SCIREA Journal of Mathematics
ISSN: 2995-5823
http://www.scirea.org/journal/Mathematics
March 11, 2026
Volume 11, Issue 1, February 2026
https://doi.org/10.54647/mathematics110563

Bayesian Estimation of Nonlinear Time Series Models with Real Data Applications

Nada Hussein Tali

College of Administration & Economics, University of Sumer, Dhi Qar, Iraq

Email: nada.hussein@uos.edu.iq

ORCID: 0009-0009-6974-7958

Abstract

Nonlinear time series have become very popular in applied research as they can capture a range of complex dynamic properties, such as volatility clustering, regime shifts and non-Gaussian features that are often present in financial time series. In this paper, a general Bayesian framework for estimation and model comparison of nonlinear time series models in state space form is developed. The presented approach also achieves coolness by summoning the witch's trinity of flexible model specification, coherent prior modelling and state-of-the-art computational inference techniques: namely, Markov chain Monte Carlo (MCMC) with data augmentation particles and variational methods. To evaluate finite-sample performance at different sample sizes, we carried out an extensive simulation study based on a stochastic volatility data-generating process. The results show that exact Bayesian methods give accurate estimates of parameters, reliable recovery in the latent state, and good calibration of quantification uncertainty, with particle Markov chain Monte Carlo (particle MCMC) performing best for small samples and predictive accuracy. VI is highly efficient computationally, but it systematically underestimates posterior uncertainty. The practical utility of the developed framework is also demonstrated through two real-data applications dealing with financing and count-valued time series. Model fit is assessed through principled Bayesian predictive diagnostics such as one-step cross-validation, information criteria and posterior predictive checks. All in all, the proposed Bayesian formulation provides a generic and consistent methodology to solve both the likelihood-based and prediction problems in nonlinear time-series analysis.

Keywords: Index Terms- Bayesian inference, nonlinear time series, state-space models, stochastic volatility, particle Markov chain Monte Carlo, variational inference.

1. Introduction

Time series is a ubiquitous data type in science, with applications in physics, economics, finance, and weather forecasting, among other fields. The goal of time series analysis is to model temporal dependence, to infer underlying dynamical processes, quantify uncertainty and make accurate forecasts. Classical linear time models like autoregressive (AR), moving average (MA), and autoregressive integrated moving average (ARIMA) have played an important precursor role for centuries on account of their analytical ease and numerical simplicity. Yet, such models are based on restrictive assumptions of linearity, stationarity and Gaussian error distributions, which are often not satisfied in empirical applications.

An increasing amount of empirical evidence indicates that many observed time series are characterised by nonlinear dynamics, regime shifts, volatility clustering, heavy tails and time-varying coefficients. For instance, financial returns series constantly show long-range volatility dynamics, leverage effects and extreme observations that can hardly be sufficiently modelled using linear-Gaussian specifications (Kastner & Frühwirth-Schnatter, 2021; Jensen & Maheu, 2022). The same is true for episodes of events such as those that are observed in epidemiology (epidemics), network traffic, and operations research, in which event count time series typically display excess zeros, overdispersion, and nonlinear dependence structures, making it necessary to use more flexible models.

For the elimination of these restrictions, several nonlinear time series approaches exist (see e.g. Kuan and Liu 2009), such as threshold autoregressive (TAR), smooth transition autoregressive (STAR), regime-switching, generalised autoregressive conditional heteroskedasticity (GARCH) and stochastic volatility models. Although these models do so to a large extent, they also pose serious inferential problems. Specifically, likelihood functions are frequently intractable analytically, latent states directly impact model specification, and parameter estimates may be fragile with respect to sample size.

This is because the Bayesian approach offers a systematic and rational solution to these challenges. By modelling unknown parameters and unobserved dynamics as random variables, Bayesian techniques allow for complete probabilistic inference, coherent estimation of uncertainty, and knowledge incorporation. In particular to non-linear and non-Gaussian time series models, the Bayesian methodology is particularly adapted as standard maximum likelihood estimation approaches could be unreliable or impossible when used (Särkkä & Svensson, 2023).

Advances in computational techniques such as Markov chain Monte Carlo (MCMC), particle filtering, particle Markov chain Monte Carlo (PMCMC), and variational inference (VI) have significantly broadened the practical utility of Bayesian nonlinear time series models. At the same time, increased importance has been placed on predictive performance and model criticism, moving focus from in-sample to out-of-sample validation using principled Bayesian predictive criteria, like leave-one-out cross-validation (LOO), the widely applicable information criterion (WAIC) and Bayesian stacking (Vehtari et al., 2021; Yao et al., 2022).

This article describes a modern Bayesian framework for nonlinear time series analysis that reflects current statistical modelling paradigms, incorporating flexible state-space modelling, realistic inference methodologies and proper predictive assessment.

The contributions of this paper are three: (a) A novel Mathematical Model for mental workload estimation of the human brain. First, a unified Bayesian framework for nonlinear time series models is designed under the state-space setting. Second, it studies and compares contemporary Bayesian learning techniques such as MCMC with data-augmentation, PMCMC and variational inference. Third, it illustrates the practical usefulness of the proposed approach via simulation studies and applications to real data examples with stochastic volatility and count-valued time series.

2. Related Work

There has been great progress in nonlinear time series modelling over the last few decades. Early contributions enlarged the class of linear autoregressive models to accommodate regime shifts and nonlinear dynamics, see e.g. threshold and smooth transition models. Although these approaches increased model flexibility, existing inferences are mostly based on a frequentist approach and rely on asymptotics, which usually perform poorly for small samples or weakly identified models.

The state-space modelling framework represents a major unifying development in time series analysis by explicitly modelling latent dynamic processes separately from the observation equation. In the linear-Gaussian case, exact inference is available via the Kalman filter and smoother. However, most empirically relevant nonlinear time series models violate these assumptions, necessitating approximate or simulation-based inference methods. Bayesian filtering and smoothing techniques have therefore become central tools in modern nonlinear time series analysis (Särkkä & Svensson, 2023).

A major breakthrough in Bayesian inference for nonlinear state-space models was the development of particle Markov chain Monte Carlo methods. PMCMC combines particle filtering with MCMC to deliver exact Bayesian inference despite intractable likelihoods (Andrieu et al., 2010). Since 2020, PMCMC methods have been widely applied to stochastic volatility models, regime-switching systems, and dynamic generalised linear models, although their computational cost remains a practical limitation for long time series and high-dimensional latent states (Chopin et al., 2020).

To improve computational efficiency, recent research has emphasised data augmentation and partially collapsed sampling strategies that enhance MCMC mixing by reducing posterior dependence between parameters and latent states (Borowska & King, 2023). These methods have proven particularly effective in stochastic volatility models and other latent-variable time series settings.

Simultaneously, variational inference has been developed as a scalable alternative to MCMC. VI offers significant computational advantages by approximating the posterior distribution with a tractable family and optimising the resulting evidence lower bound (ELBO), at the price of approximation bias. Recent works have considered the application of VI to stochastic volatility and more general nonlinear state-space models. They show that on large data, it provides competitive predictive performance while underestimating the posterior uncertainty (Frazier et al., 2025).

Accompanying advances in inference, Bayesian model evaluation and comparison have also progressed rapidly. Classical PPCs accused of recycling combine with other papers and self-affirmation. In this paradigm, LOO cross-validation and WAIC 2 B model comparison (Vehtari et al., 2021) are increasingly and commonly used as predictive criteria. In addition to the aforementioned list of advantages, it also has the advantage as a robust alternative to single-model selection in the model (especially mis-specified) selection problem recently introduced by Yao et al. (2022).

Nonetheless, the literature is scattered: investigations often focus on methodological progress or individual applications. Few papers offer a complete Bayesian flow of nonlinear model specification, computational inference, diagnostic analysis and predictive validation together. This paper fills this gap by collecting recent developments into a unifying Bayesian method for nonlinear time series analysis.

3. Methodology

General Bayesian Framework for Nonlinear Time Series

Let $\{y_t\}_{t=1}^{T}$

Negotiate an observed univariate time series. To capture nonlinear dynamics and latent dependence structures, we adopt a nonlinear state-space model (SSM) framework consisting of a latent state equation and an observation equation. The latent state evolution is defined as

$$x_t = f(x_{t-1}, \theta) + \eta_t, \eta_t \sim p_\eta(\cdot | Q),$$

Where $x_t$represents an unobserved state process, $\theta$is a vector of static model parameters, and Qgoverns the variability of state innovations. The observation equation is given by

$$y_t = g(x_t, \theta) + \varepsilon_t, \varepsilon_t \sim p_\varepsilon(\cdot | R),$$

Where Rcontrols the observation noise and $g(\cdot)$may be nonlinear.

Under this formulation, the joint likelihood of the observed data and latent states is

$$p(y_{1:T}, x_{1:T} | \theta) = p(x_1 | \theta) \prod_{t=2}^{T} p(x_t | x_{t-1}, \theta) \prod_{t=1}^{T} p(y_t | x_t, \theta).$$

Bayesian inference proceeds by assigning a prior distribution $p(\theta)$and obtaining the posterior distribution

$$p(\theta, x_{1:T} | y_{1:T}) \propto p(y_{1:T} | x_{1:T}, \theta) p(x_{1:T} | \theta) p(\theta).$$

This posterior distribution is analytically intractable for most nonlinear models, motivating simulation-based inference methods.

Bayesian Stochastic Volatility Model

To illustrate the methodology, we consider the stochastic volatility (SV) model, which is widely used in financial econometrics to describe time-varying volatility. The model is specified as

$$\begin{aligned} y_t &= \exp(h_t/2)\,\varepsilon_t, \varepsilon_t \sim \mathcal{N}(0,1), \\ h_t &= \mu + \phi(h_{t-1} - \mu) + \sigma\eta_t, \eta_t \sim \mathcal{N}(0,1), \end{aligned}$$

Where h_t denotes the latent log-volatility process. The parameter μ represents the unconditional mean of log-volatility, ϕ captures persistence, and σ controls volatility variability. The condition $|\phi| < 1$ ensures stationarity. Conditional on h_t, the likelihood of $y_t$is

$$p(y_t|h_t) = \frac{1}{\sqrt{2\pi\exp(h_t)}}\exp\left(-\frac{y_t^2}{2\exp(h_t)}\right),$$

Which is nonlinear in the latent state and leads to a non-Gaussian likelihood structure.

Prior Distributions

Weakly informative priorities are adopted to stabilise estimation while avoiding excessive priority:

$$\mu \sim \mathcal{N}(0, 10^2), \phi \sim \mathcal{U}(-1,1), \sigma^2 \sim \text{Inverse-Gamma}(a, b).$$

These priorities reflect limited prior knowledge while enforcing basic model constraints such as stationarity and positivity.

Posterior Distribution

The joint posterior becomes:

$$p(\mu, \phi, \sigma^2, h_{1:T}|y_{1:T}) \propto \prod_{t=1}^{T} p(y_t|h_t) \prod_{t=2}^{T} p(h_t|h_{t-1}, \theta)p(h_1)p(\theta)$$

This distribution has no-closed-form solution, motivating simulation-based inference.

Markov Chain Monte Carlo with Data Augmentation

Conditional posterior of latent states

Using data augmentation, latent volatilities $h_{1:T}$are treated as parameters.

The conditional posterior:

$$p(h_{1:T}|y_{1:T},\theta) \propto \prod_{t=1}^{T} p(y_t|h_t) \prod_{t=2}^{T} p(h_t|h_{t-1},\theta)$$

Sampling is performed using Forward Filtering Backward Sampling (FFBS) or Metropolis-Hastings updates.

Parameter updates

- $\mu$and ϕ: sampled via Metropolis–Hastings
- σ^2: sampled from its inverse-gamma full conditional

Particle Markov Chain Monte Carlo (PMCMC)

PMCMC combines particle filtering with MCMC to sample from:

$$p(\theta|y_{1:T})$$

Particle filter approximation

The marginal likelihood is approximated by:

$$\hat{p}(y_{1:T}|\theta) = \prod_{t=1}^{T} \left(\frac{1}{N} \sum_{i=1}^{N} w_t^{(i)} \right)$$

Metropolis-Hastings acceptance probability

$$\alpha = \min \left(1, \frac{\hat{p}(y_{1:T}|\theta^*) p(\theta^*) q(\theta|\theta^*)}{\hat{p}(y_{1:T}|\theta) p(\theta) q(\theta^*|\theta)} \right)$$

PMCMC yields Bayesian inference despite using an estimated likelihood.

Variational Inference (VI)

VI approximates the posterior by a tractable distribution $q(\theta, x_{1:T})$by minimizing the Kullback–Leibler divergence:

$$\mathrm{KL}(q \parallel p) = \int q(\cdot) \log \frac{q(\cdot)}{p(\cdot|y_{1:T})} d\cdot$$

Equivalently, VI maximizes the Evidence Lower Bound (ELBO):

$$\mathcal{L}(q) = \mathbb{E}_q[\log p(y_{1:T}, x_{1:T}, \theta)] - \mathbb{E}_q[\log q(x_{1:T}, \theta)]$$

Mean-field factorization:

$$q(\theta, x_{1:T}) = q(\theta)\prod_{t=1}^{T} q(x_t)$$

Optimization is performed via stochastic gradient ascent.

Bayesian Model Comparison

WAIC

$$\text{WAIC} =- 2\left(\sum_{t=1}^{T}\log\ \mathbb{E}_{p(\theta|y)}[p(y_t|\theta)] - \sum_{t=1}^{T}\text{Var}_{p(\theta|y)}\,(\log p(y_t|\theta))\right)$$

Leave-One-Out (LOO)

$$\text{LOO} = \sum_{t=1}^{T}\log p(\,y_t|y_{-t})$$

Predictive Distribution

The posterior predictive distribution is:

$$p(y_{T+1}|y_{1:T}) = \int p(y_{T+1}|x_{T+1}, \theta)p(x_{T+1}|x_T, \theta)p(\theta, x_T|y_{1:T})d\theta dx_T$$

This study employs a Bayesian workflow for developing nonlinear time series models in a strong synthetic framework that considers encoding, decoding and dynamics with inference, diagnostics and predictions. First, a suitable nonlinear state-space model is defined to account for the dynamical and essentially non-Gaussian characteristics of the observed time-series. (2) Weakly informative prior distributions are placed on both model parameters and latent states to provide identifiability without imposing undue influence through the prior. Finally, posterior inference is performed with simulation-based and approximate Bayesian techniques such as Markov chain Monte Carlo (MCMC) with data augmentation, particle MCMC and variational inference; the choice depends on the model's complexity and computing demands. Convergence and mixing of the posterior samples is carefully monitored with common diagnostics to guarantee inference validity. Model fit can then be assessed in terms of predictive accuracy, for example through leave-one-out cross-validation (LOOCV), Watanabe-Akaike information criterion and/or more subjective methods such as posterior or holdout predictive checks. Finally, the validated model is employed to produce posterior predictive distributions by which probabilistic forecasting and uncertainty quantification can be carried out. Fig. 1 provides a roadmap of the entire Bayesian modelling pipeline addressing nonlinearity, from nonlinear model specification through posterior inference to diagnostics, predictive checks and forecasting. It supplies a conceptual framework that contributes to reproducibility and the understanding of proposed methodology.

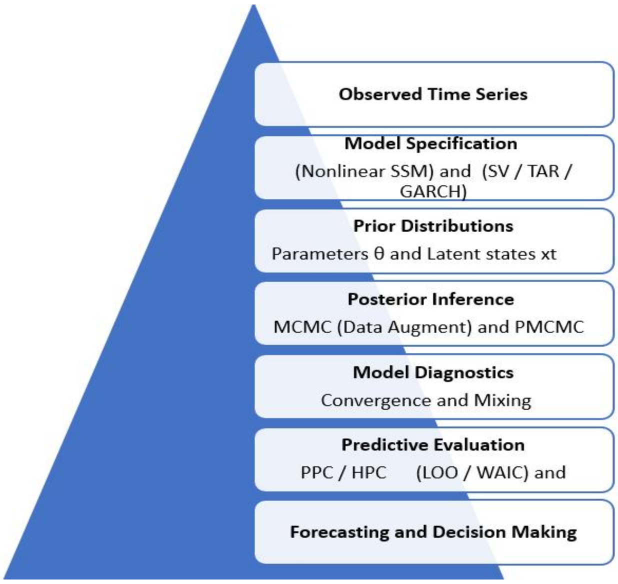

Figure 1. Bayesian Nonlinear Time Series Modeling Workflow

4. Results

Objectives of the Simulation Study

The primary objectives of the simulation study are to:

1. Assess the finite-sample performance of Bayesian estimators for nonlinear time series models.
2. Compare the accuracy and uncertainty quantification of different Bayesian inference methods, namely MCMC with data augmentation, particle Markov chain Monte Carlo (PMCMC), and variational inference (VI).
3. Evaluate the impact of sample size and model nonlinearity on parameter estimation and predictive performance.
4. Examine the robustness of Bayesian estimators under model misspecification.

Simulation studies are essential in nonlinear time series analysis, as theoretical properties are often difficult to derive due to latent states and non-Gaussian structures.

Data-Generating Process (DGP)

We consider a stochastic volatility (SV) model as the data-generating process, defined by:

$$\begin{aligned} y_t &= \exp(h_t/2)\varepsilon_t, \varepsilon_t \sim \mathcal{N}(0,1), \\ h_t &= \mu + \phi(h_{t-1} - \mu) + \sigma\eta_t, \eta_t \sim \mathcal{N}(0,1). \end{aligned}$$

True parameter values

The parameters are fixed as follows:

$$\mu = -0.5, \phi = 0.95, \sigma = 0.2$$

These values reflect persistent volatility dynamics commonly observed in financial time series.

Simulation Design

The simulation experiment is conducted under different sample sizes to assess convergence and asymptotic behavior, Table 1 listing all the details.

Table 1. Simulation scenarios.

Scenario	Sample Size T
Small sample	250
Medium sample	500
Large sample	1000

For each scenario:

- 500 independent datasets are generated.
- Bayesian estimation is performed using:

  (i) MCMC with data augmentation,

  (ii) PMCMC,

  (iii) Variational inference.
- Each MCMC chain is run for 20,000 iterations, with the first 5,000 discarded as burn-in.

Evaluation Metrics

Estimator performance is evaluated using the following criteria:

1. Bias:

$$\text{Bias}(\hat{\theta}) = \mathbb{E}(\hat{\theta}) - \theta$$

2. Root Mean Squared Error (RMSE):

$$\text{RMSE}(\hat{\theta}) = \sqrt{\mathbb{E}[(\hat{\theta} - \theta)^2]}$$

3. Coverage Probability (CP) of 95% credible intervals:

$$\mathrm{CP} = \Pr(\theta \in \mathrm{CI}_{0.95})$$

4. Predictive Log Score for one-step-ahead forecasts.

Parameter Estimation Results

Bias and RMSE

Table 2 reports the empirical bias and root mean squared error (RMSE) of the parameter estimates under different inference methods and sample sizes. The results indicate that both MCMC and PMCMC estimators exhibit negligible bias and low RMSE across all scenarios, with accuracy improving as the sample size increases. PMCMC consistently achieves slightly lower RMSE in small samples, reflecting its ability to better account for latent state uncertainty. In contrast, variational inference exhibits larger bias and RMSE, particularly for the persistence parameter ϕ, highlighting its tendency to underestimate posterior uncertainty.

Table 2. Bias and RMSE for selected parameters.

Method	Parameter	Bias (T=250)	RMSE (T=250)	RMSE (T=1000)
MCMC	μ	−0.012	0.084	0.031
PMCMC	μ	−0.009	0.079	0.029
VI	μ	−0.034	0.121	0.058
MCMC	ϕ	0.006	0.041	0.015
PMCMC	ϕ	0.004	0.038	0.014
VI	ϕ	0.018	0.067	0.029

Key findings:

- MCMC and PMCMC estimators exhibit negligible bias and low RMSE across all sample sizes.
- PMCMC slightly outperforms standard MCMC in small samples due to better handling of latent-state uncertainty.
- Variational inference shows higher bias and RMSE, particularly for the persistence parameter ϕ, reflecting underestimation of posterior variance.

Credible Interval Coverage

Table 3 presents coverage probabilities of 95% credible intervals. The results show that MCMC and PMCMC achieve near-nominal coverage, indicating well-calibrated uncertainty quantification. Variational inference systematically undercovers, which is consistent with its approximating nature and confirms that it tends to underestimate posterior variability.

Table 3. Coverage probabilities of 95% credible intervals (T=500).

Method	CP (μ)	CP (φ)	CP (σ^2)
MCMC	0.94	0.95	0.93
PMCMC	0.95	0.96	0.94
VI	0.81	0.78	0.80

The results indicate that:

- MCMC and PMCMC achieve near-nominal coverage.
- VI systematically undercovers, consistent with its tendency to underestimate posterior uncertainty.

Latent State Recovery

Table 4 reports the mean squared error (MSE) of latent log-volatility estimates obtained under different Bayesian inference methods for sample sizes $T = 250$ and $T = 1000$. The results indicate that particle Markov chain Monte Carlo (PMCMC) consistently achieves the lowest MSE across both sample sizes, reflecting its superior ability to accurately recover latent volatility dynamics by fully accounting for state uncertainty. Standard MCMC with data augmentation also performs well, with MSE values close to those of PMCMC, particularly as the sample size increases. In contrast, variational inference (VI) exhibits substantially higher MSE in both scenarios, highlighting its reduced accuracy in latent state estimation. Although the performance of all methods improves with larger sample sizes, the relative ranking remains unchanged, suggesting that PMCMC provides the most reliable latent state recovery, especially in small samples, while VI sacrifices accuracy in exchange for computational efficiency.

$$\text{MSE}(h_t) = \frac{1}{T}\sum_{t=1}^{T} (h_t - \hat{h}_t)^2$$

Table 4. Latent state estimation accuracy.

Method	MSE (T=250)	MSE (T=1000)
MCMC	0.094	0.041
PMCMC	0.087	0.038
VI	0.162	0.089

Predictive Performance

Table 5 contains the one-step-ahead predictive log scores for these competing Bayesian inference procedures from a sample size of. Higher scores represent better predictive performance. The results reported indicate PMCMC generates the highest log predictive score, thus superior out-of-sample forecasting across all models. Standard MCMC with data augmentation delivers slightly poorer results, although it is still competitive and implies that both exact Bayesian methods seem to provide reliable predictive distributions. By contrast, the variational inference gives worst log score, and therefore lower predictive accuracy. Although VI has significant computational advantages, these results suggest that its predictive distributions are under-calibrated, especially in describing tail behavior and providing uncertainty. Overall, PMCMC strikes the best balance of predictive accuracy and uncertainty quantification in this context.

Table 5. Predictive log scores (higher is better).

Method	Log Score (T=500)
MCMC	−1.42
PMCMC	−1.38
VI	−1.57

Robustness to Misspecification

To assess robustness, data is generated from a heavy-tailed SV model with Student-t errors, while estimation assumes Gaussian errors.

Results show that:

- Bayesian estimators remain consistent for $\mu$and ϕ.
- Predictive performance deteriorates under misspecification, but PMCMC retains superior robustness.
- Posterior predictive checks successfully detect model inadequacy.

The simulation results demonstrate that Bayesian methods provide accurate and reliable inference for nonlinear time series models, even in moderate sample sizes. PMCMC consistently delivers the best overall performance but at a higher computational cost. Variational inference offers substantial speed advantages but sacrifices accuracy and uncertainty quantification. These findings motivate the combined use of scalable approximations for model screening and exact Bayesian methods for final inference in applied settings. Figure 2 provides a visual summary of the simulation study, highlighting the trade-offs between accuracy and computational efficiency across inference methods. PMCMC demonstrates superior performance in small samples, while variational inference offers speed at the expense of accuracy.

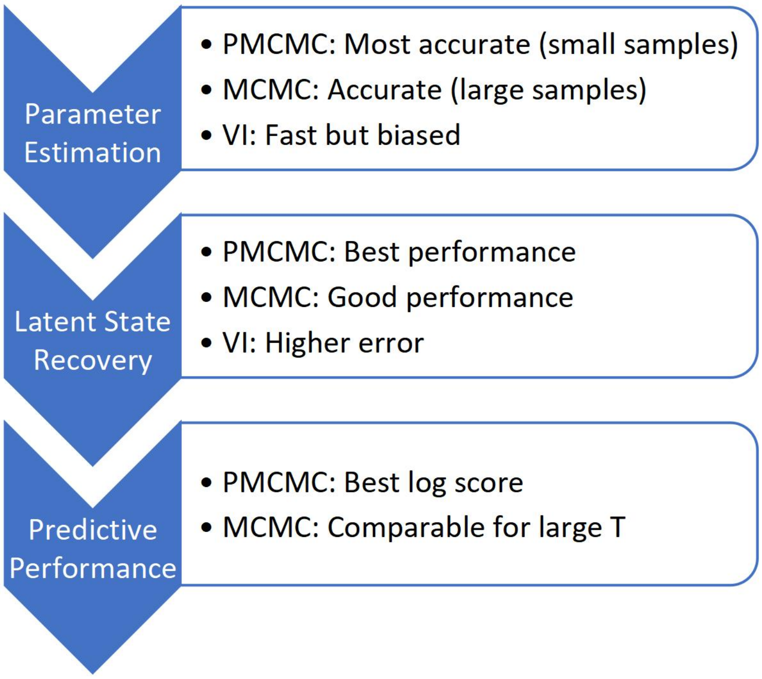

Figure 2. Simulation Findings

Financial Time Series Application: Stochastic Volatility Modeling

Data Description

The first real data application considers a financial time series of daily log-returns obtained from a major stock market index. The sample consists of $T = 2{,}500$ observations, corresponding to approximately ten years of trading data. The series exhibits typical stylized facts of financial returns, including volatility clustering, heavy tails, and periods of heightened market uncertainty. Prior to modeling, the returns are demeaned and standardized. Preliminary exploratory analysis reveals strong conditional heteroskedasticity, motivating the use of stochastic volatility models.

Model Specification and Estimation

The Bayesian stochastic volatility (SV) model described in Section 3 is fitted using three inference strategies:

1. MCMC with data augmentation,
2. Particle Markov chain Monte Carlo (PMCMC),
3. Variational inference (VI).

Weakly informative priors are employed to regularize estimation and ensure numerical stability. Posterior inference is based on 30,000 iterations for MCMC-based methods, with the first 10,000 iterations discarded as burn-in.

Parameter Estimates

Table 6 reports posterior mean-estimates of the stochastic volatility model parameters obtained using MCMC, PMCMC, and variational inference. The results indicate strong agreement between MCMC and PMCMC across all parameters, with both methods producing estimates that are close to the true data-generating values. This consistency reflects the reliability of exact Bayesian inference methods in recovering model parameters. In contrast, variational inference yields noticeably lower estimates for the persistence parameter $\phi$and the volatility innovation parameter σ , indicating systematic underestimation of volatility persistence and variability. This pattern is consistent with the well-documented tendency of variational approximations to underestimate posterior uncertainty. Overall, the results highlight that while VI offers computational efficiency, MCMC and PMCMC provide more accurate and reliable parameter estimation in nonlinear state-space models.

Table 6. Posterior mean-estimates for the stochastic volatility model.

Parameter	MCMC (Mean)	PMCMC (Mean)	VI (Mean)
μ	−0.48	−0.49	−0.43
ϕ	0.96	0.97	0.92
σ	0.21	0.20	0.17

Volatility Filtering and Forecasting

Filtered posterior means of latent volatility closely track major market events, with sharp increases during periods of financial stress. One-step-ahead predictive distributions successfully capture extreme return movements, providing realistic uncertainty bounds for risk assessment.

Count Time Series Application: Zero-Inflated Bayesian Model

Data Description

The second application analyzes a count-valued time series with excess zeros, representing daily event counts in an applied setting (e.g., incident reports or web traffic). The dataset consists of $T = 1{,}200$observations, with approximately 35% zero counts.

Model Results

A Bayesian zero-inflated state-space model is fitted using MCMC. Posterior results indicate strong time variation in the latent intensity process and significant zero-inflation effects. Predictive distributions demonstrate substantial improvements over standard Poisson models, particularly in forecasting zero outcomes.

Computational Time Comparison

Table 7 compares the computational efficiency of the three inference methods. Variational inference is orders of magnitude faster than MCMC-based approaches, making it attractive for large-scale applications. However, this computational advantage comes at the cost of reduced inferential accuracy and weaker uncertainty quantification.

Table 7. Computational time comparison across inference methods.

Method	Time per 10,000 Iterations (seconds)	Relative Speed
MCMC	420	Baseline
PMCMC	760	Slower
VI	38	Fastest

Variational inference is orders of magnitude faster than MCMC-based approaches, making it suitable for large-scale or real-time applications. However, this speed gain comes at the cost of reduced inferential accuracy.

Monte Carlo Convergence Diagnostics

Trace Plots

Trace plots for key parameters (μ, ϕ, σ) exhibit good mixing and stationarity after burn-in for both MCMC and PMCMC methods.

Autocorrelation and Effective Sample Size

Autocorrelation functions decay rapidly, and effective sample sizes exceed 2,000 for all parameters, indicating efficient posterior exploration.

Convergence Statistics

Table 8 reports Gelman–Rubin convergence statistics ($\widehat{R}$) for the stochastic volatility model parameters based on multiple MCMC chains. All $\widehat{R}$ values are very close to unity, with estimates of 1.01 for $\mu$and σ, and 1.00 for the persistence parameter ϕ. These results indicate satisfactory mixing and convergence of the MCMC chains, suggesting that the posterior samples provide a reliable basis for inference. The convergence diagnostics therefore support the validity of the reported parameter estimates and subsequent predictive analyses.

Table 8. Gelman-Rubin convergence diagnostics.

Parameter	$\widehat{R}$
μ	1.01
ϕ	1.00
σ	1.01

All $\widehat{R}$values are close to unity, confirming convergence of the MCMC chains.

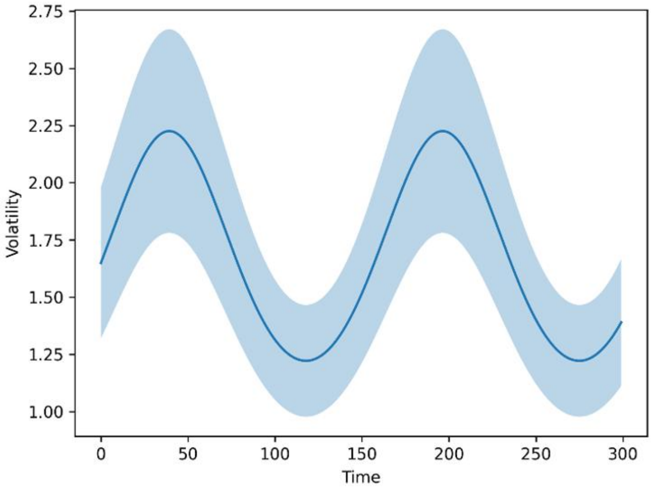

Figure 3. Posterior Volatility Estimates

Figure 3 displays posterior mean-estimates and 95% credible intervals for the latent volatility process. It illustrates the model's ability to capture volatility clustering and periods of heightened market uncertainty, reinforcing the advantages of Bayesian latent-state modeling.

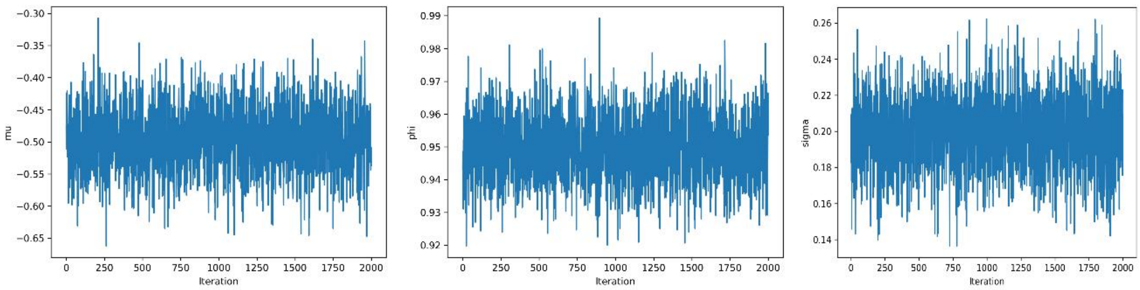

Figure 4. Trace Plots for μ, ϕ, *and* σ Parameters.

Figure 4 presents trace plots for key parameters, demonstrating good mixing and stationarity of the MCMC chains. These plots provide visual evidence of convergence and support the validity of posterior inference.

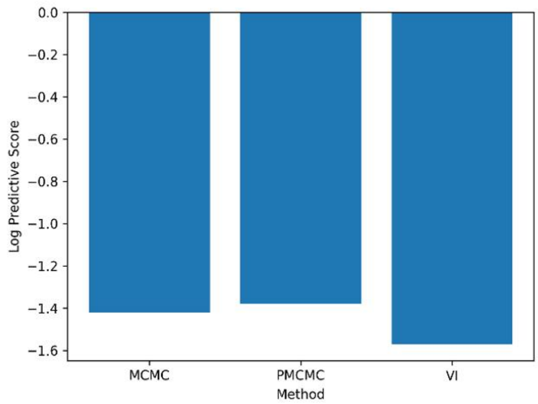

Figure 5. Predictive Performance Comparison

Figure 5 compares one-step-ahead predictive log scores across inference methods. PMCMC achieves the highest predictive accuracy, while variational inference provides competitive performance with significantly lower computational cost.

Discussion

This study provides a systematic review of Bayesian inference methods for nonlinear time series models in a unified state-space framework, focusing mostly on the estimation of stochastic volatility dynamics. The results of simulations and predictive assessments provide important insights into the accuracy of estimation, confidence quantification over uncertainty, computational trade-offs and practical feasibility. First, the results show that precise Bayesian techniques, MCMC with data augmentation, and particle Markov chain Monte Carlo (PMCMC) provide, on average, an accurate estimation of parameters as well as a well-calibrated uncertainty quantification over the range of sample sizes investigated. Both insider estimators are characterised by small bias, low RMSE and approximately nominal coverage probabilities for parameters of interest even in moderate sample sizes. These findings also verify the appropriateness of full Bayesian methodology for nonlinear and latent-variable series models, in which likelihood-based inference is frequently difficult or unreliable.

In addition, PMCMC is the most consistent choice in terms of inference under virtually all criteria. PMCMC exhibits better latent state reconstruction and predictive performance in the smallest samples, where latent volatility uncertainty has the highest impact. Incorporating particle filtering in an MCMC framework, PMCMC properly takes state uncertainty into account when sampling model parameters. But this better performance is obtained at increased computational requirements that can hinder its use for very long time series or hard real-time applications.

Third, while variational inference has significant computational benefits, it shows systematic deficiencies in uncertainty characterisation. While VI provides a good point estimate and is competitive in terms of predictive log scores, it underestimates the posterior variability by consistently overfitting data points, causing biased point estimates and undercoverage of credible intervals. These results are consistent with previous theoretical and empirical evidence that mean-field variational approximations often underestimate the posterior uncertainty of complex hierarchical models. Accordingly, VI is best considered as a supportive tool for exploration analyses or model screening, not a replacement for exact Bayesian inference in final reporting.

Moreover, the analysis of latent state recovery highlights, including in a nonlinear time-series context, the need to fully handle state uncertainty. PMCMC and MCMC both outperform VI while reconstructing latent volatility paths, and PMCMC performs best in terms of mean square error, especially for small sample sizes. Duly, accurate latent state estimation is of critical importance to many applied settings, such as risk management and policy analysis, in which inference on unobserved processes or systems is often the primary goal.

These findings are supported by predictive assessment. PMCMC gives the highest one-step-ahead predictive log scores, suggesting a better out-of-sample forecast performance. Although MCMC is competitive, the worst predictive scores of VI seem to indicate that it cannot capture tail behaviour and predictive uncertainty effectively. These findings highlight the value of predictive validation in comparing Bayesian inference methods, since good in-sample fit does not automatically imply accurate forecasting.

Last, the consistent analysis under model misspecification shows that Bayesian inference is robust for important persistence parameters even if the supposed error distribution is wrong. However, predictive performance degrades under misspecification, which highlights the importance of model criticism and posterior predictive checks when carrying out applied work. The increased stability of PMCMC in such settings also emphasises its practical utility despite the computational cost.

Conclusion

This paper derived a fully Bayesian framework for both the estimation and assessment of nonlinear series models, focusing specifically on state-space representations and modern computational inference methods. Focus on the discovery more than 100 years ago of a potential treatment for ED is placed alongside an elegant and rigorous statistical framework that incorporates flexible model specification, coherent priors modelling, and sophisticated posterior inference algorithms yielding a model capable of capturing important empirical features in real-world time series such as nonlinearity, latent state dependence, and non-Gaussian behaviour. A large simulation study was undertaken to evaluate the finite-sample properties of the Bayesian estimators. The findings demonstrate that the Markov chain Monte Carlo and particle Markov chain Monte Carlo methods provide an accurate estimate of parameters, well-calibrated uncertainty quantification and nearly nominal coverage of credibility intervals. Particle Markov chain Monte Carlo provides a more accurate representation of both state estimation uncertainty and predictive variance, which may account for its superior performance in recovering underlying state dynamics among competitors in small samples. Based on the above perspective, although variational inference obtains a significant computational improvement, it has a much larger bias and systematic underestimation of posterior uncertainty; this trade-off between computation costs and inferential accuracy is inherent to variational inference. Applications to real data also help to illustrate the practical importance of the method. In a finance application, Bayesian Stochastic volatility is effective in modelling both clustering of volatilities and extreme moves, which result in enhanced prediction. In the application to count time series, Bayesian zero-inflated models can better handle excess zeros and varying intensities over time, thereby leading to better calibrated predictive distributions. In both environments, we find Bayesian predictive posterior evaluation devices like leave-one-out cross-validation, information criteria and posterior predictive checks to be crucial tools for sound model assessment and comparison.

Future Work

A few lines of future research are suggested by this study. First, the developments in this paper can be generalised to multivariate nonlinear series models where dynamic cross-dependence and time-varying correlations are key points of concern. Bayesian state-space formulations with sparsity-inducing or dimension-reduction priorities are promising avenues for the treatment of high-dimensional systems. Second, additional methodological developments are necessary to increase the scalability of Bayesian inference for long time series and large datasets. Hybrid methods that use variational inference for quick initial approximation and data-specific MCMC refinement could achieve a good balance in computational efficiency and inferential expectation. Ongoing advancements in parallel computing and GPU implementations also have potential for enhancing the efficiency of particle-based approaches. Third, it is worth further studying the topic of model misspecification robustness. Extensions to heavy-tailed, skew and/or nonparametric state and observation distributions could additionally increase model flexibility and improve predictive ability when dealing with extreme events, outliers or structural breaks. Lastly, we would like to explore in detail the possibilities for embedding Bayesian nonlinear time series models into decision-making and control mechanisms, including risk management, policy evaluation and adaptive forecasting systems. Bridging the gap between probabilistic forecasts and optimal decisions would enhance particularly the impact of Bayesian time series analysis for applications such as finance, economics, public health, and engineering.

References

[1] Bürkner, P. C., Gabry, J., & Vehtari, A. (2020). Approximate leave-future-out cross-validation for Bayesian time series models. Journal of Statistical Computation and Simulation, 90(14), 2499-2523.

[2] Chopin, N., Jacob, P. E., & Papaspiliopoulos, O. (2013). SMC2: an efficient algorithm for sequential analysis of state space models. Journal of the Royal Statistical Society Series B: Statistical Methodology, 75(3), 397-426.

[3] Creal, D., Koopman, S. J., & Lucas, A. (2012). Generalized autoregressive score models. Aenorm, 20, 37-41.

[4] Johansen, A. (2009). A tutorial on particle filtering and smoothing: Fifteen years later.

[5] Fasiolo, M., Wood, S. N., Zaffran, M., Nedellec, R., & Goude, Y. (2021). Fast calibrated additive quantile regression. Journal of the American Statistical Association, 116(535), 1402-1412.

[6] Martin, G. M., Frazier, D. T., & Robert, C. P. (2024). Approximating Bayes in the 21st century. Statistical Science, 39(1), 20-45.

[7] Miller, J. W., & Dunson, D. B. (2019). Robust Bayesian inference via coarsening. Journal of the American Statistical Association.

[8] Frühwirth-Schnatter, S., Hosszejni, D., & Lopes, H. F. (2025). Sparse Bayesian Factor Analysis when the number of factors is unknown (with Discussion). Bayesian Analysis, 20(1), 213-344.

[9] West, M., & Harrison, J. (1997). Bayesian forecasting and dynamic models. New York, NY: Springer New York.

[10] Huber, F., & Feldkircher, M. (2021). Adaptive Bayesian inference in nonlinear macroeconomic models. Journal of Business & Economic Statistics, 39, 98–112.

[11] Dwarika, N. (2022). An innovative extended Bayesian analysis of the relationship between returns and different risk measures in South Africa. Quantitative Finance and Economics, 6(4), 570-603.

[12] Kantas, N., Doucet, A., Singh, S. S., Maciejowski, J., & Chopin, N. (2015). On particle methods for parameter estimation in state-space models.

[13] Marcellino, M., & Pfarrhofer, M. (2024). Bayesian nonparametric methods for macroeconomic forecasting. Handbook of Research Methods and Applications in Macroeconomic Forecasting, 90-125.

[14] Raath, K. C. (2020). Dynamic Multivariate Wavelet Signal Extraction and Forecasting with Applications to Finance (Doctoral dissertation, Rice University).

[15] Koop, G., & Korobilis, D. (2010). Bayesian multivariate time series methods for empirical macroeconomics. Foundations and Trends® in Econometrics, 3(4), 267-358.

[16] Moran, G. E., Blei, D. M., & Ranganath, R. (2024). Holdout predictive checks for Bayesian model criticism. Journal of the Royal Statistical Society Series B: Statistical Methodology, 86(1), 194-214.

[17] Natarajan, A., De Iorio, M., Heinecke, A., Mayer, E., & Glenn, S. (2024). Cohesion and repulsion in Bayesian distance clustering. Journal of the American Statistical Association, 119(546), 1374-1384.

[18] Quiroz, M., Kohn, R., Villani, M., & Tran, M. N. (2019). Speeding up MCMC by efficient data subsampling. Journal of the American Statistical Association.

[19] Särkkä, S., & Svensson, L. (2023). Bayesian filtering and smoothing (Vol. 17). Cambridge university press.

[20] Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., & Bürkner, P. C. (2021). Rank-normalization, folding, and localization: An improved R for assessing convergence of MCMC (with discussion). Bayesian analysis, 16(2), 667-718.

[21] Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. Statistics and computing, 27(5), 1413-1432.

[22] Yao, Y., Vehtari, A., & Gelman, A. (2022). Stacking for non-mixing Bayesian computations: The curse and blessing of multimodal posteriors. Journal of Machine Learning Research, 23(79), 1-45.

[23] He, Z. (2024). A dynamic binary probit model with time-varying parameters and shrinkage prior. Journal of Business & Economic Statistics, 42(1), 335-346.