Forecasting global monthly cotton prices: the superiority of NNAR models over traditional models

Rumita Limbu Sanwa*, Raksha Khadka and Yeon Nain Chi

Department of Agriculture, Food, and Resource Sciences, University of Maryland Eastern Shore, Princess Anne, MD, United States

Accurate forecasting of agricultural commodity prices is essential for informed decision-making by farmers, traders, and policymakers. This study evaluates and compares the predictive performance of traditional statistical and machine learning models in forecasting global monthly cotton prices. Price volatility and nonlinear patterns in cotton markets present challenges for conventional models such as the Auto Regressive Integrated Moving Average (ARIMA), which often fail to capture complex dynamics. The novelty of this research lies in systematically comparing traditional statistical models (ARIMA, ETS, STL, TBATS, Theta), machine learning models (Neural Network Auto-Regressive [NNAR]), and hybrid approaches to determine the best forecasting tool. Performance was evaluated using Root Mean Square Error (RMSE), Mean Error (ME), Mean Absolute Error (MAE), Mean Percentage Error (MPE), and Mean Absolute Percentage Error (MAPE). Results revealed that the NNAR (26, 1, 14) [12] model outperformed all models, achieving the lowest RMSE (1.16383774), MAE (0.832275572), and MAPE (1.19%), indicating high predictive accuracy and minimal bias. The 30-month forecast for cotton prices using the NNAR model indicates fluctuations between approximately $0.66 and $0.74 per pound, following a cyclical pattern without a clear long-term trend. These findings highlight the strength of advanced machine learning techniques, particularly NNAR, in capturing complex nonlinear patterns, improving forecasting reliability, and supporting effective decision-making in volatile cotton markets. This study provides practical insights for stakeholders seeking to anticipate cotton price changes and make informed decisions in the global market.

KEYWORDS

NNAR, hybrid models, cotton price forecasting, time series, predictive model

1 Introduction

Cotton (*Gossypium* spp.) is an important cash crop cultivated commercially in more than 50 countries and plays a vital role in supporting livelihoods worldwide (Khadi et al., 2010). Often referred to as "white gold," cotton contributes substantially to foreign exchange earnings in many countries (Khan et al., 2020). Its importance extends far beyond the farm; cotton is at the heart of global industries, economies, and cultures (Jabran et al., 2019). However, its price is highly variable, influenced by seasonal changes, regional differences, global market conditions, and broader economic trends. For farmers and stakeholders in the textile supply chain, such variability creates substantial uncertainty in decision-making (Darekar and Reddy, 2017).

Price volatility in cotton markets poses a serious challenge for producers, traders, and policymakers. Inaccurate price expectations can result in inefficient resource allocation, poor

risk management, and financial instability among farmers (Hudson, 1998; Wang, 2023). As a key raw material for the textile industry, cotton faces production and price uncertainties, posing challenges for farmers. Since prices are shaped by local and international supply and demand, accurate price forecasting has become increasingly important, especially following market liberalization, which has led to more frequent price fluctuations and reduced predictability in cotton farming. In this context, timely and accurate price forecasts play an important role in enhancing market efficiency and supporting the financial stability of cotton producers (Darekar and Reddy, 2017; Kumar et al., 2024).

Accurate price forecasts are crucial for strategic planning in agriculture. It works like a dynamic filter, using past price data to predict future prices. Accurate forecasts help farmers and industries plan their farming activities and budgets, which often depend on expected future prices (Chi, 2021). They enable farmers to optimize crop choices and marketing strategies, while industries and policymakers use them to ensure supply chain stability and maintain competitiveness in global markets (Sun et al., 2023). Price shifts can significantly impact not just livelihoods and food affordability, but also the broader economy and supply chains that rely on affordable raw materials (Wang, 2023). It affects not only what people pay for food, but also how farmers plan their production. Given cotton's pivotal role in the global textile industry, advancements in forecasting models contribute not only to agricultural development but also to economic stability and food security (Zhang et al., 2020). As Hudson (1998) noted, when producers have access to accurate price information, they are better able to decide what and when to produce or sell. This knowledge benefits not only the farmers, but also plant breeders, ginners, merchants, and textile mills who depend on predictable prices to manage supply chains and production costs.

2 Literature review

2.1 Traditional time series models for cotton price forecasting

Cotton prices are inherently volatile, making reliable forecasting models essential for supporting informed decision-making. The Auto Regressive Integrated Moving Average (ARIMA) model, originally proposed by Box and Jenkins (1970), has emerged as a widely used and practical statistical approach for modeling and predicting time series data that follows regular and continuous patterns. Its ability to capture trends, seasonality, and autocorrelations in historical price data makes it particularly suitable for forecasting in agricultural markets such as cotton, where structured, periodic fluctuations are common (Weng et al., 2019). However, its performance tends to decline when dealing with more complex or nonlinear price movements, as noted by Liu (2025). Similarly, Wang et al. (2013) emphasized ARIMA's strengths as a linear forecasting tool for economic time series but also acknowledged its limitations in capturing sudden market fluctuations. To address these limitations, Other linear forecasting models have also been applied to capture trends and seasonal patterns, including the Exponential Smoothing State Space (ETS) model (introduced by Hyndman et al., 2008), the Seasonal and Trend decomposition using Loess (STL) model (introduced by Cleveland et al., 1990), and the Theta model

(introduced by Assimakopoulos and Nikolopoulos, 2000). For time series with multiple or irregular seasonal patterns, the *TBATS* model (*Trigonometric, Box-Cox Transformation, ARMA Errors, Trend, and Seasonal components*) is particularly effective (De Livera, 2010; Hyndman and Athanasopoulos, 2021). It is designed to handle time series with multiple seasonal patterns, irregular cycles, or high-frequency seasonality. De Livera (2010) showed that TBATS can outperform more traditional models when dealing with such complex seasonal behavior. Later, Hyndman and Athanasopoulos (2021) expanded on its capabilities, showing how TBATS can be applied to a wide range of forecasting problems with irregular or multi-seasonal data patterns.

2.2 Machine learning approaches

To overcome these limitations of traditional linear models, researchers have increasingly turned to machine learning (ML) methods such as artificial neural networks (ANNs), which can model complex, nonlinear relationships in economic data (Patel et al., 2020; Raghav et al., 2022). Neural networks have proven to be powerful tools for time series forecasting due to their ability to model complex, nonlinear relationships. Zhang and Qi (2005) demonstrated that they outperform traditional statistical models, particularly when handling time series data with irregular fluctuations and unpredictable patterns. Reza and Debnath (2025) further demonstrated that the *Neural Network Autoregressive (NNAR)* model, first introduced by Zhang (2003), outperformed the ARIMA model in forecasting accuracy. Their study showed that NNAR achieved higher predictive accuracy, as measured by R-squared, and produced lower forecasting errors, highlighting its ability to capture complex and nonlinear patterns in time series data. Similar results were reported by Melina et al. (2024), who found that the NNAR model consistently outperformed ARIMA in price forecasting tasks. The strength of the NNAR model lies in its flexibility. Unlike traditional linear models, it does not assume linearity or normality, making it particularly effective for capturing the nonlinear dynamics often present in agricultural commodity prices. This adaptability makes NNAR valuable for improving forecast reliability in complex and rapidly changing market environments. Recent research has focused on combining neural network models with traditional forecasting methods to enhance accuracy, particularly for datasets exhibiting linear and nonlinear characteristics. Almarashi et al. (2024a) demonstrated that the NNAR model outperformed other models, achieving the best evaluation metrics in time series forecasting. Among various artificial neural network (ANN) architectures, the single hidden layer ANN is frequently used for time series forecasting due to its simplicity and effectiveness (Zhang et al., 1998). A real-world time series may exhibit both linear and nonlinear patterns.

2.3 Hybrid forecasting models

Hybrid forecasting models that combine statistical and machine learning approaches have been shown to enhance predictive performance by combining the strengths of both linear and nonlinear methods. Zhang (2003) introduced the ARIMA-ANN hybrid model, demonstrating that neural networks can effectively complement

ARIMA by capturing nonlinear structures that traditional models fail to represent. This study laid the groundwork for the development of neural network autoregressive (NNAR) models, which apply feed-forward neural networks for time series forecasting. Subsequent studies, such as Wang et al. (2013), Kumar and Thenmozhi (2014), and Babu and Reddy (2014), provided further evidence that hybrid approaches outperform standalone models across financial and economic datasets. More recently, Sun et al. (2023) reaffirmed the potential of hybrid and neural-based approaches in improving the accuracy and robustness of cotton price forecasting. Collectively, these studies support the use of NNAR models as a robust extension of ARIMA-ANN frameworks for modeling complex time series dynamics.

Another hybrid methodology combining linear and nonlinear exponential smoothing models from innovation state space (ETS) with artificial neural networks (ANN) was proposed (Panigrahi and Behera, 2017). The model demonstrated superior forecast accuracy due to their ability to capture diverse linear and nonlinear patterns, leveraging the strengths of both ETS and ANN in modeling. In another approach, Jaiswal et al. (2022) introduced a hybrid model that combines Seasonal-Trend Decomposition using Loess (STL) with an Extreme Learning Machine (ELM), called STL-ELM. This method is particularly useful for forecasting agricultural prices, which often show seasonal and nonlinear trends. Their study found that STL-ELM captured these complex patterns more accurately than traditional methods, improving forecasting performance for challenging real-world data.

Advances in statistical modeling highlight the need for robust and reliable methods in prediction and estimation. For example, Ahmad et al. (2025a) developed an auxiliary variables-based estimator to improve population distribution estimates under stratified sampling and non-response, and Ahmad et al. (2025b) proposed an unbiased ratio estimator for distribution functions in complex sampling settings. While these studies are not directly focused on agricultural markets, they underline the importance of methodological innovation in improving predictive accuracy. Inspired by this principle, the present study evaluates linear, nonlinear, and hybrid forecasting approaches to improve the reliability of global cotton price predictions.

This study seeks to address the following key research questions:

1 How accurately do traditional models such as ARIMA, ETS, Theta, and TBATS perform in forecasting global monthly cotton prices?
2 To what extent can machine learning models such as Neural Network Auto-Regressive (NNAR) improve forecasting performance compared to traditional statistical models?
3 Do hybrid models that combine statistical and machine learning techniques provide more reliable forecasts by capturing both linear and nonlinear patterns in cotton price forecasting?

The primary objectives of this research are:

- To identify and compare the most suitable time series models in forecasting monthly global cotton prices.
- To assess the effectiveness of hybrid forecasting approaches in enhancing prediction reliability for complex agricultural time series.
- To provide practical insights for cotton industry stakeholders by improving the reliability of price forecasts.

2.4 Summary and research gap

Accurate forecasting of cotton prices is essential for producers, traders, and policy makers, yet the underlying time series often exhibit a mixture of linear trends, seasonality, and nonlinear fluctuations. Forecasting methods have evolved from traditional statistical models to advanced machine-learning approaches, each with distinct strengths and limitations. Linear models such as ARIMA and ETS capture autoregressive and seasonal components effectively but tend to perform poorly on irregular or nonlinear patterns (Wang et al., 2013; Weng et al., 2019). Conversely, neural-network-based methods, particularly the Neural Network Auto-Regressive (NNAR) model, are well-suited to modelling nonlinear dynamics and have outperformed ARIMA in various applications (Reza and Debnath, 2025; Almarashi et al., 2024b). However, empirical and methodological studies indicate that NNAR and related models may underperform when a time series contains a dominant linear structure, motivating hybrid strategies that model linear and nonlinear components separately (Zhang, 2003; Khandelwal et al., 2015).

Hybrid frameworks such as ARIMA-ANN, ETS-ANN, and STL-ELM have demonstrated improved forecast accuracy in diverse fields, including economics, stock returns, and agricultural commodities (Wang et al., 2013; Kumar and Thenmozhi, 2014; Sun et al., 2023; Panigrahi and Behera, 2017; Jaiswal et al., 2022). Yet, despite these advances, few studies have systematically compared linear, nonlinear, and hybrid approaches within a single framework for cotton price forecasting. Existing comparative work (e.g., Kumar and Thenmozhi, 2014; Babu and Reddy, 2014) has focused mainly on financial datasets, limited regions, or small samples and generally omitted NNAR alongside hybrid models.

Accordingly, this study addresses this gap by systematically evaluating and comparing the forecasting performance of ARIMA, ETS, Theta, TBATS, STL, NNAR, and hybrid approaches to identify the most effective models for predicting global cotton prices.

3 Materials and methods

3.1 Data

The data used in this study, the monthly global cotton price (PCOTTINDUSDM) from January 1990 to January 2025 (Units: U. S. Cents per Pound, Not Seasonally Adjusted), is sourced from the International Monetary Fund via FRED, the Federal Reserve Bank of St. Louis.[1] The *forecast Hybrid* package in R (version 4.4.2) was used to implement a hybrid model that integrates six forecasting methods: auto. arima, ets, thetam, nnetar, stlm, and tbats. Equal weights (i.e., each model contributed equally to the final forecast) and weights based on cross-validation errors for higher accuracy were assigned to each model. The set.seed () function ensured reproducibility of results, and forecasts for the next 30 months were generated using the forecast () function. Model performance was assessed with standard error metrics, including Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Percentage Error (MPE), Mean Absolute Percentage Error (MAPE), and

---

1 https://fred.stlouisfed.org/series/PCOTTINDUSDM

Mean Error (ME). Forecasts were visualized using the plot () function to compare the predicted values across different models and hybrids. This graphical representation highlighted the relative performance of the models in capturing trends and seasonal patterns.

3.2 Methodology

We employed multiple forecasting models to model the monthly price of cotton, including a hybrid forecasting approach that combines linear and nonlinear time series models. The following models were utilized for prediction:

3.2.1 Autoregressive integrated moving average (ARIMA)

The ARIMA model is widely used for forecasting, incorporating past values (autoregression), differencing to achieve stationarity (d), and past forecast errors (moving average) (Box and Jenkins, 1970). The (d-differencing order) represents the number of differencing steps required to transform a time series into a stationary one. To identify whether the series is stationary and determine the appropriate d value, we applied the Augmented Dickey-Fuller (ADF) test proposed by Dickey and Fuller (1981). However, selecting the optimal parameters (p-autoregressive order), (d-differencing order), and (q-moving average order) can be complex and require statistical expertise. In practice, automatic forecasting for large univariate time series is often needed, especially in business contexts. According to Hyndman and Athanasopoulos (2018), the auto. Arima () function in R utilizes the Hyndman-Khandakar algorithm (Hyndman and Khandakar, 2008). This algorithm simplifies model selection by integrating unit root tests to determine differencing requirements, minimizing the corrected Akaike Information Criterion (AICc) for model selection, and applying Maximum Likelihood Estimation (MLE) to optimize parameters, ensuring the most suitable ARIMA model is identified.

The general formula for the ARIMA (p, d, q) model, based on the foundational work of Box and Jenkins (1970), is expressed as:

$$\phi_p(B)(1-B)^d y_t = C + \theta_q(B)\epsilon_t \quad (1)$$

Where:

- B is the backshift operator $B\, y_t = y_{\{t-1\}}$
- $\phi_p(B)$ represents the autoregressive (AR) polynomial of order p:

$$\phi_{p(B)} = 1 - \phi_1 B - \phi_2 B^2 - \ldots - \phi_p B^p$$

- $(1-B)^d$ represents the differencing in order d.
- C is a constant term
- $\theta_q(B)$ represents the moving average (MA) polynomial of order q:

$$\theta_{q(B)} = 1 - \theta_1 B - \theta_2 B^2 - \ldots - \theta_q B^q$$

- ϵ_t is the error term, which is assumed to be white noise.

3.2.2 Seasonal autoregressive integrated moving average model (SARIMA) model

Considering the seasonal pattern exhibited by the monthly cotton price, a seasonal process may be considered; therefore, the ARIMA model will become a Seasonal Autoregressive Integrated Moving Average (SARIMA) model. The SARIMA model extends the ARIMA model by incorporating seasonal effects, making it well-suited for analyzing such time series data. Khadka and Chi (2024) demonstrated that the SARIMA model effectively captures seasonal patterns and trends in agricultural price data.

The SARIMA model is denoted as ARIMA (p, d, q) (P, D, Q) S and has the following specification based on the backshift operator

- (p, d, q) denote the non-seasonal autoregressive order, differencing order, and moving average order.
- (P, D, Q) represent the seasonal components of autoregression, differencing, and moving average.
- s refers to the length of the seasonal cycle (e.g., s = 12 for monthly data).

This formulation enhances the model's flexibility in capturing complex patterns by incorporating seasonal and non-seasonal behaviors within the same structure (Box et al., 2015). The general form of the SARIMA model using the backshift operator B is Chang et al. (2012):

$$\Phi_P(B^s)\phi_p(B)(1-B)^d(1-B^s)^D y_t = C + \Theta_Q(B^s)\theta_q(B)\epsilon_t \quad (2)$$

Where,

$\phi_p(B)$: non-seasonal autoregressive polynomial.

$\theta_q(B)$: non-seasonal moving average polynomial.

$\Phi_P(B^s)$: seasonal autoregressive polynomial.

$\Theta_Q(B^s)$: seasonal moving average polynomial.

$(1-B)^d$: non-seasonal differencing.

$(1-B^s)^D$: seasonal differencing.

C: constant (intercept term).

ϵ_t: white noise error term

3.2.3 Neural network autoregression (NNAR) model

Artificial neural networks are forecasting techniques inspired by simplified mathematical representations of the brain's functioning, enabling them to capture complex nonlinear relationships between variables. The NNAR (Neural Network Autoregressive) model integrates traditional autoregressive (AR) models with artificial neural networks (ANN) to capture nonlinear relationships in time series data. This model utilizes lagged values of the series as input variables

for a feed-forward network, where a hidden layer introduces nonlinearity, and the output layer generates predictions for future observations. The NNAR framework is designed to detect intricate patterns in data, making it highly suitable for forecasting scenarios involving nonlinear relationships (Hyndman and Athanasopoulos, 2021).

The NNAR (p, k) model, which uses lagged values as inputs and k hidden neurons, can be represented as:

$$Y_t = f\left(\sum_{\{i=1\}}^{\{p\}} w_i Y_{\{t-i\}} + b\right) \tag{3}$$

Where:

- Y_t is the predicted value at time t,
- $f(\cdot)$ is the non-linear activation function for the hidden layer
- The sum represents the weighted combination of the lagged inputs
- w_i are the weights associated with the lagged values $Y_{\{t-i\}}$
- b is the bias term for the hidden layer.

The NNAR (p, P, k)m model is represented with p as the number of input lags, P as the seasonal lags, k as the number of neurons in the hidden layer, and m as the seasonal period length (Hyndman and Athanasopoulos, 2018). The mathematical representation of the NNAR (p, P, k) m model is given as follows:

$$y_t = f(y_{\{t-1\}}, y_{\{t-2\}}, \ldots, y_{\{t-p\}}, y_{\{t-m\}}, \\ y_{\{t-2m\}}, \ldots, y_{\{t-Pm\}}, \theta) + \epsilon_t \tag{4}$$

Here, f represents the neural network with k hidden nodes in a single layer, and ϵ_t is the residual series.

The NNAR (p, P, k) m model is implemented in R through the nnetar () function in the "forecast" package, which automatically selects appropriate values for p, P, and k based on criteria like the Akaike Information Criterion (AIC). For seasonal time series, it sets $p = 1$ and selects p based on the optimal linear model fitted to the seasonally adjusted data. The default number of hidden nodes k is calculated as $k = (p + P + 1)/2$, rounded to the nearest integer (Hyndman and Athanasopoulos, 2021).

3.2.4 Exponential smoothing state space model (ETS)

Exponential smoothing methods generate forecasts by applying exponentially decreasing weights to past observations, prioritizing more recent data (Holt, 1957; Brown, 1959; Winters, 1960). This methodology enables efficient and reliable predictions across diverse time series, making it particularly advantageous for industrial applications. As described by Hyndman et al. (2008), ETS models systematically account for various patterns in time series data while allowing for different combinations of error structures, trend components, and seasonal adjustments. The method selection is typically determined by identifying the main components of the time series, such as trend and seasonality, and how these components are incorporated into the smoothing method (e.g., additive, damped, or multiplicative). Each model includes a measurement equation representing the observed data and state equations describing how the unobserved components (level,

trend, seasonal) evolve over time. As a result, these models are referred to as state space models (Hyndman and Athanasopoulos, 2018).

The general ETS framework is given as ETS (A, N, N), ETS (M, Ad, N), etc., where:

- Error (E): Additive (A) or Multiplicative (M)
- Trend (T): None (N), Additive (A), Multiplicative (M), or Additive damped (Ad)
- Seasonality (S): None (N), Additive (A), or Multiplicative (M)

Mathematical Formulation of ETS (M, Ad, N) (Hyndman and Athanasopoulos, 2018):

$$l_t = \alpha y_t + (1-\alpha)\left(l_{\{t-1\}} + \phi\, b_{\{t-1\}}\right)$$

$$b_t = \beta\left(l_t - l_{\{t-1\}}\right) + (1-\beta) b_{\{t-1\}}$$

$$\hat{y}_{t+h} = \ell_t + h\, b_t \tag{5}$$

Where:

- l_t = level at time t,
- b_t = trend at time t,
- y_t = observed value at time t,
- α, β, ϕ = smoothing parameters
- $\hat{y}_{t+h}$ = forecasted value at time t + h
- h = forecast horizon

No seasonal equation exists since this model has no seasonal component (N). The model ETS (M, Ad, N), used in the given forecast, includes a multiplicative error structure, an additive damped trend, and no seasonality. Pegels (1969) introduced a method that incorporated a multiplicative trend. This approach was later expanded by Gardner (1985) to include methods with an additive-damped trend.

3.2.5 TBATS (trigonometric, Box-Cox, ARMA, trend, seasonal)

The BATS model (Box-Cox, ARMA, Trend, Seasonal) is a time series forecasting method for complex seasonal and trend patterns. It stabilizes variance using Box-Cox transformations, models autocorrelations with ARIMA, captures long-term trends, and incorporates seasonal components for periodic fluctuations. Additionally, it uses trigonometric functions to account for non-annual and multiple seasonalities, making it suitable for non-linear data.

The TBATS model introduced by De Livera (2010) extends Exponential Smoothing State Space Models (ETS) to handle multiple seasonality, including integer, non-integer, and dual calendar effects. It uses a state-space framework based on the innovations approach, enhancing flexibility for complex seasonal structures. TBATS can model multiple seasonal periods and long seasonal cycles, with Fourier terms enabling it to capture non-linear seasonality alongside traditional effects.

The framework offers several benefits: (i) a larger parameter space for improved forecasts (Hyndman et al., 2008), (ii) support for both nested and non-nested seasonal components; (iii) handling of nonlinear features standard in real-time series; (iv) accounting for autocorrelation in residuals; and (v) a more straightforward, more

TABLE 1 Summary statistics of the global monthly price of cotton (1990–2025).

Summary	Value
Minimum	$37.22
Maximum	$229.67
Mean	$77.857
Median	$75.93
Standard deviation	$24.794
Lower quartile	60.44
Upper quartile	86.98

efficient estimation procedure. In addition, it is shown that the TBATS models can be used to decompose complex seasonal time series into trend, seasonal, and irregular components. In decomposing time series, the trigonometric approach has several important advantages over the traditional seasonal formulation.

A TBATS (p, ω_i, ϕ, m_i) consists of:

- T: Trigonometric seasonality (Fourier terms).
- B: Box-Cox transformation.
- A: ARMA errors.
- T: Trend.
- S: Seasonal components.

The model is formulated as follows:

a. Box-Cox transformation

$$y_t^{\{(\lambda)\}} = \{\frac{\left(y_t^\lambda - 1\right)}{\lambda}, \log\left(y_t\right), if \ \lambda \neq 0 \ and \ if \ \lambda = 0\}$$

Where λ is the Box-Cox transformation parameter.

b. Local trend component

$$\ell_t = \ell_{\{t-1\}} + \phi b_{\{t-1\}} + \alpha \varepsilon_t$$

$$b_t = \phi b_{\{t-1\}} + \beta \varepsilon_t$$

Where:

- ℓ_t is the level,
- b_t is the trend,
- ϕ is the damping parameter,
- α, β smoothing parameters,
- ε_t is the error term.

c. Trigonometric seasonal component

For each seasonal period m_i, define:

$$s_{\{t,i\}} = \sum_{\{k=1\}}^{\{K_i\}} \gamma_{\{i,k\}} \cos\left(\frac{\{2\pi k\, t\}}{\{m_i\}}\right) + \omega_{\{i,k\}} \sin\left(\frac{\{2\pi k\, t\}}{\{m_i\}}\right)$$

Where:

- $s_{\{t,i\}}$: seasonal effect at time t, for seasonality i
- $\gamma_{\{i,k\}}, \omega_{\{i,k\}}$: Parameters
- m_i: seasonal period for component i
- K_i: number of harmonics for seasonality i

d. ARMA error component

$$\varepsilon_t = \sum_{\{j=1\}}^{\{p\}} \phi_j \varepsilon_{\{t-j\}} + \sum_{\{j=1\}}^{\{q\}} \theta_k \eta_{\{t-k\}} + \eta t$$

Where:

- ϕ_j are AR terms,
- θ_k are MA terms,
- ηt is white noise.

The full TBATS model combines these components:

$$y_t^{\{(\lambda)\}} = \ell_t + s_{\{t,1\}} + s_{\{t,2\}} + \ldots + s_{\{t,I\}} + \varepsilon_t \qquad (6)$$

This allows TBATS to handle multiple seasonal periods, long seasonal cycles, and damping trends, making it highly effective for complex and irregular seasonal data.

3.2.6 Theta model

The Theta model, introduced by Assimakopoulos and Nikolopoulos (2000), enhances forecasting accuracy by decomposing a time series into two or more Theta lines, which are independently extrapolated before being recombined to produce forecasts. The model's core principle is decomposing the series into linear and curvature-adjusted components. It performed well using Theta = 0 (linear trend) and Theta = 2 (enhanced local curves), especially for monthly and microeconomic data.

One common version of the formula used in this model is:

$$y_t = \theta_1 \cdot y_t + \theta_2 \cdot \{Trend\}_t \qquad (7)$$

Where:

- y_t : The original time series value at time t,
- θ_1: Theta coefficient for the original series,
- θ_2: Theta coefficient for the trend component,
- $\{Trend\}_t$: The extracted trend at time t,

- In the Theta approach, different values of θ (like $\theta = 0$, $\theta = 2$) are used to generate separate Theta lines, which are then extrapolated and combined.

3.2.7 Seasonal and trend decomposition using LOESS (STL)

It is a robust and versatile method for decomposing time series data into its fundamental components: trend, seasonal, and remainder (residual) components. Introduced by Cleveland et al. (1990), STL employs locally weighted regression (Loess) to estimate both the trend and seasonal components, allowing it to adapt to changes over time. This decomposition helps in understanding underlying patterns and improving forecasting accuracy. Unlike classical decomposition methods, STL can effectively handle missing values, changing seasonality, and outliers.

The STL method has been widely adopted across various fields due to its adaptability and effectiveness. Hyndman and Athanasopoulos (2021) highlight STL's capability to handle any type of seasonality, making it a preferred choice over classical decomposition methods. Furthermore, STL's integration into statistical software packages, such as R's forecast package, has facilitated its application in practical forecasting scenarios.

STL model formula:

$$Y_t = T_t + S_t + R_t \tag{8}$$

Where:

- Y_t is the observed value at time t.
- T_t represents the trend component.
- S_t denotes the seasonal component.
- R_t is the remainder (residual) component.

3.2.8 Hybrid forecasting models

In addition to individual models, we employed a hybrid model, combining several of the above approaches to enhance predictive accuracy. Specifically, the methodology integrates auto. Arima, ets, thetam, nnetar, stlm, and tbats models, implemented using the forecast and smooth R packages. The auto.arima function automatically selects the best ARIMA model, while ets uses exponential smoothing for trend and seasonality. The thetam model combines exponential smoothing and autoregressive components, and nnetar utilizes a neural network for autoregressive forecasting. Stlm decomposes time-series data into seasonal, trend, and remainder components, and tbats handles complex seasonalities and trends. The hybrid model aggregates the outputs from these individual models, applying equal weights to each, providing a robust and accurate forecast by leveraging the strengths of each technique. This methodology offers a comprehensive approach to forecasting time-series data with varying patterns (Hyndman and Khandakar, 2008; Wang et al., 2013).

3.3 Model evaluation metrics

In this study, the accuracy of individual and hybrid models was evaluated on the in-sample performance (i.e., using the fitted one-step-ahead forecast). Forecast performance was evaluated with five widely used error metrics: Mean Error (ME), Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and Mean Percentage Error (MPE). These measures are standard in forecasting research and are particularly useful for comparing different models on the same dataset (Hyndman and Koehler, 2006). Since each metric has its own strengths and limitations, we considered them collectively to ensure a more balanced evaluation of model performance (Perone, 2022). Model evaluation remains a critical but challenging step, as no single metric is universally appropriate across all contexts; the choice depends on data characteristics and the forecasting objectives.

The accuracy of individual and hybrid models was assessed using five commonly used error metrics:

i. Mean error (ME):

$$ME = \frac{1}{n}\sum_{t=i}^{n}\left(y_i - \widehat{y_i}\right) \tag{9}$$

ii. Mean absolute error (MAE):

$$ME = \frac{1}{n}\sum_{t=i}^{n}|\,y_i - \widehat{y_i}\,| \tag{10}$$

iii. Root mean squared error (RMSE):

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=i}^{n}\left(-\widehat{y_i}\right)\wedge 2} \tag{11}$$

iv. Mean absolute percentage error (MAPE) (Kim and Kim, 2016):

$$MAPE = \sum_{t=i}^{n}|\frac{\left\{y_i - \widehat{y_i}\right\}}{y_i}| \tag{12}$$

v. Mean percentage error (MPE):

$$MPE = \frac{100}{n}\sum_{t=i}^{n}\frac{\left\{y_i - \widehat{y_i}\right\}}{y_i} \tag{13}$$

Where n is the number of observations y_i represents the actual values and $\widehat{y_i}$ are the predicted values.

4 Results and discussion

The descriptive analysis of the data shows that the average monthly Global Price of Cotton (2025) was $77.857 U. S. Cents per

FIGURE 1
Time series plot of the monthly global price of cotton (1990/1/1−2025/1/1) (source: own work).

FIGURE 2
Additive time series decomposition showing observed data, trend, seasonal, and random components of global monthly cotton price (source: own work).
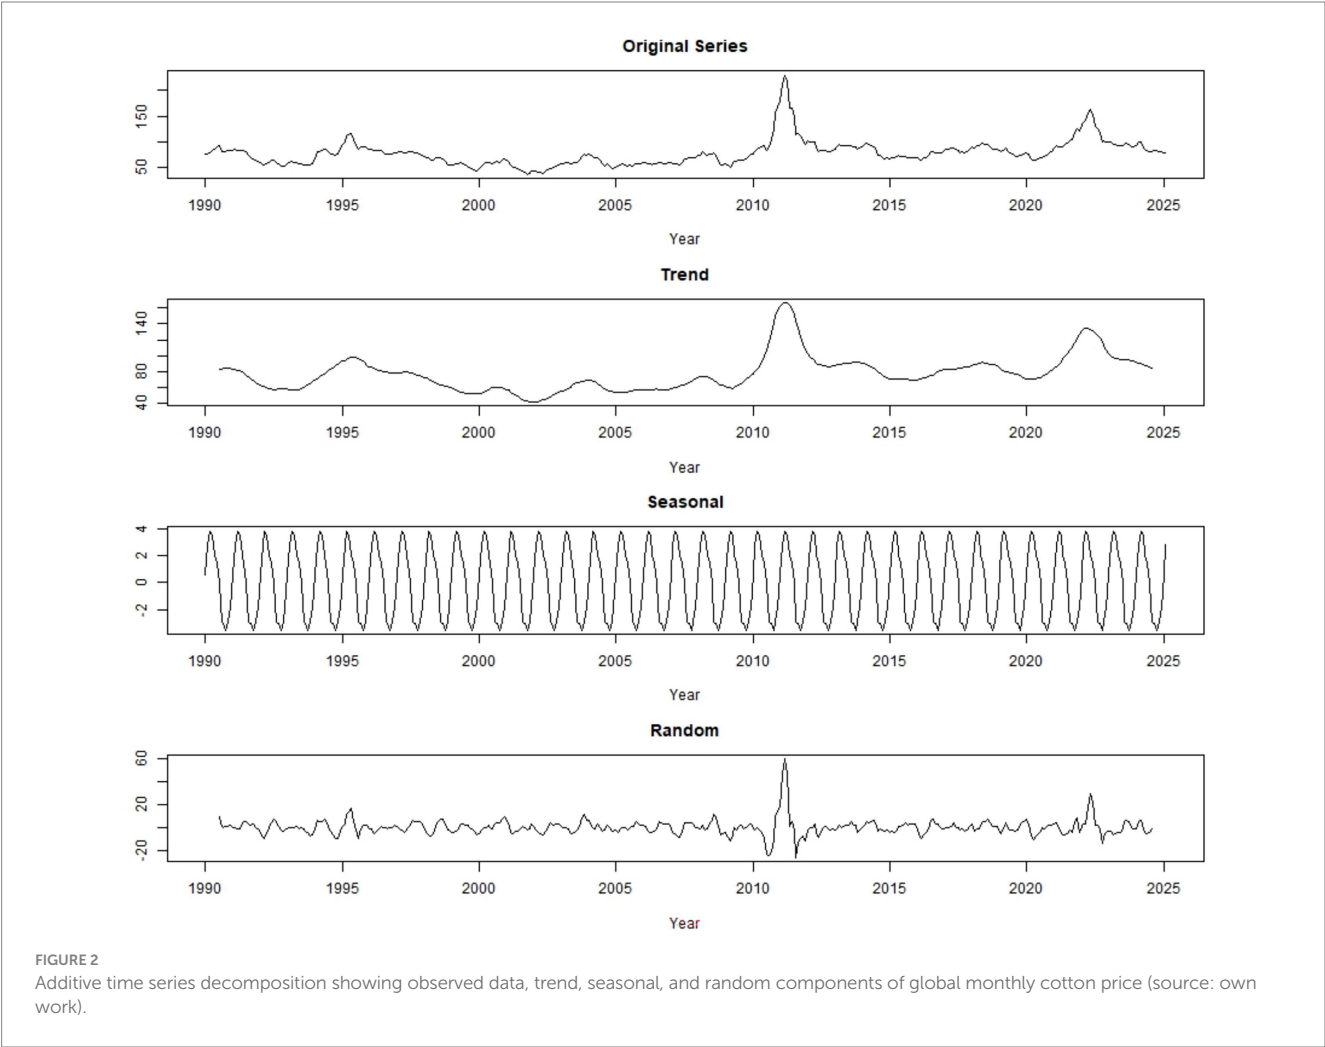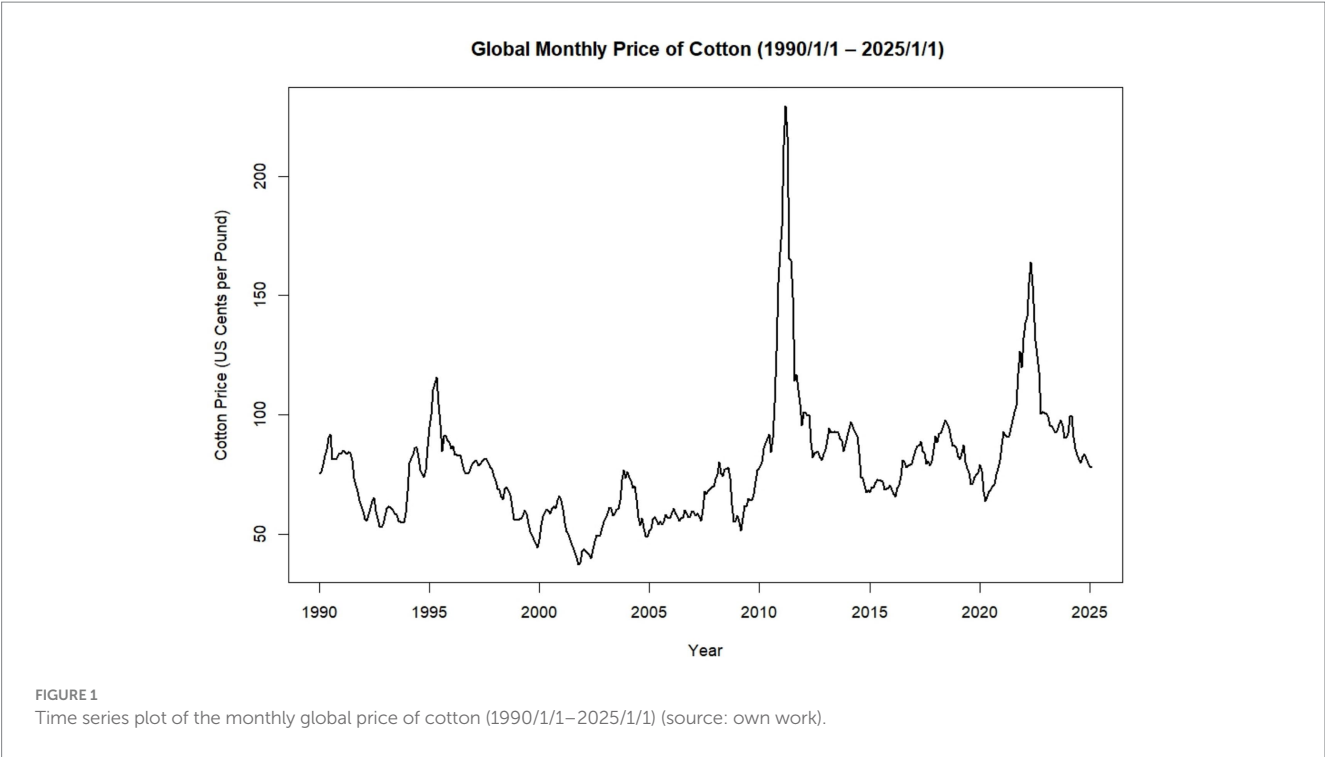
Pound with a standard deviation of $24.794. The lower and upper quartiles of the dataset were 60.44 and 86.98, respectively. The price range ranges from a minimum of $37.22 to a maximum of $229.67, with a median of $75.93. The data exhibits some degree of variability, with the minimum and maximum values indicating significant fluctuations over the observed period.

Table 1 and Figure 1 shows the trend in global monthly cotton prices from January 1990 to January 2025. Throughout this period, the time series exhibits clear fluctuations, with two major price peaks observed around 2010–11 and 2022–23. While certain intervals display relative stability, the overall pattern is characterized by variability and irregular movements in price (Figure 2).

Seasonal adjustment removes regular seasonal patterns from a time series to make the underlying trends and unexpected changes easier to see. By subtracting the seasonal effects from the original data, the adjusted series shows how values change over time without the usual ups and downs, highlighting long-term trends and unusual shifts.

Figure 3 shows the data after removing recurring seasonal patterns, allowing the underlying trends and unusual fluctuations to be more clearly observed. By eliminating regular seasonal ups and downs, the adjusted series highlights long-term movements and unexpected changes, providing a clearer view of the true dynamics in the data over time.

The forecasting performance of various models was evaluated using multiple error metrics, including Mean Error (ME), Root Mean Square Error (RMSE), Mean Absolute Error (MAE), Mean Percentage Error (MPE), and Mean Absolute Percentage Error (MAPE). The results are presented in Table 2.

4.1 Model performance

Table 2 summarizes the forecasting performance of both individual and hybrid models based on the evaluation using ME, RMSE, MAE, MPE, and MAPE. Among all models, the NNAR (26,1,14) [12] model showed the best overall performance, achieving the lowest RMSE (1.164), MAPE (1.192%), and MAE (0.832). This indicates its strong predictive accuracy and ability to effectively capture nonlinear and seasonal patterns in global cotton prices. The model used 26 lagged observations and 14 neurons in a single hidden layer, with a seasonal period of 12 to represent annual variations.

In contrast, the THETA model performed the weakest, recording the highest RMSE (5.668), MAE (3.329), and MAPE (4.014%), reflecting poor predictive capability. Among the hybrid models, Hybrid 6 (ETS × TBATS) had the largest errors (RMSE = 5.004, MAE = 2.996, MAPE = 3.625%), while Hybrid 2 (ARIMA × NNAR) showed the highest mean error (ME = 0.084), suggesting a stronger systematic bias. Although combining multiple modeling techniques, the hybrid models did not enhance forecast accuracy and generally underperformed relative to the NNAR model.

The results indicate that the global cotton price series was primarily driven by nonlinear seasonal dynamics, which were effectively captured by the NNAR model. This suggests that neural network-based approaches are well-suited for modeling commodity price behavior when nonlinearities dominate. However, hybrid models may offer additional advantages in more complex or highly volatile datasets, particularly when both linear and nonlinear
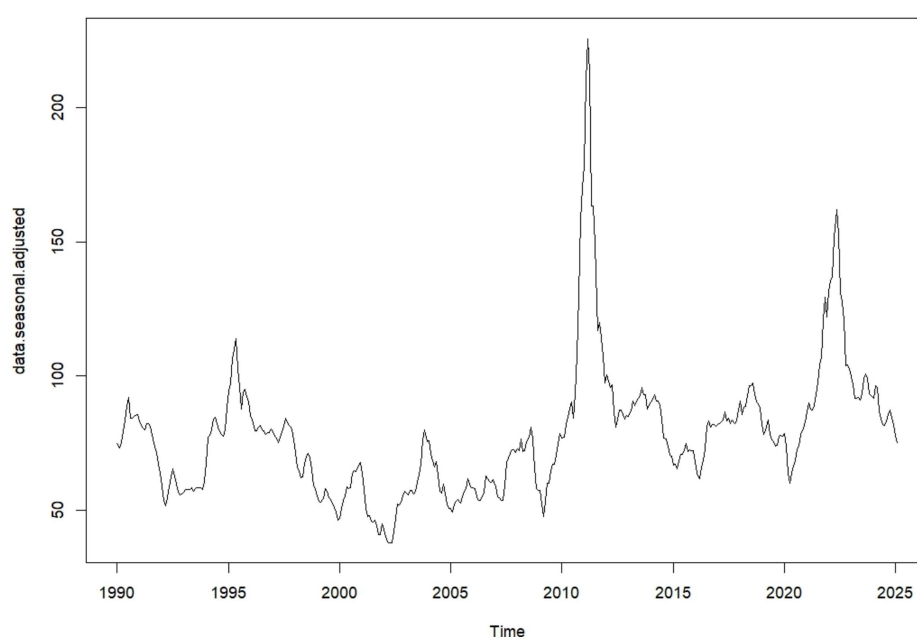

FIGURE 3
Seasonally adjusted time series plot of the global monthly cotton price (source: own work).

TABLE 2 Performance of individual and hybrid models based on multiple evaluation metrics for identifying the best-fit model.

Models	ME	RMSE	MAE	MPE	MAPE
EQUAL	0.042	4.264	2.576	−0.008	3.123
CV. ERRORS	0.043	4.290	2.591	−0.012	3.139
AUTO. ARIMA(4,1,1) (0,0,2) [12]	0.116	4.819	2.989	−0.086	3.626
ETS (M, Ad, N)	−0.014	5.182	3.044	0.017	3.711
NNAR (26,1,14) [12]	−0.008	*1.164*	*0.832*	−0.064	*1.192*
THETA	0.000	5.668	3.329	*−0.147*	4.014
STL +ETS (M, Ad, N)	−0.015	4.768	2.942	0.012	3.617
TBATS (0, 0, 0, −, −)	*−0.054*	5.007	3.047	−0.123	3.658
Hybrid 1 (ARIMA X ETS)	0.055	4.909	2.951	−0.038	3.589
Hybrid 2 (ARIMA X NNAR)	0.084	2.820	1.875	−0.041	2.336
Hybrid 3 (ARIMA X TBATS)	0.029	4.858	2.971	−0.105	3.582
Hybrid 4 (ARIMA X STL)	0.055	4.650	2.828	−0.040	3.443
Hybrid 5 (ETS X NNAR)	0.013	2.968	1.883	0.012	2.351
Hybrid 6 (ETS X TBATS)	−0.036	5.004	2.996	−0.059	3.625
Hybrid 7 (ETS X STL)	−0.014	4.916	2.946	0.014	3.596
Hybrid 8 (NNAR X TBATS)	−0.006	2.947	1.923	−0.059	2.378
Hybrid 9 (NNAR X STL)	0.013	2.875	1.893	0.015	2.377
Hybrid 10 (TBATS X STL)	−0.036	4.771	2.887	−0.061	3.493
Hybrid 11 (ARIMA X ETS X NNAR)	0.064	3.465	2.168	−0.002	2.665
Hybrid 12 (ARIMA X ETS X STL)	0.060	3.997	2.450	0.022	3.000
Hybrid 13 (ARIMA X ETS X STL)	0.033	4.791	2.869	−0.022	3.494
Hybrid 14 (ARIMA X ETS X NNAR X STL X TBATS)	0.042	4.110	2.509	0.003	3.051

The bold values of RMSE (1.164), MAPE (1.192%), and MAE (0.832) for the NNAR (26,1,14) [12] model showed minimal error and the highest accuracy among all the models, indicating the best overall performance. The italics values showed minimal error and the highest accuracy among all the models, indicating the best overall performance.

components jointly influence price movements. By combining complementary modeling techniques, hybrid approaches can improve resilience and accuracy in the presence of noise, sudden structural changes, or shifts in market regimes.

In Figure 4 left panel displays forecasts obtained by equal weighting the six models (ARIMA, ETS, ThETAM, NNETAR, STLM, TBATS), while the right panel shows forecasts using cross-validation (CV) error-based weights. Weights of individual models are shown in the top right corner of each panel, reflecting their contribution to the hybrid forecast. Equal weighting results in uniform contributions (0.167 each), whereas CV error weighting adjusts contributions based on model accuracy, emphasizing models with lower prediction error. The black line represents historically observed cotton prices, and the blue line indicates the model forecast. Compared to equal weighting, the CV-based method provides a data-driven weighting structure that may improve forecast reliability depending on model performance (Figure 5).

A. The AUTO. ARIMA(4,1,1) (0,0,2) [12] model, with a weight of 0.171 in CV-based averaging, showed moderate price fluctuations and increasing forecast uncertainty over time. It's higher error values (RMSE = 4.819, MAE = 2.989, MAPE = 3.626) compared to the NNAR model indicate less accurate predictions and struggles with nonlinear and complex patterns in cotton prices.

B. The ETS (M, Ad, N) model, assigned a weight of 0.154 in CV-based model averaging, underperformed relative to the NNAR model, as indicated by higher error metrics (RMSE = 5.182, MAE = 3.044, MAPE = 3.711). This underperformance is likely due to the model's lack of a seasonal component and it's limited ability to capture nonlinear and complex patterns present in cotton prices.

C. Among all the models, the NNAR (26,1,14) [12] model, which received a weight of 0.159 in cross-validated model averaging, emerged as the best performer. It achieved the lowest evaluation metrics: RMSE (1.164), and MAE (0.832), indicating the smallest average error magnitudes, and MAPE (1.192) showing minimal relative percentage errors, with minimal bias (ME = −0.008, MPE = −0.064). By effectively capturing nonlinear patterns and short-term seasonal fluctuations in global monthly cotton prices, the NNAR model provided stable forecasts with a slight downward adjustment.

D. The THETA model, a simple decomposition-based method, was assigned a weight of 0.179 in cross-validated model averaging. Its bias metrics (ME = 0.000, MPE = −0.147) indicate minimal systematic error; however, its overall forecast accuracy is limited, as reflected by the highest RMSE (5.668), MAE (3.329), and MAPE (4.014) among all models. These results highlight that simplistic decomposition methods are inadequate for modeling complex, nonlinear time series as compared to the NNAR model.

E. The STLM [STL + ETS (M, Ad, N)] model, assigned a weight of 0.155 in cross-validated model averaging, exhibited a minor under prediction bias (ME = −0.015, MPE = −0.123). However, its forecast accuracy was lower than that of the NNAR model, as reflected by higher error values (RMSE = 4.768, MAE = 2.942, MAPE = 3.617). This performance suggests that the STLM model is less capable of
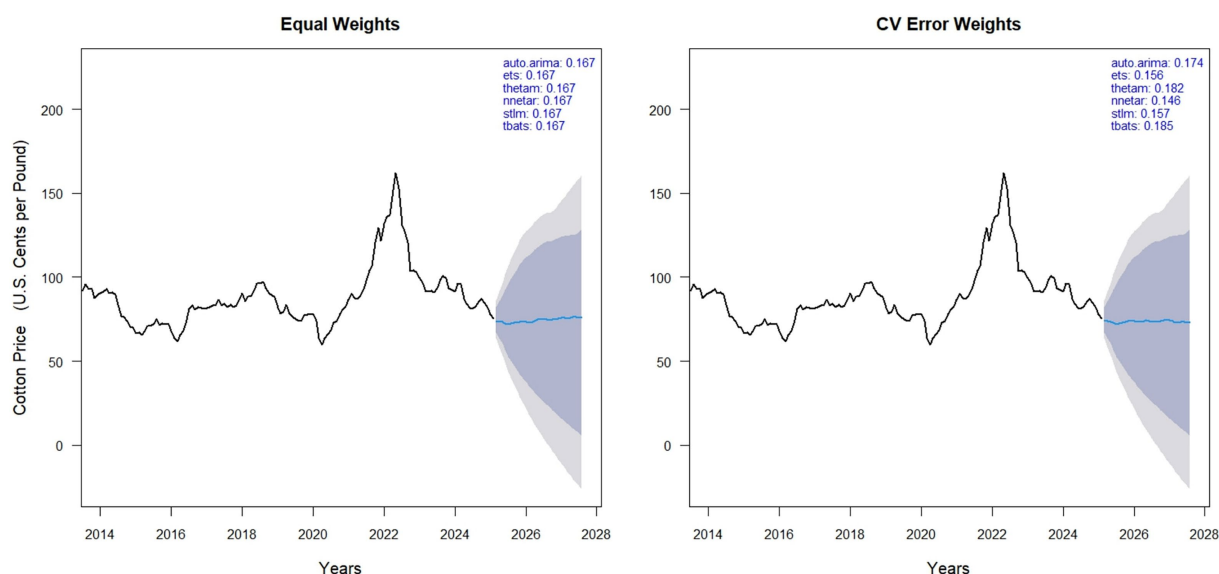
FIGURE 4
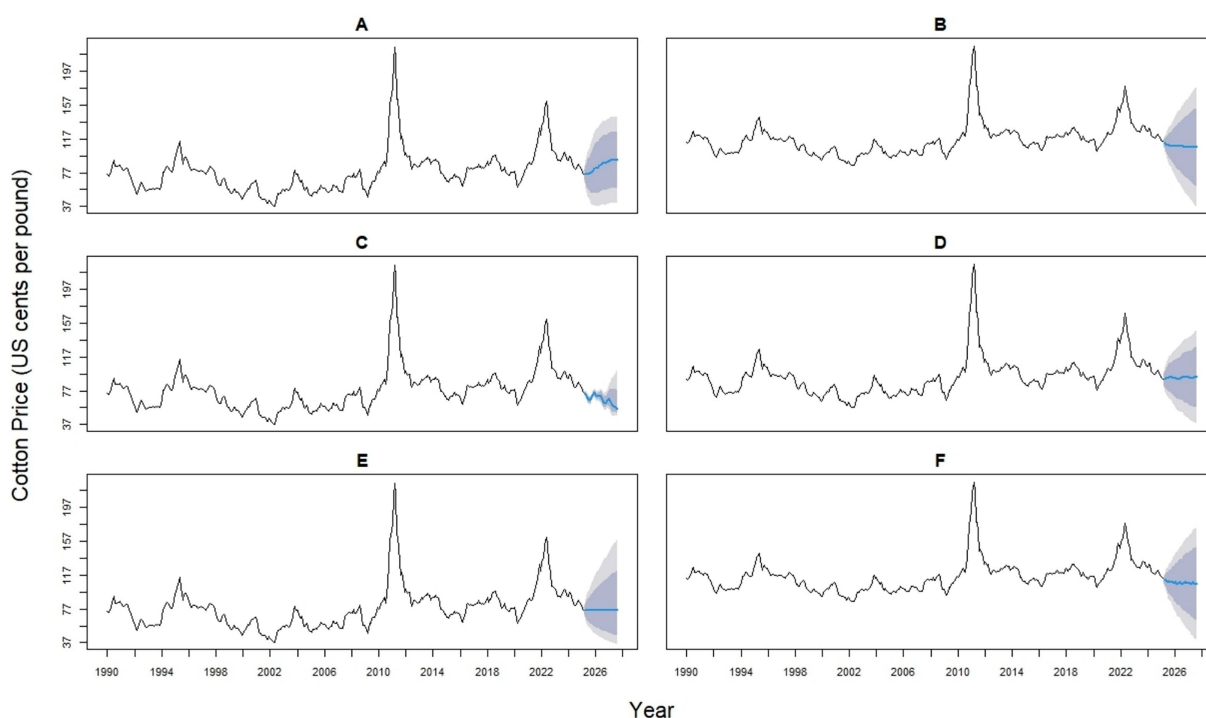Forecasts of global monthly cotton prices using model averaging approaches.

FIGURE 5
Time series forecasts of global monthly cotton prices (in U. S. cents per pound) using six forecasting models **(A–F)**. **(A)** ARIMA (4,1,1) (0,0,2) [12] model; **(B)** ETS (M, Ad, N) model; **(C)** NNAR (26,1,14) [12] model; **(D)** Theta model; **(E)** TBATS (0,0,0, −, −) model; **(F)** STL + ETS (M, Ad, N) model. The black lines represent historical observations, blue lines show model forecasts, and shaded areas indicate prediction intervals.

capturing complex nonlinear dynamics and short-term fluctuations in the global cotton price series, likely due to its sensitivity to historical seasonal spikes, which contributed to increased short-term forecast volatility.

F. The TBATS (0,0,0, −, −) model, with a cross-validated weight of 0.182, exhibited minimal bias (ME = −0.054,

MPE = −0.123). However, its relatively higher errors compared to the NNAR model (RMSE = 5.007, MAE = 3.047, MAPE = 3.658) indicate difficulty in capturing the fluctuations of cotton prices. Its lack of seasonal or complex components likely limits its overall accuracy.

4.2 Hybrid models

The 14 hybrid models represent a spectrum of statistical and machine learning integrations, designed to enhance predictive accuracy by capturing both linear and nonlinear patterns. Hybrids 1–4 combine ARIMA with ETS (0.53/0.47), NNAR (0.514/0.486), TBATS (0.485/0.515), or STL (0.525/0.475). Hybrids 5–7 pair ETS with NNAR (0.491/0.509), TBATS (0.458/0.542), or STL (0.499/0.501), while Hybrids 8–9 integrate NNAR with TBATS (0.431/0.569) or STL (0.499/0.501). Hybrid 10 combines TBATS (0.499) with the STL (0.501) model. Likewise, Hybrid 11 blends ARIMA (0.364), ETS (0.327), and NNAR (0.312). Hybrid 12 combines ARIMA (0.267), ETS (0.240), and STL (0.241). Hybrid 13 incorporates ARIMA (0.357), ETS (0.321), and STL (0.322), whereas Hybrid 14 provides the most complex structure, balancing ARIMA (0.212), ETS (0.191), NNAR (0.180), STL (0.192), and TBATS (0.225) (Figures 6, 7).

Among the hybrid models, those incorporating NNAR, particularly Hybrid 2 (ARIMA × NNAR), showed lower error values (RMSE = 2.820, MAE = 1.875, MAPE = 2.336) compared with other hybrids. However, these errors remained higher than those of the standalone NNAR model (RMSE = 1.164, MAE = 0.832, MAPE = 1.192). Therefore, the NNAR model proved to be the best-fitting model, demonstrating the lowest error metrics and the highest predictive accuracy.

In our study, the standalone NNAR model achieved lower error values than the hybrid models, highlighting its strong ability to capture nonlinear patterns in cotton price dynamics. This aligns with the findings of Taskaya-Temizel and Casey (2005), who reported that combining multiple models for forecasting does not always lead to

better accuracy and can sometimes perform worse than individual models.

However, the superior performance of NNAR in our dataset does not diminish the value of hybrids. Forecast combinations are widely recognized as a means of improving robustness and stability across different time horizons and market conditions, even when a single model performs best in a specific setting (Makridakis et al., 2018). Hybridization remains valuable for enhancing resilience under uncertainty and mitigating the risk of model-specific weaknesses. As noted by Zhang (2003), hybrid models tend to outperform single models when the series exhibits a mix of linear and nonlinear patterns, because each component is specialized in capturing distinct aspects of the data. In our case, however, the cotton price series appears to be dominated by nonlinear patterns, which the NNAR model was able to capture more effectively than the hybrid configuration.

In this study, 22 forecasting models were evaluated to assess their effectiveness in predicting monthly cotton prices. The models included both statistical and machine learning approaches, namely Autoregressive Integrated Moving Average (ARIMA), Exponential Smoothing State Space (ETS), Neural Network Autoregression (NNAR), Theta Model, TBATS, and Seasonal-Trend Decomposition using LOESS (STL). Given the complex and nonlinear dynamics of cotton prices, identifying the most accurate and reliable forecasting approaches is critical for informed decision-making by stakeholders. By systematically comparing these models, this study addresses the need to evaluate both traditional and machine learning-based methods within a unified framework, providing insights into which techniques best capture the temporal patterns of cotton markets.
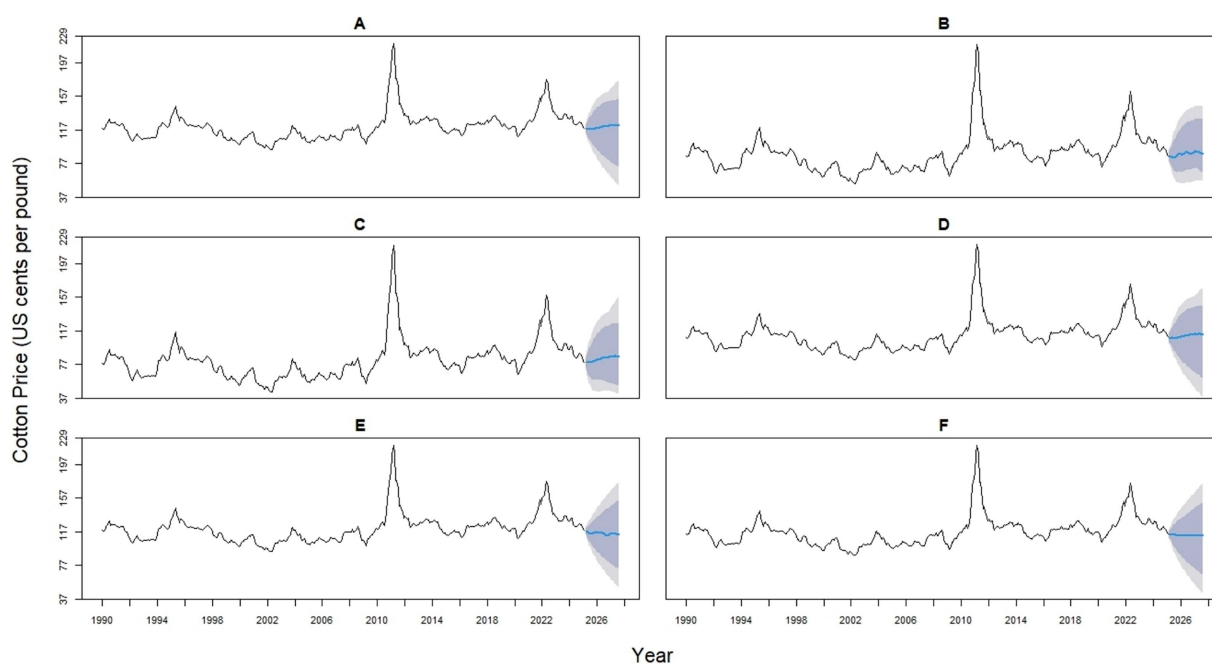

FIGURE 6
Forecast comparison of six hybrid models for monthly global cotton prices. Each panel displays one hybrid model (Hybrid 1−6, labeled **A−F**), formed by combining two base models (ARIMA, ETS, NNAR, Theta, TBATS, or STL) with varying weights. Historical prices are shown in black, and forecasts are in blue. Shaded regions represent 80 and 95% prediction intervals, illustrating the uncertainty in the forecasts.
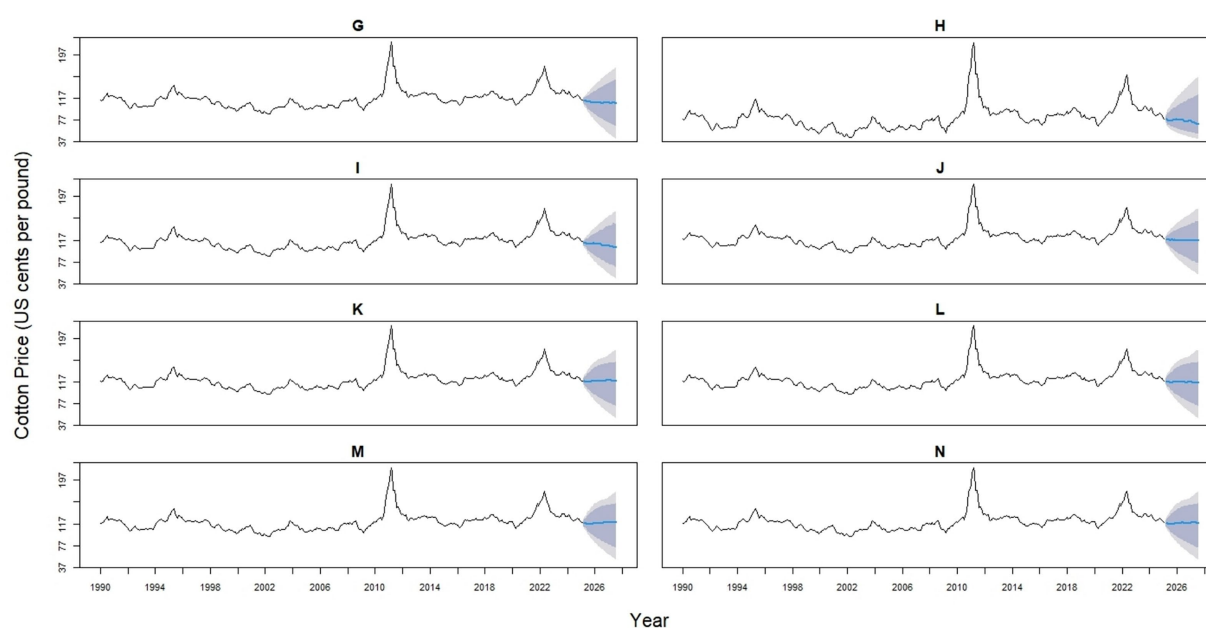
FIGURE 7
Forecast comparison of eight hybrid models for global monthly cotton prices. Each panel displays one hybrid model (Hybrid 7–14, labeled **G–N**), formed by combining two base models (ARIMA, ETS, NNAR, Theta, TBATS, or STL) with varying weights. Historical prices are shown in black, and forecasts are in blue. Shaded regions represent 80 and 95% prediction intervals, illustrating the uncertainty in the forecasts.
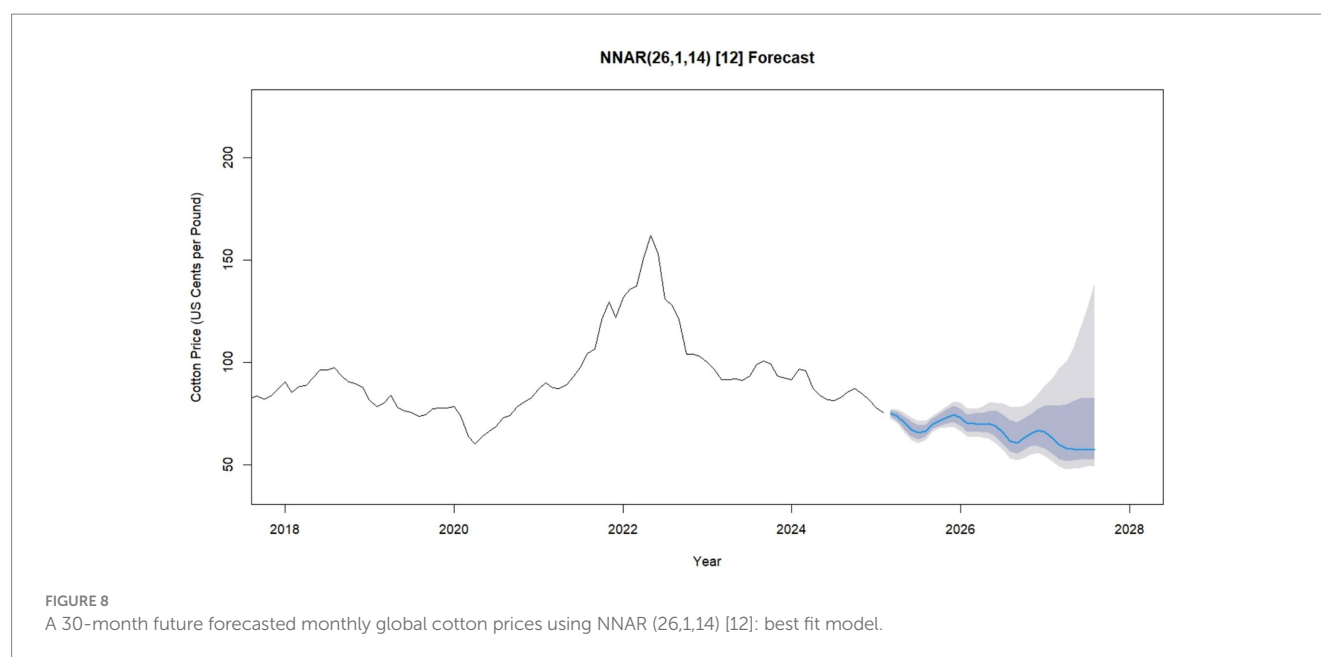
4.3 Model interpretability: best fit model

A Neural Network Autoregressive (NNAR) model with the configuration NNAR (26, 1, 14)12 was employed. This model utilized 26 lagged observations as inputs and 14 neurons in a single hidden layer to model nonlinear dependencies within the data. The seasonal period of 12 was again used to capture annual patterns. The NNAR model demonstrated superior predictive performance compared to all other models considered. It achieved the lowest values for Root Mean Square Error (RMSE) and Mean Absolute Error (MAE), indicating a high level of accuracy and minimal prediction error. Additionally, the model recorded the lowest Mean Absolute Percentage Error (MAPE) of 1.20325, underscoring its effectiveness in reducing relative forecast errors.

Figure 8 presents a 30-month forecast of cotton prices generated by the NNAR (26,1,14) [12] model. The x-axis shows time in months, while the y-axis represents predicted cotton prices in U. S. cents per pound. The model uses the previous 26 months of data as input lags, applies one seasonal difference, and includes 14 hidden neurons to capture nonlinear patterns. The forecast indicates moderate fluctuations, with prices expected to vary between 66 and 74 cents per pound. These cyclical movements reflect potential periods of price increases (peaks) and declines (troughs), capturing both seasonal and market-driven dynamics. As highlighted by Gneiting et al. (2007), effective probabilistic forecasting requires a balance between sharpness (narrowness of intervals) and calibration (ensuring that the actual values fall within the predicted range). In this context, the NNAR model provides not only point forecasts but also a clear measure of uncertainty, enabling stakeholders in the cotton sector to make informed, resilient decisions in the face of market variability.

In our study, the NNAR model's superior performance in forecasting cotton prices, as supported by its lowest RMSE, MAE, and MAPE, highlights the effectiveness of machine learning techniques in capturing the complex and nonlinear dynamics of agricultural commodity markets. This is attributed to the model's ability to uncover hidden patterns and nonlinear dependencies in the data.

Based on this, Khashei and Bijari (2010) further emphasized the utility of artificial neural networks (ANNs) in time series forecasting, highlighting their adaptability and ability to model complex relationships. This flexibility makes ANNs particularly effective for real-world forecasting problems, which often involve uncertainty and nonlinearity. Among these models, Neural Network Autoregressive (NNAR) models have gained recognition for their reliability, with Zhang (2003) emphasizing their strength in capturing nonlinear trends and seasonal patterns often observed in economic and financial time series. Several studies have demonstrated the applicability and effectiveness of the NNAR model in forecasting a wide range of commodities, validating their performance through various error metrics that assess forecasting accuracy. In the study done by Reza and Debnath (2025), NNAR (5,1,10) [12] and NNAR (3,1,10) [12] models were applied to forecast the prices of wheat and rice, two key agricultural commodities. The models effectively captured both short-term patterns and complex relationships, with performance measured using standard accuracy metrics like RMSE, MAE, MASE, and MAPE. This study further supports the practical value of NNAR models in agricultural forecasting contexts. Similarly, Raghav et al. (2022) used the NNAR (1,1) model to forecast pulse production. Although simpler in design, the model was assessed using a broad set of evaluation metrics: Mean Error (ME), RMSE, MAE, Mean Percentage Error (MPE),

FIGURE 8
A 30-month future forecasted monthly global cotton prices using NNAR (26,1,14) [12]: best fit model.

MAPE, and MASE. This comprehensive evaluation helped ensure the model's reliability in capturing the variability and irregularities typical of agricultural production data. Similar results were reported by Almarashi et al. (2024a), where the NNAR model outperformed other competing models, including ARIMA, ETS, and TBATS, based on lower values of RMSE, MAE, and MAPE. These findings suggest that the NNAR model is a more effective approach for time series forecasting compared to conventional methods. Likewise, Melina et al. (2024) applied the NNAR (1,10) model to forecast gold prices. With a single lag input and a higher number of hidden neurons, this model was well-equipped to capture complex nonlinear dependencies in volatile financial data. Performance was assessed using RMSE and MAE, metrics especially relevant in financial forecasting, where even small errors can have substantial consequences.

The robustness of the Neural Network Auto-Regressive (NNAR) model extends beyond agricultural and financial applications. For instance, Qureshi et al. (2025) showed that NNAR effectively predicted cardiovascular disease mortality, while Alshanbari et al. (2023) found it successful in forecasting COVID-19 cases and recoveries in Pakistan. These studies underscore NNAR's ability to capture complex nonlinear patterns in diverse time series datasets, suggesting its potential as a robust tool for forecasting agricultural prices, such as cotton, which frequently exhibit similar nonlinear dynamics. Furthermore, hybrid approaches that combine NNAR with nonlinear models and residual volatility modeling, as proposed by Iftikhar et al. (2025) for stock indices, could help improve forecast accuracy and reliability in commodity markets. Despite these strengths, NNAR and hybrid methods remain less transparent than traditional statistical models, emphasizing the need for further methodological refinement to improve interpretability and practical applicability in agricultural economics. Overall, these studies collectively validate the NNAR model's adaptability across different forecasting contexts. This shows that NNAR is a reliable and flexible tool for finding both simple and complex patterns in different forecasting situations.

4.4 Practical implications

Historical events illustrate the sensitivity of cotton markets to global dynamics. In Figure 1, the sharpest rise in global cotton prices happened in early 2011, reaching over $2.00 per pound, the highest in more than 100 years. This spike came after cotton demand began to recover following the 2008–2009 economic crisis. The surge was caused by a mix of factors, including a shortage of cotton and heavy trading (Agarwal and Narala, 2020). The supply shortage was worsened by poor harvests in key producing countries, extreme weather events, export restrictions, and rising global demand (Janzen et al., 2018). Weather events like floods, droughts, and heavy rains caused major disruptions in cotton production, with yields varying by 10–80% (Hoerling et al., 2013). In the U. S., the worst drought in over a century between mid-2010 and September 2011 greatly impacted cotton production (Anyamba et al., 2014). After prices rose above 200 cents per pound post-2011, the demand for cotton dropped, and people began using more Man-Made Fibers (MMF). However, when prices stabilized between 90 and 100 cents per pound from 2012 to 2014, cotton became more competitive again against MMF and other non-cotton textiles (Agarwal and Narala, 2020). This period highlights the importance of monitoring the cotton market and taking timely actions to reduce future price swings.

Most recently, another significant fluctuation in global cotton prices observed between 2020 and 2022 can largely be attributed to two major global disruptions: the COVID-19 pandemic and the Russia-Ukraine conflict. The pandemic severely disrupted global supply chains, particularly in sectors like textiles, leading to both reduced demand and increased production costs. Early in the pandemic, lockdowns and economic uncertainty caused a sharp drop in textile demand, resulting in lower cotton prices (Emran and

Schmitz, 2023). At the same time, supply chain disruptions made it harder to get important materials and caused delivery delays, which increased the cost of producing textiles (Cai and Luo, 2020). The COVID-19 pandemic resulted in shortages and disruptions in the production and transportation of raw materials (Paul et al., 2021). Despite this initial downturn, the cotton market began to recover by the third quarter of 2020 as global demand gradually resumed. Prices continued to rise through 2021, eventually peaking in June 2022. These prices movements reflect how closely cotton markets are tied to broader global dynamics and supply chain conditions (Erkencioglu et al., 2023). In addition, the Russia-Ukraine conflict introduced further instability by disrupting global trade flows and raising input and import costs. These pressures indirectly increased the cost of cotton production and added uncertainty to global cotton markets (Arndt et al., 2023). Together, these events underscore the vulnerability of agricultural commodity markets to global crises and the importance of adaptable forecasting models in such volatile contexts.

5 Conclusion

Agricultural commodity prices are shaped by complex global and local supply–demand dynamics. Timely, accurate forecasts help stakeholders make informed decisions and manage price-related risks. This study examined the predictive performance of various univariate time series forecasting models for global monthly cotton prices using data from January 1990 to January 2025. The NNAR (26,1,14) [12] model demonstrated superior accuracy in forecasting global monthly cotton prices, in comparison to several individual models (ARIMA, ETS, STL, TBATS, Theta, and NNAR), as well as hybrid models combining these methods, validated by the key evaluation metrics like RMSE, MAE, and MAPE. The NNAR model's ability to capture both linear and nonlinear trends, as well as seasonal variations, highlights its suitability for modeling complex time series data. Given that each forecasting method has distinct strengths and limitations, model selection should be guided by the nature of the data and the specific forecasting goals. In some cases, combining models may help further reduce prediction errors. Our analysis also highlights the sensitivity of global cotton prices to macroeconomic shocks, supply chain disruptions, and external events. For instance, the sharp price spike in 2011 and the volatility seen between 2020 and 2022 reflect the broader impact of such factors. These findings emphasize the need for reliable forecasting tools to support informed decision-making, risk management, and strategic planning in the agricultural sector. This study was limited to univariate time series models, without considering multivariate approaches. Future research could incorporate relevant covariates, such as weather, trade policies, and economic indicators, to improve forecasting accuracy. The use of advanced machine learning and hybrid methods, including deep learning, Random Forest, and Support Vector Regression, holds strong potential for improving predictive performance. Developing such models and aligning them with real-world applications can help the agricultural sector better anticipate market changes, enhancing price stability and resilience. Furthermore, adopting a rolling-origin (out-of-sample) evaluation approach provides additional validation of forecasting accuracy.

Data availability statement

Publicly available datasets were analyzed in this study. This data can be found at: https://fred.stlouisfed.org/series/PCOTTINDUSDM.

Author contributions

Funding

Acknowledgments

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The authors declare that no Gen AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

Agarwal, I., and Narala, A. (2020). Comparing predictive accuracy through price forecasting models in cotton. Journal of Cotton Research and Development. 34:146–157. Available at: https://www.cabidigitallibrary.org/doi/full/10.5555/20203190149

Ahmad, S., Iftikhar, H., Qureshi, M., Khan, I., Omer, A. S. A., Torres Armas, E. A., et al. (2025a). A new auxiliary variables-based estimator for population distribution function under stratified random sampling and non-response. *Sci. Rep.* 15, 1–25. doi: 10.1038/s41598-025-98246-y

Ahmad, S., Qureshi, M., Iftikhar, H., Rodrigues, P. C., and Rehman, M. Z. (2025b). An improved family of unbiased ratio estimators for a population distribution function. *AIMS Math.* 10, 1061–1084. doi: 10.3934/math.2025051

Almarashi, A. M., Daniyal, M., and Jamal, F. (2024a). A novel comparative study of NNAR approach with linear stochastic time series models in predicting tennis player's performance. *BMC Sports Sci. Med. Rehabil.* 16:28. doi: 10.1186/s13102-024-00815-7

Almarashi, A. M., Daniyal, M., and Jamal, F. (2024b). Modelling the GDP of KSA using linear and non-linear NNAR and hybrid stochastic time series models. *PLoS One* 19:e0297180. doi: 10.1371/journal.pone.0297180

Alshanbari, H. M., Iftikhar, H., Khan, F., Rind, M., Ahmad, Z., and El-Bagoury, A. A. A. H. (2023). On the implementation of the artificial neural network approach for forecasting different healthcare events. Diagnostics 13:1310. doi: 10.3390/diagnostics13071310

Anyamba, A., Small, J. L., Britch, S. C., Tucker, C. J., Pak, E. W., Reynolds, C. A., et al. (2014). Recent weather extremes and impacts on agricultural production and vector-borne disease outbreak patterns. *PLoS One* 9:e0092538. doi: 10.1371/journal.pone.0092538

Arndt, C., Diao, X., Dorosh, P., Pauw, K., and Thurlow, J. (2023). The Ukraine war and rising commodity prices: implications for developing countries. Global Food Security 36:100680. doi: 10.1016/j.gfs.2023.100680

Assimakopoulos, V., and Nikolopoulos, K. (2000). The theta model: a decomposition approach to forecasting. *Int. J. Forecast.* 16, 521–530. doi: 10.1016/S0169-2070(00)00066-2

Babu, C. N., and Reddy, B. E. (2014). A moving-average filter-based hybrid ARIMA-ANN model for forecasting time series data. *Appl. Soft Comput.* 23, 27–38. doi: 10.1016/j.asoc.2014.05.028

Box, G. E. P., and Jenkins, G. M. (1970). Time series analysis: Forecasting and control. San Francisco, CA: Holden-Day.

Box, G. E., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). Time series analysis: Forecasting and control. Hoboken, NJ: John Wiley and Sons.

Brown, R. G. (1959). Statistical forecasting for inventory control. New York, NY: McGraw-Hill.

Cai, M., and Luo, J. (2020). Influence of COVID-19 on manufacturing industry and corresponding countermeasures from supply chain perspective. *J. Shanghai Jiaotong Univ.* 25, 409–416. doi: 10.1007/s12204-020-2206-z

Chang, X., Gao, M., Wang, Y., and Hou, X. (2012). Seasonal autoregressive integrated moving average model for precipitation time series. *J. Mathematical Stat.* 8, 500–505. doi: 10.3844/jmssp.2012.500.505

Chi, Y. N. (2021). Time series forecasting of the global price of soybeans using a hybrid SARIMA and NARNN model. *Data Sci. J. Comput. Appl. Inf.* 5, 85–101. doi: 10.32734/jocai.v5.i2-5674

Cleveland, R. B., Cleveland, W. S., McRae, J. E., and Terpenning, I. J. (1990). STL: a seasonal-trend decomposition procedure based on loess. *J. Off. Stat.* 6, 3–33.

Darekar, A., and Reddy, A. A. (2017). Cotton price forecasting in major producing states. *Econ. Aff.* 62, 373–380. doi: 10.5958/0976-4666.2017.00047

De Livera, A. M. (2010). Automatic forecasting with a modified exponential smoothing state space framework. *Monash Econ. Bus. Stat. Working Papers* 10, 1–6. Available at: https://ideas.repec.org/p/msh/ebswps/2010-10.html

Dickey, D. A., and Fuller, W. A. (1981). Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica* 49, 1057–1072. doi: 10.2307/1912517

Emran, S. J., and Schmitz, A. (2023). Impact of COVID-19 (2020-2022) on cotton and garments market of Bangladesh: a small country case. *J. Agric. Food Ind. Organ.* 21, 89–98. doi: 10.1515/jafio-2022-0027

Erkencioglu, B. N., Zuhal, M., Tokel, D., and Ozyigit, I. I. (2023). Worldwide cotton production and trade during COVID-19 pandemic: an empirical analysis for a three-year observation. *Not. Bot. Horti Agrobot. Cluj-Napoca* 51, 1–10. doi: 10.15835/NBHA51413341

Gardner, E. S. (1985). Exponential smoothing: the state of the art. *J. Forecast.* 4, 1–28. doi: 10.1002/for.3980040103

Global Price of Cotton. (2025). Federal Reserve Bank of St. Louis. Available online at: https://fred.stlouisfed.org/series/PCOTTINDUSDM

Gneiting, T., Balabdaoui, F., and Raftery, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *J. Royal Stat. Soc.* 69, 243–268. doi: 10.1111/j.1467-9868.2007.00587.x

Hoerling, M., Kumar, A., Dole, R., Nielsen-Gammon, J. W., Eischeid, J., Perlwitz, J., et al. (2013). Anatomy of an extreme event. *J. Clim.* 26, 2811–2832. doi: 10.1175/JCLI-D-12-00270.1

Holt, C. E. (1957). Forecasting seasonals and trends by exponentially weighted averages (O.N.R. Memorandum no. 52). Pittsburgh, CA: Carnegie Institute of Technology.

Hudson, D. (1998). Contemporary issues: cotton market price information-how it affects the industry. Available online at: https://www.researchgate.net/publication/237123186

Hyndman, R. J., and Athanasopoulos, G. (2018). Forecasting: Principles and practice. Melbourne: OTexts.

Hyndman, R. J., and Athanasopoulos, G. (2021). Forecasting: Principles and practice. Melbourne: OTexts.

Hyndman, R. J., and Khandakar, Y. (2008). Automatic time series forecasting: the forecast package for R. *J. Stat. Softw.* 27, 1–22. doi: 10.18637/jss.v027.i03

Hyndman, R. J., and Koehler, A. B. (2006). Another look at measures of forecast accuracy. *Int. J. Forecast.* 22, 679–688. doi: 10.1016/j.ijforecast.2006.03.001

Hyndman, R. J., Koehler, A. B., Ord, J. K., and Snyder, R. D. (2008). Forecasting with exponential smoothing: The state space approach. Berlin: Springer.

Iftikhar, H., Khan, F., Torres Armas, E. A., Rodrigues, P. C., and López-Gonzales, J. L. (2025). A novel hybrid framework for forecasting stock indices based on nonlinear time series models. *Comput. Stat.* 40, 4163–4186. doi: 10.1007/s00180-025-01614-5

Jabran, K., Ul-Allah, S., Chauhan, B. S., and Bakhsh, A. (2019). "An introduction to global production trends and uses, history and evolution, and genetic and biotechnological improvements in cotton" in Cotton production. eds. K. Jabran and B. S. Chauhan (Amsterdam: Wiley), 1–22.

Jaiswal, R., Choudhary, K., and Kumar, R. R. (2022). Stl-elm: a decomposition-based hybrid model for price forecasting of agricultural commodities. *Natl. Acad. Sci. Lett.* 45, 477–480. doi: 10.1007/s40009-022-01169-9

Janzen, J. P., Smith, A., and Carter, C. A. (2018). Commodity price comovement and financial speculation: the case of cotton. *Am. J. Agric. Econ.* 100, 264–285. doi: 10.1093/ajae/aax052

Khadi, B. M., Santhy, V., and Yadav, M. S. (2010). "Cotton: an introduction" in Biotechnology in agriculture and forestry. ed. G. C. Van Kooten (Cham: Springer), 1–14.

Khadka, R., and Chi, Y. N. (2024). Forecasting the global price of corn: unveiling insights with SARIMA modelling amidst geopolitical events and market dynamics. *Am. J. Appl. Stat. Econ.* 3, 124–135. doi: 10.54536/ajase.v3i

Khan, M. A., Wahid, A., Ahmad, M., Tahir, M. T., Ahmed, M., Ahmad, S., et al. (2020). World cotton production and consumption: an overview. In: Cotton Production and Uses: Agronomy, Crop Protection, and Postharvest Technologies. Singapore: Springer. p. 1–7. doi: 10.1007/978-981-15-1472-2_1

Khandelwal, I., Adhikari, R., and Verma, G. (2015). Time series forecasting using hybrid ARIMA and ANN models based on DWT decomposition. *Procedia Comput. Sci.* 48, 173–179. doi: 10.1016/j.procs.2015.04.167

Khashei, M., and Bijari, M. (2010). An artificial neural network (p, d, q) model for time series forecasting. *Expert Syst. Appl.* 37, 479–489. doi: 10.1016/j.eswa.2009.05.044

Kim, S., and Kim, H. (2016). A new metric of absolute percentage error for intermittent demand forecasts. *Int. J. Forecast.* 32, 669–679. doi: 10.1016/j.ijforecast.2015.12.003

Kumar, A., Kumar, V., Chetna, Ghalawat, S., Kaur, J., Kumari, K., et al. (2024). Forecasting cotton (*Gossypium* spp.) prices in major Haryana markets: a time series and ARIMA approach. *Indian J. Agric. Sci.* 94, 1013–1018. doi: 10.56093/ijas.v94i9.150524

Kumar, M., and Thenmozhi, M. (2014). Trading and forecasting performance of different hybrid ARIMA-neural network models for stock returns. *Int. J. Model. Oper. Manag.* 4, 137–144. doi: 10.1504/IJMOM.2014.067361

Liu, Z. (2025). Prediction of US cotton futures price under different models. *SHS Web Conf.* 213:01017. doi: 10.1051/shsconf/202521301017

Makridakis, S., Spiliotis, E., and Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: concerns and ways forward. *PLOS ONE* 13:e0194889. doi: 10.1371/journal.pone.0194889

Melina, Sukono, Napitupulu, H., Mohamed, N., Chrisnanto, Y. H., Hadiana, A. I. D., et al. (2024). Comparative analysis of time series forecasting models using ARIMA and neural network autoregression methods. *Barekeng* 18, 2563–2576. doi: 10.30598/barekengvol18iss4pp2563-2576

Panigrahi, S., and Behera, H. S. (2017). A hybrid ETS-ANN model for time series forecasting. *Eng. Appl. Artif. Intell.* 66, 49–59. doi: 10.1016/j.engappai.2017.07.007

Patel, H., Shah, M., and Thakkar, P. (2020). Paddy price prediction using LSTM networks. *Int. J. Adv. Comput. Sci. Appl.* 11, 560–567. doi: 10.14569/IJACSA.2020.0110567

Paul, S. K., Chowdhury, P., Moktadir, M. A., and Lau, K. H. (2021). Supply chain recovery challenges in the wake of COVID-19 pandemic. *J. Bus. Res.* 136, 316–329. doi: 10.1016/j.jbusres.2021.07.056

Pegels, C. C. (1969). Exponential forecasting: some new variations. Management Science 311–315. Available at: https://www.jstor.org/stable/2628137

Perone, G. (2022). Using the SARIMA model to forecast the fourth global wave of cumulative deaths from COVID-19: evidence from 12 hard-hit big countries. *Econometrics* 10:18. doi: 10.3390/econometrics10020018

Qureshi, M., Ishaq, K., Daniyal, M., Iftikhar, H., Rehman, M. Z., and Salar, S. A. A. (2025). Forecasting cardiovascular disease mortality using artificial neural networks in Sindh, Pakistan. *BMC Public Health* 25, 1–14. doi: 10.1186/S12889-024-21187-0

Raghav, Y. S., Mishra, P., Alakkari, K. M., Singh, M., Al Khatib, A. M. G., and Balloo, R. (2022). Modelling and forecasting of pulses production in south Asian countries and its role in nutritional security. *Legum. Res.* 45, 454–461. doi: 10.18805/LRF-645

Reza, A., and Debnath, T. (2025). An approach to make comparison of ARIMA and NNAR models for forecasting price of commodities. Available online at: https://www.researchgate.net/publication/342563043

Sun, F., Meng, X., Zhang, Y., Wang, Y., Jiang, H., and Liu, P. (2023). Agricultural product price forecasting methods: a review. *Agriculture* 13, 1–18. doi: 10.3390/agriculture13091671

Taskaya-Temizel, T., and Casey, M. C. (2005). A comparative study of autoregressive neural network hybrids. *Neural Netw.* 18, 781–789. doi: 10.1016/J.NEUNET.2005.06.003

Wang, Y. (2023). Agricultural products price prediction based on improved RBF neural network model. *Appl. Artif. Intell.* 37, 1–15. doi: 10.1080/08839514.2023.2204600

Wang, L., Zou, H., Su, J., Li, L., and Chaudhry, S. (2013). An ARIMA-ANN hybrid model for time series forecasting. *Syst. Res. Behav. Sci.* 30, 244–259. doi: 10.1002/sres.2179

Weng, Y., Wang, X., Hua, J., Wang, H., Kang, M., and Wang, F. Y. (2019). Forecasting horticultural products price using ARIMA model and neural network based on a large-scale data set collected by web crawler. *IEEE Trans. Comput. Soc. Syst.* 6, 547–553. doi: 10.1109/TCSS.2019.2914499

Winters, P. R. (1960). Forecasting sales by exponentially weighted moving averages. *Manag. Sci.* 6, 324–342. doi: 10.1287/mnsc.6.3.324

Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing* 50, 159–175. doi: 10.1016/S0925-2312(01)00759-5

Zhang, D., Chen, S., Liwen, L., and Xia, Q. (2020). Forecasting agricultural commodity prices using model selection framework with time series features and forecast horizons. *IEEE Access* 8, 28197–28209. doi: 10.1109/ACCESS.2020.2971591

Zhang, G., Patuwo, B. E., and Hu, M. Y. (1998). Forecasting with artificial neural networks: the state of the art. *Int. J. Forecast.* 14, 35–62. doi: 10.1016/S0169-2070(97)00044-7

Zhang, G. P., and Qi, M. (2005). Neural network forecasting for seasonal and trend time series. *Eur. J. Oper. Res.* 160, 501–514. doi: 10.1016/j.ejor.2003.08.037