FUDMA Journal of Sciences (FJS)
ISSN online: 2616-1370
ISSN print: 2645 - 2944
Vol. 10 No. 18, September, 2026, pp 295 – 303
DOI: https://doi.org/10.33003/fjs-2026-1018-6086

Bayesian Analysis of Seasonal and Periodic Time Series: Evidence from Nigerian Monthly Rainfall, 1981–2024

Oyenuga Augustine Yomi, Olatayo Timothy Olabisi, Taiwo Abass Ishola and *Obadina Olugbenga Gabriel

Department of Statistics, Olabisi Onabanjo University, Ago-Iwoye, Ogun State, Nigeria.

*Corresponding authors' email: olugbenga.obadina@yahoo.com

ABSTRACT
Monthly rainfall in Nigeria exhibits pronounced annual seasonality, season-dependent variability, and changes in temporal dependence that may not be adequately represented by a single fixed-parameter model. This study presents a Bayesian modelling framework for monthly rainfall records from 37 Nigerian Meteorological Agency stations spanning January 1981 to December 2024, a 44-year period comprising 528 monthly values per station. The framework combines Whittle-likelihood spectral estimation, B-spline modelling of the log-spectral density, Bayesian structural time series (BSTS) decomposition, periodic autoregression, locally stationary spectral analysis, and reversible-jump Markov chain Monte Carlo for structural-break detection. Predictive distributions from complementary specifications are combined through Bayesian model averaging (BMA). Reported validation results for 2022–2024 indicate that BMA reduces root mean squared error relative to SARIMA by 8.3% at a 12-month horizon and 15.0% at a 24-month horizon. BMA has the lowest reported mean absolute percentage error at both horizons and the lowest 24-month RMSE; however, BSTS has the lowest 12-month RMSE. Structural-break summaries identify modal break locations within 1987–1989 at 22 stations and within 2001–2004 at 19 stations. These findings support the use of complementary seasonal models, particularly for longer-horizon prediction, rather than establishing uniform superiority of model averaging. Interpretation remains conditional on the reported evaluation design, and further documentation of forecast origins, model weights, zero-rainfall treatment, and posterior diagnostics is required for independent verification.

Received: 08 Sep. 2026
Accepted: 11 Sep. 2026
Published: 28 Sep. 2026

Keywords: Bayesian Inference; Nigerian Rainfall; Whittle Likelihood; Spectral Density; Structural Time Series; Periodic Autoregression; Bayesian Model Averaging

INTRODUCTION

Seasonal time series are central to climatology, agriculture, hydrology, and energy planning. Standard seasonal autoregressive integrated moving-average (SARIMA) models provide useful representations when the dependence structure is sufficiently stable after appropriate transformation and differencing. Their adequacy is less certain when seasonal amplitude, phase, or dependence changes over time. In such settings, modelling a recurring annual pattern is not equivalent to modelling an invariant annual pattern.

Nigerian rainfall provides a demanding application because seasonal behaviour varies substantially across the country. The station descriptions in the available study distinguish the relatively concentrated rainy seasons of northern locations from the longer, frequently bimodal seasonal profiles of southern locations. Middle-belt stations occupy an intermediate position. These differences concern not only the timing of rainfall, but also the shape of the seasonal cycle, the variability of monthly totals, and the persistence of deviations from the seasonal mean.

A fixed harmonic representation can describe a non-sinusoidal cycle when sufficient harmonics are included. The more difficult problem is allowing that representation to evolve without introducing excessive complexity. Similarly, a stationary autoregressive structure may describe average dependence while obscuring changes between climatic regimes. A useful analysis therefore needs to distinguish recurring seasonality, season-specific dependence, gradual evolution, and abrupt structural change.

Bayesian methods provide a coherent way to represent uncertainty about these features. Prior distributions regularise flexible models, posterior sampling propagates parameter and latent-state uncertainty, and predictive mixtures account for uncertainty across candidate specifications. These advantages do not guarantee better empirical forecasts. Their practical value must be assessed against credible benchmarks using a clearly documented out-of-sample design.

The aim of this study is to evaluate a Bayesian framework for the analysis and forecasting of seasonal and periodic time series, using Nigerian monthly rainfall as the empirical application.

The contribution is an application-oriented combination of established methods rather than a claim that one model is optimal for every station, horizon, or loss function. Particular attention is given to the distinction between the theoretical coherence of Bayesian prediction and its observed forecasting performance.

Monthly rainfall forecasts can inform seasonal water budgeting, reservoir planning, and agricultural risk assessment. Distributional forecasts are especially useful when decisions depend on uncertainty rather than solely on expected rainfall. Monthly totals, however, do not directly identify within-month rainfall onset, daily dry spells, or short-duration flood extremes. Claims about these outcomes would require higher-frequency observations and a corresponding evaluation design.

The analysis treats the 37 stations separately. It does not estimate a spatial rainfall field or explicitly incorporate ocean–atmosphere indices. Consequently, the results concern station-level temporal patterns and forecast comparisons; they do not establish physical attribution of climatic changes. The framework may be relevant to other seasonal series, but transferability beyond this application is not demonstrated here.

Related literature

Bayesian Frequency-Domain Analysis

Whittle's (1957) frequency-domain approximation makes inference computationally convenient by expressing a likelihood in terms of periodogram ordinates. Under appropriate stationarity and short-memory conditions, ordinates at distinct positive Fourier frequencies behave approximately as independent exponential variables whose means are determined by the spectral density. This provides a natural basis for Bayesian smoothing of an otherwise highly variable periodogram.

Choudhuri, Ghosal, and Roy (2004) developed Bayesian spectral-density estimation using Bernstein-polynomial priors and established consistency under specified assumptions. Their work provides an important foundation for nonparametric frequency-domain inference. It does not, by itself, establish consistency for every alternative prior or for a time-varying model. In particular, a B-spline prior with a growing basis and a locally stationary extension require their own regularity and prior-concentration arguments.

Structural Time Series and Periodic Autoregression

State-space models provide a complementary time-domain interpretation by separating trend, seasonality, and irregular variation. Conditional Gaussian structure enables efficient posterior simulation using filtering and smoothing methods, as developed by Carter and Kohn (1996). Scott and Varian (2014) demonstrate how Bayesian structural time series can combine latent components with uncertainty about model specification. In a rainfall application, an evolving seasonal state can accommodate changes in the magnitude or shape of the annual cycle without imposing a separate deterministic curve for every year.

Periodic autoregressive models instead allow autoregressive parameters and innovation variances to depend on the season. Foundational treatments include Pagano (1978) and Vecchia (1985), while Lund and Basawa (2000) develop recursive prediction and likelihood evaluation for periodic ARMA processes. These models are well suited to situations in which wet- and dry-season dependence differ systematically.

Fourier autoregressive (FAR) models offer a parsimonious representation of periodically varying coefficients through trigonometric terms. Taiwo et al. (2019) applied this approach to Nigerian rainfall, motivating the inclusion of FAR(1) as a relevant benchmark in the present comparison. BSTS, PAR, and FAR therefore address related but distinct aspects of seasonality: changing latent components, season-specific dependence, and a compact harmonic parameterisation of periodic dependence, respectively.

Time-Varying Spectra and Structural Change

Dahlhaus (1997) formalised locally stationary processes through a spectral density indexed jointly by rescaled time and frequency. This framework distinguishes gradual evolution in dependence from the stronger assumption that one spectrum describes the entire record. Rosen, Stoffer, and Wood (2009) developed a Bayesian mixture-of-splines approach to local spectral estimation, illustrating how flexible smoothing and posterior uncertainty can be combined in the time–frequency domain.

Abrupt changes require a different representation. In a piecewise stationary model, an unknown set of breakpoints partitions the record into segments with distinct parameters. Green's (1995) reversible-jump MCMC framework permits transitions between models of different dimension and is therefore applicable when the number of breaks is unknown. Smooth local evolution and abrupt segmentation should be viewed as complementary descriptions rather than interchangeable assumptions.

Forecast Combination and Uncertainty

Vosseler and Weber (2018) use Bayesian model averaging to address uncertainty in seasonal time-series forecasting, including uncertainty in lag order and structural breaks. McAlinn and West (2019) extend forecast combination through dynamic Bayesian predictive synthesis, in which relationships among predictive distributions can evolve over time. The present framework uses the simpler starting point of static model averaging across complementary specifications.

The rationale for averaging is to propagate model uncertainty, not to guarantee that the mixture outperforms every component in every realised sample. Likewise, strictly proper scoring rules justify honest distributional prediction under the distribution used to calculate expected loss; they do not establish automatic finite-sample superiority under an unknown data-generating process (Gneiting and Raftery, 2007). These distinctions are important when interpreting the rainfall results.

MATERIALS AND METHODS

Data Source, Coverage, and Preprocessing

The dataset is described as monthly rainfall totals from 37 Nigerian Meteorological Agency (NIMET) climatological stations. Coverage extends from January 1981 through December 2024, which is 44 complete calendar years and 528 monthly values per station. The completed station-by-month panel therefore contains 19,536 values. These are panel entries after the reported preprocessing, not a verified count of originally observed gauge measurements.

The station network is described as covering the Sahel and Sudan Savanna in the north, the Guinea Savanna and derived-forest transition in the middle belt, and the Humid Forest and Atlantic littoral in the south. Examples discussed in the study include Maiduguri, Kano, Abuja, Lokoja, Ibadan, Benin City, Lagos, and Port Harcourt. A complete station inventory, coordinates, and elevation data were not included in the material available for this revision.

Trace values recorded as "Tr" were assigned a value of zero, and isolated gaps of one or two months were filled by linear interpolation. The record is therefore better described as a completed monthly panel than as an uninterrupted set of original observations. The number and seasonal distribution of imputed values were not reported. Although the gaps were short, sensitivity to interpolation remains relevant because smoothing missing values can alter variability and dependence. Table 1 summarises the available data and evaluation information.

Table 1: Dataset Coverage, Preprocessing, And Validation Information

Item	Reported specification or derived count
Variable and source	Monthly rainfall totals (mm); Nigerian Meteorological Agency (NIMET).
Station network	37 stations; stations modelled separately across Nigeria's principal climatic zones.
Full period	January 1981–December 2024: 44 years and 528 monthly values per station.
Completed panel	19,536 station-month entries (37 × 528), including any imputed values.
Trace rainfall	Values recorded as "Tr" assigned zero.
Missing values	Isolated one- or two-month gaps filled by linear interpolation; counts not supplied.
Validation period	January 2022–December 2024: 36 months; reported horizons h = 12 and h = 24 months.
Pre-validation history	January 1981–December 2021: 492 months per station; exact fitting and refitting schedule not supplied.

Organisation of the Analysis

The framework comprises complementary frequency-domain and time-domain analyses. Stationary spectral estimation identifies dominant periodicities; BSTS decomposes trend and evolving seasonality; PAR represents month-specific dependence; and local spectra or segmented models describe temporal changes in dependence. Model averaging combines predictive distributions from selected specifications. Table 2 distinguishes the roles of these components, while Figure 1 provides a conceptual overview.

The available specification does not define a single joint likelihood that simultaneously couples all components. Accordingly, the framework is presented as a coordinated set of analyses and a predictive combination, rather than as one fully specified joint generative model. This distinction avoids attributing properties of an individual component to the complete modelling strategy without additional justification.

Table 2: Complementary Components of the Bayesian Modelling Framework

Component	Core representation	Analytical role
Stationary spectrum	Whittle likelihood; B-spline expansion of log spectral density.	Identify annual and harmonic structure; regularise periodogram variability.
BSTS	Local linear trend plus stochastic seasonal harmonics.	Separate trend, evolving seasonality, and irregular variation.
Bayesian PAR	Month-specific autoregressive coefficients and innovation variances.	Represent periodic dependence and seasonal heteroscedasticity.
Local spectrum	Tensor-product splines in rescaled time and frequency.	Describe gradual changes in spectral structure.
Segmented spectrum	Unknown break number and locations; reversible-jump MCMC.	Describe abrupt changes between approximately stationary segments.
BMA	Posterior-weighted mixture of PAR, BSTS, and locally stationary predictive densities.	Represent uncertainty across the stated predictive specifications.

This table summarises the stated framework, not a new fitted model. SARIMA (1, 1, 1) $(1, 1, 1)_{12}$ and FAR (1) are the classical forecast benchmarks. BMA weights and several component-specific tuning choices were not reported.

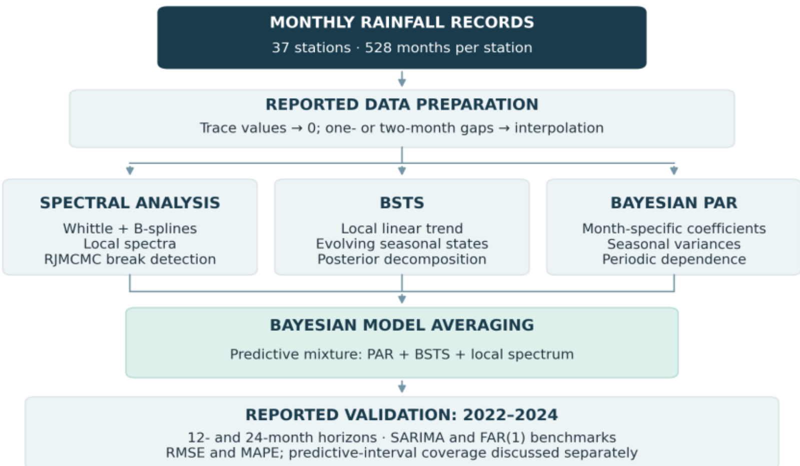

Figure 1: Conceptual organisation of the Bayesian analysis

Spectral Representation and the Whittle Approximation

For a real-valued, zero-mean stationary process X_t with autocovariance function $c(k)$, absolute summability of the autocovariances is a sufficient condition for the spectral density to exist and satisfy

$$f(\omega) = \frac{1}{2\pi}\sum_{k=-\infty}^{\infty} c(k)e^{-ik\omega}, \qquad -\pi \le \omega \le \pi. \tag{1}$$

More generally, a second-order stationary process has a spectral measure, which need not possess a density. The density-based formulation used here therefore requires additional regularity. Applied to rainfall, these expressions concern a suitably centred process or an approximately stationary segment; the exact centring, detrending, or transformation used in the empirical spectral calculations should be documented.

For monthly observations, the annual frequency and its first two harmonics are

$$\omega_A = \frac{2\pi}{12} \approx 0.524, \qquad 2\omega_A \approx 1.047, \qquad 3\omega_A \approx 1.571. \qquad (2)$$

Frequencies are measured in radians per month. Define the discrete Fourier transform and periodogram by

$$d_j = (2\pi n)^{-1/2} \sum_{t=1}^{n} X_t\, e^{-it\omega_j}, \qquad I_j = |d_j|^2, \qquad \omega_j = \frac{2\pi j}{n}. \qquad (3)$$

At distinct positive Fourier frequencies, excluding zero and the Nyquist frequency, the short-memory Gaussian approximation gives $I_j / f(\omega_j) \approx \mathrm{Exp}(1)$. Omitting constants that do not depend on the spectrum, the corresponding Whittle log-likelihood is

$$\ell_W(f) = -\sum_{j=1}^{\lfloor (n-1)/2 \rfloor} \{\log f(\omega_j) + \frac{I_j}{f(\omega_j)}\}. \qquad (4)$$

Excluding the boundary frequencies avoids assigning the interior-frequency exponential law to ordinates with different real-valued sampling distributions. The approximation also implies that the periodogram is asymptotically unbiased at regular frequencies but remains inconsistent without smoothing: its variance is approximately the square of the spectrum, so its relative variability does not vanish as the sample grows.

The exact Gaussian log-likelihood, up to an additive constant, is

$$\ell_G(f) = -\frac{1}{2}\log|\Sigma_n(f)| - \frac{1}{2} X^{\mathrm{T}} \Sigma_n(f)^{-1} X. \qquad (5)$$

Whittle inference approximates the leading frequency-domain behaviour of this likelihood under appropriate smoothness, boundedness, and dependence assumptions. No uniform, unnormalised $o_p(1)$ bound between the two log-likelihoods is established here. Stationary likelihood results also cannot be transferred automatically to a periodically correlated or locally stationary process.

Bayesian B-Spline Spectral Model

A flexible positive spectrum is obtained by modelling its logarithm. On the nonnegative frequency interval, let

$$\log f(\omega) = \sum_{k=1}^{K} \theta_k\, B_k(\omega), \qquad \theta_k \sim N(0, \tau^2), \qquad 0 \leq \omega \leq \pi. \qquad (6)$$

Here, B_k denotes a B-spline basis function, K is the basis dimension, and the spectrum is extended evenly to negative frequencies for a real-valued process. The logarithmic representation enforces positivity. Basis size, knot placement, and the prior scale jointly control flexibility and regularisation; the Gaussian coefficient prior is not, on its own, an explicit roughness-penalty specification.

Spline approximation provides a useful support argument: sufficiently rich bases can approximate a smooth log-spectrum, and a nondegenerate Gaussian prior assigns positive probability to neighbourhoods of finite coefficient vectors. Under suitable uniform bounds on the spectra, sufficiently close approximations can also make the spectral Kullback–Leibler discrepancy small:

$$D(f_0, f) = \frac{1}{4\pi} \int_{-\pi}^{\pi} \{\log \frac{f(\omega)}{f_0(\omega)} + \frac{f_0(\omega)}{f(\omega)} - 1\}\, d\omega. \qquad (7)$$

Prior support alone is not a complete posterior-consistency proof. A growing-basis construction additionally requires control of model complexity, prior concentration, and suitable tests for the dependent-data likelihood. Schwartz's (1965) consistency framework and the spectral results of Choudhuri et al. (2004) motivate such an argument, but the required conditions have not been verified for the full specification considered here. Consistency is therefore treated as a theoretical objective rather than a newly established result.

Bayesian Structural Time Series Decomposition

For a given station, the BSTS model decomposes rainfall into a local trend, an evolving seasonal component, and observation noise:

$$Y_t = \mu_t + \gamma_t + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2). \qquad (8)$$

The trend follows a local linear specification:

$$\mu_{t+1} = \mu_t + \delta_t + \eta_t, \qquad \delta_{t+1} = \delta_t + \zeta_t. \qquad (9)$$

The disturbances η_t and ζ_t are independent zero-mean Gaussian innovations with their own variance parameters. The seasonal component is represented by time-varying harmonic coefficients:

$$\gamma_t = \sum_{j=1}^{S} \{a_{jt} \cos(\frac{2\pi j t}{12}) + b_{jt} \sin(\frac{2\pi j t}{12})\}. \qquad (10)$$

The coefficients follow Gaussian random walks, allowing seasonal amplitude and phase to change gradually. The number of retained harmonics, S, must be specified; a full monthly harmonic representation requires separate treatment of the Nyquist cosine term because the corresponding sine term vanishes at integer times.

The reported computational approach uses forward filtering–backward sampling for latent states and Gibbs updates for variance parameters with inverse-gamma priors. The prior hyperparameters, initial-state distributions, retained harmonic count, chain lengths, and convergence diagnostics were not supplied. These details are necessary to reproduce the decomposition and evaluate sensitivity to prior choices.

Periodic Autoregressive Specification

A PAR model allows dependence and innovation variance to vary by calendar month. For a seasonally centred processX_t, a PAR (p) specification is

$$X_t = \sum_{k=1}^{p} \phi_k(s_t) X_{t-k} + \epsilon_t, \qquad \epsilon_t \sim N(0, \sigma_{s_t}^2). \qquad (11)$$

Here, $s_t = 1 + ((t-1) \bmod 12)$ identifies the month. There are $12p$ autoregressive coefficients and 12 innovation variances. If raw rainfall rather than a centred series is modelled, month-specific means or a suitable intercept component must also be included and documented.

Independent Gaussian priors on the autoregressive coefficients provide shrinkage. Periodic Yule–Walker estimates are used as initial values, and lag identification is based on periodic autocorrelation and partial-autocorrelation diagnostics, together with periodic AIC and BIC. Seasonal innovation variances directly represent wet- and dry-season heteroscedasticity.

Periodic stability is a property of the complete seasonal coefficient sequence, not of each coefficient considered separately. Unrestricted Gaussian priors do not automatically enforce this property. Reproducible implementation therefore requires an explicit account of stability restrictions or posterior stability checks. The season-dependent covariance structure should also not be described as an ordinary stationary spectrum obtained merely by substituting one month's coefficients into an AR spectral formula.

Local Spectra and Structural Breaks

Gradual evolution is represented by a local spectral density indexed by rescaled time$u = t/n$. A tensor-product spline model has the form

$$\log f(u, \omega) = \sum_{k=1}^{K_u} \sum_{l=1}^{K_\omega} \vartheta_{kl}\, G_k(u) B_l(\omega). \qquad (12)$$

The coefficient array is assigned a separable Gaussian prior. This construction allows the frequency profile to vary over time, with flexibility controlled by the time and frequency bases. It is related in motivation to local spline-based spectral analysis, but it is not identical to the mixture-of-splines model of Rosen et al. (2009).

Abrupt changes are instead represented by m ordered breakpoints, $b_1 < \cdots < b_m$, defining $m+1$ approximately stationary segments. A separate spectral parameter vector is assigned to each segment. Reversible-jump MCMC explores the number and locations of breaks through birth, death, and relocation proposals. Trans-dimensional acceptance probabilities must include the relevant posterior, proposal, and Jacobian terms where applicable.

The source description also allows changes in spline-knot configurations. Knot selection in a smooth surface and breakpoint selection in a segmented model are distinct operations, even if both use reversible-jump proposals. Their priors, admissible configurations, minimum segment lengths, and proposal mechanisms need separate documentation. The reported breakpoint counts are presented in Section 5.4 without reconstructing an unavailable posterior distribution over break dates.

Posterior Prediction and Model Averaging

Let $\mathcal{D}$ denote the data available at a forecast origin and let ψ_j collect the parameters and latent states of modelM_j. The model-specific predictive density is

$$p_j(y|\mathcal{D}) = \int p(y|\psi_j, \mathcal{D}, M_j)\, p(\psi_j|\mathcal{D}, M_j)\, d\psi_j. \quad (13)$$

Future innovations are simulated along with parameter and state uncertainty to obtain posterior predictive trajectories. BMA combines these distributions as

$$q(y|\mathcal{D}) = \sum_{j=1}^{J} w_j\, p_j(y|\mathcal{D}), \qquad w_j = p(M_j|\mathcal{D}), \qquad \sum_{j=1}^{J} w_j = 1. \quad (14)$$

The stated combination comprises PAR, BSTS, and locally stationary specifications. Posterior model probabilities depend on both marginal likelihoods and prior model probabilities. The original implementation describes a harmonic-mean estimate of marginal likelihoods. Because this estimator can be unstable, the numerical weights and their sensitivity require scrutiny; a stable marginal-likelihood method would be preferable in a reproducibility analysis. The available summary does not provide the weights or explain fully how the separately labelled "Nonparametric Spectral" forecast relates to the local-spectral component in BMA.

A precise decision-theoretic statement is available without claiming universal predictive dominance. For a negatively oriented strictly proper scoring rule$S$, if q is the posterior predictive distribution used to evaluate uncertainty, then

$$\mathbb{E}_q\{S(q,Y)\} \le \mathbb{E}_q\{S(g,Y)\} \quad (15)$$

For any admissible alternative predictive distribution$g$, with equality only when the distributions coincide. This follows directly from strict propriety. It is a posterior expected-loss statement, not a guarantee that BMA has the smallest realised RMSE under an unknown true process. Under squared-error loss, the predictive mean is the corresponding Bayes point forecast.

Model averaging also retains between-model uncertainty. Its predictive variance includes both within-model variances and variation among component means. It may therefore be larger than the variance of a selected component, even when its point forecast is more accurate. Uncertainty representation and variance reduction should not be conflated.

Forecast Evaluation

The reported validation period is January 2022 to December 2024, covering 36 months. The remaining period, January 1981 to December 2021, contains 492 months per station and is the available pre-validation history. Full-record descriptive analysis can use all 528 values, whereas honest forecasting requires estimation and preprocessing to use only information available at each forecast origin.

Results are reported at horizons of 12 and 24 months and are described as averages across the 37 stations. For station$r$, horizon$h$, and a collection of $N_{r,h}$ matched forecasts, the usual error measures are

$$\mathrm{RMSE}_{r,h} = \left(\frac{1}{N_{r,h}} \sum_i [Y_{r,i} - \hat{Y}_{r,i|i-h}]^2\right)^{1/2}, \quad (16)$$

And, when all denominators are positive,

$$\mathrm{MAPE}_{r,h} = \frac{100}{N_{r,h}} \sum_i \left| \frac{Y_{r,i} - \hat{Y}_{r,i|i-h}}{Y_{r,i}} \right| \quad (17)$$

RMSE is measured in millimetres and MAPE in percent. The table heading indicates an average of station-level scores, rather than a single score pooled over all station-months. However, the forecast-origin schedule, the number of forecasts per horizon, and any rolling refitting or updating were not documented. These omissions prevent independent reconstruction of the reported evaluation.

MAPE is undefined for zero rainfall and can be dominated by observations close to zero. This is particularly important because trace values were explicitly recoded to zero. The reported MAPE values are retained, but the rule used for zero and near-zero observations must be supplied before they can be interpreted as fully reproducible comparisons. MAE and scale-adjusted alternatives would provide useful complementary assessments.

The original design also mentions MAE, DIC, and WAIC, but supplies no numerical results for these quantities. They are therefore not used to support the conclusions. Predictive-interval coverage is discussed where reported, although interval widths and the associated evaluation scope are unavailable. Formal comparison of predictive distributions would additionally benefit from proper scores such as the continuous ranked probability score or an interval score.

All model comparisons below are descriptive. Neither station-level forecast errors nor were uncertainty estimates for score differences supplied, so no statistical-significance claim is made. Geographic dependence between stations should also be considered in any subsequent test or resampling procedure.

RESULTS AND DISCUSSION

Seasonal and Spectral Characteristics

The reported periodograms identify the annual frequency as the dominant feature at all 37 stations, with additional peaks near the second and third harmonics. Northern examples, particularly Maiduguri, are described as having relatively concentrated annual peaks. Southern examples, especially Port Harcourt, show broader spectral structure and a more prominent second harmonic. These qualitative patterns are consistent with differences in seasonal shape across the station network.

The smoothed spectra and posterior bands referenced in the original text were not included in the supplied PDF. Consequently, their detailed shapes, peak magnitudes, and interval widths cannot be reproduced here. In particular, a claim that uncertainty is necessarily smallest at the annual peak does not follow from the variance of the raw periodogram alone. Posterior uncertainty depends on the fitted model, smoothing, and the scale on which uncertainty is measured.

Structural Decomposition and Periodic Dependence

For Kano, the reported BSTS decomposition indicates declining trend levels during the 1980s and early 1990s, followed by partial recovery from the mid-1990s. Seasonal amplitude is described as lower during 1983–1986, while the principal wet-season peak remains concentrated around July and August. The source reports no significant residual autocorrelation through lag 24, although numerical diagnostic statistics were not provided.

Port Harcourt exhibits a different reported seasonal profile, with two peaks and a weakening of the October peak relative to the June peak after approximately 2005. However, the stated 90% credible intervals for October amplitudes in 1990–2004 and 2015–2024 overlap. The available summary therefore does not establish a clear persistent reduction. At Ibadan, residual autocorrelation at lags 6 and 12 suggests remaining semiannual and annual dependence after the basic BSTS decomposition.

Periodic correlation diagnostics suggest PAR (1) or PAR (2) for many northern and middle-belt stations, and PAR (2) or PAR (3) for southern stations. Periodic AIC and BIC reportedly agree at 34 of 37 stations; PAR (3) was selected in the three cases of disagreement. The forecast table nevertheless labels the Bayesian PAR comparator as PAR (2). That label is retained in Table 3, but a station-by-station specification list is needed to reconcile it with the reported identification procedure.

The first-lag coefficients are described as larger near the seasonal margins, particularly in months 4–5 and 10–11, and smaller during peak wet-season months. This suggests that persistence may vary with the stage of the seasonal cycle. A physical explanation of this pattern would require additional meteorological covariates rather than being established by the autoregressive coefficients alone.

Forecast Accuracy and Predictive Intervals

Table 3 reproduces the numerical forecast results supplied in the original manuscript. Figure 2 displays the same values to make horizon-specific rankings easier to compare. The relative RMSE changes in Table 4 are calculated directly from these reported, rounded scores.

Table 3: Reported Out-Of-Sample Forecast Accuracy, 2022–2024: Means across 37 Stations

Model	RMSE (mm) h = 12	MAPE (%)h = 12	RMSE (mm) h = 24	MAPE (%)h = 24
SARIMA(1,1,1)(1,1,1)$_{12}$	44.7	18.3	61.2	24.1
FAR(1)	42.1	17.1	58.8	23.0
Bayesian PAR(2)	40.3	16.4	56.1	22.1
BSTS (single model)	39.8	16.2	55.4	21.8
Nonparametric spectral	41.0	16.7	57.3	22.5
BMA	41.0	15.0	52.0	20.5

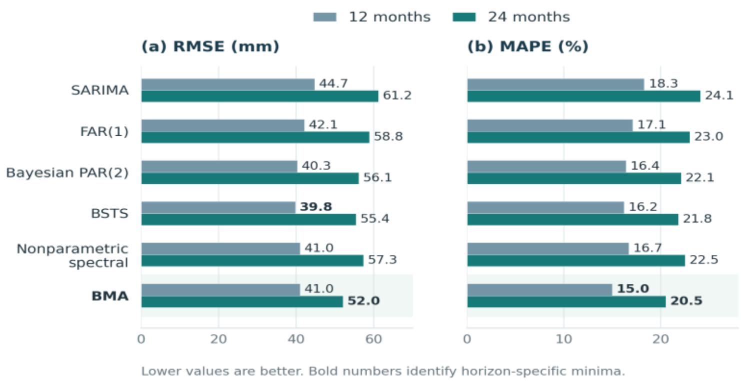

Figure 2: Forecast-Error Comparison at 12- and 24-Month Horizons

(a) RMSE in millimetres; (b) MAPE in percent. Values are transcribed from the original Table 4.1 and reproduced in Table 3. Panel scales differ. BSTS has the lowest 12-month RMSE; BMA has the lowest 24-month RMSE and the lowest reported MAPE at both horizons. No uncertainty bars are shown because none were supplied. MAPE requires clarification of zero-rainfall handling.

At the 12-month horizon, BSTS has the lowest reported RMSE, at 39.8 mm, followed by Bayesian PAR (2), at 40.3 mm. BMA and the nonparametric spectral comparator both have an RMSE of 41.0 mm, ahead of FAR (1), at 42.1 mm, and SARIMA, at 44.7 mm. BMA therefore improves on both classical benchmarks but does not outperform every individual Bayesian model under this metric.

At 24 months, BMA has the lowest reported RMSE, at 52.0 mm, compared with 55.4 mm for BSTS, 56.1 mm for Bayesian PAR(2), 57.3 mm for the spectral comparator, 58.8 mm for FAR(1), and 61.2 mm for SARIMA. Relative to SARIMA, the BMA reductions are 8.3% at 12 months and 15.0% at 24 months. Relative to FAR (1), the corresponding reductions are 2.6% and 11.6%, rather than the approximate 5% and 11% stated in the original narrative.

Table 4: RMSE Reduction from BMA Relative to Each Comparator, Calculated from Table 3

Comparator	12-month reduction (%)	24-month reduction (%)
SARIMA	8.3	15.0
FAR(1)	2.6	11.6
Bayesian PAR(2)	−1.7	7.3
BSTS	−3.0	6.1
Nonparametric spectral	0.0	9.2

Calculated as 100 × (RMSE of comparator − RMSE of BMA) / RMSE of comparator, using the rounded source values. Positive values favour BMA, negative values favour the comparator, and zero indicates a tie. These are percentage reductions, not percentage-point changes.

BMA has the lowest reported MAPE at both horizons: 15.0% at 12 months and 20.5% at 24 months. These rankings differ from the 12-month RMSE ranking because the metrics weight errors differently. They should nonetheless remain provisional until the treatment of zero rainfall is documented.

For nominal 90% predictive intervals, the source reports empirical coverage of 88.9% for BMA and 83.4% for SARIMA. These correspond to shortfalls from nominal coverage of 1.1 and 6.6 percentage points, respectively. BMA is closer to the nominal level in this summary. However, the relevant horizon, station subset, sample size, and interval widths are not specified, so the figures do not establish a complete comparison of calibration and sharpness.

The narrative describes larger model-averaging gains at coastal stations with bimodal seasonality and smaller differences at northern stations with simpler annual profiles. This is a plausible interpretation of model complementarity, but no station-level scores or regional uncertainty intervals were supplied. The geographic pattern is therefore retained as a qualitative reported finding rather than presented as a separately verified regional result.

Structural-break summaries

The most frequently reported breakpoint window is 1987–1989, containing modal posterior break locations at 22 of the 37 stations. A second window, 2001–2004, is reported at 19 stations. Figure 3 expresses these counts as both numbers and proportions of the station network: 59.5% and 51.4%, respectively. A station may contribute to both windows, so the categories are not mutually exclusive and their counts must not be added to obtain the number of affected stations.

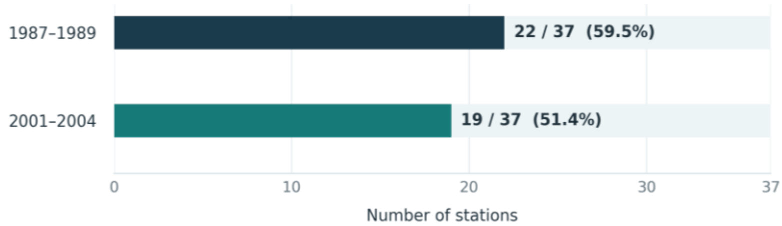

Figure 3: Reported station counts by structural-break window.

Bars show the number of stations with modal posterior break locations within each window, as described in the original Section 4.4 (PDF p. 13). Percentages are calculated using all 37 stations as the denominator. The windows may include the same stations; counts are non-exclusive and do not represent posterior probabilities. Light background bars show the full station network.

The source also discusses spectra conditional on two breaks, using the illustrative segmentation 1981–1988, 1989–2002, and 2003–2024. Annual-frequency spectral mass is described as lowest in the first segment and higher in the later segments. These common segment boundaries should not be treated as the estimated break dates of every station.

Greater annual spectral mass indicates a stronger annual component relative to the chosen spectral representation. It does not, by itself, demonstrate improved out-of-sample predictability. That interpretation would require segment-specific forecasts, a consistent comparison of signal and background variation, and uncertainty estimates. Likewise, associating the detected windows with a particular climatic mechanism would require independent climate evidence and station-history checks. The reported summaries support temporal change, but not statistically unambiguous physical attribution.

Discussion

The forecast comparison supports a horizon-dependent rather than universal advantage of Bayesian model averaging (BMA). At 24 months, BMA records an RMSE of 52.0 mm, compared with 61.2 mm for SARIMA, representing a 15.0% reduction. This finding is consistent with the methodological rationale of Vosseler and Weber (2018), who incorporated uncertainty about autoregressive orders, structural breaks, and break dates into Bayesian forecasting of seasonal series. However, their application concerned German monthly unemployment and averaging within a periodic autoregressive framework, whereas the present framework combines different representations of rainfall dynamics. The agreement therefore concerns the value of accounting for model uncertainty, not evidence that the same mechanism or magnitude of improvement applies across the two studies. Model weights and station-level errors would be needed to determine why the reported advantage is larger at 24 months.

The 12-month results qualify this agreement. BSTS and Bayesian PAR (2) achieve RMSEs of 39.8 and 40.3 mm, respectively, compared with 41.0 mm for BMA. This contrasts with the weather-ensemble application of Raftery et al. (2005), in which the BMA deterministic forecast outperformed the individual ensemble forecasts and the ensemble mean. Their short-range numerical weather forecasts and the present multi-month rainfall forecasts differ in variables, lead times, model composition, and estimation procedures; the contrast is therefore a warning against assuming universal ensemble superiority, rather than a contradiction of their findings. Within the present comparison, BMA also has the lowest reported MAPE at both horizons, illustrating that rankings depend on the loss measure, although its interpretation remains conditional on clarification of the zero-rainfall rule.

The relatively strong BSTS result can also be placed alongside Scott and Varian (2014), whose structural time-series approach represents trend and seasonal variation while incorporating contemporaneous predictors. Their economic-nowcasting framework additionally averages over predictor configurations, whereas the present benchmark is reported as a single univariate BSTS model. The comparison is therefore methodological rather than a direct replication. The present results suggest that an individual structural model remains a credible shorter-horizon competitor; they do not establish which latent component accounts for its advantage or imply that BSTS and model averaging are inherently competing approaches.

The importance of periodic structure is consistent with Taiwo et al. (2019), who selected FAR (1) using periodic information criteria and reported white-noise residuals in their Nigerian rainfall application. In the present benchmark comparison, FAR (1) likewise improves on SARIMA at both horizons: its RMSEs are 42.1 and 58.8 mm, compared with 44.7 and 61.2 mm for SARIMA. BMA nevertheless improves on FAR (1) by 2.6% at 12 months and 11.6% at 24 months. These within-study comparisons support retaining periodic dependence while examining whether additional model flexibility improves prediction. Unlike the residual adequacy reported by Taiwo et al., the remaining dependence at lags 6 and 12 in the present Ibadan BSTS residuals indicates that a basic decomposition may not remove all seasonal structure. Differences in datasets, models, and evaluation designs prevent a direct numerical ranking of the two studies.

The interval results should be interpreted against the probabilistic-forecasting literature. Raftery et al. (2005) reported that BMA improved calibration relative to their raw ensemble and produced sharper intervals than sample climatology. The present coverage figures, 88.9% for BMA and 83.4% for SARIMA at a nominal 90% level are directionally consistent with improved uncertainty representation, but missing interval widths and an unspecified coverage horizon prevent an equivalent claim about calibration and sharpness. Sloughter et al. (2007) further showed the importance of a precipitation-specific BMA formulation that distinguishes zero precipitation from positive amounts. This is directly relevant to the present Gaussian modelling assumptions and the unresolved treatment of dry months. Following Gneiting and Raftery (2007), comparisons using proper scoring rules, such as the continuous ranked probability score, would provide a stronger assessment of the full predictive distributions than RMSE or aggregate coverage alone.

Taken together, these comparisons support flexible seasonal modelling and the explicit treatment of model uncertainty, while showing that their benefits must be assessed by horizon, loss function, and diagnostic performance. The reported contribution is not that BMA invariably dominates individual models, but that its clearest RMSE advantage occurs at the longer horizon in this Nigerian rainfall comparison. This conclusion is narrower, and better supported, than a general claim of universal predictive superiority.

CONCLUSION

This study presents a Bayesian framework integrating spectral estimation, structural decomposition, month-specific autoregression and flexible temporal dynamics for analysing 44 years (1981–2024) of monthly rainfall records from 37 Nigerian stations. The reported results show that BMA reduces RMSE relative to SARIMA by 8.3% at 12 months and 15.0% at 24 months, achieves the lowest 24-month RMSE and records the lowest MAPE at both horizons, although the MAPE comparisons require clarification of zero-rainfall treatment. BSTS, however, achieves the lowest 12-month RMSE, indicating that model averaging is promising particularly for longer-horizon prediction rather than uniformly superior across metrics and horizons. Reported breakpoint windows and station-specific residual patterns further motivate models that allow seasonal dynamics to change, but physical causation, posterior consistency of the full framework and universal predictive optimality are not established by the available evidence. Limitations arising from stationwise Gaussian modelling, incomplete computational and validation documentation, and unavailable supporting outputs warrant reproducible, leakage-free rolling-origin evaluation with robust zero-rainfall handling, probabilistic scoring and sensitivity analysis, together with investigation of spatial hierarchical and non-Gaussian models, dynamic forecast combination, climate covariates and computational efficiency.

REFERENCES

Carter, C. K., & Kohn, R. (1996). Markov chain Monte Carlo in conditionally Gaussian state space models. *Biometrika, 83*(3), 589–601. https://doi.org/10.1093/biomet/83.3.589

Choudhuri, N., Ghosal, S., & Roy, A. (2004). Bayesian estimation of the spectral density of a time series. *Journal of the American Statistical Association, 99*(468), 1050–1059. https://doi.org/10.1198/016214504000000557

Dahlhaus, R. (1997). Fitting time series models to nonstationary processes. *The Annals of Statistics, 25*(1), 1–37.

Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association, 102*(477), 359–378. https://doi.org/10.1198/016214506000001437

Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika, 82*(4), 711–732.

Lund, R. B., & Basawa, I. V. (2000). Recursive prediction and likelihood evaluation for periodic ARMA models. *Journal of Time Series Analysis, 21*(1), 75–93. https://doi.org/10.1111/1467-9892.00174

McAlinn, K., & West, M. (2019). Dynamic Bayesian predictive synthesis in time series forecasting. *Journal of Econometrics, 210*(1), 155–169. https://doi.org/10.1016/j.jeconom.2018.11.010

Nigerian Meteorological Agency (NIMET). (2024). *Monthly rainfall records from 37 climatological stations in Nigeria, 1981–2024* [Data set]. NIMET. [Dataset attribution retained from the supplied manuscript; access details were not provided.]

Pagano, M. (1978). On periodic and multiple autoregressions. *The Annals of Statistics, 6*(6), 1310–1317. https://doi.org/10.1214/aos/1176344376

Rosen, O., Stoffer, D. S., & Wood, S. (2009). Local spectral analysis via a Bayesian mixture of smoothing splines. *Journal of the American Statistical Association, 104*(485), 249–262. https://doi.org/10.1198/jasa.2009.0118

Schwartz, L. (1965). On Bayes procedures. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete, 4*, 10–26.

Scott, S. L., & Varian, H. R. (2014). Predicting the present with Bayesian structural time series. *International Journal of Mathematical Modelling and Numerical Optimisation, 5*(1–2), 4–23.

Taiwo, A. I., Olatayo, T. O., Adedotun, A. F., & Adesanya, K. K. (2019). Modeling and forecasting periodic time series data with Fourier autoregressive model. *Iraqi Journal of Science, 60*(6), 1367–1373. https://doi.org/10.24996/ijs.2019.60.6.20

Vecchia, A. V. (1985). Maximum likelihood estimation for periodic autoregressive moving average models. *Technometrics, 27*(4), 375–384.

Vosseler, A., & Weber, E. (2018). Forecasting seasonal time series data: A Bayesian model averaging approach. *Computational Statistics*, *33*(4), 1733–1765. https://doi.org/10.1007/s00180-018-0801-3

Whittle, P. (1957). Curve and periodogram smoothing. *Journal of the Royal Statistical Society: Series B, 19*(1), 38–63.