Exchange Rate Risk and Forecasting

Author: Ian Wallgren

May 2022 Bachelor Thesis Department of Economics
Supervisor: Thomas Fischer

Abstract

Since the collapse of the Bretton Woods system, the system of fixed exchange rates amongst principal industrial countries, in the early 1970s, a new era began, introducing the floating exchange rate regime. Since the inception of the floating rate regime, the general interest in forecasting exchange rate movements has grown considerably. Yet, forecasting the exchange rate is extremely difficult, and it is often referred to as an impossible task to execute successfully over a long period of time and the predictions generated by forecasting models are conventionally claimed not to outperform the random walk. However, some researchers have proven to reject these claims by introducing combinations of models to forecast the Forex market, essentially utilizing the synergy between the models rather than applying them as standalone techniques.

The core of this project is designed by introducing a combination of two forecasting models, of which one focuses on traditional economic theory to forecast future values, and the other one on extrapolating historical patterns from the exchange rate data itself into the future. Three different currency pairs are analyzed on a one, two, and three year horizon, and I find promising results in favor of the predictability of the Forex market. More specifically, I find that each model contains information complementary to each other and that they therefore are fit to be applied jointly. Additionally, I present a strategy with the aim of lowering the volatility of returns, essentially decreasing the exchange rate risk involved as the investor is exposed to the currency market.

Keywords: Exchange rate risk, exchange rate forecasting, Autoregressive Integrated Moving Average (ARIMA), Uncovered Interest Rate Parity (UIRP).

Acknowledgements

I would like to express my gratitude to my supervisor, Mr. Thomas Fischer from the Department of Economics at Lund University. The core of this thesis has improved after each of our meetings, and I am very grateful to have had the pleasure of working with you.

I also want to extend my gratitude to the collaborator of this project, Ericsson AB. Specifically, I want to thank Mr. Petter Bernström at Ericsson for sparking an interest in this topic from the start and for generously lending me your time to ensure I had everything I needed to make progress throughout the course of the project.

Contents

1 Introduction

Since the collapse of the Bretton Woods system, the system of fixed exchange rates amongst principal industrial countries, in the early 1970s, a new era began, introducing the floating exchange rate regime (Melvin and Norrbin, 2012). The Forex (FX) market is the largest financial market in the world, trading trillions of dollars each day, and its decentralized marketplace is open for commerce around the clock, five days a week (Kumar, 2014).

From the introduction of the floating rate regime, the general interest in forecasting exchange rate movements has grown considerably (Lam et al., 2008). In pursuit of greater returns on investment, traders, portfolio managers, and many others continuously try to design strategies that would outperform the market. However, should the infamous efficient market hypothesis hold true, investors are suggested not to be able to generate above abnormal returns using only publicly available information (Fama, 1970). By that principle, it follows that it should not be possible to successfully and consistently forecast the exchange rate of currencies. In fact, in a pioneering study performed by Meese and Rogoff in 1983, aiming to research a number of different forecasting models at a short-to-medium term time horizon, it was concluded that all models were outperformed by the random walk (Chang and Hamori, 2020). This was also supported by Mussa in 1979, who concluded that most changes in exchange rates are unexpected and are therefore impossible to predict (Lam et al., 2008). Such findings imply that the exchange rate is more accurately predicted should the current spot rate be used in the place of traditional economic models (Frankel, 1993).

However, some researchers have proven to reject this hypothesis, showing evidence in favor of the superior predictability power some models, or combinations of models, have over the random walk (Chang and Hamori, 2020). In algorithmic trading, for example, state-of-the-art machine-learning models learn market patterns and behaviors by examining historical FX data, essentially generating signals for when to enter or exit the market (Loh et al., 2022). Empirical results from these studies often suggest that combined methods are more likely to outperform its benchmark as compared to relying on a single model (Lam et al., 2008).

Following the financial crisis of 2008, companies around the world have grown more inclined towards risk management. Yet, this is scarcely reflected in changes to formal procedures, apart from players within the financial industry and other companies of which risk management failure has been causing serious damage in the recent past (OECD, 2014). Naturally, exchange rate risk applies especially to companies whose operations partly or entirely lie across national borders and face unpredictable profits and losses due to currency fluctuations (Chen, 2021).

Ericsson AB, one of the world's leading information and communications technology companies, is a prime example of such a company. Globally, as of 2020, Ericsson recorded 232 billion SEK in net sales and had customers in 180 different countries (Ericsson, 2020). Undoubtedly, Ericsson faces challenges due to its significant currency exposure. Motivated by this situation, Ericsson granted the author of this report an opportunity to research ways of how the company might utilize different risk mitigation techniques with the purpose of decreasing its unhedged risk in the FX market.

As such, this paper aims at investigating ways to mitigate exchange rate risk by constructing a forecasting model. Specifically, I examine a traditional economic model, Uncovered Interest Rate Parity (UIRP), and a more sophisticated statistical autoregressive method, ARIMA, and their ability to forecast future exchange rates on a time horizon of one, two, and three years. The currency pairs upon which to conduct analysis are selected based on their relative frequency in Ericsson's historical transactions and constitute of: USD-GBP, USD-JPY, and USD-CNY.

I find promising results in favor of the predictability of the Forex market. Firstly, I conduct thorough analysis on each predictor separately in order to examine how well they manage to forecast each currency on a three year horizon as standalone methods. I find that ARIMA appears to outperform UIRP slightly, however, there is no apparent superiority visible as analysis for all currencies is compiled. Secondly, I find that each model contains information complementary to each other and that they therefore are fit to be applied jointly. Lastly, I present a strategy with the aim of lowering the volatility of returns, essentially decreasing the exchange rate risk involved as the investor is exposed to the currency market.

1.1 Thesis outline

The remained of this paper is organised as follows. The reader is first introduced to a more thorough explanation of the topic of exchange rate risk, primarily focusing on what we will come to know as transaction exposure. Then, the reader will be presented previous research findings which will lead up to the theoretical background of the methodological framework of this dissertation. This is followed by a brief presentation of data collection and processing.

Following this is a background on preliminaries regarding methodology, with particular emphasis on the ARIMA model, accompanied by an exposition of research methodology of the paper in general. Lastly, results are presented and analysis is conducted, and the paper is then wrapped up by a section of discussion aiming to identify potential issues with the models as well as presenting areas subject to further research.

2 Exchange rate risk

Exchange rate risk, also referred to as foreign exchange risk or simply currency risk, arises from unforeseen fluctuations in the relative price, i.e. the exchange rate, between two currencies (Kumar, 2014). As the future market to a strong degree is unpredictable, market participants that are in one way or another exposed to a foreign currency are facing an unavoidable risk, supposedly greater than often referred to. In fact, a survey including 200 CFOs as well as almost 300 treasurers showed evidence of the magnitude of the problem caused by exchange rate risk: According to 70% of the CFOs, their company had in the prior two years suffered reduced earnings caused by unhedged currency exposures, and 51% of CFOs of larger corporations claimed exchange risk to be the risk they were least well-equipped to handle (Nordea, 2021).

However, the exchange rate risk following the acquisition of a foreign currency does not merely depend on the uncertainty of its future value. Instead, in accordance with the rationale of the theory of portfolio diversification, the given risk accrued by the acquisition of the currency also depends on the quantities of the currency the investor already holds, and the correlation matrix of returns between this currency and other assets in the portfolio (Offiong et al., 2016).

Clearly, any type of exchange rate exposure a company might have, is not necessarily deemed to result in a loss. Depending on the position acquired, i.e. long or short, of the company in a certain currency and how this currency fluctuates, i.e. up or down, will eventually determine the exchange rate gain/loss. However, based on the assumption that investors are risk averse and generally require a risk premium to hold assets with uncertain future values (Thune, 2021), exchange rate fluctuations are not desirable.

The risk is often further dissected into three main components: translation exposure, transaction exposure, and operating exposure (McCarthy, 2016). For this thesis, only transaction exposure is relevant and will therefore be explained more in detail below.

2.1 Transaction exposure

Transaction exposure is expressed as the amount to which realized domestic currency values of cash flows denominated in foregein currencies are sensitive in regards to unanticipated fluctuations in the exchange rate. In other words, transaction exposure corresponds to the value of the total amount of outstanding financial obligations that were incurred to the at the time given exchange rate but due to be settled post a change in this rate (Eun and Resnick, 2007).

Currency exposure, SEK billion

Exposure currency	Sales translation exposure	Sales transaction exposure	Sales net exposure	Net external trans-action exposure [1]	Net internal transaction exposure [2]
USD [3]	76.7	36.4	113.1	−12.3	39.6
EUR	26.0	9.5	35.5	12.5	−6.4
CNY	15.3	−0.1	15.2	3.0	−8.8
INR	7.0	−0.2	6.8	0.2	−3.3
AUD	9.0	−0.5	8.5	0.2	4.3
JPY	11.8	–	11.8	0.1	9.1
BRL	2.9	–	2.9	1.0	−0.5
SAR	7.0	1.5	8.5	1.6	2.0
GBP	6.1	−0.7	5.4	−1.0	2.2

1) Net external sales and purchases in foreign currency.
2) Internal sales and purchases in foreign currency.
3) Sales transaction exposure in 2020 includes volume in the cash flow hedge of USD 517 million. Based on the outstanding cash flow hedge volume at year end, the hedged sales volume that will occur in 2021 is USD 200 million.

Figure 1: Currency exposure for Ericsson, extracted from the financial report 2020.

According to Ericsson's financial report of 2020, material cash in- and outflows are only hedged if such are highly certain. Figure 1 reveals a large translation exposure, followed by a bit more shy transaction exposure. Arguably, the transaction exposure Ericsson acquires in different currencies ought to impact returns.

However, in Ericsson's case, it is arguably easier said than done to tackle the problem of transaction exposure. To the author's knowledge, a substantial hurdle in the way of rectifying this problem is the fact that only a fraction of the outgoing tenders end up a signed deal. A tender, once it is shared with its counterpart, might be adjusted and updated back and forth, essentially extending the process with weeks, months or even years. Obviously, not being sure as to exactly when payments are due or when the deal is expected to be signed might make it difficult to undertake hedging actions with the purpose of eliminating any currency risk. Possibly, this is the reason why such a small portion of Ericsson's outstanding currency exposure is currently hedged.

However, the process of eliminating currency exposure using hedging tools might be simplified. The currencies in which Ericsson is expected to make purchases in should the contract be signed can either appreciate or depreciate in relation to the base currency used by Ericsson. Clearly, it would be foolish to hedge against a currency which is highly certain to depreciate against the base currency, since such development would essentially make the purchase less costly. Analogously, it might be a wise decision to hedge against a currency which is highly certain to appreciate in value. Therefore, should it be possible to identify whether a currency is highly likely to either appreciate or depreciate in the future, Ericsson could optimise its hedging actions based on such forecast and potentially decrease volatility of exchange rate returns.

3 Theory

3.1 Previous research

Despite the evidence in favour of the efficient market hypothesis (Malkiel, 2003), some market participants continue to rely on models ahead of the spot rate and a proportion of these have turned out to be somewhat successful, for example by including lagged values of current spot rate in their models (Frankel, 1993). In recent times, combinations of conventional economic models with a lot more sophisticated models, rooted within the realm of machine learning, was found to outperform the predictability power of a classic random walk as short-term exchange rates across 12 major currencies were tested (Amat et al., 2018). In another research, where more advanced machine learning methodologies using random forest methods, neural networks and support vectors machines, were combined with fundamental models, such as the Power Purchasing Parity, attempting to outcompete the classic random walk, some results were also deemed successful (Chang and Hamori, 2020).

Forecasting the exchange rate (especially in the short term) solely using information in the form of rates of inflation, income levels, money supplies, etc., successfully and consistently is an extremely difficult, if not impossible, task. It is argued that the current spot rate contains a lot more useful information than any other macroeconomic variables would contain combined. For example, it is described that the monetary model fails to incorporate information about unobserved elements, such as speculative bubbles. Such elements might therefore better be modeled by a lagged version of the spot rate (Frankel, 1993).

Nonetheless, traditional economic models seem to contain complementing information to the spot rate. Thus, it is argued that a superior way of predicting the future rates is therefore a combination of models, rather than using them as standalone methods (Altavilla and De Grauwe, 2006). Hence, based on the literature review of previous research, the methodology of this thesis will be two-piece. One half will be devoted to a more conventional economic model, aiming to explain future movements in the exchange rate based on macroeconomic variables, whereas the other half will focus on the time series regression model Autoregressive Integrated Moving Avergae (ARIMA), attempting to extrapolate observed historical patterns of the currency itself into the future.

To the author's knowledge, the research conducted in this project differs across a few areas from what has been done before. Firstly, none of the above mentioned reports utilised ARIMA as a predictor. Secondly, the construction of such a forecasting method as this research presents as a combination of ARIMA and more fundamental economic theory has not been found in existing literature. Also, as the content of this research to some degree is constructed for the results to be applicable for the collaborating firm Ericsson, strategies will reflect this and the assumptions for backtesting such strategies will be unique for this project. Additionally, both the composition of currencies used and the selected time horizons of the forecasts differ from what has been researched previously. Lastly, this report will contain an analysis section with the purpose of examining different versions of the methods applied and how the forecasted values react to such changes; partly by varying the sample size used to train the model, partly by altering the model's initial parameters.

3.2 Theoretical background

There exist evidence in favor of the predictability of future exchange rates yielded by a combination of models as presented earlier, where the fundamental parts heavily rely on incorporating information about what, theoretically, determines the exchange rate. Hence, the following section will briefly examine the theory of monetary forces and their impact on the market.

First of all, it should be clarified that this paper solely focuses on floating exchange rates post the collapse of the Bretton Woods system. Hence, the exchange rate is the relative price of two currencies and the fundamentals that essentially determine it are supply and demand for these currencies. However, the supply and demand for a specific currency is in turn affected by the public perception of a fair exchange rate, which is determined by several factors of both technical and fundamental nature, inflation rate and money supply (Banton, 2021). An unexpected change, or shock, in any of these factors will subsequently affect the rate of supply and demand of the currency, potentially causing fluctuations in the exchange rate (Leiva-Leon et al., 2020). With this background, it is desired to understand how these factors can be measured and if they display any correlation with the exchange rate. More specifically, this thesis will focus on the differential in the risk-free interest rates between two arbitrary economies and to what degree such relationship is able to explain the determination of the future exchange rate between those economies.

3.3 Uncovered Interest Rate Parity

Uncovered Interest Rate Parity (UIRP) is an economic theory stating that the difference in interest rate between two economies equates to the relative change in exchange rate over the same period. The UIRP is expressed as follows:

$$F_0 = S_0 \cdot \frac{1 + i_B}{1 + i_A}. \tag{1}$$

In the above equation, F_0 is the expected future exchange rate, S_0 is the spot rate, and $1 + i_A$ and $1 + i_B$ are the risk-free interest rates in country A and B, respectively (Hayes, 2021b).

The rationale behind why such a relationship is expected to hold can be demonstrated as follows. By assumption, the exchange rate between two arbitrary countries, A and B, is one-to-one, meaning that one unit of country A's currency is in the market exchanged for one unit of country B's currency. Suppose country B offers a risk-free interest rate of 10 percent, whereas the risk-free rate in country A is 5 percent. Clearly, investors would be better of should they hold all their money in the currency of country B, as this would enable an investment that generates 5 percentage points more than in country A. However, as previously explained, as the demand for country B's currency rises, so does the exchange rate. The currency of country B will grow more expensive in relation to the currency of country A. In turn, such development will decrease the demand for the currency of country B. This process will proceed until an equilibrium price in the market is achieved (Hayes, 2021b) (Frankel, 1993).

A recent study showed evidence in support of the UIRP when both the country of origin and destination are high-income economies (Orellana and Pino, 2021). In the short- to medium term, however, it is sometimes claimed that UIRP does not hold. Instead, empirical evidence has shown that the higher yielding currency often strengthens rather than weakens (Hayes, 2021b). Other research findings contradict this claim, asserting the mixed evidence for UIRP to be currency-dependent rather than horizon-dependent (Bekaert et al., 2007).

3.4 ARIMA model

In the short run, the idea that all publicly available information ought to be incorporated in the exchange rate at all times is not supported due to the phenomenon of price stickiness (Hayes, 2021a). Price stickiness is the notion that the market fails to adjust its prices quickly in accordance with discrete shifts in the broad economy, such as changes in the money supply (Gali, 2010). In practice, this implies that instead of an instantaneous change in prices of goods, the equilibrium will be attained through changes in the financial market prices (Hayes, 2021a). Therefore, any disturbance in the nominal exchange rate, which in a world of almost perfect capital mobility occurs frequently, will be reflected in the real exchange rate and die out only very slowly over time (Frankel, 1993).

As such, part of the forecasting conducted in this research will be based on regression analysis of the exchange rate itself. For regression based methods in general, time is used as an explanatory variable whereas the time series value is the dependent variable; future values of the time series are forecasted based on historical and present observations (Kotu and Deshpande, 2015). Autoregressive Integrated Moving Average, or ARIMA, is a time series forecasting methodology developed by Box and Jenkins in 1970 (Box and Jenkins, 1970), and has since become a popular model applied in time series forecasting. In practice, forecasts are obtained by extrapolating information of different trends in the data into the future.

The ARIMA model is constructed by dissecting the statistical attributes of the data and analysing these accordingly. More specifically, the model consists of the following components;

- Autoregression (AR)
- Integrated (I)
- Moving Average (MA)

The Autoregression (AR) terms indicate that the data points in the sample are regressed on its own lagged values. The Integrated (I) terms measure the number of times the data set is required to be differenced in order to attain stationarity. A stationary time series is necessary since the ARIMA is a linear regression model using its own lags as predictors, where it is important that the predictors are uncorrelated and independent of each other (Prabhakaran, 2021). Lastly, the Moving Average (MA) terms explain how observations depend on errors from the model's lagged observations (Bora, 2021).

The model, commonly denoted as ARIMA(p, d, q), consists of three parameters which are derived as follows:

- p - number of autoregressive terms.
- d - number of differences required to achieve stationarity.
- q - number of lagged forecast errors included in the prediction.

Mathematically, the equation of ARIMA can be expressed as

$$\hat{y}_t = \mu + \phi_1 y_{t-1} + ... + \phi_p y_{t-p} - \theta_1 e_{t-1} - ... - \theta_q e_{t-q}, \tag{2}$$

where the moving average parameters, following the convention introduced by Box and Jenkins, are defined to have a negative sign (Prabhakaran, 2021). In the equation, $\hat{y}_t$ is the variable to be predicted, ϕ are the autoregressive parameters, θ are the moving average parameters, and e are the error terms (Bora, 2021).

4 Data

4.1 Data description

As this project originated in consultation with telecom-company Ericsson, the exchange rates on which to deploy the forecasting on will be selected based on their relative frequency in the company's historical transactions. Also, as part of the forecasting will be achieved by utilising bond data restricted to country level, it is important not to include any currency adopted in multiple countries.

The list of currencies being analysed in this report looks as follows:

- Great Brittish Pound (GBP),
- Japanese Yen (JPY),
- Chinese Yuan (CNY).

The base currency of Ericsson (one of two, in fact) is United States Dollar (USD), and is therefore the obvious choice of base currency in this project as well. The development of each currency in relation to the USD between 2009 and 2016 is visible in Figure 2. The reader should be informed of the fact that the exchange USD-SEK has been excluded from further analysis due to restricted access to government bond data for Sweden (however, Figure 2 and Figure 15-17 in Appendix contain information of this exchange rate and has not been excluded due to technical issues).

Government bond data is publicly available, provided by Investing.com (a publicly accepted provider of financial data), with a resolution of daily basis ranging from 2009/01/01 - 2019/01/01. As the forecasting horizons for this report are 1, 2 and 3 years, these periods are the obvious choices for the bonds' yields to maturity as well. Due to the fact that the bond data extracted differed in length across countries, a problem which originates from the fact that commerce days are not identical across the markets, all data is re-sampled to a monthly basis. See Appendix Figure 8 for a visual demonstration of historical bond prices used.

Exchange rates are downloaded from Yahoo Finance, directly through a built-in Python module. The data of the exchange rates is extracted on a daily basis, also ranging from 2009/01/01 - 2019/01/01. In order to reduce some of the noise embedded in the data, exchange rates are re-sampled to a monthly basis. This alteration is also required to ensure consistency between the exchange rate data and bond data.

In order to be able to benchmark the performance of the models constructed in this research, data points occurring past 2016/01/01 will be excluded during the process of generating predictions. Instead, these data points will be used to determine the measure of fit after such predictions have been extrapolated.

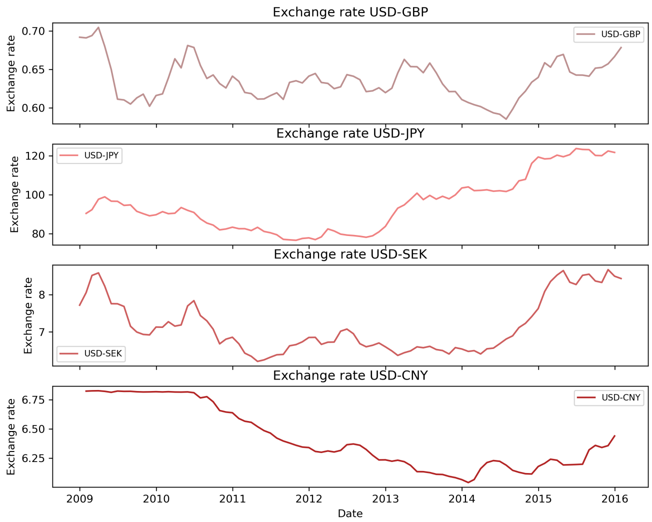

Figure 2: Exchange rates with base currency USD for GBP, JPY, SEK, and CNY between 2009-2016. Exchange rates for all periods are averaged on a monthly basis.

4.2 Data processing

Primarily, programming language Python has been utilised to construct ARIMA models as well as to backtest the different strategies presented. The rationale behind this choice is the vast selection of publicly available libraries for data handling. Mainly, the Python librabries utilised throughout this project are Numpy and Pandas. However, additional libraries for plotting, such as Matplotlib and Seaborn, have been used as well. Most of the scripts used to run the data are available on the author's Github repository [1], and can be provided upon request.

[1] https://github.com/ianwallgren

5 Methodology

As previously described, the forecasting conducted in the project will be dependent on both the historical time series of the exchange rate itself, by utilising the ARIMA model, and on a macroeconomic indicator in UIRP. Hence, this section will be divided into four parts, one which briefly describes the necessary steps to implement the ARIMA model in practice, one which describes how the forecasting by utilising the UIRP will be conducted, and eventually two short sections illustrating different strategies as combinations of these models and how such strategies are backtested.

5.1 ARIMA Forecasting

This section about ARIMA can be elaborated extensively, however, in order to keep it concise and readable, merely the main points will be covered, attempting to describe how the model is set up. These points are illustrated in the flow chart in Figure 3.

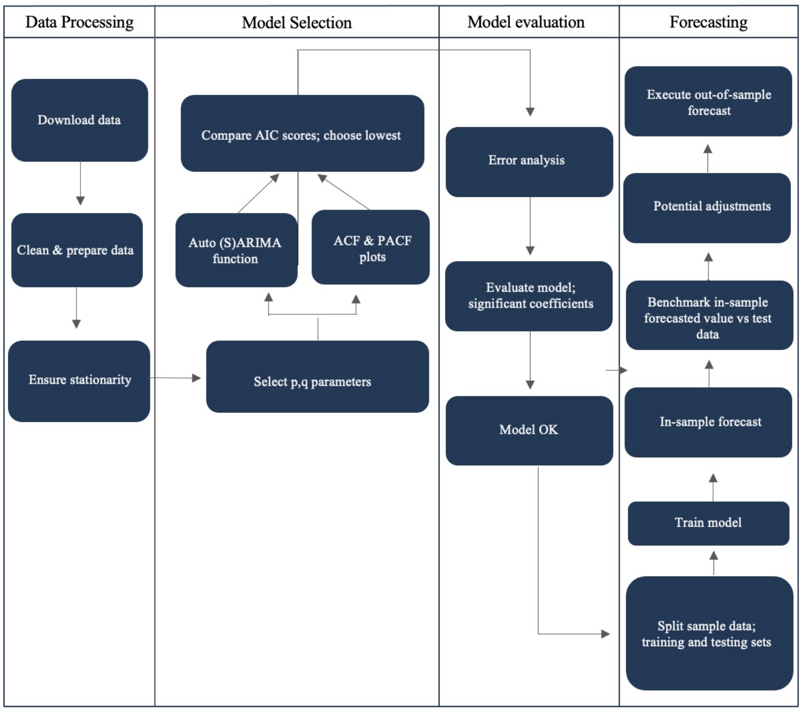

Figure 3: Flowchart describing ARIMA from data processing to model forecasting.

5.1.1 Data processsing

Firstly, desired data is extracted from Yahoo Finance and stored in Python. Next, as a necessary condition for the ARIMA model to be applied is a stationary time series to begin with, one can utilise the technique of differencing, and the different degrees of differencing are visualised below:

- $d = 0 : y_t = Y_t$
- $d = 1 : y_t = Y_t - Y_{t-1}$
- $d = 2 : y_t = Y_t - Y_{t-1} - (Y_{t-1} - Y_{t-2}) = Y_t - 2 \cdot Y_{t-1} + Y_{t-2}$

and so on.

A statistical Dickey Fuller test is used and will indicate the degrees of differencing necessary to achieve stationarity. However, merely differencing the time series does not take into account whether there is an element of seasonality in the data. Therefore, should the time series display seasonal hikes or drops after normal differencing, the issue can be rectified by applying seasonal differencing. If so, the ARIMA model is denoted SARIMA$(p, d, q)(P, D, Q)_m$, with the latter part indicating the seasonal parameters and m the length during which the seasonality is present (Prabhakaran, 2021).

As the time series is converted to be stationary, with d indicating the number of times the series was required to be differenced, the initial conditions are met and the parameters p and q for the ARIMA model can be selected.

5.1.2 Model selection

To achieve the best possible model, p and q will be selected by utilising two different techniques, of which the pair which produces the lowest Akaike Information Criterion (AIC)[2]. score will be selected, as the parameters yielding the lowest score is estimated to be the best fit for the data (Bevans, 2020). It is also important to test for the statistical significance of these parameters once they have been selected. For this step, the entire sample data is used, i.e. the data set stretching between 2009-2016.

One technique is to read the plot of the autocorrelation function (ACF) and partial autocorrelation function (PACF) of the time series. The other technique utilises a

[2] AIC is a common statistical score to evaluate how well a model fits the underlying data by measuring the degree to which variation is explained using the fewest number of independent variables (Bevans, 2020)

Exchange rate	Original
USD-GBP	(1,1,0)(0,1,1)[12]
USD-JPY	(1,1,1)(0,1,1)[12]
USD-CNY	(0,1,1)(0,1,2)[12]

Table 1: (S)ARIMA parameters for each currency pair.

built-in method.

5.1.3 Model evaluation

When the optimal parameters are chosen, the model is evaluated using Residual Sum of Squares (RSS) and by plotting the distribution of residual errors, which ought to be normally distributed with a uniform variance (Prabhakaran, 2021).

Next, the model predictions are scaled back to its original format and compared against the actual values of the time series and the root mean squared error (RMSE) is calculated as a measure of fit. Provided that the RMSE is low enough (there is no universal threshold for this measure, instead one has to determine it subjectively), the forecasting procedure may begin. If the model is not OK (see flowchart in Figure 3), we loop back to the stage of selecting the p and q parameters.

For the currency pairs analysed in this research, the following parameters were selected (see Table 1):

As explained in Section 5.1.1, for each combination of parameters depicted in Table 1, the first tuple corresponds to the ordinary ARIMA paramerters, followed by the second tuple which corresponds to the seasonal parameters and [12] to indicate the length during which the seasonality is present (since monthly resolution is used, [12] implies seasonality cycle lasts one year).

5.1.4 Forecasting

The data is then split into two different sets; a training set and a testing set. The training set will be used to train the model, whereas the testing set will be used to analyse the in-sample performance of the predictions generated by the training set. This process is an efficient way of detecting any obvious errors and correct them prior to the out-of-sample forecasting. If adjustments are made, we loop back to the stage where the model is trained. Finally, should the in-sample results have proven to be satisfactory, out-of-sample forecasting (for this research, out-of-sample corresponds to the period 2016-01-01 - 2019-01-01) is conducted. The forecasting is performed using a built-in forecasting method. The steps are visualised in the graph below.

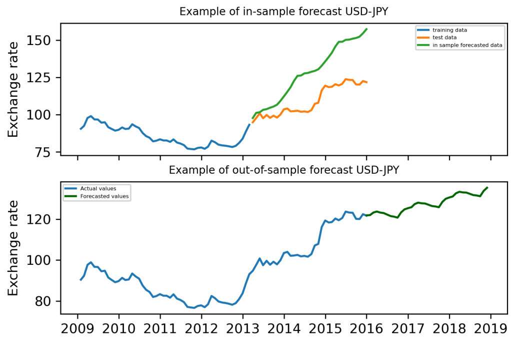

Figure 4: Example of in- and out-of-sample forecasting for USD-JPY.

5.2 UIRP Forecasting

Firstly, desired government bond data is downloaded from Investing.com. As previously described, data sets downloaded will consist of bond data with yields to maturity of 1−, 2−, and 3 years for USA, Great Britain, Japan, and China, ranging from 2009-2016.

For each time horizon, expected future exchange rate for the specific currency will be calculated by utilising Equation (1) presented previously.

5.3 Strategies

Based on the two forecasting methods presented above, five strategies will be constructed as combinations of these and backtested accordingly. These are presented below.

- Strategy 1: always enter a forward contract.
- Strategy 2: enter a forward contract if ARIMA prediction suggests the exchange rate to depreciate.
- Strategy 3: enter a forward contract if UIRP prediction suggests the exchange rate to depreciate.
- Strategy 4: enter a forward contract if either the ARIMA prediction or the UIRP prediction suggest the exchange rate to depreciate.
- Strategy 5: enter a forward contract if both the ARIMA prediction and the UIRP prediction suggest the exchange rate to depreciate.

Note that since Ericsson makes purchases by exchanging USD (base currency) for the purchasing currency, a depreciation in such exchange rate would imply that the specific currency will grow more expensive in relation to the base currency and hence, should no forward contract be entered, an exchange rate loss will occur. Analogously, should the exchange rate appreciate, an exchange rate gain will be generated.

5.4 Benchmarking

To benchmark each strategy, it will be assumed that the date of the outgoing tender is 2016/01/01, and that purchases in GBP, JPY, and CNY are expected to be made in $1-$, $2-$, and 3 years. It is also assumed that Ericsson holds enough USD at all times and will be required to exchange USD for GBP, JPY and CNY when payment is due.

An additional assumption made is that Ericsson may use today's spot rate, for each currency used for purchases in the outgoing tender, when entering contracts with subcontractors. This will be analogous to entering a forward contract for the specific currency pair without the inclusion of any transaction costs, i.e. the forward rate will equate to today's spot rate. Therefore, it should already be obvious for the reader that Strategy 1 will yield an exchange rate gain/loss of 0.

To benchmark the measure of fit between each model's predictions and the actual data points (data points between 2016-01-01 - 2019-01-01, that is), a variety of well-known evaluation metrics will be used. This will be presented more in detail in Section 6.1.

6 Results

In order to improve the reading experience and make the results more comprehensible, this section will be divided into two sections. The first section will briefly visualize the performance of the predictors as standalone methods. The reader is in this part introduced to evaluation metrics with the intention to compare the two predictors in terms of how well, for each data point forecasted, each one were able to forecast the different exchange rates. The second section will present a table of statistics and address the performance of the strategies presented above.

6.1 Evaluation metrics of predictors

For each currency, ARIMA and UIRP predictors will be evaluated based on the Mean Absolute Percentage Error (MAPE) and Root Mean Squared Error (RMSE). It is important for the reader to note that the RMSE should not be compared across currencies, as the nominal values of the currencies differ a lot and will therefore heavily impact the result of the metric. This should become quite obvious to the reader as the construction of the metric is inspected more in detail below.

$$MAPE = \frac{1}{n}\sum_{t=1}^{n}\frac{|e_t|}{|A_t|}. \tag{3}$$

In Equation (3), $|e_t|$ is the absolute difference between the actual value and the forecasted value at time t, whereas $|A_t|$ denotes the absolute value of the actual value and n the number of data points predicted. MAPE, due to the advantage of scale-independency it provides, is regarded a good measure of forecast accuracy (Kim and Kim, 2016). Still, one shall be aware that the MAPE score depends much on the data sample and the use case, but in general the rule of thumb for classifying MAPE scores is as follows (Allwright, 2021):

- $MAPE < 0.1 \rightarrow$ Very good.
- $0.1 < MAPE < 0.2 \rightarrow$ Good.
- $0.2 < MAPE < 0.5 \rightarrow$ OK.
- $MAPE > 0.5 \rightarrow$ Not good.

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}e_t^2}. \tag{4}$$

In Equation (4), e_t^2 denotes the squared value of the error at time t, and n the number of data points predicted. The rationale behind the choice of RMSE as an evaluation metric is that this measure does not treat all errors equally, instead it puts a greater emphasize on larger errors, essentially meaning that a few bad predictions will generate a very bad (large) RMSE (Vandeput, 2021). This way, the two metrics compensate each other well and provide a good foundation to evaluate the models' accuracy.

In the next section, predictions for ARIMA and UIRP, respectivelly, will be evaluated for each currency pair, starting with USD-GBP.

6.1.1 USD-GBP

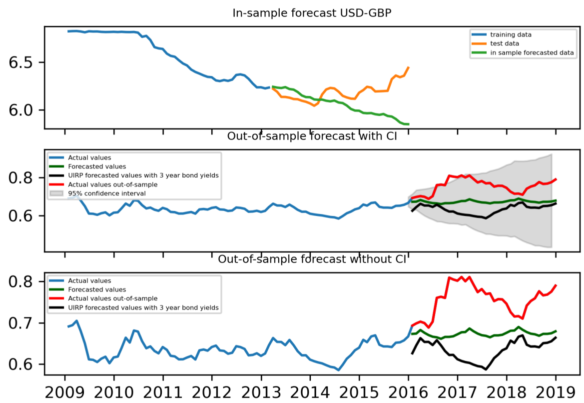

Figure 5: Plot of USD-GBP performance in- and out-of-sample. Note: Confidence interval for the second plot applies only to the ARIMA predictor and is not dependent on the UIRP predictor. NOTE: First panel (displaying in-sample forecast) is incorrect, the correct version is found in Appendix Figure 9.

As is depicted in Figure 5, the actual values observed out-of-sample were at all points greater than what was predicted by both predictors. However, predictions generated by ARIMA seem to outperform all of the predictions made by the UIRP predictor. It is also visible that the actual values fall within the confidence interval generated by the ARIMA model.

The observed values of the evaluation metrics for USD-GBP are (see Table 2):

Clearly, the above claims are confirmed by these metrics, where both MAPE and RMSE are substantially less for ARIMA than they are for UIRP. Nevertheless, both ARIMA and UIRP display a MAPE score between 0.1 and 0.2, indicating that forecasts are satisfactory.

	ARIMA	UIRP
MAPE	0.105274	0.15840
RMSE	0.08927	0.13308

Table 2: USD-GBP evaluation metrics.

	ARIMA	UIRP
MAPE	0.15487	0.088061
RMSE	17.71982	11.08651

Table 3: USD-JPY evaluation metrics.

6.1.2 USD-JPY

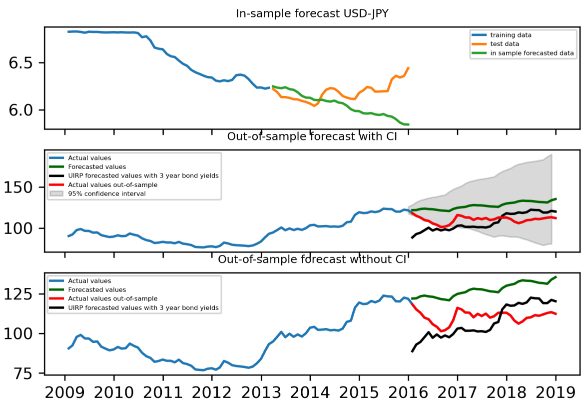

Figure 6: Plot of USD-JPY performance in- and out-of-sample. Note: Confidence interval for the second plot applies only to the ARIMA predictor and is not dependent on the UIRP predictor. NOTE: First panel (displaying in-sample forecast) is incorrect, the correct version is found in Appendix Figure 12.

For USD-JPY, it is apparent that ARIMA overestimates the exchange rate. UIRP, in turn, seems to lie relatively close to the actual values. However, as is visible in Figure 6 as well, it looks as if UIRP and the exchange rate itself are negatively correlated, even though the trend is preserved to some extent.

The observed values of the evaluation metrics for USD-JPY are (see Table 3):

Clearly, the above statistics indicate that UIRP outperforms ARIMA. As in the previous case, both models' predictions yield good MAPE scores. It is also evident that the rather large error estimate during 2017 heavily affect the RMSE score, where values of 17.72 and 11.09 are observed for ARIMA and UIRP, respectively.

	ARIMA	UIRP
MAPE	0.03361	0.046331
RMSE	0.282153	0.35392

Table 4: USD-CNY evaluation metrics.

6.1.3 USD-CNY

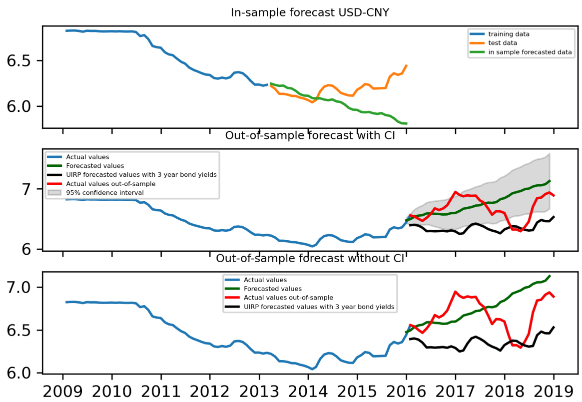

Figure 7: Plot of USD-CNY performance in- and out-of-sample. Note: Confidence interval for the second plot applies only to the ARIMA predictor and is not dependent on the UIRP predictor. NOTE: First panel (displaying in-sample forecast) is incorrect, the correct version is found in Appendix Figure 18.

Figure 7 exhibits that though the ARIMA does not seem to capture the high volatility of the observed values, it does seem to capture the overall trend of the exchange rate. As for UIRP, the predictor seems to underestimate all data points, however, as the range of which the actual values deviates from the UIRP predictions is relatively small, the predictions should not necessarily be deemed unsuccessful.

The observed values of the evaluation metrics for USD-JPY are (see Table 4):

Interpreting the calculated values above, ARIMA indeed outperforms the UIRP predictor for this currency pair. However, it should be noted that both predictors, judging by the relatively shy size of the metrics, are very successful in estimating the values. The observed MAPE scores are very low and indicate good forecasts.

6.2 Evaluation of strategies

In order to quantify the performance across the strategies, logarithmic returns will be used. The rationale behind this is simple. Clearly, calculating returns as the

	Actual rate	ARIMA	UIRP 3Y	UIRP 2Y	UIRP 1Y
Spot	0.67				
1Y	0.80	0.68			0.67
2Y	0.75	0.68		0.66	
3Y	0.79	0.68	0.66		

	Strategy 1	Strategy 2	Strategy 3	Strategy 4	Strategy 5
1Y	Hedge	No hedge	No hedge	No hedge	No hedge
2Y	Hedge	No hedge	Hedge	Hedge	No hedge
3Y	Hedge	No hedge	Hedge	Hedge	No hedge

Log returns	Strategy 1	Strategy 2	Strategy 3	Strategy 4	Strategy 5
1Y	0	0.1773	0.1773	0.1773	0.1773
2Y	0	0.1128	0	0	0.1128
3Y	0	0.1648	0	0	0.1648
Total	**0**	**0.4549**	**0.1773**	**0.1773**	**0.4549**

Table 5: USD-GBP evaluation of strategies.

increase/decrease in nominal value will yield incorrect results as the nominal values differ significantly across currencies and over time. However, the problem is not rectified in its entirety by calculating the returns in percentages - in order to compile the performance of each strategy put together all currencies, we need to be able to, for each strategy, summarise the returns for all currency pairs. Returns calculated as percentages does not have the additive property required for this, which gives reason as to why logarithmic returns will be applied instead.

6.2.1 USD-GBP

Figure 5 displays the evaluation of strategies for USD-GBP. Over the entire three year period, the exchange rate USD-GBP rose by 17.91 percent (calculated in log returns, this number corresponds to 0.1648), and was at the beginning of each year up in relation to the spot rate at 2016-01-01. Therefore, the best alternative would have been to stay with the market, not entering any forward contracts to fix the rate. Hence, Strategy 2 and Strategy 5 are the best performers for this currency as they successfully managed to avoid such fate. In contrast, Strategy 3 and Strategy 4 indicate a forward contract should be entered for two- and three years, respectively. However, the latter two strategies did manage to take advantage of the one year hike in the exchange rate without entering a forward contract, resulting in a positive total return for both strategies.

	Actual rate	ARIMA	UIRP 3Y	UIRP 2Y	UIRP 1Y
Spot	121.73				
1Y	116.01	125.42			120.92
2Y	112.97	130.65		120.54	
3Y	112.30	135.41	120.19		

	Strategy 1	Strategy 2	Strategy 3	Strategy 4	Strategy 5
1Y	Hedge	No hedge	Hedge	Hedge	No hedge
2Y	Hedge	No hedge	Hedge	Hedge	No hedge
3Y	Hedge	No hedge	Hedge	Hedge	No hedge

Log returns	Strategy 1	Strategy 2	Strategy 3	Strategy 4	Strategy 5
1Y	0	-0.0481	0	0	-0.0481
2Y	0	-0.0747	0	0	-0.0747
3Y	0	-0.0806	0	0	-0.0806
Total	**0**	**-0.2034**	**0**	**0**	**-0.2034**

Table 6: USD-JPY evaluation of strategies.

6.2.2 USD-JPY

Figure 6 displays the evaluation of strategies for USD-JPY. The overall trend for the currency pair USD-JPY was clearly negative, as the exchange rate dropped by -7.74 percent (calculated in log returns, this number corresponds to -0.0806). Even though the rate seemed to fluctuate slightly over the three year period, all periods of one, two, and three years see a depreciation in price. Looking at the strategies, only Strategy 3 and Strategy 4 were able to capture this (apart from Strategy 1, of course), fixing the rate at 121.73.

6.2.3 USD-CNY

Figure 7 displays the evaluation of strategies for USD-CNY. Even thought it drops slightly between year one and year two, the exchange rate USD-CNY, in relation to the spot rate, increases in all three cases. Hence, it is desirable not to enter a forward contract. Clearly, all strategies (apart from Strategy 1, of course) are successful and generate a positive return in accordance.

6.2.4 Compilation of all currencies

In total, after total returns for all exchange rates have been compiled, the strategies are ranked, from best to worst, as follows (see table 8):

	Actual rate	ARIMA	UIRP 3Y	UIRP 2Y	UIRP 1Y
Spot	6.44				
1Y	6.95	6.60			6.56
2Y	6.60	6.85		6.55	
3Y	6.89	7.13	6.53		

	Strategy 1	Strategy 2	Strategy 3	Strategy 4	Strategy 5
1Y	Hedge	No hedge	No hedge	No hedge	No hedge
2Y	Hedge	No hedge	No hedge	No hedge	No hedge
3Y	Hedge	No hedge	No hedge	No hedge	No hedge

Log returns	Strategy 1	Strategy 2	Strategy 3	Strategy 4	Strategy 5
1Y	0	0.0762	0.0762	0.0762	0.0762
2Y	0	0.0245	0.0245	0.0245	0.0245
3Y	0	0.0675	0.0675	0.0675	0.0675
Total	**0**	**0.1682**	**0.1682**	**0.1682**	**0.1682**

Table 7: USD-CNY evaluation of strategies.

Rank		Total log return
1	Strategy 2	0.4197
1	Strategy 5	0.4197
3	Strategy 3	0.3455
3	Strategy 4	0.3455
5	Strategy 1	0

Table 8: All currencies - evaluation of strategies.

7 Analysis

In this section, the sensitivity of the ARIMA model is analysed, followed by an analysis on the strategies presented in Section 5.3.

7.1 Sensitivity analysis

While creating a forecasting method such as the ARIMA model, the robustness of the model ought to be addressed. The robustness is a measurement of how well the model performs while parameters, variables or assumptions are altered (Kenton, 2020). A robust strategy in-sample has better chances to be successful out-of-sample. In our case, we have tested the robustness of the model by varying the selection of parameters as well as lengths of sample data.

7.1.1 Parameter uncertainty

In some cases, it is not apparent to tell if the original time series is stationary or non-stationary. If it is non-stationary, as explained earlier, it ought to be differenced at least once. However, if the d parameter is enforced to 0, even though the data might be non-stationary, predictions change a lot, partly because enforcing d to 0 affects the selection of the p and q parameters as well.

The other variable to investigate, the sample length with which the model is trained with, is more difficult to optimise. Using a lot of historical data might put too much emphasis on obsolete data, rather than more recent data which might be more relevant for the forecast. However, merely using the very recent data might be erroneous in its own way, should the exchange rate display a temporary hike or drop, for example.

The ideas above are demonstrated in Figure 9-20 in Appendix, where we compare the original ARIMA predictions against ARIMA predictions generated subsequent to altering the (p, d, q)-parameters and the sample size, respectively. By inspection, some forecasts benefit from a shorter sample size. Looking at USD-GBP, for example, the original set of training data put too little emphasis on the most recent hike, resulting in a flatter trend line of the forecasted values than what is desirable. Furthermore, altering the (p, d, q)-parameters with the constraint $d = 0$ deteriorates the predictions in (almost) all cases. The selected parameters for each case and currency are visualised in Table 9.

Exchange rate	Original	d enforced to 0	Shorter sample
USD-GBP	(1,1,0)(0,1,1)[12]	(2,0,1)(0,1,1)[12]	(1,1,0)(1,1,0)[12]
USD-JPY	(1,1,1)(0,1,1)[12]	(4,0,0)(2,1,0)[12]	(0,1,0)(2,1,0)[12]
USD-CNY	(0,1,1)(0,1,2)[12]	(2,0,1)(0,1,1)[12]	(0,1,1)(0,1,1)[12]

Table 9: (S)ARIMA parameters

7.2 Strategy analysis

Considering the rather low number of observations, especially as extracting required bond data for Sweden is not feasible, one shall be careful before jumping to conclusions. However, apart from Strategy 1, it is obvious, by construction, that Strategy 4 carries the lowest risk. Still, we observe a positive total return for the strategy. The low risk taking of Strategy 4 is the underlying reason why it does not manage to capture the hike in price during year 2 and 3 for USD-GBP, however, the decision of hedging the position during USD-JPY drop in price should be attributed to the same low risk. Strategy 2, on the other hand, produces the highest total return but does not manage to detect the depreciation in the USD-JPY rate, essentially yielding a more volatile profit and loss curve. By coincidence, Strategy 3 performs identical to Strategy 2, however, it is clear that the former strategy carries a higher risk by the way it is constructed. Therefore, it stands to reason to hold Strategy 4 in high esteem in regards to its applicability to be deployed in practice.

8 Discussion

8.1 Potential issues with the models

There are a few conditions that should raise some concern when it comes to the validity of the model. Firstly, only a handful of currencies are tested. Results might not hold should additional data sets be included in the backtesting. This has not been done because of the relatively short time frame during which the project is run. Secondly, all exchange rates, except USD-JPY, experience a generally positive trend subsequent to the date of the forecast. Although such issue could not have been prevented beforehand as an important part of the forecasting is to avoid overfitting the data, it might be rectified automatically by including additional currency pairs.

It should also be noted that the chosen parameters and time periods for training and testing data are not necessarily optimal. Clearly, one can see that the in-sample tests generally display a rather poor performance for all currencies. Still, adjusting the parameters to fit the in-sample test might cause problems of overfitting, which is usually that the chosen parameters perform well in-sample but poorly out-of-sample. Figure 18-20 in Appendix demonstrates this well; enforcing d to 0 generates a better in-sample forecast, but a worse result as data points are forecasted out-of-sample.

Lastly, it might be naive to assume that it is feasible for the investor to enter a forward contract using today's spot rate. As proposed earlier, should the market display a substantial incline in price, for example, chances are that the forward prices for one, two, and three years will exceed today's spot price.

8.2 Further research

Forecasting of exchange rates is an interesting topic and there is definitely room for improvement. First off, the potential issues with the model should be addressed. It would be interesting to see how Strategy 4 performs as it is challenged with different sample sizes and a greater number of exchange rates.

Another area of further research is to include the level of inflation while calculating future exchange rates using UIRP. Admittedly, both interest rates and the spot rate are affected by the current rate of inflation, but due to price stickiness, this does adaptation is not immediate and including the element of inflation might therefore add additional information.

Lastly, it shall be noted that no attention has been directed towards the variance-covariance matrix of the different currencies. As stated in the beginning of this thesis, the exchange rate risk does not only depend on the uncertainty of the future, but also on how much of the specific currency the investor already holds. Clearly, an investor holding several currencies has achieved a more diversified portfolio and will therefore automatically reduce her/his risk exposure. This idea of portfolio diversification could be a suitable next step in the process of creating an extensive exchange rate risk mitigation program.

9 Concluding remarks

The purpose of this thesis was to investigate whether a traditional economic model and a more sophisticated statistical autoregressive method could be combined and applied jointly for the purpose of forecasting exchange rates. More specifically, I utilise the economic theory of Uncovered Interest Rate Parity together with the statistical Autoregressive Integrated Moving Average method, with the purpose of deciding whether the exchange rate of a currency pair is expected to appreciate or depreciate in the future. The time horizons during which the exchange rates are predicted are one, two, and three years.

Even though the sample size for this research has been relatively small, I find some promising results supporting the predictability of the Forex market. Specifically, I find that a combination of these models is a good tool for mitigating the inaccuracy of the forecast, as compared to using the predictors as standalone methods, and that such an application of the forecasting models might serve as a good tool for mitigating exchange rate risk of the portfolio in general. This is demonstrated by Strategy 4 in the analysis section - a joint application of the models reduces the risk associated with the exchange rate exposure. Additionally, thorough analysis is conducted on the models separately and it is concluded that ARIMA is slightly superior to UIRP, however, the difference is not vast and an increase in sample size is required in order to assert anything with confidence.

References

Allwright, S. (2021). What is a good mape score? Available online: https://stephenallwright.com/good-mape-score/, [Accessed 18 May 2022].

Altavilla, C. and De Grauwe, P. (2006). Forecasting and combining competing models of exchange rate determination. *Applied Economics*, 42(1747):1–29.

Amat, C., Michalski, T., and Stoltz, G. (2018). Fundamentals and exchange rate forecastability with simple machine learning methods. *The Journal of Finance*, 88:1–24.

Banton, C. (2021). How are international exchange rates set? Available online: https://www.investopedia.com/ask/answers/forex/how-forex-exchange-rates-set.asp, [Accessed 20 April 2022].

Bekaert, G., Wei, M., and Xing, Y. (2007). Uncovered interest rate parity and the term structure. *Journal of International Money and Finance*, 26:1038–1069.

Bevans, R. (2020). Akaike information criterion — when how to use it. Available online: https://www.scribbr.com/statistics/akaike-information-criterion/, [Accessed 19 March 2022].

Bora, N. (2021). Understanding arima models for machine learning. Available online: www.capitalone.com/tech/machine-learning/understanding-arima-models, [Accessed 10 March 2022].

Box, G. and Jenkins, G. (1970). *Time Series Analysis: Forecasting and Control*. Holden-Day, 1 edition.

Chang, Y. and Hamori, S. (2020). The predictability of the exchange rate when combining machine learning and fundamental models. Available online: https://doi.org/10.3390/jrfm13030048, [Accessed 18 March 2022].

Chen, J. (2021). Currency risk. Available online: https://www.investopedia.com/terms/c/currencyrisk.asp, [Accessed 13 May 2022].

Ericsson (2020). Ericsson in sweden. Available online: https://www.ericsson.com/en/about-us/company-facts/ericsson-worldwide/sweden, [Accessed 17 May 2022].

Eun, C. S. and Resnick, B. G. (2007). *International Financial Management*. Mc Graw Hill, 7 edition.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417.

Frankel, J. A. (1993). *On Exchange Rates*. MIT press, 1 edition.

Gali, J. (2010). *Handbook of Monetary Economics*. Elsevier, 3 edition.

Hayes, A. (2021a). Price stickiness. Available online: https://www.investopedia.com/terms/s/sticky-down.asp, [Accessed 15 May 2022].

Hayes, A. (2021b). Uncovered interest rate parity (uip). Available online: https://www.investopedia.com/terms/u/uncoveredinterestrateparity.asp, [Accessed 13 May 2022].

Kenton, W. (2020). Robust. Available online: https://www.investopedia.com/terms/r/robust.asp, [Accessed 15 May 2022].

Kim, S. and Kim, H. (2016). A new metric of absolute percentage error intermittent demand forecasts. *International Journal of Forecasting*, 32(3):669–679.

Kotu, V. and Deshpande, B. (2015). *Predictive Analytics and Data Mining*. Elsevier, 1 edition.

Kumar, R. (2014). *Strategies of Banks and Other Financial Institutions*. Elsevier, 1 edition.

Lam, L., Fung, L., and Yu, I.-w. (2008). Comparing forecast performance of exchange rate models. *Hong Kong Monetary Authority*, 1(1):1–23.

Leiva-Leon, D., Martínez-Martín, J., and Ortega, E. (2020). Exchange rate shocks and inflation comovement in the euro area. *ECB Working Paper Series*, 1(2383):1–55.

Loh, L.-Y., Kueh, H.-K., Parikh, N.-J., Chan, H., Ho, N.-H., and Chua, M.-H. (2022). An ensembling architecture incorporating machine learning models and genetic algorithm optimization for forex trading. *FinTech*, 1(2):1–55.

Malkiel, B. G. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspective*, 17(1):59–82.

McCarthy, S. (2016). Foreign exchange operating exposure: A practical teaching approach. *Journal of Financial Education*, 42(1-2):116–136.

Melvin, M. and Norrbin, S. C. (2012). *International Money and Finance*. Imprint, 8 edition.

Nordea (2021). 5 steps to manage your currency risk. Available online: https://www.nordea.com/en/news/5-steps-to-manage-your-currency-risk, [Accessed 20 April 2022].

OECD (2014). Risk management and corporate governance. Available online: http://dx.doi.org/10.1787/9789264208636-en, [Accessed 22 April 2022].

Offiong, A. I., Riman, H. B., and Eyo, E. E. (2016). Determining optimal portfolio in a three-asset portfolio mix in nigeria. *Journal of Mathematical Finance*, 4(6):524–540. Available online: https://www.scirp.org/journal/paperinformation.aspx?paperid=71156, [Accessed 19 April 2022].

Orellana, V. and Pino, G. (2021). Uncovered interest rate parity: A gravity-panel approach. *CellPress*, 7(11):1–9.

Prabhakaran, S. (2021). Arima model - complete guide to time series forecasting in python. Available online: https://www.machinelearningplus.com/time-series/arima-model-time-series-forecasting-python/, [Accessed 15 March 2022].

Thune, K. (2021). Risk averse: Definition example. Available online: https://seekingalpha.com/article/4459543-risk-averse, [Accessed 4 May 2022].

Vandeput, N. (2021). Forecast kpis: Rmse, mae, mape bias forecasting in python. Available online: https://towardsdatascience.com/forecast-kpi-rmse-mae-mape-bias-cdc5703d242d, [Accessed 18 May 2022].

10 Appendix

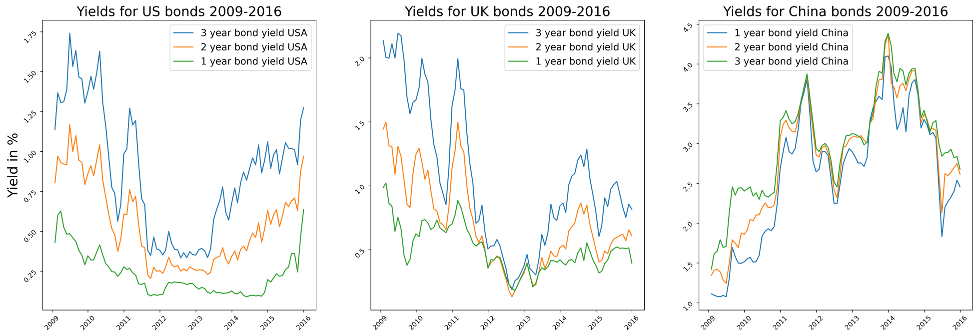

Figure 8: Example of bond yields for USA, UK and China between 2009-2016. Yields for all periods are averaged on a monthly basis. Please notice that the color code changed for the last panel, this issue could no be rectified due to technical issues.

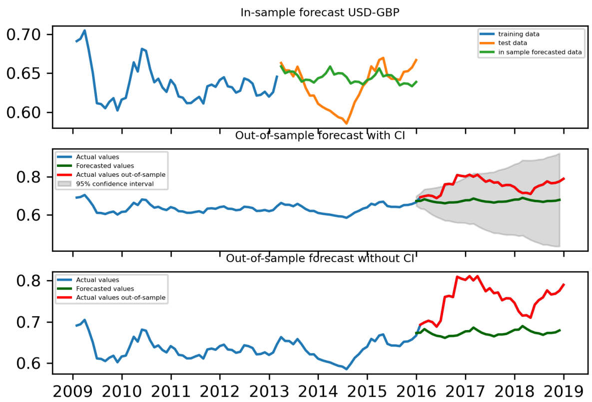

Figure 9: USD-GBP forecast. Original parameters used.

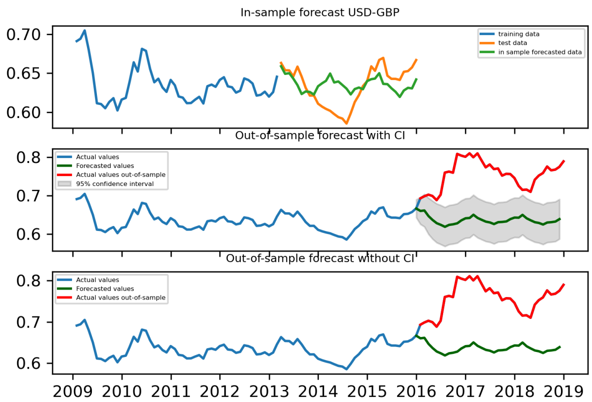

Figure 10: USD-GBP forecast. d enforced to 0.

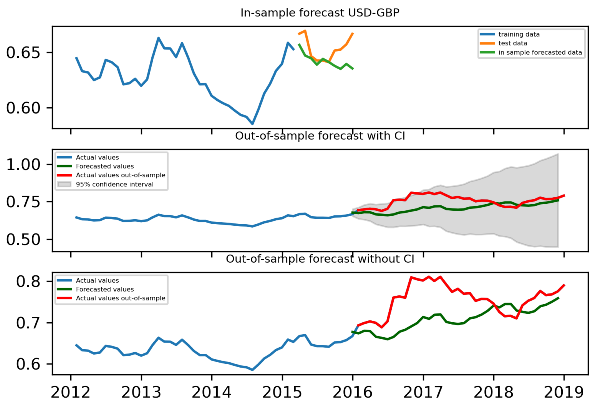

Figure 11: USD-GBP forecast. Sample size decreased by three years.

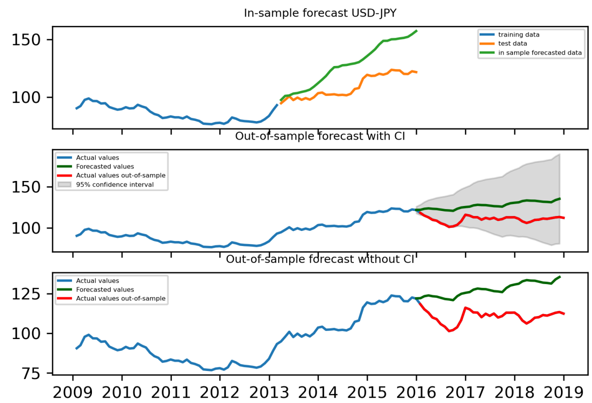

Figure 12: USD-JPY forecast. Original parameters used.

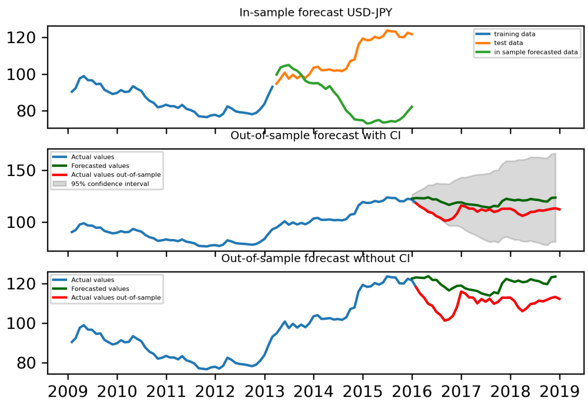

Figure 13: USD-JPY forecast. d enforced to 0.

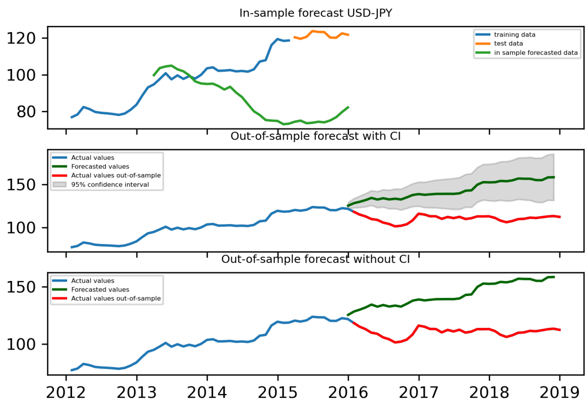

Figure 14: USD-SEK forecast. Sample size decreased by three years.

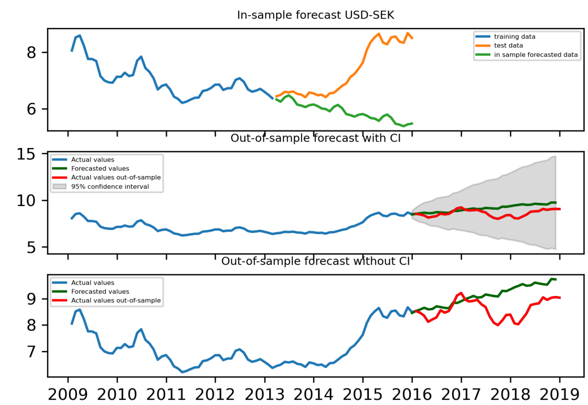

Figure 15: USD-SEK forecast. Original parameters used.

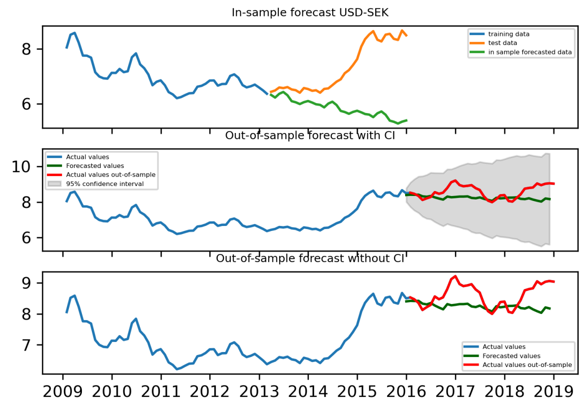

Figure 16: USD-SEK forecast. d enforced to 0.

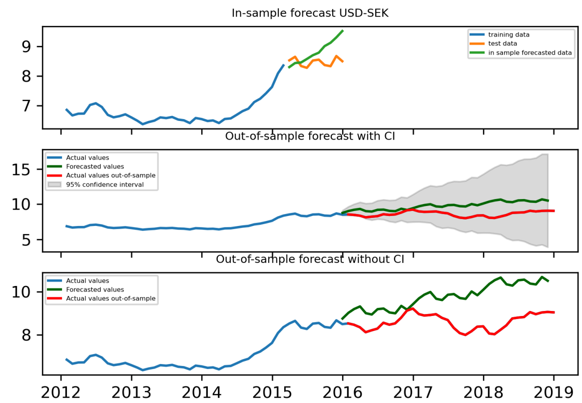

Figure 17: USD-SEK forecast. Sample size decreased by three years.

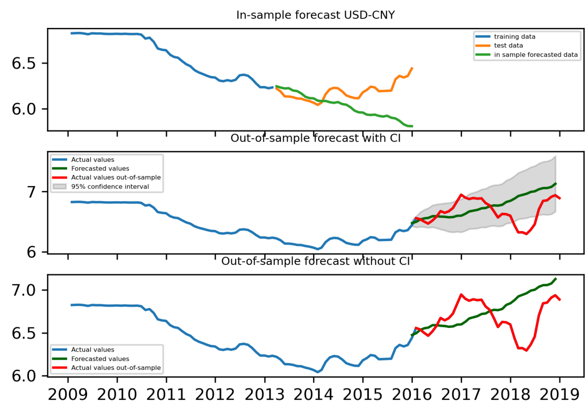

Figure 18: USD-CNY forecast. Original parameters used.

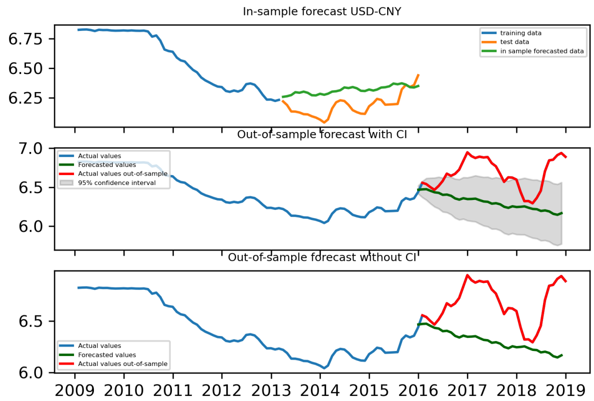

Figure 19: USD-CNY forecast. d enforced to 0.

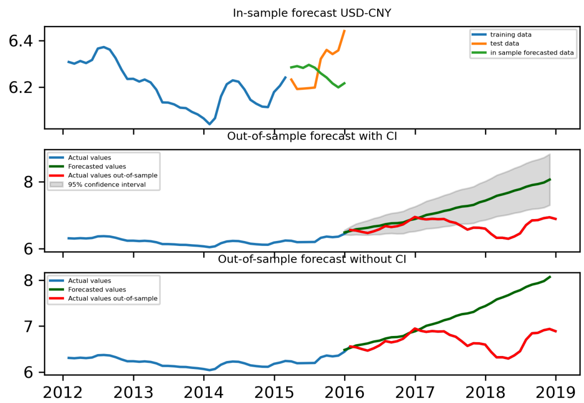

Figure 20: USD-CNY forecast. Sample size decreased by three years.