



# Spectral Subsampling MCMC for Stationary Multivariate Time Series with Applications to Vector ARTFIMA Processes

Mattias Villani<sup>a,b,\*</sup>, Matias Quiroz<sup>a,c</sup>, Robert Kohn<sup>d</sup>, Robert Salomone<sup>e</sup>

<sup>a</sup> Department of Statistics, Stockholm University, Stockholm SE-106 91, Sweden

<sup>b</sup> Department of Computer and Information Science, Linköping University, Sweden

<sup>c</sup> School of Mathematical and Physical Sciences, University of Technology Sydney, Australia

<sup>d</sup> School of Business, University of New South Wales, Australia

<sup>e</sup> Centre for Data Science, Queensland University of Technology, Australia

## ARTICLE INFO

### Article history:

Received 29 March 2021

Revised 3 October 2022

Accepted 6 October 2022

Available online 13 October 2022

### Keywords:

Bayesian

Markov chain Monte Carlo

Semi-long memory

Spectral analysis

Whittle likelihood

## ABSTRACT

A multivariate generalisation of the Whittle likelihood is used to extend spectral subsampling MCMC to stationary multivariate time series by subsampling matrix-valued periodogram observations in the frequency domain. To assess the performance of the methodology in challenging problems, a multivariate generalisation of the autoregressive tempered fractionally integrated moving average model (ARTFIMA) is introduced and some of its properties derived. Bayesian inference based on the Whittle likelihood is demonstrated to be a fast and accurate alternative to the exact time domain likelihood. Spectral subsampling is shown to provide up to two orders of magnitude additional speed-up, while retaining MCMC sampling efficiency and accuracy, compared to spectral methods using the full dataset.

© 2022 The Author(s). Published by Elsevier B.V. on behalf of EcoSta Econometrics and Statistics.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

## 1. Introduction

Recent technological developments in sensors, data storage and computing power make it possible to collect high frequency time series data at low cost; some examples are financial transaction data (Mykland and Zhang, 2012), neuroimaging data with high temporal resolution (Chen et al., 2019), sensor data from robots (Deisenroth et al., 2013) or meteorological weather stations, and GPS and smart card data used in transportation (Welch and Widita, 2019).

However, statistical analysis of time series with tens of thousands, hundreds of thousands, or even millions of data points is computationally challenging, especially when inferences are obtained with iterative methods such as Markov chain Monte Carlo (MCMC) simulation or stochastic optimisation algorithms, where the likelihood is evaluated a large number of times. It is therefore common to only use a portion of the data for inference, for example only the most recent observations, or by systematically selecting every  $k$ th data point over the study period. Such downsampling wastes valuable data, gives less precise inferences, and is not even an option when predictions are required at the original sampling frequency.

\* Corresponding author.

E-mail address: [mattias.villani@stat.su.se](mailto:mattias.villani@stat.su.se) (M. Villani).

Salomone et al. (2020) propose *spectral subsampling MCMC* to accelerate MCMC for long stationary univariate time series. They use the asymptotically motivated *Whittle likelihood* (Whittle, 1953) to approximate the likelihood of a stationary time series. The Whittle likelihood is based on the discrete Fourier transform with the important property of transforming a time series with dependent observations to *asymptotically independent* periodogram observations in the frequency domain. The key insight in Salomone et al. (2020) is that such independence makes it possible to extend subsampling MCMC approaches for independent data (Quiroz et al., 2018a; 2019; 2021; Dang et al., 2019) to univariate stationary time series by systematic subsampling of periodogram observations.

All of the applications mentioned above are naturally analysed in a multivariate setting: financial portfolios consisting of many assets, neuroimaging data simultaneously measured at multiple brain locations, meteorological data collected at several spatial locations, and so on. The multivariate aspect naturally leads to even more demanding computations in a high frequency setting.

The main contribution in this article is extending the spectral subsampling MCMC methodology to stationary *vector-valued* time series by using a multivariate version of the Whittle likelihood based on the asymptotic properties of the matrix-valued periodogram. The proposed multivariate spectral subsampling MCMC algorithm is evaluated on three challenging large-scale multivariate time series applications from meteorology, hydrology and environmental science.

The multivariate Whittle likelihood is of general interest for large-scale likelihood and Bayesian inference beyond subsampling MCMC. Even without subsampling, the Whittle likelihood is substantially more scalable to large data as it sidesteps the costly matrix inversions needed for the exact likelihood in the time domain. But the independence of periodogram observations also makes the multivariate Whittle likelihood directly useful for computing the unbiased gradient estimates from random subsets/batches of frequencies needed for large-scale variational inference (Tran et al., 2017) or maximum likelihood estimates from stochastic gradient descent algorithms (Goodfellow et al., 2016). An additional contribution of our article is that our empirical results clearly demonstrate that the posterior based on the multivariate Whittle likelihood gives an excellent approximation to the exact time domain posterior, even in complex models such as the VARMA model.

A challenge for subsampling MCMC methods is keeping the variance of the likelihood estimator small enough for the MCMC chain to mix well. Previous literature is therefore typically restricted to applications using models with a very small number of parameters, or models with a moderate number of parameters with a simple structure, such as logistic regression. To test our multivariate spectral subsampling MCMC methodology on a challenging set of problems, we introduce a new multivariate model with semi-long range dependence that extends the autoregressive tempered fractionally integrated moving average (ARTFIMA, Sabzikar et al. (2019)) to the multivariate setting. Several properties of the vector ARTFIMA model are derived, including the spectral density matrix needed for the Whittle likelihood. For all examples considered, the tempered fractional differencing is shown to improve upon the standard multivariate vector autoregressive moving average (VARMA) model in terms of the Bayesian information criterion. Our empirical examples show that spectral subsampling works well and gives a very large speed-up for time series that are long enough to make the variance reducing control variates effective. We also highlight the limitation of subsampling MCMC by showing empirically that spectral subsampling MCMC can get stuck when estimating complex multivariate models on shorter time series.

The rest of the article is organised as follows. Section 2 presents the Whittle likelihood for multivariate stationary time series. Section 3 outlines the subsampling MCMC methodology. Section 4 presents the models considered in our applications and establishes properties needed for the implementation of spectral methods. Section 5 demonstrates the methodology and evaluates the efficiency of the proposed spectral subsampling algorithm on real data. Section 6 compares our proposed vector ARTFIMA to the alternative of separately estimating univariate ARTFIMA models for each series (Sabzikar et al., 2019), both in terms of estimated parameters and in forecasting. Section 7 concludes and discusses future research. Appendix A derives some properties of the vector ARTFIMA model and Appendix B contains additional empirical results.

## 2. The Whittle likelihood for multivariate time series

Let  $\mathbf{X}_t \in \mathbb{R}^r$  be an  $r$ -variate zero mean stationary time series with absolutely summable autocovariance matrix function

$$\gamma_{\mathbf{X}}(\tau) = \text{Cov}(\mathbf{X}_t, \mathbf{X}_{t-\tau}), \text{ for } \tau \in \mathbb{Z}, \quad (1)$$

where  $\mathbb{Z}$  is the set of integers. The spectral density matrix is

$$f_{\mathbf{X}}(\omega) = \frac{1}{2\pi} \sum_{\tau=-\infty}^{\infty} \gamma_{\mathbf{X}}(\tau) \exp(-i\omega\tau), \text{ for } \omega \in (-\pi, \pi], \quad (2)$$

with the diagonal elements being the usual spectral density for each univariate time series and the off-diagonal elements are the cross-spectral densities

$$f_{jk}(\omega) = \frac{1}{2\pi} \sum_{\tau=-\infty}^{\infty} \gamma_{jk}(\tau) \exp(-i\omega\tau), \text{ for } \omega \in (-\pi, \pi], \quad (3)$$

where  $\gamma_{jk}(\tau)$  is the cross-correlation function between the  $j$ th and  $k$ th time series in  $\mathbf{X}_t$ . Since the elements of  $\gamma_{\mathbf{X}}(\tau)$  are real,  $f_{\mathbf{X}}(\omega)$  is Hermitian, i.e.  $f_{\mathbf{X}}(\omega)^H = f_{\mathbf{X}}(\omega)$ , where  $\mathbf{A}^H = (\bar{\mathbf{A}})^T$  is the conjugate transpose of a matrix and  $\bar{\mathbf{A}}$  is the matrix of complex conjugates of the elements of  $\mathbf{A}$ . Brillinger (2001)[Theorem 2.5.1] proves that  $f_{\mathbf{X}}(\omega)$  is also non-negative definite.

The discrete Fourier transform (DFT) of the multivariate time series  $\{\mathbf{X}_t\}_{t=0}^{T-1}$  is

$$J_T(\omega_k) = \sum_{t=0}^{T-1} \mathbf{X}_t \exp(-i\omega_k t), \quad (4)$$

for  $\omega_k \in \Omega_T = \{2\pi k/T \text{ for } k = -\lceil T/2 \rceil + 1, \dots, \lfloor T/2 \rfloor\}$ , the set of Fourier frequencies. Let  $\mathbf{X} \sim \text{CN}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  denote that the  $r$ -dimensional complex-valued vector  $\mathbf{X}$  follows the multivariate complex normal distribution (Brillinger, 2001, Ch. 4.2), i.e. that

$$\begin{pmatrix} \text{Re } \mathbf{X} \\ \text{Im } \mathbf{X} \end{pmatrix} \sim N_{2r} \left[ \begin{pmatrix} \text{Re } \boldsymbol{\mu} \\ \text{Im } \boldsymbol{\mu} \end{pmatrix}, \frac{1}{2} \begin{pmatrix} \text{Re } \boldsymbol{\Sigma} & -\text{Im } \boldsymbol{\Sigma} \\ \text{Im } \boldsymbol{\Sigma} & \text{Re } \boldsymbol{\Sigma} \end{pmatrix} \right],$$

where  $\boldsymbol{\Sigma}$  is a Hermitian positive definite covariance matrix. The  $J_T(\omega_k)$  are asymptotically independent over the different frequencies and (Brillinger, 2001, Theorem 4.4.1)

$$\frac{1}{\sqrt{T}} J_T(\omega_k) \sim \text{CN}(0, 2\pi f_{\mathbf{X}}(\omega_k)) \text{ as } T \rightarrow \infty, \quad (5)$$

except at  $\omega_k = 0$  and  $\omega_k = \pi$ , where instead  $(1/\sqrt{T})J_T(\omega) \sim N_r(0, 2\pi f_{\mathbf{X}}(\omega_k))$  as  $T \rightarrow \infty$ . Following Brillinger (2001), we will assume that  $f_{\mathbf{X}}(\omega)$  is non-singular.

The *periodogram* ordinate at frequency  $\omega$  is defined as

$$I_T(\omega) = (2\pi T)^{-1} J_T(\omega) J_T(\omega)^H.$$

The periodogram ordinates are therefore asymptotically independent complex Wishart distributed with one degree of freedom  $I_T(\omega) \sim \text{CW}_r(1, f_{\mathbf{X}}(\omega))$  (Brillinger, 2001, Theorem 7.2.4), except for the frequencies  $\omega_k = 0$  and  $\omega_k = \pi$  where the  $I_T(\cdot)$  instead follow a (real) Wishart distribution. The periodogram ordinates  $I_T(\omega)$  are singular matrices for  $r > 1$ . The density function of this singular Wishart distribution is derived for the real case by Uhlig (1994) with respect to the Hausdorff measure, and by Srivastava (2003) with respect the Lebesgue measure for the functionally independent elements of the matrix. The density for the complex singular Wishart distribution  $\mathbf{W} \sim \text{CW}_r(\nu, \boldsymbol{\Sigma})$  for  $\nu < r$  over the space of  $r \times r$  positive semidefinite matrices of rank  $\nu$  is derived in Ratnarajah and Vaillancourt (2005)[Theorem 3] as

$$p(\mathbf{W} | \nu, \boldsymbol{\Sigma}) = \frac{\pi^{\nu(\nu-r)}}{\Gamma_{\nu}(\nu)} \left( \prod_{j=1}^{\nu} \ell_j \right) |\boldsymbol{\Sigma}|^{-\nu} \exp(-\text{tr } \boldsymbol{\Sigma}^{-1} \mathbf{W}),$$

where  $\Gamma_{\nu}(\nu) = \pi^{\nu(\nu-1)/2} \prod_{k=1}^{\nu} \Gamma(\nu - k + 1)$ ,  $\ell_j$  is the  $j$ th eigenvalue in the reduced spectral decomposition  $\mathbf{W} = \mathbf{E}_1 \boldsymbol{\Lambda} \mathbf{E}_1^H$ , where  $\mathbf{E}_1$  is an  $r \times \nu$  complex orthogonal matrix and  $\boldsymbol{\Lambda} = \text{Diag}(\ell_1, \dots, \ell_{\nu})$ .

Consider now inference for a parametric model with parameter vector  $\boldsymbol{\theta}$  and spectral density matrix  $f_{\mathbf{X},\boldsymbol{\theta}}(\omega_k)$ . The Whittle log-likelihood exploits the asymptotic result for the periodogram and is defined using the complex singular Wishart distribution as

$$\ell_{\mathcal{W}}(\boldsymbol{\theta}) = - \sum_{\omega_k \in \tilde{\Omega}_T} \left( \log |f_{\mathbf{X},\boldsymbol{\theta}}(\omega_k)| + \text{tr} [f_{\mathbf{X},\boldsymbol{\theta}}(\omega_k)^{-1} I_T(\omega)] \right), \quad (6)$$

where  $\tilde{\Omega}_T$  is the set of Fourier frequencies with the omission of  $\omega_k = 0$  and  $\omega_k = \pi$ . The term for  $\omega_k = 0$  is not included when the time series is demeaned since then  $J_T(0) = 0$  by construction; the term for  $\omega_k = \pi$  is removed for simplicity since it has a different distribution than the other frequencies and its influence is negligible asymptotically. Note that since  $f_{\mathbf{X},\boldsymbol{\theta}}(\omega_k)$  is Hermitian for an absolutely summable stationary process (Brillinger, 2001), both  $|f_{\mathbf{X},\boldsymbol{\theta}}(\omega_k)|$  and

$$\text{tr} [f_{\mathbf{X},\boldsymbol{\theta}}(\omega_k)^{-1} I_T(\omega)] = I_T(\omega)^H f_{\mathbf{X},\boldsymbol{\theta}}(\omega_k)^{-1} I_T(\omega),$$

are real-valued.

### 3. Subsampling Markov chain Monte Carlo

Subsampling MCMC uses the framework of pseudo-marginal MCMC (Andrieu and Roberts, 2009), in which an estimator of the likelihood is used within a Metropolis-Hastings algorithm. This section briefly reviews subsampling MCMC; see Quiroz et al. (2019) and Salomone et al. (2020) for details on how pseudo-marginal methods are used for subsampling problems.

Let  $\pi(\boldsymbol{\theta}) \propto L_n(\boldsymbol{\theta})p(\boldsymbol{\theta})$  denote the posterior distribution of the model parameter  $\boldsymbol{\theta}$  from a sample of  $n$  observations with likelihood function  $L_n(\boldsymbol{\theta})$ , and prior distribution  $p(\boldsymbol{\theta})$ . Metropolis-Hastings MCMC algorithms sample iteratively from  $\pi(\boldsymbol{\theta})$  by proposing a parameter vector  $\boldsymbol{\theta}^{(j)}$  at the  $j$ th iteration from the proposal distribution  $g(\boldsymbol{\theta}^{(j)} | \boldsymbol{\theta}^{(j-1)})$  and accepting the draw with probability

$$\min \left\{ 1, \frac{L_n(\boldsymbol{\theta}^{(j)})p(\boldsymbol{\theta}^{(j)})}{L_n(\boldsymbol{\theta}^{(j-1)})p(\boldsymbol{\theta}^{(j-1)})} \cdot \frac{g(\boldsymbol{\theta}^{(j-1)} | \boldsymbol{\theta}^{(j)})}{g(\boldsymbol{\theta}^{(j)} | \boldsymbol{\theta}^{(j-1)})} \right\}. \quad (7)$$

The cost of computing the likelihood  $L_n(\theta)$  in the acceptance probability (7) is a major concern when  $n$  is large. Quiroz et al. (2019) propose speeding up MCMC for large  $n$  by replacing  $L_n(\theta)$  with an estimate  $\hat{L}(\theta, \mathbf{u})$  based on a small random subsample of  $m \ll n$  observations, where  $\mathbf{u} = (u_1, \dots, u_m)$  indexes the selected observations. Their algorithm samples  $\theta$  and  $\mathbf{u}$  jointly from an extended target distribution  $\tilde{\pi}(\theta, \mathbf{u}) \propto \hat{L}(\theta, \mathbf{u})p(\theta)p(\mathbf{u})$ . Andrieu and Roberts (2009) show that such pseudo-marginal MCMC algorithms sample from the full-data posterior  $\pi(\theta)$  if  $\hat{L}(\theta, \mathbf{u})$  is an unbiased and almost surely positive estimator of  $L_n(\theta)$ , where the unbiasedness condition is  $E_{\mathbf{u}}\hat{L}(\theta, \mathbf{u}) = \int \hat{L}(\theta, \mathbf{u})p(\mathbf{u})d\mathbf{u} = L_n(\theta)$ . Quiroz et al. (2019) use an unbiased estimator of the log-likelihood  $\hat{\ell}(\theta, \mathbf{u})$  and then debias  $\exp(\hat{\ell}(\theta, \mathbf{u}))$  to estimate the full-data likelihood. This debiasing approach does not remove all bias and the marginal distribution of  $\theta$  from their pseudo-marginal sampler is the slightly perturbed posterior

$$\eta(\theta) = \frac{\int \hat{L}(\theta, \mathbf{u})p(\mathbf{u})d\mathbf{u}}{\int \int \hat{L}(\theta, \mathbf{u})p(\mathbf{u})d\mathbf{u}p(\theta)d\theta}. \quad (8)$$

The perturbation error of  $\eta(\theta)$  in (8) is within  $O(n^{-1}m^{-2})$  distance in total variation norm of the true posterior  $\pi(\theta)$ , and the applications in Quiroz et al. (2019, 2021) and Dang et al. (2019) show negligible bias. Salomone et al. (2020) show that these results extend to subsampling for the Whittle likelihood, with the true posterior in this case based on the Whittle likelihood for the full data.

To apply subsampling MCMC, the log-likelihood needs to decompose as a sum  $\ell(\theta) = \sum_{k=1}^n \ell_k(\theta)$ ; either by assuming independent data or by using the Whittle likelihood in the frequency domain for temporally dependent data as in (6). Estimating the log-likelihood is analogous to the problem of estimating a population total in survey sampling (Quiroz et al., 2018b).

It is by now well-known that subsampling MCMC requires a likelihood estimator with small variance, otherwise the sampler tends to get stuck (Quiroz et al., 2019; 2018b). Quiroz et al. (2019) propose using the *difference estimator*

$$\hat{\ell}(\theta, \mathbf{u}) = q(\theta) + \frac{n}{m} \sum_{i=1}^m (\ell_{u_i}(\theta) - q_{u_i}(\theta)), \quad \text{with } q(\theta) = \sum_{k=1}^n q_k(\theta),$$


with control variates  $q_k(\theta)$  to reduce the variance. The second term in the first equation is an unbiased estimate of  $\ell(\theta) - q(\theta)$  and is efficient if  $q_k(\theta) \approx \ell_k(\theta)$  for  $k = 1, \dots, n$ . The control variates homogenise the log-density terms in the estimator so that the observations can be sampled by simple random sampling (Quiroz et al., 2019). Bardenet et al. (2017) propose setting  $q_k(\theta)$  equal to a second order Taylor expansion around  $\theta^*$  (e.g. the posterior mode). An important property of this control variate is that the  $q(\theta)$  term can be computed in  $O(1)$  time.

The control variate in Bardenet et al. (2017) works well when the log-density  $\ell_k(\theta)$  of each observation is approximately quadratic in the neighbourhood of  $\theta^*$  explored by the MCMC. There is, however, no guarantee that the individual log-densities  $\ell_k(\theta)$  are close to quadratic in complex models, particularly in large parameter spaces where the posterior may not be highly concentrated around  $\theta^*$ . Salomone et al. (2020) therefore propose a grouped quadratic control variate, in which observations are divided into  $G$  groups and the log-likelihood contribution for each group is approximated by a quadratic function. The idea is that the Bernstein-von Mises theorem (asymptotic normality of the posterior) suggests an approximately quadratic log-likelihood for the group, given that the number of observations in the group is large enough.  $G$  should be chosen such that there are enough observations in each group for the approximation to be accurate. Note that the sampling units in this approach are the groups, and hence there are  $G$  group log-likelihoods to subsample. For the Whittle likelihood, the order of the dataset corresponds to the order of the frequencies in the DFT (from low to high). The groups should then be formed via systematic sampling, so that each group contains periodogram ordinates over the whole frequency spectrum; Fig. 1 graphically illustrates this and the block setup which is now described.

Pseudo-marginal methods can be made much more efficient by correlating the estimators used at the numerator and denominator of the Metropolis-Hastings acceptance ratio (Deligiannidis et al., 2018; Tran et al., 2016; Quiroz et al., 2021). We use the Tran et al. (2016) approach and divide the random numbers  $\mathbf{u}$  (here the set of subsampled observations indices) into  $K$  blocks. By updating the  $\mathbf{u}$  only within one of the blocks at each MCMC iteration, Tran et al. (2016) show that the correlation between the logs of the estimators in the numerator and denominator is approximately  $1 - 1/K$ , and setting  $K$  large makes the pseudo-marginal method less sensitive to the variance of the log of the likelihood estimator; see Fig. 1 for the relation between the grouping described in the previous paragraph and the block setup. Algorithm 1 sketches the implementation of the algorithm; see Quiroz et al. (2019) for details.

#### 4. The Vector ARTFIMA process

Salomone et al. (2020) demonstrate that spectral subsampling MCMC can be successfully applied in univariate time series models with likelihood functions that are known to be non-Gaussian in the parameters due to local non-identification, such as ARMA and ARTFIMA. We explore the performance of spectral subsampling MCMC for multivariate versions of ARMA and ARTFIMA, which pose an even greater challenge since the number of parameters in multivariate models typically increase quadratically with the number of time series.



**Fig. 1.** Explaining grouping and blocking in spectral subsampling MCMC. The periodogram observations and their corresponding log-density terms  $\ell_{\omega_i}(\theta)$  (top row) are systematically divided into groups (indicated by colors) and summed  $\ell^{(g)}(\theta)$  (bottom row,  $\mathcal{G}_g$  denotes the indices belonging to group  $g$ ), so that each group spans over the whole frequency domain. The groups are then divided into blocks and only the group subsampling indicators within a subset of the blocks are updated at each MCMC iteration.

---

**Algorithm 1:** Subsampling MCMC with a random walk Metropolis proposal for  $\theta$  and a block proposal for  $\mathbf{u}$ .

---

**Input:** data  $\mathbf{y}$ , unbiased log-likelihood estimator  $\widehat{\ell}(\theta, \mathbf{u})$  based on Taylor control variates, variance estimator  $\widehat{\sigma}_{\widehat{\ell}}^2(\theta, \mathbf{u})$  of  $\widehat{\ell}(\theta, \mathbf{u})$ , prior density  $p(\theta)$ , initial value  $\theta^{(0)}$ , initial subsample  $\mathbf{u}^{(0)} = (\mathbf{u}_1^{(0)}, \dots, \mathbf{u}_K^{(0)})$  divided into  $K$  blocks, proposal covariance  $\Omega$  for  $\theta$ .

define  $\widehat{p}(\mathbf{y}|\theta, \mathbf{u}) = \exp(\widehat{\ell}(\theta, \mathbf{u}) - \frac{1}{2}\widehat{\sigma}_{\widehat{\ell}}^2(\theta, \mathbf{u}))$

**for**  $j = 1$  **to**  $N$  **do**

$\backslash\backslash$  propose the subsample  $\mathbf{u}$

    sample  $k$  uniformly from  $\{1, \dots, K\}$ , generate  $\mathbf{u}'_k \sim p(\mathbf{u}_k)$  and propose

$$\mathbf{u}' = (\mathbf{u}_1^{(j-1)}, \dots, \mathbf{u}'_k, \dots, \mathbf{u}_K^{(j-1)})$$

$\backslash\backslash$  propose the parameter  $\theta$

    generate  $\theta' \sim N(\theta^{(j-1)}, \Omega)$

    set  $\mathbf{u}^{(j)}, \theta^{(j)} \leftarrow \mathbf{u}', \theta'$  with probability

$$\alpha = \min\left(1, \frac{\widehat{p}(\mathbf{y}|\theta', \mathbf{u}')p(\theta')}{\widehat{p}(\mathbf{y}|\theta^{(j-1)}, \mathbf{u}^{(j-1)})p(\theta^{(j-1)})}\right)$$

    else set  $\mathbf{u}^{(j)}, \theta^{(j)} \leftarrow \mathbf{u}^{(j-1)}, \theta^{(j-1)}$

**end**

**Output:** autocorrelated random draws  $\theta^{(1)}, \dots, \theta^{(N)}$  from  $\eta(\theta)$  in (8).

---

#### 4.1. Vector ARMA

The vector autoregressive moving average VARMA( $p, q$ ) process is

$$\Phi(L)(\mathbf{Y}_t - \boldsymbol{\mu}) = \Theta(L)\boldsymbol{\varepsilon}_t, \quad (9)$$

where  $\{\boldsymbol{\varepsilon}_t\}_{t=1}^T$  is an iid sequence from  $N(0, \Sigma_{\varepsilon})$ ,  $\Phi(L) = I_r - \Phi_1 L - \dots - \Phi_p L^p$  and  $\Theta(L) = I_r + \Theta_1 L + \dots + \Theta_q L^q$  are the AR and MA lag matrix polynomials, respectively. We assume that the usual conditions for stationarity and invertibility of the VARMA process hold:

**Assumption 1.** The matrix polynomials  $\Phi(z)$  and  $\Theta(z)$  share no common zeros and  $|\Phi(z)| \neq 0$  and  $|\Theta(z)| \neq 0$  for  $|z| \leq 1$ .

The spectral density matrix of the VARMA( $p, q$ ) model is (Brockwell and Davis, 1991, Ch. 11)

$$f_{\mathbf{Y}}(\omega) = \frac{1}{2\pi} \Phi^{-1}(e^{-i\omega}) \Theta(e^{-i\omega}) \Sigma_{\varepsilon} \Theta^H(e^{-i\omega}) \Phi^{-H}(e^{-i\omega}), \quad (10)$$

where  $\mathbf{A}^{-H}$  denotes the inverse of the conjugate transpose of the matrix  $\mathbf{A}$ .




Fig. 2. Water velocity data.




Fig. 3. Swedish temperature data. The final 878 time points in red are used for testing the forecast accuracy of the model.

The Whittle likelihood is valid if the process is stationary, i.e. if all roots of  $\Phi(z)$  are outside of the unit circle. We therefore use the reparametrisation in [Ansley and Kohn \(1986\)](#) to map a set of unconstrained real-valued AR coefficient matrices to the set of stationary parameters. The same reparametrisation is also used on the MA parameter matrices to ensure invertibility. These reparametrisations are an additional source of non-linearity/non-Gaussianity; even a plain vector AR (VAR) model is no longer linear in the parameters.

#### 4.2. Vector ARTFIMA

[Sabzikar et al. \(2019\)](#) define the univariate ARTFIMA( $p, d, \lambda, q$ ) process for  $Y_t$  as

$$\phi(L)\Delta^{d,\lambda}(Y_t - \mu) = \vartheta(L)\varepsilon_t, \quad (11)$$

where  $\{\varepsilon_t\}_{t \in \mathbb{Z}}$  is an iid sequence of zero mean random variables with variance  $\sigma_\varepsilon^2$ ,  $\phi(L) \equiv 1 - \phi_1 L - \dots - \phi_p L^p$ , and  $\vartheta(L) \equiv 1 + \vartheta_1 L + \dots + \vartheta_q L^q$ , are the autoregressive and moving average lag polynomials, and  $L$  is the lag operator, i.e.  $L^k Y_t \equiv Y_{t-k}$ . The tempered fractional differencing operator  $\Delta^{d,\lambda}$ , where  $d \notin \mathbb{Z}$  is the fractional differencing parameter and  $\lambda \geq 0$  is the tempering parameter, is defined by the generalised binomial theorem as

$$\Delta^{d,\lambda} Y_t \equiv (1 - e^{-\lambda} L)^d Y_t = \sum_{j=0}^{\infty} b_j^{d,\lambda} Y_{t-j}, \quad (12)$$

where

$$b_j^{d,\lambda} \equiv (-1)^j \binom{d}{j} e^{-\lambda j} \quad \text{and} \quad \binom{d}{j} = \frac{\Gamma(1+d)}{\Gamma(1+d-j)j!}.$$

We follow the convention in time series of not explicitly writing out the lag operator  $L$  in differencing operators unless needed for clarity, i.e.,  $\Delta^{d,\lambda} \equiv \Delta^{d,\lambda}(L)$ .

To explain the role of the parameters  $d$  and  $\lambda$ , note that for  $\lambda = 0$  and  $d$  a non-negative integer,  $\Delta^{d,\lambda} Y_t$  reduces to simple differencing of order  $d$  and the ARTFIMA model in (11) reduces to the autoregressive integrated moving average (ARIMA) process. For  $\lambda = 0$  and fractional  $d$  we obtain the autoregressive fractionally integrated moving average (ARFIMA) model ([Granger and Joyeux, 1980](#)). The ARFIMA process is stationary and invertible for  $-0.5 < d < 0.5$ , and has long-range or long-memory dependence with an autocovariance function dying off so slowly that it is not absolutely summable. The tempering parameter  $\lambda > 0$  in ARTFIMA allows for semi-long range dependence, i.e. ARFIMA-like long range dependence for a number



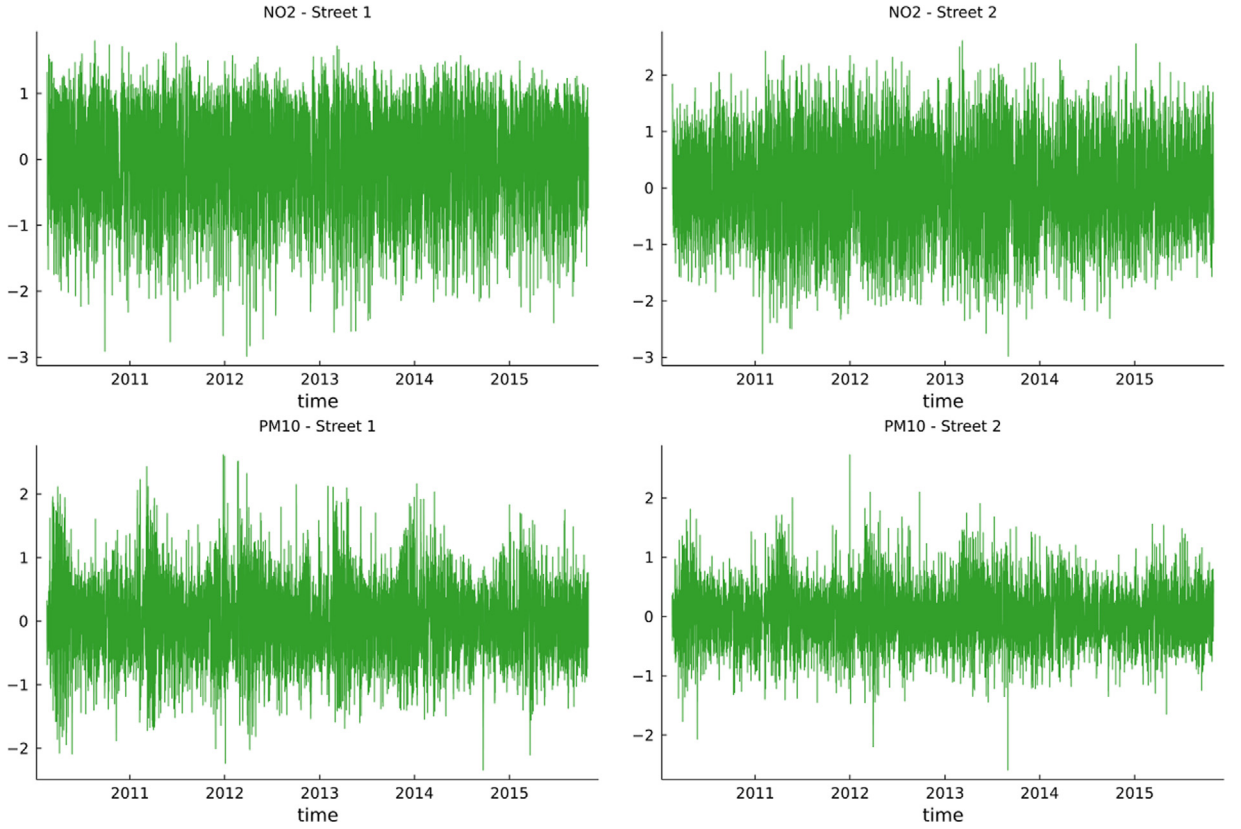


Fig. 4. Stockholm air pollution data.

of lags beyond which the autocovariances decay exponentially fast. [Sabzikar et al. \(2019\)](#) prove that the ARTFIMA process is stationary for any  $d \notin \mathbb{Z}$  if  $\lambda > 0$ , provided that the univariate case of Assumption 1 holds.

The univariate ARTFIMA process is now extended to the multivariate case, where  $\mathbf{Y}_t$  is an  $r$ -dimensional vector-valued time series. We define the vector ARTFIMA (VARTFIMA) process as

$$\Phi(L)\Delta^{\mathbf{d},\lambda}(\mathbf{Y}_t - \boldsymbol{\mu}) = \Theta(L)\boldsymbol{\varepsilon}_t, \quad (13)$$

where  $\Delta^{\mathbf{d},\lambda}$  is the multivariate tempered fractional differencing operator defined by

$$\Delta^{\mathbf{d},\lambda}\mathbf{Y}_t \equiv \left( \Delta^{d_1,\lambda_1}Y_{1,t}, \dots, \Delta^{d_r,\lambda_r}Y_{r,t} \right)^\top, \quad (14)$$

hence allowing for different fractional differences and temporal differencing for the  $r$  time series.

[Theorem 1](#) generalises the spectral density result in [Sabzikar et al. \(2019\)](#) to the vector case. Its proof is in [Appendix A](#).

**Theorem 1.** *The VARTFIMA process is causal and stationary for all  $d \notin \mathbb{Z}$  and all  $\lambda > 0$  if  $|\Phi(z)| \neq 0$  for  $|z| \leq 1$ . The spectral density matrix is*

$$f_{\mathbf{Y}}(\omega) = \frac{1}{2\pi} \Delta^{-\mathbf{d},\lambda}(e^{-i\omega}) \Phi^{-1}(e^{-i\omega}) \Theta(e^{-i\omega}) \Sigma_{\varepsilon} \Theta^H(e^{-i\omega}) \Phi^{-H}(e^{-i\omega}) \Delta^{-\mathbf{d},\lambda}(e^{-i\omega})^H,$$

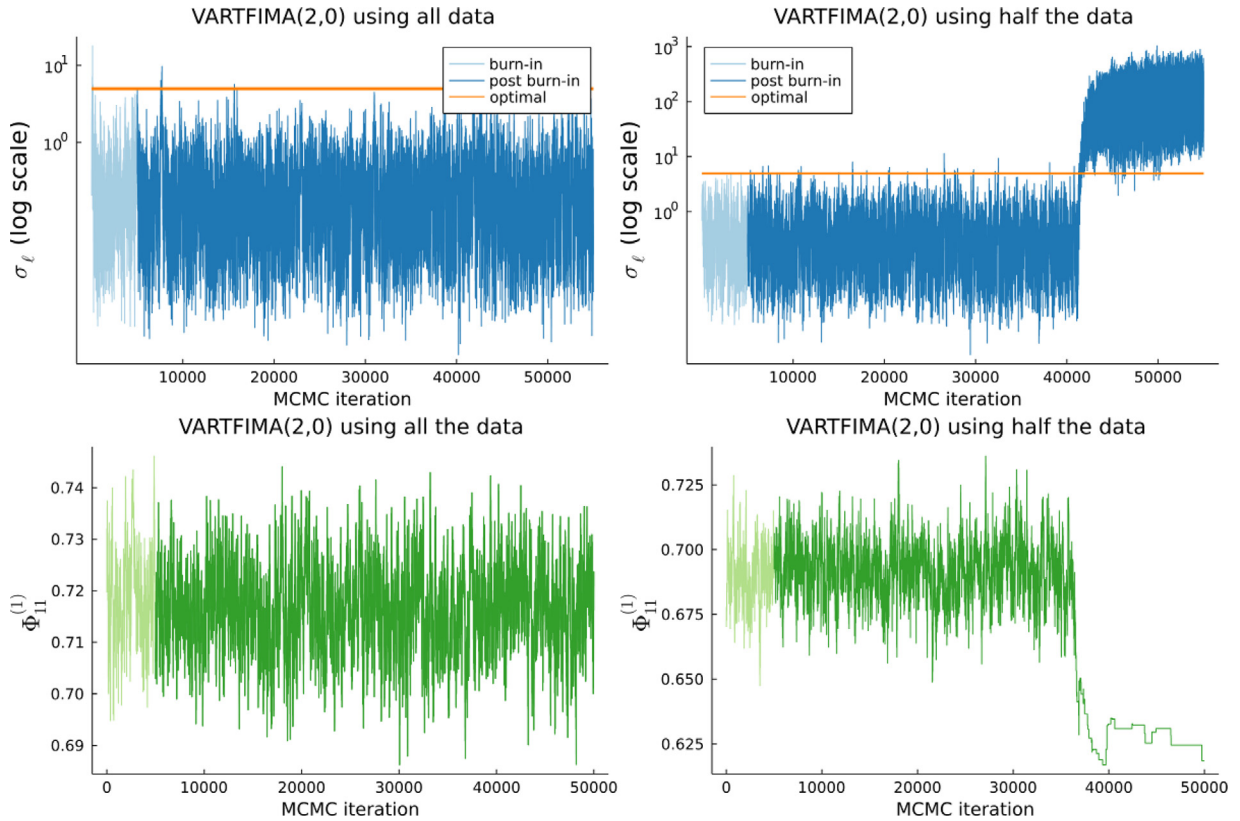
where  $\Delta^{-\mathbf{d},\lambda}(z) = \text{Diag}((1 - e^{-\lambda_1}z)^{-d_1}, \dots, (1 - e^{-\lambda_r}z)^{-d_r})$ , and  $\mathbf{A}^{-H}$  the inverse of the conjugate transpose of  $\mathbf{A}$ .

## 5. Applications

We now illustrate the proposed subsampling MCMC methodology on three long multivariate time series datasets of varying dimensions.

### 5.1. Datasets

The first dataset contains observations on mean water velocity from two measurement stations at St Clair River and Detroit River, located on opposite sides of Lake St Clair forming a connecting channel between Lake Huron and Lake Erie.



**Fig. 5.** Standard deviation of the log-likelihood estimator over the MCMC iterations for the VARTFIMA(2,0) model fitted to the Temperature dataset using all the data (top left) versus when only the last 60001 observations was used in the estimation (top right). The optimal standard deviation (Quiroz et al., 2021, Lemma S8) is marked out with an orange line. The graphs in the bottom row of the figure show the MCMC chains for the AR parameter of the first series on its own first lag,  $\Phi_{11}^{(1)}$ , for the respective dataset.

The measurements are recorded every 12th minute. The final dataset contains 130,001 observations from Jan 3, 2016 at 00:00 hours until Dec 21, 2018 at 08:00 hours for each of the two locations. Missing observations are imputed using the `na.interp` function in the R package `forecast`. Fig. 2 plots the data.

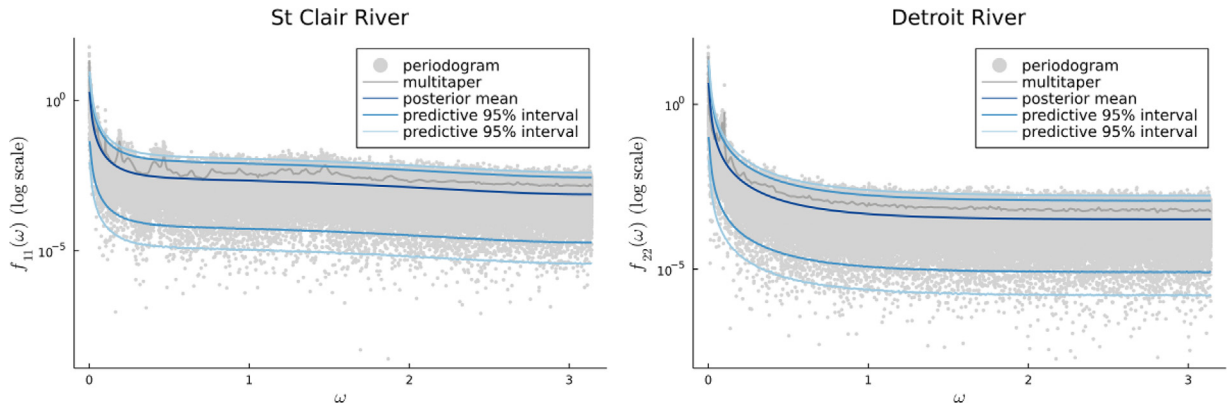
The second dataset contains 124,879 observations of Swedish temperatures at each of three airport locations (Arlanda, Bromma and Landvetter), giving a three-dimensional time series. The data are measured in an hourly scale for the time period February 1, 2008 until May 1, 2022, and are processed as follows. Each univariate series is preprocessed separately: missing observations are imputed with the value at a nearby location if available or otherwise imputed with the `na.interp` function in the R package `forecast`. Trend and seasonal components are then removed using the `mst1` function in the R package `stats` removing both a daily and an annual seasonal cycle. Fig. 3 plots the processed data. The last 878 observations are used as a test set for evaluating prediction performance in Section 6 and are therefore not used for the posterior inference. The raw data before pre-processing are plotted in Fig. B.17 in Appendix B.

The third dataset contains 50,001 observations of nitrogen dioxide (NO2) and particulate matter (PM10) pollution at two streets in central Stockholm, giving a four-dimensional time series. The data are measured hourly for the time period February 16, 2010 until October 31, 2015. The missing values are imputed with the `na.interp` function for seasonal data in R. The series are analysed in logs through the transformation  $\log(x_t - \min\{x_t\}_{t=1}^T + 1)$  to make the series more normally distributed; we subtract the minimum value and add one to ensure that the series are positive before taking logs (the two PM10 series have a few negative values and the NO2 series have some values that are very close to zero). The transformed series are finally filtered through an estimated seasonal ARIMA(0, 0, 0)(1, 0, 1)<sub>24</sub> to remove the natural 24 hour cyclic trend. The autocorrelation functions of the residuals from this particular seasonal ARIMA model indicate stationarity for all four filtered series. Fig. 4 plots the transformed and filtered data. The raw data are plotted in Fig. B.18 in Appendix B.

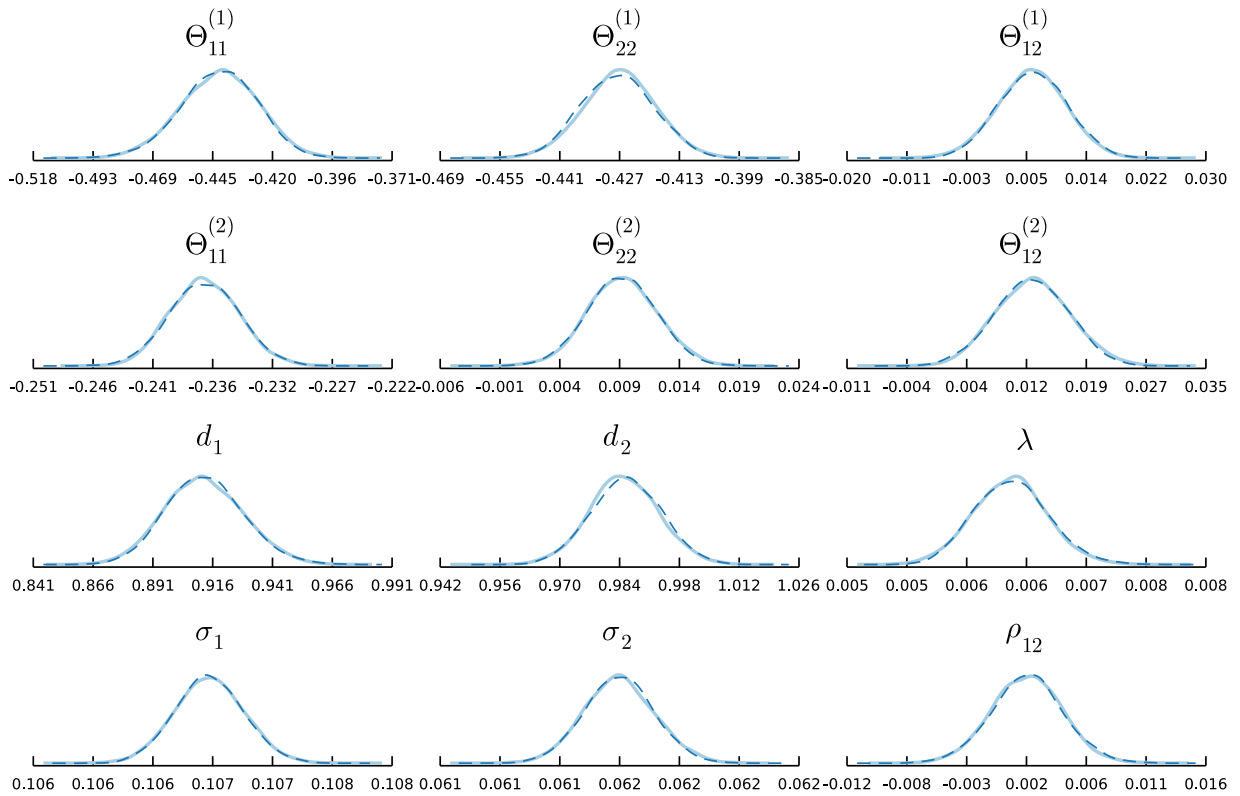
## 5.2. Algorithm, model and prior settings

The grouped control variates in Section 3 are used by dividing the observations into  $G = 1000$  groups, which gives accurate enough control variates in most of our estimated models. See Salomone et al. (2020) for a demonstration on how varying  $G$  affects the variance of the estimator. Fig. 1 illustrates how the  $g$ th group is chosen by including the  $g$ th lowest





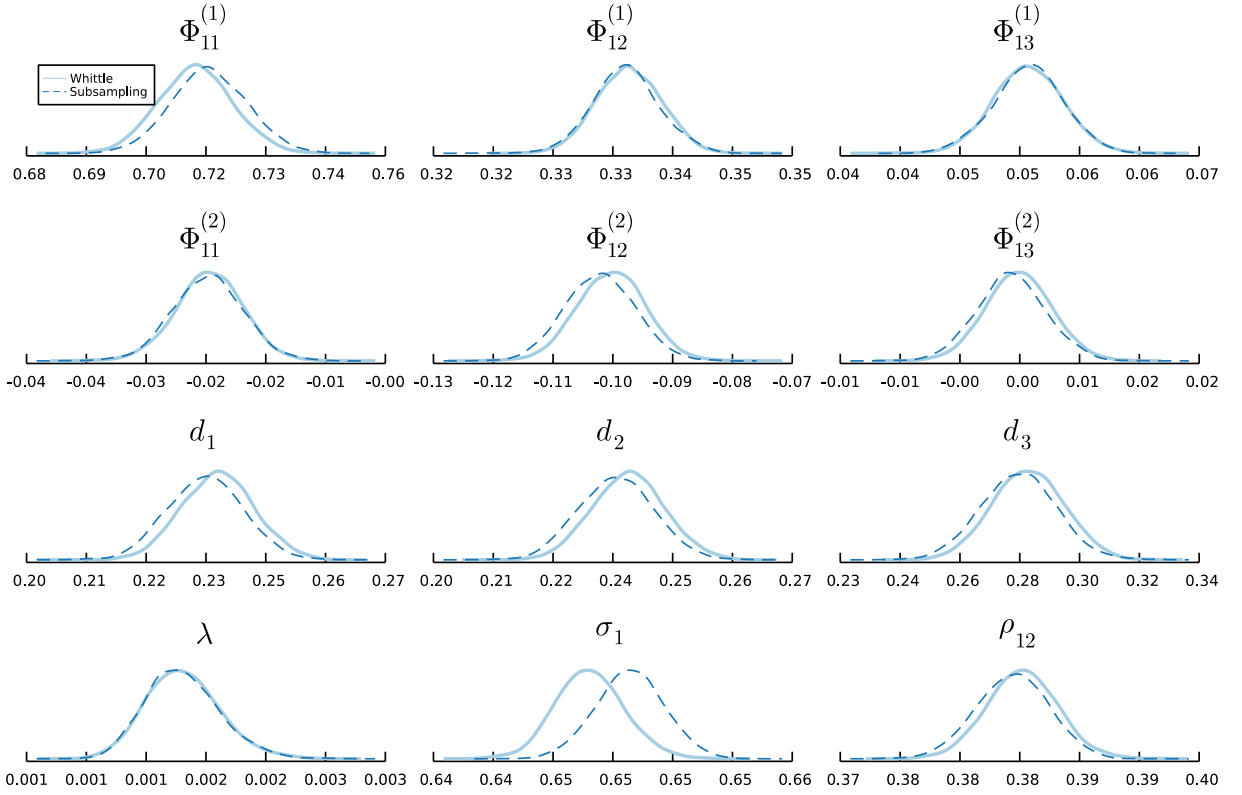
**Fig. 6.** Posterior predictive fit of the univariate periodogram data for the VARTFIMA(0,2) model fitted to the Water velocity data. The predictive intervals are obtained by simulation from the asymptotic Whittle distribution  $\mathcal{I}(\omega) \sim \text{Expon}(f_\theta(\omega))$  with parameters  $\theta$  drawn from the VARTFIMA(0,2) posterior.



**Fig. 7.** Kernel density estimates of a subset of the marginal posterior densities for the VARTFIMA(0,2) model fitted to the Water velocity data. The solid densities are from MCMC using the Whittle posterior on the whole dataset, and the dashed densities are obtained from spectral subsampling MCMC.

frequency and subsequently systematically sampling every  $G$ th frequency, ensuring that each group contains periodogram ordinates over the whole frequency spectrum. We use 10 randomly selected groups (1% of the data) to estimate the full-data likelihood in each MCMC iteration. For the block pseudo-marginal, we divide the random numbers into  $K = 10$  blocks (one group per block), resulting in a correlation between the log of the estimators of approximately 0.9 as discussed in Section 3. For all examples, we sample 55000 draws from the posterior distribution and discard 5000 draws as burn-in.

For the ARTFIMA models, we allow for different differencing parameters for each time series. The tempering parameters are restricted to be equal for each time series as the more general model with different  $\lambda$  gave estimates of  $\lambda$  that are very close across all series in a given dataset.



**Fig. 8.** Kernel density estimates of a subset of the marginal posterior densities for the VARTFIMA(2,0) model fitted to the Temperature data. The solid densities are from MCMC using the Whittle posterior on the whole dataset, and the dashed densities are obtained from spectral subsampling MCMC.

We use a Minnesota-style prior (Doan et al., 1984) for the autoregressive and moving average coefficients, which is a normal prior with diagonal covariance matrix with elements

$$v_{ij,l} = \begin{cases} (\lambda_0/l)^2, & \text{if } i = j \\ (\lambda_0\theta_0\sigma_i/l\sigma_j)^2 & \text{if } i \neq j. \end{cases}$$

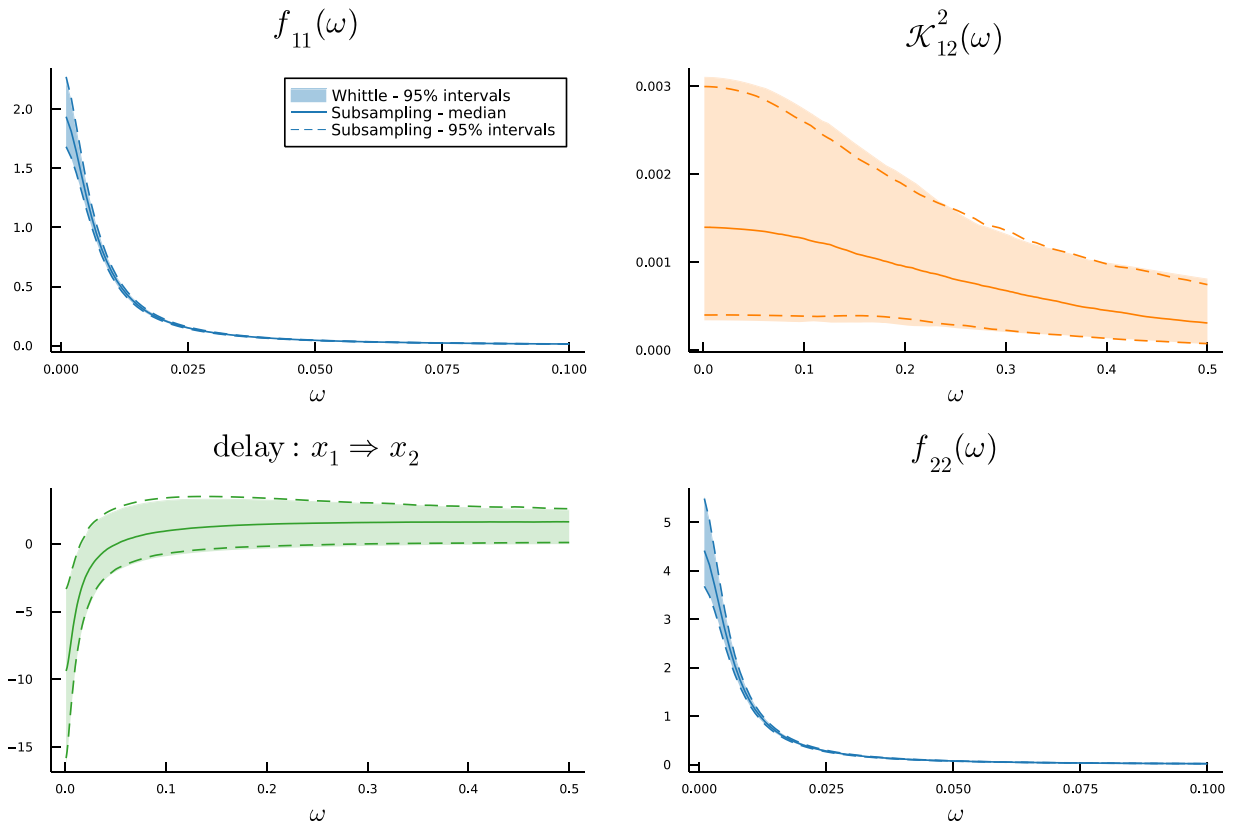
All prior means are set to zero. The hyper-parameter  $\lambda_0$  controls how tightly the coefficient of the first lag is concentrated around 0; we set  $\lambda_0 = 1$ . Note that the prior variance decreases with increasing lag length  $l$ , allowing more shrinkage of coefficients corresponding to lags further back in time. The hyper-parameter  $\theta_0$  accounts for the belief that most of the variation in each variable is accounted for by its own lags and is hence set to a value in the range  $[0,1]$ . We set  $\theta_0 = 0.2$ . The ratio  $\sigma_i^2/\sigma_j^2$  accounts for the difference in the variability of the variables. Following standard practice,  $\sigma_i^2$  is set to the residual variance computed by fitting a univariate AR model to the  $i$ th series. Note, however, that we use the above Minnesota prior on the unrestricted parameters in the Ansley and Kohn (1986) reparameterisation so the prior is acting on (a rescaled version of) the partial autocorrelation matrices instead of the original AR and MA coefficients. We parameterise the covariance matrix  $\Sigma_\varepsilon$  as a Cholesky factor with a logarithm transform on the diagonals and assign independent  $N(0, 0.1)$  priors for all elements in this parameterisation. We also use a log-transformation for the single  $\lambda$  and assign  $N(0, 0.1)$ . Finally, we assign independent priors  $d_k \sim N(0, 1)$ , for  $k = 1, \dots, r$ . We have verified that our results are robust to the choice of prior.

### 5.3. Model comparison and fit

We perform model selection using the BIC approximation of the log marginal likelihood (Kass and Raftery, 1995)

$$\log p_{\text{BIC}}(\mathbf{Y}) = \log p(\mathbf{Y}|\hat{\boldsymbol{\theta}}) - \frac{k \log n}{2},$$

where  $k$  is the number of estimated parameters,  $n$  is the length of the time series and  $\hat{\boldsymbol{\theta}}$  is the maximum likelihood estimate obtained by numerical optimisation. We fit all combinations of models for  $p+q \leq 2$  and Table 1 shows the log marginal likelihood for all models considered for the three datasets introduced in Section 5.1. The results show that models with tempered fractional differencing give a better model fit for all datasets and AR/MA orders. The two pure MA models,



**Fig. 9.** Posterior for the spectral density matrix in the VARTFIMA(0,2) model fitted to the Water velocity data. The plots on the diagonal are the marginal spectral densities for each time series. The plots above the diagonal are the squared coherence, and the plots below the diagonal are the time delays from the phase spectrum. The dashed lines in each subplot displays the posterior median and 95% credible intervals from spectral subsampling MCMC. The shaded regions are the 95% credible intervals from the Whittle posterior on the whole dataset.

VARMA(0, 1) and VARMA(0, 2), perform very poorly on the Temperature data, but improve dramatically when tempered fractional differencing is added. According to the BIC approximation of the marginal likelihood, we conclude the following:

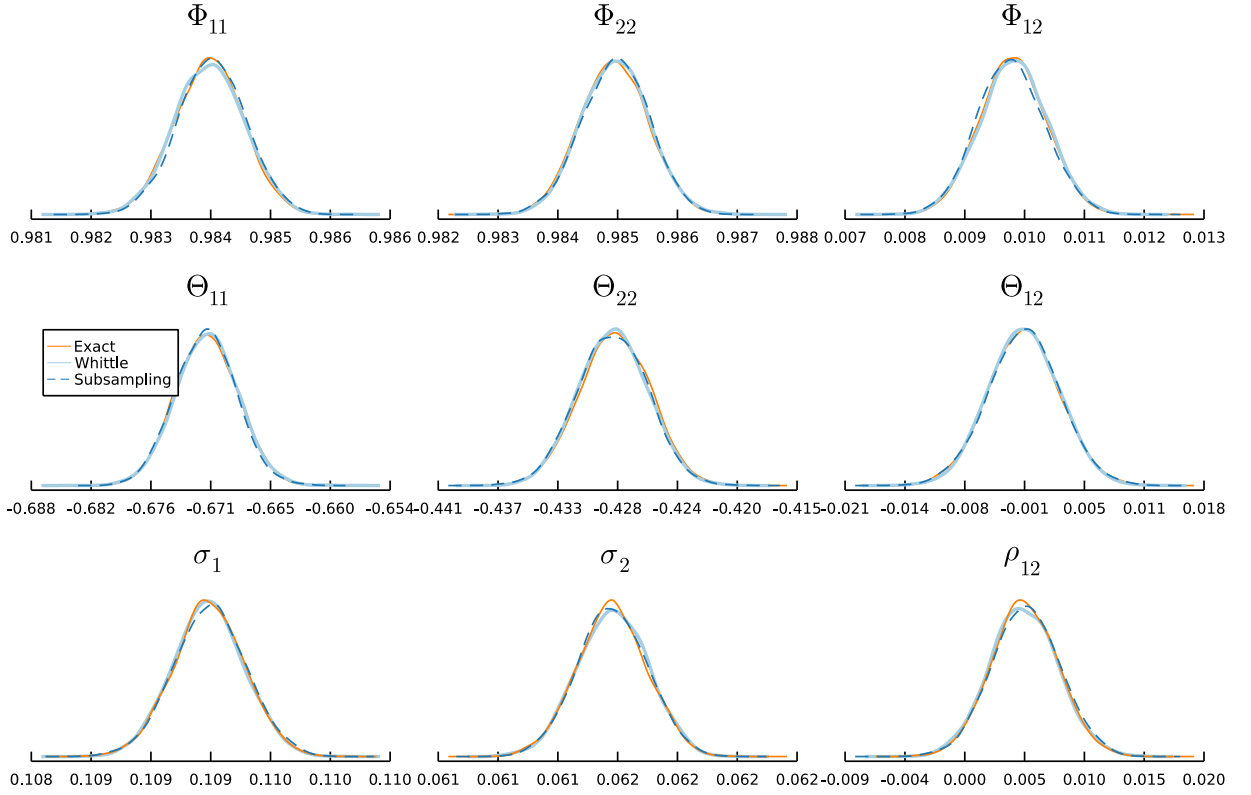
- VARTFIMA(0, 2) (14 parameters) is best for the Water velocity dataset.
- VARTFIMA(2, 0) (28 parameters) is best for the Temperature dataset.
- VARTFIMA(1, 1) and VARTFIMA(2, 0) are best for the Pollution dataset. Both models have 47 parameters.

We implemented spectral subsampling MCMC successfully in all cases except for

- Water velocity: VARMA(0,2).
- Temperature: VARMA(0,1), VARMA(0,2) and VARTFIMA(0,2).
- Pollution: VARMA(0,2), VARMA(1,1), VARTFIMA(2,0), VARTFIMA(1,1).

The reason for the occasional failure of spectral subsampling MCMC for some models is that the control variates do not reduce the variance of the likelihood estimator sufficiently, even after grouping. The situation improves for longer time series, since the control variates typically improve with more data (Quiroz et al., 2019). To illustrate this, we consider the VARTFIMA(0,2) model for the Temperature data, where spectral subsampling was successful. The top left graph in Fig. 5 shows the standard deviation of the log-likelihood estimator over the MCMC iterations, which is the key quantity determining the efficiency of pseudo-marginal algorithms (Pitt et al., 2012); the optimal standard deviation from Lemma S8 in Quiroz et al. (2021) is also indicated in the graph. The variability of the estimator does not exceed the optimal value and spectral subsampling works well, as exemplified by the good mixing of the MCMC chain in the bottom left graph of Fig. 5 for the AR parameter of the first series on its own first lag,  $\Phi_{11}^{(1)}$ . The top right graph in Fig. 5 shows the same quantity, but from spectral subsampling on a posterior based on only the last 60001 data points. The variability of the estimator is initially below the optimal value and the sampler mixes well, but the algorithm eventually moves to a region of the parameter space where the estimator is often much too variable, leading to low acceptance rates and the MCMC chain getting stuck for long spells (bottom right figure).

Since the control variates are based on a Taylor expansion around the posterior mode, they will work poorly when the posterior has multiple modes or a long ridge of fairly constant density. Such shapes are typical in models with local



**Fig. 10.** Kernel density estimates of a subset of the marginal posterior densities for the VARMA(1,1) model fitted to the Water velocity data. The solid orange densities are from MCMC on the exact time domain posterior, the solid blue densities are from MCMC using the Whittle posterior on the whole dataset, and the dashed densities are obtained from spectral subsampling MCMC.

non-identification. VARMA models are, for example, known to have local non-identification problems resulting from near cancellation of roots in the AR and MA polynomials (Chib and Greenberg, 1994), i.e. the likelihood function is flat in the direction of certain linear combination of the parameters. There are also potential identification issues in the MA part of VARMA models (Lütkepohl, 2013, Ch. 7). Adding tempered fractional differencing provides more flexibility, but also additional local identification problems since the fractional differencing parameter  $d$  becomes non-identified when  $\lambda \rightarrow \infty$ . Prior distributions alleviate these identification issues to some extent, but the posteriors in the models considered here are nevertheless very challenging. Luckily, subsampling MCMC fails on the worst fitting models in the Water velocity and Temperature data, but unfortunately also gets stuck on the two best fitting models for the Pollution data. For the Pollution dataset we will therefore only present results using MCMC on the full data without subsampling, but using both the exact time domain likelihood as well as the approximate Whittle likelihood.

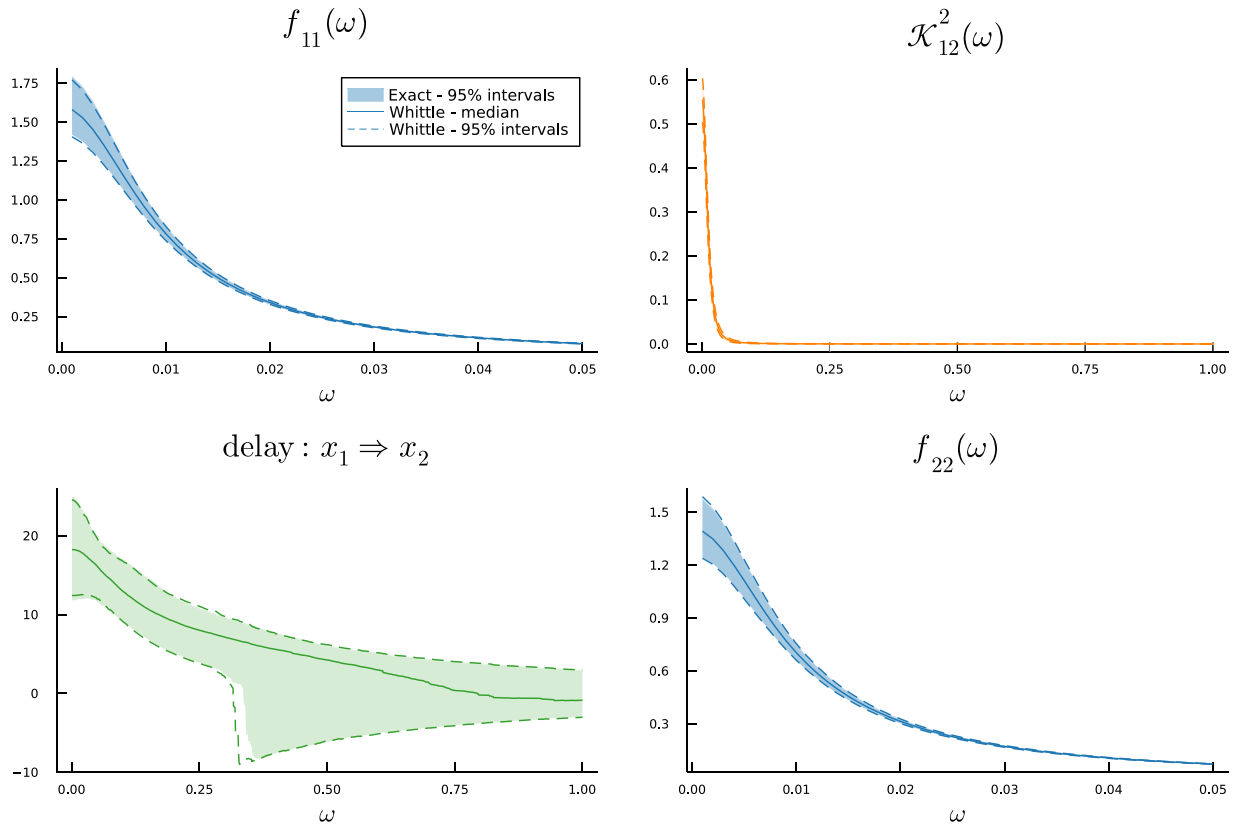
Fig. 6 assesses the fit of the VARTFIMA(0,2) for the Water velocity data by how well the predictive distribution captures the univariate periodogram data. The predictive distribution is computed by, for each posterior draw  $\theta$ , simulating 100 periodogram observations using the Whittle approximation

$$I_{j,T}(\omega) | \theta \stackrel{\text{indep}}{\sim} \text{Expon}(f_{jj,\theta}(\omega)),$$

where  $I_{j,T}$  is the periodogram data for the  $j$ th time series. A nonparametric multitaper estimate (Barbour and Parker, 2014) is also shown as a grey line in the figures. The predictive distribution in Fig. 6 captures the periodogram data quite well. The multitaper estimate suggests that there are some peaks in the spectral density not captured by the VARTFIMA(0,2) model. A VARTFIMA with higher AR and MA orders would be able to also fit those spectral bumps, but the peaks occur at different periods in the two series and may just be an artefact of the large number of observations in the dataset. Figures B.19 and B.20 in Appendix B present the same results for the two other datasets. The figures show that some seasonality seems to have survived the pre-processing of the datasets. An alternative is fitting a multivariate seasonal ARTFIMA model to the original data, but that is not pursued here.

#### 5.4. Efficiency of spectral subsampling MCMC

To measure the computational advantage of subsampling we use a performance measure that takes into account both the cost of estimating the likelihood and the inefficiency of the MCMC chain. The computational time (CT) of an algorithm



**Fig. 11.** Posterior for the spectral density matrix in the VARMA(1,1) model fitted to the Water velocity data. The plots on the diagonal are the marginal spectral densities for each time series. The plot above the diagonal is the squared coherence, and the plot below the diagonal is the phase spectrums. The dashed lines in each subplot displays the posterior median and 95% credible intervals from the Whittle posterior on the whole dataset. The shaded regions are the 95% credible intervals from MCMC on the exact time domain posterior on the whole dataset.

**Table 1**

BIC approximation of the log marginal likelihood for different models for each of the three datasets in Section 5.1. A higher value indicates a better model fit. The AR and MA columns indicate the lag order in the AR and MA component. The No TFI and TFI columns indicate if the process has tempered fractional differencing. The model with the largest marginal likelihood for each dataset is marked in bold font. Both the VARTFIMA(2,0) and VARTFIMA(1,1) are in bold font for the Pollution data since the difference between them is ‘not worth more than a bare mention’ on the modified Jeffreys’ scale of evidence in Kass and Raftery (1995).

AR	MA	Water Velocity		Temperature		Pollution	
		No TFI	TFI	No TFI	TFI	No TFI	TFI
1	0	737079	759123	327097	334122	363760	366022
0	1	588297	759457	61320	332888	306068	365658
2	0	749650	761200	335201	<b>335757</b>	365522	<b>366266</b>
0	2	621765	<b>761786</b>	93256	333948	325717	366142
1	1	758838	761305	333582	335647	365762	<b>366267</b>

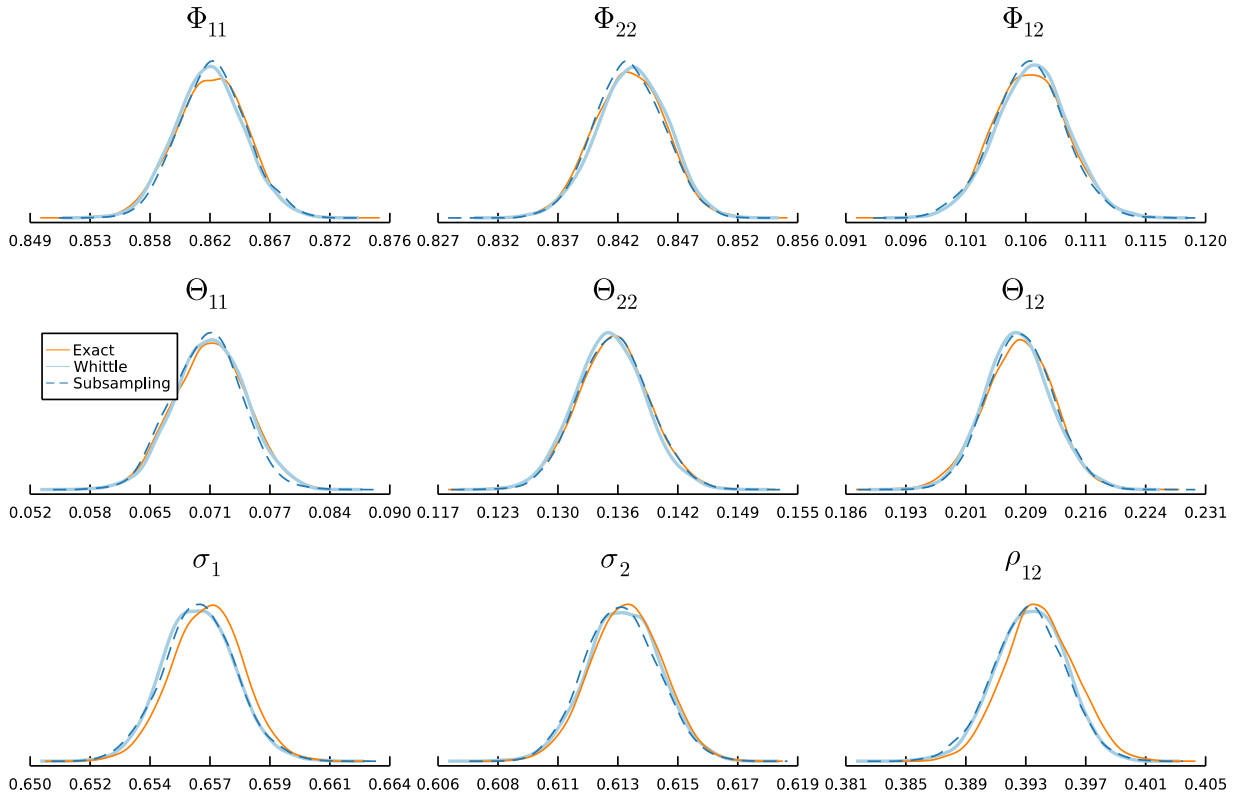
is defined as

$$\text{CT} \equiv \text{IACT} \times \text{Computing time for a single iteration}, \quad (15)$$

where  $\text{IACT} \equiv 1 + 2 \sum_{k=1}^{\infty} \rho_k$  is the integrated autocorrelation time of the MCMC chain, and  $\rho_k$  is the autocorrelation at lag  $k$  of the posterior draws. CT measures the execution time for obtaining the equivalent of a single iid draw from the posterior. To obtain an implementation independent measure, the computing time is set proportional to the number of density evaluations in a run of the algorithm, including the evaluations needed to construct the control variate for subsampling. We use the coda package (Plummer et al., 2006) in R to estimate the integrated auto-correlation time; see Quiroz et al. (2019) for more details.

Table 2 shows the relative computational time (RCT) of MCMC on the full datasets in relation to spectral subsampling MCMC. The RCT is defined as the ratio between the CT of full-data MCMC and that of spectral subsampling MCMC. Hence,





**Fig. 12.** Kernel density estimates of a subset of the marginal posterior densities for the VARMA(1,1) model fitted to the Temperature data. The solid orange densities are from MCMC on the exact time domain posterior, the solid blue densities are from MCMC using the Whittle posterior on the whole dataset, and the dashed densities are obtained from spectral subsampling MCMC.

**Table 2**

Relative computational time (CT) of comparing MCMC using the full dataset to spectral subsampling MCMC. The value 1 indicates that spectral subsampling MCMC and MCMC are equally efficient, and values larger than 1 indicate that spectral subsampling MCMC is the better algorithm. The results for the Pollution dataset are not shown as subsampling MCMC got stuck for the best models for that dataset (VARTFIMA(2,0) and VARTFIMA(1,1), see Table 1).

Dataset	Model	Min	Mean	Max
Water velocity	VARTFIMA(0,2)	87	98	125
Temperature	VARTFIMA(2,0)	68	89	114

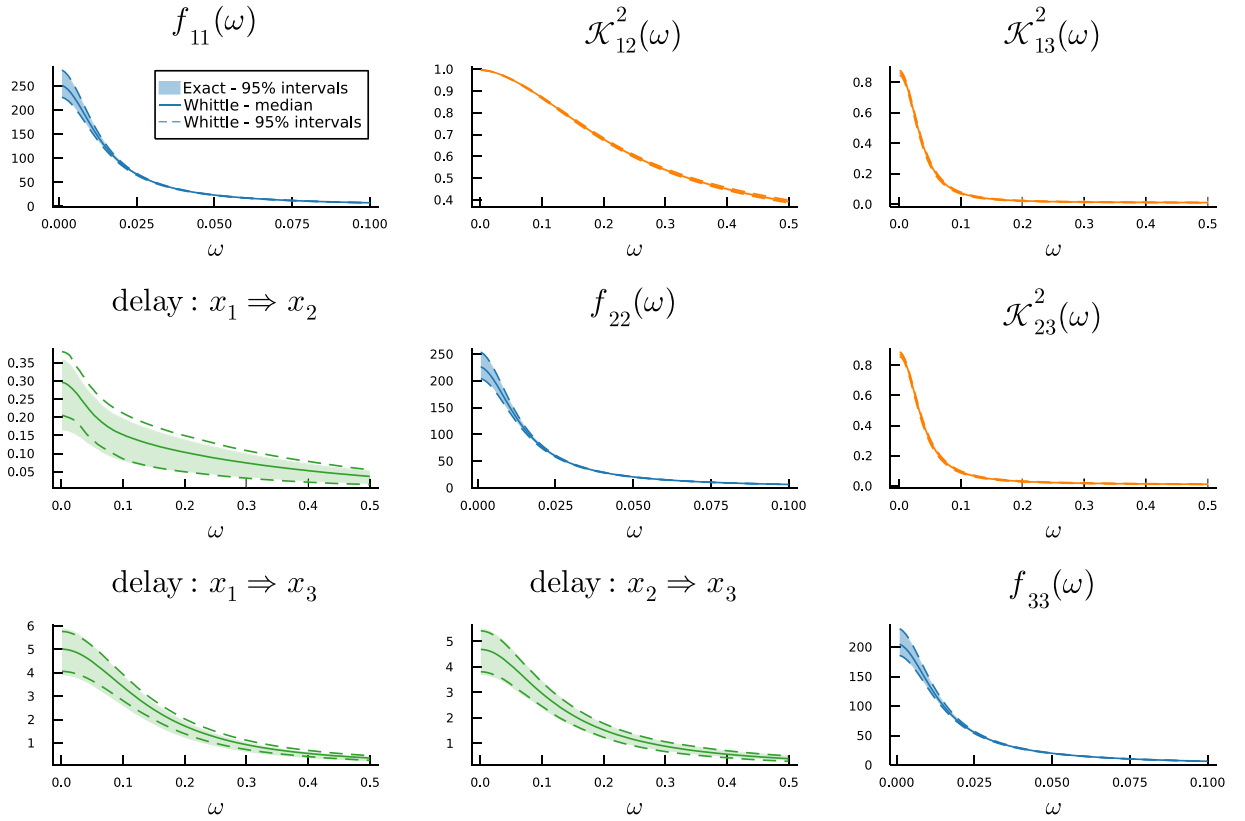
values larger than one mean that spectral subsampling is more efficient when taking into account both the computing cost and the sampling efficiency. The results show that our subsampling algorithm is between 68–125 times faster than MCMC on the full dataset with CT as the measure of computational efficiency. Similar speed-ups are observed on the Pollution dataset for the models where subsampling MCMC did not get stuck.

### 5.5. Accuracy of spectral subsampling for the Whittle posterior

This subsection explores how well spectral subsampling MCMC approximates the posterior based on the Whittle likelihood for the full data. The next section investigates how well the Whittle likelihood for the full data approximates the time domain likelihood in finite samples.

Figure 7 and 8 compare kernel density estimates of the posterior distribution of a subset of the parameters using subsampling MCMC and full-data MCMC. We conclude that the subsampling MCMC algorithm provides very similar answers, and Table 2 shows it is up to two orders of magnitude faster.

The individual parameters in time series models are seldom of practical interest. We will therefore now explore how well our method approximates the posterior of summary quantities based on the spectral density matrix that are often used in



**Fig. 13.** Posterior for the spectral density matrix in the VARMA(1,1) model fitted to the Temperature data. The plots on the diagonal are the marginal spectral densities for each time series. The plot above the diagonal is the squared coherence, and the plot below the diagonal is the phase spectrums. The dashed lines in each subplot displays the posterior median and 95% credible intervals from the Whittle posterior on the whole dataset. The shaded regions are the 95% credible intervals from MCMC on the exact time domain posterior on the whole dataset.

practical work. The *coherence* between two time series  $X_{it}$  and  $X_{jt}$  at frequency  $\omega$  is defined as

$$\mathcal{K}_{ij}(\omega) = \frac{f_{ij}(\omega)}{\sqrt{f_{ii}(\omega)f_{jj}(\omega)}}.$$

The squared coherence

$$0 \leq |\mathcal{K}_{ij}|^2 \leq 1,$$

measures the linear correlation between the pair of time series  $X_{it}$  and  $X_{jt}$  at frequency  $\omega$ . Expressing the complex-valued cross spectral density  $f_{ij}(\omega)$  in polar form as

$$f_{ij}(\omega) = |f_{ij}(\omega)| \exp(i\varphi(\omega)),$$

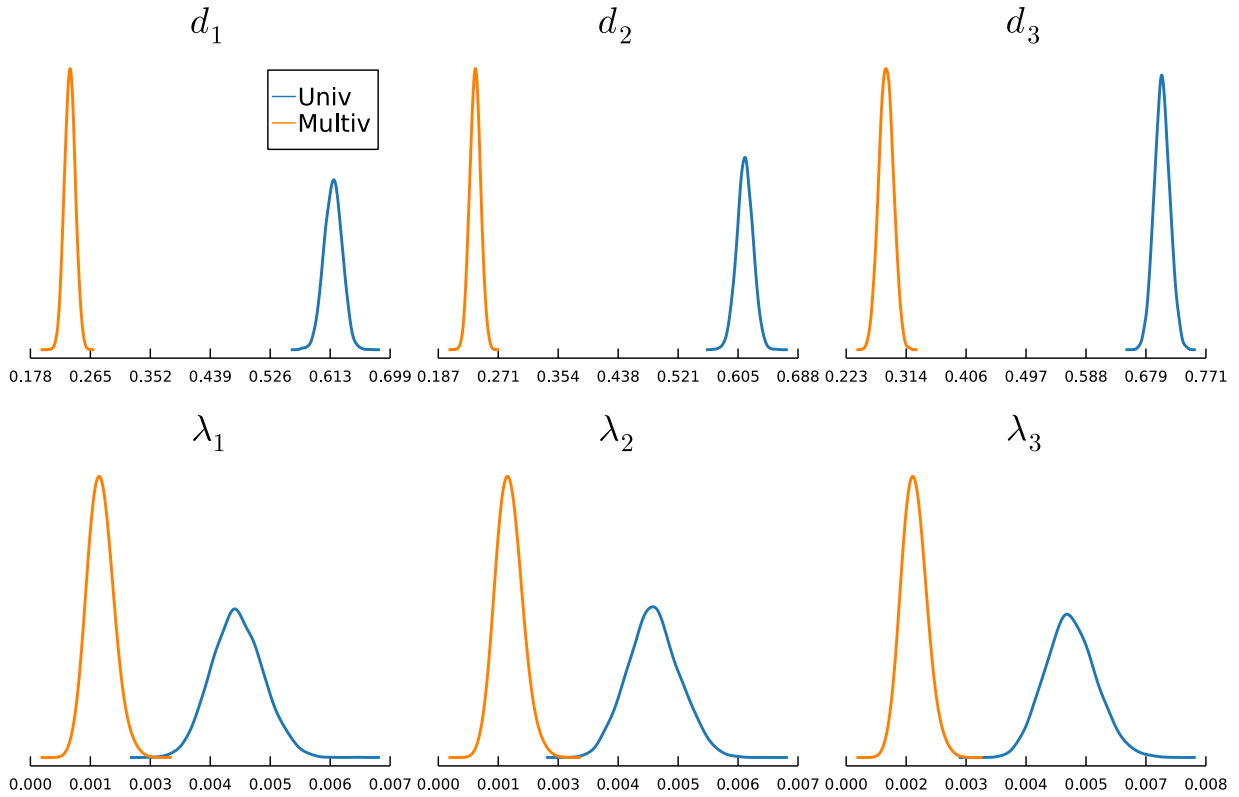
where the phase spectrum

$$\varphi_{ij}(\omega) = \arg(f_{ij}(\omega)),$$


measures the time shift of the signal at frequency  $\omega$ , the *time delay* from variable  $X_{it}$  to variable  $X_{jt}$  is measured by  $-\varphi_{ij}(\omega)/\omega$  (Wei et al., 1990). See Geweke (1982) and Ashby (2019) for connections between these frequency based measures of linear association and direction, and cross-correlation functions and Granger causality measures in the time domain.

Figure 9 plots the posterior distribution of these spectral density matrix quantities for the Water velocity dataset. The figure shows the marginal spectral densities (graphs on the diagonal) of each time series, squared coherence (above diagonal) and phase/time delay (below diagonal). Each figure displays the posterior mean and 95% credible intervals from spectral subsampling MCMC as lines. The 95% credible intervals from MCMC on the Whittle posterior based on the full dataset is shown as shaded regions. As expected from the previously shown parameter posteriors, we again see that subsampling gives virtually no distortion with respect to the Whittle posterior on the full dataset. The same is true for the Temperature data; see Figure B.21 in Appendix B.

The interpretation of the lag delay is clearest in systems without feedback loops. This is likely to be the case only in the Water velocity data where the Denver River location ( $x_2$ ) is located downstreams of the St Clair River location ( $x_1$ ). The top



**Fig. 14.** Comparing the marginal posteriors for the fractional differencing and tempering parameters in the multivariate VARTFIMA(2,0) model (blue) to univariate ARTFIMA models (orange) fitted to the Temperature data. The multivariate model is restricted to have the same  $\lambda$  for each time series and the posterior density for this common  $\lambda$  is therefore repeated in each graph on the second row in the figure. Spectral subsampling MCMC is used for both models.

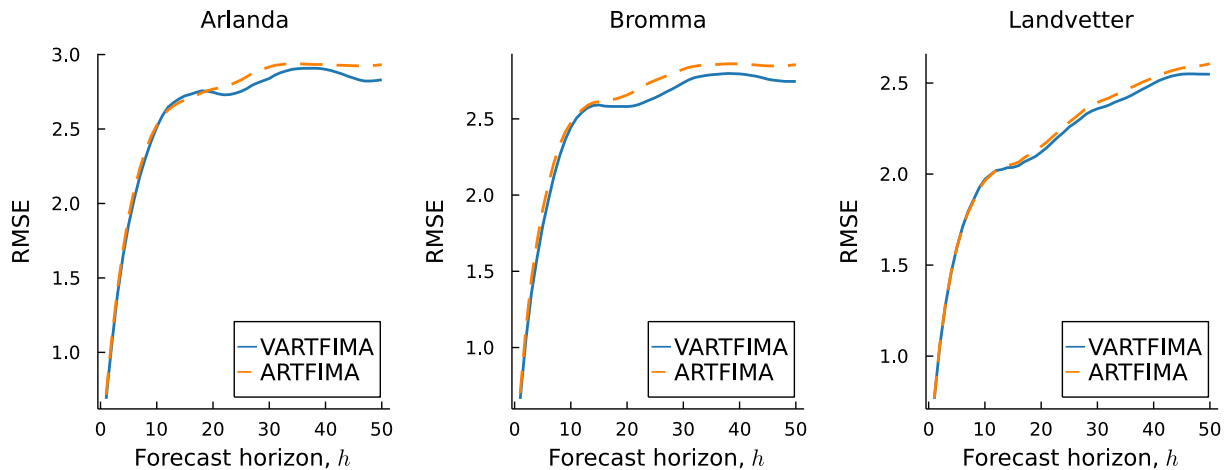


**Fig. 15.** Investigating the convergence of forecasts as  $p_*$  increases in the VARMA( $p_*$ , 0) approximation of the VARTFIMA(2, 0) model. The forecasts are all produced standing at the last time point in the training data.

right graph in Figure 9 shows, however, that there is essentially no coherence between the two locations at any frequency, so the lag delay is insignificant. This is most probably because the two locations are quite far apart (approx 100 kilometers) and separated by Lake St Clair. The VARMA(1,1) model fitted in the next subsection picks up a sizeable coherence only at the very lowest frequencies and a positive delay at those frequencies (see Figure 11), which is the expected sign since Denver River is downstream of St Clair River.

### 5.6. Accuracy of the Whittle posterior for the time domain posterior

The previous subsection shows that spectral subsampling MCMC gives very accurate approximations to the Whittle posteriors based on the full datasets. The Whittle likelihood can, however, be a poor approximation to the exact time domain



**Fig. 16.** Comparing the out-of-sample RMSE forecasting performance of the VARTFIMA(2,0) model (solid blue) to univariate ARTFIMA(2,0) models fitted to each series (dashed orange) in the Temperature data. Both models were fitted to the training data (Jan 3, 2016 – to Dec 21, 2018) and the estimates were then kept fixed throughout the forecasting period (Dec 21, 2018 – Dec 29, 2018).

likelihood (Contreras-Cristán et al., 2006), at least for shorter time series. However, the original motivation for spectral subsampling MCMC is settings where the time series are very long, and this section demonstrates that the Whittle approximation is excellent for the three datasets in Section 5.1.

The time domain likelihood is only computationally feasible for VARMA models, and hence we illustrate these results using a VARMA(1, 1) model for each dataset. Figure 10 shows the kernel density estimates of the posterior distribution of a subset of the parameters for the Water velocity data. Figure 11 plots the time domain and Whittle posteriors of the spectral density matrix based on the full Water velocity dataset. The results from the Whittle posterior are nearly indistinguishable from the exact time domain posteriors. Also, there are virtually no differences between the time domain posteriors and the Whittle posteriors in the two other datasets; see Figure 12 and 13 for the Temperature data and Figure B.23 in Appendix B for the Pollution data. In summary, we conclude from this and the previous subsection that spectral subsampling MCMC gives an accurate approximation of the posterior based on the exact time domain likelihood.

We have implemented the exact time domain likelihood for VARMA models using the state space representation in the Python package `statsmodels` and fair timing comparisons with spectral subsampling MCMC are therefore difficult, but it is well known that time domain likelihoods are substantially more costly than their frequency domain (Whittle) counterparts, even without the extra speed-up from subsampling.

## 6. Comparing ARTFIMA and VARTFIMA

Our paper focuses on the computational performance of spectral subsampling MCMC on challenging multivariate processes, and we leave a detailed analysis of the empirical performance of the VARTFIMA model to a separate paper. However, a reviewer suggested comparing the proposed multivariate VARTFIMA model with univariate ARTFIMA models (Sabzikar et al., 2019) for each series separately, and we present some results on this aspect below using the Temperature data.

First, we compute the BIC approximation of the log marginal likelihood for a univariate version of the preferred VARTFIMA(2,0) model by restricting the diagonals of all  $\Phi$  and  $\Theta$  and  $\Sigma$  to zero in the estimation, but allowing for different tempering parameter and fractional differencing for each series. This corresponds to fitting ARTFIMA models to each series separately, but still obtaining a single BIC that approximates the marginal likelihood of all time series; these BIC values are therefore comparable to the ones from the multivariate VARTFIMA models. The BIC for this univariate version is 312824 which should be compared to the BIC of 335757 for the multivariate VARTFIMA(2,0) in Table 1. The massive improvement in BIC from jointly modeling the time series in the VARTFIMA(2,0) model is expected since the series are cross-correlated; see Figure B.21.

Furthermore, the fractional differencing parameters can be quite different when modeling the time series jointly as in VARTFIMA compared to univariate ARTFIMA models. Figure 14 compares the marginal posteriors of the fractional differencing and tempering parameters in the multivariate VARTFIMA(2,0) model to univariate ARTFIMA(2,0) models in the Temperature data. The VARTFIMA model in Figure 14 is restricted to have the same  $\lambda$  in the three series; the results are quite similar when separate  $\lambda$  are used. Both the tempering and the fractional differencing parameters are substantially reduced for all three series when the multivariate model is used.

Finally, we compare the forecasting performance of VARTFIMA(2,0) model to the univariate ARTFIMA(2,0). The models are estimated on the same data as used above, i.e. the period from February 1, 2008 at 00:00 hours to March 25, 2022 at 16:00 hours, and the forecasts are evaluated over subsequent period from March 25, 2022 at 17:00 hours to May 1, 2022 at

6:00 hours, making up a test set with 878 number of observations. Since the training data is very large and the estimates are precise, we do not update the estimates as we move across the test data.

There are several ways of computing the forecasts from the VARTFIMA model. One way is by computing the autocovariance matrix function from an inverse FFT of the spectral density matrix at the estimated VARTFIMA parameter and then using the conditioning properties of the multivariate normal distribution to get the forecast, see [Hamilton \(1995\)](#) for the theory and [McLeod et al. \(2008\)](#) for an R package implementation. This is an elegant and general approach, but does not scale well to large data sets: the `artfima` package in R, which uses the [McLeod et al. \(2008\)](#) R package for forecasting, crashes when using more than 20% of the Temperature data on a 32 GB RAM Linux machine.

We take a more direct and computationally faster approach here by approximating the  $\text{VARTFIMA}(p, q)$  with a  $\text{VARMA}(p_*, q)$  for a sufficiently large  $p_*$ . The  $\text{VARTFIMA}(p, q)$  model is

$$(I_r - \Phi_1 L - \dots - \Phi_p L^p) \Delta^{\mathbf{d}, \lambda} (\mathbf{Y}_t - \boldsymbol{\mu}) = \Theta(L) \boldsymbol{\varepsilon}_t, \quad (16)$$

where from the definition of  $\Delta^{\mathbf{d}, \lambda}$  we can express

$$\Delta^{\mathbf{d}, \lambda} = I_r + B_1 L + B_2 L^2 + \dots \quad (17)$$

with  $B_j = \text{Diag}(b_j^{d_1, \lambda_1}, \dots, b_j^{d_r, \lambda_r})$  and  $b_j^{d_k, \lambda_k} \equiv (-1)^j \binom{d_k}{j} e^{-\lambda_k j}$  for  $j = 1, 2, \dots$ . We can therefore write

$$(I_r - \Phi_1 L - \dots - \Phi_p L^p) (I_r + B_1 L + B_2 L^2 + \dots) = I_r - \Pi_1 L - \Pi_2 L^2 - \dots, \quad (18)$$

where the  $\Pi_k$  can be found by equating the coefficients term by term for each  $L^k$  to obtain

$$\Pi_k = \sum_{j=1}^p \Phi_j B_{k-j} - B_k, \quad (19)$$

by defining  $B_0 = I_r$  and  $B_j = \mathbf{0}$  for  $j < 0$ . Since  $b_j^{d_k, \lambda_k} \rightarrow 0$  as  $j \rightarrow \infty$  exponentially fast eventually, if  $\lambda_k > 0$  ([Sabzikar et al., 2019](#)) we have that the elements of  $\Pi_k$  goes to zero as  $k$  increases. We can therefore truncate this infinite order lag matrix polynomial at some lag  $p_*$  to get a  $\text{VARMA}(p_*, q)$  approximation

$$\Pi(L) (\mathbf{Y}_t - \boldsymbol{\mu}) = \Theta(L) \boldsymbol{\varepsilon}_t, \quad (20)$$

where  $I_r - \Pi_1 L - \dots - \Pi_{p_*} L^{p_*}$ . Forecasts  $k$  steps ahead can now be obtained with established techniques for VARMA models ([Tsay, 2013](#)); we use the MTS package in R. [Figure 15](#) shows that  $p_* = 10$  is more than sufficient for approximating the  $\text{VARTFIMA}(2, 0)$  model used in the forecast evaluation below on the Temperature data in that the forecast do not change beyond  $p_* = 10$ .

[Figure 16](#) shows that the multivariate  $\text{VARTFIMA}(2, 0)$  produces more accurate forecasts at all horizons than univariate  $\text{ARTFIMA}(2, 0)$  models fitted to each time series separately, particularly at the longer horizons; note that a forecast horizon with  $h = 50$  corresponds to 10 hours ahead.

## 7. Conclusions

Our paper proposes a subsampling MCMC approach for stationary multivariate time series models. Using a measure which takes into account both computing cost and the statistical inefficiency of likelihood estimators, we demonstrate a speed-up factor of up to two orders of magnitude on three datasets compared to MCMC using the full dataset.

To test the proposed spectral subsampling MCMC in challenging problems, we propose a new multivariate time series model by extending the univariate ARTFIMA model to a multivariate setting. Some properties of this new model are derived, including its spectral density matrix, and our results show that the vector ARTFIMA model outperforms VARMA models in all three datasets. This suggest the VARTFIMA model as a useful model for modelling a wide range of multivariate time series and future work should investigate this more fully.

Our work demonstrates that MCMC sampling to explore the posterior based on the Whittle likelihood scales well to very long multivariate time series, and that the approximation to the exact time domain posterior is excellent. We further show that spectral subsampling can give an additional speed-up of two orders of magnitude on large time series in rather complex multivariate time series models with semi-long range dependence. The biggest challenge is to obtain good control variates in high-dimensional parameter spaces for models with potentially very non-Gaussian likelihoods. We are unable to obtain good enough control variates for the best fitting models on the shortest of the three dataset. This is probably the result of the VARMA and VARTFIMA class of models having challenging likelihood functions, especially the moving average (MA) part of the model. We demonstrate that longer time series are more amenable to subsampling than shorter ones, so spectral subsampling MCMC is a method most suitable to large scale problems, at least when it comes to VARMA/VARTFIMA models. Future research will extend the methodology to high-dimensional complex models by developing better control variates in high-dimensions that consider the local nonidentification aspects, and exploring the potential of alternative inference algorithms such as variational inference ([Blei et al., 2017](#)). Variational inference with an estimated likelihood is less sensitive to the variability in the estimator ([Tran et al., 2017](#)), and is thus an appealing alternative to explore.



## Declaration of Competing Interest

All authors declare that they have no conflict of interest regarding the submission of the manuscript to Econometrics and Statistics.

## Acknowledgment

The authors are grateful for comments from two reviewers and the associate editor that improved both the content and the presentation of the paper. Mattias Villani was supported by grants from the Swedish Research Council (2020-02846) and the Swedish Foundation for Strategic Research (Smart Systems: RIT 15-0097). Robert Kohn was partially supported by the Australian Research Council (IC190100031, DP210103873).

## Appendix A. Proofs

**Proof of Theorem 1.** Theorem 11.3.1 of Brockwell and Davis (1991) shows that the condition  $\Phi(z)$  for  $|z| \leq 1$  implies that there exists  $\epsilon > 0$  such that the matrix  $\Phi^{-1}(z)$  exists and has a power series expansion

$$\Phi^{-1}(z) = \sum_{j=0}^{\infty} \mathbf{A}_j z^j, \text{ for } |z| < 1 + \epsilon,$$

with the elements of  $\mathbf{A}_j$  being absolutely summable. Inverting the operator  $\Delta^{\mathbf{d}, \lambda}$  gives

$$\Delta^{-\mathbf{d}, \lambda} \mathbf{Y}_t = \left( \Delta^{-d_1, \lambda_1} Y_{1,t}, \dots, \Delta^{-d_r, \lambda_r} Y_{r,t} \right)^{\top},$$

where the inverted univariate operator is defined by (Sabzikar et al., 2019)

$$\Delta^{-d, \lambda} Y_t = (1 - e^{-\lambda} L)^{-d} Y_t = \sum_{j=0}^{\infty} c_j^{-d, \lambda} Y_{t-j},$$

with  $c_j^{-d, \lambda} = (-1)^j e^{-\lambda j} \binom{-d}{j}$ . Let  $\Delta^{-\mathbf{d}, \lambda}(z)$  be the  $r \times r$  diagonal matrix with  $k$ th diagonal element equal to  $(1 - e^{-\lambda_k} z)^{-d_k}$ . Direct matrix multiplication shows that the element in row  $k$ , column  $l$  of the matrix  $\Delta^{-\mathbf{d}, \lambda}(z) \Phi^{-1}(z)$  is the product

$$\left( \sum_{i=0}^{\infty} c_i^{-d_k, \lambda_k} z^i \right) \left( \sum_{j=0}^{\infty} a_j^{(kl)} z^j \right) = \sum_{v=0}^{\infty} b_v^{(kl)} z^v,$$

where  $a_j^{(kl)}$  is the element in row  $k$ , column  $l$  of  $\mathbf{A}_j$  and  $b_v^{(kl)} = \sum_{s=0}^v c_s^{-d_k, \lambda_k} a_{v-s}^{(kl)}$ . Sabzikar et al. (2019)[Proof of Theorem 2.2a] show that the sequence  $c_j^{-d, \lambda}$  is absolutely summable if  $\lambda > 0$  and  $d \notin \mathbb{Z}$ , and it is shown above that  $\sum_{j=0}^{\infty} |a_j^{(kl)}| < \infty$ . Since the product of two absolutely summable series is absolutely summable (Knopp, 1990, Ch. 4.17),  $\sum_{v=0}^{\infty} b_v^{(kl)}$  is also absolutely summable for all  $k, l$ .

Hence, by Proposition 3.1.1 in Brockwell and Davis (1991) we can apply the operator  $\Delta^{-\mathbf{d}, \lambda} \Phi^{-1}(L)$  to both sides of

$$\Phi(L) \Delta^{\mathbf{d}, \lambda} (\mathbf{Y}_t - \boldsymbol{\mu}) = \Theta(L) \boldsymbol{\varepsilon}_t \quad (\text{A.1})$$

to obtain

$$\mathbf{Y}_t = \boldsymbol{\mu} + \Psi(L) \boldsymbol{\varepsilon}_t, \quad (\text{A.2})$$

where

$$\Psi(z) \equiv \sum_{j=0}^{\infty} \Psi_j z^j = \Delta^{-\mathbf{d}, \lambda}(z) \Phi^{-1}(z) \Theta(z), \quad \text{for } |z| \leq 1,$$

and the  $\{\Psi_j\}_{j=0}^{\infty}$  are absolutely summable elementwise since  $\Theta(z)$  is a polynomial of finite order  $q$ . By Proposition 3.1.2 in Brockwell and Davis (1991) the VARTFIMA process is therefore causal and stationary.


The causal representation in (A.2) shows that  $\{\Psi_j\}_{j=0}^{\infty}$  acts as a time invariant linear filter on the iid sequence  $\{\boldsymbol{\varepsilon}_t\}_{t=0}^{\infty}$ . The spectral density of  $\mathbf{Y}_t$  therefore follows from Brockwell and Davis (1991, Theorem 11.8.3) and is

$$\begin{aligned} f_{\mathbf{Y}}(\omega) &= \Psi(e^{-i\omega}) f_{\boldsymbol{\varepsilon}}(\omega) \Psi(e^{-i\omega})^H \\ &= \frac{1}{2\pi} \Delta^{-\mathbf{d}, \lambda}(e^{-i\omega}) \Phi^{-1}(e^{-i\omega}) \Theta(e^{-i\omega}) \Sigma_{\boldsymbol{\varepsilon}} \Theta^H(e^{-i\omega}) \Phi^{-H}(e^{-i\omega}) \Delta^{-\mathbf{d}, \lambda}(e^{-i\omega})^H, \end{aligned}$$

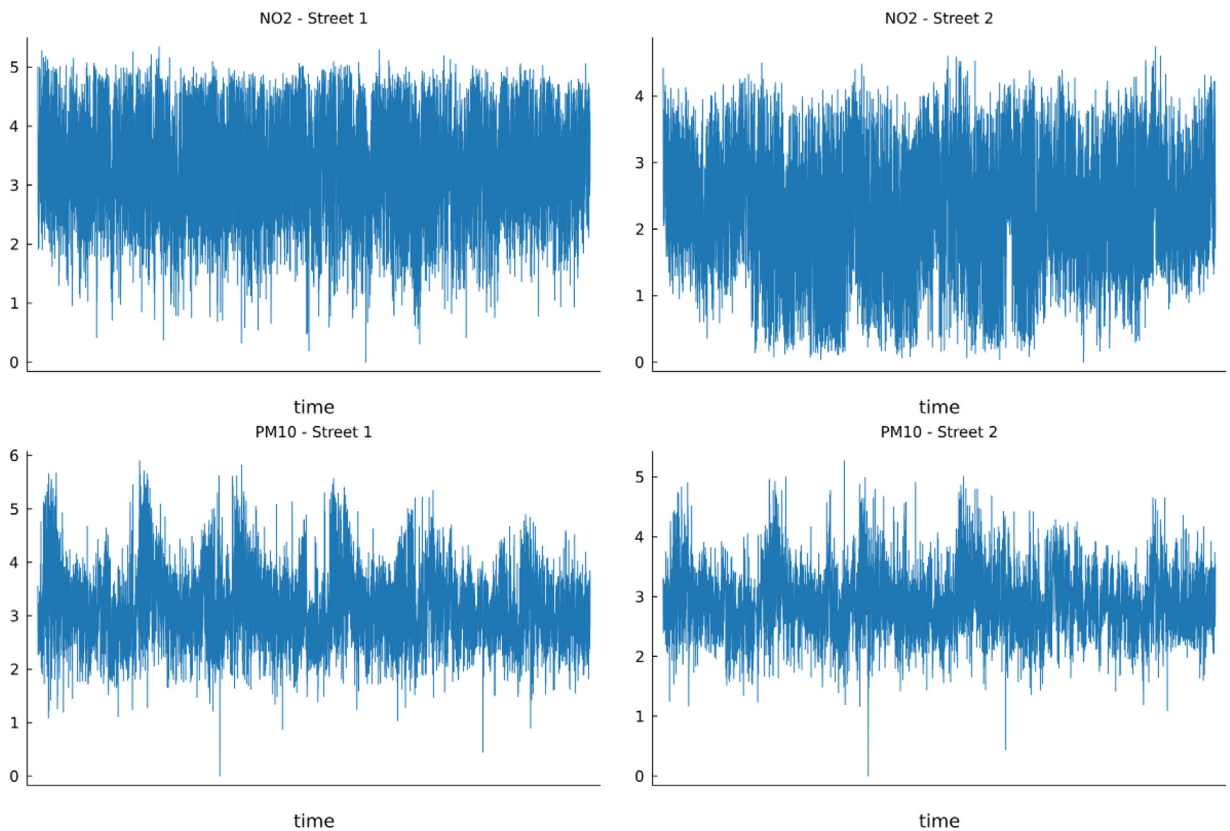
where  $f_{\boldsymbol{\varepsilon}}(\omega) = \Sigma_{\boldsymbol{\varepsilon}}$  is the spectral density matrix of the white noise process  $\{\boldsymbol{\varepsilon}_t\}$  and  $\Delta^{-\mathbf{d}, \lambda}(z) = \text{Diag}((1 - e^{-\lambda_1} z)^{-d_1}, \dots, (1 - e^{-\lambda_r} z)^{-d_r})$ .  $\square$

## Appendix B. Raw data and additional empirical results


This appendix contains graphs of the raw data before detrending and deseasoning as well as additional empirical results referred to in the main text.



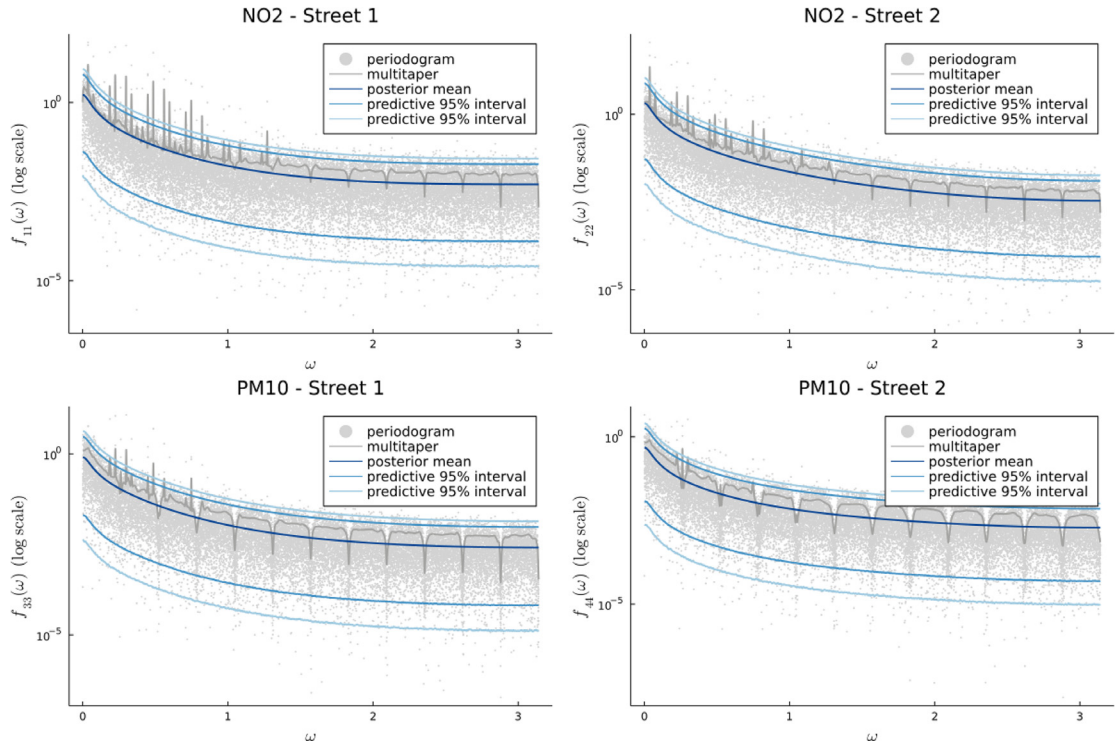
**Fig. B.17.** Temperature data after interpolation, but before deseasoning.



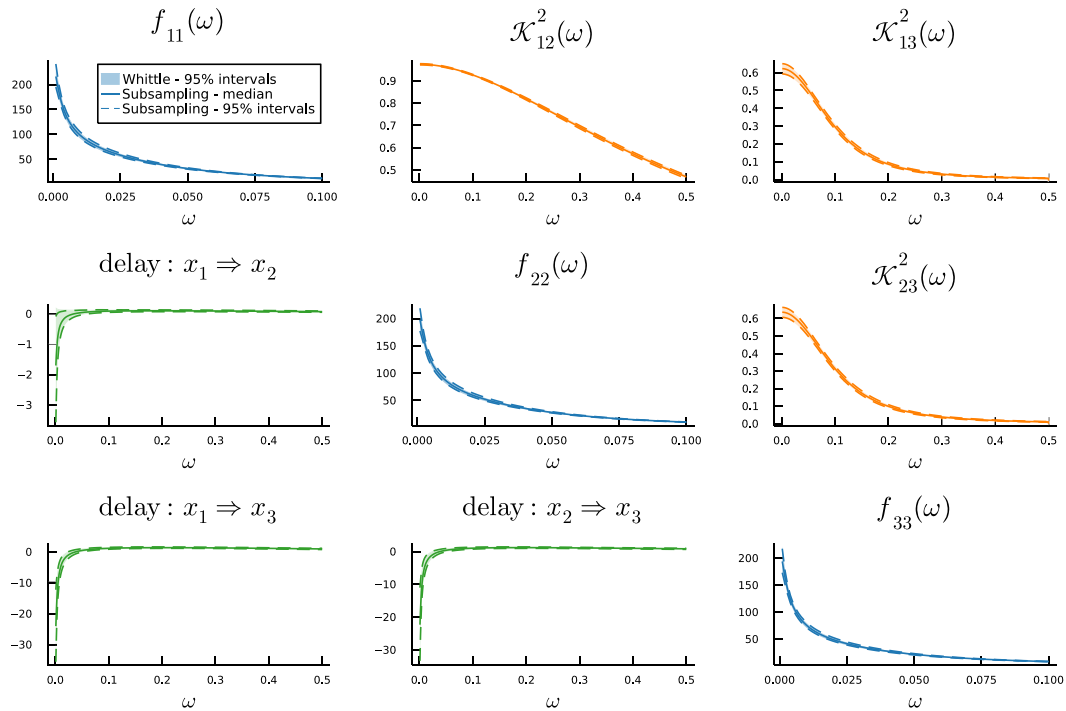
**Fig. B.18.** Stockholm air pollution data after interpolation and logarithmic transform, but before deseasoning.



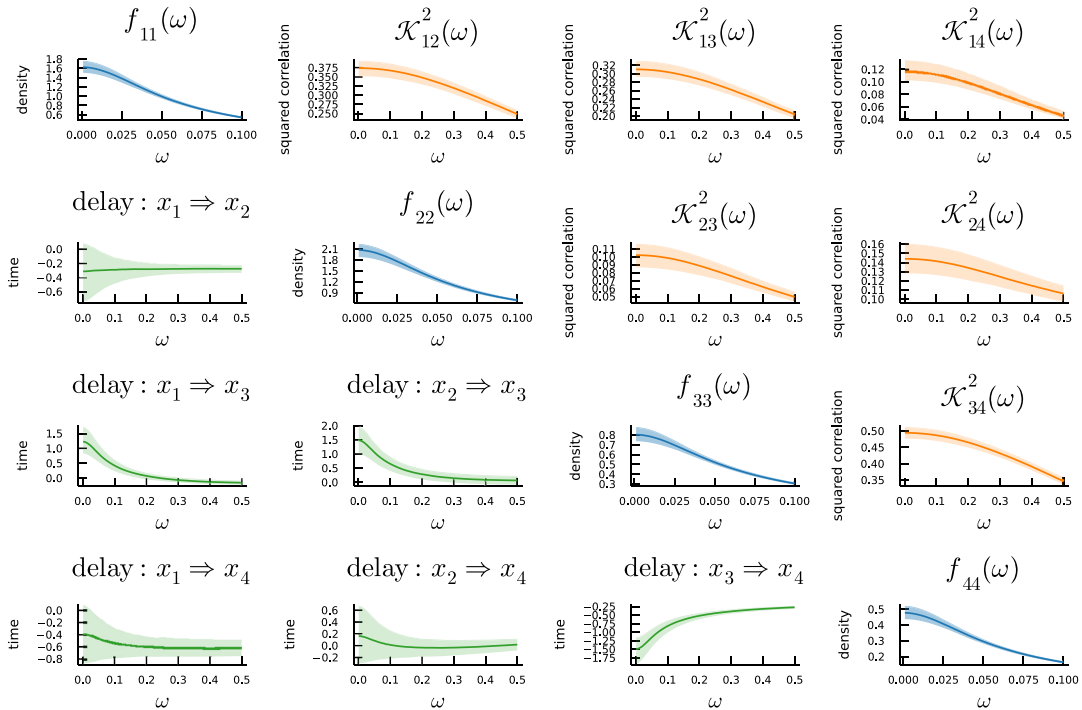
**Fig. B.19.** Posterior predictive fit of the univariate periodogram data for the VARTFIMA(2,0) model fitted to the Temperature data. The predictive intervals are obtained by simulation from the asymptotic Whittle distribution  $\mathcal{I}(\omega) \sim \text{Expon}(f_\theta(\omega))$  with parameters  $\theta$  drawn from the VARTFIMA(1,1) posterior.



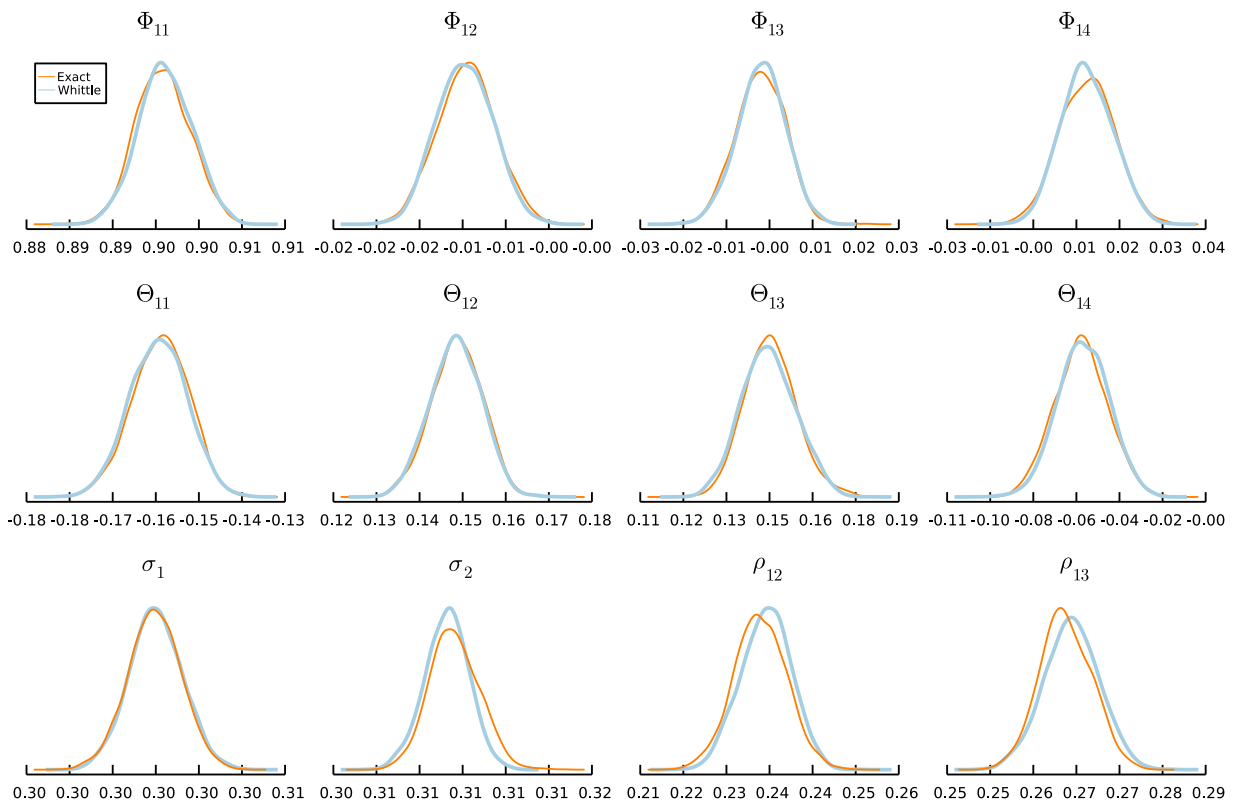
**Fig. B.20.** Posterior predictive fit of the univariate periodogram data for the VARTFIMA(2,0) model fitted to the Pollution data. The predictive intervals are obtained by simulation from the asymptotic Whittle distribution  $\mathcal{I}(\omega) \sim \text{Expon}(f_\theta(\omega))$  with parameters  $\theta$  drawn from the VARTFIMA(1,1) posterior.



**Fig. B.21.** Posterior for the spectral density matrix in the VARTFIMA(2,0) model fitted to the Temperature data. The plots on the diagonal are the marginal spectral densities for each time series. The plots above the diagonal are the squared coherences, and the plots below the diagonal are the time delays from the phase spectrum. The dashed lines in each subplot display the posterior median and 95% credible intervals from spectral subsampling MCMC. The shaded regions are the 95% credible intervals from the Whittle posterior on the whole dataset.



**Fig. B.22.** Posterior for the spectral density matrix in the VARTFIMA(2,0) model fitted to the Pollution data using the Whittle posterior on the whole dataset. The plots on the diagonal are the marginal spectral densities for each time series. The plots above the diagonal are the squared coherence, and the plots below the diagonal are the time delays from the phase spectrum. The posterior mean is marked out with a solid line and the shaded regions are the 95% credible intervals.



**Fig. B.23.** Kernel density estimates of a subset of the marginal posterior densities for the VARMA(1,1) model fitted to the Pollution data. The solid orange densities are from MCMC on the exact time domain posterior and the solid blue densities are from MCMC using the Whittle posterior on the whole dataset.

## References

- Andrieu, C., Roberts, G.O., 2009. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics* 37 (2), 697–725.
- Ansley, C.F., Kohn, R., 1986. A note on reparameterizing a vector autoregressive moving average model to enforce stationarity. *Journal of Statistical Computation and Simulation* 24 (2), 99–106.
- Ashby, F.G., 2019. *Statistical Analysis of fMRI Data*. MIT press.
- Barbour, A.J., Parker, R.L., 2014. psd: Adaptive, sine multitaper power spectral density estimation for R. *Computers & Geosciences* 63, 1–8.
- Bardenet, R., Doucet, A., Holmes, C., 2017. On Markov chain Monte Carlo methods for tall data. *The Journal of Machine Learning Research* 18 (1), 1515–1557.
- Blei, D.M., Kucukelbir, A., McAuliffe, J.D., 2017. Variational inference: A review for statisticians. *Journal of the American Statistical Association* 112 (518), 859–877.
- Brillinger, D.R., 2001. *Time Series: Data Analysis and Theory*. SIAM.
- Brockwell, P.J., Davis, R.A., 1991. *Time Series: Theory and Methods*. Springer.
- Chen, J.E., Polimeni, J.R., Bollmann, S., Glover, G.H., 2019. On the analysis of rapidly sampled fMRI data. *Neuroimage* 188, 807–820.
- Chib, S., Greenberg, E., 1994. Bayes inference in regression models with ARMA(p,q) errors. *Journal of Econometrics* 64 (1–2), 183–206.
- Contreras-Cristán, A., Gutiérrez-Peña, E., Walker, S.G., 2006. A note on Whittle's likelihood. *Communications in Statistics-Simulation and Computation* 35 (4), 857–875.
- Dang, K.-D., Quiroz, M., Kohn, R., Minh-Ngoc, T., Villani, M., 2019. Hamiltonian Monte Carlo with energy conserving subsampling. *Journal of Machine Learning Research* 20 (100), 1–31.
- Deisenroth, M.P., Fox, D., Rasmussen, C.E., 2013. Gaussian processes for data-efficient learning in robotics and control. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37 (2), 408–423.
- Deligiannidis, G., Doucet, A., Pitt, M.K., 2018. The correlated pseudomarginal method. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 80 (5), 839–870.
- Doan, T., Litterman, R., Sims, C., 1984. Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews* 3 (1), 1–100.
- Geweke, J., 1982. Measurement of linear dependence and feedback between multiple time series. *Journal of the American Statistical Association* 77 (378), 304–313.
- Goodfellow, I., Bengio, Y., Courville, A., 2016. *Deep Learning*. MIT press.
- Granger, C.W.J., Joyeux, R., 1980. An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis* 1 (1), 15–29.
- Hamilton, J.D., 1995. *Time series analysis. Economic Theory*. II, Princeton University Press, USA 625–630.
- Kass, R.E., Raftery, A.E., 1995. Bayes factors. *Journal of the American Statistical Association* 90 (430), 773–795.
- Knopp, K., 1990. *Theory and Application of Infinite Series*. Courier Corporation.
- Lütkepohl, H., 2013. *Introduction to Multiple Time Series Analysis*. Springer Science & Business Media.
- McLeod, A.I., Yu, H., Krougly, Z.L., 2008. Algorithms for linear time series analysis: With r package. *Journal of Statistical Software* 23, 1–26.
- Mykland, P.A., Zhang, L., 2012. The econometrics of high frequency data. *Statistical Methods for Stochastic Differential Equations* 124, 109.
- Pitt, M.K., dos Santos Silva, R., Giordani, P., Kohn, R., 2012. On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics* 171 (2), 134–151.



- Plummer, M., Best, N., Cowles, K., Vines, K., 2006. CODA: Convergence diagnosis and output analysis for MCMC. *R News* 6 (1), 7–11.
- Quiroz, M., Kohn, R., Villani, M., Tran, M.-N., 2019. Speeding up MCMC by efficient data subsampling. *Journal of the American Statistical Association* 114 (526), 831–843.
- Quiroz, M., Tran, M.-N., Villani, M., Kohn, R., 2018a. Speeding up MCMC by delayed acceptance and data subsampling. *Journal of Computational and Graphical Statistics* 27 (1), 12–22.
- Quiroz, M., Tran, M.-N., Villani, M., Kohn, R., Dang, K.-D., 2021. The block-Poisson estimator for optimally tuned exact subsampling MCMC. *Journal of Computational and Graphical Statistics* 30 (4), 877–888.
- Quiroz, M., Villani, M., Kohn, R., Tran, M.-N., Dang, K.-D., 2018b. Subsampling MCMC-An introduction for the survey statistician. *Sankhya A* 80 (1), 33–69.
- Ratnarajah, T., Vaillancourt, R., 2005. Complex singular Wishart matrices and applications. *Computers & Mathematics with Applications* 50 (3–4), 399–411.
- Sabzikar, F., McLeod, A.I., Meerschaert, M.M., 2019. Parameter estimation for ARTFIMA time series. *Journal of Statistical Planning and Inference* 200, 129–145.
- Salomone, R., Quiroz, M., Kohn, R., Villani, M., Tran, M.-N., 2020. Spectral subsampling MCMC for stationary time series. In: *Proceedings of the 37th international conference on machine learning*, vol. 119. PMLR, pp. 8449–8458.
- Srivastava, M.S., 2003. Singular Wishart and multivariate beta distributions. *The Annals of Statistics* 31 (5), 1537–1560.
- Tran, M.-N., Kohn, R., Quiroz, M., Villani, M., 2016. The block pseudo-marginal sampler. *arXiv preprint arXiv:1603.02485*.
- Tran, M.-N., Nott, D.J., Kohn, R., 2017. Variational Bayes with intractable likelihood. *Journal of Computational and Graphical Statistics* 26 (4), 873–882.
- Tsay, R.S., 2013. *Multivariate time series analysis: with R and financial applications*. John Wiley & Sons.
- Uhlig, H., 1994. On singular Wishart and singular multivariate beta distributions. *The Annals of Statistics* 395–405.
- Wej, William, S. W., 1990. *Time Series Univariate and Multivariate Methods*. Addison Wesley Publishing Company Inc..
- Welch, T.F., Widita, A., 2019. Big data in public transportation: A review of sources and methods. *Transport Reviews* 39 (6), 795–818.
- Whittle, P., 1953. The analysis of multiple stationary time series. *Journal of the Royal Statistical Society: Series B (Methodological)* 15 (1), 125–139.