DEVELOPING TRADING STRATEGIES UNDER THE DIRECTIONAL CHANGES FRAMEWORK

With application in the FX Market

By

Amer M. Bakhach

A thesis submitted for the degree of Doctor of Philosophy

School of Computer Science and Electronic Engineering

Centre for Computational Finance and Economic Agents (CCFEA)

University of Essex

October 2017

Acknowledgements

First and foremost, I wish to thank my supervisor, Prof. Edward Tsang. Without his guidance, support, and endless patient throughout the years of my degree this thesis would not have been possible. His expertise in computational finance and wisdom allowed me to learn a lot from him. Prof. Tsang get retired before the defense of this thesis. I was very lucky to be one of his last students before he has retired.

It is with immense gratitude that I acknowledge the support and help of my brother-in-law, Prof. Hassan El-Rifai (Sharjah University, UAE). It all starts with his advice to get a Ph.D. from UK. Furthermore, as a self-funded student I couldn't afford my tuition fees without his financial support.

I would also like to thank Dr. Wing Lon Ng (ex-lecturer at University of Essex) and Dr. Raju V.L. Chinthalapati (University of Greenwich). Their scientific notes were really inspiring.

My appreciation also goes towards my fellows Dr. Mateusz Gątkowski, Dr. Hamid Jalalian, Ran Tao, and James Chun. Our discussions have been always a great source of knowledge.

Last but not least, this thesis would not have been possible without the person who was always behind me - my wife Ola. Her unconditional support was always a source of encouragement that kept me going during the research journey.

Abstract

Directional Changes (DC) is a framework for studying price movements. Many studies have reported that the DC framework is useful to analyse financial markets. Other studies have suggested that, theoretically, a trading strategy, which exploits the full promise of the DC framework, could be astonishingly profitable. However, such strategy is yet to be discovered. In this thesis, we explore, and consequently provide a proof of, the usefulness of the DC framework as the basis of profitable trading strategies.

Existing trading strategies can be categorised into two groups: the first comprising those that rely on forecasting models; the second group comprising all other strategies. In line with the existing researches, we develop two trading strategies: the first relies on forecasting Directional Changes to decide when to trade; whereas the second strategy is based on the DC framework but uses no forecasting models at all.

This thesis comprises three original research elements:

1. We formalize the problem of forecasting the change of a trend’s direction under the DC framework. We propose a solution for the defined forecasting problem. Our solution encloses discovering a novel indicator which is based on the DC framework.
2. We develop a trading strategy, named TSFDC. TSFDC relies on the forecasting approach established in point 1. to decide when to trade.
3. We develop a second trading strategy, named DBA. DBA does not rely on any forecasting model. DBA employs a DC-based procedure to examine historical prices in order to discover profitable trading rules.

We examine the performances of TSFDC and DBA in the foreign exchange market. The results indicate that TSFDC and DBA can be highly profitable. We compare TSFDC and DBA with other DC-based trading strategies. The results suggest that none of TSFDC and DBA can outperform the other in terms of profitability and risk simultaneously.

List of Figures

List of Tables

List of Publications

- A. Bakhach, E. P. K. Tsang and H. Jalalian, "Forecasting directional changes in the FX markets," *2016 IEEE Symposium on Computational Intelligence for Financial Engineer & Economic (CIFEr' 2016)*, Athens, Greece, 2016, pp. 1-8. doi: 10.1109/SSCI.2016.7850020.

- A. Bakhach, E. P. K. Tsang, Wing Lon Ng and V. L. R. Chinthalapati, "Backlash Agent: A trading strategy based on Directional Change," *2016 IEEE Symposium on Computational Intelligence for Financial Engineer & Economic (CIFEr' 2016)*, Athens, Greece, 2016, pp. 1-9. doi: 10.1109/SSCI.2016.7850004.

- A. Bakhach, E. P. K. Tsang and V. L. R. Chinthalapati, "TSFDC: A Trading Strategy Based on Forecasting Directional Changes", submitted on 15th July 2017 to Journal of "Intelligent Systems in Accounting, Finance and Management" (under review).

Table of Contents

Part I

Introduction, Background and Literature Review

1 Introduction

In this introductory chapter, we describe the adopted rationale which was utilized to conduct this research. Firstly, we explain two concepts, namely: the Foreign Exchange (FX) market and the Directional Change (DC) framework. We then discuss thesis' motivations and objectives. Finally, the thesis structure is described succinctly.

1.1 The foreign exchange market and the Directional Changes framework

Currency trading is the act of buying and selling different world currencies. The foreign exchange (FX) market is the market on which these currencies are traded. The importance of the FX markets has developed with increased global trade, capital flows and investment. The main participants in the FX market are central banks, commercial banks, institutional investors, traders, hedge funds, corporations and retail investors [1] [2]. The objectives pursued by these participants range from pure profit generation (hedge funds, financial institutions) to hedging cash flows; from business core activities (corporations) to implementing macroeconomic and monetary policy objectives (central banks). The analysis of the FX market is a common objective of all market's participants. Institutional and retail investors are particularly interested in discovering moneymaking trading strategies for currency trading (i.e. the devising of a set of rules to indicate when to buy or sell a given currency). Many studies have been published for this purpose (e.g. [3] [4] [5] [6] [7] [8]).

Directional Changes (DC) is a technique that summarizes market prices [9] [10]. Under the DC framework the market is cast into alternating upward and downward trends. A DC trend is identified as a change in market price larger than, or equal to, a given threshold. This threshold, named *theta*, is set by the observer and usually expressed as a percentage. A DC trend ends whenever a price change of the same threshold *theta* is observed in the opposite direction. For example, a market downtrend ends when we observe a price rise of magnitude *theta*; in this case we say that the market changes its direction to an uptrend. Similarly, a market's uptrend ends when we observe a price decline of magnitude *theta*; in which case we say that the market changes its direction to a downtrend. Many studies have reported that the DC framework is useful for analysing the FX market (e.g. [11] [12] [13] [14]). A DC-based trading strategy is a model which employs the DC framework to analyse, and sometimes to forecast, price movements in order to establish profitable trading rules of when to buy or to sell a given asset. Some studies have tried to develop profitable DC-based trading strategies (e.g. [15] [16] [17]). However, the full promise of the DC framework as basis of trading strategies has not yet been completely exploited [16].

1.2 Thesis motivations and objectives

A very important, and also very attractive, research area is the design of trading strategy. This thesis is motivated by the following notes:

a. Some studies (e.g. [18] [19]) have suggested that the produced profits by an idealistic DC-based trading strategy could be of up to 1600% per year (assuming perfect foresight of market's trends under the DC context).
b. In 2017, Golub at al., [16] suggested that the full promise of the DC framework as the basis of a trading strategy is yet to be exploited [16].

Motivated by these notes, the objective of this thesis is to explore, and subsequently to provide a proof of, the usefulness of the DC framework as the basis of profitable trading strategies. To this end, we aim to develop trading strategies based on the DC framework.

Most existing trading strategies can be classified into two groups: 1) strategies that do rely on forecasting models, and 2) strategies that do not. In keeping with the existing research, this thesis proposes two trading strategies, both of which are based on the DC framework. The first one comprises a forecasting model which aims to predict the change of direction of market trend's under the DC context. The proposed trading strategy, then, uses this forecasting model to decide when to initiate a buy or sell order. Our second intended DC-based trading strategy employs no forecasting model. It examines historical prices, using a DC-based computational approach, to unveil profitable conditions of when to buy or sell a given asset.

In order to reach our stated goal certain steps must be taken, the first of which being answers to the following questions.

A. Are Directional Changes predictable?

A common objective for traders is to predict the direction of a market trend (either up or down). Based on this forecasting, the trader makes the decision to buy or sell a particular asset. In this thesis, we address the questions of: how to formulate the problem of forecasting a trend's direction under the DC framework; how to solve this problem; and, how accurate is the proposed forecasting model when compared to other existing forecasting techniques?

We answer these questions in Chapter 5. We consider the problem of whether the current trend will continue for a specific threshold of price change before the trend changes. We also propose a solution for this problem. We compare the accuracy of our approach to the traditional forecasting technique called ARIMA [20].

B. How to develop a successful trading strategy based on forecasting DC?

Even an accurate forecasting model does not necessary guarantee profit in trading. To translate accurate forecasting into profit, a trader need a trading strategy that can utilize the forecasting effectively [21]. Therefore, we need to answer the question of how to develop a successful trading strategy based on forecasting the change of a trend's direction Directional Changes of a given price series?

In Chapter 6, we present a DC-based trading strategy which relies on the forecasting approach, from question A. above, to decide when to initiate a trade. We will examine the performance of the proposed trading strategy and compare it to other DC-based trading strategies.

C. What would be a useful DC-based examination of historical prices to establish a profitable trading strategy?

Some trading strategies do not employ any forecasting models. A common approach is to examine historical price movements to discover lucrative conditions of when to buy or sell a particular asset. In this part of the thesis, we address the question of what a useful DC-based approach to examine historical market price movements might be in order to develop a profitable trading strategy from it?

In Chapter 7, we introduce a new DC-based trading strategy that does not rely on any forecasting model. Instead, it examines the historical prices of a given asset, using a DC-based approach, to discover profitable trading rules. We will examine the performance of this second proposed trading strategy and compare it to other DC-based trading strategies.

Naturally, one might ask why to introduce two trading strategies if one of them is better than the other? We answer this question in Chapter 8. We compare the performances of the two proposed trading strategies. We argue that each of them can be more attractive for different type of traders.

1.3 Thesis outline

The organization of this thesis is as follows:

Chapter 2 provides a general overview of the FX market and looks at the basic terminology of FX trading. Chapter 3 reviews some existing trading strategies in financial markets. We also list and explain some evaluation metrics that are utilized to evaluate the performance of a given trading strategy. In Chapter 4, we explain in detail the Directional Changes concept and clarify how market price movements are sampled under the DC framework. We list some studies that provide evidences as to the importance of the DC framework in analysing the FX market. We also review some trading strategies that are based on the DC concept.

In Chapter 5 we propose a formalism of the problem of forecasting the change of a trend's direction based on the DC framework. We also offer a solution for the established forecasting problem. We prove that our approach provides better accuracy than the ARIMA model. In Chapter 6 we introduce a trading strategy, named TSFDC. TSFDC relies on the forecasting model, developed in Chapter 5, to decide when to trade. We apply TSFDC to eight currency pairs. We evaluate the performance of TSFDC using a rolling window approach. We measure the profitability, risk and risk-adjusted return of TSFDC. We compare TSFDC with other DC-based trading strategies.

In Chapter 7 we present a second trading strategy, named Dynamic Backlash Agent (DBA). We clarify how DBA uses a DC-based procedure to discover profitable trading rules. The performance of DBA will be evaluated the same way as TSFDC in Chapter 6. We compare TSFDC with other DC-based trading strategies.

In Chapter 8 we compare the performances of TSFDC and DBA. The objective of Chapter 8 is to answer the question as to whether either TSFDC or DBA can simultaneously provide greater profit and less risk than the other. Finally Chapter 9 presents our conclusions, which will wrap up this thesis and propose possible future works.

2 The Foreign Exchange Market

In this chapter we provide a brief introduction to the Foreign Exchange (FX) market. We list essential vocabularies related to FX trading. Finally, we review some studies that have examined the profitability of FX trading.

2.1 Introduction

The foreign exchange (FX) market is the market on which currencies are traded. This includes all aspects of buying, selling and exchanging currencies at determined prices. In terms of volume of trading, it is by far the largest market in the world with an average daily turnover of 5.1 trillion US dollars as of April 2016 [1]. The FX market determines the exchange rates for global trade. Thus, it is critical to the support of imports and exports around the world.

The FX market is largely organized as an over-the-counter (OTC) market. In other words, there is no centralized exchange. In centralized exchange-based markets, there is a single price obtaining at any point in time – the market price. However, the FX market is a global decentralized market for the trading of currencies. In decentralized markets, by default, there is no visible common price. The FX market is the largest market of this kind. Unlike stock markets, FX trading is not dealt across a trading floor during a fixed period of several hours a day. Instead, trading is done online (e.g. via computer networks) between dealers in different trading centres around the world.

In the last decade, the study of the FX market has gained increasing interest in the literature. Some studies have focused on the relationship between the FX market and international economics (e.g. [22]), or the relationship between capital flows and trade balance (e.g. [23]). Other studies have focused on the impact of the intervention of the central banks on the FX market (e.g. the case of the Bank of Japan [24] [25] [26], the case of the Czech National Bank [27], the case of the Bank of Canada [28]). In addition, many studies have concentrated on the discovery of statistical properties (e.g. scaling laws and seasonality statistics in the FX market [14] [29] [30]). Some studies (e.g. [3] [4] [5] [6]) have focused on developing profitable trading strategies that specify when to buy or sell a given currency (i.e. FX trading).

The foreign exchange market is unique because of the following characteristics [1] [2]:

- *Market Size*: The FX market is by far the most liquid market in the world. This high liquidity has pushed transaction costs to very low levels.
- *Market Participants*: A very heterogeneous set of actors participates in the FX market (e.g. central banks, commercial banks, institutional investors, traders, corporations and retail

investors). These market participants, often, do not share the same interests when trading currencies.

- *Global Decentralized Market*: There is no specific physical centre to exchange currencies.

This chapter continues as follows: we list and explain some essential terminologies related to FX trading in Section 2.2. In Section 2.3 we review some studies those have examined how profitable the FX trading could be.

2.2 Essential terminologies for FX trading

In this section we describe some essential vocabularies related to FX trading [31]:

- *Exchange Rate*: In a typical foreign exchange transaction, a party purchases a quantity of one currency by paying with a quantity of another currency. The exchange rate represents the number of units of one currency that can be exchanged for a unit of another.
- *Currency Pair*: A currency pair is the quotation and pricing structure of currencies traded in the FX market. The value of a currency is known as a 'rate' and is determined by its comparison to another currency. For example, the currency pair quoted as 'EUR/USD' represents the number of US dollars that can be bought with one euro (see Fig. 2.1 for example).

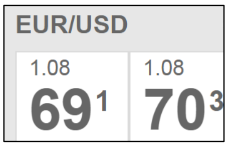

Fig. 2.1. A typical quote of the EUR/USD currency pair. The bid price is 1.08691, the ask price is 1.08703.

- *Base and Counter Currency*: For a given currency pair (e.g. EUR/USD in Fig. 2.1), the first listed currency of a currency pair (i.e. EUR) is called the base currency, and the second currency (i.e. USD) is called the counter currency. The currency pair indicates how much of the counter currency is needed to purchase one unit of the base currency. The counter currency is also referred to as the *quoted* currency.
- *Bid, Ask, and Mid-price*: The bid price represents how much of the counter currency you need in order to purchase one unit of the base currency. The ask price for the currency pair represents how much you will acquire of the counter currency for selling one unit of base currency. For example, in Fig. 2.1 below the bid price of EUR/USD is 1.08691; while the ask price is 1.08703. The mid-price is defined as the average of the bid and ask prices being quoted. For example, in Fig. 2.1 the mid-price would be: (1.08691 + 1.08703) / 2 = 1.08697.

Usually, the mid-price is utilized to illustrate the historical exchange rates of a given currency pair over a specific period (see Fig. 2.2 for example).

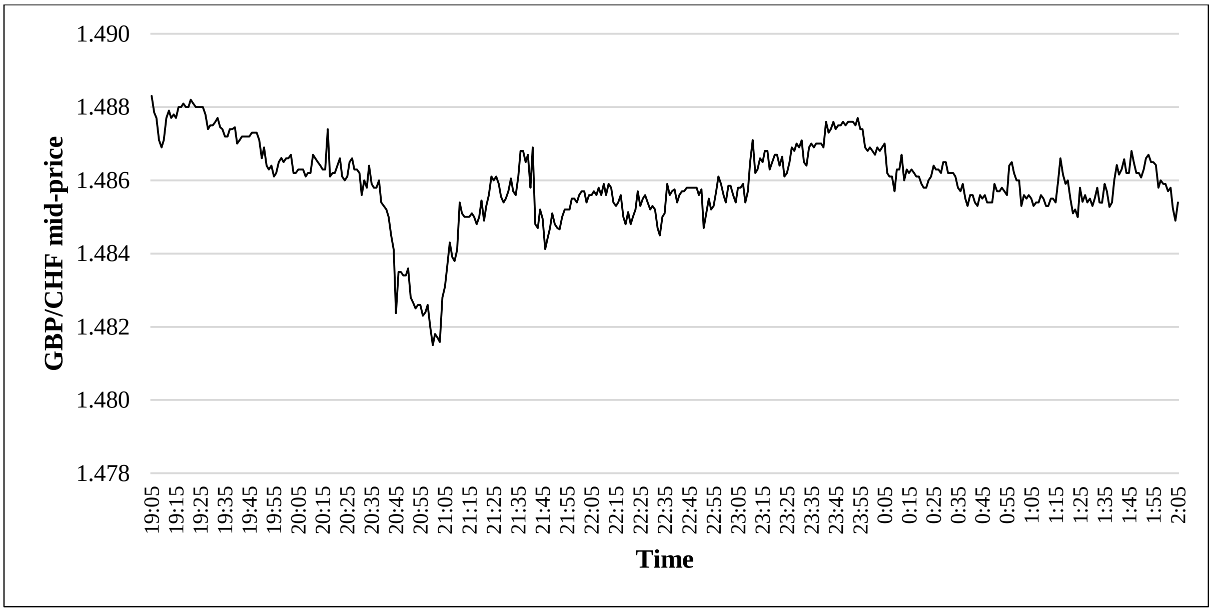

Fig. 2.2. GBP/CHF mid-prices sampled minute by minute from 1/1/2013 19:05 to 1/2/2013 02:05 (UK).

- *FX Market Maker*: A financial institution whose primary business is entering into transactions on both sides of markets and seeking profits by taking risks in these markets. Market makers set both the bid and the ask prices on their systems and display them publicly on their quote screens. The market maker buys from and sells to its investors as well as other market-makers accordingly makes earnings from the difference between the bid and the ask price. Their systems are prepared to make transactions at these prices with their customers, who range from small banks to retail FX traders.
- *Individuals and Retail FX Traders*: A retail investor is an individual investors who buy and sell securities for their personal account, and not for another company or organization. Also known as an ‘individual investor’ or ‘small investor’. An individual trader is expected to deal (i.e. buy and sell) with a market maker.
- *Transaction costs*: Transaction costs are expenses incurred when buying or selling an asset. In a financial sense, transaction costs include market maker’s commissions.
- *Transaction data*: The transaction data (or simply ‘transaction’) denote the details of one trade (a buy or sell agreement between a buyer and a seller). These details includes: a time-stamp (the time at which the trade has occurred), price (either bid or ask), order size (i.e. quantity of share/volume to sell or to buy). It’s worthy to note that several trades (buy or

sell orders) may occur within one second. The details of each trade is considered as transaction data. These data are usually referred to as 'high frequency data'.

2.3 About the profitability of FX trading

In this section we review some studies those have studied the profitability of FX trading. We noticed that most of these studies focus on a specific trading style named 'technical trading'. Typically, a technical trader tries to discover pattern in the historical price movements of a security using some *technical indicators*. Technical indicators are statistics used to measure current conditions as well as to forecast financial trends. Technical indicators are used to predict changes in market trends or price patterns in any traded asset [32] [33]. Eventually, a technical trader establishes a trading strategy (i.e. buy and sell rules) based on the discovered pattern(s). A Technical Trading Rule (TTR) is an instruction that is based on technical indicators and indicate whether the security displays a suitable behaviour to buy or to sell.

In 2013, Neely and Weller [34] studied the convenience of technical trading in the FX market. They reported that technical trading can produce profit in the FX market; especially when applied to emerging markets' currencies (e.g. Latin America). They reported that technical trading can produce better return in comparison to the risk it undertake in the FX market than it does in the S&P500. Their results suggested that it would be better not to embrace fixed technical trading rules or fixed portfolios of these rules, but rather to employ a strategy that switches between different rules and currency pairs according to past performance. Finally, they reported that technical trading in the FX market could generate profits even during financial crisis.

In 2016, Coakley, et al. [35] provided an empirical investigation of the profitability of more than 100,000 technical trading rules (TTR) in the FX market for 22 currencies pairs. They reported that technical trading can achieve an annualised returns of up to 30%.

In 2016, Hsu et al. [36] carried out an investigation of more than 20,000 technical trading rules (TTR) in the foreign exchange market, using daily data sampled over 45 years for 30 developed and emerging market currencies. They reported that technical trading can generate attractive excess returns[a]. Moreover, they concluded that these returns are not, in general, wiped out when realistic allowance is made for transaction costs; which confirm the finding of other studies (e.g. [3] [36] [37]).

[a] Excess returns are investment returns from a security or portfolio that exceed the riskless rate on a security generally perceived to be risk free, such as a certificate of deposit or a government-issued bond.

In 2017, Zarrabi et al. [3] examined the profitability of technical trading rules (TTR) in the foreign exchange market taking into account transaction costs. They consider a universe of 7,650 trading rules and six currencies: SEK, CHF, GBP, NOK, JPY and CAD. The findings indicate that technical trading could generate positive returns even during the financial crisis (e.g. between January 2007 and December 2009). Their results also suggest that an investor should update her portfolio frequently to adapt to changes in the economy rather than sticking to a specific set of TTRs; which confirm the findings of Neely and Weller [34]. They also reported that technical trading can still achieve an attractive level of risk-adjusted return after taking into account transaction costs; which conform to the deduction of Hsu et al., [36].

In 2016, Davison [38] examined the profitability of retail traders in the FX market. He considered the quarterly data collected from 19 US market makers, during the period from 1/10/2010 to 31/3/2014, and aggregated by the on-line website Finance Magnates (Finance Magnates [39]). He reported that, on average, 20% of retails traders can end up with a profitable account; which goes along with the results of Heimer and Simon [40]. Davison [38] concluded that around 40% of retails traders might expect their account to be subject to a margin call[b]. He also reported that there is no conclusive evidences that the successfulness of profitable retails traders is due to their knowledge and skills edge.

To conclude, the studies conducted in [3] [35] [36] examined the profitability of thousands of technical trading rules (TTRs). They concluded that many TTRs can generate profits in the FX market. However, Davison [38] reported that, on average, 20% of retail traders do make profit in reality. A possible reason for the inconsistency of these conclusions could be that it is not easy for most retail traders to examine several thousands of TTRs, to examine the profitability of certain trading rules, before start trading with real money. Besides, some studies (e.g. [34] [3]) reported that, in order to make profits using TTRs consistently, traders must update their TTRs often to adjust to the variations in the market rather than sticking to a particular set of TTRs. Otherwise, a trader might not make profit as he/she has expected. This necessity of updating TTRs continuously makes FX trading harder for retail traders.

[b] An investor receives a margin call from a market maker if one or more of the securities he had bought with borrowed money decreases in value past a certain point. The investor must either deposit more money in the account or sell off some of his assets.

2.4 Summary

The FX market is the market on which currencies are traded. It comprises a wide range of heterogeneous participants (e.g. central banks, retails investors). In Section 2.2, we describes some essential terminologies related to FX trading (e.g. base and counter currencies, mid-price rate). We also reviewed some studies (e.g. [3] [35] [36]) those have highlighted the profitability of FX trading (Section 2.3). Some studies (e.g. [3] [36]) have concluded that FX trading can be attractively profitable even after taking into account the transaction costs. However, other studies (e.g. [38]) warned that, in reality, most retail traders do not make profits as they might have expected.

3 Trading Strategies for Financial Markets

In this chapter, we review some of the existing trading strategies and list selected evaluation metrics to assess the performance of a trading strategy.

3.1 Introduction

A trading strategy is a set of objective 'trading rules'. Trading rules are the conditions that must be met to initiate a *buy* or *sell* order. In this chapter we review previous research into existing trading strategies. In general, these trading strategies can be classified into two categories:

1. The first consists of strategies that aim, firstly, to forecast market prices or change of trend's direction and, secondly, to create trading strategies based on the established forecasting model. The trading strategies in this category usually employ machine learning models to predict market prices or trend's direction. They, then, employ these forecasting model to decide when to initiate buy or sell orders.
2. The second category embraces trading models that do not rely on any forecasting model.

This chapter continues as follows: We review some trading strategies from the first and second categories in Sections 3.2 and 3.3 respectively. In Section 3.4 we list and explain essential evaluation metrics that aim to measure the performance of a given trading strategy. We conclude in Section 3.5.

3.2 The first category: Trading strategies based on forecasting models

Our objective in this section is not to provide an extensive literature review, but rather to provide general examples as to the approaches currently prevailing for developing trading strategies. We will not provide any detailed review of these approaches as they are not based on the DC framework. Trading strategies that are based on the DC framework will be revised later in Chapter 4.

In this section we briefly review some published trading strategies that are based on forecasting models. Generally, trading strategies under this category try, firstly, to forecast the prices or the direction of a financial market's trend. Secondly, they build trading strategies upon the established forecasting model. The following outlines some trading strategies belonging to this category:

In 2009, Li et al. [41] proposed a framework for turning point prediction. The proposed model combines chaotic dynamic analysis with an Ensemble Artificial Neural Network (EANN) model. The sought objective is to capture the non-linear and chaotic behaviour of the financial market in order to forecast potential turning points. A Genetic Algorithm (GA) module is then added to optimize predefined trading parameters to maximize the produced profit of the proposed trading strategy. They applied their forecasting, and trading strategy, to the Dow Jones Industrial Average (DJIA) index time series and TESCO stock (UK). Experimental results suggested that applying the proposed trading strategy to the TESCO stock (UK) could produce rates of return of up to 11.63% within two months.

In 2012, Huang et al. [42] proposed a methodology for stock selection using Support Vector Regression (SVR) and Genetic Algorithm (GA). They used a SVR model to predict, and classify, the profitability of stocks. This classification process includes the usage of fundamental stock criteria (e.g. share price rationality, growth, profitability, liquidity). The stocks classified as 'most profitable' are then employed to form a portfolio. On top of this model, a GA is employed for the optimization of the trading model's parameters. The reported experiment consists of building a portfolio using 30 stocks. Experimental results suggested that, in the best case, the proposed trading system can produce an annualized return of 17.57%.

In 2013, Evans et al., [6] introduced a prediction and decision making model based on Artificial Neural Networks (ANN) and Genetic Algorithms (GA) to predict the change in a market's trend direction. The dataset utilized for this research comprises 70 weeks of historical exchange rates of GBP/USD, EUR/GBP, and EUR/USD currency pairs. They reported that the proposed forecasting model achieved 72.5% prediction accuracy and the associated optimal trading strategy produced an annualized net return of 23.3%.

In 2015, Giacomel et al. [43] proposed an ANN model to predict the direction of price movements. The proposed model was tested using 18 stocks selected from the North American and the Brazilian stock markets. Experimental results suggested that the forecasting model can achieve an accuracy of up to 64%. The results suggest that the proposed trading strategy produced a profit of up to 56% in 166 trading days.

In 2016, Chourmouziadis and Chatzoglou [44] presented a trading fuzzy system. They used a mixture of four technical indicators to predict stock prices. Two of these indicators are very rarely used in research papers, namely Parabolic SAR and GANN-HiLo. They presented 16 fuzzy rules in total based on these four technical indicators. The fuzzy system assigns a weight to each rule based on its profitability during the training (in-sample) period. The experiments were conducted using daily data from the Athens Stock Exchange over a period of more than 15 years. This data was divided into bull and bear market periods. The results suggested the proposed system produces fewer losses during bear market periods and smaller gains during bull market periods compared with the buy and hold[c] strategy.

In 2016, Chen and Chen [45] proposed an intelligent pattern recognition model to predict a turning point of upward trends (i.e. bullish turning point). The proposed model uses nine technical indicators as pattern recognition factors for recognizing stock pattern. They employed the rough sets theory and genetic algorithms for forecasting the bullish turning point. Then, the authors established a trading strategy based on the proposed forecasting model. In the model verification, they evaluated the proposed model in two stock databases (TAIEX and NASDAQ). They measured the total index return percentage[d] (TIR%) that measures a gain rate in the price index for total trades. The authors reported that the proposed trading model produced a TIR% of more than 400% during the period from 18/02/1997 to 24/03/2004.

In 2016, Göçken et al. [46] presented a model to predict stock prices on the Istanbul Stock Exchange. The proposed model employs a hybrid Artificial Neural Network where the inputs are technical indicators chosen via a model that combines Harmony Search (HS) and Genetic Algorithm (GA). They established a trading strategy based on the proposed forecasting model. They applied the proposed trading strategy to Turkey's stock index BIST 100. They reported a positive return of 6.04% during 160 trading sessions.

Finally, we should note that in spite of the fact that forecasting financial time series has been a very attractive objective, many studies (e.g. [47] [48] [49] [50] [51]) do not support their forecasting model with any trading strategy. The establishment of a trading strategy is important in order to give some empirical guarantee that the proposed forecasting method can be used in a real-world situation [21].

3.3 The second category: Trading strategies with no embedded forecasting models

This category encompasses a variety of trading styles that does not rely on any forecasting model. In this section we provide three examples of trading styles that fall under this category, namely: technical trading, momentum strategy and carry trade. Keep in mind that a detailed review of these trading styles is out of the scope of this thesis as they are not based on the DC framework.

---

[c] Buy and hold is an investment strategy in which an investor buys stocks and holds them for a long period of time (a month or years), regardless of fluctuations in the market. The principle of this strategy is based on the view that in the long run financial markets give a good rate of return to investors.

[d] The total return index is a type of equity index that tracks both the capital gains of a group of stocks over time, and assumes that any cash distributions, such as dividends, are reinvested back into the index.

3.3.1 Technical trading

The first trading style, which we consider, is 'technical trading'. Typically, a technical trader analyses price charts to develop theories as to what direction the market is likely to move. This sort of analysis employs a large set of *technical indicators*. Technical indicators look to predict future price levels, or simply the general price direction, of a security by looking at past patterns. Eventually, the discovery of such pattern(s) can help in establishing trading strategies (i.e. buy and sell rules). Examples of traditional technical indicators includes: Moving Average Convergence Divergence; Average Directional Index; Relative Strength Index; Stochastic Oscillator; and Bollinger Bands [32] [33]. Developing trading strategies based on technical indicators is very common in the literature (e.g. [52] [53] [54] [55]). In this section we list some technical trading strategies.

In 2009, Watson [56] established a new approach to studying the profitability of two technical indicators, namely: *head and shoulders* and *point and figure*. He applied his approach to daily data of 4,983 stocks traded on the London Stock Exchange sampled from January 1st 1980 to December 31st 2003. He concluded that the head and shoulders pattern generated a mean excess return of 5.5% on an annual basis. He also concludes that point and figure is particularly suited to intra-day trader[e].

In 2009, Schulmeister [53] examined the profitability of 2,580 technical trading rules (TTR). He reported that the profitability of these TTRs has steadily declined since 1960, and has been unprofitable since the early 1990s when using daily data. However, when based on 30-minutes-data the same TTRs produce an average return of 7.2% per year. He reported that technical trading can be particularly profitable for intra-day trading.

In 2015, Cervelló-Royo at al. [57] proposed a risk-adjusted technical trading rule. They proposed a modified version of a technical indicator named 'flag pattern' that aims to "*strengthen the robustness of the flag pattern and its use in the design of the trading rule*" [57]. They generated 96 different configurations of trading rules and applied these trading rules to three indexes: the US Dow Jones (DJIA), the German DAX and the British FTSE. Experimental results suggest that the trading rules were able to produce returns of up to 94.9% in the period from November 26th 2004 to February 27th 2007.

[e] The name "intra-day trader" refers to a trader who opens and closes a position in a security in the same trading day.

3.3.2 Momentum strategy

The second considered trading style, which does not depend on any trading model, is ‘Momentum strategy’. In general terms, a momentum strategy consists of buying assets with high recent returns and selling assets with low recent returns results.

In 2011, a study by the Monetary and Economic Department at the Bank of International Settlement (BIS) [7] provided a broad empirical investigation concerning the profitability of momentum strategies in the FX market. The authors found that momentum portfolios are significantly skewed towards minor currencies (i.e. currencies that are not actively traded in the FX markets) that have relatively high transaction costs (sometimes these transactions are estimated as high as 50% of momentum returns). They also argued that momentum strategies may deliver higher returns in the FX markets than in stock markets.

In 2013, Daryl et al. [58] proposed a momentum strategy that embedded a security selection approach based on a new risk-return ratio criterion. They sought to create portfolios based on the introduced risk-return ratio criterion. They applied their model to the stock market index of South Korea (KOSPI 200) over the period from June 2006 to June 2012. They reported that the proposed momentum strategy produces attractive positive returns.

3.3.3 Carry trade

The carry trade is a strategy in which traders borrow a currency that has a low interest rate and use the funds to buy a different currency that is paying a higher interest rate. The FX carry trade is of major practical relevance since it represents an important investment style implemented by FX managers [59].

In 2011, Bertolini [8] examined the profitability of several carry portfolio strategies. He analyses whether different asset allocation, market-timing and money management methodologies have the potential to improve the performance of a simple carry portfolio. The experiments were directed using datasets from the G10 currency universe[f] in the period from the 1st January 1999 to the 5th March 2010. He argued that a good practice for FX traders is to adopt a broad diversification of risk indicators for carry trade timing.

In 2014, Laborda et al., [60] proposed an asset allocation strategy which aims to improve the performance of the currency carry trade where currencies are selected from the G10 currency universe. The proposed model assigns weights for long and short positions in a carry trade portfolio

[f] For more information about G10, see http://www.investopedia.com/terms/g/groupoften.asp.

dynamically. These weights is determined by a combination of financial variables that reflect variations in macroeconomic conditions and in the likelihood of crash risk across periods. They reported that the proposed asset allocation strategy produces, markedly, more return than a naive currency carry trade during the out-of-sample period between January 2009 and February 2012.

3.4 Evaluating the performance of a trading strategy

A trading strategy can be analysed on historical data to project the future performance of the strategy. This process is known as 'backtesting'. Backtesting is accomplished by reconstructing, with historical data, trades that would have occurred in the past using the rules defined by a given strategy. The result of backtesting offers statistics that can be utilized to gauge the effectiveness of the strategy. Using a rule-based trading strategy has some benefits:

- It helps remove human emotions from decision making.
- Models can be easily backtested on historical data to check their worth before taking the dive with real money.

There exist many metrics that attempt to evaluate the performance of a given trading strategy. In this thesis, we choose the following metrics to measure the performance of our planned trading strategies. These metrics have been reported as appropriate for a decent assessment ([61] [62]).

- Rates of returns: The rate of return (*RR*) symbolizes the bottom line for a trading system over a definite period of time. Let Total Profit (*TP*) represents the profitability of total trades. *TP* is computed by removing the gross loss of all losing trades from the gross profit of all winning trades (3.1). *TP* can be negative when the loss is greater than the gain. We denote by *RR* (3.2) the gain or loss on an investment over a given evaluation period expressed as a percentage of the amount invested. In (3.2) *INV* denotes the initial capital employed in investment.

$$TP = \text{sum of all profits} - \text{sum of all losses} \tag{3.1}$$

$$RR = \frac{TP}{INV} * 100 \tag{3.2}$$

- Profit factor: The profit factor is defined as the sum of profits of all profitable trades divided by the sum of losses of all losing trades for the entire trading period. This metric measures the amount of profit per unit of risk, with values greater than one signifying a profitable system.

$$Profit\ factor = \frac{\text{sum of all profits}}{\text{sum of all losses}} \tag{3.3}$$

- Max drawdown (%): The drawdown (3.4) is defined as the difference, in percentage, between the highest profit (or capital), previous to the current time point, and the current profit (or capital) value. The Maximum Drawdown (*MDD*) is the largest drawdown observed during a specific trading period. *MDD* measures the risk as the 'worst case scenario' for a trading period. This metric can help measure the amount of risk incurred by a system and determine if a system is practical. In (3.4) and (3.5), the subscript *i* denotes the time-index (i.e. time-stamp). *Current capital*$_i$ denote the value of capital at time (*i*). The *maximum capital* refers to the peak capital's value that has been reached since the beginning of trading up to time *i*. Thus, $drawdown_i$ (3.4), is interpreted as the peak-to-trough decline during a specific recorded period of an investment. Note that, based on (3.4), we have $drawdown_i \leq 0$ for all i. The *MDD* (3.5) is the minimum value among all computed $drawdown_i$. Many studies (e.g. [4] [17] [16]) have used *MDD* to measure the risk of a trading strategy. If the largest amount of money that a trader is willing to risk is less than the maximum drawdown, the trading system is not suitable for the trader.

$$drawdown_i = \frac{current\ capital_i - maximum\ capital}{maximum\ capital} \tag{3.4}$$

$$MDD = Min\ (drawdown_i) \tag{3.5}$$

- Win ratio: The win ratio is calculated by dividing the number of winning trades by the total number of trades for a specified trading period. It expresses the probability that a trade will have a positive return.

$$Win\ ratio = \frac{number\ of\ winning\ trades}{total\ number\ of\ all\ trades} \tag{3.6}$$

- Sharpe ratio [63]: The Sharpe ratio (3.7) is a measure for calculating risk-adjusted return. The basic purpose of the Sharpe ratio is to allow an investor to analyse how much greater a return he or she is obtaining in relation to the level of additional risk taken to generate that return. The Sharpe ratio can be seen as the average return earned in excess of the risk-free rate per unit of volatility or total risk. To date, it remains one of the most popular risk-adjusted performance measures due to its practical use. Some studies (e.g. [64] [65]) have reported that, despite its shortcomings, the Sharpe ratio indicates similar performance rankings to the more sophisticated performance risk-adjusted ratios (e.g. Treynor ratio [66]).

$$Sharpe\ ratio = \frac{R_p - R_f}{\sigma_p} \tag{3.7}$$

Where: R_p denotes the expected portfolio retunes; R_f is the risk-free rate; σ_p designs the standard deviation of the portfolio's returns. One intuition of this calculation is that a portfolio engaging in "zero risk" investment, such as the purchase of U.S. Treasury bills (for which the expected return is the risk-free rate), has a Sharpe ratio of exactly zero. Generally, the greater the value of the Sharpe ratio, the more attractive the risk-adjusted return.

- Sortino ratio [67]: The downside risk (3.8) is defined as the standard deviation of negative asset returns. The Sortino ratio (3.9) uses the downside risk to measure the risk associated to a given investment. In (3.9), the 'return' represents the profits generated by a given trading strategy and the 'target return' is the minimum acceptable return (MAR). In (3.8) m denote the number of the trading periods measured in days, weeks, months, ..etc. Sortino ratio represents the average return earned in excess of the risk-free rate per unit of volatility or total risk.

$$Downside\ risk = \sqrt{\frac{\sum_{i=1}^{m}(return_i - target\ return_i)^2 f(t)}{m}}; \tag{3.8}$$

Where $f(t) = \begin{cases} 1 & if\ return_i < \ target\ return_i \\ 0 & if\ return_i \geq target\ return_i \end{cases}$

$$Sortino\ ratio = (return - target\ return) \div Downside\ risk \tag{3.9}$$

Where $return = \sum_{i=1}^{m} return_i$; $target\ return = \sum_{i=1}^{m} target\ return_i$

- Beta: Beta is a measure of the volatility, or systematic risk, of a security or a portfolio, in comparison to a benchmark [68]. Beta measures how the strategy responds to a benchmark. A Beta of greater than 1 indicates that the security's price will be more volatile than the considered benchmark. For example, if an asset's Beta is 1.3, then it's theoretically 30% more volatile than the benchmark. Essentially, Beta denotes the vital trade-off between reducing risk and maximizing return. Ruppert [69] reports that (3.10) gives the estimated value of Beta (see equations (7.9) and (7.10), p. 230-231 [69]). Theoretically, the overall trading period is divided into a set of sub-trading periods. For example, let the overall trading period lasts for 12 months. It can be composed into 12 sub-trading periods each of which has a length of 1 months. Or, it can be composed into 4 sub-trading periods each of which has a length of 3 months. Let n denote the number of sub-trading periods.

$$Beta_{p,b} = \frac{\sum_i^n \left(R_b^i - \bar{R}_b\right)\left(R_p^i - \bar{R}_p\right)}{\sum \left(R_b^i - \bar{R}_b\right)^2} \quad (3.10)$$

Where, $Beta_{p,b}$ is the Beta of the portfolio p computed with reference to a benchmark b. R_p^i and R_b^i denote, respectively, the return of the portfolio and the benchmark over the i^{th} sub-trading periods. $\bar{R}_p$ and $\bar{R}_b$ are the average of the returns over the n sub-trading periods of the portfolio and the selected benchmark respectively.

- Jensen's Alpha: Jensen's Alpha (3.11) measures the trading return in excess of a security, or portfolio of securities, over the theoretical expected return [70]. Jensen's Alpha is a measure of an investment's performance on a risk-adjusted basis. A positive Jensen's Alpha of 1.0 means the fund has outperformed its benchmark index by 1%. The Jensen's Alpha is computed as:

$$Jensen's\ Alpha = R_p - R_f - Beta_{p,b} \times (R_b - R_f) \quad (3.11)$$

Where, R_p is the total return of the portfolio, R_f is the risk free rate, and $Beta_{p,b}$ is computed as in (3.10).

All of these evaluation metrics will be used later in this thesis to evaluate the performance of our proposed trading strategies as we shall describe in Chapters 6 and 7.

3.5 Summary

In this chapter, we briefly reviewed some of the existing trading strategies from the literature. We identified two categories of trading strategies. The first category contains trading strategies that employ forecasting models. Strategies under this category, usually, embed a machine learning, or artificial intelligence, model to predict market's prices or trend's direction (Section 3.2). The second category consists of those strategies that do not rely on any forecasting model at all. Under this category, we reviewed three trading styles, namely: technical trading, momentum strategy, and carry trade (Section 3.3). None of the trading strategies reviewed in this chapter is based on the directional changes framework.

In Section 3.4, we listed and explained selected evaluation metrics usually employed to evaluate the performance of a given trading strategy. All of these metrics will be used later to assess the performances of our intended trading strategies.

4 The Directional Changes Framework

Directional Changes (DC) is a framework for summarizing price movements. In this chapter, we provide a detailed explanation of the concept of DC. We review several studies that have concluded that the DC framework is useful in analysing the foreign exchange (FX) market. We also review some existing trading strategies those are based on the DC framework. Finally, we clarify the difference between the DC concept and other similar notions.

4.1 Introduction

A common way to summarize raw data in the financial markets is to first choose a time interval, and then sample raw data at fixed time points with the chosen interval; for example, hourly, daily or monthly. We call data summarized this way 'interval-based summary'. Naturally, an interval-based summary becomes a time series. A time series is a sequence of numerical data observations recorded sequentially in time [71].

The Foreign Exchange (FX) market is open 24 hours a day. Trading activities in the FX market can be affected by many factors. For instance, on the announcement of political or economic news, there tends to be a sharp rise in market trading activity in response to the news. Similarly, during weekends, trading activity has a tendency to decline [12]. Thus, an interval-based summary may not capture irregularity in traders' activities appropriately. This raises an essential need to come up with a time-framework that, adequately, captures significant price movements in financial time series beyond the notion of the interval-based summary. This need is particularly important for the analysis of high-frequency data [72].

The concept of 'intrinsic time' is an approach to studying financial time series [73]. Intrinsic time is defined by events. In this context, events are price movements considered as vital by the observer. The objective of using the event-based approach to summarize a time series is to eliminate irrelevant details of price evolution. Although there are many ways of defining events, in this thesis, we consider a specific type of event named Directional Changes (or DC for short) which was established by Guillaume et al., [9].

This chapter continues as follow: in Section 4.2, we provide a detailed explanation of how the DC concept summarizes a market's activities (as explained in Guillaume et al., [9]; Ao and Tsang [10]). In Section 4.3 we discuss some studies that have examined the DC framework's usefulness in analyzing the FX market. We review some existing DC-based trading strategies in Section 4.4.

In Section 4.5, we clarify the difference between the concept of Directional Changes framework, adopted in this thesis, and other similar notions. We conclude with Section 4.6.

4.2 Directional Changes: An introduction

In this section, we explain how market prices are summarized based on the DC concept [9] [10]. Directional changes (DC) is an approach to summarizing price changes. Under the DC framework, the market is represented as alternating uptrends and downtrends. The basic idea is that the magnitude of price change during an uptrend, or a downtrend, must be at least equal to a specific threshold *theta*. Here, *theta* is a percentage that the observer considers substantial. Any price change less than the identified threshold will not be considered as a trend when summarizing market prices.

Let us consider a market in a downtrend. Let P_{EXT} be the lowest price in this downtrend and P_c be the current price. We say that the market switches its direction from downtrend to uptrend whenever P_c becomes greater than P_{EXT} by at least *theta* (where *theta* is the threshold predetermined by the observer; usually expressed as a percentage). Similarly, if the market is in uptrend, P_{EXT} would refer to the highest price in this uptrend. We say that the market switches its direction from an uptrend to a downtrend if P_c is lower than P_{EXT} by at least *theta* (the threshold predetermined by the observer). The detection of a new uptrend or a new downtrend is a formalized inequality, as shown in (4.1).

$$\left|\frac{P_c - P_{EXT}}{P_{EXT}}\right| \geq theta \tag{4.1}$$

If (4.1) holds, then the time at which the market traded at P_{EXT} is called an ‘extreme point’ and the time at which the market trades at P_c is called a DC confirmation point, or DCC point for short. For example, in Fig. 4.2 points A, B, C, D, E, F and G symbolize the ‘extreme points’. Whereas, points $A^{0.1}$, $B^{0.1}$, $C^{0.1}$, $D^{0.1}$, $E^{0.1}$, $F^{0.1}$, and $G^{0.1}$ symbolise the ‘DCC points’. Note that whilst an extreme point is the end of one trend, it is also the start of the next trend, which has an opposite direction. An extreme point is only recognized in hindsight; precisely at the DCC point. For example, in Fig. 4.2, at point $A^{0.1}$ we confirm that point A is an extreme point. Similarly, in Fig. 4.2, at point $D^{0.1}$ we confirm that point D is an extreme point.

Under the DC framework, a trend is dissected into a DC event and an overshoot (OS) event. A DC event starts with an extreme point and ends with a DCC point. We refer to a specific DC event by its starting point, i.e. extreme point, and its DCC point. For example, in Fig. 4.2 the DC event which starts at point A and ends at point $A^{0.1}$ is denoted as $[AA^{0.1}]$. An OS event starts at the DCC point and ends at the next extreme point.

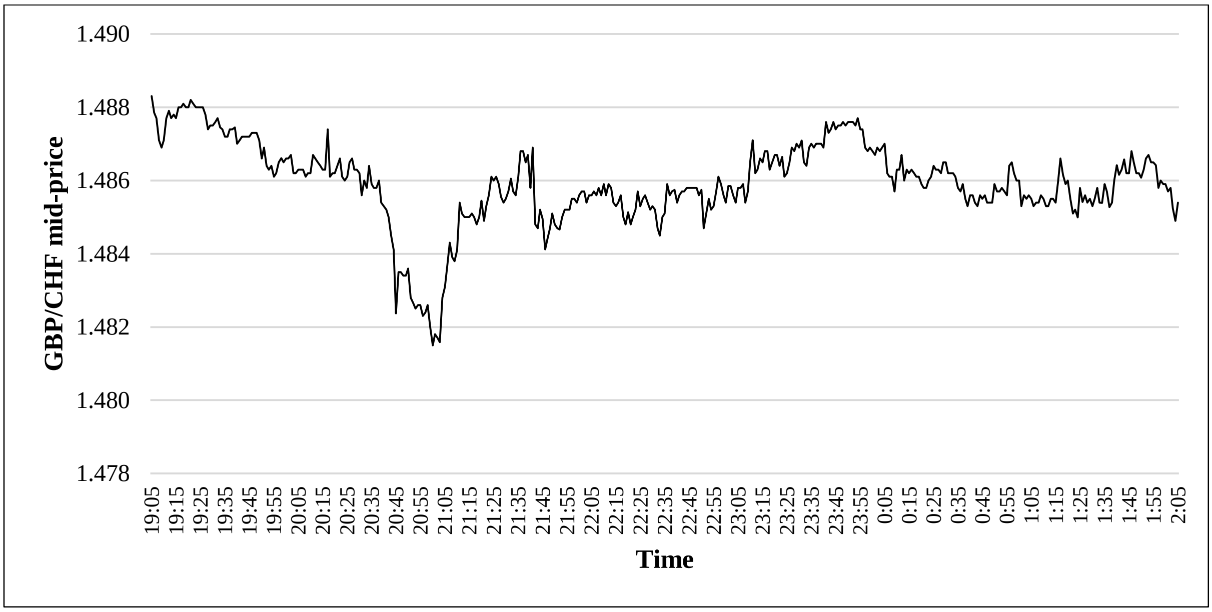

Fig. 4.1. GBP/CHF mid-prices sampled minute by minute from 1/1/2013 19:05 to 1/2/2013 02:05 (UK).

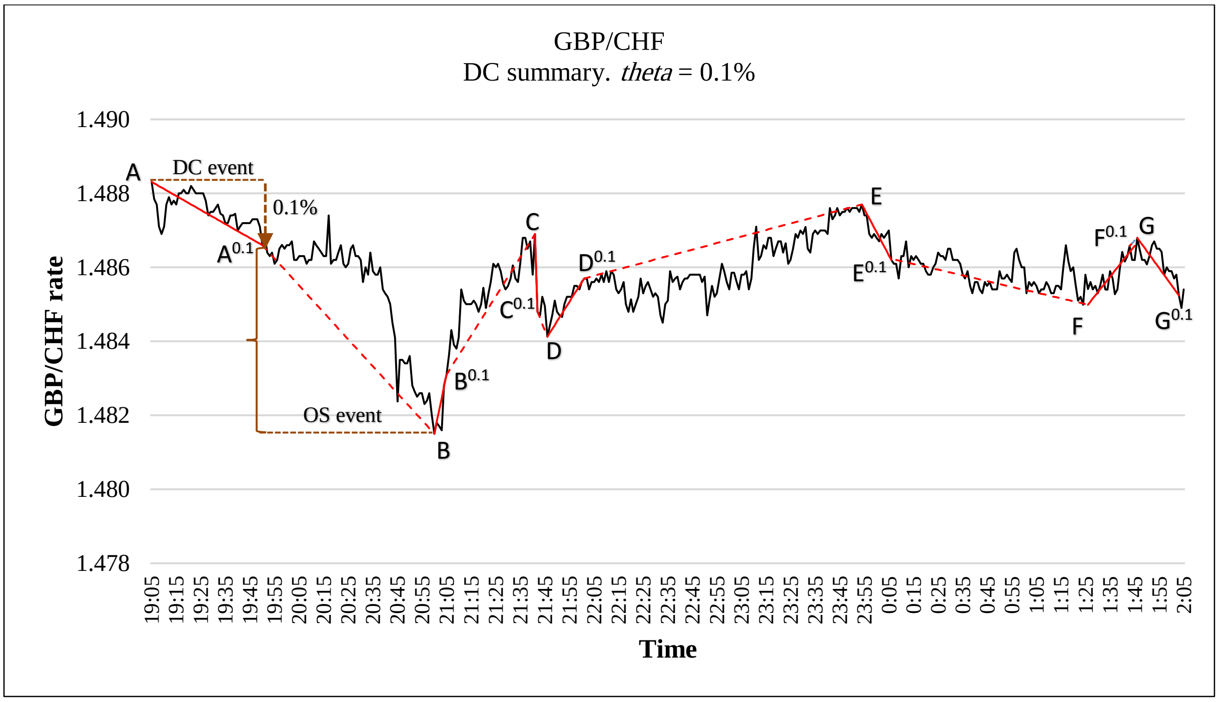

Fig. 4.2. An example of a DC-based summary of the price series shown in Fig. 4.1. Threshold *theta* = 0.10%. The black line indicates GBP/CHF mid-prices sampled minute by minute. Solid red lines represent DC events. Dashed red lines represent OS events. Each of the points A, B, C, D, E, F, G is an extreme point. Each of the points $A^{0.1}$, $B^{0.1}$, $C^{0.1}$, $D^{0.1}$, $E^{0.1}$, $F^{0.1}$, $G^{0.1}$ is a DC confirmation point (DCC point).

The DC summary of a given market is the identification of the DC and OS events, governed by the threshold *theta*. Fig. 4.2 shows an example of a DC summary. Note that for a given time series and a predetermined threshold, the DC summary is unique. However, we can produce multiple

DC summaries for the same considered price series by selecting multiple thresholds. For example, Fig. 4.2 and Fig. 4.3 illustrate two distinct DC summaries for the same price series considered in Fig. 4.1. The used thresholds are 0.10% (Fig. 4.2) and 0.20% (Fig. 4.3).

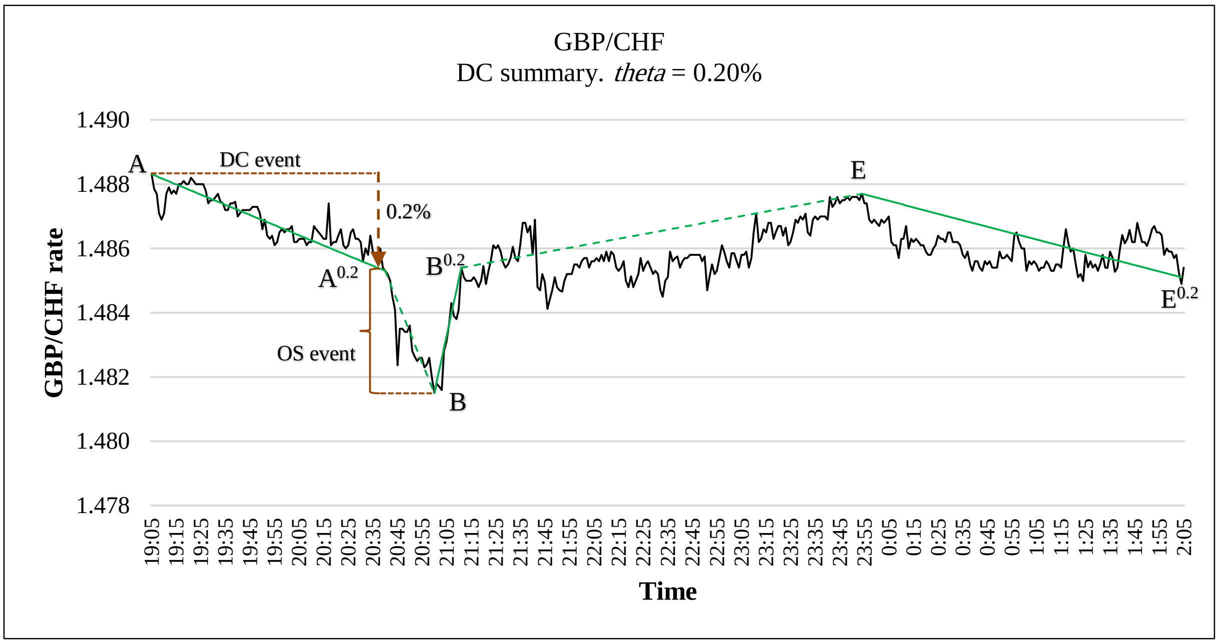

Fig. 4.3. An example of a DC-based summary of the price series shown in Fig. 4.1. *theta* = 0.20%. The black line indicates GBP/CHF mid-prices. Solid green lines represent DC events. Dashed green lines represent OS events. Each of the points A, B, E is an extreme point. Each of the points $A^{0.2}$, $B^{0.2}$, $E^{0.2}$ is a DC confirmation point.

Keep in mind that the observer should specifies the value of the DC threshold *theta*. One observer may consider 0.10% to be an important change, while another observer may consider 0.20% as important. Observers who use different thresholds will observe different DC events and trends. The chosen threshold determines what constitutes a directional change. If a greater threshold had been chosen, then fewer directional changes would have been concluded between the points in Fig. 4.1. For instance, in Fig. 4.2 the DC summary of threshold 0.10% uncovers 4 downtrends and 3 uptrends. Whereas, in Fig. 4.3 the DC summary of threshold 0.20% uncovers 2 downtrends and 1 uptrend.

In this thesis, we use some DC-based notations introduced by Tsang et al., [74]. Table 4.1 lists these notations with basic descriptions. For instance, if the market is in downtrend (uptrend), P_{EXT} would refer to the highest (lowest) price in the overshoot period and $P_{DCC\downarrow*}$ ($P_{DCC\uparrow*}$) would denote the price required to confirm a new downtrend (uptrend) of threshold *theta*. Simply put, in the case of a DC uptrend, if $P_c \leq P_{DCC\downarrow*}$ then we confirm a new downward DC event (i.e. we say that the market has changed its direction to downtrend). Similarly, in the case of a DC downtrend, if $P_c \geq P_{DCC\uparrow*}$ then we confirm a new upward DC event (i.e. we say that the market has changed its direction to uptrend). In Table 4.1, we introduce $PDCC^*$as the price required to confirm a new DC event. That is:

$$\left|\frac{PDCC^* - P_{EXT}}{P_{EXT}}\right| \geq theta \quad (4.2)$$

Table 4.1: List of some notations used in this thesis (source: Tsang et al. [74]). Appendix A provides the code of how to compute these variables.

Name / Description	**Notation**
Threshold	*theta*
Current price	P_c
Price at extreme point: price at which one trend ends and a new trend starts.	P_{EXT}
The highest price, during an uptrend's OS event, required to confirm that the market's direction has changed to downtrend (i.e. to confirm a downtrend's DC event).	$P_{DCC\downarrow*} = P_{EXT} \times (1 - theta)$
The least price, during a downtrend's OS event, required to confirm that the market's direction has changed to uptrend (i.e. to confirm an uptrend's DC event).	$P_{DCC\uparrow*} = P_{EXT} \times (1 + theta)$
$PDCC^*$is the price of the theoretical directional change confirmation point of the current trend.	$PDCC^* = P_{DCC\downarrow*}$ if the current trend is downtrend; otherwise $PDCC^* = P_{DCC\uparrow*}$.

4.3 Applying DC to analyse financial markets

In this section, we review some studies that have concluded the DC framework to be helpful in analysing the FX markets. In 2011, Glattfelder et al., [12] revealed new scaling laws (i.e. stylized facts), based on the DC concept, which uncover innovative facts in the FX market. The authors consider five years of tick-by-tick data for 13 currency pairs. In detail, 11 out of the 18 novel scaling-law relations relate to DC and OS events. Two examples of these scaling laws are: 1) on average, a DC event of threshold *theta* is followed by an OS event of same scale *theta*, and 2) on average, the OS event lasts about double the amount of time that it took for the DC event to complete (see Fig. 4.4).

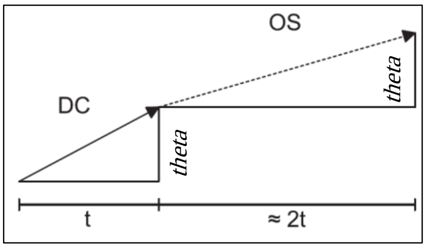

Fig. 4.4. An illustration of two scaling laws related to the DC and OS events reported in [12].

In 2012, Bisig et al. [75] presented the so-called Scale of Market Quakes (SMQ) based on the DC concept. SMQ aims to quantify FX market activity during noteworthy economic and political events declarations. For this purpose, SMQ measures the excess price moves during the OS event. The authors suggested that the SMQ model can be used in different ways. For instance, an investor can use SMQ as a tool to filter the significance of market events. The authors also suggested that SMQ can be used as an input to forecasting or trading models to identify regime shifts. They applied the SMQ model to monitor the behaviour of EUR/USD on the occasion of eight releases of nonfarm employment numbers from the Bureau of Labour Statistics (https://www.bls.gov/). They recognized a wide variety of market responses (e.g. little reaction from the market, volatile market or a drop immediately followed by a recovery [75]).

In 2013, Aloud et al. [11] analysed the statistical properties of the transactions data in the FX market using a DC-based approach. They reported the discovery of four new scaling laws holding across EUR/USD and EUR/CHF transactions. In contrast to the scaling laws presented by Glattfelder et al., [12] which focused on price movements, these new scaling laws focused on transactions data. For instance, the authors found that, on average, an OS event contains roughly twice as many transactions as a DC event.

Also in 2013, Masry [13] presented a study that deciphers FX market activity during the overshoot (OS) event based on the DC concept. She provided empirical evidence of diminishing market liquidity at the end of the overshoot period for all studied currency pairs. She found that a price overshoot stops due to more participants placing counter trend trades, a finding that is valid across all magnitudes of price movement events. She also found that small imbalances of market activity in large overshoots can modify the price trend. She also identified when the market would be vulnerable to the placement of large orders, and the impact of opening counter trend or with trend positions on price overshoots.

In 2014, Golub et al., [76] proposed a new way to measure the liquidity in the FX market based on the DC framework. Their new approach seeks to model market dynamic to predict stress in financial markets. They defined an information theoretic measurement termed liquidity that characterises the instability of price curves during the overshoot event and argue that the new metric can forecast stress in financial markets. They proposed that their model to quantify liquidity in the FX market can be used as an early warning system [76].

In 2017, Aloud and Fasli [77] presented an agent-based model which aims to reproduce, to a certain extent, the stylized facts (e.g. seasonality, scaling laws) previously discovered in the FX market transactions data by Aloud et al. [11]. The presented study examined the existence of relation between the functionality of a DC-based trading strategy and some discovered stylized fact in the FX market. They suggested that the proposed model can be utilized to help in the design of agent trading strategies and decision support systems for the FX market.

In 2017, Tsang et al., [74] presented a new approach to profiling companies and financial markets. They proposed several DC-based indicators to characterize the high-frequency price movements a given market. They suggested that these indicators help to compare markets in terms of volatility and potential profit. They concluded that information obtained through DC-based analysis and from time series complement each other.

4.4 DC-based trading strategies

Recently, some studies have tried to develop trading strategies based on the DC framework (i.e. DC-based trading strategies). In this section, we review four of these studies.

4.4.1 The 'DCT1'

In 2012, Aloud et al., [14] presented a DC-based trading strategy named Zero Intelligence Directional Change Trading (ZI-DCT0). ZI-DCT0 runs a DC summary with a threshold named 'Δ_{xDC}'. ZI-DCT0 has two trading rules:

a. It initiates a trade at the DC confirmation point of a DC event. The type of the trade could be either: counter trend (CT) or trend follow (TF)[g]. In case of CT, ZI-DCT0 opens position against market's trend. TF is the opposite case. The user must specifies the type of trades: either CT or TF.
b. ZI-DCT0 closes the position at the DC confirmation point of the succeeding DC event.

When trading with ZI-DCT0, the trader must determine two parameters:

- The type of trade: CT or TF.
- The threshold Δ_{xDC} to be used for conducting the DC summary.

In 2015, Aloud [15] presented a trading strategy called 'DCT1'. The DCT1 was presented as an updated version of ZI-DCT0. The trading rules of DCT1 are the same as ZI-DCT0 (i.e. rules a. and b. shown above); however DCT1 is designed to automatically compute the two parameters: the DC threshold Δ_{xDC}, and the type of trade (CT or TF). Firstly, the trader defines a range of

[g] A CT (contrarian) trader opens a position (i.e. makes a buy or sell order) with the expectation that the current trend will reverse. A TF (trend follower) trader opens their position with the expectation that the current trend will continue.

thresholds. Secondly, DCT1 automatically examines the profitability of each threshold, included in the specified range, using historical price data (as training set). To this end, for each threshold value, the DCT1 applies the trading rules of ZI-DCT0 from two points of view: counter trend (CT) and trend follow (TF). In other words, during the training period, the DCT1 examines the profitability of all possible combinations of: 1) threshold, included in the range, and 2) the trade type (CT or TF). DCT1 returns the threshold Δ_{xDC} and the type of trade (CT or TF) corresponding to the highest produced returns during the training period. It then uses these values to trade over the trading period.

DCT1 was tested using high frequency data of the EUR/USD currency pair. The author reported that DCT1 was able to produce a rate of return of 6.2% during a testing period of one year (with bid-ask spread being counted). The author did not report any: a) comparison to a benchmark, b) measurement of risk (e.g. *MDD*), or c) evaluation of risk-adjusted metrics (e.g. Sharpe ratio).

4.4.2 A DC-based trading strategy

In 2015, Gypteau et al., [78] presented a DC-based trading strategy. The proposed approach follows the standard tree-based Genetic Programming (GP) configuration. Each GP individual trees comprises internal and terminal nodes. The internal nodes are Boolean functions {AND, OR, NOR, XOR, NOT}.

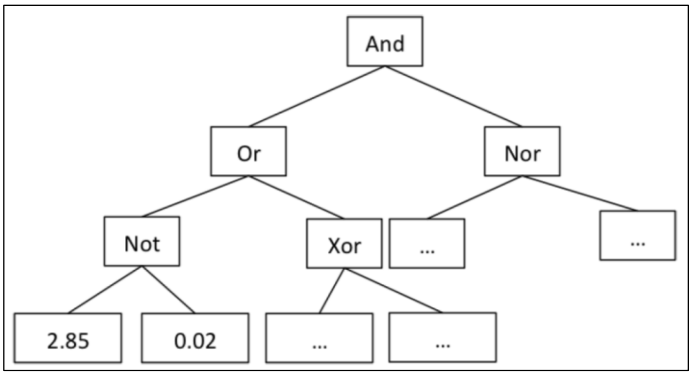

Fig. 4.5. A sample individual GP tree: internal nodes are represented by Boolean functions, while terminal nodes correspond to different DC thresholds. Given a price, terminal nodes output a Boolean value according to the DC or OS events detected. For example, if we detect a downtrend (uptrend) DC event of a DC summary of threshold 2.85%, then the left-most terminal node will be replaced with 'False' ('True'). Source Gypteau et al., [78].

The terminal nodes represent the output of DC thresholds as Boolean values: 'True' if the detected event is an upward DC event; 'False' if the detected event is a downward DC event. For example, Fig. 4.5, shown above, illustrates a sample individual GP tree. In this example, if we detect an upward (downward) DC event of threshold 2.85%, then the left-most terminal nodes would be set as 'True' ('False'). So that, for a given price, all of the terminal nodes of the GP-tree will be replaced with either 'True' or 'False'.

Each GP tree can be interpreted as a Boolean expression; the output of which is either ‘True’ or ‘False’. In summary, given a GP tree, the strategy consists of iterating over the training (in-sample) dataset and based on the output of the individual GP tree, takes the action of selling or buying a stock. At each iteration, the current price information (data point) is used as an input for each DC threshold node. Based on the detected event, the expression represented by a GP tree evaluates to a Boolean value that indicated the action to be taken: buy at the current price (True); sell at the current price (False).

In order to evaluate the output of a GP tree, the algorithm provides a price value to the terminal nodes, which enables the different thresholds to detect DC events. Based on these detected events, each terminal node is replaced by a Boolean value (‘True’ or ‘False’). Consequently, the overall Boolean expression, represented by the GP tree returns, a ‘True’ or ‘False’. This output of a GP tree is, then, translated into trading rules; with ‘True’ triggering a buy signal and ‘False’ triggering a sell signal. Thus, each GP tree denote a trading strategy.

The values of the thresholds, in the terminal nodes, are randomly chosen at the start of the algorithm. The evolution of GP consists of finding the best GP tree (i.e. the thresholds of the terminal nodes and Boolean functions of the internal nodes) which has produced the highest profit during the training period.

With respect to the evaluation of the proposed DC-based strategy, the authors applied their trading model to four markets: two stocks from the UK FTSE100 market (Barclays, Marks & Spencer), and two indices (NASDAQ, NYSE) sampled using daily closing price or index. For each market, they considered a training period of 1000 days in length to train their GP model. Then, they considered a testing (out-of-sample) period of 500 days in length for evaluation. However, the authors did not report the dates of the training and testing periods!

The authors reported only the returns of the proposed trading strategy [78]. The reported returns are less than 10% over a trading period of 500 days for each considered market. Furthermore, they did not report any: a) comparison to a benchmark, b) measurement of risk (e.g. *MDD*), or c) evaluation of risk-adjusted metrics (e.g. Sharpe ratio).

4.4.3 The ‘DC + GA’

In 2017, Kampouridis and Otero [17] proposed a DC-based trading strategy named ‘DC+GA’. DC+GA runs multiple DC summaries concurrently (using multiple thresholds). For each threshold, DC+GA calculates the average time length of each DC and OS event for every DC trend during a training (in-sample) period. DC+GA employs two variables to express the average ratio of the OS event length over the DC event length. These two variables are r_u and r_d, where r_u is the average ratio of the upwards OS event, and r_d is the average ratio of the downwards OS event. Thus, DC+GA analyses uptrends and downtrends separately. The objective is to be able to anticipate the end of an uptrend, or downtrend, (approximately) and as a result make trading decisions (buy or sell) once an OS event had reached the average ratio of r_u or r_d. Theoretically, DC+GA initiates a trade when the length of an OS event exceeds r_u or r_d.

However, the r_u and r_d ratios are just average approximations. In reality, it's expected that many times the OS event might last longer, or shorter, than the estimated average (r_u and r_d). To address this issue, the authors created two parameters, namely b_1 and b_2, which define a range of time within the OS period, where trading is allowed. For instance, if a trader expects the OS event to last for 2 hours (this expectation is based on the calculus of r_u and r_d), and assuming that the range of [b_1, b_2] is [0.9, 1.0], then this means that DC+GA is going to trade (buy or sell) at the last 10% of the 1 hours duration, i.e. in the last 6 minutes.

Recall that DC+GA runs multiple DC summaries simultaneously (using multiple thresholds) for a given currency pair. Let N_{theta} be the number of employed DC thresholds. The user/trader should chose the values of the N_{theta} thresholds. DC+GA assigns a weight to each DC threshold. For a given price observation, each threshold provides a recommendations (buy, sell or hold) based on the values of b_1, b_2, r_u and r_d. At a given time, the N_{theta} thresholds provide N_{theta} recommendations. These N_{theta} recommendations are grouped into two groups based on the produced recommendation: the first group contains the DC thresholds recommending to buy; the second group contains the DC thresholds with sell recommendations. In order to decide which recommendation (buy, sell, or hold) to adopt, DC+GA sums the weights of the thresholds of the two groups: if the sum of the weights for all thresholds recommending a buy (sell) action is greater than the sum of the weights for all thresholds recommending a sell (buy) action, then the strategy's action will be to buy (sell).

To optimize the weights of these N_{theta} thresholds and the associated trading parameters, DC+GA employs a Genetic Algorithm (GA) approach. DC+GA symbolizes a trading strategy as a GA gene. In this context, a GA's gene comprises: the weights of the N_{theta} thresholds, b_1, b_2, and Q; with Q being the order size (i.e. how much to buy or to sell). During the in-sample (training) period, the evolution of GA consists of discovering the best GA gene. The best GA gene is the one which returns the maximum profits during the training period. This best gene will be used for trading during the out-of-sample (trading) period. DC+GA employs a fitness function which aims to minimize the maximum drawdown (MDD) and maximize returns at the same time.

To evaluate the performance of DC+GA, the authors considered five currency pairs sampled with a 10-minute interval over one year. They adopted a daily-basis rolling window approach with the training period being 1 day. When examining the reported monthly returns (in Tables 5 and A1, pages 156 and 158 respectively, Kampouridis and Otero [17]) one can easily note that the proposed trading models incur losses in about 50% of the cases! The authors concluded that the proposed model "…*could not consistently return profitable strategies and thus their mean returns were negative.*"Kampouridis and Otero [17] reported the average monthly returns of applying DC+GA to five currency pairs (shown in Table 6, page 158 [17]). We note that DC+GA incurs overall losses in two out of the five cases.

As for the risk-adjusted performance, the authors did not provide any risk-adjusted measurement. However, based on the reported monthly returns (Table 5, page 158, [17]), we can compute the Sharpe ratio. If we consider a risk-free rate of 5% per annum, then we find that DC+GA will have negative Sharpe ratio in four out of the five considered currency pairs as follow:

- In the case of EUR/GBP: – 0.9
- In the case of EUR/JPY: 0.2
- In the case of EUR/USD: – 0.7
- In the case of GBP/CHF: – 0.6
- In the case of GBP/USD: – 0.1

We should finally note that the reported average *MDD* of DC+GA is less than 0.15% (measured on daily basis) in all considered cases (Table 8, [17]). We consider this value as an attractive level of the drawdown risk.

4.4.4 The 'Alpha Engine'

In 2017, Golub at al., [16] presented a DC-based trading strategy called 'Alpha Engine'. The Alpha Engine is a contrarian trading strategy. The mechanism of initialization of new positions and the management of existing positions in the market works as follow:

Initially, the Alpha Engine opens a new position, against the market trend, during the OS event when the price's change exceeds a certain threshold named 'ω'. ω is a function of the predetermined DC threshold *theta* and a parameter named α (which is governed by the money management module as we shall describe next).

$$\omega = \alpha \times theta \tag{4.3}$$

The Alpha Engine does not have an explicit stop-loss rule. Instead, it employs a sophisticated money management approach. Each time the Alpha Engine opens a new position, it names this position a 'trading agent'. The Alpha Engine is capable of opening and managing multiple positions (i.e. multiple trading agents) concurrently. When Alpha Engine opens a new position (i.e. initiates a new trading agent), it keeps managing the size of this position until it closes in a profit. The Alpha Engine increases and decreases the size of the position (i.e. the quantity of inventory held by a trading agent) as the price progresses. The basic idea is that an existing position is increased by some increment in case of a loss, bringing the average closer to the current price. For a de-cascading event, an existing position is decreased, realizing a profit.

When triggering a new trade, a trading agent must decide the '*time*' and the '*size*' of that trade. For this purpose, the Alpha Engine takes into concern two main factors:

a. *The accumulation of inventory sizes as the market price moves up and down*: the threshold ω is essentially utilized to control the time at which a trading agent should initiate a new order. More particularly, the Alpha Engine manage the parameter α to control the value of ω (4.3). The value of α is a function of the inventory size. Let I denote the overall inventory size held by all generated trading agents altogether. The authors considered I as a proxy for the market. The Alpha Engine uses the value of I to manage the parameter α; and, consequently, the threshold ω.

b. *A probability indicator, denoted as '$\mathcal{L}$'*: The value of $\mathcal{L}$ is interpreted as the probability that the trend will go up or down provided the current state. $\mathcal{L}$ is computed using a transition network of states which has two states: ω and *theta*. This transition network is designed so that in the case of unlikely price trajectory (i.e. abnormal market behavior) $\mathcal{L} \approx 0$. On the other hand, if the markets show normal behavior, i.e. no strong trend can be recognized, then $\mathcal{L} \approx 1$. The Alpha Engine uses $\mathcal{L}$ to control the size of a new order. The size of a new order increases (decreases) as $\mathcal{L}$ approaches to 1 (0). It follows from the previous description that $\mathcal{L}$ helps the trading agents not to build up large positions which they cannot unload. Besides, by slowing down the increase of the inventory of a trading agent during market's overshoots, the overall trading models experiences smaller drawdowns and better risk-adjusted performance. The introduction of $\mathcal{L}$ to analyses and models market behavior was originally described in [76].

Moreover, the Alpha Engine uses asymmetric thresholds for uptrends and downtrends. The authors found that the market is most likely to exhibit different behaviors during uptrends and downtrends. To cover this dilemma, the Alpha Engine employs two different DC thresholds (instead of just one: '*theta*'); one for uptrends ($theta_{up}$) and another for downtrends ($theta_{down}$).

Similarly, the Alpha Engine has two different de thresholds ω, the so-called ω_{up} and ω_{down}. With $\omega_{up} = \alpha_{up} \times$ *theta*$_{up}$ and $\omega_{down} = \alpha_{down} \times$ *theta*$_{down}$; with α_{up} and α_{down} are two trading parameters; the values of which rely on the inventory *I* as explained in point a. above.

The details of this money management mechanism is quite complicated. For more details about this mechanism see Golub et al., [16]. Most importantly, we should note that this money management approach is an integrated module of the Alpha Engine.

The Alpha Engine was extensively backtested using a portfolio of 23 currency rates sampled tick-by-tick over a period of eight years: from the beginning of 2006 until the beginning of 2014. Alpha Engine produces a return of 21.34% (the bid-ask spread was counted), with a maximum drawdown of 0.71% (calculated on a daily basis). The authors reported an annual Sharpe ratio (4.4) of 3.06. However, they did not specify the used risk-free rate. This is an important issue. For example, if we consider an annual risk-free rate of 5% then the Alpha Engine will have a negative Sharpe ratio. In (4.4), R_p denotes the expected portfolio retunes; R_f is the risk-free rate; σ_p designs the standard deviation of the portfolio's returns.

$$Sharpe\ ratio = \frac{R_p - R_f}{\sigma_p} \tag{4.4}$$

The authors also measured the performance of Alpha Engine using synthetic data consisting of 10 million ticks with annualized volatility of 25% generated by a geometric random walk. Amazingly, the Alpha Engine produced positive returns of about 35% in this simulation. They concluded that the Alpha Engine produces profitable results even on time series generated by a random walk. The authors made the code of Alpha Engine available online at Github [79].

4.5 Notions and concepts similar to DC

In this section, we distinguish the DC concept adopted in this thesis from other similar notions. Despite the similarity in the names, the DC concept as described in this paper is completely different from both the 'Change Direction' [80] and 'Direction– of– Change' [81]concepts. In both studies, [80] and [81] the authors used interval– based datasets (daily close value); neither a threshold *theta* was used, nor a DC event defined. Instead, they tried to forecast when a given stock index would switch its trend direction (upward or downward) at the daily closing price without measuring the magnitude of the price change. Their models aimed to answer the question: "will today's close price extend yesterday's trend?"

However, the DC concept is similar to the zigzag indicator. The zigzag approach model price movement as alternating uptrend and downtrend [82] [83] [84]. The price change during an uptrend or a downtrend must be at least equal to a specific threshold. The literature comprises another similar notion: the 'turning points'. In general, price movement can be symbolized as alternating uptrends and downtrends, separated by 'turning points'. Turning points are essentially local minimum and maximum points on a time series, or in practical terms, the peaks and troughs [41]. Turning points are the points at which the trend's direction reverse; usually for a magnitude predetermined by the observer. Turning points can be interpreted as the extreme points under the DC context.

The zigzag indicator and turning points concepts are pretty similar to the DC framework with the main difference being that a trend, under the DC methodology, is dissected into: 1) a DC event of fixed percentage equal to the selected threshold and 2) an OS event represented by the remaining part of the trend before it reverses. Such partitioning is neither part of the zigzag indicator nor of turning point model. Keep in mind that the dissection of a trend into DC and OS event, under the DC framework, has been reported to be helpful to analyse and characterize financial markets in many studies (e.g. [11] [12] [74] [75] [85]).

4.6 Summary

In this chapter, we explained the concept of Directional Changes (DC). The DC framework is an approach to summarizing prices in the financial markets. A directional change is defined by a threshold that the observer considers significant, e.g. 5%. A *theta*% directional change is basically a price change of *theta*% from the last peak or bottom price. Under the DC framework the market is seen as a series of alternated uptrends and downtrends. A trend is dissected into a DC event (of fixed threshold *theta*) and an OS event (consisting of the remaining part of the trend). In Section 4.2, we listed some important DC-based notations (e.g. $PDCC^*$, P_{EXT}) those will be used later in this thesis.

Reviewing the literature in Section 4.3, we found that many studies have concluded that the DC framework is helpful in gaining more insight into the analysis of the FX market. This comprised the discovery of new scaling laws, understanding the impact of new trades on market's trend, and measuring the impact of political and economic events on the market. We also noticed that only recently, some studies have tried to develop trading strategies based on the DC framework. We reviewed four of these studies in Section 4.4. Later in this thesis, we will compare these four DC-based trading strategies to our planned trading strategies in Chapters 6 and 7.

In this thesis we aim to explore, and consequently to provide a proof of, the usefulness of the DC framework as a foundation for successful trading strategies. It’s important to note that our planned DC-based trading strategies in this thesis are not based on any other DC-based strategy. However, some similar features may exist as we shall discuss in Chapters 6 and 7.

Part II

Thesis Contributions

5 Forecasting Directional Changes: Problem Formulation and Solution

Many studies have tried to forecasting the change of the direction of market trend. To the best of our knowledge, no study has considered this problem within the DC context. In this chapter, we study this problem under the DC framework. The central research question which we pose here is whether the current trend will continue for a specific percentage before the direction of the trend reverses.

In this chapter, we formalize this forecasting problem from the DC perspective and propose a solution. We evaluate the accuracy of our approach using eight currency pairs from the FX market. The experimental results suggest that the accuracy of the proposed forecasting model is very good; in some cases, prediction accuracy is over 80%.

5.1 Introduction

Forecasting financial time series is a common objective for financial institutions and traders. This task has proven to be very challenging [86]. Many studies have focused on the issue of next-value prediction, which entails forecasting the future value of time series at the oncoming time step, given the historical observations up to the current time. There may, however, be advantages in predicting the change of market trend's direction directly (i.e. without explicitly predicting the future value of the series). For example, traders may take decisions based on their estimation of whether the price of a particular market will rise or fall [81].

Many studies have tried to predict when a given market would switch its trend direction. These studies usually aims to answer the question: will today's close price extend yesterday's trend? In other words, these studies consider the market prediction problem as a classification problem where the question is whether the market goes up or down? Usually this problem is referred to as forecasting the change of direction. For instance, Park et al. [80] proposed a continuous Hidden Markov Model (HMM) to predict the change of direction of financial time series. They proposed to split the data, consisting of daily closing prices, into two classes with respect to change direction of next day's closing price, and train the two HMMs (one for class). The two formed HMM models are, then, employed to forecast change direction of next day's closing price. Skabar [81] presented a Bayesian multilayer perceptron model to predict the direction of the daily close value of the Australian financial index. Skabar [87] proposed another forecasting model in which he used a similarity-based classification model to predict the trend's direction of tomorrow's close price. He admitted that both models, [81] and [87], have almost equal accuracy. Giacomel et al. [43], proposed an ensemble of two ANNs to predict the direction of price movement. The proposed model was tested using two cases: the North American and the Brazilian stock markets for a total of 18 stocks. Evans et al. [6] introduced a model which combines Artificial Neural Networks (ANN) and Genetic Algorithms (GA) to predict intra-day market price direction. They employed a GA module to search the best network topology of a multiple layer perceptron (MLP) in order to improve forecasting accuracy. It's important to note that the objective of these studies is to forecast whether the next price observation will be larger, or less, than the last recorded price? In this chapter we consider a different forecasting problem as we shall describe next.

Price movements can be symbolized as alternating uptrends and downtrends, separated by 'turning points'. Turning points are essentially local minimum and maximum points on a time series or, in practical terms, the peaks and troughs [41]. Turning points are the points at which the trend's direction reverses; usually for a magnitude predetermined by the observer (similarly to the concept of 'extreme point' under the DC context as described in Section 4.5). An investor who can trade exactly at the turning points (e.g. buying at minima and selling at maxima) would gain the maximum possible profit. Therefore, a common objective for traders in the financial markets is to forecast turning points. Predicting turning points has long been a tough task in the field of time series analysis. Many machine learning models have been developed for this purpose, with the majority of cases focusing on stock markets.

For instance, Azzini et al. [83] tried to predict a turning point in the S&P500 index. Their objective was to predict the magnitude of price change of the entire trend (i.e. between two consecutive turning points) before the trend reversed. They used two models for this purpose: fuzzy logic and neural networks. Li et al. [41] proposed a framework for turning point prediction that combines chaotic dynamic analysis with a neural network. Their proposed model try to predict whether the next time step in the time series is a peak, a trough or none. El-Yaniv and Faynburd [88] proposed a model for the prediction of turning points based on support vector regression.

Many studies have concluded that the directional change (DC) framework is useful in analysing the FX market (e.g. [11] [12] [14] [75]). In this chapter, we consider the problem of forecasting the change of a trends' direction from the DC perspective. The task is to forecast whether the current trend, either uptrend or downtrend, will continue in the same direction for a specific percentage before it reverses (i.e. before the occurrence of the next extreme point). Answering this question can be useful for investment decisions. For example, it could help a trader to make a buy or a sell decision (as we shall argue in Chapter 6).

Forecasting crucially depends on the variables used. As a first attempt to tackle the proposed forecasting problem, we introduce an original DC-based independent variable. We prove that it is useful for the proposed forecasting problem. Our forecasting model, in this chapter, is novel because:

- *In term of problem formulation*: We consider the problem of 'forecasting whether the current trend will continue in the same direction for a specific percentage before it reverses' from the DC perspective. To the best of our knowledge, no previous study has considered this problem from the DC perspective.
- *In term of the proposed solution*: We will introduce an original DC-based indicator and prove that it is helpful in predicting the change of a trend's direction with very good accuracy. Most of the existing forecasting approaches use traditional technical indicators [21].

The rest of this chapter is organized as follows: we introduce a new concept named Big-Theta, which is based on the DC concept, in Section 5.2. Then we provide the formal definition of our proposed forecasting problem in Section 5.3. In Section 5.4, we present our approach to solving the introduced forecasting problem. We describe a set of experiments in Section 5.5, designed to examine the accuracy of our forecasting model. The experimental results are reported and discussed in Section 5.6. We conclude with Section 5.7.

5.2 The concept of Big-Theta

5.2.1 Big-Theta

In this section, we introduce a new concept, named Big-Theta. The notion of Big-Theta states that a DC event of threshold *BTheta* will embrace at least one DC event of a smaller threshold *STheta* (with *BTheta* > *STheta*). As explained in Chapter 4, the DC summary consists of identifying the DC and OS events, corresponding to a predetermined threshold, of a given price series. The DC summary is unique given a specific threshold. However, for the same price series, we may produce several DC summaries by using multiple thresholds. For instance Fig. 5.1 illustrates a DC summary of GBP/CHF prices series using a threshold (*STheta* = 0.10%), whereas, Fig. 5.2 illustrates a DC summary of the same GBP/CHF prices series as Fig. 5.1 but uses another threshold (*BTheta* = 0.20%). The smaller the threshold, the more DC events are recognized. For instance, in Fig. 5.1 (with a threshold of 0.10%) we observe four downtrends and three uptrends. However, in Fig. 5.2 (with a threshold of 0.20%) we observe only two downtrends and one uptrend.

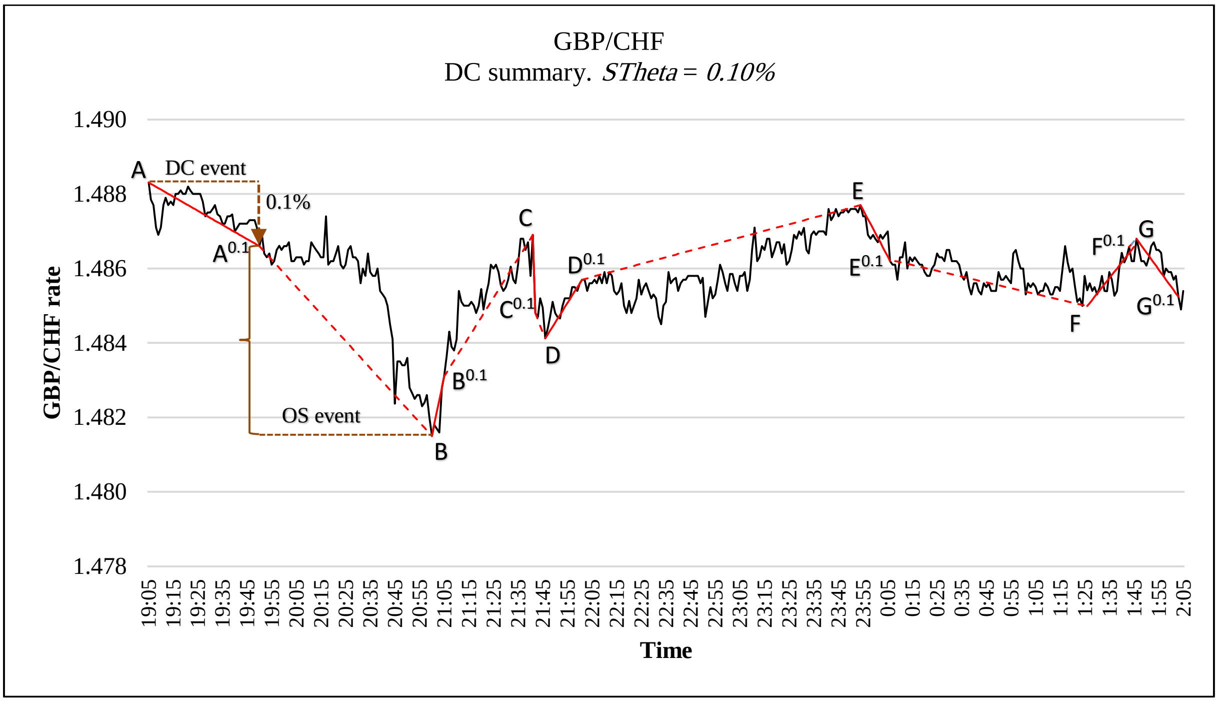

Fig. 5.1. An example of a DC-based summary. The black line indicates GBP/CHF mid-prices sampled minute by minute from 1/1/2013 19:05 to 1/2/2013 02:05. *Theta* = 0.10%. Solid red lines represent DC events. Dashed red lines represent OS events. Each of the points A, B, C..etc, represents a specific time.

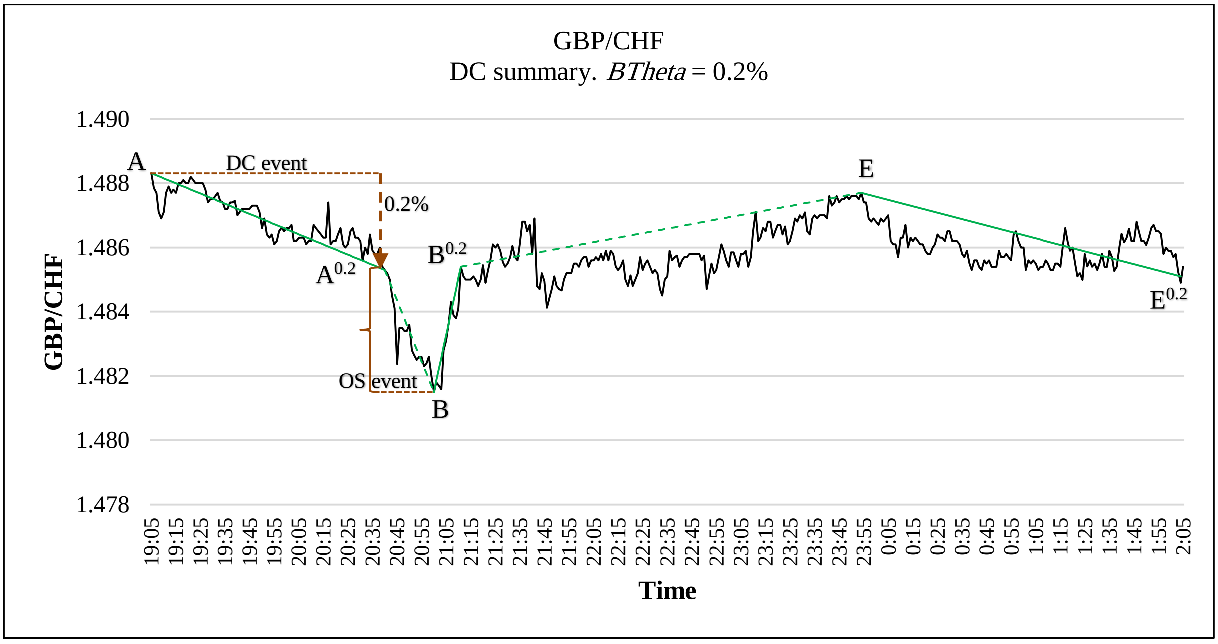

Fig. 5.2. A DC summary for GBP/CHF mid-prices sampled minute by minute from 1/1/2013 19:05 to 1/2/2013 02:05. *BTheta* = 0.20%. Solid green lines represent DC events. Dashed green lines represent OS events.

Let $EP^{0.002}$ be the set of all extreme points shown in Fig. 5.2. $EP^{0.002}$ = {A, B, E}. Similarly, let $EP^{0.001}$ be the set of all extreme points shown in Fig. 5.1. $EP^{0.001}$ = {A, B, C, D, E, F, G}. An important feature to note is that each point in $EP^{0.002}$ is also a point of $EP^{0.001}$. However, the inverse is not true. For instance, points C and D are elements of $EP^{0.001}$, but they are not elements of $EP^{0.002}$. Fig 5.3 illustrates the synchronization of the two DC summaries shown in Fig 5.1 and Fig 5.2. Fig

5.3, shown below, helps to exemplify the concept of Big-theta. It illustrates the fact that each extreme point recognized under the DC summary of threshold *BTheta* (0.20%) is also recognized as an extreme point under the DC summary of threshold *STheta* (0.10%).

By definition, the elements of $EP^{0.001}$ and $EP^{0.002}$ are sorted chronologically. For example, points A and B in $EP^{0.001}$ are, respectively, the extreme points of the 1st and 2nd trends as observed under a DC threshold of 0.10% (Fig. 5.1). In general, let EP^{BTheta} and EP^{STheta} be the sets of all extreme points of the DC summaries of thresholds *BTheta* and *STheta* respectively. An extreme point of the DC summary of threshold *BTheta*, i.e. an element of EP^{BTheta}, is also recognized as an extreme point under the DC summary of threshold *STheta*, i.e. an element of EP^{STheta}; provided that *BTheta* > *STheta*. We will use Big-theta for two purposes: 1) to formalize the problem of forecasting a trend's direction (in Section 5.3) and 2) to propose a solution (in Section 5.4).

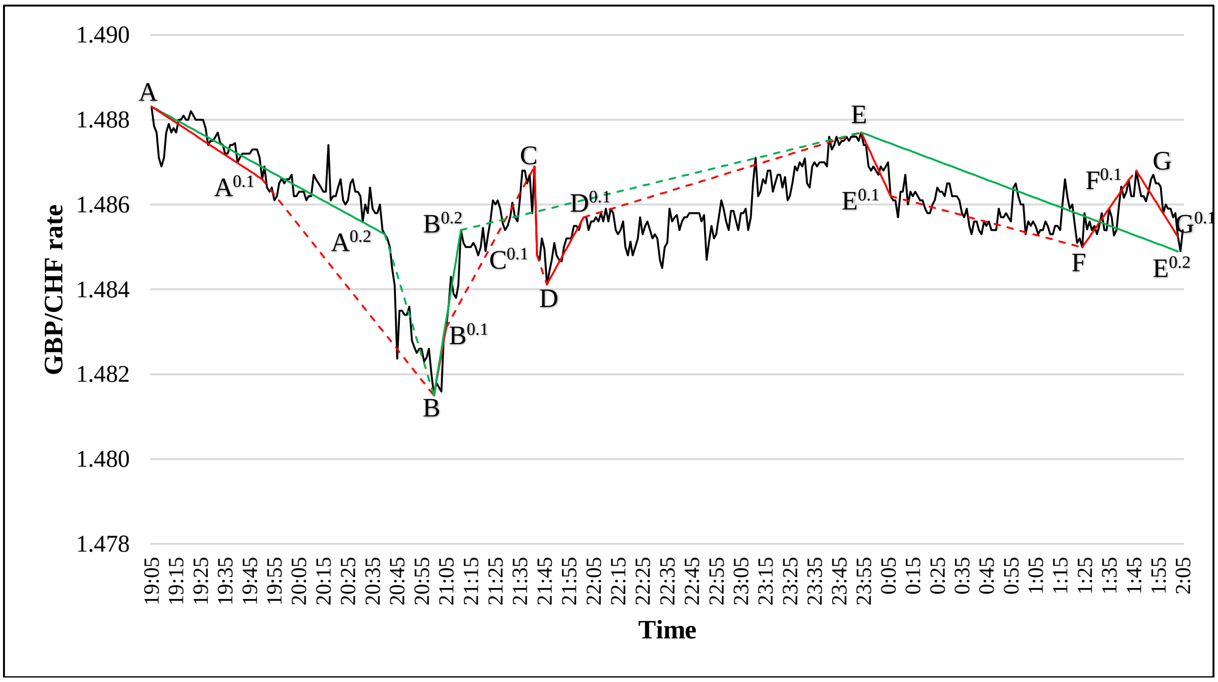

Fig. 5.3. The synchronization of the two DC summaries shown in Fig. 5.1 and Fig. 5.2.

5.2.2 The Boolean variable BBTheta

In this section, we use the concept of 'Big-Theta' to introduce a new Boolean variable named *BBTheta*. For each DC event of threshold *STheta*, we associate a value of the Boolean variable *BBTheta*. For example, let $BBTheta^1$ denote the value of *BBTheta* associated to the first DC event of threshold *STheta* which is [$AA^{0.1}$] in this case (see Fig. 5.3). In general, let $BBTheta^i$ be the value of *BBTheta* associated with the i^{th} DC event of the DC summary of threshold *STheta*. $BBTheta^i$ can be only *True* or *False*. The value of $BBTheta^i$ is defined as follows:

If the extreme point of the ith DC event of the DC summary of threshold STheta is also an extreme point of another DC event of threshold BTheta then BBTheta^{i} = True; otherwise BBTheta^{i} = False. In other words, if the ith point of EP^{STheta} is also a point of EP^{BTheta} then BBTheta^{i} = True; otherwise, BBTheta^{i} = False.

Based on Fig. 5.3, we recap that $EP^{0.001}$ = {A, B, C, D, E, F, G} and $EP^{0.002}$ = {A, B, E}. Point B is the extreme point of the second trend of threshold *STheta* = 0.10% (B is the second point in the set $EP^{0.001}$). The same point, B, is also the extreme point of another trend of threshold *BTheta* = 0.20% (point B is an element of $EP^{0.002}$). Therefore, in this case, $BBTheta^{2}$ = *True*. Similarly, D is the extreme point of the 4th trend of the DC summary of threshold *STheta* = 0.10% (D is the fourth point in the set $EP^{0.001}$). However, point D is not an extreme point of a DC event of threshold *BTheta* = 0.20% (point D is not an element of $EP^{0.002}$). Hence, $BBTheta^{4}$ = *False*. Given two DC summaries of the same price series, corresponding to two different thresholds, *STheta* and *BTheta*, we compute $BBTheta^{i}$ for each DC event of threshold *STheta*.

We notice that $BBTheta^{i}$ is '*True*' only if the prices' change between the i^{th} and $(i+1)^{th}$ extreme points listed in EP^{STheta} is larger than or equal to *BTheta*. We use Table 5.1, shown below, to explain this note. The first column from the left in Table 5.1 represents the index of the DC event of threshold *STheta* (i.e. 1st, 2nd, etc.). The column 'Extreme point' contains the extreme points resulting from the DC summary of threshold *STheta* (according to Fig. 5.3). The column 'Price at extreme point' shows the market's price at the indicated extreme point. We can catch the value of $BBTheta^{i}$ by calculating the magnitude of prices' change between the i^{th} and $(i+1)^{th}$ extreme points detected under the threshold *STheta*.

For example, to compute $BBTheta^{1}$ we calculate the price change between the prices of the 1st and 2nd extreme points shown in column 'Extreme point' (i.e. points A and B). In this example, the price change is:

$$(P_A - P_B) / P_A = (1.48831 - 1.48150) \ / \ 1.48831 = 0.00458^{h} \qquad (5.1)$$

The value of (5.1), 0.00458, is larger than *BTheta* (0.20%). Thus, $BBTheta^{1}$ = *True* as shown in column 'FBBTheta'. Similarly, to compute $BBTheta^{3}$ we calculate the price change between the prices of the 3rd and 4th extreme points shown in column 'Extreme point' (i.e. points C and D). In this example, the price change is $(1.48690 - 1.48412)\ /1.48690 = 0.00187 <$ *BTheta* (0.20%). Thus,

[h] In this example, the numbers are rounded to 5 decimal places.

BBTheta^3= *False* as shown in column 'FBBTheta'. The column 'BBTheta' embraces the set of all instances of *BBTheta*^{i}. We refer to this set as *BBTheta*.

Table 5.1: Example of DC events of threshold *STheta* and the computation of corresponding *BBTheta*^{i} based on Fig. 5.3.

DC event index (*STheta*)	Extreme point	Mid-price at extreme point	DCC point (*STheta*)	BBTheta
1	A	1.48831	$A^{0.1}$	*BBTheta*^{1}= *True*
2	B	1.48150	$B^{0.1}$	*BBTheta*^{2}= *True*
3	C	1.48690	$C^{0.1}$	*BBTheta*^{3}= *False*
4	D	1.48412	$D^{0.1}$	*BBTheta*^{4}= *False*
5	E	1.48770	$E^{0.1}$	*BBTheta*^{5}= *True*
6	F	1.48499	$F^{0.1}$	*BBTheta*^{6}= *False*
7	G	1.48680	$G^{0.1}$	*BBTheta*^{7}= *False*

5.3 Formulation of the forecasting problem

In this chapter, our task is to forecast the value of *BBTheta*. In other words, we are looking to forecast, at the DCC point of a DC event of threshold *STheta* (e.g. points $A^{0.1}$, $B^{0.1}$ from Table 5.1), whether the associated instance of *BBTheta* (shown in the column 'BBTheta' Table 5.1) is *True* or *False*. In this section, we introduce our proposed forecasting problem.

Table 5.2, shown below, simplifies the synchronization of the two DC summaries. We use Table 5.2 to provide an example of the proposed forecasting problem. Based on Table 5.2, we consider two uptrend DC events:

1. The DC event [$BB^{0.1}$] of threshold 0.10%. [$BB^{0.1}$] starts at time 21:00:00 and ends at time 21:05:00.
2. The DC event [$BB^{0.2}$] of threshold 0.20%. [$BB^{0.2}$] starts at time 21:00:00 and ends at time 21:10:00.

In column 'Point', $B^{0.1}$ denotes the DCC point of the DC event [$BB^{0.1}$], and $B^{0.2}$ denotes the DCC point of the DC event [$BB^{0.2}$]. We also note two facts:

1. Both DC events, [$BB^{0.1}$] and [$BB^{0.2}$], start at the same point B.
2. Point $B^{0.1}$ (which is observed at time 21:05:00, column 'Time') occurred before we observed point $B^{0.2}$ (at time 21:10:00).

Note that at point $B^{0.1}$ (i.e. at time 21:05:00) we can confirm that point B is the extreme point of an uptrend DC event of threshold *STheta* = 0.10%. However, at point $B^{0.1}$ we cannot confirm yet whether point B is also an extreme point of another uptrend DC event of threshold *BTheta* = 0.20% (i.e. whether *BBTheta*^{2} is *True* or *False*). At point B$^{0.2}$ (i.e. at time 21:10:00) we can confirm that point B is an extreme point of a DC event of threshold *BTheta* (i.e. *BBTheta*^{2} is *True*), but not before that. The objective, in this case, is to predict at point B$^{0.1}$ whether *BBTheta*^{2} is *True*.

Table 5.2: The synchronization of two DC summaries of GBP/CHF mid-prices sampled between 19:05:00 1/1/2013 and 00:06:00 2/1/2013. The two thresholds are: *STheta* = 0.10% and *BTheta* = 0.20%. Unnecessary minutes and prices are omitted.

Time	**Mid-price**	**DC Summary (*STheta* = 0.1%)**	**DC Summary (*BTheta* = 0.2%)**	**Point**
19:05:00	1.48831	start DC event (DOWNTREND)	start DC event (DOWNTREND)	A
………..				
19:50:00	1.48660	start OS event (DOWNTREND)		A$^{0.1}$
………..				
20:40:00	1.48530		start OS event (DOWNTREND)	A$^{0.2}$
………..				
21:00:00	1.48150	start DC event (UPTREND)	start DC event (UPTREND)	B
21:01:00	1.48180			
21:02:00	1.48170			
21:03:00	1.48159			
21:04:00	1.48280			
21:05:00	1.48310	start OS event (UPTREND)		B$^{0.1}$
21:06:00	1.48365			
21:07:00	1.48430			
21:08:00	1.48390			
21:09:00	1.48380			
21:10:00	1.48541		start OS event (UPTREND)	B$^{0.2}$
………..				
21:41:00	1.48690	start DC event (DOWNTREND)		C
21:42:00	1.48480	start OS event (DOWNTREND)		C$^{0.1}$
21:43:00	1.48470			
21:44:00	1.48520			
21:45:00	1.48495			
21:46:00	1.48412	start DC event (UPTREND)		D
………..				
22:01:00	1.48570	start OS event (UPTREND)		D$^{0.1}$
………..				
23:45:00	1.48770	start DC event (DOWNTREND)		E
………..				
00:06:00	1.48620	start OS event (DOWNTREND)		E$^{0.1}$

In other words, we want to predict at point $B^{0.1}$, whether the current uptrend will continue so that its total magnitude will reach a threshold of 0.20% (i.e. *BTheta*). In this example, [$BB^{0.1}$] is the second DC event of threshold *STheta* (see Table 5.1). Therefore, our objective is to forecast whether $BBTheta^2$ is *True*. In general, for the i^{th} DC event of threshold *STheta,* we want to predict whether the corresponding $BBTheta^i$ is *True.*

To conclude, in this chapter we propose to tackle the following forecasting problem: 'to forecast whether the current DC trend of threshold *STheta* will continue so that the total price change of this DC trend will be at least equal to *BTheta*'. This forecasting objective is shortened as to predict the Boolean variable *BBTheta*. To the best of our knowledge, no previous study has provided a similar formalization of this forecasting problem under the DC context. We believe that solving such forecasting problem under the DC framework could be the basis of a successful trading strategy (as we shall argue in Chapter 6).

5.4 Our approach to forecasting the end of a trend

In this section, we propose an approach to solving the forecasting problem presented in Section 5.3. The objective is to forecast for the i^{th} DC event of threshold *STheta* whether the corresponding $BBTheta^i$ is *True*. To this end, in this section, we introduce a novel DC-based indicator, which is also based on the concept of Big-Theta. We use the J48 procedure to make the forecast. Firstly, however, we must list some essential notations. We then introduce the novel DC-based indicator which will be sued as the independent variable. Finally, we briefly describe the adopted machine learning procedure, J48, which we will use to forecast *BBTheta*.

5.4.1 Directional Changes notations

In this section, we list some notations related to the DC framework. Table 5.3 shows these notations and provides a brief description for each. These notations are adopted from Tsang et al., [74] and have been explained in Section 4.2. In the context of this chapter, P_{EXT} denotes the price at which a new DC event starts. In Table 5.3, $PDCC^*$ denotes the price required to confirm the observation of a new DC event, either for an uptrend or a downtrend.

Table 5.3: List of some notations used in this thesis (source: Tsang et al. [74])

Name / Description	**Notation**
Threshold	*theta*
Current price	P_c
Price at extreme point: price at which one trend ends and a new trend starts.	P_{EXT}
The highest price, during an uptrend's OS event, required to confirm that the market's direction has changed to downtrend (i.e. to confirm a downtrend's DC event).	$P_{DCC\downarrow*} = P_{EXT} \times (1 - theta)$
The least price, during a downtrend's OS event, required to confirm that the market's direction has changed to uptrend (i.e. to confirm an uptrend's DC event).	$P_{DCC\uparrow*} = P_{EXT} \times (1 + theta)$
$PDCC^*$is the price of the theoretical directional change confirmation point of the current trend.	$PDCC^* = P_{DCC\downarrow*}$ if the current trend is downtrend; otherwise $PDCC^* = P_{DCC\uparrow*}$.

5.4.2 *The independent variable*

The accuracy of a forecasting model depends on the used independent variable(s). Many forecasting models rely on technical indicators to make a forecast (e.g. [6] [44] [46]). Our task is particularly difficult because, so far, no published work has provided a formal method as to how to apply existing technical indicators (e.g. Ehler Leading Indicator [89], Aroon indicator [32] RSI or ADX [90]) can be applied under the DC framework. Recently, Kampouridis and Otero [17] suggested that more research should be undertaken into defining new indicators emerging from the DC concept, in a manner similar to how technical indicators exist within traditional time series. Tsang et al., [74] introduced several DC-based indicators with the aim of profiling the financial markets. However, they did not examine the usefulness of these indicators for forecasting purposes.

In this section, we introduce a novel DC-based indicator named OSV_{BTheta}^{STheta}. OSV_{BTheta}^{STheta} is the single independent variable which we will use to forecast *BBTheta*. By definition each DC event is associated with an instance of the variable OSV_{BTheta}^{STheta}. Let $OSV_{BTheta}^{STheta\_i}$ be the instance of OSV_{BTheta}^{STheta} corresponding to the i^{th} DC event as observed under threshold *STheta*. To forecast $BBTheta^i$, of the i^{th} trend of the DC summary of threshold *STheta*, we should calculate $OSV_{BTheta}^{STheta\_i}$. We rewrite $OSV_{BTheta}^{STheta\_i}$ as *OSV*(EP_i^{STheta}, *BTheta*); with EP_i^{STheta} denoting the i^{th} extreme point of the DC summary of threshold *STheta*. Next, we will provide an example of how to calculate *OSV*(EP_i^{STheta}, *BTheta*) and then will state the general formula.

Considering the two sets of extreme points: $EP^{0.001}$ = {A, B, C, D, E, F, G} and $EP^{0.002}$ = {A, B, E} (previously defined in Section 5.2), we denote by $EP_i^{0.001}$ the i^{th} element of $EP^{0.001}$. For example: $EP_2^{0.001}$ and $EP_5^{0.001}$ represent points B and E respectively. Take the objective of predicting whether $BBTheta^2$ is *True*, at point $B^{0.1}$, we compute $OSV(EP_2^{0.001}, 0.002)$ as follows:

$$OSV_{0.002}^{0.001\_2} = OSV(EP_2^{0.001}, 0.002) = ((P_B - PDCC^{*0.002}) \,/\, PDCC^{*0.002}) / \, 0.002 \tag{5.2}$$

Where P_B is the price at point B. $PDCC^{*0.002}$ is the $PDCC^*$ computed with reference to the last confirmed DC event under threshold *BTheta*, which is, in this case, [$AA^{0.2}$]. Point A is an extreme point of a downward DC event of threhsold 0.20% (see Table 5.2). Hence, in this example, $PDCC^{*0.002} = P_A \times (1 - 0.002)$; where P_A is the price at point A.

In general, we define $OSV_{BTheta}^{STheta\_i}$ as:

$$OSV_{BTheta}^{STheta\_i} = OSV(EP_i^{STheta}, BTheta) = ((P_i^{STheta} - PDCC^{*BTheta}) / \, PDCC^{*BTheta}) / \, BTheta \tag{5.3}$$

Where EP_i^{STheta} is the i^{th} extreme point of a DC summary of threshold *STheta*. P_i^{STheta} is the price at the extreme point of the i^{th} DC event of threshold *STheta*. $PDCC^{*BTheta}$ is the $PDCC^*$ of the last confirmed DC event of threshold *BTheta*.

We provide a second example as to how to compute $OSV(EP_i^{STheta}, BTheta)$. The extreme point of the uptrend DC event [$EE^{0.1}$] is E. E is the 5th element of $EP^{0.001}$. Therefore, in this case, the objective is to predict whether $BBTheta^5$ is *True*. In this case, we should compute $OSV(EP_5^{0.001}$, 0.002) as in (5.4):

$$OSV_{0.002}^{0.001\_5} = OSV(EP_5^{0.001}, 0.002) = ((P_E - PDCC^{*0.002}) \,/\, PDCC^{*0.002}) \,/\, 0.002 \tag{5.4}$$

Where P_E is the price at point E. $PDCC^{*0.002}$ is the $PDCC^*$ computed with reference to the last confirmed extreme point of the DC summary of threshold *BTheta*, which is, in this case, [$BB^{0.2}$]. Note that [$BB^{0.2}$] is an uptrend DC event (see Table 5.2). Hence, $PDCC^{*0.002} = P_B \times (1 + 0.002)$; where P_B is the price at point B.

5.4.3 The decision tree procedure J48

In this chapter, we employ the decision tree procedure, J48, to find the relation between the two variables *BBTheta* and OSV_{BTheta}^{STheta}. J48 is the open-source Java implementation of the C4.5 algorithm [91]. J48 has three main steps. First, for each attribute λ it computes the normalized information gain ratio from splitting on λ. Let *λ_best* be the attribute with the highest normalized information gain. Second, it creates a decision node *nd* that splits on *λ_best*. Third, it recurs on the sub-lists obtained by splitting on *λ_best*, and adds those nodes as children of node *nd*. The three steps are repeated until a base case is reached.

5.5 Evaluation of our approach to forecasting DC: Experiments

In Section 5.4, we explained our approach to forecasting the change of market trend's direction under the DC context. In this section, we aim to examine the accuracy of our proposed forecasting approach. We test this approach in the FX market using eight currency pairs. We provide two sets of experiments: 1) the objective of the first set is of evaluate the accuracy of our forecasting approach, 2) the objective of the second set is to evaluate the impact of the value of *BTheta* on the accuracy of our forecasting approach. We firstly introduce a variable, named α, which we will use to measure the *True-False* imbalance in *BBTheta.*

5.5.1 Measuring the True-False imbalance

In Section 5.3 we introduced *BBTheta* as the Boolean dependent variable to be predicted. Some studies (e.g. [92]) have reported that the performance of some machine learning algorithms can be affected by the *True-False* imbalance in the dependent variable. In this section, we introduce a new variable named α. The objective of α is to measure the levels of *True-False* imbalance in the dependent variable *BBTheta*. α is measured as the fraction of *True* instances of *BBTheta.*

Let $nbTrends\_BTheta$ be the number of all trends obtained by directing a DC summary with threshold *BTheta* on a particular currency pair. Similarly, let $nbTrends\_STheta$ be the number of all trends obtained by running a DC summary with threshold *STheta*. We compute α as:

$$\alpha = \frac{nbTrends\_BTheta}{nbTrends\_STheta} \tag{5.5}$$

The value of α is interpreted as follow: if α = 0.70, then 70% of the instances of *BBTheta* are *True* and 30% are *False.* Note that, as explained in Section 4.2, the number of DC trends as observed under threshold *BTheta* is greater than the number of DC trends as observed under threshold *STheta* because *STheta* < *BTheta* (i.e. $nbTrends\ \ STheta > nbTrends\_BTheta$).

5.5.2 Experiment 5.1: Evaluating the accuracy of our forecasting approach

In order to evaluate the accuracy of our forecasting approach, we apply our forecasting approach to eight currency pairs: EUR/CHF, GBP/CHF, EUR/USD, GBP/AUD, GBP/JPY, NZD/JPY, AUD/JPY and EUR/NZD. Each currency pair is sampled minute-by-minute from 1/1/2013 to

31/7/2015 and split into in-sample and out-of-sample datasets. For each currency pair, we use the training (in-sample) set to learn the J48 decision tree before using the obtained tree to make the forecast over the testing (out-of-sample) set. The lengths of the in-sample and out-of-sample datasets are selected arbitrarily. The value of *STheta* and *BTheta* are chosen arbitrarily.

In preliminary experiments, we found that it would be better to forecast the uptrends and downtrends, of threshold *STheta*, separately. This practice — of splitting upward and downward trends for forecasting purposes — was also adopted by some studies (e.g. [80]). In this experiment, we consider, and save, the uptrends and downtrends as two independent datasets. Then, we divide each of the downtrends and uptrends into training (i.e. in-sample) and testing (i.e. out-of-sample) sets. As a benchmark, we chose to compare the accuracy of our forecasting model with the ARIMA model. The Autoregressive Integrated Moving Average model (ARIMA) has been reported in some studies (e.g. [20] [93]) as a good forecasting technique for time series. The AIRMA model has been used as a benchmark for forecasting models in many studies (e.g. [47] [94]).

5.5.3 Experiment 5.2: The impact of BTheta on the accuracy of our forecasting model

In this experiment, we aim to examine whether the accuracy of our approach can be affected by the value of *BTheta*. To this end, we consider the eight currency pairs listed in Experiment 5.1. In this experiment, *STheta* is fixed to 0.10% for each of the eight currency pairs. For each of these eight currency pairs, we apply our forecasting approach using ten different values of *BTheta* (from 0.13% to 0.22% with a step size of 0.01). For each currency pair, the training and testing periods are set to be the same as in Experiment 5.1. For each value of *BTheta*, we measure the corresponding accuracy for downtrends and uptrends separately.

We use the linear regression model to examine the impact of *BTheta* on the accuracy of our forecasting approach. Therefore, we apply the linear regression model, setting *BTheta* as the independent variable and the accuracy of our approach as the dependent variable. By analysing the *p*-value of *BTheta*, resulting from the linear regression, we can answer the question of whether *BTheta* has a significant linear impact on the accuracy of our approach.

In Section 5.5.1, we defined as ‘α’ the fraction of ‘*True*’ instances in *BBTheta*. α is employed to express the *True-False* imbalance in the dependent variable *BBTheta*. Note that the value of α depends on the value of *BTheta*. Consequently, in this experiment, by choosing ten different values of *BTheta*, we obtain ten different levels of *True-False* imbalance in the dependent variable *BBTheta* (i.e. ten different values of α). Thus, we can use the results of this experiment to study the accuracy of our forecasting approach under different levels of *True-False* imbalance.

5.6 Evaluation of our approach to forecasting DC: Results and discussion

5.6.1 Experiment 5.1: Evaluating accuracy of our forecasting approach

5.6.1.1 Experiment 5.1: Results

The objective of this experiment is to evaluate the accuracy of our approach to forecasting the change of a trend's direction, within the DC context, in the FX market. To this end, we apply our approach to eight currency pairs sampled minute-by-minute. For each currency pair, we consider the uptrends and downtrends separately, of the DC summary of threshold *STheta*. The values of the *STheta* and *BTheta* thresholds are chosen arbitrarily.

The experimental results and parameters' values are reported in Table 5.4. In Table 5.4, the column 'Currency Pair' specifies the considered currency pair. The columns 'STheta (%)' and 'BTheta (%)'denote the values of *STheta* and *BTheta* respectively. The column 'α' denotes the *True-False* imbalance resulting from the values of *STheta* and *BTheta*. The column 'Type of Trend' specifies whether the set of uptrends or downtrends, corresponding to the DC analysis of *STheta*, is in question. The columns 'Training period' and 'Testing Period' indicate the periods of the in-sample (training) and out-of-sample (testing) for each currency pair. The column 'Accuracy' shows the accuracy of our approach, computed as:

$$\text{Accuracy} = \frac{TP+TN}{N} \tag{5.6}$$

Where *N* is either the total number of upward or downward DC events (see the column 'Type of Trend' to know) obtained from running the DC summary of threshold *STheta*. *TP* is the number of correctly forecasted *True* instances of *BBTheta*. *TN* is the number of correctly forecasted *False* instances of *BBTheta*. All reported accuracies in Table 5.4 are measured for the out-of-sample period of each currency pair.

Table 5.4: The settings and results of applying our forecasting approach, and ARIMA model, to the eight currency pairs. All reported accuracies correspond to the out-of-sample testing periods.

Currency Pair	STheta (%)	BTheta (%)	α	Training Period	Testing Period	Type of Trend	Accuracy of our approach	ARIMA
EUR/CHF	0.10	0.13	0.63	From 1/1/2013 to 30/6/2015	From 1/7/2015 to 31/7/2015	Uptrends	0.81	0.59
						Downtrends	0.82	0.54
GBP/CHF	0.20	0.25	0.65	From 1/1/2013 to 30/4/2015	From 1/5/2015 to 31/7/2015	Uptrends	0.80	0.59
						Downtrends	0.82	0.58
EUR/USD	0.30	0.35	0.76	From 1/1/2013 to 31/12/2014	From 1/1/2015 to 31/7/2015	Uptrends	0.83	0.68
						Downtrends	0.85	0.70
GBP/AUD	0.10	0.13	0.51	From 1/1/2013 to 31/12/2014	From 1/1/2015 to 31/7/2015	Uptrends	0.81	0.72
						Downtrends	0.82	0.73
GBP/JPY	0.10	0.13	0.47	From 1/1/2013 to 31/12/2014	From 1/1/2015 to 31/7/2015	Uptrends	0.82	0.65
						Downtrends	0.81	0.62
NZD/JPY	0.10	0.13	0.54	From 1/1/2013 to 31/12/2014	From 1/1/2015 to 31/7/2015	Uptrends	0.81	0.59
						Downtrends	0.83	0.60
AUD/JPY	0.10	0.13	0.49	From 1/1/2013 to 31/12/2014	From 1/1/2015 to 31/7/2015	Uptrends	0.76	0.57
						Downtrends	0.76	0.58
EUR/NZD	0.10	0.13	0.45	From 1/1/2013 to 31/12/2014	From 1/1/2015 to 31/7/2015	Uptrends	0.76	0.59
						Downtrends	0.77	0.61

We then compare the accuracy of our approach with the ARIMA forecasting technique. For this purpose, we symbolize the '*True*' and '*False*' instances of *BBTheta* as '1' and '0' respectively. Then we apply ARIMA to the composed sequence of '1' and '0'. We use the function auto.arima() from the package 'forecast' of the statistical software R to predict *BBTheta*. The forecasting accuracy of the ARIMA model is reported in column 'ARIMA' in Table 5.4.

5.6.1.2 Experiment 5.1: Results' discussion

The objective of this experiment is to examine the accuracy of our forecasting approach. As can be seen in Table 5.4, for different testing periods and different selected values of *STheta* and *BTheta*, each of the obtained accuracies of our forecasting approach is above 0.75 (i.e. 75%). These results indicate that the proposed independent variable, OSV_{BTheta}^{STheta}, is very useful for forecasting *BBTheta*. The column 'ARIMA' in Table 5.4 shows the accuracy obtained by forecasting *BBTheta* using the ARIMA model. By comparing the accuracies of our approach (reported in column 'Accuracy of our approach') and the accuracy of the ARIMA technique (reported in column 'ARIMA') we notice that our approach outperforms ARIMA in all cases.

5.6.2 Experiment 5.2: The impact of BTheta on forecasting accuracy

The objective of this experiment is to examine whether the value of *BTheta* may affect the accuracy of the forecasting approach proposed in this chapter. To this end, we apply our forecasting approach to each of the considered eight currency pairs using ten different values of *BTheta*. To avoid tedious results we report the results of four currency pairs in this section. The results of the remaining four currency pairs are reported in Appendix B.

Table 5.5: Analyzing the impact of *BTheta* on the accuracy of our forecasting approach on the currency pair EUR/CHF. *STheta* is fixed to 0.10%. The testing period is 4 weeks in length. The reported accuracy corresponds to the testing (out-of-sample) period.

Uptrends of DC summary with *STheta* = 0.10%				Downtrends of DC summary with *STheta* = 0.10%		
BTheta (%)	**Accuracy**	**α**		**BTheta (%)**	**Accuracy**	**α**
0.13	0.82	0.63		0.13	0.82	0.63
0.14	0.78	0.54		0.14	0.78	0.54
0.15	0.74	0.48		0.15	0.75	0.48
0.16	0.72	0.42		0.16	0.72	0.42
0.17	0.70	0.37		0.17	0.70	0.37
0.18	0.67	0.33		0.18	0.67	0.33
0.19	0.65	0.30		0.19	0.66	0.30
0.20	0.64	0.27		0.20	0.64	0.27
0.21	0.63	0.25		0.21	0.64	0.25
0.22	0.62	0.22		0.22	0.62	0.22

5.6.2.1 Experiment 5.2: Results

The results of this experiment relating to the currency pairs EUR/CHF, GBP/CHF, EUR/USD and GBP/AUD are reported in Tables 5.5, 5.6, 5.7 and 5.8 respectively. Each table, with self-explanatory column headings, reports the results of applying our forecasting approach to the uptrends and downtrends of one currency pair. We will also use the results of this experiment to evaluate the performance of our forecasting approach under different levels of *True-False* imbalance in the dependent variable *BBTheta*.

Table 5.6: Analysing the impact of *BTheta* on the accuracy of our forecasting approach on the currency pair GBP/CHF: *STheta* is fixed to 0.10%. The testing period is 3 months. The reported accuracy corresponds to the testing (out-of-sample) period.

Uptrends of DC summary with *STheta* = 0.10%				Downtrends of DC summary with *STheta* = 0.10%		
BTheta (%)	**Accuracy**	**α**		**BTheta (%)**	**Accuracy**	**α**
0.13	0.82	0.64		0.13	0.81	0.64
0.14	0.79	0.55		0.14	0.77	0.55
0.15	0.75	0.49		0.15	0.75	0.49
0.16	0.73	0.42		0.16	0.71	0.42
0.17	0.71	0.37		0.17	0.70	0.37
0.18	0.69	0.33		0.18	0.68	0.33
0.19	0.67	0.30		0.19	0.66	0.30
0.20	0.64	0.27		0.20	0.64	0.27
0.21	0.64	0.25		0.21	0.64	0.25
0.22	0.62	0.23		0.22	0.63	0.23

Table 5.7: Analysing the impact of BTheta on the accuracy of our forecasting approach on the currency pair EUR/USD: STheta is fixed to 0.10%. The testing period is 7 months and 2 weeks. The reported accuracy corresponding to the testing (out-of-sample) period.

Uptrends of DC summary with *STheta* = 0.10%				Downtrends of DC summary with *STheta* = 0.10%		
BTheta (%)	**Accuracy**	**α**		**BTheta (%)**	**Accuracy**	**α**
0.13	0.82	0.64		0.13	0.80	0.64
0.14	0.80	0.56		0.14	0.77	0.56
0.15	0.77	0.50		0.15	0.74	0.50
0.16	0.74	0.45		0.16	0.72	0.45
0.17	0.71	0.40		0.17	0.70	0.40
0.18	0.70	0.36		0.18	0.67	0.36
0.19	0.68	0.33		0.19	0.65	0.33
0.20	0.64	0.30		0.20	0.66	0.30
0.21	0.65	0.28		0.21	0.63	0.28
0.22	0.64	0.26		0.22	0.62	0.26

Table 5.8: Analysing the impact of BTheta on the accuracy of our forecasting approach on the currency pair GBP/AUD: STheta is fixed to 0.10%. The testing period is 7 months. The reported accuracy corresponding to the testing (out-of-sample) period.

Uptrends of DC summary with *STheta* = 0.10%				Downtrends of DC summary with *STheta* = 0.10%		
BTheta (%)	**Accuracy**	**α**		**BTheta (%)**	**Accuracy**	**α**
0.13	0.81	0.51		0.13	0.82	0.51
0.14	0.78	0.49		0.14	0.78	0.49
0.15	0.75	0.47		0.15	0.75	0.47
0.16	0.72	0.45		0.16	0.72	0.45
0.17	0.70	0.41		0.17	0.70	0.41
0.18	0.68	0.37		0.18	0.68	0.37
0.19	0.67	0.34		0.19	0.67	0.34
0.20	0.66	0.30		0.20	0.65	0.30
0.21	0.65	0.28		0.21	0.64	0.28
0.22	0.63	0.26		0.22	0.63	0.26

5.6.2.2 *Experiment 5.2: Results' discussion*

The objective of this experiment was to ascertain whether the value of *BTheta* affects the accuracy of our approach. In each of the Tables 5.5 through 5.8, we note that the values in column 'Accuracy' increase as 'BTheta (%)' decreases. To statistically validate this note, we apply a linear regression model in which the column 'BTheta (%)' symbolises the independent variable and the column 'Accuracy' represents the dependent variable. We apply the linear regression model to each of these four tables, separately evaluating the uptrends and downtrends. We examine the *p*-value corresponding to *BTheta* for each linear regression analysis. The resulted *p*-value of the explanatory variable, 'BTheta (%)', is less than 0.01 in all cases. This is less than the common level of 0.05, which indicates that the value of *BTheta* can significantly impact the accuracy of our forecasting approach. Moreover, the R-square[i] (R^2), associated to the linear regression model, is greater than 0.90 in all four currency pairs (see for example Fig. 5.4 below). These results, of *p*-value and R^2, show that changes in *BTheta* are associated with changes in accuracy.

---

[i] R-squared is a statistical measure of how close the accuracies are to the fitted regression line (see Fig. 5.4 below). See https://en.wikipedia.org/wiki/Coefficient_of_determination.

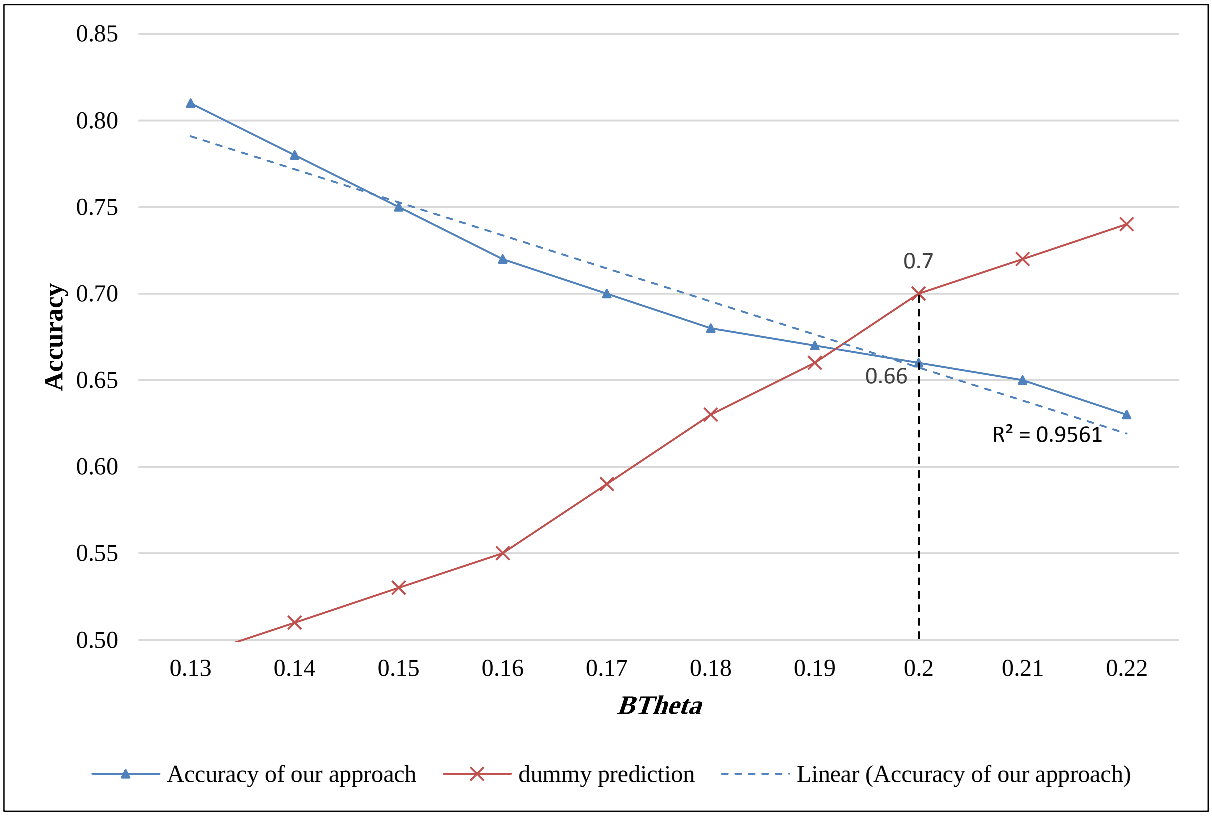

Fig. 5.4. The illustration of the variation in accuracy of forecasting the uptrends of GBP/AUD as a function of *BTheta* (see Table 5.8). The blue and orange lines denote the curves of the accuracy of our approach and the dummy prediction respectively. The blue dashed line symbolizes the linear regression line that fit most with the 'Accuracy of our approach'.

Furthermore, as stated in Section 5.5.3, the results, shown in Tables 5.5 through 5.8, also allow us to examine the performance of our proposed forecasting model under different levels of *True-False* imbalance in the dependent variable. These results highlight two points:

- The accuracy of our approach is quite good for most levels of *True-False* imbalance in the dependent variable *BBTheta*. For example, in the case of Table 5.5, we note that α ranges between 0.22 (i.e. 22% of *BBTheta* instances are *True*) and 0.63 (i.e. 63% of *BBTheta* instances are *True*). The corresponding accuracies range between 0.62 and 0.82. As for the results corresponding to GBP/CHF, shown in Table 5.6, we note that α ranges between 0.23 and 0.64. The corresponding accuracies range between 0.62 and 0.82. The results obtained based on EUR/USD are reported in Table 5.7, from which we can see the range of α is between 0.26 and 0.64. The range of accuracy is between 0.62 and 0.82. The results of GBP/AUD, shown in Table 5.8, match with the results reported in Tables 5.5 through 5.7. We consider this range of accuracy (between 0.62 and 0.82) to be fairly good.

- These results also suggest that the accuracy of our forecasting approach is reasonably consistent across the four considered currency pairs. In each table, the accuracies range between 0.62 and 0.82.

However, these results also highlight two limitations:

- In general, the accuracy of our forecasting model decreases as the difference between *BTheta* and *STheta* increases (see Fig. 5.4 for example).
- When the difference between *BTheta* and *STheta* becomes greater than a specific value, our model cannot outperform a 'dummy' predictor (which always predicts *False*). For example, in the case of GBP/AUD (Table 5.8) we note that for a *BTheta* greater than, or equal to, 0.20% (as in the case of an 'uptrend') the corresponding α becomes less than, or equal to, 0.30. In such a situation, the accuracy of the dummy prediction is expected to be larger than 0.70, whereas, in this case, the forecasting accuracy of our approach is less than, or equal to, 0.70 (see Fig. 5.4 above). The same observation holds true in the Tables 5.5 through 5.7 for any *BTheta* greater than, or equal to, 0.20%.

To conclude, in this section, we reported and analysed the results of applying our forecasting approach to four currency pairs. The results of the linear regression analysis show that *BTheta* does have a significant impact on the accuracy of our approach. We want to highlight that the analysis of the remaining four currency pairs, reported in Appendix B, supports this conclusion.

5.7 Summary and conclusion

In this chapter, we addressed the problem of forecasting the change of a trend's direction within the DC framework. Our first objective was to formalize the considered forecasting problem under the DC context. The second objective was to provide a solution for this problem.

The first contribution of this chapter was in formulating the prediction of the change of direction of market's trend under the DC framework. For this purpose we suggested to track prices movement using 2 DC thresholds: *STheta* and *BTheta*. Our task was to forecast whether a DC trend, as observed under threshold *STheta*, will continue so that its total magnitude will be at least equal to *BTheta*. We introduced a new concept named Big-Theta which originates from the DC framework. The notion of Big-Theta states that a DC event of threshold *BTheta* will embrace at least one DC event of a smaller threshold *STheta* (with *BTheta* > *STheta*). We used the concept of Big-Theta to introduce the Boolean variable named *BBTheta* (Section 5.2.1). The value of *BBTheta* denote the fact of whether the total price change of a DC trend, as observed under the threshold *STheta*, reaches *BTheta* (Section 5.3). Thus, our objective was to forecast *BBTheta*.

Our second contribution was in identifying one novel DC-based indicator as the independent variable, and in proving that it is relevant to our prediction problem. This DC-based indicator, also based on the concept of Big-Theta, is OSV_{BTheta}^{STheta}(Section 5.4). We used the machine learning procedure J48 to detect the relation between OSV_{BTheta}^{STheta} and $BBTheta$.

We examined the performance of our forecasting approach using eight currency pairs sampled minute-by-minute (Section 5.5). The results pointed out that our approach outperforms the traditional forecasting technique ARIMA (Table 5.4, Section 5.6.1).The results indicated that the accuracy of our approach ranges between 62% and 80% (Section 5.6.2). We consider this range as pretty good. However, the results also suggested that the accuracy of our approach decreases as the difference between *STheta* and *BTheta* increases. When this difference reaches a definite level, our approach is outperformed by a dummy prediction, which keeps predicting *False* (Section 5.6.2).

To conclude, we believe that this is the first attempt to forecast the change of a trend's direction under the DC-framework. Our contribution is in formulating the forecasting problem and proposing a solution. We shortened the formalization of this problem as to forecast one Boolean variable named *BBTheta*. The proposed solution comprises the discovery of a novel DC-based indicator named OSV_{BTheta}^{STheta}. We demonstrated that OSV_{BTheta}^{STheta} is helpful in forecasting *BBTheta*. We argued that our forecasting approach is more accurate than the ARIMA model and that the change of a trend's direction is predictable under the DC framework with pretty good accuracy.

6 TSFDC: A Trading Strategy Based on Forecasting Directional Changes

The previous chapter introduced an approach to forecasting the change of the direction of market's trend under the Directional Changes (DC) framework. Based on our findings in Chapter 5, this chapter aims to develop a successful trading strategy founded on the established forecasting model. In order to examine the success of this proposed trading strategy, called TSFDC, we provide several experiments using eight currency pairs from the FX market. The results suggest that TSFDC can generate returns of more than 500% within seven months. We argue that TSFDC outperforms other DC-based trading strategies.

6.1 Introduction

The objective of this thesis is to explore, and consequently to provide a proof of, the usefulness of the DC framework as the basis of profitable trading strategies. In Chapter 3, we suggested that existing trading strategies can mostly be categorised into two groups (see Sections 3.2 and 3.3). The first group contains trading strategies that are based on forecasting models (e.g. [6] [41] [42] [43] [44] [45]). The second group consists of trading strategies that do not rely on any forecasting model (e.g. [3] [56] [57] [58] [95]). In line with literature, in this thesis we aim to develop two DC-based trading strategies – one strategy belongs to the first identified group of trading strategies and the second strategy belongs to the second group.

In chapter 5, we formalized the problem of forecasting the change of trend's direction under the DC framework. In this chapter, we develop a trading strategy named 'Trading Strategy based on Forecasting DC'; henceforth TSFDC. TSFDC relies on the forecasting model developed in Chapter 5 to decide when to start a trade. We provide a set of experiments to examine the performance of TSFDC using eight currency pairs from the FX market.

The chapter continues as follows: Section 6.2 provides a brief summary of the forecasting model introduced in Chapter 5. We present TSFDC and its trading rules in Section 6.3. We discuss the selection and preparation of the used datasets in Section 6.4. The details of the experiments, conducted to evaluate the performance of TSFDC, are provided in Section 6.5. Section 6.6 reports and discusses the results of these experiments. We compare our trading strategy with other DC-based strategies in Section 6.7. Finally, we summarize the major findings of this chapter in Section 6.8.

6.2 Forecasting DC: A brief overview

In Chapter 5 we formalized a new forecasting problem under the DC framework. To formalize this objective, we tracked price changes with two thresholds simultaneously: *BTheta* and *STheta* (with *BTheta* > *STheta*; as in Fig. 6.1 below). The objective of which was to forecast whether the total price change of a DC trend, as observed under the threshold *STheta*, reaches the selected threshold of *BTheta*.

We defined a Boolean variable named *BBTheta* (Section 5.2.2). Each DC trend of threshold *STheta* is associated with a value of *BBTheta* which is *True* if, and only if, the magnitude of total price change of this trend is at least equal to *BTheta*. Our aim was to predict *BBTheta* at the DC confirmation point (DCC point) of a DC event of threshold *STheta*. For example, in Fig. 6.1 [$AA^{0.1}$] denote the first DC event observed under threshold *STheta* (0.10%). Let $BBTheta^1$ denote the value of *BBTheta* corresponding to [$AA^{0.1}$]. Point $A^{0.1}$ is the DCC point of the DC event [$AA^{0.1}$]. At $A^{0.1}$ we don't yet know whether $BBTheta^1$ is *True*. In this example, we want to forecast $BBTheta^1$ at $A^{0.1}$. Note that, in this case, at point $A^{0.2}$ we are able to confirm that $BBTheta^1$ is *True*; but not before that.

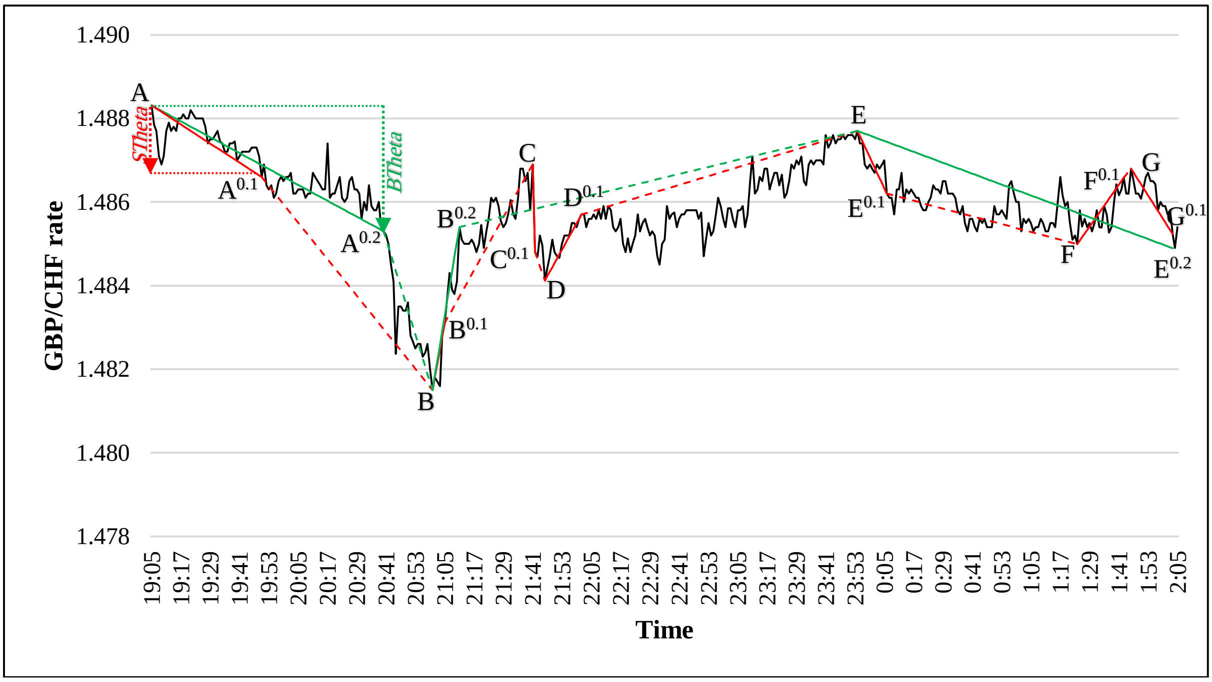

Fig. 6.1. The synchronization of two DC summaries with two thresholds: *STheta* = 0.10% (in red lines) and *BTheta* = 0.20% (in green lines) for GBP/CHF rate sampled minute by minute from 1/1/2013 19:05:00 to 1/2/2013 02:05:00.

Generally, for each DC event, of threshold *STheta*, we associate a value of *BBTheta*. In Chapter 5, we provided an approach to forecasting the value of *BBTheta* associated to each DC event of threshold *STheta* (Section 5.4). In many cases, the accuracy of our forecasting model was over 80%

(see Table 5.4, Section 5.6.1). In this chapter, our objective is to develop a successful trading strategy based on this forecasting model.

6.3 Introducing the trading strategy TSFDC

In this section we introduce a DC based trading strategy named 'Trading Strategy based on Forecasting DC' (TSFDC). TSFDC is designed as a contrarian trading strategy (i.e. TSFDC generates buy and sell signals against the market's trend) and is based on the forecasting model established in Chapter 5. We present two versions of TSFDC: TSFDC-down and TSFDC-up. The former is to be applied if the market exhibits a downward trend under the DC context, with the latter employed in the opposite case. The following explains how TSFDC-down and TSFDC-up operate.

6.3.1 TSFDC-down

TSFDC-down is only applicable when the market is in a downtrend. TSFDC-down relies on the forecasting approach presented in Chapter 5 to decide when to trigger a buy signal. Let $BBTheta^i$ be the value of *BBTheta* associated with the i^{th} DC event of threshold *STheta* (see Section 5.2.2). Let $FBBTheta^i$ denote the forecasted value of $BBTheta^i$. The value of $FBBTheta^i$ is determined based on the forecasting model described in Chapter 5. Note that we compute the value of $FBBTheta^i$ at the DCC point of the i^{th} DC event of threshold *STheta* (e.g. $FBBTheta^1$ is calculated at point $A^{0.1}$ in Fig. 6.1 above). If $FBBTheta^i$ is *True*, then we anticipate that the i^{th} DC trend, observed under threshold *STheta*, will continue, so that the total price change of this trend will be at least equal to *BTheta*. TSFDC-down relies on $FBBTheta^i$ to decide when to trigger a buy signal. More particularly, there are two conditions under which TSFDC-down generates buy signal (depending on whether $FBBTheta^i$ is *True* or *False*):

At the DCC point for the i^{th} *DC trend (STheta), we predict* $FBBTheta^i$*:*

- *Rule TSFDC-down.1 (generate buy signal):*

  If $FBBTheta^i = False$ then generate buy signal.

- *Rule TSFDC-down.2 (generate buy signal):*

  If ($FBBTheta^i = True$) and (we confirm a new DC event of threshold *BTheta*) then generate buy signal.

- *Rule TSFDC-down.3 (generate sell signal):*

  If ($P_c \geq P_{DCC\uparrow*}$) then generate sell signal.

With P_c denoting the current price and $P_{DCC\uparrow*}$ denotes the minimum prices required to confirm the occurrence of the succeeding uptrend DC event of threshold *STheta*. If the condition of *Rule*

TSFDC-down.1 is satisfied, then TSFDC-down generates a buy signal at the DCC point as observed under the threshold *STheta*. On the other hand, if both conditions of *Rule TSFDC-down.2* are fulfilled then TSFDC-down generates a buy signal at the DCC point as observed under the threshold *BTheta*. The condition of *Rule TSFDC-down.3* denotes the case under which we confirm the DCC point of a new uptrend DC event of threshold *STheta*. *Rule TSFDC-down.3* is applicable only if a buy signal has been triggered (either by *TSFDC-down.1* or *TSFDC-down.2*). *TSFDC-down.3* plays two simultaneous roles: *take-profit* and *stop-loss*. When *TSFC-down.3* triggers a sell signal, it may incur losses (hence, functioning as *stop-loss*) or generates profits (thus, working as *take-profit*).

Table 6.1, shown below, exemplifies two DC summaries with two different thresholds 0.10% (*STheta*) and 0.20% (*BTheta*). We use Table 6.1 to provide two trading scenarios that demonstrate the function of TSFDC-down's trading rules. *Scenario 1:* Consider the DC event [$AA^{0.1}$] (of threshold *STheta* = 0.10%) which starts at point A (see column 'Point', Table 6.1).

a) [$AA^{0.1}$] refers to a downward DC event of threshold 0.10% which starts at time 19:05:00 (shown in column 'Time', Table 6.1). Point $A^{0.1}$ is the DCC point of [$AA^{0.1}$] as observed at time at time 19:50:00. At point $A^{0.1}$, assume that we predict[j] *FBBTheta*^{1} is *True* (as shown in column 'FBBTheta').
b) [$AA^{0.2}$] refers to a downward DC event of threshold 0.20% which starts at time 19:05:00. Point $A^{0.2}$ is the DCC point of [$AA^{0.2}$] as observed at time 20:40:00.
c) Based on a) and b), the conditions of *Rule TSFDC-down.2* are fulfilled at point $A^{0.2}$. Thus, TSFDC-down initiates a buy signal at point $A^{0.2}$ (i.e. at time 20:40:00).
d) [$BB^{0.1}$] refers to the uptrend DC event, of the threshold 0.10%, that directly follow [$AA^{0.1}$]. At time 21:05:00, we confirm the DCC point of [$BB^{0.1}$]; which is $B^{0.1}$. Following *Rule TSFDC-down.3*, TSFDC-down will trigger a sell signal at point $B^{0.1}$.

Scenario 2: Consider the downward DC event [$CC^{0.1}$] which starts at time 21:41:00.

a) [$CC^{0.1}$] refers to a downward DC event of threshold 0.10% which starts at time 21:41:00. At point $C^{0.1}$ (at time 21:42:00) assume that we predict *FBBTheta*^{3} is *False* (as shown in column 'FBBTheta').
b) Based on a), the condition of *Rule TSFDC-down.1* holds at point $C^{0.1}$. Thus, TSFDC-down initiates a buy signal at point $C^{0.1}$.

[j] As [$AA^{0.1}$] is the first DC event in Table 6.1, our objective is to forecast the value of *BBTheta*^{1}. Here, we denote by *FBBTheta*^{1} the forecasted value of *BBTheta*^{1}.

c) [$DD^{0.1}$] refers to the upward DC event of threshold 0.10% which directly follow [$CC^{0.1}$]. At time 22:01:00, we confirm the DCC point of [$DD^{0.1}$]; which is $D^{0.1}$. Following *Rule TSFDC-down.3*, TSFDC-down will trigger a sell signal at point $D^{0.1}$.

Table 6.1: The synchronization of two DC summaries of GBP/CHF mid-prices sampled between 19:05:00 1/1/2013 and 00:06:00 2/1/2013. The two thresholds are: *STheta* = 0.10% and *BTheta* = 0.20%. Unnecessary minutes and prices are omitted. The '*True*' and '*False*' shown in column 'FBBTheta' are hypothetical (for explanation purpose only).

Time	**Mid-price**	**DC Summary (*STheta* = 0.1%)**	**DC Summary (*BTheta* = 0.2%)**	**Point**	**FBBTheta**
19:05:00	1.48831	start DC event (DOWNTREND)	start DC event (DOWNTREND)	A	
………..					
19:50:00	1.48660	start OS event (DOWNTREND)		$A^{0.1}$	*True*
………..					
20:40:00	1.48530		start OS event (DOWNTREND)	$A^{0.2}$	
………..					
21:00:00	1.48150	start DC event (UPTREND)	start DC event (UPTREND)	B	
21:01:00	1.48180				
21:02:00	1.48170				
21:03:00	1.48159				
21:04:00	1.48280				
21:05:00	1.48310	start OS event (UPTREND)		$B^{0.1}$	*True*
21:06:00	1.48365				
21:07:00	1.48430				
21:08:00	1.48390				
21:09:00	1.48380				
21:10:00	1.48541		start OS event (UPTREND)	$B^{0.2}$	
………..					
21:41:00	1.48690	start DC event (DOWNTREND)		C	
21:42:00	1.48480	start OS event (DOWNTREND)		$C^{0.1}$	*False*
21:43:00	1.48470				
21:44:00	1.48520				
21:45:00	1.48495				
21:46:00	1.48412	start DC event (UPTREND)		D	
………..					
22:01:00	1.48570	start OS event (UPTREND)		$D^{0.1}$	*False*
………..					
23:45:00	1.48770	start DC event (DOWNTREND)		E	
………..					
00:06:00	1.48620	start OS event (DOWNTREND)		$E^{0.1}$	

6.3.2 TSFDC-up

TSFDC-up is the mirror of TSFDC-down in that it is only applicable when the market exhibits an upward trend. TSFDC-up uses $FBBTheta^i$ (i.e. the forecasted value of $BBTheta^i$) to decide when to open a position. TSFDC-up relies on $FBBTheta^i$ to decide when to trigger a sell signal. More particularly, there are two conditions under which TSFDC-up generates sell signal (depending on whether $FBBTheta^i$ is *True* or *False*):

At the DCC point for the i^{th} *DC trend (STheta), we predict* $FBBTheta^i$*:*

- *Rule TSFDC-up.1 (generate sell signal):*

  If $FBBTheta^i$ = *False* then generate sell signal.

- *Rule TSFDC-up.2 (generate sell signal):*

  If ($FBBTheta^i$ = *True*) and (we confirm a new DCC point of DC event of threshold *BTheta*) then generate sell signal.

- *Rule TSFDC-up.3 (generate buy signal):*

  If ($P_c \leq P_{DCC\downarrow *}$) then generate buy signal.

Note that if the condition of *Rule TSFDC-up.1* is *True* then TSFDC-up generates a sell signal at the DCC point observed under threshold *STheta*. On the other hand, if the conditions of *Rule TSFDC-up.2* are both *True* then TSFDC-up triggers a sell signal at the DCC point observed under threshold *BTheta*. *Rule TSFDC-up.3* denotes the case under which we confirm the DCC point for a new DC downtrend of threshold *STheta*. *Rule TSFDC-up.3* is applicable only if a sell signal has been triggered (either by *TSFDC-up.1* or *TSFDC-up.2*). When TSFDC-up closes a position, it may generate profits or losses. Rule *TSFDC-up.3* has the same roles of, *taking-profits* and *stop-loss*, as *Rule TSFDC-down.3*.

We use Table 6.1, shown above, to provide two trading scenarios in demonstration of how TSFDC-up's rules are applied. *Scenario 1*: Consider the uptrend DC event [$BB^{0.1}$] (of threshold *STheta* = 0.10%):

a) [$BB^{0.1}$] refers to an upward DC event of threshold 0.10% which starts at time 21:00:00 (shown in column 'Time', Table 6.1). Point $B^{0.1}$ is the DCC point of [$BB^{0.1}$] as observed at time at time 21:05:00. At point $B^{0.1}$, assume that we predict $FBBTheta^2$ is *True* (as shown in column 'FBBTheta').
b) [$BB^{0.2}$] refers to an upward DC event of threshold 0.20% which starts at time 21:00:00. Point $B^{0.2}$ is the DCC point of [$BB^{0.2}$] as observed at time 21:10:00.

c) Based on a) and b), the conditions of *Rule TSFDC-up.2* are fulfilled at point $B^{0.2}$. Thus, TSFDC-up initiates a sell signal at point $B^{0.2}$ (i.e. at time 21:10:00).

d) [$CC^{0.1}$] refers to the uptrend DC event, of the threshold 0.10%, that directly follow [$BB^{0.1}$]. At time 21:42:00, we confirm the DCC point of [$CC^{0.1}$]; which is $C^{0.1}$. Following *Rule TSFDC-up.3*, TSFDC-up will trigger a buy signal at point $C^{0.1}$.

Scenario 2: Consider the upward DC event [$DD^{0.1}$] (of threshold *STheta* = 0.10%).

a) At time 22:01:00, at point $D^{0.1}$, assume that we predict *FBBTheta*[4] is *False* (as shown in column 'FBBTheta').

b) Based on a), the condition of *Rule TSFDC-up.1* holds at point $D^{0.1}$. Thus, TSFDC-up initiates a sell signal at point $D^{0.1}$.

c) [$EE^{0.1}$] refers to the downward DC event of threshold 0.10% which directly follow [$DD^{0.1}$]. At time 00:06:00, we confirm the DCC point [$EE^{0.1}$]; which is $E^{0.1}$. Following *Rule TSFDC-up.3*, TSFDC-up will trigger a buy signal at point $E^{0.1}$.

6.4 Preparation of the datasets and other considerations

This section provides essential notes regarding the selection and preparation of the datasets used in our experiments. When designing our experiment approach, we paid attention to some important concerns put forward by some studies (e.g. [61] [96]) that highlight serious experimental flaws presented in several published papers. In the context of our experiments, we consider the following points:

6.4.1 Data selection

Pardo [61] emphasizes the importance of backtesting (see Section 3.4 for definition of backtesting) using a set of assets with different trends. Such variation in the selected dataset will help to test the performance of the trading strategy under different market scenarios. This broadening helps in avoiding any bias towards particular patterns. In this chapter, we consider eight currency pairs, namely: EUR/CHF, GBP/CHF, EUR/USD, GBP/AUD, GBP/JPY, NZD/JPY, AUD/JPY, and EUR/NZD. The mid-prices of these currency pairs are sampled minute-by-minute during a period of 31 months between 01/01/2013 and 31/07/2015. Our focus, in this section, is to examine the variation of the trends of these currency pairs during the (out-of-sample) trading period which lasts from 1/1/2015 to 31/7/2015. The training (in-sample) period took place between 1/1/2013 and 31/12/2014. Holidays and weekends are not included in our datasets.

In this section, we investigate the variation of the trends of the selected currency pairs. Variation is important because some studies (e.g. [61]) have shown that trend changes can have a large and often negative impact on trading performance. Fig. 6.2, shown below, depict the normalized daily exchange rates of the selected eight currency pairs throughout the considered trading period of seven months (from 1/1/2015 to 31/7/2015). It provides a visual indication as to the existence of a variety of trends in our dataset over the considered trading period. The variation of the trends, as visualized in Fig. 6.2, indicate that we avoid possible bias in our experiment, which would have occurred had we only picked currency pairs with similar trends during the selected trading period.

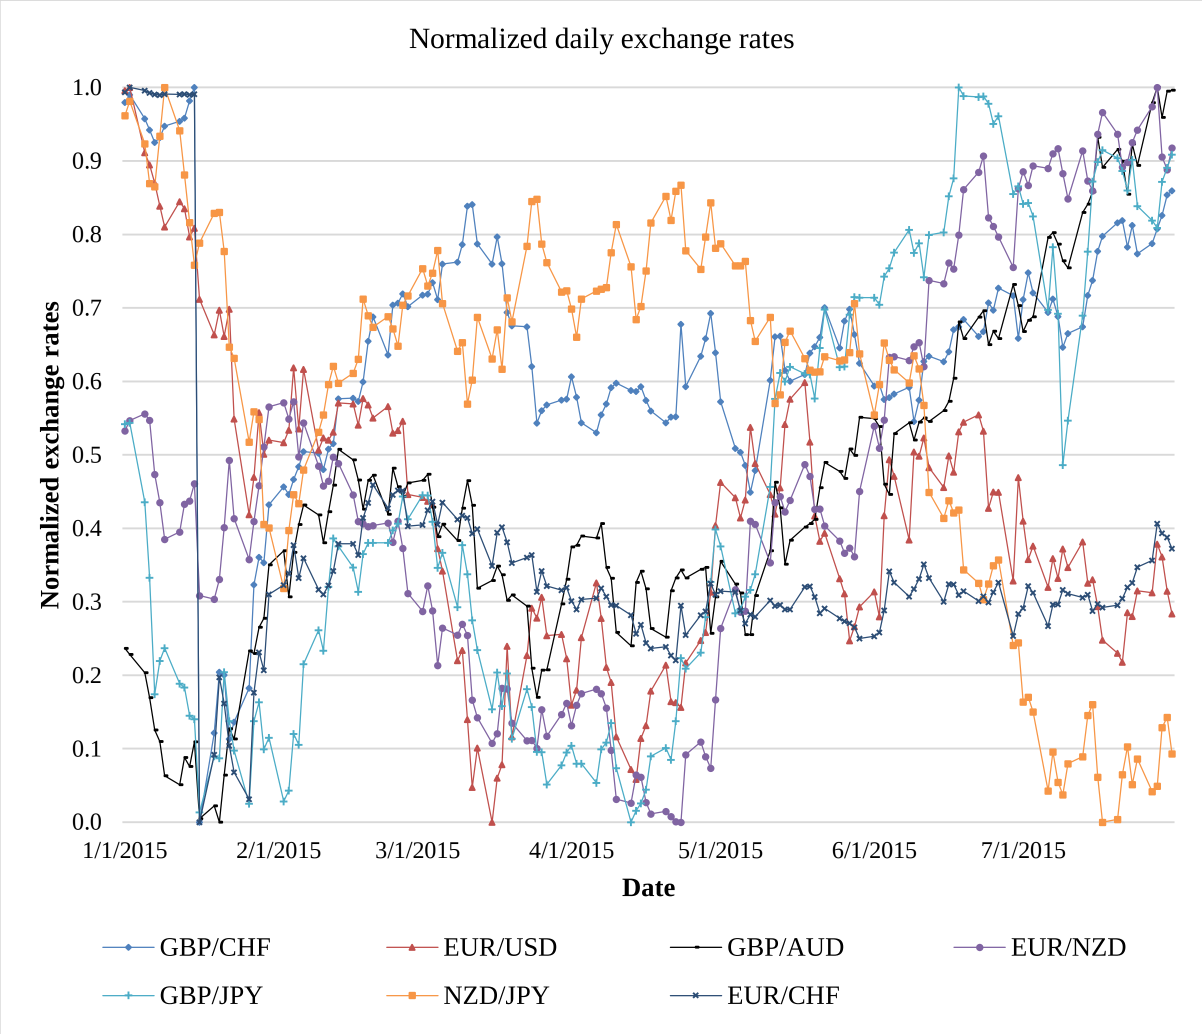

Fig. 6.2 Normalized daily exchanges rate of the 8 selected currency pairs between 1/1/2015 and 31/7/2015. This figures aims to illustrate the divergence of trends of selected currency pairs. In order to avoid excessive points, we use a daily exchange rate instead of a minute-based exchange rates.

Fig. 6.2 indicates that the selected currency pairs exhibit different trends during the trading period. The trends of the training period, considered from 1/1/2013 to 31/12/2014, was not studied as it is not specifically related to the evaluation of the performance of TSFDC during the out-of-sample period.

6.4.2 Evaluating the performance of a trading strategy

Many studies define success solely on the grounds of forecasting accuracy and win ratios, which, practically, has little value [97] [98]. Practically, an investor might be interested in other metrics that evaluate the risk and risk-adjusted performance of a given trading strategy [62] [99]. In this chapter, we evaluate the performance of TSFDC using a range of evaluation metrics such as: profit factor, maximum drawdown, Sharpe ratio, Jensen's Alpha, Beta and others (see Section 3.4). These metrics are marked as adequate for a decent evaluation of the performance of a given trading strategy [62] [61].

6.4.3 Model training and testing process

Pardo [61] suggests the adoption of a rolling window approach as being more reliable to test a trading strategy. This approach is usually used for evaluating trading systems and establishes a more rigorous and convincing methodology. This method involves splitting the data into overlapping training-applied sets and, on each cycle, moving each set forward through the time series. This methodology tends to result in more robust models due to more frequent retraining and large out-of-sample data sets (increasing training processing requirements but also resulting in models which adapt more quickly to changing market conditions). In our experiments, we train and test TSFDC using a monthly-basis rolling window as we will explain next.

6.4.4 Preparing the rolling windows

Our experiments examine eight currency pairs: EUR/CHF, GBP/CHF, EUR/USD, GBP/AUD, GBP/JPY, NZD/JPY, AUD/JPY, and EUR/NZD and consider the minute-by-minute transaction mid-prices of these currency pairs for 31 months: from 1/1/2013 to 31/7/2015. Given that the preparation process of the rolling windows for each currency pair is the same, we will use the two-step preparation of the rolling windows, explained below, for the currency pairing GBP/CHF as an example to detail our method.

6.4.4.1 Step 1: Producing DC summary for the dataset

We run the Directional Change (DC) summary on the initial dataset of GBP/CHF sampled minute-by-minute over 31 months. Section 4.2 provides a detailed description of the DC summary. In simple terms, given a threshold *STheta*, we achieve, through DC summary, the identification of all DC and OS events in the initial dataset (see Table 6.2 below). Arbitrarily, we set *STheta* = 0.10% and produce the DC summary to the initial dataset of GBP/CHF. Let *GBPCHF_DC0.1* be the output of this DC summary. Part of *GBPCHF_DC0.1* is illustrated in Table 6.2. *GBPCHF_DC0.1* comprises the date, time and the price of each observation of the initial dataset. In Table 6.2, the column ‘Event Type’ marks the occurrence of any DC or OS event that starts at the specified date and time (see Section 4.2 for more info about DC summary).

Table 6.2: An example of DC summary using GBP/CHF mid-prices sampled minute-by-minute from 21:41:00 to 22:01:00 (UK time).

Date	Time	Mid-price	Event Type
1/1/2013	21:41:00	1.48690	start DC event (DOWNTREND)
1/1/2013	21:42:00	1.48480	start OS event (DOWNTREND)
1/1/2013	21:43:00	1.48470	
1/1/2013	21:44:00	1.48520	
1/1/2013	21:45:00	1.48495	
1/1/2013	21:46:00	1.48412	start DC event (UPTREND)
1/1/2013	21:47:00	1.48440	
1/1/2013	21:48:00	1.48470	
1/1/2013	21:49:00	1.48510	
1/1/2013	21:50:00	1.48480	
1/1/2013	21:51:00	1.48470	
1/1/2013	21:52:00	1.48466	
1/1/2013	21:53:00	1.48500	
1/1/2013	21:54:00	1.48520	
1/1/2013	21:55:00	1.48520	
1/1/2013	21:56:00	1.48520	
1/1/2013	21:57:00	1.48550	
1/1/2013	21:58:00	1.48550	
1/1/2013	21:59:00	1.48540	
1/1/2013	22:00:00	1.48560	
1/1/2013	22:01:00	1.48570	start OS event (UPTREND)

6.4.4.2 *Step 2: Composing the rolling windows*

Motivated by the recommendation of Pardo [61], we use a rolling window approach (see Fig. 6.3 below) to evaluate the performance of our proposed trading strategy. As the dataset *GBPCHF_DC0.1* covers 31 months, we compose seven rolling windows — each of which comprises a training window (24 months in length) and an applied window (1 month in length).

So that the overall trading period, throughout the seven rolling windows, is seven months. The lengths of the training and applied windows are set arbitrarily. Note that we measure the length of the training and applied windows as a function of months, not as a fixed number of days. For example, the training period of the second rolling window lasts from 1/2/2013 to 31/1/2015 (i.e. 24 months). The associated applied window lasts from 1/2/2015 00:01:00 to 28/2/2015 23:59:00 (i.e. the month of February 2015). Let *GBPCHF_RWDC0.1* represent the set of these seven rolling windows. Similarly, we construct seven sets of rolling windows (one for each of the remaining currency pairs). For example, let *EURCHF_RWDC0.1* be the set of the seven rolling windows corresponding to EUR/CHF and let *EURUSD_RWDC0.1* be the set of the seven rolling windows corresponding to EUR/USD and so on. These sets are compiled in the same two steps as *GBPCHF_RWDC0.1* with a threshold *STheta* = 0.10%.

Finally, we get the following eight sets of rolling windows: *EURCHF_RWDC0.1*, *GBPCHF_RWDC0.1*, *EURUSD_RWDC0.1*, *GBPAUD_RWDC0.1*, *GBPJPY_RWDC0.1*, *NZDJPY_RWDC0.1*, *AUDJPY_RWDC0.1*, and *EURNZD_RWDC0.1*.

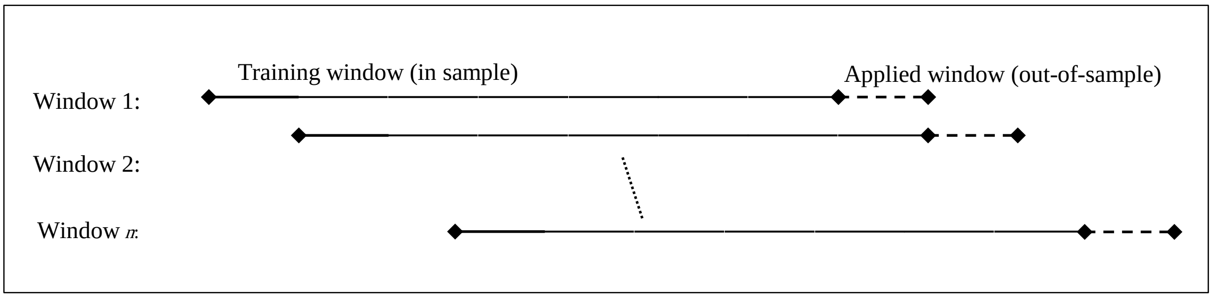

Fig. 6.3. Illustration of *n* rolling windows. The dashed lines represent the applied windows.

6.5 Evaluation of TSFDC: The experiments

In this section, we examine the performance of TSFDC. The objective is to evaluate the profitability and risk of both versions of TSFDC (i.e. TSFDC-down and TSFDC-up) using the rolling windows previously composed in Section 6.4.4. We provide the details of the experiments after describing the adopted money management approach.

6.5.1 Money management approach

We apply the money management approach to both TSFDC-down and TSFDC-up as follows. When TSFDC-down initiates a buy signal, we convert the entire capital from the counter currency to the base currency[k] (more details about counter and base currencies were provided in Section 2.2). When TSFDC-down generates a sell signal we convert the entire capital from the base currency to the counter currency. Likewise in the case of TSFDC-up. Although this sounds like a naïve approach to money management, our main objective is to prove that TSFDC is a successful trading strategy. Future works may address the development of a better money management approach.

When we operate any version of TSFDC, we make sure that no position is left open at the end of the trading period. Should we encounter an open position at the end of the trading period, then the last trades will not be considered when computing the evaluation metrics — instead, we roll back to the previous transaction. In other words, we do not count this last trade when measuring any of the evaluation metrics (previously introduced in Section 3.4). Thus, as a result of this very approach, if TSFDC opens a position it will not be able to open any other positions until the current position is closed.

In our experiments, we do not account the transaction costs. Eventually, counting transaction cost will reduce the returns of a trading strategy. However, some studies (e.g. [3] [36] [37] [100]) have concluded that counting transaction costs is not expected to have a substantial negative impact on the profitability of FX trading. Besides, some market makers (e.g. OANDA[l]) do not charge their customers for transaction costs for FX trading. Additionally, in Section 4.4 we reviewed four DC-based trading strategies ([15] [16] [17] [78]). All of these strategies did not consider the transaction costs in their experiments. Disregarding the transaction costs in our experiments serves to provide a fairly comparisons between our planned trading strategies, in this thesis, and these reviewed DC-based trading strategies.

We should also point out that we ignore the effect of 'slippage' in our trading simulations. In trading, the slippage refers to the difference between what a trader expects to pay for a trade and the actual price at which the trade is executed. Normally, the slippage happens because there might be a slight time delay between the trader initiating the trade and the time the broker receives the order. During this time delay, the price may have changed. It can either work in favour of, or against, the trader [101].

---

[k] For a given currency pairs 'X/Y', 'X' denotes the 'base currency' and 'Y' denotes the 'counter currency' (see Section 2.2 for more details about base and counter currencies). In this thesis, a 'sell' signal means that we are selling the base currency in exchange for the counter currency; whereas a 'buy' signal means that we are buying the base currency using the counter currency.

[l] OANDA: https://www.oanda.com/resources/news/pr/fxtrade03292001.

6.5.2 Experiment 6.1: Evaluation of the performance of TSFDC

The objective of this experiment is to evaluate the performance of TSFDC-down and TSFDC-up. For this purpose, we apply both versions to the eight currency pairs sampled minute-by-minute: EUR/CHF, GBP/CHF, EUR/USD, GBP/AUD, GBP/JPY, NZD/JPY, AUD/JPY, and EUR/NZD. We consider the eight sets of rolling windows: *EURCHF_RWDC0.1*, *GBPCHF_RWDC0.1*, *EURUSD_RWDC0.1…*etc. (previously composed in Section 6.4.4). For each of these eight sets, the training period of each rolling window (24 months) is utilized to train the forecasting model (developed in Chapter 5). Next, the forecasting model is employed to compute the value of *FBBTheta* (i.e. to forecast *BBTheta*) for each DC event, of threshold *STheta*, during the trading period (i.e. the associated applied window of 1 month). TSFDC uses *FBBTheta* to decide when to initiate a trade, as described in Section 6.3, during the trading period. The overall trading period of each set is seven months in length: from 1/1/2015 to 31/7/2015. For each of the eight sets, *BTheta* is fixed, arbitrarily, to 0.13%. We measure the evaluation metrics previously listed in Section 3.4 to evaluate the performance of TSFDC.

The evaluation metrics Jensen's Alpha and Beta serve to evaluate the profitability and risk of a given trading strategy, with reference to a benchmark (Section 3.4). In this thesis, we consider the buy and hold approach as our benchmark. Thus, we apply the buy and hold approach to each considered currency pair (buying at the opening price on a monthly basis; holding it over the course of the trading month, and selling at the closing price). For each currency pair, we compute the monthly returns resulting from applying the buy and hold to the specified trading periods (during the seven months: from January 2015 to July 2015). We then use these monthly returns to compute Jensen's Alpha and Beta of TSFDC. Note that although our initial datasets in this experiment (i.e. the eight currency pairs) are sampled as a time series (with a time interval of one minute), the TSFDC's trading rules (presented in Section 6.3) are based on variables (e.g. *FBBTheta*) which originate from the DC concept.

6.5.3 Experiment 6.2: Compare the return and risk of both versions of TSFDC

The objective of this experiment is to test whether there is a significant difference between the performances of TSFDC-down and TSFDC-up. To this end, we compare the return and risk of both versions of TSFDC. For simplicity, we consider the maximum drawdown (*MDD*) as a measure of risk similarly to [17] and [76]. We use the monthly rate of returns (*RR*) and maximum drawdown (*MDD*) resulted from applying both versions of TSFDC to the eight currency pairs from the previous experiment. In order to validate our test statistically, we chose to apply the non-parametric Wilcoxon test. More particularly, we apply the Wilcoxon signed rank test [102].

In this experiment, we apply the Wilcoxon test twice. Firstly, we apply the Wilcoxon test with the null hypothesis that the median difference between the two sets of monthly *RR* of TSFDC-down and TSFDC-up is zero. In this instance, we consider the monthly *RR* generated by applying TSFDC-down to the eight currency rates as the first set. This set consists of 56 observations (8 currency rates × 7 monthly *RR* for each currency rate). Similarly, the second set comprises the monthly *RR* generated by applying TSFDC-up to the eight currency rates (a total of other 56 observations). We report the details of these two sets in Appendix C.

Secondly, we seek to compare the risk of both versions of TSFDC. Taking the maximum drawdown as an indicator of risk (as in [16] [17]), we compose a first set by applying TSFDC-down to the eight currency rates. This set comprises 56 observations (8 currency rates × 7 monthly *MDD* for each currency rate). We compose a second set of monthly *MDD* data by applying TSFDC-up to the eight currency rates and apply the Wilcoxon signed rank test to each, with the null hypothesis that the median difference between the two sets of monthly *MDD* of TSFDC-down and TSFDC-up is zero (Appendix C comprises the details of these two sets).

6.6 Evaluation of TSFDC: Results and discussion

6.6.1 Experiment 6.1: Evaluation of the performance of TSFDC

The objective of this experiment is to evaluate the performance of TSFDC-down and TSFDC-up using eight currency pairs sampled minute-by-minute. To this end, we applied the two versions of TSFDC to the eight sets of rolling windows composed in Section 6.4.4. We followed the money management approach outlined in Section 6.5.1 and measured the evaluation metrics listed in Section 3.4. These evaluation metrics are:

- Rate of returns (RR): *RR* is interpreted as the gain or loss on an investment over a given evaluation period expressed as a percentage of the amount invested.
- Profit factor: It is calculated by dividing the sum of profits produced by all profitable trades by the sum of losses incurred by all losing trades. This metric measures the amount of profit per unit of risk.
- Max drawdown: It is the largest difference, in percentage, between the maximum amount (i.e. peak) and the minimum amount (i.e. through) of capital during a trading period. It measures the risk as the worst peak-to-trough decline in capital.
- Win ratio: This is the probability that a trade produces a positive return.
- Sharpe ratio: It measures the risk-adjusted return. It represents the average return earned in excess of the risk-free rate per unit of volatility.

- Sortino ratio: It denote the excess return over the risk-free rate divided by the downside semi-variance, and so it measures the return to ‘bad’ volatility.
- Jensen’s Alpha: It indicates whether a trading strategy is earning the proper return for its level of risk.
- Beta: It serves to measure the volatility, or systematic risk, of a security or a portfolio, in comparison to a benchmark.

In order to avoid tedious details, this section reports TSFDC’s general trading performance during the overall trading period for the eight currency pairs. The details of its monthly performance on these currency pairs are provided in Appendix D.

6.6.1.1 Experiment 6.1: The results

For each currency pair, we use the same values of *STheta* (0.10%) and *BTheta* (0.13%). These values are chosen arbitrarily. Bear in mind that, for each currency pair, we compose seven rolling windows. Each window comprises a trading period of one month. At the beginning of the first trading period, i.e. January 2015, both TSFDC-down and TSFDC-up start with a capital = 1,000,000[m]; this represents the initial, hypothetically, invested amount of money.

Table 6.3, shown below, reports the general performance of both versions of TSFDC during the overall trading period of seven months. The detailed monthly evaluation of applying TSFDC to these eight currency rates is provided in Appendix D. In Table 6.3, the column ‘Currency Pair’ denotes the considered currency pair. The column ‘Trading Strategy’ indicates which version of TSFDC is applied. The columns ‘RR’, ‘Profit Factor’, ‘Max Drawdown (%)’, and ‘Win Ratio’ refer to the chosen evaluation metrics. The last row in Table 6.3 is interpreted as follows: applying TSFDC-up to EUR/NZD generates a total return of 571.89% during the seven-month trading period. In this case, TSFDC-up executes 4,218 trades with an overall Win Ratio of 0.77. The maximum drawdown in capital is – 5.1%.

[m] For each currency pairs, in case of trading with TSFDC-down, we assume that we start with 1,000,000 monetary units of the counter currency. For example: in the case of EUR/CHF, we start with 1,000,000 CHF. Whereas in the case of NZD/JPY, we start with 1,000,000 JPY. However, in the case of TSFDC-up we assume that we start with 1,000,000 monetary units of the base currency.

Table 6.3: Trading performance of TSFDC-down and TSFDC-up models following the seven months out-of-sample period of the eight currency pairs.

Currency Pair	Trading Strategy	RR	Profit Factor	Total Number of Trades	Max Drawdown (%)	Win Ratio
EUR/CHF	**TSFDC-down**	84.59	1.93	2056	– 13.4	0.73
	TSFDC -up	63.03	1.83	2009	– 15.1	0.71
GBP/CHF	**TSFDC-down**	94.03	1.73	2489	– 12.1	0.72
	TSFDC -up	115.19	1.69	2531	– 10.8	0.70
EUR/USD	**TSFDC-down**	27.04	1.26	1431	– 5.0	0.65
	TSFDC -up	36.09	1.32	1453	– 5.8	0.67
GBP/AUD	**TSFDC-down**	92.63	1.86	3021	– 3.4	0.70
	TSFDC -up	63.03	1.54	2960	– 3.5	0.68
GBP/JPY	**TSFDC-down**	32.48	1.53	1585	– 4.8	0.69
	TSFDC -up	28.91	1.42	1601	– 5.7	0.69
NZD/JPY	**TSFDC-down**	183.13	2.20	3046	– 4.0	0.73
	TSFDC -up	190.73	2.08	3010	– 4.9	0.74
AUD/JPY	**TSFDC-down**	104.11	1.70	2885	– 5.0	0.71
	TSFDC -up	116.35	1.81	2860	– 5.2	0.72
EUR/NZD	**TSFDC-down**	489.13	2.98	3961	– 4.6	0.77
	TSFDC -up	571.89	2.86	4218	– 5.1	0.77

The results of monthly Rates of Return (*RR*) of applying TSFDC-down and TSFDC-up to these currencies pairs are shown in Tables 6.4 and 6.5 respectively. These *RR* will be utilized to compute the Sharpe and Sortino ratios, Jensen’s Alpha and Beta. The Sortino and Sharpe ratios of both versions of TSFDC are reported in Table 6.6. The minimum acceptable return (MAR) and the risk-free rate are set to 5% per annum. The computation of Jensen’s Alpha and Beta consists of comparing TSFDC to a particular benchmark.

Table 6.4: Monthly *RR* of applying TSFDC-down to the eight currency pairs shown in Table 6.3.

Trading period	**EUR/ CHF**	**GBP/ CHF**	**EUR/ USD**	**GBP/ AUD**	**GBP/ JPY**	**NZD/ JPY**	**AUD/ JPY**	**EUR/ NZD**
Jan 2015	4.47	13.59	1.12	19.70	7.72	19.14	15.36	24.12
Feb 2015	14.40	19.02	7.54	10.51	6.40	26.90	16.47	50.04
Mar 2015	17.59	14.96	– 0.36	10.14	4.04	19.95	10.51	49.76
Apr 2015	7.58	6.71	4.20	13.52	7.05	30.41	16.69	59.39
May 2015	13.37	9.85	5.73	15.97	8.38	24.27	25.51	79.92
Jun 2015	12.41	15.17	7.85	11.52	0.99	17.20	10.48	104.91
Jul 2015	14.77	14.73	0.96	11.27	– 2.10	45.26	9.09	120.99
Sum	84.59	94.03	27.04	92.63	32.48	183.13	104.11	489.13

Table 6.5: Monthly *RR* of applying TSFDC-up to the eight currency pairs shown in Table 6.3.

Trading period	**EUR/ CHF**	**GBP/ CHF**	**EUR/ USD**	**GBP/ AUD**	**GBP/ JPY**	**NZD/ JPY**	**AUD/ JPY**	**EUR/ NZD**
Jan 2015	4.26	31.54	6.81	13.34	11.39	26.96	21.48	26.27
Feb 2015	9.75	16.30	9.27	10.06	3.64	18.06	14.88	68.74
Mar 2015	16.87	21.67	1.69	9.09	6.00	24.06	17.30	64.56
Apr 2015	5.71	12.34	1.66	9.23	3.07	22.98	12.25	78.72
May 2015	7.61	7.59	9.67	9.51	4.11	24.92	21.15	82.81
Jun 2015	10.15	14.13	6.13	5.97	4.16	32.66	17.32	101.88
Jul 2015	8.68	11.62	0.86	5.83	– 3.46	41.09	11.97	148.91
Sum	63.03	115.19	36.09	63.03	32.37	190.73	116.35	571.89

Table 6.6: The Sortino and Sharpe ratio of the two versions of TSFDC. The math symbol ∞ denotes positive infinity.

Currency pair	**TSFDC-down**		**TSFDC-up**	
	Sortino ratio	**Sharpe ratio**	**Sortino ratio**	**Sharpe ratio**
EUR/CHF	∞	2.6	∞	1.8
GBP/CHF	∞	3.2	∞	2.0
EUR/USD	177.3	1.0	∞	1.7
GBP/AUD	∞	3.7	∞	3.4
GBP/JPY	37.2	1.1	19.9	0.9
NZD/JPY	∞	2.7	∞	3.6
AUD/JPY	∞	2.6	∞	4.2
EUR/NZD	∞	2.0	∞	2.2

In this thesis, we adopt the buy and hold approach as a benchmark. The buy and hold (B&H) approach has been used as benchmark for trading strategies' performance in many studies (e.g. [4] [43]). For each currency pair, we apply the B&H approach on a monthly basis over the considered trading period from 1/1/2015 to 31/7/2015 (seven months). Table 6.7, shown below, summarizes the monthly *RR* of applying the B&H approach to the eight currency pairs. The column 'Sum', in Table 6.7, shows the sum of all *RR* generated by applying B&H to the seven months for each considered currency pair. We use the monthly *RR* of the buy and hold method to calculate Jensen's Alpha and Beta of TSFDC. The values of Jensen's Alpha and Beta corresponding to these comparisons are reported in Table 6.8.

Table 6.7: Summary of the monthly RR (%) obtained by applying the buy and hold (B&H) approach to each of the eight considered currency pairs. The trading period is from 1/1/2015 to 31/7/2015.

	Jan	**Feb**	**Mar**	**Apr**	**May**	**Jun**	**Jul**	**Sum**
EUR/CHF	– 12.88	1.75	– 1.95	0.10	– 1.41	0.99	1.77	– 11.63
GBP/CHF	– 9.68	5.17	– 2.01	– 0.60	0.57	1.92	2.69	– 1.94
EUR/USD	– 6.48	– 1.07	– 3.66	3.96	– 2.31	1.72	– 1.38	– 9.22
GBP/AUD	2.07	1.57	– 1.42	– 0.12	2.81	1.59	5.11	11.61
GBP/JPY	5.43	4.59	– 3.73	3.34	3.32	1.34	0.81	4.24
NZD/JPY	– 9.04	6.60	– 1.14	1.60	– 2.93	– 5.41	– 1.84	– 12.16
AUD/JPY	– 7.28	3.02	– 2.26	3.49	0.49	0.27	4.48	2.21
EUR/NZD	0.54	– 5.08	– 2.54	2.38	4.43	6.12	1.79	7.64

Table 6.8: The values of Jensen's Alpha and Beta of TSFDC with reference to the buy and hold as benchmark. The values are rounded to one decimal digit.

Currency pair	**TSFDC-down**		**TSFDC-up**	
	Jensen's Alpha	**Beta**	**Jensen's Alpha**	**Beta**
EUR/CHF	11.9	0.6	7.5	0.3
GBP/CHF	11.3	0.2	9.4	– 1.3
EUR/USD	3.2	0.5	1.1	– 0.3
GBP/AUD	9.3	0.4	5.6	– 0.4
GBP/JPY	1.8	0.0	– 0.3	– 0.7
NZD/JPY	26.4	0.7	20.9	– 0.7
AUD/JPY	11.8	0.0	11.7	– 0.8
EUR/NZD	76.1	5.0	76.1	5.0

6.6.1.2 Experiment 6.1: Results' Discussion

We begin with an examination of the results obtained from the B&H (shown in Table 6.7). For each currency pair (i.e. each row), we note that the B&H approach can generate profit in some months, but it incurs losses in others. This observation indicates that none of the selected currency pairs exhibit a monotonic trend during the trading period. Besides, the numbers shown in the column ‘Sum’ (Table 6.7) show that, overall, the B&H method generates profit in four cases: GBP/AUD, GBP/JPY, AUD/JPY, and EUR/NZD (with total rate of returns, *RR*, of up to 11.61% in the case of GBP/AUD). The same column also shows that the buy and hold method incurs losses in the other four cases (with total *RR* equal to –12.16 in the case of NZD/JPY). These observations support our claim regarding the variation of the trends of the selected currency rates in Section 6.4.1.

We then examine the profitability of both versions of TSFDC. The monthly rate of returns (*RR*) reported in Tables 6.5 and 6.6 suggest that both versions of TSFDC are mostly profitable (except in very few cases; e.g. trading with TSFDC-down on EUR/USD in March 2015 when it incurred losses of – 0.36%, Table 6.4). The results in column (RR), shown in Table 6.3, suggests that TSFDC can be highly profitable (with *RR* of up to 571.89 %, as in the case of applying TSFDC-up to EUR/NZD, the last row in Table 6.3). The overall Win Ratio of TSFDC (i.e. the probability of having a winning trade) ranges between 0.77 (as in the case of applying TSFDC-down to EUR/NZD) and 0.65 (in the case of applying TSFDC-down to EUR/USD). We consider this range to be reasonably acceptable.

However, it is important to note that the profitability of TSFDC may vary largely from one currency pair to another – as demonstrated in Table 6.3 when TSFDC-up is applied to GBP/JPY and EUR/NZD. One can easily observe an important difference between the produced total *RR* (from 28.91% for GBP/JPY, compared to 571.89% for EUR/NZD). This indicates that, whilst TSFDC may generate profits in most cases, its performance may vary substantially from one currency rate to another. It follows then that a trader may want to consider other currency pairs as TSFDC may, possibly, perform better on these currencies than on those reported in this chapter.

When we inspect the risk of TSFDC, in Table 6.3, we notice that, in most cases, the maximum drawdown (*MDD*) is no worse than – 6.0% (except in two cases: EUR/CHF and GBP/CHF) — values we consider to be relatively low. Moreover, the values of the Sortino ratio, reported in Table 6.6, are, in most cases, a positive infinity (∞). This reflects the fact that the downside risk (see equation (3.5) in Section 3.4) of TSFDC is null in most of these experiments. Also, in most cases, the values of the figures in the column ‘Beta’ (indicated in Table 6.8) range between – 1.0 and 1.0. This range point out that TSFDC is, generally, less volatile than the buy and hold approach. Keep in mind that the volatility of returns is usually used as an indicator of the risk of a trading strategy [62].

Lastly, we examine the risk-adjusted performance of TSFDC. For this purpose, we consider the values of the Sharpe ratio and Jensen's Alpha shown in Tables 6.7 and 6.9 respectively. The Sharpe ratio is consistently positive (Table 6.6). A positive Sharpe ratio indicates that the TSFDC has surpassed the 5% annual risk-free rate, demonstrating that TSFDC generates worthy excess returns for each additional unit of risk it takes. The Jensen's Alpha results are, generally, consistent with the Sharpe ratio scores (though with one exception in the case of applying TSFDC-up to GBP/JPY – Table 6.8). We conclude that TSFDC earns more than enough return to compensate for the risk it took over the trading period.

We conclude from the previous analysis that TSFDC-down and TSFDC-up provide more *RR* and, in most cases, less risk than the buy and hold method. Additionally, both versions of TSFDC can be highly profitable, with *RR* of more than 400% (Table 6.3). We also argued that TSFDC can consistently deliver a positive Sharpe ratio. Finally, the established variety of the selected currency pairs in the initial dataset (Section 6.4.1) suggest that TSFDC can be profitably applied to a wide range of currency rates.

6.6.2 Experiment 6.2: Compare the return and risk of both versions of TSFDC

The objective of this experiment is to compare the return and risk of both versions of TSFDC, TSFDC-up and TSFDC-down. We consider the monthly rate of returns (*RR*) and monthly maximum drawdown (*MDD*) resulted from applying both versions of TSFDC to the eight currency pairs in the previous experiment. Firstly, we apply the Wilcoxon test with the null hypothesis that the median difference between the two sets of monthly *RR* of TSFDC-down and TSFDC-up is zero. Appendix C comprises the details of these two sets. In this case, the Wilcoxon test returns a *p*-value of 0.79. Since the *p*-value is greater than the common cut-off value 0.05, the Wilcoxon test cannot reject the null hypothesis that the median difference between the monthly *RR* for TSFDC-down and TSFDC-up is zero.

Secondly, we apply the Wilcoxon test with the null hypothesis being that the median difference between the two sets of monthly *MDD* of TSFDC-down and TSFDC-up is zero. Appendix C compiles the details of these two sets. In this case, the Wilcoxon test returned a *p*-value of 0.50. This *p*-value is greater than 0.05. Therefore, the Wilcoxon test cannot reject the null hypothesis that there is no difference between the two sets of monthly *MDD* for TSFDC-down and TSFDC-up. To conclude, the Wilcoxon tests do not suggest that the monthly *RR* and the monthly *MDD* of TSFDC-down and TSFDC-up are different.

6.7 Comparing TSFDC to other DC-based strategies

In Section 4.4, we reviewed some existing trading strategies that are based on the DC framework. In this section, we compare TSFDC with two other DC-based trading strategies: (a) the one presented by Gypteau et al., [78] and (b) the DC+GA (Kampouridis and Otero [17]). The details of these two strategies can be found in Section 4.4. In this section, we aim to compare these strategies with TSFDC.

6.7.1 The DC-based trading strategy by Gypteau et al.

In this section, we highlight the differences between TSFDC and the DC-based trading strategy presented by Gypteau et al., [78] which was revised in details in Section 4.4.2. We start with a brief recap about the functionality of this DC-based trading strategy; then we compare it to TSFDC.

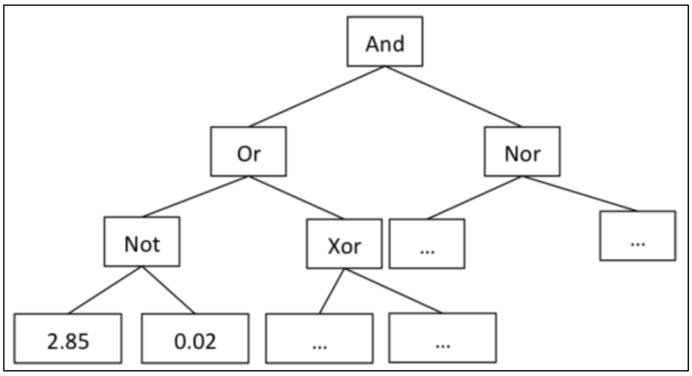

Fig. 6.4. A sample individual GP tree: internal nodes are represented by Boolean functions, while terminal nodes correspond to different DC thresholds. Given a price, terminal nodes output a Boolean value according to the DC or OS events detected. For example, if we detect a downtrend (uptrend) DC event of a DC summary of threshold 2.85%, then the left- most terminal node will be evaluated as 'False' ('True'). Source Gypteau et al., [78].

The proposed approach follows the standard tree-based Genetic Programming (GP) configuration. It runs multiple DC summaries, using different DC thresholds, concurrently. For example, Fig. 6.4, shown above, illustrates a sample individual GP tree. Each GP individual trees comprises internal and terminal nodes. The internal nodes are Boolean functions and the terminal nodes are DC thresholds. In Fig. 6.4 each threshold, shown in terminal nodes, is replaced with a 'True' or 'False' depending whether an uptrend or downtrend DC event of the stated threshold is detected. For example, in Fig. 6.4, if we detect an upward (downward) DC event of threshold 2.85% the left-most terminal nodes would be set as 'True' ('False').

These 'True' and 'False' values at the terminal nodes are, then, combined together using the Boolean functions (e.g. AND, Nor, Xor), presented in the internal nodes, to form a GP-tree (see Fig. 6.4). As such a GP tree can be interpreted as a Boolean expression; the output of which can be only True or False. This output is translated into trading rules with 'True' triggering a buy signal and 'False' triggering a sell signal. Consequently, each GP tree represents a trading strategy. The profitability of the GP tree (i.e. trading strategy) is measured based on the returns resulted from the triggered buy and sell orders over an in-sample dataset. The evolution of the objective function of the Genetic Programming (GP) aims to find the best GP-tree which yields the highest returns.

The authors applied their trading strategy to four markets: two stocks from the UK FTSE100 market (Barclays Bank and Marks & Spencer), and two international indices (NASDAQ, and NYSE). For each market, they used a training period of 1000 days to train their GP model. Then, they follow a testing period of 500 days for evaluation. Unfortunately, the authors did not report the dates of the training or testing periods! (For more details about this trading model, see Section 4.4.2).

We provide the following two comments on the study of Gypteau et al. [78]:

1. The authors stated that: "*... the proposed approach aims to find an optimal trading strategy to forecast the future price moves of a financial market*" [78]. However, having investigated the study [78], we could not find a formal representation of any forecasting problem. The authors, in [78], did not identify any dependent or independent variables. Besides, they did not report any forecasting measurements (e.g. mean squared error, accuracy). Therefore, we could not conclude that the proposed strategy, in [78], does clearly employ a forecasting model.
2. With respect to the evaluation of the proposed DC-based strategy, the authors reported only the returns of the proposed trading strategy [78]. The reported returns are less than 10% over a trading period of 500 days for each considered market. Furthermore, they did not report any a) comparison to a benchmark (e.g. buy and hold), b) measurement of risk (e.g. MDD), and c) evaluation of risk-adjusted metrics (e.g. Sharpe ratio). Therefore we believe that the reported returns are not sufficiently convincing regarding the feasibility of the proposed strategy.

In contrast, we consider TSFDC to be founded on a well-formulated forecasting model. This forecasting model, established in Chapter 5, aims to forecast the change of trend's direction under the DC framework. It has a clear objective, dependent and independent variables (see Section 5.4). By contrast, the study of Gypteau et al. [78] does not define any dependent or independent variables. Another difference is that, in contrast to the study of Gypteau et al. [78], we provided a thorough evaluation of the risk and profitability of TSFDC (Section 6.6).

We should also highlight that TSFDC and the trading strategy proposed by Gypteau et al., have different trading approaches: TSFDC forecasts the change of trend's direction to decide when to trigger a new trade. While Gypteau et al. employs a GP approach to develop an expression of Boolean functions, and several DC thresholds, which is then converted to trading rules.

Finally, we want to highlight that applying the strategy of Gypteau et al. to stock markets produced a maximum profit of less than 10% (over a trading period of more than 1 year). In Section 6.6.1, we examine the profitability of TSFDC in the FX market and concluded that it can produce rates of return of more than 400% in less than 7 months. Moreover, we want to recap that the authors in [78] evaluate the proposed trading strategy in stock market where prices are sampled on daily basis. In contrast, TSFDC was evaluated in the FX market using minute-by-minute mid-prices. Despite that the results of *RR* indicate that TSFDC is much more profitable than the strategy of Gypteau et al.; it would be better to evaluate both strategies using same dataset in order to make fair comparison regarding whether TSFDC is more profitable.

6.7.2 The DC-based trading strategy: 'DC+GA'

In this section, we compare TSFDC with the trading strategy named 'DC+GA' (Kampouridis and Otero [17]). The authors in [17] stated that their objective was "*to offer a more complete analysis on the directional changes paradigm from a financial forecasting perspective*." The details of this strategy was reviewed in Section 4.4.3. Here we briefly recap the mechanism of this strategy, then we compare it with TSFDC.

DC+GA consists of running N_{theta} DC summaries concurrently using N_{theta} thresholds. These N_{theta} thresholds are to be chosen by the trader. DC+GA uses some parameters: b_1, b_2, and Q (see Section 4.4.3 for more details about these parameters). The first two parameters (b_1 and b_2) help DC+GA to decide when to initiate a trade during an OS event. The third parameter 'Q' denote the order size. For a given market's price, each DC threshold generates a buy or sell recommendation upon the type of the detected DC event (either downward or upward). In addition, each DC threshold is assigned a 'weight'. For a given market's price, the N_{theta} DC-thresholds may produce N_{theta} recommendations. These thresholds are, then, clustered in two groups based on the proposed recommendations: the first group covers the thresholds those recommend a buy action, the second group covers those recommending a sell action. To make a buy or sell decision, DC+GA sum the weights of the thresholds belongs to each group: if the sum of the weights for all thresholds recommending a buy (sell) action is greater than the sum of the weights for all thresholds recommending a sell (buy) action, then the strategy's action will be to buy (sell).

The evolution of the GA consists of finding the best set of weights of the N_{theta} DC thresholds along with the trading parameters (b_1, b_2, and Q) that maximize the total profits during the training process. The best set of DC's thresholds, and their associated weight and trading parameters, will be used for trading during the out-of-sample trading period (see Section 4.4.3). The employed fitness function is designed so that it maximizes *RR* and minimizes the *MDD* at the same time.

A common feature between TSFDC and DC+GA is that they both analyse uptrends and downtrends separately. Though, we can identify the following differences between TSFDC and DC+GA:

- DC+GA initiates a trade when the time length of an OS event reaches certain value. Whereas, TSFDC initiates a trade when the magnitude of price's change reach certain threshold (either *STheta* or *BTheta*).
- TSFDC relies on the forecasting approach presented in Chapter 5 to decide whether to initiate a trade when a new DC event is detected. Whereas, DC+GA employs a GA module to anticipating the best *time* at which it should initiate a trade.
- TSFDC uses two DC thresholds (*STheta* and *BTheta*), whereas DC+GA takes into consideration N_{theta} DC summaries at the same time.

Kampouridis and Otero [17] reported the mean *RR* results of applying DC+GA to five currency pairs (Table 6, page 158, [17]). We note that DC+GA incurred overall losses in two out of the five considered currency pairs. Moreover, when examining the detailed monthly returns (Table 5, page 158, [17]) we note that, in most months, DC+GA reported losses. By contrast, when inspecting the monthly returns of TSFDC reported in Tables 6.5 and 6.6, we note that in the majority of cases TSFDC's monthly returns are positive. Furthermore, the overall returns of applying TSFDC to the eight currency pairs (over the trading period of seven months) are consistently positive (Table 6.3). Thus, we conclude that TSFDC is more profitable than DC+GA.

We then examine the risk-adjusted returns of DC+GA and TSFDC. The authors in [17] did not provide any risk-adjusted measurement for DC+GA. However, based on the reported monthly returns in Table 5 (page 158, [17]), we can compute the Sharpe ratio. If we consider a risk-free rate of 5% per annum, then we find that DC+GA will have a negative Sharpe ratio in four out of the five considered currency pairs (see Section 4.4.3 for details). Whereas, the results shown in Table 6.6 (Section 6.6.1) indicate that TSFDC produces a positive Sharpe ratio. Based on this analysis, we conclude that TSFDC outperforms 'DC+GA' in terms of profitability and risk-adjusted returns.

Finally, we compare the risk of TSFDC and DC+GA measured in term of *MDD*. The *MDD* of DC+GA reported in [17] is no worse than – 0.15% (Table 8, [17]) in all considered currency pairs. This is better than the *MDD* of TSFDC as reported in Table 6.3.

To conclude, by comparing the results of DC+ GA (reported in [17]) and the results of TSFDC (Section 6.6.1) we deduce that TSFDC outperforms DC+GA in terms of *RR* and risk-adjusted returns. However, the results of *MDD* suggest that DC+GA is less risky than TSFDC.

6.8 Summary and conclusion

In this chapter, our objective was to develop a successful trading strategy based on forecasting DC. Following our findings in Chapter 5 concerning forecasting the change of the direction of a DC trend, this chapter uses this forecasting model to develop a trading strategy named TSFDC. TSFDC is a contrarian trading strategy that relies on the forecasting model (summarized in Section 6.2) to decide when to generate a trade (Section 6.3). The trading rules of TSFDC was presented in Section 6.3.

The performance of TSFDC was examined using eight currency pairs. We utilized 1-minute trade records for these eight currency pairs covering the period between 1/1/2013 and 31/7/2015. We argued that these currency pairs exhibited various trends' patterns during the considered trading period of seven months (Section 6.4.1). We evaluated TSFDC using a monthly-basis rolling window approach. Each rolling window comprised 1) a training period (24 months in length), which we use to train the forecasting model developed in Chapter 5, and 2) a trading period (1 month in length) to which we applied the trading rules of TSFDC (Section 6.4.4). We utilized a set of evaluation metrics to assess the performance of TSFDC.

In our experiment, as a benchmark model, we implemented the buy and hold strategy, buying at the opening price on a monthly basis, holding it over the course of the trading month and selling at the closing price. The inclusion of this zero-intelligence benchmark model was to assess the usefulness and potential outperformance of our trading strategies in general. However, it should be noted that, like many other DC-based trading strategies (e.g. [15] [17] [78]), the transaction costs were not considered in our experiments.

The experimental results (reported in Section 6.6.1) suggest that TSFDC is successful. By examining the returns reported in Table 6.3 (Section 6.6.1), we concluded that TSFDC can be highly profitable (with a *RR* of more than 500%, as per EUR/NZD) and yet retain a reasonable level of risk (with *MDD* equal to – 5.1%). When examining the values of Jensen's Alpha (shown in Table 6.8, Section 6.6.1), we concluded that TSFDC generated promising rates of return compared to the level of risk it took in relation to the buy and hold method. From the Beta results detailed in Table 6.8 (Section 6.6.1), we see that in the majority of cases TSFDC was less, or equally, volatile than the buy and hold method. This indicates that, generally, TSFDC is less risky than the buy and hold approach. We also argued that TSFDC outperforms other DC-based trading strategies in Section 6.7.

To conclude, in this chapter we developed a DC-based trading strategy, named TSFDC, which, we believe it to be the first DC-based trading strategy that is based on a well-formulated forecasting model. As our main contribution, we argued that TSFDC is more profitable than another DC-based trading strategy (Section 6.7). The experimental results indicates that TSFDC can be highly profitable (Section 6.6.1). We examined the effectiveness of TSFDC over eight different currency rates that have different patterns. Therefore, we believe that TSFDC could be successful in a broad range of currency pairs. Despite what would be considered as experimental weaknesses (e.g. ignoring the transaction costs), we consider these results as a proof of the usefulness of the DC framework as a basis of trading strategies.

7 Backlash Agent: A Trading Strategy Based on Directional Changes

In this chapter, we introduce a trading strategy named Backlash Agent, or BA for short. BA is designed so that it does not employ any forecasting model. We evaluate the performance of BA the same way we evaluated TSFDC in Chapter 6. The results indicate that BA can generate profits of more than 300% within seven months.

7.1 Introduction

As stated in Section 1.2, the objective of this thesis is to explore, and consequently to provide a proof of, the usefulness of the DC framework as the basis of profitable trading strategies. Surveying the literature in Chapter 3, we noticed that most trading strategies can be classified into two classes based on whether they rely on forecasting models or not. In keeping with the existing research, this thesis aims to establish two trading strategies those are based on the DC framework: the first relies on a forecasting model and the second does not employ any forecasting approach. This first strategy, named TSFDC, was introduced in Chapter 6 and relies on the forecasting model previously established in Chapter 5.

This chapter develops the second trading strategy, which is also based on the DC framework, but does not rely on any forecasting model. This strategy is called Backlash Agent, or BA for short. The chapter continues as follows: Section 7.2 is a brief recap of some essential DC notations. The trading rules of BA are provided in Section 7.3. The details of the experiments conducted to examine the performance of BA are described in Section 7.4. We report and discuss the experimental results in Section 7.5. Then, we compare the performance of BA with other DC-based strategies in Section 7.6. Finally, the major findings of this chapter are summarized in Section 7.7.

7.2 DC notations

This section is essentially a revision of the DC notations previously explained in Section 4.2. These notations are adopted from Tsang et al., [74] and are listed in Table 7.1. In the context of this chapter we recap that:

- $P_{DCC\downarrow*}$: This is the price required to confirm the observation of the succeeding downtrend DC event. It is employed if the current trend is an uptrend.
- $P_{DCC\uparrow*}$: This is the price required to confirm the observation of the succeeding uptrend DC event. It is employed if the market is currently in downtrend.

- OSV: The objective of Overshoot Value (OSV) is measuring the magnitude of an overshoot event. Instead of using the absolute value of the price change, we would like this measure to be relative to the threshold, *theta*. We should note that, based on the formula provided in the last row in Table 7.1, OSV will be negative in case of downtrend and positive otherwise.

Table 7.1: List of DC notations used in this thesis (source: Tsang et al. [74]). Appendix A provides the code of how to compute these variables.

Name / Description	**Notation**
Threshold	*theta*
Current price	P_c
Price at extreme point: price at which one trend ends and a new trend starts.	P_{EXT}
The highest price, during an uptrend's OS event, required to confirm that the market's direction has changed to downtrend (i.e. to confirm a downtrend's DC event).	$P_{DCC\downarrow*} = P_{EXT} \times (1 - theta)$
The least price, during a downtrend's OS event, required to confirm that the market's direction has changed to uptrend (i.e. to confirm an uptrend's DC event).	$P_{DCC\uparrow*} = P_{EXT} \times (1 + theta)$
$PDCC^*$ is the price of the theoretical directional change confirmation point of the current trend.	$PDCC^* = P_{DCC\downarrow*}$ if the current trend is downtrend; otherwise $PDCC^* = P_{DCC\uparrow*}$.
Overshoot value (OSV) is defined at price P_c during an OS event.	$OSV = ((P_c - PDCC^*)/PDCC^*)/theta$

7.3 Backlash Agent

In this section, we present the trading rules of BA. BA is a contrarian trading strategy. It generates buy and sell signals against the market's trend. We introduce two types of BA: Static BA (SBA) and Dynamic BA (DBA). For each of SBA and DBA we provide two versions: down and up. So that in total we introduce four versions of BA: two statics (SBA-down and SBA-up), and two dynamics (DBA-down and DBA-up). We provide the trading rules of SBA-down and SBA-up in Section 7.3.1 and Section 7.3.2 respectively. The two versions of dynamic BA (i.e. DBA-down and DBA-up) will be presented in Section 7.3.3.

7.3.1 Static BA-down (SBA-down)

In this section, we introduce a trading strategy named Static BA-down, or SBA-down for short. SBA-down is only applicable when the market is in a downtrend (hence its name). SBA-down consists of two rules:

Rule SBA-down.1: (generate buy signal)

If (the current event is OS on a downtrend) and (*OSV* ≤ *down_ind*) then generate buy signal.

Rule SBA-down.2: (generate sell signal)

If ($P_c \geq P_{DCC\uparrow*}$) then generate sell signal.

In *Rule SBA-down.1*: *OSV* is the variable previously defined in Table 7.1 above; and *down_ind* is a trading parameter. In *Rule SBA-down.2*, $P_{DCC\uparrow*}$denotes the minimum price required to confirm the observation of the succeeding uptrend DC event (see Table 7.1 above). In simple terms, SBA-down generates buy signal when the Overshoot Value (*OSV*) drops below a certain threshold, *down_ind*, during a downtrend's OS event. The value of *down_ind* is the choice of the trader. SBA-down generates sell signal when the DC confirmation point of the next upward DC event is confirmed.

The condition of *Rule SBA-down.2* denote the case under which we confirm the DCC point of the next uptrend DC event of threshold *theta*. Note that *Rule SBA-down.2* is applicable only if a buy signal has been triggered. *SBA-down.2* plays two roles at the same time: *take-profit* and *stop-loss*. When *SBA-down.2* triggers a sell signal, it may incur losses (hence, functioning as *stop-loss*) or generates profits (thus, working as *take-profit*).

Table 7.2, shown below, illustrates an example of a DC summary. We use Table 7.2 to provide an example of how the trading rules of SBA-down function by examining the downward DC event [$CC^{0.1}$], of threshold 0.10%, which starts at time 21:41:00:

a) Suppose that the trader has chosen *down_ind* = – 0.45.
b) At time 21:43:00 (shown in column 'Time'), we determine that the *OSV* = – 0.48006847 (shown in column '*OSV*'), which is less than *down_ind* (– 0.45). The *OSV* is computed as indicated in Table 7.1.
c) Based on a) and b), all conditions of *Rule SBA-down.1* are fulfilled. Therefore SBA-down generates a buy signal at time 21:43:00.
d) [$DD^{0.1}$] is the upward DC event which immediately follows the downward DC event [$CC^{0.1}$]. At time 22:01:00, we confirm the DCC point of [$DD^{0.1}$] — which is $D^{0.1}$. Based on *Rule SBA-down.2*, SBA-down will generate a sell signal at time 22:01:00.

7.3.2 Static BA-up (SBA-up)

In this section, we introduce the second version of SBA named SBA-up. SBA-up is the mirror of SBA-down. SBA-up generate sell signal while the market is in an uptrend and only if the value of *OSV* exceeds a certain threshold, named *up_ind*. SBA-up generates a buy signal when a new downward DC event is observed. SBA-up consists of two rules:

Rule SBA-up.1: (generate sell signal)

If (the current event is OS on an uptrend) and ($OSV \geq up\_ind$) then generate sell signal.

Rule SBA-up.2: (generate buy signal)

If ($P_c \leq P_{DCC\downarrow *}$) then generate buy signal.

Table 7.2: An example of a DC summary of GBP/CHF mid-prices sampled minute-by-minute on 1/1/2013 from 21:00:00 to 22:01:00 (UK time). Excessive and unnecessary observation were omitted. *theta* = 0.10%. We also compute the values of *PDCC** and *OSV* as indicated in Table 7.1.

Time	Mid-price	DC Event	*PDCC**	Point	*OSV*
21:00:00	1.48150	start DC event (UPTREND)		B	
21:01:00	1.48180				
21:02:00	1.48170				
21:03:00	1.48159				
21:04:00	1.48280				
21:05:00	1.48310	start OS event (UPTREND)	1.48298150	$B^{0.1}$	0.07990659
21:06:00	1.48365				0.45078108
21:07:00	1.48430				0.88908729
21:08:00	1.48390				0.61936039
21:09:00	1.48380				0.55192867
………………………					
21:41:00	1.48690	start DC event (DOWNTREND)		C	2.64231213
21:42:00	1.48480	start OS event (DOWNTREND)	1.48541310	$C^{0.1}$	– 0.41274713
21:43:00	1.48470				– 0.48006847
21:44:00	1.48520				– 0.14346177
21:45:00	1.48495				– 0.31176512
21:46:00	1.48412	start DC event (UPTREND)		D	– 0.87053224
………………………					
22:01:00	1.48570	start OS event (UPTREND)	1.48560412	$D^{0.1}$	0.0645394

Here, $P_{DCC\downarrow *}$ denotes the highest price required to confirm the observation of the next downtrend DC event. *up_ind* is a trading parameter. The condition of *Rule SBA-up.2* indicates the case under which we confirm the DCC point of the next downtrend DC event of threshold *theta*. Note that

Rule SBA-up.2 is applicable only if a sell signal has been triggered. *Rule SBA-up.2* plays two roles at the same time: *take-profit* and *stop-loss*. When *Rule SBA-up.2* triggers a buy signal, it may incur losses (hence, functioning as *stop-loss*) or generate profits (thus, working as *take-profit*).

We use Table 7.2 above to provide an example of how the trading rules of SBA-down function, by examining the upward DC event [$BB^{0.1}$], of threshold 0.10%, which starts at time 21:00:00.

a) Suppose that the trader sets *up_ind* = 0.80.
b) At time 21:07:00, we determine that *OSV* = 0.88908729. *OSV* is larger than *up_ind* (0.80).
c) Based on a) and b) all conditions of *Rule SBA-up.1* are fulfilled and therefore SBA-up generates a sell signal at time 21:07:00.
d) [$CC^{0.1}$] is the upward DC event which immediately follows the downward DC event [$BB^{0.1}$]. At time 21:42:00, we confirm the DCC point, $C^{0.1}$, of the next downtrend DC event, which is [$CC^{0.1}$]. Based on *Rule SBA-up.2*, SBA-up will generate a buy signal.

7.3.3 Dynamic Backlash Agent

When trading with static BA, we have no hint as to how SBA-down, or SBA-up, will perform if the value of *down_ind* or *up_ind* is chosen arbitrarily. Theoretically, the investor should use his/her expertise to choose the value of the parameters *down_ind* or *up_ind*. However, in some cases, the investor may not have sufficient experience to do so. Moreover, there is no guarantee, should SBA-down perform well for a given value of *down_ind* during a trading period, *x*, that it will behave similarly during another trading period, *y*, using the same value of *down_ind*. The same note holds true for SBA-up. These facts are the motivation behind the development of the two versions of dynamic BA, namely DBA-down and DBA-up respectively.

7.3.3.1 Dynamic BA-down (DBA-down)

DBA-down comprises two stages. In the first stage, DBA-down automatically determines the value of the parameter *down_ind*. For this purpose, DBA-down applies a procedure, named FIND_DOWN_IND, to a training (i.e. in-sample) dataset to determine the value of *down_ind*. In the second stage, DBA-down uses the same two rules of SBA-down to trade over a trading, out-of-sample, dataset using the value of *down_ind* returned by FIND_DOWN_IND.

The objective of the procedure FIND_DOWN_IND is to find an appropriate value for the parameter *down_ind* that to be utilized to trade with SBA-down during the applied period. The output of the procedure FIND_DOWN_IND is one numerical variable, named *best_down_ind*. In order to determine *best_down_ind*, FIND_DOWN_IND applies the trading rules of SBA-down to the training dataset using 100 different values of *down_ind* (from – 0.01 to – 1.00, with a step size of – 0.01). For each value of *down_ind*, we compute the returns, either profits or losses, obtained by applying SBA-down to the training dataset. Thus, for a given training period we get 100 returns — one return for each distinct value of *down_ind*. We define *best_down_ind* as the value of *down_ind* under which SBA-down has generated the highest returns using the training dataset. In the second stage of DBA-down, we follow the trading rules (*SBA-down.1* and *SBA-down.2*) with the input parameter '*down_ind*' being assigned the value of *best_down_ind* to trade over the trading dataset.

7.3.3.2 Dynamic BA-up (DBA-up)

DBA-up is the dynamic version of SBA-up, as DBA-down is to SBA-down. DBA-up also has two stages, like DBA-down. The first stage consists of automatically finding an appropriate value of *up_ind*, using the training period. This is done by a procedure called FIND_UP_IND. FIND_UP_IND has the same role as FIND_DOWN_IND. FIND_UP_IND uses the training dataset to compute one numerical variable named *best_up_ind*. To determine *best_up_ind*, FIND_UP_IND applies the trading rules of SBA-up to the training dataset using 100 different values of *up_ind* (from 0.01 to 1.00, with a step size of 0.01). For each value of *up_ind*, we compute the returns, either profits or losses, obtained by applying SBA-up to the training period. Consequently, for a given training dataset we get 100 returns — one return for each value of *up_ind*. We define *best_up_ind* as the value of *up_ind* under which SBA-up has generated the highest returns during the training period. The second stage of DBA-up follows the trading rules (*SBA-up.1* and *SBA-up.2*) with the input parameter '*up_ind*' being assigned the value of *best_up_ind* to trade over the trading period.

7.4 Evaluation of the Backlash Agent: Methodology and experiments

To evaluate the performance of all versions BA, we consider the minute-by-minute mid-prices of the eight currency pairs (EUR/CHF, GBP/CHF, EUR/USD, GBP/AUD, GBP/JPY, NZD/JPY, AUD/JPY, and EUR/NZD) for 31 months: from 1/1/2013 to 31/7/2015. For each currency pair, we run the DC analysis with *theta* = 0.10%, and we compose a set of seven rolling windows. Each rolling window comprises a training window of 24 months in length and an applied (i.e. trading) period of 1 month in length. Basically, we use the same eight sets of rolling windows previously composed in Section 6.4.4; namely: *EURCHF_RWDC0.1*, *GBPCHF_RWDC0.1*, *EURUSD_RWDC0.1*, *GBPAUD_RWDC0.1*, *GBPJPY_RWDC0.1*, *NZDJPY_RWDC0.1*, *AUDJPY_RWDC0.1, and EURNZD_RWDC0.1*. See Section 6.4.4 for details of how to prepare these eight sets.

In this chapter, we provide five sets of experiments: 1) the first experiment is designed to estimate the best and the worst performance of SBA-down and SBA-up; 2) the second examine whether there are specific values of the parameters *down_ind* and *up_ind* for which SBA-down and SBA-up perform best; 3) the third evaluates the performance of DBA-down and DBA-up; 4) the fourth compare the profitability of SBA and DBA; 5) the fifth experiment aims to compare the performances of DBA-down and DBA-up.

We use the same money management approach described in Section 6.5.1 for each of these experiments. In summary: When any version of BA generates a buy or sell signal, it uses the entire capital to trade. When we apply any version of the Backlash Agent, we make sure that no position is left open at the end of the trading period. Should we encounter an open position at the end of the trading period, then the last transaction will not be considered when computing the results — instead, we rollback to the previous transaction. In other words, we do not count this last trade when measuring any of the considered evaluation metrics. Transaction costs and bid-ask spread are not counted.

7.4.1 Experiment 7.1: Evaluation of Static BA

The objective of this section is to evaluate the best and the worst performance of static BA (both versions: SBA-down and SBA-up).

7.4.1.1 Experiment 7.1.1: Estimating the best and worst RR of SBA-down

For simplicity, we consider the rate of returns (*RR*) as the primary performance indicator. *RR* is defined as the gain or loss on an investment expressed as a percentage of the amount invested (see Section 3.4).We will use the currency pair EUR/CHF to describe our approach to estimating the maximum and minimum *RR* that could be produced by applying SBA-down to EUR/CHF. More particularly, we consider the set of rolling window named *EURCHF_RWDC0.1*. The same method will apply to each of the remaining seven sets of rolling windows.

As stated in Section 7.2, static BA is not applicable unless the investor knows what values to assign to the parameters. Keep in mind that *EURCHF_RWDC0.1* includes seven applied windows. To provide a reasonable evaluation, we apply SBA-down to each applied window in *EURCHF_RWDC0.1* using 100 different values of *down_ind* (from – 0.01 to – 1.00, with a step size of – 0.01). Consequently, for each applied window we will have 100 *RR* (each *RR* corresponding to one distinct value of *down_ind*). For each applied window, we consider the maximum and the minimum generated *RR*. So that, in total we get seven maximum *RR* and seven minimum *RR*. To estimate the overall maximum *RR* of trading with SBA-down over *EURCHF_RWDC0.1*, we sum the seven maximum *RR* of these seven applied windows. This is complemented by other measures, mainly the profit factor, *MDD* and win ratio. Similarly, we apply SBA-down to the applied windows of each of the remaining seven sets of rolling windows (previously composed in Section 6.4.4) and we measure the maximum and the minimum produced *RR* of applying SBA-down to each set. In this experiment, as well as in the following experiments, we apply the money management approach described in Section 6.5.1.

7.4.1.2 *Experiment 7.1.2: Estimating the best and worst RR of SBA-up*

This experiment aims to evaluate the best and the worst performance of SBA-up. In line with the previous experiment, we apply SBA-up to each applied window in *EURCHF_RWDC0.1* using 100 different values of *up_ind* (from 0.01 to 1.00, with a step size of 0.01). Consequently, for each applied window we will have 100 *RR* (each *RR* corresponds to a distinct value of *up_ind*). *EURCHF_RWDC0.1* has seven applied periods. For each applied window we consider the maximum and the minimum generated *RR*. So that, in total we get seven maximum *RR* and seven minimum *RR*. To compute the maximum *RR* of trading with SBA-up over *EURCHF_RWDC0.1* we sum the seven maximum *RR*. We also measure additional metrics: the profit factor, *MDD* and win ratio. Similarly, we apply SBA-up to the applied windows of each of the remaining seven sets of rolling windows (previously composed in Section 6.4.4) and we measure the maximum and the minimum produced *RR* of applying SBA-up to each set.

7.4.2 Experiment 7.2: Is there one optimal value for the parameters down_ind and up_ind?

This experiment investigates whether there are specific values of the parameters, *down_ind* and *up_ind*, under which SBA-down and SBA-up will consistently produce maximum *RR*. For this purpose, we apply SBA-down and SBA-up to the eight currency pairs: EUR/CHF, GBP/CHF, EUR/USD, GBP/AUD, NZD/JPY, AUD/JPY, GBP/JPY, and EUR/NZD. In this experiment, we consider the period from 01/08/2014 to 31/07/2015 (12 months) as the trading period.

For each currency pair, for each month, we simulate 100 trades with SBA-down. For each trade, we use a different value of the *down_ind* parameter (from – 0.01 to –1.00, with a step size of – 0.01). Consequently, for each month we will have 100 returns (each return corresponds to a distinct value of *down_ind*). For each currency pair, and for each trading month, we compute the maximum *RR* generated by SBA-down. We select and report the values of the *down_ind* parameter that correspond to these maximum *RR*. In total, for each currency pair we obtain 12 values of *down_ind* that represent the best performance of SBA-down during the12 months (one value for each trading month).

We perform the same 100 trade simulations, in the same trading period (12 months) and on the same eight currency pairs, using SBA-up — each time using a different value of the *up_ind* parameter. For each currency pair, we get another 12 values of *up_ind* corresponding to the highest possible *RR* generated by SBA-up during the12 months. We analyse these values of *down_ind*, or *up_ind*, to find out whether there exists a particular value for which SBA-down, or SBA-up, will deliver the best possible performance consistently.

7.4.3 Experiment 7.3: Evaluating the performances of DBA-down and DBA-up

If choosing the value of the parameters *down_ind* or *up_ind* arbitrarily, a trader cannot have any precise perception of how good, or otherwise, would be the performance of the static BA. With this point in mind, we developed the dynamic version, DBA, as explained in Section 7.2.3. In this experiment, we aim to evaluate the performance of both versions of DBA. To this end, we apply DBA-down and DBA-up to the eight sets of rolling windows detailed in Section 6.4.4; namely, *EURCHF_RWDC0.1, GBPCHF_RWDC0.1, EURUSD_RWDC0.1, GBPAUD_RWDC0.1, GBPJPY_RWDC0.1, NZDJPY_RWDC0.1, AUDJPY_RWDC0.1, and EURNZD_RWDC0.1*.

Each window comprises: 1) a training period (of 24 months in length), and 2) a trading window (of 1 month in length). For each rolling window, the training period is utilized to find the values of the *down_ind* or *up_ind* parameters, based on the procedures FIND_DOWN_IND or FIND_UP_IND described in Section 7.3.3. Then, we use these values in the trading period associated with the specified rolling window. The performance of DBA is evaluated by measuring the metrics reported in Section 3.3.

7.4.4 Experiment 7.4: Comparing the RR of DBA and SBA

The objective of this experiment is to figure out what is the probability that DBA produces higher *RR* than SBA provided that when trading with SBA the parameters *down_ind* and *up_ind* are assigned random values. This probability will help us to evaluate the efficiency of the proposed procedures: FIND_DOWN_IND and FIND_UP_IND (Section 7.3.3) that are designed to find appropriate values for the parameters *down_ind* and *up_ind*. Note that when trading with the static versions it is the trader who must choose the values of the parameters *down_ind* and *up_ind*.

Choosing these randomly offers a way of assessing the relative performance of SBA-down, and SBA-up, against DBA-down, and DBA-up.

Consider that a trader assigns a random value to the parameter *down_ind*, or *up_ind*, when trading with SBA-down, or SBA-up. In such case, the question is: What is the probability that the dynamic BA (DBA-down or DBA-up) will produce higher returns than the static BA (SBA-down or SBA-up)? Let γ denote this probability. To compute γ, we estimate the performance of the static version using a set of randomly chosen values for input parameters *down_ind* and *up_ind*. The following provides an example of how to estimate γ based on the EUR/CHF dataset.

We simulate trading with SBA-down on *EURCHF_RWDC0.1* 10,000 times[n]. Each time, we trade with SBA-down on each applied window in *EURCHF_RWDC0.1*. Every time, and for each applied window, we assign a new random value to the parameter *down_ind*. In other words, each time that we trade with SBA-down we use seven random values of *down_ind*, each random value being ranged between – 0.01 and – 1.00 and used for one applied window. With every trading simulation, we measure the *RR* generated by SBA-down. Hence, we obtain 10,000 *RR*. Each return corresponds to one trade with SBA-down on the seven rolling windows of *EURCHF_RWDC0.1*. γ can be calculated as the fraction of how many of these 10,000 *RR* are less than the *RR* generated by the dynamic version, DBA-down, in Experiment 7.3 (Section 7.4.3). Similarly, we apply SBA-up to the applied windows of *EURCHF_RWDC0.1* 10,000 times with randomly picked values for parameter *up_ind*. Each time and for each applied window, we assign a new random value to the parameter *up_ind*. We obtain another 10,000 *RR*. Each return corresponds to one trade with SBA-up on the seven rolling windows of *EURCHF_RWDC0.1*. Again, γ is computed as the fraction of how many of these 10,000 *RR* are less than the *RR* generated by DBA-up in Experiment 7.3 (Section 7.4.3).

The entire procedure is repeated for each of the remaining seven sets of rolling windows: *GBPCHF_RWDC0.1*, *EURUSD_RWDC0.1*, *GBPAUD_RWDC0.1*, *GBPJPY_RWDC0.1*, *NZDJPY_RWDC0.1*, *AUDJPY_RWDC0.1*, and *EURNZD_RWDC0.1*. For each of these sets, we apply SBA-down and SBA-up with randomly chosen parameters, *down_ind* and *up_ind*, to each of the seven applied periods 10,000 times. Hence, we obtain 10,000 *RR* resulting from trading with SBA-down and another 10,000 *RR* resulting from trading with SBA-up. For each set of rolling

[n] In a preliminary experiment we consider various number of trading simulation to determine γ. We found that for more than 10,000 trading simulations (e.g. 13,000; 15,000) the value of γ changes only marginally (less than 0.5%).

windows, we evaluate γ as the percentage of how many of these 10,000 *RR* are less than the *RR* generated by DBA-down and DBA-up in Experiment 7.3.

7.4.5 Experiment 7.5: Comparing the returns and risk of both versions of DBA

The objective of this experiment is to test whether there is difference between the performances of DBA-down and DBA-up. To this end, we compare the returns and risk of both versions of DBA. In this experiment, we consider the monthly rate of returns (*RR*) and monthly maximum drawdown (*MDD*) resulted from applying both versions of DBA to the eight currency pairs from the Experiment 7.3. To validate our test statistically, we will apply the Wilcoxon signed rank test [102].

Initially, we compare the *RR* of DBA-down and DBA-up by composing two sets of *RR* based on the results of Experiment 7.3. The first set consists of the 56 monthly *RR* resulted from trading with DBA-down over the eight rolling windows (8 currency rates × 7 monthly *RR* for each currency rate). The second set consists of the 56 monthly *RR* obtained by applying DBA-up to the eight considered currency pairs. These two sets are presented in Appendix E. Then, we apply the non-parametric Wilcoxon test with the null hypothesis being that the median difference of the two sets of monthly *RR* is zero.

Secondly, we compare the risks of DBA-down and DBA-up. To this end, we compare the monthly *MDD* resulting from applying DBA-down and DBA-up to the eight currency rates. We compose two sets of *MDD* data. The first set contains the 56 *MDD* (8 currency rates × 7 monthly *MDD* for each currency rate) corresponding to trading with DBA-down. The second set contains the monthly *MDD* resulting from applying DBA-up to the eight currency rates (see Appendix E). We apply the Wilcoxon signed rank test to these sets with the null hypothesis being that the median difference of the two sets of monthly *MDD* is zero.

7.5 Evaluation of Backlash Agent: Results and discussion

7.5.1 Experiment 7.1: Evaluation of Static BA

7.5.1.1 Experiments 7.1.1: Evaluating the performance of SBA-down

The objective of this experiment is to estimate the best and the worst possible performance of SBA-down. For simplicity, we consider the maximum and the minimum produced *RR* as the primary indicators of the best and the worst performance respectively. We consider eight currency pairs. We compose eight sets of rolling windows (one set for each currency pair). Each set is composed of seven rolling windows (see Section 6.4.4). We apply SBA-down to the applied windows of each of set of rolling windows using 100 different values of *down_ind*. We adopt the money management approach described in Section 6.5.1. In this experiment, we are not concerned with the detailed monthly evaluation. Instead, we focus on the general performance of SBA-down during the overall seven months (i.e. the entire trading period) of each set of rolling windows. We also measure the overall profit factor, *MDD*, and win ratio.

Table 7.3 and Table 7.4 display, respectively, the best and the worst, estimated, performances of SBA-down when applied to the composed sets of rolling windows (see Section 7.3.1). These tables include the following metrics: rate of returns (*RR*), profit factor, maximum drawdown, and win ratio (see Section 3.4 for more details about these metrics). For each currency pair, at the beginning of the first applied window, i.e. January 2015, SBA-down starts with capital = 1,000,000[o]; this represents the initial, hypothetically, invested amount of money. From Table 7.3, let us consider the case of EUR/CHF. The reported results have the following interpretation: the total rate of returns (*RR*) is 80.41% as shown in column 'RR'. This represents the maximum total *RR* that could be produced by applying SBA-down to the seven applied windows of *EURCHF_RWDC0.1*.

Table 7.3: Summary of the best trading performance of the SBA-down model following the seven months out-of-sample period of the eight currency rates.

Currency Pairs	RR	Profit Factor	Total Number of Trades	Max Drawdown (%)	Win Ratio
EUR/CHF	80.41	2.30	1798	– 10.9	0.74
GBP/CHF	103.74	1.89	2539	– 10.7	0.72
EUR/USD	29.27	1.33	1935	– 2.9	0.66
GBP/AUD	75.98	1.57	2707	– 1.6	0.70
GBP/JPY	34.01	1.49	1748	– 3.9	0.69
NZD/JPY	122.65	1.71	3409	– 2.8	0.70
AUD/JPY	79.39	1.52	2861	– 3.0	0.70
EUR/NZD	395.92	2.56	3919	– 1.3	0.75

In this case, SBA-down generates 1,798 trades, as shown in column 'Total Number of Trades', with an overall win ratio of 0.74 as shown in column 'Win Ratio'. Whereas, in Table 7.4, in the case of EUR/CHF, we note that the minimum *RR*, during the trading period of seven months, is

[o] For each currency pair, trading with SBA-down, or DBA-down, we assume that we start with 1,000,000 units of counter currency. For example: in the case of EUR/CHF, we start with 1,000,000 CHF. Whereas in the case of NZD/JPY, we start with 1,000,000 JPY. However, in the case of SBA-up, or DBA-up, we assume that we start with 1,000,000 units of the base currency.

37.30%. In this case, SBA-down generates 1,167 trades with an overall win ratio of 0.73. Based on Table 7.3, in the best case, SBA-down can generate *RR* of 395% (see EUR/NZD, the last row in Table 7.3). Based on Table 7.4, in the worst case, SBA-down can generate *RR* of 4.72% (see EUR/USD, as shown in Table 7.4).

When examining the difference between the maximum and minimum *RR* produced by SBA-down, by comparing the *RR* shown in Tables 7.3 and 7.4, it is evident that this difference can be important. For example, in the case of AUD/JPY, the maximum *RR* estimated for SBA-down is 79.39% (Table 7.3). This is more than double the minimum *RR* obtained by applying SBA-down to AUD/JPY (which is 32.18 %, as reported in Table 7.4). The same note holds true for the *RR* obtained by SBA-down for all other currency rates. Keep in mind that this difference between the maximum and minimum *RR* is a result of the choice of the parameter *down_ind*. In other words, for a given currency pair, the max and min rates of return (*RR*) are obtained using two different values of *down_ind* (see Section 7.4.1). These results highlight the important impact of the *down_ind* value on the profitability of SBA-down. To conclude, SBA-down may have an attractive profitability in the best case. However, the value of *down_ind* may seriously affect the performance of SBA-down.

Table 7.4: Summary of the worst trading performance of the SBA-down model following the seven months out-of-sample period of the eight currency rates.

Currency Pairs	RR	Profit Factor	Total Number of Trades	Max Drawdown (%)	Win Ratio
EUR/CHF	37.30	2.56	1167	– 11.8	0.73
GBP/CHF	43.10	2.08	1270	– 9.5	0.72
EUR/USD	4.72	1.14	1649	– 3.8	0.62
GBP/AUD	59.87	1.51	2571	– 1.6	0.69
GBP/JPY	15.55	1.50	1290	– 4.2	0.67
NZD/JPY	54.07	1.72	2515	– 3.0	0.69
AUD/JPY	32.18	1.53	2111	– 3.2	0.68
EUR/NZD	148.13	2.56	2873	– 1.4	0.73

7.5.1.2 Experiments 7.1.2: Evaluating the performance of SBA-up

We apply SBA-up to each of the eight sets of rolling windows. Each set includes seven applied windows — the length of each is one month. For each set, and for each month of the applied windows, we use 100 different values of *up_ind*. We measure the maximum and the minimum *RR* as primary indicators of the best and worst performance of SBA-up respectively. We also measure the overall profit factor, *MDD*, and win ratio. We apply the same money management approach described in Section 6.5.1. Table 7.5 and Table 7.6 show, respectively, the estimated best and the worst performance of SBA-up when applied to the composed sets of rolling windows (see Section 7.3.2). Table 7.5 and Table 7.6 have the same interpretation as Tables 7.3 and 7.4. As in the previous experiment, we are mostly concerned with the general performance of SBA-up during the overall trading period (i.e. from 1/12015 to 31/7/2015) of each set of rolling windows.

Table 7.5: Summary of the evaluation of the best performance of applying SBA-up to the trading period of each set of rolling windows.

Currency Pairs	RR	Profit Factor	Total Number of Trades	Max Drawdown (%)	Win Ratio
EUR/CHF	64.78	2.25	1963	– 14.0	0.72
GBP/CHF	80.92	1.79	2435	– 22.6	0.71
EUR/USD	47.59	1.52	2000	– 4.1	0.69
GBP/AUD	68.86	1.64	2332	– 1.5	0.70
GBP/JPY	38.48	1.60	1545	– 1.7	0.70
NZD/JPY	200.73	2.13	3262	– 3.2	0.73
AUD/JPY	99.38	1.88	2486	– 1.9	0.72
EUR/NZD	401.50	2.75	3851	– 1.6	0.77

Table 7.6: Summary of the evaluation of the worst performance of applying SBA-up to the trading period of each set of rolling windows.

Currency Pairs	RR	Profit Factor	Total Number of Trades	Max Drawdown (%)	Win Ratio
EUR/CHF	17.77	2.20	1018	– 14.0	0.71
GBP/CHF	32.55	1.64	1313	– 23.5	0.70
EUR/USD	14.92	1.30	1709	– 2.1	0.67
GBP/AUD	26.21	1.53	1195	– 1.4	0.69
GBP/JPY	13.19	1.47	937	– 1.5	0.69
NZD/JPY	68.79	1.96	1968	– 2.8	0.72
AUD/JPY	34.06	1.73	1506	– 1.9	0.71
EUR/NZD	127.58	2.53	2334	– 1.6	0.76

For each currency pair, at the beginning of the first applied window, i.e. January 2015, SBA-up starts with capital equal to 1,000,000; this represents the initial, hypothetically, invested amount of money. From Table 7.5, using EUR/CHF, we note that the total *RR* are 64.78%. This represents the maximum possible *RR* that can be obtained by applying SBA-up to the seven applied windows of *EURCHF_RWDC0.1*. In this case, SBA-up generates 1,963 trades with an overall Win Ratio of 0.72. Whereas, in Table 7.6, again using EUR/CHF, we note that the minimum possible *RR* generated by SBA-up during the same trading period of seven months is 17.77%. In this case, SBA-up generates 1,018 trades with an overall win ratio of 0.71.

The objective of this experiment is to estimate the best and worst performances of SBA-up. Based on Table 7.5, in the best case SBA-up can generate *RR* of more than 400% (see the case of EUR/NZD, the last row in Table 7.5). Based on Table 7.6, in the worst case SBA-down can generate returns of 13.19% (see the case of GBP/JPY, as shown in Table 7.6).

When examining the difference between the maximum and minimum *RR* produced by SBA-up, by comparing the *RR* reported in Tables 7.5 and 7.6, it is clear that this difference can be important. For example, in the case of AUD/JPY, the maximum *RR* estimated for SBA-up is 99.38% (Table 7.5). This is more than double the minimum *RR* obtained by applying SBA-down to AUD/JPY (which is 34.06 %, as reported in Table 7.6). The same note holds true for the *RR* obtained by SBA-up for all other currency rates. For a given currency pair, the best and worst rates of return (*RR*) are obtained using two different values of *up_ind* (see Section 7.4.2). These results highlight the important impact of the value of *up_ind* on the profitability of SBA-up. To conclude, SBA-up may have an attractive profitability in the best case. However, the value of *up_ind* may seriously affect the performance of SBA-up.

7.5.2 Experiment 7.2: Is there one optimal value for the parameter down_ind or up_ind?

The objective of this section is to investigate whether there exists a specific value of the parameters *down_ind*, or *up_ind*, for which SBA-down, or SBA-up, will consistently generate the highest possible *RR*. We consider the same eight currency pairs as in Section 7.4 and apply SBA-down and SBA-up to each of these 100 times for a trading period of 12 months. Each time, for each month, we assign different values for the parameters *down_ind* and *up_ind* and measure the produced returns. In this experiment, our main interest is the values of the parameters *down_ind* and *up_ind* associated with the highest *RR*. Our goal is to analyse these values of these parameters. Table 7.7 shows the values of *down_ind* associated with the maximum monthly *RR* produced by SBA-down. For each currency pair (i.e. for each column in Table 7.7), the largest and the smallest values of *down_ind* are formatted in **bold** and *italic* respectively. For example, in column 'EUR/CHF' the numbers **– 0.01** and *– 0.84* denote, respectively, the largest and the smallest values of *down_ind* under this column. These **bold** and *italic* figures, for the same column, indicate the range of parameter *down_ind* in which the specified trading model performs best. Similarly, Table 7.8 shows the values of *up_ind* associated with the best monthly *RR* generated by SBA-up for each of the 12 trading months considered in this experiment.

Table 7.7: The values of *down_ind* corresponding to the highest RR generated by SBA-down for each month. For each currency pair, the figures in **bold** and *italic* indicate, respectively, the largest and the smallest values of *down_ind*.

Trading period		EUR/ CHF	GBP/ CHF	EUR/ USD	GBP/ AUD	AUD/ JPY	NZD/ JPY	GBP/ JPY	EUR/ NZD
2014	Aug	– 0.82	– 0.17	– 0.32	– 0.34	– 0.08	– 0.09	– 0.74	– 0.29
	Sep	– 0.08	– 0.04	– 0.92	**– 0.02**	– 0.15	– 0.43	– 0.62	– 0.16
	Oct	– 0.27	– 0.05	– 0.90	– 0.36	– 0.23	– 0.62	– 0.76	– 0.56
	Nov	– 0.40	*– 0.93*	– 0.07	– 0.13	– 0.45	– 0.27	– 0.53	*– 0.73*
	Dec	**– 0.01**	– 0.31	– 0.53	– 0.10	*– 0.62*	– 0.33	– 0.40	– 0.50
2015	Jan	*– 0.84*	– 0.30	*– 0.96*	– 0.36	– 0.32	*– 0.69*	– 0.25	**– 0.06**
	Feb	– 0.43	– 0.08	– 0.12	– 0.16	– 0.07	**– 0.03**	**– 0.05**	– 0.46
	Mar	– 0.01	**– 0.01**	– 0.57	*– 0.49*	– 0.11	– 0.07	– 0.05	– 0.49
	Apr	– 0.04	– 0.10	– 0.23	– 0.34	– 0.15	– 0.32	– 0.12	– 0.54
	May	– 0.07	– 0.02	– 0.38	– 0.37	– 0.16	– 0.46	– 0.22	– 0.67
	Jun	– 0.14	– 0.12	– 0.07	– 0.18	– 0.10	– 0.08	– 0.15	– 0.38
	Jul	– 0.39	– 0.02	**– 0.05**	– 0.41	**– 0.07**	– 0.13	*– 0.98*	– 0.28

Table 7.8: The values of *up_ind* corresponding to the highest *RR* generated by SBA-up for each month. Figures in **bold** and *italic* indicate, respectively, the largest and the smallest values of *up_ind* for each currency pair.

Trading period		EUR/ CHF	GBP/ CHF	EUR/ USD	GBP/ AUD	AUD /JPY	NZD /JPY	GBP/ JPY	EUR/ NZD
2014	Aug	0.06	0.03	0.41	0.01	0.13	**0.62**	0.07	0.31
	Sep	0.01	0.18	0.32	0.34	0.27	0.46	0.34	*0.15*
	Oct	0.15	0.03	0.05	0.12	0.46	0.24	0.12	0.35
	Nov	0.42	0.06	0.18	0.32	0.35	0.36	0.32	0.48
	Dec	0.36	**0.59**	0.06	0.10	0.51	0.13	0.10	0.46
2015	Jan	**0.73**	0.38	0.24	0.03	*0.03*	0.04	0.13	0.54
	Feb	0.04	0.28	0.13	**0.39**	0.23	0.02	0.29	0.43
	Mar	0.09	0.16	0.51	0.07	0.02	0.04	0.18	0.61
	Apr	0.11	0.06	0.42	0.10	**0.86**	0.16	*0.05*	0.62
	May	0.04	0.04	**0.61**	*0.01*	0.13	*0.01*	0.13	**0.71**
	Jun	*0.01*	*0.03*	0.51	0.20	0.07	0.07	0.72	0.51
	Jul	0.15	0.28	*0.04*	0.24	0.03	0.02	**0.96**	0.64

Experiment 7.2: Results' discussion

The objective of this experiment is to find out whether there exists a unique value of *down_ind* or *up_ind* for which SBA-down or SBA-up can consistently provide the best performance. By examining the **bold** and *italic* figures reported in Tables 7.7 and 7.8, we highlight the following observations:

1. Concerning SBA-down (Table 7.7): we note that SBA-down can generate maximum *RR* using either a small value or a large value of *down_ind.* For example, in the case of EUR/CHF: the maximum returns generated by SBA-down in January 2015 is obtained for *down_ind* = *– 0.84*. However, the maximum returns generated by SBA-down in December 2014 is obtained for *down_ind* = **– 0.01**. The majority of the results corresponding to the other currency pairs support this note that maximum *RR* can be attaint using either a small value or a large value of *down_ind*. For example, in the case of EUR/USD, SBA-down may generate high returns for a small value (as in January with *down_ind* = *– 0.96*) or for a large value (as in July 2015 with *down_ind* = **– 0.05**).

   In general, we note that the difference between the smallest and the largest values of *down_ind* (see numbers formatted in **bold** and *italic* for each column in Table 7.7) is more than 0.60 in most cases (the only exception is the case of GBP/AUD).
2. Concerning SBA-up (Table 7.8): we note that SBA-up can generate the best returns using either a small value or a large value of *up_ind*. For example, in the case of EUR/CHF: the maximum return generated by SBA-up in June 2015 is obtained for a low value *up_ind* = *0.01*. However, the maximum profits produced by SBA-up in January 2015 are obtained for a large *up_ind* = **0.73**. The majority of the results corresponding to the other currency pairs validate this note. For example, in the case of AUD/JPY, SBA-up may generate high returns with a small value of *up_ind* (as in January 2015 with *up_ind* = *0.03*) or for a large value of *up_ind* (as in April 2015, with *up_ind* = **0.86**).

   In general, we note that the difference between the smallest and the largest values of *up_ind* shown in bold in Table 7.8 is more than 0.50 in most cases (the only exception is the case of GBP/AUD).

To conclude, these two observations (1. and 2.) suggest that, in most cases, there is no specific value, or a tight range, for the parameters *down_ind* and *up_ind* for which SBA-down and SBA-up will exhibit the best performance consistently. This conclusion raise the need for a dynamic versions of BA.

7.5.3 *Experiment 7.3: Evaluation of the performance of DBA-down and DBA-up*

In this experiment we apply DBA-down and DBA-up to the eight sets of rolling windows (previously composed in Section 6.4.4). For each of DBA-down and DBA-up, we start with 1,000,000 monetary units as the initially invested capital. We use the same money management approach described in Section 6.5.1. Table 7.9 reports the general performance, during the overall trading period of seven months, of both versions of DBA in this experiment. The detailed monthly evaluation of applying DBA to these currency rates is provided in Appendix F.

Table 7.9: Summary of trading performance of the DBA-down and DBA-up models following the seven months out-of-sample period for the eight currency pairs.

Currency Pairs	Trading Strategy	RR	Profit Factor	Total Number of Trades	Max Drawdown (%)	Win Ratio
EUR/CHF	**DBA-down**	63.61	1.68	2008	– 11.7	0.71
	DBA -up	59.60	1.64	2105	– 14.7	0.71
GBP/CHF	**DBA-down**	73.39	1.85	2486	– 10.7	0.72
	DBA -up	77.44	1.67	2606	– 23.2	0.70
EUR/USD	**DBA-down**	12.66	1.80	1919	– 3.8	0.65
	DBA-up	25.62	1.93	2142	– 4.6	0.66
GBP/AUD	**DBA-down**	68.94	1.79	2542	– 1.7	0.69
	DBA-up	66.23	1.55	2469	– 1.8	0.69
GBP/JPY	**DBA-down**	32.77	1.90	1792	– 3.9	0.68
	DBA-up	32.07	2.20	1752	– 1.7	0.68
NZD/JPY	**DBA-down**	115.55	1.75	3194	– 2.7	0.71
	DBA-up	181.71	2.08	3196	– 3.2	0.73
AUD/JPY	**DBA-down**	73.59	1.57	2717	– 2.9	0.70
	DBA-up	87.35	1.75	2567	– 1.8	0.71
EUR/NZD	**DBA-down**	387.53	2.61	2892	– 1.6	0.75
	DBA-up	348.19	2.70	2960	– 1.7	0.75

The column ‘Currency Pair’ denotes the considered currency pair. The column ‘Trading Strategy’ indicates which version of DBA is applied. The column ‘RR’ is the total *RR*. The column ‘Profit Factor’ is calculated by dividing the sum of all generated returns by the sum of incurred losses during the overall trading period of seven months. The column ‘Max Drawdown (%)’ refers to the worst scenario measured as the worst peak-to-trough decline in capital during the trading period of seven months. The column 'Win Ratio' is the overall probability of having a wining trade. The last row in Table 7.9 is interpreted as follows: applying DBA-up to EUR/NZD generates a return of returns (*RR*) of 348.19 % during the trading period of seven months. In this case, DBA-up executes 2,960 trades with an overall Win Ratio of 0.75. The maximum drawdown in capital is – 1.7 %. The details of the monthly Rates of Return (*RR*) corresponding to applying DBA-down and DBA-up to these currencies pairs are shown below in Tables 7.10 and 7.11 respectively. These *RR* will be employed to compute the Sharpe ratio and Sortino ratio, Jensen's Alpha and Beta.

Table 7.10: Summary of monthly *RR* of trading with the DBA-down model following the seven months out-of-sample period of each of the eight currency pairs shown in Table 7.9.

Trading period	**EUR/ CHF**	**GBP/ CHF**	**EUR/ USD**	**GBP/ AUD**	**GBP/ JPY**	**NZD/ JPY**	**AUD/ JPY**	**EUR/ NZD**
Jan 2015	3.77	12.98	– 2.74	12.12	4.78	10.88	7.69	23.47
Feb 2015	13.56	12.26	2.53	10.96	6.69	20.96	12.87	43.08
Mar 2015	11.01	11.44	– 2.69	4.62	2.85	15.82	10.37	49.90
Apr 2015	5.99	6.96	3.78	6.97	7.59	19.39	10.10	45.40
May 2015	7.36	7.22	5.50	13.67	7.05	16.29	14.73	62.99
Jun 2015	8.94	10.39	5.50	6.74	5.35	11.81	9.58	73.46
Jul 2015	12.98	12.15	0.78	13.86	– 1.54	20.40	8.61	89.23
Sum	63.61	73.39	12.66	68.94	32.77	115.55	73.95	387.53

Table 7.11: Summary of monthly *RR* of trading with the DBA-up model following the seven months out-of-sample period of each of the eight currency pairs shown in Table 7.9.

Trading period	**EUR/ CHF**	**GBP/ CHF**	**EUR/ USD**	**GBP/ AUD**	**GBP/ JPY**	**NZD/ JPY**	**AUD/ JPY**	**EUR/ NZD**
Jan 2015	– 4.87	9.45	5.64	13.26	10.90	25.41	15.62	17.83
Feb 2015	10.32	12.71	4.15	8.30	3.08	17.26	10.41	42.06
Mar 2015	14.71	15.79	1.67	8.33	5.87	22.82	13.78	36.82
Apr 2015	7.37	10.31	1.20	8.28	5.65	21.32	6.73	48.94
May 2015	11.42	7.83	6.24	11.03	4.00	24.40	17.13	54.98
Jun 2015	10.21	10.07	4.77	8.63	3.96	27.64	12.15	71.73
Jul 2015	10.44	11.28	1.95	8.40	– 1.39	42.86	11.53	75.83
Sum	59.60	77.44	25.62	66.23	32.07	181.71	87.35	348.19

The Sortino and Sharpe ratios of both versions of DBA are reported in Table 7.12. The minimum acceptable return (MAR) and the risk-free rate are set to 5% per annum. The computation of Jensen's Alpha and Beta help to compare TSFDC to a particular benchmark. In this thesis, we adopt the buy and hold approach as a benchmark. For each currency pair, we apply the buy and hold approach on a monthly basis over the considered trading period from 1/1/2015 to 31/7/2015 (seven months).

Table 7.12: The Sortino ratio and Sharpe ratio of the two versions of DBA.

Currency pair	DBA-down		DBA-up	
	Sortino ratio	Sharpe ratio	Sortino ratio	Sharpe ratio
EUR/CHF	∞	2.7	32.3	1.5
GBP/CHF	∞	4.6	∞	4.6
EUR/USD	8.7	0.6	∞	1.9
GBP/AUD	∞	2.9	∞	5.2
GBP/JPY	56.1	1.3	60.8	1.3
NZD/JPY	∞	4.4	∞	3.4
AUD/JPY	∞	4.7	∞	3.9
EUR/NZD	∞	2.8	∞	2.7

Table 7.13: Summary of the monthly RR (%) obtained by applying the buy and hold strategy to each of the eight considered currency pairs. The trading period is from 1/1/2015 to 31/7/2015.

	Jan	Feb	Mar	Apr	May	Jun	Jul	Sum
EUR/CHF	– 12.88	1.75	– 1.95	0.10	– 1.41	0.99	1.77	– 11.63
GBP/CHF	– 9.68	5.17	– 2.01	– 0.60	0.57	1.92	2.69	– 1.94
EUR/USD	– 6.48	– 1.07	– 3.66	3.96	– 2.31	1.72	– 1.38	– 9.22
GBP/AUD	2.07	1.57	– 1.42	– 0.12	2.81	1.59	5.11	11.61
GBP/JPY	5.43	4.59	– 3.73	3.34	3.32	1.34	0.81	4.24
NZD/JPY	– 9.04	6.60	– 1.14	1.60	– 2.93	– 5.41	– 1.84	– 12.16
AUD/JPY	– 7.28	3.02	– 2.26	3.49	0.49	0.27	4.48	2.21
EUR/NZD	0.54	– 5.08	– 2.54	2.38	4.43	6.12	1.79	7.64

Table 7.14: The values of Jensen's Alpha and Beta of both versions of DBA with reference to the buy and hold approach as benchmark. The values are rounded to one decimal digit.

Currency pair	DBA-down		DBA-up	
	Jensen's Alpha	Beta	Jensen's Alpha	Beta
EUR/CHF	8.5	0.5	10.3	1.0
GBP/CHF	7.3	– 0.1	8.4	0.1
EUR/USD	2.0	0.7	– 0.3	– 0.2
GBP/AUD	10.0	1.7	6.8	0.2
GBP/JPY	2.5	0.0	0.32	– 0.6
NZD/JPY	52.9	0.1	49.0	0.5
AUD/JPY	8.2	0.2	8.1	– 0.5
EUR/NZD	57.3	2.7	52.0	2.9

Table 7.13 summarizes the monthly *RR* of applying the buy and hold approach to the eight currency pairs. We use these monthly *RR* to compare the performance of DBA with the buy and hold approach for each currency pair. The values of Jensen's Alpha and Beta corresponding to this comparison are reported in Table 7.14.

Experiment 7.3: Results' Discussion and Analysis

To begin, we examine the profitability of both versions of DBA. The monthly *RR* reported in Tables 7.10 and 7.11 indicate that both versions of DBA are, in most cases, profitable (except in a few cases; e.g. trading with DBA-down on EUR/CHF in January 2015, seen in Table 7.10). The total rates of return (*RR*), reported in Table 7.9, suggest that DBA can be very profitable (with *RR* of up to 387.53%; as in the case of applying DBA-down to EUR/NZD). The overall Win Ratio of DBA (i.e. the probability of having a winning trade) ranges between 0.75 (the case of applying DBA-down to EUR/NZD, see Table 7.9) and 0.65 (the case of applying DBA-down to EUR/USD, see Table 7.9). We consider this range to be reasonably acceptable.

We also noticed that the profitability of DBA can vary largely from one currency pair to another. For instance, from Table 7.9 we note that in the case of EUR/USD DBA-up generates rates of return (*RR*) of 12.66%; whereas DBA-up generates *RR* of 348.19% in the case of EUR/NZD (in the same table). This indicates that the performance of DBA may vary substantially from one currency pair to another. Which in turn suggests that a trader may want to consider other currencies, given that DBA may, possibly, perform better with these than those reported in this chapter.

We then inspect the risk of DBA. Based on the results reported in Table 7.9, we identify that, in most cases, the maximum drawdown (*MDD*) is no worse than – 4.6% (except in two cases: EUR/CHF and GB/CHF, Table 7.9). We consider these values of *MDD* to be reasonably low. Furthermore, the downside risk (Section 3.3) of DBA is null in most of these experiments, which is why most values of the Sortino ratio reported in Table 7.12 are positive infinity (denoted as ∞). Also, in most cases, the values of the figures in the 'Beta' column (indicated in Table 7.14) range between – 1.0 and 1.0. This range indicates that DBA is, generally, less volatile than the buy and hold approach. Keep in mind that the volatility of *RR* is usually used as an indicator of risk.

Lastly, we examine the risk-adjusted performance of DBA. For this purpose, we consider the values of the Sharpe ratio and Jensen's Alpha shown in Tables 7.12 and 7.14 respectively. The Sharpe ratio is consistently positive. A positive Sharpe ratio indicates that the DBA has surpassed the 5% annual risk-free rate. This fact indicates that DBA generates worthy excess returns for each additional unit of risk it takes. The Jensen's Alpha results (in Table 7.14) are, generally, consistent with the Sharpe ratio scores (though with one exception: the case of applying DBA-up to EUR/USD). We conclude that, generally, DBA earns more than enough returns to be compensated for the risk it took over the trading period.

We conclude from the previous analysis that DBA-down and DBA-up generate more returns and, in most cases, are less risky than the buy and hold method. Additionally, both versions of DBA can be highly profitable, with total *RR* of more than 300% (Table 7.9). We also argued that DBA can consistently deliver a positive Sharpe ratio. Finally, the established variety of the selected currency pairs in the initial dataset (Section 6.4.1) suggest that DBA can be profitably applied to a wide range of currency pairs.

7.5.4 Experiment 7.4: Comparing the RR of DBA and SBA

In this section we compare the *RR* of DBA to SBA (with SBA being assigned randomly picked parameters). We should mention that in Section 7.4.1 we evaluated the maximum possible *RR* of both versions of SBA. Rationally, DBA is not capable of producing higher *RR* than the estimated maximum *RR* of SBA (reported in Tables 7.3 and 7.4 respectively). Our objective in this experiment is rather to answer the question: What is the probability that the dynamic BA will produce higher *RR* than the static BA provided that the parameters of SBA (i.e. *down_ind* and *up_ind*) has been assigned random values?

To answer this question, in the case of EUR/CHF, we apply each of SBA-down and SBA-up 10,000 times to the seven applied windows of *EURCHF_RWDC0.1* using randomly picked values for the *down_ind* and *up_ind* parameters. Thus we obtain 10,000 *RR* for simulate trading with SBA-down and another 10,000 *RR* for simulating trading with SBA-up. We define γ as the fraction of how many of these 10,000 *RR* are less than the returns obtained by DBA-down and DBA-up (reported in Table 7.11 and Table 7.12). The *EURCHF_RWDC0.*1 is one out of eight sets of rolling windows composed in Section 6.4.4. We repeat the same procedure to compute γ based on each of the remaining seven sets of rolling windows.

The results of γ are shown below in Table 7.15. The number shown in the last row of column ‘EUR/USD’ is 89. This indicates that the probability that DBA-up generates higher *RR* than SBA-up (with randomly selected values of *up_ind*) is 89%. Similarly, the number shown in the last row of column ‘EUR/NZD’ is 91. This indicates that the probability that DBA-up generates higher *RR* than SBA-up (with randomly assigned values of *up_ind*) is 91%. The rest of numbers in this table are interpreted similarly.

Table 7.15: The values of the probability γ (%) for the considered currency pairs.

	EUR/ CHF	GBP/ CHF	EUR/ USD	GBP/ AUD	AUD/ JPY	NZD/ JPY	GBP/ JPY	EUR/ NZD
DBA-down vs. SBA-down	88	85	81	70	91	86	92	93
DBA-up vs. SBA-up	97	87	89	99	84	87	89	91

When examining the results in Table 7.15, we note that the probability that the DBA will produce higher *RR* than the SBA (with randomly chosen parameters) is, mostly, over 80%. We consider this probability as very good. The minimum value of γ is 70% (the case of GBP/AUD), which we consider as acceptable. We consider these results as an evidence of the efficiency of our procedures (FIND_DOWN_IND and FIND_UP_IND, Section 7.3.3) to find appropriate values of the parameters *down_ind* and *up_ind*.

7.5.5 Experiment 7.5: Compare the RR and risk of both versions of DBA

The objective of this experiment is to test whether there is a significant difference between the returns and risk of both versions of DBA. We assess the monthly rates of return (*RR*) and the *MDD* as the main indicators of the profitability and the risk respectively. Each of these sets consists of 56 observations (see Appendix E for details).

Firstly, we apply the Wilcoxon test with the null hypothesis being hypothesis that the median difference between the two sets of monthly *RR* of DBA-down and DBA-up is zero. In this test, the Wilcoxon test returns a *p*-value of 0.45. Given that this *p*-value is greater than common cut-off value 0.05, we cannot reject the Wilcoxon test's null hypothesis.

Secondly, we apply the Wilcoxon test to the two sets of monthly MDDs of DBA-down and DBA-up, the null hypothesis being that the median difference between them is zero. Appendix E comprises the details of these two sets. In this case, the Wilcoxon test returns a *p*-value of 0.57. Since the *p*-value is greater than 0.05, we could not reject the null hypothesis To conclude, the Wilcoxon tests do not suggest that the *RR* and the *MDD* of DBA-down and DBA-up are different.

7.6 DBA vs. other DC-based trading strategies

In this section, we compare DBA to two DC-based trading strategies, namely: 'DCT1' (Aloud [15]) and 'Alpha Engine' (Golub et al., [16]). The authors of these trading strategies did not claim that they employ any forecasting models. The details of these two strategies can be found in Sections 4.4.1 and 4.4.4 respectively. In this section, we aim to compare these strategies with DBA in terms of concept and performance.

7.6.1 The DC-based trading strategy: DCT1

In this section, we compare DBA with the trading strategy named 'DCT1' (Aloud [15]). The details of this strategy was reviewed in Section 4.4.1. Here we briefly recap the mechanism of this strategy, then we compare it with DBA.

DCT1 run a DC summary with a specific threshold named Δ_{xDC}. DCT1 consists of two trading rules:

- DCT1 initiates a new position (either buy or sell) at the DC confirmation point of one DC event.
- DCT1 closes this trade at the DC confirmation point of the following DC event.

Initially, the trader defines a range of thresholds. DCT1 examines this range to automatically compute: (1) the DC threshold Δ_{xDC}, and (2) the type of trade (whether contrarian or trend follower). For this purpose, DCT1 examines the profitability of each threshold in the specified range using historical price data (as training set). For each threshold value, the DCT1 will apply the above trading rules from two points of view: counter trend (CT) and trend follow (TF). Based on its produced *RR* during the training period, DCT1 returns the type of trade (CT or TF) and the threshold Δ_{xDC} corresponding to the highest produced returns. These parameters (type of trade and threshold) are then utilized to trade over the applied (out-of-sample) period.

We highlight the following differences between DBA and DCT1:

- Both versions of DBA, DBA-up and DBA-down, are contrarian. DCT1 could be either contrarian or trend follower.
- DBA triggers a new trade only if price's change during the OS event exceed certain threshold. DCT1 triggers a new trade exactly at DCC point of a DC event.

Nevertheless, DCT1 and DBA have a common feature which is: they both close trade at the DC confirmation point of the next DC event.

In term of evaluation of DCT1 and DBA we have the following notes:

- DCT1 was backtested using high frequency data of one currency pair: EUR/USD. Evaluating a trading strategy using one asset is not convincing according to Pardo [61] who emphasizes the importance of backtesting using a set of assets with different trends. In this chapter, DBA was backtested using eight currency pairs that exhibit different trends (see Section 6.4.1).

- The author reported that DCT1 was able to produce a rate of return of 6.2% during a testing period of one year (bid and ask prices being counted). The *RR* produced by DBA-up is 25.62% within seven months (Table 7.9, Section 7.5.3). These results indicate that DBA produces higher *RR* than DCT1. However, we did not count the bid and ask prices in our experiment. Therefore, it would be oppressive to make such statement.
- The author in [15] did not report any measurement of risk (e.g. *MDD*) or risk-adjusted metrics (e.g. Sharpe ratio) of DCT1. Therefore we cannot compare DCT1 with DBA from these perspectives.

7.6.2 The DC-based trading strategy: The 'Alpha Engine'

In this section, we compare DBA with the trading strategy named 'Alpha Engine' (introduced by Golub et al., [16]). The details of this strategy was reviewed in Section 4.4.4. Here we briefly recap the mechanism of this strategy, then we compare it with DBA.

The Alpha Engine consists of opening a counter-trend position when the overshot value (*OSV*) exceeds a specific threshold named 'ω':

$$\omega = \alpha \times theta \tag{7.1}$$

Where, *theta* is the employed DC threshold and α is a parameter. The value of α depends on the inventory size denoted as 'I'.

The Alpha Engine does not have an explicit stop-loss rule. Instead, it employs a sophisticated money management approach: When the Alpha Engine opens a position, it keeps increases and decreases the size of this position until a profit is reached. The increasing and decreasing of the position is designed to mitigate the accumulation of large inventory sizes during trending markets. The generation of a new trade (either buy or sell) depends on two factors:

- The inventory size denoted as 'I'; which is used to manage the value of α in (7.1). Thus, I serves to control the time at which Alpha Engine trigger a new trade.
- The size of a trade is a factor of a probability indicator (denoted as '$\mathcal{L}$'). The value of $\mathcal{L}$ is used as an estimation of the probability that the trend will move up or down provided the current state. The value of $\mathcal{L}$ is determined using a transition network model which has two states: the DC threshold *theta* and the threshold 'ω'. If the markets show normal behavior then $\mathcal{L} \approx 1$. On the other hand, in the case of abnormal market behavior $\mathcal{L} \approx 0$. The objective of $\mathcal{L}$ is to prevent the Alpha Engine from building up large positions which it cannot unload.

To summarize, the management of the position is a function of two variables: the size of inventory 'I' and the probability indicator '$\mathcal{L}$'. This approach of computing, and managing, the size of a position is an integrated part of the Alpha Engine. The Alpha Engine consider the uptrends and downtrends separately. So that it adopts two instances of the parameter ω; namely ω_{down} and ω_{up}. For more details about the mechanism of Alpha Engine see Golub et al., [16].

The Alpha Engine has three common features with Dynamic Backlash Agent (DBA):

- The positions of both trading strategies are countertrend, meaning that a price move down triggers a buy; a price up move, a sell.
- They both try to analyse the uptrends and downtrends separately.
- They both open positions when the *OSV* exceed certain *thresholds*. In the case of DBA, we have two thresholds (denoted as *down_ind* and *up_ind*). Similarly, in the case of the Alpha Engine, the authors identified two thresholds (denoted as ω_{down} and ω_{up}).

As for the differences between DBA and Alpha Engine, we have the following notes:

- The most important difference between DBA and Alpha Engine is that the former has an explicit stop-loss rule (Section 7.3.3) whereas the later does not. The money management approach employed by Alpha Engine makes it pretty complicated in comparing to DBA.
- The Alpha Engine may manage multiples positions simultaneously. Whereas, at any time DBA can have only one position opened (based on the adopted money management approach described in Section 6.5.1).
- Both DBA and Alpha Engine employ some parameters to decide when to initiate a new trade (i.e. *down_ind* and *up_ind* in case of DBA; ω_{down} and ω_{up} in case of Alpha Engine). However, they have different approaches to compute these parameters. DBA adopt a computational approach as explained in Section 7.3.3; whereas, Alpha Engine uses the size of the inventory 'I' to manage ω_{down} and ω_{up}.
- An important advantage of Alpha Engine is that the authors did not fine-tune any parameter to maximize performance. In the case of DBA, the value of DC threshold *theta* is to be set by the user. Further experiments should be done in this direction. For instance, we do not know how the value of the DC threshold *theta* may affect the performance of DBA?

The performance of the Alpha Engine was examined using a portfolio comprising 23 currency rates, sampled tick-by-tick, over a period of 8 years. The Alpha Engine produces a return of 21.34% (including bid and ask price). As can be seen in Table 7.9 (Section 7.5.3), DBA may generate total returns of more than 300% within 7 months (see the last raw in Table 7.9). These results indicate that DBA produces much higher *RR* than Alpha Engine. However, we did not count the bid and ask prices in our experiment. Therefore, it could be unfair to make such statement.

The authors in [16] reported that Alpha Engine has an annual Sharp ratio of 3.06. The Sharpe ratio is intended to measure the return earned in excess of the risk-free rate per unit of volatility (7.2). However, the authors did not specify the risk-free rate in [16]. This is a serious issue. For example, if we consider an annual risk-free rate of 5% then the total risk-free returns will be at least 40% over 8 years. In this case, the Alpha Engine will have a negative Sharpe ratio (because the risk-free rate is greater than the returns of the Alpha Engine, which is 21.34%). By contrast, the results suggest that DBA surpasses the 5% annual risk-free rate as it consistently provides a positive Sharpe ratio (see Table 7.12).

$$\text{Sharpe ratio} = \frac{R_p - R_f}{\sigma_p} \tag{7.2}$$

Where: R_p denotes the expected portfolio return and R_f is the risk-free rate. σ_p designs the standard deviation of the portfolio's returns and it's utilized to measure the volatility of the returns.

To conclude, in spite of the fact that the results indicates that DBA can produce higher *RR* than Alpha Engine; but it would be unfair to make such declaration as we did not consider the bid-ask spread when evaluating the performance of DBA (Table 7.9).

7.7 Summary and conclusion

In this chapter, we introduced a contrarian trading strategy, named BA, which is based on the DC framework. BA opens a position when the overshoot value (OSV) exceed the values of specific parameters (Section 7.2). BA has two types: Static and Dynamic. The static type of BA, SBA, relies on the expertise of the investor to set these parameters. By contrast, the dynamic type of BA, DBA, applies a DC-based computational approach to examine historical prices to automatically find appropriate values of the parameters. Then, DBA uses these values to trade over the out-of-sample (trading) period (Section 7.3). We consider DBA, the autonomous type of BA, as our original trading strategy, whereas, SBA serves to compute the best and worst possible performances of BA (Section 7.5.1).

To evaluate the performance of DBA, we adopted the same methodology employed in chapter six to assess the performance of TSFDC: We applied DBA to the eight sets of rolling windows previously composed in Section 6.4.4. Each set corresponds to one currency pair. Each set comprises a training period to which we applied the predetermined DC-based computational approach to compute the values of the parameters. We used a set of evaluation metrics to measure the profitability and risk of DBA. As a benchmark model, we implemented the standard buy and hold strategy. We also compared the performance of DBA to other DC-based trading strategies (Section 7.6).

The experimental results (reported in Section 7.4.2) suggest that DBA is successful. By examining the returns reported in Table 7.9 (Section 7.4.2), we can conclude that DBA can be highly profitable (with total *RR* of more than 380%) and yet retain an attractive level of risk (with an *MDD* equal to – 1.6%). When examining the values of Jensen's Alpha (shown in Table 7.21, Section 7.4.2), we can see that DBA generates promising returns compared to the level of risk it takes in relating to the buy and hold method. The values of Beta (Table 7.21, Section 7.4.2) point out that in the majority of cases DBA is less, or equally, volatile than the buy and hold method. We compared the DBA to two other DC-based trading strategies (Section 7.6). Despite that DBA generates higher *RR* than two other DC-based trading strategies, however we could not make definitive conclusion regarding whether DBA outperforms these strategies. This because they both considered the bid-ask spread in their experiments, but we did not in this thesis.

To conclude, in this chapter we developed a DC-based trading strategy, named DBA; which does not rely on any forecasting model. As our main contribution, we argued that DBA can be highly profitable. We examined the effectiveness of DBA over eight different currency pairs that have different patterns. Therefore, we believe that DBA could be successful in a broad range of currency pairs. Despite what would be considered as experimental weaknesses (e.g. ignoring the transaction costs), we argue that these results provide an evidence as to the usefulness of the DC framework as a basis of trading strategies.

8 Comparing TSFDC with DBA

In this thesis, we presented two trading strategies TSFDC (Chapter 6) and DBA (Chapter 7). In this chapter, we aim to compare the performances of TSFDC and DBA. The objective is to find out whether one of them outperforms the other in every aspect. More particularly, we focus on three aspects: profitability, drawdown and risk-adjusted returns. We rely on the results of the experiments organized in Chapters 6 and 7 to compare TSFDC and DBA.

We start this chapter with a brief summary of the two trading strategies, TSFDC and DBA. We then list the adopted metrics that will be utilized to compare TSFDC and DBA. Next, we summarize the results of our experiments (carried out in Chapters 6 and 7). Finally, we compare TSFDC and DBA using these results.

8.1 Introduction

In Chapter 5, we presented a forecasting model which aims to predict the change of the direction of market's trend under the DC framework. In Chapter 6, we introduced a trading strategy named TSFDC, which is based on the proposed forecasting model. TSFDC uses the historical prices of a given currency pair as an in-sample dataset to train the forecasting model presented in Chapter 5. It then relies on the formed prediction model to decide when to trigger a buy or a sell signal during the out-of-sample (i.e. trading) period.

In Chapter 7, we introduced a trading strategy named DBA. In contrast to TSFDC, DBA does not employ any forecasting model. DBA initiates a trade when the magnitude of price change exceeds specific parameters. DBA run a predefined procedure, which examines historical (in-sample) prices using a DC-based approach, to determine the value of these parameters. Then, DBA uses these values of these parameters to decide when to start a trade during the out-of-sample (i.e. trading) period.

Both TSFDC and DBA are contrarian strategies. We evaluated both strategies using the same methodology and datasets. We considered eight currency pairs from the FX market sampled minute-by-minute. For each currency pair, we composed seven rolling windows (see Section 6.4.4). We applied both strategies to these rolling windows. We concluded that both strategies, TSFDC and DBA, outperform the buy and hold approach. We also argued that both strategies can be highly profitable.

In this chapter, we compare the performances of TSFDC and DBA with the objective of studying whether one of them outperforms the other in every aspect. Mainly, we focus on three fundamental aspects: profitability, maximum drawdown and risk-adjusted returns. For this purpose, we use the results of the experiments organized in Chapters 6 and 7.

8.2 Comparing the performances of TSFDC and DBA: Criteria of comparison

In this section, we list the metrics that will be considered to compare the performance of TSFDC and DBA. The detailed description of these metrics has been provided in Section 3.4. In this section, we provide a recap of each metric. These metrics are selected to represent three aspects:

- *Profitability*: We consider the 'Rate of Return (*RR*)' as the main metric to evaluate the profitability of a trading model. Let Total Profit (*TP*) represents the overall losses or gains during the entire trading period. We define RR as the gain or loss on an expressed as a percentage of the amount invested. In (8.1) *INV* denotes the initial capital employed in the investment.

$$RR = \frac{TP}{INV} * 100 \tag{8.1}$$

- *Maximum Drawdown*: We use the Maximum Drawdown (*MDD*) to measure the risk of a trading strategy (as in [16] [17]). The *MDD* measures the risk as the worst-case-scenario of a given trading strategy. In (8.2), the subscript *I* denotes the time-index (i.e. time-stamp). *Current capital*$_i$ denotes the amount of capital counted at time (*i*). The *maximum capital* refers to the peak capital's value that has been reached since the beginning of trading up to time *i*. Thus, $drawdown_i$, (8.2), is interpreted as the peak-to-trough decline in capital during the period of an investment. Note that, based on (8.2), we have $drawdown_i \leq 0$ for all *i*. The *MDD* (8.3) is the minimum value among all computed $drawdown_i$.

$$drawdown_i = \frac{current\ capital_i - maximum\ capital}{maximum\ capital} \tag{8.2}$$

$$MDD = Min\ (drawdown_i) \tag{8.3}$$

- *Risk-adjusted return*: Under this aspect, we consider the 'Sharpe ratio' [68]. The Sharpe ratio is the average return earned in excess of the risk-free rate per unit of volatility. The formula to calculate the Sharpe ratio is:

$$\text{Sharpe ratio} = \frac{R_p - R_f}{\sigma_p} \tag{8.4}$$

Where: R_p denotes the expected portfolio returns; R_f is the risk-free rate; σ_p designs the standard deviation of the portfolio's returns.

8.3 Comparing the performances of TSFDC and DBA: The results

In this section, we summarize the results of the evaluations of TSFDC and DBA (from Chapters 6 and 7). More particularly, we consider the results corresponding to the metrics of the three aspects listed in the previous section. The results of each aspect are summarized in one table. For example, Table 8.1 summarizes the results of the *RR* of both strategies TSFDC and DBA. Similarly, Table 8.2 shows the results *MDD*, and Table 8.3 shows the results of the Sharpe ratio. The last row, of each of these tables, denotes the average of the results of each trading model for the selected metrics. Although not statistically significant, comparing these averages for both strategies can provide a general indication of the superiority of one of them, if any.

In these tables, for each currency pair (i.e. each row), one number is formatted in **bold**. This formatting is to highlight the best performance among the four trading models: TSFDC-down, TSFDC-up, DBA-down and DBA-up. For example, in Table 8.1, let us take the results of the currency pair EUR/CHF. The number '**84.59**' is formatted in **bold**; which implies that the best *RR* is **84.59%** which was produced by TSFDC-down (as shown in the corresponding column's header). The same interpretation applies for the remaining rows (i.e. currency pairs) of Table 8.1.

Table 8.1: Comparing the profitability, measured as 'Rate of returns (*RR*)', of TSFDC and DBA. For each currency pair, the **bold** figure represents the best performance across the four strategies: TSFDC-down, TSFDC-up, DBA-down, and DBA-up.

Currency pair	TSFDC-down	TSFDC-up	DBA-down	DBA-up
EUR/CHF	**84.59**	63.03	63.61	59.60
GBP/CHF	94.03	**115.19**	73.39	77.44
EUR/USD	27.04	**36.09**	12.66	25.62
GBP/AUD	**92.63**	63.03	68.94	66.23
GBP/JPY	32.48	28.91	**32.77**	32.07
NZD/JPY	**190.73**	183.13	115.55	181.71
AUD/JPY	104.11	**116.35**	73.59	87.35
EUR/NZD	489.13	**571.89**	387.53	348.19
Average *RR*	139.34	**147.20**	103.51	109.78

Likewise, for each currency pair (i.e. each row) shown in Tables 8.2 and 8.3, the number formatted in **bold** denote the supremacy of a trading strategy under the specified metric. In the next section, we focus on the figures formatted in **bold** to compare the performances of TSFDC and DBA.

Table 8.2: Comparing the Maximum Drawdown (*MDD*) of TSFDC and DBA. For each currency pair, the **bold** figure represents the best performance across the four strategies.

Currency pair	TSFDC-down	TSFDC-up	DBA-down	DBA-up
EUR/CHF	– 13.4	– 15.1	**– 11.7**	– 14.7
GBP/CHF	– 12.1	– 10.8	**–10.7**	– 23.2
EUR/USD	– 5.0	– 5.8	**– 3.8**	– 4.6
GBP/AUD	– 3.4	– 3.5	**– 1.7**	– 1.8
GBP/JPY	– 4.8	– 5.7	– 3.9	**– 1.7**
NZD/JPY	– 4.9	– 4.0	**– 2.7**	– 3.2
AUD/JPY	– 5.0	– 5.2	– 2.9	**– 1.8**
EUR/NZD	– 4.6	– 5.1	**– 1.6**	– 1.7
Average *MDD*	– 6.7	– 6.9	**– 4.9**	– 6.6

Table 8.3: Comparing the risk-adjusted return, in terms of the Sharpe Ratio, of TSFDC and DBA. The risk-free rate is 5% per year. For each currency pair, the **bold** figure represents the best performance across the four strategies.

Currency pair	TSFDC-down	TSFDC-up	DBA-down	DBA-up
EUR/CHF	**2.6**	2.1	1.3	2.4
GBP/CHF	3.2	2.0	4.1	**4.1**
EUR/USD	1.0	1.3	0.4	**1.6**
GBP/AUD	3.7	3.3	2.5	**4.6**
GBP/JPY	1.1	0.8	**1.3**	1.1
NZD/JPY	2.7	3.6	**4.0**	3.1
AUD/JPY	3.6	**4.2**	4.1	3.5
EUR/NZD	2.0	2.2	**2.5**	2.4
Average *Sharpe ratio*	2.5	2.4	2.5	**2.9**

8.4 Comparing TSFDC and DBA

8.4.1 In term of profitability

In this section, we analyze the rate of returns *RR* results shown in Table 8.1. The analysis of the **bold** figures in Table 8.1 suggests that TSFDC generates more *RR* than DBA in 7 out of 8 currency pairs (with only the exception of GBP/JPY, when DBA-down produces a higher *RR*). The averages of the *RR* (shown in the last row of Table 8.1) indicate that TSFDC generates markedly higher returns than DBA. For instance, the average *RR* of TSFDC-up over the eight currencies rate is

147.20%, whereas neither DBA-down nor DBA-up has an average *RR* of more than 110% (the last row in Table 8.1). These observations suggest that TSFDC is more profitable than DBA.

8.4.2 In term of maximum drawdown

In this section, we compare the estimated *MDD* of TSFDC and DBA. The analysis of the **bold** figures in Table 8.2 indicate that TSFDC has a worse *MDD* than DBA in all cases. Although purely indicative, the averages of MDD's results (the last row in Table 8.2) indicate that both versions of DBA has better *MDD* than both versions of TSFDC. Some studies (e.g. [17] [16] [4]) consider the maximum drawdown *MDD* as a metric to measure the risk of a trading strategy. Thus, we conclude from the results of *MDD* that DBA is more advantageous than TSFDC in terms of risk.

8.4.3 In term of risk-adjusted performance

In this section, we compare the risk-adjusted returns of TSFDC and DBA. When we examine the values of Sharpe's ratio (in Table 8.3), we note that DBA provides a greater Sharpe ratio in 6 out of 8 currency pairs (see **bold** figures in Table 8.3). In general, we note that the average Sharpe ratios of DBA-up is larger than the averages Sharpe ratios of both versions of TSFDC (as shown in the last row of Table 8.3). However we also note that the average Sharpe ratios of DBA-down and TSFDC-down are both equal to 2.5. Therefore, we do not consider the supremacy of DBA over TSFDC, in term of risk-adjusted returns, as considerable.

8.5 Conclusion

In this thesis, we introduced two trading strategies based on the DC framework, namely TSFDC (Chapter 6) and DBA (Chapter 7). The former employs a forecasting approach to decide when to trade while the latter does not. In this chapter, we compared the performances of TSFDC and DBA with the objective of finding out whether either of these strategies outperforms the other in every aspect. Principally, we considered three aspects: profitability (measured as rate of returns *RR*), maximum drawdown *MDD* (used as a measure of risk) and risk-adjusted return (as measured by the Sharpe ratio). The comparisons carried out in this chapter indicates that TSFDC is more profitable than DBA. However, DBA is less risky than TSFDC. We also observe that DBA marginally outperforms TSFDC in terms of risk-adjusted returns. We conclude that neither TSFDC nor DBA outperforms the other in all aspects. These results conform to the Modern Portfolio Theory (MPT) of Markowitz [103] which states that to generate more profit, an investor must undertake higher risk. With TSFDC being more profitable but riskier than DBA, choosing which model to implement relies on the level of risk the investor is willing to withstand.

Part III

Concluding Remarks

9 Conclusions

The Directional Change (DC) Framework is an approach to study price movements in financial markets. Many studies have reported that the DC framework is helpful in analysing the price movements and traders' behaviors in the FX market. Some studies have tried to develop trading strategies based on the DC framework. This study set out to explore, and consequently to provide a proof of, the potential of the Directional Changes framework as the basis of profitable trading strategies. This chapter provides a summary of the thesis, points out its contributions, and discusses possible future research work.

9.1 Summary

The DC framework is an event-based technique to summarize price movements in financial market. Under the DC framework the market is cast into alternating upward and downward trends. A trend is identified as a change in market price larger than, or equal to, a given threshold. This threshold, named *theta*, is set by the observer and usually expressed as a percentage. In Section 4.3 we reviewed some studies (e.g. [11] [12] [13] [74]) those have pointed out to the usefulness of the DC framework to analyze price movements in the FX market. It was also reported that an ideal DC-based trading strategy could be amazingly profitable; nonetheless, the full promise of the DC framework for developing trading strategies has not been completely exploited [16] [19]. In Chapter 3, we sorted existing trading strategies into two groups: 1) strategies that embed forecasting approaches (e.g. [6] [41] [42] [43] [44] [45] [46]); and 2) strategies that do not rely on any forecasting model (e.g. [7] [8] [53] [56] [57] [58] [60]).

Our intended aim of this research was to explore, and consequently to provide a proof of, the convenience of the Directional Changes (DC) framework as a basis of profitable trading strategies. To attaint our stated objective, and in line with existing research, we developed two DC-based trading strategies: one strategy, named TSFDC, which is based on forecasting DC (Chapter 6); and a second strategy, named DBA, which is based on the DC framework but does not employ any forecasting method (Chapter 7). We examine the performances of TSFDC and DBA in the foreign exchange (FX) market using the same methodology and datasets.

In this chapter, we summarize the functionalities of TSFDC and DBA. We also highlight the differences between our proposed trading strategies, TSFDC and DBA, and some existing DC-based trading strategies. Finally, we list our contributions and suggest future researches.

9.2 In a nutshell: TSFDC and DBA

9.2.1 TSFDC: A trading strategy based on forecasting Directional Changes

In Chapter 6 we introduced our first DC-based trading strategy, named TSFDC. TSFDC was designed as a forecasting-based trading strategy. Forecasting the change of trend's direction in financial time series is a common problem (e.g. [41] [80] [81] [87]). However, we noticed that this problem had not been formalized under the DC context. Therefore, as a first step, we provided a formalization of the problem of forecasting the change of a trend's direction under the DC framework (Section 5.2.2). To this end, we tracked price movements using two DC thresholds: *STheta* and *BTheta*. We formalized the problem as 'to forecast whether the magnitude of total price change of a DC trend, as observed under *STheta*, will be at least equal to *BTheta* before the trend changes'.

We also discovered an original DC-based indicator, named OSV_{BTheta}^{STheta}, and we selected an appropriate machine learning procedure to propose one solution for the established forecasting problem (Section 5.4.1). We applied our forecasting model to eight currency pairs from the foreign exchange market. The experimental results suggested that the accuracy of our prediction model range between 62% and 82% outperforming the traditional ARIMA technique (Section 5.6.1). These results indicate that our proposed indicator, OSV_{BTheta}^{STheta}, is useful for forecasting purpose under the DC framework.

The second step consisted of employing the established forecasting model to develop a trading strategy named 'TSFDC' (Chapter 6). TSFDC relies on this forecasting model to decide when to initiate a new trade. To evaluate the performance of TSFDC, we applied it to eight currency pairs, using a monthly-based rolling windows approach, for an overall out-of-sample trading period of seven months. Experimental results suggested that TSFDC can be highly profitable (Section 6.6). We also argued that TSFDC outperforms another DC-based trading strategy (Section 6.7).

9.2.2 DBA: The second DC-based trading strategy

The second trading strategy, named DBA, was introduced in Chapter 7. The objective was to develop a successful DC-based trading strategy that does not rely on any forecasting model. DBA opens a position when the overshoot value exceeds a particular *threshold*. DBA examines the historical prices using a DC-based computational approach to determine this *threshold.* To evaluate the performance of DBA, we followed the same experimental methodology and utilized the same datasets previously adopted to evaluate the performance of TSFDC in Chapter 6. Experimental results suggested that DBA can be highly profitable (Section 7.5.3).

It's worthy to highlight an important difference between TSFDC and DBA: In contrast to DBA, TSFDC relies on a forecasting model which: 1) has a clearly-defined dependent and independent variables and 2) employs a machine learning procedure to predict the dependent variable (Section 5.4). Thus, in contrast to DBA, we consider TSFDC to be a forecasting-based trading strategy.

A comparison between the performances of TSFDC and DBA was carried out in Chapter 8. The objective was to find out whether either TSFDC or DBA could outperform the other. This comparison focused on the three aspects: profitability, drawdown and risk-adjusted returns. The results suggested that, in general, TSFDC generates higher returns than DBA (Section 8.3). However, the results suggested that DBA has better maximum drawdown than TSFDC (Section 8.3). The results also indicated that DBA has a slightly better risk-adjusted performance than TSFDC (Section 8.3). We concluded that none of DBA and TSFDC could outperform the other in every aspect. These results suggest that each of DBA and TSFDC could be an attractive choice for different types of traders. Choosing which strategy to adopt, TSFDC or DBA, would depend on the level of risk the trader is willing to undertake.

Despite what would be considered as defects in our experiments (e.g. ignoring the transaction costs), we argue that the results of the evaluation of the performances of TSFDC (Section 6.6.1) and DBA (Section 6.5.3) support our objective as to provide a proof of the usefulness of the DC framework as a basis of profitable trading strategies.

9.3 Comparing TSFDC and DBA with other DC-based trading strategies

In Section 4.4 we reviewed some existing DC-based trading strategies. In Chapters 6 and 7 we compared TSFDC and DBA to some of these trading strategies. In this section we review the differences between our proposed strategies, TSFDC and DBA, and other existing DC-based trading strategies.

9.3.1 Comparing TSFDC with other DC-based trading strategies

In Section 6.7, we compared TSFDC to two DC-based trading strategies proposed by Gypteau et al., [78] and Kampouridis and Otero [17]. The reason of chosen these two trading strategies is that the authors in both studies, [17] and [78], claimed that they proposed trading strategies that employed forecasting models. In this section we summarize these comparisons.

1. In Section 6.7.1 we compared TSFDC to the trading strategy presented by Gypteau et al., [78]. Here, we recap the following differences:

- TSFDC is founded on the well-articulated forecasting approach established in Chapter 5 which has clearly identified dependent and independent variables. Despite the fact that the authors in [78] declared that they "*…aims to find an optimal trading strategy to forecast the future price moves of a financial market"*; they did not identify any dependent or independent variables.
- TSFDC relies on forecasting the change of the direction of market's trend to decide when to start a new trade. Whereas, the trading strategy by Gypteau et al., [78] was presented as a GP-tree. This GP-tree includes multiples DC thresholds. The detection of DC events of these thresholds is interpreted as 'True' and 'False' values. Based on the detected event(s), the expression represented by a GP tree evaluates to a Boolean value that indicated the action (either buy or sell) to be taken.
- In Section 6.7.1 we argued that TSFDC can, probably, generated higher *RR* than the trading strategy introduced by Gypteau et al., [78].

2. We compared the trading strategy named DC+GA presented by Kampouridis and Otero [17] with TSFDC in Section 6.7.2. Here, we recap the following remarks:
   - TSFDC has different trading rules of when to start or end a trade than DC+GA: For instance, TSFDC initiates a trade either on the DCC points of *STheta* or *BTheta*. In practice terms, TSFDC focuses on magnitude of price change to decide when to start a trade. Whilst, DC+GA initiates a trade when the *time* length of OS event lasts longer than a specific time-parameter (see Section 6.7.1 for details).
   - TSFDC relies on the forecasting approach presented in Chapter 5 to decide when to trigger a new trade. Whereas, DC+ GA employs a GA module to anticipating the best *time* at which it should initiate a trade.
   - TSFDC uses two DC thresholds; whilst DC+GA may consider up to N_{theta} thresholds to decide when to initiate a trade.
   - The authors in [17] claimed that their objective was "*to offer a more complete analysis on the directional changes paradigm from a financial forecasting perspective*." However, in contrast to our forecasting approach established in Chapter 5, they did not identify any dependent or independent variables!
   - We compared the results of TSFDC and DC+GA in Section 6.7.2. We concluded that TSFDC outperforms DC+GA in terms of produced *RR* and risk-adjusted returns.
   - However, the results of *MDD* suggest that DC+GA is less risky than TSFDC.

- A common feature between TSFDC and DC+GA is that they both try to analyse the uptrends and downtrends separately.

9.3.2 *Comparing DBA with other DC-based trading strategies*

In Section 7.6 we compared DBA to two other DC-based trading strategies, namely 'DCT1' [15] and 'the Alpha Engine' [16]. The authors of DCT1 and Alpha Engine did not claim that their proposed trading strategies employ any forecasting model. In this section we briefly recap the differences and similarities between these trading strategies and DBA.

1. As for the differences between DBA and DCT1 [15], we have the following comments:
   - DBA is a contrarian strategy. Whereas DCT1 can be either a contrarian or trend's follower.
   - DBA triggers a new trade only if price's change during the OS event exceed certain threshold. DCT1 do not have such 'threshold'. DCT1 triggers a new trade when a DC event is confirmed (i.e. at the DCC point).
   - Nevertheless, DCT1 and DBA have a common feature which is: they both close trade at the next DC confirmation point.
   - According to our experiments, DBA produces higher *RR* than DCT1. However we could not confirm that DBA outperforms DCT1 as, in contrast to the experiment in Aloud [15], we did not consider the bid-ask spread in this thesis.

2. As for the differences between DBA and Alpha Engine [16], we have the following notes:
   - The most important difference between DBA and Alpha Engine is that the former has an explicit stop-loss rule (Section 7.3.3) whereas the later does not. Alpha Engine employs a sophisticated money management approach. The Alpha Engine uses a transition network model to control the size of each new trade (Section 4.4.4).
   - It derives from the previous point that the Alpha Engine may manage multiples positions simultaneously. Whereas, at any time DBA can have only one position opened.
   - DBA employs a computational approach to decide the *OSV* at which it should make a new trade. Whereas, the Alpha Engine take into consideration the total amount of inventory to decide the value of *OSV* at which it should make a new trade.
   - An important advantage of Alpha Engine is that the authors did not fine-tune the parameters to maximize performance. In the case of DBA, the user must specify the DC threshold *theta.* Further experiments should be done to examine how the value *theta* may affect the performance of DBA?

- According to our experiments, DBA produces higher RR than Alpha Engine. However, we could not confirm that DBA outperforms the Alpha Engine. This is because the authors in [16] did count the bid-ask spread, but we did not (Section 7.5.2).

Nevertheless, we can note some similarity between DBA and Alpha Engine:

- They both trigger contrarian trades.
- They both open positions during the overshoot when the price's change reach a specific threshold.
- They both try to analyse the uptrends and downtrends separately.

9.4 Contributions

This thesis contributes toward providing evidence as to the potential of the DC framework as the foundation of trading strategies. The major contributions of this work can be summarized as follows:

- We formulated the problem of forecasting the change of trend's direction under the DC framework (Chapter 5). The objective was to forecast whether the current DC trend, of threshold *STheta*, will continue so that its total price change will reach another threshold named *BTheta* (Section 5.3). This objective was shortened as to predict one Boolean variable named *BBTheta*.
- The second contribution is discovering a useful DC-based indicator named OSV_{BTheta}^{STheta}. We proved that this indicator is helpful in forecasting the change of the direction of a market's trend under the DC context. We used this indicator to establish a forecasting model that have a pretty good accuracy ranging between 62% and 82% (Section 5.6.2). We also proved that our forecasting model has better accuracy than the ARIMA model (Table 5.4, Section 5.6.1).
- We employed the proposed forecasting model to develop a successful trading strategy, named TSFDC (Chapter 6). We argued that TSFDC outperforms other DC-based trading strategies (Section 6.7). The results of the preliminary tests suggested that TSFDC could produce positive Sharpe ratio consistently (Section 6.6.1).
- We presented a second trading strategy, named DBA, which is based on the DC concept but does not rely on any forecasting model (Chapter 7). DBA follows a computational approach to examine the historical prices to discovering profitable trading rules of when to initiate a trade. We argued that DBA can be highly profitable (Sections 7.6.1). The

results of the preliminary tests suggested that DBA could produce positive Sharpe ratio consistently (Section 7.5.3).

The comparison of TSFDC and DBA, carried out in Chapter 8, suggested that TSFDC produces more profits than DBA; but, DBA is less risky than TSFDC. Therefore, each of them can be more advantageous for different types of traders, based on the level of risk the trader is willing to withstand (Section 8.4).

To conclude, the objective of this thesis was to explore, and consequently to provide a proof of, the usefulness of the DC framework as the basis of profitable trading strategies. Despite some experimental flaws (e.g. ignoring the transaction costs), the results of the evaluation of the performances of our proposed trading strategies, TSFDC (Section 6.6.1) and DBA (Section 7.5.3), support our stated objective. The results of rate of return (*RR*) generated by TSFDC and DBA is much less than the estimated maximum annual *RR* that could be possibly generated by a DC-based trading strategy (which is 1600% [19]), but, in our opinion, a vital step in the right direction.

9.5 Future works

In this thesis, we introduced two DC-based trading strategies: TSFDC and DBA. We believe that both strategies can be further improved in many ways.

9.5.1 Money management: Controlling order size

In this thesis we focused on discovering profitable trading rules under the DC framework. However, a trading system must consider two other essential parts: risk control and money management [33]. Money management refers to the actual size of the trade to be initiated [86]. Some studies (e.g. [61] [99]) reported that models that do not take into consideration effective money management decisions can lead to sub-optimal solutions. In this thesis we adopted a naïve approach of money management (previously described in Section 6.5.1). Thus, the overall performances of TSFDC and DBA can be improved by developing a good money management module. For this purpose, a good objective would be to relate the sizing of a new trade to periodic patterns of market activity. In other words, to discover the time at which TSFDC or DBA would mostly be profitable and, then, use this discovery to decide the size of a new trade. Aloud el al., [19] reported that periodic patterns do exist under the DC framework. For example, Fig. 9.1, shown below, reports the number of events of two DC thresholds (0.03% and 0.10%) in different time periods of the 5th, 7th and 9th January 2009 in EUR/CHF mid-price time series. This figure pinpoints tow important observations: (a) the same periods of time with the same threshold size on different days may contain a different number of events, and (b) with the same threshold size, some periods on the same day have more events than others [19].

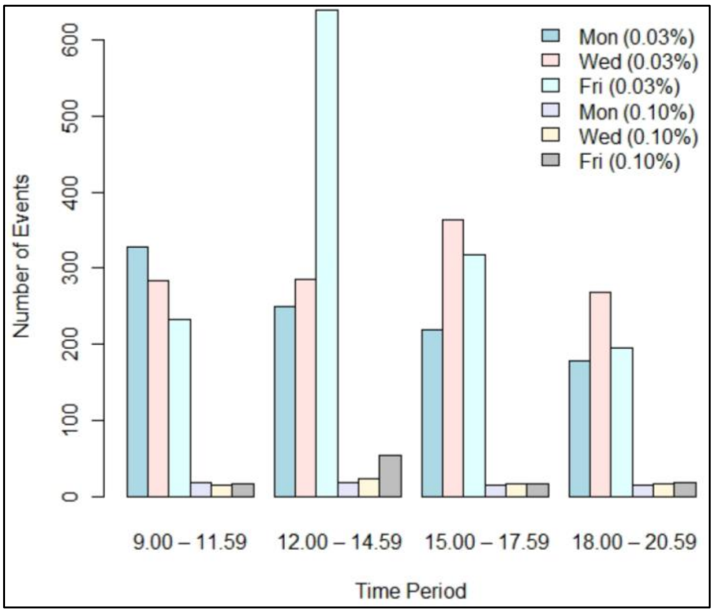

Fig. 9.1 Number of DC events of threshold 0.03% and 0.10% in different periods of the 5th (Monday), 7th (Wednesday) and 9th (Friday) January 2009 in EUR/CHF mid-price time series. Source Aloud et al. [19]

Based on these observations, a DC-based trading strategy will probably have different performances during different time periods. As a future work, we suggest to analyze the returns of TSFDC and DBA as function of time period (similarly to Fig. 9.1). In other words, we suggest to discover a relationship between time periods (i.e. hours of the day, days of the week) and the generated returns of each trade triggered by TSFDC and DBA. For this purpose, we can examine the existence of 'association rules' between the returns of all trades and time periods. Association rules can be utilized to discover and analyze the existence of strong rules among several variables, in databases, using some measures of interestingness [104]. Some machine learning algorithms (e.g. Apriori algorithm [105], OPUS search algorithm [106]) could be useful for such task. Then, the discovered association rules will be utilized to establish a function which determines the size of a trade.

9.5.2 *Identifying favorable markets conditions*

The experimental results reported in Chapters 6 and 7 showed that the performances of TSFDC and DBA can vary substantially from one currency pair to another. Knowing markets' characteristics for which TSFDC and DBA perform best is an interesting topic. This can be achieved by applying the DC-based market profiling approach introduced by Tsang et al., [74]. They proposed a set of DC-based indicators that aim to characterize price dynamics (e.g. volatility, fluctuation, and maximum possible returns over the specified period) of a given market. They suggested that the proposed indicators can help a trader to decide which market to trade in (e.g. normal market condition, stress market condition).

The performances of TSFDC and DBA was tested using a rolling window approach (Section 6.5.1). In this context, we can consider the training period of a rolling window as the profiling period (i.e. we compute the profiling indictors based on the dataset of training period of each rolling window). We then measure selected evaluation metrics (e.g. rate of returns RR, maximum drawdown MDD) of trading with TSFDC and DBA during the associated trading period of the same window. Table 9.1, shown below, illustrates our idea. The columns 'TMV', 'R', and 'T' are profiling indicators identified in Tsang et al., [74]. They will be utilized to characterize a given market during the training period of a rolling window. Whereas, the columns 'RR' and 'MDD' symbolize the performance of TSFDC, or DBA, during the corresponding trading period.

The objective is to find a relation between these profiling indicators and the selected evaluation metrics. The establishment of such relationship will be useful to anticipate the performance of TSFDC and DBA during the trading period. The examination of such relationship could be done using many machine learning algorithms. For example, if we consider *RR* as a set of qualitative elements (e.g. 'profitable', 'unprofitable') then the problem of finding such relation becomes a classification problem which can be solved using algorithms such as C5.0 and J48graft. On the other hand, if we measures *RR* as numbers (e.g. 2.1%, –1.5%), then we may use other algorithms such as M5P to examine that relation. In both ways, we will be able to decide whether a specific market is 'favorable' or 'unfavorable' for trading with TSFDC or DBA. Such market classification will allow us to allocate our capital more efficiently.

Table 9.1. An illustration of profiling indicators (computed based on a training period) and evaluation metrics (computed based on the associated trading period) of the same rolling window. The column named '….' symbolises other profiling indicators presented in the study of Tsang et al. [74].

Market profiling during training period				**Evaluation of TSFDC and DBA during trading period**	
TMV	**R**	**T**	**….**	**RR**	**MDD**

References

[1] Monetary and Economic Department, "Foreign exchange turnover in April 2016," Triennial Central Bank Survey, Bank for International Settlements BIS, 2016.

[2] M. R. King, C. Osler and D. Rime, "Foreign Exchange Market Structure, Players And Evolution," Working Paper 10/2011, Norges Bank, Norway, 2011.

[3] N. Zarrabia, S. Snaithb and J. Coakleya, "FX technical trading rules can be profitable sometimes!," *International Review of Financial Analysis,* vol. 49, pp. 113–127, 2017.

[4] B. J. Vanstone, T. Hahn and G. Finnie, "Developing High-Frequency Foreign Exchange Trading Systems," in *25th Australasian Finance and Banking Conference*, Sydney, 2012. Available at SSRN: https://ssrn.com/abstract=2132390 or http://dx.doi.org/10.2139/ssrn.2132390

[5] A. Mahdi, A. Hussain and D. Al-Jumeily, "Adaptive neural network model using the immune system for financial time series forecasting," in *Computational Intelligence, Modelling and Simulation, 2009. CSSim '09. International Conference on*, 2009. https://doi.org/10.1109/CSSim.2009.57.

[6] C. Evans, K. Pappasa and F. Xhafab, "Utilizing artificial neural networks and genetic algorithms to build an algotrading model for intra-day foreign exchange speculation," *Mathematical and Computer Modelling,* vol. 58, no. 5-6, pp. 1249-1266, 2013.

[7] L. Menkhoff, L. Sarno, M. Schmeling and A. Schrimpf, "Currency Momentum Strategies," Working Papers No 366, Monetary and Economic Department, BIS, 2011.

[8] L. Bertolini, "Trading Foreign Exchange Carry Portfolios," Ph.D. thesis, Cass Business School, City University, London, 2011.

[9] D. Guillaume, M. Dacorogna, R. Davé, U. Müller, R. Olsen and O. Pictet, "From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets," *Finance and stochastic,* vol. 1, no. 2, pp. 95-129, 1997.

[10] H. Ao and E. Tsang, “Capturing Market Movements with Directional,” Working Paper WP069-13, Centre for Computational Finance & Economic Agents (CCFEA), University of Essex, Colchester, 2013.

[11] M. Aloud, M. Fasli, E. Tsang, A. Dupuis and R. Olsen, “Stylized Facts of the FX Market Transactions Data: An Empirical Study,” *Journal of Finance and Investment Analysis,* vol. 2, pp. 145-183, 2013.

[12] J. Glattfelder, A. Dupuis and R. Olsen, “Patterns in high-frequency FX data: Discovery of 12 empirical scaling laws,” *Quantitative Finance,* vol. 11, no. 4, pp. 599-614, 2011.

[13] S. Masry, “Event-Based Microscopic Analysis of the FX Market,” Ph.D. thesis, Centre for Computational Finance & Economic Agents (CCFEA), University of Essex, Colchester, 2013.

[14] M. Aloud, E. K. Tsang and R. Olsen, “Modelling the FX market traders’ behaviour: an agent-based approach,” in *Simulation in Computational Finance and Economics: Tools and Emerging Applications,* B. Alexandrova-Kabadjova, S. Martinez-Jaramillo, A. Garcia-Almanza and E. Tsang, Eds., Hershey, PA, IGI Global, pp. 303-338, 2012.

[15] M. E. Aloud, “Directional-Change Event Trading Strategy: Profit-Maximizing Learning Strategy,” In: COGNITIVE 2015: *The Seventh International Conference on Advanced Cognitive Technologies and Applications.* IARIA. pp.123-129.

[16] A. Golub, J. B. Glattfelder and R. B. Olsen, “The Alpha Engine: Designing an Automated Trading Algorithm,” in *High-Performance Computing in Finance: Problems, Methods, and Solutions,* E. Vynckier, J. Kanniainen, J. Keane and M. A. H. Dempster, Eds., CRC Press, 2017. SSRN: https://ssrn.com/abstract=2951348.

[17] M. Kampouridis and F. E. Otero, “Evolving trading strategies using directional changes,” *Expert Systems With Applications,* vol. 73, pp. 145-160, 2017.

[18] A. Dupuis and R. Olsen, “High Frequency Finance: Using Scaling Laws to Build Trading Models,” in *Handbook of Exchange Rates,* J. James et al, Ed., NJ,USA, Wiley, 2012, pp. 563-582.

[19] M. Aloud, E. P. K. Tsang, R. Olsen and A. Dupuis, “A Directional-Change Events Approach for Studying Financial Time Series,” *Economics Papers, Economics: The Open-Access, Open-Assessment E-Journal,* 6 (2012-36): 1–17. http://dx.doi.org/10.5018/economics-ejournal.ja.2012-36

[20] T. C. Mills, Time Series Techniques for Economists, Cambridge University Press, UK, 1990.

[21] R. C. Cavalcante, R. C. Brasileiro, V. L. Souza, J. P. Nobrega and A. L. Oliveira, “Computational intelligence and financial markets: A survey and future directions,” *Expert Systems with Applications,* vol. 55, pp. 194-211, 2016.

[22] O. Kodongoa and K. Ojah, “The dynamic relation between foreign exchange rates and international portfolio flows: Evidence from Africa's capital markets,” *International Review of Economics & Finance,* vol. 24, pp. 71-87, 2012.

[23] O. Kodongoa and K. Ojah, “Real exchange rates, trade balance and capital flows in Africa,” *Journal of Economics and Business,* vol. 66, pp. 22-46, 2013.

[24] M. Frenkela, C. Pierdzioch and G. Stadtmann, “The effects of Japanese foreign exchange market interventions on the yen/U.S. dollar exchange rate volatility,” *International Review of Economics & Finance,* vol. 14, no. 1, pp. 27-39, 2005.

[25] G. Galati, W. Melick and M. Micu, “Foreign exchange market intervention and expectations: The yen/dollar exchange rate,” *Journal of International Money and Finance,* vol. 24, no. 6, pp. 982-1011, 2005.

[26] T. Hoshikawa, “The effect of intervention frequency on the foreign exchange market: The Japanese experience,” *Journal of International Money and Finance,* vol. 27, no. 4, pp. 547-559, 2008.

[27] A. Scalia, “Is foreign exchange intervention effective? Some microanalytical evidence from the Czech Republic,” *Journal of International Money and Finance,* vol. 27, no. 4, pp. 529-546, 2008.

[28] R. Fatum, “Daily effects of foreign exchange intervention: Evidence from official Bank of Canada data,” *Journal of International Money and Finance,* vol. 27, no. 3, pp. 438-454, 2008.

[29] T. Ito and Y. Hashimoto, “Intraday seasonality in activities of the foreign exchange markets: Evidence from the electronic broking system,” *Journal of the Japanese and International Economies,* vol. 20, pp. 637-664, 2006.

[30] J. Cotter and K. Dowd, “Intra-day seasonality in foreign exchange market transactions,” *International Review of Economics & Finance,* vol. 19, no. 2, pp. 287-294, 2010.

[31] J. James, I. W. Marsh and L. Sarno, Handbook of Exchange Rates, New Jersey: John Wiley & Sons, Inc., 2012.

[32] T. Chande and S. Kroll, The new technical trader, New York, USA: John Wiley & Sons, 1994.

[33] T. S. Chande, Beyond technical analysis: How to develop and implement a winning trading system (2nd ed), John Wiley & Sons, 1999.

[34] C. J. Neely and P. A. Weller, “Lessons from the Evolution of Foreign Exchange Trading Strategies,” *Journal of Banking & Finance,* vol. 37, pp. 3783–3798, 2013.

[35] J. Coakley, M. Marzano and J. Nankervis, “How profitable are FX technical trading rules?,” *International Review of Financial Analysis,* vol. 45, pp. 273-282, 2016.

[36] P.-H. Hsu, M. P. Taylor and Z. Wang, “Technical trading: Is it still beating the foreign exchange market?,” *Journal of International Economics,* vol. 102, pp. 188-208, 2016.

[37] M. Qi and Y. Wu, “Technical trading-rule profitability, data snooping, and reality check: Evidence from the foreign exchange market,” *Journal of Money, Credit and Banking,* vol. 38, pp. 2135-2158, 2006.

[38] C. J. Davison, “The Retail FX Trader: Random Trading and the Negative Sum Game,” 5 January 2016. [Online]. Available: http://ssrn.com/abstract=2711214. [Accessed 5 August 2017].

[39] F. Magnates, “Finance Magnates,” Finance Magnates, 10 August 2017. [Online]. Available: http://www.financemagnates.com/forex/. [Accessed 10 August 2017].

[40] R. Z. Heimer and D. Simon, “Facebook Finance: How Social Interaction Propagates Active Investing,” working paper no. 15-22, Federal Reserve Bank of Cleveland, Cleveland, 2015.

[41] X. Li, Z. Deng and J. Luo, “Trading strategy design in financial investment through a turning points prediction scheme,” *Expert System With Application,* vol. 36, pp. 7818-7826, 2009.

[42] C.-F. Huang, “A hybrid stock selection model using genetic algorithms and support vector regression,” *Applied Soft Computing,* vol. 12, no. 2, pp. 807-818, 2012.

[43] F. Giacomel, R. Galante and A. Pereira, “An Algorithmic Trading Agent based on a Neural Network Ensemble: a Case of Study in North American and Brazilian Stock Markets,” in *IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology,* IEEE Computer Society, vol. 02, pp. 230-233, 2015, doi:10.1109/WI-IAT.2015.43, 2015.

[44] K. Chourmouziadis and P. D. Chatzoglou, “An intelligent short term stock trading fuzzy system for assisting investors in portfolio management,” *Expert Systems with Applications,* vol. 43, pp. 298-311, 2016.

[45] T. L. Chen and F. Y. Chen, “An intelligent pattern recognition model for supporting investment decisions in stock market,” *International Journal of Information Sciences,* vol. 346, pp. 261-274, 2016.

[46] M. Göçkena, M. Özçalıcıb, A. Borua and A. T. Dosdo˘gru, “Integrating metaheuristics and Artificial Neural Networks for improved stock price prediction,” *Expert Systems with Applications,* vol. 44, p. 320–331, 2016.

[47] F. Bahramy and S. F. Crone, "Forecasting foreign exchange rates using Support Vector Regression," in *IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFEr),* Singapore, 2013. https://doi.org/10.1109/CIFEr.2013.6611694.

[48] L. Szu-Yin, C. Chi-Hua and L. Chi-Chun, "Currency Exchange Rates Prediction based on Linear Regression Analysis Using Cloud Computing," *International Journal of Grid and Distributed Computing,* vol. 6, no. 2, pp. 1-9, 2013.

[49] V. Pacelli, "Forecasting Exchange Rates: a Comparative Analysis," *International Journal of Business and Social Science,* vol. 3, no. 10, pp. 145-156, 2012.

[50] J. L. Ticknor, "A Bayesian regularized artificial neural network for stock market forecasting," *Expert Systems with Applications,* vol. 40, no. 14, pp. 5501-5506, 2013.

[51] B. Aluko, D. Smonou, M. Kampouridis and E. Tsang, "Combining different meta-heuristics to improve the predictability of a financial forecasting algorithm," in *IEEE Computational Intelligence for Financial Engineering & Finance (CIFEr),* London, UK, 2014. https://doi.org/10.1109/CIFEr.2014.6924092.

[52] J. Stanković, I. Marković and M. Stojanović, "Investment Strategy Optimization Using Technical Analysis and Predictive Modeling in Emerging Markets," *Procedia Economics and Finance,* vol. 19, pp. 51-62, 2015.

[53] S. Schulmeister, "Profitability of technical stock trading: Has it moved from daily to intraday data?," *Review of Financial Economics,* vol. 18, pp. 190-201, 2009.

[54] A. Boris S. and J. A. Doukas, "Is Technical Analysis Profitable for Individual Currency Traders?," *Journal of Portfolio Management,* vol. 39, pp. 142-150, 2012.

[55] U. Bonenkamp, C. Homburg and A. Kempf, "Fundamental Information in Technical Trading Strategies," *journal of Business Finance & Accounting,* vol. 38, no. 7-8, pp. 842-860, 2011.

[56] T. D. Watson, "On the profitability of technical trading," Ph.D. thesis, Department of Economics and Finance, Durham Business School, Durham university, Durham, 2009.

[57] R. Cervelló-Royo, F. Guijarro and K. Michniuk, “Stock market trading rule based on pattern recognition and technical analysis: Forecasting the DJIA index with intraday data,” *Expert Systems with Applications,* vol. 4, no. 14, pp. 5963-5975, 2015.

[58] C. R. M. Daryl, L. K. J. Shawn and C. H. Y. Sabrina, “Return and Risk-Return Ratio Based Momentum Strategies: A Fresh Perspective,” *Journal of Finance and Investment Analysis,* vol. 2, no. 1, pp. 1-13, 2013.

[59] M. T. Pojarliev and R. M. Levich., “Do professional currency managers beat the benchmark?,” *Financial Analysts Journal,* vol. 64, no. 5, pp. 18-32, 2008.

[60] J. Laborda, R. Laborda and J. Olmo, “Optimal currency carry trade strategies,” *International Review of Economics and Finance,* vol. 33, pp. 52-66, 2014.

[61] R. Pardo, The evalution and optimization of trading strategy, NJ, USA: John Wiley & Sons, 2011.

[62] I. Aldridge, High-Frequency Trading, New Jersey, USA.: John Wiley & Sons, Wiley Trading series, 2013.

[63] W. F. Sharpe, “Asset allocation: Management Style and Performance Measurement,” *Journal of Portfolio Management,* vol. Winter, pp. 7-19, 1992.

[64] J. Prokop, “Further evidence on the role of ratio choice in hedge fund performance evaluation,” *Journal of Finance and Investment Analysis,* vol. 1, no. 3, pp. 181–195, 2012.

[65] B. R. Auer and F. Schuhmacher, “Robust evidence on the similarity of sharpe ratio and drawdown-based hedge fund performance rankings,” *Journal of international financial markets, institutions and money,* vol. 24, pp. 153-165, 2013.

[66] J. L. Treynor, “How to rate management of investment funds,” *Harvard Business Review,* vol. 43, no. 1, pp. 63-75, 1965.

[67] F. Sortino and R. van der Meer, “Downside risk,” *Journal of Portfolio Management,* vol. 17, no. 4, pp. 27-31, 1991.

[68] W. F. Sharpe, “Capital asset prices: a theory of market equilibrium under conditions of risk,” *Journal of Finance,* vol. 19, no. 3, pp. 425-442, 1964.

[69] D. Ruppert, Statistics and Finance, an Introduction, NEW YORK,USA, Springer Sciences + Business Media, 2004.

[70] M. Jensen, “The Performance of Mutual Funds in the Period 1945-1964,” *Journal of Finance* , vol. 23, no. 2, pp. 389-416, 1967.

[71] P. J. Brockwell and R. A. Davis, Time series: Theory and methods (2nd ed), Springer-Verlag, NEW YORK, USA, 1991.

[72] M. M. Dacorogna, C. Gen, R. Muller, R. B. Olsen and O. V. Pictet, An introduction to high-frequency finance, San Diego, Ca, USA, Academic Press, 2001.

[73] B. B. Mandelbrot and H. W. Taylor, “On the Distribution of Stock Price Differences,” *Operations Research,* vol. 15, pp. 1057-1062, 1967.

[74] E. P. K. Tsang, R. Tao, A. Serguieva and S. Ma, “Profiling Financial Market Dynamics under Directional Changes,” *Quantitative finance,* vol. 17, no. 2, pp. 217-225, 2017.

[75] T. Bisig, A. Dupuis, V. Impagliazzo and R. Olsen, “The scale of market quake,” *Quantitative Finance,* vol. 12, no. 4, pp. 501-508, 2012.

[76] A. Golub, G. Chliamovitch, A. Dupuis and B. Chopard, “Multiscale representation of high frequency market liquidity,” *Algorithmic Finance,* vol. 5, no. 1, pp. 3-19, 2016.

[77] M. Aloud and M. Fasli, “Exploring Trading Strategies and their Effects in the Foreign Exchange Market,” *Computational Intelligence,* vol. 33, no. 2, pp. 280-307, 2017.

[78] J. Gypteau, F. Otero and M. Kampouridis, “Directional Changes based trading strategy with genetci programming,” in: Mora, A. M. and Squillero, G. eds. EvoApplications, EvoStar 2015. Springer, pp. 1-12.

[79] A. V. Golub, “www.github.com,” Lykke Corp, 1 April 2017. [Online]. Available: https://github.com/AntonVonGolub/Code/blob/master/code.java. [Accessed 20 August 2017].

[80] S.-H. Park, J.-H. Lee, J.-W. Song and T.-S. Park, “Forecasting Change Directions for Financial Time Series Using Hidden Markov Model,” in *Rough Sets and Knowledge Technology*, vol. 5589, Springer Berlin Heidelberg, 2009, pp. 184-191.

[81] A. Skabar, “Direction-of-Change Financial Time Series Forecasting using Bayesian Learning for MLPs,” in *Proceedings of the World Congress on Engineering*, London, UK. vol.2. International Association of Engineers, 2008, pp. 1160-1165.

[82] A. Sklarew, Techniques of a Professional Commodity Chart Analyst, New York, USA: Commodity Research Bureau, 1980.

[83] A. Azzini, C. d. C. Pereira and A. G. Tettamanzi, “Predicting Turning Points in Financial Markets with Fuzzy-Evolutionary and Neuro-Evolutionary Modeling,” in *Applications of Evolutionary Computing, Lecture Notes in Computer Science, vol 5484.*, Berlin, Heidelberg, Springer-Verlag, 2009, pp. 213-222.

[84] S. Raftopoulos, “The Zigzag Trend Indicator,” *Technical Analysis of Stocks and Commodities*, vol. 21, no. 11, pp. 26–32, 2003.

[85] S. Masry and E. Tsang, “Simulating market clearance dynamics under a simple event calculus market model,” in *3rd Computer Science and Electronic Engineering Conference*, Colchester, UK, 2011. https://doi.org/10.1109/CEEC.2011.5995835.

[86] B. Vanstone and G. Finnie, “An empirical methodology for developing stock market trading systems using artificial neural networks.,” *Expert Systems with Applications,* vol. 36, no. 3, pp. 6668-6680, 2009.

[87] A. Skabar, “Direction-of-Change Financial Time Series Forecasting using a Similarity-Based Classification Model,” *Journal of Forecasting,* vol. 32, no. 5, pp. 409-422, 2013.

[88] R. El-Yaniv and A. Faynburd, "Autoregressive short-term prediction of turning points using support vector regression," 24 Sep 2012. [Online]. Available: https://arxiv.org/pdf/1209.0127v2. [Accessed 26 Nov 2016].

[89] J. Ehlers, MESA and Trading Market Cycles: Forecasting and Trading Strategies from the Creator of MESA, New York, USA: John Wiley & Sons, 2002.

[90] J. Wilder, New Concepts in Technical Trading Systems, N.C., USA: Hunter Publishing Company, 1978.

[91] J. Quinlan, C4.5: Programs for Machine Learning, San Francisco, CA,USA: Morgan Kaufmann Publishers Inc., 1993.

[92] K. McCarthy, B. Zabar and G. Weiss, "Does cost-sensitive learning beat sampling for classifying rare classes?," in *Proceedings of the 1st international workshop on Utility-based data mining*, ACM, New York, USA, pp. 69-77, 2005.

[93] D. Asteriou and S. G. Hall, ARIMA Models and the Box–Jenkins Methodology, in: *Applied Econometric*s, *A Modern Approach*, Palgrave McMillan, New york, USA, 2007, pp. 229–247.

[94] R. de A. Araújo, A. L. Oliveira and S. Meira, "A hybrid model for high-frequency stock market forecasting," *Expert Systems with Applications*, vol. 42, pp. 4081-4096, 2015.

[95] J. Grant, "Trading Strategies in Futures Markets," Ph.D. thesis; Business School, Imperial College London, London, 2015.

[96] M. Kearns, A. Kulesza, and Y. Nevmyvaka, "Empirical limitations on High-Frequency Trading Profitability," *Journal of Trading*, vol. 5, no. 4, pp. 50–62, 2010. Available at SSRN: https://ssrn.com/abstract=1678758 or http://dx.doi.org/10.2139/ssrn.1678758.

[97] E. de Faria, M. P. Albuquerque, J. Gonzalez, J. Cavalcante and M. P. Albuquerque, "Predicting the Brazilian stock market through neural networks and adaptive exponential smoothing methods," *Expert Systems with Applications*, vol. 36, no. 10, pp. 12506–12509, 2009.

[98] B. Krollner, B. Vanstone and G. Finnie, “Financial time series forecasting with machine learning techniques: A survey,” in *European Symposium on Artificial Neural Networks: Computational and Machine Learning*, Bruges, Belgium, pp. 25-30, 2010.

[99] V. Vella and W. L. Ng, “A Dynamic Fuzzy Money Management Approach for Controlling the Intraday Risk-Adjusted Performance of AI Trading Algorithms,” *Intelligent Systems in Accounting, Finance and Management,* vol. 22, no. 2, pp. 153-178, 2015.

[100] P.H. Hsu and M. P. Taylor, “Forty Years, Thirty Currencies and 21,000 Trading Rules: A Large-Scale, Data-Snooping Robust Analysis of Technical Trading in the Foreign Exchange Market,” 5 May 2013. [Online]. Available: Available at SSRN: https://ssrn.com/abstract=2517125. [Accessed 30 June 2017].

[101] S. Bohn, “The slippage paradox,” 14 March 2011. [Online]. Available: http://arxiv.org/pdf/1103.2214. [Accessed 22 June 2017].

[102] F. Wilcoxon, “ Individual comparisons by ranking methods,” *Biometrics Bulletin* , vol. 1, pp. 80-83, 1945.

[103] H. Markowitz, “Portfolio Selection,” *The Journal of Finance,* vol. 7, no. 1, pp. 77-91, 1952.

[104] G. Piatetsky-Shapiro, “Discovery, analysis, and presentation of strong rules,” in *Knowledge Discovery in Databases*, G. Piatetsky-Shapiro and W. J. Frawley, Eds., Cambridge, MA., AAAI/MIT Press, 1991.

[105] R. Agrawal and R. Srikant, “Fast algorithms for mining association rules,” in *Proceeding 20th International Conference on Very Large Data Bases, VLDB,* Santiago, Chile, September 1994. Morgan Kaufmann, CA, USA, pp. 487-499. 1994

[106] G. I. Webb, “OPUS: An Efficient Admissible Algorithm for Unordered Search,” *Journal of Artificial Intelligence Research,* vol. 3, pp. 431-465, 1995.

Appendix A: R-Code to Detect Directional Changes

In this appendix, we provide the R code, named ‘*DCSummary’,* which produces the DC summary of a particular prices series, given a threshold *theta*, as explained previously in Section 4.2. The texts placed after the ‘#’ are comments and not part of the code.

DCSummary.r

 # In the code below, the variable ‘prices’ denotes the vector of price series. The codes of loading prices from given file is irrelevant and, therefore, omitted.

```
l = length(prices) # ‘prices’ denotes the vector of prices. l denotes the number of prices’ #observation in the prices series.
x_ext_index=1
while (i< l)
{
    if (mode <  1)# mode is downtrend
    {
        if (prices [i]< x_ext)
            {
            x_ext = prices [i]
            x_ext_index = i
            is_double_ext = 0
            }
            else if (((prices [i]– x_ext)/x_ext)>= theta)
            {
            nb_up=nb_up + 1
                    if (is_double_ext < 1)
                    {Event[x_ext_index] = "(start EXT UP)"
                    }
                    else
                    {Event[x_ext_index] = "(OS DOWN & start EXT UP)"
                    }
            Event[i] = "(start OS UP)"
            OS_up_OS_down_indicator[i] = 1
            x_os_index = i
            DCC = prices [x_ext_index]*(1+ theta)
            DCCs[i] = DCC
            OSV_OS[i]= (( prices [i]– DCC)/DCC)/theta
            x_ext_index = i
            x_ext= prices [i]
            mode =1
```

```
                    is_double_ext =1
                    }
        }
        else if( mode > 0) # mode is uptrend
        {
                    if (prices [i] >x_ext)
                    {
                    x_ext = prices [i]
                    x_ext_index = i
                    is_double_ext = 0
                    }
                    else if (((prices [i]– x_ext)/x_ext) <= – theta)
                    {
                    nb_down=nb_down+ 1
                                    if (is_double_ext < 1)
                                    {Event[x_ext_index]= "(start EXT DOWN )"
                                    }
                                    else
                                    {Event[x_ext_index] = "(OS UP & start EXT DOWN)"
                                    OS_up_OS_down_indicator[x_ext_index] = 1
                                    }
                    Event[i] = "(start OS DOWN)"
                    downTrendID = downTrendID + 1
                    DCC = prices [x_ext_index] * (1– theta)
                    trace_DCC = DCC
                    DCCs[i] = DCC
                    OSV_OS[i] = (( prices [i]– DCC)/DCC)/theta
                    x_os_index = i
                    x_ext_index = i
                    x_ext = prices [i]
                    is_double_ext =1
                    mode = 0
                    }
        }
i = i+ 1 # proceed with the next price’s observation
}
DCSummary = data.frame(prices, EventType=Event, DCC_Prices=DCCs, OSV=OSV_OS)
```

At the end of the above R code, the *dataframe* named *DCSummary* will comprise four vectors:

- ‘*prices*’: this is the initial price series,
- ‘*EventType*’: it comprises all DC and OS events detected
- ‘*DCC_Prices*’: it denote the price required to confirm the detection of a new DC event of the specified threshold *theta*.
- ‘*OSV*’: the overshoot values computed at the DCC point of each DC event.

The *DCC* prices and the *OSV* are computed at the *DCC* point of each DC event. See Table 4.1 in Chapter 4 for more details regarding the theory of Directional Changes’ summary.

Appendix B: The Impact of *BTheta* on the Accuracy of our Forecasting Model

This appendix lists the results of Experiment 5.2 '*The Impact of BTheta on the Accuracy of our Forecasting Model*' (presented in Section 5.6.2) for the remaining four currency pairs: GBP/JPY, NZD/JPY, AUD/JPY, and EUR/NZD. *STheta* is fixed to 0.10%. The reported accuracy corresponding to the testing (out-of-sample) period. For each of these currency pairs, the testing period is 7 months.

For each of these tables, we apply the linear regression model to examine the impact of *BTheta* on the accuracy of our approach. The resulted p-values for all cases are always above the common level of 0.05. This indicates that *BTheta* has significant impact on the accuracy of our approach. We also note that the accuracy of our approach is quite good for most levels of True-False imbalance (α). In each table, the accuracies range between 0.62 and 0.82; which conforms to the conclusion reported in Section 5.6.2.

Table B.1: Analyzing the impact of *BTheta* on the accuracy of our forecasting approach. The case of GBP/JPY.

Uptrends of DC summary with *STheta* = 0.10%				Downtrends of DC summary with *STheta* = 0.10%		
BTheta (%)	**Accuracy**	**α**		**BTheta (%)**	**Accuracy**	**α**
0.13	0.81	0.64		0.13	0.82	0.64
0.14	0.77	0.55		0.14	0.76	0.55
0.15	0.73	0.49		0.15	0.73	0.49
0.16	0.71	0.43		0.16	0.71	0.43
0.17	0.68	0.38		0.17	0.69	0.38
0.18	0.66	0.34		0.18	0.67	0.34
0.19	0.64	0.31		0.19	0.64	0.31
0.20	0.63	0.28		0.20	0.62	0.28
0.21	0.62	0.26		0.21	0.61	0.26
0.22	0.61	0.23		0.22	0.60	0.23

Table B.2: Analyzing the impact of *BTheta* on the accuracy of our forecasting approach. The case of NZD/JPY.

Uptrends of DC summary with *STheta* = 0.10%				Downtrends of DC summary with *STheta* = 0.10%		
BTheta (%)	**Accuracy**	**α**		**BTheta (%)**	**Accuracy**	**α**
0.13	0.82	0.63		0.13	0.82	0.63
0.14	0.78	0.54		0.14	0.78	0.54
0.15	0.74	0.48		0.15	0.74	0.48
0.16	0.72	0.42		0.16	0.72	0.42
0.17	0.70	0.37		0.17	0.70	0.37
0.18	0.67	0.33		0.18	0.67	0.33
0.19	0.65	0.30		0.19	0.65	0.30
0.20	0.64	0.27		0.20	0.64	0.27
0.21	0.63	0.25		0.21	0.63	0.25
0.22	0.62	0.22		0.22	0.62	0.22

Table B.3: Analyzing the impact of value of *BTheta* to the accuracy of our forecasting approach. The case of AUD/JPY.

Uptrends of DC summary with STheta = 0.10%				Downtrends of DC summary with STheta = 0.10%		
BTheta (%)	**Accuracy**	**α**		**BTheta (%)**	**Accuracy**	**α**
0.13	0.79	0.56		0.13	0.79	0.56
0.14	0.78	0.53		0.14	0.77	0.53
0.15	0.75	0.51		0.15	0.76	0.51
0.16	0.70	0.49		0.16	0.70	0.49
0.17	0.68	0.48		0.17	0.69	0.48
0.18	0.66	0.45		0.18	0.66	0.45
0.19	0.65	0.42		0.19	0.66	0.42
0.20	0.64	0.35		0.20	0.64	0.35
0.21	0.64	0.31		0.21	0.65	0.31
0.22	0.63	0.28		0.22	0.63	0.28

Table B.4: Analyzing the impact of *BTheta* on the accuracy of our forecasting approach. The case of EUR/NZD.

Uptrends of DC summary with STheta = 0.10%				Downtrends of DC summary with STheta = 0.10%		
BTheta (%)	**Accuracy**	**α**		**BTheta (%)**	**Accuracy**	**α**
0.13	0.82	0.63		0.13	0.82	0.63
0.14	0.78	0.55		0.14	0.78	0.55
0.15	0.74	0.50		0.15	0.74	0.50
0.16	0.72	0.47		0.16	0.72	0.47
0.17	0.70	0.40		0.17	0.70	0.40
0.18	0.67	0.39		0.18	0.67	0.39
0.19	0.65	0.33		0.19	0.65	0.33
0.20	0.64	0.29		0.20	0.64	0.29
0.21	0.63	0.27		0.21	0.63	0.27
0.22	0.62	0.25		0.22	0.62	0.25

Appendix C: Comparing the Return and Risk of TSFDC-down and TSFDC-up

In Experiment 6.2 (Section 6.6.2), we aimed to test whether the TSFDC-down and TSFDC-up has different profitability and risk. We consider the monthly rate of returns (*RR*) as indicator of profitability. The risk is measured as *MDD*. In the following table we summarize the results of monthly *RR* and *MDD* obtained by applying TSFDC-down and TSFDC-up to the eight currency pairs based on the experiments organized in Section 6.6.2.

As can be noted in Table D.1, shown below, we have two sets of monthly *RR* shown under the column named ‘RR’: one for TSFDC-down and the other is for TSFDC-up. Each set encompasses 56 observations. We apply the non-parametric Wilcoxon test with the null hypothesis being that the median difference between these two sets is zero. Similarly, Table D.1 identifies two sets of monthly *MDD* shown under the column named ‘MDD’: one for TSFDC-down and the other is for TSFDC-up. Each set encompasses 56 observations. We apply the non-parametric Wilcoxon test with the null hypothesis being that the median difference between these two sets is zero.

Table D.1: summary of monthly rate of returns (RR) and MDD of TSFDC-down and TSFDC-up based on Experiment 6.2

Observation number	Currency pairs	Trading Month	RR			MDD	
			TSFDC–down	TSFDC–up		TSFDC–down	TSFDC–up
1	EUR/CHF	**Jan**	4.47	4.26		– 13.4	– 15.1
2		**Feb**	14.40	9.75		– 1.4	– 2.5
3		**Mar**	17.59	16.87		– 0.6	– 3.4
4		**Apr**	7.58	5.71		– 0.7	– 2.8
5		**May**	13.37	7.61		– 0.7	– 1.5
6		**Jun**	12.41	10.15		– 1.4	– 3.3
7		**Jul**	14.77	8.68		– 0.6	– 1.8
8	GBP/CHF	**Jan**	13.59	31.54		– 12.1	– 10.8
9		**Feb**	19.02	16.30		– 2.7	– 3.9
10		**Mar**	14.96	21.67		– 2.9	– 3.8
11		**Apr**	6.71	12.34		– 2.5	– 4.0
12		**May**	9.85	7.59		– 2.9	– 3.1
13		**Jun**	15.17	14.13		– 3.7	– 4.1
14		**Jul**	14.73	11.62		– 2.2	– 2.8
15	EUR/USD	**Jan**	1.12	6.81		– 4.2	– 5.8
16		**Feb**	7.54	9.27		– 3.1	– 3.9
17		**Mar**	– 0.36	1.69		– 5.0	– 4.8
18		**Apr**	4.20	1.66		– 2.9	– 3.9
19		**May**	5.73	9.67		– 3.3	– 2.5
20		**Jun**	7.85	6.13		– 3.7	– 2.8
21		**Jul**	0.96	0.86		– 3.4	– 3.0

Table D.1 (*continued*): summary of monthly rate of returns (RR) and MDD of TSFDC-down and TSFDC-up based on Experiment 6.2

Observation number	**Currency pairs**	**Trading Month**	**RR**			**MDD**	
			DBA–down	**DBA–up**		**DBA–down**	**DBA–up**
22	**GBP/AUD**	**Jan**	19.70	13.34		– 2.83	– 1.36
23		**Feb**	10.51	10.06		– 3.18	– 3.52
24		**Mar**	10.14	9.09		– 1.53	– 1.56
25		**Apr**	13.52	9.23		– 1.14	– 2.39
26		**May**	15.97	9.51		– 0.84	– 1.39
27		**Jun**	11.52	5.97		– 1.25	– 1.29
28		**Jul**	11.27	5.83		– 3.35	– 1.91
29	**GBP/JPY**	**Jan**	7.72	11.39		– 4.8	– 4.2
30		**Feb**	6.40	3.64		– 3.8	– 3.2
31		**Mar**	4.04	6.00		– 2.8	– 5.7
32		**Apr**	7.05	3.07		– 4.7	– 2.9
33		**May**	8.38	4.11		– 3.5	– 1.9
34		**Jun**	0.99	4.16		– 4.1	– 3.0
35		**Jul**	– 2.10	– 3.46		– 3.1	– 3.7
36	**NZD/JPY**	**Jan**	19.14	26.96		– 2.6	– 4.0
37		**Feb**	26.90	18.06		– 3.2	– 3.0
38		**Mar**	19.95	24.06		– 4.9	– 2.2
39		**Apr**	30.41	22.98		– 2.8	– 2.9
40		**May**	24.27	24.92		– 3.1	– 2.4
41		**Jun**	17.20	32.66		– 2.6	– 3.0
42		**Jul**	45.26	41.09		– 3.1	– 2.2
43	**AUD/JPY**	**Jan**	15.36	21.48		– 5.0	– 2.3
44		**Feb**	16.47	14.88		– 3.2	– 2.3
45		**Mar**	10.51	17.30		– 2.9	– 4.2
46		**Apr**	16.69	12.25		– 2.8	– 5.2
47		**May**	25.51	21.15		– 2.1	– 2.6
48		**Jun**	10.48	17.32		– 3.6	– 3.5
49		**Jul**	9.09	11.97		– 3.1	– 3.1
50	**EUR/NZD**	**Jan**	24.12	26.27		– 1.2	– 5.1
51		**Feb**	50.04	68.74		– 4.6	– 2.7
52		**Mar**	49.76	64.56		– 2.1	– 3.9
53		**Apr**	59.39	78.72		– 2.8	– 1.9
54		**May**	79.92	82.81		– 3.0	– 2.9
55		**Jun**	104.91	101.88		– 2.8	– 2.9
56		**Jul**	120.99	148.91		– 2.9	– 2.6

Appendix D: The Monthly Evaluation of Performances of TSFDC-down and TSFDC-up

This appendix comprises the details of monthly evaluation of the performance of trading with TSFDC-down and TSFDC-up over the eight currency pairs reported in Section 6.6.1.

Table C.1: Monthly trading performance of the TSFDC-down model following the seven months out-of-sample period of *EURCHF_RWDC0.1*.

Applied Window	RR	Profit Factor	Total Number of Trades	MDD (%)	Win Ratio
Jan 2015	4.47	1.12	576	– 13.4	0.66
Feb 2015	14.40	2.37	351	– 1.4	0.73
Mar 2015	17.59	2.87	329	– 0.6	0.75
Apr 2015	7.58	1.85	191	– 0.7	0.71
May 2015	13.37	3.26	187	– 0.7	0.76
Jun 2015	12.41	1.94	245	– 1.4	0.77
Jul 2015	14.77	3.65	177	– 0.6	0.80
Sum	84.59	– –	2056	– –	– –

Table C.2: Monthly trading performance of the TSFDC-up model following the seven months out-of-sample period of *EURCHF_RWDC0.1*.

Applied Window	RR	Profit Factor	Total Number of Trades	MDD (%)	Win Ratio
Jan 2015	4.26	1.13	582	– 15.1	0.65
Feb 2015	9.75	2.01	353	– 2.5	0.69
Mar 2015	16.87	3.99	320	– 3.4	0.78
Apr 2015	5.71	1.73	191	– 2.8	0.73
May 2015	7.61	2.25	174	– 1.5	0.75
Jun 2015	10.15	2.00	247	– 3.3	0.68
Jul 2015	8.68	3.09	142	– 1.8	0.79
Sum	63.03	– –	2009	– –	– –

Table C.3: Monthly trading performance of the TSFDC-down model following the seven months out-of-sample period of *GBPCHF_RWDC0.1*.

Applied Window	RR	Profit Factor	Total Number of Trades	MDD (%)	Win Ratio
Jan 2015	13.59	1.30	697	– 12.1	0.69
Feb 2015	19.02	2.95	382	– 2.7	0.75
Mar 2015	14.96	2.06	359	– 2.9	0.72
Apr 2015	6.71	1.44	258	– 2.5	0.73
May 2015	9.85	1.58	264	– 2.9	0.70
Jun 2015	15.17	1.96	290	– 3.7	0.72
Jul 2015	14.73	2.24	239	– 2.2	0.72
Sum	94.03	– –	2489	– –	– –

Table C.4: Monthly trading performance of the TSFDC-up model following the seven months out-of-sample period of *GBPCHF_RWDC0.1*.

Applied Window	RR	Profit Factor	Total Number of Trades	MDD (%)	Win Ratio
Jan 2015	31.54	1.51	696	– 10.8	0.67
Feb 2015	16.30	1.95	397	– 3.9	0.72
Mar 2015	21.67	2.48	367	– 3.8	0.74
Apr 2015	12.34	1.78	258	– 4.0	0.72
May 2015	7.59	1.38	273	– 3.1	0.67
Jun 2015	14.13	1.74	298	– 4.1	0.72
Jul 2015	11.62	1.67	242	– 2.8	0.69
Sum	115.19	– –	2531	– –	– –

Table C.5: Monthly trading performance of the TSFDC-down model following the seven months out-of-sample period of *EURUSD_RWDC0.1*

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	1.12	1.08	200	– 4.2	0.66
Feb 2015	7.54	1.79	177	– 3.1	0.69
Mar 2015	– 0.36	0.98	234	– 5.0	0.64
Apr 2015	4.20	1.31	195	– 2.9	0.65
May 2015	5.73	1.37	210	– 3.3	0.63
Jun 2015	7.85	1.40	267	– 3.7	0.68
Jul 2015	0.96	1.07	148	– 3.4	0.62
Sum	27.04	– –	1431	– –	– –

Table C.6: Monthly trading performance of the TSFDC-up model following the seven months out-of-sample period of *EURUSD_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	6.81	1.6	211	– 5.8	0.68
Feb 2015	9.27	2.08	180	– 3.9	0.71
Mar 2015	1.69	1.08	233	– 4.8	0.63
Apr 2015	1.66	1.09	198	– 3.9	0.64
May 2015	9.67	1.66	233	– 2.5	0.70
Jun 2015	6.13	1.27	262	– 2.8	0.68
Jul 2015	0.86	1.06	136	– 3.0	0.63
Sum	36.09	– –	1453	– –	– –

Table C.7: Monthly trading performance of the TSFDC-down model following the seven months out-of-sample period of *GBPAUD_RWDC0.1*

Applied Window	RR	Profit Factor	Total Number of Trades	MDD (%)	Win Ratio
Jan 2015	19.70	2.29	442	– 2.83	0.72
Feb 2015	10.51	1.63	431	– 3.18	0.69
Mar 2015	10.14	1.59	436	– 1.53	0.68
Apr 2015	13.52	1.88	444	– 1.14	0.69
May 2015	15.97	2.12	446	– 0.84	0.72
Jun 2015	11.52	1.68	446	– 1.25	0.68
Jul 2015	11.27	1.90	376	– 3.35	0.74
Sum	92.63	– –	3021	– –	– –

Table C.8: Monthly trading performance of the TSFDC-up model following the seven months out-of-sample period of *GBPAUD_RWDC0.1*

Applied Window	RR	Profit Factor	Total Number of Trades	MDD (%)	Win Ratio
Jan 2015	13.34	1.83	445	– 1.36	0.71
Feb 2015	10.06	1.63	431	– 3.52	0.71
Mar 2015	9.09	1.56	426	– 1.56	0.67
Apr 2015	9.23	1.56	437	– 2.39	0.67
May 2015	9.51	1.54	435	– 1.39	0.69
Jun 2015	5.97	1.29	449	– 1.29	0.66
Jul 2015	5.83	1.43	337	– 1.91	0.68
Sum	63.03	– –	2960	– –	– –

Table C.9: Monthly trading performance of the TSFDC-up model following the seven months out-of-sample period of *GBPCHF_RWDC0.1*.

Applied Window	RR	Profit Factor	Total Number of Trades	MDD (%)	Win Ratio
Jan 2015	7.72	1.49	401	– 4.8	0.68
Feb 2015	6.40	1.86	212	– 3.8	0.72
Mar 2015	4.04	1.43	229	– 2.8	0.63
Apr 2015	7.05	2.0	203	– 4.7	0.74
May 2015	8.38	2.14	206	– 3.5	0.72
Jun 2015	0.99	1.51	179	– 4.1	0.68
Jul 2015	– 2.10	0.81	155	– 3.1	0.61
Sum	32.48	– –	1585	– –	– –

Table C.10: Monthly trading performance of the TSFDC-up model following the seven months out-of-sample period of *GBPCHF_RWDC0.1*.

Applied Window	RR	Profit Factor	Total Number of Trades	MDD (%)	Win Ratio
Jan 2015	11.39	1.98	389	– 4.2	0.69
Feb 2015	3.64	1.36	223	– 3.2	0.69
Mar 2015	6.00	1.73	223	– 5.7	0.72
Apr 2015	3.07	1.32	194	– 2.9	0.69
May 2015	4.11	1.41	213	– 1.9	0.71
Jun 2015	4.16	1.55	198	– 3.0	0.67
Jul 2015	– 3.46	0.7	161	– 3.7	0.60
Sum	28.91	– –	1601	– –	– –

Table C.11: Monthly trading performance of the TSFDC-down model following the seven months out-of-sample period of *NZDJPY_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	19.14	1.75	546	– 2.6	0.68
Feb 2015	26.90	2.87	443	– 3.2	0.79
Mar 2015	19.95	1.88	420	– 4.9	0.74
Apr 2015	30.41	2.93	380	– 2.8	0.78
May 2015	24.27	2.15	376	– 3.1	0.74
Jun 2015	17.20	1.56	358	– 2.6	0.70
Jul 2015	45.26	2.08	523	– 3.1	0.72
Sum	183.13	– –	3046	– –	– –

Table C.12: Monthly trading performance of the TSFDC-up model following the seven months out-of-sample period of *NZDJPY_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	26.96	2.28	553	– 4.0	0.71
Feb 2015	18.06	1.98	442	– 3.0	0.72
Mar 2015	24.06	2.23	431	– 2.2	0.76
Apr 2015	22.98	2.36	376	– 2.9	0.72
May 2015	24.92	225	364	– 2.4	0.74
Jun 2015	32.66	2.54	349	– 3.0	0.76
Jul 2015	41.09	2.00	495	– 2.2	0.73
Sum	190.73	– –	3010	– –	– –

Table C.13: Monthly trading performance of the TSFDC-down model following the seven months out-of-sample period of *AUDJPY_RWDC0.1*.

Applied Window	RR	Profit Factor	Total Number of Trades	MDD (%)	Win Ratio
Jan 2015	15.36	1.54	613	– 5.0	0.71
Feb 2015	16.47	1.99	434	– 3.2	0.73
Mar 2015	10.51	1.60	380	– 2.9	0.66
Apr 2015	16.69	2.11	361	– 2.8	0.71
May 2015	25.51	2.71	368	– 2.1	0.77
Jun 2015	10.48	1.45	318	– 3.6	0.69
Jul 2015	9.09	1.26	411	– 3.1	0.69
Sum	104.11	– –	2885	– –	– –

Table C.14: Monthly trading performance of the TSFDC-up model following the seven months out-of-sample period of *AUDJPY_RWDC0.1*.

Applied Window	RR	Profit Factor	Total Number of Trades	MDD (%)	Win Ratio
Jan 2015	21.48	1.96	606	– 2.3	0.72
Feb 2015	14.88	1.82	428	– 2.3	0.72
Mar 2015	17.30	2.15	385	– 4.2	0.75
Apr 2015	12.25	1.63	366	– 5.2	0.68
May 2015	21.15	2.38	355	– 2.6	0.71
Jun 2015	17.32	1.92	314	– 3.5	0.71
Jul 2015	11.97	1.35	406	– 3.1	0.69
Sum	116.35	– –	2860	– –	– –

Table C.15: Monthly trading performance of the TSFDC-down model following the seven months out-of-sample period of *EURNZD_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	24.12	2.41	614	– 1.2	0.74
Feb 2015	50.04	5.0	600	– 4.6	0.79
Mar 2015	49.76	3.76	570	– 2.1	0.75
Apr 2015	59.39	3.80	534	– 2.8	0.77
May 2015	79.92	3.29	540	– 3.0	0.77
Jun 2015	104.91	3.47	557	– 2.8	0.75
Jul 2015	120.99	2.85	546	– 2.9	0.76
Sum	489.13	– –	3961	– –	– –

Table C.16: Monthly trading performance of the TSFDC-up model following the seven months out-of-sample period of *EURNZD_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	26.27	2.47	639	– 5.1	0.75
Feb 2015	68.74	6.97	640	– 2.7	0.82
Mar 2015	64.56	3.49	602	– 3.9	0.80
Apr 2015	78.72	3.36	574	– 1.9	0.77
May 2015	82.81	3.04	580	– 2.9	0.74
Jun 2015	101.88	2.67	591	– 2.9	0.75
Jul 2015	148.91	3.15	592	– 2.6	0.74
Sum	571.89	– –	4218	– –	– –

Appendix E: Comparing the Return and Risk of DBA-down and DBA-up

In Experiment 7.3 (Section 7.4.3), we aimed to test whether the DBA-down and DBA-up do have different *RR* and *MDD*. In Table E.1, shown below, we summarize the results of monthly *RR* and *MDD* obtained by applying DBA-down and DBA-up to the eight currency pairs based on the experiments carried out in Section 7.4.2. As can be noted in Table E.1, we have two sets of monthly *RR* shown under the column named 'RR': one for DBA-down and the other is for DBA-up. Each set encompasses 56 observations. We apply the non-parametric Wilcoxon test with the null hypothesis being that the median difference between these two sets is zero. Similarly, Table E.1 identifies two sets of monthly *MDD* shown under the column named 'MDD': one for DBA-down and the other is for DBA-up. Each set encompasses 56 observations. We apply the non-parametric Wilcoxon test with the null hypothesis being that the median difference between these two sets is zero.

Table E.1: Summary of monthly RR and MDD of DBA-down and DBA-up based on Experiment 7.3 (Section 7.5.3)

Observation number	**Currency pairs**	**Trading Month (2015)**	**RR**			**MDD**	
			DBA–down	**DBA–up**		**DBA–down**	**DBA–up**
1	**EUR/CHF**	**Jan**	3.77	– 4.87		– 11.7	– 14.7
2		**Feb**	13.56	10.32		– 0.9	– 8.6
3		**Mar**	11.01	14.71		– 0.8	– 0.4
4		**Apr**	5.99	7.37		– 0.4	– 1.7
5		**May**	7.36	11.42		– 0.6	– 0.3
6		**Jun**	8.94	10.21		– 1.3	– 0.9
7		**Jul**	12.98	10.44		– 0.6	– 0.8
8	**GBP/CHF**	**Jan**	12.98	9.45		– 10.7	– 23.2
9		**Feb**	12.26	12.71		– 0.6	– 1.0
10		**Mar**	11.44	15.79		– 0.9	– 0.7
11		**Apr**	6.96	10.31		– 1.1	– 2.0
12		**May**	7.22	7.83		– 2.1	– 2.1
13		**Jun**	10.39	10.07		– 0.9	– 1.0
14		**Jul**	12.15	11.28		– 0.8	– 0.8

Table E.1 (*continued*): Summary of monthly profits and MDD of DBA-down and DBA-up based on Experiment 7.3

Observation number	Currency pairs	Trading Month	RR			MDD	
			DBA–down	DBA–up		DBA–down	DBA–up
15	EUR/USD	Jan	– 2.74	5.64		– 3.6	– 1.4
16		Feb	2.53	4.15		– 1.0	– 0.9
17		Mar	– 2.69	1.67		– 3.8	– 4.6
18		Apr	3.78	1.20		– 2.2	– 2.8
19		May	5.50	6.24		– 1.0	– 1.2
20		Jun	5.50	4.77		– 1.4	– 1.0
21		Jul	0.78	1.95		– 1.2	– 2.4
22	GBP/AUD	Jan	12.12	13.26		– 0.8	– 1.5
23		Feb	10.96	8.30		– 1.1	– 1.6
24		Mar	4.62	8.33		– 1.7	– 0.6
25		Apr	6.97	8.28		– 1.1	– 1.2
26		May	13.67	11.03		– 1.2	– 1.0
27		Jun	6.74	8.63		– 1.4	– 1.0
28		Jul	13.86	8.40		– 1.2	– 1.8
29	GBP/JPY	Jan	4.78	10.90		– 3.9	– 0.9
30		Feb	6.69	3.08		– 0.7	– 1.1
31		Mar	2.85	5.87		– 1.0	– 1.0
32		Apr	7.59	5.65		– 0.6	– 0.6
33		May	7.05	4.00		– 0.7	– 1.6
34		Jun	5.35	3.96		– 1.4	– 0.9
35		Jul	– 1.54	– 1.39		– 2.5	– 1.7
36	NZD/JPY	Jan	10.88	25.41		– 2.7	– 1.0
37		Feb	20.96	17.26		– 1.4	– 3.2
38		Mar	15.82	22.82		– 1.0	– 1.3
39		Apr	19.39	21.32		– 1.1	– 1.2
40		May	16.29	24.40		– 1.2	– 1.4
41		Jun	11.81	27.64		– 1.5	– 1.1
42		Jul	20.40	42.86		– 2.3	– 1.8
43	AUD/JPY	Jan	7.69	15.62		– 2.9	– 1.2
44		Feb	12.87	10.41		– 1.9	– 1.7
45		Mar	10.37	13.78		– 0.7	– 0.9
46		Apr	10.10	6.73		– 1.2	– 1.4
47		May	14.73	17.13		– 0.7	– 0.9
48		Jun	9.58	12.15		– 2.4	– 1.2
49		Jul	8.61	11.53		– 2.7	– 1.8
50	EUR/NZD	Jan	23.47	17.83		– 1.1	– 1.7
51		Feb	43.08	42.06		– 0.8	– 1.1
52		Mar	49.90	36.82		– 1.0	– 1.0
53		Apr	45.40	48.94		– 0.9	– 0.8
54		May	62.99	54.98		– 0.9	– 1.0
55		Jun	73.46	71.73		– 1.6	– 0.8
56		Jul	89.23	75.83		– 0.7	– 0.9

Appendix F: The Monthly Evaluation of Performances of DBA-down and DBA-up

This appendix comprises the details of monthly evaluation of the performance of trading with DBA-down and DBA-up over the eight currency pairs reported in Experiment 7.3 (Section 7.5.3).

Table F.1: Monthly trading performance of the DBA-down model following the seven months out-of-sample period of *EURCHF_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	3.77	1.09	516	– 11.7	0.67
Feb 2015	13.56	2.45	348	– 0.9	0.74
Mar 2015	11.01	2.12	311	– 0.8	0.69
Apr 2015	5.99	1.86	186	– 0.4	0.71
May 2015	7.36	2.03	192	– 0.6	0.73
Jun 2015	8.94	1.83	237	– 1.3	0.76
Jul 2015	12.98	2.78	218	– 0.6	0.78
Sum	63.61	– –	2008	– –	– –

Table F.2: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *EURCHF_RWDC0.1*.

Applied Window	**RR**	**Profit factor**	**Total number of trades**	**MDD (%)**	**Win Ratio**
Jan 2015	– 4.87	0.88	504	– 14.7	0.63
Feb 2015	10.32	1.95	372	– 8.6	0.68
Mar 2015	14.71	3.21	325	– 0.4	0.74
Apr 2015	7.37	1.91	209	– 1.7	0.76
May 2015	11.42	3.14	201	– 0.3	0.78
Jun 2015	10.21	1.87	265	– 0.9	0.73
Jul 2015	10.44	2.08	229	– 0.8	0.73
Sum	59.60	– –	2105	– –	– –

Table F.3: Monthly trading performance of the DBA-down model following the seven months out-of-sample period of *GBPCHF_RWDC0.1*.

Applied window	**RR**	**Profit factor**	**Total number of trades**	**MDD (%)**	**Win Ratio**
Jan 2015	12.98	1.41	641	– 10.7	0.68
Feb 2015	12.26	2.66	386	– 0.6	0.75
Mar 2015	11.44	2.02	359	– 0.9	0.73
Apr 2015	6.96	1.61	267	– 1.1	0.73
May 2015	7.22	1.34	279	– 2.1	0.69
Jun 2015	10.39	1.78	280	– 0.9	0.74
Jul 2015	12.15	2.12	274	– 0.8	0.70
Sum	73.39	– –	2486	– –	– –

Table F.4: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *GBPCHF_RWDC0.1*.

Applied window	**RR**	**Profit factor**	**Total number of trades**	**MDD (%)**	**Win Ratio**
Jan 2015	9.45	1.14	639	– 23.2	0.67
Feb 2015	12.71	1.86	405	– 1.0	0.72
Mar 2015	15.79	2.31	369	– 0.7	0.75
Apr 2015	10.31	1.74	280	– 2.0	0.73
May 2015	7.83	1.44	297	– 2.1	0.66
Jun 2015	10.07	1.57	307	– 1.0	0.71
Jul 2015	11.28	1.65	309	– 0.8	0.69
Sum	77.44	– –	2606	– –	– –

Table F.5: Monthly trading performance of the DBA-down model following the seven months out-of-sample period of *EURUSD_RWDC0.1*.

Applied window	**RR**	**Profit factor**	**Total number of trades**	**MDD (%)**	**Win Ratio**
Jan 2015	– 2.74	0.83	301	– 3.6	0.62
Feb 2015	2.53	1.30	221	– 1.0	0.64
Mar 2015	– 2.69	0.84	313	– 3.8	0.63
Apr 2015	3.78	1.31	295	– 2.2	0.65
May 2015	5.50	1.55	283	– 1.0	0.68
Jun 2015	5.50	1.43	284	– 1.4	0.68
Jul 2015	0.78	1.07	222	– 1.2	0.60
Sum	12.66	– –	1919	– –	– –

Table F.6: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *EURUSD_RWDC0.1*.

Applied window	**RR**	**Profit factor**	**Total number of trades**	**MDD (%)**	**Win Ratio**
Jan 2015	5.64	1.46	303	– 1.4	0.66
Feb 2015	4.15	1.49	216	– 0.9	0.67
Mar 2015	1.67	1.09	355	– 4.6	0.66
Apr 2015	1.20	1.01	337	– 2.8	0.63
May 2015	6.24	1.53	319	– 1.2	0.69
Jun 2015	4.77	1.26	360	– 1.0	0.66
Jul 2015	1.95	1.13	252	– 2.4	0.66
Sum	25.62	– –	2142	– –	– –

Table F.7: Monthly trading performance of the DBA-down model following the seven months out-of-sample period of *GBPAUD_RWDC0.1*.

Applied window	**RR**	**Profit factor**	**Total number of trades**	**MDD (%)**	**Win Ratio**
Jan 2015	12.12	1.84	433	– 0.8	0.70
Feb 2015	10.96	1.74	360	– 1.1	0.72
Mar 2015	4.62	.35	295	– 1.7	0.66
Apr 2015	6.97	1.49	371	– 1.1	0.68
May 2015	13.67	2.16	355	– 1.2	0.69
Jun 2015	6.74	1.37	319	– 1.4	0.68
Jul 2015	13.86	1.65	409	– 1.2	0.72
Sum	68.94	– –	2542	– –	– –

Table F.8: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *GBPAUD_RWDC0.1*.

Applied window	**RR**	**Profit factor**	**Total number of trades**	**MDD (%)**	**Win Ratio**
Jan 2015	13.26	1.80	439	– 1.5	0.68
Feb 2015	8.30	1.62	343	– 1.6	0.72
Mar 2015	8.33	1.71	279	– 0.6	0.70
Apr 2015	8.28	1.48	351	– 1.2	0.66
May 2015	11.03	1.66	334	– 1.0	0.73
Jun 2015	8.63	1.47	322	– 1.0	0.68
Jul 2015	8.40	1.32	401	– 1.8	0.67
Sum	66.23	– –	2469	– –	– –

Table F.9: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *GBPJPY_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	4.78	1.28	414	– 3.9	0.67
Feb 2015	6.69	1.82	246	– 0.7	0.74
Mar 2015	2.85	1.28	254	– 1.0	0.64
Apr 2015	7.59	2.02	240	– 0.6	0.74
May 2015	7.05	1.81	228	– 0.7	0.70
Jun 2015	5.35	1.58	212	– 1.4	0.71
Jul 2015	– 1.54	0.88	198	– 2.5	0.58
Sum	32.77	– –	1792	– –	– –

Table F.10: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *GBPJPY_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	10.90	1.87	410	– 0.9	0.67
Feb 2015	3.08	1.28	242	– 1.1	0.70
Mar 2015	5.87	1.69	240	– 1.0	0.72
Apr 2015	5.65	1.53	237	– 0.6	0.72
May 2015	4.00	1.37	225	– 1.6	0.69
Jun 2015	3.96	1.38	207	– 0.9	0.65
Jul 2015	– 1.39	0.89	191	– 1.7	0.60
Sum	32.07	– –	1752	– –	– –

Table F.11: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *NZDJPY_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	10.88	1.43	557	– 2.7	0.67
Feb 2015	20.96	2.41	472	– 1.4	0.77
Mar 2015	15.82	1.74	464	– 1.0	0.74
Apr 2015	19.39	2.13	401	– 1.1	0.73
May 2015	16.29	1.76	400	– 1.2	0.71
Jun 2015	11.81	1.48	372	– 1.5	0.68
Jul 2015	20.40	1.56	528	– 2.3	0.67
Sum	115.55	– –	3194	– –	– –

Table F.12: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *NZDJPY_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	25.41	2.21	556	– 1.0	0.71
Feb 2015	17.26	1.84	471	– 3.2	0.73
Mar 2015	22.82	2.13	457	– 1.3	0.74
Apr 2015	21.32	2.11	402	– 1.2	0.73
May 2015	24.40	2.11	394	– 1.4	0.73
Jun 2015	27.64	2.19	375	– 1.1	0.72
Jul 2015	42.86	1.94	541	– 1.8	0.72
Sum	181.71	– –	3196	– –	– –

Table F.13: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *AUDJPY_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	7.69	1.31	526	– 2.9	0.69
Feb 2015	12.87	1.86	409	– 1.9	0.74
Mar 2015	10.37	1.70	368	– 0.7	0.68
Apr 2015	10.10	1.71	342	– 1.2	0.67
May 2015	14.73	2.17	335	– 0.7	0.76
Jun 2015	9.58	1.50	304	– 2.4	0.71
Jul 2015	8.61	1.30	433	– 2.7	0.68
Sum	73.59	– –	2717	– –	– –

Table F.14: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *AUDJPY_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	15.62	1.82	515	– 1.2	0.71
Feb 2015	10.41	1.65	383	– 1.7	0.70
Mar 2015	13.78	2.10	341	– 0.9	0.74
Apr 2015	6.73	1.44	319	– 1.4	0.66
May 2015	17.13	2.50	324	– 0.9	0.73
Jun 2015	12.15	1.78	294	– 1.2	0.70
Jul 2015	11.53	1.42	391	– 1.8	0.71
Sum	87.35	– –	2567	– –	– –

Table F.15: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *EURNZD_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	23.47	2.37	597	– 1.1	0.72
Feb 2015	43.08	5.27	547	– 0.8	0.79
Mar 2015	49.90	3.43	543	– 1.0	0.74
Apr 2015	45.40	3.08	502	– 0.9	0.76
May 2015	62.99	3.28	543	– 0.9	0.75
Jun 2015	73.46	2.54	546	– 1.6	0.75
Jul 2015	89.23	3.10	549	– 0.7	0.75
Sum	387.53	– –	3827	– –	– –

Table F.16: Monthly trading performance of the DBA-up model following the seven months out-of-sample period of *EURNZD_RWDC0.1*.

Applied Window	**RR**	**Profit Factor**	**Total Number of Trades**	**MDD (%)**	**Win Ratio**
Jan 2015	17.83	2.07	587	– 1.7	0.72
Feb 2015	42.06	4.14	564	– 1.1	0.79
Mar 2015	36.82	3.10	542	– 1.0	0.74
Apr 2015	48.94	3.58	402	– 0.8	0.76
May 2015	54.98	2.85	525	– 1.0	0.75
Jun 2015	71.73	2.85	539	– 0.8	0.73
Jul 2015	75.83	2.47	540	– 0.9	0.74
Sum	348.19	– –	3699	– –	– –