

# Embarrassingly parallel sequential Markov-chain Monte Carlo for large sets of time series

ROBERTO CASARIN, RADU V. CRAIU\*, AND FABRIZIO LEISEN

---

Bayesian computation crucially relies on Markov chain Monte Carlo (MCMC) algorithms. In the case of massive data sets, running the Metropolis-Hastings sampler to draw from the posterior distribution becomes prohibitive due to the large number of likelihood terms that need to be calculated at each iteration. In order to perform Bayesian inference for a large set of time series, we consider an algorithm that combines “divide and conquer” ideas previously used to design MCMC algorithms for big data with a sequential MCMC strategy. The performance of the method is illustrated using a large set of financial data.

AMS 2000 SUBJECT CLASSIFICATIONS: Primary 62C40, 37M10; secondary 62M99.

KEYWORDS AND PHRASES: Big Data, Panel of time series, Parallel Monte Carlo, Sequential Markov-chain Monte Carlo.

---

## 1. INTRODUCTION

There is little doubt that one of the main challenges brought on by the advent of Big Data in Bayesian statistics is to develop Markov chain Monte Carlo (MCMC) algorithms for sampling a posterior distribution derived from a very large sample. While MCMC has become the default tool to study posterior distributions when they are not available in closed form, many commonly used sampling algorithms, e.g. the Metropolis-Hastings samplers, can become computationally prohibitive when a large number of likelihood calculations are needed at each iteration.

In recent years we have witnessed a large research effort devoted to dividing the MCMC computational load among a number of available processors and recombining the results with as little loss in statistical efficiency as possible. For instance, the approaches developed in [27] and [25] divide the available data in smaller batches and sample the resulting *partial posteriors* obtained from each batch of data. They propose different methods to combine the resulting draws so that the efficiency of the resulting Monte Carlo estimator is close to the one that would have been obtained if the full data posterior had been sampled. In the consensus MCMC

of [27] this is achieved via reweighting the partial posterior samples, while the embarrassingly parallel approach of [25] relies on kernel density approximations of the partial posteriors to produce an approximation of the full one. In [31] the authors propose a refined recombination strategy based on the Weierstrass transformation of all partial posteriors.

While dividing the whole data into batches can be done easily when the data are independent, one must proceed cautiously when the data exhibit long range dependencies, as is the case in time series. In such cases, simply splitting time series into blocks can lead to poor estimates of the parameters. Instead, one can sometimes bypass the computational load by sequentially updating the posterior over time [see, for instance, 18].

Sequential sampling may be improved when combined with parallel and possibly interacting Monte Carlo methods that were used elsewhere, e.g. for parallel adaptive MCMC [12], for interacting MTM [9], for population Monte Carlo [7], for Sequential Monte Carlo [16, 15] and for massively parallel computing [23].

Sequential estimation is useful in many applied contexts such as on-line inference of econometric models for both out-of-sample and in-sample analyses. However, sequential estimation is a challenging issue in Bayesian analysis due to the computational cost of the numerical procedures for posterior approximation. Moreover, the computational cost rapidly increases with the dimension of the model and the inferential task becomes impossible. In this sense our paper contributes to the recent expansion of the literature on the use of Central Processing Unit (CPU) and Graphics Processing Unit (GPU) parallel computing in econometrics (e.g., see [17], [29], [13], [14], [28]).

Our contribution here is to consider possible ways to combine strategies following the work of [25] with sequential MCMC in order to address difficulties that appear when studying large panel time series data models. Analyses of panel time series data are frequently encountered in the econometrics literature, as discussed in the review papers of [6] and [20]. However, the use of latent variables for time series panel models in combination with a Bayesian inference approach [21, 22, 3] can be quite challenging, due to the computational burden required for the latent variable estimation. These challenges motivate this work in which the Bayesian stochastic volatility model recently proposed in [32] and discussed in [8] is adapted to the context of large panel of time series.

---

arXiv: 1512.01496

\*Corresponding author.

In Sections 2 and 3 we introduce, respectively, the issues related to sequential sampling in models with latent variables and present our algorithm. Section 4 contains the simulation studies and the real data analysis. The paper closes with conclusions and future directions.

## 2. POSTERIOR DISTRIBUTION FACTORIZATION

Consider a time series sample  $\mathbf{y}_t = (y_{t1}, \dots, y_{tm}) \in Y \subset \mathbb{R}^m$ ,  $t = 1, \dots, T$  with probability density function (pdf)  $h_t(\mathbf{y}_t|\boldsymbol{\theta})$  where  $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^p$  is a parameter vector. Of interest is sampling from the posterior distribution of  $\boldsymbol{\theta}$ . We are considering here the case in which the latter task is made easier if a data augmentation approach is adopted. We introduce auxiliary variables  $\mathbf{x}_t \in X \subset \mathbb{R}^n$ ,  $1 \leq t \leq T$  that exhibit Markovian serial dependence, i.e. each  $\mathbf{x}_t$  has pdf  $g(\mathbf{x}_t|\mathbf{x}_{t-1}, \boldsymbol{\theta})$ . If  $f(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta})$  is the conditional pdf of  $\mathbf{y}_t$  given  $\mathbf{x}_t$  and  $\boldsymbol{\theta}$ , then for prior distribution  $\pi(\boldsymbol{\theta})$  we obtain the joint posterior of the parameters and the latent variables as

$$\pi(\boldsymbol{\theta}, \mathbf{x}_{1:T}|\mathbf{y}_{1:T}) = \frac{1}{Z} \prod_{t=1}^T f(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta})g(\mathbf{x}_t|\mathbf{x}_{t-1}, \boldsymbol{\theta})\pi(\boldsymbol{\theta})$$

where  $Z$  is the normalizing constant of  $\pi(\boldsymbol{\theta}, \mathbf{x}_{1:T}|\mathbf{y}_{1:T})$ ,  $\mathbf{x}_{1:T} = \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$  and  $\mathbf{y}_{1:T} = \{\mathbf{y}_1, \dots, \mathbf{y}_T\}$ .

Henceforth, we assume that the time series data has panel structure such that if we consider all the data collected up to time  $t$ ,  $\mathbf{y}_{1:t}$ , then it is possible to partition them into  $M$  blocks of size  $K$  each,

$$(1) \quad \mathbf{y}_{1:t} = \bigcup_{i=1}^M \mathbf{y}_{1:t}^{(i)}$$

where the  $i$ th block  $\mathbf{y}_{1:t}^{(i)}$  contains the measurements up to time  $t$  for the components  $k_{i-1}+1, k_{i-1}+2, \dots, k_i$  of  $\mathbf{y}_{1:t}$  (for notational simplicity we set  $k_i = K \times i$  but other allocations are possible), i.e. for all  $0 \leq i \leq M$

$$\mathbf{y}_{1:t}^{(i)} = \bigcup_{j=k_{i-1}+1}^{k_i} \{y_{1j}, \dots, y_{tj}\} := \bigcup_{j=k_{i-1}+1}^{k_i} \mathbf{y}_{1:t,j}$$

Each set  $\mathbf{y}_{1:t}^{(i)}$  of the partition contains the  $i$ th panel of dependent components of  $\mathbf{y}_t$ . An important assumption of the model is that, conditional on a parameter value  $\boldsymbol{\theta}$ , the components in two partition sets are independent, i.e.

$$(2) \quad \mathbf{y}_{1:t}^{(i)} \perp \mathbf{y}_{1:t}^{(i')},$$

for any  $i \neq i'$ .

Corresponding to the partition (1) there is an equivalent partition of the auxiliary variables

$$(3) \quad \mathbf{x}_{1:t} = \bigcup_{i=1}^M \mathbf{x}_{1:t}^{(i)}$$

where the components of  $\mathbf{x}_{1:t}^{(i)}$  correspond to the components

of  $\mathbf{y}_t$  included in  $\mathbf{y}_{1:t}^{(i)}$ . A second crucial assumption for the validity of our algorithm is the independence of the auxiliary variables contained in two elements of the partition (3), i.e.

$$(4) \quad \mathbf{x}_{1:t}^{(i)} \perp \mathbf{x}_{1:t}^{(i')},$$

for all  $1 \leq i \neq i' \leq M$ . Finally, we assume that  $\mathbf{y}_{1:t}^{(i)}$  depends only on those auxiliary variables included in  $\mathbf{x}_{1:t}^{(i)}$ , i.e. given the latter we have

$$(5) \quad \mathbf{y}_{1:t}^{(i)} \perp \mathbf{x}_{1:t}^{(i')}.$$

These assumptions are not unusual in the context of dynamic panel data models as they are used for theoretical derivations in [2], [4] and [5] as well as in applications [see, for instance, 1, 21, 22, 3].

The basic principle underlying our approach is to learn sequentially over time using different cross-sectional data blocks. Our algorithm relies on samples from the the joint posterior distribution of  $\boldsymbol{\theta}$  and  $\mathbf{x}_{1:t}^{(i)}$ ,  $\pi_{it}(\boldsymbol{\theta}, \mathbf{x}_{1:t}^{(i)})$ , conditional on the sub-sample  $\mathbf{y}_{1:t}^{(i)}$

$$(6) \quad \begin{aligned} \pi_{it}(\boldsymbol{\theta}, \mathbf{x}_{1:t}^{(i)}) &= \pi(\boldsymbol{\theta}, \mathbf{x}_{1:t}^{(i)}|\mathbf{y}_{1:t}^{(i)}) \\ &= \frac{1}{Z_{it}} \prod_{s=1}^t f(\mathbf{y}_s^{(i)}|\mathbf{x}_s^{(i)}, \boldsymbol{\theta})g(\mathbf{x}_s^{(i)}|\mathbf{x}_{s-1}^{(i)}, \boldsymbol{\theta})\pi(\boldsymbol{\theta})^{1/M} \end{aligned}$$

where  $\mathbf{y}_s^{(i)} = \{y_{sk_{i-1}+1}, \dots, y_{sk_i}\}$ ,  $\mathbf{x}_s^{(i)} = \{x_{sk_{i-1}+1}, \dots, x_{sk_i}\}$  and  $Z_{it}$  is the  $i$ -th block normalizing constant.

Using the assumptions (2) and (4) we obtain

$$(7) \quad \begin{aligned} \pi(\boldsymbol{\theta}, \mathbf{x}_{1:t}|\mathbf{y}_{1:t}) &\propto \pi(\boldsymbol{\theta})f(\mathbf{y}_{1:t}|\mathbf{x}_{1:t}, \boldsymbol{\theta})g(\mathbf{x}_{1:t}|\boldsymbol{\theta}) = \\ &= [\pi(\boldsymbol{\theta})^{1/M}]^M \prod_{i=1}^M f(\mathbf{y}_{1:t}^{(i)}|\mathbf{x}_{1:t}^{(i)}, \boldsymbol{\theta})g(\mathbf{x}_{1:t}^{(i)}|\boldsymbol{\theta}) \\ &\propto \prod_{i=1}^M \pi_{it}(\boldsymbol{\theta}, \mathbf{x}_{1:t}^{(i)}|\mathbf{y}_{1:t}^{(i)}) \end{aligned}$$

From (7) we can infer that the type of factorization of the posterior distribution used in [25] holds in this case for every  $t$  since

$$(8) \quad \begin{aligned} \pi(\boldsymbol{\theta}|\mathbf{y}_{1:t}) &\propto \int \pi(\boldsymbol{\theta}, \mathbf{x}_{1:t}|\mathbf{y}_{1:t})d\mathbf{x}_{1:t} \\ &\propto \int \dots \int \prod_{i=1}^M \pi_{it}(\boldsymbol{\theta}, \mathbf{x}_{1:t}^{(i)}|\mathbf{y}_{1:t}^{(i)})d\mathbf{x}_{1:t}^{(1)} \dots d\mathbf{x}_{1:t}^{(M)} \\ &= \prod_{i=1}^M \pi_{it}(\boldsymbol{\theta}|\mathbf{y}_{1:t}^{(i)}). \end{aligned}$$

## 3. EMBARRASSINGLY PARALLEL SMCMC

So far we have discussed the factorization of the posterior distribution based on the panel-type structure of the data. In this section we show how the algorithm handles the serial dependence in the data and samples from  $\{\pi_{it} : 1 \leq t \leq T\}$

for each  $i \in \{1, \dots, M\}$  using a sequential MCMC strategy. For expository purposes we assume that the parameter estimates are updated every time a new observation becomes available, but in the application we will update the estimates after every  $J$ th observation is collected.

Let us define  $\boldsymbol{\lambda}_t = (\boldsymbol{\theta}, \mathbf{x}_{1:t})$ ,  $t \in \mathbb{N}$ , the time sequence of augmented parameter vectors with non-decreasing dimension  $d_t = d_{t-1} + d$ ,  $t \geq 1$ . In order to take advantage of the partition described in the previous section we also introduce  $\boldsymbol{\lambda}_t^{(i)} = (\boldsymbol{\theta}, \mathbf{x}_{1:t}^{(i)})$ , a parameter vector of dimension  $d_t^{(i)}$ . Since the augmented parameter vector can then be partitioned as  $\boldsymbol{\lambda}_t = (\boldsymbol{\lambda}_{t-1}, \mathbf{x}_t)$  and  $\boldsymbol{\lambda}_t^{(i)} = (\boldsymbol{\lambda}_{t-1}^{(i)}, \mathbf{x}_t^{(i)})$  for all  $1 \leq i \leq M$ , the correctness of the algorithms relies on the compatibility condition for the priors on  $\boldsymbol{\lambda}_t^{(i)}$  at each time  $t$  and for each  $i$ th block of data. Specifically, we assume that if the prior is  $p_t^{(i)}(\boldsymbol{\lambda}_t^{(i)}) = p(\boldsymbol{\theta})p(\mathbf{x}_{1:t}^{(i)}|\boldsymbol{\theta})$ , where  $p(\mathbf{x}_{1:t}^{(i)}|\boldsymbol{\theta}) = \prod_{t=1}^T g(\mathbf{x}_t^{(i)}|\mathbf{x}_{t-1}^{(i)}, \boldsymbol{\theta})$ , then it satisfies the compatibility condition

$$(9) \quad p_t^{(i)}(\boldsymbol{\lambda}_t^{(i)}) = \int p_{t+1}^{(i)}(\boldsymbol{\lambda}_t^{(i)}, \mathbf{x}_{t+1}^{(i)}) d\mathbf{x}_{t+1}^{(i)},$$

for all  $i = 1, \dots, M$  and  $t = 1, \dots, T-1$ .

Our embarrassing SMCMC algorithm iterates over time and data blocks. At each time  $t = 1, \dots, T$ , the algorithm consists of two steps. In the first step, for each data block  $\mathbf{y}_{1:t}^{(i)}$ ,  $i = 1, \dots, M$  we use  $L$  parallel SMCMC chains, each of which yields  $n_t$  samples from  $\pi_{it}(\boldsymbol{\lambda})$ , i.e. we generate  $\boldsymbol{\lambda}_{it}^{(l,j)}$ ,  $j = 1, \dots, n_t$ , and  $l = 1, \dots, L$  from  $\pi_{it}(\boldsymbol{\lambda})$ . Based on all samples  $\boldsymbol{\theta}_{it}^{(l,j)}$ ,  $j = 1, \dots, n_t$ , and  $l = 1, \dots, L$  we produce the kernel density estimates  $\hat{\pi}_{it}(\boldsymbol{\theta})$  of the marginal sub-posteriors  $\pi_{it}(\boldsymbol{\theta})$ .

In the second step we take advantage of the factorization (7) and the posterior  $\pi(\boldsymbol{\theta}|\mathbf{y}_{1:T})$  is approximated by combining the approximate sub-posteriors  $\hat{\pi}_{iT}(\boldsymbol{\theta})$ ,  $i = 1, \dots, M$ . Samples from this distribution can be obtained by applying the asymptotically exact posterior sampling procedure detailed in Algorithm 1 of [25]. It is worth noting that (7) holds for any  $t = 1, \dots, T$  so, if needed, one can approximate  $\pi(\boldsymbol{\theta}|\mathbf{y}_{1:t})$  at intermediate times  $t \in \{1, \dots, T\}$  by combining the  $\hat{\pi}_{it}$ 's. However,  $\pi(\boldsymbol{\theta}|\mathbf{y}_{1:t})$  is not used directly in the final posterior  $\pi(\boldsymbol{\theta}|\mathbf{y}_{1:T})$  so if one is interested only in the latter then the merging of partial posteriors is performed only at time  $T$ .

The pseudocode of the proposed EP-SMCMC is given in Algorithm 1 and the details of the SMCMC and of the merge step are detailed in the following sections.

### 3.1 Sequential MCMC

In this section we discuss the construction of the SMCMC samplers that are used in the first step of the EP-SMCMC. To simplify the notation we drop the index  $i$  indicative of the data block. In the SMCMC algorithm a population of  $L$  parallel inhomogeneous Markov chains are used to generate the samples  $\boldsymbol{\lambda}_t^{(l,j)}$  with  $j = 1, \dots, n_t$ ,  $l = 1, \dots, L$  and

---

### Algorithm 1 Embarrassingly Parallel SMCMC (EP-SMCMC)

---

For  $t = 1, 2, \dots, T$

1. For  $i = 1, \dots, M$  draw  $\boldsymbol{\lambda}_{it}^{(l,j)}$ ,  $j = 1, \dots, n_t$ , and  $l = 1, \dots, L$  from  $\pi_{it}(\boldsymbol{\lambda})$  by using the SMCMC transition.
  2. When needed (usually at time  $t = T$ ):
    - (a) Compute the kernel density estimate  $\hat{\pi}_{it}(\boldsymbol{\theta})$  of the marginal sub-posteriors  $\pi_{it}(\boldsymbol{\theta})$  by using the samples  $\boldsymbol{\theta}_{it}^{(l,j)}$ ,  $j = 1, \dots, n_t$ , and  $l = 1, \dots, L$ .
    - (b) approximate the posterior  $\pi(\boldsymbol{\theta}|\mathbf{y}_{1:t})$  by combining the approximate sub-posteriors  $\hat{\pi}_{it}(\boldsymbol{\theta})$ ,  $i = 1, \dots, M$ .
- 

$t = 1, \dots, T$  from the sequence of posterior distributions  $\pi_t$ ,  $t = 1, \dots, T$ . Each Markov chain of the population is defined by a sequence of transition kernels  $K_t(\boldsymbol{\lambda}, A)$ ,  $t \in \mathbb{N}$ , that are operators from  $(\mathbb{R}_{t-1}^d, \mathcal{B}(\mathbb{R}_{t-1}^d))$  to  $(\mathbb{R}_t^d, \mathcal{B}(\mathbb{R}_t^d))$ , such that  $K_t(\boldsymbol{\lambda}, \cdot)$  is a probability measure for all  $\boldsymbol{\lambda} \in \mathbb{R}^{d_{t-1}}$ , and  $K_t(\cdot, A)$  is measurable for all  $A \in \mathcal{B}(\mathbb{R}^{d_t})$ .

The kernel  $K_t(\boldsymbol{\lambda}, A)$  has  $\pi_t$  as stationary distribution and results from the composition of a jumping kernel,  $J_t$  and a transition kernel,  $T_t$ , that is

$$K_t(\boldsymbol{\lambda}, A) = J_t \circ T_t^{n_t}(\boldsymbol{\lambda}, A) = \int_{\mathbb{R}^{d_t}} J_t(\boldsymbol{\lambda}, d\boldsymbol{\lambda}') T_t^{n_t}(\boldsymbol{\lambda}', A)$$

where the fixed dimension transition is defined as

$$T_t^{n_t}(\boldsymbol{\lambda}, A) = T_t \circ T_t^{n_t-1}(\boldsymbol{\lambda}, A) = \int_{\mathbb{R}^{d_t}} T_t(\boldsymbol{\lambda}, d\boldsymbol{\lambda}') T_t^{n_t-1}(\boldsymbol{\lambda}', A)$$

with  $n_t \in \mathbb{N}$ , and  $T^0 = Id$  is the identity kernel. We assume that the jumping kernel satisfies

$$J_{t+1}(\boldsymbol{\lambda}_t, \boldsymbol{\lambda}_{t+1}) = J_{t+1}(\boldsymbol{\lambda}_t, \mathbf{x}_{t+1}) \delta_{\boldsymbol{\lambda}_t}(\tilde{\boldsymbol{\lambda}}_t),$$

where  $\boldsymbol{\lambda}_{t+1} = (\tilde{\boldsymbol{\lambda}}_t, \mathbf{x}_{t+1})$  and  $J_{t+1}(\boldsymbol{\lambda}_t, \mathbf{x}_{t+1}) = J_{t+1}(\boldsymbol{\lambda}_t, (\boldsymbol{\lambda}_t, \mathbf{x}_{t+1}))$ . This condition ensures that the error propagation through the jumping kernel can be controlled over the SMCMC iterations.

In order to apply the SMCMC one need to specify the transition kernel  $T_{t+1}$  and the jumping kernel  $J_{t+1}$  at the iteration  $t+1$ . The transition kernel  $T_t$  at the iteration  $t$  allows each parallel chain to explore the sample space of a given dimension,  $d_t$ , and to generate samples  $\boldsymbol{\lambda}_t^{(l,j)}$ , from the posterior distribution  $\pi_t$ . The jumping kernel  $J_{t+1}$  at the iteration  $t+1$ , allows the chains to go from a space of dimension  $d_t$  to one of dimension  $d_{t+1}$ .

### 3.2 Merge step

The merge step relies on the following approximation of the posterior distribution

$$(10) \quad \pi_t(\boldsymbol{\theta}) = \prod_{i=1}^M \hat{\pi}_{it}(\boldsymbol{\theta})$$

where

$$\begin{aligned} \hat{\pi}_{it}(\boldsymbol{\theta}) &= \frac{1}{N_t} \sum_{l=1}^L \sum_{j=1}^{n_t} \frac{1}{h^p} K \left( \frac{\|\boldsymbol{\theta} - \boldsymbol{\theta}_{it}^{(l,j)}\|}{h} \right) \\ (11) \quad &= \frac{1}{N_t} \sum_{j_i=1}^{N_t} \frac{1}{h^p} K \left( \frac{\|\boldsymbol{\theta} - \boldsymbol{\theta}_{it}^{(j_i)}\|}{h} \right) \end{aligned}$$

where  $\boldsymbol{\theta}_{it}^{(k)} = \boldsymbol{\theta}_{it}^{(l,j)}$  with  $l = k \operatorname{div} L + 1$  and  $j = k \operatorname{mod} L$ ,  $N_t = n_t L$ , and  $K$  is a Gaussian kernel with bandwidth parameter  $h$ . Following the embarrassing MCMC approach the posterior distribution can be written as

$$(12) \quad \pi_t(\boldsymbol{\theta}) \propto \sum_{j_1=1}^{N_t} \dots \sum_{j_m=1}^{N_t} w_j \cdot \frac{1}{h^p} K \left( \frac{\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}_j}{\bar{h}} \right),$$

where  $\bar{h} = h/\sqrt{M}$ , and

$$\bar{\boldsymbol{\theta}}_j = \frac{1}{M} \sum_{i=1}^M \boldsymbol{\theta}_{it}^{(j_i)}, \quad w_j = \prod_{i=1}^M \frac{1}{h^p} K \left( \frac{\|\boldsymbol{\theta}_{it}^{(j_i)} - \bar{\boldsymbol{\theta}}_j\|}{h} \right).$$

### 3.3 Parameter tuning

As suggested by [18], the number of iterations of the Sequential MCMC sampler at each time  $n_t$  are chosen accordingly to the correlation across the parallel chains. We let the number of iterations at the iteration  $t$ ,  $n_t$ , be the smallest integer  $s$  such that  $r_t(s) \leq 1 - \epsilon$ , where  $r_t(s)$  is the rate function associated with the transition kernel  $T_t$  and  $\epsilon$  is a given threshold level. An upper bound for the rate function is provided by chain autocorrelation function at the  $s$ -th lag. It can be estimated sequentially using the output of all the parallel chains:  $\hat{r}_t(s) = \max\{\hat{r}(s, j), j = 1, \dots, (n_t + p)\}$  where

$$\begin{aligned} \hat{r}(s, j) &= \left[ \sum_{l=1}^L \left( \lambda_j^{(s+1,t,l)} - \bar{\lambda}_j^{(s+1,t)} \right) \left( \lambda_j^{(1,t,l)} - \bar{\lambda}_j^{(1,t)} \right) \right] : \\ &: \left\{ \left[ \sum_{l=1}^L \left( \lambda_j^{(s+1,t,l)} - \bar{\lambda}_j^{(s+1,t)} \right)^2 \right]^{\frac{1}{2}} \times \right. \\ (13) \quad &\times \left. \left[ \sum_{l=1}^L \left( \lambda_j^{(1,t,l)} - \bar{\lambda}_j^{(1,t)} \right)^2 \right]^{\frac{1}{2}} \right\}, \end{aligned}$$

with  $\lambda_j^{(s,t,l)}$  the  $j$ -th element of the vector  $\boldsymbol{\lambda}^{(s,t,l)}$  of the parameters  $\boldsymbol{\gamma}^{(l)}$  and the latent states generated up to time  $t$  by the  $l$ -th chain at the  $s$ -th iteration.  $\bar{\lambda}_j^{(s,t)} = L^{-1} \sum_{l=1}^L \lambda_j^{(s,t,l)}$  is the average of the draws over the  $L$  parallel chains.

## 4. NUMERICAL EXPERIMENTS

### 4.1 Simulated data

We consider a time series model in which  $y_{it}$  is the realized variance for the  $i$ -th log-return series at time  $t$ , and  $x_{it}$




Figure 1. Samples  $y_{it}$  (in log-scale),  $i = 1, \dots, m$ ,  $t = 1, \dots, T$ , simulated from the state-space model in Eq. 14–15, with  $T = 1,000$ ,  $m = 1,000$  and parameter setting  $\lambda = 0.7$ ,  $\kappa = 3.8$  and  $\nu = 10$ .

is the latent stochastic volatility process. In order to capture the time variations in the volatility of the series, we consider the exponential family state space model for positive observations recently proposed in [32] and extend it to the context of panel data. The model can be mathematically described then as:

$$(14) \quad y_{it}|x_{it} \sim \mathcal{G}a(\kappa/2, \kappa x_{it}/2)$$

$$(15) \quad x_{it}|x_{it-1} \sim x_{it-1} \psi_{it}/\lambda, \quad \psi_{it} \sim \mathcal{B}e(\nu/2, \kappa/2)$$

for  $t = 1, \dots, T$ , where  $\mathcal{G}a(a, b)$  denotes the gamma distribution and  $\mathcal{B}e(a, b)$  the beta distribution of the first type.

We generate 1,000 time series of 1,000 observations each, and obtain a dataset of 1 million observations. In Fig. 1 we illustrate such a simulated dataset. Inference for a nonlinear latent variable model with this large a sample via MCMC sampling can be computationally challenging. In the simulation we set  $\lambda = 0.7$ ,  $\kappa = 3.8$  and  $\nu = 10$ , with initial condition  $x_{0i} = 10, \forall i$ . For each series we have generated 3,000 realizations, but the initial 2,000 samples were discarded so that dependence on initial conditions can be considered negligible.

We aim to estimate the common parameters  $\lambda$ ,  $n$  and  $\kappa$  and assume a uniform prior distribution for  $\lambda$ , i.e.  $\lambda \sim \mathcal{U}[0, 1]$  and proper vague prior distributions for  $\nu$  and  $\kappa$ , that is  $\nu \sim \mathcal{G}a(0.5, 0.5)$  and  $\kappa \sim \mathcal{G}a(0.5, 0.5)$  truncated on the region  $\{(\nu, \kappa) : \nu > \kappa - 1\}$ .

In order to apply the EP-SMCMC algorithm we assume that the  $m$  series are split into  $M$  blocks of  $K = m/M$  series each. The updates are performed every  $J$  observations, so the total number of updates is  $n = T/J$ . Let us denote with the column vector  $\mathbf{u}_{1:t} = (u_s, \dots, u_t)'$  a collection of variables  $u_r$  with  $r = s, \dots, t$ . We define the  $i$ -th block of series and the  $i$ -th block of latent variables as the  $(t \times K)$ -matrices  $Y_{it} = (\mathbf{y}_{(i-1)K+1,1:t}, \dots, \mathbf{y}_{iK,1:t})$  and  $X_{it} = (\mathbf{x}_{(i-1)K+1,1:t}, \dots, \mathbf{x}_{iK,1:t})$ , respectively, with  $\mathbf{y}_{j,1:t} = (y_{j1}, \dots, y_{jt})'$ , and  $\mathbf{x}_{j,1:t} = (x_{j1}, \dots, x_{jt})'$ . Then, the complete-data likelihood function at time  $t$  for the  $i$ -th block is

(16)

$$L(Y_{it}, X_{it}|\boldsymbol{\theta}) = \prod_{j=k_{i-1}+1}^{k_i} \prod_{s=1}^t \frac{1}{\Gamma(\frac{\kappa}{2})} \left(\frac{\kappa x_{js}}{2}\right)^{\kappa/2} y_{js}^{\frac{\kappa}{2}-1}$$

(17)

$$\exp\left(-\frac{\kappa x_{js}}{2} y_{js}\right) \left(\frac{\lambda x_{js}}{x_{js-1}}\right)^{\frac{\kappa}{2}-1} \left(1 - \frac{\lambda x_{js}}{x_{js-1}}\right)^{\frac{\kappa}{2}-1} \frac{\lambda}{x_{js-1}}$$

where  $\boldsymbol{\theta} = (\lambda, \kappa, \nu)$ . Then the sub-sample posterior, based on the complete-data likelihood function of the  $i$ -th block is

$$\begin{aligned} \pi(\boldsymbol{\theta}, X_{it}|Y_{it}) &\propto L(Y_{it}, X_{it}|\boldsymbol{\theta})(\nu + \kappa)^{-\frac{M}{2}} \times \\ &\times \exp\left(-\frac{M(\nu + \kappa)}{2}\right) \mathbb{I}(\nu > \kappa - 1) \end{aligned}$$

At time  $t$ , and for the  $i$ -th block, the  $l$ -th SMCMC parallel chain has the transition kernel of a Gibbs sampler which iterates over the following steps:

- 1.1 generate  $\boldsymbol{\theta}^{(l,j)}$  from  $f(\boldsymbol{\theta}|Y_{it}, X_{it}^{(l,j-1)})$
- 1.2 generate  $X_{it}^{(l,j)}$  from  $f(X_{it}|\boldsymbol{\theta}^{(l,j)}, Y_{it})$

with  $j = 1, \dots, n_t$ , and  $X_{it}^{(l,0)} = ((X_{it-1}^{(l,n_t-1)})', (x_{k_{i-1}+1,t}^{(l,1)}, \dots, x_{k_i,t}^{(l,1)})')'$  is a  $(t \times K)$ -dim matrix where the  $t$ -th row elements drawn from the jumping kernel at time  $t-1$ . At time  $t+1$ , as a new observation become available, the dimension of the state space for the  $i$ -th block SMCMC chains increase from  $d_t$  to  $d_{t+1} = d_t + J$ . We choose as jumping kernel of the  $l$ -th parallel chain to be the transition kernel of a Gibbs sampler with the following full conditional distribution

2.  $x_{kt+1}^{(l,1)} \sim f(x_{kt+1}|X_{it}^{(l,j)}, X_{it}^{(l,j)}, \boldsymbol{\theta}^{(l,j)})$ , with  $k = k_{i-1} + 1, \dots, k_i$

where  $j = n_t$ . The details of the sampling procedures for the three steps are given in Appendix A.

In the simulation example we compare our EP-SMCMC with a MCMC repeated sequentially over time and a Sequential Monte Carlo (SMC) with unknown parameters. For the EP-SMCMC we used  $L = 10$  parallel chains and a number of iterations  $n_t$  close to 50, on average. For the MCMC we use a multi-move Gibbs sampler where the latent states are sampled in one block by applying a forward-filtering backward sampling (FFBS) procedure. The filtering and smoothing relationships are given in Appendix C. In order to update the parameters we have used a Metropolis-Hastings step. The MCMC chain of our multi-move sampler is mixing quite well due to the FFBS, and also the Metropolis step has acceptance rates about 0.3, which is a good rate for many models as argued in [26].

In our MCMC analysis we considered two cases. The first one is based on samples of 1,000 iterations after a burn-in phase of 500 iterations in order to have a computational complexity similar to the one of the EP-SMCMC. The second one is based on samples of 25,000 iterations after a

burn-in phase of 20,000 iterations and is used to have reliable MCMC estimates of the parameter posterior distribution based on the whole sample.

For the SMC we also consider the regularized auxiliary particle filter (RAPF) combined with the embarrassingly parallel algorithm to obtain a EP-RAPF. In Appendix B we present the computational details of running the r-APF for our stochastic volatility model. We refer to [24] for the definition of regularized particle filter and to [10] for a comparison of different regularized filters for stochastic volatility models.

When we compare the EP-SMCMC, MCMC and EP-RAPF algorithms, we use as a measure of efficiency the mean square errors for the parameters  $\kappa$ ,  $\nu$  and  $\lambda$  after the last time-block of data has been processed. The mean square errors have been estimated using 40 independent runs of each algorithm. The same set of data has been used to reduce the standard error of the MSE estimates. The experiments have been conducted on a cluster multiprocessor system with 4 nodes; each node comprises of four Xeon E5-4610 v2 2.3GHz CPUs, with 8 cores, 256GB ECC PC3-12800R RAM, Ethernet 10Gbit, 20TB hard disk system with Linux. The algorithms have been implemented in Matlab (see [30]) and the parallel computing makes use of the Matlab parallel computing toolbox.

The structure of the EP-SMCMC sampler allows for sequential data acquisition and for parallel estimation based on different cross-sectional blocks. Thus, the MSE obtained after processing the last block is given in Table 1 for different dimensions of the time blocks  $J$  (rows) and of the cross section blocks  $M$  (columns). Fig. 2 shows the posterior approximation obtained from one run of the EP-SMCMC on the dataset shown in Fig. 1. From our experiments we find that the parameters that are most difficult to estimate are  $\kappa$  and  $\nu$ , whereas  $\lambda$  has lower MSEs regardless of the choice of block size and type of algorithm. For the EP-SMCMC sampler a larger size  $J$  of the time blocks reduces the MSE, possibly due to a reduction in the propagation of the approximation error over the iterations. The behaviour of the MSE with respect to the cross-sectional block size  $K$  is not monotonic. The MSE initially decreases with  $K$ , but for larger  $K$  it increases. From our experiments, values of  $K$  between 20 and 40 yield the best results. The  $\kappa$  and  $\nu$  MSEs for the MCMC are higher than their EP-SMCMC counterparts, likely due to the lack of convergence of the MCMC in the 1,500 iterations per time block. The large MSEs for the EP-RAPF are due to artificial noise introduced by the regularization step for the parameters, and also to the hidden state estimation. In the EP-RAPF implementation we have used 1,500 particles and the artificial noise is necessary to avoid degeneration of their weights. Note that within each time block the EP-RAPF is using the filtered states instead of the smoothed states. A combination of the RAPF with a MH step or a particle MCMC would improve the performance of the EP-RAPF algorithm in the estimation of both

Table 1. Mean square error for the parameters  $\kappa$ ,  $\nu$  and  $\lambda$ , for different size of the time blocks,  $J$ , and of the cross-sectional blocks,  $K$ , for the EP-SMCMC and a sequence of MCMC (1,500 iterations) started each time a block of observations is acquired and EP-RAPF with 1,500 particles. The average standard deviation, across algorithms and experiments, of the  $\kappa$ ,  $\nu$  and  $\lambda$  MSE estimates is 0.3, 1.13 and 0.00001, respectively

Parameter $\kappa$ MSE							
$J$	EP-SMCMC					MCMC	EP-RAPF
	( $K = 10$ )	( $K = 20$ )	( $K = 30$ )	( $K = 40$ )	( $K = 50$ )		
50	3.91	2.37	2.07	4.81	7.02	29.68	1.14
100	2.09	1.68	2.11	3.06	6.75	32.33	0.90
200	3.07	1.68	1.20	0.77	3.98	25.06	6.90
Parameter $\nu$ MSE							
$m$	EP-SMCMC					MCMC	EP-RAPF
	( $K = 10$ )	( $K = 20$ )	( $K = 30$ )	( $K = 40$ )	( $K = 50$ )		
50	45.04	48.52	39.73	53.55	51.90	78.20	31.18
100	28.41	29.14	20.52	39.01	43.52	56.19	101.14
200	29.83	21.01	11.92	28.54	39.07	58.23	53.12
Parameter $\lambda$ MSE							
$m$	EP-SMCMC					MCMC	EP-RAPF
	( $K = 10$ )	( $K = 20$ )	( $K = 30$ )	( $K = 40$ )	( $K = 50$ )		
50	0.02	0.01	0.01	0.01	0.01	0.01	0.02
100	0.02	0.01	0.01	0.01	0.01	0.01	0.01
200	0.01	0.01	0.01	0.01	0.01	0.01	0.01

Table 2. Computing time, in hours, for fixed cross-sectional block size  $K = 40$ , for different size of the time blocks  $J$  and different number of parallel CPU cores  $r$ , for the EP-SMCMC and a sequence of MCMC (1,500 iterations) started each time a block of observations is acquired

$J$	EP-SMCMC					MCMC
	( $r = 10$ )	( $r = 20$ )	( $r = 30$ )	( $r = 40$ )	( $r = 50$ )	
50	10.87	6.66	2.66	1.64	1.01	49.39
100	5.13	2.45	1.88	0.70	0.47	28.10
200	2.57	1.43	0.71	0.32	0.19	15.27

states and parameters, at a price of increasing computing. We leave the further study of this issue for a future communication.

Also, we compare EP-SMCMC, MCMC and EP-RAPF in terms of computing time. For the EP-SMCMC we consider cross-sectional blocks of size  $K = 30$ . The computing times are given in Table 2 for different time acquisition rates  $J$  (rows) and different CPUs  $r$  working in parallel (columns). We conclude that the EP-SMCMC implemented on cluster multiprocessor system can be up to 80 times faster than the standard MCMC. These results are in line with the ones obtained in previous studies on parallel Monte Carlo methods ([11], [23], [19]). The calculation in all our experiments have been carried out using double precision. If an application allows for a lower degree of precision, then single precision calculation can lead to large gains in computing time as documented by [23] in the context of GPU parallel computing.

## 4.2 Real data

We consider a panel of 12,933 assets of the US stock market and collect prices at a daily frequency from 29 December 2000 to 22 August 2014, which yields 3,562 time observations. Then we compute logarithmic returns for all stocks and obtain a dataset of 46,067,346 observations.

In order to control for the liquidity of the assets and consequently for long sequences of zero returns, we impose that each stock has been traded a number of days corresponding to at least 40% of the sample size. Also, we focus on the last part of the sample which includes observations from 8 February 2013 to 22 August 2014. After cleaning the dataset we obtain 6,799 series and 400 time observations, and the size of the dataset reduces to 2,846,038. The original and the cleaned datasets are given in Fig. 3.

In order to capture the time variations in the volatility of the series we consider the panel state space model presented in the previous section. We apply our EP-SMCMC sampler to the panel of 6,799 time series. The data are acquired sequentially over time in blocks of 100 time observations and each panel consists of  $K = 40$  series. At each point in time, for each parameter  $k$ ,  $\nu$  and  $\lambda$  we obtain 500 posterior densities from each parallel SMCMC chain (see C.1–C.3 in Appendix C an example for  $t = 100, 200, 300, 400$ ).

Figure 4 shows the sequential posterior inference after the merge step of the embarrassingly parallel SMCMC algorithm is applied to the output of the SMCMC samples. At each point in time we obtain an approximated posterior density for the whole cross-section from the embarrassingly parallel step (see solid lines in Figure 4 for  $t = 100, 200, 300, 400$ ).

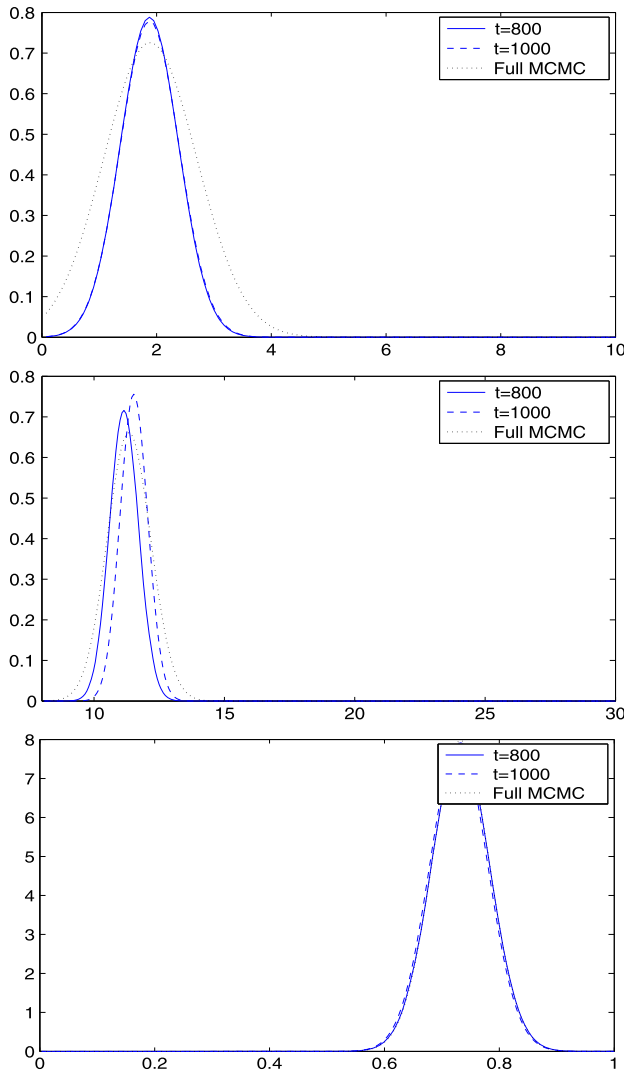


Figure 2. Estimation of the posterior densities of the parameters  $\kappa$  (top panel),  $\nu$  (middle panel) and  $\lambda$  (bottom panel). The EP-SMCMC estimates are obtained at iterations 800 (solid line) and 1000 (dashed line) when  $K = 30$  and  $J = 200$ . The full posterior densities estimates are obtained from a MCMC with 25,000 iterations (dotted line).

The approximation of the posterior produced by the EP-SMCMC is close to the approximation based on a MCMC analysis of the full posterior. This can be seen in Figure 4 where the solid line shows the EP-SMCMC approximation at the last update ( $t = 400$ ) and the dashed line represents the full-sample estimate based on a standard MCMC with 25,000 iterations after a burn-in period of 20,000 iterations. The posterior mean and posterior quantiles approximated by EP-SMCMC are given in Tab. 3. In order to approximate the high posterior density (HPD) intervals and the posterior mean we generate samples from the posterior distribution given in Eq. 8 by applying the independent Metropolis within Gibbs algorithm given in [25].




Figure 3. Quantiles at the 5% and 90% (gray area) and mean (solid line) of the cross-sectional daily log-return distribution. Returns for all the 12,933 assets (top panel) of the US stock market, from 29 December 2000 to 22 August 2014 and for a subsample of 6,799 assets (bottom panel) from 8 February to 22 August 2014.

Table 3. EP-SMCMC and MCMC approximation of the posterior mean ( $\hat{\theta}$ ) and the 95% high probability density region HPD given by the 2.5% and 97.5% inter-quantile inveral ( $q_{0.025}, q_{0.975}$ )

$\theta$	EP-SMCMC		MCMC	
	$\hat{\theta}$	HPD	$\hat{\theta}$	HPD
$\kappa$	0.66	(0.61,0.71)	0.54	(0.31,0.58)
$\nu$	0.87	(0.79,0.95)	0.91	(0.86,1.01)
$\lambda$	0.98	(0.97,0.99)	0.96	(0.94,0.97)

## 5. CONCLUSION

We propose a new MCMC algorithm which combines embarrassingly parallel MCMC and Sequential MCMC. The algorithm is developed for data that exhibit dependent patterns, in particular for large sets of time series for which a standard MCMC-based analysis would be very slow. Here we take advantage of the independence between the unit-specific observations and latent variables of the panel to partition the data and factor the full posterior in a product of partial posteriors. In the absence of clear independent panel units, an interesting and difficult question concerns alternative strategies to divide the data and combine the partial posterior samples.

It is apparent that the development of novel MCMC algorithms for big data is evolving rapidly. While “divide and

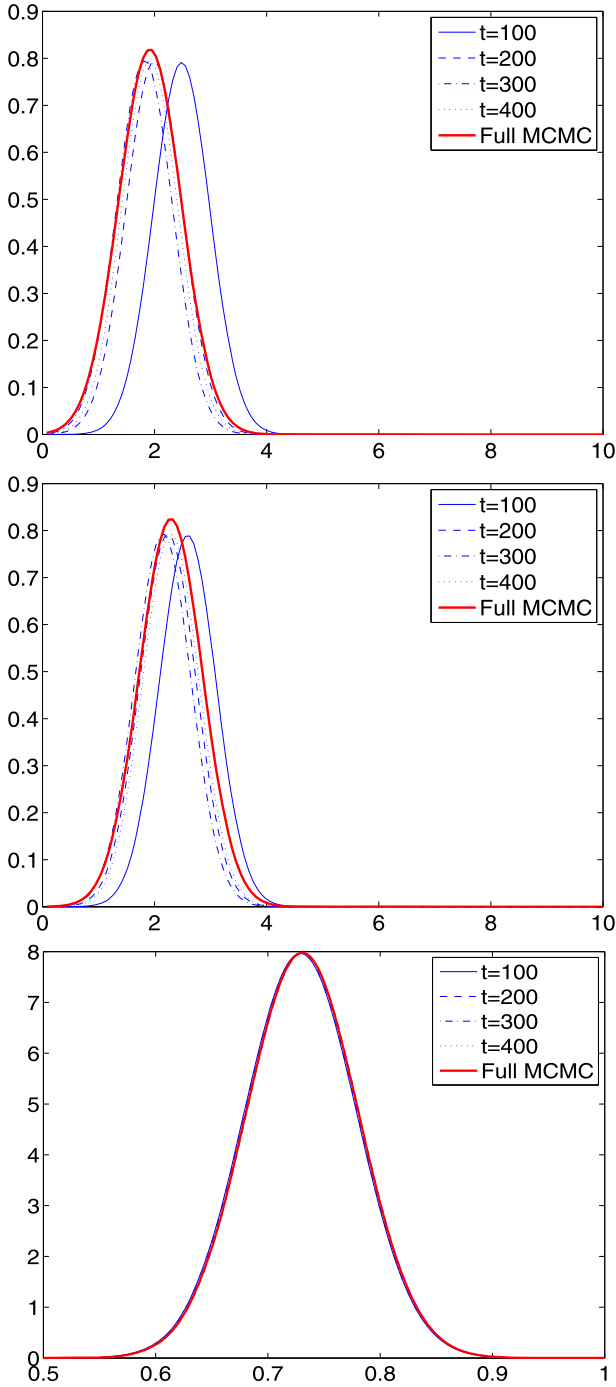


Figure 4. Sequential estimation of the posterior densities of the parameter  $k$  (top panel),  $\nu$  (middle panel) and  $\lambda$  (bottom panel) at different points in time  $t$ . The whole sample and pooled data posterior densities of the parameter are given by the dashed line.

conquer” strategies continue to develop, one must devise techniques to handle the additional approximations that are introduced by the current existing methods, including EP-SMCMC. Quantifying and controlling the error introduced

by these approximations remains central to the success of MCMC for Big Data.

## ACKNOWLEDGEMENT

We thank the Editor, the Associate Editor and three referees for comments and suggestions that have greatly improved the paper. This research used the SCSCF multiprocessor cluster system at University Ca’ Foscari of Venice. Roberto Casarin’s research is supported by funding from the European Union, Seventh Framework Programme FP7/2007-2013 under grant agreement SYRTO-SSH-2012-320270, by the Institut Europlace of Finance, “Systemic Risk grant”, the Global Risk Institute in Financial Services, the Louis Bachelier Institute, “Systemic Risk Research Initiative”, and by the Italian Ministry of Education, University and Research (MIUR) PRIN 2010-11 grant MISURA. Radu Craiu’s research is supported by the Natural Sciences and Engineering Research Council of Canada. Fabrizio Leisen’s research is supported by the European Community’s Seventh Framework Programme FP7/2007-2013 under grant agreement number 630677.

## APPENDIX A. COMPUTATIONAL DETAILS

### A.1 Transition kernel

As regards the Step 1.1. of the transition kernel, the distribution  $\pi(\boldsymbol{\theta}|X_{it}^{(l,j)}, Y_{it}) \propto \pi(\boldsymbol{\theta}, X_{it}^{l,j}|Y_{it})$  is not tractable and we applied a Metropolis-Hastings. At the  $j$ -th iteration of the  $l$ -th SMCMC chain we generate the MH proposal from a Gaussian random walk on the transformed parameter space  $\tilde{\theta}_1 = \log(\kappa)$ ,  $\tilde{\theta}_2 = \log(\nu)$  and  $\tilde{\theta}_3 = -\log[(1-\lambda)/\lambda]$ .

In the Step 1.2 of the transition kernel, we exploit the tractability of the state space model and apply a multi-move Gibbs sampler, where the hidden states  $\mathbf{x}_{i1:t}$  are updated in one step. By applying Proposition 1 in [32] with  $m = 1$ , one gets the following filtered, and prediction distributions

$$(A.1) \quad x_{it}|\boldsymbol{\theta}, Y_{it} \sim \mathcal{G}a((\kappa + k)/2, \kappa\sigma_{it}^2/2)$$

$$(A.2) \quad x_{it+1}|\boldsymbol{\theta}, Y_{it} \sim \mathcal{G}a(\kappa/2, \kappa\sigma_{it}^2/2/\lambda)$$

where  $\sigma_{it}^2 = y_{it} + \lambda\sigma_{it-1}^2$ , and the backward smoothed distribution

$$(A.3) \quad x_{it}|\boldsymbol{\theta}, x_{it+1}, Y_{it} \sim \lambda x_{it+1} + z_{it+1}, \quad z_{it+1} \sim \mathcal{G}a(\kappa/2, (\kappa\sigma_{it}^2)/2).$$

which is used to generate  $X_{it}^{(l,j)}$  given  $\boldsymbol{\theta}^{(l,j)}$  and  $Y_{it}$ .

### A.2 Jumping kernel

As regard to the jumping kernel, when  $d_{t+1} = d_t + 1$ , it is given by a Gibbs sampler transition kernel with full conditional distribution

$$(A.4) \quad x_{it+1}|\boldsymbol{\theta}, Y_{i,t+1} \sim \mathcal{G}a((\kappa + k)/2, \kappa\sigma_{it+1}^2/2)$$



## APPENDIX B. COMPUTATIONAL DETAILS FOR THE EP-RAPF

Let us consider the following reparametrization  $\theta = (\log(\kappa), \log(\nu - \kappa + 1), -\log[(1 - \lambda)/\lambda])$ . Given the initial sets of weighted random samples  $\{\mathbf{x}_{it_0}^j, \theta_{it_0}^j, w_{it_0}^j\}_{j=1}^N$ ,  $i = 1, \dots, M$  the regularized APF performs the following steps, for  $t_0 < t \leq T - 1$ ,  $i = 1, \dots, M$  and  $j = 1, \dots, N$ :

(i) Simulate  $r_i^j \sim q(r_i) \propto \sum_{l=1}^N w_{it}^l \delta_l(r_i)$  where

$$w_{it}^l \propto w_{it-1}^l \prod_{k=k_{i-1}+1}^{k_i} \mathcal{G}a(y_{kt+1} | \kappa_{it}^{r_i}/2, \kappa_{it}^{r_i} \mu_{kt+1}^{r_i}/2)$$

with

$$\mu_{kt+1}^r = \frac{x_{kt}^r \nu_{it}^r}{\lambda_{it}^r \nu_{it}^r + \kappa_{it}^r}.$$

(ii) Simulate  $\theta_{it+1}^j \sim \mathcal{N}(a\theta_{it}^j + (1-a)\bar{\theta}_{it}, h^2 V_{it})$  where  $V_{it}$  and  $\bar{\theta}_{it}$  are the empirical variance matrix and the empirical mean respectively and  $a \in [0, 1]$  and  $h^2 = (1 - a^2)$ ,

(iii) Simulate  $x_{kt+1}^j \sim x_{kt}^{r_i^j} \psi_{it}^j / \lambda_{it+1}^j$ , with  $\psi_{kt}^j \sim \mathcal{B}e(\nu_{it}^j/2, \kappa_{it}^j/2)$  for  $k = k_{i-1} + 1, \dots, k_i$

(iv) Update the weights

$$\omega_{it+1}^j \propto \prod_{k=k_{i-1}+1}^{k_i} \frac{\mathcal{G}a(y_{kt+1} | \kappa_{it}^j/2, \kappa_{it}^j x_{kt+1}^j/2)}{\mathcal{G}a(y_{kt+1} | \kappa_{it}^{r_i^j}/2, \kappa_{it}^{r_i^j} \mu_{kt+1}^{r_i^j}/2)}$$

(v) If  $\text{ESS}_{t+1} < \varepsilon$ , simulate  $\{\mathbf{x}_{it+1}^j, \theta_{it+1}^j\}_{j=1}^N$  from  $\{\mathbf{x}_{it+1}^j, \theta_{it+1}^j, \omega_{it+1}^j\}_{j=1}^N$  and set  $w_{it+1}^j = 1/N$ . Otherwise set  $w_{it+1}^j = \omega_{it+1}^j$

where

$$\text{ESS}_t = \frac{N}{1 + N \sum_{i=1}^N \left( \omega_t^i - N^{-1} \sum_{i=1}^N \omega_t^i \right)^2 / \left( \sum_{i=1}^N \omega_t^i \right)^2}.$$

is the effective sample size.

In the merge step of our EP-RAPF, the particle set  $\{\theta_{it}^j, \omega_{it}^j\}_{j=1}^N$  is used to build the following approximation of the posterior distribution

$$(B.1) \quad \pi_t(\theta) = \prod_{i=1}^M \hat{\pi}_{it}(\theta)$$

by applying the embarrassingly parallel algorithm, as in the EP-SMCMC.

## APPENDIX C. SMCMC OUTPUT

Figures C.1–C.3 show an example of sequential SMCMC approximation of the posterior densities of the parameters  $k$ ,  $\nu$  and  $\lambda$  for the different blocks of observations (different lines in each plot) and at different point in time  $t = 100, 200, 300, 400$  (different plots) for our panel of

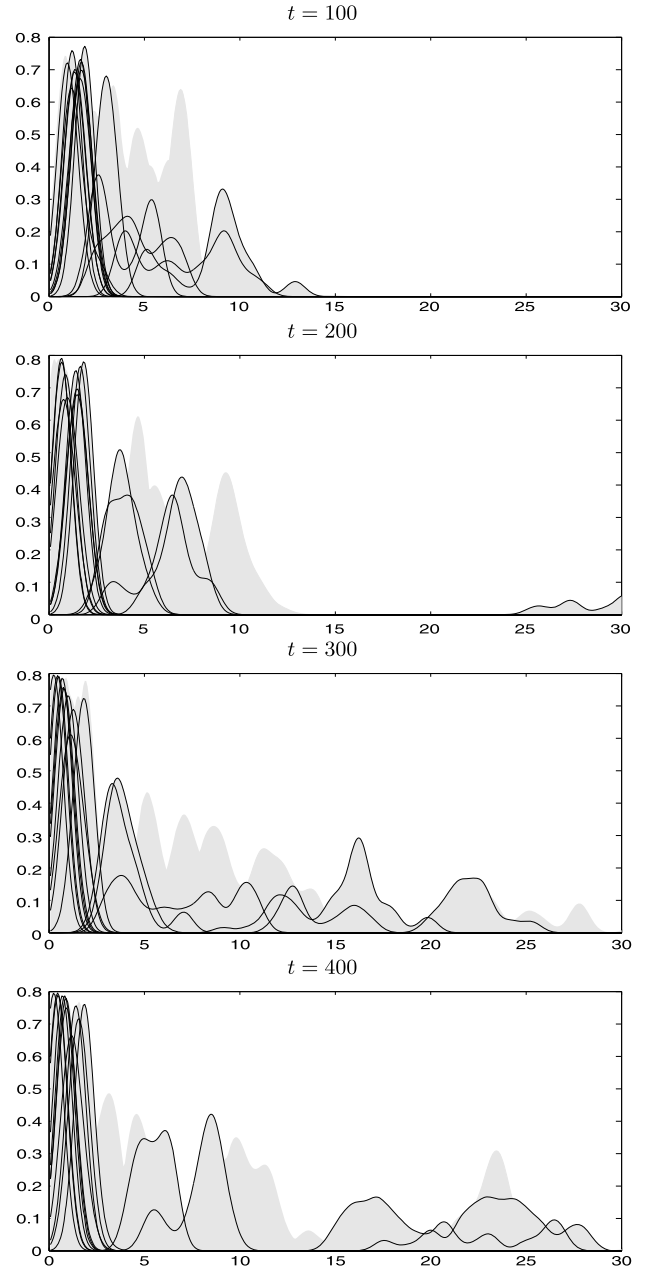


Figure C.1. Sequential estimation of the posterior densities of the parameter  $k$ , for the different blocks of observations (different lines) and at different point in time  $t$  (different plots). For expository purposes the  $\lfloor m/K \rfloor = 169$  lines have been subsampled and the grey area represents the area below the envelope of the  $N$  densities.

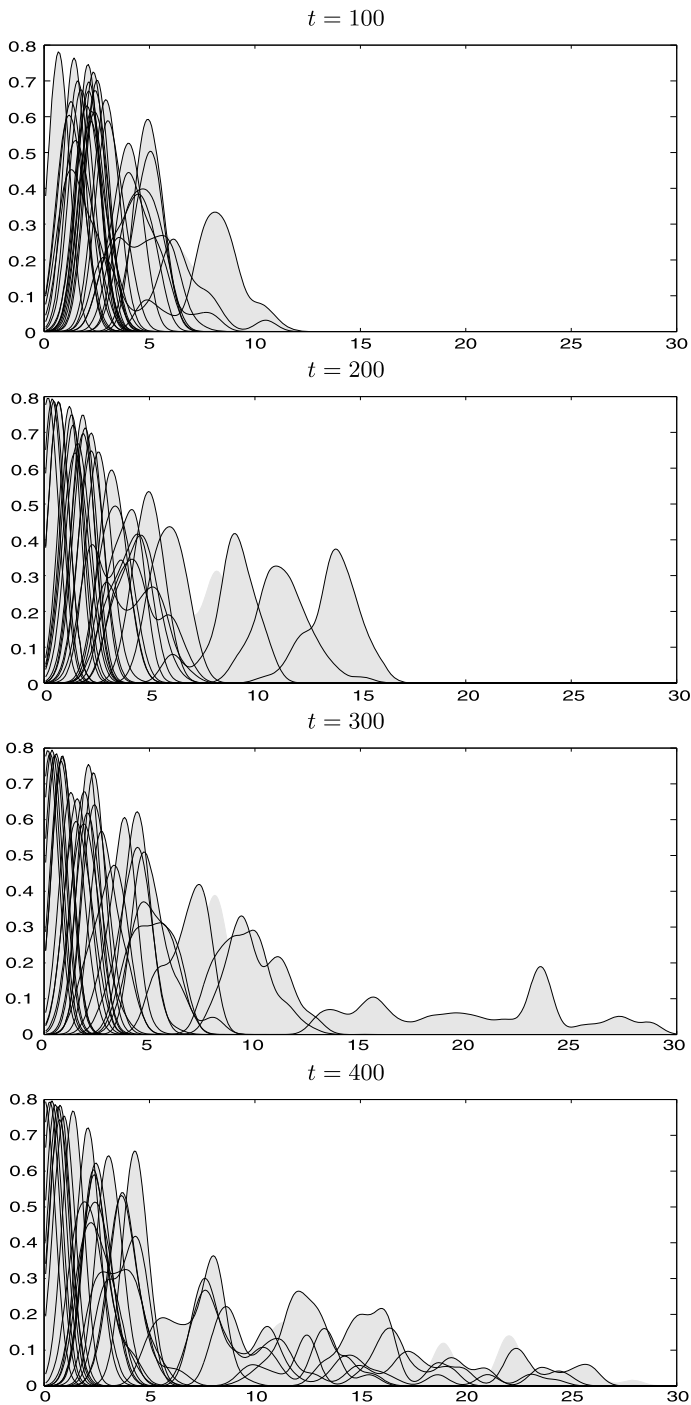


Figure C.2. Sequential estimation of the posterior densities of the parameter  $\nu$ , for the different blocks of observations and at different point in time  $t$ . For expository purposes the  $\lfloor m/K \rfloor = 169$  lines have been subsampled and the grey area represents the area below the envelope of the  $N$  densities.




Figure C.3. Sequential estimation of the posterior densities of the parameter  $\lambda$ , for different blocks of observations (different lines) and at different points in time  $t = 100, 200, 300, 400$ . For expository purposes the  $\lfloor m/K \rfloor = 169$  lines have been subsampled and the grey area represents the area below the envelope of the  $N$  densities.

$m = 6,799$  time series. We consider  $\lfloor m/M \rfloor = 169$  cross-sectional blocks with  $M = 40$  observations each.

Received 21 February 2015

## REFERENCES

- [1] ABREVAYA, J. and SHEN, S. (2014). Estimation of censored panel-data models with slope heterogeneity. *Journal of Applied Econometrics* **29** 523–548. [MR3253200](#)
- [2] ANDERSON, T. W. and HSIAO, C. (1982). Formulation and estimation of dynamic models using panel data. *Journal of Econometrics* **18** 47–82. [MR0661665](#)
- [3] BILLIO, M., CASARIN, R., RAVAZZOLO, R. and VAN DIJK, H. K. (2015). Interactions between eurozone and US booms and busts: A Bayesian panel Markov-switching VAR model. *Journal of Applied Econometrics* forthcoming.
- [4] BLUNDELL, R. and BOND, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics* **87** 115–143.
- [5] BONEVA, L., LINTON, O. and VOGT, M. (2015). A semiparametric model for heterogeneous panel data with fixed effects. *Journal of Econometrics* **188** 327–345. [MR3383213](#)
- [6] CANOVA, F. and CICCARELLI, M. (2013). Panel vector autoregressive models: A survey. In *VAR Models in Macroeconomics: New Developments and Applications: Essays in Honor of Christopher A. Sims* (THOMAS L. K., FOMBY B. and MURPHY A., eds.).
- [7] CAPPE, O., GUILLIN, A., MARIN, J. M. and ROBERT, C. P. (2004). Population Monte Carlo. *J. Comput. Graph. Statist.* **13** 907–929. [MR2109057](#)
- [8] CASARIN, R. (2014). A comment on a tractable state-space model for symmetric positive-definite matrices. *Bayesian Analysis* **9** 793–804. [MR3293954](#)
- [9] CASARIN, R., CRAIU, R. V. and LEISEN, F. (2013). Interacting multiple try algorithms with different proposal distributions. *Statistics and Computing* **23** 185–200. [MR3016937](#)
- [10] CASARIN, R. and MARIN, J. M. (2009). Online data processing: Comparison of Bayesian regularized particle filters. *Electronic Journal of Statistics* **3** 239–258. [MR2495838](#)
- [11] CASARIN, R., GRASSI, S., RAVAZZOLO, F. and VAN DIJK, H. K. (2015). Parallel sequential Monte Carlo for efficient density combination: The DeCo Matlab toolbox. *Journal of Statistical Software* **68** 1–30.
- [12] CRAIU, R. V., ROSENTHAL, J. S. and YANG, C. (2009). Learn from thy neighbor: Parallel-chain adaptive and regional MCMC. *Journal of the American Statistical Association* **104** 1454–1466. [MR2750572](#)
- [13] CREEL, M. (2005). User-friendly parallel computations with econometric examples. *Computational Economics* **26** 107–128.
- [14] CREEL, M. and GOFFE, W. L. (2008). Multi-core CPUs, clusters, and grid computing: A tutorial. *Computational Economics* **32** 353–382.
- [15] DEL MORAL, P. (2004). *Feynman-Kac Formulae. Genealogical and Interacting Particle Systems with Applications*. Springer. [MR2044973](#)
- [16] DEL MORAL, P. and MICLO, L. (2000). Branching and interacting particle systems approximations of Feynman-Kac formulae with applications to non linear filtering. In *Séminaire de Probabilités XXXIV*. Lecture Notes in Mathematics, No. 1729. 1–145. Springer. [MR1768060](#)
- [17] DOORNIK, J. A., HENDRY, D. F. and SHEPHARD, N. (2002). Computationally-intensive econometrics using a distributed matrix-programming language. *Philosophical Transactions of the Royal Society of London, Series A* **360** 1245–1266. [MR1917073](#)
- [18] DUNSON, D. and YANG, Y. (2013). Sequential Markov Chain Monte Carlo. ArXiv: <http://arxiv.org/abs/1308.3861>.
- [19] GEWEKE, J. and DURHAM, G. (2012). Massively Parallel Sequential Monte Carlo for Bayesian Inference Working Papers, National Bureau of Economic Research, Inc.
- [20] HSIAO, C. (2015). Challenges for panel financial analysis. In *Econometrics of Risk, Studies in Computational Intelligence* (V.-N. HUYNH, V. KREINOVICH, S. SRIBOONCHITTA and K. SURIYA, eds.) 3–15.
- [21] KAUFMANN, S. (2010). Dating and forecasting turning points by Bayesian clustering with dynamic structure: A suggestion with an application to Austrian data. *Journal of Applied Econometrics* **25** 309–344. [MR2758637](#)
- [22] KAUFMANN, S. (2015). K-state switching models with time-varying transition distributions – Does loan growth signal stronger effects of variables on inflation? *Journal of Econometrics* **187** 82–94. [MR3347296](#)
- [23] LEE, A., CHRISTOPHER, Y., GILES, M. B., DOUCET, A. and HOLMES, C. C. (2010). On the utility of graphic cards to perform massively parallel simulation with advanced Monte Carlo methods. *Journal of Computational and Graphical Statistics* **19** 769–789.
- [24] LIU, J. S. and WEST, M. (2001). Combined parameter and state estimation in simulation based filtering. In *Sequential Monte Carlo Methods in Practice* (DOUCET A., DE FREITAS N. AND GORDON N., eds.) Springer-Verlag. [MR1847783](#)
- [25] NEISWANGER, W., Y., Y. and XING, E. P. (2014). Asymptotically exact, embarrassingly parallel MCMC. In *Proceedings of the 30th International Conference on Conference on Uncertainty in Artificial Intelligence*.
- [26] ROBERTS, G. O. and ROSENTHAL, J. S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statist. Sci.* **16** 351–367. [MR1888450 \(2003a:65006\)](#)
- [27] SCOTT, S. L., BLOCKER, A. W. and BONASSI, F. V. (2013). Bayes and big data: The consensus Monte Carlo algorithm. In *EFaB-Bayes 250 conference* **16**.
- [28] SUCHARD, M., HOLMES, C. and WEST, M. (2010). Some of the What?, Why?, How?, Who? and Where? of graphics processing unit computing for Bayesian Analysis. *Bulletin of the International Society for Bayesian Analysis* **17** 12–16.
- [29] SWANN, C. A. (2002). Maximum likelihood estimation using parallel computing: An introduction to MPI. *Computational Economics* **19** 145–178.
- [30] THE MATHWORKS, INC. (2011). MATLAB – The Language of Technical Computing, Version R2011b The MathWorks, Inc., Natick, Massachusetts.
- [31] WANG, X. and DUNSON, D. (2013). Parallel MCMC via Weierstrass sampler. ArXiv: <http://arxiv.org/abs/1312.4605>.
- [32] WINDLE, J. and CARVALHO, Y. (2014). A tractable state-space model for symmetric positive-definite matrices. *Bayesian Analysis* **9** 759–792. [MR3293953](#)

Roberto Casarin  
Department of Economics  
University Ca'Foscari of Venice  
Venice  
Italy  
E-mail address: [r.casarin@unive.it](mailto:r.casarin@unive.it)

Radu V. Craiu  
Department of Statistical Sciences  
University of Toronto  
Toronto  
Canada  
E-mail address: [craiu@utstat.toronto.edu](mailto:craiu@utstat.toronto.edu)

Fabrizio Leisen  
School of Mathematics, Statistics and Actuarial Science  
University of Kent  
Canterbury  
United Kingdom  
E-mail address: [F.Leisen@kent.ac.uk](mailto:F.Leisen@kent.ac.uk)