Journal Of Economics and Management | VOLUME: 10/ N°: 01 (2026), PP 1-8
P-ISSN: 2572-0090 E-ISSN: 2710-799X

Comparative Analysis of Multi-Layer Perceptron Architectures for USD/DZD Exchange Rate Forecasting

Ahmed Benlaria [1*]
[1] Department of Finance and Accounting, Ahmed Draia University, Adrar (Algeria),
ahmed.benlaria@univ-adrar.edu.dz

Received: 02/03/2026 **Accepted: 01/05/2026** **Published: 10/06/2026**

Abstract :
Exchange rate forecasting is a cornerstone of international finance; however, it remains a significant challenge due to the nonstationarity and nonlinearity of financial markets. This study investigates the optimisation of deep neural network (DNN) architectures for predicting the USD/DZD exchange rate, a currency pair that is less frequently studied in emerging markets. Utilising daily data from the International Monetary Fund (IMF) spanning May 2010 to May 2024, this study employs systematic feature engineering, including lagged values and moving averages, to capture short- and medium-term market dynamics. We evaluated various multilayer perceptron (MLP) configurations, ranging from single-hidden-layer models to deeper, multi-layer structures. The results indicate that while all tested architectures achieved high predictive accuracy, a modest two-layer configuration—specifically ANN ([10, 5])—yielded the lowest error metrics, with a Mean Squared Error (MSE) of 3.4452×10^{-6} and a Root Mean Squared Error (RMSE) of 0.0018561. These findings suggest that increasing network complexity beyond a certain threshold offers diminishing returns when effective feature engineering is applied. This study underscores the practical utility of optimised DNNs for financial stakeholders and policymakers seeking to mitigate currency risk and enhance capital allocation in volatile economic environments.
Keywords: Exchange Rate Forecasting; Deep Neural Networks (DNN), USD/DZD; Multi-Layer Perceptron (MLP).
JEL Classification Codes: C45; F31; G17; C53.

I. Introduction

Exchange rate forecasting plays a pivotal role in international finance, influencing decision-making, risk management, and the formulation of effective economic policies in a globalised economy. Accurate forecasting enables financial managers and investors to anticipate future exchange rate movements, informing decisions on foreign investments, hedging strategies, and capital allocation (Eun, Resnick, & Chuluun, 2021). For firms engaged in cross-border trade, reliable predictions of currency fluctuations are essential for mitigating currency risk exposure through the judicious use of derivatives, such as forwards and options (Avdjiev, Bruno, Koch, & Shin, 2019). Beyond corporate interests, policymakers also rely on exchange rate forecasts to guide monetary policy decisions, ensuring economic stability and growth, even under conditions of uncertainty (Chen et al., 2020).

Notwithstanding their importance, exchange rate forecasting efforts are often hindered by the inherent complexity and volatility of the financial markets. Market conditions frequently exhibit non-stationarity—where statistical properties evolve over time—and intricate interdependencies that cannot be fully captured by traditional linear models. These complexities heighten the challenge of accurately predicting future movements, as rapid and often unpredictable fluctuations require continuous adaptation in both trading strategies and policy responses (Antonakakis, Chatziantoniou, & Gabauer, 2019; Kukacka & Kristoufek, 2020). Central banks, for instance, must remain vigilant and adaptable when crafting monetary policies in an increasingly intricate international financial landscape (Carney, 2019).

* Corresponding author.

In essence, exchange rate forecasting is both critical and challenging. To navigate this complex terrain, advanced modeling techniques that can handle nonlinearities, capture evolving market dynamics, and provide robust predictive performance are imperative.

The need for robust, accurate, and automated prediction models has grown as currency markets become more volatile and interconnected. Existing forecasting methods, including traditional econometric and shallow machine-learning approaches, often struggle with the intricacies of nonstationary data and the complex interplay among global financial factors. Deep neural networks (DNNs) offer a promising avenue for improvement, as their multilayered structures can approximate complex nonlinear functions, potentially yielding more reliable and precise forecasts.

Against this backdrop, our research is guided by the following questions: How can we leverage deep learning architectures to enhance the predictive accuracy of exchange rate forecasts? Which architectural configurations and feature engineering techniques offer the most substantial improvements in model performance? Answering these questions will help fill the gap in the literature where systematic assessments of different DNN architectures for currency forecasting remain limited.

The primary objective of this study is to optimize deep learning architectures for predicting the USD/Algerian Dinar exchange rate. We explore various neural network configurations—including different numbers of hidden layers and neurons—and systematically engineer input features, such as lagged values and moving averages, to more effectively capture underlying market dynamics.

This work contributes to the existing literature in several ways. First, it presents a comprehensive evaluation of multiple DNN architectures with varying degrees of complexity, providing insights into the trade-offs between model depth and predictive performance. Second, it demonstrates the effectiveness of systematic feature engineering within a deep learning framework, showing how carefully crafted inputs can enhance forecasting accuracy. Finally, by focusing on the USD/Algerian Dinar exchange rate, this study broadens the scope of currency forecasting research beyond the more commonly examined major currency pairs, thus offering valuable findings for researchers, practitioners, and policymakers operating in less frequently studied markets.

II. Literature Review

Early forecasting efforts relied heavily on traditional statistical models, particularly ARIMA and GARCH, to predict exchange rate movements. ARIMA models leverage autoregressive and moving average components to handle non-stationary data, while GARCH models focus on capturing volatility clustering in financial time series (Abreu, Souza, & Oliveira, 2019). Although these methods provide foundational tools for exchange rate analysis, they often struggle with nonlinear dynamics. Recent research suggests that incorporating nonlinearity, as in GARCH-NN frameworks, can enhance forecast accuracy (Charef, 2024). Furthermore, findings by Girgin (2023) highlight the sensitivity of ARIMA-GARCH approaches to specific distributional assumptions, underscoring the need for more robust methods that can adapt to volatile and complex financial environments.

Over the last decade, machine learning and shallow neural network approaches have gained traction in currency forecasting. While machine learning techniques, such as linear regression or support vector regression, offer interpretability and can efficiently model linear patterns (Mohammadi, 2019), they often fail to capture the nonlinear relationships inherent in financial markets (Nazareth & Reddy, 2023). Similarly, early neural networks applied to financial time series faced limitations in modeling intricate market structures and were prone to overfitting. These constraints highlighted the necessity for more sophisticated architectures that can learn richer representations and better approximate nonlinear dependencies.

Deep neural networks (DNNs) address these limitations by providing deeper architectures and more complex feature extraction capabilities. Research indicates that deep learning models outperform traditional statistical and shallow machine learning methods in forecasting market trends (Nguyen et al., 2024). They can leverage extensive datasets more effectively, discovering latent patterns that elude simpler models (Pandya, 2024). Moreover, DNNs demonstrate improved generalization by capturing complex interactions, leading to more robust and accurate predictions (Ozbayoglu, Gudelek, & Sezer, 2020).

Despite growing evidence that deep learning techniques can significantly enhance exchange rate forecasting, several gaps remain. Many studies focus on major currency pairs, overlooking currencies from emerging markets, where the complexities and nonlinearities may differ. In addition, the literature lacks a systematic assessment of how varying deep network architectures—differing in depth, layer composition, and feature engineering methods—affect forecasting performance. Few works provide comprehensive guidance on selecting and optimizing these architectures, leaving practitioners and researchers with limited insights into the configurations that yield the most reliable and accurate forecasts. Addressing these gaps is essential to fully harnessing the potential of deep learning for exchange rate prediction, particularly for under-studied currency pairs like USD/DZD.

III. Methodology

1. Data Collection and Preprocessing

This study utilizes daily USD/Algerian Dinar (DZD) exchange rate data spanning from May 24, 2010, to May 3, 2024, obtained from the International Monetary Fund (IMF) database. The chosen time frame provides extensive coverage of multiple market cycles, capturing various economic conditions and facilitating a robust modeling framework. Daily frequency ensures sufficient granularity to identify both short-term fluctuations and longer-term trends.

Prior to model training, a series of preprocessing steps were applied to the data:

2. Handling Missing Values:

Observations with missing values, particularly those introduced through the creation of lagged features, were removed. Eliminating incomplete entries helps maintain data integrity and ensures that model training is not compromised by gaps or inconsistencies.

3. Normalization:

To ensure that all input features contribute proportionally to the learning process, the exchange rates were normalized to a [0, 1] range using the following formula:

$$normalized_X = \frac{X - min(X)}{max(X) - min(X)}$$

Normalization stabilizes model training, aids in faster convergence, and avoids potential issues arising from differing feature scales [18][19].

4. Data Splitting:

The normalized dataset was divided into three subsets: training, validation, and testing. The training set was used to fit the model parameters, the validation set guided model selection and hyperparameter tuning, and the testing set provided an unbiased evaluation of the model's generalization capabilities.

5. Feature Engineering

Effective feature engineering is vital for capturing the underlying temporal dependencies and patterns in exchange rate series. In this study, several engineered features were introduced to improve the model's predictive power:

5.1. Lagged Values:

Lag features, such as lag7 and lag14, represent the normalized exchange rates from 7 and 14 days prior, respectively. By incorporating lagged values, the model can learn short-term dependencies, recognizing how recent price movements influence future trends [20]. This approach aids in capturing autocorrelations and identifying recurring patterns in the data.

5.2. Moving Averages:

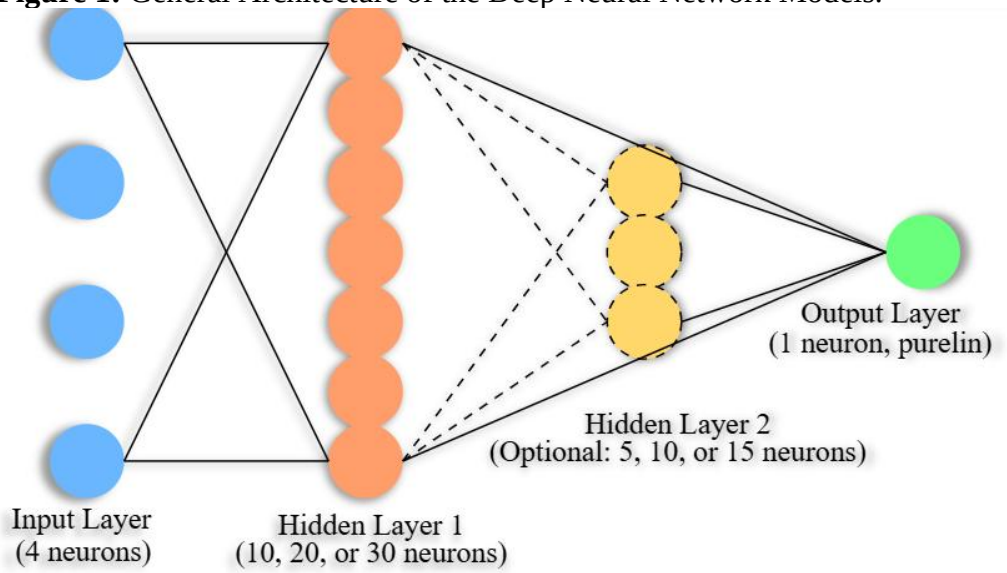
To capture medium-term trends and smooth out short-term volatility, moving average features (ma7 and ma14) were computed. These represent the 7-day and 14-day moving averages of the normalized exchange rates. Moving averages highlight underlying market tendencies and reduce the noise introduced by daily fluctuations, providing the model with a clearer view of broader directional shifts [21].

Together, these feature engineering techniques—lagged values and moving averages—enable the model to better discern both immediate market signals and more persistent trends, enhancing its ability to produce accurate and reliable exchange rate forecasts.

6. Model Architectures and Training Procedures

The study experimented with a range of neural network configurations to evaluate how varying depths, layer sizes, and activation functions influenced predictive performance. These models are structured to process engineered features through a hierarchical learning process.

Figure 1: General Architecture of the Deep Neural Network Models.



- **Input Layer:** Comprises 4 neurons corresponding to the engineered features (lag7, lag14, ma7, and ma14).
- **Hidden Layers:** The configurations scale from a single hidden layer (10, 20, or 30 neurons) to dual-layer architectures where a second optional layer (5, 10, or 15 neurons) is introduced to capture deeper nonlinear dependencies.
- **Output Layer:** A single neuron using a linear activation function (*purelin*) to produce the final continuous-valued exchange rate forecast.

Initial attempts included relatively simple feedforward networks (Multilayer Perceptrons, MLPs) with a single hidden layer, gradually scaling up to deeper architectures featuring multiple hidden layers and an increasing number of neurons. The choice of architectures followed two main principles: model simplicity and scalability. Simpler MLPs served as baseline models, while deeper and more complex architectures allowed the model to learn richer representations, potentially capturing nonlinear dependencies that shallow networks might miss.

7. Model Selection Criteria:

The choice of architectures followed two main principles: model simplicity and scalability. Simpler MLPs served as baseline models, offering insights into whether increased complexity was necessary to capture the underlying patterns in the data. Deeper and more complex architectures allowed the model to learn richer representations, potentially capturing nonlinear dependencies that shallow networks might miss. Comparing these models provided an understanding of the trade-offs between model complexity, training time, and performance.

8. Training Procedures:

All models were trained using the Mean Squared Error (MSE) as the loss function, given its widespread use in regression tasks and its suitability for measuring prediction accuracy in continuous-valued outputs. The Adam optimizer, known for its robustness and efficiency in handling noisy gradients, was employed to update model weights iteratively. To prevent overfitting—an especially critical concern in deep architectures—regularization techniques were implemented. Bayesian

regularization, for instance, integrates a complexity penalty directly into the training process, leading to smoother solutions and reduced overfitting. Additionally, dropout layers were considered in some architectures to randomly deactivate a portion of neurons during training, promoting better generalization by preventing excessive reliance on specific connections.

9. Hyperparameter Tuning and Validation

Optimizing hyperparameters—such as learning rates, number of neurons in each layer, or dropout probabilities—is essential to achieving robust performance. A systematic approach was adopted to identify the ideal hyperparameter settings. Techniques like grid search, which exhaustively tests predefined parameter sets, and Bayesian optimization, which guides the search process more intelligently by modeling the performance surface, were considered. This combination allowed for a balance between thorough exploration and computational efficiency.

The validation set played a pivotal role in guiding model selection and hyperparameter tuning. By monitoring validation performance rather than just training accuracy, the risk of overfitting was minimized. Models that showed promising performance on the validation set were retained for final testing, while those that overfit—evidenced by improved training metrics but deteriorating validation metrics—were discarded or reconfigured. This iterative tuning process ensured that the final chosen architecture and hyperparameters reflected genuine predictive ability rather than artifacts of over-training on the historical dataset.

IV. Results

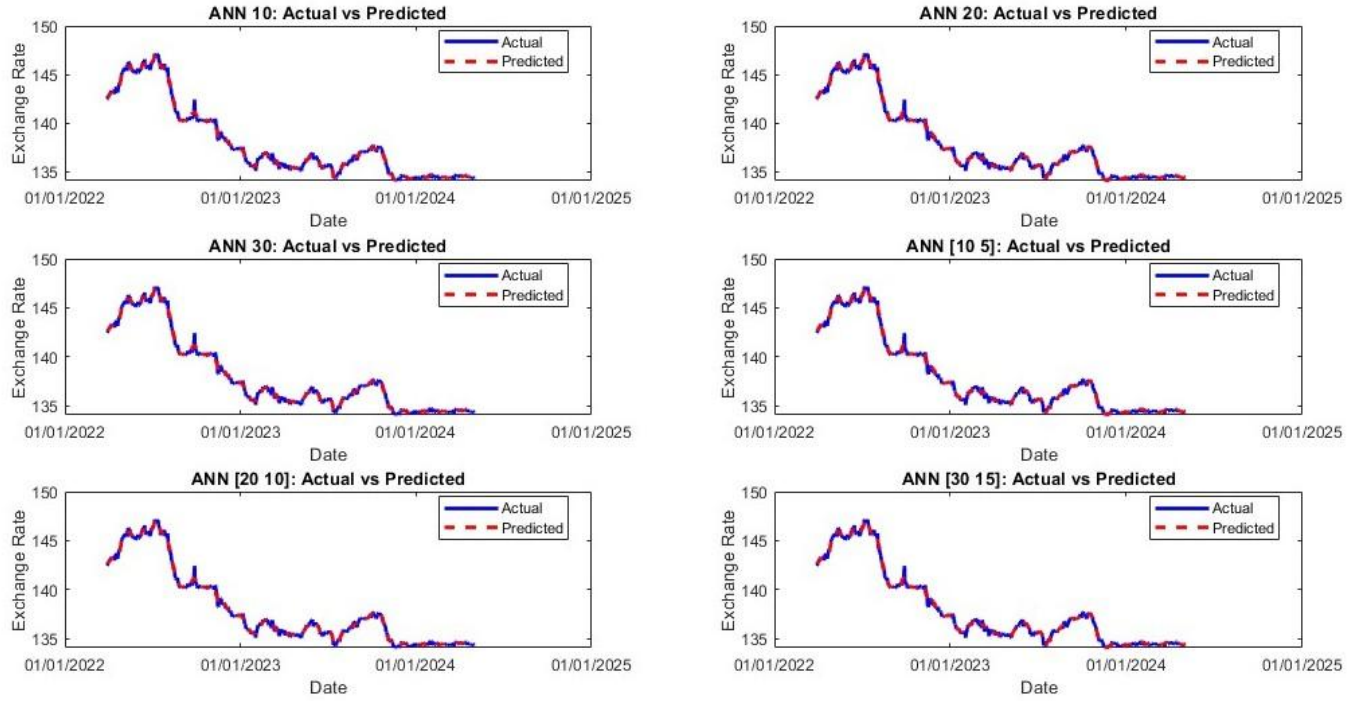
1. Performance Comparison of Architectures

Table 1 presents the performance of various ANN architectures evaluated in this study. The architectures ranged from single-hidden-layer MLPs with 10 neurons to deeper models, such as those with multiple hidden layers and increasing neuron counts (e.g., [10, 5], [20, 10], [30, 15]). Performance metrics included the Mean Squared Error (MSE) and Root Mean Squared Error (RMSE) on the validation set.

Table 1: Performance Metrics Across Evaluated ANN Architectures

Model Architecture	MSE $\times 10^{-6}$	RMSE
ANN (10)	3.4830	0.0018663
ANN (20)	3.4517	0.0018579
ANN (30)	3.4658	0.0018617
ANN ([10, 5])	3.4452	0.0018561
ANN ([20, 10])	3.4658	0.0018617
ANN ([30, 15])	3.4588	0.0018598

Figure 2: Visual Comparison of Actual vs. Predicted USD/DZD Exchange Rates.



As illustrated in Table 1 and Figure 1, all architectures achieved comparably low error values, reflecting strong predictive capabilities. Notably, the two-layer model with architecture [10, 5] yielded a marginally better performance than other configurations, achieving the lowest MSE and RMSE among the evaluated models. The visual results in Figure 1 confirm this statistical finding, as the red dashed line (predicted) almost perfectly overlays the blue solid line (actual) across the entire test period for all configurations.

This improvement, though slight, suggests that adding a modest second layer can enhance the network's ability to capture nuanced patterns without substantially increasing complexity. While minor variations in RMSE were statistically significant, the differences were practically small, indicating that all tested networks performed at a high level of accuracy. Confidence intervals constructed around the RMSE values further confirmed the narrow range of variation across models, underscoring the consistency in predictive performance.

2. Model Convergence and Training Efficiency

An analysis of training curves revealed that all models converged stably, with no indications of severe overfitting or divergence. Simpler architectures, such as the single-layer ANN (10), generally converged slightly faster, reaching stable performance levels in fewer epochs. In contrast, deeper networks with additional hidden layers or more neurons, like ANN ([30, 15]), required slightly longer training times to achieve similar performance metrics.

Despite the added complexity, deeper architectures did not present significant stability issues. The inclusion of regularization techniques, such as Bayesian regularization and, in some cases, dropout, ensured that networks maintained steady learning progress. The computational cost of training deeper models was marginally higher, reflected in more training epochs needed to reach minimal validation loss and the additional time for processing a larger number of parameters.

However, the incremental computational overhead was relatively small given modern hardware capabilities. The differences in training duration amounted to a matter of seconds or minutes rather than hours, suggesting that even more complex models can be feasibly trained for exchange rate forecasting tasks. In practice, the slight increase in computational effort may be justified by the marginal performance gains observed in more nuanced architectures, particularly for applications where even a small improvement in predictive accuracy can yield tangible benefits in financial decision-making and risk management.

V. Discussion

1. Insights into Architectural Choices

The results indicate that even relatively simple deep neural network (DNN) architectures can deliver high accuracy in forecasting the USD/Algerian Dinar exchange rate. Models like ANN ([10, 5]) consistently outperformed or matched more complex configurations, suggesting that increasing network depth and neuron count beyond a certain point may yield diminishing returns. This finding aligns with the notion that while deeper architectures are capable of capturing more nuanced patterns, such complexity is not always necessary—especially when effective feature engineering (e.g., lagged values and moving averages) and proper normalization techniques are employed. The trade-off, therefore, is between complexity and accuracy gains. Adding layers and parameters can marginally improve performance, but the computational overhead and prolonged training time may not justify the incremental benefits. Practitioners may find that simpler, well-tuned architectures strike an optimal balance, maintaining efficiency without sacrificing predictive fidelity.

2. Real-World Implications

The robust performance of the tested architectures carries significant value for financial stakeholders. Accurate exchange rate forecasts can enhance currency hedging strategies, enabling firms and investors to mitigate exposure to volatile exchange movements. In scenarios involving international investments, better predictions of future currency fluctuations allow for more informed capital allocation decisions. Additionally, automated trading systems and risk assessment frameworks can leverage these improved forecasting models, incorporating them into algorithmic strategies that react swiftly to market changes. By employing models that are not only accurate but also computationally manageable, real-time and near-real-time trading and risk management applications become more feasible, potentially improving returns and reducing unexpected losses in global financial markets.

3. Limitations and Potential Improvements

While this study demonstrates the efficacy of DNN-based forecasting for the USD/DZD exchange rate, several limitations should be noted. First, the focus on a single currency pair may limit the generalizability of the findings. Different currency pairs, especially those with distinct market characteristics or lower liquidity, may respond differently to the same architectures and features. Similarly, the time frame under consideration, although extensive, may not capture all possible market conditions, including rare crisis events or regime shifts that could challenge the model's stability.

Future research could address these limitations by broadening the scope of currency pairs and including macroeconomic indicators—such as interest rates, inflation rates, and geopolitical factors—to enhance model inputs and potentially boost predictive accuracy. Moreover, exploring more advanced architectures like Long Short-Term Memory (LSTM) networks, Transformers, or Graph Neural Networks (GNNs) could uncover additional improvements. Transfer learning, wherein models trained on more frequently traded currency pairs are adapted for less liquid or emerging market pairs, represents another promising avenue. By pursuing these enhancements, researchers and practitioners can continue refining exchange rate forecasting models, driving further improvements in both predictive performance and real-world applicability.

VI. Conclusion

This study has demonstrated that deep neural network (DNN) architectures, coupled with effective preprocessing and feature engineering, can significantly enhance the accuracy of exchange rate forecasts. By incorporating lagged values and moving averages, the models successfully captured both short-term dependencies and medium-term trends in the USD/Algerian Dinar exchange rate. While deeper and more complex architectures offered marginal improvements in certain cases, simpler multi-layer perceptrons—particularly those with a modest increase in depth—yielded comparable or even superior predictive performance. The results underscore that optimizing network complexity, rather than simply increasing it, is critical to achieving high accuracy without incurring unnecessary computational overhead. In essence, a thoughtfully selected and tuned DNN

configuration can substantially improve forecasting accuracy, offering valuable insights for financial decision-makers and risk managers.

Building on the promising outcomes observed here, future research could explore more advanced and domain-specific architectures, such as recurrent neural networks (e.g., LSTMs) or transformer-based models, to determine whether their ability to capture complex temporal dynamics further enhances predictive capability. Examining additional currency pairs, especially those subject to different market conditions or lower liquidity, would help validate the model's generalizability. Moreover, integrating macroeconomic indicators, geopolitical factors, and other relevant domain knowledge could improve the robustness and interpretability of the forecasts. Finally, leveraging larger and more diverse datasets, as well as employing techniques like transfer learning, may not only refine predictive accuracy but also broaden the scope of real-world financial applications, supporting more informed economic decision-making at both corporate and policy levels.

References:

Abreu, R. J., Souza, R. M., & Oliveira, J. G. (2019). Applying singular spectrum analysis and ARIMA-GARCH for forecasting EUR/USD exchange rate. RAM. Revista de Administração Mackenzie, 20, eRAMF190146.

Antonakakis, N., Chatziantoniou, I., & Gabauer, D. (2019). Cryptocurrency market contagion: Market uncertainty, market complexity, and dynamic portfolios. Journal of International Financial Markets, Institutions and Money, 61, 37–51.

Avdjiev, S., Bruno, V., Koch, C., & Shin, H. S. (2019). The dollar exchange rate as a global risk factor: evidence from investment. IMF Economic Review, 67, 151–173.

Carney, M. (2019, August). The growing challenges for monetary policy in the current international monetary and financial system. In Remarks at the Jackson Hole symposium (Vol. 23, pp. 377–411).

Charef, F. (2024). Exchange rate forecasting: Nonlinear garch-nn modeling approach. Annals of Data Science, 11(3), 947–957.

Chen, L., Du, Z., & Hu, Z. (2020). Impact of economic policy uncertainty on exchange rate volatility of China. Finance Research Letters, 32, 101266.

Eun, C. S., Resnick, B. G., & Chuluun, T. (2021). International financial management. McGraw-Hill.

Girgin, Y. (2023). Modeling Exchange Rate Volatility Using ARMA-GARCH Aproach with Non-Gaussian Distributions (Master's thesis, Middle East Technical University).

Kukacka, J., & Kristoufek, L. (2020). Do 'complex'financial models really lead to complex dynamics? Agent-based models and multifractality. Journal of Economic Dynamics and Control, 113, 103855.

Mohammadi, W. A. Z. I. R. (2019). Currency exchange rate forecasting using machine learning techniques. Graduate School of Applied Sciences. Near East University.

Nazareth, N., & Reddy, Y. V. R. (2023). Financial applications of machine learning: A literature review. Expert Systems with Applications, 219, 119640.

Nguyen, P. D., Thao, N. N., Kim Chi, D. T., Nguyen, H. C., Mach, B. N., & Nguyen, T. Q. (2024). Deep learning-based predictive models for forex market trends: Practical implementation and performance evaluation. Science Progress, 107(3), 00368504241275370.

Ozbayoglu, A. M., Gudelek, M. U., & Sezer, O. B. (2020). Deep learning for financial applications: A survey. Applied soft computing, 93, 106384.

Pandya, J. B. (2024). DEEP LEARNING APPROACH FOR STOCK MARKET TREND PREDICTION AND PATTERN FINDING (Doctoral dissertation, GUJARAT TECHNOLOGICAL UNIVERSITY AHMEDABAD).