

# Bayesian Median Autoregression for Robust Time Series Forecasting

Zijian Zeng and Meng Li

Department of Statistics, Rice University

## Abstract

We develop a Bayesian median autoregressive (BayesMAR) model for time series forecasting. The proposed method utilizes time-varying quantile regression at the median, favorably inheriting the robustness of median regression in contrast to the widely used mean-based methods. Motivated by a working Laplace likelihood approach in Bayesian quantile regression, BayesMAR adopts a parametric model bearing the same structure of autoregressive models by altering the Gaussian error to Laplace, leading to a simple, robust, and interpretable modeling strategy for time series forecasting. We estimate model parameters by Markov chain Monte Carlo. Bayesian model averaging is used to account for model uncertainty including the uncertainty in the autoregressive order, in addition to a Bayesian model selection approach. The proposed methods are illustrated using simulations and real data applications. An application to U.S. macroeconomic data forecasting shows that BayesMAR leads to favorable and often superior predictive performance compared to the selected mean-based alternatives under various loss functions that encompass both point and probabilistic forecasts. The proposed methods are generic and can be used to complement a rich class of methods that build on the autoregressive models.

## 1 Introduction

Time series forecasting is a long-standing problem in econometrics and statistics, where the overwhelming focus has been on mean-based models (Prado & West, 2010; Hyndman & Athanasopoulos, 2018). Although conditional means are the optimal forecast under the squared error loss, complex characteristics violating model assumptions that are present in real data may hamper the predictive performance. Flexible nonparametric methods (Ferraty & Vieu, 2006; Fan & Yao, 2008) building on minimal assumptions are an appealing remedy; however, they often have disadvantages in terms of not only interpretability but also in involving large numbers of parameters that often lead to daunting computation and communication issues. We were thus motivated by an attempt to propose a simple, interpretable, and principled strategy that can improve upon mean-based models in real-world time series forecasting.

There is a rich literature on robust time series forecasting including categorizing outliers (Fox, 1972; Akouemo & Poinelli, 2014), adjusting autoregressive (AR) models to offset effects of outliers (Chen & Liu, 1993a,b), exponential smoothing and Holt-Winters seasonal methods to M-estimation (Croux et al., 2008), weighted forecasts (Jose & Winkler, 2008), and detecting structural changes (Qu, 2008; Oka & Qu, 2011; Casini & Perron, 2019), just to name a few. As a nonlinear alternative, the median is more robust than the mean. While the application of median-based

methods in time series at least dates back to 1974 when John Tukey introduced the running median method (Tukey, 1974), there is surprisingly little work to comprehensively investigate the modeling, fitting, and uncertainty quantification of a median-based model in the context of time series forecasting, and in particular, how it compares with state-of-the-art mean-based methods using real data and under various loss functions.

In the quantile regression literature that encompasses the median as a special case (Koenker & Bassett Jr, 1978), Koenker & Xiao (2006) proposes quantile autoregression (QAR) models which depict the conditional distributions of the response more comprehensively at various quantile levels. Engle & Manganelli (2004) proposed the conditional autoregressive value at risk (CAViaR) model for risk management at extreme quantile levels, and this model has been followed by many others (Giacomini & Komunjer, 2005; Chen & So, 2006; Geweke & Keane, 2007; Gerlach et al., 2011). However, they were developed either for general quantile levels or extreme quantile levels and have not been applied to time series forecasting. Moreover, QAR and CAViaR are semiparametric approaches, resorting to minimizing the check loss function for estimation and necessitating non-trivial modifications to likelihood-based order selection criteria when the order is unknown.

In this paper, we propose a simple strategy by extending the traditional AR model to a median AR model (MAR) for time series forecasting. The AR model is arguably one of the most popular methods in time series, serving as the building block for other models such as generalized autoregressive conditional heteroskedasticity models (GARCH, Bollerslev, 1986) and time-varying vector autoregressive models (TV-VAR Primiceri, 2005). The proposed method utilizes time-varying quantile regression but focuses on the median, favorably inheriting the robustness of median regression in contrast to the widely used mean-based methods. It relies on parametric assumptions, and this aids interpretation and enables convenient uncertainty quantification and propagation through a principled Bayesian framework. Numerical experiments using U.S. macroeconomic data show that this simple MAR approach leads to favorable and often superior predictive performances compared to selected state-of-the-art mean-based methods that are much more complicated in nature. It is remarkable that the comparison is made not only under the absolute but also the squared error loss for point forecasts and is extended to probabilistic forecasts. The proposed methods are generic and can be used to complement any methods that build on the AR models by altering the Gaussian error assumption therein.

The MAR model has a close connection with the working asymmetric Laplace likelihood approach in Bayesian quantile regression. The working likelihood formalized by Yu & Moyeed (2001) has recently gained increasing attention (Gerlach et al., 2011; Sriram et al., 2016; Youngman, 2018; Liu et al., 2020). It provides a principled and convenient framework to quantify uncertainties in the parameter estimation eliminating the challenging task of estimating the unknown conditional density functions that are required by the conventional quantile regression for inference (Yang et al., 2016). Although the asymmetric Laplace likelihood is generally not the true data generating likelihood, a pragmatic view to support its use is that the maximum a posterior estimates resemble the usual quantile regression estimates which optimizes the check loss function. Theoretically, the posterior distribution of parameters concentrates on what minimizes the Kullback-Leibler divergence with respect to the true data-generating models (Kleijn & van der Vaart, 2006). Unlike quantile regression where one may vary the quantile levels, the median is the primary quantile level of interest in time series forecasting. Therefore, the MAR model uses the Laplace distribution as the likelihood, alleviating the concern that the working likelihood is not a valid likelihood when considering multiple quantile levels. This fully parametric model enables routine posterior sampling. We estimate model parameters by Markov chain Monte Carlo, and propose using a Bayesian

model averaging (BMA) approach (Hoeting et al., 1999) to propagate the model uncertainty in the unknown autoregressive order in addition to a Bayesian model selection strategy.

The rest of the paper is organized as follows. Section 2 introduces the MAR model, estimation, and forecasting procedure. In Section 3 we conduct simulations to compare the proposed approach with other competitive methods to assess parameter estimation. Section 4 consists of a variety of applications using real-world economic data. Section 5 concludes the paper.

## 2 Methods

### 2.1 Median Autoregressive (MAR) model

We propose a median autoregressive (MAR) model for time series forecasting, as an alternative to the widely used mean-based models. Suppose we observe a vector of time series  $\mathbf{y} = (y_1, \dots, y_T)$ . A MAR model with order  $p$ , denoted as MAR( $p$ ), first assumes a time-varying quantile regression structure

$$y_t = \beta_0 + \beta_1 y_{t-1} + \beta_2 y_{t-2} + \dots + \beta_p y_{t-p} + \epsilon_t = \mathbf{y}'_{t-1} \boldsymbol{\beta} + \epsilon_t, \quad (1)$$

where  $\mathbf{y}_{t-1} = (1, y_{t-1}, \dots, y_{t-p})'$ ,  $\boldsymbol{\beta} = (\beta_0, \beta_1, \dots, \beta_p)'$  is vector of unknown coefficient, and  $\epsilon_t$  is random error with *median* 0. Model (1) is *semiparametric* as the error distribution is left unspecified other than the constraint of possessing a zero median. The classic AR model with a given order  $p$ , denoted as AR( $p$ ), assumes a Gaussian error distribution with mean 0 and standard deviation  $\sigma$ , i.e.,  $\epsilon_t \sim N(0, \sigma^2)$ . The MAR model further assumes  $\epsilon_t \sim \text{Laplace}(0, 2\tau)$  whose probability density function is

$$f(x; 0, \tau) = \frac{1}{4\tau} \exp\left(-\frac{|x|}{2\tau}\right), \quad (2)$$

where  $\tau > 0$  is a scale parameter. The Laplace error assumption combined with the semiparametric structure in Equation (1) yielding the following likelihood function for the MAR model

$$L(\boldsymbol{\beta}, \tau) = f(\mathbf{y}|\boldsymbol{\beta}, \tau) \propto \tau^{-(T-p)} \exp\left\{-\frac{1}{2\tau} \sum_{t=p+1}^T |y_t - \mathbf{y}'_{t-1} \boldsymbol{\beta}|\right\}, \quad (3)$$

which is *parametric*. The parametric assumption in (2) aids interpretation and enables convenient uncertainty quantification and propagation through a principled Bayesian framework. In addition, the autoregressive structure in the MAR model resembles the widely used AR model provides enormous flexibility and potential to complement the rich literature that builds on the AR model.

The use of Laplace distributions is common in literature the Bayesian quantile regression where the  $\theta$ th quantile of the error in Equation (1) is assumed to be 0. For general  $\theta \in (0, 1)$ , a working likelihood method adopts the asymmetric Laplace distribution  $AL(\mu, \tau, \theta)$  as the error distribution, which has the probability density function

$$f(x; \mu, \tau, \theta) = \theta(1 - \theta)\tau^{-1} \exp\left\{-\tau^{-1}(x - \mu)(\theta - \mathbb{1}[x < \mu])\right\},$$

where  $\mu$  is a location parameter and  $\mathbb{1}(\cdot)$  is the indicator function. The asymmetric Laplace distribution reduces to the Laplace distribution at the median by setting  $\mu = 0$  and  $\theta = 0.5$ . In view of this intimate connection with Bayesian quantile regression as well as the subsequent Bayesian estimation and prediction, we also refer to the MAR model with Laplace errors given by (1) and (2) as Bayesian MAR, or BayesMAR, and use it exchangeably with the MAR model.

## 2.2 Prior specification and posterior sampling at a given $p$

We use an uninformative prior for  $\beta$  and the Jeffreys prior for  $\tau$ , namely,

$$\pi(\beta) \propto 1, \quad \pi(\tau) \propto \tau^{-1}.$$

The posterior distribution of  $(\beta, \tau)$  is

$$\pi(\beta, \tau | \mathbf{y}) \propto \pi(\beta, \tau) f(\mathbf{y} | \beta, \tau) \propto \tau^{-(T-p+1)} \exp \left\{ -\tau^{-1} \sum_{t=p+1}^T \frac{1}{2} |y_t - \mathbf{y}'_{t-1} \beta| \right\},$$

which is proper (Choi & Hobert, 2013). The regression coefficients  $\beta$  are instrumental for time series forecasting, and we derive their marginal posterior distributions by integrating out  $\tau$  for efficient sampling:

$$\begin{aligned} \pi(\beta | \mathbf{y}) &= \int \pi(\beta, \tau | \mathbf{y}) d\tau \propto \int \tau^{-(T-p+1)} \exp \left\{ -\tau^{-1} \sum_{t=p+1}^T \frac{1}{2} |y_t - \mathbf{y}'_{t-1} \beta| \right\} d\tau \\ &\propto \left\{ \sum_{t=p+1}^T \frac{1}{2} |y_t - \mathbf{y}'_{t-1} \beta| \right\}^{-(T-p)}. \end{aligned} \quad (4)$$

The posterior sampling of  $\beta$  proceeds by Markov chain Monte Carlo (MCMC) via the Metropolis-Hastings (MH) algorithm (Metropolis et al., 1953; Hastings, 1970):

1. At iteration  $i$ , draw a candidate sample from the proposal  $\beta_j^* = \beta_j^{i-1} + a \cdot u_j$  independently for  $j = 1, \dots, p$ , where  $\beta_j^{i-1}$  is the value of  $\beta_j$  at iteration  $i-1$ ,  $u_j$  follows a Uniform( $-0.1, 0.1$ ) distribution and the scalar  $a$  controls the step size of each move;
2. Accept  $\beta^*$  as  $\beta^i$  with probability  $p = \min \{1, \pi(\beta^* | \mathbf{y}) / \pi(\beta^{i-1} | \mathbf{y})\}$ . Otherwise, set  $\beta^i = \beta^{i-1}$ .

We tune the parameter  $a$  in Step 1 such that the final acceptance rate is between 20% and 50% (Gelman et al., 1996). In addition, we have also implemented Gaussian and heavy-tailed student- $t$  proposals, which are recommended by Gerlach et al. (2011) when studying extreme quantile levels under the CAViaR model. We did not observe empirical advantages of using such proposals over a uniform proposal under the MAR model, suggesting that one may choose more flexible proposals for median regression. For all experiments, we use 40,000 MCMC samples with 25,000 burn-ins, initialize  $\beta$  randomly within the unit interval, and use the posterior mean as the Bayes estimate of  $\beta$ . Trace plots indicate the MCMC samples converge quickly, mostly within thousands of iterations.

The proposed Bayesian approach provides a convenient method for density forecasting beyond

point forecasts. The 1-step ahead predictive density conditional on  $\mathbf{y}_{1:T} = (y_1, \dots, y_T)$  is

$$\begin{aligned}
p(y_{T+1} | \mathbf{y}_{1:T}) &= \int \int p(y_{T+1}, \boldsymbol{\beta}, \tau | \mathbf{y}_{1:T}) d\boldsymbol{\beta} d\tau \\
&= \int \int p(y_{T+1} | \boldsymbol{\beta}, \tau, \mathbf{y}_{1:T}) \cdot \pi(\boldsymbol{\beta}, \tau | \mathbf{y}_{1:T}) d\boldsymbol{\beta} d\tau \\
&\propto \int \int \tau^{-1} \exp\left(-\frac{1}{2\tau} |y_{T+1} - \mathbf{y}'_T \boldsymbol{\beta}|\right) \cdot \tau^{-(T-p+1)} \exp\left\{-\tau^{-1} \sum_{t=p+1}^T \frac{1}{2} |y_t - \mathbf{y}'_{t-1} \boldsymbol{\beta}|\right\} d\boldsymbol{\beta} d\tau \\
&\propto \int_{\mathbb{R}^{p+1}} \left[ \sum_{t=p+1}^{T+1} |y_t - \mathbf{y}'_{t-1} \boldsymbol{\beta}| \right]^{-(T-p+1)} d\boldsymbol{\beta}.
\end{aligned}$$

For large  $T$  the integrand  $\left[ \sum_{t=p+1}^T |y_t - \mathbf{y}'_{t-1} \boldsymbol{\beta}| \right]^{-(T-p+1)}$  may quickly decay to zero, leading to considerable numerical errors in direct evaluation of this integral. Alternatively, the MCMC samples of  $(\boldsymbol{\beta}, \tau)$  allow a convenient sampling strategy to approximate the predictive density. In particular, we draw samples of  $y_{T+1}$  from the Laplace likelihood  $p(y_{T+1} | \boldsymbol{\beta}, \tau, \mathbf{y}_{1:T})$  conditional on each posterior sample of  $(\boldsymbol{\beta}, \tau)$ . This strategy easily generalizes to  $q$ -step ahead predictive densities for any  $q \geq 2$  by drawing samples jointly for  $(y_{T+1}, \dots, y_{T+q})$  through iterative conditional distributions, which are all Laplace distributions.

### 2.3 Order selection and Bayesian model averaging

The order  $p$  in MAR( $p$ ) is typically unknown. We address the problem of unknown  $p$  in the Bayesian framework by putting a prior on  $p$ . In practice, we can usually specify a maximum order; otherwise, a  $p$  that is too large hampers the interpretability. We endow the order  $p$  with a uniform prior on  $\{1, 2, \dots, K\}$  with  $K$  being the specified maximum order. Then the posterior distribution of  $p$  in the prior support is

$$\pi(p | \mathbf{y}) \propto \pi(\mathbf{y} | p) \cdot \pi(p),$$

where  $\pi(\mathbf{y} | p) = \int_{\mathbb{R}^+} \int_{\mathbb{R}^{p+1}} \pi(\mathbf{y}, \boldsymbol{\beta}, \tau | p) \pi(\boldsymbol{\beta}) \pi(\tau) d\boldsymbol{\beta} d\tau$  is the marginal likelihood of  $p$ . The order  $p$  can be selected by using the maximum a posteriori (MAP) estimate:

$$\hat{p} = \underset{p \in \{1, \dots, K\}}{\operatorname{argmax}} \pi(p | \mathbf{y}).$$

For time series forecasting, a more appealing perspective is to use Bayesian model averaging (BMA) to propagate uncertainties in the model space, i.e.,

$$\hat{y}_{T+q} = \sum_{p=1}^K \pi(p | \mathbf{y}) \hat{y}_{T+q}^{(p)} \quad \text{and} \quad p(y_{T+q} | \mathbf{y}_{1:T}) = \sum_{p=1}^K \pi(p | \mathbf{y}) p(y_{T+q}^{(p)} | \mathbf{y}_{1:T}),$$

where  $\hat{y}_{T+q}^{(p)}$  and  $p(y_{T+q}^{(p)} | \mathbf{y}_{1:T})$  are the  $q$ -step ahead point prediction and predictive density of  $y_{T+q}$  under order  $p$ , respectively.

The main challenge in implementing MAP and BMA lies in prior specifications of model parameters and the evaluation of the marginal likelihood at given  $p$ . To date, a consensus on

the default choice for prior specifications in the context of model selection is still lacking, and one needs to be cautious about using improper priors—which are typically the default choice for a given model—in view of prior sensitivity and the Jeffreys-Lindley paradox. We refer interested readers to the rich literature on BMA, e.g., Cameron & Pettitt (2014); Robert (2001); Yao et al. (2018); Li & Dunson (2020), to name a few. We here resort to an approximation to  $\pi(p \mid \mathbf{y})$  using the Bayesian information criterion or BIC (Kass & Raftery, 1995; Neath & Cavanaugh, 2012), which is appealing as it eliminates the need to deal with prior specification and approximates Bayes factors reasonably well in certain cases (Kass & Wasserman, 1995). We observe that the BIC-based implementation tends to choose the oracle order with large probability in simulations.

Letting  $(\hat{\boldsymbol{\beta}}_{\text{MLE}}, \hat{\tau}_{\text{MLE}})$  be the maximum likelihood estimates (MLEs) of  $(\boldsymbol{\beta}, \tau)$ , then the BIC of  $\text{MAR}(p)$  is

$$\text{BIC}_p = (p + 2) \log(n) - 2 \log(L(\hat{\boldsymbol{\beta}}_{\text{MLE}}, \hat{\tau}_{\text{MLE}})), \quad (5)$$

where  $n$  is the sample size. We approximate  $\pi(p \mid \mathbf{y})$  by  $\exp\{-\text{BIC}_p/2\}$  up to multiplicative constants, leading to aggregated predictions

$$\hat{y} = \sum_{p=1}^K \pi(p \mid \mathbf{y}) \hat{y}^{(p)} \approx \sum_{p=1}^K \frac{\exp\{-\text{BIC}_p/2\}}{\sum_{i=1}^K \exp\{-\text{BIC}_i/2\}} \hat{y}^{(p)}. \quad (6)$$

It turns out that we can calculate the MLEs  $(\hat{\boldsymbol{\beta}}_{\text{MLE}}, \hat{\tau}_{\text{MLE}})$  efficiently. To see this, first note

$$(\hat{\boldsymbol{\beta}}_{\text{MLE}}, \hat{\tau}_{\text{MLE}}) = \underset{\boldsymbol{\beta}, \tau}{\text{argmax}} \tau^{-(T-p)} \exp \left\{ -\tau^{-1} \sum_{t=p+1}^T \frac{1}{2} |y_t - \mathbf{y}'_{t-1} \boldsymbol{\beta}| \right\}.$$

For any  $\tau > 0$ , the likelihood function  $L(\boldsymbol{\beta}, \tau)$  in Equation (3) attains its maximum at

$$\hat{\boldsymbol{\beta}}_{\text{MLE}} = \underset{\boldsymbol{\beta}}{\text{argmax}} \left\{ \sum_{t=p+1}^T \frac{1}{2} |y_t - \mathbf{y}'_{t-1} \boldsymbol{\beta}| \right\}^{-(T-p)} = \underset{\boldsymbol{\beta}}{\text{argmin}} \sum_{t=p+1}^T \frac{1}{2} |y_t - \mathbf{y}'_{t-1} \boldsymbol{\beta}|,$$

provided  $T > p$ . This corresponds to the estimators of minimizing absolute error in median regression, which can be efficiently solved by linear programming (Koenker, 2005). An analytical solution of  $\hat{\tau}_{\text{MLE}}$  is available through a Gamma kernel:

$$\begin{aligned} \hat{\tau}_{\text{MLE}} &= \underset{\tau}{\text{argmax}} \tau^{-(T-p)} \exp \left\{ -\tau^{-1} \frac{1}{2} \sum_{t=p+1}^T |y_t - \mathbf{y}'_{t-1} \hat{\boldsymbol{\beta}}_{\text{MLE}}| \right\} \\ &= \frac{\frac{1}{2} \sum_{t=p+1}^T |y_t - \mathbf{y}'_{t-1} \hat{\boldsymbol{\beta}}_{\text{MLE}}|}{T - p + 1}. \end{aligned}$$

Before substituting  $\hat{\boldsymbol{\beta}}_{\text{MLE}}$  and  $\hat{\tau}_{\text{MLE}}$  into Equation (5), we notice that both the likelihood function and sample size depend on the order  $p$ . To reconcile various sample sizes at different orders, we use the last  $T - K$  samples to evaluate the likelihood function for any  $p$ . Consequently,  $\text{BIC}_p$  is given by

$$\text{BIC}_p = (p + 2) \log(T - K) - 2 \log \left[ (4\hat{\tau}_{\text{MLE}})^{-(T-K)} \exp \left( -\frac{1}{2\hat{\tau}_{\text{MLE}}} \sum_{t=K+1}^T |y_t - \mathbf{y}'_{t-1} \hat{\boldsymbol{\beta}}_{\text{MLE}}| \right) \right].$$

The same methods to approach the unknown order and provide predictive densities apply to AR, where the likelihood function changes to Gaussian and the MLE of  $\beta$  that resembles the least square estimates has a simple closed-form expression.

### 3 Simulation

In this section, we conduct simulations to assess the performances of BayesMAR with mean-based methods, focusing on parameter estimation under various model assumptions. To this end, we choose the AR model and the generalized autoregressive conditional heteroscedasticity model (GARCH), and defer predictive comparisons and more recent mean-based methods to real data application in Section 4.

We generate data according to the model

$$y_t = \beta_0 + \beta_1 y_{t-1} + \beta_2 y_{t-2} + \varepsilon_t,$$

where  $\beta = (\beta_0, \beta_1, \beta_2) = (0.3, 0.75, -0.35)$ . We consider two scenarios depending on the distribution of  $\varepsilon_t$ : Gaussian error where  $\varepsilon_t \sim N(0, 1)$  and Laplace error where  $\varepsilon_t \sim \text{Laplace}(0, 1)$ , which correspond to the model assumptions of the AR and BayesMAR models, respectively.

For each error assumption, we generate 200 observations and replicate such simulation 100 times. We estimate AR models via a Bayesian procedure with priors  $\pi(\beta) \propto 1$  and  $\pi(\sigma) \propto \sigma^{-1}$ . We also conduct the maximum likelihood estimation (Gardner et al., 1980) for AR using the ‘*arima*’ function in the R package **stats**, which leads to almost identical performance and is thus not reported here. As such, we use AR and BayesAR exchangeably throughout this paper. We use AR( $p$ )-GARCH(1,1) when implementing GARCH, i.e.,

$$\begin{aligned} y_t &= \beta_0 + \beta_1 y_{t-1} + \beta_2 y_{t-2} + \dots + \beta_p y_{t-p} + e_t; \\ e_t &= \sqrt{h_t} \eta_t, \quad \eta_t \sim N(0, 1); \\ h_t &= \omega_0 + \alpha_1 e_{t-1}^2 + \alpha_2 h_{t-1}. \end{aligned}$$

We fit the model using the R package **fGarch**, where all parameters are estimated by quasi-maximum likelihood (Bollerslev & Wooldridge, 1992). In addition, we implement the Quantile Autoregression (QAR) method proposed by Koenker & Xiao (2006) using the R package **quantreg**, to compare its finite sample performance with BayesMAR.

We assess estimates of  $(\beta_0, \beta_1, \beta_2)$  by each method based on mean squared error (MSE). For a generic parameter  $\theta$ , letting the estimate be  $\hat{\theta}_i$  in the  $i$ th simulation and  $\bar{\theta} = \frac{1}{100} \sum_{i=1}^{100} \hat{\theta}_i$ , then the MSE and its standard error are

$$\text{MSE} = \frac{1}{100} \sum_{i=1}^{100} (\hat{\theta}_i - \theta)^2, \quad \text{SE}_{\text{MSE}} = \frac{1}{10} \sqrt{\text{sample variance of } \{(\hat{\theta}_i - \theta)^2\}_{i=1}^{100}}.$$

We first provide the true order  $p = 2$  to all models and compare their performances. Table 1 reports the MSE of all methods. We can see that all methods benefit from a correctly specified error distribution: AR and GARCH have the smallest MSEs under Gaussian error, while BayesMAR and QAR have smaller MSEs when data are generated from Laplace distributions. However, BayesMAR appears to suffer less than AR from model misspecification; for example, the increase of MSE of  $\beta_2$  under Gaussian error from AR to BayesMAR is 0.13, which is within two standard errors, while

Table 1: MSE of all methods under Gaussian error and Laplace error. Standard errors are reported below MSEs. All summaries have been multiplied by  $10^2$ .

Models	Error	Gaussian			Laplace		
		$\beta_0$	$\beta_1$	$\beta_2$	$\beta_0$	$\beta_1$	$\beta_2$
BayesMAR	MSE	1.07	0.56	0.63	0.73	0.27	0.19
	SE	0.13	0.07	0.07	0.11	0.04	0.03
QAR	MSE	1.21	0.62	0.75	0.75	0.27	0.19
	SE	0.15	0.08	0.08	0.11	0.05	0.02
AR	MSE	0.77	0.41	0.50	1.08	0.50	0.43
	SE	0.09	0.05	0.08	0.16	0.06	0.06
GARCH	MSE	0.76	0.41	0.51	1.08	0.50	0.40
	SE	0.09	0.05	0.08	0.16	0.06	0.05

AR doubles the MSE of  $\beta_2$ , and so is beyond three standard errors of BayesMAR under Laplace error. It is reassuring that BayesMAR gives either the same or better MSEs than QAR in all cases, although all differences are within one standard error. This finite sample performance is consistent with the findings in Gerlach et al. (2011) when comparing sampling-based Bayesian approaches with optimization-based counterparts for extreme quantile levels.

We next investigate the selection of the unknown order  $p$  in BayesMAR using the BIC approach described in Section 2.3. We provide a large upper bound  $K = 20$  for the order  $p$ . Figure 1 plots the distribution of  $p$  in both scenarios and suggest that the selected orders almost always concentrate around the oracle value  $p = 2$ , even when the model is misspecified under Gaussian error. The overall accuracy across all simulations to select  $p = 2$  using MAP is 95% for normal errors and 98% for Laplace errors.

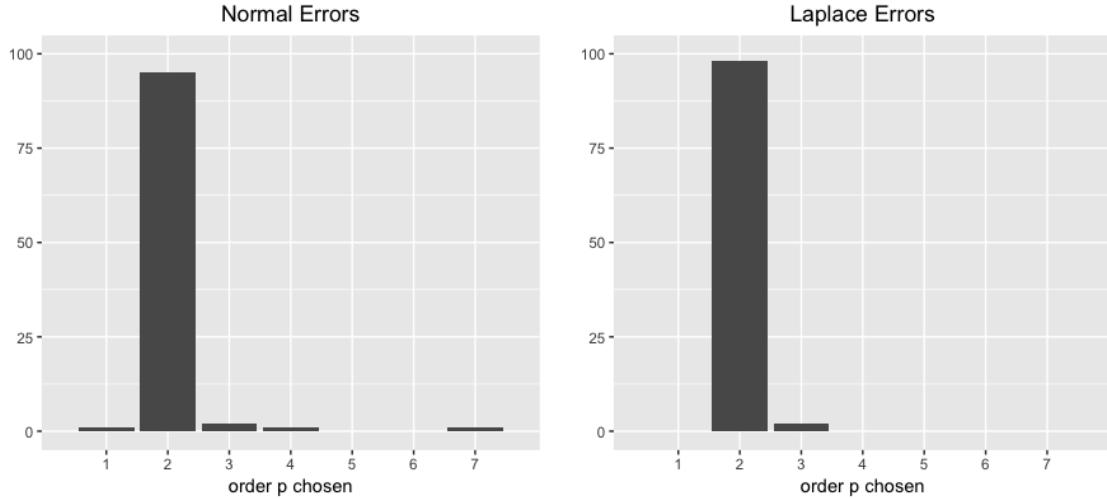


Figure 1: Distributions of selected orders using BIC.

## 4 Real Data Application

In this section, we compare the predictive performances of the proposed methods to that of mean-based methods using various real-world data. We use three economic series from Federal Reserve Economic Data: the quarterly data Producer Price Index for all commodities (FRED: U.S. Bureau of Labor Statistics, 2018), 3-Month Treasury Bill: secondary market Rate (FRED: Board of Governors of the Federal Reserve System (US), 2018), and Unemployment Rate (FRED: Organization for Economic Co-operation and Development, 2018), coded as PPI, TBR, and UR, respectively. Each time, the series ranges from 1968Q3 to 2018Q2, containing 200 observations. The unemployment rate data is seasonally adjusted by the method of seasonal-trend decomposition using Loess (Cleveland et al., 1990). The three time series ( $y_t$ ) and the lagged data of order one ( $y_t - y_{t-1}$ ) are plotted in Figure 2. These three data sets have distinct patterns: the lagged PPI tends to be more stable before the crisis in the year of 2008, in contrast to the substantial fluctuation since then; the lagged TBR has more dramatic changes in earlier periods than in later periods; and the lagged UR appears to contain several extreme values, while a periodic pattern may still persist even after seasonal adjustment. These complex characteristics of the data enable a comparison between model-based methods when there is no guarantee for model assumptions to hold.

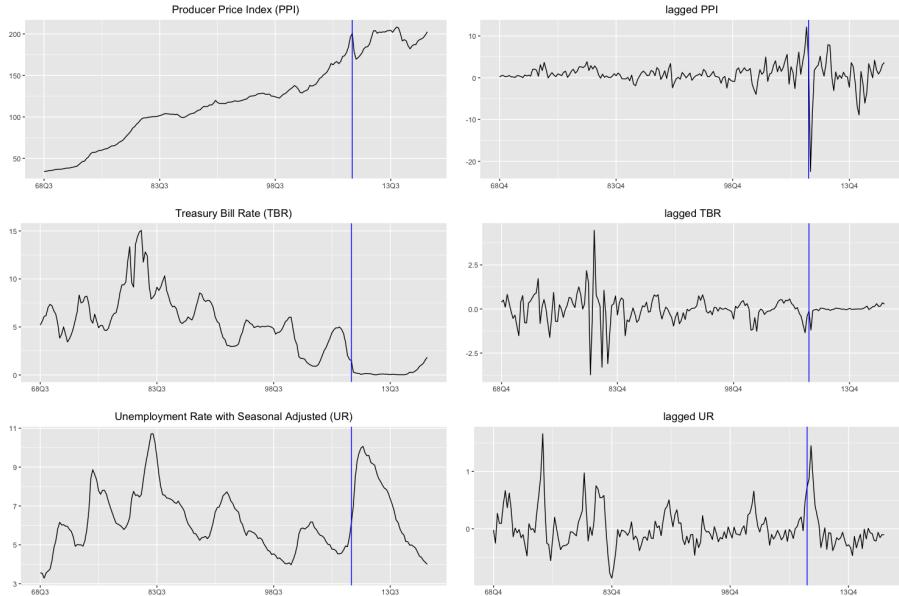


Figure 2: Three time series (PPI, TBR, and UR) and the lagged data. The blue line is  $t_0 = 08Q3$ .

In addition to AR and GARCH, we implement selected methods proposed in research on the dynamic linear model. In particular, we consider the time-varying vector AR (TV-VAR) model proposed by Nakajima & West (2013), which links time-varying parameters to latent threshold processes and achieves state-of-the-art predictive performance in selected applications. A vector with a length of three that stacks the three time series  $y_t^{(3)}$ , expands the dynamic linear model by

utilizing a latent threshold vector  $\mathbf{d} = (d_1, \dots, d_k)$  with  $d_i \geq 0$  for  $i = 1, \dots, k$  :

$$\begin{aligned}\mathbf{y}_t^{(3)} &= \mathbf{c}_t + \mathbf{B}_{1t}\mathbf{y}_{t-1}^{(3)} + \dots + \mathbf{B}_{pt}\mathbf{y}_{t-p}^{(3)} + \mathbf{u}_t, \quad \mathbf{u}_t \sim N(\mathbf{0}, \Sigma_t), \quad \mathbf{b}_t = (b_{1t}, \dots, b_{kt})', \\ b_{it} &= \beta_{it}s_{it} \text{ with } s_{it} = \mathbb{1}(|\beta_{it}| \geq d_i), \\ \beta_t &= \boldsymbol{\mu}_\beta + \Phi_\beta(\beta_{t-1} - \boldsymbol{\mu}_\beta) + \boldsymbol{\eta}_{\beta_t}, \quad \boldsymbol{\eta}_{\beta_t} \sim N(\mathbf{0}, \mathbf{V}_\beta),\end{aligned}$$

where  $\mathbf{c}_t$  is the  $3 \times 1$  vector of time-varying intercepts,  $\mathbf{B}_{jt}$  is the  $3 \times 3$  matrix of time-varying coefficients at lag  $j$ ,  $\mathbf{b}_t$  stacks  $\mathbf{c}_t$  and  $\mathbf{B}_{jt}$  by rows and by order  $j$  from 1 to  $p$ ,  $k = 3(1 + 3p)$ , and  $\beta_t = (\beta_{1t}, \dots, \beta_{kt})$  is a latent time-varying parameter vector whose dynamic updates are specified by  $(\boldsymbol{\mu}_\beta, \Phi_\beta, \mathbf{V}_\beta)$  via a standard VAR model. The latent threshold vector  $\mathbf{d}$  shrinks the time-varying coefficients  $\beta_t$  to zero, leading to dynamic sparsity and improved prediction. We implement two versions of TV-VAR: TV-VAR without latent threshold (NT) and TV-VAR with latent threshold (LT). Both methods are applicable to multivariate time series, and we adapt the OxMetrics code provided in Nakajima & West (2013) to the vector observations consisting of the three time series.

For BayesMAR, we use BayesMAR-BMA to denote the Bayesian model averaging strategy and BayesMAR-MAP when the MAP estimate of  $p$  is used, and we adopt the same convention for AR, i.e., BayesAR-BMA and BayesAR-MAP, which are implemented similarly to BayesMAR but with the Gaussian likelihood. GARCH(1,1) uses the same autoregressive order  $p$  as chosen in BayesAR-MAP; although orders other than (1, 1) can be used, we observe that the GARCH(1,1) consistently led to the best predictive performances of GARCH in our real data application. Since there are no immediately available model selection methods for NT or LT, we run the order from 1 to 5, and then choose the optimal results to favor NT and LT under each criterion. We observe that LT considerably outperformed NT in all scenarios, and thus only present the results for LT.

We use recursive out-of-sample forecasting to assess each method after  $t_0=2008Q3$ . In particular, we fit each model with data up to quarter  $t$  and conduct an  $h$ -step-ahead prediction for  $h = 1, 2, 3, 4$ . Then we move one period ahead and repeat the same procedure until we reach  $t = T$ . All methods are applied to the lagged data of order one to remove local trends and forecast the changes that lead to predictions of  $y_t$ .

We calculate the root MSE (RMSE) and mean absolute error (MAE) from  $t_0$  to  $T$  to compare performance of each method, i.e.,

$$\text{RMSE} = \sqrt{\frac{1}{35} \sum_{i=1}^{35} (y_{t_0+i} - \hat{y}_{t_0+i})^2}, \quad \text{MAE} = \frac{1}{35} \sum_{i=1}^{35} |y_{t_0+i} - \hat{y}_{t_0+i}|.$$

We additionally calculate the the relative change using BayesMAR-BMA as the reference to ease comparison, that is,

$$\text{Relative change} = \left[ \frac{\text{RMSE (or MAE) of each method}}{\text{RMSE (or MAE) of BayesMAR-BMA}} - 1 \right] \times 100\%.$$

Table 2 reports the RMSEs of all methods for the three data sets. The results indicate that BayesMAR outperforms alternatives in nearly all scenarios. In particular, the RMSEs of BayesMAR-BMA is uniformly smaller than its AR counterpart. Compared to BayesMAR-BMA, the methods of AR, GARCH, and LT increase the RMSE by 10% to 89.2% for TBR, and 3.0% to 35.6% for UR, with two exceptions for the UR when compared to GARCH at three- and four-steps-ahead predictions where GARCH slightly reduces the RMSE by 0.3% and 2.8%, respectively. Although

the gap in the prediction for PPI is smaller (between 0.8% to 8.9%), all competing methods have a larger RMSE than BayesMAR-BMA. GARCH accounts for comprehensive variance structures and LT dynamically updates regression coefficients, which are arguably powerful methods with considerable complexity; it is remarkable that the proposed simple BayesMAR leads to favorable and often superior performance using real-world data.

Table 3 reports the MAEs of all methods, suggesting similar observations as those made from Table 2. The proposed BayesMAR methods give the smallest MAEs in nearly all scenarios, with one exception for the UR when compared to GARCH.

The RMSEs of BayesMAR-BMA and BayesMAR-MAP are close to each other with at most 1.1% (PPI), -0.6% (TBR), 2.3% (UR) relative differences, and similar observations hold for MAEs. This is partly caused by highly concentrated weights of model orders when using BIC for order selection. In particular, we find BayesMAR-BMA puts most weight on the order selected in BayesMAR-MAP, leading to minimal differences between the two variants of BayesMAR. Comparing BayesAR-BMA and BayesAR-MAP leads to the same conclusion for AR.

Figure 3 compares up to four-steps-ahead absolute predictive errors at each  $t$  from  $t_0$  to  $T$  given by the three static models, MAR, AR, and GARCH, which provides insights into understanding the performance of BayesMAR. We choose AR and GARCH and implement MAR and AR using MAP such that all methods in the figure build on similar model structures. We can see that AR yields a large prediction error at the beginning (PPI), which is substantially reduced by GARCH, which incorporates heterogeneous variance structures. It is reassuring that the proposed BayesMAR method, which uses a simple constant variance structure, achieves the same predictive gain (for PPI) or even further improvements (for TBR and UR). At later time points when the time series stabilizes without usually large deviations, BayesMAR tends to perform similarly to AR and GARCH. Since BayesMAR is a parametric model bearing the same model structure as AR, these comparisons suggest that further performance gains may be possible by following the rich literature that extends AR to more advanced models such as GARCH and dynamic models, by simply altering the error assumption from Gaussian to Laplace.

We plot the 95% credible intervals of AR and MAR in Figure 4, both using the BMA version. For both AR and MAR, four-steps-ahead predictive intervals appear wider than one-step-ahead predictive intervals. This makes sense, as more uncertainty propagates, and this is particularly the case for the UR data. MAR is more robust than AR at the early stage of the PPI time series. Since an informative comparison between MAR and AR through this visualization seems difficult, we next turn to numerical comparisons.

The proposed method provides density forecasts in addition to point forecasts; see Section 2. We evaluate the probabilistic forecasts produced by the proposed method and AR through the continuous ranked probability score (CRPS) (Gneiting et al., 2007), which is implemented in the R package **scoringRules**. We can see that the proposed MAR model produces a considerably smaller CRPS than AR for the TBR and UR data, while being close to AR for PPI. This suggests that the advantage of MAR may extend to probabilistic forecasts beyond point forecasts (see Table 4).

Table 2: Forecasting Performance for U.S. macroeconomic data after 2007-2008 financial crisis: RMSEs for  $h$ -step ahead prediction for  $h = 1, 2, 3, 4$ . The relative change uses BayesMAR-BMA as the baseline. RMSEs for Treasury Bill Rate and Unemployment Rate have been multiplied by 10.

Producer Price Index								
Models	RMSE				Relative change (%)			
	1	2	3	4	1	2	3	4
BayesMAR-BMA	3.18	5.94	7.74	9.24	-	-	-	-
BayesMAR-MAP	3.21	5.94	7.69	9.13	0.7	0.1	-0.7	-1.1
BayesAR-BMA	3.37	6.24	8.04	9.39	6.0	5.1	3.8	1.6
BayesAR-MAP	3.38	6.26	8.07	9.43	6.0	5.4	4.2	2.1
GARCH	3.21	6.05	7.85	9.31	0.8	1.9	1.4	0.8
LT	3.47	6.06	7.94	9.45	8.9	2.0	2.5	2.3
Treasury Bill Rate								
Models	RMSE				Relative change (%)			
	1	2	3	4	1	2	3	4
BayesMAR-BMA	0.91	1.49	2.16	2.94	-	-	-	-
BayesMAR-MAP	0.90	1.48	2.16	2.93	-0.9	-0.1	-0.1	-0.1
BayesAR-BMA	1.63	2.10	2.99	4.16	79.7	41.6	38.2	41.7
BayesAR-MAP	1.71	2.37	3.31	4.56	89.2	59.3	53.1	55.1
GARCH	1.66	2.02	2.63	3.4	83.6	35.9	21.7	18.5
LT	1.00	1.87	2.66	3.55	10.3	25.6	23.0	21.0
Unemployment Rate								
Models	RMSE				Relative change (%)			
	1	2	3	4	1	2	3	4
BayesMAR-BMA	1.76	3.07	4.34	5.85	-	-	-	-
BayesMAR-MAP	1.77	3.11	4.42	5.94	0.5	1.2	1.8	1.6
BayesAR-BMA	1.88	3.37	4.88	6.88	7.2	9.5	12.3	17.6
BayesAR-MAP	1.88	3.42	5.11	7.18	6.8	11.3	17.5	22.8
GARCH	1.86	3.17	4.33	5.69	5.6	3.0	-0.3	-2.8
LT	2.40	4.17	5.53	7.17	36.7	35.6	27.4	22.7

Table 3: Forecasting Performance for U.S. macroeconomic data after 2007-2008 financial crisis: MAEs for  $h$ -step ahead prediction for  $h = 1, 2, 3, 4$ . The relative change uses BayesMAR-BMA as the baseline. MAEs for Treasury Bill Rate and Unemployment Rate have been multiplied by 10.

Producer Price Index								
Models	MAE				Relative change (%)			
	1	2	3	4	1	2	3	4
BayesMAR-BMA	2.50	4.44	5.90	7.26	-	-	-	-
BayesMAR-MAP	2.50	4.44	5.84	7.18	0.2	0.1	-1.0	-1.1
BayesAR-BMA	2.69	4.49	6.05	7.33	7.8	1.2	2.5	1.0
BayesAR-MAP	2.70	4.54	6.05	7.37	8.0	2.2	2.5	1.6
GARCH	2.54	4.59	5.98	7.29	1.8	3.5	1.4	0.4
LT	2.72	4.47	6.29	7.45	8.9	0.8	6.6	2.6
Treasury Bill Rate								
Models	MAE				Relative change (%)			
	1	2	3	4	1	2	3	4
BayesMAR-BMA	0.56	1.05	1.56	2.10	-	-	-	-
BayesMAR-MAP	0.56	1.05	1.56	2.10	-0.6	-0.1	0.0	0.0
BayesAR-BMA	1.09	1.61	2.25	3.04	95.5	52.7	44.2	44.6
BayesAR-MAP	1.14	1.80	2.52	3.43	103.2	70.5	61.6	63.1
GARCH	0.99	1.43	2.01	2.68	77.8	35.7	28.7	27.6
LT	0.65	1.23	1.78	2.37	16.0	16.2	13.9	13.0
Unemployment Rate								
Models	MAE				Relative change (%)			
	1	2	3	4	1	2	3	4
BayesMAR-BMA	1.24	2.30	3.49	4.85	-	-	-	-
BayesMAR-MAP	1.26	2.34	3.57	4.95	1.2	1.4	2.3	2.1
BayesAR-BMA	1.47	2.61	3.91	5.73	18.1	13.1	12.2	18.2
BayesAR-MAP	1.49	2.69	4.16	5.94	19.8	16.9	19.2	22.4
GARCH	1.27	2.12	3.16	4.27	2.4	-8.1	-9.5	-11.9
LT	1.66	3.26	4.63	6.32	33.4	41.7	32.6	30.4

Table 4: Forecasting Performance for U.S. macroeconomic data after 2007-2008 financial crisis: CRPS for  $h$ -step ahead prediction for  $h = 1, 2, 3, 4$ . The relative change uses BayesMAR-BMA as the baseline. CRPS for Treasury Bill Rate and Unemployment Rate have been multiplied by 10.

Producer Price Index									
Models	CRPS				Relative change (%)				
	1	2	3	4	1	2	3	4	
BayesMAR-BMA	1.93	2.12	2.02	2.01	-	-	-	-	
BayesMAR-MAP	1.95	2.13	2.00	2.00	1.1	0.8	-0.8	-0.8	
BayesAR-BMA	1.93	2.08	2.13	2.05	0.1	-1.8	5.7	2.1	
BayesAR-MAP	1.93	2.08	2.14	2.06	0.3	-1.5	6.1	2.6	
Treasury Bill Rate									
Models	CRPS				Relative change (%)				
	1	2	3	4	1	2	3	4	
BayesMAR-BMA	1.22	1.37	1.43	1.45	-	-	-	-	
BayesMAR-MAP	1.22	1.37	1.42	1.45	0.0	0.1	-0.3	0.0	
BayesAR-BMA	1.80	1.89	1.92	1.97	48.4	37.5	34.2	35.6	
BayesAR-MAP	1.81	1.88	1.91	1.97	48.5	37.3	33.8	35.7	
Unemployment Rate									
Models	CRPS				Relative change (%)				
	1	2	3	4	1	2	3	4	
BayesMAR-BMA	0.93	0.97	1.04	1.16	-	-	-	-	
BayesMAR-MAP	0.93	0.98	1.05	1.16	0.7	0.7	0.7	0.5	
BayesAR-BMA	1.04	1.14	1.22	1.42	12.6	17.0	17.9	22.2	
BayesAR-MAP	1.04	1.13	1.26	1.45	11.9	16.5	21.1	24.7	

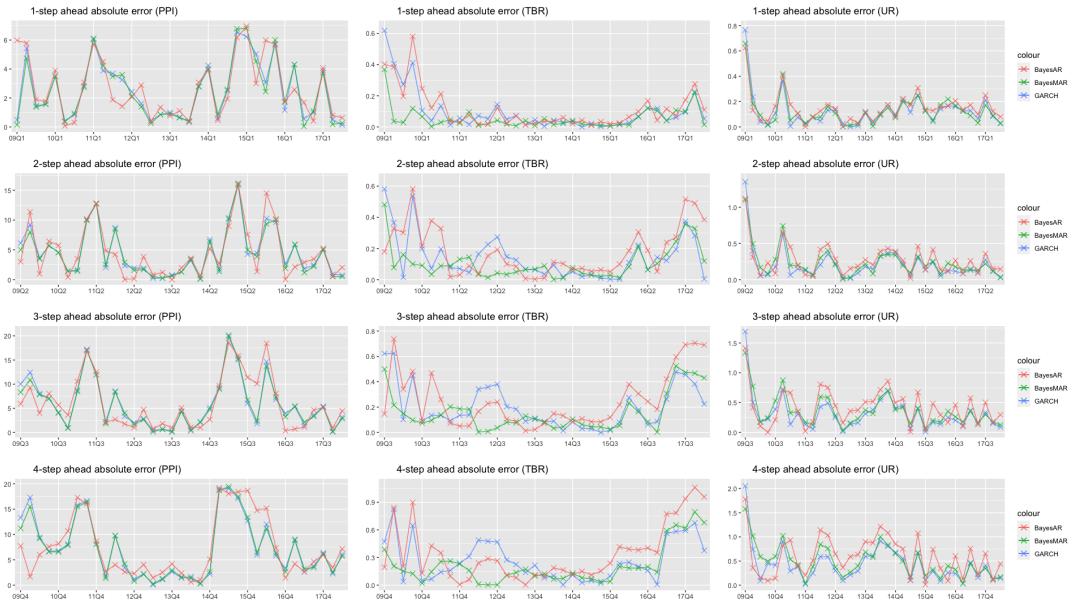


Figure 3: Absolute  $h$ -step ahead predictive errors at each period from 2008Q3 for  $h = 1, 2, 3, 4$ . Orders of BayesMAR and BayesAR are selected by the BIC.

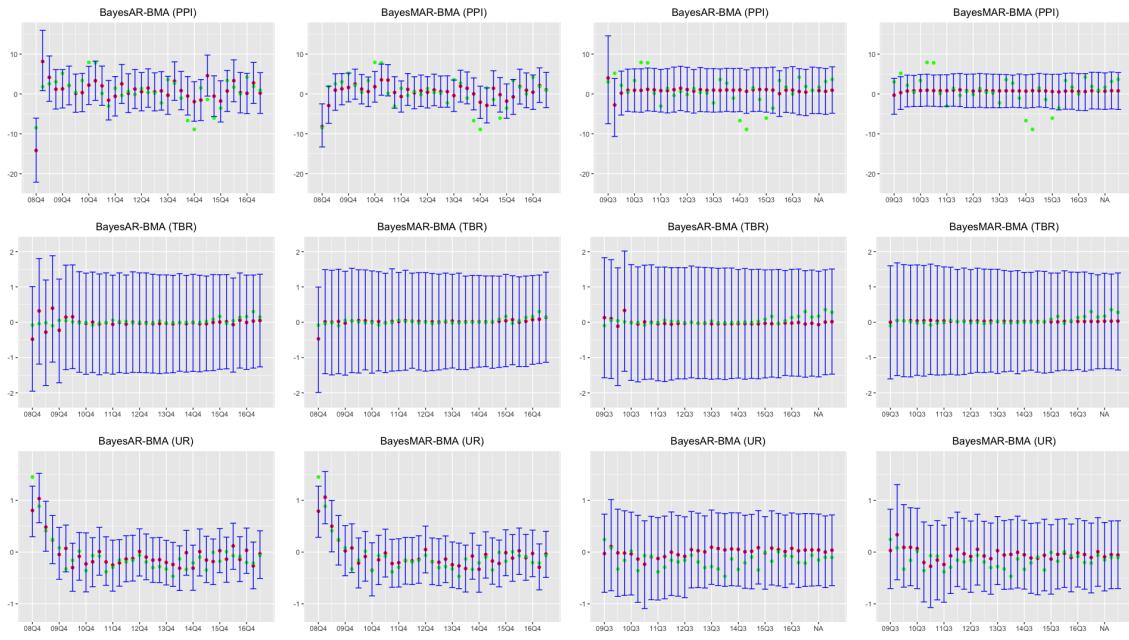


Figure 4: 1-step ahead (the first two columns) and 4-step ahead (the last two columns) 95% credible intervals of MAR and AR using BMA. True observations are points marked green, and predictions are marked red. Blue bars are the 95% credible intervals.

## 5 Concluding remarks

This article proposes a Bayesian median autoregressive (BayesMAR) model for robust time series forecasting. The proposed method has close connections with time-varying quantile regression. BayesMAR adopts a parametric model bearing the same structure as AR models by altering the Gaussian error to Laplace, leading to a simple, robust, and interpretable modeling strategy for time series forecasting with principled uncertainty quantification through Bayesian model averaging and Bayesian model selection. Real data applications using U.S. macroeconomic data show that BayesMAR leads to favorable and often superior predictive performances compared to the selected state-of-the-art mean-based alternatives under various loss functions that encompass both point and probabilistic forecasts.

BayesMAR enjoys practical benefits as a technical tool to introduce robustness and improve predictions. In addition, since the autoregressive structure in BayesMAR resembles the widely used AR model, BayesMAR can be used to complement a rich class of methods that build on the AR model. The AR model is arguably one of the most popular methods in time series, serving as the building block for other models such as GARCH and TV-VAR in research on the dynamic linear model. The proposed MAR model shows the potential for further performance gains by following the rich literature that extends AR to more advanced models with Laplace error assumptions rather than Gaussian ones.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

This work was partially supported by the grant DMS-2015569 from the National Science Foundation.

## Data availability

The R code to implement the proposed methods is publicly available at <https://github.com/xylimeng/BayesMAR>.

## References

Akouemo, H. N., & Povinelli, R. J. (2014). Time series outlier detection and imputation. In *Proceedings of IEEE PES General Meeting | Conference & Exposition* (pp. 1–5).

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31, 307–327.

Bollerslev, T., & Wooldridge, J. M. (1992). Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews*, 11, 143–172.

Cameron, E., & Pettitt, A. (2014). Recursive pathways to marginal likelihood estimation with prior-sensitivity analysis. *Statistical Science*, 29, 397–419.

Casini, A., & Perron, P. (2019). Structural breaks in time series. *Oxford Research Encyclopedia of Economics and Finance*, .

Chen, C., & Liu, L.-M. (1993a). Forecasting time series with outliers. *Journal of Forecasting*, 12, 13–35.

Chen, C., & Liu, L.-M. (1993b). Joint estimation of model parameters and outlier effects in time series. *Journal of the American Statistical Association*, 88, 284–297.

Chen, C. W., & So, M. K. (2006). On a threshold heteroscedastic model. *International Journal of Forecasting*, 22, 73–89.

Choi, H. M., & Hobert, J. P. (2013). Analysis of mcmc algorithms for bayesian linear regression with laplace errors. *Journal of Multivariate Analysis*, 117, 32–40.

Cleveland, R. B., Cleveland, W. S., McRae, J. E., & Terpenning, I. (1990). Stl: A seasonal-trend decomposition. *Journal of official statistics*, 6, 3–73.

Croux, C., Gelper, S., & Fried, R. (2008). Computational aspects of robust holt-winters smoothing based on m-estimation. *Applications of Mathematics*, 53, 163.

Engle, R. F., & Manganelli, S. (2004). Caviar: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics*, 22, 367–381.

Fan, J., & Yao, Q. (2008). *Nonlinear Time Series: Nonparametric and Parametric Methods*. Springer Science & Business Media.

Ferraty, F., & Vieu, P. (2006). *Nonparametric Functional Data Analysis: Theory and Practice*. Springer Science & Business Media.

Fox, A. J. (1972). Outliers in time series. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34, 350–363.

FRED: Board of Governors of the Federal Reserve System (US) (2018). 3-Month Treasury Bill: Secondary Market Rate. Retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/TB3MS>, Aug 09, 2018.

FRED: Organization for Economic Co-operation and Development (2018). Unemployment Rate: Aged 15-64: All Persons for the United States. Retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/LRUN64TTUSQ156N>, Aug 09, 2018.

FRED: U.S. Bureau of Labor Statistics (2018). Producer Price Index for All Commodities. Retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/PPIACO>, Aug 09, 2018.

Gardner, G., Harvey, A. C., & Phillips, G. D. A. (1980). An algorithm for exact maximum likelihood estimation of autoregressive-moving average models by means of Kalman filtering. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 29, 311–322.

Gelman, A., Roberts, G. O., & Gilks, W. R. (1996). Efficient metropolis jumping rules. In J. M. Bernardo, J. O. Berger, A. P. Dawid, & A. F. M. Smith (Eds.), *Bayesian Statistics 5* (pp. 599–607). Oxford.

Gerlach, R. H., Chen, C. W., & Chan, N. Y. C. (2011). Bayesian time-varying quantile forecasting for value-at-risk in financial markets. *Journal of Business & Economic Statistics*, 29, 481–492.

Geweke, J., & Keane, M. (2007). Smoothly mixing regressions. *Journal of Econometrics*, 138, 252–290.

Giacomini, R., & Komunjer, I. (2005). Evaluation and combination of conditional quantile forecasts. *Journal of Business & Economic Statistics*, 23, 416–431.

Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69, 243–268.

Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57, 97–109.

Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. (1999). Bayesian model averaging: a tutorial (with discussion). *Statistical Science*, 14, 382–417.

Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and Practice*. OTexts.

Jose, V. R. R., & Winkler, R. L. (2008). Simple robust averages of forecasts: Some empirical results. *International Journal of Forecasting*, 24, 163–169.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773–795.

Kass, R. E., & Wasserman, L. (1995). A reference bayesian test for nested hypotheses and its relationship to the schwarz criterion. *Journal of the American Statistical Association*, 90, 928–934.

Kleijn, B. J. K., & van der Vaart, A. W. (2006). Misspecification in infinite-dimensional Bayesian statistics. *The Annals of Statistics*, 34, 837–877.

Koenker, R. (2005). *Quantile regression*. Cambridge University Press: Cambridge, UK.

Koenker, R., & Bassett Jr, G. (1978). Regression quantiles. *Econometrica: Journal of the Econometric Society*, (pp. 33–50).

Koenker, R., & Xiao, Z. (2006). Quantile autoregression. *Journal of the American Statistical Association*, 101, 980–990.

Li, M., & Dunson, D. B. (2020). Comparing and weighting imperfect models using d-probabilities. *Journal of the American Statistical Association*, 115, 1349–1360.

Liu, Y., Li, M., & Morris, J. S. (2020). Function-on-scalar quantile regression with application to mass spectrometry proteomics data. *The Annals of Applied Statistics*, 14, 521–541.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21, 1087–1092.

Nakajima, J., & West, M. (2013). Bayesian analysis of latent threshold dynamic models. *Journal of Business & Economic Statistics*, 31, 151–164.

Neath, A. A., & Cavanaugh, J. E. (2012). The bayesian information criterion: background, derivation, and applications. *Wiley Interdisciplinary Reviews: Computational Statistics*, 4, 199–203.

Oka, T., & Qu, Z. (2011). Estimating structural changes in regression quantiles. *Journal of Econometrics*, 162, 248–267.

Prado, R., & West, M. (2010). *Time Series: Modeling, Computation, and Inference*. Chapman and Hall/CRC.

Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies*, 72, 821–852.

Qu, Z. (2008). Testing for structural change in regression quantiles. *Journal of Econometrics*, 146, 170–184.

Robert, C. P. (2001). *The Bayesian Choice*. Springer.

Sriram, K., Shi, P., & Ghosh, P. (2016). A bayesian quantile regression model for insurance company costs data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 179, 177–202.

Tukey, J. W. (1974). Nonlinear (nonsuperposable) methods for smoothing data. *CONGRESS RECORD (EASCO)*, (p. 673).

Yang, Y., Wang, H. J., & He, X. (2016). Posterior inference in bayesian quantile regression with asymmetric laplace likelihood. *International Statistical Review*, 84, 327–344.

Yao, Y., Vehtari, A., Simpson, D., & Gelman, A. (2018). Using stacking to average bayesian predictive distributions. *Bayesian Analysis*, 13, 917–1007.

Youngman, B. D. (2018). Generalized additive models for exceedances of high thresholds with an application to return level estimation for us wind gusts. *Journal of the American Statistical Association*, (pp. 1–31). To appear.

Yu, K., & Moyeed, R. A. (2001). Bayesian quantile regression. *Statistics & Probability Letters*, 54, 437–447.