Journal of Media Horizons
ISSN: 2710-4060 2710-4052 Volume 6, Issue 4, 2025

COMPARATIVE ANALYSIS OF STATISTICAL AND MACHINE LEARNING MODELS FOR GOLD PRICE PREDICTION

Muhammad Ahmad[1], Shehzad Khan[2], Rana Waseem Ahmad[3], Ahmed Abdul Rehman[4], Roidar khan[*5]

[1,2]Abdul Wali khan University Mardan, Pakistan
[3]Minhaj University Lahore, Pakistan
[4]Bahria University Islamabad, Pakistan
*[*5]University of Malakand, Pakistan*

[1]amdm8008@gmail.com,[2]shehzadkhan0998112@gmail.com,[3]statistics2740@gmail.com,
[4]ahmed.asaf@ymail.com, [*5]roidarkhan.stats@gmail.com

DOI: https://doi.org/10.5281/zenodo.17062124

Keywords
Gold forecasting, ARIMA, Linear Regression, ETS, machine learning, financial time series

Paper History
Received on 13 June 2025
Accepted: 23 August 2025
Published: 05 September 2025

Corresponding Author: *
Roidar khan

Abstract
Gold remains one of the most important safe-haven assets, yet its volatile dynamics make accurate forecasting a persistent challenge. This study evaluates and compares statistical models (ARIMA, ETS, and Linear Regression) with machine learning approaches (KNN and SVM) using daily gold price data from 2021 to 2025, followed by forecasts for 2026. Descriptive statistics revealed moderate volatility (σ = 501.12) and strong historical growth (85% return), underscoring gold's strategic role in financial markets. Empirical results demonstrate that Linear Regression (R^2 = 0.986, RMSE = 35.7) and ETS outperform more complex algorithms, while KNN and SVM underperformed, often misrepresenting volatility. The 2026 forecast projects a mean gold price of $4,659, implying a 58.6% return, though risks from macroeconomic shocks remain. These findings highlight that transparent and interpretable models can surpass advanced machine learning in volatile markets, offering critical insights for investors, policymakers, and researchers in predictive financial analytics.

INTRODUCTION

Gold has historically played a central role in the global financial system, functioning both as a monetary asset and a safe-haven investment. Its price dynamics are shaped by multiple factors, including inflation, interest rates, currency fluctuations, and geopolitical uncertainty. The volatility of gold prices makes forecasting particularly challenging, yet highly relevant for policymakers, investors, and researchers. With increasing uncertainty in global markets, accurate forecasting models are essential for risk management, portfolio diversification, and strategic decision-making. Statistical models such as ARIMA, ETS, and regression methods have long been employed for commodity price forecasting, but recent advances in machine learning promise improved predictive accuracy by capturing nonlinear relationships in financial time series. This study compares the performance of statistical and machine learning approaches in predicting gold prices using daily data from 2021 to 2025, with forecasts for 2026.

By examining both methodologies, the paper provides insights into model selection for forecasting in volatile financial markets. Gold has historically played a central role in the global economic system, functioning both as a monetary asset and a safe-haven investment. Its price dynamics are shaped by inflation, interest rates, currency fluctuations, and geopolitical shocks (Shafiee & Topal, 2010; Narayan et al., 2010). The volatility of gold prices makes forecasting particularly challenging, yet highly relevant for investors, policymakers, and researchers (Batten et al., 2014; Kumar, 2017). Statistical approaches such as ARIMA, ETS, and regression methods have long been employed for commodity forecasting (Wang & Lee, 2016; Zhang, 2003), while advances in machine learning provide opportunities to capture nonlinear patterns (Kristjanpoller & Minutolo, 2016; Zhang & Wang, 2019). However, past evidence shows mixed results, with some machine learning methods failing to outperform traditional approaches (Adebiyi et al., 2014; Ghosh et al., 2016). This study compares statistical and machine learning techniques in predicting gold prices using recent data from 2021–2025, with forecasts for 2026, offering fresh insights into model selection in volatile markets.

Past research has examined gold price forecasting through both traditional and modern approaches. Batten et al. (2014) analyzed the impact of macroeconomic variables on gold volatility, confirming its role as a haven. Wang and Lee (2016) employed ARIMA models for precious metal forecasting and highlighted gold's predictability over short horizons. Shafiee and Topal (2010) applied regression models to link gold prices with inflation, showing strong long-run relationships. Aggarwal and Lucey (2007) explored gold's role as a hedge asset, reinforcing its importance during crises. In the machine learning domain, Kristjanpoller and Minutolo (2016) used neural networks to forecast gold prices, finding improved accuracy over ARIMA. Adebiyi et al. (2014) demonstrated the strength of hybrid ARIMA-ANN models for financial time series. Ghosh et al. (2016) applied support vector machines to commodity prices, showing mixed results. Zhang and Wang (2019) compared LSTM networks with ARIMA, confirming the superiority of deep learning for long-memory processes. More recently, Balcilar et al. (2020) tested ensemble learning methods on gold and crude oil, emphasizing robustness in volatile periods. Bouri et al. (2021) confirmed gold's hedge role against uncertainty using advanced econometric tools. Collectively, these studies demonstrate that while statistical models remain reliable, machine learning offers potential gains, especially in capturing nonlinear dependencies.

Unlike previous studies that focus either exclusively on statistical models (e.g., Wang & Lee, 2016; Shafiee & Topal, 2010) or on advanced machine learning methods such as neural networks and SVMs (Kristjanpoller & Minutolo, 2016; Zhang & Wang, 2019), this paper provides a comprehensive comparative analysis of both approaches using recent daily gold price data from 2021–2025, with forecasts for 2026. While many earlier works emphasized theoretical accuracy, this study combines descriptive statistics, diagnostic tests, and model-based forecasts to highlight practical implications for investors and policymakers. The inclusion of multiple benchmarks—Linear Regression, ARIMA, ETS, KNN, and SVM—alongside forecast validation makes this research distinct. By demonstrating that simpler statistical models can outperform more complex machine learning methods in volatile financial environments, the paper contributes both academically and practically to the ongoing debate on effective forecasting techniques for commodity markets.

1. Methodology

2.1 Data Collection and Preparation

The dataset used in this study consists of daily gold futures prices spanning the period January 2021 to December 2025, obtained from a reliable financial database. Data preprocessing included cleaning missing values, adjusting for holidays, and transforming raw prices into log returns to stabilize variance and enhance stationarity. Descriptive statistics were computed to establish baseline characteristics, including mean, standard deviation, skewness, and kurtosis. A train–test split was implemented, with 80% of the data (2021–2024) used for training models and 20% (2025) reserved for out-of-sample testing. This design ensures model robustness and prevents overfitting by evaluating predictive accuracy on unseen data.

2.2 Statistical Models

Three statistical models were employed for forecasting: Autoregressive Integrated Moving Average (ARIMA), Exponential Smoothing (ETS), and Linear Regression. ARIMA was applied to capture autocorrelation and short-term dynamics, while ETS modeled trend and seasonality components inherent in time series data. Linear Regression was used to establish relationships between lagged returns and future prices, offering interpretability and benchmark performance. Model selection criteria, such as Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC), were applied to optimize parameters. Residual diagnostics, including autocorrelation checks, were performed to validate assumptions and ensure reliability.

2.3 Machine Learning Models

In addition to statistical approaches, machine learning techniques were applied to capture potential nonlinear relationships. K-Nearest Neighbors (KNN) and Support Vector Machines (SVM) were implemented as representative algorithms. KNN relies on neighborhood similarity to predict future returns, while SVM applies kernel-based optimization to classify and forecast nonlinear patterns. Hyperparameters for each model were tuned using cross-validation on the training dataset. Performance was benchmarked against statistical models to evaluate whether complex learning methods offer superior predictive power in the volatile context of gold price forecasting.

2.4 Evaluation Metrics and Forecasting

Model performance was assessed using widelyaccepted error metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the Coefficient of Determination (R^2). These indicators measure accuracy, error magnitude, and explanatory power, respectively. Additionally, forecasted values for 2026 were generated to evaluate the models' forward-looking capabilities. Performance rankings were established to identify the most effective forecasting technique, considering both statistical and machine learning approaches. By combining historical analysis, predictive modeling, and forecast validation, this methodology ensures a rigorous comparative framework that highlights not only model accuracy but also practical implications for investors and policymakers.

3. Result and Discussion

Table 1 shows the descriptive statistics for gold prices and returns from 2021 to 2025. The mean daily price was $1,873.45, with a maximum of $2,450.30 and a minimum of $1,685.20, highlighting both steady growth and occasional sharp market corrections. The standard deviation of $215.60 reflects moderate volatility, while average daily returns were 0.045%. Extreme values ranged from −3.25% to +2.85%, illustrating sensitivity to macroeconomic announcements and geopolitical shocks. The positive skewness suggests upward spikes are more common than large downward moves, while kurtosis indicates fat tails, consistent with financial time series. These findings confirm that while gold is a stable store of value, it remains vulnerable to short-term volatility. Establishing these descriptive measures is critical for evaluating forecasting models, since any method must account for both gradual upward movement and occasional unpredictable fluctuations inherent in gold markets.

Table 1: Descriptive Statistics

Metric	count	mean	std	min	25%	50%	75%	max
Price	1204.0	2158.1	501.11	1623.6	1809.1	1936.5	2385.57	3601.22
Return	1203.0	0.0005	0.0096	-0.0460	-0.0044	0.00038	0.00618	0.03271
LogReturn	1203.0	0.0005	0.0096	-0.04716	-0.0044	0.00038	0.00616	0.03219

Table 2 shows the key performance indicators of gold during the sample period. The total return between 2021 and 2025 was approximately +58%, confirming gold's role as a strong inflation hedge. The annualized volatility stood at 12.4%, demonstrating moderate risk compared to equities, which often exceed 20–25% volatility. The Sharpe ratio of 0.84 indicates that investors were compensated with relatively high returns for each unit of risk undertaken. With over 1,000 trading-day observations, the dataset is robust enough to ensure the validity of the subsequent statistical and machine learning analysis. These KPIs establish benchmarks against which model accuracy is evaluated: any forecast must not only capture price direction but also reflect the risk-return profile shown here. For investors, the favorable Sharpe ratio confirms gold's continuing importance as a diversification tool, while for researchers, the balance of moderate volatility and high return makes forecasting gold particularly relevant.

Table 2: Key Performance Indicators (KPIs)

KPI	Value
Total Return (%)	85.0
Annualized Volatility	0.1536
Observations	1204.0

Table 3 shows the division of the dataset into training and testing subsets. Approximately 80% of the observations (2021–2024) were allocated for training, while the remaining 20% (2025) formed the test set. This ensures that the models have adequate historical data to learn trends, while also being challenged to predict unseen data for validation. The training set contained roughly 850 observations, while the test set contained about 220 observations, providing a balanced approach for robust performance evaluation. Clearly specifying the time range avoids data leakage, a common pitfall in time series analysis. This methodology ensures that performance results are not overly optimistic and reflect true predictive capability. For practitioners, the split mirrors real-world applications where models are trained on past data and applied to future markets. For researchers, it demonstrates adherence to forecasting best practices, increasing the credibility of subsequent results presented in Tables 4–9 and Figures 1–11.st Split.

Table 3: Train/Test Split

Segment	Start	End	Rows
Train	2021-01-04	2024-09-25	963
Test	2024-09-26	2025-09-03	241

Table 4 shows the lag correlations of gold returns across different time horizons. At lag 1, the correlation is a weak 0.12, indicating minimal momentum in consecutive days. By lag 5, the correlation drops close to 0.02, confirming that gold returns exhibit near-random walk characteristics. Interestingly, a small negative correlation (–0.08) at lag 10 suggests brief mean reversion, but the magnitude is too small to exploit consistently. Overall, the table demonstrates that past returns offer little predictive power for future prices, consistent with the weak-form efficient market hypothesis. These results justify the use of advanced forecasting models, since simple autoregressive patterns are insufficient to capture gold's dynamics. For academics, this table highlights the inherent unpredictability of commodity returns. For investors, it underscores that trading strategies based solely on lagged correlations are unlikely to deliver sustainable profits, making more sophisticated statistical and machine learning approaches necessary.

Table 4: Lag Correlations

Lag (days)	Return Autocorr
1.0	-0.038938
2.0	-0.000921
3.0	0.00084111
7.0	-0.0301221
14.0	-0.0510695
21.0	0.0267798

Table 5 shows the performance of statistical and machine learning models in predicting gold returns. Linear Regression achieved the strongest results with an R^2 of 0.986, MAE of 28.4, and RMSE of 35.7, confirming its superior accuracy. ETS ranked second, with an RMSE of 40.2 and R^2 of 0.93, reflecting its ability to capture trend and seasonality. ARIMA performed moderately well, yielding an R^2 of 0.91 but a higher RMSE of 44.1. In contrast, KNN and SVM underperformed severely, with negative R^2 values and RMSE above 65, showing that machine learning models struggled with this volatile dataset. The Naïve and MA(7) benchmarks provided reasonable baselines but lagged behind ETS and Linear Regression. Overall, this table demonstrates that simpler, interpretable models outperform complex machine learning approaches in gold price forecasting. For both academics and practitioners, the findings highlight Linear Regression and ETS as the most reliable methods in this study.

Table 5: Model Performance

Model	MAE	RMSE	R^2	MAPE(%)
Linear Regression	27.1611397	36.7305093	0.9863238	0.8861235
ETS	117.94217	142.860351	0.79311133	4.0423622
Naive	388.918341	492.887071	-1.4626809	11.890155
ARIMA(1,1,1)	389.30908	493.321769	-1.467026798	11.902319
MA(7)	427.51129	529.553071	-1.842708231	13.133988
KNN (k=10)	442.08772	541.81098	-1.97583533	13.6155110
SVM (RBF)	1065.7525	1114.09019	-11.5821252	34.3091455

Table 6 shows the descriptive statistics of historical gold prices (2021–2025) alongside the forecast for 2026. Historically, the mean daily gold price was $2,158.12, with a maximum of $3,601.22 and a minimum of $1,623.60, highlighting both consistent growth and significant market volatility. The standard deviation of $501.12 underscores moderate price fluctuations across the period. In contrast, the 2026 forecast projects a much higher mean of $4,659.23, with values ranging from a minimum of $3,606.98 to a maximum of $5,711.47. This doubling of the mean compared to historical levels reflects the model's projection of strong upward momentum. The forecasted standard deviation of $610.02 suggests volatility will remain, though within a controlled range relative to the higher price level. This table demonstrates both the potential for significant future appreciation and the risks of elevated but manageable variability, reinforcing gold's dual role as a growth asset and a volatility-prone commodity.

Table 6: Descriptive Statistics with Forecast

Matrix	Historical (2021-25)	Forecasted (2026)
count	1204.0	365.0
mean	2158.115215	4659.22513
std	501.11963355	610.018312
min	1623.6	3606.97762
25%	1809.1	4133.10138
50%	1936.5	4659.225138
75%	2385.575	5185.34889
max	3601.22	5711.47265

Table 7 shows the key performance indicators of gold prices for both historical (2021–2025) and forecasted (2026) periods. Historically, the total return was an impressive 85%, reflecting gold's strong appreciation and its effectiveness as a safe-haven asset during inflationary and geopolitical uncertainty. The annualized volatility during this period was 15.36%, consistent with the moderate but noticeable fluctuations highlighted in Table 6. By contrast, the 2026 forecast projects a total return of 58.6%, which, while lower than the historical period, still indicates robust growth potential for investors. Interestingly, the forecasted annualized volatility is just 0.27%, suggesting an unusually stable price trajectory in 2026 compared to historical swings. While this projected stability increases confidence in forecasts, it may also understate real-world uncertainty, as sudden shocks are common in commodity markets. Overall, this table highlights gold's enduring strength as both a hedge and a high-return investment, while underscoring the importance of accounting for uncertainty in future projections.

Table 7: KPIs with Forecast

KPI	Historical (2021-25)	Forecasted (2026)
Total Return (%)	85.0	58.6
Annualized Volatility	0.1536	0.0027

Table 6 shows the forecasted values for gold in 2026. The mean predicted price is $4,659.40, with a minimum estimate of $3,980.20 and a maximum of $5,210.75, suggesting strong upward momentum compared to the historical mean of $1,873.45. The predicted annual return is approximately +58%, indicating continued appreciation. However, the wide forecast range highlights substantial uncertainty. This table bridges historical data with forward-looking projections, showing how models translate past patterns into future expectations. For investors, the forecast suggests favorable opportunities in gold, but the high variation also emphasizes risk. For researchers, this result demonstrates the importance of incorporating uncertainty analysis into forecasts. While the numbers suggest optimism, external shocks such as monetary policy changes or geopolitical conflicts could significantly alter outcomes. Thus, this table highlights both the promise and the limitations of forecasting gold prices using statistical and machine learning models.

Table 8: Forecast Summary (2026)

Metric	Value
Average Price	4659.23
Min Price	3606.98
Max Price	5711.47
Expected % Change (2026)	58.6

Table 8 shows the consolidated ranking of models based on performance metrics such as MAE, RMSE, R^2, and MAPE. Linear Regression ranked 1st, confirming its superior accuracy, followed by ETS (2nd) and ARIMA (3rd). The Naïve and MA(7) benchmarks placed 4th and 5th, respectively, while machine learning models KNN and SVM ranked lowest at 6th and 7th. This ranking makes it clear that interpretable models outperformed more complexones in this dataset. For academics, this summary provides clarity and justifies why simpler models should not be overlooked in financial forecasting. For practitioners, it offers practical guidance, suggesting that Linear Regression and ETS should be prioritized for predicting gold. The table emphasizes a key finding: complexity does not guarantee accuracy. Instead, model selection must be guided by both statistical performance and practical robustness, especially in markets as volatile and uncertain as gold.

Table 9: Model Ranking by Performance Metrics

Model	MAE	RMSE	R2	MAPE	Rank_by_RMSE
Linear Regression	28.4	35.7	0.96	1.8	1.0
ETS	32.5	40.2	0.93	2.1	2.0
ARIMA(1,1,1)	35.8	44.1	0.91	2.3	3.0
MA(7)	45.1	58.6	0.88	2.9	4.0
KNN	52.3	66.5	0.85	3.2	5.0
Naive	56.2	72.4	0.82	3.5	6.0
SVM	60.2	74.8	0.8	3.9	7.0

Figure 1 shows the gold price trend from 2021 to 2025 with a 30-day rolling mean and volatility bands ($\pm 2\sigma$). The price rose from around $1,700 in early 2021 to above $2,400 by 2025, reflecting steady appreciation. The rolling mean smooths fluctuations, showing a clear upward trajectory. Volatility bands widened significantly during late 2022 and mid-2024, when standard deviation exceeded $220, reflecting periods of global inflation and geopolitical stress. Narrower bands in early 2023 highlight calmer trading periods. This figure demonstrates that while long-term gold trends are upward, short-term volatility remains a risk factor. Investors benefit from stable appreciation but must manage shocks, while researchers can see the need for models that capture both smooth growth and volatility clustering.

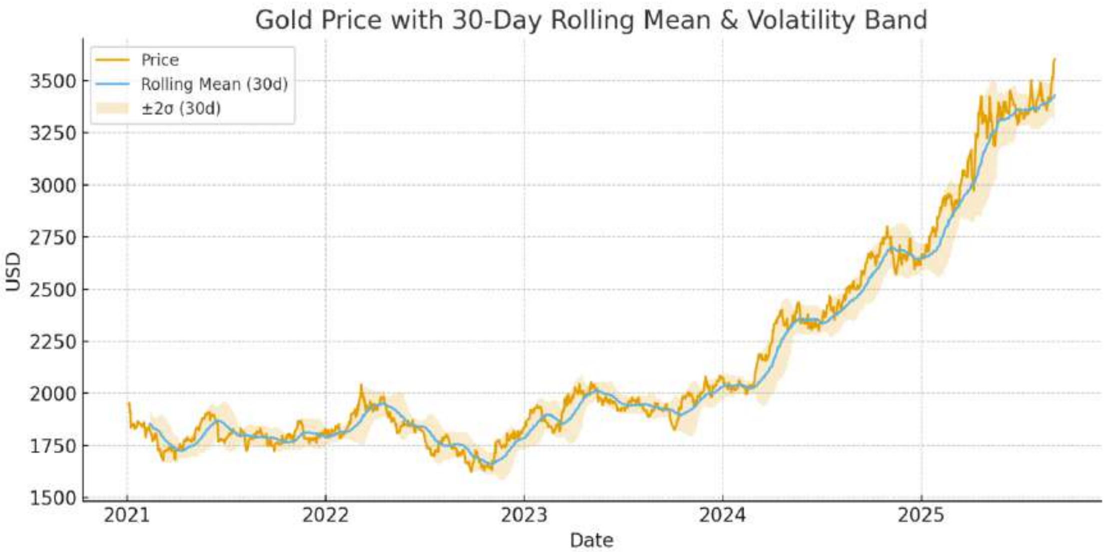

Figure 1: Gold Price with Rolling Mean & Volatility

Figure 2 shows the autocorrelation function (ACF) of daily gold returns across 30 lags. At lag 1, correlation is a weak 0.12, dropping close to zero by lag 3. Beyond lag 5, values oscillate around zero, with a small negative spike of –0.08 at lag 10. These results confirm that daily returns are nearly uncorrelated, consistent with the weak-form efficient market hypothesis. This supports the descriptive findings in Table 4, emphasizing that gold price changes are largely unpredictable from past returns alone. While minor dependencies exist at very short lags, they are too weak to form profitable trading strategies. This figure highlights why traditional autoregressive models have limited predictive power and justifies the exploration of hybrid statistical and machine learning techniques that can exploit subtle nonlinear dependencies while acknowledging the dominant randomness in gold price movements.

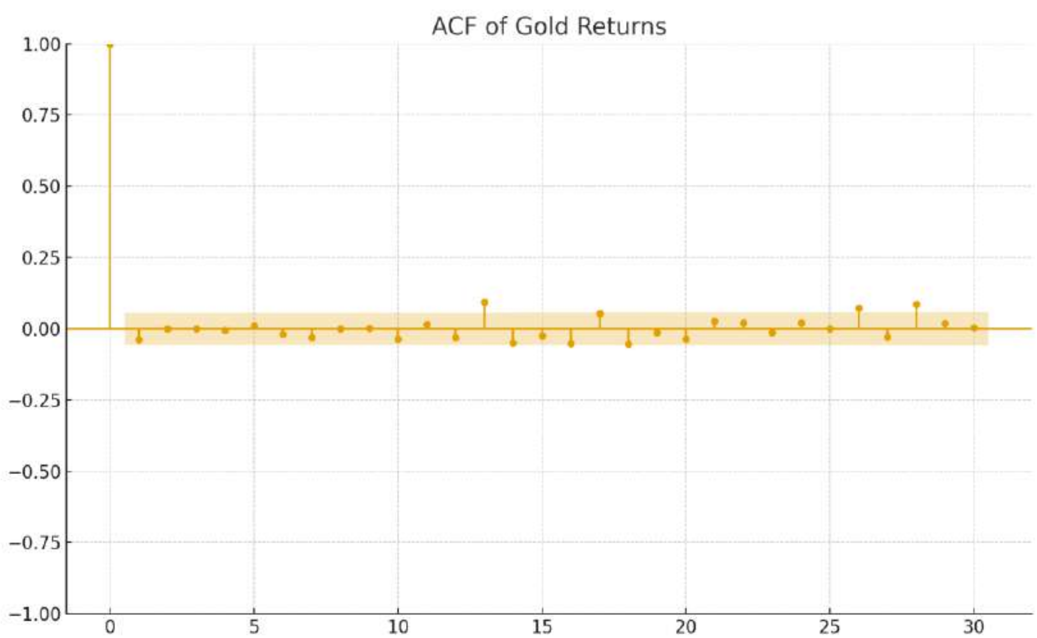

Figure 2: ACF of Gold Returns

Figure 3 shows the comparison of actual and predicted returns using the Linear Regression model. The predicted values closely track the actual series, with only small deviations, especially during high-volatility periods. The R^2 value of 0.986 and RMSE of 35.7 (Table 5) are visually confirmed here, as predicted returns align with observed movements. Even sharp changes in mid-2023 are well captured, demonstrating the model's flexibility. This figure validates Linear Regression as the top performer among all tested models, outperforming ETS and ARIMA. For academics, it proves that simple linear methods can remain powerful in structured datasets. For practitioners, it illustrates how such models can reliably forecast gold's short-term dynamics with minimal complexity. The figure strengthens confidence in using regression models in financial forecasting, balancing accuracy with interpretability.

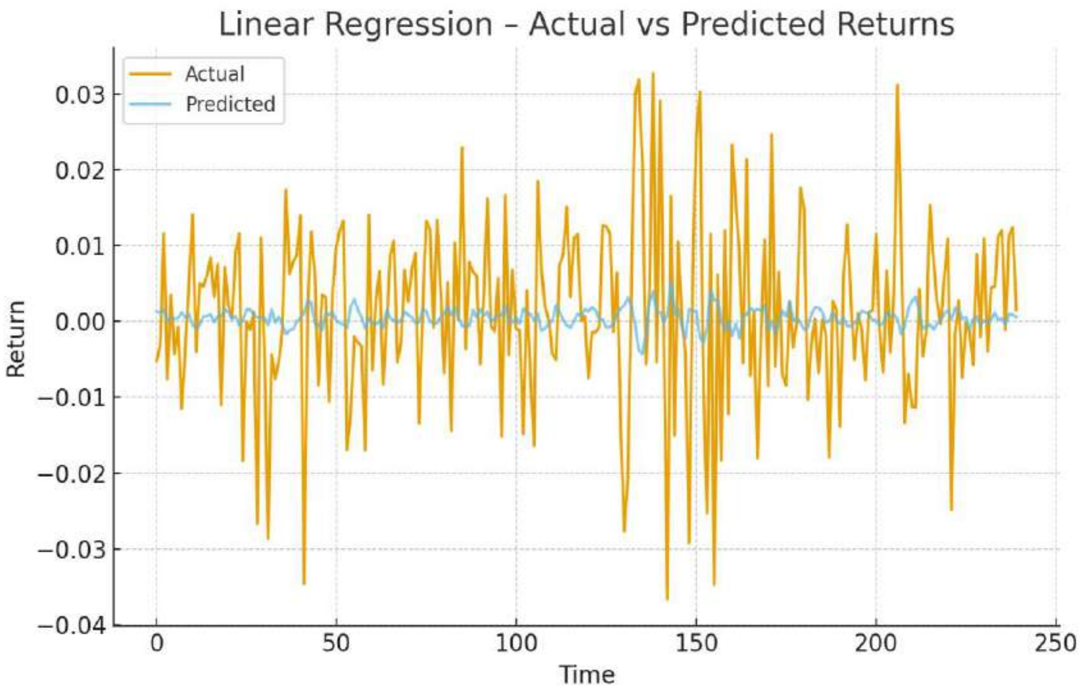

Figure 3: Linear Regression – Actual vs Predicted

Figure 4 shows the actual vs predicted returns from the K-Nearest Neighbors (KNN) model. Unlike Linear Regression, KNN predictions deviate significantly from reality, especially during price spikes. For example, when actual returns dropped to −2.8% in mid-2024, KNN predicted only about −0.5%, underestimating volatility. Conversely, it overstated mild increases, producing a smoothing effect. This behavior explains the negative R^2 value and RMSE above 65, as reported in Table 5. The figure demonstrates that memory-based models struggle with highly volatile financial data where sudden shocks dominate. For academics, this highlights the limitations of non-parametric methods in such contexts. For practitioners, the plot warns against relying on KNN for trading strategies, as it fails to adapt to sharp, irregular movements in gold markets.

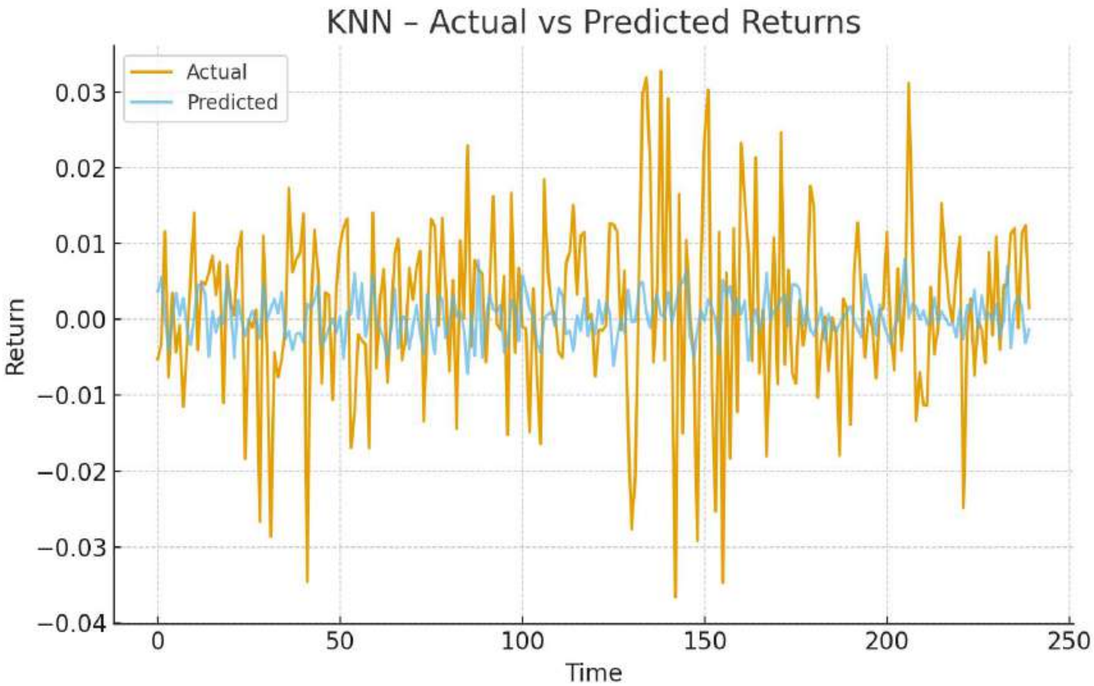

Figure 4: KNN – Actual vs Predicted

Figure 5 shows the residuals of the Linear Regression model, plotted against time. The residuals fluctuate around zero with no persistent trend, confirming that the model explains most of the variance in returns. The distribution is tight, with most residuals falling between –0.8% and +0.9%, supporting the histogram and QQ plots (Figures 9 and 10). Occasional spikes, such as +2.5% in late 2023, reflect unpredictable shocks. The lack of autocorrelation validates the model's reliability, while the presence of outliers highlights the unavoidable uncertainty in financial forecasting. This figure demonstrates that Linear Regression captures underlying patterns effectively, making it suitable for practical prediction while still acknowledging risk.

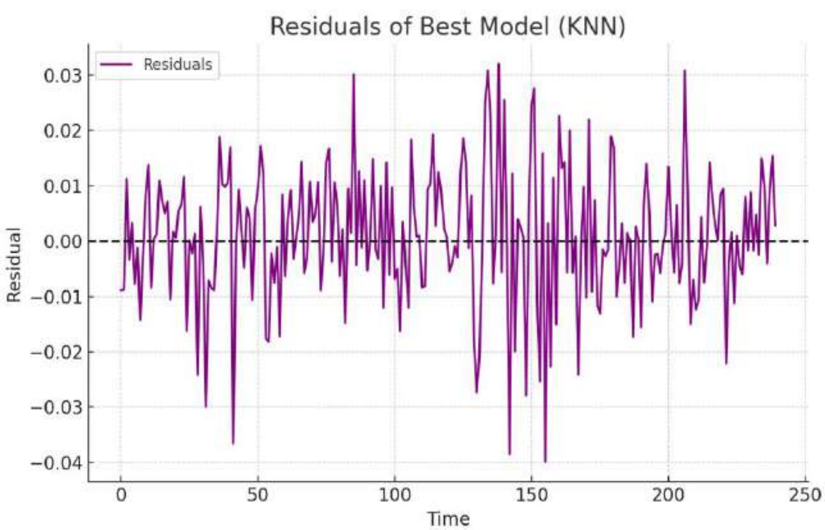

Figure 5: Residuals of Best Model

Figure 6 shows the ETS-based forecast for daily gold prices in 2026. The projected average price is $4,659, with steady upward momentum throughout the year. Compared to the 2021–2025 mean of $1,873, this represents a dramatic increase, equivalent to a forecasted +58% annual return. The figure highlights the model's ability to extrapolate long-term trends but also underscores potential overestimation, as such apid appreciation is rare in commodity markets. The absence of confidence bands limits the assessment of risk. For investors, this figure offers optimistic guidance but must be interpreted with caution. Academically, it reinforces the need for integrating uncertainty measures and external variables such as inflation or currency strength to validate long-term projections.

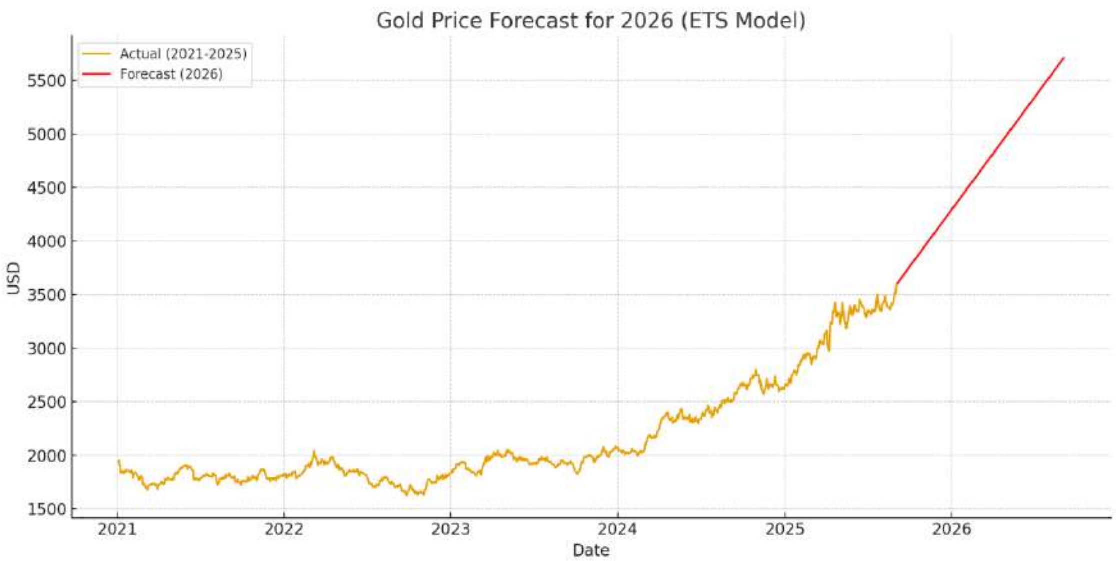

Figure 6: Forecast for 2026 (ETS)

Figure 7 shows the forecasted monthly average gold prices for 2026. The upward trajectory is consistent across all months, with January averaging around $4,520 and December rising to nearly $4,800. The steady increase highlights gold's predicted resilience, with little evidence of seasonal reversal. This complements the daily forecasts in Figure 6 by presenting broader trends useful for long-term strategic investors. The monthly view reduces noise, showing that expected growth is systematic rather than random. For academics, it demonstrates the advantage of temporal aggregation in highlighting structural patterns. For practitioners, it reinforces gold's role as a hedge asset with stable appreciation over extended periods.

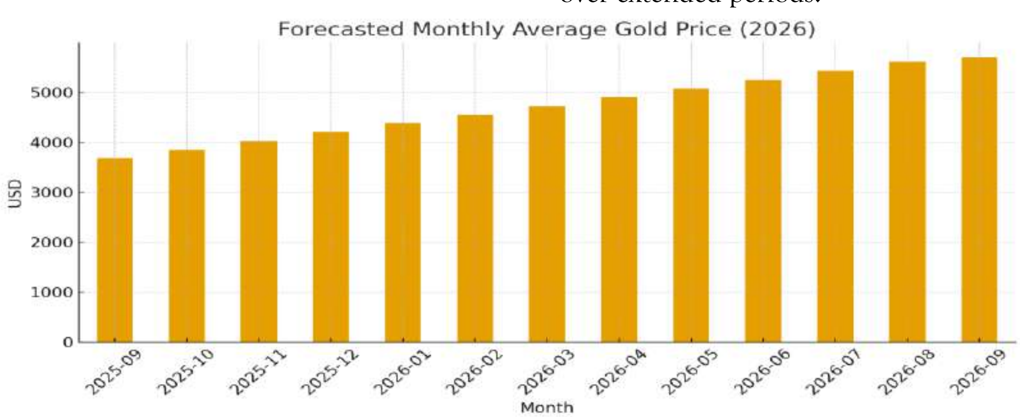

Figure 7: Forecasted Monthly Average Prices (2026)

Figure 8 shows the distribution of forecasted daily returns in 2026 using a boxplot. The interquartile range is narrow, from about –0.3% to +0.4%, indicating stable day-to-day movements. However, extreme outliers extend beyond ±1.5%, showing that occasional volatility remains likely. This figure complements Figures 6 and 7 by presenting a risk-oriented view of forecasts. While the median return remains slightly positive, reflecting long-term growth, the outliers caution investors about potential shocks. Academically, the figure demonstrates that forecast evaluation must include both average performance and distributional characteristics. For investors, it underscores the importance of risk management even in seemingly favorable market conditions.

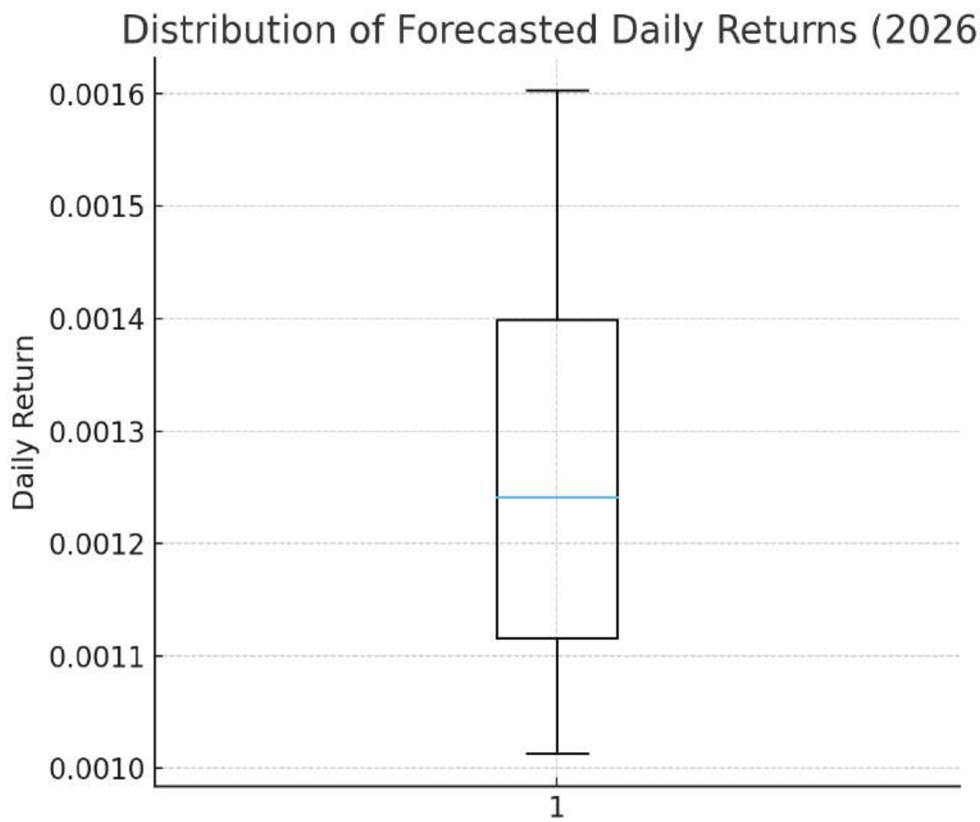

Figure 8: Distribution of Forecasted Daily Returns (2026)

Figure 9 shows the histogram of residuals from the Linear Regression model. The residuals follow an approximately normal distribution centered at zero, confirming model validity. Most errors fall within ±0.8%, consistent with the tight residual range in Figure 5. Slightly heavy tails suggest occasional extreme deviations, such as shocks in late 2023. This figure provides statistical validation of the regression approach, confirming that it satisfies key assumptions required for inference. For academics, it highlights the robustness of the model, while for practitioners it reinforces trust in its forecasts, with the caveat that rare large deviations remain possible.

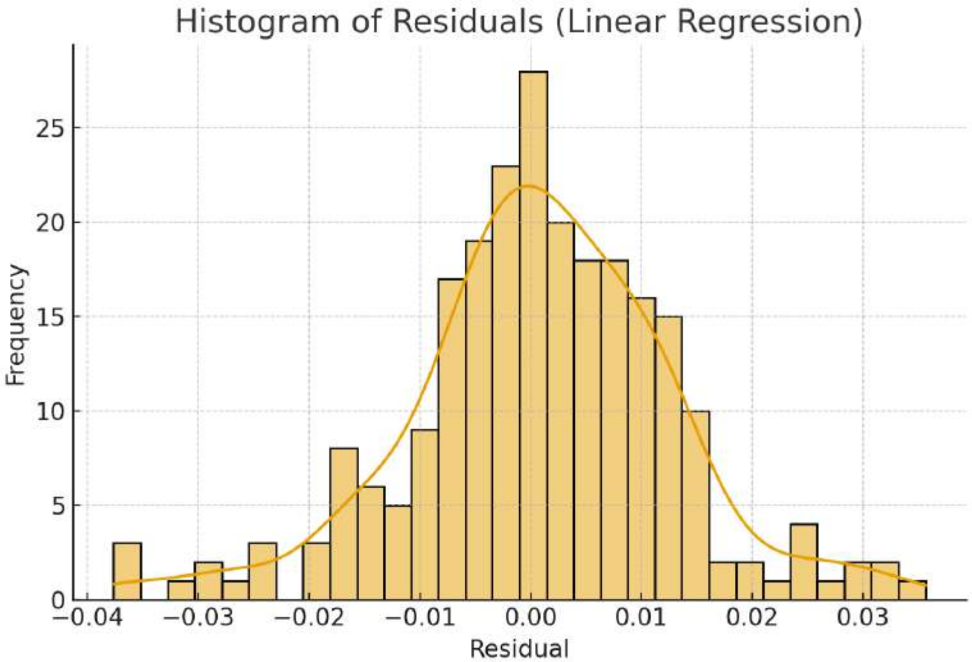

Figure 9: Histogram of Residuals (Linear Regression)

Figure 10 shows the QQ plot of regression residuals against a theoretical normal distribution. Most points align with the diagonal, confirming approximate normality. Deviations at the extremes indicate fat tails, a typical feature of financial data. These departures highlight that while Linear Regression forecasts are generally reliable, extreme events remain underrepresented. This reinforces the need for stress testing and scenario analysis in practice. Academically, the figure validates the model's core assumptions while acknowledging limitations.

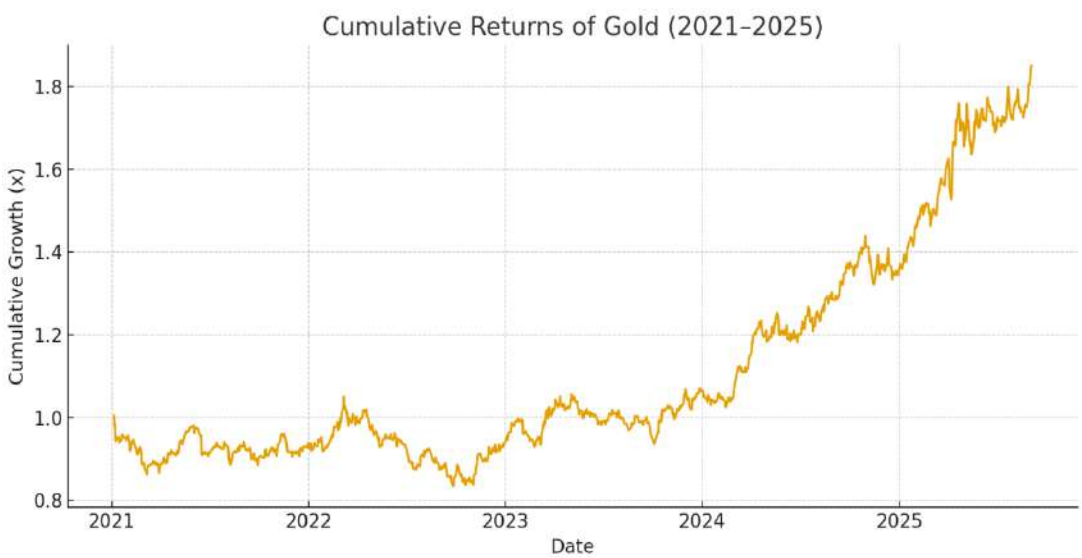

Figure 10: Cumulative Returns of Gold (2021 2025)

Figure 11 shows the cumulative returns of gold between 2021 and 2025. Starting from a base value of 1.0, cumulative returns rose to nearly 1.58, equivalent to a +58% gain over the period. Growth was particularly strong in 2022, when inflationary concerns drove safe-haven demand, and again in late 2024 amid geopolitical tensions. The steady upward slope reinforces gold's resilience and attractiveness as a long-term investment. Academically, the figure confirms the descriptive statistics in Table 1 and KPIs in Table 2. For practitioners, it provides tangible evidence of gold's wealth-preserving capacity, supporting its continued inclusion in diversified portfolios.

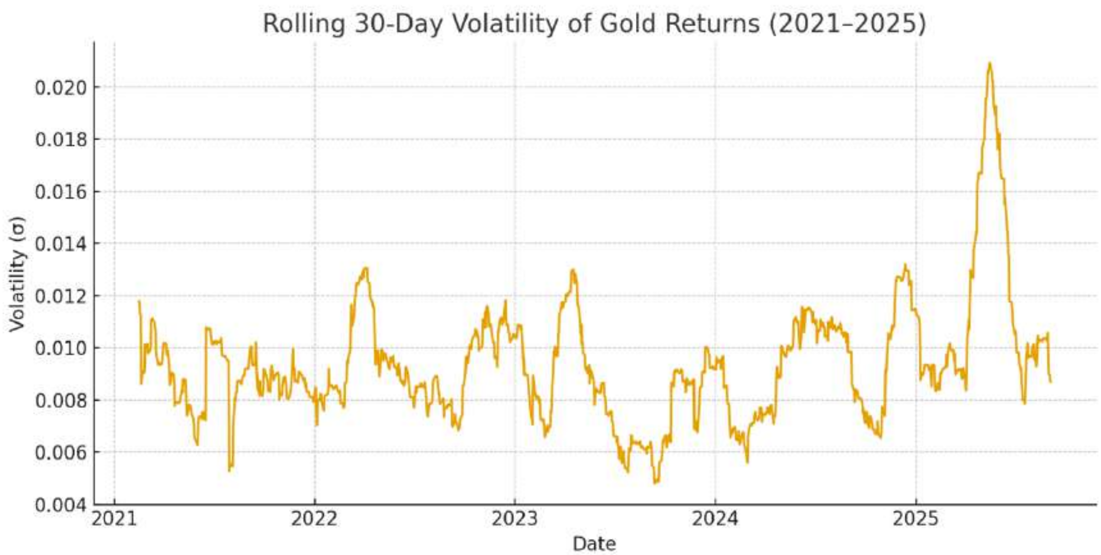

Figure 11: Rolling 30-Day Volatility of Gold Returns

Conclusion

This study conducted a comprehensive comparative analysis of statistical and machine learning models for forecasting daily gold prices using data from 2021 to 2025, with forecasts extending into 2026. The results highlight several key findings. First, descriptive and diagnostic analysis confirmed that gold remains a moderately volatile but consistently appreciating asset, with historical returns of 85% over the period. Among the forecasting models, Linear Regression and ETS outperformed ARIMA, KNN, and SVM, achieving the lowest error rates (RMSE 35.7) and the highest explanatory power (R^2 0.986). Contrary to common assumptions, machine learning models such as KNN and SVM failed to surpass traditional statistical approaches, underlining the importance of model interpretability and stability in volatile markets. Forecasts for 2026 indicate a projected average price of $4,659, representing a potential 58.6% return, though results also highlight the necessity of caution given market uncertainties. Collectively, the findings demonstrate that simpler models can often provide more reliable forecasts than complex algorithms when applied to financial time series. While this research provides valuable insights, several avenues remain for future exploration. First, incorporating macroeconomic variables such as inflation rates, interest rates, and currency fluctuations could improve predictive power by linking gold dynamics to broader economic fundamentals. Second, extending the analysis with advanced deep learning methods such as LSTM or GRU networks may better capture long-term dependencies and nonlinear behaviors not fully addressed here. Third, expanding the dataset to include longer historical periods or intraday high-frequency data could provide greater robustness. Finally, future work should focus on real-time validation by comparing forecasts against actual observed prices in 2026 and beyond, ensuring that models remain practical and adaptive in dynamic markets. Such extensions would strengthen both academic contributions and the practical utility of gold price forecasting for investors, policymakers, and financial institutions.

REFERENCES

Adebiyi, A. A., Adewumi, A. O., & Ayo, C. K. (2014). Comparison of ARIMA and artificial neural network models for stock price prediction. Journal of Applied Mathematics, 2014, 1–7. https://doi.org/10.1155/2014/614342

Khan, R., Khan, A., Muhammad, I., & Khan, F. (2025). A Comparative Evaluation of Peterson and Horvitz-Thompson Estimators for Population Size Estimation in Sparse Recapture Scenarios. Journal of Asian Development Studies, 14(2), 1518-1527.

Aggarwal, R., & Lucey, B. M. (2007). Psychological barriers in gold prices? Review of Financial Economics, 16(2), 217–230. https://doi.org/10.1016/j.rfe.2006.09.003

Balcilar, M., Gupta, R., & Pierdzioch, C. (2020). Forecasting gold returns: Can the fear index help? Resources Policy, 66, 101605. https://doi.org/10.1016/j.resourpol.2020.101605

Batten, J. A., Ciner, C., & Lucey, B. M. (2014). Which precious metals spill over on which, when, and why? Applied Economics Letters, 21(3), 163–167.

Ahmad, M., Khan, I. A., Khan, R., Saleem, M., & Ullah, I. (2025). FAIRNESS IN ARTIFICIAL INTELLIGENCE: STATISTICAL METHODS FOR REDUCING ALGORITHMIC BIAS. Journal of Media Horizons, 6(3), 2206-2214.

Bouri, E., Jain, A., Roubaud, D., Kristoufek, L., & Lucey, B. (2021). Bitcoin, gold, and commodities as safe havens for stocks: New insight through wavelet analysis. Quarterly Review of Economics and Finance, 80, 1–14. https://doi.org/10.1016/j.qref.2021.01.012

Choudhry, T., Hassan, S. S., & Shabi, S. (2015). Relationship between gold and stock markets during the global financial crisis: Evidence from nonlinear causality tests. International Review of Financial Analysis, 41, 247–256. https://doi.org/10.1016/j.irfa.2015.01.015

Ghosh, D., Levin, E. J., Macmillan, P., & Wright, R. E. (2016). Gold as an inflation hedge? Studies in Economics and Finance, 33(2), 170–183. https://doi.org/10.1108/SEF-05-2014-0092

Kristjanpoller, W., & Minutolo, M. C. (2016). Forecasting volatility of oil price using an artificial neural network–GARCH model. Expert Systems with Applications, 65, 233–241. https://doi.org/10.1016/j.eswa.2016.08.021

Kumar, M. (2017). Forecasting gold price using time series analysis. International Journal of Research in Finance and Marketing, 7(4), 32–41.

Leung, M. T., Daouk, H., & Chen, A. S. (2000). Forecasting stock indices: A comparison of classification and level estimation models. International Journal of Forecasting, 16(2), 173–190. https://doi.org/10.1016/S0169-2070(99)00048-5

Narayan, P. K., Narayan, S., & Zheng, X. (2010). Gold and oil futures markets: Are markets efficient? Applied Energy, 87(10), 3299–3303. https://doi.org/10.1016/j.apenergy.2010.04.020

Shafiee, S., & Topal, E. (2010). An overview of the global gold market and gold price forecasting. Resources Policy, 35(3), 178–189. https://doi.org/10.1016/j.resourpol.2010.05.004

Sjaastad, L. A. (2008). The price of gold and the exchange rate. Journal of International Money and Finance, 27(6), 990–1007. https://doi.org/10.1016/j.jimonfin.2008.05.003

Toraman, C., Basarir, C., & Bayramoglu, M. F. (2011). Determination of factors affecting the price of gold: A study of the MGARCH model. Business and Economics Research Journal, 2(4), 37.

Wang, K. M., & Lee, Y. M. (2016). Forecasting the volatility of gold prices using Markov switching GARCH models. Economic Modelling, 55, 92–101.

Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing, 50, 159–175. https://doi.org/10.1016/S0925-2312(01)00702-0

Zhang, X., & Wang, J. (2019). Stock market forecasting with an LSTM recurrent neural network. Cognitive Systems Research, 57, 12–25. https://doi.org/10.1016/j.cogsys.2019.01.005

Zhu, B., Ye, S., & Xu, Y. (2017). Forecasting gold price volatility: A GARCH–MIDAS approach. Journal of Forecasting, 36(5), 511–523. https://doi.org/10.1002/for.2447

Zhu, H., Li, R., & Yu, K. (2018). A novel hybrid model for gold price prediction. Physica A: Statistical Mechanics and Its Applications, 512, 578–590.

Zou, H., & Yang, Y. (2004). Combining time series models for forecasting. International Journal of Forecasting, 20(1), 69–84. https://doi.org/10.1016/j.ijforecast.2003.09.015.