Directional Change Trading Strategies in the Foreign Exchange Markets

by

Sorin Voicu

August, 2012

Supervised by:

Dr Wing Lon Ng

MSc High Frequency Finance and Trading

University of Essex

Centre for Computational Finance and Economic Agents (CCFEA)

Acknowledgements

First and foremost, I would like to extend my sincerest gratitude to my supervisor, Dr Wing Lon Ng, for his continued diligence, patience and unconditional support. You have been both a mentor and a friend.

I am grateful to Dr. Richard B. Olsen for sharing his vast commercial and research experience with me. His pertinent comments and 'little ironies' helped me to gain more clarity in my approach. Also I would like to thank Olsen & Associates and OANDA Ltd for providing me with the all too precious high-frequency empirical data used in this research paper.

Finally, I would like to thank all CCFEA professors for sharing their knowledge and advice both in and out of the classroom.

Abstract

In a recent paper, Glattfelder et al. (2011) identify 12 scaling laws in high-frequency foreign exchange data and suggest that these may be used to build robust, profitable trading strategies. This dissertation aims to verify some of the claimed laws on recent tick data, across six currency pairs and propose a portfolio of original automated trading strategies that make use of some of these scaling laws in order to trade profitable. We show the benefits of using this approach over the equal weights 'naive' buy-and-hold and mean-variance portfolio, both in terms of returns as well as in terms of performance measures like the Sharpe and Sortino ratios.

Table of Contents

Table of Exhibits

1. Introduction

Over the last decade financial markets have seen a steady increase in the use of information technology as shown in Exhibit 1. The removal of emotions from the trading process, faster access to exchanges, the consistency of returns and the exploitation of newly apparent market inefficiencies were only just a few of the promises electronic trading offered to investors (Aldridge 2010).

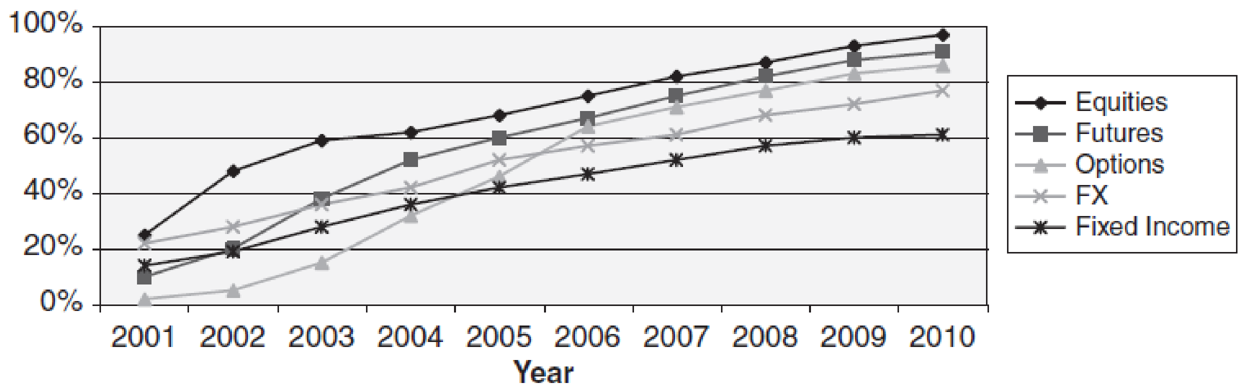

Exhibit 1: Adoption of electronic trading capabilities by asset class

Source: Aldridge (2010), page 10

In terms of investor preference however, the foreign exchange market exceeds by far the other asset classes by daily turnover. Exhibit 2 provides a snapshot of the recent evolution of the global foreign exchange market in terms of daily average turnover.

The increasing availability of high-frequency financial data, recently well in excess of 50,000 data points per day in spot FX markets (Glattfelder, Dupuis and Olsen 2011), is a direct artefact of increasing electronic trading activity. This makes the study of the price formation process and market microstructure certainly more tractable and statistically more significant. The purpose of this dissertation is to gain additional insight into the statistical properties of high-frequency foreign exchange data, to propose a trading strategy that makes use of these properties and to discuss the empirical results of a portfolio of currency pairs.

	Daily average turnover[1]	Growth since 2007 (%)
Global FX market	1,854	25%
By instrument		
Spot	697	108%
Outright forwards	228	84%
Currency options	135	27%
Currency swaps	793	-14%

Exhibit 2: Daily average foreign exchange market turnover in April 2010

Source: Bank of International Settlements (2010), page 6

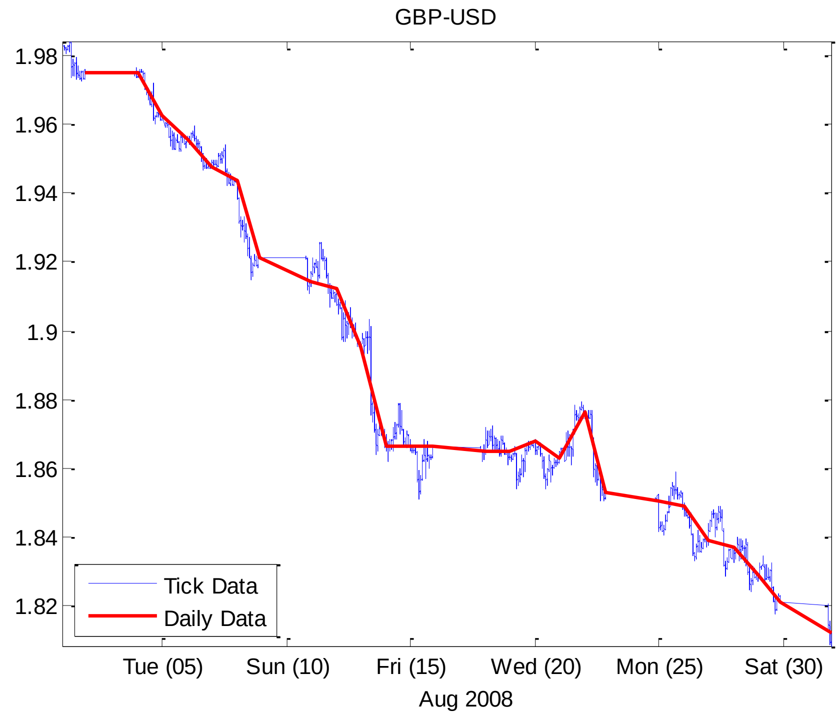

Exhibit 3: The fractal nature of tick-by-tick data

Consider Exhibit 3 which shows tick-by-tick and daily price data for the GBP-USD ranging from 1st – 31st of August 2008. The daily series contains 31 observations while the tick series contains 837,917 observations – a whopping multiple of 27,029. Although both series follow each other

[1] in billions of USD, adjusted for local double-counting. Each leg of a foreign currency transaction other than the US dollar leg has been converted into original currency amounts at average current April exchange rates and then reconverted into US dollar amounts at average April 2010 exchange rates.

quite closely, it is obvious that the tick data exhibits more volatility. This in turn yields a price ‘coastline’ – i.e. the length of all the price moves - of 2,410% compared to only 10% for daily data. Intuitively, this property of high-frequency data suggests that trading at higher frequencies will increase the potential profit margin manyfold than at lower frequencies, even only when a fraction of the coastline is traded. The search for an efficient way to achieve this goal brings us to another property of tick data – its time non-homogeneity.

Each tick (i.e. transaction price/quote) is formed at varying time intervals. This poses significant challenges in terms of research. Until recently, most economic and financial time series models have dealt with equally spaced, lower frequency data. For this issue, there have been mainly two approaches.

The first approach was to homogenize the high-frequency time series by means of interpolation (Dacorogna, et al. 2001, Engle and Russell 2010). However this is not flawless. By changing the filter (i.e. from piece-wise constant to cubic spline for example), the time series would differ significantly and therefore the statistical properties would not be directly comparable. Bauwens and Hautsch (2009) further argued that by aggregating and interpolating tick data amongst fixed or predetermined time intervals, important information about the market microstructure and trader behaviour is lost. For example, by analyzing a filtered time series, the differences between a ‘fast market’ - i.e. a period of high price volatility and a high number of orders or a high volume of trading - and a ‘slow market’ - i.e. a period of relatively low price volatility and low number of orders or a low volume of trading - would be blurred.

These arguments led to the development of point processes (Bauwens and Hautsch 2009). Beside the actual tick information, point processes study the inter-event waiting times (durations) and their frequency. Quote durations, price durations, volume durations have successfully enabled the modelling of intraday volatility, liquidity and transactions costs.

One other area of research that analyses high-frequency time series from the perspective of fractal theory was initiated by Mandelbrot (1963). His pivotal work has inspired others to search for empirical patterns in market data – namely scaling laws that would enhance the understanding of how markets work. A scaling law establishes a mathematical relationship between two variables that holds true over multiple orders of magnitude. For example, one of the most reported scaling laws in the foreign exchange markets (Müller, et al. 1990, Galluccio, et al. 1997, Dacorogna, et al. 2001, Di Matteo, Aste and Dacorogna 2005) is the relationship between the average absolute price change and the time interval of its occurrence:

$$|\overline{\Delta x}| = c\Delta t^k \tag{1}$$

where $|\overline{\Delta x}|$ is the average absolute price change over the time interval Δt; c and k are scaling law constants.

Guillaume, et al. (1997) discovered another scaling law that relates the number of rising and falling price moves of a certain size (threshold) with that respective threshold. Although this discovery is acknowledged, there is no consensus on what drives this behaviour.

Going back to point processes, Dupuis A. and Olsen R. B. argue that time measured as physical time is unable to grasp the information about the state of the market – i.e. a minute in a 'fast market' is not the same as one minute in a 'slow market'. Instead they propose the radical approach of event-based or 'intrinsic time', and they build on the research of their predecessors by making the case of 12 empirical scaling laws in the foreign exchange markets (Dupuis and Olsen 2012, Glattfelder, Dupuis and Olsen 2011).

The intrinsic time framework dissects the time series not based on homogenous physical time, but based on market events where the direction of the trend alternates (Exhibit 4).

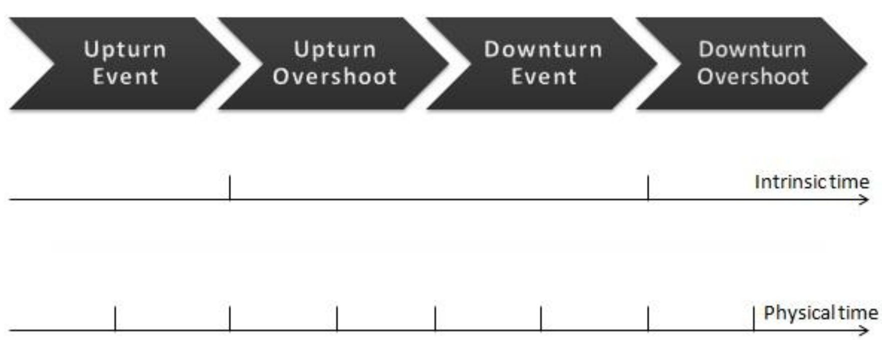

Exhibit 4: Event-based time

Source: Adaptation after Tsang, E.P.K – Directional changes, Definitions (2010)

These directional change events are identified by changes in price of a given threshold value (i.e. net return) set ex-ante. These can be upturn events or downturn events. Once a directional change event has been confirmed an overshoot event has begun. Overshoots continue the trend identified by the directional changes. So, if an upward event has been confirmed, an upward overshoot follows and vice versa. An overshoot event is confirmed when the opposite directional change is confirmed. With each directional change event, the intrinsic time 'ticks' one unit.

The benefits of this approach in the analysis of high-frequency data are manifold. Firstly, it can be applied to non-homogenous time series without the need for further data transformations. Secondly, multiple directional change thresholds can be applied at the same time for the same tick-by-tick data (Exhibit 5). And thirdly, it seamlessly captures the level of market activity at any one time.

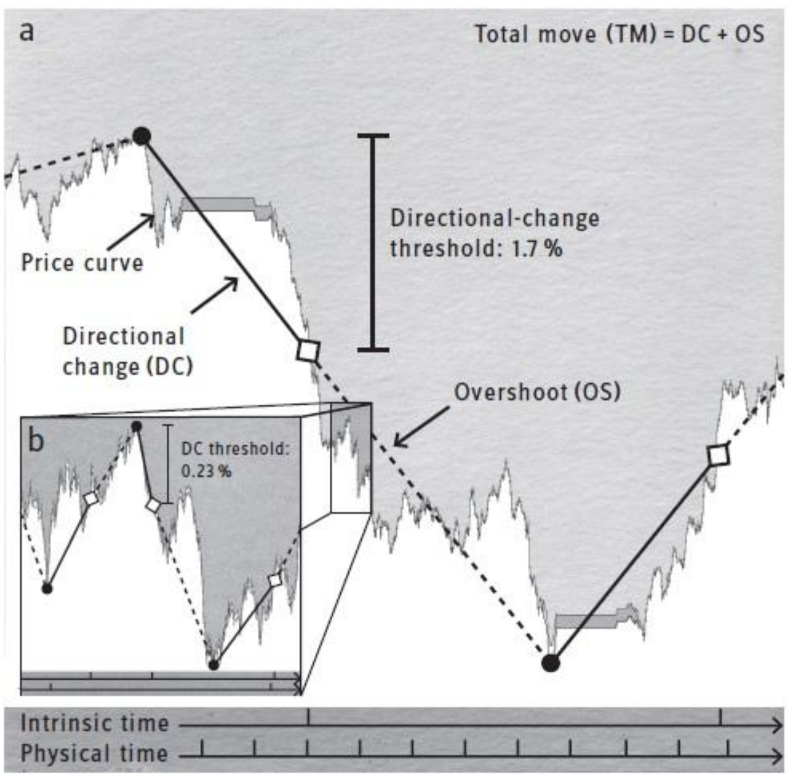

Exhibit 5: Intrinsic time framework

Source: Glattfelder, Dupuis and Olsen (2011), page 4

The downside is that only a portion of the whole price coastline is traded given a certain threshold. Realistically, even if the threshold value is set to a minimum, one has to take into account transaction costs in order to trade profitably. Glattfelder et al. (2011) argue that due to the whooping value of the price coastline, one has to trade less than 1% of that to obtain a profit of 6% per annum or more, net of transaction costs.

The coastline length is the result of the particular microstructure of the foreign exchange market. Ghashghaie et al. (1996) look into what they define as ‘turbulent cascades’. Firstly, because of fierce competition amongst market makers for the best execution price, they post very tight bid-ask spreads compared to the daily average price move. This implies that every market maker has to update the quote such that its exposure is limited. If a market maker does not do this, then a small imbalance between buy and sell orders will result in a large exposure that could possibly turn into a loss. Thus, even a small order can change the quote.

Secondly, the vast majority of foreign exchange players are leveraged. This allows one to trade on margin. However, as even small orders can cause price moves, these do not affect only the market maker, but also other players in the market; especially those who are close to their margin limits. The ‘turbulent cascades’ are therefore the result of one price jump that triggers the closure of positions for some market players. In turn, these closures aggravate the price move in the exact same direction and further margin calls are triggered. Indeed, annual price moves of 10-20% in the foreign exchange market are a rule rather than an exception.

Given the particular microstructure and behaviour of FX markets, Glattfelder et al. (2011) have discovered 12 scaling laws that hold true across 13 currency pairs. They provide a valuable grip on understanding the ‘physics’ of the price moves. They link together concepts like the length of the overshoot, the length of the directional change, the number of directional changes and the number of overshoot ticks with the threshold value, to name just a few.

1.1 Summary

High-frequency data offers increased profit opportunities compared to data at lower frequencies. Due to specific microstructure constructs, tick-by-tick foreign exchange data exhibits increased price volatility when compared to other asset classes. One way to understand and to account for this behaviour in a robust manner, for trading purposes, is by using scaling laws. An interesting extension of the intrinsic time research is the application of a trading algorithm built on scaling laws to a portfolio of currency pairs.

Therefore, the aim of this research paper is twofold. Firstly, it is to verify that some of the scaling laws discovered previously still hold for recent high-frequency data sets. And secondly, it is to use the confirmed scaling parameters to devise a portfolio of directional change strategies across six currency pairs and across multiple thresholds. The performance of this portfolio will then be compared to that of a buy-and-hold and mean-variance portfolios.

In Section 2 we will review our choice of data and we will discuss in detail the methodology put in place in devising the trading strategy. Section 3 details the portfolio performance and contrasts the empirical results with that of the benchmark portfolios. Section 4 concludes and discusses future paths for research.

2. Methodology

Our first aim is to find relationships amongst the scaling laws. We can then derive averaged parameters that would perform in a robust manner across all the tested currency pairs. This will simplify the development of the algorithm and will extend its applicability to out-of-sample data. As with Dupuis et al. (2012), our trading model is countertrend, i.e. a negative price move triggers a buy, while a positive price move triggers a sell. We devise a multi-level trading model that makes use of the intrinsic time framework together with the scaling law relationships we will confirm.

Subsection 2.1 presents the scaling laws we are interested in while Subsections 2.2 - 2.4 describe in full detail each of the levels of the trading algorithm.

2.1 The scaling laws

Let x_t be defined as the price at time t and Δx_t as the price change, i.e. net return:

$$\frac{x_t - x_{t-1}}{x_{t-1}} \tag{2}$$

Also, let $\langle \cdot \rangle$ be the averaging operator defined as:

$$\langle x \rangle = \frac{1}{n}\sum_{j=1}^{n} x_j \tag{3}$$

Glattfelder, Dupuis and Olsen (2011) have shown that 12 scaling laws exist in high-frequency foreign exchange data. Of the 12 scaling laws, for the purposes of our trading strategy, we are interested in testing for confirmation the following:

i. $\langle \Delta x_{os} \rangle = a_{x,os}(\Delta x_{dc})^{b_{x,os}}$, i.e. the average overshoot move as a function of the directional change threshold;

ii. $\langle \Delta t_{dc} \rangle = a_{t,dc}(\Delta x_{dc})^{b_{t,dc}}$, i.e. the average time of the directional change as a function of the directional change threshold;

iii. $\langle \Delta t_{os} \rangle = a_{t,os}(\Delta x_{dc})^{b_{t,os}}$, i.e. the average time of the overshoot as a function of the directional change threshold;

iv. $\langle N(\Delta x_{dc}) \rangle = a_{N,dc}(\Delta x_{dc})^{b_{N,dc}}$, i.e. the average directional change tick count as a function of the directional change threshold;

v. $\langle N(\Delta x_{os}) \rangle = a_{N,os}(\Delta x_{dc})^{b_{N,os}}$, i.e. the average overshoot tick count as a function of the directional change threshold;

2.2 Level 1 – the "coastline trader"[2]

The "coastline trader" is a phrase coined by Dupuis et al. (2012) to describe a process that exploits the price moves, up or down, of certain pre-specified threshold λ. A coastline trader is initialized only after a price overshoot of threshold λ has been observed. This approach is motivated by the scaling relationship (i).

Let $\tau > 0$ be the age of the process expressed as the number of events. An event is a price change of threshold λ. Also, let L_τ be the length of the overshoot and e_τ be the exposure of the coastline trader at τ. Initially, the process has an exposure of

$$e_1 = G(L_1 = 1) \tag{4}$$

where G is the function describing the position increments. In our implementation,

$$G = 2^{L_\tau - 1} * 1000 \tag{5}$$

When a new price move of λ in the same direction of the initial overshoot occurs, the length counter is incremented $L_{\tau+1} = L_\tau + 1$ and the trader increases its exposure to

$$e_{\tau+1} = e_\tau + G(L_{\tau+1}) \tag{6}$$

This practice of increasing a position size while in a 'trend' with the expectation that a trend reversal will take place is also known in technical analysis literature as "pyramiding" (Schwager 1996). For every position increase, the entry price is saved also. The entry price is the market maker's ask price, if the trader is long, or the market maker's bid price if the trader is short.

[2] For the algorithm implementation we have used the following software packages: MATLAB R2012a and SQL Server 2008 R2. From the scripting point of view, we have used procedural as well as object-oriented programming for MATLAB, while for SQL Server we have used SQL – for simpler queries and Transact-SQL in the form of stored procedures and functions to automate in a transparent manner the diverse data transformation tasks.

When a new price move of ω·λ in the opposite direction of the initial overshoot occurs, the trader decreases its exposure to

$$e_{\tau+1} = e_\tau - G(L_\tau) \tag{7}$$

and the length counter is decremented $L_{\tau+1} = L_\tau - 1$. ω represents the target profit multiplier and for the purposes of this research it is set rather aggressively to 1.5. The reason behind this choice is that the scaling laws suggest that, on average, each directional change is followed by an overshoot of equal magnitude. Moreover, this overshoot is followed by the opposite directional change and its respective overshoot. By setting the ω=1.5 we expect to reasonable maximize our profit per trader.

The target profit can be set to any positive value. It is worth pointing out that the bigger the value the longer it will take for the trader to realize the target profit. The profit is determined with the following formula:

$$\pi_{\tau+1} = \left(x_{mode,\tau+1} - x_{init,L_\tau}\right) * G(L_\tau) * mode \tag{8}$$

where $mode \in \{-1,1\}$ is the variable describing the upward or the downward overshoot move respectively; $x_{init,L_\tau}$is the entry price of the last position and $x_{mode,\tau+1}$is the current bid or ask price depending on whether the variable mode is 1 or -1. Finally, the coastline trader closes itself when the length counter $L_\tau = 0$.

2.3 Level 2 – the "directional change engine"

Dealing with the coastline traders directly can be quite challenging from the programming point of view. Thus we have devised an algorithm that manages for us the creation and performance of individual coastline processes. We will call this algorithm a "directional change engine". The input parameters for the engine are: the total allocated capital, the operating threshold value and the currency pair.

For the purpose of this research we have chosen the following threshold values: 0.05%, 0.10%, 0.22%, 0.56%, 0.96% and 2%. The values are uniformly distributed in the log-log space. The logarithmic step is always 0.7378. Furthermore, we have chosen only 6 values in order to keep the computational speed within acceptable limits using retail hardware configurations. The ultimate goal is to have a fully functional 36 directional change engines (6 currency pairs times 6 thresholds each).

For each tick value the directional change engine checks if a directional move has taken place given the before mentioned threshold parameter. If the price move is not consistent with a directional move, the tick price information is ignored. However, if the price move is consistent with a directional move of at least threshold λ, the engine saves the tick price in a temporary variable.

From this point on, according to the scaling laws based on the intrinsic time framework, we can expect an overshoot, on average, of similar magnitude. On one hand, this could mean that an overshoot of less than λ may occur. Implicitly, an overshoot event has ended when a new directional change event has been confirmed. In this case, the above mentioned temporary variable is reset to current directional change tick information.

On the other hand, an overshoot of equal or more than λ may occur. In this case, the temporary variable is deleted and a new coastline trader is initialised with the latest tick price information. From this point on, the coastline trader is self-managed as explained in Section 2.1.

The relationship between the coastline traders and the directional change engine does not stop here. For every directional change followed by an overshoot of similar magnitude, the engine initiates an additional coastline trader, if there is enough unused capital. Similarly, each of the coastline traders cannot add to their positions unless there is remaining unused capital. This constraint is important in the sense that one cannot invest more than 100% of the capital available.

A directional change engine can have multiple coastline traders 'under management'. Once a coastline trader takes profit, the capital value of the engine is increased accordingly. The big advantage with this approach is that the capital amount is shared amongst the coastline traders. This means that once a trader has taken profit and released that proportion of capital, the overall capital value is updated and, should the right conditions be in place, is available to other traders to invest; traders that might otherwise be unable to build up their position, or traders that may be prevented from being initialised in the first place.

2.4 Level 3 – portfolio optimisation

In his pivotal work, Markowitz H. (1952) revealed a model that looked at investments assets as a whole rather than on an individual basis. The purpose was to combine the innate risk aversion of investors with their preference for high expected returns. His definition of an 'efficient portfolio' can be interpreted in two ways. Firstly, a portfolio is efficient when it minimises the volatility given a desired target return; secondly, when it maximises the expected return given a target level of risk (volatility). All feasible portfolios in the expected return-risk universe following these two definitions describe a curve. This curve is called the 'efficient frontier' (Exhibit 6).

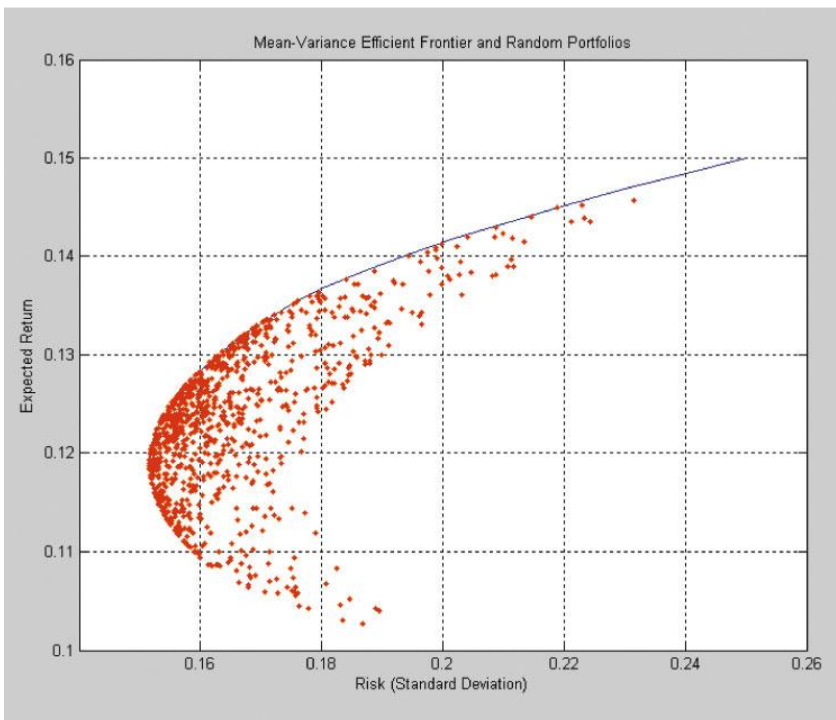

Exhibit 6: Illustration of mean-variance efficient frontier

Source: Mathworks.co.uk[3]

An investor wants to construct the portfolio in such a way that he will maximize his wealth at the end of the holding period. Thus he takes into account the first two moments of the distribution of the portfolio return – the mean and the variance. His utility function is an increasing function of the expected return. As the investors are assumed to be risk averse, this utility function is

[3] http://www.mathworks.co.uk/company/newsletters/articles/developing-portfolio-optimization-models.html

concave. Markowitz' mean-variance analysis is completely consistent with the expected utility maximisation, given that over short periods of time it is safe to assume that the returns are normally distributed.

The solution to the problem described above can be found either by maximisation of expected returns or by minimisation of portfolio variance. The problem can be formalised as:

$$\min_{w} \frac{1}{2} w' \Sigma w \qquad (9)$$
$$I'w = 1, R'w = \mu$$

where Σ is the variance-covariance matrix of returns, I is the identity matrix, R is the expected returns vector, w is the weights vector and μ is the portfolio return. Furthermore, the condition that $w_i \geq 0$ applies in the case of long-only portfolios.

Initially, all our 36 directional change engines are allocated equal amounts of capital. Because the average duration, from initiation to closure, of a coastline trader varies from a few days to a few months we have decided to consider returns on a monthly basis, rather than on a lower time scale. Taking this into account the portfolio optimisation algorithm functions on a rolling window principle. It redistributes the weights every calendar month, given the previous 12 monthly returns.

a. Performance indicators

Beside the self explanatory monthly returns, we have also implemented 2 additional portfolio performance measures: Sharpe ratio and Sortino ratio.

The Sharpe ratio is a measure of risk-adjusted performance of an investment. It is calculated by extracting the risk-free rate from the expected rate of return and dividing the result by the standard deviation of returns:

$$SR = \frac{E(r) - r_f}{\sigma} \qquad (10)$$

The Sharpe ratio measures whether the returns are due to skilled decisions or due to excess risk (Fabozzi 2003). A negative Sharpe ratio indicates that a risk-free asset would have outperformed the analysed investment. The greater the Sharpe ratio value means the less additional risk is taken for an additional unit of return.

In the case of a portfolio of assets, the expected return is defined as:

$$E(r_p) = \sum_i E(r_i) * w_i \tag{11}$$

where *w* is the portfolio weights vector with the property that

$$\sum_i w_i = 1 \tag{12}$$

and $E(r_i)$ is the expected individual returns over the reported period.

Similarly, in the case of a portfolio, the standard deviation is defined as:

$$\sigma_p = \sqrt{\sum_i \sum_j w_i w_j \sigma_i \sigma_j \rho_{ij}} \tag{13}$$

where w_i is the weight of the i-th asset, σ_i is the standard deviation of the i-th asset and ρ_{ij} is the correlation coefficient between the returns of asset *i* and asset *j*.

Another performance metric used in this research is the Sortino ratio. It is similar to the Sharpe ratio in the sense that it measures the risk-adjusted return also. The difference is that penalises only those returns that fall below a user specified target (required rate of return) while the Sharpe ratio penalises both the upside and the downside volatility equally. The Sortino ratio is defined as:

$$SoR = \frac{E(r) - mar}{\sigma_d} \tag{14}$$

where *E(r)* = expected return, *mar* = minimum acceptable return and σ_d = standard deviation of negative asset returns, or the square root of the return distribution's lower partial moment of degree 2.

$$\sigma_d = \sqrt{\int_{-\infty}^{mar} (mar - x)^2 \, f(x) \, dx} \tag{15}$$

$f(\cdot)$ is the probability density function of the returns. So, the ratio can be seen as the rate of return in excess of the investor's minimum acceptable return per unit of downside risk (Fabozzi 2003). For the purposes of comparability between the two ratios, we will consider *mar* and r_f to be equal.

b. Portfolio optimisation implementation

Given the complexity of this project we have approached the portfolio optimisation in three stages. The first stage implies setting the allocation weights for each currency pair. The second stage implies setting the allocation weights for each currency-threshold combination. And the third step is the calculation the global capital weights by multiplying the two previous vectors so that the total sum equals to 1.

For each of the first two steps we have used the quadratic programming technique[4]. Again, the estimation of the weights is based on a window of 12 months, rolling each month after the first year.

Another aspect worth mentioning are the upper and lower bounds we have considered. For the first and second steps – the lower bounds are 8% and 10% respectively and the upper bounds are 20% and 40% respectively. These settings imply a minimum of 0.8% and a maximum of 8% for any currency-threshold combination. These values were chosen on one hand to ensure that no one or few strategies get most of the capital, thus defeating whole the purpose of having a diversified portfolio. On the other hand, sufficient minimum capital had to be provided even to the least performing strategies to ensure future opportunities could be acted upon.

It is worth noting that the minimum and maximum weight values should be chosen with regards to the number of active strategies (in our case 36), the total investment capital (in our case 3,600,000 – the monetary units will be discussed later) and the position increment function (5).

[4] we have used the function *quadprog* available in the standard MATLAB libraries

c. Portfolio rebalancing

After each portfolio optimisation, the capital weights will change, so the directional change engines and implicitly the coastline traders should reflect this change accordingly. We have identified multiple possible cases and we will handle them differently as follows.

The first case is when the all of the weights for all the thresholds, for a certain currency pair have been increased. We have dealt with this by updating the capital value for the respective directional change engines with the new weights multiplied with the latest total capital value.

The second case is when some of the weights for the thresholds, for a certain currency pair have been decreased, while the other have remained constant or have been increased. We have approached this by redistributing the coastline traders in a proportional manner amongst the different thresholds, but only within the same currency pair. It would have made little sense shorting the positions for which the weights have shrunk and thus incurring a loss, only to allocate remaining capital to the same currency pair for future re-investments. The order in which the reallocation is made is in the ascending order of the magnitude of the threshold. The reason behind this is that the coastline traders with lower thresholds have a shorter duration from initiation until closure than the coastline traders with higher thresholds. We thus minimise the time until realised potential profits are made.

The third case is when all the weights for all the thresholds, for a certain currency pair have been decreased. We have chosen to take a loss by exiting positions proportional to the weights' differential.

For all the above cases, the capital amount parameter for each directional change engine is updated accordingly.

d. Rebased price data

In order to have a clearer picture of the relative price moves amongst all currency pairs we have decided to rebase all the price series so that each mid-price would start from the value of 1 at 1st January 2007 as shown in Exhibit 8 (page 20). The percentage returns of both original bid and ask prices are kept unchanged. This approach has multiple benefits.

Firstly, we only need to know only the rebasing factors in order to perform conversions between the currency pairs. Thus, it relieves us from managing a ‘high-frequency’ exchange rate for all the pairs in question.

Secondly, by creating an equivalent ‘virtual’ currency, each of the currency pairs can be seen as an investment asset which yields the same return as the original priced one. Moreover, the capital can be expressed in only one equivalent ‘currency’ and the complexity and speed of the portfolio allocation algorithm can be greatly improved.

In the next Section we will have a closer look at the empirical results of this trading model.

3. Analysis of empirical results

This Section introduces the tick data used, presents the estimated scaling law parameters and discusses in detail the trading performance both in absolute as well as in relative terms to the benchmark portfolios.

3.1 Empirical data

The dataset used in this paper consists of 6 currency pairs spanning over 5 years and 3 months – from 1st Jan 2007 to 30th Mar 2012. The original time series are displayed in Exhibit 7 while the rebased series are displayed in Exhibit 8. The following currency pairs are considered with the number of ticks given in parenthesis: EUR_CHF (62,850,265), EUR_USD (63,823,640), GBP_USD (44,348,615), USD_CAD (30,688,870), USD_CHF (38,547,403) and USD_JPY (39,636,700). The different number of ticks is due to the varying liquidity of each spot product (Glattfelder, Dupuis and Olsen 2011). The reason behind choosing these currency pairs is that they account for 65.7% of daily transacted volume across all spot FX market (Bank of England 2012).

The same data spanning over 4 years – from 1st Jan 2007 to 31st Dec 2010 will be used to test some of the scaling laws discovered by Glattfelder et al. These scaling laws will further be used as a framework to develop the trading algorithm. The number of ticks for each currency is allocated as follows: EUR_CHF (49,736,183), EUR_USD (48,018,314), GBP_USD (36,515,702), USD_CAD (22,645,706), USD_CHF (33,591,718) and USD_JPY (34,559,965).

3.2 Estimation of scaling law parameters

The estimation dataset consists of a timestamp, a bid price and a ask price column. The mid-price is determined using the formula below:

$$x_t = (bid_t + ask_t)/2 \tag{16}$$

No other transformations are applied to the raw data. The mid-price data is then used in the procedure outlined in Section 1 to dissect the total price movement into directional changes and overshoots. Also, we selected 10 measurement points for the laws proportional to price thresholds. The range is between 0.01% and 5% in logarithmic steps, so that in log-space the

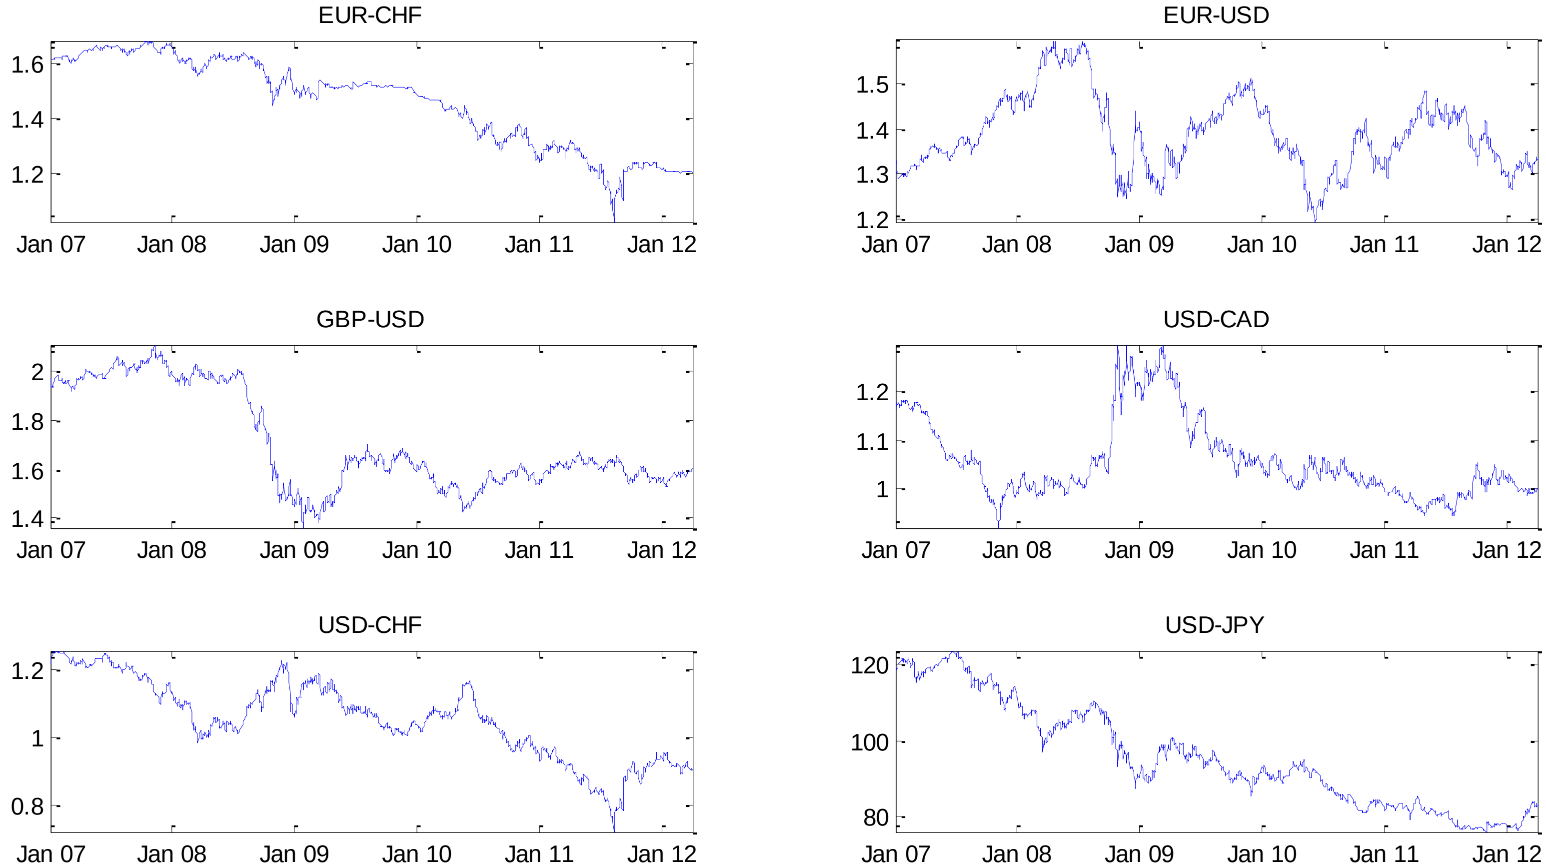

Exhibit 7: Original mid-price series

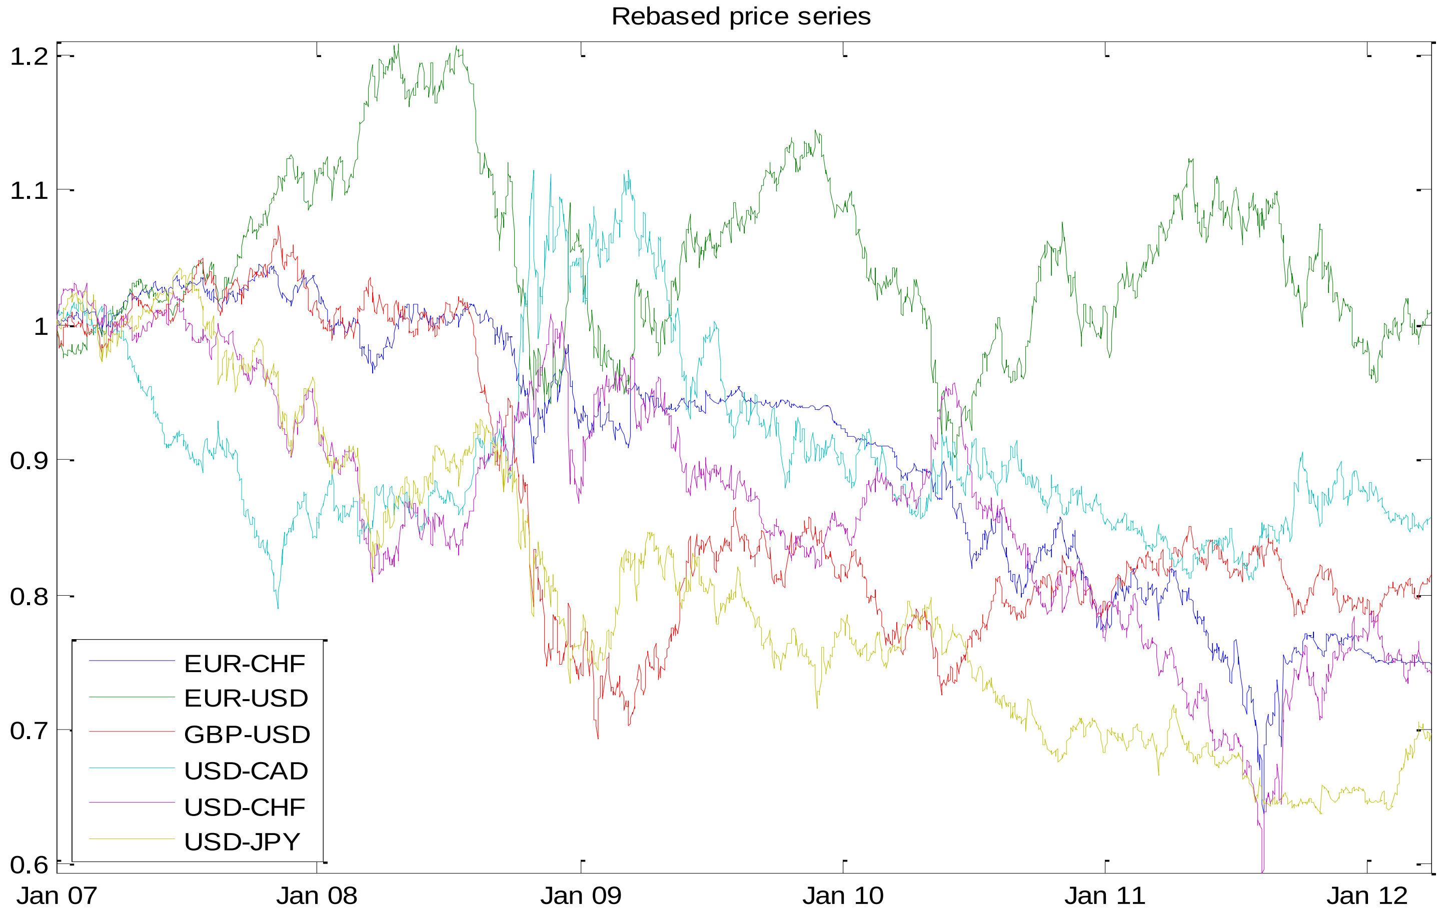

Exhibit 8: Rebased mid-price series

difference between consecutive threshold values is equal – in this case 0.2999. Exhibit 9 quantifies the number of directional changes.

Δx_{dc} (%)	EUR-CHF	EUR-USD	GBP-USD	USD-CAD	USD-CHF	USD-JPY
0.01%	3,192,853	2,786,694	3,153,176	3,845,752	1,922,624	2,602,003
0.02%	1,058,427	1,139,302	1,031,328	1,234,191	629,378	717,089
0.04%	313,048	368,034	316,182	364,919	196,812	169,265
0.08%	90,249	106,457	90,019	104,245	57,800	37,841
0.16%	24,599	28,633	23,690	27,406	15,662	8,158
0.32%	6,251	7,065	5,822	6,919	4,062	1,724
0.63%	1,584	1,808	1,441	1,744	1,004	406
1.26%	419	430	358	466	240	98
2.51%	88	100	88	114	72	28
5.00%	28	30	18	29	14	7

Exhibit 9: Number of directional changes for a period between 1st Jan 2007 and 31st Dec 2010 for different directional change thresholds

The scaling laws have a constant parameter (i.e. $a_{x,os}$) and a scaling exponent (i.e. $b_{x,os}$). By taking the logarithm of both sides of the scaling law, the power relationship is transformed into a linear equation of slope $b_{x,os}$ and intercept $log\ (a_{x,os})$. Standard MATLAB OLS regression is

$$log\ (\langle \Delta x_{os} \rangle) = log\ (a_{x,os}) + b_{x,os} \cdot log\ (\langle \Delta x_{dc} \rangle) \tag{17}$$

used to estimate the individual parameters and their standard errors. In addition, the adjusted R^2 and MSE measures are calculated to determine the goodness of fit of the linear model.

Exhibit 10 and Exhibit 11 confirm the existence of a scaling law. The constant average parameter across all currency pairs of 0.9993 (i.e. $10^{-0.0003}$) and the scaling factor of 1.0096 gives an approximate scaling relationship of:

$$\langle \Delta x_{os} \rangle \approx \langle \Delta x_{dc} \rangle \tag{18}$$

The law $\langle \Delta x_{dc} \rangle = \Delta x_{dc}$ is considered self evident. Furthermore, Exhibit 22, Exhibit 24, Exhibit 26 and Exhibit 28 in the Appendix prove the existence of another two important relationships:

$$\langle \Delta t_{os} \rangle \approx 2 \langle \Delta t_{dc} \rangle \tag{19}$$

$$(a^{av}{}_{t,dc} \approx 6774.37\ days, a^{av}{}_{t,os} \approx 13776.28\ days, b^{av}{}_{t,dc} \approx 1.91, b^{av}{}_{t,os} \approx 1.95)$$

$$\langle N(\Delta x_{os})\rangle \approx 2\langle N(\Delta x_{dc})\rangle \tag{20}$$

$$(a^{av}{}_{x,dc} \approx 189{,}346{,}287\ ticks, a^{av}{}_{x,os} \approx 337{,}455{,}524\ ticks, b^{av}{}_{x,dc} \approx 1.95, b^{av}{}_{x,os} \approx 1.94)$$

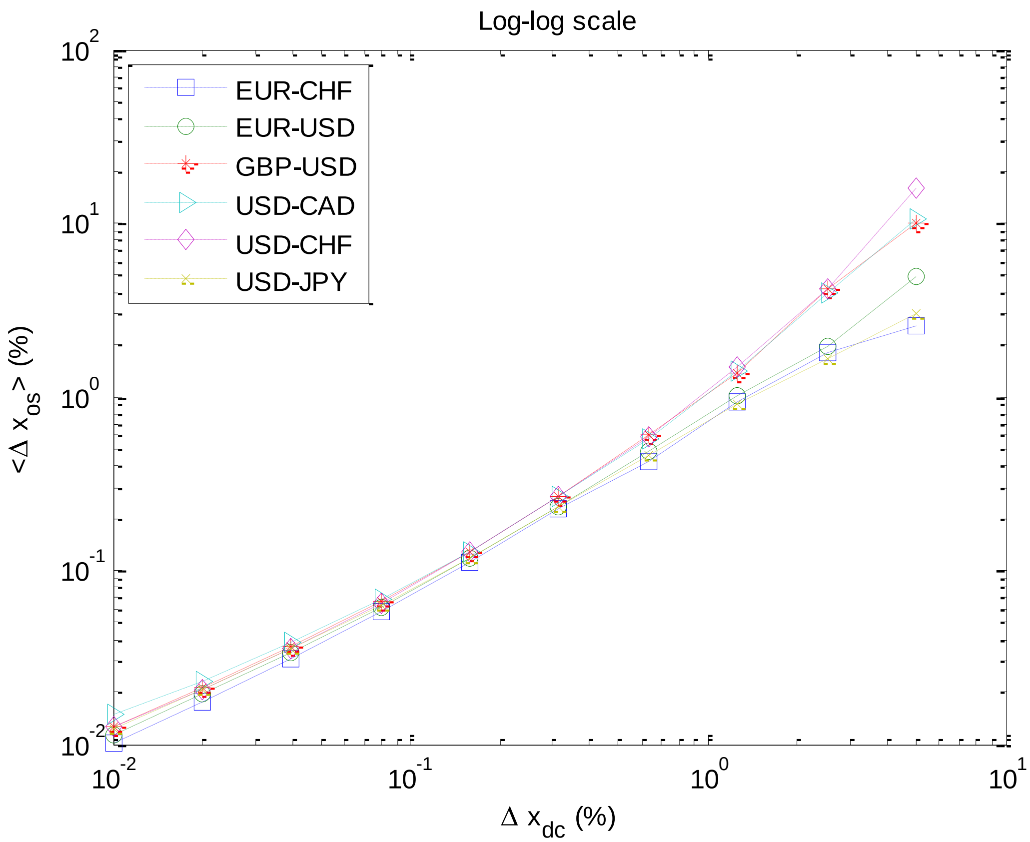

Exhibit 10: Average overshoot move vs. directional change threshold scaling law

In other words, a directional change of say 0.16% that takes on average 2,591 seconds, or 43 minutes and 681 ticks to unfold is followed by an overshoot of 0.15% that takes on average 4,170 seconds, or 96 minutes and 1,268 ticks to be confirmed.

	Slope	**s. e.**	**Intercept**	**s. e.**	**Adj. R^2**	**MSE**
EUR-CHF	0.9319	1.69E-02	-0.3131	4.70E-02	0.9971	2.11E-03
EUR-USD	0.9719	2.24E-02	-0.1467	6.25E-02	0.9952	3.74E-03
GBP-USD	1.0737	4.47E-02	0.2350	1.24E-01	0.9846	1.48E-02
USD-CAD	1.0568	5.10E-02	0.2010	1.42E-01	0.9795	1.93E-02
USD-CHF	1.1194	5.86E-02	0.3751	1.63E-01	0.9759	2.55E-02
USD-JPY	0.9042	1.25E-02	-0.3533	3.49E-02	0.9983	1.16E-03
	1.0096	**3.43E-02**	**-0.0003**	**9.57E-02**		

Exhibit 11: Estimated scaling law (i) parameter values

These findings are similar with those of Glattfelder, Dupuis and Olsen (2011) measured for a period between 1st December 2002 and 1st December 2007.

3.3 Trading performance

Having confirmed the above mentioned scaling law relationships we are now ready to implement them into the trading algorithm, more specifically into the Levels 1 and 2 – the “coastline trader” and the “directional change engine”. Level 2 of the algorithm manages the creation of the respective coastline traders. When the target profit has been reached, the coastline trader takes profit and deletes itself from memory.

The trading period covers 5 years and 3 months – from 1st Jan 2007 to 30th Mar 2012. There are 6 currency pairs with 6 thresholds each. We will consider them to be a portfolio of 36 separate strategies, each with its own individual capital allowance. At 1st Jan 2007, each strategy will be allocated an equal amount of capital of 100,000 monetary units. Moreover, in order to increase the speed and reduce the complexity of the algorithm, we have rebased all the price series so that the first tick will be equal to 1. The first tick value of each original currency pair represents the rebasing factor.

For the period between 1st Jan 2007 and 31st Dec 2007, all the strategies keep their initial equal weights configuration. After that, for each month, the portfolio is rebalanced following a 12-month rolling window approach. All profit is reinvested.

Exhibit 12 shows the average duration of a coastline trader for each currency pair – threshold strategy. The coastline trader has a more complex behaviour than simply entering or exiting a position once a directional change of a certain threshold has been confirmed - as described in Section 2.2. Interestingly, it seems from the table values that the average duration of the coastline trader is in itself a ‘scaling law’.

According to Section 2.4c we can refer to these durations as ‘unaltered’. Once the portfolio rebalancing starts, some of the active traders will be allocated to a different threshold of the same currency pair in order to minimise loss. After the ‘altered’ trader has closed all of its exposure, the counter of the latest threshold-currency pair is incremented. Therefore, from now on, the average coastline trader duration values are ‘altered’.

	Coastline trader average duration (days)						
		Threshold					
		0.05%	**0.10%**	**0.22%**	**0.46%**	**0.96%**	**2.00%**
Currency pair	**GBP_USD**	3.5	8.3	12.7	31.5	53.2	77.7
	USD_CAD	5.5	6.6	9.5	18.3	40.3	-
	USD_CHF	4.0	5.9	12.0	24.3	48.3	52.1
	USD_JPY	3.6	7.2	14.3	28.8	50.5	76.6
	EUR_USD	4.4	6.3	20.4	25.2	24.5	-
	EUR_CHF	6.4	10.8	23.9	34.1	49.4	-

Exhibit 12: Average 'life' of a coastline trader in days for the period 1st Jan - 31st Dec 2007

Exhibit 13 shows the 'altered' duration values. As expected, the previous duration 'scaling law' is no longer apparent.

	Coastline trader average duration (days)						
		Threshold					
		0.05%	**0.10%**	**0.22%**	**0.46%**	**0.96%**	**2.00%**
Currency pair	**GBP_USD**	4.4	10.2	14.9	41.7	80.2	156.0
	USD_CAD	2.8	4.4	7.0	35.4	112.2	89.4
	USD_CHF	3.2	7.2	22.5	31.8	60.8	73.4
	USD_JPY	3.0	6.7	11.0	52.3	45.1	102.0
	EUR_USD	4.7	11.4	15.5	28.0	81.7	102.7
	EUR_CHF	5.1	7.4	12.8	43.5	57.6	147.0

Exhibit 13: Average 'life' of a coastline trader in days for the period 1st Jan 2007 – 31st Mar 2012

The cumulative trading performance for each of the 6 thresholds, for the GBP-USD currency pair is exemplified in Exhibit 14. The graphs for all the currency pairs are shown in the Appendix - Exhibit 29. It is worth noting two aspects. First, the equity curves seem only to increase. This is not entirely true. The coastline traders will only close a position once a target profit has been achieved. However, the only time traders make a loss is when portfolio rebalancing takes place. During the entire tick data period, the monthly loss due to reallocation of capital varies from 0 - 1.52%.

Second, during the first year, some of the equity curves seem to reach a plateau. This happens because the respective coastline traders have invested the entire allocated capital and the prices are trending strongly in one direction, so there is no opportunity to achieve the target profit by means of a directional change. The traders are effectively in 'stalemate' and the invested capital is not making returns. In order to make use more efficiently of the available capital, every month a portfolio rebalancing takes place. The most capital is allocated to the most profitable strategy in the last 12-month rolling window. In order to minimise rebalancing costs, the coastline traders of the same currency pair are redistributed according to the new portfolio weights.

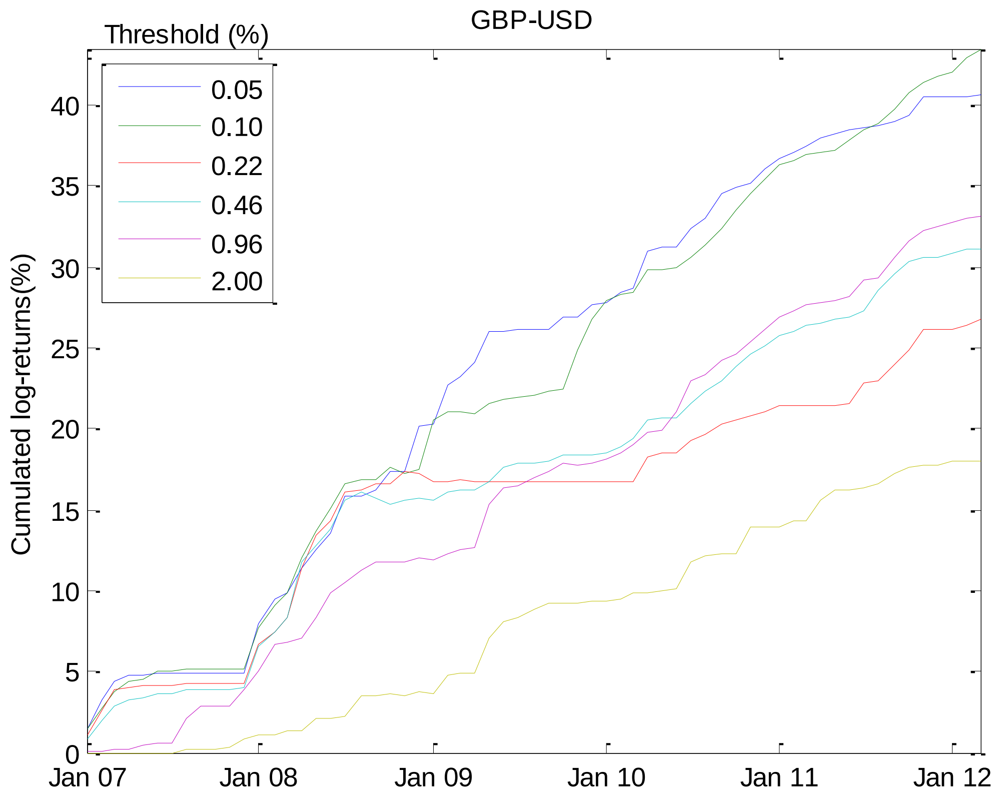

Exhibit 14: Cumulative performance of the trading strategies for GBP-USD

Exhibit 16 displays a cumulated performance of 35.16% at the end of the evaluation period, with an annualised Sharpe ratio of 4.83 and annualised Sortino ratio of 153.72. While these results are certainly impressive, we should also refer to the portfolio weights. Exhibit 15 shows the evolution of the aggregated weights per currency pair. For the detailed evolution of the currency pair – threshold weights please see Exhibit 17. Our prudent approach in implementing constraints into the portfolio reallocation algorithm – described in Section 2.4c – limited the reallocation costs while allowing promising strategies to increase their investment.

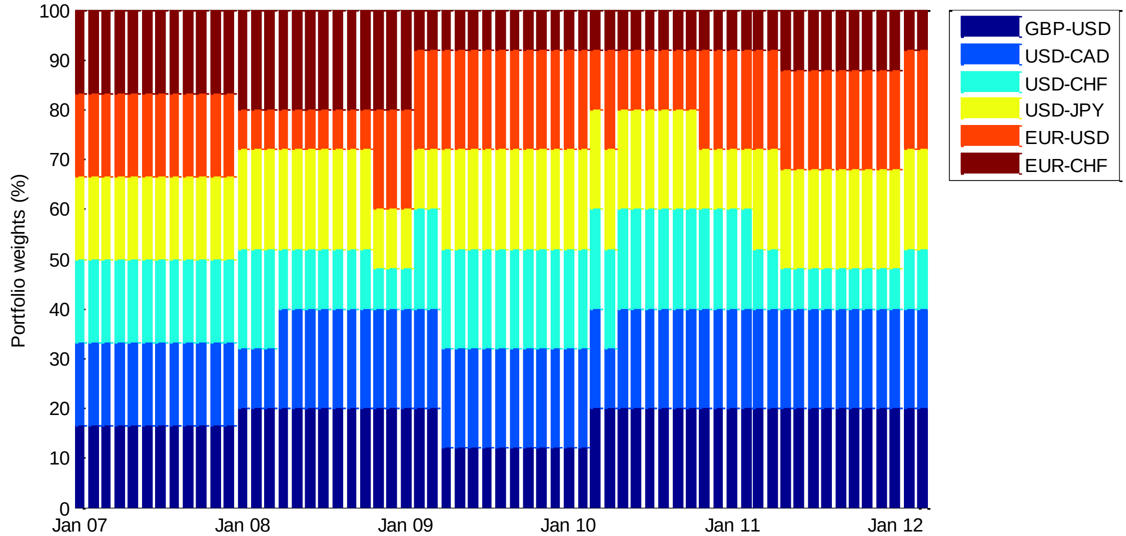

Exhibit 15: Aggregated portfolio weights for each currency pair

3.4 Benchmark comparison

It is worth getting a sense of our strategy performance relative to other benchmarks. We have chosen two models for this purpose: the equal weights 'naive' buy-and-hold portfolio and the mean-variance portfolio.

Exhibit 17 captures graphically the cumulated performance of the three portfolios. To maintain comparability amongst the portfolios, the mean-variance portfolio acts like the equal-weights portfolio for the entire year of 2007. After that, it is rebalanced every month taking the previous 12-month rolling window approach. Unlike our directional change portfolio weights, the mean-variance portfolio weights are not constrained. Both benchmarks portfolios are long-only.

Regarding performance indicators, the directional change portfolio outperforms the benchmarks both in terms of Sharpe ratio (Exhibit 19) as well as in terms of Sortino ratio (Exhibit 20) for each of the tested years.

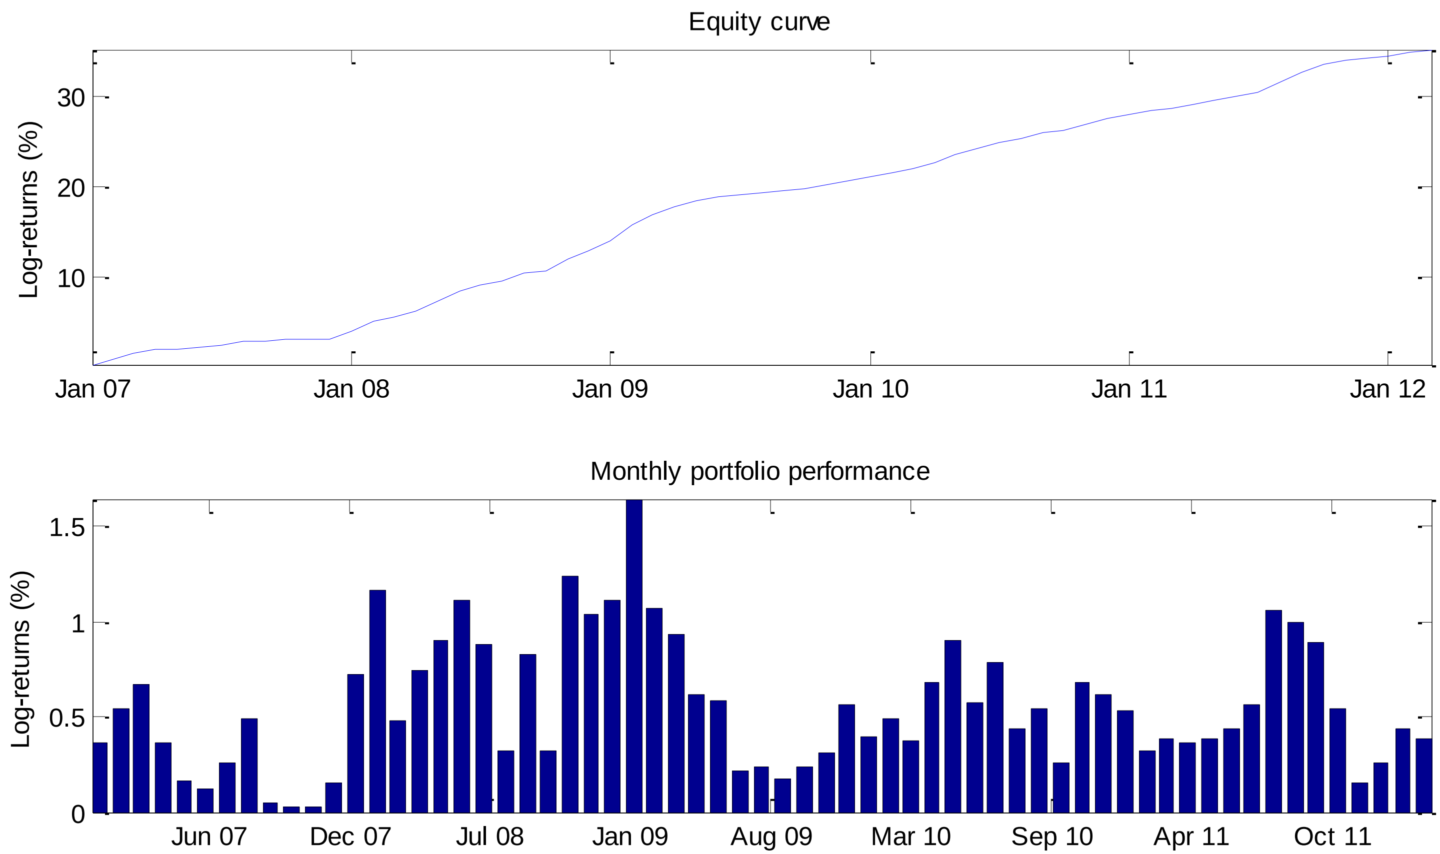

Exhibit 16: Overall performance of the directional change portfolio

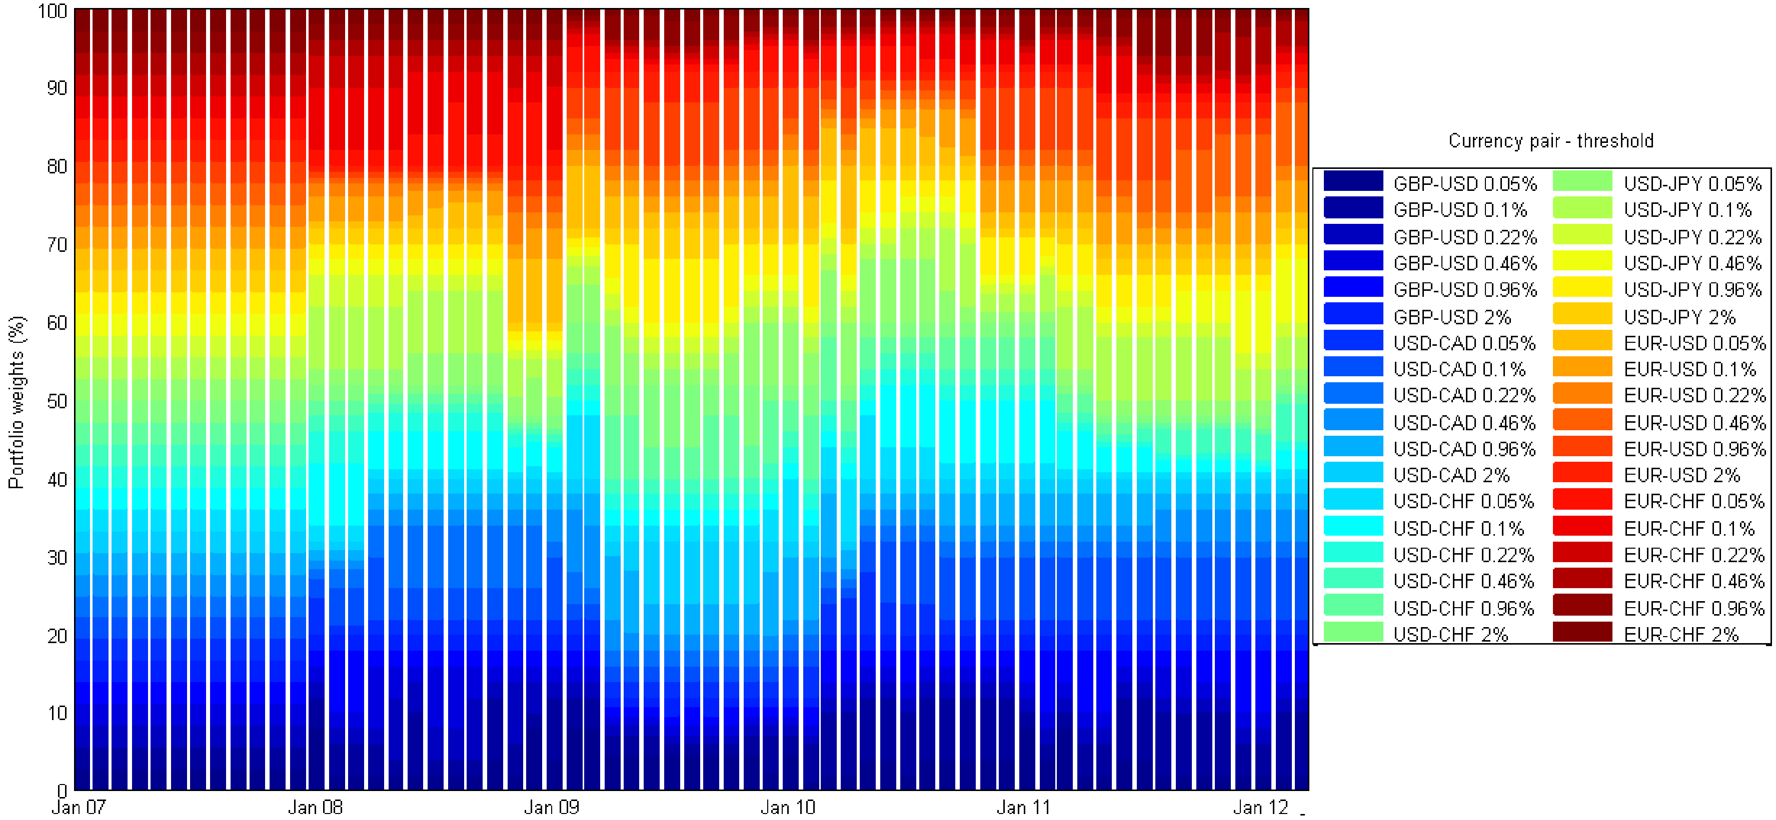

Exhibit 17: Evolution of the directional change portfolio weights

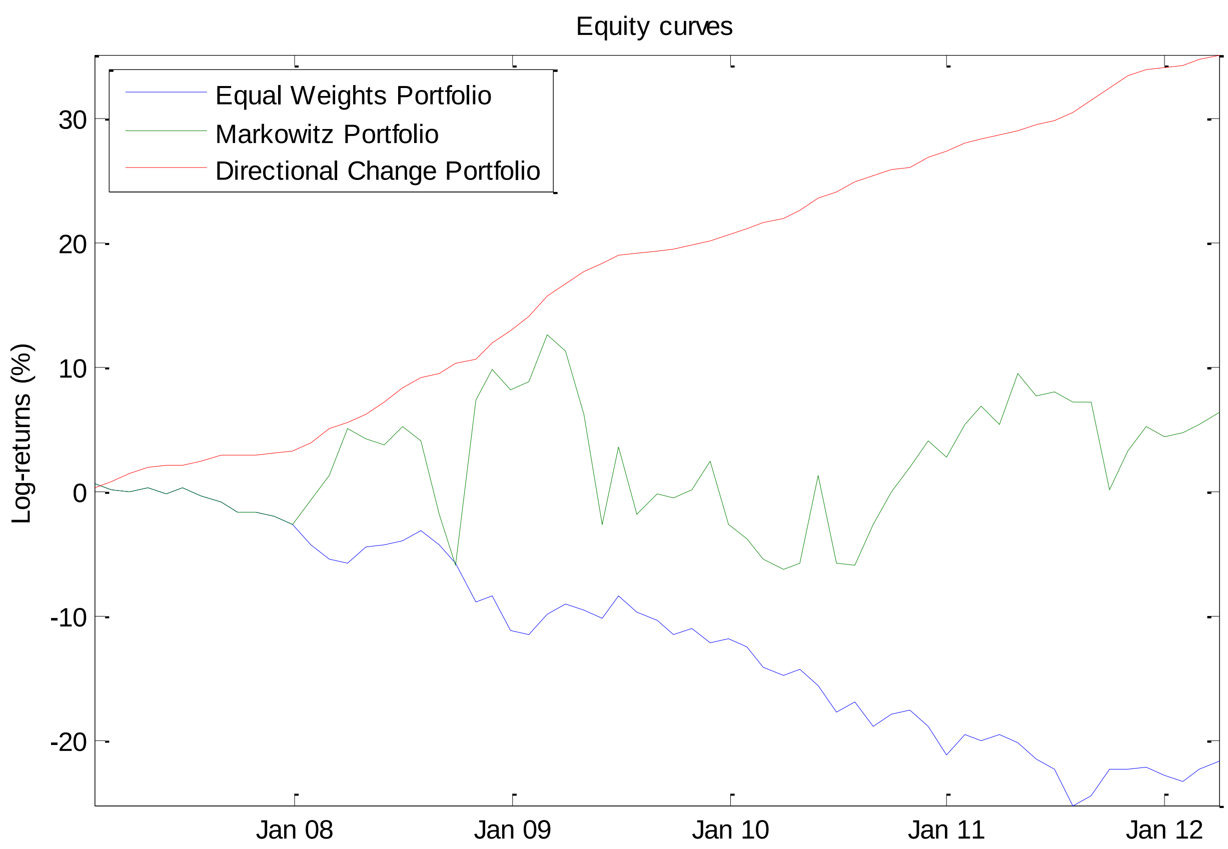

Exhibit 18: Directional change portfolio performance relative to mean-variance and equal weights buy-and-hold portfolios

Annualised Sharpe Ratio[5]			
Year	Equal-weights buy and hold portfolio	Mean-variance portfolio	Directional change portfolio
2007	-2.18	-2.18	2.97
2008	-1.95	1.32	8.15
2009	-0.44	0.10	4.70
2010	-2.52	-1.05	9.43
2011	-0.54	-0.20	5.30
2012	1.09	-1.89	10.33

Exhibit 19: Annualised Sharpe ratio benchmark comparison

Annualised Sortino Ratio[6]			
Year	Equal-weights buy and hold portfolio	Mean-variance portfolio	Directional change portfolio
2007	-0.78	-0.78	17.54
2008	-1.02	4.01	107.85
2009	-0.44	0.11	63.82
2010	-1.64	-1.47	47.81
2011	-1.37	-0.05	44.04
2012	2.42	-38.73	14.56

Exhibit 20: Annualised Sortino ratio benchmark comparison

A detailed month-by-month illustration of the realised Sharpe and Sortino ratios are given in the Appendix - Exhibit 30 and Exhibit 31. For comparability reasons, the risk-free rate and the minimum acceptable return rate are set at 1% per annum.

[5] Computed given a risk-free rate of 1% per annum
[6] Computed given a minimum acceptable return rate (MAR) of 1% per annum

4. Conclusion

This dissertation was built on the results of Glattfelder, Dupuis and Olsen (2011) in exploring scaling law relationships in high-frequency foreign exchange data. Our findings confirm five of their results on a recent time frame. We used these scaling laws to derive 3 relationships:

- The directional change move is on average equal to the overshoot move;
- The number of ticks in a directional change move is on average equal to half the number of ticks in an overshoot move;
- The duration of a directional change move is on average equal to half the duration of an overshoot move.

We have used these relationships as building blocks in devising an original, multi-level trading strategy across multiple currency pairs for the period from 1st Jan 2007 to 31st March 2012.

We have compared the empirical results with the benchmark portfolios in terms of returns, as well as in terms of Sharpe and Sortino ratios. Our directional change portfolio outperformed both benchmarks by a significant margin. Our results are even more significant if we consider that the period 2007-2009 has been the peak of the most recent financial crisis.

We conclude suggesting a few possible paths for further development:

- Currently, our algorithm uses averaged scaling parameters across all 6 currency pairs. One might use the individual scaling parameters to check for additional increase in performance.
- This multi-level approach towards the strategy makes it manageable to code and to separate distinct logic interactions amongst levels. However, it makes it increasingly difficult to conduct sensitivity analysis on the vast number of parameters involved – especially when large tick-by-tick data is involved. One might rewrite it such that a principal component analysis on the significance of the parameters can be conducted.
- Last, but not least, the importance of finding new scaling laws that trade profitably should not be underestimated.

5. Bibliography

Aldridge, I. *High-Frequency Trading - A practical guide to Algorithmic Strategies and Trading Systems.* New Jersey: Wiley, 2010.

Bank of England. “Markets.” April 2012. http://www.bankofengland.co.uk/markets/Documents/forex/fxjsc/fxturndata120730.xls (accessed August 20, 2012).

Bank of International Settlements. “BIS triennial survey of Foreign Exchange and Over-the-Counter Interest Rate Derivatives Markets in April 2010.” 1 September 2010. http://www.bankofengland.co.uk/statistics/bis-survey/fxotcsum10.pdf (accessed June 16, 2012).

Bauwens, L., and N. Hautsch. “Modelling financial high frequency data using point processes.” In *Handbook of financial time series*, by T. G. Anderson and et al., 953-976. Berlin: Springer, 2009.

Dacorogna, M. M., R. Gencay, U. A. Müller, R. B. Olsen, and O. V. Pictet. *An introduction to high-frequency finance.* San Diego (CA): Academic Press, 2001.

Di Matteo, T., T. Aste, and M. M. Dacorogna. “Long term memories of developed and emerging markets: using the scaling analysis to characterize their stage of development.” *Journal of Banking and Finance* 29, no. 4 (2005): 827-851.

Dupuis, A., and R. B. Olsen. “High Frequency Finance: Using Scaling Laws to Build Trading Models.” In *Handbook of Exchange Rates*, by J. James and et al., 563-582. Wiley, 2012.

Engle, R. F., and J. R. Russell. “Analysis of high frequency financial data.” In *Handbook of financial econometrics*, by Y. Ait-Sahalia and L. P. Hansen, 383-424. Amsterdam: Elsevier, 2010.

Fabozzi, Frank J. *Investment performance measurement.* New Jersey: John Wiley & Sons, 2003.

Galluccio, S., G. Caldarelli, M. Marsili, and Y. C. Zhang. “Scaling in currency exchange.” *Physica A* 245 (1997): 423-436.

Ghashghaie, S., P. Talkner, W. Breymann, J. Peinke, and Y. Dodge. “Turbulent cascades in foreign exchange markets.” *Nature* 381 (1996): 767-770.

Glattfelder, J. B., A. Dupuis, and R. B. Olsen. “Patterns in high-frequency FX data: Discovery of 12 empirical scaling laws.” *Quantitative Finance* 11 (2011): 599-614.

Guillaume, D. M., M. M. Dacorogna, R. D. Dave, U. A. Müller, R. B. Olsen, and O. V. Pictet. “From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets.” *Finance and Stochastics* 1 (1997): 95-129.

Mandelbrot, B. B. “The variaton of certain speculative prices.” *The Journal of Business* 36 (1963): 394-419.

Müller, U. A., M. M. Dacorogna, R. B. Olsen, O. V. Pictet, M. Schwarz, and C. Morgenegg. “Statistical study of foreign exchange rates, empirical evidence of a price change scaling law and intraday analysis.” *Journal of Banking and Finance* 14 (1990): 1189-1208.

Schwager, J. D. *Schwager on Futures. Technical Analysis.* New York: Wiley, 1996.

Van Benthem, J. *The logic of time : a model-theoretic investigation into the varieties of temporal ontology and temporal discourse.* Boston: D. Reidel Publishing Co, 1983.

6. Appendix

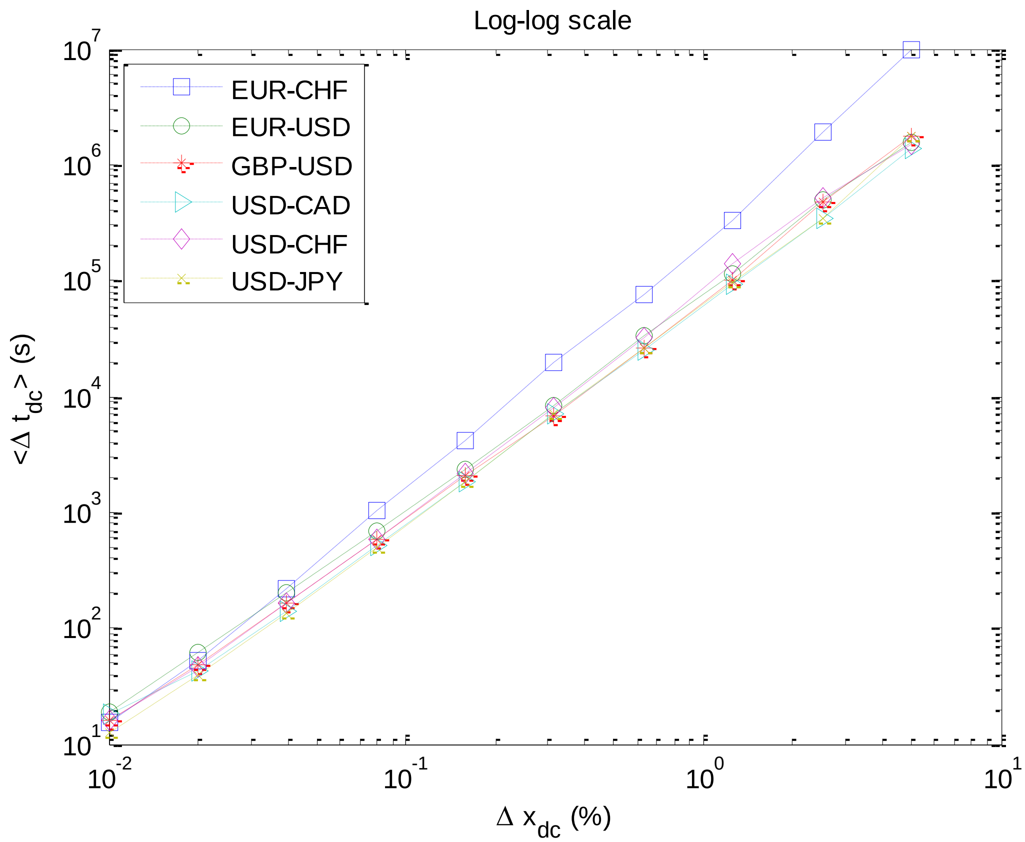

Exhibit 21: Average time of the directional change vs. directional change threshold scaling law

	Slope	s. e.	Intercept	s. e.	Adj. R^2	MSE
EUR-CHF	2.1498	2.54E-02	4.7446	7.08E-02	0.9987	4.78E-03
EUR-USD	1.8383	1.75E-02	3.6331	4.87E-02	0.9992	2.27E-03
GBP-USD	1.8765	2.50E-02	3.6754	6.98E-02	0.9984	4.65E-03
USD-CAD	1.8384	2.53E-02	3.5261	7.04E-02	0.9983	4.73E-03
USD-CHF	1.8877	2.05E-02	3.7275	5.72E-02	0.9989	3.12E-03
USD-JPY	1.8986	1.98E-02	3.6785	5.51E-02	0.9990	2.90E-03
	1.9149	**2.22E-02**	**3.8309**	**6.20E-02**		

Exhibit 22: Estimated scaling law (ii) parameter values

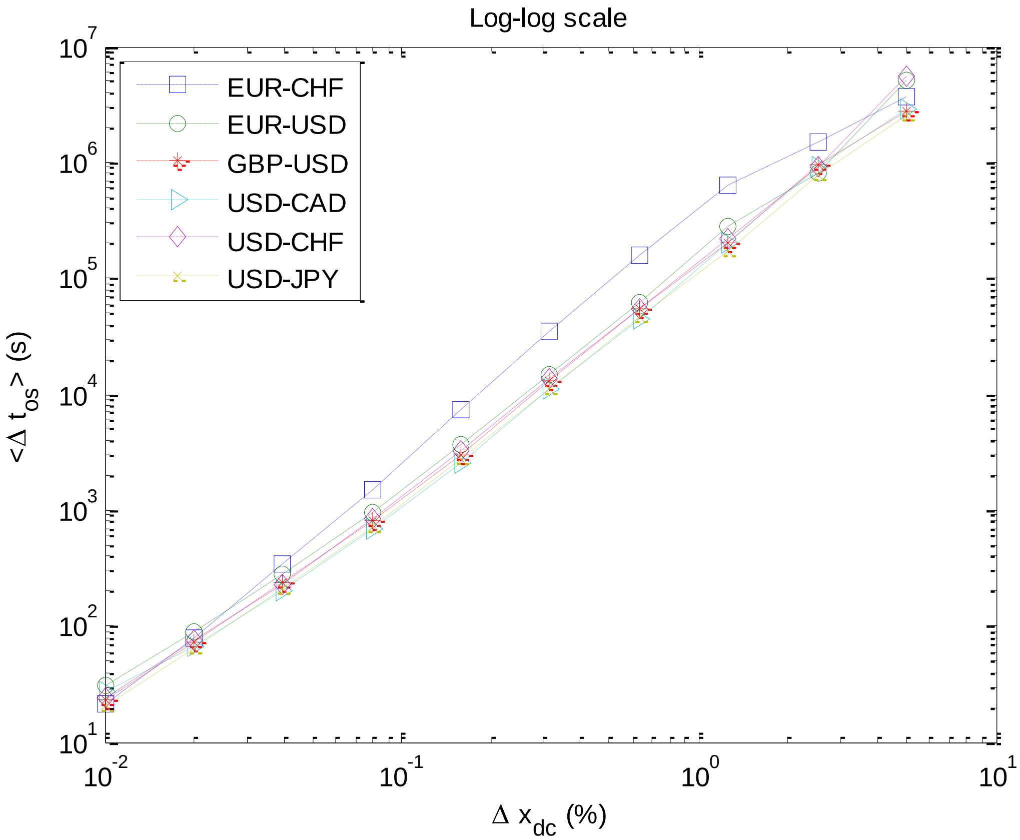

Exhibit 23: Average time of the overshoot vs. directional change threshold scaling law

	Slope	s. e.	Intercept	s. e.	Adj. R^2	MSE
EUR-CHF	2.0233	5.58E-02	4.5338	1.55E-01	0.9932	2.31E-02
EUR-USD	1.9382	3.66E-02	4.1459	1.02E-01	0.9968	9.95E-03
GBP-USD	1.9232	2.55E-02	4.0169	7.11E-02	0.9984	4.83E-03
USD-CAD	1.9233	4.21E-02	3.9822	1.17E-01	0.9957	1.32E-02
USD-CHF	1.9764	3.77E-02	4.1988	1.05E-01	0.9967	1.06E-02
USD-JPY	1.9231	2.24E-02	3.9572	6.24E-02	0.9988	3.72E-03
	1.9512	**3.67E-02**	**4.1391**	**1.02E-01**		

Exhibit 24: Estimated scaling law (iii) parameter values

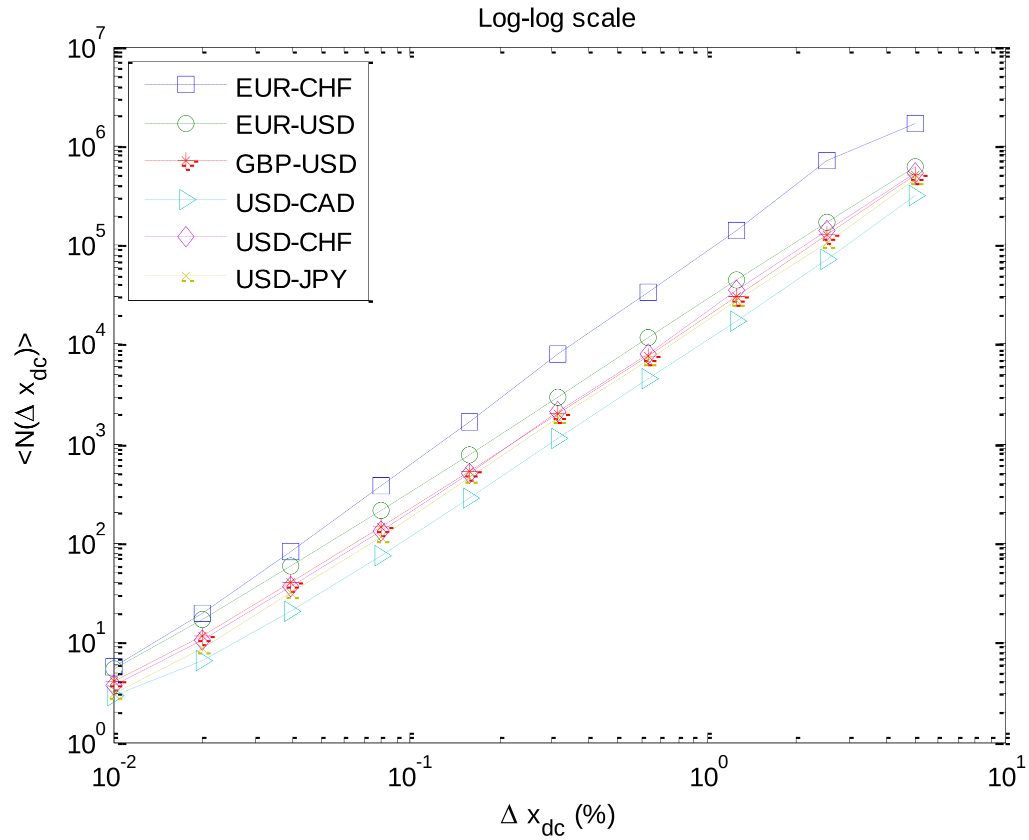

Exhibit 25: Average directional change tick count vs. directional change threshold scaling law

	Slope	s. e.	Intercept	s. e.	Adj. R^2	MSE
EUR-CHF	2.1019	2.81E-02	9.1154	7.84E-02	0.9984	5.87E-03
EUR-USD	1.8903	1.46E-02	8.2265	4.07E-02	0.9995	1.58E-03
GBP-USD	1.9093	2.14E-02	8.1259	5.96E-02	0.9989	3.39E-03
USD-CAD	1.9030	3.71E-02	7.8638	1.03E-01	0.9966	1.02E-02
USD-CHF	1.9398	2.25E-02	8.2027	6.27E-02	0.9988	3.76E-03
USD-JPY	1.9382	2.00E-02	8.1292	5.58E-02	0.9990	2.98E-03
	1.9471	**2.40E-02**	**8.2773**	**6.68E-02**		

Exhibit 26: Estimated scaling law (iv) parameter values

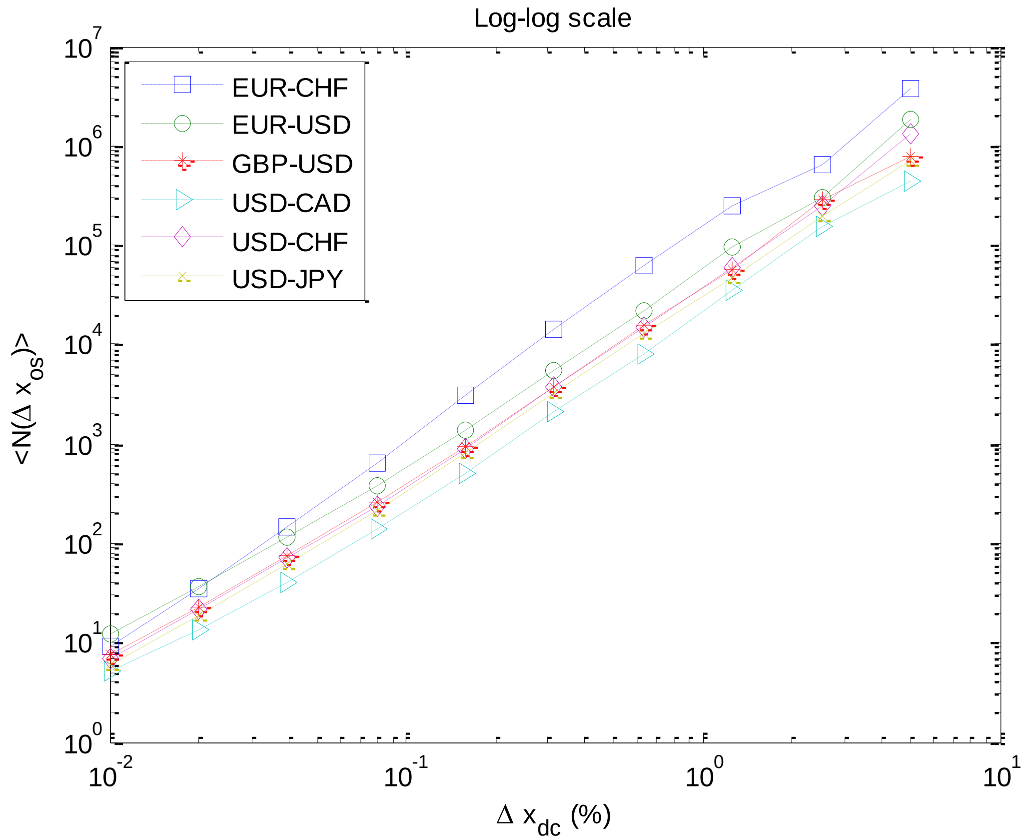

Exhibit 27: Average overshoot tick count vs. directional change threshold scaling law

	Slope	s. e.	Intercept	s. e.	Adj. R^2	MSE
EUR-CHF	2.0882	2.65E-02	9.3040	7.38E-02	0.9986	5.20E-03
EUR-USD	1.9105	3.55E-02	8.5850	9.90E-02	0.9969	9.37E-03
GBP-USD	1.9034	2.54E-02	8.3834	7.08E-02	0.9984	4.79E-03
USD-CAD	1.8870	3.57E-02	8.0929	9.95E-02	0.9968	9.47E-03
USD-CHF	1.9514	3.37E-02	8.5124	9.39E-02	0.9973	8.43E-03
USD-JPY	1.9016	1.90E-02	8.2916	5.30E-02	0.9991	2.68E-03
	1.9403	**2.93E-02**	**8.5282**	**8.17E-02**		

Exhibit 28: Estimated scaling law (v) parameter values

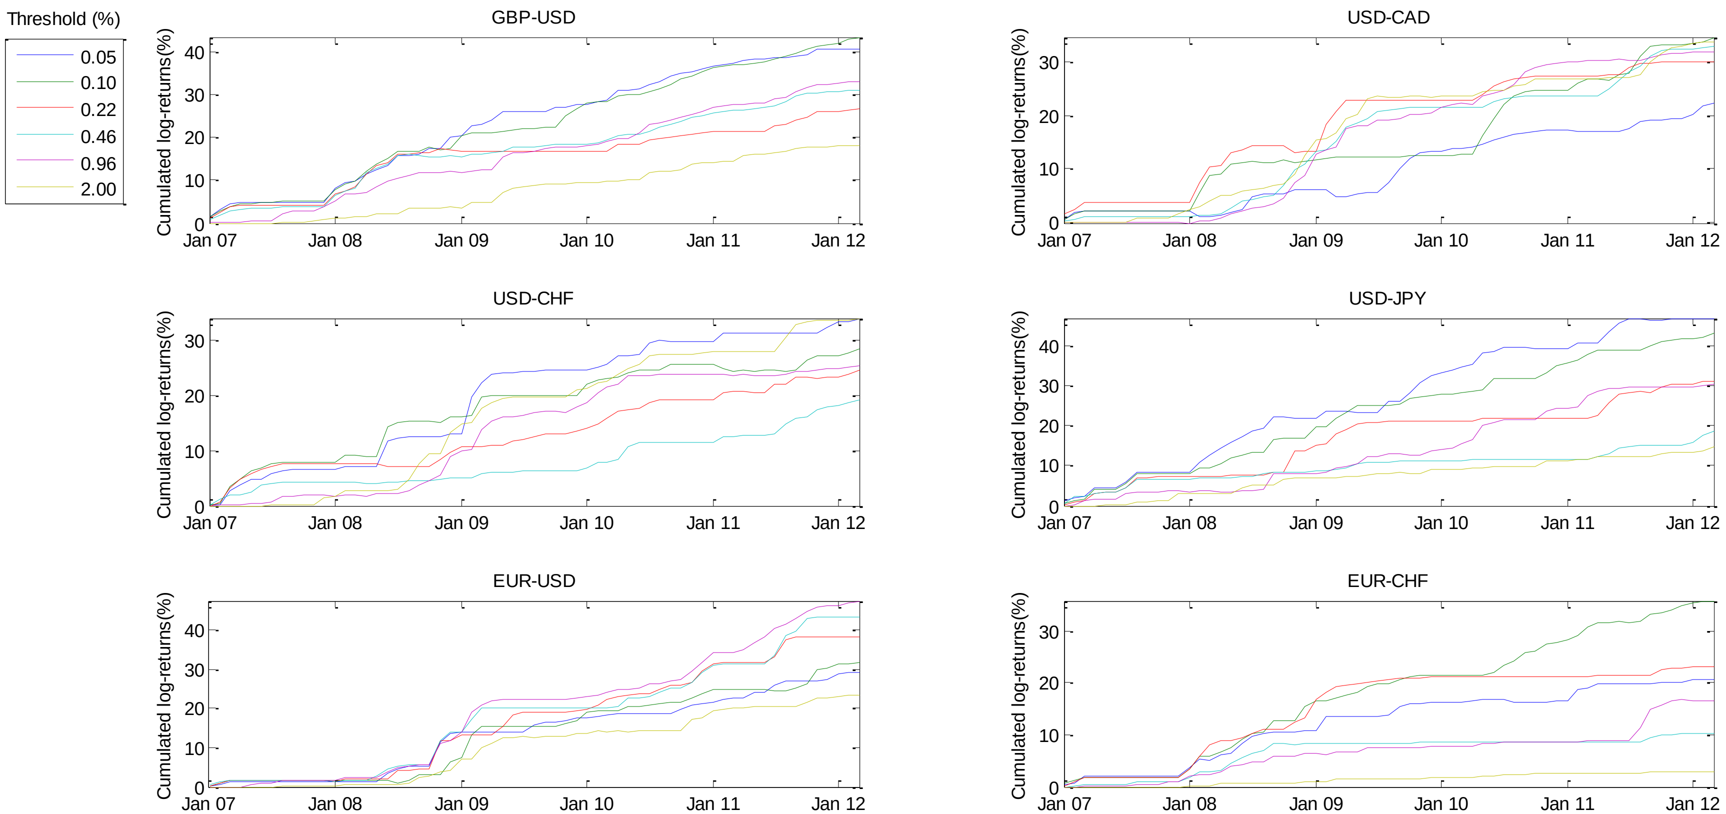

Exhibit 29: Performance of individual directional change strategies

Annualised Sharpe Ratio[7]																		
Currency pair	**GBP-USD**						**USD-CAD**						**USD-CHF**					
Threshold	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**
2007	1.85	2.41	1.72	2.25	1.36	-2.61	0.82	1.13	1.49	0.44	-15.48	-0.07	2.34	2.30	2.63	2.43	1.27	-14.53
2008	3.78	3.83	3.83	2.83	3.61	1.30	1.14	1.89	1.79	2.56	2.41	3.68	1.15	1.23	-0.05	-0.50	2.02	2.23
2009	2.17	1.94	-3.35	1.98	1.89	2.15	1.64	0.03	1.68	3.53	2.45	2.71	1.65	0.91	2.20	0.54	2.09	2.06
2010	2.99	4.32	2.18	4.88	4.11	1.90	3.91	2.54	2.09	1.15	3.42	2.11	1.77	2.54	3.57	1.99	2.42	3.19
2011	3.78	5.49	2.44	3.71	4.23	2.28	0.88	2.33	1.23	3.44	1.11	1.83	0.34	0.33	1.56	2.49	-0.30	1.60
2012	-104.70	3.79	0.77	12.27	14.98	0.91	8.64	2.60	-9.39	-5.45	-9.30	17.34	2.06	1.13	1.59	5.60	2.40	-2.36

Annualised Sharpe Ratio																		
Currency pair	**USD-JPY**						**EUR-USD**						**EUR-CHF**					
Threshold	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**
2007	2.64	2.35	2.38	2.47	1.89	0.58	0.44	0.71	0.68	0.91	0.93	-14.21	0.86	1.06	0.98	0.17	0.13	0.00
2008	3.81	2.51	1.12	0.98	0.96	1.74	1.62	0.46	1.34	1.58	1.64	1.56	3.37	3.49	3.68	3.82	2.84	-0.25
2009	2.07	2.94	2.58	1.85	1.86	0.45	1.17	1.50	1.91	1.35	1.94	2.30	1.45	3.89	2.01	-7.47	0.44	-0.42
2010	2.02	1.78	-0.45	-2.74	2.53	0.95	1.75	2.23	4.62	2.40	3.31	1.07	-1.31	2.55	-12.35	0.00	-0.11	0.05
2011	2.11	3.36	2.28	1.76	1.71	1.07	2.52	1.44	1.81	2.46	5.26	2.29	1.21	2.76	0.63	0.50	1.82	-3.02
2012	1.26	15.59	2.70	5.32	1.39	1.52	2.89	3.51	0.00	-14.06	2.56	6.30	1.43	122.24	14.86	-5.05	-9.16	-11.92

Exhibit 30: Individual strategies' annualised Sharpe ratio

[7] Computed given a risk-free rate of 1% per annum

Annualised Sortino Ratio[8]

Currency pair	GBP-USD						USD-CAD						USD-CHF					
Threshold	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**
2007	22.00	26.43	19.30	19.90	13.18	-2.38	4.85	5.72	12.98	1.21	-3.39	-0.23	45.03	61.88	48.13	29.02	7.91	-3.38
2008	102.2	23.24	149.2	17.60	56.99	6.45	2.83	12.70	6.31	25.98	24.72	NaN	20.18	11.73	-0.12	-0.86	10.08	36.61
2009	47.44	26.53	-2.47	9.46	21.86	21.21	3.91	0.06	40.97	1087.4	79.08	62.47	84.65	14.40	16.20	2.14	29.15	23.61
2010	110.3	78.93	27.31	101.3	NaN	37.84	14.51	75.88	16.43	5.77	18.36	15.86	13.14	40.86	37.90	24.82	30.76	51.92
2011	NaN	161.0	27.92	NaN	NaN	22.54	3.15	44.37	8.82	47.27	2.34	26.82	1.74	0.64	12.23	28.59	-0.44	19.43
2012	-3.46	NaN	2.23	NaN	NaN	2.89	NaN	NaN	-3.35	-3.16	-3.35	NaN	25.54	4.23	9.11	NaN	219.7	-2.41

Annualised Sortino Ratio

Currency pair	USD-JPY						EUR-USD						EUR-CHF					
Threshold	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**	**0.05 %**	**0.10 %**	**0.22 %**	**0.46 %**	**0.96 %**	**2.00 %**
2007	45.83	46.58	46.03	48.78	16.79	1.90	1.32	2.92	2.93	3.48	3.75	-3.37	4.84	4.65	4.90	0.46	0.38	-3.46
2008	21.09	20.50	30.68	2.71	6.39	16.84	43.17	0.97	28.56	23.38	29.16	13.75	51.08	75.05	86.91	26.96	23.83	-0.59
2009	17.86	54.84	40.92	19.24	8.26	0.74	6.24	42.00	35.59	23.03	47.55	54.25	18.67	19.50	95.68	-3.17	1.41	-1.01
2010	13.11	40.77	-0.79	-2.32	115.0	7.36	13.28	18.81	79.84	31.09	21.12	11.62	-1.34	28.33	-3.35	-3.46	-0.27	0.10
2011	29.77	32.99	39.89	12.08	13.53	5.43	30.17	11.41	48.43	81.45	260.9	47.99	17.41	18.71	2.86	2.11	32.23	-2.70
2012	5.18	NaN	NaN	NaN	6.39	8.03	NaN	NaN	-3.46	-3.41	NaN	NaN	6.84	NaN	NaN	-3.12	-3.35	-3.39

Exhibit 31: Individual strategies' annualised Sortino ratio

[8] Computed given a minimum acceptable return rate (MAR) of 1% per annum; NaN – not a number