forecasting

MDPI

Article

Modeling Positive Seasonal Time Series with Dynamic Precision: The Generalized BPSARMA Model

Kleber H. Santos and **Francisco Cribari-Neto** *

Departamento de Estatística, Universidade Federal de Pernambuco, Recife 50670-901, PE, Brazil; kleber.henrique@ufpe.br
* Correspondence: francisco.cribari@ufpe.br

Highlights

What are the main findings?

- We introduce the generalized BPSARMA model, a new framework for modeling positive seasonal time series that allows the precision parameter to evolve dynamically over time.
- In several empirical applications involving reservoir inflow data, the proposed model often provides competitive or superior out-of-sample forecasting accuracy relative to alternative beta prime-based models, SARIMA, and Holt–Winters exponential smoothing.

What are the implications of the main findings?

- Allowing the precision parameter to vary over time provides a flexible way to capture changing variability in positive time series, potentially improving forecast performance.
- The proposed approach offers a practical forecasting tool for environmental and hydrological applications where the data are strictly positive and display seasonal patterns.

Abstract

This paper proposes a generalized seasonal beta prime autoregressive moving average model with dynamic precision, denoted by BPSARMA, for modeling and forecasting positive-valued seasonal time series. The proposed framework extends the generalized BPARMA model by incorporating stochastic seasonal dynamics in the conditional mean through seasonal autoregressive and moving average components while allowing a flexible autoregressive structure for the conditional precision parameter, thereby accommodating time-varying uncertainty. The model also allows the inclusion of covariates and deterministic seasonal regressors. Parameter estimation is carried out by conditional maximum likelihood, and the main inferential and diagnostic tools are discussed. Monte Carlo simulations are conducted to examine the finite-sample behavior of the estimators and associated inference procedures. The practical usefulness of the proposed approach is illustrated through hydro-environmental time series applications, where its forecasting performance is evaluated using both in-sample and out-of-sample predictive measures. The empirical results indicate that the BPSARMA specification often provides competitive or superior forecasting accuracy relative to competing models, highlighting its usefulness for modeling and prediction in positive seasonal time series.

Keywords: beta prime distribution; dynamic precision; forecasting; Monte Carlo simulation; positive time series; seasonality

Check for updates

Academic Editor: Jun A. Zhang

Received: 17 April 2026
Revised: 22 June 2026
Accepted: 23 June 2026
Published: 24 June 2026

1. Introduction

Time series assuming values on the positive real line arise naturally in many environmental and hydrological applications. Examples include river discharge, reservoir inflows, rainfall accumulation, pollutant concentrations, and wind speeds. Such variables typically display strong right-skewness and heteroscedastic behavior, with variability increasing with the level of the series. In addition, many environmental processes are subject to pronounced seasonal fluctuations driven by climatic cycles. These features often render classical Gaussian time series models inadequate for capturing both the distributional characteristics and the dynamic structure of the data.

In response to these challenges, increasing attention has been devoted to the development of observation-driven time series models based on non-Gaussian distributions defined on restricted domains. In particular, models built upon flexible continuous distributions supported on the positive real line have proven useful for describing environmental and hydrological processes. Such models offer an important practical advantage: they naturally produce forecasts that respect the support of the data, ensuring that predicted values remain in $(0, \infty)$. Among the distributions suitable for this purpose, the beta prime distribution provides a versatile framework for modeling strictly positive data, as it accommodates a wide range of distributional shapes and allows the variance to vary with the mean in a natural way.

Dynamic models for time series taking values in $(0, 1)$ have been extensively studied in the literature. The βARMA (beta autoregressive moving average) model for such series was introduced by [1] and later extended to accommodate time-varying precision by [2]. A modification of the model to incorporate stochastic seasonality while maintaining fixed precision was proposed by [3]. These models have been successfully applied to the modeling and forecasting of hydro-environmental time series; see, for instance, [4,5].

Recent advances in forecasting have also focused on nonlinear and regime-switching autoregressive structures, as well as Bayesian forecasting approaches. For example, Ref. [6] proposed a double-banded-threshold mixture autoregressive model for capturing complex nonlinear dynamics, while [7] developed Bayesian forecasting procedures for logistic mixture double autoregressive models with explanatory variables. These contributions highlight the growing interest in flexible forecasting frameworks capable of accommodating nonlinear behavior and time-varying dynamics. In contrast, the forecasting framework proposed in this paper is specifically designed for positive-valued seasonal time series and is based on a conditional beta prime distribution with dynamic mean and precision structures. The present contribution therefore complements these recent developments by addressing a different class of forecasting problems involving continuous positive-valued seasonal data.

The beta prime distribution, which is defined on the positive real line, is closely related to the beta distribution, arising as the distribution of the odds transformation of a beta random variable. This connection makes it a natural candidate for modeling positive-valued stochastic processes. In this paper, we develop a dynamic modeling framework based on the beta prime distribution, with particular emphasis on hydro-environmental time series exhibiting pronounced seasonal behavior.

The generalized BPARMA (beta prime autoregressive moving average) class of models for positive-valued time series, based on the beta prime distribution, was introduced by [8]. In that formulation, the conditional mean follows an autoregressive moving-average structure, while the precision parameter evolves dynamically through a first-order autoregressive specification. An important feature of this formulation is that both parameters of the conditional distribution evolve over time through separate dynamic submodels. As a result, temporal changes in the conditional density are governed not only by the

https://doi.org/10.3390/forecast8040053

mean but also by the precision parameter, allowing the model to capture time-varying variability in a flexible way. In the generalized BPARMA framework, seasonal patterns may be represented through deterministic regressors such as harmonic components and seasonal indicator variables.

Although this approach has proven useful in practice, many environmental time series exhibit seasonal dependence structures that are more complex than those captured by deterministic regressors alone. In particular, environmental processes often display seasonal dynamics that arise from stochastic dependence across seasons, in addition to deterministic periodic components. Moreover, the evolution of the variability of the series may require richer specifications for the precision parameter than a first-order autoregressive structure.

Motivated by these considerations, this article proposes an extension of the generalized BPARMA class that explicitly incorporates seasonal dynamics and a more flexible specification for the precision component. In the proposed model, the conditional mean follows a seasonal autoregressive moving-average structure in addition to the usual nonseasonal dynamics. At the same time, the precision parameter is allowed to follow an autoregressive process of general order, thereby expanding the range of dynamic heteroscedastic patterns that can be captured by the model. The resulting class, which we refer to as the BPSARMA model, also allows the inclusion of nonstochastic regressors, such as harmonic components and indicator variables, providing additional flexibility for representing deterministic seasonal patterns and structural changes.

Seasonality in environmental time series is often represented using one of several alternative approaches, such as seasonal dynamic components, harmonic regressors, or seasonal indicator variables. These strategies are frequently viewed as competing modeling choices. In practice, however, environmental processes may exhibit seasonal patterns that are too complex to be adequately represented by a single mechanism. In this article, we show that these approaches can be used jointly within a unified modeling framework. By extending the generalized BPARMA model to incorporate seasonal autoregressive and moving-average dynamics, the proposed specification allows deterministic seasonal regressors and stochastic seasonal dynamics to coexist in the same model. This combination proves particularly effective for the hydro-environmental series analyzed in this study, which exhibit strong and multifaceted seasonal fluctuations.

Parameter estimation is carried out by conditional maximum likelihood. Closed-form expressions are obtained for the conditional score vector and the conditional Fisher information matrix, facilitating numerical optimization and statistical inference. Diagnostic tools suitable for assessing model adequacy are also discussed. In addition, a Monte Carlo simulation study is conducted to examine the finite-sample performance of the proposed estimators.

The empirical usefulness of the proposed modeling framework is illustrated through five hydro-environmental applications involving monthly inflows to reservoirs associated with hydroelectric power plants located in the five geographic regions of Brazil. Brazil is a country of continental dimensions and is subject to diverse climatic regimes. Consequently, the analyzed series display substantial variability in their temporal behavior while sharing a common feature of strong seasonal fluctuations. Some series also present additional complexities, such as prolonged drought periods or structural changes in their level and variability. In our analysis, seasonal behavior is represented through a combination of three complementary mechanisms: seasonal dynamic components, harmonic regressors, and indicator variables. This strategy allows the model to capture both stochastic and deterministic aspects of seasonal variation.

The proposed model yields high-quality fits and accurate out-of-sample forecasts for all series considered. Its predictive performance is compared with that of several alternative

https://doi.org/10.3390/forecast8040053

approaches, including the BPSARMA model with fixed precision, the generalized BPARMA model, Gaussian SARIMA (seasonal autoregressive integrated moving average) models, and the Holt–Winters exponential smoothing algorithm. Overall, the results indicate that the proposed specification provides a flexible and effective framework for modeling positive-valued environmental time series exhibiting strong seasonal behavior and time-varying variability.

It is important to distinguish the present setting from the extensive literature on integer-valued time series, including threshold and count autoregressive models such as those developed by [9–11]. Those models are designed for discrete-valued count processes and rely on probabilistic mechanisms tailored to integer-valued data, such as thinning operators or count distributions. By contrast, the hydro-environmental series analyzed in this paper consist of continuous positive measurements. This distinction motivates the use of a beta prime conditional distribution together with dynamic structures for both the conditional mean and conditional precision.

The remainder of the paper is organized as follows: Section 2 introduces the proposed model and discusses its main properties. Section 3 presents the conditional maximum likelihood estimation procedure and derives the score vector and Fisher information matrix. Model selection, diagnostic tools, and forecasting are discussed in Section 4. Section 5 reports the results of the Monte Carlo simulation study. Section 6 presents the empirical applications to hydro-environmental time series. Finally, Section 7 concludes the paper.

2. A Seasonal Dynamic Extension of the Generalized BPARMA Model

In this section, we introduce the proposed seasonal dynamic extension of the generalized beta prime autoregressive moving average model. We begin with a brief review of the beta prime distribution and the generalized BPARMA framework of [8]. Building on this foundation, we develop a more flexible specification that incorporates stochastic seasonal dynamics in the conditional mean and allows for higher-order autoregressive structures in the precision submodel. These extensions broaden the range of temporal dependence patterns that can be captured within the beta prime time series framework.

2.1. The Beta Prime Distribution

The beta prime distribution was introduced by [12,13]. Also known as the inverse beta or beta type II distribution, it has support on the positive real line. The beta prime distribution arises as the distribution of the odds ratio associated with the beta distribution: if X follows a beta distribution, then $X/(1-X)$ follows a beta prime distribution. Equivalently, it can be expressed as the distribution of the ratio of two independent random variables following standard gamma distributions with unit scale parameter.

Later, Ref. [14] proposed a reparameterization in terms of a mean parameter μ and a precision parameter ϕ, which facilitates interpretation and makes the distribution particularly convenient for regression and time series modeling. Let Y be a random variable following a beta prime distribution, denoted by $Y \sim BP(\mu, \phi)$. The cumulative distribution function and probability density function are given, respectively, by

$$F(y;\mu,\phi) = I_{y/(1+y)}(\mu(1+\phi), \phi+2)$$

and

$$f(y;\mu,\phi) = \frac{y^{\mu(\phi+1)-1}(1+y)^{-[\mu(\phi+1)+\phi+2]}}{\mathcal{B}(\mu(1+\phi), \phi+2)},$$

$y > 0$, where $\mu, \phi > 0$, $I_{y/(1+y)}(\mu(1+\phi), \phi+2) = \mathcal{B}_{y/(1+y)}(\mu(1+\phi), \phi+2)/\mathcal{B}(\mu(1+\phi), \phi+2)$ denotes the regularized beta function,

$$\mathcal{B}_{y/(1+y)}(\mu(1+\phi),\phi+2) = \int_0^{y/(1+y)} t^{\mu(1+\phi)-1}(1-t)^{\phi+1}dt$$

is the incomplete beta function, and $\mathcal{B}(\mu(1+\phi),\phi+2) = \Gamma(\mu(1+\phi))\Gamma(\phi+2)/\Gamma(\mu(1+\phi)+\phi+2)$ is the beta function, with $\Gamma(\cdot)$ denoting the gamma function. The mean and variance of Y are, respectively,

$$\mathbb{E}(Y) = \mu \quad \text{and} \quad \text{Var}(Y) = \frac{\mu(1+\mu)}{\phi}.$$

For all admissible values of μ and ϕ, the beta prime distribution is positively skewed. The skewness coefficient depends on both the mean and precision parameters and is given by

$$\psi_1 = \frac{2(1+\phi)(1+2\mu)}{\phi-1}\sqrt{\frac{\phi}{\mu(1+\mu)(1+\phi)^2}},$$

for $\phi > 1$. It is worth noting that in several traditional two-parameter distributions with support on the positive real line, the skewness coefficient depends only on the corresponding precision parameter, as is the case for the gamma distribution ($2\phi^{1/2}$) and the reparameterized Birnbaum–Saunders distribution ($4(3\phi+11)(2\phi+5)^{-3/2}$). As discussed by [14], the beta prime distribution can exhibit substantially higher levels of skewness than those typically observed for the gamma and inverse Gaussian distributions.

2.2. The Generalized BPARMA Model

Ref. [8] introduced a dynamic model for variables taking values on the positive real line based on the beta prime distribution, referred to as the generalized BPARMA model. In this framework, the authors extended the classical formulation commonly adopted in dynamic models by incorporating a parsimonious dynamic submodel for the precision parameter. Specifically, an autoregressive and moving average structure was considered for the conditional mean, while a first-order autoregressive structure was specified for the conditional precision of the process. As a result, both parameters indexing the distribution evolve over time, providing greater modeling flexibility and allowing the conditional density to change shape throughout the time series.

Let $\{Y_t\}_{t=1}^n$ be a time series of random variables such that, for $t \in \{1,\ldots,n\}$, the conditional distribution of Y_t given the past information set $\mathcal{F}_{t-1}$, with $\mathcal{F}_{t-1} := \sigma\{Y_{t-1}, Y_{t-2}, \ldots\}$ denoting the σ-field generated by the information available up to time $t-1$, is beta prime with mean μ_t and precision parameter ϕ_t, that is, $Y_t \mid \mathcal{F}_{t-1} \sim BP(\mu_t,\phi_t)$. The conditional density of $Y_t \mid \mathcal{F}_{t-1}$ is

$$f(y \mid \mathcal{F}_{t-1};\mu_t,\phi_t) = \frac{y^{\mu_t(\phi_t+1)-1}(1+y)^{-[\mu_t(\phi_t+1)+\phi_t+2]}}{\mathcal{B}(\mu_t(1+\phi_t),\phi_t+2)}, \quad y > 0, \tag{1}$$

with $\mu_t, \phi_t > 0$. The conditional mean and conditional variance of Y_t are given by $\mathbb{E}(Y_t \mid \mathcal{F}_{t-1}) = \mu_t$ and $\text{Var}(Y_t \mid \mathcal{F}_{t-1}) = \mu_t(1+\mu_t)/\phi_t$, respectively.

The generalized BPARMA(p,q) model consists of two dynamic structures and is defined as

$$g_1(\mu_t) = \alpha_1 + \boldsymbol{x}_t^\top\boldsymbol{\beta} + \sum_{i=1}^{p}\varphi_i\left[g_1(Y_{t-i}) - \boldsymbol{x}_{t-i}^\top\boldsymbol{\beta}\right] + \sum_{j=1}^{q}\theta_j r_{t-j}$$

and

$$g_2(\phi_t) = \alpha_2 + \delta z_{t-1},$$

where $\alpha_1 \in \mathbb{R}$ and $\alpha_2 \in \mathbb{R}$ are scalar parameters, $\boldsymbol{x}_t$ denotes an l-vector of nonstochastic regressors at time t, and $\boldsymbol{\beta} := (\beta_1, \ldots, \beta_l)^\top$ is the corresponding vector of coefficients. The parameters $\varphi_1, \ldots, \varphi_p$ correspond to autoregressive terms, while $\theta_1, \ldots, \theta_q$ represent moving average parameters. The term $r_t := g_1(Y_t) - g_1(\mu_t)$ denotes the random error on the predictor scale, and $z_{t-1} := Y_{t-1}/Y_{t-2}$ corresponds to the ratio between consecutive lagged observations. Moreover, $p, q \in \mathbb{N}_0 := \{0, 1, 2, \ldots\}$ denote the autoregressive and moving-average orders, respectively, with $p + q \geq 1$. The functions $g_1 : \mathbb{R}_+ \to \mathbb{R}$ and $g_2 : \mathbb{R}_+ \to \mathbb{R}$ are strictly monotonic and twice differentiable link functions, with inverses $g_1^{-1} : \mathbb{R} \to \mathbb{R}_+$ and $g_2^{-1} : \mathbb{R} \to \mathbb{R}_+$ also being twice differentiable.

When the series exhibits seasonal fluctuations, harmonic regressors are commonly used to represent periodic patterns. In this case, one may set $\boldsymbol{x}_t = (\sin(2\pi t/S), \cos(2\pi t/S))^\top$, where S denotes the seasonal period. These regressors provide a parsimonious way of modeling deterministic seasonal cycles through sinusoidal functions, allowing the model to capture smooth periodic variations in the mean structure of the series. However, this specification assumes that the amplitude and phase of the seasonal pattern remain constant over time. This approach was employed in [8] in the analysis of water flow series from hydroelectric reservoirs.

2.3. The Generalized BPSARMA Model

We now introduce a seasonal dynamic extension of the generalized BPARMA model that increases its flexibility for modeling positive-valued time series. The proposed formulation incorporates stochastic seasonal autoregressive and moving average components in the conditional mean structure while allowing the precision parameter to follow an autoregressive process of arbitrary order. These features enable the model to capture evolving seasonal patterns and the evolution of the conditional precision.

The proposed generalized beta prime seasonal autoregressive moving average model, denoted by $\text{BPSARMA}_\mu(p_1, q) \times (P, Q)_S + \text{AR}_\phi(p_2)$, is defined by

$$\Phi(B^S)\varphi(B)\left[g_1(Y_t) - \boldsymbol{x}_t^\top \boldsymbol{\beta}\right] = \alpha_1 + \Theta(B^S)\theta(B)r_t \tag{2}$$

and

$$g_2(\phi_t) = \alpha_2 + \sum_{k=1}^{p_2} \delta_k z_{t-k} =: \eta_{2t}, \tag{3}$$

where B denotes the backshift operator such that, for a nonnegative integer b, $B^b g_1(Y_t) = g_1(Y_{t-b})$. Furthermore, $\varphi(B) := 1 - \varphi_1 B - \varphi_2 B^2 - \cdots - \varphi_{p_1} B^{p_1}$ denotes the nonseasonal autoregressive polynomial, while $\theta(B) := 1 - \theta_1 B - \theta_2 B^2 - \cdots - \theta_q B^q$ corresponds to the nonseasonal moving average polynomial. The seasonal autoregressive polynomial is $\Phi(B^S) := 1 - \Phi_1 B^S - \Phi_2 B^{2S} - \cdots - \Phi_P B^{PS}$, whereas $\Theta(B^S) := 1 - \Theta_1 B^S - \Theta_2 B^{2S} - \cdots - \Theta_Q B^{QS}$ denotes the seasonal moving average polynomial. Here, $p_1, q, P, Q, p_2 \in \mathbb{N}_0$. The seasonal period of the time series is represented by S and $z_{t-k} := (Y_{t-k} - Y_{t-k-1})/Y_{t-k-1}$.

Although Equations (2) and (3) are defined for general link functions, throughout this paper, we employ the logarithmic link in both submodels, i.e., $g_1(u) = \log(u)$ and $g_2(u) = \log(u)$. This choice is natural for positive-valued time series, since it ensures that both the conditional mean μ_t and the precision parameter ϕ_t remain positive. Furthermore, the transformation allows the dynamic dependence structure to be modeled on the log scale, where autoregressive and moving-average effects can be interpreted in terms of relative changes rather than absolute changes.

The regressors z_{t-k} in the precision submodel are defined as relative increments of the response variable rather than as simple ratios of consecutive observations. This choice preserves the dimensionless nature of the precision covariates, which is important because the precision parameter ϕ_t must be free of measurement units. Indeed, since $\mathrm{Var}(Y_t \mid \mathcal{F}_{t-1}) = \mu_t(1+\mu_t)/\phi_t$, a dimensionless ϕ_t ensures that the conditional variance has the appropriate squared measurement units of the response.

In addition, the adopted specification offers practical and interpretative advantages. First, the relative increment $z_{t-k} = (Y_{t-k} - Y_{t-k-1})/Y_{t-k-1}$ can be interpreted as a local growth rate, with zero indicating stability, positive values indicating increases, and negative values indicating decreases. This yields a more natural interpretation of the effect of recent changes on the precision dynamics. Second, unlike simple ratios, relative increments are centered around zero, which improves the interpretation of the intercept and may reduce collinearity with it. Finally, because precision is intended to capture changes in variability, using relative increments provides a more direct link between recent fluctuations in the series and the temporal evolution of the precision parameter. Preliminary empirical and simulation experiments with alternative ratio-based formulations consistently indicated better overall performance for the specification adopted here.

Equation (2) of the BPSARMA$_\mu(p_1, q) \times (P, Q)_S + \mathrm{AR}_\phi(p_2)$ model can be written equivalently as

$$\begin{aligned} g_1(\mu_t) = {} & \alpha_1 + \boldsymbol{x}_t^\top \boldsymbol{\beta} + \sum_{i=1}^{p_1} \varphi_i \left[g_1(Y_{t-i}) - \boldsymbol{x}_{t-i}^\top \boldsymbol{\beta} \right] + \sum_{I=1}^{P} \Phi_I \left[g_1(Y_{t-IS}) - \boldsymbol{x}_{t-IS}^\top \boldsymbol{\beta} \right] \\ & - \sum_{i=1}^{p_1} \sum_{I=1}^{P} \varphi_i \Phi_I \left[g_1(Y_{t-(i+IS)}) - \boldsymbol{x}_{t-(i+IS)}^\top \boldsymbol{\beta} \right] - \sum_{j=1}^{q} \theta_j r_{t-j} - \sum_{J=1}^{Q} \Theta_J r_{t-JS} \\ & + \sum_{j=1}^{q} \sum_{J=1}^{Q} \theta_j \Theta_J r_{t-(j+JS)} =: \eta_{1t}. \end{aligned}$$

In what follows, we write $\boldsymbol{\varphi} := (\varphi_1, \ldots, \varphi_{p_1})^\top$, $\boldsymbol{\theta} := (\theta_1, \ldots, \theta_q)^\top$, $\boldsymbol{\Phi} := (\Phi_1, \ldots, \Phi_P)^\top$, $\boldsymbol{\Theta} := (\Theta_1, \ldots, \Theta_Q)^\top$, and $\boldsymbol{\delta} := (\delta_1, \ldots, \delta_{p_2})^\top$.

This representation makes explicit the dynamic specification adopted for the conditional mean on the transformed scale. In particular, it shows that $g_1(\mu_t)$ evolves through a combination of nonseasonal and seasonal autoregressive components together with nonseasonal and seasonal moving-average terms. Expressing the model in this form clarifies its relationship with the generalized BPARMA structure while making the different dynamic contributions to the conditional mean more transparent. Moreover, it highlights how the proposed specification extends the original BPARMA framework by incorporating stochastic seasonal dynamics in the conditional mean while simultaneously allowing for richer temporal dependence in the precision submodel.

The dynamic structure specified for the conditional mean is similar to that adopted in the βSARMA (beta seasonal autoregressive moving average) model and in the MKSARMAX (modified Kumaraswamy seasonal autoregressive moving average with exogenous regressors) model; see, respectively, [3,15]. It is important to emphasize, however, that those models were developed for random variables with support restricted to the unit interval. In contrast, the generalized BPSARMA model proposed in this paper is suitable for random variables defined on the positive real line.

An important feature of the proposed specification is that seasonal dynamics are modeled directly through seasonal autoregressive and moving average components. This approach allows seasonal patterns to evolve dynamically over time, rather than being restricted to fixed deterministic cycles. In many empirical applications, seasonality is introduced through harmonic regressors, which impose periodic patterns with constant amplitude and phase. While such representations are convenient, they assume that the seasonal structure remains stable throughout the sample.

In practice, however, seasonal patterns often evolve over time. This is particularly common in hydro-environmental time series, where seasonal fluctuations may change in magnitude, persistence, or timing due to climatic variability, hydrological regulation, or other environmental factors. Capturing these evolving seasonal dynamics is therefore important for an adequate statistical description of the data. The stochastic seasonal structure adopted in the proposed model accommodates such features by allowing the magnitude and persistence of seasonal effects to adjust according to the underlying dynamics of the process. As a result, the model provides a more flexible representation of seasonal behavior than approaches based solely on deterministic harmonic components.

It is also worth noting that the proposed formulation does not preclude the use of deterministic seasonal regressors. Because the model continues to allow the inclusion of covariates in the linear predictor through $\boldsymbol{x}_t^\top \boldsymbol{\beta}$, harmonic regressors can be incorporated whenever appropriate. Consequently, the framework accommodates hybrid specifications that combine deterministic seasonal effects with stochastic seasonal dynamics.

In addition to incorporating a dynamic structure capable of capturing seasonality in the conditional mean submodel, the proposed model further generalizes the BPARMA formulation in the precision submodel; see Equation (3). The model introduced by [8] adopts a first-order autoregressive dynamic specification for the conditional precision. By contrast, in the formulation developed here, the precision submodel is specified through an autoregressive structure of order p_2. This extension provides greater flexibility for modeling the temporal dynamics of the conditional precision, which reflects the level of uncertainty associated with the response variable.

Fluctuations in the precision parameter translate directly into changes in the conditional variance of the process. By allowing a higher-order autoregressive structure, the proposed specification allows the persistence and propagation of uncertainty levels to be captured more accurately over time. As a result, the model can accommodate richer patterns of temporal variation in conditional variability, yielding a more faithful representation of the evolving uncertainty underlying the observed series.

The model with fixed precision can be viewed as a particular case of the proposed formulation. It is obtained by adopting the identity link for g_2 and imposing $\delta_k = 0$, for $k \in \{1, \ldots, p_2\}$. In this case, the precision parameter remains constant over time. This nested specification highlights that the proposed framework encompasses both time-varying and constant-precision models as special cases.

For βARMA models with fixed precision, consistency and asymptotic normality of the conditional maximum likelihood estimator can be established under conditions analogous to those used in Gaussian ARMA models, namely, requiring the autoregressive and moving-average polynomials to have no common roots and all roots to lie outside the unit circle; see [16]. However, such results do not directly extend to the present generalized BPSARMA formulation, which combines stochastic seasonal components, dynamic precision governed by a higher-order autoregressive structure, and nonlinear link functions in both submodels. In particular, the interaction between seasonal and nonseasonal autoregressive polynomials and the simultaneous evolution of ϕ_t preclude a direct application of the root-restriction arguments used in the simpler setting.

Because the generalized BPSARMA model is specified through conditional mean and precision recursions, inference and forecasting are developed conditionally on the past information set $\mathcal{F}_{t-1}$. The usual notion of stationarity retains its conventional meaning as a property of the unconditional joint distribution of the process; however, since the estimation framework is entirely conditional and no analytical characterization of the admissible parameter space is currently available for this model class, stationarity is not operationally enforced during estimation. Instead, numerical stability of the conditional recursions serves as the relevant practical criterion.

No explicit stationarity or invertibility constraints are therefore imposed during the numerical maximization of the conditional log-likelihood. As a pragmatic diagnostic strategy, similarly to [17], autoregressive or moving-average parameter estimates with unusually large magnitudes (particularly exceeding one in absolute value) may be regarded as potential indicators of numerical instability in the conditional recursions for η_{1t}. In the empirical applications presented in this paper, no such signs of instability were detected. We emphasize, however, that this is only an ad hoc diagnostic criterion: unstable behavior in the recursions may still arise even when all estimated coefficients lie within $(-1, 1)$, and, conversely, estimates outside this range do not necessarily imply divergence of the fitted process. As in the related GARMA and dynamic beta regression literature, inference and forecasting are grounded in the conditional specification of the model.

3. Conditional Likelihood Inference

Let $\{Y_t\}_{t=1}^n$ denote a realization from a BPSARMA$_\mu(p_1, q) \times (P, Q)_S + \text{AR}_\phi(p_2)$ stochastic process, and let $\boldsymbol{\gamma} := (\alpha_1, \boldsymbol{\beta}^\top, \boldsymbol{\varphi}^\top, \boldsymbol{\theta}^\top, \boldsymbol{\Phi}^\top, \boldsymbol{\Theta}^\top, \alpha_2, \boldsymbol{\delta}^\top)^\top$ denote the w-dimensional parameter vector, where $w := p_1 + q + P + Q + \nu + p_2 + 2$. Let $m := \max\{p_1 + PS,\ q + QS,\ p_2 + 1\}$, which represents the minimum number of initial observations required for the conditional likelihood evaluation. The conditional log-likelihood function can be written as

$$\ell = \ell(\boldsymbol{\gamma}) := \sum_{t=m+1}^{n} \log\big(f(Y_t \mid \mathcal{F}_{t-1}; \mu_t, \phi_t)\big) = \sum_{t=m+1}^{n} \ell_t(\mu_t, \phi_t). \tag{4}$$

where

$$\begin{aligned}\ell_t(\mu_t, \phi_t) :=& [\mu_t(1+\phi_t) - 1]\log(y_t) - [\mu_t(1+\phi_t) + \phi_t + 2]\log(1+y_t) \\ &- \log(\Gamma(\mu_t(1+\phi_t))) - \log(\Gamma(\phi_t+2)) + \log(\Gamma(\mu_t(1+\phi_t) + \phi_t + 2)).\end{aligned}$$

We now derive the conditional score vector and the conditional Fisher information matrix, obtained from the first- and second-order derivatives of Equation (4), respectively. To this end, we partition the parameter vector as $\boldsymbol{\gamma} := (\boldsymbol{\gamma_1}^\top, \boldsymbol{\gamma_2}^\top)^\top$, where

$$\boldsymbol{\gamma_1} := (\alpha_1, \boldsymbol{\beta}^\top, \boldsymbol{\varphi}^\top, \boldsymbol{\theta}^\top, \boldsymbol{\Phi}^\top, \boldsymbol{\Theta}^\top)^\top \quad \text{and} \quad \boldsymbol{\gamma_2} := (\alpha_2, \boldsymbol{\delta}^\top)^\top.$$

As before, we consider the error defined on the predictor scale. In what follows, primes denote derivatives.

3.1. Conditional Score Vector

The conditional score vector is obtained from the first-order derivatives of the conditional log-likelihood function given in Equation (4) with respect to the components of the parameter vector $\boldsymbol{\gamma}$.

Let $\boldsymbol{\kappa_1} := (\kappa_{1m+1}, \ldots, \kappa_{1n})^\top$ and $\boldsymbol{\kappa_2} := (\kappa_{2m+1}, \ldots, \kappa_{2n})^\top$ denote the vectors of partial scores associated with the mean and precision submodels, respectively. We also define the diagonal matrices $\boldsymbol{A_1} := \text{diag}\Big\{\frac{1}{g_1'(\mu_{m+1})}, \ldots, \frac{1}{g_1'(\mu_n)}\Big\}$ and $\boldsymbol{A_2} := \text{diag}\Big\{\frac{1}{g_2'(\phi_{m+1})}, \ldots, \frac{1}{g_2'(\phi_n)}\Big\}$, as well as the vectors $\boldsymbol{z} := (z_m, \ldots, z_{n-1})^\top$ and $\boldsymbol{\nu} := \Big(\frac{\partial \eta_{1m+1}}{\partial \alpha_1}, \ldots, \frac{\partial \eta_{1n}}{\partial \alpha_1}\Big)^\top$. Additionally, let $\boldsymbol{B}$, $\boldsymbol{C_1}$, $\boldsymbol{C_2}$, $\boldsymbol{D_1}$, and $\boldsymbol{D_2}$ be matrices of dimensions $(n-m) \times v$, $(n-m) \times p_1$, $(n-m) \times P$, $(n-m) \times q$, and $(n-m) \times Q$, respectively, with corresponding (i, j) elements

$$B_{i,j} := \frac{\partial \eta_{1i+m}}{\partial \beta_j},\ C_{1i,j} := \frac{\partial \eta_{1i+m}}{\partial \varphi_j},\ C_{2i,j} := \frac{\partial \eta_{1i+m}}{\partial \Phi_j},$$

$$D_{1i,j} := \frac{\partial \eta_{1i+m}}{\partial \theta_j}, \quad \text{and} \quad D_{2i,j} := \frac{\partial \eta_{1i+m}}{\partial \Theta_j}.$$

The score vector is given by

$$\boldsymbol{U}(\gamma) := (U_{\alpha_1}(\gamma), \boldsymbol{U}_{\boldsymbol{\beta}}(\gamma)^\top, \boldsymbol{U}_{\boldsymbol{\varphi}}(\gamma)^\top, \boldsymbol{U}_{\boldsymbol{\theta}}(\gamma)^\top, \boldsymbol{U}_{\boldsymbol{\Phi}}(\gamma)^\top, \boldsymbol{U}_{\boldsymbol{\Theta}}(\gamma)^\top, U_{\alpha_2}(\gamma), U_{\delta}(\gamma))^\top,$$

with components

$$\begin{aligned}
U_{\alpha_1}(\gamma) :=& \boldsymbol{\nu}^\top \boldsymbol{A_1}\boldsymbol{\kappa_1}, \ \boldsymbol{U}_{\boldsymbol{\beta}}(\gamma) := \boldsymbol{B}^\top \boldsymbol{A_1}\boldsymbol{\kappa_1}, \ \boldsymbol{U}_{\boldsymbol{\varphi}}(\gamma) := \boldsymbol{C_1}^\top \boldsymbol{A_1}\boldsymbol{\kappa_1}, \\
\boldsymbol{U}_{\boldsymbol{\theta}}(\gamma) :=& \boldsymbol{D_1}^\top \boldsymbol{A_1}\boldsymbol{\kappa_1}, \ \boldsymbol{U}_{\boldsymbol{\Phi}}(\gamma) := \boldsymbol{C_2}^\top \boldsymbol{A_1}\boldsymbol{\kappa_1}, \ \boldsymbol{U}_{\boldsymbol{\Theta}}(\gamma) := \boldsymbol{D_2}^\top \boldsymbol{A_1}\boldsymbol{\kappa_1}, \\
U_{\alpha_2}(\gamma) :=& \boldsymbol{1}_{\boldsymbol{n}}^\top \boldsymbol{A_2}\boldsymbol{\kappa_2}, \ \text{and} \ U_{\delta}(\gamma) := \boldsymbol{z}^\top \boldsymbol{A_2}\boldsymbol{\kappa_2},
\end{aligned}$$

where $\boldsymbol{1}_{\boldsymbol{n}}$ denotes an $(n-m)\times 1$ vector of ones. Further details are provided in Appendix A.

The conditional maximum likelihood estimator (CMLE) of γ, denoted by $\hat{\gamma}$, is obtained as the solution of the nonlinear system of equations $\boldsymbol{U}(\gamma) = \boldsymbol{0}$, where $\boldsymbol{0}$ denotes a null vector of dimension w. However, this system does not admit a closed-form solution. Consequently, the CMLEs are obtained through iterative computational procedures based on the maximization of the conditional log-likelihood function using Newton or quasi-Newton nonlinear optimization algorithms. For further details, see [18,19].

To implement the iterative procedure, initial values for η_{1t} and η_{2t} must be specified. For $t > m$, the quantities η_{1t} and η_{2t} are computed using the derivatives with respect to the parameters of the mean and precision submodels, respectively.

3.2. Conditional Information Matrix

To obtain the conditional Fisher information matrix, $\boldsymbol{K}(\gamma)$, it is necessary to compute the second-order derivatives of the conditional log-likelihood given in Equation (4) and their expected values. These results allow the computation of the standard errors of the CMLEs, the construction of asymptotic confidence intervals, and the implementation of hypothesis tests. The matrix $\boldsymbol{K}(\gamma)$ corresponds to the conditional Fisher information matrix, obtained as the expectation of the negative Hessian of the conditional log-likelihood.

Let $\boldsymbol{K_1} := \operatorname{diag}\{\lambda_{1m+1}, \ldots, \lambda_{1n}\}$, $\boldsymbol{K_2} := \operatorname{diag}\{\lambda_{2m+1}, \ldots, \lambda_{2n}\}$ and $\boldsymbol{K_3} := \operatorname{diag}\{\lambda_{3m+1}, \ldots, \lambda_{3n}\}$. The conditional Fisher information matrix associated with the parameter vector γ can be expressed as

$$\boldsymbol{K} = \boldsymbol{K}(\gamma) = \begin{bmatrix}
K_{\alpha_1,\alpha_1} & \boldsymbol{K}_{\alpha_1,\boldsymbol{\beta}} & \boldsymbol{K}_{\alpha_1,\boldsymbol{\varphi}} & \boldsymbol{K}_{\alpha_1,\boldsymbol{\theta}} & \boldsymbol{K}_{\alpha_1,\boldsymbol{\Phi}} & \boldsymbol{K}_{\alpha_1,\boldsymbol{\Theta}} & K_{\alpha_1,\alpha_2} & K_{\alpha_1,\delta} \\
\boldsymbol{K}_{\boldsymbol{\beta},\alpha_1} & \boldsymbol{K}_{\boldsymbol{\beta},\boldsymbol{\beta}} & \boldsymbol{K}_{\boldsymbol{\beta},\boldsymbol{\varphi}} & \boldsymbol{K}_{\boldsymbol{\beta},\boldsymbol{\theta}} & \boldsymbol{K}_{\boldsymbol{\beta},\boldsymbol{\Phi}} & \boldsymbol{K}_{\boldsymbol{\beta},\boldsymbol{\Theta}} & \boldsymbol{K}_{\boldsymbol{\beta},\alpha_2} & \boldsymbol{K}_{\boldsymbol{\beta},\delta} \\
\boldsymbol{K}_{\boldsymbol{\varphi},\alpha_1} & \boldsymbol{K}_{\boldsymbol{\varphi},\boldsymbol{\beta}} & \boldsymbol{K}_{\boldsymbol{\varphi},\boldsymbol{\varphi}} & \boldsymbol{K}_{\boldsymbol{\varphi},\boldsymbol{\theta}} & \boldsymbol{K}_{\boldsymbol{\varphi},\boldsymbol{\Phi}} & \boldsymbol{K}_{\boldsymbol{\varphi},\boldsymbol{\Theta}} & \boldsymbol{K}_{\boldsymbol{\varphi},\alpha_2} & \boldsymbol{K}_{\boldsymbol{\varphi},\delta} \\
\boldsymbol{K}_{\boldsymbol{\theta},\alpha_1} & \boldsymbol{K}_{\boldsymbol{\theta},\boldsymbol{\beta}} & \boldsymbol{K}_{\boldsymbol{\theta},\boldsymbol{\varphi}} & \boldsymbol{K}_{\boldsymbol{\theta},\boldsymbol{\theta}} & \boldsymbol{K}_{\boldsymbol{\theta},\boldsymbol{\Phi}} & \boldsymbol{K}_{\boldsymbol{\theta},\boldsymbol{\Theta}} & \boldsymbol{K}_{\boldsymbol{\theta},\alpha_2} & \boldsymbol{K}_{\boldsymbol{\theta},\delta} \\
\boldsymbol{K}_{\boldsymbol{\Phi},\alpha_1} & \boldsymbol{K}_{\boldsymbol{\Phi},\boldsymbol{\beta}} & \boldsymbol{K}_{\boldsymbol{\Phi},\boldsymbol{\varphi}} & \boldsymbol{K}_{\boldsymbol{\Phi},\boldsymbol{\theta}} & \boldsymbol{K}_{\boldsymbol{\Phi},\boldsymbol{\Phi}} & \boldsymbol{K}_{\boldsymbol{\Phi},\boldsymbol{\Theta}} & \boldsymbol{K}_{\boldsymbol{\Phi},\alpha_2} & \boldsymbol{K}_{\boldsymbol{\Phi},\delta} \\
\boldsymbol{K}_{\boldsymbol{\Theta},\alpha_1} & \boldsymbol{K}_{\boldsymbol{\Theta},\boldsymbol{\beta}} & \boldsymbol{K}_{\boldsymbol{\Theta},\boldsymbol{\varphi}} & \boldsymbol{K}_{\boldsymbol{\Theta},\boldsymbol{\theta}} & \boldsymbol{K}_{\boldsymbol{\Theta},\boldsymbol{\Phi}} & \boldsymbol{K}_{\boldsymbol{\Theta},\boldsymbol{\Theta}} & \boldsymbol{K}_{\boldsymbol{\Theta},\alpha_2} & \boldsymbol{K}_{\boldsymbol{\Theta},\delta} \\
K_{\alpha_2,\alpha_1} & \boldsymbol{K}_{\alpha_2,\boldsymbol{\beta}} & \boldsymbol{K}_{\alpha_2,\boldsymbol{\varphi}} & \boldsymbol{K}_{\alpha_2,\boldsymbol{\theta}} & \boldsymbol{K}_{\alpha_2,\boldsymbol{\Phi}} & \boldsymbol{K}_{\alpha_2,\boldsymbol{\Theta}} & K_{\alpha_2,\alpha_2} & K_{\alpha_2,\delta} \\
K_{\delta,\alpha_1} & \boldsymbol{K}_{\delta,\boldsymbol{\beta}} & \boldsymbol{K}_{\delta,\boldsymbol{\varphi}} & \boldsymbol{K}_{\delta,\boldsymbol{\theta}} & \boldsymbol{K}_{\delta,\boldsymbol{\Phi}} & \boldsymbol{K}_{\delta,\boldsymbol{\Theta}} & K_{\delta,\alpha_2} & K_{\delta,\delta}
\end{bmatrix},$$

where $K_{\alpha_1,\alpha_1} := -\boldsymbol{\nu}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{\nu}$, $\boldsymbol{K}_{\alpha_1,\boldsymbol{\beta}} = \boldsymbol{K}_{\boldsymbol{\beta},\alpha_1}^\top := -\boldsymbol{\nu}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{B}$, $\boldsymbol{K}_{\alpha_1,\boldsymbol{\varphi}} = \boldsymbol{K}_{\boldsymbol{\varphi},\alpha_1}^\top := -\boldsymbol{\nu}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{C_1}$, $\boldsymbol{K}_{\alpha_1,\boldsymbol{\theta}} = \boldsymbol{K}_{\boldsymbol{\theta},\alpha_1}^\top := -\boldsymbol{\nu}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{D_1}$, $\boldsymbol{K}_{\alpha_1,\boldsymbol{\Phi}} = \boldsymbol{K}_{\boldsymbol{\Phi},\alpha_1}^\top := -\boldsymbol{\nu}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{C_2}$, $\boldsymbol{K}_{\alpha_1,\boldsymbol{\Theta}} = \boldsymbol{K}_{\boldsymbol{\Theta},\alpha_1}^\top := -\boldsymbol{\nu}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{D_2}$, $K_{\alpha_1,\alpha_2} = K_{\alpha_2,\alpha_1} := -\boldsymbol{\nu}^\top \boldsymbol{A_2}\boldsymbol{K_3}\boldsymbol{A_1}\boldsymbol{1_n}$, $K_{\alpha_1,\delta} = K_{\delta,\alpha_1}^\top := -\boldsymbol{\nu}^\top \boldsymbol{A_2}\boldsymbol{K_3}\boldsymbol{A_1}\boldsymbol{z}$, $\boldsymbol{K}_{\boldsymbol{\beta},\boldsymbol{\beta}} := -\boldsymbol{B}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{B}$, $\boldsymbol{K}_{\boldsymbol{\beta},\boldsymbol{\varphi}} = \boldsymbol{K}_{\boldsymbol{\varphi},\boldsymbol{\beta}}^\top := -\boldsymbol{B}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{C_1}$, $\boldsymbol{K}_{\boldsymbol{\beta},\boldsymbol{\theta}} = \boldsymbol{K}_{\boldsymbol{\theta},\boldsymbol{\beta}}^\top := -\boldsymbol{B}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{D_1}$, $\boldsymbol{K}_{\boldsymbol{\beta},\boldsymbol{\Phi}} = \boldsymbol{K}_{\boldsymbol{\Phi},\boldsymbol{\beta}}^\top := -\boldsymbol{B}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{C_2}$, $\boldsymbol{K}_{\boldsymbol{\beta},\boldsymbol{\Theta}} = \boldsymbol{K}_{\boldsymbol{\Theta},\boldsymbol{\beta}}^\top := -\boldsymbol{B}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{D_2}$, $\boldsymbol{K}_{\boldsymbol{\beta},\alpha_2} = \boldsymbol{K}_{\alpha_2,\boldsymbol{\beta}}^\top := -\boldsymbol{B}^\top \boldsymbol{A_2}\boldsymbol{K_3}\boldsymbol{A_1}\boldsymbol{1_n}$, $\boldsymbol{K}_{\boldsymbol{\beta},\delta} = \boldsymbol{K}_{\delta,\boldsymbol{\beta}}^\top := -\boldsymbol{B}^\top \boldsymbol{A_2}\boldsymbol{K_3}\boldsymbol{A_1}\boldsymbol{z}$, $\boldsymbol{K}_{\boldsymbol{\varphi},\boldsymbol{\varphi}} := -\boldsymbol{C_1}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{C_1}$, $\boldsymbol{K}_{\boldsymbol{\varphi},\boldsymbol{\theta}} = \boldsymbol{K}_{\boldsymbol{\theta},\boldsymbol{\varphi}}^\top := -\boldsymbol{C_1}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{D_1}$, $\boldsymbol{K}_{\boldsymbol{\varphi},\boldsymbol{\Phi}} = \boldsymbol{K}_{\boldsymbol{\Phi},\boldsymbol{\varphi}}^\top := -\boldsymbol{C_1}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{C_2}$, $\boldsymbol{K}_{\boldsymbol{\varphi},\boldsymbol{\Theta}} = \boldsymbol{K}_{\boldsymbol{\Theta},\boldsymbol{\varphi}}^\top := -\boldsymbol{C_1}^\top \boldsymbol{K_1}\boldsymbol{A_1^2}\boldsymbol{D_2}$, $\boldsymbol{K}_{\boldsymbol{\varphi},\alpha_2} =$

$\boldsymbol{K}_{\alpha_2,\boldsymbol{\varphi}}^\top := -\boldsymbol{C}_1^\top \boldsymbol{A}_2\boldsymbol{K}_3\boldsymbol{A}_1\boldsymbol{1}_n$, $\boldsymbol{K}_{\boldsymbol{\varphi},\delta} = \boldsymbol{K}_{\delta,\boldsymbol{\varphi}}^\top := -\boldsymbol{C}_1^\top \boldsymbol{A}_2\boldsymbol{K}_3\boldsymbol{A}_1\boldsymbol{z}$, $\boldsymbol{K}_{\boldsymbol{\theta},\boldsymbol{\theta}} := -\boldsymbol{D}_1^\top \boldsymbol{K}_1\boldsymbol{A}_1^2\boldsymbol{D}_1$, $\boldsymbol{K}_{\boldsymbol{\theta},\boldsymbol{\Phi}} = \boldsymbol{K}_{\boldsymbol{\Phi},\boldsymbol{\theta}}^\top := -\boldsymbol{D}_1^\top \boldsymbol{K}_1\boldsymbol{A}_1^2\boldsymbol{C}_2$, $\boldsymbol{K}_{\boldsymbol{\theta},\boldsymbol{\Theta}} = \boldsymbol{K}_{\boldsymbol{\Theta},\boldsymbol{\theta}}^\top := -\boldsymbol{D}_1^\top \boldsymbol{K}_1\boldsymbol{A}_1^2\boldsymbol{D}_2$, $\boldsymbol{K}_{\boldsymbol{\theta},\alpha_2} = \boldsymbol{K}_{\alpha_2,\boldsymbol{\theta}}^\top := -\boldsymbol{D}_1^\top \boldsymbol{A}_2\boldsymbol{K}_3\boldsymbol{A}_1\boldsymbol{1}_n$, $\boldsymbol{K}_{\boldsymbol{\theta},\delta} = \boldsymbol{K}_{\delta,\boldsymbol{\theta}}^\top := -\boldsymbol{D}_1^\top \boldsymbol{A}_2\boldsymbol{K}_3\boldsymbol{A}_1\boldsymbol{z}$, $\boldsymbol{K}_{\boldsymbol{\Phi},\boldsymbol{\Phi}} := -\boldsymbol{C}_2^\top \boldsymbol{K}_1\boldsymbol{A}_1^2\boldsymbol{C}_2$, $\boldsymbol{K}_{\boldsymbol{\Phi},\boldsymbol{\Theta}} = \boldsymbol{K}_{\boldsymbol{\Theta},\boldsymbol{\Phi}}^\top := -\boldsymbol{C}_2^\top \boldsymbol{K}_1\boldsymbol{A}_1^2\boldsymbol{D}_2$, $\boldsymbol{K}_{\boldsymbol{\Phi},\alpha_2} = \boldsymbol{K}_{\alpha_2,\boldsymbol{\Phi}}^\top := -\boldsymbol{C}_2^\top \boldsymbol{A}_2\boldsymbol{K}_3\boldsymbol{A}_1\boldsymbol{1}_n$, $\boldsymbol{K}_{\boldsymbol{\Phi},\delta} = \boldsymbol{K}_{\delta,\boldsymbol{\Phi}}^\top := -\boldsymbol{C}_2^\top \boldsymbol{A}_2\boldsymbol{K}_3\boldsymbol{A}_1\boldsymbol{z}$, $\boldsymbol{K}_{\boldsymbol{\Theta},\boldsymbol{\Theta}} := -\boldsymbol{D}_2^\top \boldsymbol{K}_1\boldsymbol{A}_1^2\boldsymbol{D}_2$, $\boldsymbol{K}_{\boldsymbol{\Theta},\alpha_2} = \boldsymbol{K}_{\alpha_2,\boldsymbol{\Theta}}^\top := -\boldsymbol{D}_2^\top \boldsymbol{A}_2\boldsymbol{K}_3\boldsymbol{A}_1\boldsymbol{1}_n$, $\boldsymbol{K}_{\boldsymbol{\Theta},\delta} = \boldsymbol{K}_{\delta,\boldsymbol{\Theta}}^\top := -\boldsymbol{D}_2^\top \boldsymbol{A}_2\boldsymbol{K}_3\boldsymbol{A}_1\boldsymbol{z}$, $K_{\alpha_2,\alpha_2} := -\boldsymbol{1}_n^\top \boldsymbol{K}_2\boldsymbol{A}_2^2\boldsymbol{1}_n$, $\boldsymbol{K}_{\alpha_2,\delta} = \boldsymbol{K}_{\delta,\alpha_2}^\top := -\boldsymbol{1}_n^\top \boldsymbol{K}_2\boldsymbol{A}_2^2\boldsymbol{z}$ and $\boldsymbol{K}_{\delta,\delta} := -\boldsymbol{z}^\top \boldsymbol{K}_2\boldsymbol{A}_2^2\boldsymbol{z}$. For details, see Appendix A.

Under the usual regularity conditions and for sufficiently large n, the CMLE $\hat{\gamma}$ is consistent and asymptotically normally distributed [20]. In particular,

$$\hat{\gamma} \sim \mathcal{N}_w\left(\gamma, \boldsymbol{K}^{-1}(\gamma)\right),$$

approximately, where $\mathcal{N}_w$ denotes the w-dimensional normal distribution, and $\hat{\gamma}$ is the CMLE of γ.

3.3. Confidence Intervals and Hypothesis Testing

Let γ_s denote the s-th component of the parameter vector γ, $\hat{\gamma}_s$ the corresponding conditional maximum likelihood estimator contained in $\hat{\gamma}$, and $K(\hat{\gamma})^{ss}$ the s-th element of the main diagonal of the inverse Fisher information matrix evaluated at $\hat{\gamma}$, for $s \in \{1, \ldots, w\}$. From the asymptotic normality of the CMLE of γ, it follows that

$$\frac{\hat{\gamma}_s - \gamma_s}{\sqrt{K(\hat{\gamma})^{ss}}} \sim \mathcal{N}(0,1),$$

approximately. This result can be used to construct asymptotic confidence intervals as well as to perform hypothesis tests on the model parameters.

An asymptotic confidence interval for γ_s, with confidence level $100(1-\varepsilon)\%$, $0 < \varepsilon < 1$, is given by

$$\left[\hat{\gamma}_s - z_{1-\varepsilon/2}\sqrt{K(\hat{\gamma})^{ss}}, \hat{\gamma}_s + z_{1-\varepsilon/2}\sqrt{K(\hat{\gamma})^{ss}}\right],$$

where $z_{1-\varepsilon/2}$ denotes the $(1-\varepsilon/2)$ quantile of the standard normal distribution. For further details on asymptotic confidence intervals, see [21].

The asymptotic normality of $\hat{\gamma}$ can also be used to perform hypothesis tests on the parameters indexing the model. Suppose we wish to test $\mathcal{H}_0\colon \gamma_s = \gamma_s^{(0)}$ against $\mathcal{H}_1\colon \gamma_s \neq \gamma_s^{(0)}$. The z test statistic is given by

$$z := \frac{\hat{\gamma}_s - \gamma_s^{(0)}}{\sqrt{K(\hat{\gamma})^{ss}}}.$$

Under $\mathcal{H}_0$ and for sufficiently large n, z is approximately standard normally distributed. Thus, for a given significance level ε, $\mathcal{H}_0$ is rejected if $|z| > z_{1-\varepsilon/2}$.

When the interest lies in simultaneously testing multiple parametric restrictions, more general hypothesis tests may be employed, such as the likelihood ratio test [22], the score test [23], and the Wald test [24]. Under $\mathcal{H}_0$ and for sufficiently large samples, the statistics associated with these tests are asymptotically χ^2-distributed, with degrees of freedom equal to the number of imposed restrictions.

4. Model Selection, Diagnostic Analysis, and Forecasting

In this section, we present the model selection criteria, diagnostic tools, and forecasting procedures adopted for the generalized BPSARMA model. The diagnostic analysis aims to assess whether the fitted model adequately captures the dynamic structure present in the data. If this condition is satisfied, the model can be regarded as suitable for producing reliable out-of-sample forecasts.

4.1. Model Selection

Information criteria are widely used in model selection, particularly when comparing competing specifications within the same model class. For selecting the generalized BPSARMA model, we adopt the modified Bayesian information criterion (MBIC) proposed by [3]. This modified criterion is defined as

$$\mathrm{MBIC} := -2\hat{\ell} \times \frac{n}{n-m} + w\log(n),$$

where $\hat{\ell}$ denotes the maximized value of the log-likelihood function. The authors also considered a modified version of the Akaike Information Criterion (MAIC).

4.2. Residuals

Residual analysis constitutes a fundamental step in assessing the adequacy of the estimated model to the observed data. To investigate the fit of the generalized BPSARMA model, we employ the quantile residual, which is widely used in the literature on regression and time series models. This residual is defined as

$$\hat{r}_t^q := \Phi^{-1}\big(F(Y_t \mid \mathcal{F}_{t-1}; \hat{\mu}_t, \hat{\phi}_t)\big),$$

where $\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution and $\Phi^{-1}(\cdot)$ its corresponding quantile function, and $F(Y_t \mid \mathcal{F}_{t-1}; \hat{\mu}_t, \hat{\phi}_t)$ represents the conditional cumulative distribution function of the beta prime distribution, conditional on the past information set $\mathcal{F}_{t-1}$, evaluated at Y_t. Under correct model specification, the quantile residual $\hat{r}_t^q$ is asymptotically standard normally distributed [25].

4.3. Portmanteau Tests

Portmanteau-type tests are widely used to assess the adequacy of a fitted model. This assessment is based on jointly verifying whether the first v autocorrelations of the residuals equal zero; that is, the null hypothesis $\mathcal{H}_0\colon \rho_1 = \cdots = \rho_v = 0$ is tested. One of the most commonly used test statistics is the one proposed by [26]. An alternative is the statistic proposed by [27], which is based on the partial autocorrelations of the residuals.

The model is considered adequate if the residual autocorrelations (or partial autocorrelations) do not jointly differ significantly from zero. Otherwise, the fitted specification may be regarded as inadequate, and the fitted model is rejected. Under the null hypothesis, these statistics are asymptotically χ^2-distributed with $v-(p+q+P+Q)$ degrees of freedom. Ref. [3] suggests adopting $v = \max\{10, 2S\}$.

4.4. Forecasting

In-sample forecasts for the conditional mean of Y_t, with $t \in \{m+1, \ldots, n\}$, are given by the fitted conditional means $\hat{\mu}_t$. These forecasts are constructed by setting $\hat{r}_t = 0$ for $t \in \{1, \ldots, m\}$ and, for $t \in \{m+1, \ldots, n\}$, taking $\hat{r}_t = g_1(Y_t) - g_1(\hat{\mu}_t)$. Additionally, the unknown parameters in Equation (2) are replaced by their conditional maximum likelihood estimates, after which the inverse link function is applied. Thus, the in-sample forecast $\hat{\mu}_t$, for $t \in \{m+1, \ldots, n\}$, is given by

$$\begin{aligned}\hat{\mu}_t = g_1^{-1}\Big(&\hat{\alpha}_1 + \boldsymbol{x}_t^\top\hat{\boldsymbol{\beta}} + \sum_{i=1}^{p_1}\hat{\varphi}_i\big[g_1(Y_{t-i}) - \boldsymbol{x}_{t-i}^\top\hat{\boldsymbol{\beta}}\big] + \sum_{I=1}^{P}\hat{\Phi}_I\big[g_1(Y_{t-IS}) - \boldsymbol{x}_{t-IS}^\top\hat{\boldsymbol{\beta}}\big] \\ &- \sum_{i=1}^{p_1}\sum_{I=1}^{P}\hat{\varphi}_i\hat{\Phi}_I\big[g_1(Y_{t-(i+IS)}) - \boldsymbol{x}_{t-(i+IS)}^\top\hat{\boldsymbol{\beta}}\big] - \sum_{j=1}^{q}\hat{\theta}_j\hat{r}_{t-j} - \sum_{J=1}^{Q}\hat{\Theta}_J\hat{r}_{t-JS}\end{aligned}$$

$$+\sum_{j=1}^{q}\sum_{J=1}^{Q}\hat{\theta}_j\hat{\Theta}_J\hat{r}_{t-(j+JS)}\Big).$$

Similarly, out-of-sample forecasts can also be obtained. The h-step-ahead forecast of the conditional mean, for $h \in \{1,2,\dots\}$, is denoted by $\hat{\mu}_{n+h}$ and is given by

$$\begin{aligned}\hat{\mu}_{n+h} = g_1^{-1}\Big(&\hat{\alpha}_1 + \boldsymbol{x}_{n+h}^{\top}\hat{\boldsymbol{\beta}} + \sum_{i=1}^{p_1}\hat{\varphi}_i\big[g_1(Y_{n+h-i}) - \boldsymbol{x}_{n+h-i}^{\top}\hat{\boldsymbol{\beta}}\big] + \sum_{I=1}^{P}\hat{\Phi}_I\big[g_1(Y_{n+h-IS})\\ &- \boldsymbol{x}_{n+h-IS}^{\top}\hat{\boldsymbol{\beta}}\big] - \sum_{i=1}^{p_1}\sum_{I=1}^{P}\hat{\varphi}_i\hat{\Phi}_I\big[g_1(Y_{n+h-(i+IS)}) - \boldsymbol{x}_{n+h-(i+IS)}^{\top}\hat{\boldsymbol{\beta}}\big] - \sum_{j=1}^{q}\hat{\theta}_j\hat{r}_{n+h-j}\\ &- \sum_{J=1}^{Q}\hat{\Theta}_J\hat{r}_{n+h-JS} + \sum_{j=1}^{q}\sum_{J=1}^{Q}\hat{\theta}_j\hat{\Theta}_J\hat{r}_{n+h-(j+JS)}\Big),\end{aligned}$$

where

$$[g_1(Y_t)] = \begin{cases} g_1(\hat{\mu}_t), & \text{if} \quad t > n,\\ g_1(Y_t), & \text{if} \quad t \le n.\end{cases}$$

5. Simulation Results

In this section, we present Monte Carlo simulation evidence on the finite-sample properties of the CMLEs of the parameters indexing the $\text{BPSARMA}_{\mu}(p_1,q)\times(P,Q)_S + \text{AR}_{\phi}(p_2)$ model. We consider four sample sizes, namely $n \in \{100, 200, 500, 1000\}$. In each replication, a burn-in period of length 1000 was used; that is, time series of length $n + 1000$ were generated, and the first 1000 observations were discarded in order to mitigate the influence of initial conditions on the simulated sample paths. The resulting series are denoted by $y_1,\dots,y_n$. Each series was generated from the generalized BPSARMA model specified in (2) and (3), using the logarithmic link function in both submodels.

Parameter estimation was performed by numerically maximizing the conditional log-likelihood function. For this purpose, we employed the quasi-Newton limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm with box constraints (L-BFGS-B), using analytical first derivatives. The initial parameter values were specified as follows: (i) the initial values for α_1, $\varphi_1,\dots,\varphi_{p_1}$, and $\Phi_1,\dots,\Phi_P$ were obtained from ordinary least squares estimates of the linear regression of $g_1(y_t)$ on $g_1(y_{t-1}),\dots,g_1(y_{t-p_1})$ and $g_1(y_{t-(p_1+1)}),\dots,g_1(y_{t-(p_1+P)})$, and in this case, the initial value of α_1 corresponds to the intercept of the regression; (ii) the parameters $\theta_1,\dots,\theta_q$, and $\Theta_1,\dots,\Theta_Q$ were initialized at zero; (iii) the initial value of α_2 was obtained as

$$g_2\!\left(n^{-1}\sum_{t=1}^{n} y_t(1 + n^{-1}\sum_{t=1}^{n} y_t)/s^2\right);$$

where s^2 is the sample variance of $y_1,\dots,y_n$; and, (iv) finally, the coefficients of the precision submodel, $\delta_1,\dots,\delta_{p_2}$, were initialized at -1, providing simple feasible starting values for the numerical optimization.

The initialization strategy was designed to provide starting values that are consistent with the dynamic structure of the model while maintaining numerical stability during optimization. In particular, the initial values of α_1, $\varphi_1,\dots,\varphi_{p_1}$, and $\Phi_1,\dots,\Phi_P$ were obtained from a linear approximation to the conditional mean dynamics on the link-function scale. The moving-average parameters were initialized at zero, a common choice in likelihood-based estimation of ARMA-type models, since reliable preliminary estimates of these parameters are generally unavailable. The initial value of α_2 was obtained by matching the sample moments to the variance structure implied by the beta prime distribution.

The coefficients of the precision submodel were initialized at -1, providing a simple and feasible starting configuration while allowing for an initial negative association between recent relative changes in the series and the precision dynamics, which is often observed in preliminary empirical analyses. Alternative initialization strategies, including different fixed starting values for the moving-average parameters and for the coefficients of the precision submodel, were also explored in preliminary numerical experiments. The adopted scheme was retained because it consistently yielded stable convergence and satisfactory optimization performance while remaining computationally simple.

The simulations were carried out using the statistical computing environment R [28]. A total of 5000 Monte Carlo replications were conducted, and no convergence failures occurred during the numerical maximization of the log-likelihood function. We adopt the logarithmic link function for both the mean and precision submodels.

In the first simulation scenario, the data were generated from the BPSARMA$_\mu(1,1) \times (1,1)_{12} + \mathrm{AR}_\phi(2)$ model. The parameter values were set to $\alpha_1 = 0.13$, $\varphi_1 = 0.65$, $\theta_1 = -0.08$, $\Phi_1 = 0.95$, $\Theta_1 = 0.68$, $\alpha_2 = 3.17$, $\delta_1 = -1.39$, and $\delta_2 = -1.20$. These values correspond to the point estimates obtained from fitting the BPSARMA$_\mu(1,1) \times (1,1)_{12} + \mathrm{AR}_\phi(2)$ model to real data on inflow to the Estreito reservoir and therefore represent a realistic empirical configuration.

Table 1 reports, for each estimator, the estimated mean (Mean), standard deviation (SD), bias (Bias), and mean squared error (MSE). As expected, both the biases and the mean squared errors decrease as the sample size increases, providing evidence of the consistency of the CMLEs. It is also observed that the parameters φ, θ, Φ, and Θ tend to be underestimated. This behavior is consistent with the results reported for the βSARMA model in [3].

The tendency toward negative bias in the estimators of φ_1, θ_1, Φ_1, and Θ_1 is likely a finite-sample effect associated with the estimation of dynamic dependence parameters in ARMA-type models. Similar behavior has been reported for the βSARMA model by [3]. The effect is particularly pronounced for the seasonal parameters, whose estimation relies on observations separated by one seasonal cycle. Consequently, the effective amount of information available for estimating Φ_1 and Θ_1 is smaller than that available for the corresponding nonseasonal parameters, leading to larger finite-sample biases. As expected, these biases decrease steadily as the sample size increases.

In addition, the biases of the estimators of the nonseasonal parameters, φ_1 and θ_1, decrease more rapidly than those of the seasonal parameters, Φ_1 and Θ_1. This finding corroborates the discussion in [29], according to which the estimators of seasonal parameters tend to exhibit larger biases than the remaining estimators. For example, when $n = 100$, the estimated bias of $\hat{\varphi}_1$ ($\hat{\theta}_1$) is -0.0470 (-0.0216), whereas that of $\hat{\Phi}_1$ ($\hat{\Theta}_1$) is -0.1216 (-0.1555). Overall, the results indicate that the parameters of the generalized BPSARMA model can be estimated with good accuracy, particularly as the sample size increases.

A different pattern is observed for the intercept parameter α_1, whose estimator exhibits a comparatively larger finite-sample bias. This behavior may be attributed, at least in part, to the fact that the intercept is estimated jointly with the autoregressive, moving-average, seasonal, and precision components and therefore reflects the cumulative effect of finite-sample estimation errors in the remaining parameters. Nevertheless, the bias decreases rapidly as the sample size increases, falling from 0.3652 when $n = 100$ to 0.0330 when $n = 1000$, which provides further evidence of the consistency of the CMLE.

Table 1. Monte Carlo results for the CMLEs under the BPSARMA$_{\mu}(1,1) \times (1,1)_{12} + \text{AR}_{\phi}(2)$ model: mean (Mean), standard deviation (SD), bias (Bias), and mean squared error (MSE); first scenario.

	n	Parameter							
		α_1 **0.13**	φ_1 **0.65**	θ_1 **−0.08**	Φ_1 **0.95**	Θ_1 **0.68**	α_2 **3.17**	δ_1 **−1.39**	δ_2 **−1.20**
Mean	100	0.4952	0.6030	−0.1016	0.8284	0.5245	3.1818	−1.3790	−1.1361
	200	0.2897	0.6291	−0.0894	0.8922	0.5949	3.1669	−1.3887	−1.1704
	500	0.1928	0.6426	−0.0821	0.9262	0.6407	3.1642	−1.3834	−1.1817
	1000	0.1630	0.6466	−0.0810	0.9373	0.6579	3.1657	−1.3895	−1.1888
SD	100	0.5007	0.1346	0.1655	0.1407	0.1853	0.1846	0.6467	0.6316
	200	0.1894	0.0801	0.1014	0.0639	0.0950	0.1226	0.3908	0.3719
	500	0.0687	0.0463	0.0599	0.0252	0.0463	0.0842	0.2259	0.2168
	1000	0.0378	0.0321	0.0415	0.0148	0.0298	0.0608	0.1534	0.1502
Bias	100	0.3652	−0.0470	−0.0216	−0.1216	−0.1555	0.0118	0.0110	0.0639
	200	0.1597	−0.0209	−0.0094	−0.0578	−0.0851	−0.0031	0.0013	0.0296
	500	0.0628	−0.0074	−0.0021	−0.0238	−0.0393	−0.0058	0.0066	0.0183
	1000	0.0330	−0.0034	−0.0010	−0.0127	−0.0221	−0.0043	0.0005	0.0112
MSE	100	0.3841	0.0203	0.0279	0.0346	0.0585	0.0342	0.4184	0.4030
	200	0.0614	0.0069	0.0104	0.0074	0.0163	0.0150	0.1528	0.1392
	500	0.0087	0.0022	0.0036	0.0012	0.0037	0.0071	0.0511	0.0474
	1000	0.0025	0.0010	0.0017	0.0004	0.0014	0.0037	0.0235	0.0227

We next present four additional simulation scenarios aimed at evaluating the finite-sample performance of inference for the parameters indexing the proposed model. For these additional scenarios, we considered only the sample sizes $n \in \{200, 1000\}$. These sample sizes were selected to represent moderate- and large-sample settings, allowing us to evaluate the rate at which the asymptotic properties of the estimators become evident. In addition to the point-estimation results, we report the empirical coverage rates (CRs) of the asymptotic 95% confidence intervals. The adopted specification is the BPSARMA$_{\mu}(1,1) \times (0,1)_{12} + \text{AR}_{\phi}(2)$ model, under different combinations of the parameter values Θ_1 and δ_1, allowing us to investigate distinct levels of seasonal dependence as well as different dynamic regimes in the variable-precision component. Specifically, the parameter values were set to $\alpha_1 = 1.69$, $\varphi_1 = 0.64$, $\theta_1 = -0.21$, $\Theta_1 \in \{-0.46, 0.46\}$, $\alpha_2 = 1.84$, $\delta_1 \in \{-0.37, -2.5\}$, and $\delta_2 = -0.63$.

The results for the four additional simulation scenarios are reported in Tables 2–5. Overall, the point estimators display satisfactory finite-sample performance. In general, the means of the estimates remain close to the corresponding true parameter values, even for the smaller sample size ($n = 200$), indicating relatively small biases. Moreover, these biases systematically decrease as the sample size increases. A substantial reduction in the standard deviations is also observed as n increases. For nearly all parameters, the variability of the estimators when $n = 1000$ is less than half of that observed for $n = 200$, reflecting considerable gains in precision. Consequently, the mean squared errors decrease markedly across the sample sizes. Taken together, these results indicate that increasing the amount of sample information contributes simultaneously to reducing both bias and variability, which is consistent with the asymptotic behavior expected for the CMLEs and provides favorable empirical support for their use in moderate and large samples.

A comparison across the different simulation scenarios reveals that the estimators of the parameters associated with the precision dynamics (δ_1 and δ_2) generally exhibit the largest standard deviations and mean squared errors, particularly in the scenarios where $|\delta_1|$ assumes larger values. This suggests that stronger precision dynamics make the estimation of these parameters more challenging in finite samples. Nevertheless, increasing the sample size leads to a substantial reduction in both estimator variability and mean squared error, indicating that these effects are progressively mitigated as more information becomes available.

The coverage rates of the 95% confidence intervals are broadly consistent with the behavior predicted by the asymptotic theory. For $n = 1000$, the empirical coverage rates are very close to the nominal level for all parameters and in all scenarios, generally ranging between 0.94 and slightly above 0.95. For $n = 200$, a mild undercoverage is observed for some parameters, particularly those associated with the conditional mean dynamics, with coverage rates ranging approximately from 0.92 to 0.94. However, these deviations from the nominal level are small in magnitude and become progressively less pronounced as the sample size increases, providing empirical support for the adequacy of the asymptotic approximations used for interval estimation.

Table 2. Monte Carlo results for the CMLEs under the BPSARMA$_\mu(1,1) \times (0,1)_{12} + \text{AR}_\phi(2)$ model: mean (Mean), standard deviation (SD), bias (Bias), mean squared error (MSE), and coverage rate (CR); second scenario.

	n	Parameter						
		α_1 **1.69**	φ_1 **0.64**	θ_1 **−0.21**	Θ_1 **−0.46**	α_2 **1.84**	δ_1 **−0.37**	δ_2 **−0.63**
Mean	200	1.7749	0.6178	−0.2193	−0.4491	1.8611	−0.3489	−0.6038
	1000	1.7027	0.6365	−0.2106	−0.4578	1.8433	−0.3650	−0.6260
SD	200	0.3721	0.0854	0.1095	0.0685	0.1433	0.2892	0.2903
	1000	0.1497	0.0341	0.0453	0.0271	0.0600	0.1156	0.1128
Bias	200	0.0849	−0.0222	−0.0093	0.0109	0.0211	0.0211	0.0262
	1000	0.0127	−0.0035	−0.0006	0.0022	0.0033	0.0050	0.0040
MSE	200	0.1457	0.0078	0.0121	0.0048	0.0210	0.0841	0.0850
	1000	0.0226	0.0012	0.0021	0.0007	0.0036	0.0134	0.0127
CR	200	0.9315	0.9374	0.9233	0.9276	0.9244	0.9557	0.9542
	1000	0.9480	0.9492	0.9478	0.9451	0.9506	0.9513	0.9508

Table 3. Monte Carlo results for the CMLEs under the BPSARMA$_\mu(1,1) \times (0,1)_{12} + \text{AR}_\phi(2)$ model: mean (Mean), standard deviation (SD), bias (Bias), mean squared error (MSE), and coverage rate (CR); third scenario.

	n	Parameter						
		α_1 **1.69**	φ_1 **0.64**	θ_1 **−0.21**	Θ_1 **0.46**	α_2 **1.84**	δ_1 **−0.37**	δ_2 **−0.63**
Mean	200	1.7954	0.6170	−0.2206	0.4601	1.8577	−0.3486	−0.5930
	1000	1.7092	0.6358	−0.2111	0.4593	1.8426	−0.3627	−0.6213
SD	200	0.3814	0.0837	0.1082	0.0653	0.1430	0.2924	0.2857
	1000	0.1568	0.0343	0.0451	0.0265	0.0602	0.1156	0.1133
Bias	200	0.1054	−0.0230	−0.0106	0.0001	0.0177	0.0214	0.0370
	1000	0.0192	−0.0042	−0.0011	−0.0007	0.0026	0.0073	0.0087
MSE	200	0.1566	0.0075	0.0118	0.0043	0.0208	0.0859	0.0830
	1000	0.0250	0.0012	0.0020	0.0007	0.0036	0.0134	0.0129
CR	200	0.9295	0.9436	0.9347	0.9305	0.9221	0.9555	0.9538
	1000	0.9471	0.9498	0.9461	0.9333	0.9496	0.9516	0.9502

Table 4. Monte Carlo results for the CMLEs under the BPSARMA$_\mu(1,1) \times (0,1)_{12} + \text{AR}_\phi(2)$ model: mean (Mean), standard deviation (SD), bias (Bias), mean squared error (MSE), and coverage rate (CR); fourth scenario.

	n	Parameter						
		α_1 **1.69**	φ_1 **0.64**	θ_1 **−0.21**	Θ_1 **−0.46**	α_2 **1.84**	δ_1 **−2.50**	δ_2 **−0.63**
Mean	200	1.7439	0.6324	−0.2113	−0.4465	1.8490	−2.5185	−0.6179
	1000	1.6974	0.6375	−0.2095	−0.4576	1.8409	−2.5017	−0.6293
SD	200	0.3056	0.0722	0.0776	0.0558	0.1555	0.4423	0.2702
	1000	0.1177	0.0275	0.0305	0.0220	0.0638	0.1756	0.1016
Bias	200	0.0539	−0.0167	−0.0013	0.0135	0.0090	−0.0185	0.0121
	1000	0.0074	−0.0025	0.0005	0.0024	0.0009	−0.0017	0.0007
MSE	200	0.0963	0.0055	0.0060	0.0033	0.0243	0.1959	0.0732
	1000	0.0139	0.0008	0.0009	0.0005	0.0041	0.0308	0.0103
CR	200	0.9184	0.9382	0.9128	0.9287	0.9049	0.9533	0.9541
	1000	0.9430	0.9489	0.9471	0.9466	0.9453	0.9497	0.9509

Table 5. Monte Carlo results for the CMLEs under the BPSARMA$_\mu(1,1) \times (0,1)_{12} + \text{AR}_\phi(2)$ model: mean (Mean), standard deviation (SD), bias (Bias), mean squared error (MSE) and coverage rate (CR); fifth scenario.

	n	Parameter						
		α_1 **1.69**	φ_1 **0.64**	θ_1 **−0.21**	Θ_1 **0.46**	α_2 **1.84**	δ_1 **−2.50**	δ_2 **−0.63**
Mean	200	1.7747	0.6214	−0.2107	0.4562	1.8342	−2.4506	−0.5867
	1000	1.7037	0.6370	−0.2093	0.4591	1.8382	−2.4839	−0.6200
SD	200	0.3219	0.0713	0.0797	0.0471	0.1634	0.4495	0.2747
	1000	0.1243	0.0274	0.0312	0.0184	0.0662	0.1756	0.1002
Bias	200	0.0847	−0.0186	−0.0007	−0.0038	−0.0058	0.0494	0.0433
	1000	0.0137	−0.0030	0.0007	−0.0009	−0.0018	0.0161	0.0100
MSE	200	0.1108	0.0054	0.0064	0.0022	0.0267	0.2045	0.0773
	1000	0.0156	0.0008	0.0010	0.0003	0.0044	0.0311	0.0101
CR	200	0.9224	0.9338	0.9231	0.9216	0.9160	0.9532	0.9521
	1000	0.9425	0.9457	0.9444	0.9408	0.9436	0.9512	0.9508

In summary, the simulation study provides evidence that the conditional maximum likelihood estimators, both point and interval, perform satisfactorily in finite samples. The point estimators generally exhibit small biases, decreasing variability as the sample size increases, and empirical coverage rates close to the nominal level. The results also reinforce the consistency of the point estimators. Overall, the inferential procedures based on the asymptotic normal approximation show adequate performance, particularly for moderate and large sample sizes, yielding confidence intervals with coverage probabilities compatible with those theoretically expected.

6. Hydro-Environmental Applications and Out-of-Sample Forecasting

In this section, we analyze the monthly mean inflow to the reservoirs of five Brazilian hydroelectric power plants. The plants were selected so as to represent the five major geographic regions of the country: Estreito (North), Sobradinho (Northeast), Manso (Central-West), Camargos (Southeast), and Itaipu (South). Reservoir inflow corresponds to the volume of water entering the reservoir from upstream sources and constitutes a key indicator for water resource planning and management. Figure 1 presents the time series plots for the five inflow series: Manso (top panel), Camargos (middle-left panel), Estreito (middle-right panel), Sobradinho (bottom-left panel), and Itaipu (bottom-right panel). All series display pronounced seasonal fluctuations. In this context, and with the aim of improving the efficiency of water resource utilization, the primary objective of this empirical application is to obtain accurate forecasts of future inflow values.

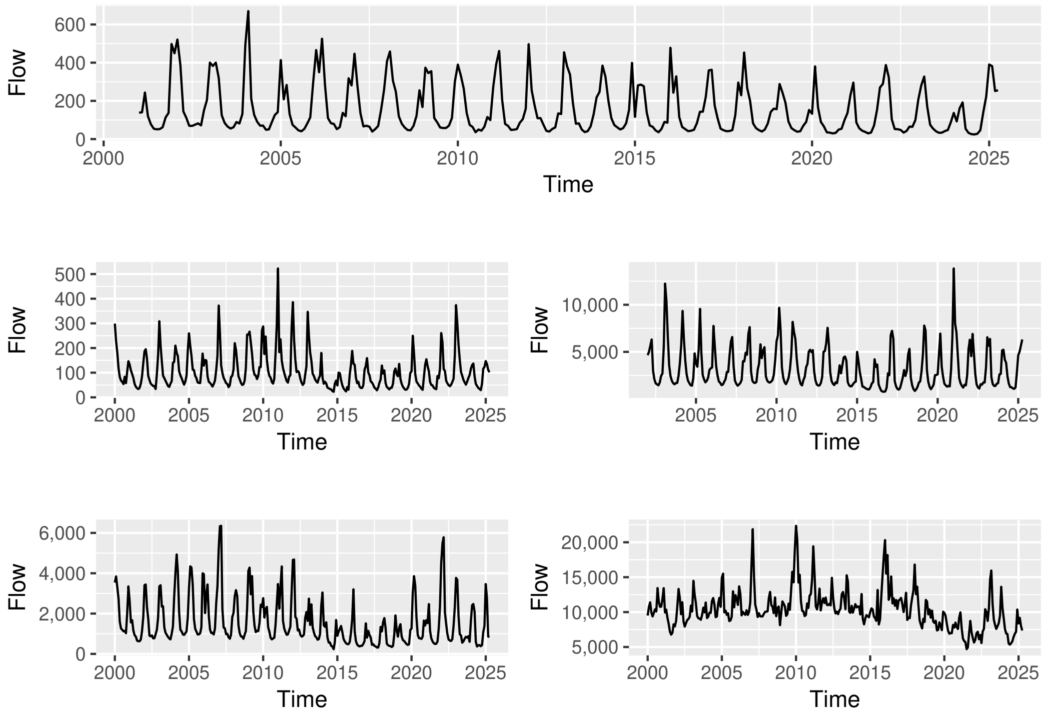

Figure 1. Monthly inflow time series for the reservoirs of Manso (**top panel**), Camargos (**middle-left panel**), Estreito (**middle-right panel**), Sobradinho (**bottom-left panel**), and Itaipu (**bottom-right panel**).

The data were obtained from the Brazilian National Electric System Operator (ONS, https://www.ons.org.br, accessed on 16 March 2026) and are expressed in cubic meters per second ($\mathrm{m}^3\,\mathrm{s}^{-1}$). In all five applications, October 2025 was adopted as the cutoff date. The starting dates vary across reservoirs because, for each case, the longest available historical series was used so as to maximize the information available for estimation. The six most recent observations (from May 2025 to October 2025) were removed from the sample prior to the modeling stage and subsequently used to evaluate the out-of-sample predictive performance of the competing models. The withheld values are: Manso {91.6229, 56.7200, 45.3168, 31.6503, 32.9113, 37.5645}; Camargos {63.0226, 58.7770, 44.1452, 36.0968, 36.1980, 33.3113}; Estreito {3024.2884, 2095.7993, 1656.8719, 1544.5168, 1453.4987, 1743.4229}; Sobradinho {692.5806, 516.6667, 485.8065, 434.8387, 458.0000, 469.3548}; and Itaipu {7066.511, 7084.345, 6073.487, 5890.801, 6760.856, 7012.301}.

Generalized BPSARMA model selection was conducted by considering all combinations of orders $p_1, q \in \{0, \ldots, 4\}$, $P, Q \in \{0, 1, 2\}$, and $p_2 \in \{1, \ldots, 4\}$. For all fitted models, logarithmic link functions were adopted for both the mean and precision submodels. In addition, exogenous regressors were incorporated, including harmonic components: $\boldsymbol{x}_t = (\sin(2\pi t/12), \cos(2\pi t/12))^\top$, for $t \in \{1, \ldots, n\}$. During the modeling process, the inclusion of harmonic regressors in the generalized BPSARMA specification improved the goodness of fit and enhanced out-of-sample forecasting accuracy. Seasonal dummy variables were also considered in order to further reinforce the modeling of seasonal fluctuations.

The final model selection was based on the MBIC information criterion, the statistical significance of the parameters (assessed through z tests), inspection of the autocorrelation function (ACF) and partial autocorrelation function (PACF) of the residuals, and the Ljung–Box and Monti Portmanteau tests. Whenever components of the conditional mean or precision submodel (including regressors and both nonseasonal and seasonal autoregressive and moving-average terms) were not statistically significant at the 5% level, they were removed from the specification. When more than one non-significant component was present, elimination proceeded sequentially in decreasing order of their p-values.

In order to assess out-of-sample predictive performance, in addition to the proposed model, we also considered the BPSARMA model with fixed precision. We further included the following competing approaches: the generalized BPARMA model with harmonic regressors [8], the Gaussian SARIMA model, and the Holt–Winters exponential smoothing algorithm with both additive and multiplicative seasonality. Analogously to the generalized BPSARMA specification, the BPSARMA model with fixed precision was fitted including harmonic regressors and seasonal dummy variables, with model selection based on the MBIC criterion. The generalized BPARMA model was estimated using the same harmonic regressors and dummy variables, with specification chosen according to the Bayesian Information Criterion (BIC). The SARIMA specification was selected using the `auto.arima` function from the *forecast* package in the R statistical computing environment, adopting the arguments `ic = ‘‘bic’’`, `d = 0`, and `allowdrift = FALSE`, thereby allowing models with a linear trend component to be considered. The Holt–Winters method was fitted using the `hw` function from the same package.

The competing methods were selected to assess different aspects of the proposed forecasting framework. The BPSARMA model with fixed precision allows us to isolate the contribution of dynamically modeling the precision parameter. The generalized BPARMA model provides a comparison with an existing beta prime-based forecasting approach and therefore allows us to assess the gains obtained by incorporating dynamic precision and a richer seasonal structure. The Gaussian SARIMA and Holt–Winters methods were included for a different purpose. Both are widely used forecasting tools in applied time-series analysis and are frequently employed through automated specification and forecasting procedures. Our objective is therefore not to compare the proposed model with fully customized or manually enhanced versions of all competing approaches (for example, by augmenting SARIMA specifications with additional exogenous seasonal regressors), but rather to evaluate whether the additional flexibility of the generalized BPSARMA framework translates into improved predictive performance relative to forecasting methods routinely adopted in practice. Taken together, these comparisons allow us to assess the forecasting performance of the proposed model from two complementary perspectives: first, relative to closely related beta prime-based alternatives under comparable modeling structures, and second, relative to widely used benchmark forecasting methodologies, as they are commonly implemented in practice.

The comparison of the models in the context of out-of-sample forecasting was carried out using standard predictive accuracy measures, namely the relative root mean squared error (RRMSE) and the mean absolute relative error (MARE). These measures are defined as

$$\text{RRMSE} = \sqrt{\frac{1}{h}\sum_{j=1}^{h}\left(\frac{Y_{n+j}-\hat{\mu}_{n+j}}{Y_{n+j}}\right)^2} \quad \text{and} \quad \text{MARE} = \frac{1}{h}\sum_{j=1}^{h}\left|\frac{Y_{n+j}-\hat{\mu}_{n+j}}{Y_{n+j}}\right|,$$

for $h \in \{1,2,\ldots\}$, where $\hat{\mu}_{n+j}$ denotes the j-step-ahead forecast of Y_{n+j}.

6.1. Manso Water Reservoir Inflow

The Manso Hydroelectric Power Plant, located on the Manso River in the state of Mato Grosso, was built through a partnership between the private sector, represented by the PROMAN consortium—formed by the companies Odebrecht, Servix, and Pesa—and FURNAS, a company administered by Eletronorte (Centrais Elétricas do Norte do Brasil S/A). With an installed capacity of 210 MW, the reservoir covers an area of approximately 427 km^2, extending across the municipalities of Chapada dos Guimarães and Nova Brasilândia. In addition to its role in electricity generation, the reservoir plays an important role in regulating the flood and drought regimes of the Cuiabá River, thereby contributing to the mitigation of socioeconomic impacts in the region.

Monthly inflow data for the Manso reservoir from January 2001 to October 2025 were considered, yielding an initial sample of $n = 298$ observations. For the purpose of evaluating predictive performance, the last six observations were reserved for out-of-sample validation. Consequently, the sample effectively used for parameter estimation consists of $n = 292$ observations.

Table 6 presents descriptive statistics for the time series. The average monthly inflow to the Manso reservoir during the period under analysis is $154.40\,\mathrm{m}^3\,\mathrm{s}^{-1}$. The smallest and largest recorded values ($24.41\,\mathrm{m}^3\,\mathrm{s}^{-1}$ and $670.21\,\mathrm{m}^3\,\mathrm{s}^{-1}$) occurred in August 2024 and February 2024, respectively.

Table 6. Descriptive statistics for the Manso reservoir inflow series.

Min.	Max.	Median	Mean	Std. Dev.	Skewness	Kurtosis
24.41	670.21	99.61	154.40	128.97	1.27	3.79

In order to improve the representation of these seasonal fluctuations, we included an indicator variable for the months from January to April. Specifically, x_{3t} is defined to take the value 1 for these months and 0 otherwise. The inclusion of this regressor led to noticeable improvements in model fit.

In addition, the series appears to exhibit a possible change in its mean level over time. Considering data up to 2014, the average monthly inflow is $170.16\,\mathrm{m}^3\,\mathrm{s}^{-1}$, whereas for the period beginning in 2015, the mean decreases to $133.05\,\mathrm{m}^3\,\mathrm{s}^{-1}$. To account for this structural change, we introduced a second indicator variable, denoted by x_{4t}, defined as 1 for the period from January 2001 to December 2014, and 0 for the subsequent period, that is, from January 2015 to the end of the series.

The selected specification was $\mathrm{BPSARMA}_\mu(1,0) \times (2,0)_{12} + \mathrm{AR}_\phi(4)$, including harmonic regressors $x_{1t} = \sin(2\pi t/12)$ and $x_{2t} = \cos(2\pi t/12)$, as well as the two dummy variables described above, x_{3t} and x_{4t}. However, the z tests did not reject the null hypotheses for Φ_2, δ_3, and δ_4 at the 5% level. The corresponding terms were then dropped from the model.

The inclusion of harmonic regressors and seasonal dummy variables was motivated by empirical evidence obtained during the model-building process. It is important to note that the baseline specification already incorporates seasonal dynamics through the seasonal autoregressive and seasonal moving-average dynamics of the generalized BPSARMA model. Therefore, the purpose of introducing harmonic regressors and seasonal dummy variables was not to create seasonality in the model, but rather to determine whether additional deterministic seasonal information could further improve model fit and forecasting performance.

To assess their contribution, we compared three nested specifications sharing the same dynamic structure. Model 1 included only the seasonal dynamic component. Model 2 augmented this specification with harmonic regressors, while Model 3 additionally included seasonal dummy variables. In each case, non-significant terms were sequentially removed at the 5% significance level, and the model was subsequently refitted. The MBIC values of the resulting final models were 2986.9196, 2875.0054, and 2849.7574, respectively. The correlation between the observed and fitted values increased from 0.7544 in Model 1 to 0.8368 in Model 2 and to 0.8580 in Model 3. These results indicate that the seasonal dynamic structure alone does not fully capture the seasonal behavior of the series. Harmonic regressors provide an improvement by representing smooth deterministic seasonal fluctuations, whereas seasonal dummy variables yield an additional gain by accommodating more localized seasonal effects. Furthermore, the specification combining all three mechanisms produced fitted values that more accurately reproduced the observed seasonal peaks. These findings support the view that seasonal dynamic components, harmonic regressors, and seasonal dummy variables act as complementary rather than competing tools for modeling seasonality.

As previously explained, at the model selection stage, we considered all combinations with $p_1, q \in \{0, \ldots, 4\}$, $P, Q \in \{0, 1, 2\}$, and $p_2 \in \{1, \ldots, 4\}$, yielding a total of 800 candidate specifications. Model comparison was performed using the MBIC criterion. The five specifications with the lowest MBIC values were $\text{BPSARMA}_\mu(1,0) \times (2,0)_{12} + \text{AR}_\phi(4)$ (MBIC = 2832.7263), $\text{BPSARMA}_\mu(1,2) \times (0,2)_{12} + \text{AR}_\phi(4)$ (MBIC = 2832.9054), $\text{BPSARMA}_\mu(1,2) \times (0,2)_{12} + \text{AR}_\phi(3)$ (MBIC = 2833.9145), $\text{BPSARMA}_\mu(2,0) \times (2,0)_{12} + \text{AR}_\phi(4)$ (MBIC = 2834.0015), and $\text{BPSARMA}_\mu(4,0) \times (2,0)_{12} + \text{AR}_\phi(4)$ (MBIC = 2834.1581). The MBIC values of the best-performing specifications were very close, indicating that several model structures provide a comparable fit to the data. Nevertheless, the $\text{BPSARMA}_\mu(1,0) \times (2,0)_{12} + \text{AR}_\phi(4)$ model achieved the smallest MBIC value and was therefore selected as the initial specification. Subsequently, statistical significance was assessed through z tests. The seasonal parameter Φ_2 was removed due to lack of significance (p-value = 0.4034). After refitting the model, the parameters δ_4 and δ_3 were sequentially excluded because their corresponding p-values were 0.0705 and 0.0821, respectively. Following this sequential elimination procedure, all remaining parameters were statistically significant at the 5% level, yielding the final model reported below.

Table 7 reports the parameter estimates, their corresponding standard errors, z statistics, and p-values, together with selected diagnostic measures for the fitted model. The number of lags (v) used in the Portmanteau tests was set equal to twice the seasonal period. The null hypothesis of no residual autocorrelation is not rejected by either the Ljung–Box or the Monti tests at conventional significance levels. In addition, the Jarque–Bera test for normality applied to the quantile residuals provides no evidence against normality at the 10% significance level. Overall, these results suggest that there is no indication of misspecification in the fitted model.

Figure 2 presents diagnostic plots for the fitted generalized BPSARMA model. The upper panels display the residual correlogram and partial correlogram. All residual autocorrelations and partial autocorrelations lie within the 95% asymptotic confidence bands (horizontal dashed lines), indicating the absence of serial correlation in the residuals and corroborating the conclusions drawn from the Portmanteau tests. These graphical findings are in agreement with the Ljung–Box and Monti test results reported in Table 7. Since both tests are applied to the quantile residuals, the large p-values provide no evidence against the null hypothesis of residual serial independence.

The bottom-left panel shows the residuals plotted over time. The residuals fluctuate randomly around zero, with no evidence of systematic patterns. The bottom-right panel presents the quantile–quantile (Q–Q) plot, where the empirical quantiles of the residuals are compared with the theoretical quantiles of the normal distribution. The close alignment of the points with the reference line suggests that the normal approximation is adequate. These graphical findings are consistent with the results of the Jarque–Bera normality test.

Table 7. Parameter estimates and diagnostic statistics for the fitted generalized BPSARMA model; Manso reservoir inflow.

Parameter	Estimate	Std. Error	z Statistic	p-Value
α_1	2.0796	0.0513	40.5231	<0.0001
φ_1	0.4816	0.0390	12.3361	<0.0001
Φ_1	0.1160	0.0472	2.4603	0.0139
β_1	0.5348	0.0585	9.1405	<0.0001
β_2	0.7472	0.0599	12.4750	<0.0001
β_3	0.3259	0.0599	5.4373	<0.0001
β_4	0.2301	0.1000	2.3020	0.0213
α_2	2.8001	0.1958	14.3017	<0.0001
δ_1	−0.8461	0.1995	−4.2419	<0.0001
δ_2	−0.8900	0.2545	−3.4967	0.0005

MBIC = 2849.7575
Ljung–Box statistic ($v = 24$): $Q_{LB} = 16.1940$ (p-value = 0.7045)
Monti statistic ($v = 24$): $Q_M = 18.9020$ (p-value = 0.5282)
Jarque–Bera statistic: $JB = 3.4234$ (p-value = 0.1806)

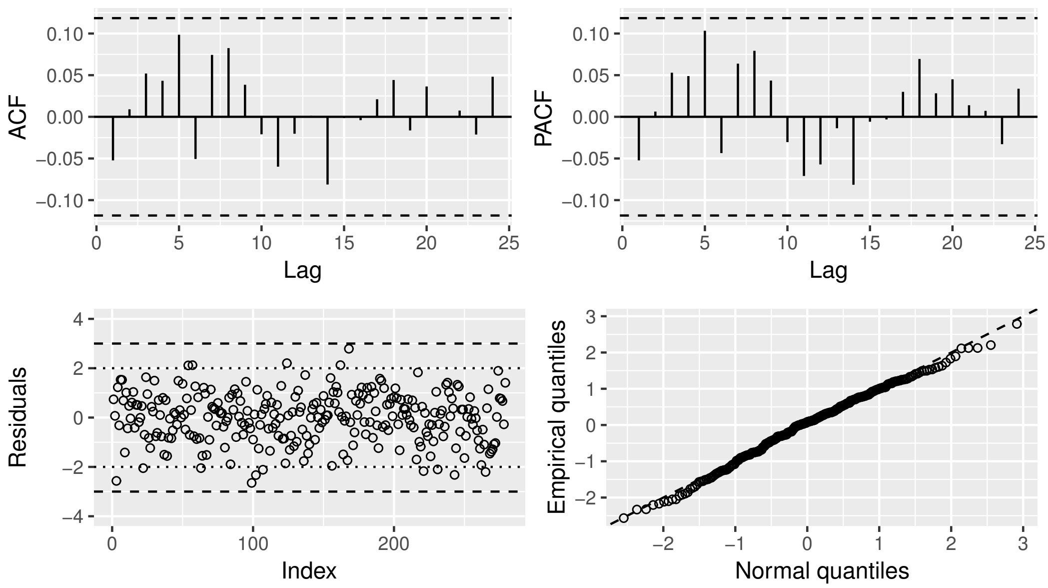

Figure 2. Residual autocorrelation function (**top-left**), residual partial autocorrelation function (**top-right**), residuals over time (**bottom-left**), and normal Q–Q plot of the residuals (**bottom-right**); fitted generalized BPSARMA model for the Manso reservoir inflow series.

Overall, both the graphical diagnostics and the statistical tests indicate that the fitted model provides an adequate description of the data and is therefore suitable for forecasting purposes. Figure 3 displays the observed time series (gray line) together with the fitted values obtained from the generalized BPSARMA model (black line). The fitted values closely track the dynamics of the observed series, indicating good agreement between observed and model-implied values. For completeness, we also evaluated Pearson-type standardized residuals. The resulting residual ACF and PACF plots showed no significant autocorrelations, and the corresponding Ljung–Box and Monti tests led to the same conclusions as those obtained from the quantile residuals. The results are therefore omitted for brevity.

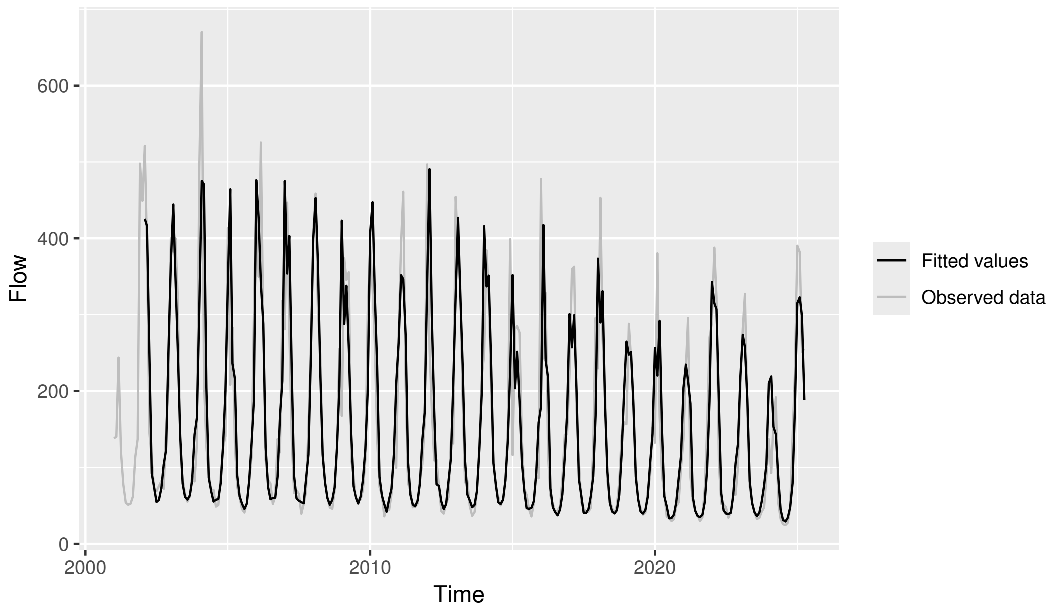

Figure 3. Observed time series (gray line) and fitted values from the generalized BPSARMA model (black line); Manso reservoir inflow.

After establishing that the generalized BPSARMA model provides a satisfactory fit to the time series, we proceed with an initial comparison with the BPSARMA model with fixed precision, that is, the specification obtained by imposing $\delta = \mathbf{0}$, where $\mathbf{0}$ denotes the p_2-dimensional null vector, and by adopting the identity link function for the precision submodel.

Following the model-building procedure described earlier, an initial specification was selected based on the MBIC criterion. This yielded the model $\text{BPSARMA}_\mu(2,0) \times (1,0)_{12}$ with regressors x_{1t}, x_{2t}, x_{3t}, and x_{4t}. Subsequently, individual z tests were used to assess the statistical significance of the regression coefficients. The parameters β_3 and β_4, associated with x_{3t} and x_{4t}, respectively, were removed because their corresponding p-values exceeded the 5% significance level. Thus, the final fixed-precision specification retained only the statistically significant regressors.

Once the fixed-precision benchmark model had been defined, we assessed whether allowing the precision parameter to vary dynamically led to a statistically significant improvement in model fit. For this purpose, we performed a likelihood ratio test comparing the fixed-precision model and the generalized model with dynamic precision. The null hypothesis was $\mathcal{H}_0\colon \delta = \mathbf{0}$, corresponding to constant precision, against the alternative $\mathcal{H}_1\colon \delta \neq \mathbf{0}$, corresponding to time-varying precision. The resulting p-value was smaller than 0.0001, providing strong evidence against the fixed-precision assumption and indicating that modeling the precision dynamically yields a significantly better fit.

In addition, the models were compared using information criteria. The generalized BPSARMA model yielded MAIC (MBIC) values of 2802.7180 (2843.1620), whereas the fixed-precision model produced MAIC (MBIC) values of 2898.9717 (2924.7089). Both criteria therefore favor the model with time-varying precision. Overall, the results provide clear evidence in favor of a specification with dynamic precision.

The inclusion of the variables z_{t-k} in the precision submodel aims to allow the conditional precision to adapt to the recent relative dynamics of the series. Unlike the previous formulation, which used only ratios of consecutive observations, the present specification is based on relative increments, $z_{t-k} = (Y_{t-k} - Y_{t-k-1})/Y_{t-k-1}$, thereby distinguishing not only the magnitude but also the direction of recent changes in the process. This provides greater interpretability and flexibility, since increases and decreases in the series may have asymmetric effects on the local variability. In practical terms, this specification allows the model to capture situations in which abrupt increases or decreases in the observed process are associated with distinct levels of uncertainty, leading to a more responsive dynamic representation of the precision parameter.

With respect to the variables z_{t-k}, $k \in 1, 2$, three situations can be distinguished: (i) $z_{t-k} < 0$, which occurs when $y_{t-k} < y_{t-k-1}$; (ii) $z_{t-k} > 0$, that is, when $y_{t-k} > y_{t-k-1}$; and (iii) $z_{t-k} = 0$, in which case $y_{t-k} = y_{t-k-1}$. In light of the estimates of α_2, δ_1, and δ_2 reported in Table 7, it follows that in situation (i), the estimated precision tends to increase, whereas in situation (ii), it tends to decrease. In other words, when the value of the series declines (for instance, due to seasonal fluctuations), the estimated precision tends to increase; conversely, when the series rises, the precision tends to decrease.

To illustrate this behavior, consider observation $t = 258$, for which the estimated precision is $\hat{\phi}_t = 45.7621$. It should be noted that the precision at time t depends on past information, in this case on the observed values at times 255, 256, and 257, which are 323.0406, 104.2377, and 52.4281, respectively. Thus, in a decreasing segment of the series, a relatively large value of the estimated precision is obtained. In contrast, consider observation $t = 182$, for which $\hat{\phi}_t = 0.3582$. The observed values of the series at times 179, 180, and 181 are 90.1160, 85.7374, and 477.9087, respectively. In this case, characterized by an increasing pattern in the series, the estimated precision is substantially smaller.

The temporal evolution of the estimated precision is displayed in Figure 4. The pattern is approximately the inverse of that observed in Figure 3, reinforcing the finding that the precision tends to increase during decreasing segments of the series, whereas it tends to decline during increasing segments. In Figure 4, the dashed horizontal line represents the fixed precision estimate from the BPSARMA model with constant precision, equal to 9.2040, whereas the mean precision obtained from the generalized model is 17.0884. Moreover, approximately 68% of the precision estimates produced by the generalized model exceed the fixed precision estimate.

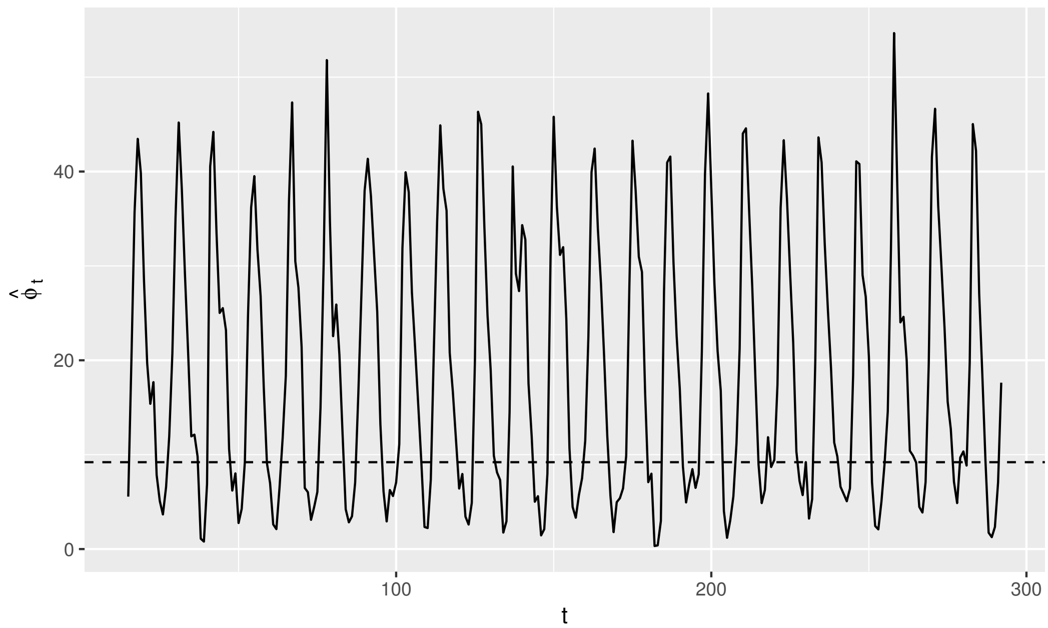

Figure 4. Estimated conditional precisions from the generalized BPSARMA model (solid line) and fixed precision estimate from the constant-precision BPSARMA model (dashed horizontal line); Manso reservoir inflow.

We now turn to the evaluation of out-of-sample forecasts. In addition to the generalized BPSARMA model and the BPSARMA model with fixed precision, we also consider the generalized BPARMA model with harmonic regressors and dummy variables. Furthermore, we include as benchmark methods the Gaussian SARIMA model and the Holt–Winters exponential smoothing algorithm with additive seasonality. The generalized BPARMA model selected with harmonic regressors and dummy variables was the BPARMA$(4,0)$ model. However, the autoregressive terms of orders 2 and 3 were not statistically significant at the 5% level and were consequently removed from the final specification. The selected SARIMA specification was the $(1,0,1) \times (1,1,1)_{12}$ model.

Table 8 reports the forecast accuracy measures for prediction horizons ranging from one to six steps ahead, with the best results highlighted in bold. Overall, the generalized BPSARMA model consistently outperforms the competing models according to both accuracy measures across all considered horizons. To illustrate this superior performance, consider the case $h = 6$. The forecasts generated by the generalized BPSARMA model yield an RRMSE of 0.3007, which is lower than that of the competing approaches. The second-best performance is obtained by the generalized BPARMA model, with an RRMSE of 0.4052, indicating a noticeable loss in predictive accuracy when compared to the proposed model.

Table 8. Out-of-sample forecast accuracy measures for different forecasting horizons; Manso reservoir inflow.

	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$
			RRMSE			
Generalized BPSARMA	**0.0642**	**0.0456**	**0.0761**	**0.0999**	**0.1564**	**0.3007**
Generalized BPSARMA fixed prec.	0.2899	0.2803	0.2389	0.3022	0.3815	0.5550
Generalized BPARMA	0.1369	0.1087	0.0964	0.1598	0.2455	0.4052
SARIMA	0.1348	0.2574	0.2904	0.3864	0.4322	0.5520
Holt–Winters additive	0.1020	0.1923	0.1935	0.3023	0.3421	0.4809
Holt-Winters multiplicative	0.0875	0.1535	0.2067	0.3137	0.4543	0.5909
			MARE			
Generalized BPSARMA	**0.0642**	**0.0354**	**0.0619**	**0.0839**	**0.1246**	**0.2119**
Generalized BPSARMA fixed prec.	0.2899	0.2801	0.2263	0.2799	0.3443	0.4634
Generalized BPARMA	0.1369	0.1034	0.0906	0.1361	0.1981	0.3029
SARIMA	0.1348	0.2364	0.2734	0.3517	0.3974	0.4888
Holt–Winters additive	0.1020	0.1770	0.1833	0.2633	0.3043	0.4029
Holt-Winters multiplicative	0.0875	0.1431	0.1903	0.2715	0.3770	0.4860

As an additional sensitivity analysis, we fitted a SARIMAX specification to the Manso inflow series using the same harmonic regressors and seasonal dummy variables adopted in the beta prime-based models. Although the additional regressors were individually significant at the 5% level according to z tests, the resulting SARIMAX model did not yield empirical gains over the benchmark SARIMA specification. In particular, the SARIMA model achieved lower information criteria (AIC $=$ 3194.5660; BIC $=$ 3212.7399) than the SARIMAX model (AIC $=$ 3322.5551; BIC $=$ 3355.6459), and slightly higher in-sample correlation between observed and fitted values (0.85 versus 0.84). In the out-of-sample comparison, the SARIMAX forecasts were more accurate than those of SARIMA only for the horizons $h = 1$ and $h = 2$, while SARIMA outperformed SARIMAX for the remaining forecasting horizons according to both RRMSE and MARE. These results suggest that the benchmark SARIMA specification used in this study provides a reasonable practical reference.

To further assess forecasting performance, we conducted a rolling-window evaluation. The model structures selected in the full-sample analysis were kept fixed throughout this exercise, and only the model parameters were re-estimated at each forecasting origin. Starting with the first 263 observations, each model was estimated and used to generate a one-step-ahead forecast for the next observation. The estimation sample was then expanded by one observation, and the procedure was repeated until the first 298 observations had been used for estimation. This resulted in 36 pseudo-out-of-sample one-step-ahead forecasts for each competing model, corresponding to the last three seasonal cycles of the series. For brevity, the Holt–Winters method was not included in this supplementary rolling-window analysis.

Rather than focusing on a particular forecast accuracy measure, we recorded, at each forecasting origin, which model produced the smallest absolute prediction error. The generalized BPSARMA model produced the most accurate one-step-ahead forecast in 75% of the comparisons against the BPSARMA model with fixed precision, in 72% of the comparisons against the generalized BPARMA model, and in 83% of the comparisons against the SARIMA model. When all four competing models were considered simultaneously, the generalized BPSARMA model yielded the smallest absolute prediction error in 55% of the forecasting exercises.

These results provide additional evidence that the forecasting gains observed in the six-step-ahead out-of-sample evaluation reported previously, which was based on forecasts generated from the end of the series, are not restricted to a single forecasting origin. However, this rolling-window exercise is intended only as a supplementary robustness check and has a limited scope, as it is restricted to a single empirical application and one-step-ahead forecasts. Therefore, it should not be viewed as a replacement for a systematic rolling-window comparison across all applications and forecast horizons. The main forecasting assessment in this paper remains the multi-step holdout evaluation conducted for all empirical series.

For the sake of brevity and to avoid repetition, the remaining applications are presented more concisely, emphasizing the selected specifications and the six-step-ahead out-of-sample forecasting performance based on the holdout observations at the end of each series.

6.2. Camargos Water Reservoir Inflow

The Camargos Hydroelectric Power Plant is located in the municipality of Itutinga, in the state of Minas Gerais, Brazil, on the Grande River. The plant began operations in 1960 and was the first facility integrated into the Furnas system (Furnas Centrais Elétricas S/A). Its installed capacity is 45 MW. The reservoir covers an area of $73.35\,\mathrm{km}^2$ and has a total storage capacity of 792 million cubic meters of water. In addition to electricity generation, the Camargos plant plays an important role in regulating river flows, supporting the operation of other facilities within the Furnas Hydroelectric Complex. It also contributes to the regulation of the water level of Furnas Lake, helping maintain the operational stability of the Brazilian National Interconnected Power System and enhancing the reliability of electricity supply.

The Camargos reservoir data series used in this study begins in January 2000 and comprises $n = 310$ observations. The last six observations were reserved for forecast evaluation, leaving $n = 304$ observations for model estimation. The average monthly inflow is $104.85\,\mathrm{m}^3\,\mathrm{s}^{-1}$. The minimum value occurred in October 2014, at $22.42\,\mathrm{m}^3\,\mathrm{s}^{-1}$, whereas the maximum was recorded in January 2011, reaching $522.81\,\mathrm{m}^3\,\mathrm{s}^{-1}$. Higher inflows typically occur between December and March, reflecting the seasonal pattern of the series. In 2014, the region experienced a severe drought, during which the lowest value in the series was recorded; see [30]. To account for this exogenous shock, a dummy variable was included in the model specification, denoted by x_{3t}, defined as $x_{3t} = 1$ for observations in 2014 and $x_{3t} = 0$ otherwise. Additionally, a seasonal dummy variable, denoted by x_{4t}, was included to capture the period of lower monthly inflow: specifically, $x_{4t} = 1$ for observations corresponding to the months from May to September, and $x_{4t} = 0$ otherwise.

The selected model was BPSARMA$\mu(2,2) \times (1,1)_{12} + \mathrm{AR}\phi(2)$, which incorporates harmonic regressors to capture seasonal patterns. However, no clear seasonal behavior was observed during 2014, a period associated with a severe drought episode. To accommodate this feature, the harmonic terms were interacted with the indicator function $(1 - x_{3t})$, allowing the seasonal component represented by the harmonic regressors to be effectively switched off during the drought period. This specification yields the regressors $x_{1t}(1 - x_{3t})$ and $x_{2t}(1 - x_{3t})$. In addition, the model includes the dummy variables x_{3t} and x_{4t} to account for structural changes related to the drought period and other exogenous influences. The first-order nonseasonal autoregressive term was removed from the model due to lack of statistical significance at the 5% level. As before, model adequacy was assessed using the Ljung–Box and Monti tests. The p-values were 0.5967 and 0.6366, respectively. Hence, there is no evidence of serial residual correlation. The distributional assumption of the residuals was evaluated using the Jarque–Bera test based on quantile residuals, which resulted in a p-value of 0.4162. We thus conclude that there is no evidence against the normality of the residuals and, consequently, of model misspecification.

For the evaluation of out-of-sample forecasts, we also considered a BPSARMA specification with fixed precision. The selected specification for the latter was BPSARMA$_\mu(1,1) \times (1,1)_{12}$, incorporating the regressors x_{1t}, x_{2t}, and x_{3t}. Additionally, a generalized BPARMA model including harmonic regressors and the dummy variable was fitted, for which the specification BPARMA$(2,0)$ was selected. We also considered a SARIMA model of order $(0,0,1) \times (0,1,1)_{12}$ as a benchmark.

Table 9 reports the forecast accuracy measures for prediction horizons ranging from one to six steps ahead. According to the RRMSE criterion, the varying-precision BPSARMA model exhibits superior predictive performance for the horizons $h = 1, 3, 4$, and 5. In contrast, for $h = 2$ and 6, the Holt–Winters multiplicative method outperforms the competing models. When the evaluation is based on the MARE metric, the proposed model provides the best performance for horizons from $h = 1$ to 5. For $h = 6$, the Holt–Winters multiplicative method achieves the lowest MARE, followed by the varying-precision BPSARMA model. To illustrate these results, consider the case $h = 3$. The RRMSE associated with the forecasts produced by the proposed model is 0.0482. The second- and third-best performances are obtained by the Holt–Winters multiplicative method and the SARIMA model, with RRMSE values of 0.0767 and 0.1386, respectively. These values exceed that of the proposed model by more than 50%, highlighting its superior predictive accuracy.

Table 9. Out-of-sample forecast accuracy measures for different forecasting horizons; Camargos reservoir inflow.

	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$
			RRMSE			
Generalized BPSARMA	**0.0065**	0.0426	**0.0482**	**0.0978**	**0.1028**	0.3448
Generalized BPSARMA fixed prec.	0.8901	0.9153	0.9266	0.9324	0.9352	0.9367
Generalized BPARMA	0.2285	0.1617	0.1497	0.2099	0.2965	0.5401
SARIMA	0.2066	0.1466	0.1386	0.1587	0.1563	0.3187
Holt–Winters additive	0.1571	0.2320	0.2629	0.2885	0.3069	0.2803
Holt–Winters multiplicative	0.0440	**0.0415**	0.0767	0.1000	0.1082	**0.2300**
			MARE			
Generalized BPSARMA	**0.0065**	**0.0332**	**0.0414**	**0.0753**	**0.0844**	0.2058
Generalized BPSARMA fixed prec.	0.8901	0.9149	0.9262	0.9321	0.9349	0.9364
Generalized BPARMA	0.2285	0.1182	0.1195	0.1722	0.2404	0.3911
SARIMA	0.2066	0.1124	0.1153	0.1383	0.1399	0.2329
Holt–Winters additive	0.1571	0.2226	0.2536	0.2788	0.2973	0.2519
Holt–Winters multiplicative	0.0440	0.0414	0.0673	0.0879	0.0976	**0.1661**

6.3. Estreito Water Reservoir Inflow

The Estreito Hydroelectric Power Plant is located on the Tocantins River, on the border between the states of Maranhão and Tocantins, encompassing the municipalities of Estreito (MA), Aguiarnópolis (TO), and Palmeiras do Tocantins (TO). The plant is currently operated by the Estreito Energia Consortium (CESTE), formed by the companies Engie, Vale, Alcoa, and InterCement. It has an installed capacity of 1087 MW, a reservoir area of 555 km^2, and a normal upstream water level of 156 m.

For model estimation, the effective sample covers the period from January 2002 to April 2025. The interval from May to October 2025 was reserved for the evaluation of out-of-sample forecasts. The average monthly inflow during the observed period was 3238.79 m^3 s^{-1}. The series exhibits a range of 13,094.86 m^3 s^{-1}, with a minimum value of 762.61 m^3 s^{-1} and a maximum value of 13,857.47 m^3 s^{-1}, recorded in September 2016 and January 2021, respectively. The series displays a well-defined seasonal pattern, with higher monthly inflow values typically occurring from January to April. To better capture these seasonal fluctuations, a dummy variable for these months was introduced, denoted by x_{3t}.

The model initially selected was the BPSARMA$_\mu(2,0) \times (2,1)_{12} + \text{AR}_\phi(3)$ specification, incorporating the two harmonic regressors x_{1t} and x_{2t}, as well as the dummy variable x_{3t}. However, the dynamic terms associated with the parameters φ_2 and Φ_2 were removed from the model because they were not statistically significant at the 5% level.

The Ljung–Box and Monti tests yielded p-values of 0.2583 and 0.1337, respectively, providing no evidence of serial autocorrelation in the residuals. The normality of the residuals was evaluated using the Jarque–Bera test, which produced a p-value of 0.7158. Taken together, these results indicate no evidence of misspecification in the fitted model.

For the evaluation of out-of-sample forecasts, alternative models were considered in addition to the proposed specification. First, a BPSARMA model with fixed precision was fitted, for which the selected specification was BPSARMA$_\mu(1,2) \times (1,2)_{12}$, also incorporating harmonic regressors and the dummy variable. In addition, a generalized BPARMA model of order $(1,4)$ was estimated, likewise including these regressors. The SARIMA model selected for comparison was $(1,0,0) \times (2,1,0)_{12}$.

Table 10 reports the forecast accuracy measures for prediction horizons ranging from one to six steps ahead. According to all the metrics considered, the BPSARMA model with varying precision exhibits superior predictive performance for the horizons $h = 1,2,5$, and 6. For the horizons $h = 3$ and 4, the additive Holt–Winters method provides the most accurate forecasts, followed by the BPSARMA model with varying precision.

Table 10. Out-of-sample forecast accuracy measures for different forecasting horizons; Estreito reservoir inflow.

	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$
			RRMSE			
Generalized BPSARMA	**0.0589**	**0.0425**	0.0521	0.0765	**0.0690**	**0.0671**
Generalized BPSARMA fixed prec.	0.1132	0.1889	0.1979	0.1997	0.1794	0.1743
Generalized BPARMA	0.0869	0.0647	0.0865	0.1150	0.1104	0.1050
SARIMA	0.0950	0.0754	0.0868	0.1054	0.0944	0.1159
Holt–Winters additive	0.0650	0.0472	**0.0431**	**0.0384**	0.0839	0.0824
Holt–Winters multiplicative	0.1052	0.1346	0.1100	0.1580	0.1414	0.1302
			MARE			
Generalized BPSARMA	**0.0589**	**0.0354**	0.0460	0.0654	**0.0562**	**0.0563**
Generalized BPSARMA fixed prec.	0.1132	0.1776	0.1900	0.1938	0.1623	0.1596
Generalized BPARMA	0.0869	0.0579	0.0781	0.1022	0.0997	0.0951
SARIMA	0.0950	0.0717	0.0832	0.0993	0.0813	0.0994
Holt–Winters additive	0.0650	0.0400	**0.0378**	**0.0329**	0.0605	0.0628
Holt–Winters multiplicative	0.1052	0.1319	0.0897	0.1304	0.1060	0.0952

6.4. Sobradinho Water Reservoir Inflow

The Sobradinho reservoir is located in the state of Bahia, on the São Francisco River. It extends for approximately 320 km and has a water surface area of about 4214 km^2. Its storage capacity reaches 34.1 billion cubic meters when the water level is at 392.5 m, making it one of the largest artificial lakes in the world. The hydroelectric plant associated with the reservoir has an installed capacity of 1050.3 MW. The reservoir is currently operated by Companhia Hidro Elétrica do São Francisco (CHESF).

The time series of inflow to the Sobradinho reservoir spans the period from January 2000 to October 2025. For modeling purposes, the effective sample considered extends up to April 2025, totaling $n = 304$ monthly observations. The mean inflow during the analyzed period was 1585.87 $\text{m}^3\,\text{s}^{-1}$, with a minimum value of 234.19 $\text{m}^3\,\text{s}^{-1}$, observed in October 2014, and a maximum value of 6352.29 $\text{m}^3\,\text{s}^{-1}$, recorded in March 2007.

An exploratory inspection of the series suggests a possible level shift between 2013 and 2020, reflected in changes in both the mean level and the variability of the inflow. In particular, during this interval, there is a noticeable reduction in the magnitude of peak inflows, along with a greater concentration of observations at lower levels compared with the preceding and subsequent periods. This behavior may be associated with hydrological, climatic, or operational factors that temporarily affected the dynamics of the series. To account for this feature, the indicator variable x_{3t} was included in the model specification. This variable takes the value 1 for the period from 2013 to 2020 and 0 otherwise, allowing the model to explicitly capture this temporary effect.

We selected the model $\text{BPSARMA}_\mu(1,0) \times (1,1)_{12} + \text{AR}_\phi(3)$, which includes harmonic regressors (x_{1t} and x_{2t}) together with the indicator variable x_{3t}. The p-values of the Ljung–Box and Monti tests were 0.5565 and 0.2580, respectively, providing no statistical evidence of residual serial autocorrelation.

Regarding the normality of the residuals, the Jarque–Bera test rejected the null hypothesis of normality at the 5% significance level. An inspection of the quantile residuals revealed that the departure from normality is primarily driven by three observations with unusually large absolute residuals. These residuals are associated with abrupt and atypical fluctuations in the inflow series rather than with systematic patterns affecting a broader portion of the sample. In particular, the two largest absolute residuals occurred during periods in which the series would typically exhibit increasing inflow levels; however, in October 2014 and December 2023, the inflow dropped by more than 30% relative to the previous month. Likewise, the third largest residual occurred in May 2024, when the inflow decreased by more than 70% compared with the previous month. Such abrupt changes are highly unusual in the context of the observed hydrological dynamics. Importantly, despite this departure from normality, the residual ACF and PACF plots do not indicate serial dependence, and both the Ljung–Box and Monti tests fail to reject the null hypothesis of residual independence. Therefore, the available diagnostic evidence suggests that the fitted model adequately captures the dynamic dependence structure of the series, although a few extreme observations generate heavier tails than those expected under the assumed model.

As before, competing models were also considered. First, a BPSARMA model with a constant precision parameter was fitted. The selected specification was $\text{BPSARMA}_\mu(1,2) \times (1,2)_{12}$, including the harmonic regressors x_{1t} and x_{2t} as well as the seasonal dummy variable. In addition, a generalized BPARMA model incorporating the same harmonic regressors and dummy variable was estimated, yielding the specification $\text{BPARMA}(2,3)$. As a classical time series benchmark, a SARIMA model of order $(2,0,0) \times (0,1,1)_{12}$ was also fitted.

Table 11 reports the forecast accuracy measures for horizons ranging from one to six steps ahead. According to both the RRMSE and MARE metrics, the proposed model yields more accurate forecasts than the competing models across all forecast horizons considered.

Table 11. Out-of-sample forecast accuracy measures for different forecasting horizons; Sobradinho reservoir inflow.

	$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$
			RRMSE			
Generalized BPSARMA	**0.0905**	**0.0685**	**0.0565**	**0.0982**	**0.1182**	**0.1969**
Generalized BPSARMA fixed prec.	0.0952	0.0940	0.0853	0.1449	0.1934	0.3072
Generalized BPARMA	0.3639	0.3376	0.2876	0.2589	0.2548	0.3133
SARIMA	0.2032	0.1441	0.1394	0.2074	0.2202	0.2463
Holt–Winters additive	1.2575	1.6198	1.8455	2.0193	2.1079	2.1265
Holt–Winters multiplicative	0.3243	0.3247	0.3425	0.3371	0.3461	0.3404
			MARE			
Generalized BPSARMA	**0.0905**	**0.0625**	**0.0463**	**0.0773**	**0.0972**	**0.1482**
Generalized BPSARMA fixed prec.	0.0952	0.0940	0.0841	0.1254	0.1645	0.2397
Generalized BPARMA	0.3639	0.3365	0.2716	0.2391	0.2388	0.2847
SARIMA	0.2032	0.1095	0.1161	0.1714	0.1902	0.2166
Holt–Winters additive	1.2575	1.5861	1.8006	1.9675	2.0600	2.0863
Holt–Winters multiplicative	0.3243	0.3247	0.3417	0.3364	0.3451	0.3392

6.5. Itaipu Water Reservoir Inflow

The Itaipu reservoir extends for approximately 170 km along the Paraná River, located in the state of Paraná, Brazil. It has an inundated area of about 1350 km^2 and a storage capacity of approximately 29 billion cubic meters. Its energy production index is estimated at 10.4 MW km^{-2}. Consequently, the Itaipu reservoir constitutes a strategic asset not only for electricity generation but also for maintaining ecological balance and supporting the economic development of the surrounding region. In this context, continuous monitoring of the reservoir is essential to ensure both the structural integrity of the dam and the environmental preservation of its surroundings.

The inflow series for the Itaipu reservoir was observed from January 2000 onward, yielding an effective sample of $n = 304$ observations. The mean monthly inflow was 10,405.97 m^3 s^{-1}. The smallest value recorded was 4688.48 m^3 s^{-1} in July 2021, whereas the largest occurred in January 2010, reaching 22,356.95 m^3 s^{-1}.

Visible structural changes can be identified in the behavior of the inflow over time, indicating the presence of distinct regimes. In particular, the initial portion of the sample is characterized by a higher mean level and a greater occurrence of extreme peaks, whereas later periods exhibit a reduction in the mean level together with changes in variability. To account for these features, we include two indicator variables in the model.

The first variable, x_{3t}, is defined as $x_{3t} = 1$ during the first 19 years of the sample (2000–2018) and $x_{3t} = 0$ thereafter, with the aim of capturing the initial regime characterized by higher mean inflow levels. The second variable, x_{4t}, is defined as $x_{4t} = 1$ from 2023 onward and $x_{4t} = 0$ in the preceding periods in order to capture a possible more recent regime associated with additional changes in both the level and the variability of the series. The inclusion of these variables allows regime shifts to be modeled explicitly. Moreover, these dummies contribute to improving model fit and predictive accuracy by accommodating long-run variations that are not captured solely by the autoregressive, seasonal, and harmonic components. When generating out-of-sample forecasts, we set

$x_{3t} = 0$ and $x_{4t} = 1$, thereby assuming that the most recent regime of the series persists beyond the sample period.

We selected the model $\text{BPSARMA}_{\mu}(1,0) \times (2,0)_{12} + \text{AR}_{\phi}(2)$, incorporating the harmonic regressors, x_{1t} and x_{2t}, as well as the indicator variables x_{3t} and x_{4t}. However, the seasonal autoregressive component of order one was not statistically significant at the 5% level and was therefore removed from the model. A diagnostic analysis was conducted to assess the adequacy of the specification. The p-values of the Ljung–Box and Monti tests were 0.8049 and 0.7424, respectively, indicating no evidence of serial correlation in the residuals.

As in the modeling of the Sobradinho series presented in Section 6.4, the normality of the residuals was initially rejected by the Jarque–Bera test at the 5% significance level. Nevertheless, the residual diagnostics indicated no evidence of remaining serial dependence, with the Portmanteau Ljung–Box and Monti tests yielding p-values of 0.8049 and 0.7424, respectively. A more detailed inspection revealed that two residuals with relatively large absolute values were exerting a substantial influence on the outcome of the normality test. These points are associated with atypical and abrupt fluctuations in the time series. Specifically, they occur during periods of declining inflow; however, in June 2013 and July 2015, the inflow increased by more than 40% relative to the immediately preceding month, representing unusually sharp reversals in the local behavior of the series.

In the out-of-sample forecasting evaluation, we also considered a BPSARMA model with fixed precision. In this case, the selected specification was $\text{BPSARMA}_{\mu}(1,0) \times (1,1)_{12}$, including the two harmonic regressors and the indicator variables x_{3t} and x_{4t}. Keeping the same set of harmonic regressors and dummy variables, we also fitted a generalized BPARMA model, for which the selected specification was $\text{BPARMA}(4,1)$. Additionally, a SARIMA model of order $(1,1,1) \times (0,0,2)_{12}$ was fitted to the data.

Table 12 presents the out-of-sample forecast accuracy measures for horizons ranging from one to six steps ahead. Considering all evaluation metrics jointly, the proposed model exhibits superior predictive performance for horizons from three to six steps ahead, consistently yielding the lowest values among the models considered. For the one-step-ahead horizon, the BPSARMA model with fixed precision shows the best predictive performance. In contrast, for the two-step-ahead horizon, the SARIMA model achieves relatively better performance compared with the other models analyzed.

Table 12. Measures of forecast accuracy for different forecast horizons; Itaipu reservoir inflow.

	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$
			RRMSE			
Generalized BPSARMA	0.0382	0.0695	**0.0593**	**0.0646**	**0.0584**	**0.0536**
Generalized BPSARMA fixed prec.	**0.0097**	0.0337	0.1220	0.1950	0.2029	0.2216
Generalized BPARMA	0.0276	0.0201	0.0879	0.1266	0.1212	0.1222
SARIMA	0.0153	**0.0164**	0.1151	0.1650	0.1561	0.1482
Holt–Winters additive	0.1196	0.0961	0.0793	0.0698	0.0784	0.0723
Holt–Winters multiplicative	0.0441	0.0336	0.0697	0.1126	0.1043	0.1097
			MARE			
Generalized BPSARMA	0.0382	0.0644	**0.0540**	**0.0595**	**0.0513**	**0.0447**
Generalized BPSARMA fixed prec.	**0.0097**	0.0282	0.0815	0.1371	0.1565	0.1814
Generalized BPARMA	0.0276	**0.0172**	0.0570	0.0897	0.0912	0.0975
SARIMA	0.0153	0.0164	0.0709	0.1143	0.1141	0.1116
Holt–Winters additive	0.1196	0.0920	0.0705	0.0600	0.0695	0.0617
Holt–Winters multiplicative	0.0441	0.0309	0.0549	0.0854	0.0804	0.0897

7. Final Remarks

This article develops a flexible modeling framework for positive-valued time series exhibiting pronounced seasonal behavior. The proposed specification extends the generalized BPARMA class by incorporating seasonal autoregressive and moving-average dynamics in the conditional mean structure while allowing for a richer dynamic specification in the precision submodel. In particular, the precision parameter is permitted to follow autoregressive dynamics of arbitrary order, which enhances the model's ability to capture persistence in conditional variability. The specification also accommodates the inclusion of exogenous regressors, harmonic terms, and seasonal indicator variables, providing multiple mechanisms for representing seasonal patterns.

Inference is conducted by conditional maximum likelihood estimation. Analytical expressions for the conditional score vector and the conditional Fisher information matrix are derived, which simplifies numerical implementation and facilitates standard likelihood-based inference. Diagnostic tools are also discussed to support model checking and adequacy assessment. A Monte Carlo simulation study investigates the finite-sample behavior of the estimators and indicates satisfactory performance in sample sizes typically encountered in applied work.

The empirical analysis examines five hydro-environmental time series corresponding to the monthly inflows of reservoirs associated with hydroelectric power plants located in the five geographic regions of Brazil. These series provide a demanding empirical setting due to their pronounced seasonal patterns and the heterogeneity arising from the country's diverse climatic regimes. In the empirical applications, seasonality is represented through a combination of seasonal dynamics, harmonic components, and seasonal indicator variables. This integrated specification proves effective in capturing the complex seasonal structures present in the data.

Forecasts generated by the proposed model are evaluated against those obtained from several alternative approaches, including the BPARMA model with fixed precision, the generalized BPARMA model, a Gaussian SARIMA specification, and the Holt–Winters exponential smoothing method. Across the series analyzed, the proposed model typically achieves superior predictive accuracy, highlighting the benefits of jointly modeling seasonal dynamics and time-varying precision within the beta prime framework.

Overall, the results suggest that the proposed model constitutes a useful tool for the analysis and forecasting of positive seasonal time series. In the context of hydroelectric systems, accurate modeling and forecasting of reservoir inflows are essential for operational planning and water resource management. More broadly, the modeling framework introduced here may be applied to a wide range of environmental, hydrological, and economic time series characterized by positive support and strong seasonal variation.

Author Contributions: Conceptualization, K.H.S. and F.C.-N.; Methodology, K.H.S. and F.C.-N.; Software, K.H.S.; Validation, K.H.S. and F.C.-N.; Formal Analysis, K.H.S. and F.C.-N.; Investigation, K.H.S. and F.C.-N.; Resources, K.H.S. and F.C.-N.; Data Curation, K.H.S. and F.C.-N.; Writing—Original Draft Preparation, K.H.S. and F.C.-N.; Writing—Review & Editing, K.H.S. and F.C.-N.; Visualization, K.H.S.; Supervision, F.C.-N.; Project Administration, F.C.-N.; Funding Acquisition, F.C.-N. All authors have read and agreed to the published version of the manuscript.

Funding: This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior—Brasil (CAPES)—Finance Code 001. We also gratefully acknowledge partial financial support from Conselho Nacional de Desenvolvimento Científico e Tecnológico–CNPq (grant 304646/2023-7).

Data Availability Statement: Publicly available datasets were analyzed in this study. The data can be obtained from the the Brazilian National Electric System Operator (ONS, https://www.ons.org.br, accessed on 16 March 2026).

Acknowledgments: The authors are grateful to three anonymous reviewers for their careful reading of the manuscript and for their constructive comments and suggestions, which helped improve the quality and clarity of the paper. During the preparation of this manuscript, the authors used ChatGPT (OpenAI, GPT-5.5 version) to assist with language editing, text revision, and improvements in readability. The authors carefully reviewed and edited all generated content and take full responsibility for the final version of the manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

BIC	Bayesian Information Criterion
βARMA	Beta Autoregressive Moving Average
βSARMA	Beta Seasonal Autoregressive Moving Average
BPARMA	Beta Prime Autoregressive Moving Average
BPSARMA	Seasonal Beta Prime Autoregressive Moving Average
CMLE	Conditional Maximum Likelihood Estimador
MAIC	Modified Akaike Information Criterion
MBIC	Modified Bayesian Information Criterion
MARE	Mean Absolute Relative Error
MKSARMAX	Modified Kumaraswamy Seasonal Autoregressive Moving Average with Exogenous Regressors
MSE	Mean Squared Error
RRMSE	Relative Root Mean Squared Error
SARIMA	Seasonal Autoregressive Moving Average
SD	Standard Deviation

Appendix A

We now present the quantities required to obtain the conditional score vector. To compute the first derivative of the conditional log-likelihood with respect to the ith element of $\boldsymbol{\gamma_1}$, with $i \in \{1, \ldots, p_1 + q + P + Q + \nu + 1\}$, the chain rule is used. We have that

$$\frac{\partial \ell}{\partial \gamma_{1i}} = \sum_{t=m+1}^{n} \frac{\partial \ell_t(\mu_t, \phi_t)}{\partial \mu_t} \frac{d\mu_t}{d\eta_{1t}} \frac{\partial \eta_{1t}}{\partial \gamma_{1i}},$$

where

$$\frac{\partial \ell_t(\mu_t, \phi_t)}{\partial \mu_t} = (1 + \phi_t)(Y_t^* - \mu_t^*) =: \kappa_{1t} \tag{A1}$$

and

$$\frac{d\mu_t}{d\eta_{1t}} = \frac{1}{g_1'(\mu_t)}.$$

Here, $Y_t^* := \log(Y_t/(1 + Y_t))$ and $\mu_t^* := \psi(a_t) - \psi(a_t + b_t)$, with $a_t := \mu_t(1 + \phi_t)$ and $b_t := \phi_t + 2$, and $\psi(\cdot)$ denotes the digamma function. Recall that $g_1(\mu_t) = \eta_{1t}$. Therefore,

$$\frac{\partial \ell}{\partial \gamma_{1i}} = \sum_{t=m+1}^{n} \frac{\kappa_{1t}}{g_1'(\mu_t)} \frac{\partial \eta_{1t}}{\partial \gamma_{1i}}.$$

We next obtain closed-form expressions for the derivatives of η_{1t} with respect to each component of the parameter vector $\boldsymbol{\gamma_1}$. These derivatives have a recursive structure and are given by

$$
\begin{aligned}
\frac{\partial \eta_{1t}}{\partial \alpha_1} &= 1 + \sum_{j=1}^{q} \theta_j \frac{\partial \eta_{1t-j}}{\partial \alpha_1} + \sum_{J=1}^{Q} \Theta_J \frac{\partial \eta_{1t-JS}}{\partial \alpha_1} - \sum_{j=1}^{q} \sum_{J=1}^{Q} \theta_j \Theta_J \frac{\partial \eta_{1t-(j+JS)}}{\partial \alpha_1}, \\
\frac{\partial \eta_{1t}}{\partial \beta_c} &= x_{tc} - \sum_{i=1}^{p_1} \varphi_i x_{(t-i)c} - \sum_{I=1}^{P} \Phi_I x_{(t-IS)c} + \sum_{i=1}^{p_1} \sum_{I=1}^{P} \varphi_i \Phi_I x_{[t-(i+IS)]c} \\
&\quad + \sum_{j=1}^{q} \theta_j \frac{\partial \eta_{1t-j}}{\partial \beta_c} + \sum_{J=1}^{Q} \Theta_J \frac{\partial \eta_{1t-JS}}{\partial \beta_c} - \sum_{j=1}^{q} \sum_{J=1}^{Q} \theta_j \Theta_J \frac{\partial \eta_{1t-(j+JS)}}{\partial \beta_c}, \\
\frac{\partial \eta_{1t}}{\partial \varphi_i} &= g_1(Y_{t-i}) - \boldsymbol{x}_{t-i}^{\top} \boldsymbol{\beta} - \sum_{I=1}^{P} \Phi_I [g_1(Y_{t-(i+IS)}) - \boldsymbol{x}_{t-(i+IS)}^{\top} \boldsymbol{\beta}] + \sum_{j=1}^{q} \theta_j \frac{\partial \eta_{1t-j}}{\partial \varphi_i} \\
&\quad + \sum_{J=1}^{Q} \Theta_J \frac{\partial \eta_{1t-JS}}{\partial \varphi_i} - \sum_{j=1}^{q} \sum_{J=1}^{Q} \theta_j \Theta_J \frac{\partial \eta_{1t-(j+JS)}}{\partial \varphi_i}, \\
\frac{\partial \eta_{1t}}{\partial \theta_j} &= -r_{t-j} + \sum_{J=1}^{Q} \Theta_J r_{t-(j+JS)} + \sum_{j=1}^{q} \theta_j \frac{\partial \eta_{1t-j}}{\partial \theta_j} + \sum_{J=1}^{Q} \Theta_J \frac{\partial \eta_{1t-JS}}{\partial \theta_j} \\
&\quad - \sum_{j=1}^{q} \sum_{J=1}^{Q} \theta_j \Theta_J \frac{\partial \eta_{1t-(j+JS)}}{\partial \theta_j}, \\
\frac{\partial \eta_{1t}}{\partial \Phi_I} &= g_1(Y_{t-IS}) - \boldsymbol{x}_{t-IS}^{\top} \boldsymbol{\beta} - \sum_{i=1}^{p_1} \varphi_i [g_1(Y_{t-(i+IS)}) - \boldsymbol{x}_{t-(i+IS)}^{\top} \boldsymbol{\beta}] + \sum_{j=1}^{q} \theta_j \frac{\partial \eta_{1t-j}}{\partial \Phi_I} \\
&\quad + \sum_{J=1}^{Q} \Theta_J \frac{\partial \eta_{1t-JS}}{\partial \Phi_I} - \sum_{j=1}^{q} \sum_{J=1}^{Q} \theta_j \Theta_J \frac{\partial \eta_{1t-(j+JS)}}{\partial \Phi_I}, \\
\frac{\partial \eta_{1t}}{\partial \Theta_J} &= -r_{t-JS} + \sum_{j=1}^{q} \theta_j r_{t-(j+JS)} + \sum_{j=1}^{q} \theta_j \frac{\partial \eta_{1t-j}}{\partial \Theta_J} + \sum_{J=1}^{Q} \Theta_J \frac{\partial \eta_{1t-JS}}{\partial \Theta_J} \\
&\quad - \sum_{j=1}^{q} \sum_{J=1}^{Q} \theta_j \Theta_J \frac{\partial \eta_{1t-(j+JS)}}{\partial \Theta_J},
\end{aligned}
$$

for $c \in \{1, \ldots, l\}$, $i \in \{1, \ldots, p_1\}$, $j \in \{1, \ldots, q\}$, $I \in \{1, \ldots, P\}$, and $J \in \{1, \ldots, Q\}$.

In the absence of moving average components, whether seasonal or nonseasonal, the quantity η_{1t} and its partial derivatives can be evaluated directly, without the need for recursion.

Next, we obtain a closed-form expression for the derivative of the conditional log-likelihood function with respect to the i-th element of the parameter vector γ_2, with $i \in \{1, \ldots, p_2 + 1\}$. We have that

$$\frac{\partial \ell}{\partial \gamma_2} = \sum_{t=m+1}^{n} \frac{\partial \ell_t(\mu_t, \phi_t)}{\partial \phi_t} \frac{d\phi_t}{d\eta_{2t}} \frac{\partial \eta_{2t}}{\partial \gamma_2},$$

where

$$\frac{\partial \ell_t(\mu_t, \phi_t)}{\partial \phi_t} = Y_t^{\dagger} - \mu_t^{\dagger} =: \kappa_{2t} \tag{A2}$$

and

$$\frac{d\phi_t}{d\eta_{2t}} = \frac{1}{g_2'(\phi_t)}.$$

Here, $Y_t^{\dagger} := \mu_t \log(Y_t) - (1 + \mu_t) \log(1 + Y_t)$ e $\mu_t^{\dagger} := \mu_t(\mu_t^{*} - \Delta_t/\mu_t)$, with $\Delta_t := \psi(a_t + b_t - \psi(b_t)$. Recall that $g_2(\phi_t) = \eta_{2t}$ and, hence,

$$\frac{\partial \ell}{\partial \gamma_2} = \sum_{t=m+1}^{n} \frac{\kappa_{2t}}{g_2'(\phi_t)} \frac{\partial \eta_{2t}}{\partial \gamma_2}.$$

The derivatives of η_{2t} with respect to the components of $\boldsymbol{\gamma_2}$ are

$$\frac{\partial \eta_{2t}}{\partial \alpha_2} = 1 \quad \text{and} \quad \frac{\partial \eta_{2t}}{\partial \delta_k} = z_{t-k}.$$

Next, we present the quantities required to obtain the conditional information matrix. The second derivatives of the conditional log-likelihood function, for $i, j \in \{1, \ldots, w\}$, are

$$\begin{aligned}\frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \gamma_{1i}\partial \gamma_{1j}} &= \sum_{t=m+1}^{n} \frac{\partial}{\partial \mu_t}\left(\frac{\partial \ell_t(\mu_t,\phi_t)}{\partial \mu_t}\frac{d\mu_t}{d\eta_{1t}}\frac{\partial \eta_{1t}}{\partial \gamma_{1j}}\right)\frac{d\mu_t}{d\eta_{1t}}\frac{\partial \eta_{1t}}{\partial \gamma_{1i}} \\ &= \sum_{t=m+1}^{n} \frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \mu_t^2}\frac{d\mu_t}{d\eta_{1t}}\frac{\partial \eta_{1t}}{\partial \gamma_{1j}} + \frac{\partial \ell_t(\mu_t,\phi_t)}{\partial \mu_t}\frac{\partial}{\partial \mu_t}\left(\frac{d\mu_t}{d\eta_{1t}}\frac{\partial \eta_{1t}}{\partial \gamma_{1j}}\right)\frac{d\mu_t}{d\eta_{1t}}\frac{\partial \eta_{1t}}{\partial \gamma_{1i}},\end{aligned}$$

$$\begin{aligned}\frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \gamma_{2i}\partial \gamma_{2j}} &= \sum_{t=m+1}^{n} \frac{\partial}{\partial \phi_t}\left(\frac{\partial \ell_t(\mu_t,\phi_t)}{\partial \phi_t}\frac{d\phi_t}{d\eta_{2t}}\frac{\partial \eta_{2t}}{\partial \gamma_{2j}}\right)\frac{d\phi_t}{d\eta_{2t}}\frac{\partial \eta_{2t}}{\partial \gamma_{2i}} \\ &= \sum_{t=m+1}^{n} \frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \phi_t^2}\frac{d\phi_t}{d\eta_{2t}}\frac{\partial \eta_{2t}}{\partial \gamma_{2j}} + \frac{\partial \ell_t(\mu_t,\phi_t)}{\partial \phi_t}\frac{\partial}{\partial \phi_t}\left(\frac{d\phi_t}{d\eta_{2t}}\frac{\partial \eta_{2t}}{\partial \gamma_{2j}}\right)\frac{d\phi_t}{d\eta_{2t}}\frac{\partial \eta_{2t}}{\partial \gamma_{2i}},\end{aligned}$$

and

$$\begin{aligned}\frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \gamma_{1i}\partial \gamma_{2j}} &= \sum_{t=m+1}^{n} \frac{\partial}{\partial \phi_t}\left(\frac{\partial \ell_t(\mu_t,\phi_t)}{\partial \mu_t}\frac{d\mu_t}{d\eta_{1t}}\frac{\partial \eta_{1t}}{\partial \gamma_{2j}}\right)\frac{d\phi_t}{d\eta_{2t}}\frac{\partial \eta_{2t}}{\partial \gamma_{1i}} \\ &= \sum_{t=m+1}^{n} \frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \phi_t \partial \mu_t}\frac{d\mu_t}{d\eta_{1t}}\frac{\partial \eta_{1t}}{\partial \gamma_{2j}} + \frac{\partial \ell_t(\mu_t,\phi_t)}{\partial \mu_t}\frac{\partial}{\partial \phi_t}\left(\frac{d\mu_t}{d\eta_{1t}}\frac{\partial \eta_{1t}}{\partial \gamma_{2j}}\right)\frac{d\phi_t}{d\eta_{2t}}\frac{\partial \eta_{2t}}{\partial \gamma_{1i}}.\end{aligned}$$

Under the usual regularity conditions, it holds that $\mathbb{E}(\partial \ell_t(\mu_t,\phi_t)/\partial \mu_t \mid \mathcal{F}_{t-1}) = 0$ and $\mathbb{E}(\partial \ell_t(\mu_t,\phi_t)/\partial \phi_t \mid \mathcal{F}_{t-1}) = 0$. This follows from the fact that the conditional expectations of Equations (A1) and (A2), given $\mathcal{F}_{t-1}$, are equal to μ_t^* and $\mu_t^\dagger$, respectively, almost surely. Therefore,

$$\mathbb{E}\left(\frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \gamma_{1i}\partial \gamma_{1j}}\Big|\mathcal{F}_{t-1}\right) = \sum_{t=m+1}^{n}\left[\mathbb{E}\left(\frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \mu_t^2}\Big|\mathcal{F}_{t-1}\right)\right]\left(\frac{d\mu_t}{d\eta_{1t}}\right)^2\frac{\partial \eta_{1t}}{\partial \gamma_{1j}}\frac{\partial \eta_{1t}}{\partial \gamma_{1i}}, \tag{A3}$$

$$\mathbb{E}\left(\frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \gamma_{2i}\partial \gamma_{2j}}\Big|\mathcal{F}_{t-1}\right) = \sum_{t=m+1}^{n}\left[\mathbb{E}\left(\frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \phi_t^2}\Big|\mathcal{F}_{t-1}\right)\right]\left(\frac{d\phi_t}{d\eta_{2t}}\right)^2\frac{\partial \eta_{2t}}{\partial \gamma_{2j}}\frac{\partial \eta_{2t}}{\partial \gamma_{2i}}, \tag{A4}$$

and

$$\mathbb{E}\left(\frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \gamma_{1i}\partial \gamma_{2j}}\Big|\mathcal{F}_{t-1}\right) = \sum_{t=m+1}^{n}\left[\mathbb{E}\left(\frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \phi_t \partial \mu_t}\Big|\mathcal{F}_{t-1}\right)\right]\left(\frac{d\mu_t}{d\eta_{1t}}\right)\left(\frac{d\phi_t}{d\eta_{2t}}\right)\frac{\partial \eta_{1t}}{\partial \gamma_{2j}}\frac{\partial \eta_{2t}}{\partial \gamma_{1i}}. \tag{A5}$$

Taking the second derivatives of (A1) and (A2) with respect to μ_t and ϕ_t, respectively, we obtain

$$\frac{\partial^2 \ell_t(\mu_t,\phi_t)}{\partial \mu_t^2} = -(1+\phi_t)^2\left[\psi'(a_t) - \psi'(a_t+b_t)\right] =: \lambda_{1t} \tag{A6}$$

and

$$\frac{\partial^2 \ell_t(\mu_t, \phi_t)}{\partial \phi_t^2} = -\mu_t^2 \psi'(a_t) + (1 + \mu_t)^2 \psi'(a_t + b_t) - \psi'(b_t) =: \lambda_{2t}, \tag{A7}$$

where $\psi'(\cdot)$ denotes the trigamma function. Plugging (A6) into (A3) and (A7) into (A4), we obtain, respectively,

$$\mathbb{E}\left(\frac{\partial^2 \ell_t(\mu_t, \phi_t)}{\partial \gamma_{1i} \partial \gamma_{1j}} \Big| \mathcal{F}_{t-1} \right) = - \sum_{t=m+1}^{n} \frac{\lambda_{1t}}{(g_1'(\mu_t))^2} \frac{\partial \eta_{1t}}{\partial \gamma_{1j}} \frac{\partial \eta_{1t}}{\partial \gamma_{1i}}$$

and

$$\mathbb{E}\left(\frac{\partial^2 \ell_t(\mu_t, \phi_t)}{\partial \gamma_{2i} \partial \gamma_{2j}} \Big| \mathcal{F}_{t-1} \right) = - \sum_{t=m+1}^{n} \frac{\lambda_{2t}}{(g_2'(\phi_t))^2} \frac{\partial \eta_{2t}}{\partial \gamma_{2j}} \frac{\partial \eta_{2t}}{\partial \gamma_{2i}}.$$

The cross-derivative is obtained by differentiating (A1) with respect to ϕ_t:

$$\frac{\partial^2 \ell_t(\mu_t, \phi_t)}{\partial \phi_t \partial \mu_t} = Y_t^* - \mu_t^* + (1 + \phi_t)\{\psi'(a_t + b_t) + \mu_t[\psi'(a_t + b_t) - \psi'(a_t)]\} =: \lambda_{3t}. \tag{A8}$$

As noted earlier, $\mathbb{E}(Y_t \mid \mathcal{F}_{t-1}) = \mu_t$ almost surely. Plugging (A8) into (A5), we arrive at

$$\mathbb{E}\left(\frac{\partial^2 \ell_t(\mu_t, \phi_t)}{\partial \gamma_{1i} \partial \gamma_{2j}} \Big| \mathcal{F}_{t-1} \right) = \sum_{t=m+1}^{n} \frac{\lambda_{3t}}{g_1'(\mu_t) g_2'(\phi_t)} \frac{\partial \eta_{2t}}{\partial \gamma_{2j}} \frac{\partial \eta_{1t}}{\partial \gamma_{1i}}.$$

References

1. Rocha, A.V.; Cribari-Neto, F. Beta autoregressive moving average models. *TEST* **2009**, *18*, 529–545. Erratum in *TEST* **2017**, *26*, 451–459. [CrossRef]
2. Scher, V.T.; Cribari-Neto, F.; Bayer, F.M. Generalized βARMA model for double bounded time series forecasting. *Int. J. Forecast.* **2024**, *40*, 721–734. [CrossRef]
3. Bayer, F.M.; Cintra, R.J.; Cribari-Neto, F. Beta seasonal autoregressive moving average models. *J. Stat. Comput. Simul.* **2018**, *88*, 2961–2981. [CrossRef]
4. Costa, E.; Cribari-Neto, F.; Scher, V.T. Test inferences and link function selection in dynamic beta modeling of seasonal hydro-environmental time series with temporary abnormal regimes. *J. Hydrol.* **2024**, *638*, 131489. [CrossRef]
5. Shad, M.; Sharma, Y.D.; Narula, P. Forecasting Southwest Indian monsoon rainfall using the beta seasonal autoregressive moving average (βSARMA) model. *Pure Appl. Geophys.* **2023**, *180*, 405–419. [CrossRef]
6. Li, A.; Yang, K. A novel double-banded-threshold mixture autoregressive model. *Stat. Pap.* **2025**, *66*, 140. [CrossRef]
7. Li, H.; Zhang, Q.; Yang, K. Bayesian forecasting for a logistic mixture double autoregressive model. *J. Forecast.* **2026**, *45*, 1665–1680. [CrossRef]
8. Santos, K.H.; Cribari-Neto, F. A varying precision beta prime autoregressive moving average model with application to water flow data. *Environmetrics* **2024**, *35*, e2886. [CrossRef]
9. Yang, K.; Zhao, Y.; Li, H.; Wang, D. On bivariate threshold Poisson integer-valued autoregressive processes. *Metrika* **2023**, *86*, 931–963. [CrossRef]
10. Yang, K.; Xu, N.; Li, H.; Zhao, Y.; Dong, X. Multivariate threshold integer-valued autoregressive processes with explanatory variables. *Appl. Math. Model.* **2023**, *124*, 142–166. [CrossRef]
11. Weiß, C.H.; Zhu, F. Tobit models for count time series. *Scand. J. Stat.* **2025**, *52*, 381–415. [CrossRef]
12. Keeping, E.S. *Introduction to Statistical Inference*; Dover Publications: New York, NY, USA, 1962.
13. McDonald, J.B. Some generalized functions for the size distribution of income. *Econometrica* **1984**, *52*, 647–663. [CrossRef]
14. Bourguignon, M.; Santos-Neto, M.; Castro, M. A new regression model for positive random variables with skewed and long tail. *Metron* **2021**, *79*, 33–55. [CrossRef]
15. Armanini Stefanan, A.; Sagrillo, M.; Palm, B.G.; Bayer, F.M. Modified Kumaraswamy seasonal autoregressive moving average models with exogenous regressors for double-bounded hydro-environmental data. *PLoS ONE* **2025**, *20*, e0324721. [CrossRef] [PubMed]
16. Scher, V.T.; Cribari-Neto, F.; Pumi, G.; Bayer, F.M. Goodness-of-fit tests for βARMA hydrological time series modeling. *Environmetrics* **2020**, *31*, e2607. [CrossRef]

17. Cribari-Neto, F.; Costa, E.; Fonseca, R.V. Numerical stability enhancements in beta autoregressive moving average model estimation. *Braz. J. Probab. Stat.* **2025**, *39*, 410–437. [CrossRef]
18. Press, W.H.; Teukolsky, S.A.; Vetterling, W.T.; Flannery, B.P. *Numerical Recipes in C*, 2nd ed.; Cambridge University Press: New York, NY, USA, 1992.
19. Nocedal, J.; Wright, S.J. *Numerical Optimization*, 2nd ed.; Springer: New York, NY, USA, 2006.
20. Andersen, E.B. Asymptotic properties of conditional maximum-likelihood estimators. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **1970**, *32*, 283–301. [CrossRef]
21. Pawitan, Y. *In All Likelihood: Statistical Modelling and Inference Using Likelihood*; Oxford University Press: Oxford, UK, 2001.
22. Neyman, J.; Pearson, E.S. On the use and interpretation of certain test criteria for purposes of statistical inference: Part I. *Biometrika* **1928**, *20*, 175–240. [CrossRef]
23. Rao, C.R. Large sample tests of statistical hypotheses concerning several parameters with applications to problems of estimation. In *Proceedings of the Mathematical Proceedings of the Cambridge Philosophical Society*; Cambridge University Press: New York, NY, USA, 1948; Volume 44, pp. 50–57.
24. Wald, A. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Trans. Am. Math. Soc.* **1943**, *54*, 426–482. [CrossRef]
25. Dunn, P.K.; Smyth, G.K. Randomized quantile residuals. *J. Comput. Graph. Stat.* **1996**, *5*, 236–244. [CrossRef]
26. Ljung, G.M.; Box, G.E.P. On a measure of lack of fit in time series models. *Biometrika* **1978**, *65*, 297–303. [CrossRef]
27. Monti, A.C. A proposal for a residual autocorrelation test in linear models. *Biometrika* **1994**, *81*, 776–780. [CrossRef]
28. R Core Team. *R: A Language and Environment for Statistical Computing*; R Foundation for Statistical Computing: Vienna, Austria, 2026.
29. Melo, M.; Alencar, A. Conway–Maxwell–Poisson seasonal autoregressive moving average model. *J. Stat. Comput. Simul.* **2022**, *92*, 283–299. [CrossRef]
30. Finke, K.; Jiménez-Esteve, B.; Taschetto, A.S.; Ummenhofer, C.C.; Bumke, K.; Domeisen, D.I.V. Revisiting remote drivers of the 2014 drought in South-Eastern Brazil. *Clim. Dyn.* **2020**, *55*, 3197–3211. [CrossRef] [PubMed]