

# Mathematical Modelling and Forecasting of Inflation Dynamics in Nigeria Using ARIMA, Prophet, and Markov Chain Models

**Tombotamunoa W. J. LAWSON<sup>1</sup>, Alhassan, C. J.<sup>2</sup>**

<sup>1</sup>Department of Mathematics, Ignatius Ajuru University of Education, Rumuolumeni, Port Harcourt, Nigeria

<sup>2</sup>Department of Mathematics, Edo State University, Iyamho, Nigeria

**Corresponding author:** [tombotamunoa.lawson@iaue.edu.ng](mailto:tombotamunoa.lawson@iaue.edu.ng)

**Abstract.** *This study develops and compares three mathematical frameworks for modelling monthly inflation dynamics in Nigeria: an autoregressive integrated moving-average model, a Prophet additive forecasting model, and a finite-state Markov chain. The analysis uses 281 monthly observations of headline inflation covering January 2003 to May 2026. Preliminary tests indicate that the level series is non-stationary, while first differencing produces a stationary process. A systematic information-criterion search selected an ARIMA(2,1,3) model. A twelve-month holdout experiment showed that ARIMA outperformed Prophet, with mean absolute error, root mean squared error, and mean absolute percentage error of 6.69, 7.60, and 41.08%, respectively, compared with 16.16, 16.92, and 95.39% for Prophet. The relatively large errors reflect a major structural discontinuity near the end of the sample and demonstrate the difficulty of extrapolating inflation under changing measurement and macroeconomic regimes. The Markov model classified inflation as low, moderate, or high and revealed strong state persistence, with self-transition probabilities of 0.913, 0.945, and 0.936. Its stationary distribution assigns probabilities of 0.246, 0.586, and 0.168 to the low, moderate, and high regimes. The combined evidence shows that ARIMA is more effective for short-run numerical forecasting, Prophet is useful for decomposable trend-seasonal representation but is vulnerable to abrupt breaks, and the Markov chain provides interpretable regime probabilities. The study recommends ensemble forecasting, explicit structural-break treatment, and periodic model re-estimation for policy-oriented inflation monitoring in Nigeria.*

**Keywords:** inflation; Nigeria; ARIMA; Prophet; Markov chain; time series; mathematical modelling; forecasting.

## 1. Introduction

Inflation is a central macroeconomic variable because it affects the purchasing power of households, the real value of savings, production costs, interest-rate decisions, exchange-rate expectations, and the allocation of investment. In an economy in which inflation is persistent and occasionally subject to abrupt changes, accurate modelling is difficult but important. Forecasting methods must represent temporal dependence, trend, seasonal behaviour, and possible movement between qualitatively different inflation regimes.

Mathematical time-series models provide a disciplined way to describe these dynamics. The autoregressive integrated moving-average (ARIMA) framework [2, 5] represents a transformed stationary series through lagged observations and lagged shocks. It is parsimonious, statistically interpretable, and widely used for short-horizon forecasting. Prophet [10] represents a time series as the sum of a trend, recurring seasonal components, event effects, and an error term. Its piecewise

trend is useful for gradual changes in growth, although abrupt structural shifts can still reduce forecast accuracy. A Markov chain [8] does not forecast a single numerical value directly. Instead, it describes the probability that inflation moves from one regime to another, thereby providing information about persistence and long-run state occupancy.

The purpose of this paper is to formulate, estimate, and compare ARIMA, Prophet, and Markov chain models for monthly headline inflation in Nigeria. The specific objectives are to: (i) examine the statistical properties of the inflation series; (ii) select an ARIMA specification using information criteria; (iii) estimate a Prophet additive model with piecewise-linear trend and annual Fourier seasonality; (iv) construct a three-state Markov chain for low, moderate, and high inflation; (v) compare out-of-sample forecast accuracy; and (vi) generate twelve-month numerical and regime forecasts.

The paper contributes a unified modelling framework in which numerical forecasts and regime probabilities are studied together. This is useful because point forecasts describe the expected magnitude of inflation, whereas transition probabilities describe the likelihood of remaining in or moving between inflation states.

## 2. Data and Preliminary Analysis

### 2.1. Data structure


The dataset used in this study contains monthly Nigerian inflation indicators. The principal variable used in this study is headline year-on-year inflation. Food and core inflation are retained for descriptive comparison. After chronological ordering, the sample consists of  $N = 281$  observations from January 2003 to May 2026.

Let  $Y_t$  denote headline year-on-year inflation in month  $t$ . The sample mean, standard deviation, minimum, median, and maximum are shown in Table 1.

**Table 1:** Descriptive statistics of monthly headline inflation


Statistic	Value (%)
Number of observations	281
Mean	14.51
Standard deviation	6.59
Minimum	3.00
First quartile	10.00
Median	12.80
Third quartile	17.33
Maximum	34.80

Figure 1 compares headline, food, and core inflation. The three series generally co-move, but food inflation frequently displays larger peaks. The headline series also exhibits long episodes of persistence and a pronounced discontinuity near the end of the sample.



**Figure 1:** Monthly Nigerian headline, food, and core inflation indicators.

Figure 2 shows the rolling mean and rolling standard deviation. Their evolution indicates that both the centre and dispersion of the process vary over time, which motivates differencing and careful validation.



**Figure 2:** Headline inflation with twelve-month rolling mean and standard deviation.

## 2.2. Stationarity tests

The augmented Dickey–Fuller test [3] evaluates the null hypothesis that a unit root is present. For a regression of the form

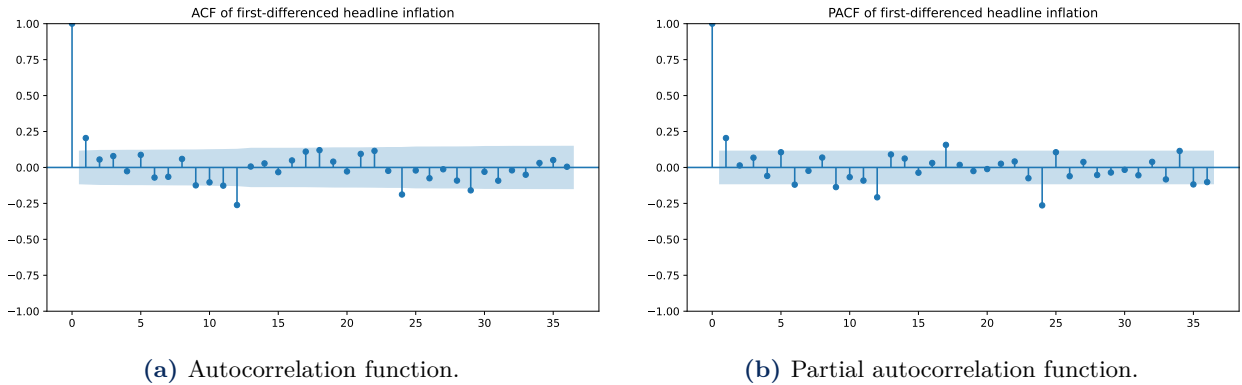
$$\Delta Y_t = \alpha + \gamma Y_{t-1} + \sum_{i=1}^k \beta_i \Delta Y_{t-i} + \varepsilon_t, \quad (1)$$

the null hypothesis is  $H_0 : \gamma = 0$ . The Kwiatkowski–Phillips–Schmidt–Shin test [6] uses stationarity as its null hypothesis and therefore provides a useful complementary diagnostic.

**Table 2:** Stationarity test results

Series	Test statistic	$p$ -value	Decision at 5%
Level, ADF	-1.721	0.421	Fail to reject unit root
First difference, ADF	-3.333	0.013	Reject unit root
Level, KPSS	0.990	0.010	Reject stationarity
First difference, KPSS	0.036	> 0.10	Fail to reject stationarity

Both tests support the conclusion that the level series is non-stationary and that the first-differenced series is stationary. The autocorrelation and partial autocorrelation functions of  $\Delta Y_t$  are displayed in Figures 3a and 3b.



**Figure 3:** Correlation diagnostics for first-differenced headline inflation.

## 3. Mathematical Models

### 3.1. ARIMA model

An ARIMA( $p, d, q$ ) model is written compactly as

$$\phi(B)(1 - B)^d Y_t = c + \theta(B)\varepsilon_t, \quad (2)$$

where  $B$  is the backward-shift operator,  $BY_t = Y_{t-1}$ ,

$$\phi(B) = 1 - \phi_1 B - \phi_2 B^2 - \dots - \phi_p B^p, \quad (3)$$

$$\theta(B) = 1 + \theta_1 B + \theta_2 B^2 + \dots + \theta_q B^q, \quad (4)$$

and  $\varepsilon_t$  is a zero-mean white-noise process with variance  $\sigma_\varepsilon^2$ .

Candidate models with  $0 \leq p, q \leq 3$  and  $d \in \{0, 1\}$  were estimated on the training sample. The Akaike information criterion [1]

$$AIC = -2\ell + 2k \quad (5)$$

and Bayesian information criterion [9]

$$\text{BIC} = -2\ell + k \log n \quad (6)$$

were used to balance model fit and complexity. The selected model was ARIMA(2,1,3), with training-sample AIC = 963.74 and BIC = 988.87.

### 3.2. Prophet model

The Prophet specification decomposes the series as

$$Y(t) = g(t) + s(t) + h(t) + \varepsilon_t, \quad (7)$$

where  $g(t)$  is a trend,  $s(t)$  is seasonality,  $h(t)$  represents event or holiday effects, and  $\varepsilon_t$  is an error term. No separate holiday calendar was introduced in this study, so  $h(t) = 0$ .

A piecewise-linear trend with changepoints  $s_j$  is represented by

$$g(t) = \left( k + \sum_{j=1}^J \delta_j \mathbf{1}_{\{t > s_j\}} \right) t + \left( m + \sum_{j=1}^J \gamma_j \mathbf{1}_{\{t > s_j\}} \right), \quad (8)$$

where  $k$  and  $m$  are the initial growth rate and intercept, and  $\delta_j$  denotes the change in growth after changepoint  $s_j$ . Annual seasonality is represented by a Fourier series,

$$s(t) = \sum_{n=1}^{N_f} \left[ a_n \cos\left(\frac{2\pi n t}{12}\right) + b_n \sin\left(\frac{2\pi n t}{12}\right) \right]. \quad (9)$$

Ten candidate changepoints were distributed over the first 80% of the training period, five Fourier harmonics were used, and ridge regularisation controlled instability in the trend coefficients.

### 3.3. Markov chain model

Headline inflation is classified into three states:

$$S_t = \begin{cases} 1, & Y_t < 10, \\ 2, & 10 \leq Y_t \leq 20, \\ 3, & Y_t > 20. \end{cases} \quad (10)$$

These correspond to low, moderate, and high inflation. The process is assumed to satisfy the first-order Markov property

$$\Pr(S_{t+1} = j \mid S_t = i, S_{t-1}, \dots, S_1) = \Pr(S_{t+1} = j \mid S_t = i) = p_{ij}. \quad (11)$$

The transition probabilities are estimated by

$$\hat{p}_{ij} = \frac{n_{ij}}{\sum_{j=1}^3 n_{ij}}, \quad (12)$$

where  $n_{ij}$  is the observed number of transitions from state  $i$  to state  $j$ . The  $h$ -step state distribution is

$$\boldsymbol{\pi}_{t+h} = \boldsymbol{\pi}_t P^h. \quad (13)$$

A stationary distribution  $\boldsymbol{\pi}$  satisfies

$$\boldsymbol{\pi} = \boldsymbol{\pi} P, \quad \sum_{i=1}^3 \pi_i = 1. \quad (14)$$

### 3.4. Forecast evaluation

The final twelve observations were reserved for out-of-sample evaluation. For actual values  $y_t$  and forecasts  $\hat{y}_t$ , the evaluation measures are

$$\text{MAE} = \frac{1}{n} \sum_{t=1}^n |y_t - \hat{y}_t|, \quad (15)$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{t=1}^n (y_t - \hat{y}_t)^2}, \quad (16)$$

$$\text{MAPE} = \frac{100}{n} \sum_{t=1}^n \left| \frac{y_t - \hat{y}_t}{y_t} \right|. \quad (17)$$

## 4. Results


### 4.1. Out-of-sample comparison

Table 3 reports twelve-month holdout accuracy. ARIMA produces substantially smaller errors than Prophet for every measure.

**Table 3:** Twelve-month out-of-sample forecast accuracy

Model	MAE	RMSE	MAPE (%)
ARIMA(2,1,3)	6.69	7.60	41.08
Prophet	16.16	16.92	95.39

Figure 4 shows that both models over-predict inflation after the abrupt decline in the holdout period. ARIMA adjusts more rapidly because it places greater weight on recent autoregressive and moving-average information. Prophet's smooth piecewise trend extrapolates the earlier high-inflation trajectory and therefore responds more slowly to the break.



**Figure 4:** Observed and forecast headline inflation in the twelve-month holdout period.

The ARIMA residual Ljung–Box statistic [7] at lag 12 is 20.05 with  $p = 0.066$ , so the null hypothesis of no residual autocorrelation is not rejected at 5%. The Jarque–Bera statistic is 1166.28 with  $p < 0.001$ , indicating non-normal residuals. Consequently, the numerical forecasts are informative, but Gaussian prediction intervals should be interpreted cautiously.

#### 4.2. Markov regime estimates


The transition count matrix is

$$N = \begin{pmatrix} 63 & 6 & 0 \\ 6 & 155 & 3 \\ 0 & 3 & 44 \end{pmatrix}, \quad (18)$$

and the estimated transition matrix is

$$\hat{P} = \begin{pmatrix} 0.913 & 0.087 & 0.000 \\ 0.037 & 0.945 & 0.018 \\ 0.000 & 0.064 & 0.936 \end{pmatrix}. \quad (19)$$

The diagonal elements are large, showing strong persistence in every regime. Direct one-month movement from low to high inflation or from high to low inflation is absent in the sample. Movement usually passes through the moderate state.



**Figure 5:** Estimated Markov transition probabilities.


The stationary distribution is

$$\boldsymbol{\pi} = (0.246, 0.586, 0.168), \quad (20)$$

which implies that, under a stable transition mechanism, the process spends approximately 24.6% of the time in the low state, 58.6% in the moderate state, and 16.8% in the high state. Starting from the moderate state observed in May 2026, the twelve-month distribution is

$$\boldsymbol{\pi}_{12} = (0.215, 0.663, 0.122). \quad (21)$$

Thus, the Markov model assigns the greatest probability to moderate inflation over the forecast horizon.




**Figure 6:** Evolution of forecast probabilities for the three inflation states.

### 4.3. Twelve-month numerical forecasts

Table 4 provides forecasts from June 2026 to May 2027. The ARIMA intervals widen with horizon, reflecting accumulated uncertainty. Prophet produces a smoother trajectory because of its trend-seasonal structure.

**Table 4:** Twelve-month headline inflation forecasts

Month	ARIMA	ARIMA lower	ARIMA upper	Prophet
Jun 2026	15.49	12.73	18.25	27.64
Jul 2026	16.49	11.92	21.06	27.70
Aug 2026	16.03	9.90	22.17	27.77
Sep 2026	15.75	8.48	23.01	27.99
Oct 2026	16.59	8.42	24.75	28.19
Nov 2026	15.89	6.75	25.03	28.28
Dec 2026	16.05	6.16	25.95	28.34
Jan 2027	16.57	5.96	27.18	28.33
Feb 2027	15.86	4.49	27.23	28.36
Mar 2027	16.33	4.37	28.30	28.72
Apr 2027	16.47	3.88	29.07	28.71
May 2027	15.94	2.73	29.15	28.92



**Figure 7:** Twelve-month ARIMA and Prophet forecasts based on the complete sample.

## 5. Discussion

The stationarity tests establish that inflation should not be modelled directly as a stable level process over the entire sample. First differencing removes the stochastic trend sufficiently for ARIMA estimation. The selected ARIMA(2,1,3) model captures short-run serial dependence more effectively than the Prophet specification in the holdout experiment.

Nevertheless, the forecast errors are large in absolute and percentage terms. This does not imply that time-series modelling is uninformative. Rather, it indicates that the final portion of the sample contains a break that cannot be inferred reliably from earlier observations alone. Both methods are backward-looking and therefore cannot anticipate a discontinuity caused by a methodological revision, policy intervention, exchange-rate regime change, supply shock, or another external event unless such information is explicitly included.

The Prophet model is designed to accommodate trend changes through changepoints. However, regularised piecewise trends remain gradual unless the training data contain a comparable break. In this application, that smoothness becomes a disadvantage. ARIMA adapts faster because recent errors enter through moving-average terms, but it also remains sensitive to the choice of estimation window.

The Markov model adds a different form of information. Its high diagonal probabilities imply that inflation regimes are persistent. Once inflation enters a high regime, the estimated probability of remaining high in the next month is approximately 0.936. Similarly, moderate inflation has a self-transition probability of 0.945. This persistence means that regime changes should be treated as important events rather than routine month-to-month fluctuations.

For practical forecasting, the three methods should not be viewed as mutually exclusive. ARIMA can provide short-run point forecasts, Prophet can provide a transparent trend-seasonal decomposition, and the Markov chain can quantify regime risk. A combined decision framework could report the ARIMA point forecast together with the Markov probability that inflation will be low, moderate, or high.

## 6. Conclusion

This paper modelled Nigerian monthly headline inflation using ARIMA, Prophet, and a three-state Markov chain. The analysis covered 281 monthly observations from January 2003 to May 2026. The level series was non-stationary, while the first difference was stationary. An AIC-based search selected ARIMA(2,1,3). In a twelve-month holdout test, ARIMA outperformed Prophet, although both models experienced large errors because of an abrupt break near the end of the sample.

The Markov chain showed strong inflation-regime persistence. Its long-run distribution places the greatest mass on moderate inflation, and its twelve-month forecast from May 2026 also favours the moderate state. These findings support four recommendations. First, inflation models should be re-estimated frequently. Second, structural-break indicators and exogenous variables should be incorporated when available. Third, rolling-window estimation should be compared with expanding-window estimation. Fourth, point forecasts should be accompanied by regime probabilities and prediction intervals.

Future research may extend the analysis to seasonal ARIMA, ARIMAX models with exchange rate and money-supply variables, Bayesian structural time-series models, threshold autoregression, hidden Markov models, and weighted forecast combinations.

## Funding

The authors received no specific funding for this work.

## Data Availability

The analysis is based on monthly Nigerian inflation indicators covering January 2003 to May 2026.

## Conflict of Interest


The authors declare no conflict of interest.

## Author Contributions

Tombotamunoa W. J. Lawson conceived the study, developed the mathematical models, conducted the analysis, and prepared the initial manuscript. Alhassan, C. J. contributed to model interpretation, validation, and revision of the manuscript. Both authors reviewed and approved the final manuscript.

## References

- [1] Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6), 716–723.
- [2] Box, G. E. P., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). *Time Series Analysis: Forecasting and Control* (5th ed.). Wiley.
- [3] Dickey, D. A., and Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366), 427–431.
- [4] Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.
- [5] Hyndman, R. J., and Athanasopoulos, G. (2021). *Forecasting: Principles and Practice* (3rd ed.). OTexts.
- [6] Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54(1-3), 159–178.



- 
- [7] Ljung, G. M., and Box, G. E. P. (1978). On a measure of lack of fit in time series models. *Biometrika*, 65(2), 297–303.
- [8] Norris, J. R. (1998). *Markov Chains*. Cambridge University Press.
- [9] Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 461–464.
- [10] Taylor, S. J., and Letham, B. (2018). Forecasting at scale. *The American Statistician*, 72(1), 37–45.
- [11] Tsay, R. S. (2010). *Analysis of Financial Time Series* (3rd ed.). Wiley.