

Bayesian Feature Selection in Joint Quantile Time Series Analysis*

Ning Ning[†]

Abstract. Quantile feature selection over correlated multivariate time series data has always been a methodological challenge and is an open problem. In this paper, we propose a general Bayesian dimension reduction methodology for feature selection in high-dimensional joint quantile time series analysis, under the name of the quantile feature selection time series (QFSTS) model. The QFSTS model is a general structural time series model, where each component yields an additive contribution to the time series modeling with direct interpretations. Its flexibility is compound in the sense that users can add/deduct components for each time series and each time series can have its own specific valued components of different sizes. Feature selection is conducted in the quantile regression component, where each time series has its own pool of contemporaneous external predictors allowing nowcasting. Bayesian methodology in extending feature selection to the quantile time series research area is developed using multivariate asymmetric Laplace distribution, spike-and-slab prior setup, the Metropolis-Hastings algorithm, and the Bayesian model averaging technique, all implemented consistently in the Bayesian paradigm. The QFSTS model requires small datasets to train and converges fast. Extensive examinations confirmed that the QFSTS model has superior performance in feature selection, parameter estimation, and forecast.

Keywords: Bayesian inference, quantile feature selection, dimension reduction, multivariate time series analysis.

MSC2020 subject classifications: Primary 62F15, 62M10; secondary 62H86.

1 Introduction

We firstly give the background and motivation in Section 1.1 and then state our contributions in Section 1.2, followed with the organization of the paper in Section 1.3.

1.1 Background and Motivation

In the era of “big data”, electronic devices are now involved in many social activities, and can capture, store, manipulate, and analyze vast caches of such data. Conventional statistical and econometric techniques become increasingly inadequate to deal with “big data” problems, such as the curse of dimensionality, large storage requirement, and high computational cost (see, e.g. Ning et al., 2021). Feature selection, also known as variable selection, attribute selection, or variable subset selection, is a machine learning technique

*The research of Ning Ning was partially supported by the Seed Fund Grant Award at Texas A&M University.

[†]Department of Statistics, Texas A&M University, College Station, TX, USA, patning@tamu.edu

for dimensionality reduction (see, e.g. Lamnisis et al., 2009). It aims to choose a small subset of the relevant features (variables, predictors) for use in model construction, from the original ones by removing irrelevant, redundant, or noisy features. It usually leads to better learning performance in terms of higher learning accuracy, lower computational cost, and better model interpretability. It has also been proven to be an effective and efficient way of shortening training times, improving data’s compatibility with a learning model class, and encoding inherent symmetries present in the input space (see, e.g. Griffin and Steel, 2021).

Time series forecasting as one of the most applied data science techniques have been used extensively in finance, supply chain management, production and inventory planning, etc. However, as the time dimension adds additional information, time series problems are more difficult to handle compared to other prediction tasks (Kalli and Griffin, 2014). Scott and Varian (2014, 2015) developed the Bayesian Structural Time Series (BSTS) model, which can be used for time series forecasting, nowcasting, inferring causal relationships, etc. Qiu et al. (2018) proposed the multivariate Bayesian structural time series (MBSTS) model, which extended the BSTS model to multivariate target time series with various components. The BSTS/MBSTS model has wide applications (see, e.g., Jammalamadaka et al., 2019; Qiu et al., 2020 and the references therein).

One can think of the future value of a univariate time series as a random variable whose distribution is the forecast distribution. Analogously, the future values of multivariate time series can be thought of as correlated random variables with the joint forecast distribution. The usual time series forecast is often to predict the mean or the median of the (joint) forecast distribution, which can be categorized as a point forecast. The (joint) quantile forecast is appealing in many economic applications (Ley and Steel, 2007), such as value at risk in the finance industry in order to develop a strategy for trading and/or hedging purposes. Recently, Griffin and Mitrodima (2022) proposed a Bayesian quantile time series model for asset returns which beautifully handled formal Bayesian inference on quantiles. It is an open question so far on how to incorporate the feature selection technique in joint quantile time series analysis, which is our goal of this paper.

1.2 Our Contributions

In this paper, we propose a general Bayesian methodology for feature selection in joint quantile analysis with multivariate correlated time series data, under the name of quantile feature selection time series (QFSTS) model. Our contributions are four-fold:

- **General model structure.** First, the QFSTS model is a structural time series model, which further belongs to state space models. All state components are assembled independently, and each component yields an additive contribution to the time series modeling with direct interpretations. The flexibility of the model allows users to include different components for each target series, such as one time series has a trend component and a regression component while another correlated time series only has a regression component. Given that our main goal is

to illustrate the feature selection methodology in joint quantile time series analysis, we focus on a trend component, a seasonal component, a regression component, and an error term (2.1) without loss of generality.

Second, the QFSTS model allows nowcasting and each time series can have its own pool of contemporaneous external predictors. Nowcasting is to forecast a current value instead of a future value (Banbura et al., 2011). Two scenarios for using nowcasting are the following: First, many economic time series are reported infrequently such as on a monthly or quarterly basis, despite being theoretically observable on finer time scales; Second, time series are also frequently revised after they are first reported, as new information becomes available. An effective nowcasting model considers both the past behavior of the series being modeled and the values of more easily observed contemporaneous signals. All predictors in the regression component of the QFSTS model can be contemporaneous with a known lag by shifting the corresponding predictors in time. The QFSTS model allows each time series to have its own pool of predictors, for example one time series has 50 predictors and another has 20 different predictors. Feature selection is conducted among each times series' pool of predictors while still considering the correlations among different times series.

- **Methodology.** First, the QFSTS model is the first on introducing the multivariate asymmetric Laplace (MAL) distribution to time series analysis. The idea of Bayesian quantile regression employing a likelihood function that is based on the asymmetric Laplace (AL) distribution, was introduced in the classical work by Yu and Moyeed (2001). The use of the AL distribution is proved as a very natural and effective way of modeling Bayesian quantile regression, irrespective of the original distribution of the data (Chen et al., 2013). In Section 2.2, we rigorously establish the explicit linkage between the MAL distribution and joint quantile regression, by setting the specific forms of parameters of the MAL distributed error term.

Second, the QFSTS model achieves feature selection in joint quantile time series analysis. It is very challenging to conduct feature selection on correlated time series where each has its own predictors, while when it comes to selecting features in quantile regression, it is much more challenging in that one has to take quantiles into consideration instead of the mean. Further, the Bayesian quantile regression coefficients depend on the quantile of interest (Yu and Moyeed, 2001). That is, for different quantiles, the coefficient of any predictor varies. The QFSTS model uses the Gibbs sampler for quantile feature selection by means of the classical spike-and-slab prior setup (George and McCulloch, 1997; Madigan and Raftery, 1994), and the Metropolis-Hastings algorithm. Bayesian variable selection can be performed using mixture priors with a spike and a slab component for the effects subject to selection, where the spike prior concentrates its mass at values close to zero allowing shrinkage of small effects to zero, and the slab prior has its mass spread over a wide range of plausible values for the regression coefficients. As the spike is concentrated at zero, variable selection is based on the probability of assigning the corresponding regression effect to the slab component. These posterior inclusion probabilities can be determined by MCMC sampling.

Third, the QFSTS model introduces the Bayesian model averaging technique (Hoeting et al., 1999) in joint quantile time series analysis. Bayesian model averaging provides a coherent mechanism to handle model uncertainty, by averaging the predicted quantile values from all the models generated in model training (Fernandez et al., 2001; Steel, 2020). In this way, we commit neither to any particular set of predictors which helps avoid an arbitrary selection, nor to point estimates of their coefficients which prevents overfitting.

- **Excellent performance.** Under the challenging situation that each time series has its own pool of predictors made of both continuous and categorical covariates, the QFSTS model provides very accurate feature selection and parameter estimation results, thoroughly examined on simulated data, with different quantile values, different size of datasets, and different correlation values among multivariate time series. Extensive analyses confirmed that the QFSTS model outperforms the “auto.arima” function in ten steps of one-step-ahead quantile forecast consistently. The “auto.arima” function automatically fits the autoregressive integrated moving average with regression (ARIMAX) model that is the most classical and popular time series model.
- **Wide applicability.** First, because of the feature selection function embedded in the QFSTS model, it is applicable in proper applications that desire higher learning accuracy, lower computational cost, better model interpretability, and other benefits generated by feature selection. Second, because of the Bayesian paradigm embedded in the QFSTS model, it is applicable in proper applications that require Bayesian inference and learning. Third, because of the QFSTS model’s general model structure, it is applicable in more applications besides the finance applications demonstrated in this paper. At last, the QFSTS model requires small datasets to train and converges fast.

1.3 Organization of the Paper

The rest of the paper proceeds as follows: In Section 2, we introduce the QFSTS model by explaining its model structure and then illustrating the linkage between the specific design of the QFSTS model with multivariate quantile prediction. In Section 3, we explain the methodology of the QFSTS model, by firstly writing the system in the distribution equivalence matrix form, and then providing the prior setup and posterior formulas. In Section 4, we provide the model training algorithm and the joint quantile time series forecast algorithm. In Section 5, we demonstrate the model performance with simulated data generated by three-dimensional target time series, and fully examine the QFSTS model’s ability in selecting the correct variables, accuracy in parameter estimation, and strength in forecast, with different quantiles and time series correlations. In Section 6, we conclude and remark. In the Supplementary Material (Ning, 2023), we conduct further numerical analysis with simulated data and demonstrate the model performance with empirical data. Throughout the paper, the symbol τ , with or without subscripts, will represent the quantile level.

2 The QFSTS Model

In this section, we introduce the QFSTS model, by firstly explaining its model structure in Section 2.1, and then illustrating the linkage between the specific design of the QFSTS model with multivariate quantile prediction in Section 2.2.

2.1 Model Structure

The QFSTS model is a structural time series model, which is constructed by components with direct interpretations, as follows:

$$\tilde{y}_t = \tilde{\mu}_t + \tilde{\iota}_t + \tilde{\xi}_t + \tilde{\epsilon}_t, \quad (2.1)$$

where t stands for a specific time point in $\{1, \dots, n\}$, the m -dimensional vector \tilde{y}_t represents target time series where $m > 1$, the m -dimensional vector $\tilde{\mu}_t$ represents the linear trend component, the m -dimensional vector $\tilde{\iota}_t$ represents the seasonal component, the m -dimensional vector $\tilde{\xi}_t$ represents the regression component, and the m -dimensional vector $\tilde{\epsilon}_t$ represents the observation error term. Since structural time series models belong to state space models, the QFSTS model then belongs to state space models. In the following, we illustrate each time series components:

- The linear trend component $\tilde{\mu}_t = [\mu_t^{(1)}, \dots, \mu_t^{(m)}]^T$ of the QFSTS model is a generalization of the local linear trend model, in the following form:

$$\mu_{t+1}^{(i)} = \mu_t^{(i)} + \delta_t^{(i)} + u_t^{(i)}, \quad [u_t^{(1)}, \dots, u_t^{(m)}]^T \stackrel{\text{iid}}{\sim} N_m(0, \Sigma_{\tilde{\mu}}), \quad (2.2)$$

$$\delta_{t+1}^{(i)} = D_i + \lambda_i(\delta_t^{(i)} - D_i) + v_t^{(i)}, \quad [v_t^{(1)}, \dots, v_t^{(m)}]^T \stackrel{\text{iid}}{\sim} N_m(0, \Sigma_{\tilde{\delta}}). \quad (2.3)$$

The QFSTS model allows each target series to have its own specific linear trend component. Here, for the i -th target series where $i \in \{1, \dots, m\}$, $\mu_t^{(i)}$ models the current “level” of the trend; $\delta_t^{(i)}$ models the current “slope” of the trend which is the expected increase in $\mu_t^{(i)}$ between time point t and time point $t + 1$; the parameter D_i models the “mean reversion” effect, i.e., a dynamic can be very unstable in the short run but stable in the long run; the parameter $\lambda_i \in [0, 1]$ is the learning rate at which the local trend is updated.

- The seasonal component $\tilde{\iota}_t = [\iota_t^{(1)}, \dots, \iota_t^{(m)}]^T$ of the QFSTS model in the following form is one frequently used model:

$$\iota_{t+1}^{(i)} = - \sum_{k=0}^{S_i-2} \iota_{t-k}^{(i)} + w_t^{(i)}, \quad [w_t^{(1)}, \dots, w_t^{(m)}]^T \stackrel{\text{iid}}{\sim} N_m(0, \Sigma_{\iota}), \quad (2.4)$$

where S_i represents the number of seasons for $y^{(i)}$. The QFSTS model allows each target series to have its own specific seasonal component. For instance, we might include a seasonal component with $S_i = 7$ to capture day-of-the-week effect for target series $y^{(i)}$, and $S_j = 30$ indicating day-of-the-month effect for another target series $y^{(j)}$ when modeling daily data.

- The regression component $\tilde{\xi}_t = [\xi_t^{(1)}, \dots, \xi_t^{(m)}]^T$ of the QFSTS model is written as follows:

$$\xi_t^{(i)} = \beta_i^T x_t^{(i)}. \quad (2.5)$$

The QFSTS model allows each target series to have its own specific regression component. Here for $i \in \{1, \dots, m\}$, $x_t^{(i)} = [x_{t1}^{(i)}, \dots, x_{tk_i}^{(i)}]^T$ is the pool of all available k_i predictors at time t for the i -th target series, for example $k_i = 30$ for the i -th target series and $k_j = 20$ for the j -th target series where $j \neq i$; $\beta_i = [\beta_{i1}, \dots, \beta_{ik_i}]^T$ represents corresponding static regression coefficients for the i -th target series.

- The vector of observation error terms $\tilde{\epsilon}_t = [\epsilon_t^{(1)}, \dots, \epsilon_t^{(m)}]^T$ follows the multivariate asymmetric Laplace (MAL) distribution

$$\tilde{\epsilon}_t \stackrel{\text{iid}}{\sim} \mathcal{AL}_m(\phi_\epsilon, \Sigma_\epsilon), \quad (2.6)$$

where ϕ_ϵ is a m -dimensional vector of means for each variable, and Σ_ϵ is a $m \times m$ -dimensional non-negative definite symmetric matrix of variance-covariance. When $\phi_\epsilon = 0$, the distribution $\mathcal{AL}_m(0, \Sigma_\epsilon)$ is the symmetric multivariate Laplace distribution. We refer interested readers to Section 6.2 of Kotz et al. (2012) for further explanations of the MAL distribution.

2.2 The MAL Distribution and Joint Quantile Regression

In this section, we establish the linkage between the MAL distribution and joint quantile regression, by setting the specific forms of ϕ_ϵ and Σ_ϵ in the $\mathcal{AL}_m(\phi_\epsilon, \Sigma_\epsilon)$ distribution, which is given in (2.15).

Firstly, we subtract the trend component and the seasonal component from the multivariate target time series and denote

$$\tilde{z}_t = [z_t^{(1)}, \dots, z_t^{(m)}]^T = \tilde{y}_t - \tilde{\mu}_t - \tilde{l}_t.$$

Then (2.1) becomes

$$\tilde{z}_t = \tilde{\xi}_t + \tilde{\epsilon}_t, \quad \tilde{\epsilon}_t \stackrel{\text{iid}}{\sim} \mathcal{AL}_m(\phi_\epsilon, \Sigma_\epsilon). \quad (2.7)$$

By Theorem 6.3.1 in Section 6.3 of Kotz et al. (2012), $\tilde{\epsilon}_t$ can be represented as follows, in the distribution equivalence sense,

$$\tilde{\epsilon}_t = \phi_\epsilon W + W^{1/2} \tilde{e}_t, \quad W \stackrel{\text{iid}}{\sim} \text{Exp}(1), \quad \tilde{e}_t = [e_t^{(1)}, \dots, e_t^{(m)}]^T \stackrel{\text{iid}}{\sim} N_m(0, \Sigma_\epsilon), \quad (2.8)$$

where W and \tilde{e}_t are independent, therefore we can rewrite (2.7) as

$$\tilde{z}_t = \tilde{\xi}_t + \phi_\epsilon W + W^{1/2} \tilde{e}_t, \quad (2.9)$$

in the distribution equivalence sense.

The specific forms of ϕ_ϵ and Σ_ϵ in the $\mathcal{AL}_m(\phi_\epsilon, \Sigma_\epsilon)$ distribution which link the MAL distribution with joint quantile regression, is based on the corresponding theory in the univariate case. Now, we summarize some of the findings in Yu and Moyeed (2001). The standard linear regression model is given by

$$y_t^u = (\mathbf{x}_t)^T \boldsymbol{\beta} + \epsilon_t^u,$$

where the superscript “ u ” indicates the univariate case and bold symbols represent vectors. Here, \mathbf{x}_t is the vector of regressors; $\boldsymbol{\beta}$ is the vector of corresponding coefficients; the error term ϵ_t^u has zero mean and finite constant variance, but it is not necessary to specify its distribution as it can take any form. Suppose that the p -th ($0 < p < 1$) quantile of the error term ϵ_t^u is the value q_p , such that $\mathbb{P}(\epsilon_t^u < q_p) = p$. The p -th conditional quantile of y_t^u given \mathbf{x}_t is then simply

$$q_p(y_t^u | \mathbf{x}_t) = (\mathbf{x}_t)^T \boldsymbol{\beta}_p, \quad (2.10)$$

where $\boldsymbol{\beta}_p$ is a vector of coefficients dependent on p .

The p -th regression quantile is defined as the solution $\widehat{\boldsymbol{\beta}}_p$ to the quantile regression minimization problem

$$\min_{\boldsymbol{\beta}} \sum_t \rho_p \left(y_t^u - (\mathbf{x}_t)^T \boldsymbol{\beta} \right),$$

where $\rho_p(\cdot)$ is a loss function that has robust properties (see, Huber, 2004) given by

$$\rho_p(u) = \frac{|u| + (2p - 1)u}{2}. \quad (2.11)$$

Yu and Moyeed (2001) showed that the minimization of the above loss function is exactly equivalent to the maximization of a likelihood function formed by combining independently distributed asymmetric Laplace (AL) densities, and the univariate AL distribution provides a direct connection between the quantile regression minimization problem and the maximum likelihood estimation.

A random variable U is said to follow the univariate AL distribution $\mathcal{AL}(\theta_{\text{loc}}, \theta_{\text{sca}}, p)$ if its probability density function is given by

$$f_p(u) = \frac{p(1-p)}{\theta_{\text{sca}}} \exp \left\{ -\rho_p \left(\frac{u - \theta_{\text{loc}}}{\theta_{\text{sca}}} \right) \right\},$$

where $\rho_p(\cdot)$ is given in (2.11), θ_{loc} is the location parameter, and θ_{sca} is the scale parameter. The mean and the variance of U are given by

$$\mathbb{E}(U) = \theta_{\text{loc}} + \theta_{\text{sca}} \frac{1 - 2p}{p(1-p)}, \quad \text{Var}(U) = \theta_{\text{sca}}^2 \frac{1 - 2p + 2p^2}{p^2(1-p)^2}.$$

Now we get back to (2.9) and investigate its univariate case of the i -th target series for $i \in \{1, 2, \dots, m\}$:

$$z_t^{(i)} = \xi_t^{(i)} + \phi_\epsilon^{(i)} W + W^{1/2} e_t^{(i)}, \quad e_t^{(i)} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, (\Sigma_\epsilon)_{ii}). \quad (2.12)$$

Suppose we are interested in the τ_i -th ($0 < \tau_i < 1$) quantile and then we need $z_t^{(i)}$ to satisfy the univariate AL distribution $\mathcal{AL}(\xi_t^{(i)}, \phi_i, \tau_i)$.

In order to obtain

$$\mathbb{E}(z_t^{(i)}) = \xi_t^{(i)} + \phi_i \frac{1 - 2\tau_i}{\tau_i(1 - \tau_i)},$$

it suffices to set

$$\phi_\epsilon^{(i)} = \phi_i \frac{1 - 2\tau_i}{\tau_i(1 - \tau_i)} \quad (2.13)$$

since

$$\mathbb{E}(\phi_\epsilon^{(i)} W) = \phi_\epsilon^{(i)} \quad \text{and} \quad \mathbb{E}(W^{1/2} e_t^{(i)}) = 0.$$

Furthermore, in order to obtain

$$\text{Var}(z_t^{(i)}) = (\phi_i)^2 \frac{1 - 2\tau_i + 2\tau_i^2}{\tau_i(1 - \tau_i)},$$

it suffices to set

$$(\Sigma_\epsilon)_{ii} = (\phi_i)^2 \frac{2}{\tau_i(1 - \tau_i)}, \quad (2.14)$$

since then

$$\begin{aligned} \text{Var}(z_t^{(i)}) &= \text{Var}(\phi_\epsilon^{(i)} W + W^{1/2} e_t^{(i)}) \\ &= \left(\phi_i \frac{1 - 2\tau_i}{\tau_i(1 - \tau_i)} \right)^2 + (\phi_i)^2 \frac{2}{\tau_i(1 - \tau_i)} \\ &= (\phi_i)^2 \frac{1 - 2\tau_i + 2\tau_i^2}{\tau_i(1 - \tau_i)}. \end{aligned}$$

At last, in order to meet the univariate requirements (2.13) and (2.14), we can take

$$\phi_\epsilon = \Phi \tilde{\phi}_\tau \quad \text{and} \quad \Sigma_\epsilon = \Phi \Sigma_\tau \Phi = \Phi (\Psi_\tau \Sigma_{\text{corr}} \Psi_\tau) \Phi, \quad (2.15)$$

where $\tilde{\phi}_\tau = \left[\frac{1-2\tau_1}{\tau_1(1-\tau_1)}, \frac{1-2\tau_2}{\tau_2(1-\tau_2)}, \dots, \frac{1-2\tau_m}{\tau_m(1-\tau_m)} \right]^T$, Σ_{corr} is the correlation matrix,

$$\Phi = \begin{bmatrix} \phi_1 & 0 & \dots & 0 \\ 0 & \phi_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \phi_m \end{bmatrix}, \quad \Psi_\tau = \begin{bmatrix} \sqrt{\frac{2}{\tau_1(1-\tau_1)}} & 0 & \dots & 0 \\ 0 & \sqrt{\frac{2}{\tau_2(1-\tau_2)}} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sqrt{\frac{2}{\tau_m(1-\tau_m)}} \end{bmatrix}.$$

3 Methodology

In this section, we explain the methodology of the QFSTS model. In Section 3.1, we write the system in the distribution equivalence matrix form to facilitate derivations following. In Section 3.2, we provide the prior setups including the spike-and-slab for variable selection in this Bayesian paradigm. In Section 3.3, we derive all necessary conditional probabilities, in order to implement the classical Gibbs sampler in model training.

3.1 The Distribution Equivalence Matrix Form

Recall that by (2.9) we have

$$\tilde{z}_t = \tilde{\xi}_t + \phi_\epsilon W + W^{1/2} \tilde{e}_t.$$

Here, $\tilde{z}_t = [z_t^{(1)}, \dots, z_t^{(m)}]^T$ is the m -dimensional vector generated by subtracting the trend component and the seasonal component from the multivariate target time series; $\tilde{\xi}_t = [\xi_t^{(1)}, \dots, \xi_t^{(m)}]^T$ is the m -dimensional regression component such that $\xi_t^{(i)} = \beta_i^T x_t^{(i)}$, where $\beta_i = [\beta_{i,1}, \dots, \beta_{i,k_i}]^T$ is the collection of regression coefficients for the i -th response variable, and $x_t^{(i)} = [x_{t1}^{(i)}, \dots, x_{tk_i}^{(i)}]^T$ is the pool of all available k_i predictors at time t for the i -th target series; ϕ_ϵ is the m -dimensional vector of means for the m -dimensional error term.

Now, we define the $n \times m$ matrix Z as

$$Z = [\tilde{z}_1, \dots, \tilde{z}_n]^T = \left[\begin{bmatrix} z_1^{(1)} \\ \vdots \\ z_1^{(m)} \end{bmatrix}, \dots, \begin{bmatrix} z_n^{(1)} \\ \vdots \\ z_n^{(m)} \end{bmatrix} \right]^T = \begin{bmatrix} [z_1^{(1)} & \dots & z_1^{(m)}] \\ \vdots \\ [z_n^{(1)} & \dots & z_n^{(m)}] \end{bmatrix},$$

and then define the mn -dimensional vector \tilde{Z} as

$$\tilde{Z} = \text{vec}(Z) = \left[[z_1^{(1)} \quad \dots \quad z_n^{(1)}], \dots, [z_1^{(m)} \quad \dots \quad z_n^{(m)}] \right]^T.$$

Analogously, we define the mn -dimensional vector \tilde{E} as

$$\tilde{E} = \text{vec}(E) \quad \text{where} \quad E = [\tilde{\epsilon}_1, \dots, \tilde{\epsilon}_n]^T,$$

and define the mn -dimensional vector $\tilde{\Phi}_\epsilon$ as

$$\tilde{\Phi}_\epsilon = \text{vec}(\Phi_\epsilon) \quad \text{where} \quad \Phi_\epsilon = [\phi_\epsilon, \dots, \phi_\epsilon]^T.$$

Define β as the K -dimensional ($K = \sum_{i=1}^m k_i$) vector of regression coefficients

$$\beta = [\beta_1^T, \dots, \beta_m^T]^T \quad \text{where} \quad \beta_i = [\beta_{i,1}, \dots, \beta_{i,k_i}]^T.$$

Define X_i as the $n \times k_i$ matrix representing all observations of k_i possible predictors for the i -th response variable

$$X_i = \left[(x_1^{(i)})^T, \dots, (x_n^{(i)})^T \right]^T = \begin{bmatrix} x_{1,1}^{(i)} & x_{1,2}^{(i)} & x_{1,3}^{(i)} & \dots & x_{1,k_i}^{(i)} \\ x_{2,1}^{(i)} & x_{2,2}^{(i)} & x_{2,i,3}^{(i)} & \dots & x_{2,k_i}^{(i)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x_{n,1}^{(i)} & x_{n,2}^{(i)} & x_{n,3}^{(i)} & \dots & x_{n,k_i}^{(i)} \end{bmatrix}, \quad (3.1)$$

and further define X as the $mn \times K$ -dimensional predictor matrix

$$X = \begin{bmatrix} X_1 & 0 & 0 & \dots & 0 \\ 0 & X_2 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & X_m \end{bmatrix}. \quad (3.2)$$

Then we have the following expression, in the distribution equivalence sense,

$$\tilde{Z} = X\beta + \tilde{\Phi}_\epsilon W + W^{1/2} \tilde{E}. \quad (3.3)$$

3.2 Prior Distributions

The spike-and-slab prior setup is a Bayesian variable selection technique. To specify spike prior, a vector of $K = \sum_{i=1}^m k_i$ indicator variables

$$\gamma = [\gamma_{1,1}, \dots, \gamma_{1,k_1}, \gamma_{2,1}, \dots, \gamma_{2,k_2}, \dots, \gamma_{m,1}, \dots, \gamma_{m,k_m}]$$

is introduced according to the rule:

$$\gamma_{i,k} = \begin{cases} 1 & \text{if } \beta_{i,k} \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Regressor indicators are assumed to be independent Bernoulli variables:

$$p(\gamma|W) = \prod_{i=1}^m \prod_{k=1}^{k_i} \pi_{i,k}^{\gamma_{i,k}} (1 - \pi_{i,k})^{1-\gamma_{i,k}}, \quad 0 \leq \pi_{i,k} = p(\gamma_{i,k}) \leq 1, \quad (3.4)$$

where $\pi_{i,k}$ is the prior inclusion probability. Equation (3.4) is often simplified by setting $\pi_{i,k} = \pi_i$, if prior information of specific predictors on response variables is not available. One could further simplify by setting $\pi_i = q_i/k_i$, where k_i is the total number of candidate predictors for the i -th target series and q_i nonzero expected predictors given by researchers. When there is sufficient prior information, assigning subjectively determined values to $\pi_{i,k}$ might provide more robust results. One could also force certain variables to be excluded or included by setting $\pi_{i,j}$ as 0 or 1. By default and in all the experimental examinations following, we set $\pi_{ik} = 0.5$ for all $i \in \{1, \dots, m\}$ and all $k \in \{1, \dots, k_i\}$.

We use a simple conventional prior specification which makes β and Σ_ϵ conditionally independent

$$p(\beta, \Sigma_\tau | \gamma) = p(\beta | \gamma) \times p(\Sigma_\tau | \gamma), \quad (3.5)$$

where

$$\beta | \gamma \stackrel{\text{iid}}{\sim} N_K(b_\gamma, A_\gamma^{-1}), \quad \Sigma_\tau | \gamma \stackrel{\text{iid}}{\sim} IW_m(v_0, V_0). \quad (3.6)$$

Equation (3.6) is the slab prior because, conditional on γ , one can choose the prior parameters to make it only very weakly informative and close to flat. $N_K(b_\gamma, A_\gamma^{-1})$ stands for the K -dimensional multivariate normal distribution, b_γ is the vector of prior means, and A_γ is the full-model prior information matrix. One can set $A_\gamma = \kappa X_\gamma^T X_\gamma / n$ where κ is the number of observations worth of weight on the prior mean vector b_γ . $IW_m(v_0, V_0)$ stands for the m -dimensional inverse Wishart (IW) distribution, where v_0 is the number of degrees of freedom and V_0 is a $m \times m$ scale matrix. One can ask analysts for an expected R^2 , and a number of observations worth of weight v_0 which must be greater than the dimension of \tilde{y}_t plus one, and set

$$V_0 = (v_0 - m - 1)(1 - R^2)\Sigma_y,$$

where Σ_y is the variance-covariance matrix for multiple target time series. For simplicity, we set $b_\gamma = 0$, $\kappa = 0.01$, $R^2 = 0.8$, and $v_0 = 5$ in all the experimental examinations following. We acknowledge that the IW prior setup may not necessarily conform to the specific form presented in (2.15), which is one example demonstrating the connection to the univariate case. However, considering the difficulty in estimating the correlation matrix for multivariate time series, we opt for the IW prior as a simpler alternative. Consequently, it is important to note that our modeling approach may not be optimal, and there is potential for further improvements to be made over our methodology.

Since we are going to use the Metropolis-Hastings Algorithm (see Chapter 6.3.1 of Robert and Casella, 2010) to learn the distribution of Φ , we allow the prior distribution of its elements to be any distribution that is proportional to 1. The prior distributions of variance-covariance matrices in the trend component and the seasonal component are set as the inverse Wishart distribution

$$\Sigma_\alpha \stackrel{\text{iid}}{\sim} IW_m(\nu_\alpha, V_\alpha), \quad \alpha \in \{\tilde{\mu}, \tilde{\delta}, \tilde{\nu}\}. \quad (3.7)$$

For simplicity, we set $\nu_\alpha = V_\alpha = 0.01$ in all the experimental examinations following.

3.3 Posterior Conditional Distributions

In order to implement the classical Gibbs sampler in this multivariate setting, we derive all necessary conditional probabilities of \tilde{Z} , β , Φ , Σ_τ , γ , and W . The full likelihood function under model assumptions is given by

$$p(\tilde{Z}, \beta, \Phi, \Sigma_\tau, \gamma, W) = p(\tilde{Z} | \beta, \Phi, \Sigma_\tau, \gamma, W) \times p(\beta, \Sigma_\tau | \gamma) \times p(\gamma) \times p(W) \times p(\Phi).$$

Then, by (3.3)–(3.6) and the setup of $p(\Phi)$, we have that

$$p(\tilde{Z}, \beta, \Phi, \Sigma_\tau, \gamma, W)$$

$$\begin{aligned}
& \propto p(\tilde{Z}|\beta, \Phi, \Sigma_\tau, \gamma, W) \times p(\beta|\gamma, W) \times p(\Sigma_\tau|\gamma, W) \times p(\gamma) \times p(W) \\
& \propto |W\Phi\Sigma_\tau\Phi|^{-n/2} \exp\left(-\frac{1}{2W}(\tilde{Z}-X_\gamma\beta_\gamma-\tilde{\Phi}_\epsilon W)^T((\Phi\Sigma_\tau\Phi)^{-1} \otimes I_n)(\tilde{Z}-X_\gamma\beta_\gamma-\tilde{\Phi}_\epsilon W)\right) \\
& \quad \times |A_\gamma|^{1/2} \exp\left(-\frac{1}{2}(\beta_\gamma-b_\gamma)^T A_\gamma(\beta_\gamma-b_\gamma)\right) |\Sigma_\tau|^{-(v_0+m+1)/2} \exp\left(-\frac{1}{2} \text{tr}(V_0\Sigma_\tau^{-1})\right) \\
& \quad \times p(\gamma) \times p(W), \tag{3.8}
\end{aligned}$$

where $|\cdot|$ stands for the determinant of a matrix, \otimes is the Kronecker product, and $\text{tr}(\cdot)$ represents the trace of a matrix.

Posterior Conditional Distribution of β

To facilitate derivation, we firstly transform

$$\tilde{Z} = X\beta + \tilde{\Phi}_\epsilon W + W^{1/2}\tilde{E},$$

where

$$\tilde{E} = \text{vec}(E) = \text{vec}([\tilde{\epsilon}_1, \dots, \tilde{\epsilon}_n]^T), \quad \tilde{\epsilon}_t \stackrel{\text{iid}}{\sim} N_m(0, \Sigma_\epsilon = \Phi\Sigma_\tau\Phi),$$

to a system with uncorrelated errors using the Cholesky decomposition of Σ_τ ,

$$\Sigma_\tau = U^T U, \quad \text{i.e. } (U^{-1})^T \Sigma_\tau U^{-1} = I. \tag{3.9}$$

Thus we have the transformed system with uncorrelated errors:

$$\hat{Z} = \hat{X}\beta + \hat{\Phi}_\epsilon W + W^{1/2}\hat{E}, \tag{3.10}$$

where

$$\begin{aligned}
\hat{Z} &= (((U\Phi)^{-1})^T \otimes I_n)\tilde{Z}, & \hat{X} &= (((U\Phi)^{-1})^T \otimes I_n)X, \\
\hat{\Phi}_\epsilon &= (((U\Phi)^{-1})^T \otimes I_n)\tilde{\Phi}_\epsilon, & \hat{E} &= (((U\Phi)^{-1})^T \otimes I_n)\tilde{E}.
\end{aligned} \tag{3.11}$$

For the following term in the first exponential in (3.8), we have

$$\begin{aligned}
& (\tilde{Z}-X_\gamma\beta_\gamma-\tilde{\Phi}_\epsilon W)^T((\Phi\Sigma_\tau\Phi)^{-1} \otimes I_n)(\tilde{Z}-X_\gamma\beta_\gamma-\tilde{\Phi}_\epsilon W) \\
& = (\tilde{Z}-X_\gamma\beta_\gamma-\tilde{\Phi}_\epsilon W)^T([U\Phi]^T U\Phi)^{-1} \otimes I_n)(\tilde{Z}-X_\gamma\beta_\gamma-\tilde{\Phi}_\epsilon W) \\
& = (\tilde{Z}-X_\gamma\beta_\gamma-\tilde{\Phi}_\epsilon W)^T((U\Phi)^{-1} \otimes I_n) \times (((U\Phi)^{-1})^T \otimes I_n)(\tilde{Z}-X_\gamma\beta_\gamma-\tilde{\Phi}_\epsilon W) \\
& = (\hat{Z}-\hat{X}_\gamma\beta_\gamma-\hat{\Phi}_\epsilon W)^T(\hat{Z}-\hat{X}_\gamma\beta_\gamma-\hat{\Phi}_\epsilon W). \tag{3.12}
\end{aligned}$$

The full conditional distribution of β can be expressed as:

$$\begin{aligned}
p(\beta|\hat{Z}, \Phi, \Sigma_\epsilon, \gamma, W) & \propto \exp\left(-\frac{1}{2}W^{-1}(\hat{Z}-\hat{X}_\gamma\beta_\gamma-\hat{\Phi}_\epsilon W)^T(\hat{Z}-\hat{X}_\gamma\beta_\gamma-\hat{\Phi}_\epsilon W)\right) \\
& \quad \times \exp\left(-\frac{1}{2}(\beta_\gamma-b_\gamma)^T A_\gamma(\beta_\gamma-b_\gamma)\right).
\end{aligned}$$

Terms in the above exponential can be written as

$$\begin{aligned}
& W^{-1}(\widehat{Z} - \widehat{X}_\gamma \beta_\gamma - \widehat{\Phi}_\epsilon W)^T (\widehat{Z} - \widehat{X}_\gamma \beta_\gamma - \widehat{\Phi}_\epsilon W) + (\beta_\gamma - b_\gamma)^T A_\gamma (\beta_\gamma - b_\gamma) \\
&= \beta_\gamma^T (W^{-1} \widehat{X}_\gamma^T \widehat{X}_\gamma + A_\gamma) \beta_\gamma - \beta_\gamma^T (W^{-1} \widehat{X}_\gamma^T \widehat{Z} - \widehat{X}_\gamma^T \widehat{\Phi}_\epsilon + A_\gamma b_\gamma) \\
&\quad - (W^{-1} \widehat{X}_\gamma^T \widehat{Z} - \widehat{X}_\gamma^T \widehat{\Phi}_\epsilon + A_\gamma b_\gamma)^T \beta_\gamma + W^{-1} (\widehat{Z} - \widehat{\Phi}_\epsilon W)^T (\widehat{Z} - \widehat{\Phi}_\epsilon W) + b_\gamma^T A_\gamma b_\gamma \\
&= (\beta_\gamma - \bar{\beta}_\gamma)^T (W^{-1} \widehat{X}_\gamma^T \widehat{X}_\gamma + A_\gamma) (\beta_\gamma - \bar{\beta}_\gamma) + W^{-1} (\widehat{Z} - \widehat{\Phi}_\epsilon W)^T (\widehat{Z} - \widehat{\Phi}_\epsilon W) \\
&\quad + b_\gamma^T A_\gamma b_\gamma - (\bar{\beta}_\gamma)^T (W^{-1} \widehat{X}_\gamma^T \widehat{X}_\gamma + A_\gamma) \bar{\beta}_\gamma, \tag{3.13}
\end{aligned}$$

where

$$\bar{\beta}_\gamma = (W^{-1} \widehat{X}_\gamma^T \widehat{X}_\gamma + A_\gamma)^{-1} (W^{-1} \widehat{X}_\gamma^T \widehat{Z} - \widehat{X}_\gamma^T \widehat{\Phi}_\epsilon + A_\gamma b_\gamma).$$

Therefore, β is still conditionally multivariate normal distributed

$$\beta | \widehat{Z}, \Phi, \Sigma_\tau, \gamma, W \stackrel{\text{iid}}{\sim} N_K(\bar{\beta}_\gamma, (W^{-1} \widehat{X}_\gamma^T \widehat{X}_\gamma + A_\gamma)^{-1}). \tag{3.14}$$

Posterior Conditional Distribution of Σ_τ

Recalling that X_i is the $n \times k_i$ -dimensional matrix given in (3.1), define the $n \times K$ -dimensional ($K = \sum_{i=1}^m k_i$) matrix X_γ^* as

$$X_\gamma^* = [X_1, X_2, \dots, X_m]_\gamma.$$

Define the $K \times m$ -dimensional matrix B_γ as

$$B_\gamma = \begin{bmatrix} \beta_1 & 0 & 0 & \dots & 0 \\ 0 & \beta_2 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \beta_m \end{bmatrix}_\gamma, \quad \beta_i = \begin{bmatrix} \beta_{i,1} \\ \vdots \\ \beta_{i,k_i} \end{bmatrix},$$

where β_i is the k_i -dimensional vector containing the collection of regression coefficients for the i -th response series. For the reason that trace is invariant under cyclic permutations, from (3.8), we know that

$$\begin{aligned}
& (\widetilde{Z} - X_\gamma \beta_\gamma - \widetilde{\Phi}_\epsilon W)^T ((\Phi \Sigma_\tau \Phi) \otimes I_n)^{-1} (\widetilde{Z} - X_\gamma \beta_\gamma - \widetilde{\Phi}_\epsilon W) \\
&= \text{vec} \left(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W \right)^T ((\Phi \Sigma_\tau \Phi)^{-1} \otimes I_n) \text{vec} \left(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W \right) \\
&= \text{tr} \left((Z - X_\gamma^* B_\gamma - \Phi_\epsilon W)^T (Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \Sigma_\tau^{-1} \Phi^{-1} \right) \\
&= \text{tr} \left(\left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right] \Sigma_\tau^{-1} \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right]^T \right) \\
&= \text{tr} \left(\left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right]^T \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right] \Sigma_\tau^{-1} \right),
\end{aligned}$$

and then we have

$$\begin{aligned}
& p(\Sigma_\tau | \tilde{Z}, \Phi, \beta, \gamma, W) \\
& \propto |\Sigma_\tau|^{-(n+v_0+m+1)/2} \exp \left(-\frac{1}{2} \operatorname{tr} \left(\frac{1}{W} \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right]^T \right. \right. \\
& \quad \left. \left. \times \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right] \Sigma_\tau^{-1} + V_0 \Sigma_\tau^{-1} \right) \right).
\end{aligned}$$

That is, the posterior conditional distribution of Σ_τ is in the invert Wishart form

$$\begin{aligned}
& \Sigma_\tau | \tilde{Z}, \beta, \Phi, \gamma, W \\
& \stackrel{\text{iid}}{\sim} IW_m \left(v_0 + n, \frac{1}{W} \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right]^T \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right] + V_0 \right).
\end{aligned} \tag{3.15}$$

Posterior Conditional Distribution of Φ

Recall that by the Cholesky decomposition in (3.9) we have that $\Sigma_\tau = U^T U$. Further recall that by (2.15) we have that the m -dimensional vector $\phi_\epsilon = \Phi \tilde{\phi}_\tau$ where Φ is a $m \times m$ -dimensional diagonal matrix and $\tilde{\phi}_\tau$ is a m -dimensional vector, and then we can write the $n \times m$ -dimensional matrix Φ_ϵ as

$$\Phi_\epsilon = [\phi_\epsilon, \dots, \phi_\epsilon]^T = \tilde{\Phi}_\tau \Phi, \quad \text{where } \tilde{\Phi}_\tau = [\tilde{\phi}_\tau, \dots, \tilde{\phi}_\tau]^T.$$

Then by (3.8) we have

$$\begin{aligned}
& p(\Phi | \tilde{Z}, \beta, \Sigma_\tau, \gamma, W) \\
& \propto |\Phi|^{-n} \exp \left(-\frac{1}{2} \operatorname{tr} \left(\frac{1}{W} \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right]^T \right. \right. \\
& \quad \left. \left. \times \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right] \Sigma_\tau^{-1} \right) \right) \\
& \propto |\Phi|^{-n} \exp \left(-\frac{1}{2} \operatorname{tr} \left(\frac{1}{W} \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right]^T \right. \right. \\
& \quad \left. \left. \times \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} \right] U^{-1} (U^{-1})^T \right) \right) \\
& \propto |\Phi|^{-n} \exp \left(-\frac{1}{2} \operatorname{tr} \left(\frac{1}{W} \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} U^{-1} \right]^T \right. \right. \\
& \quad \left. \left. \times \left[(Z - X_\gamma^* B_\gamma - \Phi_\epsilon W) \Phi^{-1} U^{-1} \right] \right) \right)
\end{aligned}$$

$$\begin{aligned}
& \propto |\Phi|^{-n} \exp\left(-\frac{1}{2} \operatorname{tr}\left(\frac{1}{W}\left[(Z - X_\gamma^* B_\gamma - \tilde{\Phi}_\tau \Phi W)\Phi^{-1} U^{-1}\right]^T\right.\right. \\
& \quad \left.\left. \times \left[(Z - X_\gamma^* B_\gamma - \tilde{\Phi}_\tau \Phi W)\Phi^{-1} U^{-1}\right]\right)\right) \\
& \propto |\Phi|^{-n} \exp\left(-\frac{1}{2} \operatorname{tr}\left(\frac{1}{W}\left[(Z - X_\gamma^* B_\gamma)\Phi^{-1} U^{-1} - \tilde{\Phi}_\tau U^{-1} W\right]^T\right.\right. \\
& \quad \left.\left. \times \left[(Z - X_\gamma^* B_\gamma)\Phi^{-1} U^{-1} - \tilde{\Phi}_\tau U^{-1} W\right]\right)\right). \quad (3.16)
\end{aligned}$$

Posterior Conditional Distribution of γ

By (3.8) and (3.13), we know that

$$\begin{aligned}
& p(\tilde{Z}, \beta, \Phi, \Sigma_\tau, \gamma, W) \\
& \propto \exp\left(-\frac{1}{2} \left[(\beta - \bar{\beta}_\gamma)^T (W^{-1} \hat{X}_\gamma^T \hat{X}_\gamma + A_\gamma)(\beta - \bar{\beta}_\gamma) + W^{-1} (\hat{Z} - \hat{\Phi}_\epsilon W)^T (\hat{Z} - \hat{\Phi}_\epsilon W)\right]\right) \\
& \quad \times |A_\gamma|^{1/2} \exp\left(-\frac{1}{2} \left[\operatorname{tr}(V_0 \Sigma_\epsilon^{-1}) + b_\gamma^T A_\gamma b_\gamma - (\bar{\beta}_\gamma)^T (W^{-1} \hat{X}_\gamma^T \hat{X}_\gamma + A_\gamma) \bar{\beta}_\gamma\right]\right) p(\gamma) \\
& \quad \times |W \Phi \Sigma_\tau \Phi|^{-n/2} |\Sigma_\tau|^{-(v_0+m+1)/2},
\end{aligned}$$

where

$$\bar{\beta}_\gamma = (W^{-1} \hat{X}_\gamma^T \hat{X}_\gamma + A_\gamma)^{-1} (W^{-1} \hat{X}_\gamma^T \hat{Z} - \hat{X}_\gamma^T \hat{\Phi}_\epsilon + A_\gamma b_\gamma).$$

Furthermore, by the fact that

$$\beta | \hat{Z}, \Phi, \Sigma_\tau, \gamma, W \stackrel{\text{iid}}{\sim} N_K(\bar{\beta}_\gamma, (W^{-1} \hat{X}_\gamma^T \hat{X}_\gamma + A_\gamma)^{-1}),$$

we have

$$\int_{-\infty}^{\infty} \exp\left(-\frac{1}{2} (\beta - \bar{\beta}_\gamma)^T (W^{-1} \hat{X}_\gamma^T \hat{X}_\gamma + A_\gamma) (\beta - \bar{\beta}_\gamma)\right) d\beta \propto |W^{-1} \hat{X}_\gamma^T \hat{X}_\gamma + A_\gamma|^{-1/2},$$

and then

$$\begin{aligned}
& p(\tilde{Z}, \Phi, \Sigma_\tau, \gamma, W) \\
& = \int_{-\infty}^{\infty} p(\tilde{Z}, \beta, \Phi, \Sigma_\tau, \gamma, W) d\beta \\
& \propto \exp\left(-\frac{1}{2} \left[W^{-1} (\hat{Z} - \hat{\Phi}_\epsilon W)^T (\hat{Z} - \hat{\Phi}_\epsilon W) - \Xi_\gamma^T (W^{-1} \hat{X}_\gamma^T \hat{X}_\gamma + A_\gamma)^{-1} \Xi_\gamma\right]\right) \\
& \quad \times |A_\gamma|^{1/2} \exp\left(-\frac{1}{2} \left[\operatorname{tr}(V_0 \Sigma_\epsilon^{-1}) + b_\gamma^T A_\gamma b_\gamma\right]\right) |W^{-1} \hat{X}_\gamma^T \hat{X}_\gamma + A_\gamma|^{-1/2} p(\gamma) \\
& \quad \times |W \Phi \Sigma_\tau \Phi|^{-n/2} |\Sigma_\tau|^{-(v_0+m+1)/2} p(W),
\end{aligned}$$

where

$$\Xi_\gamma = (W^{-1}\widehat{X}_\gamma^T\widehat{Z} - \widehat{X}_\gamma^T\widehat{\Phi}_\epsilon + A_\gamma b_\gamma).$$

Therefore, the posterior conditional distribution of γ is given by

$$\begin{aligned} p(\gamma|\widetilde{Z}, \Phi, \Sigma_\tau, W) &\propto \exp\left(-\frac{1}{2}\left[b_\gamma^T A_\gamma b_\gamma - \Xi_\gamma^T (W^{-1}\widehat{X}_\gamma^T\widehat{X}_\gamma + A_\gamma)^{-1}\Xi_\gamma\right]\right) \\ &\quad \times |A_\gamma|^{1/2}|W^{-1}\widehat{X}_\gamma^T\widehat{X}_\gamma + A_\gamma|^{-1/2}p(\gamma). \end{aligned} \quad (3.17)$$

Posterior Conditional Distribution of W

Recall that the generalized inverse Gaussian distribution (GIG) is a three-parameter family of continuous probability distributions with probability density function (see page 1 of Jorgensen, 2012)

$$f(x) = \frac{(a/b)^{p/2}}{2K_p(\sqrt{ab})} x^{(p-1)} e^{-(ax+b/x)/2}, \quad x > 0,$$

where K_p is a modified Bessel function of the second kind, $a > 0$, $b > 0$, and p is a real parameter. By (3.8) and (3.12), we have that

$$\begin{aligned} p(W|\widetilde{Z}, \beta, \widetilde{\Phi}, \Sigma_\epsilon, \gamma) \\ \propto |W|^{-n/2} \exp\left(-\frac{1}{2W}(\widehat{Z} - \widehat{X}_\gamma\beta_\gamma - \widehat{\Phi}_\epsilon W)^T(\widehat{Z} - \widehat{X}_\gamma\beta_\gamma - \widehat{\Phi}_\epsilon W) - W\right), \end{aligned}$$

based on which,

$$\begin{aligned} W|\widetilde{Z}, \beta, \Sigma_\epsilon, \gamma &\stackrel{\text{iid}}{\sim} \text{GIG}(a, b, p), \\ a = 2 + \widehat{\Phi}_\epsilon^T\widehat{\Phi}_\epsilon, \quad b = (\widehat{Z} - \widehat{X}_\gamma\beta_\gamma)^T(\widehat{Z} - \widehat{X}_\gamma\beta_\gamma), \quad p = 1 - n/2. \end{aligned} \quad (3.18)$$

Posterior Conditional Distribution of Σ_α

Next we need to derive conditional posterior distribution of Σ_α where $\alpha \in \{\widetilde{\mu}, \widetilde{\delta}, \widetilde{\iota}\}$ in the trend component and the seasonal component. Similarly, as the posterior conditional distribution of Σ_τ in the invert Wishart form in (3.15), the posterior distribution of Σ_α is conditionally inverse Wishart distributed

$$\Sigma_\alpha|\widetilde{Y}, \alpha, W \stackrel{\text{iid}}{\sim} IW_m\left(\nu_\alpha + n, V_\alpha + \frac{1}{W}AA^T\right), \quad \alpha \in \{\widetilde{\mu}, \widetilde{\delta}, \widetilde{\iota}\}, \quad (3.19)$$

where A is the matrix of a collection of residues of each time series component.

4 Algorithms

Gibbs sampling is a Markov chain Monte Carlo (MCMC) algorithm for obtaining a sequence of observations, which are approximated from a specified multivariate probability distribution. MCMC methods are to construct a Markov chain that has the

desired distribution as its equilibrium distribution. One can draw samples of the desired distribution by discarding the initial MCMC steps as “burn-in”, since the quality of samples is an increasing function of the number of steps. In Algorithm 1, the posterior distributions of the model are simulated by Gibbs sampling approach, in the way that looping through the 7 steps yields a sequence of draws $\theta = (\alpha, \Sigma_\alpha, \beta, \Phi, \Sigma_\tau, \gamma, W)$ where $\alpha \in \{\tilde{\mu}, \tilde{\delta}, \tilde{\iota}\}$, from a Markov chain with the stationary probability distribution $p(\theta|Y)$ which is the posterior distribution of θ given Y .

Algorithm 1 Model Training.

Time series state components

- 1: Draw the latent state α from $p(\alpha|\tilde{Y}, \Sigma_\alpha, \beta, \Phi, \Sigma_\tau, \gamma, W)$ where $\alpha \in \{\tilde{\mu}, \tilde{\delta}, \tilde{\iota}\}$, using the posterior simulation algorithm from Durbin and Koopman (2002).
- 2: Draw time series state component parameters Σ_α from $\Sigma_\alpha \stackrel{\text{iid}}{\sim} p(\Sigma_\alpha|\tilde{Y}, \alpha, W)$ based on the inverse Wishart distribution in (3.19).

Quantile regression component

- 3: Loop over i in a random order, draw each $\gamma_i|\gamma_{-i}, \tilde{Z}, \Phi, \Sigma_\tau, W$, namely simulating $\gamma \stackrel{\text{iid}}{\sim} p(\gamma|\tilde{Z}, \Phi, \Sigma_\tau, W)$ in (3.17), using the stochastic search variable selection (SSVS) algorithm from George and McCulloch (1997).
- 4: Draw β from $\beta \stackrel{\text{iid}}{\sim} p(\beta|\tilde{Z}, \Phi, \Sigma_\tau, \gamma, W)$ based on the multivariate normal distribution in (3.14).

Error term

- 5: Draw Σ_τ from $\Sigma_\tau \stackrel{\text{iid}}{\sim} p(\Sigma_\tau|\tilde{Z}, \beta, \Phi, \gamma, W)$ based on the inverse Wishart distribution in (3.15).
 - 6: Draw Φ based on $p(\Phi|\tilde{Z}, \beta, \Sigma_\tau, \gamma, W)$ in (3.16) using the Metropolis-Hastings Algorithm.
 - 7: Draw W from $W \stackrel{\text{iid}}{\sim} p(W|\tilde{Z}, \beta, \Sigma_\tau, \gamma)$ based on the generalized inverse Gaussian distribution in (3.18).
-

Algorithm 2 Joint Quantile Predictions.

- 1: Draw the next trend component $\alpha_{t+1} = (\tilde{\mu}_{t+1}, \tilde{\delta}_{t+1}, \tilde{\iota}_{t+1})$, given current trend component $\alpha_t = (\tilde{\mu}_t, \tilde{\delta}_t, \tilde{\iota}_t)$ and variance-covariance parameters $\Sigma_\alpha = (\Sigma_{\tilde{\mu}}, \Sigma_{\tilde{\delta}}, \Sigma_{\tilde{\iota}})$, by (2.2) and (2.3).
 - 2: Based on indicator variable γ , compute the regression component $\tilde{\xi}_{t+1}$ given the information about predictors at time $t + 1$, by (2.5).
 - 3: Draw a random error $\tilde{\epsilon}_{t+1}$ in the multivariate asymmetric Laplace distribution by (2.6), whose mean and variance are generated by expressions given in (2.15).
 - 4: Sum up $\tilde{\mu}_{t+1}$, $\tilde{\iota}_{t+1}$, $\tilde{\xi}_{t+1}$, and $\tilde{\epsilon}_{t+1}$ to generate predictions, by (2.1).
 - 5: Sum up all the generated predictions and divide by the total number of effective MCMC iterations to generate the joint quantile predictions.
-

Given draws of model parameters and latent states from their posterior distributions, we can draw samples from the posterior predictive distribution

$$p(\widehat{Y}|Y) = \int p(\widehat{Y}|\theta)p(\theta|Y)d\theta,$$

where \widehat{Y} represents the set of values to forecast. Here, the posterior predictive distribution is not conditioned on parameter estimates or the inclusion/exclusion of predictors, all of which have been integrated out. Algorithm 2 conducts joint quantile prediction, where forecasts are generated by the Bayesian model averaging approach which provides a coherent mechanism to handle model uncertainty, by averaging the predicted values from all the models generated in the MCMC model training. Through Bayesian model averaging, we commit neither to any particular set of predictors which helps avoid an arbitrary selection, nor to point estimates of their coefficients which prevents overfitting.

5 Model Performance with Simulated Data

In this section, we demonstrate the model performance with simulated data generated by three-dimensional target time series given in Section 5.1, in terms of selecting the correct variables and accuracy in parameter estimation in Section 5.2, and forecast performance of the model with different quantiles and different time series correlations in Section 5.3. The AL likelihood is known to be very restrictive in modeling the underlying error distributions and unlikely to be the true data generating likelihood (Yang et al., 2016), so it is primarily used as a working likelihood in the Bayesian quantile regression literature. Even though its posterior consistency for coefficient estimation has been established by Sriram et al. (2013) under model misspecification, the validity of using AL likelihood for prediction purpose still remains questionable. Hence, we use error terms simulated by the AL likelihood for coefficient estimation and by the Gaussian distribution for response prediction.

5.1 Generated Data

The simulated data is generated by the following three-dimensional model (i.e., $m = 3$)

$$\tilde{y}_t = \tilde{\mu}_t + \tilde{v}_t + B^T \tilde{x}_t + \tilde{\epsilon}_t, \quad (5.1)$$

where each time series has its own trend component, seasonal component, and regression component. We will mainly use the dataset size $n = 500$. The trend component $\tilde{\mu}_t$ is generated as follows:

$$\tilde{\mu}_{t+1} = \begin{bmatrix} \mu_{t+1}^{(1)} \\ \mu_{t+1}^{(2)} \\ \mu_{t+1}^{(3)} \end{bmatrix} = \begin{bmatrix} \mu_t^{(1)} \\ \mu_t^{(2)} \\ \mu_t^{(3)} \end{bmatrix} + \begin{bmatrix} \delta_t^{(1)} \\ \delta_t^{(2)} \\ \delta_t^{(3)} \end{bmatrix} + \begin{bmatrix} u_t^{(1)} \\ u_t^{(2)} \\ u_t^{(3)} \end{bmatrix}, \quad \begin{bmatrix} u_t^{(1)} \\ u_t^{(2)} \\ u_t^{(3)} \end{bmatrix} \stackrel{\text{iid}}{\sim} N_3 \left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right).$$

where its slope is generated as

$$\begin{bmatrix} \delta_t^{(1)} \\ \delta_t^{(2)} \\ \delta_t^{(3)} \end{bmatrix} = \begin{bmatrix} 0.04 + 0.6(\delta_{t-1}^{(1)} - 0.04) \\ 0.05 + 0.3(\delta_{t-1}^{(2)} - 0.05) \\ 0.02 + 0.1(\delta_{t-1}^{(3)} - 0.02) \end{bmatrix} + \begin{bmatrix} v_t^{(1)} \\ v_t^{(2)} \\ v_t^{(3)} \end{bmatrix}, \quad \begin{bmatrix} v_t^{(1)} \\ v_t^{(2)} \\ v_t^{(3)} \end{bmatrix} \stackrel{\text{iid}}{\sim} N_3 \left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \right).$$

The seasonal component \tilde{l}_t is generated as follows:

$$\tilde{l}_{t+1} = \begin{bmatrix} l_{t+1}^{(1)} \\ l_{t+1}^{(2)} \\ l_{t+1}^{(3)} \end{bmatrix} = \begin{bmatrix} -\sum_{k=0}^{100} l_{t-k}^{(1)} \\ -\sum_{k=0}^{70} l_{t-k}^{(2)} \\ -\sum_{k=0}^{40} l_{t-k}^{(3)} \end{bmatrix} + \begin{bmatrix} w_t^{(1)} \\ w_t^{(2)} \\ w_t^{(3)} \end{bmatrix}, \quad \begin{bmatrix} w_t^{(1)} \\ w_t^{(2)} \\ w_t^{(3)} \end{bmatrix} \stackrel{\text{iid}}{\sim} N_3 \left(\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 0.5 & 0 & 0 \\ 0 & 0.5 & 0 \\ 0 & 0 & 0.5 \end{bmatrix} \right).$$

The regression component $B^T \tilde{x}_t$ is generated with 8 explanatory variables, at least one of which has no effect on each target series with zero regression coefficient, as follows:

$$B = \begin{bmatrix} 2 & 4 & -3.5 & -2 & 0 & 0 & -1.6 & 0 \\ 3 & 0 & 2.5 & -3 & 0 & -1.5 & 0 & 2 \\ -2.5 & 0 & -2 & -1 & 3 & 2 & 0 & 4 \end{bmatrix}^T,$$

$$\tilde{x}_t = [x_{t1} \ x_{t2} \ x_{t3} \ x_{t4} \ x_{t5} \ x_{t6} \ x_{t7} \ x_{t8}]^T,$$

$$x_{t1} \stackrel{\text{iid}}{\sim} \mathcal{N}(5, 2), \quad x_{t2} \stackrel{\text{iid}}{\sim} \text{Pois}(10), \quad x_{t3} \stackrel{\text{iid}}{\sim} \text{Pois}(5), \quad x_{t4} \stackrel{\text{iid}}{\sim} \mathcal{N}(-2, 5),$$

$$x_{t5} \stackrel{\text{iid}}{\sim} \mathcal{N}(-5, 2), \quad x_{t6} \stackrel{\text{iid}}{\sim} \text{Pois}(15), \quad x_{t7} \stackrel{\text{iid}}{\sim} \text{Pois}(20), \quad x_{t8} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 10).$$

To examine the accuracy of variable selection and parameter estimation in Section 5.2, we use the simulated data generated by the following MAL distributed error term $\tilde{\epsilon}_t \stackrel{\text{iid}}{\sim} \mathcal{AL}_m(\phi_\epsilon, \Sigma_\epsilon)$ according to (2.15) where, for $\tau = (\tau_1, \tau_2, \tau_3)$,

$$\begin{aligned} \phi_\epsilon &= \Phi \tilde{\phi}_\tau, \quad \Sigma_\epsilon = \Phi \Sigma_\tau \Phi = \Phi (\Psi_\tau \Sigma_{\text{corr}} \Psi_\tau) \Phi, \\ \Sigma_{\text{corr}} &= \begin{bmatrix} 1 & \rho & \rho \\ \rho & 1 & \rho \\ \rho & \rho & 1 \end{bmatrix}, \quad \Phi = \begin{bmatrix} 0.7 & 0 & 0 \\ 0 & 0.6 & 0 \\ 0 & 0 & 0.9 \end{bmatrix}, \\ \tilde{\phi}_\tau &= \begin{bmatrix} \frac{1-2\tau_1}{\tau_1(1-\tau_1)} \\ \frac{1-2\tau_2}{\tau_2(1-\tau_2)} \\ \frac{1-2\tau_3}{\tau_3(1-\tau_3)} \end{bmatrix}, \quad \text{and} \quad \Psi_\tau = \begin{bmatrix} \sqrt{\frac{2}{\tau_1(1-\tau_1)}} & 0 & 0 \\ 0 & \sqrt{\frac{2}{\tau_2(1-\tau_2)}} & 0 \\ 0 & 0 & \sqrt{\frac{2}{\tau_3(1-\tau_3)}} \end{bmatrix}. \end{aligned} \tag{5.2}$$

To examine the prediction strength in Section 5.3, we use the simulated data generated by the following normal distributed error term

$$\tilde{\epsilon}_t \stackrel{\text{iid}}{\sim} N_3 \left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho & \rho \\ \rho & 1 & \rho \\ \rho & \rho & 1 \end{bmatrix} \right). \tag{5.3}$$

5.2 Model Training Performance

In this section, we are going to demonstrate the superior feature selection performances of the QFSTS model with small datasets, using only 400 MCMC iterations including 200 discarded as burn-in. The left three plots in Figure 1 provide the feature selection results for a dataset of 500 observations, generated by Model (5.1) with quantile $\tau = (0.9, 0.9, 0.9)$ and pairwise correlation 0.7. The threshold inclusion probability was set as 0.8, i.e., $\geq 80\%$ times a predictor was selected out of the $(400 - 200)$ MCMC iterations. We can see that the selected features exactly match the model setup, where the value 1 means a feature was selected all the time out of the $(400 - 200)$ MCMC iterations. The signs of selected variables also exactly match the model setup, and were marked with red for positive and blue for negative. The right three plots of Figure 1 reveal that, for each target series, only in a very small portion out of the $(400 - 200)$ MCMC iterations, the model selected more variables. Further analyses with different quantiles, correlations, and inclusion probabilities are provided in the Supplementary Material (Ning, 2023).

Figure 2 demonstrates the fast convergence and superior parameter estimation performance, with datasets generated by Model (5.1) with quantile $\tau = (0.9, 0.9, 0.9)$, pairwise correlation $\rho = 0.7$, and 7 dataset sizes (100, 200, 300, 400, 500, 600, 700), using only 400 MCMC iterations including 200 discarded as burn-in. The left three plots provide the normalized estimation errors calculated as the

$$|(\text{estimated value} - \text{true value})/\text{true value}|,$$

and the right three plots provide the standard deviations of estimation. We can see that both the estimation errors and estimation standard deviations decrease fast as the sample size increases. Similar superior model training performances for different quantiles and correlations are provided in the Supplementary Material (Ning, 2023).

5.3 Forecast Performance

Quantile time series forecasting is the prediction of the distribution of a future value of a time series. It is much more challenging than the time series mean or median forecast which is already difficult given the additional time information. The QFSTS model is a Monte Carlo-based algorithm for quantile prediction. The Monte Carlo samples generate the empirical distribution, whose mean is the quantile prediction. The most well-known quantile forecast algorithm that is publicly available, is the “auto.arima” function for univariate time series analysis, in the “forecast” R package (Hyndman and Khandakar, 2008). “auto.arima” automatically fits the best ARIMAX Model, which is the most classical and popular time series model, according to either AIC, AICc, or BIC value. Setting the “biasadj” option in the “auto.arima” function to “FALSE”, whose default value is “TRUE” for mean prediction, gives the quantile prediction. There are only 4 quantile values possible: 2.5%, 10%, 90%, and 97.5%. Therefore, based on these 4 quantile values, we analyze the QFSTS model’s forecast performances. Figure 3 reports the forecast performances for three-dimensional target time series datasets of 500 observations, generated by (5.1) with fixed pairwise correlation $\rho = 0.7$ but different quantiles: $\tau = (0.025, 0.025, 0.025)$, $\tau = (0.1, 0.1, 0.1)$, $\tau = (0.9, 0.9, 0.9)$, and

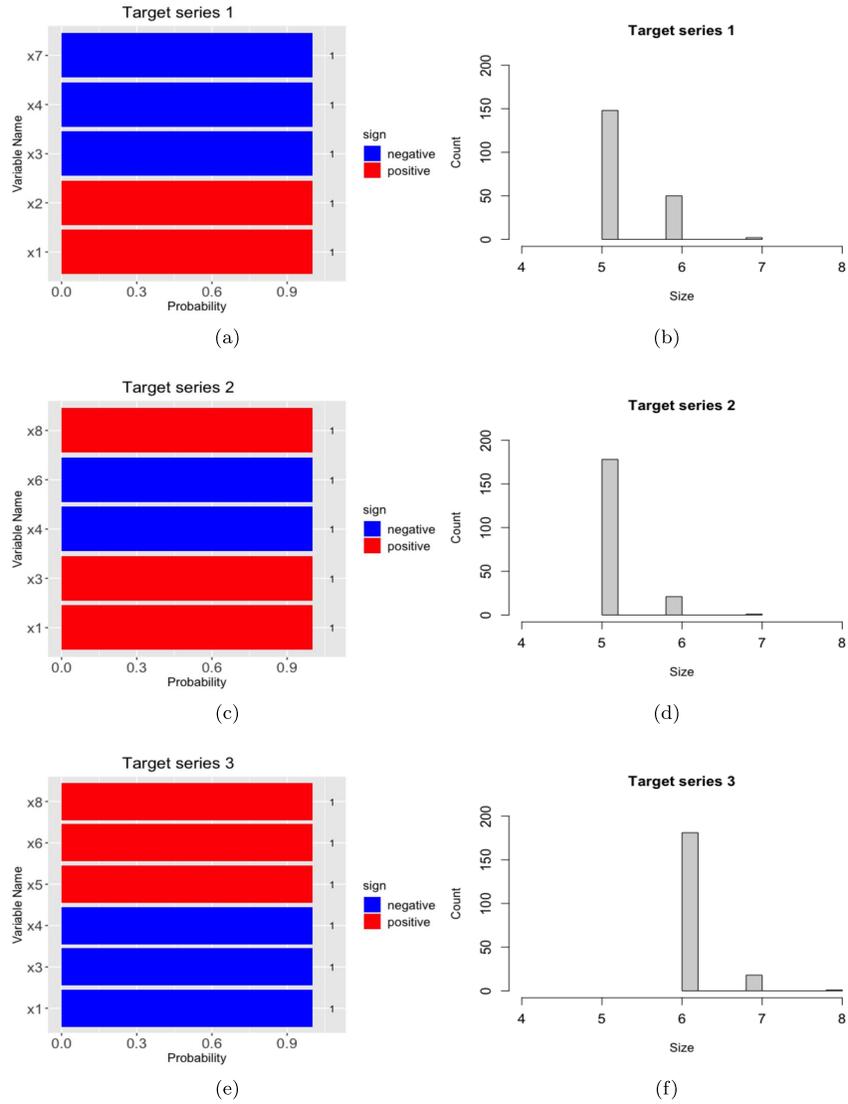


Figure 1: Feature selection results with simulated data. The 3-dimensional time series dataset of 500 observations is generated by Model (5.1) with quantile $\tau = (0.9, 0.9, 0.9)$ and pairwise correlation 0.7. The threshold inclusion probability was set as 0.8. Model training used 400 MCMC iterations including 200 discarded as burn-in. The left three plots ((a), (c), and (e)) provide the feature selection results whose X-axis represents the empirical probability (for example, 1 means that variable was selected 200 times out of the 200 MCMC iterations after burn-in), and whose Y-axis represents the variables selected. The right three plots ((b), (d), and (f)) report the count distribution out of these 200 iterations (Y-axis), where the X-axis stands for the count of selected variables (for example, 5 in (b) corresponds to the 5 variable selected in (a)).

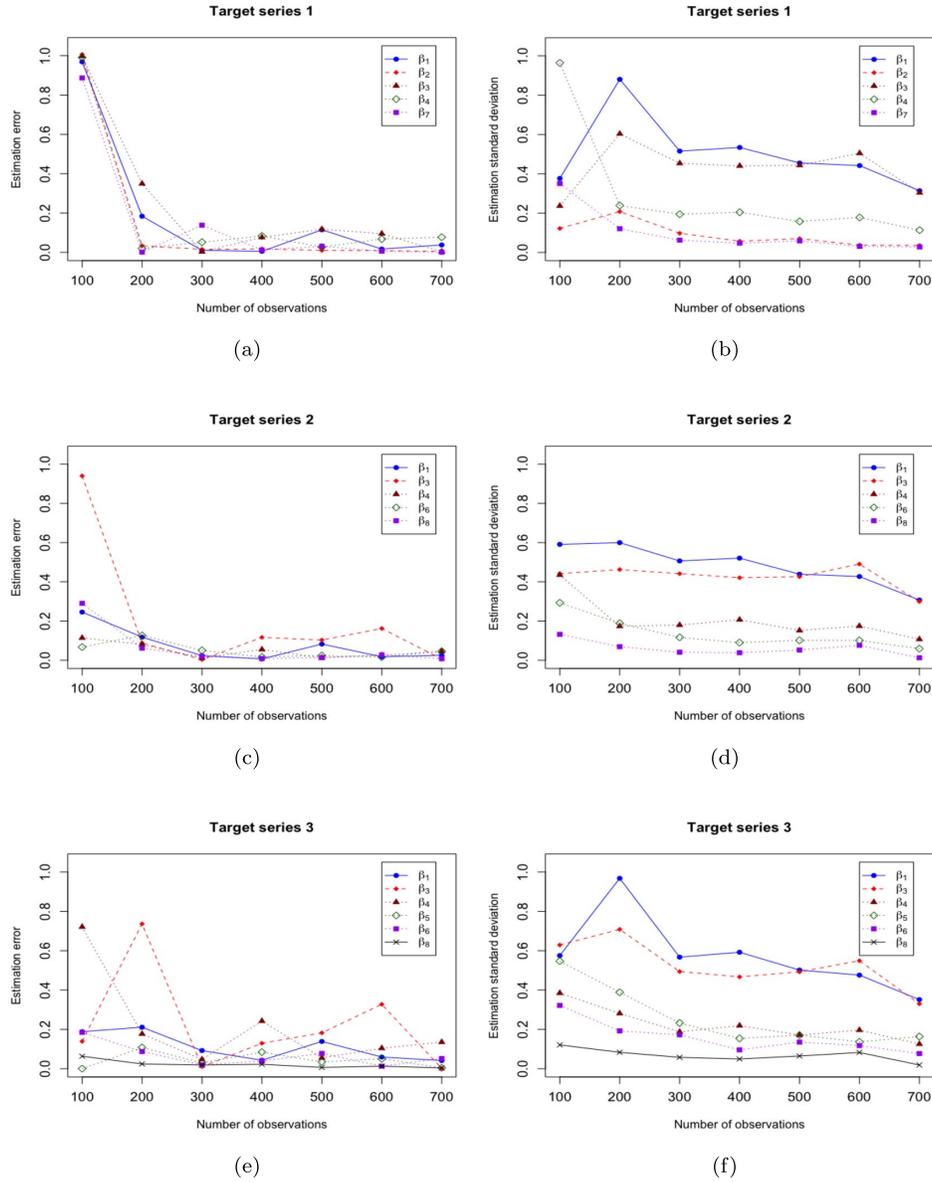


Figure 2: Estimation errors and standard deviations of regression coefficients. The 3-dimensional time series datasets were generated by Model (5.1) with quantile $\tau = (0.9, 0.9, 0.9)$, pairwise correlation $\rho = 0.7$, and 7 dataset sizes (100, 200, 300, 400, 500, 600, 700), using only 400 MCMC iterations including 200 discarded as burn-in. The left three plots ((a), (c), and (e)) provide the estimation errors of regression coefficients, and the right three plots ((b), (d), and (f)) provide the standard deviations of regression coefficients, for target series 1, 2, and 3, respectively.

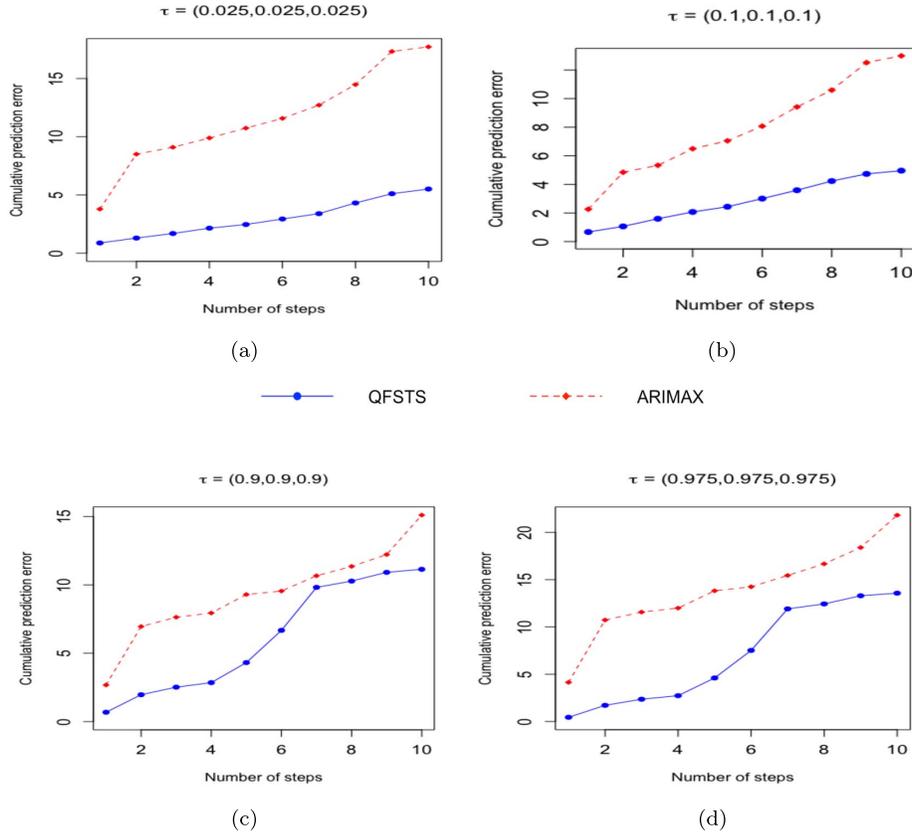


Figure 3: Cumulative absolute one-step-ahead prediction errors with simulated data for different quantiles. The red colored line is the performance of the ARIMAX model generated by “auto.arima” function, and the blue colored line is the performance of the QFSTS model. Three-dimensional time series datasets of 500 observations were generated by Model (5.1) with pairwise correlation $\rho = 0.7$. Model training used 400 MCMC iterations including 200 discarded as burn-in. Prediction error is measured by the quantile loss function given in (2.11). Plots ((a), (b), (c), and (d)) provide cumulative prediction errors for quantiles $\tau = (0.025, 0.025, 0.025)$, $\tau = (0.1, 0.1, 0.1)$, $\tau = (0.9, 0.9, 0.9)$, and $\tau = (0.975, 0.975, 0.975)$, respectively.

$\tau = (0.975, 0.975, 0.975)$. The cumulative prediction error is calculated accumulatively according to the quantile loss function given in (2.11), where this standard approach can also be seen in Chen et al. (2013). We can see that the QFSTS model outperforms “auto.arima” consistently in the tens steps of one-step ahead forecast.

6 Conclusion

In this paper, we have proposed the QFSTS model for joint quantile analysis of correlated time series with dimension $m > 1$. The correlation matrix of multivariate time series is usually very hard to estimate, while users of QFSTS do not need to evaluate the correlation among series. Hence, running QFSTS for multiple series is simpler than running it for time series one by one, with fewer user efforts involved. The IW prior is commonly used as the prior for the covariance matrix. In the distributional sense, the covariance matrix in the current setting is proportional to the correlation matrix. Hence, given the challenging situation that we allow each time series to have its own time series components (trend and season) and its own pool of predictors, we used the IW prior for simplicity. For univariate time series (the variance is just one constant and there is no correlation anymore), the current prior setup is not optimal but its proper design is out of the scope of this paper.

QFSTS achieves accurate quantile feature selection utilizing the spike-and-slab Bayesian prior setup, which uses the Bernoulli prior distribution for indicator variables. If no specific prior information on the initial inclusion probabilities of particular variables is available, the Bernoulli prior distribution is a common default choice. Variant spike-and-slab modeling can be used instead, such as spike-and-slab being normal distributions (George and McCulloch, 1993) or scale mixtures of normals (Ishwaran and Rao, 2005). The QFSTS model uses the standard Metropolis-Hastings algorithm in the MCMC model training while its speed and performance may be further improved with advanced Metropolis-Hastings algorithms (Banterle et al., 2019; Sherlock et al., 2017; Atchadé and Perron, 2005; Atchadé et al., 2011). As the first multivariate time series model with joint quantile feature selection, the QFSTS model sheds light on this new research area and outperforms the classical ARIMAX time series model consistently.

Probabilistic forecasts play a fundamental role in addressing the inherent uncertainty in data, supporting decision-making under uncertainty, evaluating and refining forecasting models, and enabling informed decision-making in various domains (Gneiting et al., 2007). Our methodology is motivated by the asymmetric Laplace working likelihood. It can also be extended to represent functional coefficients using basis representations, allowing for the borrowing of strength from nearby locations and incorporating a global-local shrinkage prior on the basis coefficients to achieve adaptive regularization (Liu et al., 2020a,b). Also motivated by a working Laplace likelihood approach, the Bayesian median autoregressive model proposed in Zeng and Li (2021) adopts a parametric model bearing the same structure as autoregressive models by altering the Gaussian error to Laplace, leading to a simple, robust, and interpretable modeling strategy for time series forecasting. Lastly, we acknowledge that the QFSTS model produces point joint quantile estimation, rather than joint interval predictions which is a challenging task that we leave for future research.

Supplementary Material

Supplement to “Bayesian Feature Selection in Joint Quantile Time Series Analysis” (DOI: [10.1214/23-BA1401SUPP](https://doi.org/10.1214/23-BA1401SUPP); .pdf).

References

- Atchadé, Y. F. and Perron, F. (2005). “Improving on the independent Metropolis-Hastings algorithm.” *Statistica Sinica*, 3–18. [MR2125717](#). 24
- Atchadé, Y. F., Roberts, G. O., and Rosenthal, J. S. (2011). “Towards optimal scaling of Metropolis-coupled Markov chain Monte Carlo.” *Statistics and Computing*, 21(4): 555–568. [MR2826692](#). doi: <https://doi.org/10.1007/s11222-010-9192-1>. 24
- Banbura, M., Giannone, D., and Reichlin, L. (2011). “Nowcasting.” *The Oxford handbook of economic forecasting*. [MR3204193](#). 3
- Banterle, M., Grazian, C., Lee, A., and Robert, C. P. (2019). “Accelerating Metropolis-Hastings algorithms by delayed acceptance.” *Foundations of Data Science*, 1(2): 103. [MR4617837](#). doi: <https://doi.org/10.3934/fods.2019005>. 24
- Chen, C. W., Dunson, D. B., Reed, C., and Yu, K. (2013). “Bayesian variable selection in quantile regression.” *Statistics and its Interface*, 6(2): 261–274. [MR3066690](#). doi: <https://doi.org/10.4310/SII.2013.v6.n2.a9>. 3, 23
- Durbin, J. and Koopman, S. J. (2002). “A simple and efficient simulation smoother for state space time series analysis.” *Biometrika*, 89(3): 603–616. [MR1929166](#). doi: <https://doi.org/10.1093/biomet/89.3.603>. 17
- Fernandez, C., Ley, E., and Steel, M. F. (2001). “Benchmark priors for Bayesian model averaging.” *Journal of Econometrics*, 100(2): 381–427. [MR1820410](#). doi: [https://doi.org/10.1016/S0304-4076\(00\)00076-2](https://doi.org/10.1016/S0304-4076(00)00076-2). 4
- George, E. I. and McCulloch, R. E. (1993). “Variable selection via Gibbs sampling.” *Journal of the American Statistical Association*, 88(423): 881–889. 24
- George, E. I. and McCulloch, R. E. (1997). “Approaches for Bayesian variable selection.” *Statistica sinica*, 339–373. 3, 17
- Gneiting, T., Balabdaoui, F., and Raftery, A. E. (2007). “Probabilistic forecasts, calibration and sharpness.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(2): 243–268. [MR2325275](#). doi: <https://doi.org/10.1111/j.1467-9868.2007.00587.x>. 24
- Griffin, J. E. and Mitrodima, G. (2022). “A Bayesian quantile time series model for asset returns.” *Journal of Business & Economic Statistics*, 40(1): 16–27. [MR4356554](#). doi: <https://doi.org/10.1080/07350015.2020.1766470>. 2
- Griffin, J. E. and Steel, M. F. (2021). “Adaptive Computational Methods for Bayesian Variable Selection.” In *Handbook of Bayesian Variable Selection*, 109–130. Chapman and Hall/CRC. 2
- Hoeting, J. A., Madigan, D., Raftery, A. E., and Volinsky, C. T. (1999). “Bayesian model averaging: a tutorial.” *Statistical science*, 382–401. [MR1765176](#). doi: <https://doi.org/10.1214/ss/1009212519>. 4
- Huber, P. J. (2004). *Robust statistics*, volume 523. John Wiley & Sons. [MR0606374](#). 7

- Hyndman, R. J. and Khandakar, Y. (2008). “Automatic Time Series Forecasting: The forecast Package for R.” *The Annals of Statistics*, 27(3). 20
- Ishwaran, H. and Rao, J. S. (2005). “Spike and slab variable selection: Frequentist and Bayesian strategies.” *The Annals of Statistics*, 33(2): 730–773. MR2163158. doi: <https://doi.org/10.1214/009053604000001147>. 24
- Jammalamadaka, S. R., Qiu, J., and Ning, N. (2019). “Predicting a Stock Portfolio with the Multivariate Bayesian Structural Time Series Model: Do News or Emotions Matter?” *International Journal of Artificial Intelligence*, 17(2): 81–104. MR3899770. 2
- Jorgensen, B. (2012). *Statistical properties of the generalized inverse Gaussian distribution*, volume 9. Springer Science & Business Media. MR0648107. 16
- Kalli, M. and Griffin, J. E. (2014). “Time-varying sparsity in dynamic regression models.” *Journal of Econometrics*, 178(2): 779–793. MR3144682. doi: <https://doi.org/10.1016/j.jeconom.2013.10.012>. 2
- Kotz, S., Kozubowski, T., and Podgorski, K. (2012). *The Laplace distribution and generalizations: A revisit with applications to communications, economics, engineering, and finance*. Springer Science & Business Media. MR1935481. doi: <https://doi.org/10.1007/978-1-4612-0173-1>. 6
- Lamnisos, D., Griffin, J. E., and Steel, M. F. (2009). “Transdimensional sampling algorithms for Bayesian variable selection in classification problems with many more variables than observations.” *Journal of Computational and Graphical Statistics*, 18(3): 592–612. MR2751643. doi: <https://doi.org/10.1198/jcgs.2009.08027>. 2
- Ley, E. and Steel, M. F. (2007). “Jointness in Bayesian variable selection with applications to growth regression.” *Journal of Macroeconomics*, 29(3): 476–493. 2
- Liu, Y., Li, M., and Morris, J. S. (2020a). “Function-on-scalar quantile regression with application to mass spectrometry proteomics data.” *The Annals of Applied Statistics*, 14(2): 521 – 541. MR4117818. doi: <https://doi.org/10.1214/19-AOAS1319>. 24
- Liu, Y., Li, M., and Morris, J. S. (2020b). “On function-on-scalar quantile regression.” *arXiv preprint arXiv:2002.03355*. 24
- Madigan, D. and Raftery, A. E. (1994). “Model selection and accounting for model uncertainty in graphical models using Occam’s window.” *Journal of the American Statistical Association*, 89(428): 1535–1546. 3
- Ning, N., Ionides, E. L., and Ritov, Y. (2021). “Scalable Monte Carlo inference and rescaled local asymptotic normality.” *Bernoulli*, 27(4): 2532–2555. MR4303894. doi: <https://doi.org/10.3150/20-BEJ1321>. 1
- Ning, N. (2023). “Supplementary Material for “Bayesian feature selection in joint quantile time series analysis”.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/23-BA1401SUPP>. 4, 20
- Qiu, J., Jammalamadaka, S. R., and Ning, N. (2018). “Multivariate Bayesian structural

- time series model.” *The Journal of Machine Learning Research*, 19(1): 2744–2776. MR3899770. 2
- Qiu, J., Jammalamadaka, S. R., and Ning, N. (2020). “Multivariate time series analysis from a Bayesian machine learning perspective.” *Annals of Mathematics and Artificial Intelligence*, 88(10): 1061–1082. MR4156189. doi: <https://doi.org/10.1007/s10472-020-09710-6>. 2
- Robert, C. P. and Casella, G. (2010). *Introducing Monte Carlo Methods with R*, volume 18. Springer. MR2572239. doi: <https://doi.org/10.1007/978-1-4419-1576-4>. 11
- Scott, S. L. and Varian, H. R. (2014). “Predicting the present with Bayesian structural time series.” *International Journal of Mathematical Modelling and Numerical Optimisation*, 5(1-2): 4–23. 2
- Scott, S. L. and Varian, H. R. (2015). “Bayesian variable selection for nowcasting economic time series.” In *Economic Analysis of the Digital Economy*, 119–135. University of Chicago Press. 2
- Sherlock, C., Thiery, A. H., and Lee, A. (2017). “Pseudo-marginal Metropolis–Hastings sampling using averages of unbiased estimators.” *Biometrika*, 104(3): 727–734. MR3694593. doi: <https://doi.org/10.1093/biomet/asx031>. 24
- Sriram, K., Ramamoorthi, R., and Ghosh, P. (2013). “Posterior consistency of Bayesian quantile regression based on the misspecified asymmetric Laplace density.” MR3066950. doi: <https://doi.org/10.1214/13-BA817>. 18
- Steel, M. F. (2020). “Model averaging and its use in economics.” *Journal of Economic Literature*, 58(3): 644–719. 4
- Yang, Y., Wang, H. J., and He, X. (2016). “Posterior inference in Bayesian quantile regression with asymmetric Laplace likelihood.” *International Statistical Review*, 84(3): 327–344. MR3580414. doi: <https://doi.org/10.1111/insr.12114>. 18
- Yu, K. and Moyeed, R. A. (2001). “Bayesian quantile regression.” *Statistics & Probability Letters*, 54(4): 437–447. MR1861390. doi: [https://doi.org/10.1016/S0167-7152\(01\)00124-9](https://doi.org/10.1016/S0167-7152(01)00124-9). 3, 7
- Zeng, Z. and Li, M. (2021). “Bayesian median autoregression for robust time series forecasting.” *International Journal of Forecasting*, 37(2): 1000–1010. 24

Acknowledgments

The author would like to thank three anonymous reviewers and the Associate Editor for their very constructive comments and efforts on this work, which greatly improved the quality of this paper.