution

$$(\eta_{it} | \eta_{i,t-1}, \tilde{\alpha}_i, \rho_i, \tau_i^2) \sim N(\tilde{\alpha}_i + \rho_i \eta_{i,t-1}, \tau_i^2), \tag{3}$$

for $i = 1, \ldots, p$. This is one of the simplest but certainly the most used specification in the literature (Jacquier et al., 1994).

Similarly, the q-dimensional vector of common factors' log-volatilities, $\boldsymbol{\lambda}_t = (\lambda_{1t}, \ldots, \lambda_{qt})'$, is modeled by a multivariate first-order autoregressive (VAR) model

$$(\boldsymbol{\lambda}_t | \boldsymbol{\lambda}_{t-1}, \boldsymbol{\alpha}, \boldsymbol{\phi}, \mathbf{U}) \sim N(\boldsymbol{\alpha} + \boldsymbol{\phi} \boldsymbol{\lambda}_{t-1}; \mathbf{U}), \tag{4}$$

for $\lambda_{it} = \log h_{it}$, $i = 1, \ldots, q$, $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_q)'$, and $\boldsymbol{\phi} = \text{diag}(\phi_1, \ldots, \phi_q)$, with correlated innovations characterized by the non-diagonal matrix $\mathbf{U}$ (Aguilar and West, 2000). When $\mathbf{U}$ is a diagonal matrix, the above multivariate model is reduced to q univariate conditionally independent autoregressive models (Pitt and Shephard, 1999).

3. Proposed extensions

In this section we argue in favor of two very important extensions to the FSV model. In the first case, $\boldsymbol{\beta}$ is replaced by $\boldsymbol{\beta}_t$, i.e., the factor loadings matrix is allowed to be time-varying. As will be discussed below, this extension allows more flexible correlation structures to be entertained. In the second case, Markovian jumps are allowed in the levels of the common factors' log-volatilities. These jumps allow for fast track of possible abrupt changes in the variance and covariance structures while addressing the issue of high persistency commonly present in financial data.

3.1. Time-varying loadings

We first consider the case of time-varying loadings. To setup the evolution for the unconstrained elements of $\boldsymbol{\beta}_t$ we stack them up in a $d = pq - q(q-1)/2$ dimensional vector $\tilde{\boldsymbol{\beta}}_t = (\beta_{21,t}, \beta_{31,t}, \ldots, \beta_{p,q,t})$, that follows a first-order autoregression evolution,

$$(\tilde{\boldsymbol{\beta}}_t | \tilde{\boldsymbol{\beta}}_{t-1}, \boldsymbol{\zeta}, \boldsymbol{\Theta}, \mathbf{W}_t) \sim N(\boldsymbol{\zeta} + \boldsymbol{\Theta} \tilde{\boldsymbol{\beta}}_{t-1}; \mathbf{W}_t) \tag{5}$$

with $\boldsymbol{\zeta} = (\zeta_1, \ldots, \zeta_d)$, $\boldsymbol{\Theta} = \text{diag}(\theta_1, \ldots, \theta_d)$. We will further assume that the evolution matrices $\mathbf{W}_t$ are completely specified by a single and known discount factor $\delta \in (0, 1)$ (West and Harrison, 1997) that measures the decay or loss of information incurred in the evolution mechanism.

Our main motivation for time-dependent loadings is to permit changes in covariances that are not exclusively associated to changes in the individual factor variances. For illustrative purpose, let us consider the simple scenario of a one-FSV model. From Eq. (2), the unconditional covariance between any two variables y_{it} and y_{jt} is given by

$$cov(y_{it}, y_{jt}) = \beta_{i1t}\beta_{j1t}h_{1t},$$

which will be time varying when either β_{i1t} or β_{j1t} or h_{1t} or any combination of them are time varying. In the fixed loadings context, time changes in covariances are exclusively due to changes in the common factor's variance, h_{1t}. As a consequence, it can be easily seen that the correlation between y_{it} and y_{jt} is

$$\rho_{ijt} = \frac{\beta_{i1t}\beta_{j1t}}{\sqrt{(\beta_{i1t}^2 + \sigma_{it}^2/h_{1t})(\beta_{j1t}^2 + \sigma_{jt}^2/h_{1t})}};$$

so a few possible sources for increase (decrease) in the correlation are: (i) decrease (increase) in either σ_{it} or σ_{jt}, (ii) increase (decrease) in h_{1t} or (iii) increase (decrease) in either β_{i1t} or β_{j1t}. The third case is only possible under our generalization. The larger (smaller) the β_{i1t} is, the more (less) important the common factor is in explaining the covariance (correlation) between y_{it} and y_{jt} without implications to the temporal behavior of σ_{it}, σ_{jt} or h_{1t}. The same rationale applies to the general q-FSV model.

It is worth noting that by allowing time-varying structures in the loadings matrix we are extending previous model structures while maintaining the model fully identified over time by the same constraints described above (Aguilar and West, 2000). Aguilar and West's (2000) model is a particular case when $\delta = 1.0$, while Pitt and Shephard's model is recovered when $u_{ij} = 0$ for $i \neq j$ and $\delta = 1.0$. Additionally, traditional factor models are recovered by setting $U = 0$ and letting $\delta = 1.0$ (Geweke and Zhou, 1996).

3.2. Markov switching SV

As in the SV model (Eq. (3)), the Markov switching SV (MSSV) model assumes that $f_{it}|\lambda_{it} \sim N(0, \exp(\lambda_{it}))$, but now the log-volatility, λ_{it}, follows a first order autoregressive process with jumps in the level

$$(\lambda_{it}|\lambda_{i,t-1}, \xi, s_t) \sim N(\alpha_{s_t} + \phi\lambda_{t-1}, u_{ii}), \tag{6}$$

where the regime indicator s_t follows a k-state first order Markovian process,

$$Pr(s_t = j|s_{t-1} = i) = p_{ij} \quad \text{for } i, j = 1, \ldots, k, \tag{7}$$

with $\alpha = (\alpha_1, \ldots, \alpha_k)$ and $\xi = (\alpha, \phi, \sigma^2)$. Eqs. (6) and (7) define a state-space model with two state vectors, namely $\lambda = (\lambda_1, ..., \lambda_T)$ and the discrete state $s = (s_1, \ldots, s_T)$. To avoid identification problems, we adopt the following reparametrization for α_{s_t}:

$$\alpha_{s_t} = \gamma_1 + \sum_{j=1}^{k}\gamma_j I_{jt}, \tag{8}$$

where $I_{jt} = 1$ if $s_t \geqslant j$ and zero otherwise, $\gamma_1 \in \Re$ and $\gamma_i > 0$ for $i > 1$. Finally, standard SV models can be recovered by forcing $k = 1$ in the MSSV models.

MSSV models will be incorporated in modeling the volatility process of the common factors. The idea behind this modification is an attempt to capture structural changes in the volatility processes in order to avoid the overestimation of the persistency parameter ϕ that is commonly observed in both univariate and multivariate SV literature (So et al., 1998; Carvalho and Lopes, 2007). This extension also allows for a better understanding of abrupt changes in volatilities due to economic forces and extreme market events.

In this model we assume a diagonal structure in U so that the level switches in each of the factors' volatilities are independent across factors, implying unrelated structural changes in the volatility processes.

4. Estimation

4.1. Prior information

In order to complete the requisites for the Bayesian analysis, prior distributions must be defined and we will restrict our analysis to conditionally conjugate prior distributions to facilitate the already complicated posterior analysis.

To begin with, the prior distribution for the unconstrained loadings at time $t = 1$ is $\tilde{\beta}_0 \sim N(m_0, C_0)$, with known hyperparameters m_0 and C_0. The prior distributions for components of ζ and the nonzero components of Θ are $\zeta_j \sim N(\zeta_{0j}, C_{0j})$ and $\theta_j \sim N(\theta_{0j}, V_{0j})$, respectively, for $j = 1, \ldots, d$. For the parameters that define both the factors's log-volatilities, α, ϕ, U, and the idiosyncratic factors's log-volatilities, $\tilde{\alpha}, \rho, \tau$ we follow Aguilar and West's (2000) suggestions. They assume independent normal priors for the univariate terms of α and $\tilde{\alpha}$ and independent truncated normal priors for the terms in ϕ and ρ. Inverted Wishart and inverted gamma distributions are assigned to U and each of the τ^2s. More specifically, $U \sim IW(r_0, r_0 R_0)$ and $\tau_i^2 \sim IG(v_{0i}/2, v_{0i}\tau_{0i}^2/2)$, for $i = 1, \ldots, p$ with r_0, R_0, v_{0i} and τ_{0i}^2 given hyperparameters.

Similar priors are used for the common parameters defining the MSSV. To complete the prior specification we assign $\gamma_1 \sim N(\gamma_{10}, C_{\gamma_1})$, $\gamma_i \sim TN_{(0,\infty)}(\gamma_{i0}, C_{\gamma_i})$ for and $p_i \sim Dir(u_{i0})$, for $p_i = (p_{i1}, \ldots, p_{ik})$ and $i = 1, \ldots, k$. In our examples the hyperparameters were chosen to represent fairly vague prior information.

4.2. MCMC posterior inference

Conditional on all information up to time T, posterior inference on a general q-FSV model can be done through a straightforward MCMC algorithm, as it appears in Geweke and Zhou (1996), Aguilar (1998) and Lopes (2000). The conditional independence structure of dynamic linear models together with *forward filtering backward sampling* (FFBS) (Carter and Kohn, 1994) and a finite mixture-of-normals approximation to log-square returns (Kim et al., 1998) create an efficient framework where samples from full conditional posterior distributions can be easily obtained.

In order to incorporate the proposed extensions to the general model, we have to first modify the sampling step for the unconstrained loadings $\tilde{\beta}_t$ (for $t = 1, \ldots, T$), that now vary in time, and also define a way to explore the full conditional distribution of the regime switching vector s_i for each factor.

Sampling $\tilde{\beta}_t$ is performed by constructing the following standard multivariate normal dynamic linear model:

$$z_t | \tilde{\beta}_t \sim N(F_t' \tilde{\beta}_t, \Sigma_t), \tag{9}$$

$$\tilde{\beta}_t | \tilde{\beta}_{t-1} \sim N(\zeta + \Theta \tilde{\beta}_{t-1}, W_t), \tag{10}$$

where $z_t = (z_{1t}, \ldots, z_{p-1,t})$ is defined by $z_{it} = y_{i+1,t} - f_{i+1,t}$ for $i = 1, \ldots, q - 1$ and $z_{it} = y_{i+1,t}$ for $i = q, \ldots, p - 1$, and $F_t = diag(F_{1t}, \ldots, F_{p-1,t})$ with $F_{it} = (f_{1t}, \ldots, f_{it})$ if $i = 1, \ldots, q-1$ and $F_{it} = (f_{1t}, \ldots, f_{qt})$ if $i = q, \ldots, p-1$. Completing the model with the discount factor δ, forward updating equations and retrospective recurrences for $\tilde{\beta}_t$ are available so a FFBS algorithm can be implemented to generate samples from the desired full conditional posterior distribution.

In the MSSV extension, the main required modification is sampling the discrete state vector of regimes $s_i = (s_{i1}, \ldots, s_{iT})$. Without loss of generality, let $q = 1$ so the full conditional distribution of s will only depend on λ, the vector of log-volatilities for the common factor. The joint, full conditional distribution of s can be written as

$$p(s | \lambda, \xi) = p(s_T | \lambda, \xi) \prod_{t=1}^{T-1} p(s_t | s_{t-1}, \lambda, \xi), \tag{11}$$

where

$$p(s_t | s_{t-1}, \lambda, \xi) \propto p(s_{t+1} | s_t) p(s_t | \lambda, \xi). \tag{12}$$

The above derivation is a direct result of the Markovian property of s and suggests that in order to generate a sample from the entire vector one should first sample $(s_T | \lambda, \xi)$ followed by samples from $(s_t | s_{t+1}, \lambda, \xi)$ for all $t = 1, \ldots, T - 1$, characterizing yet another FFBS scheme. The only missing piece is how to determine $p(s_t | \lambda, \xi)$ (*forward filtering* step) for all times, and the solution is provided by Hamilton's basic filter (Hamilton and Susmel, 1994). Details of general MCMC algorithm for MSSV models are found in So et al. (1998).

5. Illustrations

5.1. Daily exchange rate returns

To illustrate the time-varying loadings extension we analyze the returns on weekday closing spot prices for six currencies relative to the US dollar (as in Aguilar and West, 2000); the German Mark(DEM), British Pound(GBP), Japanese Yen(JPY), French Franc(FRF), Canadian Dollar(CAD) and Spanish Peseta(ESP). To keep the analysis comparable with Aguilar and West (2000) we only use the first 1000 observations ranging from 1/1/1992 to 10/31/1995.

Following Aguilar and West (2000), a $q = 3$ factor model was implemented with relatively vague priors for all model parameters. The prior for $\tilde{\alpha}_i$ is N(0, 25), while $\rho_i \sim T\mathrm{N}_{[0,1)}(0, 10)$. The elements of $\tilde{\boldsymbol{\beta}}_0$ are normally distributed with zero mean and unit variance. Finally, the priors for the FSV regression parameters, α_i and ϕ_i, are N(0, 25) and $2Be(20; 1.5) - 1$, respectively. For $\boldsymbol{U}$ we chose $\boldsymbol{R}_0 = 0.0015\boldsymbol{I}_3$ and $r_0 = 20$, in Aguilar and West's notation. Other combinations of $\boldsymbol{R}_0$ and r_0 were tested resulting in most of the parameters being unaffected, apart from those concerning the SV equations. The discount factor δ was set equal to 0.9975 representing slight moves on the factor loadings. Similar results were achieved with lower values for the discount factor, such as 0.99. The MCMC was run for 35,000 iterations, the last 5000 being used to perform the analysis. Different starting values were tried as well as different MCMC burn-in lengths. In general the results were pretty much the same, with the chain converging, in practical terms, after 20,000 iterations.

Inference for parameters defining the SV equations for both the common factor variances and for the idiosyncratic variances are summarized in Table 1. Graphical summaries are presented in Figs. 1 and 2. Fig. 1 shows the posterior means for the factor loadings through time while Fig. 2 displays the proportion of variation of the time-series explained by each of the common factors and what is left for the specific, idiosyncratic noise.

An interesting observation that highlights the importance of time-varying loadings in the context of this example is the change in the explanatory power of factor 1 (Fig. 1), the "European factor" on the British Pound. The final months of 1992 marks the withdrawal of Great Britain from the European Union exchange-rate agreement (ERM), fact that is captured in our analysis by changes in the British loading in factor 1 and emphasized by the changes in the percentage of variation of the British Pound explained by factors 1 and 2 (Fig. 2). If temporal changes on the factor loadings were not allowed, the only way the model could capture this change in Great Britain's monetary policy would be by a "shock" on the idiosyncratic variation of the Pound, reducing, in turn, the predictive ability of the latent factor structure.

Table 1

Exchange rates: posterior means and standard deviations (in parenthesis) for the fixed parameters defining the stochastic volatility equations for both common and specific factors

|	α_i	ϕ_i	u_{ii}
Common factors			
Factor 1	−10.191(0.376)	0.978(0.008)	0.053(0.015)
Factor 2	−11.828(1.351)	0.991(0.004)	0.060(0.016)
Factor 3	−10.942(0.866)	0.988(0.005)	0.078(0.024)
Country	$\tilde{\alpha}_i$	ρ_i	τ_i^2
Specific factors			
DEM	−13.665(1.507)	0.999(0.001)	0.001(0.000)
GBP	−15.446(0.212)	0.843(0.105)	0.000(0.000)
JPY	−13.758(0.323)	0.879(0.093)	0.000(0.000)
FRF	−14.351(0.411)	0.922(0.023)	0.648(0.181)
CAD	−11.651(0.190)	0.992(0.005)	0.001(0.000)
ESP	−11.967(0.248)	0.937(0.020)	0.201(0.061)

Fig. 1. Exchange rate: retrospective posterior means for the unconstrained elements of $\boldsymbol{\beta}_t$ when the idiosyncrasies have stochastic volatility structures (1/1/1992–10/31/1995).

5.2. Latin American stock returns

We now illustrate the MSSV generalization in an extended version of the data set in Lopes and Migon (2002). The data represent returns on weekday closing spot prices for four Latin American stock markets: the Brazilian Indice Bovespa (IBOVESPA), the Mexican Indice de Precios y Cotaziones (MEXBOL), the Argentinean Indice Merval (MERVAL), and the Chilean Indice de Precios Selectivos de Acciones (IPSA). The series are observed daily from January, 3rd 1994 to May, 26th 2005, comprising a total of 2974 observations. This period includes a set of international currency crises that had direct impact on Latin American markets (Carvalho and Lopes, 2007), generating higher levels of uncertainty and consequently higher levels of volatility. The goal of this example is to show how the FSV model with jumps

Fig. 2. Proportion of the time-series variances explained by each of the factors (common and specific), when the idiosyncrasies have stochastic volatility structures (1/1/1992–10/31/1995).

in the factor's volatilities is able to isolate these moments of crises reducing the estimated persistency of common volatilities.

Table 2 compares the posterior means from an traditional one-FSV model to posterior means from an FSV with 2 regimes in the common volatility process. Note that under the model with $k = 2$ the persistency parameter ϕ is likely to be smaller in line with conclusions drawn in Carvalho and Lopes (2007). The model with $k = 2$ estimates two unconditional means for the log-volatility process that corresponds to times of high and low risk in the market (bottom

Table 2
Latin American indices: posterior means for the fixed parameters defining the stochastic volatility equations for both common and specific factors with and without Markov switching

	FSV			FSV+MSSV			
	$\tilde{\alpha}$	ρ	τ^2	$\tilde{\alpha}$		ρ	τ^2
IBOVESPA	−0.202	0.980	0.040	−0.284	—	0.971	0.047
MEXBOL	−0.440	0.959	0.065	−0.434	—	0.957	0.051
MERVAL	−0.409	0.959	0.083	−0.508	—	0.947	0.068
IPSA	−0.600	0.947	0.058	−0.765	—	0.932	0.071
	α	ϕ	u	α_1	α_2	ϕ	u
Factor	−0.305	0.971	0.067	−0.951	−0.588	0.912	0.090
$E(\lambda_t)$	−10.517			−10.807	−6.682		

$E(\lambda_t)$ is the unconditional expectation of the common factor's log-volatility.

Fig. 3. Latin American stock returns: common factor volatilities when FSV and the FSV+MSSV models are implemented. The vertical lines represent situations where regime variable s_t equals two (alternative regimes).

of Table 2). More specifically, the posterior mean of the unconditional standard deviation of the common factor in the FSV model is roughly the same as the one obtained for the low-volatility regime in the FSV+MSSV; however, the factor is on a high-volatility state around 6% of the time, in which the unconditional standard deviation is about eight times higher. This allows the volatilities to react "faster" once a regime switch is identified, which is highlighted by Fig. 3 that compares FSV and FSV+MSSV common factor's volatilities.

The FSV+MSSV model captures market crashes (Asia Flue, Russian Cold and Brazilian Fever) in an instantaneous way due to the structural change caused by the regime shift. Unlike traditional FSV models, the switch mechanism avoids possible underestimation of volatilities in highly uncertain periods while allowing a faster return to "normality" after that crisis has passed (So et al., 1998).

6. Final remarks

In this paper we extended the standard FSV model in two important directions. Allowing for time-varying loadings helps capturing more complex dynamics in the returns variances while keeping the parsimony inherited from factor

analysis. Additional flexibility is added when structural changes are incorporated in the model for the common factors' log-volatilities, addressing the issue of high persistency commonly encountered in financial time-series. Methodological aspects and applicability of these extensions were discussed and MCMC schemes were designed to enable posterior inference. Two fairly realistic applications are revisited to illustrate and emphasize the practical relevance of the proposed extensions.

Throughout this work, posterior inference is performed based on the entire dataset. However, sequential estimation can be carried out by combining ideas from Aguilar and West (2000), Lopes (2000) and Carvalho and Lopes (2007) creating a sequential Monte Carlo filter (Doucet et al., 2001). Details of these inferential procedures and applications are found in Lopes and Carvalho (2005).

Part of our current and future developments are, amongst others, (i) the joint modeling of the discrete regime states for multiple common factors, (ii) the specification of more structured priors for the factor loadings matrix, possibly over time (Lopes et al., 2006; West, 2003), and (iii) model assessment (Lopes and West, 2004; Carvalho et al., 2005). Finally, we intend to compare our methodology with current GARCH-like time-varying covariance models, such as Bollerslev's (1990) constant conditional correlation model, Engle (2002) dynamic conditional correlation model, Der Weide's (2002) Generalized orthogonal GARCH and Vrontos et al.'s (2003) Full-factor GARCH model, to name a few.

References

Aguilar, O., 1998. Latent structure in Bayesian multivariate time-series models. Ph.D. Thesis, Duke University.

Aguilar, O., West, M., 2000. Bayesian dynamic factor models and variance matrix discounting for portfolio allocation. J. Business Econom. Statist. 18, 338–357.

Bollerslev, T., 1990. Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. Rev. Econom. Statist. 72, 498–505.

Carter, C., Kohn, R., 1994. On Gibbs sampling for state space models. Biometrika 81, 541–553.

Carvalho, C.M., Lopes, H.F., 2007. Simulation-based sequential analysis of Markov switching stochastic volatility models. Computational Statistics and Data Analysis 51, 4526–4542.

Carvalho, C.M., Wang, Q., Lucas, J., Chang, J., Nevins, J.R., West, M., 2005. High-dimensional sparse factor models and latent factor regression. ISDS Discussion Paper.

Chib, S., Nardari, F., Shephard, N., 2006. Analysis of high dimensional multivariate stochastic volatility models. J. Econometrics 134, 341–371.

Doucet, A., de Freitas, N., Gordon, N., 2001. Sequential Monte Carlo Methods in Practice. Springer, New York.

Engle, R., 2002. Dynamic conditional correlation: a simple class of multivariage GARCH models. J. Business Econom. Statist. 20, 339–350.

Frühwirth-Schnatter, S., 1994. Data augmentation and dynamic linear models. J. Time Ser. Anal. 15, 183–202.

Geweke, J.F., Zhou, G., 1996. Measuring the pricing error of the arbitrage pricing theory. Rev. Financial Stud. 9, 557–587.

Hamilton, J.D., Susmel, R., 1994. Autoregressive conditional heteroskedasticity and changes in regime. J. Econometrics 64, 307–333.

Harvey, A.C., Ruiz, E., Shephard, N., 1994. Multivariate stochastic variance models. Rev. Econom. Stud. 61, 247–264.

Jacquier, E., Polson, N.G., Rossi, P.E., 1994. Bayesian analysis of stochastic volatility models. J. Business Econom. Statist. 12, 371–389.

Jensen, M.J., 2004. Semiparametric Bayesian inference of long-memory stochastic volatility models. J. Time Ser. Anal. 25, 895–922.

Kim, S.N., Shephard, N., Chib, S., 1998. Stochastic volatility: likelihood inference and comparison with ARCH models. Rev. Econom. Stud. 65, 361–393.

Lopes, H.F., 2000. Bayesian analysis in latent factor and longitudinal models. Ph.D. Thesis, Duke University.

Lopes, H.F., Carvalho, C.M., 2005. On generalizations of factor stochastic volatility models. Technical Report, Graduate School of Business, University of Chicago.

Lopes, H.F., Migon, H.S., 2002. Comovements and contagion in emergent markets: stock indexes volatilities. In: Gatsonis, C., Kass, R.E., Carriquiry, A.L., Gelman, A., Verdinelli, I., Pauler, D., Higdon, D. (Eds.), Case Studies in Bayesian Statistics, vol. 6, pp. 287–302.

Lopes, H.F., Salazar, E., Gamerman, D., 2006. Spatial dynamic factor analysis. Technical Report, Graduate School of Business, University of Chicago.

Lopes, H.F., West, M., 2004. Bayesian model assessment in factor analysis. Statistica Sinica 14, 41–67.

Pitt, M.K., Shephard, N., 1999. Time varying covariances: a factor stochastic volatility approach (with discussion). In: Bernardo, J.M., Berger, J.O., Dawid, A.P., Smith, A.F.M. (Eds.), Bayesian Statistics, vol. 6, Oxford University Press, London, pp. 547–570.

So, M.K.P., Lam, K., Li, W.K., 1998. A stochastic volatility model with Markov switching. J. Business Econom. Statist. 16, 244–253.

Van Der Weide, R., 2002. A multivarite generalized orthogonal GARCH model. J. Appl. Econometrics 17, 549–564.

Vrontos, I.D., Dellaportas, P., Politis, D.N., 2003. A full-factor multivariate GARCH model. Econometrics J. 6, 312–334.

West, M., 2003. Bayesian factor regression models in the "large *p*, small *n*" paradigm. In: Bernardo, J.M., Bayarri, M.J., Berger, J.O., Dawid, A.P., Heckerman, D., Smith, A.F.M., West, M. (Eds.), Bayesian Statistics, vol. 7 Oxford University Press, Oxford, pp. 723-732.

West, M., Harrison, P.J., 1997. Bayesian Forecasting and Dynamic Model. Springer, New York.