University of Sussex

A University of Sussex PhD thesis

Available online via Sussex Research Online:

http://sro.sussex.ac.uk/

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details

FORECASTING EXCHANGE RATES: AN APPLICATION TO THE DAILY HIGH AND LOW

Nima Shahroozi

Submitted for the degree of Doctor of Philosophy

University of Sussex

Supervisor: Dr. Qi Tang

February 2017

Declaration

I hereby declare that this thesis has not been and will not be submitted in whole or in part to another University for the award of any other degree. I also declare that this thesis was composed by myself and that the work contained therein is my own, except where explicitly stated otherwise.

Signature:

Nima Shahroozi

UNIVERSITY OF SUSSEX

NIMA SHAHROOZI, SUBMITTED FOR DOCTOR OF PHILOSOPHY

<u>FORECASTING EXCHANGE RATES :</u> <u>AN APPLICATION TO THE DAILY HIGH AND LOW</u>

Abstract

In this thesis, we study the behaviour and forecastability of exchange rates. Most of the existing literature on the forecasting of exchange rates concentrates on the end of the day price, commonly known as the 'close' price. Meese and Rogoff [30] show that this price tends to follow the naive random walk model, which implies that the best forecast for the next period is the current observed value. Instead, we study the dynamics and the predictability of the daily high and low prices using real-world data for the currency pairs GBP/USD, EUR/USD and AUD/USD. The daily high and low are the maximum and minimum prices reached for each 24-hour period by the currency pairs. We find strong evidence that the daily close prices lag these highs and lows. We use this knowledge to build an autoregressive distributed lag (ARDL) rolling regression model that produces one day ahead out-of-sample forecasts of these high and low prices. We also build an algorithm that uses already existing dynamic regression methods to correct for the autocorrelation often observed in time-series data. The window size used for the estimation of our model parameters is very important due to the nature of time-series data. We propose an empirical method to find the best suitable window size for the estimation of these parameters. The out-of-sample predictability of our regression models is compared to a few benchmark models by using a number of different

performance measures. We show that our models outperform these benchmark models in terms of their forecasting ability of high and low prices. Furthermore, a triggering method is developed for trading exchange rates using a saturation-reset linear feedback controller. First, we test our triggering method on an idealized market model, for which we propose a stochastic process. We then apply this triggering method to real-world data in order to study its performance. Finally, we construct trading strategies that combine these methods with our out-of-sample forecasts.

Acknowledgements

First and foremost I would like to thank my supervisor Dr Qi Tang for all his help, patience and support throughout my entire PhD studies at the University of Sussex. It has been an absolute pleasure working with him as his PhD student and this journey would have not been possible without his advice and assistance. I am also very thankful for all his efforts to proofread my thesis and helping me right to the final day of this journey. I would also like to thank my second supervisor Dr Bertram During for all his advice and comments during the most stressful of times. A special thank you goes to all the staff and faculty of the Mathematics department at Sussex. I have had the pleasure of being part of this amazing family at Sussex since my undergraduate studies in 2007 and have spent the best parts of my life with them. Additionally, I would like to thank all of my colleagues and friends at the University for always supporting me. I greatly appreciate all the technical and non-technical help Giannis, Haidar and Christof provided me throughout this journey. Also a very special thank you to Roxanna for always being there for me and keeping me calm. Finally, I would like to express my deepest gratitude to my family for their care and support. I am so thankful to have such amazing parents, brother and sister in my life. None of this would have been possible without the love, care and support of my parents for which I am always grateful. They have always been there for me no matter what and have motivated me when I needed it the most.

Contents

1	1 Introduction				
	1.1	Introduction to the FX Market & Daily "Highs" and "Lows"	1		
	1.2	Literature Review	3		
2	Ben	chmark Models	8		
	2.1	Introduction	8		
	2.2	The Naive Model	9		
	2.3	Auto Regressive (AR) Process	9		
	2.4	Moving Average (MA) Process	10		
	2.5	Stationarity	10		
	2.6	Augmented Dicky-Fuller Test	13		
	2.7	Autocovariance (AC) & Autocorrelation Function (ACF)	14		
	2.8	Partial Autocorrelation Function (PACF)	16		
	2.9	ARMA Process: The Box-Jenkins Methodology	18		
	2.10	Ljung-Box Q Test	19		
	2.11	Forecasting Accuracy Measures	19		
	2.12	Data	21		
	2.13	Results	22		
3	Tim	e-Varying ARDL Model & Dynamic Regression	31		

	3.1	Introduction	31
	3.2	Cross-Correlation & Sample Cross-Correlation	32
	3.3	Time-Varying ARDL Model	32
	3.4	Parameter Estimation:Ordinary Least Squares	37
	3.5	Main Assumptions	39
	3.6	Statistical Tests & Model Validation	41
		3.6.1 t-test & F-test	41
		3.6.2 Choosing the Correct Sample Size	42
		3.6.3 White Test	43
		3.6.4 Variance Inflation Factor (VIF)	44
		3.6.5 HAC Estimators	44
		3.6.6 Bayesian Information Criterion (BIC)	46
	3.7	Dynamic Regression	46
	3.8	Results	49
	3.9	Extension	56
4	Trig	gering Method for Exchange Rate Trading Via Feedback Control	62
	4.1	Introduction	62
	4.2	Saturation-Reset Linear Feedback Controller	64
	4.3	The Trade Triggering Method	66
	4.4	Idealized Market Model: A Stochastic Model for the Price Evolution of High,	
		Low and Close	70
	4.5	Simulation Study	74
		4.5.1 Discretization	75
		4.5.2 Varying The Volatility of Highs σ_h	78
		4.5.3 Varying The 'Level' θ of Range $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots$	80
		4.5.4 Varying The 'Volatility' v of Range	83

		4.5.5	Varying The 'Speed' κ of Range $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots$	86					
	4.6	Real-W	Norld Application	88					
		4.6.1	Maximum Drawdown	89					
		4.6.2	Back-testing With Arbitrary Parameters	89					
		4.6.3	Back-Testing Using Optimisation	92					
	4.7	Combi	ining Forecasts and Triggering Method	94					
		4.7.1	Strategy 1	95					
		4.7.2	Strategy 2	100					
		4.7.3	Strategy 3	104					
		4.7.4	Strategy 4	109					
	4.8	Comp	aring Performances:Risk-free, FTSE100 and S&P500	114					
Aj	Appendices 1								
\mathbf{A}	Dat	a Sour	ces	125					
В	Drif	t and	Volatility Estimates of Highs and Lows	126					
С	C Estimation of the Parameters in the CIR Process								
D	D Matlab Codes 13								

List of Figures

2.1	SACF of 500 simulated observations from a $MA(1)$ process	16
2.2	SPACF of 500 simulated observations from an $AR(1)$ process	17
2.3	Distribution of prices and log returns: GBP/USD Highs, $N=5979$	23
2.4	SACF and SPACF of log returns of GBP/USD Highs, $N=5979$	24
2.5	SACF and SPACF of log returns of GBP/USD Lows, $N=5979$	24
2.6	SACF and SPACF of log returns of AUD/USD Highs, $N=5206$	25
2.7	SACF and SPACF of log returns of AUD/USD Lows, $N=5206$	25
2.8	SACF and SPACF of log returns of EUR/USD Highs, $N=3649$	25
2.9	SACF and SPACF of log returns of EUR/USD Lows, $N=3649$ \hdots	26
3.1	SCC between log returns of GBP/USD Highs and Lows, $N=5979$	34
3.2	SCC between log returns of GBP/USD Highs and Close, $N=5979$ \ldots .	34
3.3	SCC between log returns of GBP/USD Lows and Close, $N=5979$	34
3.4	SCC between log returns of AUD/USD Highs and Lows, $N=5206$ \hfill	35
3.5	SCC between log returns of AUD/USD Highs and Close, $N=5206$ $\ .$	35
3.6	SCC between log returns of AUD/USD Lows and Close, $N=5206$	35
3.7	SCC between log returns of EUR/USD Highs and Lows, $N=3649$	36
3.8	SCC between log returns of EUR/USD Highs and Close, $N=3649$ \hdots	36
3.9	SCC between log returns of EUR/USD Lows and Close, $N=3649$	36
3.10	Distribution of actual V forecasted log returns: GBP/USD Highs \ldots .	59

3.11	Distribution of actual V forecasted log returns: GBP/USD Lows	60
3.12	Distribution of actual V forecasted log returns: EUR/USD Highs \ldots .	60
3.13	Distribution of actual V forecasted log returns: EUR/USD Lows	60
3.14	Distribution of actual V forecasted log returns: AUD/USD Highs \ldots .	61
3.15	Distribution of actual V forecasted log returns: AUD/USD Lows	61
4.1	R_t -AUD/USD	73
4.2	R_t -GBP/USD	73
4.3	R_t -EUR/USD	73
4.4	Histogram of final account value for high volatility case	79
4.5	Histogram of final account value for medium volatility case	79
4.6	Histogram of final account value for low volatility case	80
4.7	Histogram of final account value for medium volatility-high θ case	81
4.8	Histogram of final account value for medium volatility-medium θ case $\ . \ . \ .$	82
4.9	Histogram of final account value for medium volatility-low θ case	82
4.10	Histogram of final account value for medium volatility-high v case \ldots .	84
4.11	Histogram of final account value for medium volatility-medium v case	84
4.12	Histogram of final account value for medium volatility-low v case \ldots .	85
4.13	Histogram of final account value for medium volatility-high κ case $\ . \ . \ .$	86
4.14	Histogram of final account value for medium volatility-medium κ case	87
4.15	Histogram of final account value for medium volatility-low κ case \ldots .	87
4.16	Account value V_t for GBP/USD : Arbitrary parameters	90
4.17	Account value V_t for AUD/USD : Arbitrary parameters	90
4.18	Account value V_t for EUR/USD : Arbitrary parameters	91
4.19	Account value V_t for GBP/USD : optimised parameters	93
4.20	Account value V_t for AUD/USD : optimised parameters	93
4.21	Account value V_t for EUR/USD : optimised parameters	93

4.22 Account value V_t for GBP/USD : Strategy 1 using arbitrary parameters . . . 96 4.23 Account value V_t for AUD/USD : Strategy 1 using arbitrary parameters . . . 96 4.24 Account value V_t for EUR/USD : Strategy 1 using arbitrary parameters . . . 97 4.25 Account value V_t for GBP/USD : Strategy 1 using optimised parameters . . 98 4.26 Account value V_t for AUD/USD : Strategy 1 using optimised parameters . . 98 4.27 Account value V_t for EUR/USD : Strategy 1 using optimised parameters 99 4.28 Account value V_t for GBP/USD : Strategy 2 using arbitrary parameters . . . 101 4.29 Account value V_t for AUD/USD : Strategy 2 using arbitrary parameters . . . 101 4.30 Account value V_t for EUR/USD : Strategy 2 using arbitrary parameters . . . 1024.31 Account value V_t for GBP/USD : Strategy 2 using optimised parameters . . 1034.32 Account value V_t for AUD/USD : Strategy 2 using optimised parameters . . 1034.33 Account value V_t for EUR/USD : Strategy 2 using optimised parameters . . 1044.34 Account value V_t for GBP/USD : Strategy 3 using arbitrary parameters . . . 1054.35 Account value V_t for AUD/USD : Strategy 3 using arbitrary parameters . . . 1064.36 Account value V_t for EUR/USD : Strategy 3 using arbitrary parameters . . . 1064.37 Account value V_t for GBP/USD : Strategy 3 using optimised parameters . . 1074.38 Account value V_t for AUD/USD : Strategy 3 using optimised parameters . . 1084.39 Account value V_t for EUR/USD : Strategy 3 using optimised parameters . . 108 4.40 Account value V_t for GBP/USD : Strategy 4 using arbitrary parameters . . . 1104.41 Account value V_t for AUD/USD : Strategy 4 using arbitrary parameters . . . 1104.42 Account value V_t for EUR/USD : Strategy 4 using arbitrary parameters . . . 111 4.43 Account value V_t for GBP/USD : Strategy 4 using optimised parameters . . 1124.44 Account value V_t for AUD/USD : Strategy 4 using optimised parameters . . 1124.45 Account value V_t for EUR/USD : Strategy 4 using optimised parameters . . 1134.46 Account value V_t for 3 months US T-bills : Buy-and-hold strategy 114

4.48	Account value V_t for	r S&l	P500 :	Bu	y-ano	l-hol	d st	rate	gy	 	•	 •	 		115
B.1	$\hat{\mu_{h,t}}$ of AUD/USD .									 			 	•	126
B.2	$\hat{\mu_{l,t}}$ of AUD/USD $% \hat{\mu_{l,t}}$.							• •		 			 	•	127
B.3	$\hat{\sigma_{h,t}}$ of AUD/USD .							•••		 			 	•	127
B.4	$\hat{\sigma_{l,t}}$ of AUD/USD $% \hat{\sigma_{l,t}}$.							• •		 			 	•	127
B.5	$\hat{\mu_{h,t}}$ of GBP/USD $% \hat{\mu_{h,t}}$.							• •		 			 	•	128
B.6	$\hat{\mu_{l,t}}$ of GBP/USD $% \hat{\mu_{l,t}}$.							• •		 			 	•	128
B.7	$\hat{\sigma_{h,t}}$ of GBP/USD $% \hat{\sigma_{h,t}}$.							• •		 			 	•	128
B.8	$\hat{\sigma_{l,t}}$ of GBP/USD $$.							• •		 			 	•	129
B.9	$\hat{\mu_{h,t}}$ of EUR/USD $% \hat{\mu_{h,t}}$.							• •		 			 	•	129
B.10	$\hat{\mu_{l,t}}$ of EUR/USD $% \hat{\mu_{l,t}}$.							• •		 			 	•	129
B.11	$\hat{\sigma_{h,t}}$ of EUR/USD .		•••							 			 	•	130
B.12	$\hat{\sigma_{l,t}}$ of EUR/USD .									 			 	•	130

List of Tables

$_{2.1}$	ADF test results for price series $\{p\}$, and return series $\{r\}$, of all variables for	
	all data sets at 5% significant level \ldots	23
2.2	Benchmark models forecasting results : AUD/USD-Highs	27
2.3	Benchmark models forecasting results : AUD/USD-Lows	27
2.4	Benchmark models forecasting results : EUR/USD-Highs	27
2.5	Benchmark models forecasting results : EUR/USD-Lows	28
2.6	Benchmark models forecasting results : GBP/USD-Highs	28
2.7	Benchmark models forecasting results : GBP/USD-Lows	28
2.8	LBQ test fails to accept H_0 of no autocorrelation % of time in (N-501) regres-	
	sions at 5% significant level: Benchmark models \hdots	29
3.1	VIF values	49
3.2	Correlation between variables EUR/USD	49
3.3	Correlation between variables AUD/USD	40
		49
3.4	Correlation between variables GBP/USD	49 50
3.4 3.5	Correlation between variables GBP/USD Diagnostic results-ARDL model: GBP/USD Highs	49 50 50
3.4 3.5 3.6	Correlation between variables GBP/USD Correlation between variables GBP/USD Diagnostic results-ARDL model: GBP/USD Highs Correlation between variables GBP/USD Diagnostic results-ARDL model: GBP/USD Lows Correlation between variables GBP/USD	 49 50 50 51
 3.4 3.5 3.6 3.7 	Correlation between variables GBP/USD	 49 50 50 51 51
 3.4 3.5 3.6 3.7 3.8 	Correlation between variables GBP/USD	 49 50 50 51 51 51

3.10	Diagnostic results-ARDL model: AUD/USD Lows	52
3.11	Forecasting results: AUD/USD Highs	54
3.12	Forecasting results: AUD/USD Lows	54
3.13	Forecasting results: EUR/USD Highs	55
3.14	Forecasting results: EUR/USD Lows	55
3.15	Forecasting results: GBP/USD Lows	55
3.16	Forecasting results: GBP/USD Highs	55
3.17	Modified models forecasting results: GBP/USD Highs	57
3.18	Modified models forecasting results: GBP/USD Lows	57
3.19	Modified models forecasting results: EUR/USD Highs	57
3.20	Modified models forecasting results: EUR/USD Lows	57
3.21	Modified models forecasting results: AUD/USD Highs	57
3.22	Modified models forecasting results: AUD/USD Lows	58
3.23	Forecast RMSE: When previous days' range falls into different quantiles	59
4.1	Generalised Hurst Exponent values of R_t	74
4.2	Idealized market account value analysis : Cases with varying volatility	80
4.3	Idealized market account value analysis : Cases with varying θ	82
4.4	Idealized market account value analysis : Cases with varying v	85
4.5	Idealized market account value analysis : Cases with varying κ	87
4.6	Account value performance for all currency pairs using arbitrary parameter	
	values	91
4.7	Account value performance for all currency pairs using optimised parameter	
	values	94
4.8	Account value performance of Strategy 1 for all currency pairs using arbitrary	
	parameter values	97

4.9	Account value performance of Strategy 1 for all currency pairs using optimised	
	parameter values	99
4.10	Account value performance of Strategy 2 for all currency pairs using arbitrary	
	parameter values	102
4.11	Account value performance of Strategy 2 for all currency pairs using optimised	
	parameter values	104
4.12	Account value performance of Strategy 3 for all currency pairs using arbitrary	
	parameter values	106
4.13	Account value performance of Strategy 3 for all currency pairs using optimised	
	parameter values	108
4.14	Account value performance of Strategy 4 for all currency pairs using arbitrary	
	parameter values	111
4.15	Account value performance of Strategy 4 for all currency pairs using optimised	
	parameter values	113

Chapter 1

Introduction

1.1 Introduction to the FX Market & Daily "Highs" and "Lows"

The foreign exchange market, usually referred to as FX or Forex, is the market where exchange rates are determined and traded. These exchange rates are the prices of one currency quoted in terms of another currency. This market has been studied for many years, as its price movements could affect economic development and international trade. Therefore, governments, central banks, international companies and financial traders closely monitor it, as it is well known that the FX markets can:

(1.) be manipulated and controlled during times of setting fiscal and monetary policies;

(2.) help companies hedge the risk they face due to movements in currency prices;

(3.) help financial traders develop new trading systems to maximise returns and profits.

According to the latest survey taken by the BIS in April 2013, almost 5.3 trillion dollars are traded daily in this market, making it one of the most liquid markets in the financial industry [40]. The enormity of this market, coupled with the fact that it operates 24 hours a day (with the exception of weekends), motivated us to choose this market as the focal point of our research. In this thesis, we look at the FX market in detail and aim to test the predictability of exchange rates. Traditionally, when speaking of financial asset prices (stocks, currencies, etc.), the common price considered has been the "close" price of the asset for that defined time-frame (1 week, 1 day, 1 hour, etc.). However, in this thesis, our interest lies in the "high" and "low" prices of exchange rates rather than in the "close" price. We test the predictability power of these price levels using real-world data. These highs and lows can be interpreted as the highest (maximum) and the lowest (minimum) prices recorded by that asset during the time frame in question.

Our first reason for choosing highs and lows over the closing price is the Efficient Market Hypothesis (EMH), as it argues against the predictability of spot (close) prices. The extensive research carried out in this field tends to agree with EMH and with the fact that exchange rate close prices are hard to predict and fail to beat the random walk model. We will discuss this in more detail in our literature review.

Our second reason for choosing these prices over the closing price is that they can be very informative and give greater insight into the market's behaviour. For example, if considering daily high and lows, then these prices show the highest and the lowest prices recorded for that asset during the defined 24 hour window, and their linear difference is known as each day's trading range. In [37], the author argues that estimating volatility using range-based methods rather than more traditional return-based methods is actually more efficient. Therefore, if these prices are forecasted and known a priori, then they can be used as

(1.) Buy/Sell levels to make profitable trades, and

(2.) a good estimate of future volatility.

1.2 Literature Review

As already mentioned, our main reason for not using close prices is based on the Efficient Market Hypothesis (EMH) initially discussed in [17] and [18]. An efficient market is described as a market in which actual prices at any given time, given the available information, are a very good estimate of intrinsic values. Therefore, when news or new information becomes available, it is reflected in the prices straight away, making it impossible for market analysts to achieve returns greater than those by holding a randomly selected portfolio of assets. [29] This essentially means that past price movements and patterns are not an indicator of future price movements, and that these future prices cannot be predicted using technical analysis or even fundamental analysis techniques. EMH states that these prices follow a random walk process where subsequent price changes are a random departure from the previous ones and are unpredictable. There have been many efforts made in the literature to argue against EMH, and many models have been developed to forecast exchange rate close prices. [30] discussed in one of their most famous classic papers that the best forecasting model for exchange rates is the random walk and that their structural exchange rate models could not beat the out-of-sample forecasting ability of the random walk. The exchange rate models they used for comparison with the random walk are the so called Frenkel-Bilson, Dornbusch-Frankel and Hooper-Morton. These structural models are drawn from economic theories and generally use variables such as inflation, interest rates, trade balances, the unemployment rate, etc. to model exchange rates and subsequently make out-of-sample forecasts. The study by [10] used the same approach and concluded that random walk is still the best predictor for these prices, especially for shorter time horizons. They used the same structural models but also imposed error correction terms and fit them using both parametric and non-parametric approaches. The study by [1] also used the same techniques but included time-series AR(1)and AR(2) models, which seemed to beat the random walk in their study. However, [44] claims that there had been a mistake in computation by [1], and when they applied the same techniques to the same data, they obtained results suggesting the superiority of the random walk model over these models. Some other studies that use these equilibrium models include [22] and [42] to name a few.

Other approaches used for forecasting exchange rates in the literature are the so-called Vector Auto Regression (VAR) models. [24] analysed the predictability of the Full-VAR (FVAR), Bayesian-VAR (BVAR) and Mixed-VAR (MVAR) models when applied to different exchange rates. Their study lead to mixed conclusions. For example, they found that whilst BVAR and MVAR had more forecasting accuracy than their FVAR counterpart, the results were not consistent for all pairs of currencies tested. [8] used BVAR as well as a Bayesian Vector Error Correction model to forecast Asian exchange rates. He showed in his study how these models outperformed the random walk model from a forecasting basis and suggested the use of such models for the more volatile Asian economies.

Other famous approaches used in the literature for forecasting exchange rates involved the use of non-linear methods such as Artificial Neural Networks (NN), fuzzy models and so on. These methods use machine learning and pattern recognition techniques to forecast rates. Some of the papers that discuss these models include [46], [35], [15], [27] and [28]. Although these studies show some seemingly interesting results, they are negated by [34], who examined the accuracy of artificial neural networks predictions compared to linear time-series models such as ARCH and GARCH. He found the black-box approach to be less accurate than the GARCH and ARCH models in the context of exchange rate predictability.

So far we have presented the main studies and arguments against the predictability of exchange rates. However, all the studies covered to this point are in fact "close" price predictors and use the spot price at the end of day (or at a certain predetermined set time) to carry out their research. Our aim is to discuss the predictability of "highs" and "lows" and to test whether the Efficient Market Hypothesis also holds for these prices. We found very little effort to have been made in the literature for the modelling of highs and lows, with [6], [41] and [9] being the only studies found of this type at the time of this writing. In [41], the authors tested the predictability of the high and low prices of Forex data. They fit an error correction model (ECM) to these prices to capture the co-integration relationship between them and used this to obtain out-of-sample forecasts. Their results were shown to be adequate, but the mean squared errors obtained seem to be much higher than those discussed in our work, which we will analyse in more detail in the coming chapters. [41] was the only study we found to forecast highs and lows of FX data rather than stock market data. In [6], the researchers extended the ECM approach and modelled the highs and lows of stock prices using a fractional vector autoregressive model with error correction (FVECM). Their motivation was that the range (the difference between the high and the low) displayed long term memory and therefore extended the ECM model by capturing this using fractional autoregressive techniques. The author in [9] also models the highs and lows of US stock market indices by an Error Correction Model. However, in this study, the author gave no indication of the forecasting ability of these models, and the regular accuracy measurements are not included for model comparison.

Our aim in this thesis is to study the predictability of these highs and lows. For this purpose we introduce different regression models for daily highs and lows. We test the forecasting ability of our models and compare them with some benchmark models widely known in the world of forecasting. We find that our models not only beat the random walk model but also have much more predictability power when the results are compared to those of [41] . We should note that our focus is solely on next day predictions. Although predicting further into the future might be of interest to some long-term investors, this would increase the forecasting uncertainty, which is not desirable. Next day predictions of highs and lows give the trader an insight into the following day's trading range. These predicted prices may be interpreted as support and resistance levels or upper and lower bands for which the trader can develop a strategy to make profitable daily trades. Similar studies that are interested in next day forecasting have been carried out in the electricity markets and their prices. Papers such as [26], [19], [33] and [16] all aim to forecast next day prices in the Californian and the Spanish electricity markets. They use regression, Time-Series ARIMA models, transfer functions and genetic algorithms for this purpose. Of course, the Forex and electricity markets are completely different in their nature and structure, but the idea of forecasting day-ahead prices and the time-series techniques needed are somewhat similar.

Furthermore, we extend our research to the field of control theory. We study the new paradigm described by [3] for the trading of equities. They introduced a saturation-reset linear feedback controller that determines the amount invested in stock over a given period of time. The back-testing of their model is first carried out on a set of synthetic prices, called the idealized market or the idealized price model. This idealized market serves as their building block and the first test to check the profitability of a strategy before backtesting on real-world data, which is a lot more expensive and time consuming. However, the assumptions made for their idealized market model are far from reality, specifically when one considers the stock market. We believe these assumptions are more reflective of the Forex market, and therefore such a controller is more suited to this market. For example, there are no transaction costs involved when trading currencies, unlike stocks. Therefore, by not taking these costs into account for the trading of equities, the trading performance of a model could be over-exaggerated. We discuss these assumptions and their implications in greater detail in Chapter 4.

In order to apply the controller to our FX data, we first build a triggering method. This is an extension to the work of [21], who introduced a triggering method that serves as a signal to enter/exit a trade. We use the high and low of FX rates to build our triggering method. The triggering method in [21] is based on the estimation of the drift of the close prices. They form a confidence interval for a pre-determined significance level that is used to decide whether the stock is trending upwards, downwards or neither. We use a similar approach, but we incorporate the daily highs and lows of exchange rate prices into our method. We use the information of both of these daily prices to determine the drift and hence the trend of the exchange rate. These double estimations serve as a signal for the controller to trigger a 'Buy', 'Sell' or 'No trade'. We test our triggering method on an idealized synthetic price model first before back-testing on real-world data. The price model we introduce is a two-factor stochastic process for the price of the highs and the range between the highs and the lows. Our model also includes the price of the lows and the close, which are all simulated using Monte-Carlo simulation techniques. We complete our back-testing on a set of real-world data for three different currency pairs. We conclude our work by combining the out-of-sample forecasts made for the highs and lows and our triggering method to build dynamic trading strategies.

Chapter 2

Benchmark Models

2.1 Introduction

In this chapter we will discuss the so-called benchmark models we will use as a reference to compare with our proposed forecasting models for the rest of this thesis. These benchmark models are widely used in the world of forecasting. In order to validate a newly built forecasting model, one has to first check its forecasting ability against these benchmark models with some given real-world data. The first model we consider is the "Naive Model", which in loose terms is no different from purely "guessing" what future prices may be by assuming a completely random structure. We then define some more complicated yet still simple processes such as autoregression, moving averages and a mixture of both. In order to explain these processes and carry out the correct analysis on our data, we need to define some major time-series concepts such as stationarity and autocorrelation functions that are also parts of this chapter. We will also define the accuracy measures that we use for the rest of this thesis. These functions will help us assess the forecastability of different methods and measure their accuracy for comparison with each other to establish the so-called more accurate and efficient models. We then conclude the chapter by presenting the results for our benchmark models.

2.2 The Naive Model

This is the most simple model in finance, as it's based on the theory of the random walk in which the best forecast for the next period's price is the last period's actual price, such that:

$$P_{f,t} = P_{a,t-1} \tag{2.1}$$

where $P_{f,t}$ is the price forecasted for time t and $P_{a,t-1}$ is the actual realised price at time t-1. Although this model may seem very basic, in finance, and specifically in the field of forecasting, it is regarded as the standard benchmark model. As we have discussed in our literature review, beating this model has been the hot topic of interest for exchange rates for many years. Hence for our discussions to be valid, we need to test the predictability of this model on our data and then show that our models perform better than this naive model. We follow tradition and take this model as our first benchmark model.

2.3 Auto Regressive (AR) Process

An autoregressive process, commonly known as an AR(p), is defined as a process where its present value at time t is linearly dependent on its past p values. This process can be defined by

$$X_t = c + \sum_{i=1}^p \theta_i X_{t-i} + \epsilon_t \tag{2.2}$$

where c is the model constant, θ_i , for i = 1..., p are the model parameters and ϵ is the model residual with $\mathbb{E}(\epsilon) = 0$ and $Var(\epsilon) = \sigma^2$ [5]. The following AR(1) and AR(2) processes form two of our benchmark models.

$$AR(1): r_{x,t} = c + \theta_1 r_{x,t-1} + \epsilon_t (2.3)$$

$$AR(2): r_{x,t} = c + \theta_1 r_{x,t-1} + \theta_2 r_{x,t-2} + \epsilon_t (2.4)$$

2.4 Moving Average (MA) Process

A moving average process, commonly known as an MA(q), is defined as a process where its present value at time t is linearly dependent on its last q random shocks. An MA(q) process can be defined by

$$X_t = \sum_{i=1}^{q} \phi_i \epsilon_{t-i} + \epsilon_t \tag{2.5}$$

where $\phi_i, \forall i = 1..., q$ are the model parameters and ϵ are the model residuals with $\mathbb{E}(\epsilon) = 0$ and $Var(\epsilon) = \sigma^2$ [5].

2.5 Stationarity

When performing time series and regression analyses on data, we require the past to be representative of the future so that we can estimate models and produce forecasts. This is the main concept behind stationarity. Fitting regression and time-series models to nonstationary data can lead to spurious and meaningless results. A time-series, $\{X_t : t = 1, 2, ..\}$ is said to be stationary if its probability distribution does not change over time. That is, if the joint distribution of $(X_{t+1}, X_{t+2}, X_{t+3}, ..., X_{t+T})$ does not depend on t, regardless of the value of T. In other words, the sequence $\{X_t : t = 1, 2, ..\}$ is identically distributed. However, since stationarity is an aspect of the underlying process rather than a single realisation, it can be fairly complex to determine whether the data we have collected is stationary or not. Therefore, a weaker form of stationarity known as covariance stationary suffices. A process is said to be second order or covariance stationary if its expectation and variance are constant and its covariance between X_t and X_{t+k} is only dependent on k and not on the location of the initial time period, t.

For a stochastic process X_t , this can be formally shown as [7]:

$$E[X_t] = \mu \quad \forall t$$

$$Var(X_t) = \gamma_0 < \infty \quad \forall t$$

$$Cov(X_t, X_{t-k}) = \gamma_k \quad \forall t, \forall k$$
(2.6)

As an example, assume that the process X_t follows the random walk model(Naive) then we can show this as:

$$X_t = X_{t-1} + \epsilon_t \tag{2.7}$$

Since ϵ_t in the above equation is white noise $(\epsilon \sim N(0, \sigma^2))$, it is assumed to be uncorrelated with X_{t-1} , that is $E(\epsilon_t | X_{t-1}, X_{t-2}, ...) = 0$. Therefore, one can show the $Var(X_t)$ as:

$$Var(X_t) = Var(X_{t-1}) + Var(\epsilon_t) \Leftrightarrow Var(X_t) = Var(X_{t-1}) + \sigma^2$$
(2.8)

In order for X_t to be a stationary process we need $Var(X_t) = Var(X_{t-1})$, so that the variance of the process is not dependent on time t. However, as we can see in (2.8), $Var(X_t)$ is increasing at each time step with σ^2 . Therefore, the random walk process X_t is non-sationary as long as $Var(\epsilon_t) \neq 0$. We can extend this example by assuming X_t follows an AR(1) process, so that:

$$X_t = c + \theta_1 X_{t-1} + \epsilon_t \tag{2.9}$$

Hence, it is easy to see that for example $X_{t-1} = c + \theta_1 X_{t-2} + \epsilon_{t-1}$ and $X_{t-2} = c + \theta_1 X_{t-3} + \epsilon_{t-2}$ which by substitution and some simplification we can show that $X_t = c(\theta_1^2 + \theta_1) + \theta_1 \epsilon_{t-1} + \theta_1^2 \epsilon_{t-2} + \epsilon_t + \theta_1^3 X_{t-3}$. Therefore, by indefinitely continuing to substitute all X_{t-k-1} back into X_{t-k} , for k = 0,1,2,..., we can show X_t as [45]:

$$X_t = c \sum_{i=0}^{\infty} \theta_1^i + \sum_{i=0}^{\infty} \theta_1^i \epsilon_{t-i}$$
(2.10)

When $|\theta_1| < 1$, this equation converges to:

$$X_t = \frac{c}{1-\theta_1} + \sum_{i=0}^{\infty} \theta_1^i \epsilon_{t-i}$$
(2.11)

therefore, when $|\theta_1| < 1$, we can show the moments of X_t as:

$$E[X_{t}] = \frac{c}{1-\theta_{1}} + \sum_{i=0}^{\infty} \theta_{1}^{i} E[\epsilon_{t-i}] = \frac{c}{1-\theta_{1}}$$

$$Var(X_{t}) = Var(\sum_{i=0}^{\infty} \theta_{1}^{i} \epsilon_{t-i}) = \frac{\sigma^{2}}{1-\theta_{1}^{2}}$$

$$Cov(X_{t}, X_{t-k}) = E[X_{t}X_{t-k}] = \frac{\sigma^{2}\theta_{1}^{k}}{1-\theta_{1}^{2}}$$
(2.12)

Therefore, we have established that the process in (2.9) is stationary when $|\theta_1| < 1$. However when $|\theta_1| \ge 1$, then the infinite sum in (2.10) will not converge and therefore we would not be able to achieve stationarity. When $|\theta_1| \ge 1$ then the process in (2.9) is said to contain a unit root. We can turn this process into a stationary one by differencing the process in time. A process which contains a unit root and needs to be differenced *d* times to become stationary is said to be of integrated order *d*, denoted as I(d). For example the random walk model in (2.7) is of integrated order 1, I(1). Since by differencing this process once we can achieve a stationary process, $\Delta X_t = (X_t - X_{t-1}) = X_{t-1} + \epsilon_t - X_{t-1} = \epsilon_t$ which is the white noise process. The simplest way to test for unit root for the time series X_t is to consider the process in (2.9). Then one can form the following hypothesis test, with the null hypothesis being the existence of a unit root in X_t :

$$H_0: \theta = 1 \tag{2.13}$$
$$H_1: \theta < 1$$

This is the basis for the Dicky-Fuller test of a unit root, which further extends to the Augmented Dicky-Fuller test that we use to test for stationarity of our data [14].

2.6 Augmented Dicky-Fuller Test

An extension to (2.9) could be made by subtracting each side of the equation by X_{t-1} and letting $\rho = \theta - 1$, which leads to

$$(X_t - X_{t-1}) = \Delta X_t = c + \rho X_{t-1} + \epsilon_t$$
(2.14)

The hypothesis test remains exactly the same as (2.13). However, in terms of ρ , it can be shown as

$$H_0: \rho = 0 \tag{2.15}$$
$$H_1: \rho < 0$$

One can perform a simple t-test with a test statistic of $t = \frac{\hat{\rho}}{s:e(\hat{\rho})}$ to decide on the rejection or the acceptance of H_0 ; $\hat{\rho}$ is the estimated ρ , and $\hat{s.e}(\hat{\rho})$ is the estimated standard error of this parameter. Usually when performing a one-sided t-test for the hypothesis of type (2.15), we require $t < -t_{CV,n,\alpha}$ in order to reject the null and accept H_1 , where $t_{CV,n,\alpha}$ is the critical value drawn from the t-distribution for sample size n at the significance level α . However, under the Dicky-Fuller framework, this critical value is not drawn from the t-distribution and instead is drawn from what is called the Dicky-Fuller distribution [13]. The critical values in this table are much larger in absolute value than their t-distribution counterparts, making the rejection of (H_0) much more difficult. This prevents us from wrongly concluding against the presence of unit root in data. For example, the conventional critical value drawn from the t-distribution for a sample size of 100 at the 5% significance level is -1.96. Yet under the Dicky-Fuller distribution for models of type (2.14), this is in the region of -2.89. The Augmented Dicky Fuller (ADF) test, first discussed by [38], extends (2.14) to the following equation so that it also includes the lagged values of X_t :

$$\Delta X_t = c + \rho X_{t-1} + \beta t + \sum_{i=1}^k \delta_i \Delta X_{t-i} + \epsilon_t$$
(2.16)

where β is the time trend coefficient and k is the lag order chosen. This ensures that ΔX_t in (2.14) are uncorrelated, and it also captures the possibility that X_t may be characterised by a higher order autoregressive process than the one used in (2.14).

2.7 Autocovariance (AC) & Autocorrelation Function (ACF)

In statistics, the general second moment of a stochastic process X_t is defined as the covariance between X_t and X_{t+k} for different values of t and k. For a stationary process with a finite constant mean (first moment), such that $E[X_t] = \mu$, this could be defined as:

$$Cov[X_t, X_{t+k}] = E[(X_t - \mu)(X_{t+k} - \mu)] = \gamma_k, \quad for \quad k = 0, 1, 2, \dots$$
(2.17)

This is called the autocovariance of the series. It reduces to the variance (σ^2) when k = 0. The set of these autocovariance coefficients denoted by γ_k , $\forall k = 0, 1, 2, ...$ form the autocovariance function of that series. If these autocovariance coefficients are standardised,

one can obtain the autocorrelation coefficients of the series such that:

$$\rho_k = \frac{\gamma_k}{\gamma_0}, \quad for \quad k = 0, 1, 2, \dots$$
(2.18)

These coefficients measure the correlation between X_t and its past lagged values X_{t-k} . The set of these ρ_k 's constitute what is referred to as the autocorrelation function, or the ACF. When using time-series data, the sample autocovariance $(\hat{\gamma}_k)$ and the sample autocorrelation coefficients $(\hat{\rho}_k)$ of the realised time series, for example, $x_1, x_2, ..., x_n$, can be computed by the following two equations respectively.

$$\hat{\gamma_k} = \frac{\sum_{t=1}^{n-k} (x_t - \bar{x})(x_{t+k} - \bar{x})}{n}, \quad for \quad k = 0, 1, 2, \dots$$
(2.19)

$$\hat{\rho_k} = \frac{\hat{\gamma_k}}{\hat{\gamma_0}}, \quad at \quad each \quad lag \quad k.$$
(2.20)

When the $\hat{\rho}_k$ s are plotted against k = 0, 1, ..., they form what is referred to as the sample autocorrelation function. This plot is a very helpful tool used for analysing and identifying patterns in data, particularly when identifying the order q of MA(q) models. The reason for this is that the theoretical ACF of MA(q) processes only shows a significant correlation up to lag q, and therefore the sample ACF can be plotted for any given data to identify this order [7]. As an example, see Figure (2.1), where we have simulated 500 observations from a MA(1) model and plotted the autocorrelation function.



Figure 2.1: SACF of 500 simulated observations from a MA(1) process

The blue horizontal lines in Figure 2.1 are two standard errors away from zero, which indicates whether the autocorrelations at each single lag k are significantly different from zero at the 95% confidence level. These standard errors can be approximated by

$$s.e(\hat{\rho}_k) = \hat{\sigma}(\hat{\rho}_k) = \sqrt{\frac{(1+2\sum_{i=1}^{k-1}\rho_i^2)}{n}}$$
(2.21)

where n is the number of observations used in fitting. This figure displays the cut-off of autocorrelation coefficients after lag 1, meaning that the process only has a significant correlation between its present value and its previous lag, which is an agreement with the fact that q = 1.

2.8 Partial Autocorrelation Function (PACF)

In this section we will introduce another tool that is primarily used for identifying the order p of AR(p) models. The following regression is used to compute the partial autocorrelation coefficients, and more importantly, their sample counterparts $\hat{\pi}_{kk}$, for different lags k:

$$X_t = \pi_{k1} X_{t-1} + \pi_{k2} X_{t-2} + \dots + \pi_{kk} X_{t-k}$$
(2.22)

As is evident from (2.22), the last regression coefficient is considered as (π_{kk}) . This value shows the correlation between X_t and X_{t+k} after accounting for the correlation at other lags. [5] If we plot these sample partial autocorrelation coefficients $\hat{\pi}_{kk}$ against k = 0, 1, ...,then the resulting correlogram is called the sample partial autocorrelation function. The theoretical PACF of an AR(p) process cuts off after lag p. This means that the PACF only shows a significant correlation up to lag p for this type of model at the desired significance level. This is why it is such a useful tool for identifying the order of these models. In Figure 2.2, we show the SPACF of 500 simulated observations from an AR(1) model that exhibit this property very clearly.



Figure 2.2: SPACF of 500 simulated observations from an AR(1) process

The horizontal lines in Figure 2.2 are two standard errors away from zero, which shows the significance of each correlation at each lag at the 95% confidence level. This standard error is very simply approximated as

$$s.e(\hat{\pi_{kk}}) = \hat{\sigma}(\hat{\pi_{kk}}) = \frac{1}{\sqrt{n}}$$
(2.23)

As can be observed from Figure 2.2, the sample partial autocorrelation coefficients are not significantly different from zero at the 5% level beyond lag 1. This is in agreement with the fact that p = 1 in our example.

2.9 ARMA Process: The Box-Jenkins Methodology

A mixture of both an autoregressive process of order p as described by (2.2) and a moving average process of order q explained by (2.5) is called an ARMA(p,q) process. It is shown by

$$X_{t} = c + \sum_{i=1}^{p} \theta_{i} X_{t-i} + \sum_{j=1}^{q} \phi_{j} \epsilon_{t-j} + \epsilon_{t}$$
(2.24)

The autoregressive moving average *ARMA* methodology developed by Box and Jenkins [5], has enormous popularity in many research areas. This is a recursive algorithm that consists of three main steps, 1.Identification, 2.Estimation and 3.Verification, that eventually lead to forecasting. This methodology requires a lot of knowledge and expertise, especially in the "Identification" step:

Step 1. The forecaster identifies which model fits the data best. In ARMA(p,q) models, this translates into identifying the order of the model, or in other words, the values of p and q. The main tools used in this identification stage are the SACF and the SPACF that were defined in Sections 2.7 and 2.8, respectively.

Step 2. Estimating the parameters of the model identified from Step 1. These are regressionlike parameters that can be estimated using least squares or maximum-likelihood estimation methods.

Step 3. Following the estimation of the parameters of the model, the residuals of the model are verified to evaluate whether the fitted model is adequate to describe the dynamics of the time-series. In this step, the residuals are checked to see whether they are white noise and are uncorrelated through time. We do so by performing a Ljung-Box Q (LBQ) test on the residuals, as described in the following section.

We use an ARMA(1,1) as one of our benchmark models. This can be shown by

$$ARMA(1,1): r_{x,t} = c + \theta_1 r_{x,t-1} + \phi_1 \epsilon_{t-1} + \epsilon_t (2.25)$$

2.10 Ljung-Box Q Test

This is the test we use for identifying autocorrelation in the model residuals. This test jointly assesses the presence of autocorrelations at individual lags.

$$H_0: \rho_1 = \dots = \rho_k = 0$$

$$H_1: \rho_i \neq 0 \qquad \text{for at least one } i \in \{1, \dots, k\}.$$
(2.26)

The test hypothesis is given by (2.26), where the null hypothesis states that errors are not serially correlated with each other and the observed correlations up to lag k are significantly no different from zero. The test statistic is given by

$$Q = n(n+2)\sum_{m=1}^{k} \frac{\hat{\rho}_m^2}{n-m}$$
(2.27)

where n is the sample size and $\hat{\rho}_m$ are the sample autocorrelations at lag m, which we defined in (2.18). Q under the null follows a χ_k^2 for which the critical values can be obtained from the χ^2 table [25].

2.11 Forecasting Accuracy Measures

All different forecasting models may and most likely will produce different results. Therefore in this section we define which measures we use to compare the accuracy of each model. We have chosen measures commonly used in the field, which are:

Mean Squared Error

This is one of the most commonly used measures in forecasting. This measure squares the errors and thus gives more weight to large deviations. If we assume X_f to be a vector of

length n of forecasted values and X_a a vector of its corresponding actual values, then we can describe the mean square error (MSE) to be

$$MSE = \frac{\sum_{t=1}^{n} (X_{a,t} - X_{f,t})^2}{n}$$
(2.28)

Root Mean Squared Error

This is the square root of MSE, and the only reason we choose this measure is because it is in the same units as the measured variable and therefore can be interpreted directly.

$$RMSE = \sqrt{MSE} \tag{2.29}$$

Mean Absolute Error

This measure is different from its MSE counterpart in that its underlying loss function is linear rather than quadratic. The MAE for a vector of n forecasts, for example X_f and their actual value counterparts X_a , can be computed by

$$MAE = \frac{\sum_{t=1}^{n} |X_{a,t} - X_{f,t}|}{n}$$
(2.30)

In the case of all three of the measures defined above, the values obtained from each measure for each forecasting value are compared, and the model that corresponds to the smallest value in these measures is considered to be the more accurate one relative to the other forecasting models.

Theil's U
This statistic measures the forecasting ability of the specified model compared to the random walk (pure guessing). If the value of Theil's U statistic is less than 1, then we can conclude that our forecasting model is statistically better than guessing the future. If the Theil's U statistic is equal to 1 then we can conclude there is no difference between the results of the forecasting method and the random walk. If the Theil's U statistic is greater than 1, then we can conclude that our forecasting model performs poorly, and better forecasts can be obtained by using the naive model (2.1). Theil's U statistic can be computed by

$$U = \sqrt{\frac{\sum_{t=1}^{n-1} \left(\frac{x_{f,t+1} - x_{a,t+1}}{x_{a,t}}\right)^2}{\sum_{t=1}^{n-1} \left(\frac{x_{a,t+1} - x_{a,t}}{x_{a,t}}\right)^2}}$$
(2.31)

where $x_{f,t}$ represents the forecasted values at t, and $x_{a,t}$ is the actual value at time t. Note that n is the sample size.

2.12 Data

In this section we describe all the FX data used in this thesis. Our study focuses on three of the most liquid and major currency pairs: GBP/USD, EUR/USD and AUD/USD. We have chosen these pairs not only due to their size but also to their time-zone and geographical differences. Although the FX market operates 24 hours a day, different time zones between Australia and Europe means that AUD predominantly experiences more trading activity during Australian trading hours rather than during European trading hours, and vice-versa. The other reasoning for choosing these pairs was because of their volatility and the daily range they experience. Whilst GBP/USD and EUR/USD are considered to be highly volatile pairs with a large average daily range, AUD/USD is somewhat less erratic and on average covers a smaller range during a given trading day. Therefore we chose to look at these specific currencies to see whether any differences or similarities could be experienced between the structure of their highs and lows, and whether our forecasting models are consistent and produce plausible results under these different conditions. Also, to keep things comparable and consistent, we have chosen all of our currency pairs to be quoted in USD.

GBP/USD:

Dates from: 02/01/1990 to: 31/12/2012

Number of observations, N: 5979

AUD/USD:

Dates from: 02/01/1993 to: 31/12/2012

Number of observations, N: 5206

EUR/USD:

Dates from: 04/01/1999 to: 31/12/2012

Number of observations, N: 3649

As noted previously, the Forex market operates 24 hours a day and we therefore have to specify which time zone our daily data corresponds to. Each day (24 hour cycle) in all of the exchange rate data used in the rest of this thesis starts from 22:00 UK time (17:00 New York time), when the market is considered to close and open simultaneously. Note that the FX market closes on Friday 22:00 UK time and reopens on Sunday 22:00 UK time. The time in-between can be regarded as the weekend.

2.13 Results

In this section we present the results from the models and tests mentioned in this chapter. We first carried out an ADF test to check for the stationarity of the variables. The following table shows the results of this test at the 5 % significance levels for all three currency pairs used in this thesis.

	p_h	r_h	p_l	r_l	p_c	r_c
AUD/USD	H_0	H_1	H_0	H_1	H_0	H_1
EUR/USD	H_0	H_1	H_0	H_1	H_0	H_1
GBP/USD	H_0	H_1	H_0	H_1	H_0	H_1

Table 2.1: ADF test results for price series $\{p\}$, and return series $\{r\}$, of all variables for all data sets at 5% significant level

As it is evident from Table 2.1, the prices for highs, lows and close (p_h, p_l, p_c) fail to reject H_0 across all these exchange rates. This confirms the existence of a unit root in prices and therefore proves them to be non-stationary. However, the same table shows that log returns of these prices reject H_0 and are therefore stationary under the framework of the ADF test. These log returns can be easily computed using

$$r_t = \log(p_t) - \log(p_{t-1}) \tag{2.32}$$

These log returns also show a much more stable and bell shaped distribution when compared to their price counterparts. This can be seen in Figure 2.3, where the histogram of prices and log returns of the highs of GBP/USD are plotted as examples.



Figure 2.3: Distribution of prices and log returns: GBP/USD Highs, N = 5979

The stability of the distribution of log returns coupled with the ADF test results imply that we use the log returns instead of prices when estimating the parameters of each model. In fact, we multiply the log returns by 100 and use percentage returns when estimating the models to increase accuracy, as their original values can be very small. It has to be noted that we can easily transform these returns back to prices once the forecasts have been made so that the performance measures can be interpreted in units of price.

To carry out the Box-Jenkins model building approach, we first plot the SACF and SPACF for all our data sets. These are shown in the following figures:



Figure 2.4: SACF and SPACF of log returns of GBP/USD Highs, N = 5979



Figure 2.5: SACF and SPACF of log returns of GBP/USD Lows, N = 5979



Figure 2.6: SACF and SPACF of log returns of AUD/USD Highs, N = 5206



Figure 2.7: SACF and SPACF of log returns of AUD/USD Lows, N = 5206



Figure 2.8: SACF and SPACF of log returns of EUR/USD Highs, N = 3649



Figure 2.9: SACF and SPACF of log returns of EUR/USD Lows, N = 3649

As can be observed from Figures 2.4-2.9, the general pattern across all of them tends to be a cut-off of autocorrelation and partial autocorrelation coefficients after lag 1 for both the SACF and SPACF. This means that both the sample autocorrelation and partial correlation functions show no significance beyond lag 1. Therefore, under the framework of Box-Jenkins, this could suggest some type of a mixed *ARMA* model.

As this section is meant to form our benchmark models, we will choose the simplest model of this kind, an ARMA(1,1). We also choose to estimate an AR(1) and an AR(2) model for our high and lows. We do this for two reasons:

1. The SPACF showing possible signs of such models, and

2. The authors in [1] also considers these models when forecasting exchange rates (albeit for close prices). Therefore, as previously mentioned, the following process, along with the naive model, form our benchmark models

$$AR(1): r_{x,t} = c + \theta_1 r_{x,t-1} + \epsilon_t (2.33)$$

$$AR(2): r_{x,t} = c + \theta_1 r_{x,t-1} + \theta_2 r_{x,t-2} + \epsilon_t (2.34)$$

$$ARMA(1,1): r_{x,t} = c + \theta_1 r_{x,t-1} + \phi_1 \epsilon_{t-1} + \epsilon_t (2.35)$$

where c is the model constant, θ_i , for i = 1, 2 and ϕ_1 are the model parameters and ϵ is the

model residuals with $\mathbb{E}(\epsilon) = 0$ and $Var(\epsilon) = \sigma^2$.

We estimate these processes (excluding the naive model, which does not need any estimation) on a rolling-daily sub-sample of size 500 (we discuss our reasoning for choosing n=500 in the next Chapter). That is, at the end of each day, we drop the oldest observation in the sample and add the latest one to keep the sample size at 500 before forecasting the prices on the next day, and so on. This results in (N - 501) regressions to be calculated, and consequently the same number of out-of-sample forecasts for each exchange rate data of size N to be presented. Note the value of N for each of our data sets is given in Section 2.12. The out-of-sample forecasting performance of these benchmark models for all of our currency pairs are given in the following tables:

	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00003	0.00002	0.00003	0.00003
RMSE	0.0050	0.0050	0.0050	0.0050
MAE	0.0034	0.0034	0.0034	0.0035
Theil U	-	0.9955	0.9979	1.0030

Table 2.2: Benchmark models forecasting results : AUD/USD-Highs

	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00003	0.00003	0.00003	0.00003
RMSE	0.0056	0.0056	0.0056	0.0056
MAE	0.0037	0.0037	0.0037	0.0037
Theil U	-	0.9995	1.0040	1.0048

Table 2.3: Benchmark models forecasting results : AUD/USD-Lows

	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00005	0.00005	0.00005	0.00005
RMSE	0.0073	0.0073	0.0073	0.0073
MAE	0.0052	0.0052	0.0052	0.0052
Theil U	-	0.9942	0.9958	0.9983

Table 2.4: Benchmark models forecasting results : EUR/USD-Highs

	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00005	0.00005	0.00005	0.00005
RMSE	0.0072	0.0071	0.0071	0.0071
MAE	0.0052	0.0051	0.0051	0.0051
Theil U	-	0.9868	0.9874	0.9878

Table 2.5: Benchmark models forecasting results : EUR/USD-Lows

	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00007	0.00007	0.00007	0.00007
RMSE	0.0086	0.0085	0.0085	0.0085
MAE	0.0060	0.0060	0.0060	0.0060
Theil U	-	0.9842	0.9851	0.9862

Table 2.6: Benchmark models forecasting results : GBP/USD-Highs

	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00008	0.00007	0.00007	0.00007
RMSE	0.0087	0.0086	0.0086	0.0086
MAE	0.0061	0.0061	0.0061	0.0061
Theil U	-	0.9874	0.9879	0.9894

Table 2.7: Benchmark models forecasting results : GBP/USD-Lows

Table 2.3 clearly shows a U statistic greater than 1 for AR(2) forecasts of AUD/USD lows. This table, alongside Table 2.2, demonstrates the same property for ARMA(1, 1)forecasts of both AUD/USD highs and lows. This suggests that the mentioned processes perform worse than the naive model and therefore cannot beat the random walk. The AR(1)process forecasts, however, show U statistics of less than 1, which means that this model performs slightly better than the naive model, but this difference is so small that it is not even reflected in the MSE, RMSE and MAE results. In fact, this is also spotted in Tables 2.4-2.7, which show the performance results of EUR/USD and GBP/USD high and lows, respectively. We observe that EUR/USD and GBP/USD out-of-sample forecasts using the mentioned models all show U statistics of less than 1, with the AR(1) forecasts showing the best performance amongst all. However, these values are all still very close to 1, so although they may perform slightly better than the naive model, they do not significantly improve the forecasts when compared to the naive model. We also look at the MSE, RMSE and MAE for the EUR/USD and GBP/USD highs and lows. These results are also in agreement with Theil's U results, and at times show an improvement of 1% in out-of-sample forecasting performance of time-series models (2.33)-(2.35) when compared to the naive process. We don't consider this to be a great improvement.

Continuing under the Box-Jenkins framework, we need to verify the credibility of our benchmark models (2.33)-(2.35). We do this by performing a Ljung-Box Q test on the fitted model residuals to check for autocorrelations and to see whether they agree with their assumption of white noise. As previously mentioned, for the highs and lows of each currency pair, the fitting process consists of 500 data points. This results in obtaining the same number of in-sample residuals $\hat{\epsilon}$. Therefore to check for autocorrelations in the residuals, each time we estimate the model at time t, we perform a LBQ test on $\hat{\epsilon}_i$ for i = t - 499, ..., t. Once we perform this operation on our whole sample (N) for each currency pair's high and low, we obtain N - 501 LBQ test results for each of them. We are only interested in periods where the LBQ test fails to accept H_0 of no autocorrelation in the errors at the 5% significance level. If the LBQ test rejects H_0 in (2.26), then we record this, and once the whole sample has been tested, we can see how many times during the N - 501 tests this has occurred. The following table shows this test results for all of our currency pairs in percentage terms.

	AR(1)	AR(2)	ARMA(1,1)
AUD/USD High	13.90	12.48	6.74
AUD/USD Low	13.20	11.33	11.54
EUR/USD High	10.39	8.07	8.39
EUR/USD Low	10.51	7.75	7.91
GBP/USD High	15.55	8.85	9.47
GBP/USD Low	20.39	15.41	16.50

Table 2.8: LBQ test fails to accept H_0 of no autocorrelation % of time in (N-501) regressions at 5% significant level: Benchmark models

Table 2.8 can be interpreted as the periods, in N-501 estimations, that our benchmark models (2.33)-(2.35) have failed to fully describe our data, as there exists some autocorrelation in their errors. We can observe in Table 2.8 that all models do suffer from some type of autocorrelation in their errors during some time-periods. However, as long as this value is fairly low, it should not completely discredit our benchmark models. The table shows that the LBQ test failed to accept the null of no autocorrelation mostly for the AR(1) model, with values showing the existence of autocorrelations between 14% to around 20% of the time. The AR(2) and ARMA(1,1) models exhibited much less autocorrelation, with values mostly under 10%, with the exception of the GBP/USD lows, implying that these models may be a better fit for our data, although their forecasting results are slightly worse than those of AR(1).

In this chapter we presented all of our benchmark models and their forecasting ability. The results show that the time-series models used did not improve the naive model in terms of their out-of-sample forecasts by a great deal. However, our aim is to show that the naive model can be beaten when forecasting the high and low of exchange rates with better models. In the next chapter, we present our own models for this purpose. We compare the results from our models to the ones obtained in this chapter and show how this naive model can be improved when considering this type of data .

Chapter 3

Time-Varying ARDL Model & Dynamic Regression

3.1 Introduction

In this chapter we introduce two of the models that we have developed to predict the next day's highs and lows of the exchange rate data. The first model introduced is an autoregressive distributed lag (ARDL) model with time-varying parameters. We illustrate the assumptions that are necessary for this model to be valid and discuss how to interpret the model in case these assumptions are violated. This eventually leads to our second model, an algorithm we have developed to deal with the shortfalls of the first model and to extend its forecasting ability. This is based on the concept of dynamic regression (DR), which was introduced by [36]. We also present all the appropriate statistical tests used to validate each model. Once the models are validated and are statistically sound, we can measure their forecasting power using the various measures defined in the previous chapter. However, first we recall the concept of cross-correlation, as it forms the basis of our motivation for developing the first model.

3.2 Cross-Correlation & Sample Cross-Correlation

In this section we describe the sample cross-correlations as they are used to produce correlograms for the analysis of our variables. The sample cross-correlations measure the correlation between two time-series observations x_t and y_t at different lags k. This could be used to determine if one variable is leading the other and whether there is a significant relationship between the variables at other lags. If we assume the time-series observations are both stationary, with a sample mean of (\bar{x}, \bar{y}) , then we can calculate their sample cross-covariance function as

$$C_{xy}(k) = \begin{cases} \frac{\sum_{t=1}^{n-k} (x_t - \bar{x})(y_{t+k} - \bar{y})}{n} & \text{for } k = 0, 1, \cdots, n-1\\ \frac{\sum_{t=1-k}^{n} (x_t - \bar{x})(y_{t+k} - \bar{y})}{n} & \text{for } k = -1, \cdots, -n+1 \end{cases}$$
(3.1)

The sample cross correlations at each lag are computed by dividing the cross covariances by the product of the standard deviations of the two series, so that

$$\rho_{xy}(k) = \frac{C_{xy}(k)}{\sqrt{C_{xx}(0)C_{yy}(0)}}$$
(3.2)

where $C_{xx}(0)$ and $C_{yy}(0)$ are the sample variances of x and y respectively.

3.3 Time-Varying ARDL Model

We introduce the following Time-Varying ARDL model for the highs and lows of exchange rates:

$$r_{h,t} = \beta_{0,t} + \beta_{h,t}r_{h,t-1} + \beta_{l,t}r_{l,t-1} + \beta_{c,t}r_{c,t-1} + \epsilon_t$$
(3.3)

$$r_{l,t} = \alpha_{0,t} + \alpha_{h,t}r_{h,t-1} + \alpha_{l,t}r_{l,t-1} + \alpha_{c,t}r_{c,t-1} + \epsilon_t$$
(3.4)

where $r_{h,t}$, $r_{l,t}$ and $r_{c,t}$ denote the log returns of the "Highs", "Lows" and "Close" at time t, respectively. ϵ_t are the model residuals at time t. One can observe from the proposed models in (3.3) and (3.4) that although they describe different dependent variables, they share the same regressors. These independent variables in each model consist of the log returns of the close, high and low, all at lag 1. In Model (3.3), the autoregressive component is the $r_{h,t-1}$, whilst the distributed lag variables are $r_{l,t-1}$ and $r_{c,t-1}$. This changes for Model (3.4), as $r_{l,t-1}$ becomes the autoregressive variable and $r_{h,t-1}$ becomes a distributed lag component. The first motivation for choosing these variables was to build a model that was purely described by its own price action and price information. The daily high, low and close prices can be very informative and describe a lot of the daily trading activities surrounding the asset. The high and the low supply information regarding the maximum and minimum prices reached for the day could have occurred at any time during that day, whilst the close price is considered a spot price that is recorded at the same time every day. Therefore it can be directly compared to its previous values and can indicate how the market has moved relative to its direct previous values. We have to note that due to the 24 hour operation of the FX market, almost always the closing price at t equals the opening price at time t+1. Therefore we only consider the close prices here. To form a more solid reasoning for choosing these variables, we look at the cross-correlations between each of them, which are shown in the following figures:



Figure 3.1: SCC between log returns of GBP/USD Highs and Lows, N = 5979



Figure 3.2: SCC between log returns of GBP/USD Highs and Close, N = 5979



Figure 3.3: SCC between log returns of GBP/USD Lows and Close, N = 5979



Figure 3.4: SCC between log returns of AUD/USD Highs and Lows, N = 5206



Figure 3.5: SCC between log returns of AUD/USD Highs and Close, N = 5206



Figure 3.6: SCC between log returns of AUD/USD Lows and Close, N = 5206



Figure 3.7: SCC between log returns of EUR/USD Highs and Lows, N = 3649



Figure 3.8: SCC between log returns of EUR/USD Highs and Close, N = 3649



Figure 3.9: SCC between log returns of EUR/USD Lows and Close, N = 3649

The blue horizontal lines in Figures 3.1-3.9 are two standard deviations away from zero, which indicates whether the cross-correlations at each lag are significantly different from zero at the 95% confidence level. These figures demonstrate that across all currency pairs, the results are consistent. We can observe from these figures that the close returns have a significantly large correlation with both the highs and the lows at lag -1. This can be interpreted as the close returns leading both their high and low counterparts. A different structure is spotted in the cross-correlations between the highs and the lows. We can see these cross-correlations in Figures 3.1, 3.4 and 3.7 for GBP/USD, AUD/USD and EUR/USD respectively. The figures mentioned show there is a significantly high correlation between r_l and r_h at both lags 1 and -1, which could mean that both series lead and lag each other. These results are in agreement with our earlier logic and therefore lead to the proposed models in (3.3) and (3.4). The other property of this model is that it is time-varying, and therefore the models parameters $\beta_{j,t}$ and $\alpha_{j,t}$ are considered to change over time (hence the t subscript) and are estimated on a rolling time window basis. This means that once new observations become available at the end of the day, the models re-estimate their parameters by dropping the oldest observation in the sample and adding the latest observation in the sample from the other end of the time-series, keeping the sample size fixed. This ensures that the model parameters are up to date and the sample is only representative of the latest n observations.

3.4 Parameter Estimation:Ordinary Least Squares

We use the method of OLS to estimate the coefficients $\beta_{0,t}$, $\beta_{h,t}$, $\beta_{l,t}$, $\beta_{c,t}$ in (3.3) and $\alpha_{0,t}$, $\alpha_{h,t}$, $\alpha_{l,t}$, $\alpha_{c,t}$ in (3.4). The OLS estimator minimizes the sum of squared prediction errors, for example this can be shown for (3.3) to be:

$$\sum_{i=2}^{n+1} (r_{h,i} - \hat{\beta}_0 + \hat{\beta}_h r_{h,i-1} + \hat{\beta}_l r_{l,i-1} + \hat{\beta}_c r_{c,i-1})^2$$
(3.5)

where $\hat{\beta}_j$ are the OLS estimators of β_j for j = 0, h, l, c. Therefore, we minimize (3.5) by taking it's derivative with respect to the *jth* regression coefficient β_j and setting it equal to zero, such that:

$$\frac{\partial}{\partial\beta_j} \sum_{i=2}^{n+1} (r_{h,i} - \hat{\beta}_0 + \hat{\beta}_h r_{h,i-1} + \hat{\beta}_l r_{l,i-1} + \hat{\beta}_c r_{c,i-1})^2 = -2 \sum_{i=2}^{n+1} r_{j,i-1} (r_{h,i} - \hat{\beta}_0 + \hat{\beta}_h r_{h,i-1} + \hat{\beta}_l r_{l,i-1} + \hat{\beta}_c r_{c,i-1}) = 0$$
(3.6)

for j = 0, h, l, c. Where, for $j = 0, r_{0,i-1} = 1, \quad \forall i$.

We can simplify the above equation in matrix form, but first to avoid confusion in notation we let $\mathbf{Y} = r_{h,i}, \forall i = 2, ..., n + 1$ to be the vector of the dependent variables of size n, and $\mathbf{X} = [1, r_{h,i-1}, r_{l,i-1}, r_{c,i-1}], \forall i = 2, ..., n+1$ to be the $n \times 4$ matrix containing the independent variables. Therefore, from (3.6) it follows that:

$$-2\mathbf{X}'(\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}) = \mathbf{0}_{4\times 1}$$

$$\Leftrightarrow \quad \mathbf{X}'\mathbf{Y} = \mathbf{X}'\mathbf{X}\hat{\boldsymbol{\beta}}$$
(3.7)
$$\Leftrightarrow \quad \hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$$

where $\hat{\boldsymbol{\beta}}$ is a 4 dimensional vector containing the OLS estimates $\hat{\beta}_0, \hat{\beta}_h, \hat{\beta}_l, \hat{\beta}_c$. Similarly we obtain the coefficients estimates $\hat{\boldsymbol{\alpha}}$ for (3.4):

$$\hat{\boldsymbol{\alpha}} = (\boldsymbol{Z}'\boldsymbol{Z})^{-1}\boldsymbol{Z}'\boldsymbol{L}$$
(3.8)

where $\boldsymbol{L} = r_{l,i}, \forall i = 2, ..., n + 1$ is the vector of the dependent variable and

 $\mathbf{Z} = [1, r_{l,i-1}, r_{h,i-1}, r_{c,i-1}], \forall i = 2, ..., n + 1$ is the $n \times 4$ matrix containing the independent variables.

3.5 Main Assumptions

In this section, we state the necessary assumptions for the proposed ARDL model. These are general time-series regression assumptions that should not be violated for this type of model. [45]

Assumption 1. Linearity and weak dependence

Firstly, we have to assume that the time series $r_{h,t}$, $r_{l,t}$ are both stationary and weakly dependent and follow the linear models in (3.3) and (3.4), respectively.

Assumption 2. Zero conditional mean

This assumption states that the errors are uncorrelated with the explanatory variables at all times t. We assume that these explanatory variables are contemporaneously exogenous, which can be expressed by $E(\epsilon_t | r_{h,t}, r_{l,t}, r_{c,t}) = 0$

Assumption 3. No perfect collinearity

Collinearity exists in the sample when the independent variables $r_{h,t-1}$, $r_{l,t-1}$, $r_{c,t-1}$ are constant or are a perfect linear combination of each other. If perfect collinearity is spotted between the independent variables, then this could lead to spurious and meaningless results. In the coming sections we will present the techniques we use to test for multicollinearity in our data.

Assumption 4. Homoskedasticity of errors

Time-series samples are said to be homoskedastic if they have a finite constant variance σ^2 across the whole sample, whilst if this is not the case, they are called heteroskedastic. In this assumption we require the error terms to be homoskedastic and therefore $Var(\epsilon_t | r_{h,t}, r_{l,t}, r_{c,t}) = \sigma^2$. However, time-series data usually show signs of heteroskedasticity, which results in the residuals from the regression exhibiting this property. In the coming sections, we will discuss the appropriate tests that can be carried out to test for homoskedasticity of errors and how to rectify the model if heteroskedasticity is apparent.

Assumption 5. No autocorrelation

This is the assumption of white noise that the errors from the model have to represent. This ensures that the errors are not correlated with each other through time, and therefore $E(\epsilon_t \epsilon_k | r_{h,t} r_{l,t} r_{c,t}, r_{h,k} r_{l,k} r_{c,k}) = 0, \forall t \neq k$. If the model residuals were in fact auto-correlated, then it would suggest that the error terms are not completely random, and some variations and the dynamics of the dependent variables are not fully captured by the model and are therefore present in the errors. However, it is highly likely for time series-regression models, especially models of type (3.3) and (3.4) that contain autoregressive and distributed lag components, to demonstrate autocorrelation in the errors even if the underlying model has no mis-specification. Therefore, in the coming sections we show how to perform the Ljung Box Q test to correctly test for the existence of autocorrelation in these types of models. We will also discuss the remedies used for these types of violations a bit later.

In [45], the author shows that for ordinary least square estimates to be consistent, we only need Assumptions 1-3 to be satisfied. Assumptions 4 and 5, however, have to be met to make sure the usual t-test and F-tests used in the regression analysis can be carried out. In the next section we will introduce the Heteroskedastic and Autocorrelation Consistent (HAC) estimators needed in case Assumptions 4 and 5 are violated.

3.6 Statistical Tests & Model Validation

In this section we present all the statistical tests used for the validation of our models. We first present the basic and widely known t and F-statistics. We then go on to describe some of the problems faced in choosing the correct sub-sample sizes when dealing with this type of data and show a method to overcome possible difficulties. The VIF, White and LBQ tests are then defined for checking the model violation of Assumptions 3 to 5, respectively. Finally, HAC estimators are introduced as a remedy for these violations.

3.6.1 t-test & F-test

A two-sided t-test on each of the regression coefficients of model (3.3) can be carried out with the following null and alternative hypothesis. The same tests are also carried out for the coefficients of (3.4), where β is replaced by α .

$$H_0: \beta_0 = 0 \tag{3.9}$$
$$H_1: \beta_0 \neq 0.$$

$$H_0: \beta_h = 0 \tag{3.10}$$
$$H_1: \beta_h \neq 0.$$

$$H_0: \beta_l = 0 \tag{3.11}$$
$$H_1: \beta_l \neq 0.$$

$$H_0: \beta_c = 0 \tag{3.12}$$
$$H_1: \beta_c \neq 0.$$

Under the null hypothesis, coefficients β_i , (i = 0, h, l, c) and α_i , (i = 0, h, l, c) are statistically no different from zero at the required significance level. This means that variable *i* can be dropped from the model, as it has no significant beneficial effect on the model. As suggested by its name, the test statistic under the null follows the t-distribution with n - 4 degrees of freedom.

If we want to test the significance of the overall model in (3.3), then we can form the following hypothesis where the significance of the model parameters are tested collectively rather than individually.

$$H_0: \beta_0 = \beta_h = \beta_l = \beta_c = 0$$

$$H_1: \beta_i \neq 0 \quad \text{for at least one } i \in \{0, h, l, c\}.$$
(3.13)

This is known as an F-test, and as the name suggests, the test statistic under the null follows a F(3, n - 4). The same test is carried out for the parameters in (3.4)

3.6.2 Choosing the Correct Sample Size

In our study, we are solely focusing on the out-of-sample predictability of these models, and therefore have to split our data samples described in Section 2.12 into smaller sub-samples. This ensures we have enough data to evaluate the forecasting ability of our models. Also, as previously mentioned, our models use a constant sample size whilst rolling daily to update their parameters. Therefore, the choice of this sub-sample size is very important, as the number of observations used in the estimation procedure could have a great effect on the obtained results. This is especially the case for financial time-series data, as they usually suffer from heavy tails in their distribution, implying that for any appropriate statistical test to be valid when applied to this type of data, the sample size has to be large enough to capture these characteristics. For example, if we consider some disjoint sub-samples of size n, then when n is too small, the moments of each sub-sample can be different and not representative of the full data set. This can lead to meaningless test results and conclusions. Therefore, if we assume that our data sets represent the full population for each time-series, then it would make sense that smaller sub-samples drawn from each population should statistically be in agreement with them. For example, if our models are estimated using the whole population N and they demonstrate that all parameters are significant at the 95% confidence level, then this result should also be reflected in their sub-samples. Hence to ensure the correct n is chosen as the size of the sub-samples, we performed a comparative assessment on our data. We carried out a t-and F-test on regression models (3.3) and (3.4) estimated by different values of n = 20, 60, 120, 250 and 500 and compared the results to determine n. The test results are presented and explained in the results section.

3.6.3 White Test

$$\hat{\epsilon}^{2} = \beta_{0} + \beta_{1} r_{h,t-1} + \beta_{2} r_{l,t-1} + \beta_{3} r_{c,t-1} + \beta_{4} r_{h,t-1}^{2} + \beta_{4} r_{l,t-1}^{2} + \beta_{6} r_{c,t-1}^{2} + \beta_{7} r_{h,t-1} r_{l,t-1} + \beta_{8} r_{l,t-1} r_{c,t-1} + \beta_{9} r_{c,t-1} r_{h,t-1}$$

$$(3.14)$$

This is a test of homoskedasticity, which is based on the auxiliary regression shown in (3.14), with squared residuals as dependent variables and regressors given by the regressors of the initial model, their squares and their cross-products [43]. The White test has the advantage over other homoskedasticity tests of not assuming a specific form of heteroskedasticity.

$$H_{0}: \beta_{0} = ... = \beta_{i} = 0$$

$$H_{1}: \beta_{i} \neq 0 \qquad \text{for at least one } i \in \{0, 1, 2, \cdots, 9\}.$$
(3.15)

The hypothesis of this test is given by (3.15), where the null states that the errors are homoskedastic, with the alternative hypothesis being the heteroskedasticity of the residuals. The test statistic is given by $n.R^2$, where n represents the sample size and R^2 is the regression's coefficient of determination. The test statistic follows a χ^2_{10} .

3.6.4 Variance Inflation Factor (VIF)

The variance inflation factor (VIF) is used to detect perfect collinearity between independent variables [31]. This test is based on a rule of thumb, where test result values of less than 2 are considered adequate to conclude against the presence of collinearity [31]. On the other hand, VIF values greater than 10 confirm severe multicollinearity between the variables. To compute the VIF, we first have to regress each independent variable against the others. For our model in (3.3) this can be shown by

$$r_h = \beta_0 + \beta_2 r_l + \beta_3 r_c + z_h \tag{3.16}$$

$$r_l = \beta_0 + \beta_1 r_h + \beta_3 r_c + z_l \tag{3.17}$$

$$r_c = \beta_0 + \beta_1 r_h + \beta_2 r_l + z_c \tag{3.18}$$

where z_h , z_l and z_c are the regression errors. Once the regressions (3.16)-(3.18) have been estimated using ordinary least squares, their R^2 can be used to compute the VIF of each corresponding variable, such that

$$VIF_{i} = \frac{1}{1 - R_{i}^{2}} \tag{3.19}$$

where *i* corresponds to each variable, i = h, l, c. This test is also carried out for the model in (3.4).

3.6.5 HAC Estimators

When heteroskedasticity and autocorrelation are present in the errors, the standard errors obtained under OLS are no longer consistent, and therefore the t- and F-tests are invalid. However, this can be treated using heteroskedastic and autocorrelation consistent estimators, as defined in [45], which provide new estimates of the standard errors of model coefficients β_i and $\alpha_i, (i = 0, h, l, c)$. As an example we show the calculations for HAC estimators of β_i . This starts by estimating the model (3.3) by ordinary least squares. We denote the standard errors for each coefficient as $s.e_{OLS}(\hat{\beta}_i)$, which can be computed by the following equation for each r_h, r_l and r_c as

$$s.e_{OLS}(\hat{\beta}_i) = \sqrt{\frac{\frac{1}{n-4}\sum_{t=1}^n \hat{\epsilon}_t^2}{\sum_{t=1}^n (r_{i,t} - \bar{r}_i)^2}}$$
(3.20)

where for $i = h, l, c, \bar{r_i}$ represents the sample mean. We then estimate the standard error $\hat{\sigma}$ of the models by

$$\hat{\sigma} = \sqrt{\frac{\sum_{t=1}^{n} \hat{\epsilon}_t^2}{n-4}} \tag{3.21}$$

Lastly, we estimate the residuals z_t from regressing each independent variable r_i against each other as shown in (3.16)-(3.18). Once these values are obtained, we can use

$$\hat{v} = \sum_{t=1}^{n} \hat{a}_t^2 + 2\sum_{h=1}^{g} [1 - \frac{h}{g+1}] (\sum_{t=h+1}^{n} \hat{a}_t \hat{a}_{t-h})$$
(3.22)

where

$$\hat{a}_t = \hat{z}_t \hat{\epsilon}_t, \quad t = 1, \cdots, n. \tag{3.23}$$

This leads to the serial-correlation robust standard errors of each β_i , which can be computed by

$$s.e(\hat{\beta}_i) = \left[\frac{s.e_{OLS}(\hat{\beta}_i)}{\hat{\sigma}}\right]^2 \sqrt{\hat{v}}$$
(3.24)

In (3.22), g is called the truncation/bandwidth parameter. It takes an integer value that can be pre-specified or automatically estimated as discussed in [2]. This parameter controls how many terms are included to correct for autocorrelation. We use the Newey-West method, as discussed in their paper [32], such that $g = 4 \times (n/100)^{2/9}$. Note that g has to be rounded to the nearest integer.

3.6.6 Bayesian Information Criterion (BIC)

The Bayesian Information Criterion (BIC) was first introduced by [39]. The BIC is a type of model-selection criterion used to search over a select number of candidate models in order to find the best model for a given set of data. The computation of BIC is based on the empirical log-likelihood function. This model-selection criterion can be shown as

$$BIC = -2lnL + pln(n) \tag{3.25}$$

where

L is the maximized likelihood function of each of the candidate models M_i ,

n is the sample size used for fitting the models, and

p is the number of independent parameters estimated in the model.

The model M_i that produces the lowest value of BIC is considered to be the best-suited model based on this criterion. The main advantage of this criterion is that it avoids overparameterisation of the model by adding the penalty term pln(n). We use the BIC to help us choose the best-fitted model when we introduce our dynamic regression algorithm later on in the chapter.

3.7 Dynamic Regression

So far, we have discussed the necessary assumptions that need to hold for our ARDL models to be valid and meaningful. However, due to the nature of time-series data, Assumption 4 and Assumption 5 in Section 3.5 can often be violated for these types of time-series regression. To overcome this violation, we introduced the HAC estimators which provide new estimates for the standard errors of the estimated parameters, under the violation of these assumptions. This was done so we can obtain new values of t- and F- statistics. In this section we discuss a different approach that can be used to deal with auto correlation in the errors. We introduce an algorithm that uses dynamic regression from [36], in cases where our tests show the existence of autocorrelation in our model errors.

We have to note that we still carry out the appropriate tests introduced earlier on in the chapter to validate that Assumption 1-4 from Section 3.5 still hold.

If serial correlation is present in the residuals of the original regression, then by modelling these residuals using an ARMA(p,q), we can extend our model to one with no serial correlations. As an example, let's assume the errors in our regression in (3.3) for the forecasting of highs are autocorrelated. First we recall this regression model:

$$r_{h,t} = \beta_{0,t} + \beta_{h,t}r_{h,t-1} + \beta_{l,t}r_{l,t-1} + \beta_{c,t}r_{c,t-1} + \epsilon_t \tag{3.26}$$

If the errors (ϵ_t) in this regression are auto-correlated, such that:

$$E(\epsilon_t \epsilon_k | r_{h,t} r_{l,t} r_{c,t}, r_{h,k} r_{l,k} r_{c,k}) \neq 0 \text{, for at least one } k \neq t$$
(3.27)

Then we can model these errors themselves as an ARMA(p,q) process, described in Equation (2.24) and extend the original ARDL model in (3.3). For example, if these autocorrelated errors ϵ_t follow an ARMA(1, 1) process, then we can show this as:

$$\epsilon_t = \theta_1 \epsilon_{t-1} + \phi_1 u_{t-1} + u_t \tag{3.28}$$

Where, u_t is a random shock (white noise).

Now that we have modelled the errors of the model themselves we can replace this back

into(3.3), to have:

$$r_{h,t} = \beta_{0,t} + \beta_{1,t}r_{h,t-1} + \beta_{2,t}r_{l,t-1} + \beta_{3,t}r_{c,t-1} + \epsilon_t =$$

$$\beta_{0,t} + \beta_{1,t}r_{h,t-1} + \beta_{2,t}r_{l,t-1} + \beta_{3,t}r_{c,t-1} + \theta_1\epsilon_{t-1} + \phi_1u_{t-1} + u_t$$
(3.29)

where we use OLS to estimate the model parameters as explained in Section 3.4.

Using this idea, we developed the following algorithm for our work:

Step 1. Estimate the models in (3.3) and (3.4) using a sub-sample of size 500.

Step 2. Perform a LBQ test on the model residuals to check for autocorrelations as described in Section 2.10. If the LBQ test results show no autocorrelation in the model residuals, then go to Step 5; otherwise, go to Step 3.

Step 3. In the presence of autocorrelations, estimate an ARMA(p,q) model for (ϵ_t) for all combinations of $p = 0, 1, \dots, 5$ and $q = 0, 1, \dots, 5$, totalling 35 models (as we exclude the (0,0) case). In this step we are modelling the residuals themselves as an ARMA process to deal with the existence of autocorrelation. This is based on the theory of dynamic regression as discussed by [36]

Step 4. Compute the BIC using (3.25) for all 35 fitted models, and choose the p and q values that correspond to the lowest BIC result. This step will make sure we choose the correct fitted model for our auto-correlated residuals.

Step 5. Forecast the day ahead dependent variables (tomorrow's highs and lows).

Step 6. Drop the oldest observation in the sub-sample in exchange for the latest observation and go back to Step 1.

When we perform this algorithm on our historical data, it will result in performing N-n-1 regressions. Please note that N is the size of our data-sets, which was defined in Section 2.12.

3.8 Results

In this section we present the results and the analysis of all the models mentioned in the previous sections.

As we discussed in Section 3.5, one of the main assumptions of our ARDL model is that there is no multicollinearity between our variables. To test for this, we computed the VIF values for all variables across all of our currency pairs. These values can be found in the following table:

	Highs	Lows	Close
EUR/USD	1.6884	1.6592	1.6307
GBP/USD	1.7760	1.7667	1.7841
AUD/USD	1.6379	1.7293	1.7392

Table 3.1: VIF values

It can be observed in Table 3.1 that the VIF values do not exceed the suggested rule of thumb value of 2 cited by [31]. This means that our variable do not exhibit any perfect multicollinearity. [31] also suggests that a correlation matrix between the regression variables can be used to check for multicollinearity. He claims that values greater than 0.75 can be considered a sign of severe multicollinearity. We show the correlations matrix for variables of each currency pair in the following tables:

	Highs	Lows	Close
Highs	1.0000	0.5664	0.5558
Lows	0.5664	1.0000	0.5447
Close	0.5558	0.5447	1.0000

Table 3.2: Correlation between variables EUR/USD

	Highs	Lows	Close
Highs	1.0000	0.5546	0.5582
Lows	0.5546	1.0000	0.5899
Close	0.5582	0.5899	1.0000

Table 3.3: Correlation between variables AUD/USD

	Highs	Lows	Close
Highs	1.0000	0.5864	0.5919
Lows	0.5864	1.0000	0.5890
Close	0.5919	0.5890	1.0000

Table 3.4: Correlation between variables GBP/USD

Tables 3.2-3.4 demonstrate that the correlations between all of our variables are around 0.5, supporting the assumption of no multicollinearity in our data.

In Section 3.6.2 we discussed the importance of choosing the right sample size when performing a regression on time-series data. The following 6 tables show some regression diagnostic results for a variety of sample sizes, with n ranging from n = 20 to n = N, where N represents the full sample values of N = 5979 for GBP/USD, N = 5206 for AUD/USD and N = 3649 for EUR/USD.

	n=20	n=60	n=120	n=250	n=500	N
t-test (β_h)	79.39	46.11	15.21	2.97	0.00	0.00
t-test (β_l)	91.98	79.30	60.05	28.42	15.57	0.00
t-test (β_c)	39.54	3.55	0.00	0.00	0.00	0.00
t-test-HAC (β_h)	69.59	41.64	15.12	2.58	0.00	0.00
t-test-HAC (β_l)	85.16	74.96	57.53	33.00	23.86	0.00
t-test-HAC (β_c)	31.29	2.38	0.00	0.00	0.00	0.00
F-test	42.41	3.30	0.00	0.00	0.00	0.00
LBQ Test	4.95	6.00	5.39	7.70	20.88	100.00
White Test	0.00	4.17	7.77	17.56	48.87	100.00

Table 3.5: Diagnostic results-ARDL model: GBP/USD Highs

	n=20	n=60	n=120	n=250	n=500	N
t-test (α_h)	90.68	77.85	64.71	38.51	14.42	0.00
t-test (α_l)	78.13	42.80	12.77	0.00	0.00	0.00
t-test (α_c)	34.79	4.11	1.52	0.00	0.00	0.00
t-test-HAC (α_h)	84.02	74.35	62.07	47.50	27.05	0.00
t-test-HAC (α_l)	67.22	34.20	12.05	0.73	0.00	0.00
t-test-HAC (α_c)	28.40	4.70	1.50	0.00	0.00	0.00
F-test	37.46	0.44	0.00	0.00	0.00	0.00
LBQ Test	6.11	6.61	8.72	17.28	39.69	100.00
White Test	0.00	3.67	7.68	21.46	42.92	100.00

Table 3.6: Diagnostic results-ARDL model: GBP/USD Lows

	n=20	n=60	n=120	n=250	n=500	Ν
t-test (β_h)	70.70	19.37	3.88	0.00	0.00	0.00
t-test (β_l)	87.73	63.29	38.97	16.42	2.60	0.00
t-test (β_c)	30.65	1.03	0.00	0.00	0.00	0.00
t-test-HAC (β_h)	56.15	19.12	4.65	0.00	0.00	0.00
t-test-HAC (β_l)	80.02	57.89	40.39	19.04	6.00	0.00
t-test-HAC (β_c)	24.59	0.70	0.00	0.00	0.00	0.00
F-test	33.57	0.47	0.00	0.00	0.00	0.00
LBQ Test	4.96	2.29	2.47	7.09	15.91	100.00
White Test	0.00	3.12	5.07	14.21	36.31	100.00

Table 3.7: Diagnostic results-ARDL model: EUR/USD Highs

	n=20	n=60	n=120	n=250	n=500	Ν
t-test (α_h)	93.08	79.24	59.89	27.52	8.10	0.00
t-test (α_l)	80.04	38.71	6.83	0.00	0.00	0.00
t-test (α_c)	28.06	0.86	0.00	0.00	0.00	0.00
t-test-HAC (α_h)	86.27	74.94	57.88	30.14	21.03	0.00
t-test-HAC (α_l)	67.64	28.29	6.07	0.00	0.00	0.00
t-test-HAC (α_c)	21.89	0.67	0.00	0.00	0.00	0.00
F-test	31.26	0.20	0.00	0.00	0.00	0.00
LBQ Test	6.12	6.55	7.26	12.36	21.79	100.00
White Test	0.00	2.40	5.50	20.69	37.23	100.00

Table 3.8: Diagnostic results-ARDL model: EUR/USD Lows

	n=20	n=60	n=120	n=250	n=500	N
t-test (β_h)	75.93	34.01	6.16	0.34	0.26	0.00
t-test (β_l)	88.95	71.04	48.46	25.11	13.67	0.00
t-test (β_c)	37.86	1.81	0.00	0.00	0.00	0.00
t-test-HAC (β_h)	67.71	30.18	7.26	1.72	1.06	0.00
t-test-HAC (β_l)	81.56	64.86	46.06	26.50	16.62	0.00
t-test-HAC (β_c)	31.51	1.59	0.00	0.00	0.00	0.00
F-test	38.98	0.74	0.00	0.00	0.00	0.00
LBQ Test	4.94	7.23	10.56	17.24	36.60	100.00
White Test	0.00	8.80	20.28	39.35	61.85	100.00

Table 3.9: Diagnostic results-ARDL model: AUD/USD Highs

	n=20	n=60	n=120	n=250	n=500	Ν
t-test (α_h)	94.19	85.56	71.62	49.12	22.87	0.00
t-test (α_l)	75.89	34.29	9.46	5.53	0.00	0.00
t-test (α_c)	33.89	1.40	0.00	0.00	0.00	0.00
t-test-HAC (α_h)	87.54	83.50	71.33	59.52	31.71	0.00
t-test-HAC (α_l)	65.75	30.09	12.78	9.59	1.32	0.00
t-test-HAC (α_c)	26.42	1.90	0.00	0.00	0.00	0.00
F-test	36.78	0.19	0.00	0.00	0.00	0.00
LBQ Test	4.94	5.29	6.96	10.58	29.65	100.00
White Test	0.00	5.81	12.98	37.17	58.47	100.00

Table 3.10: Diagnostic results-ARDL model: AUD/USD Lows

Tables 3.5 - 3.10 each show the amount of times the t-test and the HAC counterparts have failed to show the significance of a particular variable. The tables mentioned also show how many times the F-test has failed to reject the null, as displayed in (3.13). As can be observed, we also show the amount of times the LBQ test has failed to show no autocorrelation in the residuals. Similarly, we demonstrate the number of times the white test has failed to show homoskedasticity in the residuals. All the values in these tables are in percentage terms, and all the tests are carried out at the 5% significance level. The values in these tables are calculated in the following way: we perform N-n-1 regressions on each exchange rate's high and low using our proposed model in Section 3.3. We then perform a t-test and an HAC t-test on each variable as well as a F-test for each regression. If the t-tests, the HAC t-tests and F-tests for each regression fail to reject their corresponding null hypotheses at the 5% significant level, then we will record that period as, for example, H = H + 1, and otherwise as H = H + 0. Once we have done this for all the N - n - 1 regressions, we use $\left(\frac{H}{N-n-1}\right) \times 100$ to see how frequently we fail to fully validate our models. We also perform an LBQ test and a White test on the in-sample errors for each regression for which we record the number of times that these tests have failed to reject their corresponding null hypothesis at the 95% confidence level. This is done in the same way, and the results can be interpreted as the existence of heteroskedasticity and serial correlation in the model residuals.

As can be observed from Tables 3.5 - 3.10, the full sample N shows that our model is statistically sound, as it shows rejection of the null hypothesis of our t- and F-tests. This result is consistent across all currency pairs and their corresponding high and lows. However, the full sample also shows that the residuals of the model suffer from both serial correlation and heteroskedasticy, which is not desirable. This is the reason we have also computed the HAC standard errors and subsequently performed a t-test using these standard errors. These HAC t-test also rejects the null of no statistical significance of variables at the 5% level for the full sample. Therefore any sub-sample we choose should be statistically as close to this full population as possible to be a correct representation of our data. As we expected, by increasing the number of data points used in estimation, we achieve statistical results closer to those of the full population. When n = 20, we fail to show the significance of the model variables, in some cases up to almost 95% of the time, which is of course an indication that the sample size is not representative of the full population. As we increase the sample size to 60, 120, 250 and 500, we show that the statistical significance of our models is improved. Signs of serial correlation and heteroskedasticity in the errors also increase as the sample size increases. This is in agreement with the full sample results, which is what we are hoping to achieve. We also see that when n = 60, 120, 250, 500, the F-test, which tests for the significance of the model parameters collectively, has rejected the null after each regression.

This is a very promising result.

The t-test results in the tables show that when the dependent variables in our model are the highs $r_{h,t}$, then $r_{l,t-1}$ s fails to show significance for some of our N - n - 1 regressions. The same behaviour is spotted when $r_{l,t}$ s are the dependent variables being forecasted, as $r_{h,t}$ s fail to show significance for some of our N - n - 1 regressions. However, as we have already mentioned, this is not the case for the full sample. Therefore, by choosing n = 500, we have significantly improved on this effect so that it occurs less frequently than it does when considering smaller sample sizes. Although increasing the sample size to even greater than 500 would make these values closer to the full sample, since we are only concentrating on the day ahead predictions, we also like to keep our sample size as small as possible. We therefore accept n = 500 to be adequate enough for this purpose, as it does demonstrate in most cases the same characteristics of the full sample.

The following tables, (3.3) for highs, and (3.4) for the lows, demonstrate the forecasting results for our models. We also show the forecasting results of the algorithm we have developed using the concept of dynamic regression (DR) as discussed in Section 3.7.

	ARDL	DR-Algorithm	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00002	0.00002	0.00003	0.00002	0.00003	0.00003
RMSE	0.0040	0.0040	0.0050	0.0050	0.0050	0.0050
MAE	0.0028	0.0028	0.0034	0.0034	0.0034	0.0035
Theils U	0.8034	0.7909	-	0.9955	0.9979	1.0030

Table 3.11: Forecasting results: AUD/USD Highs

	ARDL	DR-Algorithm	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00002	0.00002	0.00003	0.00003	0.00003	0.00003
RMSE	0.0046	0.0046	0.0056	0.0056	0.0056	0.0056
MAE	0.0031	0.0030	0.0037	0.0037	0.0037	0.0037
Theils U	0.8273	0.8222	-	0.9995	1.0040	1.0048

Table 3.12: Forecasting results: AUD/USD Lows

	ARDL	DR-Algorithm	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00003	0.00003	0.00005	0.00005	0.00005	0.00005
RMSE	0.0057	0.0056	0.0073	0.0073	0.0073	0.0073
MAE	0.0042	0.0042	0.0052	0.0052	0.0052	0.0052
Theils U	0.7722	0.7648	-	0.9942	0.9958	0.9983

Table 3.13: Forecasting results: EUR/USD Highs

	ARDL	DR-Algorithm	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00003	0.00003	0.00005	0.00005	0.00005	0.00005
RMSE	0.0056	0.0055	0.0072	0.0071	0.0071	0.0071
MAE	0.0041	0.0040	0.0052	0.0051	0.0051	0.0051
Theils U	0.7794	0.7625	-	0.9868	0.9874	0.9878

Table 3.14: Forecasting results: EUR/USD Lows

	ARDL	DR-Algorithm	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00005	0.00005	0.00008	0.00007	0.00007	0.00007
RMSE	0.0071	0.0069	0.0087	0.0086	0.0086	0.0086
MAE	0.0050	0.0048	0.0061	0.0061	0.0061	0.0061
Theils U	0.8077	0.7907	-	0.9874	0.9879	0.9894

Table 3.15: Forecasting results: GBP/USD Lows

	ARDL	DR-Algorithm	Naive	AR(1)	AR(2)	ARMA(1,1)
MSE	0.00005	0.00004	0.00007	0.00007	0.00007	0.00007
RMSE	0.0068	0.0066	0.0086	0.0085	0.0085	0.0085
MAE	0.0049	0.0048	0.0060	0.0060	0.0060	0.0060
Theils U	0.7855	0.7644	-	0.9842	0.9851	0.9862

Table 3.16: Forecasting results: GBP/USD Highs

As we can see from Tables 3.11-3.16, we have improved the forecasting results of our benchmark models significantly for all of our exchange rates. Our proposed ARDL regression shows an increase in the accuracy over the naive model by 18-24% for all exchange rates and significantly beats all of our benchmark models from the previous chapter. As we have already discussed in our literature review, one of the only papers that we found to attempt to forecast the high and lows of exchange rate data was [41]. They show the MSE of their forecasting models to be 0.005 for the GBP/USD high and lows. By contrast, we obtained an MSE of 0.00005 for the same asset, proving the superiority of our models.

The RMSE and the MAE can both be interpreted in the units of the currency in question, and as we can see, these values are shown to be very small for all of our rates, ranging between 0.0040 and 0.0071. This again confirms the robustness of our models.

The DR-Algorithm shows improved forecasts of the ARDL model on the basis of their Theil's U statistic for all of our data sets. However, in terms of the other accuracy measures, it tends to have the same accuracy power as the ARDL regression. This model performs the best for GBP/USD highs, where it increases the accuracy of its ARDL counterpart by around 2-3% when looking at their Theil's U and RMSE statistics. However, the DR algorithm shows very little improvement compared to our ARDL model when applied to AUD/USD, as it seems to only improve the Theil U value by less than 1%, which is not a great deal but still plausible.

3.9 Extension

We recall the two regression models built for the purpose of forecasting the highs and lows of exchange rates:

$$r_{h,t} = \beta_{0,t} + \beta_{h,t}r_{h,t-1} + \beta_{l,t}r_{l,t-1} + \beta_{c,t}r_{c,t-1} + \epsilon_t$$
(3.30)

$$r_{l,t} = \alpha_0 + \alpha_{h,t} r_{h,t-1} + \alpha_{l,t} r_{l,t-1} + \alpha_{c,t} r_{c,t-1} + \epsilon_t$$
(3.31)

In the following tables, we show the forecasting accuracy of the models when $r_{l,t-1}$ and $r_{c,t-1}$ are dropped from (3.30). We also show the performance measures for times when $r_{h,t-1}$ and $r_{c,t-1}$ are dropped from (3.31).
	Original	Lows Removed	Close Removed
MSE	0.00005	0.00005	0.00006
RMSE	0.0068	0.0069	0.0079
MAE	0.0049	0.0050	0.0057
Theils U	0.7855	0.7971	0.9214

Table 3.17: Modified models forecasting results: GBP/USD Highs

	Original	Highs Removed	Close Removed
MSE	0.00005	0.00005	0.00007
RMSE	0.0068	0.0071	0.0082
MAE	0.0049	0.0050	0.0058
Theils U	0.7855	0.8133	0.9444

Table 3.18: Modified models forecasting results: GBP/USD Lows

	Original	Lows Removed	Close Removed
MSE	0.00003	0.00003	0.00005
RMSE	0.0056	0.0058	0.0068
MAE	0.0041	0.0043	0.0050
Theils U	0.7794	0.7914	0.9250

Table 3.19: Modified models forecasting results: EUR/USD Highs

	Original	Highs Removed	Close Removed
MSE	0.00003	0.00003	0.00005
RMSE	0.0056	0.0056	0.0067
MAE	0.0041	0.0041	0.0049
Theils U	0.7794	0.7859	0.9368

Table 3.20: Modified models forecasting results: EUR/USD Lows

	Original	Lows Removed	Close Removed
MSE	0.00002	0.00002	0.00002
RMSE	0.0040	0.0041	0.0047
MAE	0.0028	0.0029	0.0033
Theils U	0.8034	0.8131	0.9292

Table 3.21: Modified models forecasting results: AUD/USD Highs

	Original	Highs Removed	Close Removed
MSE	0.00002	0.00002	0.00003
RMSE	0.0046	0.0047	0.0053
MAE	0.0031	0.0031	0.0036
Theils U	0.8273	0.8361	0.9529

Table 3.22: Modified models forecasting results: AUD/USD Lows

When analysing the forecasting accuracy of each exchange rate, we see that our original models are the best performers. This is specifically evident in the Theil's U statistic, which shows the lowest value when the original model was used across all of our currency pairs. These tables demonstrate that when lows are removed for the forecasting of highs, and viceversa, the resulting MAE and RMSE are slightly worse than the original model, but not by a great amount. However, we see the importance of using the close when forecasting highs and lows, as the removal of this variable has the greatest effect on forecasting accuracy.

We can therefore conclude that the removal of the close log returns from (3.30) and (3.31) has a significant effect on the forecasting accuracy of the models. The removal of $r_{l,t-1}$ from (3.30) and of $r_{h,t-1}$ from (3.31) does affect forecasting accuracy, but not as much as the removal of the close. However, the original model performs best within the context of forecasting.

Further on, we analyse the forecasting accuracy of our models during different periods with different market volatilities. This is done by first computing the daily range for all of our sample data. This is simply calculated by $P_{h,t} - P_{l,t}$, where $P_{h,t}$ is the high price for day tand $P_{l,t}$ is the corresponding low price of that day. We then subcategorised volatility into 4 different quantiles. We then measured the forecasting accuracy of the models depending on the quantile the preceding day's range fell into. So depending on which quantile the range at t - 1 fell into, we measure the forecasting accuracy at t. The forecasting results are all shown in the table below, where Q1, Q2, Q3 and Q4 correspond to when the range at t - 1was in the first quantile, in the second quantile, and so on.

	Q1	Q2	Q3	Q4
GBP/USD Hi	0.0053	0.0060	0.0064	0.0090
GBP/USD Lo	0.0051	0.0062	0.0064	0.0099
EUR/USD Hi	0.0046	0.0050	0.0054	0.0072
EUR/USD Lo	0.0043	0.0047	0.0054	0.0072
AUD/USD Hi	0.0027	0.0028	0.0038	0.0058
AUD/USD Lo	0.0027	0.0034	0.0039	0.0070

Table 3.23: Forecast RMSE: When previous days' range falls into different quantiles

The above table clearly shows that as the range (volatility) increases, the proceeding days' forecasts tend to suffer from larger errors. This shows the model has more forecasting power when exchange markets are less volatile, where the high and low are relatively close to each other. This finding could possibly help with the improvement of the model, as one can measure the day's volatility and act accordingly to control for larger deviations between the highs and the lows.

In the figures below, we show the sample probability distribution of log returns of our exchange rates. These figures show both the actual and forecasted values against each other for comparison. It is very clear from all tables across all currency pairs that the forecasted model has a higher peak. The distribution of the forecasted values tend to shift more towards the right when considering the highs, and the opposite is true when forecasting lows.



Figure 3.10: Distribution of actual V forecasted log returns: GBP/USD Highs



Figure 3.11: Distribution of actual V forecasted log returns: GBP/USD Lows



Figure 3.12: Distribution of actual V forecasted log returns: EUR/USD Highs



Figure 3.13: Distribution of actual V forecasted log returns: EUR/USD Lows



Figure 3.14: Distribution of actual V forecasted log returns: AUD/USD Highs



Figure 3.15: Distribution of actual V forecasted log returns: AUD/USD Lows

In this chapter we presented two models for FX data highs and lows. The first model presented was a regression model for each of the highs and lows, whilst the second model was an algorithm to overcome the problem of autocorrelation in our first model. This was based on the concept of dynamic regression developed by [36]. The forecasting results show that our models beat the random walk and all the benchmark models from the previous chapter. This means that we have shown that the highs and lows of exchange rates are in fact predictable, unlike the close prices, as suggested by EMH and all the research in that area.

Chapter 4

Triggering Method for Exchange Rate Trading Via Feedback Control

4.1 Introduction

In this Chapter we aim to apply the saturation-reset linear feedback controller proposed by [3] to exchange rates . The authors of [3] used control theory to build a linear feedback control system for the trading of stocks. This controller determines the amount invested in stock each time a trade is being made. [21] proposed a triggering method that signals a buy or sell trade in which the saturation-reset linear feedback controller is then used to obtain the amount invested before entering the trade. These trading strategies are first tested on an idealized market. This serves as a filter to see whether a strategy is worthy of being tested on real-world data. This is due to the fact that back-testing on real-world data can be very time consuming as well as financially expensive. This idealized market mentioned in [3], has the following assumptions:

1. Continuous and costless trading : Assume the trader reacts to observed prices immediately and there are no costs for trading (eg. transaction costs, stamp duty tax, etc.)

2. Continuously differentiable prices : Assume prices are continuously differentiable in the time interval of interest; there are no price gaps.

3. Perfect liquidity : Assume there is no gap between the bid and the ask price.

4. Trader as price-taker : The trader has no influence on the price by entering a trade. The above assumptions set our motivation for applying the saturation-reset linear feedback controller to the trading of exchange rates rather than to stocks as used by [3]. These assumptions are far from reality, specifically when discussing the stock market, due to the facts that:

1. Stock trading incurs transaction fees.

2. Stamp duty tax may have to be paid for holding stocks in certain countries.

3. Big trading volumes could potentially move the stock price, which makes it difficult for the trader to be the price-taker.

4. The stock market is only open for a certain number of hours during the day. This, coupled with the size of the market, means prices usually have gaps and are therefore not continuous.

The transaction and stamp duty tax costs alone make the application of such controllers to the FX market a lot more appealing. The performance results obtained by trading stocks have no value unless transaction costs have been taken into account. These fees could wipe out all potential profits that back-testing may have suggested. However, when foreign exchange prices are considered, the above assumptions are much more reflective of their real-world behaviour. That is:

1. There are no transaction fees or stamp duties to be paid when trading currencies.

2. The enormity of the foreign exchange market also means that the trader has a higher chance of being the price-taker. Therefore, a large transaction made by a trader or a fund in this market would probably still not have the same immediate effect on the price that it might have in the stock market.

3. The FX market is highly liquid and is open 24 hours a day, which means prices have a higher probability of being continuously differentiable.

It has to be noted that we have neglected the interest and margin assumptions mentioned by [3]. We assume no money is being borrowed from the broker and no money is to be invested at the risk-free rate. Therefore, we assume all the money is in the trading account and cannot be invested elsewhere, even if it is not used for trading. We also assume that no position is held open overnight. Therefore there is no interest to be earned or paid by holding one currency relative to the other.

We continue our work in this chapter by defining the controller as proposed by [21]. We then expand this triggering method and build a new method to incorporate the daily high and low prices. To test this method, we propose a complete stochastic model for the prices of the highs, lows, of their range, and of the close price. This stochastic model forms the idealized market we tested our method on. We then provide the results of our triggering method for both the idealized market and real-world data. We conclude the chapter by combining the out-of-sample forecasts obtained by our ARDL model in the previous chapter and our triggering method to illustrate a real trading strategy.

4.2 Saturation-Reset Linear Feedback Controller

In this section we show the feedback controller as described in [3]. We assume trading is taking place at time $t \in [0, T]$. We denote:

- **1.** The account value at time t as V_t , with initial value V_0
- **2.** The amount invested in asset at time t as I_t , with initial value I_0
- **3.** The trading gain at time t as g_t
- **4.** The exchange rate price at time t as p_t
- 5. The maximum amount the trader is willing to invest, also known as the saturation limit,

as I_{max} .

This controller uses a pure gain K to increase or decrease the amount invested depending on how the strategy is performing subject to the saturation limit. For example, when the strategy is doing well and the trading gains are positive, the amount invested in the asset is increased, and vice-versa when the gains are negative. This can be shown by

$$\frac{dI}{dt} = K \frac{dg}{dt} \tag{4.1}$$

The equation above holds for when $I_t < I_{max}$, that is, before we have reached our saturation limit. However, during saturation, the amount invested becomes $I_t = I_{max}$. The incremental trading gains are given by:

$$dg = \frac{dp}{p}I\tag{4.2}$$

The incremental contributions to the account value are:

$$dV = dg + m(V - I)dt (4.3)$$

where m can be interpreted as both the risk-free interest rate and the brokerage's marginal rate. However, as we assume no margin or investment at the risk-free rate, in our case, we let m = 0. Therefore, by using (4.2), we can simplify (4.3) to:

$$dV = dg = \frac{dp}{p}I \tag{4.4}$$

We let ρ_t denote the daily returns of the exchange rate at time t such that:

$$\rho_t = \frac{p_t - p_{t-1}}{p_{t-1}} \tag{4.5}$$

Therefore, from (4.1), (4.2) and (4.3), we can show how the trading gain, account value and amount invested are updated daily as prices evolve from time t to t + 1:

$$g_{t+1} = g_t + \rho_t I_t$$

$$V_{t+1} = V_t + \rho_t I_t$$

$$I_{t+1} = I_t + K(g_{t+1} - g_t)$$
(4.6)

where, $I_t, K > 0$.

The system in (4.6) is only valid when we are buying at the exchange rate, also known as going long. However when we are going short (selling the asset), the previous amount invested, I_t , and the controller gain, K, both have to be negative. For example, if $\rho_t < 0$, meaning the price had decreased, then for a short trade it has to be that $g_{t+1} - g_t > 0$, reflecting a profit. However, as can be observed in (4.6), this can only be ensured by letting $I_t, K < 0$.

Since the system also incorporates a saturation limit I_{max} , we can show I_{t+1} as:

$$I_{t+1} = \min\{I_t + K(g_{t+1} - g_t), I_{max}\}$$
(4.7)

$$I_{t+1} = max\{-I_t - K(g_{t+1} - g_t), -I_{max}\}$$
(4.8)

Equation (4.7) is the updated equation for I when we are long, and (4.8) applies to periods that we sell the asset.

4.3 The Trade Triggering Method

In this section we will introduce our triggering method that determines whether we buy (go long) or sell (go short) the currency pair. The triggering method is built to be used in the feedback controller by [21] that was explained in Section 4.2.

The triggering method we introduce is similar to the one proposed by [21] that uses daily close prices to trigger a trading signal, determining whether to buy, sell or hold (no trade). We propose adapting their techniques to the high and low prices. We believe this gives us an advantage, since by having two separate price series to analyse, we can make a more informed decision about price movement. This can be regarded as an extra confirmation to trigger a trade in a certain direction. The method is based on estimating the drift ($\hat{\mu}$) and volatility ($\hat{\sigma}$) of the highs and lows on a daily rolling basis. To be more specific, for an estimation window of size *n*, these values are estimated at time *t* as:

$$\hat{\mu_{j,t}} = \frac{1}{ndt} \sum_{i=1}^{n} \rho_{j,t-i}$$
(4.9)

$$\hat{\sigma_{j,t}} = \sqrt{\frac{1}{n-1} \sum_{i=1}^{n} (\rho_{j,t-i} - \hat{\mu_{j,t}})^2}$$
(4.10)

where j = h represents the highs, j = l represents the lows and $\rho_{j,t}$ can be obtained by (4.5). dt is the incremental time interval and is taken to be 1/252, representing 252 trading days in a year. This means that the results obtained for the estimates of the drift ($\hat{\mu}$) and volatility ($\hat{\sigma}$) are annualized.

We have provided the estimated values, $\mu_{j,t}$ and $\sigma_{j,t}$ for all of our data in Appendix B. As it can be observed from Figures B.1-B.12, the estimated values are not completely constant through time as they exhibit shocks at some points. Therefore, one could argue against the stationarity of the underlying process. However, as we have discussed in Section 2.5, we have carried a unit root test and have proven that our data contains no unit root with 95% confidence. This suffices for the purpose of our work as a small degree of non-stationarity will almost remain due to the nature of financial time-series data and the fact that they exhibit shocks from time to time. Aside, from the few spikes in these figures we can see that for the rest of the times, these estimated values look fairly stable and constant.

We form a $(1 - \alpha)100\%$ confidence interval for the estimate of the drift of high and low $\mu_{j,t}$ such that

$$[\hat{\mu}_{j,t} - t_{\alpha/2}(n-1)\frac{\hat{\sigma}_{j,t}}{\sqrt{n}}, \hat{\mu}_{j,t} + t_{\alpha/2}(n-1)\frac{\hat{\sigma}_{j,t}}{\sqrt{n}}]$$
(4.11)

where $t_{\alpha/2}(n-1)$ is the $100(1-\alpha/2)$ percentile of the t-distribution with n-1 degrees of freedom. The interval in (4.11) shows the upper and lower limits of the estimate $\mu_{j,t}^{i}$. We denote the lower limit as L_{t}^{j} and the upper limit as U_{t}^{j} . We carry out this procedure for both the high and the low, meaning at each time step t + k, for k = 0, 1, 2..., we obtain the upper and lower confidence levels using the past n observations. This results in L_{t}^{h} and U_{t}^{h} as the lower and upper confidence limits of the drift of the highs, and L_{t}^{l} and U_{t}^{l} form their low price counterparts.

Once the daily confidence limits have been calculated for that day, we use the following triggering method to initiate a trade:

1. Buy / Going long: Buy if both the lower confidence limits are above zero, that is, when L_t^h , $L_t^l > 0$. This indicates the high and low are both drifting upwards and therefore the trend is up.

2. Sell / Going short: Sell if both the upper confidence limits are below zero, that is, when U_t^h , $U_t^l < 0$. This indicates the high and low are both drifting downwards and therefore the trend is down.

3. No trade: If the conditions for either of the above cases have not been met, then there is no signal, and no trading takes place that day.

Our trading method uses these confidence limits as a confirmation of whether the price is trending upwards or downwards. We first apply this method to our proposed idealized market, which is shown in the next section.

4.4 Idealized Market Model: A Stochastic Model for the Price Evolution of High, Low and Close

In this section we propose a stochastic model to which we apply our triggering method. This forms our so-called idealized market model. As standard practice, we study the idealized market's performance before applying our triggering methods to real-world data. The idealized market considered by [21] is a geometric brownian motion that they consider using the close price , that is:

$$dp_{c,t} = \mu_c p_{c,t} dt + \sigma_c p_{c,t} dW \tag{4.12}$$

Where $p_{c,t}$ is the close price at time t,

 μ_c is the drift of p_c , and

 σ_c is the volatility of p_c .

Our objective is to propose a synthetic price model for the highs and lows. However, in order to correctly simulate high and low prices, we need to make sure that:

$$p_{h,t} > p_{l,t} \qquad \forall t \tag{4.13}$$

Of course, if we just assume the two processes follow a stochastic process such as the geometric brownian motion in (4.12), such a condition cannot be guaranteed to be satisfied. We take advantage of the fact that the difference between p_h and p_l forms the range R at each time-step, such that:

$$R_t = p_{h,t} - p_{l,t} (4.14)$$

However, instead of assuming the range to be the linear difference between the high and low, we use the log difference of these prices instead as used by [37] for estimating the volatility. The log difference also demonstrates a more stable distribution. This is shown by:

$$R_t = ln(p_{h,t}) - ln(p_{l,t})$$
(4.15)

Therefore, by assuming that $p_{h,t}$ and R_t are two correlated stochastic processes, we propose the following model for them:

$$dp_{h,t} = \mu_h p_{h,t} dt + \sigma_h p_{h,t} dW^1$$

$$dR_t = \kappa (\theta - R_t) dt + \upsilon \sqrt{R_t} dW^2$$
(4.16)

where $p_{h,t}$ is the price of high of exchange rate at time t and follows a geometric brownian motion,

- μ_h is the drift of p_h ,
- σ_h is the volatility of p_h ,

 R_t is the Range as defined in (4.15) at time t and follows the CIR process of [11],

v is the volatility of R,

 κ denotes the speed of the mean reversion for the range process R_t , and

 θ is the long run average range also known as the long run level.

 dW^1 and dW^2 are correlated wiener processes such that:

$$E[dW^1dW^2] = \rho dt \tag{4.17}$$

Using (4.15) and (4.16), we can recover the price of low $p_{l,t}$ such that:

$$p_{l,t} = \exp(\ln(p_{h,t}) - R_t)$$
(4.18)

This ensures that the condition in (4.13) is satisfied and the simulated high and low prices are in the correct order.

For the purpose of trading, we also need to simulate close prices for each day that will be used as our point of entry/exit of daily trades. This is the price used to calculate ρ_t in (4.5), which is then subsequently used to update the controller in (4.6). The close prices $p_{c,t}$ for each day have to lie between their high and low counterparts such that $p_{l,t} \leq p_{c,t} \leq p_{h,t}$. We can ensure this by setting

$$p_{c,t} = \alpha p_{l,t} + (1 - \alpha) p_{h,t} \tag{4.19}$$

where α is a random number generated from a uniform distribution over [0,1], $\alpha \in U[0,1]$.

We use this proposed model to simulate prices and track the performance of our triggering method using the controller.

Justification and Estimation of the Process R_t

We have chosen a CIR process of the type in (4.16) to represent the log- daily range. This continuous process ensures the positivity of daily range as negative values would be mean-ingless. We have chosen a mean-reverting process to be representative of this daily range. As one can observe from the figures below, the range certainly has mean reverting properties. Figures 4.1-4.3 show how the log daily range for our data series fluctuate around their long term mean.



Figure 4.1: R_t -AUD/USD



Figure 4.2: R_t -GBP/USD



Figure 4.3: R_t -EUR/USD

We also use the Generalised Hurst Exponent to confirm our visual justification. The Generalised Hurst Exponent H_q examines if some statistical properties of a time-series scale with their observation time period. For a time-series X_t during the observation period T, the Hurst exponent can be defined by the scaling properties of the time-series structural function, $K_{q\tau}$ [12] :

$$K_{q\tau} = \frac{E[|X_{t+\tau} - X_t|^q]}{E[|X_t|^q]} \sim \tau^{qH_q}$$
(4.20)

Where, τ is the time-lag and q > 0. The general idea here is that $H_1 = 0.5$ demonstrates a Geometric Brownian Motion process. Where as $H_1 < 0.5$ demonstrates a mean-reverting series and $H_1 > 0.5$ implies a trending series. We estimated the H_1 for the log-daily range of our series and obtained the following values:

	H_1
R_t of AUD/USD	0.0637
R_t of GBP/USD	0.0487
R_t of EUR/USD	0.0384

Table 4.1: Generalised Hurst Exponent values of R_t

The Hurst Exponent values in Table 4.1 are much smaller than 0.5 which demonstrate R_t for our time-series data have mean-reverting properties. This result combined with Figures 4.1-4.3 justifies our use of the CIR process to represent R_t .

Interested readers may refer to Appendix C where we have included the MLE and OLS estimation methods of the CIR process for completeness.

4.5 Simulation Study

In this section we use Monte-Carlo simulation to simulate the prices for the processes in (4.16), (4.18) and (4.19) over the time interval [0, T].

4.5.1 Discretization

We first need to simulate $p_{h,t}$ and R_t from the processes in (4.16). To do so, we discretize the time interval into equally spaced time increments so that $0 = t_1 < t_2 < ... < t_k = T$, where $dt = t_i - t_{i-1}$. We first consider the GBM process in (4.16) for the evolution of the highs. This process can also be demonstrated in integral form as:

$$p_{h,t+dt} = p_{h,t} + \int_{t}^{t+dt} \mu_h(P_{h,u}, u) du + \int_{t}^{t+dt} \sigma_h(p_{h,u}, u) dW_u^1$$
(4.21)

Using Euler discretization we can approximate both these integrals using the left-point rule which produces:

$$\int_{t}^{t+dt} \mu_{h}(P_{h,u}, u) du \approx \mu_{h}(P_{h,t}, t) dt$$

$$\int_{t}^{t+dt} \sigma_{h}(p_{h,u}, u) dW_{u}^{1} \approx \sigma_{h}(p_{h,t}, t) \sqrt{dt} Z_{1}$$
(4.22)

where Z_1 is a random number generated from the standard normal distribution N(0, 1). Since W_t is the Wiener process, $W_{t+dt} - W_t \sim N(0, dt)$. Therefore $W_{t+dt} - W_t$ are identical in distribution with $\sqrt{dt}Z_1$ and are replaced in the discretized version. The CIR process for R_t in (4.16) can also have the following form:

$$R_{t+dt} = R_t + \int_t^{t+dt} \kappa(\theta - R_u) du + \int_t^{t+dt} v \sqrt{R_u du} dW_u^2$$
(4.23)

We approximate these integrals using the left-point rule to obtain:

$$\int_{t}^{t+dt} \kappa(\theta - R_u) du \approx \kappa(\theta - R_t) dt$$

$$\int_{t+dt}^{t+dt} \upsilon \sqrt{R_u du} dW_u^2 \approx \upsilon \sqrt{R_t dt} Z_2$$
(4.24)

where Z_2 is a random number generated from the standard normal distribution N(0, 1). However, since we assume:

$$E[dW^1dW^2] = \rho dt \tag{4.25}$$

 Z_2 in (4.24) is replaced by $(\rho Z_1 + \sqrt{1 - \rho^2} Z_2)$ to account for this correlation structure. Therefore, the discretized version of (4.16) using the Euler scheme can be shown as

$$p_{h,t+dt} = p_{h,t} + \mu_h p_{h,t} dt + \sigma_h p_{h,t} \sqrt{dt Z_1}$$

$$R_{t+dt} = R_t + \kappa (\theta - R_t) dt + \upsilon \sqrt{R_t dt} (\rho Z_1 + \sqrt{1 - \rho^2} Z_2)$$
(4.26)

The CIR process of R_t in (4.16) only yields positive results in continuous time. However, the discretized version in (4.26) may yield negative values in discrete time which is not desirable. To ensure that only obtain positive values of R_t we use the reflection scheme which replaces R_t with $|R_t|$, that is:

$$p_{h,t+dt} = p_{h,t} + \mu_h p_{h,t} dt + \sigma_h p_{h,t} \sqrt{dt Z_1}$$

$$R_{t+dt} = |R_t + \kappa(\theta - R_t) dt + \upsilon \sqrt{R_t dt} (\rho Z_1 + \sqrt{1 - \rho^2} Z_2)|$$
(4.27)

We investigate the convergence of our methods to ensure that our Euler discretized processes in (4.27) converge to the continuous processes in (4.16).

Definition 4.1. Given a sequence of discrete time approximations $\hat{X(0)}, \hat{X(h)}, ...,$ we say

 \hat{X} converges strongly to the SDE X if it converges in L^1 [20]:

$$\lim_{n \to \infty} E[|\hat{X(nh)} - X(T)|] = 0$$
(4.28)

Therefore, we say that \hat{X} has a strong order of convergence $\lambda > 0$ if

$$E[|\hat{X(nh)} - X(T)|] \le ch^{\lambda}$$
(4.29)

for some constant c and h sufficiently small.

Similarly we say that \hat{X} has a weak order of convergence β if for any $2(\beta + 1)$ continuously differentiable function g of polynomial growth, it holds that

$$|E[g(X(nh))] - E[g(X(T))]| \le ch^{\beta}$$
(4.30)

In general for any SDE of type:

$$dX_t = a(X, t)dt + b(X, t)dW$$

The Euler scheme has a weak order of convergence $\beta = 1$ and a strong order of convergence $\lambda = 0.5$ if:

1. Functions a() and b() are four times continously differentiable and their first derivatives are bounded.

2. They do not grow too fast with parameters.

However, the CIR process for R_t has Non-Lipschitz diffusion coefficients, as this process has non-negative values and contains square root coefficients. In our discretization of this process we used the reflection(symmetrised) Euler Scheme to ensure the positivity of this process in discrete time. The authors in [4] show in Theorem 2.2 on page 3 of their paper, how such a scheme for this types of process has a strong convergence rate of $O(\sqrt{h})$ and therefore $\lambda = 0.5$.

Therefore, we can use (4.27) to produce values for $p_{h,t+dt}$ and R_{t+dt} at time t + dt. We can then obtain results for $p_{l,t+dt}$ and $p_{c,t+dt}$ by plugging the simulated values into (4.18) and (4.19).

Our simulation starts at time t = 0, where we simulate the prices for the next 200 days, up to T = 200. We use the first n = 100 simulated values for estimating the confidence limit L_t^h , L_t^l , U_t^h and U_t^l using (4.11), setting $\alpha = 0.1$. Trading starts at time t = 100, for which we set the initial account value, $V_0 = 10,000$, and the initial amount invested, $I_0 = \lambda V_0$. We set $\lambda = 0.5$, and therefore, $I_0 = 5000$ with the controller gain K = 1. After each trading day we re-estimate the confidence limits by keeping n = 100. We do so by removing the oldest value in the data and replacing it with the newly observed value.

We apply this method to 10,000 simulated paths of 200 days each. Rather than choosing mu_h to be a constant, we let mu_h be a random number generated from the normal distribution with a mean of ($\mu = 0$) and a standard deviation of ($\sigma = 0.5$), i.e. $N(0, 0.5^2)$.

We set the model parameters to be $\theta = 0.01$, v = 0.05, $\kappa = 0.5$ and $\rho = 0.5$. The initial values for the simulation are chosen as $p_{h,0} = 1.1100$ and $R_0 = 0.0050$.

4.5.2 Varying The Volatility of Highs σ_h

To carry out the simulation, we look at 3 different cases for the volatility of the highs, σ_h , High, Medium or Low volatility.

Case 1: High Volatility- For this case we let σ_h be generated from the uniform distribution over the interval [0, 0.9].

Case 2: Medium Volatility- For this case we let σ_h be generated from the uniform distribution over the interval [0, 0.5].

Case 3: Low Volatility- For this case we let σ_h be generated from the uniform distribution

over the interval [0, 0.1].

The following figures show the histogram of the final account value, V_{200} , for all 10,000 paths for each of the cases in our idealized market.



Figure 4.4: Histogram of final account value for high volatility case



Figure 4.5: Histogram of final account value for medium volatility case



Figure 4.6: Histogram of final account value for low volatility case

As can be observed from Figures 4.4-4.6, the volatility of the highs has a great impact on the account value when using our triggering method to signal trades. When volatility is low, we observe significantly high positive returns; as the volatility becomes higher, the number of positive returns gets lower. This is also evident in the following table, where we show the maximum and minimum account values achieved for each case along with the percentage of the positive returns.

	High- σ_h Case	Medium- σ_h Case	Low- σ_h Case
Max Account Value	24188	20942	15634
Min Account Value	5798	6687	9311
Percentage of positive returns	60.99	71.40	92.79

Table 4.2: Idealized market account value analysis : Cases with varying volatility

4.5.3 Varying The 'Level' θ of Range

In this section we are interested in seeing how the account value changes by varying θ in (4.27). This is the long run value to which the range R_t reverts. We want to test whether our trading system performs differently for currencies with higher or lower long ranges. We

set the volatility of the highs σ_h to follow the medium case discussed in the previous section. We set the following 3 cases for θ

Case 1: High θ - For this case, we let θ be generated from the uniform distribution over the interval [0, 0.025].

Case 2: Medium θ - For this case, we let θ be generated from the uniform distribution over the interval [0, 0.01].

Case 3: Low θ - For this case, we let θ be generated from the uniform distribution over the interval [0, 0.005].



Figure 4.7: Histogram of final account value for medium volatility-high θ case



Figure 4.8: Histogram of final account value for medium volatility-medium θ case



Figure 4.9: Histogram of final account value for medium volatility-low θ case

	High- θ Case	Medium- θ Case	Low- θ Case
Max Account Value	19433	18870	21968
Min Account Value	6391	6624	7014
Percentage of positive returns	71.05	71.36	72.06

Table 4.3: Idealized market account value analysis : Cases with varying θ

The histogram of the account values for all the cases can be seen in Figures 4.7-4.9. As

is evident from these figures, the maximums and minimums obtained still show significant differences between cases, but only by a small amount. We observe from Table 4.3 that the different cases for θ do not actually impact the trading performance of the model by a great deal. The low θ case shows the best performance in terms of the maximum and minimum account value reached. The greatest number of positive returns were also achieved under this case, albeit by only 1% more than those achieved under the high θ case. These fairly close results between cases mean the triggering method should achieve similar results for currency pairs with different long run ranges.

4.5.4 Varying The 'Volatility' v of Range

So far we have observed that the volatility of highs, σ_h , has a direct impact on the trading results in our method. However, if we keep the volatility at the medium level case, we test to see whether the volatility of R_t , v has any direct effect on the performance of the system. The 3 cases for varying v are given as:

Case 1: High v: For this case, we let v be generated from the uniform distribution over the interval [0, 0.1].

Case 2: Medium v: For this case, we let v be generated from the uniform distribution over the interval [0, 0.05].

Case 3: Low v: For this case, we let v be generated from the uniform distribution over the interval [0, 0.01].



Figure 4.10: Histogram of final account value for medium volatility-high υ case



Figure 4.11: Histogram of final account value for medium volatility-medium v case



Figure 4.12: Histogram of final account value for medium volatility-low v case

	High- v Case	Medium- v Case	Low- v Case
Max Account Value	21724	20410	17803
Min Account Value	7068	6807	7102
Percentage of positive returns	71.11	70.92	71.01

Table 4.4: Idealized market account value analysis : Cases with varying v

Figures 4.10-4.12 demonstrate the fact that v has only a small impact on the performance results. All three cases show very close results for the distribution of V_{200} . This fact is also confirmed in Table 4.4, where the percentage of positive returns achieved under all three cases is shown to be fairly indifferent (70.92%-71.11%). However, the maximum account value reached seems to be decreasing as R_t becomes less volatile. The low v case shows the lowest maximum account value and the highest minimum account value reached between cases. As v gets higher, the account has the potential of achieving higher values, but it also becomes more risky as the minimum account value gets lower for higher v cases.

4.5.5 Varying The 'Speed' κ of Range

We finally study the speed of the mean reversion parameter κ of R_t . The three cases considered for this parameter are as follows:

Case 1: High κ : For this case, we let κ be generated from the uniform distribution over the interval [0, 0.9].

Case 2: Medium κ : For this case, we let κ be generated from the uniform distribution over the interval [0, 0.5].

Case 3: Low κ : For this case, we let κ be generated from the uniform distribution over the interval [0, 0.1].



These simulations are carried out for when σ_h follows its medium volatility case.

Figure 4.13: Histogram of final account value for medium volatility-high κ case



Figure 4.14: Histogram of final account value for medium volatility-medium κ case



Figure 4.15: Histogram of final account value for medium volatility-low κ case

	High- κ Case	Medium- κ Case	Low- κ Case
Max Account Value	21315	20093	20515
Min Account Value	6864	6858	6934
Percentage of positive returns	71.56	70.59	71.22

Table 4.5: Idealized market account value analysis : Cases with varying κ

The above Figures, coupled with Table 4.5, show that κ has no major effect on the

trading performance of the system. The maximum and minimum account value results seem to be mixed for the 3 different cases, with the high κ case showing the highest maximum account value. The percentages of the positive returns achieved for each case are also very close to each other. The high κ case achieved 71.56% positive returns, and the medium κ case was the worse performer, with 70.59% positive returns.

So far, we have shown that using our idealized market, the triggering method works pretty well in terms of producing positive returns under all conditions. The parameters of the range process, R_t , have no significant impact on trading performance. However, the underlying volatility of the highs $p_{h,t}$, denoted as σ_h , can directly effect the performance of the system. As volatility decreases, the system seems to produce more positive returns, whereas increasing volatility means the possibility of obtaining higher returns but with less probability. The study of the random cases seems to indicate lower volatility gives better performance mathematically. This coincides, interestingly, with news reports that the hedge fund industry, in general, produces overall higher returns in lower volatility environments.

4.6 Real-World Application

In this section we apply the triggering method introduced in Section 4.3 to our real-world data for EUR/USD, GBP/USD and AUD/USD. The data used covers rich and diverse market scenarios, which is ideal for back-testing to ensure reliable results. We recall from Section 2.12, that our data covers the following dates:

GBP/USD:From 02/01/1990 to 31/12/2012

AUD/USD:From 02/01/1993 to 31/12/2012

EUR/USD:From 04/01/1999 to 31/12/2012

We use the saturation reset controller from Section 4.2 to back-test our trading method.

We first study the performance using the same parameters we used when trading in our idealized market. We then build a more dynamic system wherein by optimisation we obtain the desired parameter values of the system before studying the trading performance.

4.6.1 Maximum Drawdown

Maximum Drawdown is a type of risk metric widely used in the world of finance. It can be defined as the maximum cumulative loss from a peak to the following trough in any given period. We use Maximum Drawdown on the account value of each strategy to analyse the risks associated with it. This can be shown by

$$MDD = \max_{i,j} \frac{V_{i \in (0,T)} - V_{j \in (i,T)}}{V_i}$$
(4.31)

where V_t is the account value at time $t \in (0, T)$

4.6.2 Back-testing With Arbitrary Parameters

In this section, we use the same arbitrary parameter values for our system that we used previously, so that:

Initial account value, $V_0 = 10000$,

Initial Investment amount, $I_0 = 5000$,

Maximum amount invested, i.e. the saturation limit, $I_{max} = 10000$,

Controller gain, K = 1,

Significance level used in(4.11), $\alpha = 0.1$,

Using daily data, we annualize our estimates, dt = 1/252 and

Estimation window, n = 100.



Figure 4.16: Account value V_t for GBP/USD : Arbitrary parameters



Figure 4.17: Account value V_t for AUD/USD : Arbitrary parameters



Figure 4.18: Account value V_t for EUR/USD : Arbitrary parameters

	GBP/USD	AUD/USD	EUR/USD
Max Account Value	12536	15523	15239
Min Account Value	8599	9061	9428
Maximum Drawdown (%)	25.26	23.58	10.23

Table 4.6: Account value performance for all currency pairs using arbitrary parameter values

We observe from Figure 4.16 that GBP/USD initially performs very poorly under our framework. During the first half of the data (approximately 3000 days), the controller seems to be incurring losses and we reach an account value of 8599 from the initial 10,000. However, the system recovers well considering the initial losses, as it reaches a maximum of 12,536 during the second half of the data. This is shown in Table 4.6, where we also observe that GBP/USD has recorded a maximum drawdown of 25.26% during this time-period, which we consider to be fairly high. From the same table, we observe that trading AUD/USD has also obtained a very high maximum drawdown value of 23.58%. However from Figure 4.17, it is evident that the system performs much better for this currency pair than for GBP/USD. We see that initially the model incurs small losses and reaches a minimum of 9061 before becoming profitable and reaching a maximum of 15,523 during the tested time period. Figure 4.18 demonstrates the fact that EUR/USD is the best performer out of the three currency

pairs we consider when using our triggering method for trading. We observe a consistent upwards pattern in the account value for EUR/USD after carrying out the back-testing. The minimum account value reached was 9428, from which the account value grew to a maximum of 15,239. The maximum drawdown of 10.23% achieved by EUR/USD is by far the lowest among the three currency pairs, confirming it as our best performer under this scheme.

4.6.3 Back-Testing Using Optimisation

In this section we use optimisation to obtain the model parameters from the system by maximising the final account value V_{end} .

We concentrate on obtaining the optimised values for the parameters K, α and dt, whilst keeping all other parameters as before, i.e.

Initial account value, $V_0 = 10,000$,

Initial Investment amount, $I_0 = 5000$,

Maximum amount invested, i.e. the saturation limit, $I_{max} = 10,000$ and

Estimation window in (4.11), n = 100.

We use Matlab's optimisation toolbox with the following upper and lower bounds constraints for our parameters:

$$K \in [1, 10]$$

 $\alpha \in [0.005, 0.3]$

 $dt \in [1/252, 1]$

We carry out the optimisation on a rolling window size of 240 days (approximately 1 trading year). Each day, upon the arrival of new information, we re-estimate these parameters whilst keeping everything else fixed. The figures below show the account value for the trading of AUD/USD, GBP/USD and EUR/USD using this approach.


Figure 4.19: Account value V_t for GBP/USD : optimised parameters



Figure 4.20: Account value V_t for AUD/USD : optimised parameters



Figure 4.21: Account value V_t for EUR/USD : optimised parameters

	GBP/USD	AUD/USD	EUR/USD
Max Account Value	10429	12292	15250
Min Account Value	8628	9602	9749
Maximum Drawdown (%)	17.27	17.48	16.05

Table 4.7: Account value performance for all currency pairs using optimised parameter values

We observe from Table 4.7 that the optimisation has reduced the maximum drawdown for GBP/USD and AUD/USD compared to those of Table 4.6. However, the opposite is true for EUR/USD, as the maximum drawdown is higher when trading was carried out using optimisation. We observe from Figure 4.21 that EUR/USD still has the best performance amongst our currency pairs even after optimisation. The triggering method seems to do well when trading this asset. However, the opposite is observed for GBP/USD. Figure 4.19 shows that after optimisation, GBP/USD fails to reach profitability, even though the maximum drawdown is reduced. The controller fails to recover from its initial losses as it does so well in Figure 4.16 when arbitrary parameters were used. AUD/USD shows a very average performance for the initial 4000 days, from which the account value suddenly spikes before wiping out the profits made.

4.7 Combining Forecasts and Triggering Method

In this section we make use of the out-of-sample forecasts made for the highs $p_{h,t}$ and the lows $p_{l,t}$ using our ARDL model in Section 3.3. We propose four different trading strategies that combine our out-of-sample forecasts and the triggering method in Section 4.3. We back-test each strategy on our three currency pairs with both arbitrary and optimised parameters. It is worth noting that the data we used still covers the dates mentioned in Section 2.12. However as the first 500 days of our data sets are used for obtaining our out-of-sample forecasts, we have 500 fewer days we can back-test on.

4.7.1 Strategy 1

Our first strategy is a modification of the triggering method from Section 4.3 so that it also includes the confidence interval for the estimated drift of forecasted highs and lows. The strategy for each day t can be explained as:

Step 1. Obtain the upper and lower estimates, U_t^j and L_t^j of $\hat{\mu}_{j,t}$, using (4.11).

Step 2. Repeat Step 1 but for the estimates of $\mu_{fj,t+1}$, where fj denotes our out-of-sample forecasts for highs(j = h) and lows (j = l). This results in the upper and lower estimates U_{t+1}^{fj} and L_{t+1}^{fj} for the day ahead, t + 1.

Step 3. Initiate a trade based on:

Buy/Going long: Signal buy if all four of our lower confidence limits are above zero $L_t^h, L_t^l, L_{t+1}^{fh}, L_{t+1}^{fl} > 0.$

Sell/Going Short: Signal sell if all four of our upper confidence limits are below zero $U_t^h, U_t^l, U_{t+1}^{fh}, U_{t+1}^{fl} < 0.$

No trade: If the conditions for buying or selling have not been met, then there is no signal and no trading is done that day.

We carry out these steps on a rolling daily basis, using the past n = 100 observations as well the past n = 100 out-of-sample forecasts.

Arbitrary Parameters

Firstly, we back-test this strategy on our set of arbitrary parameters given in Section 4.6.2.



Figure 4.22: Account value V_t for GBP/USD : Strategy 1 using arbitrary parameters



Figure 4.23: Account value V_t for AUD/USD : Strategy 1 using arbitrary parameters



Figure 4.24: Account value V_t for EUR/USD : Strategy 1 using arbitrary parameters

	GBP/USD	AUD/USD	EUR/USD
Max Account Value	11005	15251	15902
Min Account Value	8111	9366	9717
Maximum Drawdown (%)	23.46	24.41	9.95

Table 4.8: Account value performance of Strategy 1 for all currency pairs using arbitrary parameter values

Figures 4.22-4.24 show the account value V_t performance for each of our currency pairs. GBP/USD is the worst performer under this system, with the account value showing an initial loss from 10,000 to 8111 before reaching profitability. Table 4.8 confirms this by demonstrating the high maximum drawdown of 23.46% that GBP/USD suffered under this strategy. The AUD/USD results, however, show an even higher maximum drawdown on the account, but with a much higher maximum account value reached by the strategy. The best results were obtained when trading EUR/USD. The account value reaches a low of 9717 from its initial 10,000 before reaching a maximum of almost 16,000. The maximum drawdown is also under 10%, which is significantly lower than those of AUD/USD and GBP/USD.

Optimised Parameters

The optimisation techniques and objectives used here for the back-testing of strategy 1 are the same as those stated in Section 4.6.3. We show the account value performance of each currency in the below figures and table.



Figure 4.25: Account value V_t for GBP/USD : Strategy 1 using optimised parameters



Figure 4.26: Account value V_t for AUD/USD : Strategy 1 using optimised parameters



Figure 4.27: Account value V_t for EUR/USD : Strategy 1 using optimised parameters

	GBP/USD	AUD/USD	EUR/USD
Max Account Value	10333	16702	15674
Min Account Value	8478	9198	9892
Maximum Drawdown (%)	17.95	20.89	11.63

Table 4.9: Account value performance of Strategy 1 for all currency pairs using optimised parameter values

Optimisation reduced the maximum drawdown suffered by the GBP/USD and AU-D/USD accounts by a significant amount when compared to those of the arbitrary parameters. However, the opposite is true for EUR/USD, as the maximum drawdown has been increased by almost 2%. The optimisation has evidently improved the performance of AU-D/USD for Strategy 1. The account value shows it to have reached a high of 16,702 in Figure 4.26 for this currency pair. However, Figure 4.25 clearly shows that Strategy 1, specifically after optimisation, has worsened the performance of our account value when trading GBP/USD. The currency pair EUR/USD still shows strong trading performance under this strategy, even though the maximum drawdown and the maximum account value reached have both worsened when compared to those of the arbitrary parameters. We observe in Figure 4.27 that the initial losses the account value suffered are less severe when optimisation had been carried out for the trading of this pair under this strategy.

4.7.2 Strategy 2

Our second strategy still makes use of the triggering method from Section 4.3. However, we no longer initiate our trades at the close price. Instead we use the out-of-sample forecasts as our points of entry into the market. This Strategy for each day t can be described as : **Buy/Going long:** If $L_t^h, L_t^l > 0$, then we initiate a long trade at our forecasted lows $p_{fl,t+1}$. This can only occur if our forecasted lows for time t + 1 are actually reached by the market at time t + 1, $p_{fl,t+1} > p_{l,t+1}$.

Sell/Going Short: If $U_t^h, U_t^l < 0$, then we initiate a short trade at our forecasted highs $p_{fh,t+1}$. This can only occur if our forecasted highs $p_{fh,t+1}$ for time t+1 are actually reached by the market at time t+1, $p_{fh,t+1} < p_{h,t+1}$.

No trade: If the conditions for buying or selling have not been met, then there is no signal, and no trading is done that day.

Arbitrary Parameters



Figure 4.28: Account value V_t for GBP/USD : Strategy 2 using arbitrary parameters



Figure 4.29: Account value V_t for AUD/USD : Strategy 2 using arbitrary parameters



Figure 4.30: Account value V_t for EUR/USD : Strategy 2 using arbitrary parameters

	GBP/USD	AUD/USD	EUR/USD
Max Account Value	10284	10831	10423
Min Account Value	9018	8905	9675
Maximum Drawdown (%)	12.31	17.79	7.18

Table 4.10: Account value performance of Strategy 2 for all currency pairs using arbitrary parameter values

Figures 4.22-4.24 and Table 4.10 clearly show that Strategy 2 is not a profitable strategy when using the arbitrary parameters for our system. The maximum drawdown of the strategy seems to be fairly low when compared to our other strategies. However, the account value demonstrates a choppy performance with a tendency for higher negative returns.

Optimised Parameters

We apply our optimisation methods from Section 4.6.3 to strategy 2 to see whether it will improve its poor performance using the arbitrary parameters.



Figure 4.31: Account value V_t for GBP/USD : Strategy 2 using optimised parameters



Figure 4.32: Account value V_t for AUD/USD : Strategy 2 using optimised parameters



Figure 4.33: Account value V_t for EUR/USD : Strategy 2 using optimised parameters

	GBP/USD	AUD/USD	EUR/USD
Max Account Value	10000	10595	10204
Min Account Value	9195	9639	9567
Maximum Drawdown (%)	8.04	8.36	6.24

Table 4.11: Account value performance of Strategy 2 for all currency pairs using optimised parameter values

As is evident from Table 4.11 and Figures 4.31-4.33, the optimisation carried out has improved the maximum drawdown suffered by the account value under this strategy. However, clearly this strategy fails to produce any indication that it is useful when trading these currency pairs.

4.7.3 Strategy 3

The third strategy we propose is one very similar to Strategy 2. But instead of using forecasted highs and lows as the point of trade entry, we use them as the point of exit. This can be described as:

Buy/Going long: If $L_t^h, L_t^l > 0$, then we initiate a long trade at the open of the day. If our forecasted highs for time t + 1 are actually reached by the market at time t + 1, so that

 $p_{fh,t+1} < p_{h,t+1}$, we exit our position at $p_{fh,t+1}$. However, if $p_{fh,t+1} > p_{h,t+1}$, then as before, we close our position at the close of the day.

Sell/Going Short: If $U_t^h, U_t^l < 0$, then we initiate a short trade at the open of the day. If our forecasted lows $p_{fl,t+1}$ are actually reached by the market at time t + 1, so that $p_{fl,t+1} > p_{l,t+1}$, we close our trade at $p_{fl,t+1}$. However, if $p_{fl,t+1} < p_{l,t+1}$, then as before, we close our position at the close of the day.

No trade: If the conditions for buying or selling have not been met, then there is no signal, and no trading is done that day.

Arbitrary Parameters



Figure 4.34: Account value V_t for GBP/USD : Strategy 3 using arbitrary parameters



Figure 4.35: Account value V_t for AUD/USD : Strategy 3 using arbitrary parameters



Figure 4.36: Account value V_t for EUR/USD : Strategy 3 using arbitrary parameters

	GBP/USD	AUD/USD	EUR/USD
Max Account Value	10441	13312	13872
Min Account Value	8367	8769	9940
Maximum Drawdown (%)	16.79	19.6	8.89

Table 4.12: Account value performance of Strategy 3 for all currency pairs using arbitrary parameter values

Figures 4.34-4.36 show the back-testing results of our currency pairs for Strategy 3. These figures demonstrate the account value when arbitrary parameters were used. We observe that GBP/USD initially accumulates negative returns for approximately the first 2000 days. However, the account value then seems to recover well; it subsequently reaches a maximum of 10,441. AUD/USD also replicates the performance of GBP/USD for the first 2000 days, albeit with smaller losses. After that, the account value surges upwards for the following 2000 days before suffering a high maximum drawdown of 19.6% of the account value. We observe again that EUR/USD shows a strong back-testing performance, with a low maximum drawdown of 8.89%.

Optimised Parameters

We apply the daily rolling optimisation routine to Strategy 3 and obtain the following results



Figure 4.37: Account value V_t for GBP/USD : Strategy 3 using optimised parameters



Figure 4.38: Account value V_t for AUD/USD : Strategy 3 using optimised parameters



Figure 4.39: Account value V_t for EUR/USD : Strategy 3 using optimised parameters

	GBP/USD	AUD/USD	EUR/USD
Max Account Value	10000	11558	14973
Min Account Value	8612	9529	9929
Maximum Drawdown (%)	13.87	9.82	12.34

Table 4.13: Account value performance of Strategy 3 for all currency pairs using optimised parameter values

After the optimisation, the GBP/USD account value fails to go above the initial account value of $V_0 = 10,000$ during the whole back-testing period. The optimisation has significantly

decreased the maximum drawdown suffered on the AUD/USD account value. As we observe from Table 4.13, the maximum drawdown for this pair is 9.82%, compared to the 19.6% we showed in Table 4.12. However, this reduction has been costly, as the maximum account value reached is far less than those of the arbitrary parameters. The opposite is true for the account value of EUR/USD. Figure 4.39 shows that the optimisation routine has improved the performance of the EUR/USD account. Even though the maximum drawdown is slightly higher, the account value only reaches a minimum of 9929 and increases gradually to a maximum of almost 15,000.

4.7.4 Strategy 4

This strategy is a modification of Strategy 1 in which we formed the confidence intervals for both the realised observations and the forecasted high and low drift estimates. For this strategy we only form the confidence intervals for the forecasts rather than the observations. This can be explained as:

Step 1. Obtain the upper and lower estimates, U_{t+1}^{fj} and L_{t+1}^{fj} of $\mu_{fj,t+1}$ using (4.11).

Step 2. Initiate a trade based on:

Buy/Going long: Buy if both of our lower confidence limits are above zero $L_{t+1}^{fh}, L_{t+1}^{fl} > 0$. **Sell/Going Short:** Sell if both of our upper confidence limits are below zero $U_{t+1}^{fh}, U_{t+1}^{fl} < 0$. **No trade:** If the conditions for buying or selling have not been met, then there is no signal, and no trading is done that day.

Arbitrary Parameters



Figure 4.40: Account value V_t for GBP/USD : Strategy 4 using arbitrary parameters



Figure 4.41: Account value V_t for AUD/USD : Strategy 4 using arbitrary parameters



Figure 4.42: Account value V_t for EUR/USD : Strategy 4 using arbitrary parameters

	GBP/USD	AUD/USD	EUR/USD
Max Account Value	10684	15304	15699
Min Account Value	7790	9308	9621
Maximum Drawdown (%)	27.08	24.1	9.87

Table 4.14: Account value performance of Strategy 4 for all currency pairs using arbitrary parameter values

It is evident from Table 4.14 that the GBP/USD and AUD/USD accounts under this strategy, when used with the arbitrary parameters, suffer from high maximum drawdowns. AUD/USD and EUR/USD both show promising results, whereas the GBP/USD account value shows a similar pattern to those achieved under the other strategies. Figure 4.40 shows the account value initially accumulating large losses before slightly recovering. Figure 4.42 demonstrates that the EUR/USD account value increases steadily over time and reaches a maximum of 15,699 during the back-testing period.

Optimised Parameters

Finally, we carry out the daily optimisation routine to back-test strategy 4 on our data sets.



Figure 4.43: Account value V_t for GBP/USD : Strategy 4 using optimised parameters



Figure 4.44: Account value V_t for AUD/USD : Strategy 4 using optimised parameters



Figure 4.45: Account value V_t for EUR/USD : Strategy 4 using optimised parameters

	GBP/USD	AUD/USD	EUR/USD
Max Account Value	10333	14122	15663
Min Account Value	8689	9555	9797
Maximum Drawdown (%)	15.90	15.48	10.26

Table 4.15: Account value performance of Strategy 4 for all currency pairs using optimised parameter values

The optimisation has clearly decreased the maximum drawdown for GBP/USD and AUD/USD, as can be seen in Table 4.15 when compared to Table 4.14. EUR/USD still shows a low maximum drawdown of 10.26% even though it is slightly higher than when we used the arbitrary parameters to back-test this strategy. Once again, the GBP/USD account value shows disappointing results. Figure 4.43 clearly shows the account value of GBP/USD reaches a minimum of 8689 and fails to recover. EUR/USD shows consistent performance once again, while Figure 4.44 demonstrates that the AUD/USD account value remains within a tight range before it starts accumulating high positive returns after the initial 3500 days.

4.8 Comparing Performances:Risk-free, FTSE100 and S&P500

In this section we compare the performance of our proposed trigerring method and different trading strategies, to a standard buy-and-hold account value performance of the benchmark US risk-free rate: the 3 month T-bills and 2 different benchmark stock market indices: The FTSE100 and the S&P500. The data used have the following properties:

S&P 500:

Daily data covering from: 04/01/1999 to: 31/12/2012

FTSE100:

Daily data covering from: 03/01/2001 to: 31/12/2012

3 Months T-bill:

Annual data from: 04/01/1999 to: 31/12/2012

We have chosen the mentioned dates as they represent the same dates as our EUR/USD data making it simpler to compare the performances. This is with the exception of FTSE100 as we did not have access to data before 03/01/2001. We also have to point out that our 3month T-bill data is based on annual returns rather than daily ones that we use elsewhere. We show the performance of these benchmarks in the figures below, based on a Buy-and-hold strategy.



Figure 4.46: Account value V_t for 3 months US T-bills : Buy-and-hold strategy



Figure 4.47: Account value V_t for FTSE100 : Buy-and-hold strategy



Figure 4.48: Account value V_t for S&P500 : Buy-and-hold strategy

As it can be observed from the above figures, the benchmark risky assets would have performed much worse during the same period as our EUR/USD data. Figures 4.47 and 4.48 show that during this period an investment of 10,000 would have merely made any profits as the account value would have hardly gone above the initial 10,000. The return on the risk free investment is also interesting, we can see in Figure 4.46 that this risk-free account value would have grown to around 13,500. As we showed earlier, our EUR/USD would have performed much better during the same period proving that our systems and strategies(apart from Strategy 2) could potentially yield much greater returns than the risk-free return. AU-D/USD has also performed better under Strategy 1 than these benchmarks, confirming the superiority of our trading performance to an extent. However, the risk-free return performs better than most of our GBP/USD account values under our strategies. Therefore, this confirms our findings up to now that our proposed trading systems and strategies work better for certain assets than others. The reason behind this could be investigated further in the future.

In this chapter we used a saturation-reset linear feedback controller to perform exchange rate trades. We proposed a triggering method for the controller and we first studied its performance on a synthetic price model. We then used historical daily data for GBP/USD, AUD/USD and EUR/USD to study its performance on real-world data. Our data sets covered between 10 to 20 years of observations, which included the 2008 financial crisis period. Having such a rich data set that covers diverse market scenarios ensures that our back-testing results are meaningful. We combined the out-of-sample forecasts obtained for the highs and lows with the triggering method proposed to build diverse trading strategies. For comparison, we back-tested each strategy using both arbitrary parameters and optimal parameters obtained by carrying out a daily optimisation routine on the data. The results obtained were mixed for each currency pair. Overall, the EUR/USD account showed the most consistent back-testing results with positive returns. AUD/USD also showed promising results, in particular with strategies 1 and 4. However, the same strategies, when applied to GBP/USD, showed a choppy trading account value with high negative returns. We have shown that our forecasts and triggering method serve as good building blocks for developing trading strategies. However, these strategies need to be carefully back-tested for each currency pair in question, as evidently each pair maintains different market dynamics that cause them to perform differently under each strategy.

Bibliography

- Don Alexander and Lee R Thomas. Monetary/asset models of exchange rate determination: How well have they performed in the 1980's? International Journal of Forecasting, 3(1):53-64, 1987.
- [2] Donald WK Andrews. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica: Journal of the Econometric Society*, pages 817–858, 1991.
- [3] B Ross Barmish. On trading of equities: a robust control paradigm. In Proceedings of the IFAC World Congress, Seoul, Korea, volume 1, pages 1621–1626, 2008.
- [4] Abdel Berkaoui, Mireille Bossy, and Awa Diop. Euler scheme for sdes with non-lipschitz diffusion coefficient: strong convergence. ESAIM: Probability and Statistics, 12:1–11, 2008.
- [5] George E.P. Box and Gwilym M Jenkins. *Time series analysis:forecasting and control.* Holden-day, 1976.
- [6] Massimiliano Caporin, Angelo Ranaldo, and Paolo Santucci De Magistris. On the predictability of stock prices: A case for high and low prices. Journal of Banking & Finance, 37(12):5132–5146, 2013.
- [7] Chris Chatfield. *Time-series forecasting*. CRC Press, 2000.

- [8] An-Sing Chen and Mark T Leung. A bayesian vector error correction model for forecasting exchange rates. computers & operations research, 30(6):887–900, 2003.
- [9] Yan-Leung Cheung, Yin-Wong Cheung, and Alan TK Wan. A high-low model of daily stock price ranges. *Journal of Forecasting*, 28(2):103–119, 2009.
- [10] Menzie D Chinn and Richard A Meese. Banking on currency forecasts: How predictable is change in money? *Journal of International Economics*, 38(1):161–178, 1995.
- [11] John C Cox, Jonathan E Ingersoll, and Stephen A Ross. A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407, 1985.
- [12] Tiziana Di Matteo, Tomaso Aste, and Michel M Dacorogna. Long-term memories of developed and emerging markets: Using the scaling analysis to characterize their stage of development. Journal of Banking & Finance, 29(4):827–851, 2005.
- [13] David A Dickey and Wayne A Fuller. Distribution of the estimators for autoregressive time series with a unit root. Journal of the American statistical association, 74(366a):427–431, 1979.
- [14] David A Dickey and Wayne A Fuller. Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica: Journal of the Econometric Society*, pages 1057–1072, 1981.
- [15] Mona R El Shazly and Hassan E El Shazly. Comparing the forecasting performance of neural networks and forward exchange rates. *Journal of Multinational Financial Management*, 7(4):345–356, 1997.
- [16] Rosario Espínola, Javier Contreras, Francisco J Nogales, and Antonio J Conejo. Dayahead electricity price forecasting based on time series models: a comparison. Proc. of PSCC02, 2005.

- [17] Eugene F Fama. The behavior of stock-market prices. *Journal of business*, pages 34–105, 1965.
- [18] Eugene F Fama. Efficient capital markets: A review of theory and empirical work*. The journal of Finance, 25(2):383–417, 1970.
- [19] Reinaldo C Garcia, Javier Contreras, Marco Van Akkeren, and João Batista C Garcia. A garch forecasting model to predict day-ahead electricity prices. *Power Systems, IEEE Transactions on*, 20(2):867–874, 2005.
- [20] Stefano M Iacus. Numerical methods for sde. In Simulation and Inference for Stochastic Differential Equations, pages 61–107. Springer, 2008.
- [21] S Iwarere and B Ross Barmish. A confidence interval triggering method for stock trading via feedback control. In American Control Conference (ACC), 2010, pages 6910–6916. IEEE, 2010.
- [22] Bong-Han Kim, Hong-Kee Kim, and Keun-Yeob Oh. The purchasing power parity of southeast asian currencies: A time-varying coefficient approach. *Economic Modelling*, 26(1):96–106, 2009.
- [23] Kamil Kladıvko. Maximum likelihood estimation of the cox-ingersoll-ross process: the matlab implementation. *Technical Computing Prague*, 2007.
- [24] Te-Ru Liu, Mary E Gerlow, and Scott H Irwin. The performance of alternative var models in forecasting exchange rates. *International Journal of Forecasting*, 10(3):419– 433, 1994.
- [25] Greta M Ljung and George EP Box. On a measure of lack of fit in time series models. Biometrika, 65(2):297–303, 1978.

- [26] Alicia Troncoso Lora, Jesús Riquelme Santos, José Riquelme Santos, Antonio Gómez Expósito, and José Luís Martínez Ramos. A comparison of two techniques for nextday electricity price forecasting. In *Intelligent Data Engineering and Automated LearningIDEAL 2002*, pages 384–390. Springer, 2002.
- [27] Thomas H Lubecke, Kyung Doo Nam, Robert E Markland, and Chuck CY Kwok. Combining foreign exchange rate forecasts using neural networks. *Global Finance Journal*, 9(1):5–27, 1998.
- [28] Babita Majhi, Minakhi Rout, Ritanjali Majhi, Ganapati Panda, and Peter J Fleming. New robust forecasting models for exchange rates prediction. *Expert Systems with Applications*, 39(16):12658–12670, 2012.
- [29] Burton G Malkiel. The efficient market hypothesis and its critics. Journal of economic perspectives, pages 59–82, 2003.
- [30] Richard A Meese and Kenneth Rogoff. Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of international economics*, 14(1):3–24, 1983.
- [31] Johnathan Mun. Modeling risk: Applying Monte Carlo risk simulation, strategic real options, stochastic forecasting, and portfolio optimization, volume 580. John Wiley & Sons, 2010.
- [32] Whitney Newey and Kenneth West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–08, 1987.
- [33] Francisco J Nogales, Javier Contreras, Antonio J Conejo, and Rosario Espínola. Forecasting next-day electricity prices by time series models. *Power Systems, IEEE Transactions on*, 17(2):342–348, 2002.

- [34] Vincenzo Pacelli. Forecasting exchange rates: A comparative analysis. International Journal of Business and Social Science, 3(10):145–156, 2012.
- [35] Chakradhara Panda and V Narasimhan. Forecasting exchange rate better with artificial neural network. *Journal of Policy Modeling*, 29(2):227–236, 2007.
- [36] Alan Pankratz. Forecasting with dynamic regression models. John Wiley & Sons, 1991.
- [37] Michael Parkinson. The extreme value method for estimating the variance of the rate of return. *Journal of Business*, pages 61–65, 1980.
- [38] Said E Said and David A Dickey. Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika*, 71(3):599–607, 1984.
- [39] Gideon Schwarz et al. Estimating the dimension of a model. The annals of statistics, 6(2):461–464, 1978.
- [40] Bank For International Settlements. Triennial central bank survey foreign exchange turnover in april 2013 : preliminary global results.
- [41] Dukascopy Swiss Forex Bank & Marketplace. Forecasting high and low prices: Cointegration. 2014.
- [42] Mark P Taylor. Purchasing Power Parity and Real Exchange Rates. Routledge, 2013.
- [43] Halbert White. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica: Journal of the Econometric Society*, pages 817–838, 1980.
- [44] Christian CP Wolff. Models of exchange rates: A comparison of forecasting results. International journal of forecasting, 4(4):605–607, 1988.

- [45] Jeffrey Wooldridge. Introductory econometrics: A modern approach. Cengage Learning, 2012.
- [46] Lean Yu, Shouyang Wang, and Kin Keung Lai. A novel nonlinear ensemble forecasting model incorporating glar and ann for foreign exchange rates. *Computers & Operations Research*, 32(10):2523–2541, 2005.

Appendices

Appendix A

Data Sources

Our Forex data of EUR/USD, GBP/USD and AUD/USD were obtained from esignal.com through the paid subscription services.

The 3 month T-bill data used in Chapter 4 was obtained from the US Treasury website www.treasury.gov.

The S&P500 data was collected from Yahoo Finance.

The FTSE100 data was obtained from Investing.com.

Appendix B

Drift and Volatility Estimates of Highs and Lows



Figure B.1: $\hat{\mu_{h,t}}$ of AUD/USD



Figure B.2: $\hat{\mu_{l,t}}$ of AUD/USD



Figure B.3: $\hat{\sigma_{h,t}}$ of AUD/USD



Figure B.4: $\hat{\sigma_{l,t}}$ of AUD/USD



Figure B.5: $\hat{\mu_{h,t}}$ of GBP/USD



Figure B.6: $\hat{\mu_{l,t}}$ of GBP/USD



Figure B.7: $\hat{\sigma_{h,t}}$ of GBP/USD


Figure B.8: $\hat{\sigma_{l,t}}$ of GBP/USD



Figure B.9: $\hat{\mu_{h,t}}$ of EUR/USD



Figure B.10: $\hat{\mu_{l,t}}$ of EUR/USD



Figure B.11: $\hat{\sigma_{h,t}}$ of EUR/USD



Figure B.12: $\hat{\sigma_{l,t}}$ of EUR/USD

Appendix C

Estimation of the Parameters in the CIR Process

Here we show two different methods for the estimation of model parameters in the CIR process. We have shown in (4.16) that we assume R_t to follow the mean-reverting CIR model. To estimate these model parameters we can use both OLS or MLE.

OLS Approach

To estimate the model parameters (θ, κ, v) of the CIR process in (4.16) using OLS, we use the discretized version which is:

$$\frac{R_{t+dt} - R_t}{\sqrt{R_t}} = \frac{\kappa\theta dt}{\sqrt{R_t}} - \kappa\sqrt{R_t}dt + \upsilon\epsilon_t \tag{C.1}$$

Therefore, the estimates $\hat{\theta}$ and $\hat{\kappa}$ can be obtained by minimising the OLS objective function:

$$\arg\min_{\theta\kappa} \sum_{i=1}^{N-1} \left(\frac{R_{i+1} - R_i}{\sqrt{R_i}} - \frac{\kappa\theta dt}{\sqrt{R_i}} + \kappa\sqrt{R_i} dt \right)$$
(C.2)

And the diffusion parameter v is found to be the standard deviation of the residuals: $\hat{v} = \sqrt{\frac{Var(\epsilon_t)}{dt}}.$

MLE Approach

To carry out a Maximum Likelihood Estimation of the parameters $\Psi = (\theta, \kappa, \upsilon)$ for the CIR process in (4.16), we first need to consider it's transition density [23]. This transition density has the form:

$$p(R_{t+dt}|R_t;\Psi,dt) = ce^{-uv}(\frac{u}{v})^{\frac{q}{2}}I_q(2\sqrt{uv})$$
(C.3)

Where

 $c = \frac{2\kappa}{v^2(1 - e^{\kappa dt})},$ $u = cR_t e^{-\kappa dt},$ $v = cR_{t+dt},$ $q = \frac{2\kappa\theta}{v^2} - 1$

and

 ${\cal I}_q$ is modified Bessel function of the first kind of order q.

Therefore, the log-likelihood function for this time-series with N observations is:

$$\ln L(\Psi) = \sum_{i=1}^{N-1} \ln p(R_{i+1}|R_i;\Psi,dt) =$$

$$(N-1)\ln c + \sum_{i=1}^{N-1} \left[-u_i - v_{i+1} + 0.5q\ln(\frac{v_{i+1}}{u_i}) + \ln(I_q(2\sqrt{u_iv_{i+1}}))\right]$$
(C.4)

Where $u_i = cR_i e^{-\kappa dt}$ and $v_{i+1} = cR_{i+1}$. Therefore, by maximising the log-likelihood function (C.4) we can obtain the maximum likelihood estimates of parameters $\hat{\Psi} = (\hat{\theta}, \hat{\kappa}, \hat{v})$. Such that:

$$\hat{\Psi} = \arg\max_{\Psi} \ln L(\Psi) \tag{C.5}$$

Appendix D

Matlab Codes

Below is the matlab code used to carry out the regression 3.3 for the forecasting of highs. The diagnostic tests carried out are also included in the same code. The same code can be modified to estimate α in 3.4 and carry out the diagnostic tests.

x % Log-returns of high, low and close-Lagged by 1-day

hp % High in form of price

yh % Log-Returns of highs-Not Lagged

n %Sample size= n+1

Pforecast = [];

for i = 1: length(x) - n - 2

X=[ones(n+1,1) x(i:n+i,:)]; % Independent Variables

Y=(yh(i:n+i,1)); % Dependent Variable

[B, BINT, R, RINT, STATS] = regress(Y,X); %regression

N=[1 x(n+i+1,:)];% Todays Independent Variables

Pyest=hp(n+i); %Yesterdays Dependent Variable in form of Price

fcast = (N*B)/100; %forecasted value in form of return

Pfcast=(Pyest*exp(fcast)); % forecasted value in form of price

Actual=hp(n+1+i);% Todays actual price

Err=(Actual-Pfcast); % out of sample Errors

 $\operatorname{counter} = 0;$

% test whether P-value of regression is sig at 5% level

if STATS(1,3) > 0.05

```
counter = counter + 1;
```

 $\quad \text{end} \quad$

%LBQ test for autocorrelation

counter1=0;

h = l b q t e s t (R);

if h==1

counter1 = counter1 + 1;

 $\quad \text{end} \quad$

%White test for homoscedasticity

 $\operatorname{counter} 2 = 0;$

 $r\!=\!\!R.\,\hat{}\,2\,;$

 $[B2,BINT2,R2,RINT3,STATS3] = regress(r,[X X(:,2:4).^2 X(:,2).*X(:,3)...$

X(:,2).*X(:,4) X(:,3).*X(:,4)]);

 $Test_stat = (n+1)*STATS3(1);$

CriVal=chi2inv(0.975,10);

 $if Test_stat>CriVal$

counter 2 = counter 2 + 1;

 $\quad \text{end} \quad$

 $\% Standard \ error \ for \ Beta \ \& \ t-test \ for \ Beta$

```
\operatorname{counter3}=0;
```

 $\operatorname{counter} 4 = 0;$

 $\operatorname{counter} 5 = 0;$

 $s_2 = (sum(r)/(n-3));$

 $va=s_2*inv(X'*X);$

se = sqrt(diag(va));

Ttest=zeros(4,1);

for j=1:4

Ttest(j,1) = B(j,1) / se(j,1);

end

Tval=tinv(0.975, n-3);

if abs(Ttest(2,1)) < Tval

$$counter3 = counter3 + 1;$$

 $\quad \text{end} \quad$

if abs(Ttest(3,1)) < Tval

```
counter4 = counter4 + 1;
```

 $\quad \text{end} \quad$

if abs(Ttest(4,1)) < Tval

```
counter5 = counter5 + 1;
```

 $\quad \text{end} \quad$

 $\%\!H\!A\!C$ estimator for se and t-Test

counter6 = 0;

 $\operatorname{counter7}=0;$

counter 8 = 0;

[EstCoeffCov, seHAC, coeff] = hac(X(:,2:4), Y, 'display', 'off');

Ttest1 = coeff./seHAC;

if abs(Ttest1(2,1)) < Tval

counter6 = counter6 + 1;

 $\quad \text{end} \quad$

if abs(Ttest1(3,1)) < Tval

counter7 = counter7 + 1;

end

if abs(Ttest1(4,1)) < Tval

counter8 = counter8 + 1;

 $\quad \text{end} \quad$

%Store all values in Pforecast vector after each loop

Pforecast=[Pforecast; counter1 counter2 counter3 counter4...

counter5 counter6 counter7 counter8 counter9 Err Pfcast Actual];

 end

Below, is our Matlab code for the computation of our triggering method and the computation of Saturation-Reset Linear Feedback Controller.

function $[g, V, I] = Controller_Saturation (alpha, K, dT, rho, H, L, N)$

%H is the vector containing the Log-return of highs

%L is the vector containing the Log-return of lows

%N is the sample size used for the computation of confidence intervals %alpha= significance level used for computing the confidence intervals %K= Controller gain

%dT= incremental time interval, eg for annualized consider 1/252

%rho= daily returns

%preallocate the output variables

g=zeros(length(H)-N,1); %Daily Trading Gains

V=zeros(length(H)-N,1); %Daily Account Value

I=zeros(length(H)-N,1); %Daily Amount Invested

%Initial Values

V(1) = 10000;

I(1) = 0.5 * V(1);

g(1)=0;

%Saturation Limit

Imax=10000;

for i=1:length(H)-N-1

%Compute Mean Drift of highs

 $mubar_H=mean(H(i:N+i,1))*1/dT;$

% Compute volatility of highs

 $sigmabar_H=sqrt(1/N*sum((H(i:N+i,1)-mubar_H).^2));$

% compute upper bound of confidence level for mean drift of highs

 $U_H=mubar_H+((tinv(1-alpha/2,N)*sigmabar_H)/sqrt(N+1));$

% compute Lower bound of confidence level for mean drift of highs

 $L_H=mubar_H - ((tinv(1-alpha/2,N)*sigmabar_H)/sqrt(N+1));$

%Compute Mean Drift of lows

 $mubar_L=mean(L(i:N+i))*1/dT;$

% Compute volatility of lows

 $sigmabar_L=sqrt(1/N*sum((L(i:N+i,1)-mubar_L).^2));$

% compute upper bound of confidence level for mean drift of lows

 $U_L=mubar_L+((tinv(1-alpha/2,N)*sigmabar_L)/sqrt(N+1));$

% compute Lower bound of confidence level for mean drift of lows

 $L_L=mubar_L - ((tinv(1-alpha/2,N)*sigmabar_L)/sqrt(N+1));$

% Go Short: if both upper bounds are below zero

if U_H<=0 && U_L<=0

g(i+1)=g(i)+(rho(i+N+1)*(-abs(I(i))));

$$V(i+1)=V(i)+(rho(i+N+1)*(-abs(I(i))));$$

$$I(i+1) = max(-abs(I(i)) - K*(g(i+1)-g(i)), -Imax);$$

 $\% Going \ Long:$ if both upper bounds are above zero

elseif L_H>=0 && L_L>=0

$$g(i+1)=g(i)+rho(i+N+1)*abs(I(i));$$

$$V(i+1)=V(i)+rho(i+N+1)*abs(I(i));$$

$$I(i+1) = \min(abs(I(i)) + K*(g(i+1)-g(i)), Imax);$$

else

 end

 ${\rm end}$

 $\quad \text{end} \quad$