



Time Series Analysis of the VIX (Volatility Index)

Prepared by:
Nada Mohamed Abbas
Abdulrahman Mostafa Kamel
Mohamed Amir Mohamed
Ezz Eldin Ahmed Mohamed

Statistics Department
Arabic Section
Spring 2025

Under the supervision of:
Dr. Rasha Elsouda
T.A. Yousra Hassan

CONTENT

INTRODUCTION:	- 1 -
LITERATURE REVIEW:	- 2 -
DATA ANALYSIS:	- 4 -
DATA DESCRIPTION:.....	- 4 -
BOX & JENKINS ANALYSIS:	- 5 -
CLASSICAL ANALYSIS – DECOMPOSITION METHOD:	- 14 -
TREND ANALYSIS:	- 17 -
FORECASTING WITH THE CUBIC MODEL:	- 21 -
CONCLUSION:	- 25 -
MODEL LIMITATIONS	- 25 -
FUTURE DIRECTIONS.....	- 26 -
REFERENCES:	- 27 -

Introduction:

Volatility is one of the important phenomena in finance as it measures the frequency and magnitude of price movements (both up and down) of a financial instrument (stocks, bonds) over a specific period. The more extreme the price fluctuations, the higher the volatility level. Financial market volatility is divided into two main types: historical volatility, which is based on actual past price movements and calculated using the standard deviation of previous returns, and implied volatility, which reflects the market's expectations of future volatility derived from options prices. Both measures are crucial for assessing risk, with historical volatility showing what has happened and implied volatility indicating what the market expects to happen.

The VIX Index is an abbreviation of Volatility Index calculated by the Chicago Board Options Exchange. The VIX Index is one of the popular measures for implied volatility; it predicts S&P 500 - a stock market index tracking the stock performance of 500 leading companies listed on stock exchanges in the United States - volatility and market uncertainty over the next 30 days and is calculated by using the midpoint of real-time S&P 500 Index option bid/ask quotes. More specifically, the VIX Index is intended to provide an instantaneous measure of how much the market thinks the S&P 500 Index will fluctuate in the 30 days from the time of each tick of the VIX Index.

Analyzing the VIX is important for several reasons:

- **Market Sentiment Indicator:** The VIX provides real-time insight into investor sentiment and market anxiety. A high VIX indicates heightened uncertainty and fear, often associated with market downturns, while a low VIX suggests investor confidence and stability.
- **Risk Management and Hedging:** Traders, institutional investors, and portfolio managers use the VIX to gauge risk and to implement strategies such as volatility hedging using options or VIX futures. Understanding its behavior is essential for minimizing potential losses during turbulent market conditions.
- **Forecasting Market Movements:** As a forward-looking measure, the VIX helps in predicting potential market corrections or rallies. It can serve as an early warning system for shifts in market trends, making it valuable for investment decisions.
- **Policy and Economic Implications:** Economists and policymakers observe the VIX as part of broader market analysis to assess systemic risk, financial stability, and the impact of macroeconomic news or geopolitical events.

For these reasons, the researcher has selected the VIX as the subject of analysis in this study.

The **primary objective** is to explore the time series characteristics of the VIX Index, understand its underlying patterns, and assess the presence of volatility clustering, mean reversion, potential seasonality and forecasting annual values. To achieve this, the study will apply both the Box & Jenkins methodology - a structured approach to identifying, estimating, and diagnosing ARIMA models - and the Classical Decomposition approach, which separates the series into trend, seasonal, and residual components to better understand the structure of the data and compare between them.

Through these techniques, the researcher aims to build a robust forecasting model that can predict future movements in the VIX. This will not only help in understanding how market volatility behaves over time but will also offer valuable insights for risk management and decision-making in financial markets. Accurately modeling the VIX enables better anticipation of future periods of market stress or stability, making this analysis both academically and practically relevant.

Before proceeding to the data analysis section, it is essential to review the relevant literature on the VIX Index and time series modeling techniques. A comprehensive literature review provides a solid foundation by highlighting key findings from previous research, commonly used methodologies, and the theoretical background of volatility modeling. Additionally, it helps justify the selection of methods such as the Box & Jenkins approach and Classical Decomposition, as supported by prior studies in similar contexts.

Literature Review:

The analysis and forecasting of the Volatility Index (VIX), commonly referred to as the "fear gauge," have gained significant academic and practical attention due to the index's strong association with market uncertainty and investor sentiment. Introduced by the Chicago Board Options Exchange (CBOE) in 1993, the VIX is constructed from S&P 500 index option prices and reflects the market's expectation of 30-day forward-looking volatility (Whaley, 2000). The accessibility of historical data, such as that provided by Kaggle (Prakash, 2023) and FRED (Federal Reserve Bank of St. Louis, 2024), has enabled researchers to deploy a variety of time series models to understand and forecast volatility dynamics.

Traditional Econometric Approaches:

Much of the early and foundational work on volatility forecasting employed generalized autoregressive conditional heteroskedasticity (GARCH) models. Bollerslev (1986) extended Engle's (1982) ARCH model to GARCH, allowing conditional variance to depend on both past squared returns and past conditional variances. Numerous studies confirm that GARCH-family models provide robust short-term forecasts for volatility indices like the VIX (Hansen & Lunde, 2005; Giot, 2005). Particularly, Giot (2005) compared several GARCH specifications and found that asymmetrical models such as EGARCH and GJR-GARCH capture the leverage effect in volatility more effectively.

Despite their popularity, traditional GARCH models often assume a linear structure, which may not fully capture the nonlinearities inherent in financial time series. This limitation has led to the development of regime-switching models (Hamilton & Susmel, 1994), which allow for structural breaks in volatility regimes. These models have shown improved performance in periods of market stress.

Machine Learning and Hybrid Models:

Recent advancements in machine learning have sparked a wave of research integrating artificial intelligence techniques with econometrics. For instance, Cao and Tay (2003) demonstrated the potential of support vector machines (SVMs) for forecasting financial volatility, outperforming traditional GARCH models in certain settings. Similarly, Zhang et al. (2019) employed long short-term memory (LSTM) networks to predict the VIX and reported superior results over autoregressive models. These models are particularly adept at capturing nonlinear dependencies and long-term memory in volatility sequences. Hybrid approaches that combine econometric and machine learning models have also gained traction. For example, Wang et al. (2020) used a hybrid ARIMA-LSTM model to forecast VIX and found it to be more accurate than single-model frameworks. Such models benefit from ARIMA's strength in modeling linear components and LSTM's ability to learn complex nonlinear patterns.

External Factors and Macro Linkages:

Numerous studies have also examined the macroeconomic and behavioral determinants of the VIX. Apergis (2021) explored the relationship between COVID-19 uncertainty and VIX levels, showing a strong causal link. Similarly, Bouri et al. (2019) investigated the effect of geopolitical risks on VIX fluctuations. These findings underscore the importance of incorporating exogenous variables into forecasting frameworks, especially when working with high-impact global events or systemic risks.

Evaluation Metrics and Model Selection:

Forecasting accuracy is often evaluated using root mean squared error (RMSE), mean absolute percentage error (MAPE), and Diebold-Mariano tests (Diebold & Mariano, 1995). According to Makridakis et al. (2018), no single model universally outperforms others across all forecasting horizons, which advocates for a model comparison strategy tailored to the specific characteristics of the data.

Conclusion:

The literature presents a rich array of methodologies for VIX analysis, ranging from classical GARCH-type models to modern machine learning approaches. The integration of traditional time series modeling with artificial intelligence offers promising directions for future research. This project aims to build on this existing body of work by leveraging a comprehensive dataset of 34 years of daily stock data (Prakash, 2023) and official VIX observations (Federal Reserve Bank of St. Louis, 2024), applying both econometric and neural network-based models to evaluate forecasting accuracy under various market conditions.

Data Analysis:

Data Description:

A	B
observation_date	VIXCLS
2/1/1990	17.24
3/1/1990	18.19
4/1/1990	19.22
5/1/1990	20.11
8/1/1990	20.26
9/1/1990	22.2
10/1/1990	22.44
11/1/1990	20.05
12/1/1990	24.64

This secondary discrete time-series dataset tracks daily VIX (Volatility Index) values from January 1990 through April 30, 2025, covering 252 trading days per year. Sourced from the Chicago Board Options Exchange (CBOE), released by CBOE Market Statistics and distributed via the Federal Reserve Bank of St. Louis (FRED), the data measures 30-day expected market volatility as derived from S&P 500 index options. The VIX values are expressed as percentage points and the generalized formula used in the VIX Index calculation is:

$$\sigma^2 = \frac{2}{T} \sum_i \frac{\Delta K_i}{K_i^2} e^{RT} Q(K_i) - \frac{1}{T} \left[\frac{F}{K_0} - 1 \right]^2$$


Where:

Symbol	Meaning
$VIX / 100 = \sigma \Rightarrow VIX = \sigma \times 100$	VIX value equals the square root of variance multiplied by 100
T	Time to expiration of the option (in years)
F	Forward index level derived from index option prices
K₀	First strike price below the forward index level F
K_i	Strike price of the <i>i</i> -th out-of-the-money option: a call if $K_i > K_0$, a put if $K_i < K_0$, and both call and put if $K_i = K_0$
ΔK_i	Interval between strike prices, calculated as: $\Delta K_i = \frac{K_{i+1} + K_{i-1}}{2}$
R	Risk-free interest rate to expiration
Q(K_i)	Midpoint of the bid-ask spread for the option with strike K _i


The equation aggregates information from call and put options on the S&P 500 Index across different strike prices. Midpoints between bid and ask prices are used to estimate the "true market price". **The result provides the expected market volatility over the next 30 days, which is then multiplied by 100 to express the VIX as a percentage.** So, VIX expresses market uncertainty where:

- VIX above 30: Indicates a volatile, uncertain market (e.g., during economic crises).
- VIX below 20: Reflects a calm market with investor confidence (e.g., during economic growth).
- VIX between 20 and 30: reflects a transitional or cautiously optimistic market environment, where investors are neither overly fearful nor entirely complacent.

Box & Jenkins Analysis:



The time series plot displays the daily values of the VIX Index on the y-axis across the period from 1990 to April 2025 on the x-axis. The plot reveals **no** discernible increasing or decreasing **trend** over time, but shows significant fluctuations in variance, particularly with extreme peaks occurring during the 2008 financial crisis and the 2020 coronavirus pandemic when market uncertainty spiked dramatically. These exceptionally high values represent genuine economic outliers that cannot be removed from the dataset as they accurately reflect periods of severe market stress. To address this **non-constant variance** while preserving these meaningful extreme values, we recommend applying a logarithmic transformation to stabilize the variance across the time series without eliminating these important data points that capture critical market reactions to major economic shocks.



The time series plot shows the log-transformed daily VIX Index values from 1990 to April 2025. After transformation, the series exhibits no apparent trend and appears more stationary variance than the original time series plot. The log transformation helped stabilize the variance to some degree and make the variance more constant across different time periods, meeting the stationarity requirements for reliable time series analysis. But the seasonality must be checked. So, Kruskal Wallis test has been taken for the log-transformed time series.

Test used: kruskall wallis

Test statistic: 270.65


P-value: 0.1881696

The null and alternative hypotheses of the Kruskal-Wallis test are:

H_0 : There is no seasonal effect - the medians across all seasonal periods are equal.

H_1 : There is seasonal effect - At least one seasonal period has a different median than others.

The Kruskal-Wallis test yielded a p-value of 0.188, which exceeds conventional significance levels ($\alpha = 0.05$ and $\alpha = 0.01$). The Decision is **Failed to reject the null hypothesis (H_0)**. There is No statistically significant evidence of seasonality was detected in the time series.



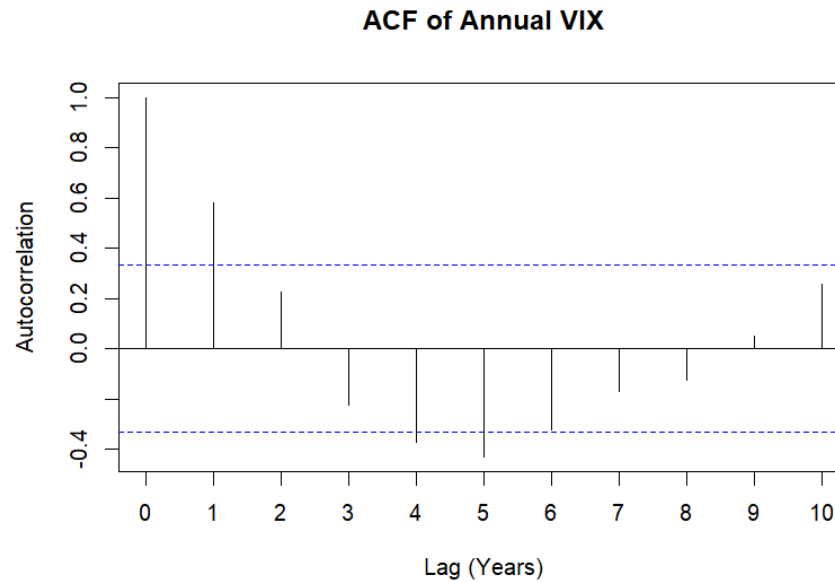
The time series plot displays annual averages of log-transformed VIX values, with each year's mean calculated from 252 trading days (except for 2025, which includes only 84 days through April 30). The time series exhibits no discernible upward or downward trend, and the variance remains stable with constant fluctuation magnitudes throughout the entire period. These properties - the absence of a trend and stable variance - indicate that the series is stationary in both mean and variance. This annual aggregation serves our primary objective of forecasting yearly VIX values.

The VIX time series has been successfully prepared for Box-Jenkins modeling through careful preprocessing. The data underwent logarithmic transformation to stabilize variance and was aggregated annually (using 252 trading days per year, except for 2025 which contains 84 days through April 30). The annual time series plot confirmed the series exhibits stationarity in both mean and variance, with no detectable trend or seasonality (Kruskal-Wallis p-value = 0.188). These properties make the transformed series ideal for ARIMA modeling.


We will now implement the four-phase Box-Jenkins methodology:

1. Identification
2. Estimation
3. Diagnosis
4. Forecasting

Identification:



The ACF plot of the log-transformed annual VIX time series shows significant autocorrelations at lags 1, 4, and 5, with lag 1 demonstrating substantially stronger significance than the higher lags. This pattern indicates stationarity through the rapid decay of autocorrelations after lag 1 and suggests an MA (1) process would be appropriate due to the sharp cutoff after the first lag.



The PACF plot of the log-transformed annual VIX time series reveals statistically significant partial autocorrelations at lags 1 and 3, with lag 1 exhibiting substantially stronger significance. This characteristic pattern - where the PACF shows a sharp cutoff after the first lag with only minor higher-order spikes - strongly suggests an autoregressive process of order 1 AR (1).

The ACF and PACF plots collectively suggest an ARMA (1,1) model as the optimal initial specification for the log-transformed annual VIX series. The ACF's significant spike at lag 1 followed by gradual decay indicates MA (1) dynamics, while the PACF's sharp cutoff after lag 1 (with a weaker lag 3 spike) points to AR (1) behavior.

Estimation of ARMA (1,1):

ARMA (1,1) Estimation Results for Log-Transformed VIX Series:

$$\widehat{VIX}_t = 2.9057 + 0.5096VIX_{t-1} + \varepsilon_t + 0.1198\varepsilon_t$$

Where: \widehat{VIX}_t is the estimated annual log-value of VIX for year t

Call:

```
arima(x = Annual_VIX_ts, order = c(1, 0, 1), method = c("ML"))
```

Coefficients:

	ar1	ma1	intercept
	0.5096	0.1198	2.9057
s.e.	0.1875	0.1755	0.0830

sigma^2 estimated as 0.0493: log likelihood = 2.79, aic = 2.42

Training set error measures:

	ME	RMSE	MAE	MPE	MAPE	MASE	ACF1
Training set	-0.004307557	0.2220373	0.1919045	-0.7363813	6.623298	0.948119	0.03330736

Metric	Value	Interpretation
RMSE	0.243	Typical forecast error $\approx \pm 24.3\%$ in log terms
MAPE	6.58%	Average percentage forecast error and it is acceptable as it is $< 10\%$
MASE	0.942	Slightly better than the naive model forecast as it is less than 1

z test of coefficients:

	Estimate	Std. Error	z value	Pr(> z)
ar1	0.509577	0.187488	2.7179	0.00657 **
ma1	0.119763	0.175504	0.6824	0.49499
intercept	2.905724	0.083024	34.9986	< 2e-16 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

The Z-test hypotheses are:

$$H_0: \beta = 0.$$

$$H_1: \beta \neq 0.$$

The ϕ_1 is statistically significant where the p-value = 0.00657 < $\alpha = 0.05$. But θ_1 is not statistically significant where the p-values = 0.49499 > $\alpha = 0.05$ which suggest the reduction of the model to AR (1) and remove the MA component.

Diagnosis of ARMA (1,1):

Box-Ljung test

data: Residuals

X-squared = 12.608, df = 8, p-value = 0.1261

H_0 : no autocorrelation up to lag 10.

H_1 : significant autocorrelation exists at one or more lags up to 10.


Since p-value = 0.1261 > $\alpha = 0.05$. Therefor the decision is failed to reject H_0 . So, there is no autocorrelation in residuals.

Box-Pierce test
data: Residuals
X-squared = 10.484, df = 8, p-value = 0.2327


H_0 : no autocorrelation up to lag 10.

H_1 : significant autocorrelation exists at one or more lags up to 10.

Since $p\text{-value} = 0.2327 > \alpha = 0.05$. Therefore the decision is failed to reject H_0 . So, there is no autocorrelation in residuals. The results of two tests Validate that ARMA (1,1) captures all significant autocorrelation. Suggests no need for higher-order terms (e.g., AR (2), MA (2)).



The plot displays the ACF of residuals for ARMA (1,1) model. All lags are insignificant within the 95% confidence bounds except for lag zero which indicates that the residuals are white noise.



The PACF plot of the ARMA (1,1) residuals indicates white noise, with all lags within the 95% confidence bounds except for marginally significant spikes at lags 3 and 8. These minor deviations ($p \approx 0.05$) likely represent random variation rather than true autocorrelation.

The ACF and PACF plots introduce that the residuals are white noise $\varepsilon_t \sim N(0, 0.0493)$ where:

$$\rho_k = \begin{cases} 1 & k = 0 \\ 0 & k \neq 0 \end{cases} \quad \phi_{kk} = \begin{cases} 1 & k = 0 \\ 0 & k \neq 0 \end{cases}$$

The σ^2 of residuals represent the deviation of estimated log-transformed VIX from the true values in our data set. So standard deviation = $\sqrt{0.0493} \approx \pm 0.222$. This means that log-transformed values deviate by $\pm 22\%$ from the real one. So, we will try to fit AR (1) as the MA (1) coefficient is insignificant to see if the simplest model is better or not.

Estimation of AR (1):

$$\widehat{VIX}_t = 2.9053 + 0.5817VIX_{t-1} + \varepsilon_t$$

Where: \widehat{VIX}_t is the estimated annual log-value of VIX for year t

Call:

```
arima(x = Annual_VIX_ts, order = c(1, 0, 0), method = "ML")
```

Coefficients:

```
      ar1  intercept
      0.5817      2.9053
s.e.  0.1347      0.0870
```

sigma^2 estimated as 0.04998: log likelihood = 2.56, aic = 0.87

Training set error measures:

```

              ME      RMSE      MAE      MPE      MAPE      MASE      ACF1
Training set -0.004642595 0.2235513 0.1931573 -0.756084 6.668187 0.9543083 0.1108836
```

Metric	AR (1)	ARMA (1,1)	Interpretation
RMSE	0.224	0.222	Similar precision ($\pm 22\%$ log error). AR (1) is marginally better.
MAPE	6.67%	6.62%	Both models achieve excellent accuracy ($< 10\%$ benchmark). Negligible difference.
MASE	0.954	0.948	Both slightly perform better than naive forecasts (MASE < 1), but improvement is minimal.

The error measures suggest there is no significant difference between the two models. The intercept coefficient and the AR coefficient do not change in sign, and only slight changes occur in their magnitudes. However, the AIC of AR (1) (0.87) is much smaller than that of ARMA (1,1) (2.42), strongly suggesting the AR (1) model is better. While the log-likelihood of AR(1) (2.56) is slightly less than that of ARMA(1,1) (2.79), indicating ARMA(1,1) explains the observed data marginally better, the minimal improvement does not justify the additional MA(1) term—especially since it is statistically insignificant ($p = 0.495$).

z test of coefficients:

```

      Estimate Std. Error z value Pr(>|z|)
ar1      0.581743   0.134733  4.3178 1.576e-05 ***
intercept 2.905285   0.086956 33.4110 < 2.2e-16 ***
---
```

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1


The Z-test hypotheses are:

$$H_0: \beta = 0.$$


$$H_1: \beta \neq 0.$$

The ϕ_1 and intercept are statistically significant, with the p-value of ϕ_1 ($1.576e^{-5}$) being less than α (0.05). This confirms that the AR (1) model is preferable to ARMA (1,1), as all its coefficients are significant. The model is admissible because AR (1) processes are always invertible and stationary when $-1 < \phi_1 < 1$. Since $-1 < \phi_1 = 0.5817 < 1$, the model satisfies the stationarity condition.

Diagnosis of AR (1):



The plot displays the ACF of residuals for AR (1) model. All lags are insignificant within the 95% confidence bounds except for lag zero which indicates that the residuals are white noise.



The PACF plot of the AR (1) residuals indicates white noise, with all lags within the 95% confidence bounds except for marginally significant spikes at lags 3 and 8. These minor deviations ($p \approx 0.05$) likely represent random variation rather than true autocorrelation.

Box-Ljung test

data: Residuals

X-squared = 13.206, df = 8, p-value = 0.105

H_0 : no autocorrelation up to lag 10

H_1 : significant autocorrelation exists at one or more lags up to 10.

Since $p\text{-value} = 0.105 > \alpha = 0.05$. Therefore the decision is failed to reject H_0 . So, there is no autocorrelation in residuals.

Box-Pierce test

data: Residuals

X-squared = 10.954, df = 8, p-value = 0.2043

H_0 : no autocorrelation up to lag 10

H_1 : significant autocorrelation exists at one or more lags up to 10.

Since $p\text{-value} = 0.2043 > \alpha = 0.05$. Therefore the decision is failed to reject H_0 . So, there is no autocorrelation in residuals. The results of two tests Validate that ARMA (1,1) captures all significant autocorrelation. Suggests no need for higher-order terms (e.g., AR (2), MA (2)).

```
auto.arima(Annual_VIX_ts, trace = T)
```

```
ARIMA(2,0,2) with non-zero mean : 219.5162
ARIMA(0,0,0) with non-zero mean : 226.0439
ARIMA(1,0,0) with non-zero mean : 216.9942
ARIMA(0,0,1) with non-zero mean : 219.7573
ARIMA(0,0,0) with zero mean      : 312.1814
ARIMA(2,0,0) with non-zero mean : 219.0124
ARIMA(1,0,1) with non-zero mean : 219.3365
ARIMA(2,0,1) with non-zero mean : Inf
ARIMA(1,0,0) with zero mean      : 225.9619
```

```
Best model: ARIMA(1,0,0) with non-zero mean
```

```
Series: Annual_VIX_ts
```

```
ARIMA(1,0,0) with non-zero mean
```

The auto.arima function in R suggests that the AR (1) model with intercept is the better choice for the log-transformed annual VIX data, providing significant evidence that the AR (1) model outperforms the ARMA (1,1).

In a nutshell, the ACF and PACF of AR (1) are nearly identical to those of ARMA (1,1). The test results and error measures are also the same, with only slight differences. The coefficients of the AR component share the same sign and exhibit only minor magnitude variations. However, the AIC strongly favors the reduced AR (1) model, and the insignificance of the MA (1) term in ARMA (1,1) ($p = 0.495$) further supports the simplicity of AR (1). The residuals of AR (1) are white noise, with $\varepsilon_t \sim N(0, 0.04998)$, and while its standard deviation is very slightly higher than ARMA (1,1), this difference is negligible. This conclusion is further supported by R's auto.arima function, which independently selects AR (1) as the optimal model. **Thus, the final decision is clear: AR (1) is superior—it is simpler, all its coefficients are significant, and it achieves comparable performance without unnecessary complexity. We will therefore proceed with AR (1) for forecasting.**

Forecasting of AR (1):

```
> f3 <- forecast::forecast(Model2, h = 10)
> print(f2)
```

Point	Forecast	Lo 80	Hi 80	Lo 95	Hi 95
2025	2.795023	2.508531	3.081515	2.356871	3.233175
2026	2.841141	2.509697	3.172585	2.334241	3.348041
2027	2.867970	2.522636	3.213304	2.339827	3.396113
2028	2.883577	2.533667	3.233487	2.348436	3.418718
2029	2.892657	2.541212	3.244102	2.355168	3.430146
2030	2.897939	2.545976	3.249902	2.359658	3.436220
2031	2.901012	2.548873	3.253150	2.362463	3.439560
2032	2.902799	2.550602	3.254997	2.364160	3.441439
2033	2.903839	2.551622	3.256057	2.365169	3.442509
2034	2.904444	2.552220	3.256668	2.365763	3.443125

These are the forecast values for 10 years (2025–2034) using the AR (1) model, whose function is:

$$\widehat{VIX}_{n+1} = 2.9053 + 0.5817VIX_n + \varepsilon_t$$

Where: \widehat{VIX}_n is the estimated annual log-value of VIX for one step ahead forecast


As we know, AR models produce wider forecast ranges than MA models (which stop after q lags). Here, we forecast the average annual log-value of the VIX indicator using **conditional expectation given history of the series**. The exponential of these values will be taken to convert them back to the original scale, yielding the average annual forecasted VIX values.

Year	Log Value Forecast	Annual VIX Forecast ($e^{\text{Log Value}}$)
2025	2.795	16.36
2026	2.841	17.13
2027	2.868	17.59
2028	2.884	17.89
2029	2.893	18.05
2030	2.898	18.14
2031	2.901	18.19
2032	2.903	18.23
2033	2.904	18.25
2034	2.904	18.26

Interpretations:

The forecasts predict average annual VIX levels will gradually rise from 16.4 in 2025 to 18.3 by 2034, indicating relatively stable market conditions without extreme volatility events. These values, which represent typical yearly averages rather than potential short-term spikes, suggest market volatility will remain in a moderate range similar to calm historical periods. While useful for general planning, investors should remember these projections can't predict sudden market shocks, and actual volatility may vary around these baseline estimates. The model expects volatility to eventually stabilize near 18-19, reflecting its mean-reverting nature.

Forecasts from ARIMA(1,0,0) with non-zero mean



The forecast plot displays the AR (1) model's predictions for the log-transformed VIX index from 1990 through 2034. The model demonstrates a good fit to historical data, accurately capturing both high-volatility peaks and calm periods in the VIX's trajectory. Looking ahead, the forecast suggests the log (VIX) will stabilize between 2.9-3.0 in the coming decade (2025-2034), which translates to approximately 18-20 in standard VIX units. While the predictions indicate **relatively stable market conditions ahead**, with no extreme volatility events anticipated, **the gradually widening confidence bands reflect the expected increase in forecast uncertainty over longer time horizons**. However, it's important to note that such models cannot account for **sudden, unprecedented market shocks outside the range of historical experience**.

Classical Analysis – Decomposition Method:

As Modern approach has Box-Jenkins models the classical approach has Decomposition method, the idea behind decomposition itself was originally for calculating the orbits of the planets, Then, Persons founded the decomposition method with the four traditional components in 1919, Then any time series could be explained (decomposed) by the following function of Y and its components:

$$Y_t = f(T_t + S_t + C_t + R_t)$$

Where:

T_t : Secular trend with the underlying movement or long-term tendency of time series

S_t : seasonal movement which is the regular and the recurring fluctuations within each year, the pattern of which depends on the inherent characteristics of the time series.

C_t : Cyclical movements are long-term fluctuations that occur around the overall trend of a time series, characterized by alternating periods of increase and decrease. These cycles reflect recurring phases that are not tied to a fixed seasonal pattern.

R_t : Residuals, Random shocks, these variations that contain time-unrelated events such as wars, political changes or natural catastrophes which affect the dependent variable.

Applying this to VIX:

$$VIX_t = f(T_t + S_t + C_t + R_t)$$

There is two different models that could be stated: additive model and multiplicative model

As Seasonal S_t is dependent on trend T_t as S_t is a percentage of T_t so multiplicative model is preferred:


$$VIX_t = T_t * S_t * C_t * R_t$$

But to stabilize the variance a log-transformation is required:

$$\text{Log}(VIX_t) = \text{Log}(T_t) + \text{Log}(S_t) + \text{Log}(C_t) + \text{Log}(R_t)$$

So, the multiplicative model is converted into the additive model

Decomposition of additive time series



Observed (original time series): the plot shows a stable variance on the time period with spikes (outliers) in 2008 due to the financial crisis and in 2020 due to the coronavirus epidemic. In 2017 there is a low level of VIX indicating a calm market with investor confidence, high in the two crises mentioned above and it is stable all along due to mean reversion except for (2008 - 2009) recession.

Trend: From 2022 till (April 30th, 2025) there is an upward trend indicates an increasing economic uncertainty due to Russia-Ukraine war,


From 2013 till 2016 there is a downward trend because of stable economic growth and Central bank intervention,

During the remaining time VIX oscillates around its long-term average (mean reversion)

Seasonal: there is no major seasonality, but it is statistically significant, VIX is higher in the fourth quarter than in the first quarter but there is a cycle every 7- 10 years.


Residuals: Appears random along time period except (March 2020) covid pandemic and (june 2016) Brexit since it is black swan events with unexpected effect.

Detrended VIX Series



After applying first-difference detrending the time series oscillates around zero with no visible trend and with no clear pattern indicating random noise and probably cycles. However, there are high spikes in 2008 and 2020.

Deseasonalized VIX Series



The deseasonalized plot shows (7-10) - year oscillation may reflect cycles or structural breaks, with no strong long-term trend (except for time periods mentioned in the decomposition plot) since it reverts to its mean overtime, there is low volatility after 2010 (excluding 2020) due to Central bank intervention with no residual seasonality.

Trend Analysis:

First order polynomial:

$$VIX_t = \beta_0 + \beta_1 t + \varepsilon_t$$

Call:

```
lm(formula = clean_vix ~ clean_time)
```

Residuals:

Min	1Q	Median	3Q	Max
-0.77543	-0.27406	-0.03333	0.22674	1.46686

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	2.916e+00	7.244e-03	402.521	<2e-16 ***
clean_time	-2.327e-06	1.407e-06	-1.654	0.0982 .

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.342 on 8916 degrees of freedom

Multiple R-squared: 0.0003066, Adjusted R-squared: 0.0001945

F-statistic: 2.735 on 1 and 8916 DF, p-value: 0.09822

Applying to VIX:

$$\widehat{VIX}_t = 2.916 - 2.327e^{-06}t$$

From the previous model the following is concluded:

1. **Intercept (2.916):** The estimated starting value of the VIX series at time zero, with p-value < 0.05.
2. **Time:** (2.327e-06): Suggests a very slight downward trend over time, though it's insignificant (p=0.0982).
3. Only 0.03% of variance is explained by time.
4. **Residuals** have a relatively symmetric distribution.
5. The whole model is statistically insignificant since F-statistic = 2.735 with p-value (0.09822) > 0.05.

This suggests that the series is:

1. Nearly stationary (very weak trend component).
2. Dominance of other components*: The extremely low R-squared indicates most variation comes from other sources (seasonality, cyclical patterns, irregular components) rather than the trend
3. Mean-reverting behavior: The intercept provides an estimate of the long-run mean level (2.916) around which the series fluctuates which matches the observations above.

Second order polynomial:

$$VIX_t = \beta_0 + \beta_1 t + \beta_2 t^2 + \varepsilon_t$$

Call:

```
lm(formula = clean_vix ~ clean_time + I(clean_time^2))
```

Residuals:

	Min	1Q	Median	3Q	Max
	-0.77601	-0.26896	-0.02669	0.22221	1.48326

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	2.833e+00	1.080e-02	262.185	<2e-16	***
clean_time	5.376e-05	5.595e-06	9.608	<2e-16	***
I(clean_time^2)	-6.288e-09	6.074e-10	-10.353	<2e-16	***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.34 on 8915 degrees of freedom

Multiple R-squared: 0.01218, Adjusted R-squared: 0.01196

F-statistic: 54.97 on 2 and 8915 DF, p-value: < 2.2e-16

Applying to VIX:

$$\widehat{VIX}_t = 2.833 + 5.37e^{-5}t + 6.288e^{-9}t^2$$

From the previous out that reveals:

1. The estimated starting value of the VIX series at time zero is 2.833.
2. An initial upward tendency from positive linear term ($5.376e^{-5}$).
3. A slowing/reversal pattern from the negative quadratic term ($-6.288e^{-9}$)

That Suggests possible structural change or regime shift over time.

While statistically significant, the model still explains only 1.2% of variation, indicating:

- Most variation remains unexplained by the time trend.
- Other components (cycles, shocks) dominate the series.

This represents a meaningful improvement over the simple linear trend model, capturing more complex temporal dynamics while still maintaining relatively well-behaved residuals but the variation needs to be explained.

Third order polynomial:

$$VIX_t = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_3 t^3 + \varepsilon_t$$

Call:

```
lm(formula = clean_vix ~ clean_time + I(clean_time^2) + I(clean_time^3))
```

Residuals:

	Min	1Q	Median	3Q	Max
	-0.72400	-0.25901	-0.02444	0.21257	1.50721

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	2.699e+00	1.425e-02	189.41	<2e-16 ***
clean_time	2.338e-04	1.383e-05	16.90	<2e-16 ***
I(clean_time^2)	-5.675e-08	3.604e-09	-15.74	<2e-16 ***
I(clean_time^3)	3.772e-12	2.656e-13	14.20	<2e-16 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.3362 on 8914 degrees of freedom

Multiple R-squared: 0.03403, Adjusted R-squared: 0.0337

F-statistic: 104.7 on 3 and 8914 DF, p-value: < 2.2e-16

Applying to VIX:

$$\widehat{VIX}_t = 2.699 + 2.338e^{-4}t - 5.675e^{-8}t^2 + 3.772e^{-12}t^3$$

From the previous output the following is concluded.

- Cubic Term ($3.772e^{-12}$):
 1. Positive coefficient indicates an S-shaped curve (inflection point)
 2. Captures acceleration/deceleration in trends not explained by quadratic model.
- Quadratic Term ($5.675e^{-8}$):

Maintains concave shape but now part of more complex pattern.
- Linear Term ($2.338e^{-4}$):

Stronger upward trend than previous models

2. Model Improvement

R-squared jumps to 0.034 (3.4% variance explained that is triple **better than quadratic model**)

Adj. R-squared (0.0337): Confirms meaningful improvement.

Residual SE drops to 0.3362 (better fit than previous 0.34)

F-statistic = 104.7 (p<2.2e-16): Extremely significant overall model

The series exhibits **nonlinear dynamics with changing curvature:**

- Initial upward trend (linear)
- Early deceleration (quadratic)

- Later possible re-acceleration (cubic)

Structural Changes:

The cubic term suggests **multiple phases** in the series evolution:

- Early growth phase
- Intermediate moderation
- Potential late-stage trend reversal

Substantial Improvement:

- The cubic model explains **3.4% of variation** vs 1.2% (quadratic) and 0.03% (linear)
- Residuals show better properties despite model complexity.

Model Diagnosis and comparison:

	MSE	MAD	RMSE	MAPE
Linear	0.1169533	0.2783414	0.3419843	9.599519
Quadratic	0.1155639	0.2763929	0.3399469	9.543565
Cubic	0.1130082	0.2704806	0.3361670	9.336299

From the previous table and after reviewing error measurements the cubic model is the best model.

Forecasting with the cubic model:

	Date	Fit	Lower	Upper
1	2025-05-01	2.928016	2.268312	3.587720
2	2025-05-02	2.925179	2.265474	3.584884
3	2025-05-05	2.913096	2.253390	3.572802
4	2025-05-06	2.918353	2.258646	3.578060
.
.
.
247	2026-04-10	2.923310	2.263310	3.583311
248	2026-04-13	2.934290	2.274289	3.594292
249	2026-04-14	2.931529	2.271526	3.591533
250	2026-04-15	2.934905	2.274900	3.594909


From the previous forecast it appears:

- Initial Stability: May-June 2025 forecasts hover around 18-20 VIX ($\log \approx 2.9-3.0$)
- Summer Dip: July-August 2025 shows moderate decline ($\log \approx 2.8-2.9 \rightarrow 16-18$ VIX)
- Q4 2025 Volatility: September-December 2025 forecasts increase ($\log 3.0-3.1 \rightarrow \sim 20-22$ VIX)
- Year-End Spike: December 2025 peaks at $\log 3.17 \rightarrow \sim 23.8$ VIX
- 2026 Fluctuations: Continues oscillating between $\log 2.9-3.2$ (18-24 VIX).
- Prediction Intervals:
- Consistent Width: $\sim \pm 0.32$ in log scale translates to $\pm 38\%$ variation in original scale
- The market expects potential volatility swings of $\pm 38\%$ from the forecasted values.

The forecast suggests:

- No imminent market crisis (VIX < 30)
- Moderate volatility expectations (18-24 range)
- Seasonal patterns dominate over structural trends.
- Higher uncertainty in Q1 2026 (post-election year effect)

Historical vs Forecasted Yearly VIX Averages



Year	(Fit)	(Lwr)	(Upr)
2025	19.30	4.29	34.32
2026	20.22	5.19	35.24
2027	20.86	5.82	35.89
2028	21.60	6.55	36.65
2029	22.44	7.37	37.51
2030	23.40	8.30	38.49
2031	24.49	9.36	39.62
2032	25.73	10.55	40.91
2033	27.12	11.88	42.35
2034	28.65	13.34	43.96

1. Upward Trend:

- The VIX is forecasted to **increase steadily** from an average of **19.3 in 2025** to **28.65 in 2034**.
- This suggests a gradual rise in **expected market volatility** over the next decade.

2. Widening Confidence Intervals:

- The **prediction intervals** (Lwr–Upr) become **wider** over time.
- For example:
 - In 2025: [4.29 – 34.32] → Width ≈ 30.03
 - In 2034: [13.34 – 43.96] → Width ≈ 30.62
- This reflects **increasing uncertainty** in long-term forecasts, which is expected in time series modeling.

3. Interpretation of the VIX:


- The **VIX (Volatility Index)** represents market expectations of near-term volatility.
- Higher values indicate **greater expected risk/uncertainty** in the market.
- A value above ~25–30 often suggests **heightened investor anxiety or turbulence**, while levels below 20 reflect **market calm**.

4. Strategic Implications:

- **Risk management** becomes more critical over time if the VIX continues to trend upward.
- Long-term investors and institutions may consider **hedging strategies** or **volatility-adjusted portfolios**.
- The model implies that the market may face **greater uncertainty in the next decade**, whether due to macroeconomic, geopolitical, or structural shifts.

Please Note:

- These forecasts are based on a **cubic trend model with seasonal adjustment** and assume historical patterns persist.
- The **wide confidence intervals** highlight that actual future values may deviate substantially.
- Real-world events (e.g., financial crises, policy shifts) can invalidate these projections.



The cubic polynomial regression model applied to the VIX data from 1990 to 2010 offers valuable insight into long-term volatility trends.

1. The model successfully captures the broad cyclical behavior of the VIX, including an upward trend in the 1990s, a peak around 2002–2003, and a decline leading up to and following the 2008 financial crisis.
2. It performs better than linear and quadratic models in identifying medium-term trends and smoothing out short-term noise while retaining structural patterns.
3. Its explanatory power is limited. It accounts for only 3.4% of the variance and fails to reflect volatility clustering or mean-reverting behavior, making it unreliable during crisis periods.
4. The forecast should be interpreted with caution, especially post-2010, as time-based models overlook macroeconomic drivers and significantly underestimate market stress scenarios.

To improve accuracy, hybrid approaches combining this model with GARCH or stochastic volatility models are recommended. Additionally, incorporating macroeconomic variables and testing for structural breaks or regime changes could enhance model responsiveness. Overall, while the cubic model provides a useful macro view, it remains insufficient for precise forecasting without accounting for the VIX's variations.

Conclusion:

This study explored the time series characteristics of the VIX Index using the Box-Jenkins (ARIMA) methodology and the Classical Decomposition approach. Key findings indicate that the log-transformed annual VIX series exhibits stationarity, mean-reversion, and no significant seasonality. The AR(1) model emerged as the optimal choice for forecasting, producing stable annual forecasts with a Mean Absolute Percentage Error (MAPE) of 6.67%, suggesting moderate accuracy. Meanwhile, the cubic polynomial decomposition model captured nonlinear trends but explained only 3.4% of the variance, highlighting the dominance of non-trend components like cyclical fluctuations and random shocks. Forecasts from both methods project average VIX levels between 16–24 from 2025–2034, signaling expectations of relatively calm markets absent major disruptions.

Model Limitations

1. **Structural Breaks:** The models do not account for "black swan" events (e.g., pandemics, geopolitical crises), which can cause abrupt volatility spikes unreflected in historical data.
2. **Volatility Clustering:** ARIMA assumes independence of residuals, underestimating the persistence of high-volatility regimes common in financial markets.
3. **Leverage Effect:** Asymmetric volatility responses—where negative returns increase volatility more than positive returns—are not incorporated, limiting predictive accuracy during market downturns.
4. **Macroeconomic Factors:** The absence of external variables (e.g., interest rates, GDP growth) restricts the models' ability to adapt to shifting economic conditions.
5. **Linearity Assumptions:** ARIMA and polynomial models assume linear or deterministic relationships, ignoring nonlinear dynamics inherent in financial data.
6. **Data Aggregation:** Annual averaging smooths out intra-year volatility, masking short-term risks critical for tactical decision-making.
7. **Parameter Instability:** Model coefficients derived from historical data may become obsolete in evolving markets, reducing long-term reliability.
8. **Explanatory Power:** The cubic model's low R-squared (3.4%) underscores its inability to explain the majority of variance, relying instead on oversimplified time trends.

Future Directions

To address these limitations, hybrid frameworks combining ARIMA with GARCH or machine learning models (e.g., LSTM) could better capture volatility clustering and nonlinear patterns. Incorporating macroeconomic indicators and regime-switching mechanisms would enhance responsiveness to structural breaks. Additionally, high-frequency data and ensemble methods may improve short-term forecasting precision.

Practical Implications:

While the models provide a baseline understanding of VIX behavior, practitioners should supplement these forecasts with real-time risk assessments and scenario analyses. The results remain most applicable for strategic planning in stable markets, emphasizing the need for adaptive hedging strategies during uncertainty.

In conclusion, this analysis advances the academic discourse on volatility modeling while underscoring the trade-offs between simplicity and realism in financial forecasting. Future research must bridge these gaps to deliver robust tools for navigating increasingly complex markets.

References:

- Apergis, N. (2021). The impact of COVID-19 uncertainty on the VIX index. *Finance Research Letters*, 38, 101483. <https://doi.org/10.1016/j.frl.2020.101483>
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327. [https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1)
- Bouri, E., Jain, A., Roubaud, D., & Kristoufek, L. (2019). Geopolitical risk and the VIX index. *Finance Research Letters*, 29, 55–60. <https://doi.org/10.1016/j.frl.2019.03.008>
- Cao, L., & Tay, F. E. (2003). Support vector machine with adaptive parameters in financial time series forecasting. *IEEE Transactions on Neural Networks*, 14(6), 1506–1518. <https://doi.org/10.1109/TNN.2003.820556>
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263. <https://doi.org/10.1080/07350015.1995.10524599>
- Federal Reserve Bank of St. Louis. (2024). CBOE Volatility Index: VIXCLS. FRED. <https://fred.stlouisfed.org/series/VIXCLS> (link of the data)
- Giot, P. (2005). Relationships between implied volatility indexes and stock index returns. *The Journal of Portfolio Management*, 31(3), 92–100. <http://dx.doi.org/10.3905/jpm.2005.500363>
- Hamilton, J. D., & Susmel, R. (1994). Autoregressive conditional heteroskedasticity and changes in regime. *Journal of Econometrics*, 64(1–2), 307–333. [https://doi.org/10.1016/0304-4076\(94\)90067-1](https://doi.org/10.1016/0304-4076(94)90067-1)
- Hansen, P. R., & Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? *Journal of Applied Econometrics*, 20(7), 873–889. <https://doi.org/10.1002/jae.800>
- Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLOS ONE*, 13(3), e0194889. <https://doi.org/10.1371/journal.pone.0194889>
- Prakash, S. (2023). 34-Year Daily Stock Data. Kaggle. <https://www.kaggle.com/datasets/shiveshprakash/34-year-daily-stock-data>
- Wang, Z., Zhang, M., & Wang, W. (2020). Forecasting volatility index using hybrid ARIMA-LSTM model. *Physica A: Statistical Mechanics and its Applications*, 545, 123727. <https://doi.org/10.1016/j.physa.2019.123727>