FORECASTING FOREIGN EXCHANGE RATE WITH MACHINE LEARNING TECHNIQUES

A THESIS SUBMITTED TO
THE GRADUATE SCHOOL OF NATURAL AND APPLIED SCIENCES
OF
MIDDLE EAST TECHNICAL UNIVERSITY

BY

ROHAT KÜÇÜK

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR
THE DEGREE OF MASTER OF SCIENCE
IN
STATISTICS

SEPTEMBER 2023

Approval of the thesis:

FORECASTING FOREIGN EXCHANGE RATE WITH MACHINE LEARNING TECHNIQUES

submitted by **ROHAT KÜÇÜK** in partial fulfillment of the requirements for the degree of **Master of Science in Statistics Department, Middle East Technical University** by,

Prof. Dr. Halil KALIPÇILAR
Dean, Graduate School of **Natural and Applied Sciences** ________________

Prof. Dr. Özlem İLK DAĞ
Head of Department, **Statistics** ________________

Assoc. Prof. Dr. Fulya GÖKALP YAVUZ
Supervisor, **Statistics, METU** ________________

Examining Committee Members:

Prof. Dr. Ceylan TALU YOZGATLIGİL
Statistics, METU ________________

Assoc. Prof. Dr. Fulya GÖKALP YAVUZ
Statistics, METU ________________

Prof. Dr. Seher Nur SÜLKÜ
Econometrics, Ankara Haci Bayram Veli University ________________

Date:07.09.2023

I hereby declare that all information in this document has been obtained and presented in accordance with academic rules and ethical conduct. I also declare that, as required by these rules and conduct, I have fully cited and referenced all material and results that are not original to this work.

Name, Surname: ROHAT KÜÇÜK

Signature :

ABSTRACT

FORECASTING FOREIGN EXCHANGE RATE WITH MACHINE LEARNING TECHNIQUES

KÜÇÜK, ROHAT
M.S., Department of Statistics
Supervisor: Assoc. Prof. Dr. Fulya GÖKALP YAVUZ

September 2023, 66 pages

Since the collapse of the Bretton Woods system, international macroeconomics has grappled with the challenges of explaining exchange rate behavior under flexible regimes. The exchange rate disconnect puzzle underscores the elusive relationship between exchange rates and their underlying macroeconomic determinants. Drawing on this, groundbraking work by Meese and Rogoff (1983), have posited the potential superiority of the simple random walk model over complex economic theories in forecasting exchange rates. To further probe into this, our study employs a diverse set of models—ARIMA, Extreme Gradient Boosting, Support Vector Regression, Long Short Term Memory, Convolutional Neural Network, and notably, an ARIMA-LSTM Hybrid—with the objective of capturing both the linear and non-linear dynamics inherent in the data. We examine the forecasting landscape for the Russian Ruble, Euro, and Turkish Lira over time horizons of 3, 7, 14, and 30 days. Utilizing models grounded in classical economic theories such as the Uncovered Interest Rate Parity, Purchasing Power Parity, and the Monetary Model and contrasted their performance with models that did not incorporate these theoretical constructs. The effectiveness of these models was assessed based on Root Mean Square Error, Mean Absolute Percentage Error, and Theil's U metrics. Our findings underscore the nuanced role of economic variables in forecasting, the absence of a universally optimal model, and the consistent phenomenon of benchmark Random Walk models holding their own, particularly with the Euro. Conclusively, the realm of forex forecasting is underscored as multifaceted, emphasizing the need of a diverse, evolving toolkit in navigating the ever-shifting global financial landscape.

Keywords: Exchange Rates, Forecasting, Machine Learning, Random Walk, Time Series

ÖZ

MAKİNE ÖĞRENMESİ TEKNİKLERİ İLE KUR TAHMİNİ

KÜÇÜK, ROHAT
Yüksek Lisans, İstatistik Bölümü
Tez Yöneticisi: Doç. Dr. Fulya GÖKALP YAVUZ

Eylül 2023 , 66 sayfa

Bretton Woods sisteminin çöküşü ile birlikte uluslararası iktisat, esnek döviz kuru rejimi altında kur davranışlarını açıklamakla ilgili zorluklarla karşı karşıya kalmıştır. Döviz kuru bağlantısızlık bilmecesi, döviz kurları ile altındaki makroekonomik belirleyiciler arasındaki zayıf ilişkiyi vurgulamaktadır. Bununla birlikte, Meese ve Rogoff(1983) tarafından yapılan önemli çalışmalar, döviz kurlarını tahmin etmede karmaşık ekonomik teorilere göre basit Rassal Yürüyüş modelinin kur tahmininde güçlü bir model olduğunu göstermiştir. Konuyu derinlemesine incelemek için bu çalışma, Otoregresif Tamamlanmış Hareketli Ortalama, Aşırı Gradyan Arttırılmış Ağaçlar, Destek Vektör Regresyonu, Uzun Kısa Süreli Bellek, Evrişimli Sinir Ağları modelleri ile beraber Otoregresif Tamamlanmış Hareketli Ortalama-Uzun Kısa Süreli Bellek Hibriti doğrusal ve doğrusal olmayan zaman serisi dinamiğini yakalaması için kullanılmıştır. Araştırmaya konu olan Rus Rublesi, Euro ve Türk Lirasının 3, 7, 14 ve 30 günlük tahminleri incelenmiştir. Açık Faiz Oranı Paritesi, Satın Alma Gücü Paritesi ve Para Modeli gibi klasik ekonomik teorilere dayalı modelleri kullanarak bu teorik yapıları içermeyen modellerle performansları karşılaştırılmıştır. Bu modellerin performansı Karesel Ortalama Hata, Ortalama Mutlak Yüzdesel Hata ve

Theil-U metrikleri ile değerlendirilmiştir. Araştırma çıktıları, ekonomik değişkenlerin kur tahmininde kesin bir etkisinin olmadığını, evrensel olarak optimal bir modelin olmadığını ve özellikle Avro ile ilgili, Rassal Yürüyüşün tutarlı bir model olduğunu vurgulamaktadır. Sonuç olarak, sürekli değişen global finans ortamında doğru bir kur tahmini yapabilmek için çeşitli ve sürekli gelişen bir araç setine ihtiyaç duyulduğu kanıtlanmıştır.

Anahtar Kelimeler: Kur Oranları, Makine Öğrenmesi, Tahminleme, Rassal Yürüyüş, Zaman Serisi

To my family

ACKNOWLEDGMENTS

First and foremost, I wish to convey my deepest gratitude to my thesis advisor, Assoc. Prof. Dr. Fulya Gökalp Yavuz. Her patience, consistent guidance, and dedication have been the linchpin of this research. The wisdom she bestowed continually steered me in the right direction. Without her relentless support and encouragement, this thesis would not have reached its completion.

I would like to thank the examining committee members; Prof. Dr. Ceylan Talu Yozgatlıgil and Prof. Dr. Seher Nur Sülkü for their invaluable feedback, and critical assessments.

I must also express my heartfelt appreciation to Burcu Sarı. Beyond academic support, she provided the motivation I needed during the most challenging times, ensuring I stayed on track. My profound thanks go to Anna Gordeeva. Her unwavering faith in my potential has played a pivotal role in the realization of this research. More than just a steadfast pillar, her insights and encouragement deeply influenced this work, keeping me centered throughout the process.

Lastly, I dedicate this work to my cherished family: my mother, Gülseven Küçük; my father, Cafer Küçük; and my brother, Can Emre Küçük. Their unwavering love, encouragement, and faith in me have been the bedrock upon which all my endeavors stand.

TABLE OF CONTENTS

LIST OF TABLES

TABLES

LIST OF FIGURES

FIGURES

LIST OF ABBREVIATIONS

ACF	Autocorrelation Function
ADF	Augmented Dickey-Fuller
AIC	Akaike Information Criteriea
AR	Autoregressive
ARIMA	Autoregressive Integrated Moving Average
BIC	Bayesian Information Criterion
BOP	Balance of Payments
BEER	Behavioral Equilibrium Exchange Rate
CBRT	Central Bank of the Republic of Türkiye
CNN	Convolutional Neural Network
Coef. Var.	Coefficient of Variation
CPI	Consumer Price Index
DCN	Deep Convolotuional networks
ECB	European Central Bank
EUR	Euro
FRED	Federal Reserve Economic Data
FX	Foreign Exchange
GDP	Gross Domestic Product
LSTM	Long Short-Term Memory
MA	Moving Average
MAPE	Mean Absolute Percentage Error
MB	Macroeconomic Balance
MSE	Mean Squared Error
MM	Monetary Model

OECD	Organisation for Economic Co-operation and Development
PACF	Partial Autocorrelation Function
PI	Production Index
PPI	Producer Price Index
PPP	Purchasing Power Parity
PSOSVR	Particle Swarm Optimization Support Vector Regression
FSPSOSVR	Feature Selection Particle Swarm Optimization Support Vector Regression
RBF	Radial Basis Function
RNN	Recurrent Neural Network
RMSE	Root Mean Square Error
RUB	Russian Rubble
ReLU	Rectified Linear Unit
SVR	Support Vector Regression
TANH	Hyperbolic Tangent
TRY	Turkish Lira
UIRP	Uncovered Interest Rate Parity
VAR	Vector Autoregression
XGB	Extreme Gradient Boosting

CHAPTER 1

INTRODUCTION

Forecasting, in simple terms, is all about predicting what is going to happen next. This idea becomes even more interesting and important when we start predicting things related to money, like exchange rates. Foreign exchange (FX) markets play a crucial role in the global economy, with trillions of dollars being traded daily[3]. Accurate forecasting of exchange rates is essential for various market participants, such as central banks, investors, and multinational corporations[4]. Central banks across the globe and other financial institutions monitor the fluctuations in exchange rate closely. Even though the exchange rate is not a direct policy target, changes in rates significantly affect medium-term inflation expectations via their impact on import prices and broader economic forces. Exchange rate fluctuations can create uncertainty that hinders global commerce and investment, highlighting the importance of comprehending how prices are formed and accurately predicting exchange rate movements for both investors and policymakers. According the Menkoff and Taylor (2007), accurate exchange rate predictions are pivotal aspect of any successful trading strategy[5]. Moreover, central banks recognise the persistent changes in nominal exchange rates may lead to significant consequences for overall economic outlook. However, precise prediction of exchange rates is a challenging job due to complex dynamics of its nature as exchange rates are influenced by a wide range of factors, such as economic indicators, political events, market trends, and investor sentiment. This enigma widely recognized in the realm of global finance and has a connection with the prominent exchange rate disconnect puzzle. The forecasting conundrum owes its origins to the groundbreaking work of Meese and Rogoff in 1983, which concluded that traditional exchange rate models frequently fall short of random walk approaches when predicting future currency values beyond their training data[6]. The puzzle refers to the surprising lack of correlation between exchange rates and almost any macroeconomic aggregates, despite extensive theoretical frameworks anticipating a strong association.

Considering the challenging nature of the problem, numerous econometric and machine learning methodologies have been proposed to confront this task. These includes traditional time series models such as Autoregressive (AR), Moving Average (MA), and Autoregressive Integrated Moving Average (ARIMA) models. The ARIMA model which combines the strengths of AR and MA models has gained broad recognition in financial forecasting because of its capability to handle non-stationary time series data[7]. In recent years the advent of machine learning has led to the development of techniques for financial forecasting including Support Vector Regression (SVR), Extreme Gradient Boosting (XGBoost), Long Short-Term Memory networks (LSTM), and Convolutional Neural Networks (CNN).

SVR, an application of SVM designed for regression tasks, has demonstrated its proficiency in handling with high-dimensional spaces and avoiding overfitting[8]. XGBoost, is an sophisticated implementation of gradient boosting algorithm and has proven to perform well in numerous machine learning (ML) competitions[9]. Deep learning models such as LSTM and CNN, have also been employed. LSTM is a type of recurrent neural network (RNN), are particularly well-suited for sequence prediction challenges owing to their ability to retain past information[10]. CNNs, traditionally used in image processing, have been applied to time series forecasting and have yielded encouraging results thanks to their capacity to identify local patterns and their robustness to noise and shift[11]. Hybrid models that combine traditional econometric model with ML algorithms have been suggested to capitalize on the advantages of both approaches. One such model is the ARIMA-LSTM, which marries ARIMA's talent for capturing linear relationships with LSTM's proficiency in modelling complex, nonlinear patterns[12].

Equilibrium exchange rate models has significant contribution to understanding and forecasting exchange rate dynamics. Models such as Purchasing Power Parity (PPP), the Monetary Model (MM), Uncovered Interest Rate Parity (UIRP), Behavioral Equilibrium Exchange Rate (BEER), and Macroeconomic Balance (MB) are widely used approaches. PPP posits that the exchange rate between two nations ought to be aligned so that identical sets of goods and services cost the same in both countries, thereby ensuring purchasing power parity[13]. This foundational principle serves as a benchmark for evaluating the values of currencies and understanding how they impact the prices of internationally traded goods. According to MM, the exchange rate is influenced by the relative supply and demand across two countries, as well as anticipated variations in inflation and real interest rates[14]. This framework takes into account the dynamic interplay between monetary forces and the resulting effects on currency valuations. The UIRP theory suggest that the expected return on an investment in an domestic market should be equivalent to the expected return on an investment on an foreign market, taking into account any changes in exchange rates[15]. Therefore investors should not be able to earn risk-free profit by borrowing in one currency and investing in another currency. BEER model is a dynamic framework that considers both short-term and long term economic factors, as well as concept of the equilibrium errors. This approach aims to capture the impact of temporary shocks as well as more structural changes on exchange rates[16]. By incorporating, the BEER model offers a comprehensive perspective on exchange rate determination, enabling policy makers and analysts to better understand how economic forces shape currency values over time. The MB approach, focuses on identifying equilibrium exchange rate that would emerge when economies operate at their full potential and maintain sustainable current account positions. This methodology emphasizes the interplay between macroeconomic variables such as GDP, inflation and trade balances, and how they influence the value of currencies. By assuming that economies tend towards a state of balance the MB model provides insights into the underlying drivers of exchange rate movements and helps users assess the alignment of currencies with their fundamental values[17].

In this study, we consider three major currencies against US dollar - the Russian ruble (RUB), Turkish Lira (TRY) and Euro (EUR) - testing the performances of forecasting model in the multiple exchange rate markets. Daily data from 2004 to 2023 used to forecasts horizons of 3,7,14 and 30. Data is gathered from Central Bank of Republic of the Türkiye (CBRT), Federal Reserve Economic Data (FRED), Yahoo Finance, World Bank data, OECD. Models are evaluated at two separated groups. First group consists of models that used just time variables to predict exchange rate and in second group same model used but this time economic variables along with the time variables are used to predict. We considered eight models including Random Walk, Random Walk with drift, SARIMAX, XGBoost, Support Vector Regression, Long Short-Term Memory, Convolutional Neural Network and hybrid SARIMAX-LSTM. From the economic theory PPP, UIRP and the MM are structured to apply in the models. Accuracy measurements of the forecasting models are Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE) and Theils U. Analysis are conducted using R Studio with version 2023.06.1 and Jupyterlab with version 3.4.3.

The study unfolds across five chapters. The initial chapter sets the stage by introducing the main objectives and the scope of the research. The subsequent chapter, Chapter 2, delves into a comprehensive literature review, highlighting prior research and findings that align with the study's subject. Moving forward, Chapter 3 elucidates the theoretical foundations of the models employed, accompanied by a introduction to the dataset. This chapter then seamlessly transitions into the application process of the models and end up with the presentation of the empirical analysis results. The study then draws to a close with final reflections and conclusions in the concluding chapter.

CHAPTER 2

LITERATURE REVIEW

At the forefront of this extensive literature is the pioneering work of Meese and Rogoff in 1983. In this section, we provide a brief overview of some key studies that have contributed to the development of exchange rate models. In the initial phases, forecasting methodologies predominantly revolved around univariate time-series models. Key models of this era included the Random Walk and the ARIMA, which employed straightforward statistical techniques. The mid-part of the 20th century witnessed a paradigm shift in exchange rate modeling. The Keynesian framework emerged as a notable example, establishing itself as a widely accepted model. This transformative period can be attributed to the seminal works of Fleming (1962)[18] and Mundell (1962)[19]. Their contributions were, in essence, extensions and enhancements of ideas originally posited by Meade in 1951[20]. The dissolution of the Bretton Woods system initiated an era marked by intense deliberation and re-evaluation of prevailing models. Within this context, the classic Mundell-Fleming framework saw substantial revisions. These modifications were aimed at incorporating asset market-based expectations, considering both flexible[21] and sticky price[22] contexts. Although there were significant developments in the theoretical understanding of exchange rate dynamics during this period, practical applications of these models frequently encountered difficulties in providing dependable empirical evidence, particularly when it came to out of sample forecasting. This issue was highlighted even further by Meese and Rogoff's (1983) research, which revealed the constraints of the commonly used exchange rate models in terms of their ability to accurately predict future trends. Ince (2014), used both PPP and Taylor rule fundamentals to forecast quarterly exchange rates for nine Organization for Economic Cooperation and Development (OECD) countries. The author found that the PPP model demonstrates greater predictive ability at longer time horizons, whereas the Taylor rule model shows higher accuracy at shorter time horizons[23]. This suggests that both models have their strengths and weaknesses depending on the specific context and time frame under consideration. In a research by Pfahler(2021), advanced ML techniques were employed to discern the predictive power of macroeconomic fundamentals in the realm of exchange rate forecasting. Specifically, the study utilized artificial neural networks, characterized by a multilayer perceptron framework, in tandem with gradient boosted decision trees via XGBoost. The examination was conducted on monthly data derived from OECD countries, with the underpinning macroeconomic fundamentals being rooted in the PPP, UIRP, and MM theoretical paradigms. Both classification and regression paradigms were assessed. The findings underscored the inherent capabilities of ML methodologies in predicting exchange rate fluctuations. However, the precise contribution of macroeconomic fundamentals, particularly when together with time variables, remains uncertain[24]. Neghab et al. (2023) conducts an in-depth analysis of the daily dynamics of the Canadian exchange rate, with a particular focus on elucidating the impact of crude oil prices and other key macroeconomic indicators, such as the Producer Price Index (PPI) and money supply. Employing tree-based models, this research demonstrates their superior performance over other models on an average performance scale[25]. According to a study conducted by the European Central Bank (ECB), several equilibrium exchange rate models, including the PPP model, the Behavioral Equilibrium Exchange Rate (BEER) model, and the Macro-economic Balance (MB) model, were examined for EUR and USD using quarterly data. The findings revealed that, on average, real exchange rates tend to gradually converge towards their equilibria when defined by either the PPP or BEER models. Moreover, the study highlighted that the predictive power of these models stems primarily from the mean-reverting properties of real exchange rates, rather than the connection between exchange rates and macroeconomic fundamentals. The study indicates that the best method for predicting currency fluctuations involves expecting a slow return to the real rates average value. Real and nominal exchange rates are identified as key fundamentals for successful forecasting. However, other economic factors, such as current account balances, may exert some influence on the short-term trajectory of exchange rate adjustments but have only limited long-term implications[26]. Among econometric models AR have been widely applied to predict various exchange rates and have demonstrated currency-specific effectiveness, producing good results for currencies like the Hungarian Forint and Korean Won but failing to perform well for the dollar against the Euro [27]. ARIMA is a strong statistical model that can compete ML models. For instance, Babu and Reddy (2015) utilized ANN and fuzzy neuron models to forecast Indian currency prices. Although they compared the results with ARIMA and ARIMA demonstrated superior performance[28]. The rise of the digital age saw the emergence of ML techniques as potential replacements for traditional statistical methodologies in exchange rate forecasting. These ML models, such as Neural Networks and SVM, signified a significant change in the landscape of exchange rate forecasting, enabling the handling of complex and non-linear relationships. Plakandaras et al. (2015) conducted a comparison of Neural Networks and SVM for forecasting multiple exchange rates, including those of the Euro, Japanese Yen, and Australlian Dollar. Their findings showed that hybrid models are competitive in exchange rate forecasting[29]. A related study employed a multi-variable predictive regression model, dubbed the "Kitchen Sink" regression, which integrates various macroeconomic indicators to forecast monthly exchange rates of nine major currencies, including the Australian dollar, Canadian dollar, Swiss franc, and British pound. This innovative approach leverages the elastic-net shrinkage estimator to efficiently extract predictive information from multiple economic fundamentals, thereby enhancing the variance-bias tradeoff in forecasting. The investigation found that the Kitchen Sink regression yielded reliable predictions, demonstrating the effectiveness of this integrated model in capturing complex relationships among exchange rates and underlying macroeconomic factors[30]. Yasir et al. (2019) examined the daily exchange rate prediction of the United Kingdom, Hong Kong, and Pakistan in terms of US dollars, incorporating sentiment analysis of major local and global events for each country. The sentiments derived from these events were utilized as input parameters in a predictive deep learning model, which was contrasted with linear and support vector regression (SVR) models. The findings revealed that the deep learning model outperformed the conventional approaches, demonstrating the potential of incorporating sentiment analysis into exchange rate prediction. Furthermore, the study discovered that events in the United States (US) had a significant impact on the currency exchange rates of Hong Kong, Pakistan, and the UK, particularly when social media sentiment was taken into account. These results underscore the vulnerability of these countries' currencies to major international events[31]. According to Shen et al. (2021), the Feature Selection Particle Swarm Optimization Support Vector Regression (FSPSOSVR) algorithm was utilized to predict exchange rates of seven major currencies, including the Euro, Japanese Yen, and Chinese Renminbi. The performance of the FSPSOSVR algorithm was compared to six alternative models, namely the Random Walk, Exponential Smoothing, Autoregressive Integrated Moving Average, Seasonal ARIMA , Support Vector Regression, and Partial Support Vector Regression (PSOSVR). The results indicated that the FSPSOSVR algorithm outperformed all other models in terms of prediction accuracy[32]. Deep Learning methods got popular with the advance in the field. Galeschuk and Mukherjee (2017a,b) explored the use of Deep Convolotuional networks (DCN), SVR, and ANN to predict exchange rates for both emerging and advanced economies. Their results indicated that the DCN model yielded substiantially higher predictive accuracy than traditional econometric models, such as ARIMA and exponential smoothing[3]. Another comparison between econometric and ML models is put forward by Parot et al (2019) with a hybrid model that combines the strenghts of Artificial Neural Networks and Vector autoregression (VAR) for predicting Euro exchange rate returns. Results showed models with ML algorithms outperformed traditional models[33]. In a recent study, the daily returns of five prominent exchange rate pairs were scrutinized, namely: the Japanese yen, Pound sterling, Euro, Swiss franc, and Canadian dollar, all benchmarked against the US Dollar. Furthermore, the research ventured into the realm of trading strategies, offering a comparative analysis between the buy-and-sell strategy, derived from the most profitable model, and the traditional buy-and-hold approach. Notably, the LSTM model showcased superior predictive capabilities, consistently outperforming its counterparts and emerging as the most profitable model in the set[34]. While ML models offer promising solution for exchange rate forecasting, they are not immune to criticism. Transparency and interpretability problem of black-box models always an issue in the area. Moreover, these models tendency to overfit data can result in sub-optimal performance when applied to new, unseen data, sparkling further debate within the academic community[35].

CHAPTER 3

METHODOLOGY

Accurately predicting future trends is crucial across various fields, and exchange rate is no exception. Statistical and ML methods are deployed for forecasting mission. In this section of the study, models that used for the forecasting exchange rates are introduced. Moreover, model evaluation and comparison processes are explained.

3.1 Random Walk

A random walk is a type of time series model where each successive observation is directly related to the previous one, but with a random element that can cause it to either increase or decrease by a certain amount. In other words, the current observation is based on the previous one, but with an added layer of uncertainty that allows for unexpected changes. It can be defined as given below

$$\mathbf{y}_t = \mathbf{y}_{t-1} + \mathbf{w}_t, \tag{31}$$

where $\mathbf{y}_t$ is the value of series at time t and $\mathbf{w}_t$ is the error term , which is assumed to be white noise with zero mean and constant variance. When forecasting, the random walk forecaster is a naive. It can be shown as

$$\hat{\mathbf{y}}_{t+h} = \mathbf{y}_t, \tag{32}$$

where h is the forecast horizon.

3.2 Weighted Bootstrap Random Walk

The Random Walk model with drift includes a drift parameter that allows for the possibility of a systematic bias or trend in the data. In other words, the model assumes that the random movements of the process are not completely independent, but rather have a slight tendency to move in one direction over time. Instead we propose The Weighted Bootstrap Random Walk (WBRW) model to leverage past observations by attributing higher importance to more recent data points when forecasting future values. This method acknowledges that the time series' recent behavior might have a stronger influence on its immediate future than older observations, due to possible non-stationarities or changing dynamics. The methodology consists of the following steps, assign weights to past observations, with more recent data points receiving higher weights. Using these weights, draw k samples with replacement from the past observations, simulating a form of importance sampling. Compute the forecast for $t+1$ as the mean if these k sampled values.

$$\mathbf{y}_{t+1} = \frac{1}{k}\sum_{j=1}^{k} y_{sample_j}, \tag{33}$$

where y_{sample_j} represents the j-th sampled observation. The weights can be determined using various strategies, but in this methodology, they are set to emphasize recent observations with linear approach.

This model is particularly useful in situations where past behavior is indicative of future performance, but with a pronounced emphasis on recent events. By considering the weighted behavior of the time series, the WBRW model can provide more nuanced and potentially more accurate forecasts than traditional random walk models.

3.3 ARIMA

ARIMA model, acclaimed for its straightforwardness and user-friendliness, enjoys popularity in exchange rate forecasting. Represented as a linear composition of prior observations and stochastic disturbances, it is designed as the use the capabilities of autoregressive and moving average components to project future exchange rate. In the ARIMA(p,d,q) model, p stands for the number of autoregressive terms, d denotes the number of differences reuqired to stabilize the series, and q represents the number of moving average terms. This models offers flexibility in selecting either purely AR terms, MA terms or a combination of both, known as ARMA modelds, allowing it to accommodate diverse data patterns. ARIMA models rely heavily on the assumption that the input time series is stationary, meaning its mean and variance remain constant throughout the observed period. By leveraging this property, the model can extract useful insights from the data and accurately forecast future values. In other words, the stationarity of the time series allows the ARIMA model to identify patterns and trends in the data and generate predictions based on those patterns. The ARIMA model defined as

$$\mathbf{y}_t = \alpha + \sum_{i=1}^{p} \beta_i \mathbf{y}_{t-i} + \sum_{j=1}^{q} \gamma_j \epsilon_{t-j} + \epsilon_t, \tag{34}$$

here $\mathbf{y}_t$ is the time series α is the intercept term β_i are the coefficients for the AR terms, γ_j are the coefficients for the MA terms, ϵ_t is a sequence of uncorrelated random variables with zero mean and constant variance and ϵ_{t-j} are the lagged error terms. The ARIMA modelling process consists of three stages, as suggested by Box and Jenkins(1976), identifying the appropriate model estimating its parameters, and verifying its accuracy through diagnostic checks. First, the model identifications stage involves examining the autocorrelation and partial autocorrelation characteristics of the time series data to determine the most suitable ARIMA model. Next, various models are estimated, and the best one chosen using a selection criteria such as the Akaike Information Criteriea (AIC).

3.4 Support Vector Regression

Support Vector Machines were first proposed by Vapnik et al.(1999) as a means to address classification issues[36]. Later, Drucker modified the SVM algorithm to create SVR, which is tailored for regression tasks[37]. The primary goal of SVR is to discover a nonlinear mapping that connects input training data to a higher-dimensional feature space. By doing so, it enables a more versatile and accurate modeling of intricate, nonlinear relationships, such as those encountered in exchange rate movements. Estimation function is defined

$$f(x) = \mathbf{w}\boldsymbol{\phi}(x) + b, \tag{35}$$

where x is the entire input space, $\mathbf{w}$ is the weight vector for the input vector $\boldsymbol{\phi}$ is the kernel function and b is the bias term.

The Objective of SVR is to minimize the following cost function Q

$$Q = \frac{1}{2}|w|^2 + C\sum_{i=1}^{N} L_\epsilon(x_i, y_i, f), \tag{36}$$

$$\text{subject to constraints} \begin{cases} y_i - wx_i - b \leq \epsilon + \xi_i, \\ wx_i + b - y_i \leq \epsilon + \xi_i^*, \\ \xi_i, \xi_i^* \geq 0. \end{cases} \tag{37}$$

The cost function consists of three components: The values ξ_i, ξ_i^* for ensures compliance with the constraint mentioned in the first and second expression, a penalty term C that penalizes errors exceeding ϵ using an ϵ-insensitive loss function L_ϵ for each training example, and constraints defined by the error term. This approach helps to improve the generalization ability of the learned model. The cost function optimization yields the SVR decision function, which can be expressed as

$$f(x) = \sum_{i=1}^{N}(\alpha_i - \alpha_i^*)K(\mathbf{x}_i, \mathbf{x}) + b, \tag{38}$$

where α_i and α_i* are the Lagrange multipliers, and $K(x_i, x_j)$ is a kernel function that defines a nonlinear decision boundary in the SVR input space.[38]

The Gaussian Radial Basis Function (RBF) is a popular kernel that enables nonlinear mapping between input and high-dimensional spaces, making it suitable for tackling complex problems. The Gaussian RBF kernel is defined as:

$$K(\mathbf{x}_i, \mathbf{x}) = e^{-s\|\mathbf{x}_i - \mathbf{x}\|^2}, \tag{39}$$

where s is the scaling factor for the kernel.

3.5 Extreme Gradient Boosting

XGB is a powerful predictive algorithm renowned for its rapidity and precision. It employs the technique of 'boosting,' where numerous simple learners are consolidated to construct a robust one. The XGB algorithm is a highly efficient implementation of gradient boosting that balances system performance and ML capabilities[9]. It has become a popular choice in various ML competitions and real-world scenarios due to its robustness and quick execution times. Boosting is a versatile technique for enhancing the accuracy of any ML algorithm. Its fundamental principle involves combining multiple models sequentially, with each subsequent model focusing on the errors made by the previous ones. At every iteration, a new weak learner is trained to correct the mistakes of the existing ensemble, leading to improved overall performance[39]. In mathematical terms, the model's predictions can be expressed as the sum of individual tree functions

$$\hat{y}_i = \sum_{k=1}^{K} f_k(x_i),\ f_k \in \Phi, \tag{310}$$

where K is the number of trees, and f_k represents the function in the functional space Φ. The goal of XGBoost is to minimize an objective function that encompasses both loss and regularization components. The objective function can be defined as following

$$Q = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k), \quad f_k \in \Phi, \tag{311}$$

Here, l symbolizes the loss function that measures the discrepancy between the predicted and real values, Q is the objective function and Ω is a regularization term that penalizes model complexity to avert overfitting[40].

Traditional optimization methods in Euclidean space cannot be used to find the solution because the model includes functions as parameters. To improve performance, additional training is necessary. The algorithm iteratively searches through the function space to find a novel function that can boost the target function, and then integrates it into the existing set of functions. The objective is to discover the function that will optimize the given objectives. If $\hat{y}_i^{(t)}$ is the prediction for the i-th instance at the t-th iteration, the task is to find the function that will optimize the objectives as follows

$$
\begin{aligned}
Q &= \sum_{i=1}^{n} l(y_i, \hat{y}_i)^{(t)} + \sum_{k=1}^{K} \Omega(f_k) \\
&= \sum_{i=1}^{n} l(y_i, \hat{y}_i)^{(t-1)} + f_t(x_i)) + \sum_{k=1}^{K} \Omega(f_k).
\end{aligned}
\tag{312}
$$

XGBoost applies second-order Taylor polynomial expansion to approximate the optimization process. This approach leads to an iterative optimization problem, where each iteration refines the model by considering gradients and Hessians. The objective function during the t-th iteration can be expressed as

$$
Q^{(t)} \approx \sum_{i=1}^{n} \left[l(y_i, \hat{y}_i^{t-1}) + g_i f_t(x_i) + \frac{1}{2} h_i f_t^2(x_i) \right] + \sum_{k=1}^{K} \Omega(f_k), \tag{313}
$$

where g_i and h_i represent gradient and hessian, respectively. The constant terms can be eliminated to obtain an approximated objective function at step t

$$
Q^{(t)} \approx \sum_{i=1}^{n} \left(g_i f_t(x_i) + \frac{1}{2} h_i f_t^2(x_i) \right) + \Omega(f_t). \tag{314}
$$

Let q be a mapping that assigns an index from $1, 2, \ldots, T$ to each input, and let ω be a vector that represents the scores of each region, where the function of ω is

$$
\begin{aligned}
f_t(x) &= \omega_q(x) \omega \in \mathbb{R}^T, \\
q &: \mathbb{R}^d \rightarrow \{1, 2, \ldots, T\}
\end{aligned}
\tag{315}
$$

In this function class, there is a regression tree that utilizes the decision tree structure on xi, and q represents this tree. Additionally, the authors can choose alternative forms of q to incorporate prior knowledge into the regions of interest. Furthermore, the authors must define the complexity of the function Ω as follows

$$
\Omega(f)_t = \frac{1}{2} \lambda \sum_{j=1}^{T} \omega_j^2 + \gamma T, \tag{316}
$$

where λ and γ are the regularization parameters. Once we apply certain mathematical operations to the objective function, the gain function used to evaluate and select the best split point for a leaf node becomes defined as the following one

$$
\text{Gain} = \frac{1}{2} \left(\frac{\left(\sum_{i \in I_L} g_i\right)^2}{\sum_{i \in I_L} h_i + \lambda} + \frac{\left(\sum_{i \in I_R} g_i\right)^2}{\sum_{i \in I_R} h_i + \lambda} - \frac{\left(\sum_{i \in I} g_i\right)^2}{\sum_{i \in I} h_i + \lambda} \right) - \gamma. \tag{317}
$$

The equation involves three terms representing the scores of the left, right, and original leaf nodes, respectively. Additionally, there is a fourth term, γ, which serves as a regularization factor for the new leaf node[41].

3.6 Long Short Term Memory

Long Short-Term Memory (LSTM) networks, a type of Recurrent Neural Network (RNN), were designed to resolve issues present in traditional RNNs, such as the vanishing and exploding gradient problems[42]. To address this issue, Hochreiter and Schmidhuber (1997) proposed a novel approach called LSTM networks[10]. LSTMs have undergone continuous improvements and have become a crucial component in the modeling of sequential data[43]. An LSTM is a neural network that is made up of three layers: an input layer, a hidden layer with memory cells, and an output layer. The number of neurons in the input layer corresponds to the number of features in the data, while the size of the output layer is determined by the task that the LSTM is being used for (e.g., binary classification). The core functionality of the LSTM resides in the hidden layer, specifically in the memory cells. Each memory cell contains three gates; the forget gate, the input gate, and the output gate as shown in Figure 3.1. The forget gate decides which information to discard from the cell state, the input gate controls which information to store in the cell state, and the output gate determines which information to reveal as the output. In other words, the LSTM uses the gates to control the flow of information in and out of the memory cell. This allows the LSTM to learn long-term dependencies in the data, which is essential for tasks such as financial forecasting.

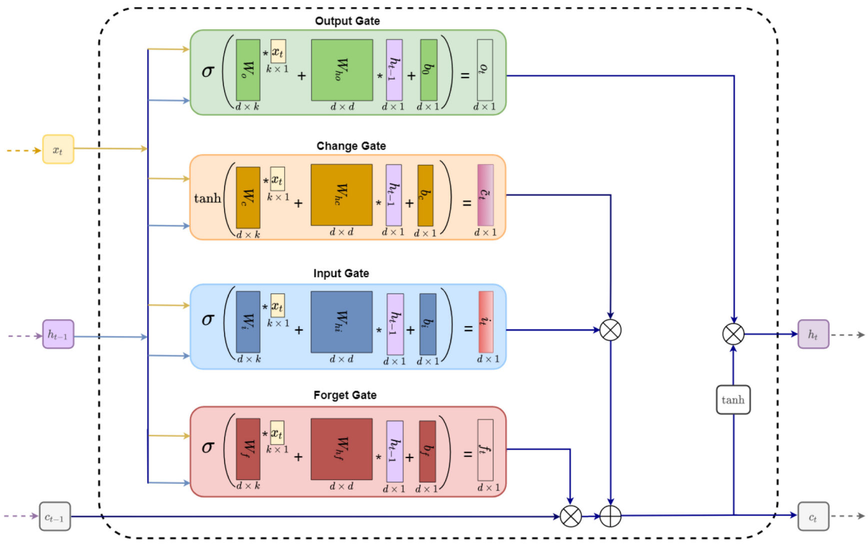

Figure 3.1: Long short-term memory (LSTM) architecture[1]

The forget gate decides which information from the previous cell state to keep and which information to discard. This is important because it allows the LSTM to forget irrelevant information and focus on the most important information. The input gate controls which new information is stored in the cell state. This is important because it allows the LSTM to learn new information and update its understanding of the data. The output gate determines which information from the cell state is revealed as the output. This is important because it allows the LSTM to generate predictions or classifications based on the information that it has learned. At each time step t, the memory cell processes the current input x_t and the previous time step's output h_{t-1}, using them to compute the following

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f). \tag{318}$$

The forget gate decides which parts of the previous cell state c_{t-1} to discard. It takes into account both the previous cell state h_{t-1} and the current input x_t and produces a value between 0 and 1 for each element of c_{t-1}. This measure reflects the significance of each component, where higher values indicate more crucial elements that should be preserved, and lower values suggest less essential elements that could be discarded. Here σ denotes the sigmoid function, W_f and b_f are the weight matrices and bias for the forget gate, respectively.

The input gate determines which new information to add to the cell state. It evaluates the current input x_t and decides which aspects of the input are most relevant to the current task, selecting and integrating this information into the cell state[44]

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i), \tag{319}$$

Next, generate a set of potential new values, $\tilde{c}_t$, that could be included in the state

$$\tilde{c}_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_c). \tag{320}$$

Now, update the old cell state, c_{t-1} into the new cell state c_t as following

$$c_t = f_t \cdot c_{t-1} + i_t \cdot \tilde{c}_t. \tag{321}$$

The forget gate's output f_t is multiplied element-wise with the previous state, producing a weighted sum of the past state. This weighted sum is then added to the product of the input gate's output i_t and the candidate values $\tilde{c}_t$, creating a hybrid state that combines both the past and present information. The output gate determines which components of the cell state will be exposed or presented outside of the LSTM network. It selects specific parts of the hybrid state produced by the forget gate and the input gate, allowing the network to communicate its internal state to the external world

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o). \tag{322}$$

The LSTM output h_t is computed by applying the hyperbolic tangent (tanh) to the cell state, and then multiplying the result by the output gate. Output gate can be shown as

$$h_t = o_t \cdot \tanh(c_t). \tag{323}$$

This process allows the LSTM network to selectively emphasize certain aspects of the cell state when generating its output.

3.7 Convolutional Neural Network

Convolutional Neural Networks (CNN) proposed by Kunihiko Fukushima (1998)[45], have shown great promise in predicting future values of time series data by discovering hidden patterns and characteristics within the sequences. In this context, we will delve into how CNNs work to extract meaningful insights from time series data. The core aspect of a CNN is the convolution operation, where specialized filters (or kernels) are applied to the input time series data. This process helps identify important patterns and features within the sequence. To perform a convolution on discrete one-dimensional signals f and g, it can be represented as a mathematical combination of the two signals using filters[11]. Defined as follows

$$(f * g)(i) = \sum_j f(j)g(i-j). \tag{324}$$

The selection of filters, typically optimized through training, enables the CNN to discern distinctive aspects or irregularities within the time series data. These filters allow the network to focus on particular patterns or features that might otherwise go unnoticed, making it more effective at analyzing and predicting complex time series phenomena.

Convolutional layers form the foundation of CNNs for time series analysis. These layers are tasked with identifying and extracting relevant spatial or temporal features from the raw input data. To accomplish this, the convolutional layers employ filters that are slid across the input, effectively transforming the data into feature maps or matrices. Following this transformation, the output often undergoes a nonlinear activation function like the Rectified Linear Unit (ReLU), thereby amplifying the model's ability to recognize intricate patterns within the data. In a convolutional layer, filters with weights w_h are applied across the input time series x, scanning each element of the series multiple times. This process produces a feature map a, where each element a_i captures information about the presence of certain patterns or features in the input data at position i

$$a(i, h) = (w_h * x)(i) = \sum_j w_h(j)x(i - j). \tag{325}$$

The suffix h denotes the different filters used in the layer, and the weight matrices w_h determine the strength of each filter's contribution to the output feature map.

After the convolutional layers, pooling layers play a crucial role in streamlining the extracted features, making the model more resilient to computational complexity and improving its overall robustness. Pooling techniques like maximum pooling or average pooling distill the salient information from the convolved features, creating a condensed representation that preserves only the most critical details. This abstraction enables the model to generalize better and adapt to new situations[46].

The fully connected layers can be seen in Figure 3.2, act as a conduit between the richly detailed feature space and the final outputs, facilitating tasks such as regression or classification. These layers synthesize the knowledge gained from earlier layers to produce accurate predictions or categorizations. To enhance performance, activation functions like the sigmoid or softmax may be utilized, particularly when dealing with multiple classes.

In the context of time series forecasting, CNNs are carefully crafted to identify and absorb translation-invariant features, which is crucial for accurately capturing long-term patterns, periodic fluctuations, or cyclic behaviors within the data. By focusing on these invariant features, CNNs can provide more reliable predictions and insights into the underlying dynamics of the time series.A common CNN architecture for one-

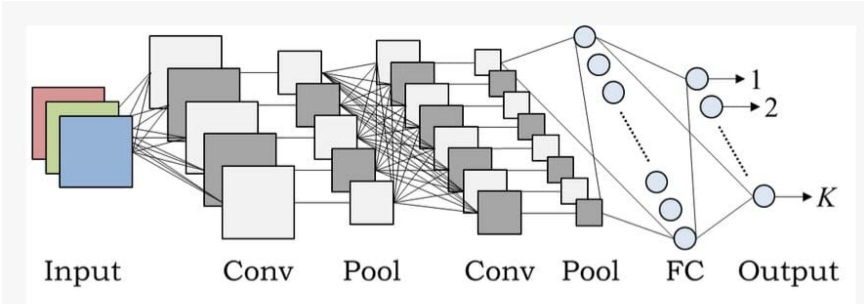

Figure 3.2: Typical Convolutional Neural Network architecture[2]

dimensional time series data consists of a set of convolutional layers, followed by pooling layers, and finally, fully connected layers. Each convolutional layer takes the input time series x and applies a set of filters to extract local features. The filters are learned during training and are applied in a sliding window fashion to scan the entire time series. The output of the convolutional layer is then passed through a pooling layer to reduce the spatial dimensions of the data and improve generalization. Finally, the pooled output is fed into fully connected layers to make predictions based on the learned features

$$a(i,h) = (w_h * x)(i) = \sum_j w_h(j)x(i-j), \tag{326}$$

where w_h it the weight matrix for the h-th filter. The modular design of CNNs provides considerable freedom in selecting appropriate layers and filter combinations, allowing for customized feature extraction. Once the network has been trained, its output is fine-tuned for the specific forecasting task at hand, with the weights refined to minimize the difference between the predicted and actual future values of the time series. This adaptability makes CNNs a powerful tool for tackling diverse time series forecasting challenges. CNNs offer a versatile and potent platform for time series forecasting, leveraging convolutional operations, pooling, and fully connected layers to decipher intrinsic temporal patterns and correlations. By harnessing these techniques, CNNs can adaptively model complex time series data, yielding accurate predictions and insights into dynamic systems.

3.8 Hybrid Method

Time series forecasting frequently involves working with data that exhibits both linear and nonlinear patterns[47]. To effectively capture these complex behaviors, combining traditional statistical methods with cutting-edge ML techniques can be extremely valuable. One potential solution is a hybrid ARIMA-LSTM model, which seeks to simultaneously account for both the linear and nonlinear elements within the time series data. The initial step in this combined strategy is to pinpoint and represent the linear connections existing in the time series data. For this purpose, we rely on the ARIMA model, a well-known statistical technique that effectively captures linear dependences in single-variable time series data. Specifically, we use ARIMA to extract the linear portion of the time series y_t, which can be mathematically represented as

$$L_t = ARIMA(y_t). \tag{327}$$

Following the extraction of the linear component, we calculate the residuals to uncover any remaining nonlinear patterns hidden within the data. The residuals are defined as

$$e_t = y_t - \hat{L}_t. \tag{328}$$

The residuals left after extracting the linear component may contain intricate nonlinear structures that exceed the capabilities of the ARIMA model. To address this limitation, we incorporate Long Short-Term Memory (LSTM) networks, a specialized type of Recurrent Neural Network, to capture the complex temporal patterns and long-term dependencies within the sequential data. LSTMs are ideally suited for modeling these nonlinear residuals, enabling the discovery of subtle yet critical patterns in the data

$$e_t = LSTM(e_{t-1}, e_{t-2}, ..., e_{t-n}) + \epsilon_t, \tag{329}$$

where ϵ_t represent the random error, and n is the window size for considering past residuals.

Finally, we integrate the linear and nonlinear components to obtain the final forecast values. The comprehensive forecast is represented as:

$$\hat{y}_t = \hat{L}_t + \hat{N}_t \tag{330}$$

where $\hat{N}_t$ is the forecasted nonlinear component obtained from the LSTM.

The proposed hybrid ARIMA-LSTM model offers a more sophisticated and effective forecasting approach by synergizing the advantages of both linear and nonlinear modeling techniques. The ARIMA model accurately captures linear trends and seasonality, while the LSTM model expertly handles complex nonlinear patterns within the data. By integrating these two stages, the hybrid model aims to provide more accurate and reliable forecasts across various time series challenges. This innovative methodology recognizes the diversified nature of real-world time series data and supplies a versatile framework that adapts to the distinct features of the specific dataset being analyzed.

3.9 Performance Measures

The performance of a predictive model can be evaluated using a range of quantitative metrics that offer insights into different aspects of forecasting accuracy, bias, and reliability. These metrics enable assessments of the model's ability to accurately predict future outcomes, its tendency towards systematic error or bias, and its consistency and stability over time.

3.9.1 Root Mean Square Error

The Root Mean Squared Error (RMSE) is a widely utilized statistical metric that assesses the difference between predicted values and actual or observed values. RMSE is derived from the Mean Squared Error (MSE), which represents the average of the squared differences between the predicted and observed values. The RMSE is computed by taking the square root of the MSE, as indicated by the following formula

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)^2}, \tag{331}$$

where t is the number of observations, y_t is the actual value for the t-th observation, $\hat{y}_t$ is the predicted value for the i-th observation.

The RMSE is expressed in the same units as the dependent variable, allowing for easy interpretation and comparison. A lower RMSE value suggests a better fit, as it implies that the model predictions are, on average, nearer to the actual values. In other words, a lower RMSE indicates that the model is performing well and providing more accurate predictions. The RMSE is preferred for its sensitivity to significant errors, as the squaring operation magnifies the impact of substantial deviations. This makes RMSE particularly relevant in situations where large mistakes are particularly undesirable.

3.9.2 The Mean Absolute Percentage Error

The Mean Absolute Percentage Error (MAPE) calculates the absolute percentage difference between the actual and predicted values, averaged across all data points. The formula for MAPE is given by

$$MAPE = \frac{100\%}{n} \sum_{t=1}^{n} \left| \frac{y_t - \hat{y}_t}{y_t} \right|, \tag{332}$$

where n is the number of data points, y_t is the actual value for the t-th observation, $\hat{y}_t$ is the predicted value for the t-th observation.

The MAPE has several benefits, including its scalability, as it can be used to compare forecasts across different scales and units. Additionally, it is often expressed as a percentage, with lower values indicating a better fit of the predicted values to the actual ones. However, it's important to note that MAPE has some limitations, particularly when dealing with zero or near-zero actual values. In such cases, the MAPE can become disproportionately large or undefined, leading to potentially misleading interpretations.

3.9.3 Theil's Inequality Coefficients

There are two commonly used forms of Theil's Inequality Coefficient, which can be represented as U_1 and U_2. These coefficients provide different insights into the forecasting process and serve unique purposes.

Theil's U_1 Coefficient, introduced by Theil in 1958, U_1 is a measure of forecast accuracy (Theil, 1958, pp 31-42)[48]. The formula for U_1 is

$$U_1 = \frac{\sqrt{\frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)^2}}{\sqrt{\frac{1}{n}\sum_{t=1}^{n} y_t^2} + \sqrt{\frac{1}{n}\sum_{t=1}^{n} \hat{y}_t^2}}, \tag{333}$$

U_1 ranges between 0 and 1, with 0 representing a perfect forecast.

Theil's U_2 Coefficient: Proposed by Theil in 1966, U_2 is a measure of forecast quality (Theil, 1966, chapter 2)[48]. The formula for U_2 is

$$U_2 = \frac{\sqrt{\sum_{t=1}^{n-1}(\frac{\hat{y}_{t+1} - y_{t+1}}{y_t})^2}}{\sqrt{\sum_{t=1}^{n-1}(\frac{\hat{y}_{t+1} - y_t}{y_t})^2}}, \tag{334}$$

where n is the number of data points, y_t is the actual value for the t-th observation, $\hat{y}_t$ is the predicted value for the t-th observation. U_2 also has a lower bound of 0, but unlike U_1, 1 represents a naive forecast. If the coefficient is lower than 1, it signifies a better forecast than the naive method, and if the coefficient is greater than 1, it denotes a less accurate forecast than the naive method. These two coefficients offer distinct perspectives on the quality of time series forecasting and can be used in tandem to provide a comprehensive evaluation of a given model performance.

CHAPTER 4

ANALYSIS

In this section, analysis steps will be explained with details. It will begin by familiarizing ourselves with the data, recognizing its importance in shaping the foundation of our research. Next, we proceed to prepare the data for analysis, employing crucial preprocessing techniques to ensure its quality and consistency. By addressing any irregularities, missing values, or outliers, we enhance the accuracy and reliability of our subsequent explorations. Our attention then turns to a detailed examination of the time series data, where we carefully investigate subsets from each frequency range, using both visual and numerical representations. This approach enables us to uncover the underlying structures and patterns within the data, paving the way for more advanced analyses. While the primary focus was on the test datasets, it's essential to note that the training set also underwent rigorous evaluation. We ensured that the models were robust and not prone to issues commonly encountered in time series forecasting, such as overfitting. Through meticulous cross-validation and hyperparameter tuning, the training phase was crafted to ensure generalizability and reliability in out-of-sample predictions. We chose to prioritize the presentation of results on validation and test datasets in this research, as these offer a more transparent evaluation of the model's forecasting capabilities in unseen scenarios. Subsequent to this initial examination, we engage in rigorous empirical analysis, applying appropriate statistical methods to each frequency subset. Then the findings are presented using both visualizations and numerical summaries, allowing for a comprehensive understanding of the results.

4.1 Data Introduction

This study examines the Russian Ruble (RUB), Turkish Lira (TRY), and Euro (EUR) - relative to the US Dollar (USD). Specifically, daily data spanning from 2005 to 2023-07 is collected. To capture the complex interactions between these currencies and economic variables, this study employs a multivariate framework that incorporates relevant external variables. These external variables include Consumer Price Index (CPI), M3 Money Supply, and Gross Domestic Product (GDP) for each country under consideration. Multiple sources used to gather the data, including the Organisation for Economic Co-operation and Development (OECD), Federal Reserve Economic Data (FRED), Investing.com, and Yahoo Finance. Descriptice statistics of the currencies are shown below.

Table 4.1: Descriptive Statistics

Measurement	**EUR**	**RUB**	**TRY**
Total Observation	4780	4780	4780
Null Count	0	0	0
Min	0.63	23.17	1.14
Max	1.04	138.97	25.99
Range	0.42	115.80	24.85
Sum	3893.16	224705.29	20273.37
Median	0.81	35.50	2.14
Mean	0.81	47.01	4.24
Var.	0.01	385.60	20.95
Std. Dev.	0.08	19.64	4.58
Coef. Var.	0.10	0.42	1.08

Each currency having 4780 observations in total. For the EUR, the values ranged from 0.63 to 1.04, with a mean of 0.81. The standard deviation is relatively small (0.08), indicating a low level of volatility in the Dollar/Euro exchange rate during the observation period. The RUB has a wider range, with minimum and maximum values of 23.17 and 138.97 respectively. The average value of the Ruble/Dollar exchange rate was 47.01, with a larger standard deviation of 19.64, signifying higher volatility compared to the EUR. The TRY also demonstrated high volatility, reflected by a standard deviation of 4.58. The USD/TRY exchange rate values ranged from 1.14 to 25.99, with an average value of 4.24. It is also worth noting the Coefficient of Variation (Coef. Var.) for each currency. This measure, defined as the ratio of the standard deviation to the mean, provides a standardized measure of dispersion. The EUR had the lowest Coef. Var. (0.10), indicating the least relative variability. The RUB and TRY had higher Coef. Var. of 0.42 and 1.08 respectively, reflecting greater relative variability in their exchange rates. The line graphs for the currencies are shown below.

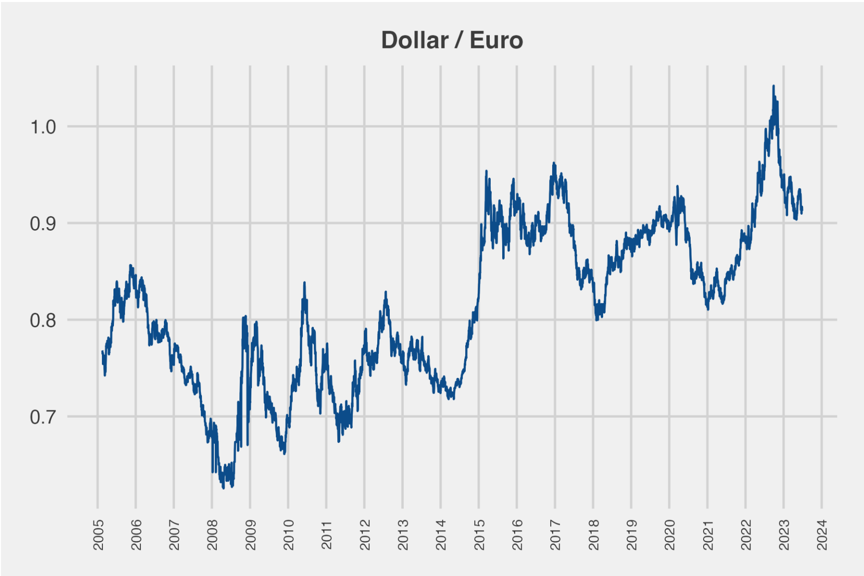

Figure 4.1: Euro against Dollar

The USD/EUR exchange rate, as depicted in Figure 4.1, exhibits characteristics of a random walk process, fluctuating up and down around a certain level without a clear upward or downward trend. This kind of behavior is typical for exchange rates in economies with stable monetary policies and low inflation rates.

On the other hand, as illustrated in Figure 4.2, the USD/TRY rate seems to have a steady but slow upward trend for the most part, interrupted by sudden spikes, particularly after 2022. This pattern could be due to a combination of inflation and economic policy changes, among other factors.

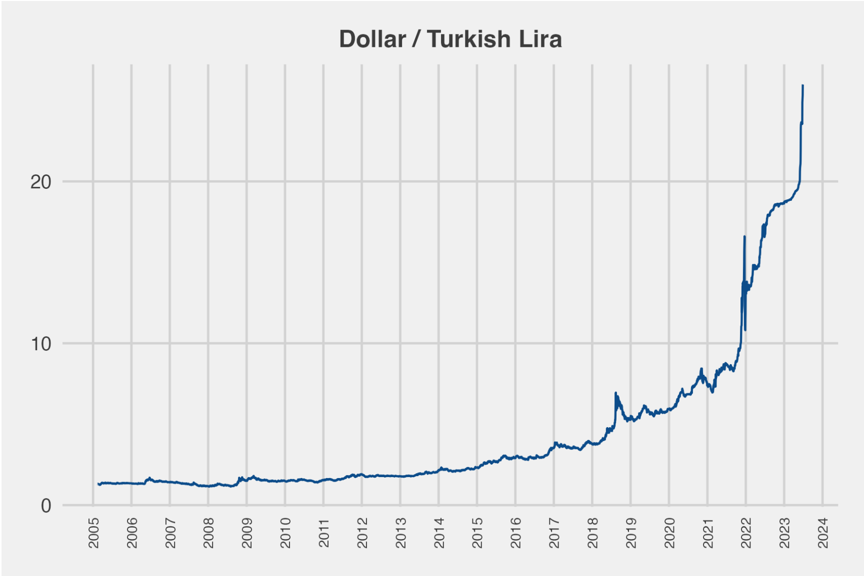

Figure 4.2: Turkish Lira against Dollar

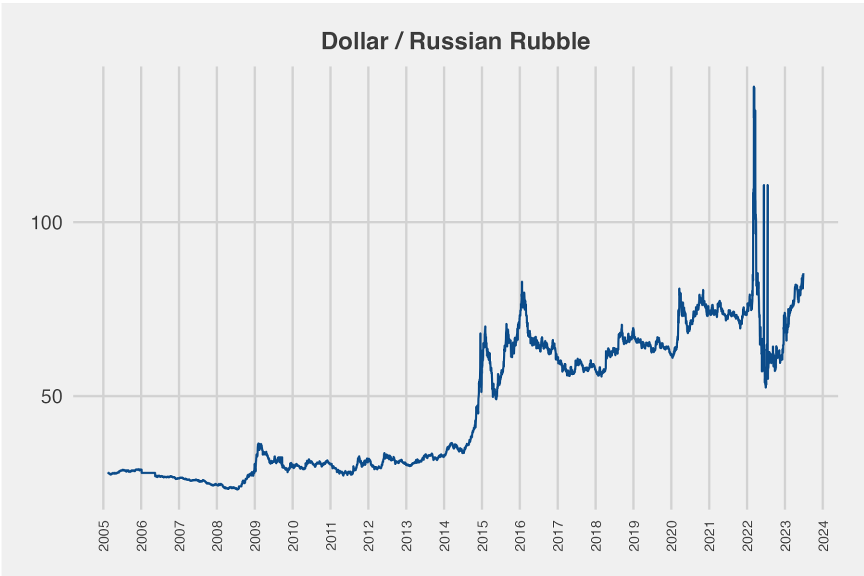

Figure 4.3: Russian Rubble against Dollar

The USD/RUB exchange rate, as presented in Figure 4.3, exhibits a pattern that is similar to the USD/TRY, characterized by a general upward trend and more pronounced fluctuations. This could indicate a higher level of economic uncertainty and volatility in Russia, influenced by various macroeconomic factors such as oil prices and international sanctions. In all three cases, the series appear to be non-stationary, with their mean and variance changing over time. This non-stationarity is important to take into account when modelling these series, as many time series models assume stationarity[49].

4.2 Data Preprocessing

In our analysis, we adopt a methodology that acknowledges the unique characteristics of financial time series data. Specifically, we work with raw data consisting of foreign exchange rates, which naturally incorporate various influences such as economic, social, and political factors. Therefore, we do not perform outlier detection or removal on the dataset. This decision is in line with the characteristics of financial time series data, which frequently exhibit unexpected shifts, spikes, or so-called 'outliers.' These apparent anomalies may actually be the important aspects of the underlying system, as suggested by Tsay (2005) and Zivot & Wang (2006)[50]. For instance, these 'outliers' could represent significant events like central bank interventions, political developments, or rapid changes in investor sentiment.

4.2.1 Stationarity

As it mentioned earlier we identified a clear trend component in the series, suggesting non-stationarity. To formally test the stationarity of our series, Augmented Dickey-Fuller (ADF) test applied. The ADF test is a popular tool in econometric analysis to detect the presence of unit roots and, thereby, non-stationarity in time series data (Dickey and Fuller, 1979)[51]. As anticipated, the results of the ADF test confirmed the non-stationary nature of our series. Differencing is applied to our data. Differencing is a transformation that can help to stabilize the mean of a time series by removing changes in the level of a time series, and thus eliminate trend and seasonality[52]. Following the differencing procedure ADF test conducted again on the differenced series. The results confirmed that our transformation was successful, and we obtained a stationary time series.

4.2.2 Time Delay Embedding

To decide time delay embedding autocorrelation function (ACF) and partial autocorrelation function (PACF) plots are checked and observed that the autocorrelation cut off sharply after one lag. Thus, this was an indication that using a lag of one would sufficiently capture the temporal dependencies present in the data. Nonetheless, to prevent leakage, lags are used with the forecast horizon. For example, if we are predicting the currency value at time t for a forecast horizon is 3, we would use the currency value at time t-3 as a predictor in our model.

4.2.3 Fourier Transform

To prepare the data, we used a mathematical method called Fourier Transforms. This technique helps us identify repeating patterns in the data by breaking it down into different frequency components. It is similar to taking a complex signal and representing it as a combination of simple sine and cosine waves, which makes it easier to understand the underlying patterns.

$$F(\omega) = \int_{-\infty}^{\infty} f(t) e^{-i\omega t} dt, \tag{41}$$

where $F(\omega)$ is the Fourier transform of the function $f(t)$ and $f(t)$ is the function in the time domain to be transformed, ω is the frequency variable, $e^{-i\omega t}$ is the basis function. By transforming the data into the frequency domain, the aim is to capture the complex, cyclical dynamics inherent in the currency exchange rate data.

4.3 Model Implementation

In the upcoming section, we delve deeper into the structure of our forecasting approach, which divides into two main components. Firstly, we rely on a combination of lags of the relevant currencies and Fourier terms to capture historical patterns and cyclical fluctuations in exchange rates. This approach enables us to leverage both empirical data and theoretical underpinnings to inform our predictions. Secondly, we integrate three well-established economic models - PPP, UIRP, and MM- into our framework. By fusing these models with historical data and cyclical factors, our schema aims to achieve a balanced blend of empirical insights and sound theoretical principles to forecast currency exchange rates effectively.

The fundamental value of the exchange rate f_t is given by the difference between the domestic and foreign interest rates. We can rewrite this as

$$UIRP : f_t = i_t - i_t^*, \tag{42}$$

where i_t and i_t^* are the domestic and foreign interest rates. Under PPP, the fundamental value of the exchange rate f_t is the difference between the log of domestic and foreign price levels. We can rewrite this as

$$PPP : f_t = \log(P_t) - \log(P_t^*), \tag{43}$$

where (P_t) and (P_t^*) are the domestic and foreign price levels. In the MM, the fundamental value of the exchange rate f_t is the difference between the log of domestic and foreign money supplies, adjusted by the difference between the logs of the domestic and foreign real incomes. We can write this as

$$MM : f_t = [\log(M_t) - \log(M_t^*)] - [\log(Y_t) - \log(Y_t^*)], \tag{44}$$

where (M_t) and (M_t^*) are the domestic and foreign money supplies, (Y_t) and (Y_t^*) are the domestic and foreign real incomes, respectively.

The overall model for the exchange rate, taking into account these fundamentals, can be written as

$$c_t = \beta_0 + \beta_i f_t + e_t. \tag{45}$$

Here, c_t is the difference of the exchange rate at time t, f_t is the fundamentals (calculated according to the structural model used), e_t is the regression error which is assumed to be a white noise, and β_0 and β_i parameters to be estimated.

The forecasting was conducted over four distinct horizons, specifically targeting 3-day, 7-day, 14-day, and 30-day periods. These intervals were selected to provide insights into short-term and medium-term trends in the currency exchange rates, thereby enabling a comprehensive understanding of the market dynamics across various time frames.

4.3.1 SARIMA

The *autoarima* function in Python, part of the *pmdarima* library, provides a comparable model fitting process for time series data. This function, conducts automatic differencing to remove any present stochastic trends and make the series appropriate for ARIMA model suggestions. The function identifies the optimal order of the Autoregressive (AR) and Moving Average (MA) components by minimizing the Akaike Information Criterion (AIC) or the Bayesian Information Criterion (BIC). These metrics are given by the following formulas.

$$AIC = 2k - 2ln(\hat{L}) \tag{46}$$

$$BIC = ln(n)k - 2ln(\hat{L}) \tag{47}$$

The algorithm selects the most suitable model based on three criteria: the length of the series (n), the number of parameters to be estimated (k), and the maximum value of the likelihood function for the observed data (L). To do this, it uses a statistical method called Maximum Likelihood Estimation (MLE) which finds the best set of model parameters by minimizing the likelihood calculated from the observed data.

The autoarima function effectively manages seasonal data by incorporating seasonal unit root testing and differencing into its analysis. When dealing with series that display seasonality, the function automatically identifies and removes any existing seasonal unit roots before estimating the appropriate seasonal AR and MA orders using the AIC or BIC criterion.

4.3.2 Support Vector Regression

We built the SVR model using the widely adopted SVM framework. Our implementation adheres to the conventional practice of ϵ-insensitive regression for regression tasks. Two crucial hyperparameters in the SVR model, C and γ, require careful tuning for optimal performance. These parameters control the penalty term and the influence radius of individual training examples, respectively. Using the Bayesian optimization technique integrated within the *skopt.BayesSearchCV* module, we meticulously search for the ideal values of these hyperparameters that balance model complexity and fitting accuracy to the training data, as measured by the cross-validated MSE.

4.3.3 XGBoost

To implement the XGBoost model, we leveraged the *xgboost* package in Python, a well-established and potent gradient boosting framework renowned for its speed and accuracy. The model employs a combination of decision trees and pursues an iterative approach, enhancing performance by targeting the mistakes made during preceding iterations. The model is configured and trained with careful consideration of the hyperparameters that influence the learning process. To optimize these hyperparameters and to achieve the best predictive performance, Bayesian Optimization is employed.

4.3.4 LSTM

When working with *Keras* library, it's essential to provide the input data in a 3D tensor format, consisting of features, time, samples. Samples represent the number of observations being processed in batches, while time indicate the lags. Features describe the number of predictors. Before training, the data undergoes scaling, resulting in a 3D array format for both the train and test data. The LSTM model's architecture consists of an input layer, an LSTM layer, a densely connected hidden layer with and a ReLU activation function, and a linear output layer. An Exponential Decay learning rate scheduler is applied to gradually reduce the learning rate during training, aiding in convergence. The performance of the model was optimized using *Optuna*, a state-of-the-art open-source optimization library. Specifically, Optuna was used to tune the hyperparameters of the LSTM layer, including the number of hidden layers and the size of each layer.

4.3.5 CNN

The Convolutional Neural Networks (CNN) model was built using the Keras library, following the same architecture as the LSTM model described previously. The primary distinction between the two models lies in the usage of the Conv1D argument.

4.3.6 ARIMA - LSTM Hybrid

Hybrid forecasting model structred combining the strengths of ARIMA and LSTM. The ARIMA model was first fitted to the data, and the residuals were extracted for further processing. These residuals were then fed into an LSTM model, which was trained to predict future errors of ARIMA model over similar time horizons. The predictions generated by both models were then combined to produce the final forecasts. This hybrid approach allowed us to leverage the ability of ARIMA to capture seasonality and nonlinear patterns in the data, while also harnessing the power of LSTM to handle complex temporal dependencies and non-stationarities.

In the subsequent subsection, due to the extensive nature of our analysis with over 300 models, visualizing each one is not practical. As such, we present the results of the models in tabulated form and provide graphical representations for the top three models corresponding to each currency.

4.4 Emprical Analysis

This section presents a comparative analysis of the out-of-sample performance for the exchange rates of Euro, Turkish Lira, and Russian Ruble against the United States Dollar. The analysis encompasses forecast horizon times of 3, 7, 14, and 30 days, and it investigates the efficacy of predictive models under two distinct structures. Specifically, we consider models that incorporate economic variables and those that do not. The performance evaluation involves a comprehensive set of metrics, including the Root Mean Square Error, the Mean Absolute Percentage Error, and Theil's U statistic. These measurements offer a robust understanding of the predictive accuracy of the models, providing insights into both the magnitude and direction of the forecasting errors. Tables detailing these results will be presented and are designed to facilitate a nuanced comparison between the models with different underlying structures. In Figure 4.4, we present a comprehensive flowchart detailing the empirical analysis undertaken in this study.

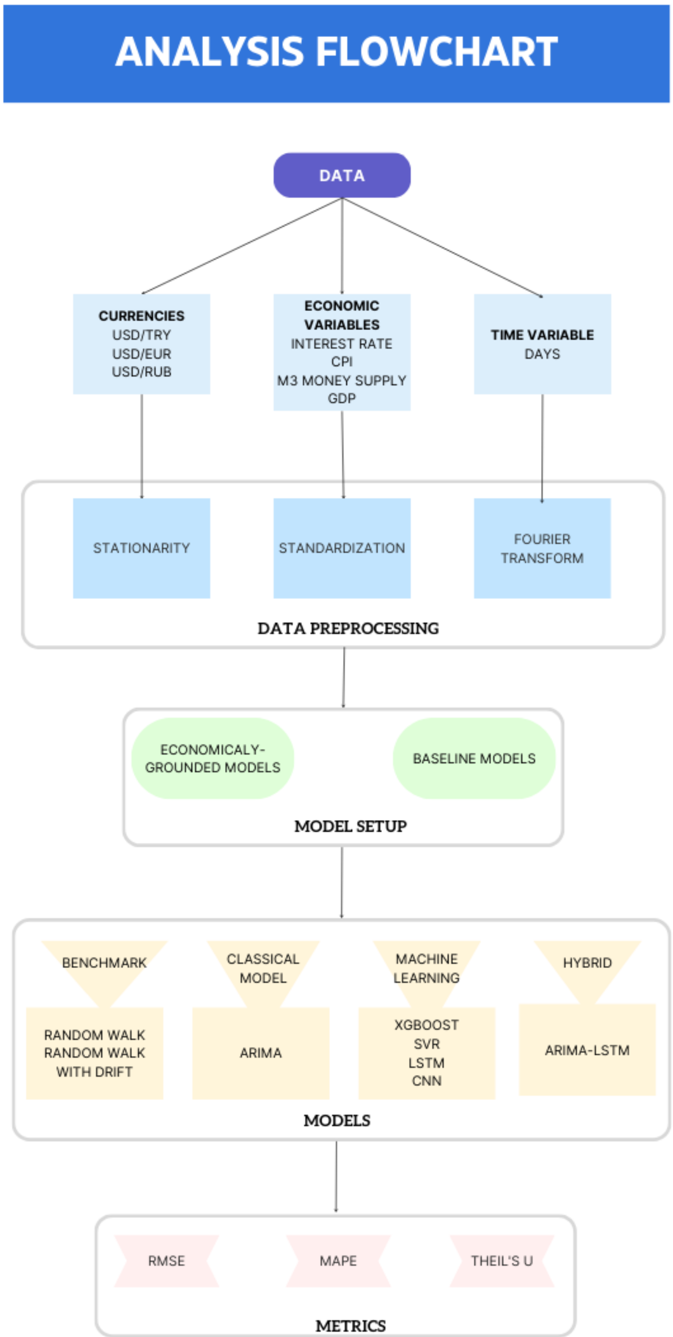

Figure 4.4: Analysis Flowchart

4.4.1 3 Day Ahead

In our analysis encompassing a three-day forecast horizon for the Turkish Lira among the myriad of models evaluated, the ARIMA models conspicuously dominate the metrics for TRY can be seen in Figure 4.5, underscoring their robustness in this specific context.

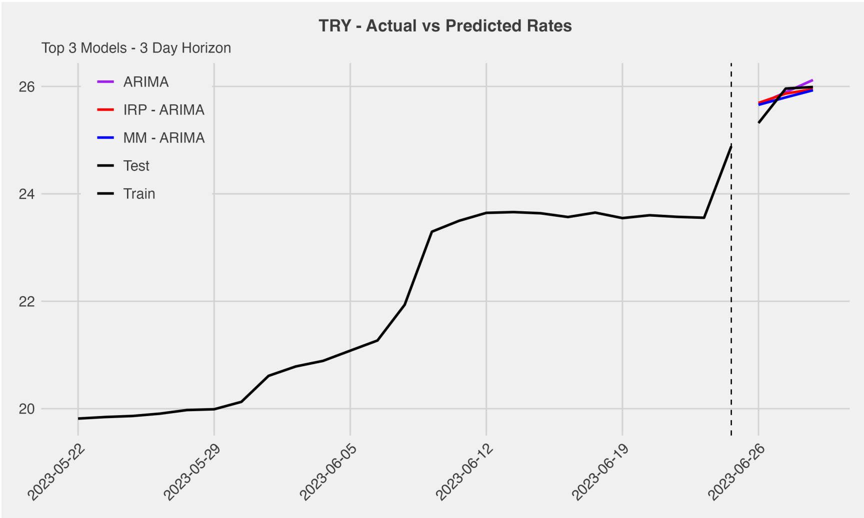

Figure 4.5: Top 3 Models for 3-Day Forecast: Turkish Lira

Within this subset of models, the univariate ARIMA technique, combined with the UIRP and MM, emerges as the most effective. Notably, the models accurately captured the upward trend in the currency, reflecting their ability to adapt to changing market conditions. Only noticeable distinction is subtle differences in trend intensity exist among the models.

In the context of the three-day forecast horizon for the Euro (EUR), Figure 4.6 reveals that the top-performing models are IRP-LSTM, IRP-CNN, and PPP-CNN.

Despite their strong overall performance, these models struggled to accurately predict the sudden drop in the EUR's value. Specifically, the PPP-CNN model, which combines the PPP theory with CNN, was particularly inadequate in capturing the downward trend.

Predicting the three-day forecast of the RUB is rendered complex due to the ongoing conflict in Ukraine, introducing an additional variable that complicates the task.

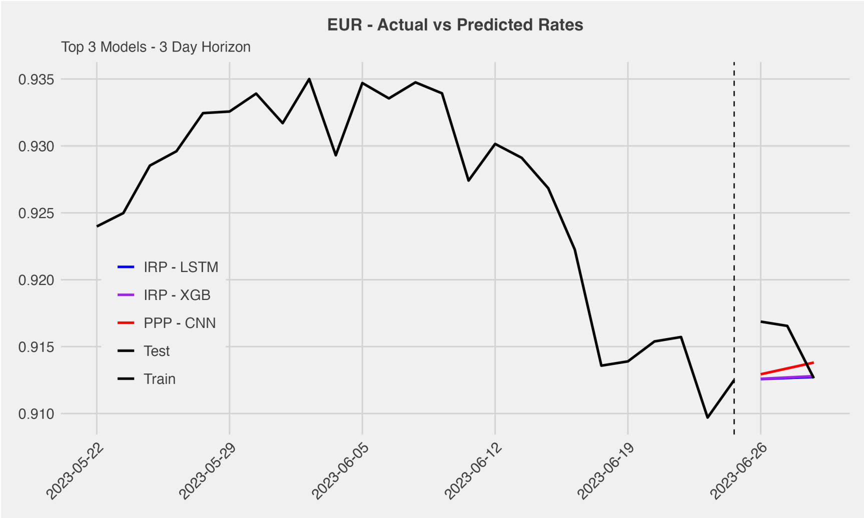

Figure 4.6: Top 3 Models for 3-Day Forecast: Euro

The looming threat of economic sanctions and the potential impact of military escalation on the Russian economy add layers of uncertainty to the currency's outlook, emphasizing the importance of continuous monitoring and adaptable forecasting.

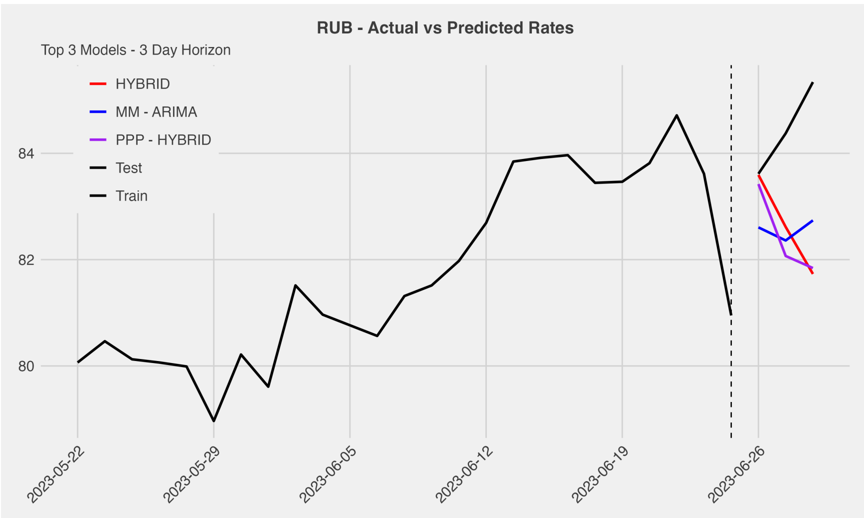

Figure 4.7: Top 3 Models for 3-Day Forecast: Russian Rubble

Figure 4.7 shows that among most successful models Hybrid (ARIMA-LSTM), MM-Arima, and PPP-Hybrid, the latter two exhibit downward trends in their predictions, whereas the former displays more accurate predictions despite deviating from the actual values.

Results of the 3 day ahead forecast using economical models shown in Table 4.2. Examining the results for the RUB, the ARIMA model coupled with the monetary model emerges as a standout. It offers the lowest RMSE, MAPE, TheilsU1, and TheilsU2. Similarly, the XGB model's integration with the IRP model showcases commendable results. The harmonizing of an ensemble technique like gradient boosting with economic indicators like IRP suggests a synergistic relationship when predicting RUB's movements. Shifting focus to the Euro, the XGB model,shines, outperforming other models across most criteria and showing remarkable consistency across the IRP, PPP, and MM models with close second CNN.

The evaluation for TRY shows the ARIMA model continues its trend of stable performance, holding its ground even when intertwined with the PPP. The LSTM model, however, offers intriguing insights.

An examination of Table. 4.3 reveals that the majority of forecasting models exhibit only marginal improvements or losses in accuracy when economic variables are removed. The Hybrid model is the standout performer for RUB when excluding economic variables. The XGB model remains the top choice for EUR. While not leading, maintains decent numbers across the board, especially in TheilsU1 (0.002).

Table 4.2: 3-Day Horizon Forecast Performance with Economic Models

Currency	Criteria	ARIMA			XGBoost			SVR			LSTM			CNN			HYBRID			WBRW	RW
		UIRP	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**-**	**-**
TRY	RMSE	0.224	0.232	**0.22**	0.6	0.747	0.546	0.951	0.864	0.863	0.898	0.898	0.898	0.903	0.903	0.899	0.892	1.117	1.201	0.924	0.912
	MAPE	**0.666**	0.881	0.736	2.169	2.717	1.972	3.479	3.163	3.158	3.272	3.271	3.272	3.289	3.29	3.275	3.414	4.138	4.581	3.367	3.319
	Theil's U_1	**0.004**	0.005	**0.004**	0.012	0.015	0.011	0.019	0.017	0.017	0.018	0.018	0.018	0.018	0.018	0.018	0.017	0.022	0.023	0.018	0.018
	Theil's U_2	**0.158**	0.433	0.272	1.544	1.916	1.405	2.427	2.204	2.2	2.3	2.3	2.3	2.312	2.313	2.303	1.793	2.816	2.187	2.367	2.335
EUR	RMSE	0.004	0.004	0.004	**0.003**	**0.003**	**0.003**	0.015	0.015	0.016	**0.003**	0.008	0.004	0.005	**0.003**	**0.003**	0.101	0.141	0.014	**0.003**	**0.003**
	MAPE	0.477	0.455	0.393	**0.301**	**0.301**	0.306	1.623	1.624	1.691	**0.301**	0.844	0.439	0.552	**0.301**	0.302	11.07	15.407	1.472	0.311	0.325
	Theil's U_1	**0.002**	**0.002**	**0.002**	**0.002**	**0.002**	**0.002**	0.008	0.008	0.009	**0.002**	0.004	**0.002**	0.003	**0.002**	**0.002**	0.059	0.083	0.007	**0.002**	**0.002**
	Theil's U_2	1.623	1.511	1.395	0.982	1.005	1.019	5.577	5.755	6.377	0.993	2.962	1.414	1.819	**0.863**	1.025	36.457	50.725	4.929	1.033	1.067
RUB	RMSE	4.144	4.176	**1.987**	3.356	3.466	3.463	3.528	3.522	3.569	3.539	3.538	3.52	3.498	3.538	4.03	3.123	2.422	26.095	3.564	3.691
	MAPE	4.59	4.624	**2.214**	3.89	4.004	4.011	4.095	4.089	4.139	4.102	4.101	4.081	4.058	4.096	4.694	3.364	2.355	30.86	4.13	4.26
	Theil's U_1	0.025	0.025	**0.012**	0.02	0.021	0.021	0.021	0.021	0.022	0.021	0.021	0.021	0.021	0.021	0.024	0.019	0.015	0.134	0.022	0.022
	Theil's U_2	5.615	5.658	**2.679**	4.274	4.449	4.41	4.474	4.466	4.536	4.501	4.498	4.472	4.443	4.505	5.071	4.303	3.409	29.764	4.533	4.73
AVG	RMSE	1.457	1.471	**0.737**	1.32	1.405	1.337	1.498	1.467	1.483	1.48	1.481	1.474	1.469	1.481	1.644	1.372	1.227	9.103	1.497	1.535
	MAPE	1.911	1.987	**1.114**	2.12	2.341	2.096	3.066	2.959	2.996	2.558	2.739	2.597	2.633	2.562	2.757	5.949	7.3	12.304	2.603	2.635
	Theil's U_1	0.01	0.011	**0.006**	0.011	0.013	0.011	0.016	0.015	0.016	0.014	0.014	0.014	0.014	0.014	0.015	0.032	0.04	0.055	0.014	0.014
	Theil's U_2	2.465	2.534	**1.449**	2.267	2.457	2.278	4.159	4.142	4.371	2.598	3.253	2.729	2.858	2.56	2.8	14.184	18.983	12.293	2.644	2.711

Table 4.3: 3-Day Horizon Forecast Performance using Time Variables Only

Currency	Criteria	ARIMA	XGBoost	SVR	LSTM	CNN	HYBRID	WBRW	RW
TRY	RMSE	**0.217**	0.877	0.923	0.898	0.898	1.064	0.912	0.924
	MAPE	**0.71**	3.206	3.373	3.271	3.272	4.081	3.319	3.367
	Theil's U_1	**0.004**	0.017	0.018	0.018	0.018	0.02	0.018	0.018
	Theil's U_2	**0.223**	2.24	2.359	2.3	2.301	2.016	2.335	2.367
EUR	RMSE	0.004	**0.003**	0.015	0.005	0.004	0.083	**0.003**	**0.003**
	MAPE	0.399	**0.307**	1.597	0.568	0.458	9.052	0.325	0.311
	Theil's U_1	**0.002**	**0.002**	0.008	0.003	**0.002**	0.047	**0.002**	**0.002**
	Theil's U_2	1.446	**1.023**	5.536	1.879	1.479	29.879	1.067	1.033
RUB	RMSE	4.152	3.361	3.503	3.525	3.548	**2.318**	3.691	3.564
	MAPE	4.599	3.908	4.07	4.086	4.106	**2.114**	4.26	4.13
	Theil's U_1	0.025	0.02	0.021	0.021	0.021	**0.014**	0.022	0.022
	Theil's U_2	5.625	4.259	4.435	4.479	4.521	**3.263**	4.73	4.533
AVG	RMSE	1.458	1.414	1.48	1.476	1.483	**1.155**	1.535	1.497
	MAPE	**1.903**	2.474	3.013	2.642	2.612	5.082	2.635	2.603
	Theil's U_1	**0.01**	0.013	0.016	0.014	0.014	0.027	0.014	0.014
	Theil's U_2	**2.431**	2.507	4.11	2.886	2.767	11.719	2.711	2.644

For the TRY, the ARIMA model appears to be the best pick, irrespective of the inclusion of economic variables. Its performance remains relatively stable. Random Walk benchmarks in shows competitiveness in EUR. A closer inspection of the data indicates that the superiority of specific models varies depending on the currency in question. These observations underscore the importance of carefully considering the suitability of individual models for each currency market, rather than assuming a universal optimal approach.

4.4.2 7 Day Ahead

ARIMA and its Hybrid dominated 7-day forecasting horizon for TRY. It can visually inspect the predicted values against the actual ones in Figure 4.8, and notice that all models have generally followed the pattern of the realized values.

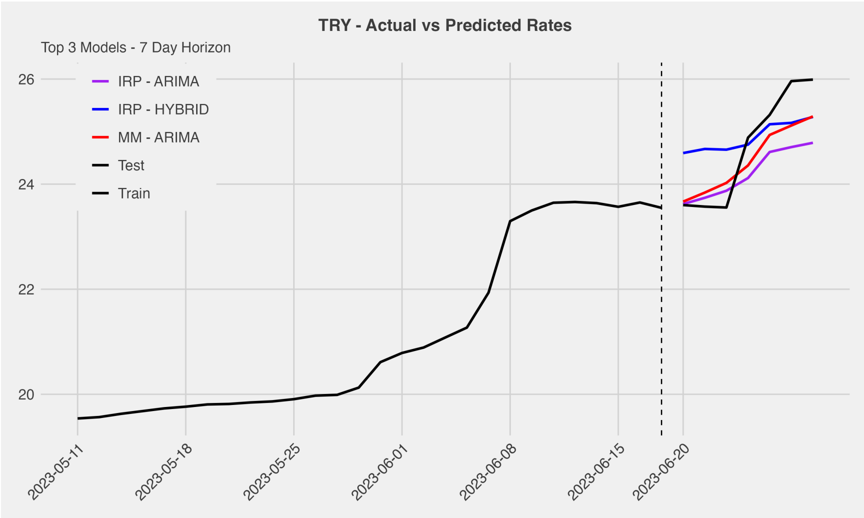

Figure 4.8: Top 3 Models for 7-Day Forecast: Turkish Lira

Using IRP, LSTM and XGB algorithms alongside with Monetary model-XGB constructed three top-performing models for euro exchange rate forecasting. These models are depicted in Figure 4.9. Our findings indicate that while the models did not perfectly capture the variability of the currency, they exhibited overall accurate predictions. Furthermore, our analysis revealed that the predictions made by the three models were largely consistent with one another.

For the Russian Ruble, as shown in Figure 4.10, the XGB and LSTM models demonstrate commendable performance for the Russian Ruble. Specifically, the IRP-XGB, PPP-LSTM, and the base LSTM models emerge as the most precise in forecasting. However, it should be noted that these models did not capture the magnitude of changes but rather tracked minor fluctuations.

In an extensive examination of the results presented in Table 4.4, several insightful observations arise concerning the effectiveness of different predictive models across distinct currency forecasting contexts. For the Russian Ruble, the XGB models con-

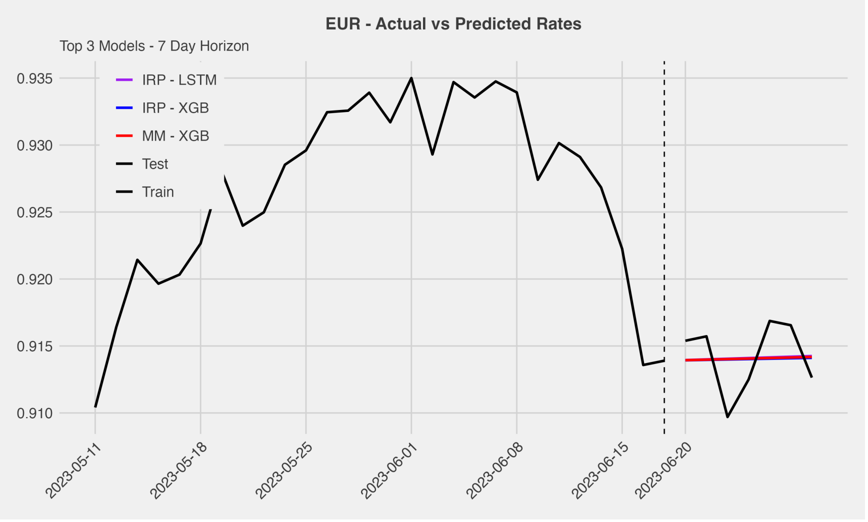

Figure 4.9: Top 3 Models for 7-Day Forecast: Euro

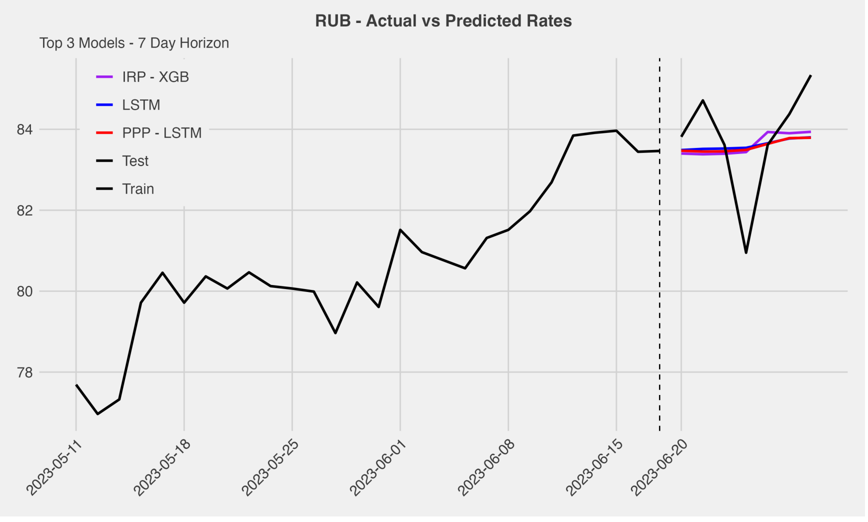

Figure 4.10: Top 3 Models for 7-Day Forecast: Russian Rubble

sistently stand out in terms of their forecasting precision, as evidenced by their lowest RMSE, MAPE, and TheilsU values across IRP, PPP, and MM principles. Notably, within this dominance, the XGB model associated with IRP emerges as the most robust, asserting itself as the preeminent model for a 7-day forecast of RUB.

Turning our attention to the Euro (EUR), XGB models, on the whole, remain impressive, typically registering the lowest RMSE and MAPE scores, which underscores their aptitude for EUR forecasting. Delving deeper into relative accuracy measures, the CNN model informed by IRP principles distinguishes itself as the second best model. ML models, especially XGB and CNN, demonstrate an ability to adapt the Euro better then statistical model like ARIMA. However, it's worth highlighting the surprisingly commendable performance of the Random Walk and Weighted Bootstrap Random Walk models in this context. These relatively simpler models not only demonstrate robustness but, at times, rival and even surpass the predictive prowess of more sophisticated models. This suggests that sometimes the nuanced behavior of a currency like the EUR might be captured by models that capitalize on inherent stochastic tendencies rather than intricate feature-based mechanisms.

However, the narrative shifts when we consider the Turkish Lira. Here, the ARIMA models, and especially the ARIMA informed by MM, rise to prominence. They not only exhibit reduced RMSE scores but also achieve commendable TheilsU1 scores. This performance accentuates the precision and reliability inherent to the ARIMA model in this specific forecasting context. Contrasting this with the patterns observed for RUB and EUR, where ML models were ascendant, it becomes evident that the traditional ARIMA model, especially when complemented with IRP, retains significant efficacy for predicting fluctuations in the TRY. This divergence underscores a pivotal point: the forecasting landscape is complex, and the success of a model often hinges on the specific characteristics intrinsic to a particular currency.

Broadening the scope and reflecting on the entire output, a few overarching themes emerge. The prowess of the XGB model is apparent, especially in the context of RUB and EUR. However, the sterling performance of ARIMA for TRY serves as a strong reminder that model superiority isn't universal. Different currencies, with their unique economic and geopolitical underpinnings, may resonate better with specific models, highlighting the need for a tailored approach in exchange rate forecasting.

Examining the results delineated in 4.4, The introduction of economic models improved the forecasting accuracy of the ARIMA model for all three examined currencies significantly. This enhancement positions the ARIMA model as a notably better performer for the RUB and TRY. Interestingly, while both the LSTM and XGB models exhibited consistent performances in the absence of economic variables, their proficiency did not witness a significant decrese upon the integration of these variables. For the Euro, the introduction of economic models appears to have an adverse impact on the SVR model, leading to a degradation in its forecasting efficacy. The performance of the Random Walk models underscores the indispensable role of benchmark models in exchange rate forecasting. Their exemplary success, especially when predicting the Euro's trajectory and their respectable outcomes for the Ruble, highlights the assertion that, on occasions, simpler models may rival or even eclipse the sophistication of their more complex counterparts. A noteworthy discovery refer to the potential of hybrid models, as exemplified by their results for the Turkish Lira. This underscores an essential insight: compound the strengths of diverse models can, at times, conclude in a forecasting methodology that surpasses its individual components.

Table 4.4: 7-Day Horizon Forecast Performance with Economic Models

Currency	Criteria	ARIMA			XGBoost			SVR			LSTM			CNN			HYBRID			WBRW	RW
		UIRP	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**-**	**-**
TRY	RMSE	0.779	1.398	**0.525**	1.426	1.456	1.416	1.358	0.839	1.186	1.492	1.492	1.492	1.502	1.502	1.496	0.809	1.542	0.995	1.544	1.514
	MAPE	2.499	4.274	**1.856**	4.237	4.272	4.18	4.079	3.15	3.638	4.327	4.327	4.327	4.354	4.354	4.337	2.937	4.978	3.43	4.485	4.39
	Theil's U_1	0.016	0.029	**0.011**	0.029	0.03	0.029	0.028	0.017	0.024	0.031	0.031	0.031	0.031	0.031	0.031	0.016	0.032	0.02	0.032	0.031
	Theil's U_2	1.273	2.281	**0.867**	2.32	2.37	2.306	2.21	1.34	1.932	2.43	2.43	2.43	2.446	2.446	2.437	1.22	2.513	1.551	2.516	2.465
EUR	RMSE	0.003	0.004	0.005	**0.002**	**0.002**	**0.002**	0.019	0.02	0.021	**0.002**	0.011	0.004	0.006	**0.002**	0.003	0.145	0.089	0.028	**0.002**	0.003
	MAPE	0.32	0.407	0.52	**0.247**	**0.247**	**0.247**	1.731	1.885	1.975	**0.247**	1.097	0.39	0.535	**0.247**	0.26	15.827	9.748	3.07	**0.247**	0.249
	Theil's U_1	0.002	0.002	0.003	**0.001**	**0.001**	**0.001**	0.01	0.011	0.012	**0.001**	0.006	0.002	0.003	**0.001**	0.002	0.086	0.051	0.016	**0.001**	**0.001**
	Theil's U_2	1.01	1.163	1.508	0.713	0.713	**0.712**	5.609	6.012	6.249	**0.712**	3.382	1.177	1.667	0.718	0.936	40.027	24.623	7.821	0.716	0.733
RUB	RMSE	3.998	4.007	1.381	**1.223**	1.293	1.303	1.327	1.333	1.346	1.309	1.251	1.296	1.29	1.303	3.38	2.997	5.572	25.181	1.331	1.378
	MAPE	3.909	3.915	1.178	1.14	1.159	**1.083**	1.234	1.249	1.287	1.179	1.111	1.128	1.138	1.161	3.713	2.806	4.558	29.409	1.233	1.25
	Theil's U_1	0.024	0.024	0.008	**0.007**	0.008	0.008	0.008	0.008	0.008	0.008	**0.007**	0.008	0.008	0.008	0.021	0.018	0.034	0.131	0.008	0.008
	Theil's U_2	2.477	2.482	0.844	**0.752**	0.797	0.809	0.818	0.821	0.829	0.807	0.771	0.799	0.797	0.803	1.948	1.857	3.443	15.184	0.82	0.849
AVG	RMSE	1.593	1.803	**0.637**	0.884	0.917	0.907	0.901	0.731	0.851	0.934	0.918	0.931	0.933	0.936	1.626	1.317	2.401	8.735	0.959	0.965
	MAPE	2.243	2.865	**1.185**	1.875	1.893	1.837	2.348	2.095	2.3	1.918	2.178	1.948	2.009	1.921	2.77	7.19	6.428	11.97	1.988	1.963
	Theil's U_1	0.014	0.018	**0.007**	0.012	0.013	0.013	0.015	0.012	0.015	0.013	0.015	0.014	0.014	0.013	0.018	0.04	0.039	0.056	0.014	0.013
	Theil's U_2	1.587	1.975	**1.073**	1.262	1.293	1.276	2.879	2.724	3.003	1.316	2.194	1.469	1.637	1.322	1.774	14.368	10.193	8.185	1.351	1.349

Table 4.5: 7-Day Horizon Forecast Performance using Time Variables Only

Currency	**Criteria**	**ARIMA**	**XGBoost**	**SVR**	**LSTM**	**CNN**	**HYBRID**	**WBRW**	**RW**
TRY	RMSE	1.338	1.466	1.409	1.492	1.492	**1.227**	1.514	1.544
	MAPE	4.08	4.32	4.15	4.327	4.327	**3.539**	4.39	4.485
	Theil's U_1	0.028	0.03	0.029	0.031	0.031	**0.025**	0.031	0.032
	Theil's U_2	2.183	2.385	2.294	2.43	2.431	**2**	2.465	2.516
EUR	RMSE	0.003	**0.002**	0.018	0.003	0.007	0.075	0.003	**0.002**
	MAPE	0.283	**0.247**	1.619	**0.247**	0.662	8.156	0.249	**0.247**
	Theil's U_1	0.002	**0.001**	0.01	**0.001**	0.004	0.043	**0.001**	**0.001**
	Theil's U_2	0.815	**0.714**	5.316	0.802	2.11	20.569	0.733	0.716
RUB	RMSE	4.021	1.483	1.323	**1.254**	1.302	1.778	1.378	1.331
	MAPE	3.933	1.517	1.226	**1.097**	1.2	1.705	1.25	1.233
	Theil's U_1	0.025	0.02	0.021	0.021	0.021	**0.014**	0.022	0.022
	Theil's U_2	5.625	4.259	4.435	4.479	4.521	**3.263**	4.73	4.533
AVG	RMSE	1.458	1.414	1.48	1.476	1.483	**1.155**	1.535	1.497
	MAPE	**1.903**	2.474	3.013	2.642	2.612	5.082	2.635	2.603
	Theil's U_1	0.024	0.009	0.008	**0.007**	0.008	0.011	0.008	0.008
	Theil's U_2	2.491	0.911	0.815	**0.773**	0.803	1.056	0.849	0.82

4.4.3 14 Day Ahead

Once again, the ARIMA model proves to be an optimal choice for forecasting the Turkish Lira. The IRP-ARIMA, PPP-ARIMA, and the base ARIMA model all exhibit patterns that closely align with the realized values. They successfully captured the upward trend, and towards the final days, they also identified the subsequent decline as showcased in Figure 4.8.

In the context of the Euro (EUR), the IRP-CNN, standalone CNN, and MM-LSTM models demonstrated superior performance. Notably, the CNN model, underpinned by the Interest Rate Parity theory, accurately forecasted the final day's value,as evidenced in Figure 4.9. However, while these models largely exhibited a linear trend, the EUR manifested more pronounced fluctuations throughout the period.

For the Russian Ruble (RUB), the IRP-XGB, MM-LSTM, and PP-XGB models emerged as the leading performers, as illustrated in Figure 4.10. These models adeptly captured the upward trend. However, they were not entirely aligned with the abrupt

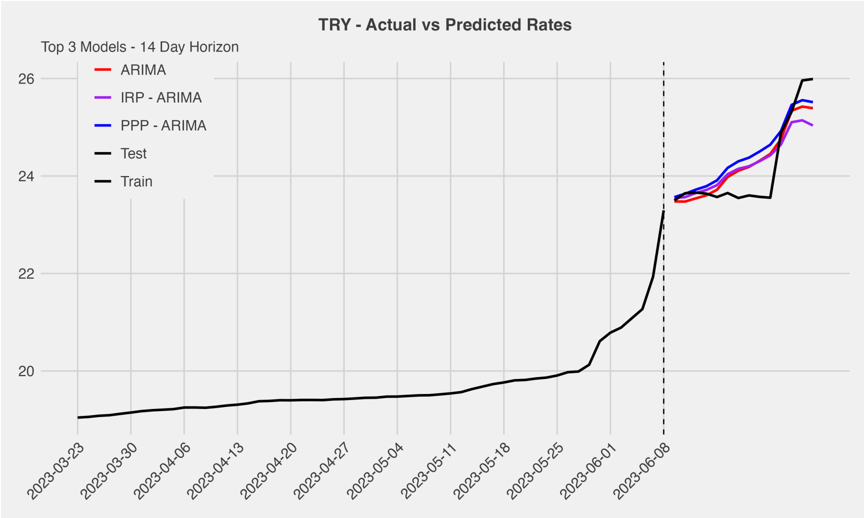

Figure 4.11: Top 3 Models for 14-Day Forecast: Turkish Lira

descent and subsequent recovery observed in the data.

The results delineated in Table 4.6 present forecasting predictions for the currencies RUB, EUR, and TRY over a 14-day horizon. For the RUB currency, the LSTM model, in particular LSTM-MM, demonstrates superior efficacy, recording the lowest RMSE value of 0.997. Interestingly, the MM model enhances the predictive accuracy of the LSTM model for the ruble, especially when compared against the UIRP and PPP models. In the case of EUR, the CNN model integrated with the IRP exhibits the most commendable performance, registering the lowest RMSE of 0.009. While the XGB models, namely IRP-XGB and PPP-XGB, are noteworthy contenders, characterized by low RMSE and MAPE values, the CNN-IRP establishes its preeminence in TheilsU metrics. It's also pertinent to note the sustained competitiveness of the Random Walk models for this 14-day horizon. For the TRY currency, the traditional ARIMA model, specifically ARIMA-IRP, move up as the preeminent model with an RMSE of 0.529. This underscores the notion that, notwithstanding the supremacy of machine learning models in predicting RUB and EUR, the ARIMA model retains its pertinence and efficacy in specific scenarios. Within the purview of the 14-day horizon, while LSTM manifests as the preferred model for RUB and CNN for EUR, ARIMA distinctly outperforms for TRY.

Upon integrating economic variables into the forecasting models, the outcome of ef-

Table 4.6: 14-Day Horizon Forecast Performance with Economic Models

Currency	**Criteria**	**ARIMA**			**XGBoost**			**SVR**			**LSTM**			**CNN**			**HYBRID**			**WBRW**	**RW**
		UIRP	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**-**	**-**
TRY	RMSE	**0.529**	0.537	5.117	1.224	1.184	1.179	0.812	0.637	0.636	1.16	1.159	1.158	1.176	1.177	1.161	0.809	0.845	0.564	1.251	1.201
	MAPE	1.722	**1.714**	20.61	3.279	2.851	3.019	2.307	2.164	2.152	3.026	3.024	3.016	3.095	3.098	3.032	3.05	2.999	1.789	3.404	3.191
	Theil's U_1	**0.011**	**0.011**	0.118	0.026	0.025	0.025	0.017	0.013	0.013	0.024	0.024	0.024	0.025	0.025	0.024	0.017	0.018	0.012	0.026	0.025
	Theil's U_2	**1.248**	1.3	11.85	2.798	2.696	2.691	1.856	1.505	1.496	2.65	2.649	2.646	2.688	2.69	2.654	1.918	1.911	1.369	2.861	2.747
EUR	RMSE	0.021	0.02	0.022	0.016	0.016	0.017	0.026	0.028	0.029	0.017	0.021	0.009	**0.006**	0.013	0.027	0.098	0.078	0.028	0.017	0.017
	MAPE	2.099	1.981	2.109	1.648	1.615	1.69	2.357	2.59	2.718	1.712	2.081	0.849	**0.482**	1.292	2.714	10.67	8.444	2.916	1.657	1.658
	Theil's U_1	0.011	0.011	0.012	0.009	0.009	0.009	0.014	0.015	0.016	0.009	0.011	0.005	**0.003**	0.007	0.015	0.057	0.044	0.016	0.009	0.009
	Theil's U_2	5.732	5.412	5.941	4.461	4.372	4.576	6.943	7.615	7.993	4.636	5.623	2.397	**1.503**	3.503	7.406	25.482	20.168	7.095	4.484	4.467
RUB	RMSE	1.865	1.943	2.511	1.688	1.455	3.2	2.332	2.154	2.487	2.248	2.252	**0.997**	2.109	2.127	3.2	2.487	2.718	6.051	2.296	2.447
	MAPE	1.838	1.919	2.555	1.867	1.565	3.443	2.558	2.37	2.706	2.476	2.477	**0.777**	2.334	2.356	3.604	2.351	2.516	6.666	2.522	2.621
	Theil's U_1	0.011	0.012	0.015	0.01	0.009	0.019	0.014	0.013	0.015	0.014	0.014	**0.006**	0.013	0.013	0.02	0.015	0.016	0.035	0.014	0.015
	Theil's U_2	1.541	1.607	2.05	1.392	1.201	2.643	1.926	1.78	2.055	1.857	1.861	**0.826**	1.742	1.757	2.629	2.063	2.255	4.637	1.897	2.023
AVG	RMSE	0.805	0.833	2.55	0.976	0.885	1.465	1.057	0.94	1.051	1.142	1.144	**0.721**	1.097	1.106	1.463	1.131	1.214	2.214	1.188	1.222
	MAPE	1.886	1.871	8.425	2.265	2.01	2.717	2.407	2.375	2.525	2.405	2.527	**1.547**	1.97	2.249	3.117	5.357	4.653	3.79	2.528	2.49
	Theil's U_1	**0.011**	**0.011**	0.048	0.015	0.014	0.018	0.015	0.014	0.015	0.016	0.016	0.012	0.014	0.015	0.02	0.03	0.026	0.021	0.016	0.016
	Theil's U_2	2.84	2.773	6.614	2.884	2.756	3.303	3.575	3.633	3.848	3.048	3.378	**1.956**	1.978	2.65	4.23	9.821	8.111	4.367	3.081	3.079

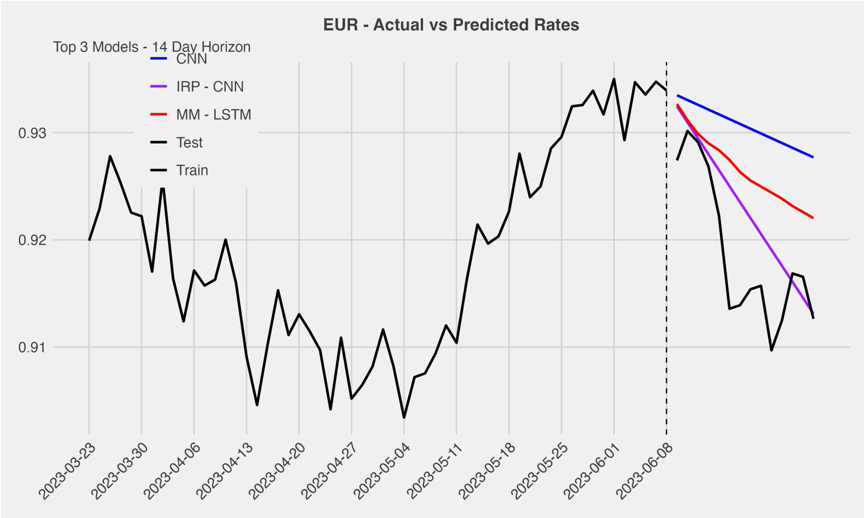

Figure 4.12: Top 3 Models for 14-Day Forecast: Euro

fects varied in nature. For instance, a positive effect in forecasting accuracy was observed in models such as MM-LSTM for the Russian Ruble and IRP-ARIMA for the Turkish Lira. Within the context of RUB forecasting, the ARIMA model emerged preeminent. A thorough assessment across all evaluation metrics positions both the ARIMA and CNN models as superior, suggesting that these models adeptly encapsulate the underlying dynamics of the RUB, especially over a 14-day forecasting horizon.

For the Euro, the CNN model distinguished itself, achieving an exceptionally low RMSE value of 0.006. This performance is notably superior compared to other models, with the next best, the XGB, recording an RMSE of 0.017. Although traditional forecasting models, such as ARIMA and Random Walk, yielded comparable RMSE and TheilsU metrics, they were somewhat eclipsed by the unparalleled accuracy exhibited by the CNN model.

In the context of the Turkish Lira, the ARIMA model showcased its prowess with an RMSE of 0.451, positioning it as the most efficacious forecasting tool. Despite the appeal of more modern algorithms, ARIMA's efficacy cannot be overlooked, especially for the TRY.

In summation, while advanced machine learning models like CNN demonstrate commendable forecasting accuracy for certain currencies, it is imperative to recognize and

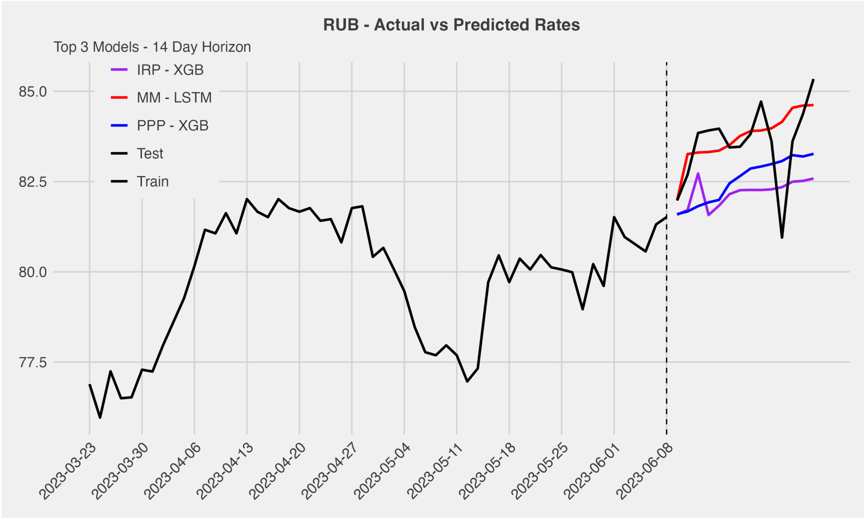

Figure 4.13: Top 3 Models for 14-Day Forecast: Russian Rubble

appreciate the continued relevance and efficacy of traditional models such as ARIMA, the utility of which can vary based on the specific currency and associated contextual nuances.

Table 4.7: 14-Day Horizon Forecast Performance using Time Variables Only

Currency	Criteria	ARIMA	XGBoost	SVR	LSTM	CNN	HYBRID	WBRW	RW
TRY	RMSE	**0.451**	1.236	1.034	1.159	1.16	0.894	1.201	1.251
	MAPE	**1.449**	3.23	2.542	3.023	3.027	3.429	3.191	3.404
	Theil's U_1	**0.009**	0.026	0.022	0.024	0.024	0.018	0.025	0.026
	Theil's U_2	**1.081**	2.823	2.357	2.649	2.651	2.101	2.747	2.861
EUR	RMSE	0.023	0.017	0.025	0.021	**0.006**	0.011	0.017	0.017
	MAPE	2.312	1.655	2.294	2.081	**0.494**	1.104	1.658	1.657
	Theil's U_1	0.013	0.009	0.014	0.011	**0.003**	0.006	0.009	0.009
	Theil's U_2	6.355	4.479	6.76	5.623	**1.545**	2.861	4.467	4.484
RUB	RMSE	**1.821**	2.167	2.512	2.179	2.04	3.099	2.447	2.296
	MAPE	**1.791**	2.347	2.71	2.408	2.257	3.119	2.621	2.522
	Theil's U_1	**0.011**	0.013	0.015	0.013	0.012	0.018	0.015	0.014
	Theil's U_2	**1.504**	1.788	2.076	1.8	1.685	2.435	2.023	1.897
AVG	RMSE	**0.765**	1.14	1.19	1.12	1.069	1.335	1.222	1.188
	MAPE	**1.851**	2.411	2.515	2.504	1.926	2.551	2.49	2.528
	Theil's U_1	**0.011**	0.016	0.017	0.016	0.013	0.014	0.016	0.016
	Theil's U_2	2.98	3.03	3.731	3.357	**1.96**	2.466	3.079	3.081

4.4.4 30 Day Ahead

The Turkish Lira has experienced significant volatility in the test set and interest rate models with ARIMA have been found to be strong forecasters, as illustrated in Figure 4.14. While these models anticipated an increase in the exchange rate, the actual surge in TRY surpassed their projections considerably, escalating from below 20 to 26 within a 30-day span Figure 4.14. The models had forecasted a rate closer to 22 USD/TRY. Notably, all three models exhibited closely aligned predictions.

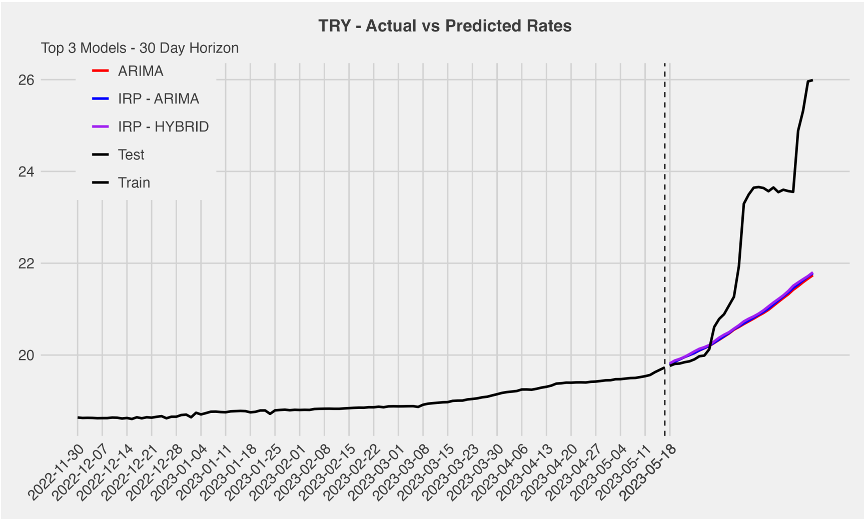

Figure 4.14: Top 3 Models for 30-Day Forecast: Turkish Lira

For the Euro, a diverse set of models emerges as the top performers. The EUR exhibits regular fluctuations, though its variations are less strong compared to other examined currencies, as illustrated in Figure 4.15. The leading three models capturing the essence of EUR's 30-day movement are the IRP-LSTM, PPP-ARIMA, and PPP-XGB. While the LSTM and XGB models displayed trajectories that are relatively linear and closely aligned, the ARIMA model demonstrated greater variability in its predictions, potentially reflecting the currency's inherent volatility to a greater extent.

The Russian Ruble forecasts exhibited the performance with the MM-CNN, PPP-XGB, and foundational XGB models, which successfully captured the prevailing upward trend. Specifically, the XGB models demonstrated an impressive ability to catch the movements that were in close proximity to actual observations, as depicted in Fig-

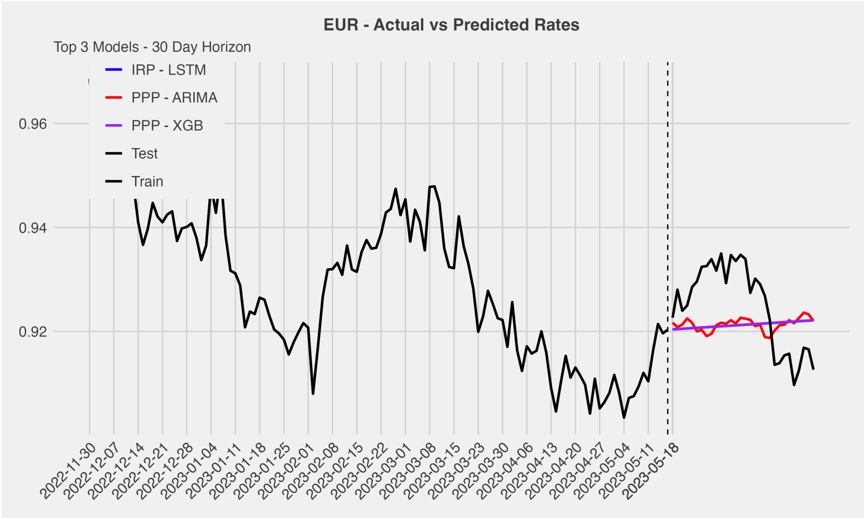

Figure 4.15: Top 3 Models for 30-Day Forecast: Euro

ure 4.16. It is worth noting that the Ruble experienced significant fluctuations during this period, ranging from approximately 80 to slightly above 85, yet our models accounted for these challenges in predicting extreme movements, anticipating a peak value around 83.

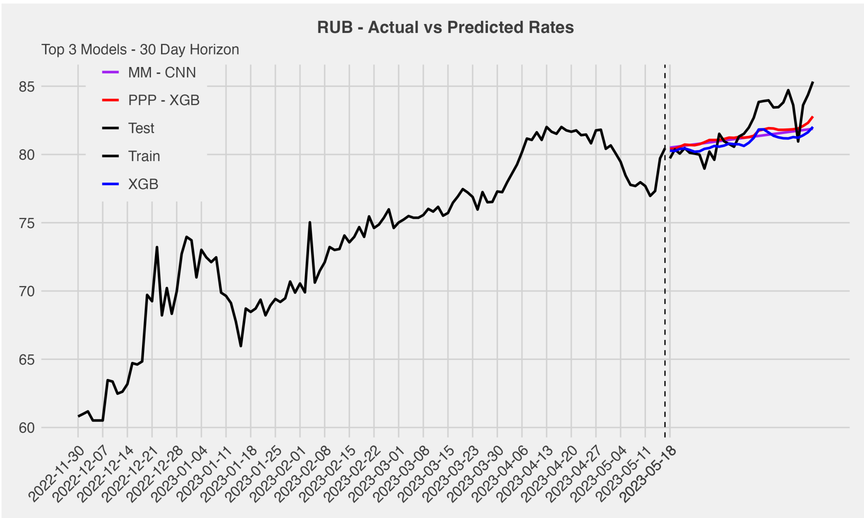

Figure 4.16: Top 3 Models for 30-Day Forecast: Russian Rubble

In the study's 30-day forecasting horizon, the results related to economic models are drawn in Table 4.8. Within the array of models assessed for the Russian Ruble (RUB), the XGB series conspicuously emerges as preeminent. Of particular note, the XGB model, when underpinned by the IRP principle, registers the most small

RMSE and MAPE metrics. Further supports its superior forecasting capability, the XGB models also record the lowest TheilsU1 values, underscoring their unparalleled relative accuracy. This supremacy of the XGB model persists in the context of the Euro, demonstrating its adeptness in forecasting the various trajectories of both the Ruble and the Euro.

Transitioning to the Turkish Lira (TRY), the hybrid of ARIMA-LSTM model asserts itself as the prime forecaster. More specifically, when this hybrid model is paired with the IRP principle, it surfaces as the most efficacious instrument for forecasting the Lira, as evidenced by the furnished metrics. Drawing overarching conclusions from these findings, it can be posited that the XGB model exhibits robust performance over lengthier horizons. Concurrently, the ARIMA model retains its indispensability, especially for the TRY. Furthermore, the IRP principle appears to be an instrumental predictor in this extended timespan.

The performance of various models, based solely on time variables, is shown in Table 4.9. When focusing on the Russian Ruble and assessing it through the prism of RMSE, MAPE, Theil's U_1 and Theil's U_2 metrics, the XGB model unfailingly establishes itself as the dominant forecaster. While both the SVR and the LSTM models present themselves as formidable contenders, underlined by their closely aligned metric values, neither achieves the benchmark set by XGB. Contrarily, the Hybrid model's performance, within this context, leaves much to be desired. Its notably elevated error rate, especially as manifest in the MAPE metric, underscores its limitations in forecasting the Ruble.

Shifting the lens to the Euro, the Random Walk model emerges as the frontrunner across most evaluated metrics, with the sole exception of MAPE. This ascendency might insinuate the Euro's inherent mean-reverting tendencies, underscoring the potential dynamism of its movement. Complementing this narrative, the XGB and LSTM models also stake their claims, showcasing commendable efficacy in capturing the Euro's trajectory.

In the landscape of the Turkish Lira, the respectable ARIMA model carves out its supremacy, particularly when viewed through most of the aforementioned metrics, admitting with an exception in the context of Theil's U_1. The Hybrid model, intrigu-

Table 4.8: 30-Day Horizon Forecast Performance with Economic Models

Currency	Criteria	ARIMA			XGBoost			SVR			LSTM			CNN			HYBRID			WBRW	RW
		UIRP	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**UIRP**	**PPP**	**MM**	**-**	**-**
TRY	RMSE	2.116	2.424	5.058	3.135	2.94	2.905	2.82	3.084	2.903	3.255	3.254	3.254	3.023	3.044	3.253	**2.094**	2.437	4.88	3.216	3.116
	MAPE	6.703	7.676	19.506	10.218	9.428	9.314	8.987	10.028	9.321	10.682	10.681	10.681	9.786	9.865	10.677	**6.646**	7.714	18.439	10.533	10.154
	Theil's U_1	**0.049**	0.057	0.126	0.074	0.069	0.069	0.066	0.073	0.069	0.077	0.077	0.077	0.072	0.072	0.077	**0.049**	0.057	0.121	0.076	0.074
	Theil's U_2	4.816	5.518	11.712	7.15	6.697	6.618	6.42	7.034	6.616	7.428	7.427	7.427	6.893	6.94	7.424	**4.764**	5.549	11.281	7.339	7.109
EUR	RMSE	0.01	**0.009**	0.028	0.01	**0.009**	**0.009**	0.084	0.08	0.086	**0.009**	0.014	0.029	0.023	**0.009**	0.023	0.144	0.097	0.021	**0.009**	**0.009**
	MAPE	0.986	0.914	2.611	**0.913**	0.917	0.916	8.281	7.909	8.417	0.917	1.137	3.023	2.361	0.924	1.777	15.503	10.456	2.108	0.915	0.931
	Theil's U_1	**0.005**	**0.005**	0.015	**0.005**	**0.005**	**0.005**	0.047	0.045	0.048	**0.005**	0.008	0.016	0.013	**0.005**	0.012	0.084	0.055	0.011	**0.005**	**0.005**
	Theil's U_2	2.809	**2.548**	7.6	2.682	2.556	2.561	23.55	22.475	23.971	2.556	3.982	8.217	6.423	2.566	6.531	39.432	26.642	5.662	2.584	2.621
RUB	RMSE	21.647	23.322	222.534	2.15	**1.406**	3.179	2.337	2.368	2.594	2.051	1.971	6.551	1.941	1.9	1.612	25.452	23.278	25.278	2.291	2.443
	MAPE	22.702	24.27	271.921	1.984	**1.408**	2.934	2.135	2.164	2.37	1.875	1.81	6.685	1.788	1.749	1.559	27.057	23.629	27.988	2.092	2.234
	Theil's U_1	0.118	0.126	0.576	0.013	**0.009**	0.02	0.014	0.015	0.016	0.013	0.012	0.041	0.012	0.012	0.01	0.136	0.126	0.135	0.014	0.015
	Theil's U_2	21.869	23.552	224.395	2.155	**1.414**	3.194	2.345	2.376	2.603	2.057	1.976	6.609	1.945	1.904	1.616	25.742	23.485	25.579	2.298	2.451
AVG	RMSE	7.924	8.585	75.873	1.765	**1.452**	2.031	1.747	1.844	1.861	1.772	1.746	3.278	1.662	1.651	1.629	9.23	8.604	10.06	1.839	1.856
	MAPE	10.13	10.953	98.013	4.372	**3.918**	4.388	6.468	6.7	6.703	4.491	4.543	6.796	4.645	4.179	4.671	16.402	13.933	16.178	4.513	4.44
	Theil's U_1	0.057	0.063	0.239	0.031	**0.028**	0.031	0.042	0.044	0.044	0.032	0.032	0.045	0.032	0.03	0.033	0.09	0.079	0.089	0.032	0.031
	Theil's U_2	9.831	10.539	81.236	3.996	**3.556**	4.124	10.772	10.628	11.063	4.014	4.462	7.418	5.087	3.803	5.19	23.313	18.559	14.174	4.074	4.06

ingly, does not lag far behind, with its results paralleling closely to those of ARIMA.

Upon consolidating these insights, a multifaceted understanding emerges: while the XGB model, with its renowned versatility and prowess in diverse predictive tasks, frequently occupies a prominent position, it does not hold a universal monopoly. Different currencies, encapsulating their unique economic and financial intricacies, necessitate tailored forecasting approaches. This accentuates the imperative of an exhaustive model evaluation tailored to each specific currency, ensuring optimal forecasting accuracy.

Table 4.9: 30-Day Horizon Forecast Performance using Time Variables Only

Currency	**Criteria**	**ARIMA**	**XGBoost**	**SVR**	**LSTM**	**CNN**	**HYBRID**	**WBRW**	**RW**
TRY	RMSE	**2.14**	3.031	2.92	3.254	3.254	2.151	3.116	3.216
	MAPE	**6.777**	9.761	9.347	10.681	10.681	6.785	10.154	10.533
	Theil's U_1	**0.05**	0.072	0.069	0.077	0.077	**0.05**	0.074	0.076
	Theil's U_2	**4.87**	6.905	6.652	7.427	7.427	4.896	7.109	7.339
EUR	RMSE	0.011	0.01	0.083	0.014	0.029	0.109	**0.009**	**0.009**
	MAPE	0.993	**0.913**	8.212	1.137	2.973	11.714	0.931	0.915
	Theil's U_1	0.006	**0.005**	0.047	0.008	0.016	0.062	**0.005**	**0.005**
	Theil's U_2	3.07	2.682	23.367	3.982	8.114	29.855	2.621	**2.584**
RUB	RMSE	21.281	**1.642**	2.362	1.943	2.361	23.821	2.443	2.291
	MAPE	22.351	**1.509**	2.159	1.788	2.17	25.317	2.234	2.092
	Theil's U_1	0.116	**0.01**	0.015	0.012	0.015	0.128	0.015	0.014
	Theil's U_2	21.501	**1.647**	2.369	1.947	2.368	24.086	2.451	2.298
AVG	RMSE	7.811	**1.561**	1.788	1.737	1.881	8.694	1.856	1.839
	MAPE	10.04	**4.061**	6.573	4.535	5.275	14.605	4.44	4.513
	Theil's U_1	0.057	**0.029**	0.044	0.032	0.036	0.08	0.031	0.032
	Theil's U_2	9.814	**3.745**	10.796	4.452	5.97	19.612	4.06	4.074

CHAPTER 5

CONCLUSION

In the realm of international finance, forecasting currency behavior presents its own set of complexities. A notable investigation by Meese and Rogoff (1983) unveiled an unexpected finding - the simple Random Walk model, devoid of the complexities of economic theories, might outpace the latter in terms of predictive accuracy for exchange rates. Expanding on this earlier notion, the concept of the ’exchange rate disconnect puzzle’ was subsequently highlighted by Rogoff in 1996. This puzzle underscores a counterintuitive phenomenon: despite being influenced by a plethora of macroeconomic forces, exchange rates often show a fragile connection to these very determinants[53]. Motivated by these insights, our study delves deep into the predictive landscape of exchange rates for three distinct currencies: the Russian Ruble, Euro, and Turkish Lira, spanning time horizons of 3, 7, 14, and 30 days. By using models anchored in economic theories, specifically Uncovered UIRP, PPP, and the MM, against models that eschew these variables, our objective was clear - to analyze the efficacy and relevance of these theoretical constructs in the real-world domain of exchange rate forecasting. As we distill our extensive analysis, this conclusion underscores the key patterns, insights, and takeaways that emerged from our rigorous examination.

At the core of our research lies an in-depth evaluation of economic variables that stem from foundational theories such as UIRP, PPP, and MM. While the integration of these economic principles may pave the way for richer insights, it is crucial to note that their incorporation does not inherently translate to enhanced forecasting precision. Such an observation underscores the complex dance of financial forecasting where a confluence of domain expertise and sophisticated machine learning methodologies doesn't invariably yield superior results. An illustrative case in point is the ARIMA model applied to the TRY: while its forecasting accuracy experienced an increase for 7 and 14-day horizons upon incorporating economic variables, the 30-day forecast saw a decline in its performance. Interestingly, the introduction of these economic variables often led to only marginal variances in results. This finding also in line with Pfahler's (2021) observations, emphasizes that this integration doesn't inherently translate to enhanced forecasting precision[24]. A case to highlight this is the Euro's SVR model, where the introduction of economic facets typically resulted in diminished predictive capabilities. Moreover, it's evident that no single economic framework emerges as universally superior; its efficacy is contingent upon the specific currency, accompanying model, and chosen forecasting horizon. For instance, in the realm of 7-day forecasting for the Ruble, XGB, when paired with UIRP, outperformed other combinations, whereas the LSTM model, synergized with PPP, secured the second spot. However, over a 14-day horizon, the LSTM model, this time harmonized with MM, took the lead.

One of the most prominent takeaways from our study is the absence of a universally optimal model for all scenarios. While the XGB model, celebrated for its adaptability, consistently stood out in predicting exchange rates for the RUB and EUR—particularly over longer forecasting periods, the Turkish Lira showcased a proclivity for more traditional forecasting tools, with ARIMA maintaining a significant upper hand. Such disparity underscores the characteristic intrinsic to each currency, emphasizing that reliance on solely cutting-edge or contemporary models might not always be the best strategy. Rather, the selection of models ought to be grounded in context-specific considerations. The variations in model efficacy across diverse currencies underscore a vital aspect: each currency, attached with its own set of economic, political, and historical complexities, mandates sophisticated forecasting strategies. Given that currencies are shaped by a multiple of global, regional, and domestic forces, it becomes important to deploy models that can adeptly navigate these distinct dynamics. The existing literature on currency price prediction has also yielded mixed results, with some studies favoring traditional time series models like ARIMA (A. Hadjixenophontos & C. Christodoulou, 2017; Babu & Reddy, 2015) while others advocating for machine learning approaches (Shen et al., 2021;

Galeschuk & Mukherjee, 2017a, b; Parot et al 2019). Thus, a nuanced evaluation, tailored to the singular attributes of each currency, remains indispensable.

In the vast landscape of model complexities, the benchmark Random Walk models have exhibited an unexpected tenacity and reliability, at certain cases even outpacing their more sophisticated counterparts, particularly evident in the case of the Euro. This observation aligns with findings from a study conducted by the European Central Bank (ECB), which revealed that the real exchange rates for currencies like the Euro and US Dollar tend to move towards their equilibrium levels over time, as predicted by economic models such as PPP and BEER[26]. Such outcomes suggest that the behavior of the currency might lean more towards inherent stochastic processes rather than being predominantly determined by external features. The proficiency of Random Walk models may also allude to the potential mean-reverting nature of certain currencies, with the Euro standing as a prime example. The Euro's resilience might stem from its representation of a more mature economy, which is relatively resistant to global volatilities. Therefore, in the long run, when come against another currency with robust economic foundations like the USD, the Euro may gravitate back towards its average value. This not only highlights the intricate dynamics of the forex market but also reinforces the indispensable role of Random Walk as a benchmark in currency rate forecasting.

In conclusion, our research has delved deeply into the complex interplay between various forecasting models and their applicability to specific currencies. While certain models, notably XGB and ARIMA, have distinguished themselves in several scenarios, it's paramount to recognize that the domain of foreign exchange rate forecasting remains multifaceted and diverse. As the dynamics of the global economy continually shift and evolve, the quest for the perfect forecasting model will undoubtedly persist. This inherent complexity underscores the necessity of utilizing a broad array of models, constantly assessing their performance, and adjusting strategies in accordance with the ever-changing land of international finance.

REFERENCES

[1] H. N. Bhandari, B. Rimal, N. R. Pokhrel, R. Rimal, K. R. Dahal, and R. K. Khatri, "Predicting stock market index using lstm," *Machine Learning with Applications*, vol. 9, p. 100320, 2022.

[2] A. Hidaka and T. Kurita, "Consecutive dimensionality reduction by canonical correlation analysis for visualization of convolutional neural networks," in *Proceedings of the ISCIE international symposium on stochastic systems theory and its applications*, vol. 2017, pp. 160–167, The ISCIE Symposium on Stochastic Systems Theory and Its Applications, 2017.

[3] S. Galeshchuk and S. Mukherjee, "Deep networks for predicting direction of change in foreign exchange rates," *Intelligent Systems in Accounting, Finance and Management*, vol. 24, no. 4, pp. 100–110, 2017.

[4] G. L. Kaminsky and C. M. Reinhart, "The twin crises: the causes of banking and balance-of-payments problems," *American economic review*, vol. 89, no. 3, pp. 473–500, 1999.

[5] L. Menkhoff and M. P. Taylor, "The obstinate passion of foreign exchange professionals: technical analysis," *Journal of Economic Literature*, vol. 45, no. 4, pp. 936–972, 2007.

[6] R. A. Meese and K. Rogoff, "Empirical exchange rate models of the seventies: Do they fit out of sample?," *Journal of international economics*, vol. 14, no. 1-2, pp. 3–24, 1983.

[7] G. Box and G. Jenkins, *Time Series Analysis: Forecasting and Control*. Holden-Day series in time series analysis and digital processing, Holden-Day, 1970.

[8] H. Drucker, C. Burges, L. Kaufman, A. Smola, and V. Vapnik, "Support vector regression machines," *Adv Neural Inform Process Syst*, vol. 28, pp. 779–784, 01 1997.

[9] T. Chen and C. Guestrin, “Xgboost: A scalable tree boosting system,” in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, 2016.

[10] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[11] A. Borovykh, S. Bohte, and C. W. Oosterlee, “Conditional time series forecasting with convolutional neural networks,” *arXiv preprint arXiv:1703.04691*, 2017.

[12] E. Dave, A. Leonardo, M. Jeanice, and N. Hanafiah, “Forecasting indonesia exports using a hybrid model arima-lstm,” *Procedia Computer Science*, vol. 179, pp. 480–487, 2021.

[13] A. M. Taylor and M. P. Taylor, “The purchasing power parity debate,” *Journal of economic perspectives*, vol. 18, no. 4, pp. 135–158, 2002.

[14] J. A. Frenkel, “Inflation and the formation of expectations,” *Journal of Monetary Economics*, vol. 1, no. 4, pp. 403–421, 1976.

[15] E. F. Fama, “Forward and spot exchange rates,” *Journal of monetary economics*, vol. 14, no. 3, pp. 319–338, 1984.

[16] P. B. Clark and R. MacDonald, “Exchange rates and economic fundamentals: a methodological comparison of beers and feers,” in *Equilibrium exchange rates*, pp. 285–322, Springer, 1999.

[17] J. Williamson, *The political economy of policy reform.* Peterson Institute, 1994.

[18] J. M. Fleming, “Domestic financial policies under fixed and under floating exchange rates (politiques finacierieures interieures avec un systeme de taux de change fixe et avec un systeme de taux de change fluctuant)(politica financiera interna bajo sistemas de tipos de cambio fijos o de tipos de cambio fluctuantes),” *Staff Papers-International Monetary Fund*, pp. 369–380, 1962.

[19] R. A. Mundell, “The appropriate use of monetary and fiscal policy under fixed exchange rates,” *IMF Staff Papers*, vol. 9, no. 1, pp. 70–79, 1962.

[20] J. E. Meade, *The theory of international economic policy*, vol. 2. Oxford University Press, 1951.

[21] M. Mussa, "The exchange rate, the balance of payments and monetary and fiscal policy under a regime of controlled floating," in *Flexible Exchange Rates and Stabilization Policy*, pp. 97–116, Springer, 1977.

[22] C. A. Wilson, "Anticipated shocks and exchange rate dynamics," *Journal of Political Economy*, vol. 87, no. 3, pp. 639–647, 1979.

[23] O. Ince, "Forecasting exchange rates out-of-sample with panel methods and real-time data," *Journal of International Money and Finance*, vol. 43, pp. 1–18, 2014.

[24] J. F. Pfahler, "Exchange rate forecasting with advanced machine learning methods," *Journal of Risk and Financial Management*, vol. 15, no. 1, p. 2, 2021.

[25] D. P. Neghab, M. Cevik, and M. Wahab, "Explaining exchange rate forecasts with macroeconomic fundamentals using interpretive machine learning," *arXiv preprint arXiv:2303.16149*, 2023.

[26] M. Ca'Zorzi, A. Cap, A. Mijakovic, and M. Rubaszek, "The predictive power of equilibrium exchange rate models," *Available at SSRN 3516749*, 2020.

[27] A. Hadjixenophontos and C. Christodoulou-Volos, "Predictability of foreign exchange rates with the ar (1) model," *Journal of Applied Finance & Banking*, vol. 7, no. 4, pp. 39–58, 2017.

[28] A. Babu and S. Reddy, "Exchange rate forecasting using arima," *Neural Network and Fuzzy Neuron, Journal of Stock & Forex Trading*, vol. 4, no. 3, pp. 01–05, 2015.

[29] V. Plakandaras, T. Papadimitriou, and P. Gogas, "Forecasting daily and monthly exchange rates with machine learning techniques," *Journal of Forecasting*, vol. 34, no. 7, pp. 560–573, 2015.

[30] J. Li, I. Tsiakas, and W. Wang, "Predicting exchange rates out of sample: Can economic fundamentals beat the random walk?," *Journal of Financial Econometrics*, vol. 13, no. 2, pp. 293–341, 2015.

[31] M. Yasir, M. Y. Durrani, S. Afzal, M. Maqsood, F. Aadil, I. Mehmood, and S. Rho, "An intelligent event-sentiment-based daily foreign exchange rate forecasting system," *Applied Sciences*, vol. 9, no. 15, p. 2980, 2019.

[32] M.-L. Shen, C.-F. Lee, H.-H. Liu, P.-Y. Chang, and C.-H. Yang, "An effective hybrid approach for forecasting currency exchange rates," *Sustainability*, vol. 13, no. 5, p. 2761, 2021.

[33] A. Parot, K. Michell, and W. D. Kristjanpoller, "Using artificial neural networks to forecast exchange rate, including var-vecm residual analysis and prediction linear combination," *Intelligent Systems in Accounting, Finance and Management*, vol. 26, no. 1, pp. 3–15, 2019.

[34] P. Maneejuk and W. Srichaikul, "Forecasting foreign exchange markets: further evidence using machine learning models," *Soft Computing*, vol. 25, no. 12, pp. 7887–7898, 2021.

[35] X. Ying, "An overview of overfitting and its solutions," in *Journal of physics: Conference series*, vol. 1168, p. 022022, IOP Publishing, 2019.

[36] V. Vapnik, *The nature of statistical learning theory*. Springer science & business media, 1999.

[37] H. Drucker, C. J. Burges, L. Kaufman, A. Smola, and V. Vapnik, "Support vector regression machines," *Advances in neural information processing systems*, vol. 9, 1996.

[38] U. Thissen, R. Van Brakel, A. De Weijer, W. Melssen, and L. Buydens, "Using support vector machines for time series prediction," *Chemometrics and intelligent laboratory systems*, vol. 69, no. 1-2, pp. 35–49, 2003.

[39] Y. Freund, R. E. Schapire, *et al.*, "Experiments with a new boosting algorithm," in *icml*, vol. 96, pp. 148–156, Citeseer, 1996.

[40] P. Li and J.-S. Zhang, "A new hybrid method for china's energy supply security forecasting based on arima and xgboost," *Energies*, vol. 11, no. 7, p. 1687, 2018.

[41] W. Wang, Y. Shi, G. Lyu, and W. Deng, "Electricity consumption prediction

using xgboost based on discrete wavelet transform," *DEStech Transactions on Computer Science and Engineering*, 2017.

[42] Y. Bengio, P. Simard, and P. Frasconi, "Learning long-term dependencies with gradient descent is difficult," *IEEE transactions on neural networks*, vol. 5, no. 2, pp. 157–166, 1994.

[43] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, "Lstm: A search space odyssey," *IEEE transactions on neural networks and learning systems*, vol. 28, no. 10, pp. 2222–2232, 2016.

[44] C. Fjellström, "Long short-term memory neural network for financial time series," in *2022 IEEE International Conference on Big Data (Big Data)*, pp. 3496–3504, IEEE, 2022.

[45] K. Fukushima, "Neocognitron: A hierarchical neural network capable of visual pattern recognition," *Neural networks*, vol. 1, no. 2, pp. 119–130, 1988.

[46] I. E. Livieris, E. Pintelas, and P. Pintelas, "A cnn–lstm model for gold price time-series forecasting," *Neural computing and applications*, vol. 32, pp. 17351–17360, 2020.

[47] G. P. Zhang, "Time series forecasting using a hybrid arima and neural network model," *Neurocomputing*, vol. 50, pp. 159–175, 2003.

[48] S. Cook, "Forecast evaluation using theil's inequality coefficients," *Retrieved October*, vol. 10, p. 2020, 2019.

[49] R. J. Hyndman and G. Athanasopoulos, *Forecasting: principles and practice*. OTexts, 2018.

[50] R. S. Tsay, *Analysis of financial time series*. John wiley & sons, 2005.

[51] D. A. Dickey and W. A. Fuller, "Distribution of the estimators for autoregressive time series with a unit root," *Journal of the American statistical association*, vol. 74, no. 366a, pp. 427–431, 1979.

[52] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.

[53] K. Rogoff, “The purchasing power parity puzzle,” *Journal of Economic literature*, vol. 34, no. 2, pp. 647–668, 1996.