



City Research Online

City St George's, University of London

Citation: Richard Maria Kos, Hartwig (2015). Momentum Effects: Essays on Trading Rule Returns in G10 Currency Pairs. (Unpublished Doctoral thesis, City University London)

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/13700/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Momentum Effects: Essays on Trading Rule Returns in G10 Currency Pairs

Chapter 1:
Momentum Effects: G10 Currency Return Survivals

Chapter 2:
Momentum Effects: G10 Currency Return Survivals,
Implications for Trading Rules

Chapter 3:
Momentum Effects: Dissecting Generic G10 Trading Rule
Returns

by

Hartwig Richard Maria Kos

Submitted in accordance with the requirements for the degree of a

Doctor of Philosophy

Cass Business School; City University

Department of Finance

June 2015

I. Abstract.....	11
II. General Introduction	12
III. Contribution to Literature	21
IV. Chapter 1: Momentum Effects: G10 Currency Return Survivals	23
A. Outline	24
1. Academic Background	24
a) Central Bank Activity.....	26
b) Data Snooping	27
c) Systematic Risk Taking	29
d) Trading Rules as Risk Factors.....	30
e) Cross Sectional Analysis of Trading Rules	31
2. Motivation of the Chapter and Main Contributions	33
B. Data and Methodology	35
1. Data, Return and Moving Average Calculations	35
a) Data Description	35
b) Return Calculation	37
c) Moving Average Signal Calculations.....	44
2. Construction of Survivorship Curves	45
C. Survivorship Analysis versus Runs Test.....	53
D. Sensitivity of Survivorship Curves to External Factors.	63
E. Resampling Techniques and Testing for Market Efficiency.....	69
1. Distributional characteristics in the underlying exchange rates.....	70
2. Serial Correlation and Heteroskedasticity in the underlying exchange rates	72
F. Empirical Results	75
1. Survivorship Curves and log-rank Test Results.....	75
2. Results for the Full Sample Period.....	78
3. Results for Sub-Sample Periods One to Nine	87
G. Trading Rule Implementation	91
H. Link between Survival Analysis and Trading Rule Results.	99
I. Conclusions	102
V. Chapter 2: Momentum Effects: G10 Currency Return Survivals:	
Implications for Trading Rules.....	105
A. Outline	106
1. Academic Background	106
a) Trading Rules in Equity Markets	106
b) Trading Rules in Foreign Exchange Markets.....	111
c) Other Explanations for Trading Rule Profitability	116
2. Motivation of the Chapter and Main Contributions	119
B. Data and Methodology	121
1. Data, Return and Moving Average Calculations	121
2. Product Limit Estimator	123
3. White's Data Snooper	124
4. Trading Rules to be Evaluated and Link to Survival Curves.....	128
5. Back-Test of Best Trading Rules	133
a) Construction of Composite Trading Rules	134
b) Accounting for Implicit Transaction Cost.....	135
a) Calculation of Breakeven Transaction Cost Levels	136

C. Empirical Evaluation	137
1. Best Absolute Data Snooping Proven Trading Rules.....	137
2. Best Relative Data Snooping Proven Trading Rules.....	145
D. Conclusions.....	156
VI. Chapter 3:	162
A. Outline.....	163
1. Academic Background	163
a) Technical Trading Rules and Systematic Risk Taking.....	164
b) Technical Trading Rules as Systematic Risk Taking.....	167
c) Technical Trading Rules and Multi Factor Models	169
2. Motivation of the Chapter and Main Contributions	173
B. Data and Methodology.....	174
1. Data, Return and Moving Average Calculations.....	174
2. Specification of Systematic Risk Factors	177
a) Factor 1: Trend	177
b) Factor 2: Momentum	178
c) Factor 3: Carry	178
d) Factor 4: Value	179
e) Factor 5: Risk Aversion	179
f) Factor 6: Volatility	180
g) Correlations across Factors.....	181
3. Specification of the Multivariate Factor Model	182
C. Empirical Evaluation	185
D. Conclusions.....	193
VII. Overall Conclusion:	196
VIII. REFERENCES:.....	202
IX. APPENDICES: Chapter 1	214
A. APPENDIX 1: Overview of FX Trading Rule Literature.....	214
B. APPENDIX 2: Descriptive Statistics Based on Log Returns	220
1. Descriptive Statistics (LOG Base Currency Returns)	220
2. Descriptive Statistics (LOG 3M T-bill Interest Rate Adj. Currency Returns)	221
3. Descriptive Statistics (LOG O/N Rate Interest Rate Adj. Currency Returns).....	222
C. APPENDIX 3: Trading Rule Analysis including Bid/Ask Spreads	223
1. Dataset 2: Descriptive Statistics (Simple Base Currency Returns)	223
2. Dataset 2: Descriptive Statistics (3M T-bill Adj. Currency Returns).....	224
3. Dataset 2: Descriptive Statistics (O/N Rate Adj. Currency Returns)	225
4. Dataset 2: Differences in IR adjusted returns (3M T-bill vs. Overnight rate).....	226
5. Dataset 2: BVTC; Enhanced vs. BM Trading Rules (Across Currencies).....	227
D. APPENDIX 4: Sensitivity of Survival times based on 36 Sub-Samples.....	228
1. Time Series Analysis of Average Survival vs External Factors.....	228
E. APPENDIX 5: GARCH (1,1) Parameters for Resampling Simulation.....	229
1. GARCH (1,1) Parameters.....	229
F. APPENDIX 6: Trading Rule Results across Base Currencies.....	230
1. Median Log Rank Values; Positive Signals; Simple Resampling.....	230
2. Median Log Rank Values; Negative Signals; Simple Resampling	231
3. Median Log Rank Values; Positive Signals; GARCH (1,1) Resampling	232
4. Median Log Rank Values; Negative Signals; GARCH (1,1) Resampling.....	233

5. Breakeven Transaction Cost; Enhanced vs. Benchmark Trading Rules	234
6. Abs. log-rank test Results vs. Benchmark Trading Rule Results	235
7. Abs. log-rank test Results vs. Enhanced Trading Rule Results	236
X. APPENDICES: Chapter 3	237
A. APPENDIX 1: OLS Regression on Systematic Risk Factors	237
1. Output of the SR1/LR5 Regression.....	237
B. APPENDIX 2: Test for Cross-Sectional Heteroskedasticity	238
1. Test results for equality of variance across currency pairs, by trading rule param.	238
C. APPENDIX 3: Correlation of OLS Residuals.....	239
1. Correlation matrix of the residuals of the SR1/LR5 OLS regression, from Appendix 1	239
D. APPENDIX 4: Test of Endogeneity of Systematic Factors	240
1. Median P-value of the Durbin-Wu-Hausman test across all analysed currency pairs	240

Chapter 1: Momentum Effects: G10 Currency Return Survivals

FIGURE 1-1: FX CAPTIAL CONTROLS AND EXCHANGE RATE REGIMES	36
TABLE 1-1: DESCRIPTIVE STATISTICS (SIMPLE BASE CURRENCY RETURNS).....	39
TABLE 1-2: DESCRIPTIVE STATISTICS (SIMPLE 3M T-BILL INTEREST RATE ADJ. RETURNS)	40
TABLE 1-3: DESCRIPTIVE STATISTICS (SIMPLE O/N RATE INTEREST RATE ADJ. RETURNS).....	41
TABLE 1-4: DIFF IN MEAN AND CORREL. BETWEEN IR ADJ. RETURNS (3M T-BILL VS. O/N RATE).....	43
FIGURE 1-2: MOVING AVERAGE COMBINATIONS.....	44
FIGURE 1-3: GRAPHICAL DESCRIPTION OF DURATION DATA	46
FIGURE 1-4: GRAPHICAL DESCRIPTION OF DURATION DATA.....	46
FIGURE 1-5: RESULTS FROM THE HYPOTHETICAL START-UP SURVEY	49
FIGURE 1-6: GRAPH. REPRESENTATION OF THE PLE FOR THE START-UP SURVEY DATA EXAMPLE.....	51
FIGURE 1-7: BASIC RUNS TEST ANALYSIS; NO ASSUMPTION ABOUT DIRECTION	56
FIGURE 1-8: AVERAGE LIFE; RUNS TEST FOR POSITIVE AND NEGATIVE RETURNS	59
FIGURE 1-9: USDGBP TRADING RULE; AVERAGE LIFE TEST FOR POSITIVE RETURNS.....	61
FIGURE 1-10: SR1/LR10 TRADING RULE: AVERAGE SURVIVAL TIME VS. KURTOSIS	65
FIGURE 1-11: SR1/LR10 TRADING RULE: AVERAGE SURVIVAL TIME VS. KURTOSIS	65
FIGURE 1-12: CROSS SECTIONAL ANALYSIS OF AVERAGE SURVIVAL VS EXTERNAL FACTORS	66
FIGURE 1-13: TIME SERIES ANALYSIS OF AVERAGE SURVIVAL VS EXTERNAL FACTORS	68

FIGURE 1-14:	
TEST FOR AUTOCORREL AND HETEROSKD IN FX RETURNS	71
FIGURE 1-15:	
EMPIRICAL PLE CURVE; POSITIVE RETURNS; USDGBP; SR1/LR10	76
FIGURE 1-16	
SIMULATED PLE POSITIVE RETURNS; USDGBP; SR1/LR10	77
FIGURE 1-17:	
LOG RANK TEST RESULTS, HEAT MAP FOR POS MAVG SIGNALS (SIMPLE RESAMPLING)	79
FIGURE 1-18:	
LOG RANK TEST RESULTS, HEAT MAP FOR NEG MAVG SIGNALS (SIMPLE RESAMPLING)	80
FIGURE 1-19:	
LOG RANK TEST RESULTS, HEAT MAP FOR POS MAVG SIGNALS (GARCH(1,1) RESAMPLING)	81
FIGURE 1-20:	
LOG RANK TEST RESULTS, HEAT MAP FOR NEG MAVG SIGNALS (GARCH(1,1) RESAMPLING)	82
FIGURE 1-21:	
LOG RANK TEST RESULTS, HEAT MAP FOR POS MAVG SIGNALS (BLOCK RESAMPLING)	83
FIGURE 1-22:	
LOG RANK TEST RESULTS, HEAT MAP FOR NEG MAVG SIGNALS (BLOCK RESAMPLING)	84
FIGURE 1-23:	
LOG RANK TEST RESULTS, HEAT MAP FOR MED VALUES OF POS. AND NEG. MAVG SIGNALS	86
FIGURE 1-24:	
SUB-SAMPLE ANALYSIS; MEDIAN LOG RANK VALUES; POSITIVE SIGNALS	90
FIGURE 1-25:	
SUB-SAMPLE ANALYSIS; MEDIAN LOG RANK VALUES; NEGATIVE SIGNALS	90
FIGURE 1-26:	
EXPOSURE LEVEL OF THE COMP/BM TRADING RULE FOR USDGBP	91
FIGURE 1-27:	
EXPOSURE LEVEL OF THE ENH AND THE BM TRADING RULE FOR USDGBP	93
FIGURE 1-28:	
USDGBP BENCHMARK TRADING RULE, ASSUMED TRANSACTIONS	95
FIGURE 1-29:	
BVTC; ENHANCED VS. BENCHMARK TRADING RULES (ACROSS CURRENCIES)	98
FIGURE 1-30:	
BVTC; ENHANCED VS. BENCHMARK TRADING RULES (ACROSS SUB-SAMPLES)	99

FIGURE 1-31:	
ABS. LOG-RANK TEST RESULTS VS. BENCHMARK TRADING RULE RESULTS.....	100
FIGURE 1-32:	
ABS. LOG-RANK TEST RESULTS VS. ENHANCED TRADING RULE RESULTS	101

Chapter 2: Momentum Effects: G10 Currency Return Survivals, Implications for Trading Rules

FIGURE 2-1: GRAPHICAL DESCRIPTION OF MOVING AVERAGE SIGNALS	122
FIGURE 2-2: PRODUCT LIMIT ESTIMATOR AND CONDITIONAL SURVIVAL PROBABILITIES OF THE PLE	124
FIGURE 2-3: BEST ABSOLUTE DATA SNOOPING PROVEN TRADING RULES FOR ALL SUB-SAMPLES	138
FIGURE 2-4: TRADING RULE PARAMETERISATION AMONGST TOP 10% OF TRADING RULES	140
FIGURE 2-5: OCCUR. OF TRADING RULE ENHANCEMENTS AMONGST TOP 10% OF TRADING RULES	141
FIGURE 2-6: CURRENCY PAIRS AMONGST TOP 10% OF TRADING RULES	143
FIGURE 2-7: BVTC; BEST DATA SNOOPING PROVEN TRADING RULES (ABSOLUTE EVALUATION)	144
FIGURE 2-8A: DIFFERENCES IN MEAN P-VALUES OF THE ENHANCED AND STANDARD TRADING RULE	147
FIGURE 2-8B: DIFFERENCES IN MEAN P-VALUES OF THE ENHANCED AND STANDARD TRADING RULE	148
FIGURE 2-8C: DIFFERENCES IN MEAN P-VALUES OF THE ENHANCED AND STANDARD TRADING RULE	149
FIGURE 2-9: BVTC; BEST DATA SNOOPING PROVEN TRADING RULES (RELATIVE EVALUATION).....	152
FIGURE 2-10: BVTC; FIRST AND SECOND ENHANCEMENT VERSUS BENCHMARK.....	153

Chapter 3: Momentum effects: Dissecting generic G10 Trading Rule Returns

FIGURE 3-1: MOVING AVERAGE COMBINATIONS.....	177
FIGURE 3-2: CORRELATION ACROSS RISK FACTORS	181
FIGURE 3-3: GMM ESTIMATION RESULTS	187
FIGURE 3-4: GMM EST. RESULTS, CONSTANT AND SENSITIVITIES OF TR. RULE RET. TO RISK FACTORS ...	190

I herewith grant powers of discretion to the University Librarian to allow the thesis to be copied in whole or in part without further reference to the author. This permission covers only single copies made for study purposes, subject to normal conditions of acknowledgement.

Abstract:

Chapter 1:

Momentum Effects: G10 Currency Return Survivals

The chapter analyses momentum effects in G10 currencies. For each of the currency crosses within the G10 universe the chapter models the “survival” probabilities of trading signals obtained from a wide set of dual crossover moving average combinations. The application of statistical tools that stem from survival time analysis sheds light on the subject of market efficiency within the currency market. Empirical momentum signals from shorter-term trading rules outlive respective benchmark signals, while longer-term moving average crossover signals have lower life expectancy than theory would suggest. Furthermore, a trading strategy constructed from a subset of short-term moving average signals exhibits clear outperformance over a trading strategy that is generically composed from all moving average crossover signals. This outperformance persists over time.

Chapter 2:

Momentum Effects: G10 Currency Return Survivals, Implications for Trading Rules

The chapter models survival probabilities of positive and negative momentum signals that are obtained from a wide set of dual crossover moving average combinations for all G10 cross currency pairs. The results of this survival analysis are used to create trading rule enhancements that aim to outperform generic dual crossover moving average trading signals. The trading rule enhancements are assessed, by applying White’s (1999) “data snooper”. The results suggest that there is scope for trading rule enhancements to outperform generic trading rules. Moreover, results present strong evidence for Lo’s (2004) Adaptive Market Hypothesis.

Chapter 3:

Momentum effects: Dissecting Generic G10 Trading Rule Returns

The chapter builds on the work of Pojarliev and Levich (2008, 2010), who dissect the returns of active currency managers by applying a multiple ordinary least squares (OLS) regression to currency fund returns. Where the chapter differs is in the specification of the dependent variable, which is in the context of the present chapter a set of trading rule parameterisations that are applied to a broad range of currency pairs. The results of this chapter suggest that there is some alpha embedded in the returns of technical trading rules. The chapter also establishes a comparatively strong positive, statistically significant link between the risk factors Trend, Momentum, Risk Aversion. The results of the chapter clearly indicate that shorter-term moving averages exhibit less systematic exposure than longer term moving averages. Other factors such as Carry, Value and Volatility have a considerably less pronounced relationship; only few factor sensitivities are statistically significant. Moreover, the results also indicate that systematic risk exposures of trend following trading strategies change with small adjustments in the design of trading rules.

II. General Introduction

Modern finance is dominated by the assumption of capital markets being priced efficiently. Present day option and asset pricing is built upon premises laid out by the “Chicago School of Thought” more than four decades ago. Albeit, the concept of market efficiency has been challenged persistently for some time now, with the majority of studies focusing on the pricing of equity markets, other asset classes have received comparatively less academic attention. This thesis looks to add to the academic literature that analyses momentum effects within the foreign exchange space. The traditional approach of detecting momentum effects as introduced by DeBondt and Thaler (1985, 1987), Jegadeesh and Titman (1993, 2001) and Rouwenhorst (1998) for equities, and Menkhoff, Sarno, Schmeling and Schrimpf (2011) for foreign exchange markets, is a filtering approach, whereby portfolios are built on the basis of historic performance of underlying assets analysed. The approach introduced in this thesis follows a somewhat different intuition. It analyses the returns generated from simple technical trading rules. Each of the three chapters of this thesis looks at trading rule returns from a slightly different perspective.

The first chapter of this thesis introduces an alternative methodology to detect market inefficiencies, based on statistical methods from lifetime statistics. The chapter analyses a dataset of daily closing prices of G10 currencies, including all of their cross currency pairs. The dataset upon which the analysis is based spans from 4th January 1974 to 31st December 2009 and contains 9025 trading days. The methodology applied is a variation of survivorship analysis, which in the context of financial data is an extension of the concept of runs tests. Both methodologies aim to compare the probability of occurrence of positive or negative return streams in an empirical time series with a theoretically derived probability. Survivorship analysis, however, is superior to the runs test, as it allows for flexibility in the definition of the benchmark processes and the time series that is investigated. The chapter assesses the lifetime characteristics of 39 moving average combinations across all currency and cross currency pairs in the data sample. Moreover, the chapter extends the work of Jochum (2000) and Kos and Todorovic (2008), who focus on equity returns in their analysis. The key findings of this chapter can be summarised as follows.

First, various exchange rates have empirical patterns that cannot be explained by any benchmark process. This suggests that empirical momentum signals live either longer or shorter than their respective benchmark signals. In the case of moving average crossover signals that utilise a set of very short-term moving average combinations, empirical signals outlive what is suggested by theory. Long-term moving average crossover signals, on the other hand, have lower life expectancy than theory would suggest. The results from a sub-sample analysis suggest that most of the deviations from market efficiency are deteriorating over time, up until the point where all momentum signals exhibit survival times that are statistically equivalent to what is suggested by benchmark processes. Second, when implementing trading rules on the same set of moving average crossover signals, it becomes evident that the profitability of a generic trading rule that incorporates all moving average signals deteriorates continuously up to the point where the trading rule becomes unprofitable. Furthermore, a trading strategy that is constructed from a subset of moving average signals, namely shorter-term moving average signals, shows clear outperformance over a trading strategy that is generically composed from all moving average crossover signals. This outperformance persists over time.

The second chapter of this thesis extends the previous chapter, which introduces survivorship analysis as an alternative methodology for detecting market inefficiencies. While the first chapter presents a simple trading strategy that outperforms a generic trading strategy, the aim of the chapter is not to search for a superior trading rule. This is where the aim of the second chapter lies. Survivorship analysis provides a wide range of information about the historic survival pattern of moving average trading signals, which can be used to establish the best exit points of a trading strategy. The key objective of this chapter is to assess whether trading rule enhancements that utilise information derived from lifetime analysis can generate returns that are superior to the returns generated from equivalent strategies that don't use such enhancements. The statistical validity of enhanced trading rule returns is ascertained by applying White's (2000) data snooping methodology. The chapter investigates four trading strategies, based on the results of the survivorship analysis. The first enhancement weighs the exposure of the trading strategy according to the historic conditional periodic survival probability of trading rule signals, as derived from calculating survivorship curves from historic data. The conditional periodic survival probability tends to start at a medium level. It then increases over time and falls off thereafter. The second trading rule enhancement weighs the periodical strategy exposure according to the unconditional survival

probability of historic trading rule signals. Hence, the exposure level decreases over time. The other two trading strategies are variations of the first two trading strategies, which aim to reduce the impact of transaction costs.

The first key conclusion that can be drawn from the second chapter is the fact that the profitability of generic trading rules diminishes over time. Moreover, the results also indicate that during the early years of the data sample, when general trading rule profitability is high, the first trading rule enhancement is able to add some value, while the second trading rule enhancement doesn't. However, this changes in the latter parts of the sample period where the first enhancement fails to add value, while the second enhancement performs strongly. The results of the chapter indicate that trading rule returns exhibit two distinct regimes, suggesting that foreign exchange markets have changed over time. This observation goes hand in hand with Lo's (2004) Adaptive Market Hypothesis. Simple trading rule strategies, which were once profitable, fail to deliver positive returns in more recent years, as market participants arbitrage the excess returns away. The results of this chapter suggest that trading strategies that are enhanced by applying survival probabilities to the exposure levels of returns are able to add value versus a standard trading rule.

The third chapter analyses the drivers of trading rule returns more closely. It aims to shed light on whether the returns derived from applying generic technical trading rules embed any compensation for systematic risk taking. Many studies that look at returns from technical trading rules merely point out that trading strategies are implemented on a long-short basis, therefore they tend to have market covariance levels that are close to zero. The chapter proposes to analyse simple trading rule returns for systematic risk factors in a broader way, thereby eliminating some of the criticisms of earlier studies. Assessing whether trading rule returns are a compensation for risk taking is undoubtedly an academically valid path to follow. However, as pointed out by Neely and Weller (2011) it is heavily dependent on the construction of a convincing model for the risk premium.

In the spirit of Lo's (2004) Adaptive Market Hypothesis market participants have to continuously adapt to a changing market environment in the foreign exchange space. As a consequence they have learned to exploit the reinforcing link between trading rules and market trends as proposed by Schulmeister (2006), or they have learned to bear the risks associated with carry strategies as suggested by Brunnermeier, Petersen and Nagel

(2008). Therefore, exploiting relationships such as trend and carry have in fact become a legitimate way of harvesting risk premia, and indeed many of the newer studies in the field of foreign exchange markets such as Pojarliev and Levich (2008, 2010), treat anomalies such as trend or carry, and many others, as a risk premium strategy in their own right.

Given these developments in the foreign exchange space, this chapter looks to assess technical trading rule returns against a set of multiple risk factors such as Trend, Momentum, Carry, Valuation, Risk Aversion and Volatility. The chapter represents an extension of the work proposed by Pojarliev and Levich (2008, 2010), whereby a wider universe of factors is used. The main difference lies in the fact that the proposed chapter looks to analyse systematic risk exposures of simple technical trading rules as opposed to the returns of active currency managers.

The results make it evident that factors such as Trend and Momentum and Risk Aversion have a relatively strong positive and statistically significant impact on trading rule returns. It should be noted, however, that this is less the case for shorter term moving averages, while as longer term moving averages exhibit more systematic exposure. Other factors such as Carry, Value and Volatility have a considerably less pronounced relationship to trading rule returns. The results of this chapter make a strong case for the fact that at least a part of the returns from technical trading rules are driven by systematic factors. Paired with the finding that shorter term moving averages exhibit higher levels of alpha, the results in this chapter would suggest that shorter term moving averages are less affected by systematic risk factors than it is the case for longer moving averages, reaffirming the findings of the first and second chapter.

When looking at all of the results from the three chapters in combination, a series of observations can be made. One of the overall conclusions that can be drawn from this thesis is the fact that the profitability of generic trading rules diminishes over time. This goes for deviations from market efficiency. As mentioned earlier, this is not the case for the returns of shorter-term moving averages, which remain generally higher, even when more generic trading rules fail to perform. Both results should not come as a surprise.

With regards to the deterioration in generic trading rule performance one could point towards increased competition amongst trend following trading strategies, as well as a general change in foreign exchange markets. Indeed the market environment within the foreign exchange space has changed considerably in recent years. The amount of assets

under management in systematic trading strategies has quadrupled in the last thirty years¹. Moreover, the composition of market participants and also the way currency markets are accessed have changed in recent years. Since 1998 the Bank of International Settlements publishes a triennial survey of foreign exchange and derivatives market activity. Besides this survey a series of working papers are published that shed light on the drivers in the change of trading volume. Galati and Melvin (2004) highlight the significant growth in the participation of Hedge Funds in particular trend following strategies, which have grown considerably. They also make the point that the landscape of Hedge Funds has changed as well. Systematic trading funds that entered the market more recently are typically smaller than the trend following funds that had been there before. They also use algorithms that are much shorter-term in their nature than what has been used before. Galati and Heath (2007) reiterate the aspect of Hedge Fund participation in their review of the years from 2004 to 2007. Moreover they also point towards the algorithmic trading as one of the key sources of turnover within foreign exchange markets. King and Rime (2010) estimate that, within foreign exchange markets high, frequency trading takes up to 25% of the volume of all spot transactions worldwide. Following the intuition of Lo's (2004) Adaptive Market Hypothesis, markets behave in an evolutionary fashion in the sense that they continuously change and asset prices are driven by the nature and preferences of market participants. Hence, as a specific investment style becomes popular, the profitability of that particular investment style deteriorates as new market participants jump on the bandwagon.

With respect to the consistent performance of shorter-term moving average rules, one could put forward the lack of a short-term valuation framework for exchange rates. Unlike other assets such as equities or bonds, exchange rates cannot easily be priced upon the principals of fundamental fair values. Most stocks for instance have dividend streams. Hence, the dividend yield gives a timely signal to investors, indicating whether a stock is cheap or expensive. If the dividend yield is very high then the stock is cheap and investors will start buying the stock and by doing so the price of a stock increases. This brings down the dividend yield to a point where the stock is no longer deemed to be cheap. The reverse happens if the dividend yield is too low. While exchange rates do have some valuation metrics, such as the purchasing power parity. The relationship to these metrics is rather loose. The purchasing power parity follows the logic of the law of one price. This means that under the assumption of no trade barriers, equivalent

¹ See: http://www.barclayhedge.com/research/indices/cta/mum/CTA_Fund_Industry.html

goods have to be priced equivalently. Hence, any discrepancy between price levels in two countries is adjusted via the exchange rate. This absolute framework can also be expressed in relative terms whereby any differential in inflation rates between countries has to be reflected in the exchange rate of the two countries. However, for various reasons such as the fact that there might be trade barriers between countries, that there might be differences in investor preferences for countries, or the fact that some goods may not be tradable in some countries, exchange rates can stay over or undervalued against their purchasing power parity for several years, in some cases even decades. Given this inability of the purchasing power parity to capture short-term movements, investors rely on other parity relationship such as the interest rate parity, to gauge short-term movements in exchange rates. The interest rate parity follows the same logic of the law of one price, applied to financial assets. Hence, if interest rates in one country are higher than interest rates in another country, the exchange rate of this country should automatically appreciate and vice versa. While this approach has biases in its own right, it provides a better proxy for the direction of movements in exchange rates than the purchasing power parity. The main disadvantage of this approach is the fact that the interest rate differential is to a great part set by central bank policy, which does not necessarily change as a result of the valuation of an exchange rate. Hence, while the dividend yield of a stock comes down as the stock price goes up, an appreciation of a currency is never automatically linked to a narrowing in the interest rate differential. The fact that this link between valuation and interest rate differential is missing makes it difficult for investors to assess whether exchange rates are short-term over or undervalued. The profound implication of this missing anchoring device is the fact that exchange rates are prone to trend much more than other financial assets. This is particularly beneficial for short-term moving average trading rules, given that they can adapt very quickly to trend changes, allowing these trading rules to benefit from very sharp reversals and relatively high day to day volatility, potentially explaining the strong results of short-term moving average trading rules.

One of the key observations of the second chapter is the fact that during the early years of the data sample, when the general level of trading rule profitability is high, the scope for trading rule enhancements to outperform is somewhat limited. This, however, changes as the level of general trading rule profitability deteriorates. Namely, one of the trading rule enhancements proposed in the second chapter of this thesis exhibits results that point towards differing market regimes within the observation time period. The

trading rule enhancement mentioned weighs its exposures according to the historical unconditional survival probability of moving average crossover trading rule signals. Therefore, given the fact that the absolute survival probability decreases rapidly over time, the level of exposure decreases rapidly over time as well. Hence, the strategy is shorter-term in its nature relative to a moving average trading rule that does not apply such exposure adjustment. In the early years of the data sample, when the general trading rule profitability is high, the mentioned trading rule enhancement severely underperforms a generic trading strategy. However, that changes in the latter parts of the sample period when simple trading rules fail to deliver positive returns. The results of the second chapter clearly point to a regime change in foreign exchange markets. Moreover, the results also indicate that the regime change seems to be in favour of technical trading rules that are shorter-term in their nature.

These observations underpin the framework of an evolutionary market environment as proposed by Lo (2004). Lo's (2004) Adaptive Market Hypothesis suggests that while the returns that were available from applying simple technical trading rules have widely been arbitrated away, other opportunities in form of trading rule enhancements arise. As pointed out earlier, Lo's (2004) Adaptive Market Hypothesis follows an evolutionary concept of market equilibrium as opposed to a steady state assumption. In that sense it incorporates a continuously changing market environment in which market inefficiencies arise, which are subsequently arbitrated away as market participants become aware of them. This suggests that the Adaptive Market Hypothesis is a more complex framework than Fama's (1969) Efficient Market Hypothesis. The former captures the dynamics of market cycles in as far as it allows different assets to be subject to different levels of efficiency at any given time. Lo (2004) argues that asset prices are driven by the nature and preferences of market participants. He also points out that there are different groups of investors have very distinct investment patterns and investment preferences. If one or many of these investment "species" find interest in one specific asset, the pricing of this asset becomes more efficient and vice versa. As a consequence, investment strategies can undergo different stages in which they show different levels of profitability. The results of the second chapter indicate that.

Summing up the results of the first two chapters. Chapter one suggests that the profitability of a generic trading rule deteriorates continuously up to the point where the trading rule becomes unprofitable. This however is not the case for shorter-term moving average signals, which show good results, even when the broader trading rule universe

fails to deliver positive returns. In light of the fact that the deviations from normality of the survival rates of shorter-term moving average trading rules dissipate over time, and also in light of the fact that trading rule enhancements, which focus on shortening exposure times, only start performing strongly at a time when deviations from market efficiency, as established in chapter one, are supposedly negligible. The persistence in performance of shorter-term moving average trading signals begs the question to which degree systematic risk factors have an impact on trading rule profitability.

The results of the third chapter shed light on this aspect. Similarly to the first chapter, the third chapter indicates that shorter-term moving averages deliver higher alpha than longer-term moving averages. Moreover, as indicated in the first chapter, returns from shorter-term moving averages can also be less explained by systematic risk adjustments than it is the case for longer-term moving averages. When looking at the results of the third chapter, it becomes evident that very short term focused trading rules such as the SR1/LR5 rule are only influenced by the factor Risk Aversion in a statistically significant way. Other factor such as Trend and Momentum do exhibit positive sensitivities to the very short-term technical trading rules analysed, but not in a statistically significant way. Given that the Risk Aversion factor is a purely heuristic factor, and that it impacts very short term focussed trading rules is a valuable insight into market psychology. The fact that shorter-term moving average trading rules systematically benefit from continuously switching between positive and negative exposures could mean that the combined effect of the lack of an anchoring device within the short-term valuation framework of foreign exchange markets and the increased use of heuristics to create short term valuation frameworks, might well explain this statistical relationship. Hence, one might deduct that some of the very short term trading rule returns are genuinely driven by human spirit and market inefficiency as opposed to risk taking. Nonetheless, when looking at some of the other results of the third chapter it becomes evident that many other short-term trading rules exhibit statistically significant relationships with systematic risk factors. Moreover, the results also indicate that risk exposures also vary depending on small changes in parameterisations and design of trading rules.

While the first chapter provides strong evidence that trading rule returns are to a great extent driven by market inefficiencies during the first part of the sample period. During the latter parts of the sample period, particularly the returns of shorter-term moving average trading rules, as shown in chapter two, exhibit much better results. Moreover,

the fact that shorter-term moving averages, except for the very short term trading rules such as the SR1/LR5 trading rule, are also subject to a fair degree of systematic risk taking would indicate that systematic risk factors have at least some impact on trading rule returns.

III. Contribution to Literature

The first chapter follows Jochum (2000) and Kos and Todorovic (2008) closely. Jochum (2000) bases his study on daily closing prices of the main equity indices in the United States, the United Kingdom and Switzerland. The sample period used spans from 1973 to 1997 and comprises 6523 data points. His study indicates that the returns of the three indices analysed exhibit microstructural pattern that cannot be explained by benchmark processes. Namely that positive and negative momentum streams live longer than what theory would suggest. Kos and Todorovic (2008) confirm Jochum's (2000) findings. Their study is also based on daily data. However, for S&P sector indices and corresponding ETFs from 1998 to 2006. The results of the study suggest that various sectors show significant deviations from normality, which can be exploited by simple trading rules. Both of these papers compare empirically observed sequences of positive or negative returns to sequences derived from simulated return time series. The simulated return time series are designed to replicate the return generating process of asset prices under the assumption of market efficiency. Under such assumption, the sequences of the empirical time series should not systematically outlive or desist earlier than the sequences of simulated return time series. The intuition of both papers is very similar to "Runs Tests" introduced by Fama (1965). The present chapter extends Jochum (2000) and Kos and Todorovic (2008) in various ways. Firstly, it applies the methodology to foreign exchange markets, which has not been done before. Secondly, while previous papers analyse simple return time series, the present chapter looks at trading signals derived from dual crossover moving average signals, which tend to live longer and provide more meaningful results. Moreover, the chapter uses a longer data sample, which allows sub-sampling. Finally, most of the results in the present chapter are based on non-parametric simulation methods, mitigating the risk of parameter estimation errors.

The second chapter of this thesis uses White's (2000) data snooping methodology to assess whether the information provided by survivorship analysis can be used to design trading rule enhancements that outperform generic trading rules. Survivorship analysis as introduced in the first chapter offers a wide range of information about historic survival pattern of moving average trading rules. This information can be used to enhance trading rules by altering exposure levels of the trading strategy in line with historic survival probabilities. Moreover, it can also be used to identify best exit points

for a trading signal. While Qi and Wu (2006) apply White's (2000) data snooping methodology within the context of foreign exchange markets, the academic contribution of this chapter can be summarised as follows. Firstly, the chapter analyses the performance of enhancements of moving average crossover trading rules, as opposed to picking the best trading rule out of a heterogeneous universe of trading rules. Secondly, the chapter undertakes an extended analysis of sub-samples, facilitating the analysis of persistency in performance of single trading rules. This has been done in previous studies as well, however, to a much lesser extent. Finally, the chapter proposes to look at the results of White's data snooping test in a relative context as opposed to an absolute context. The results of the sub-sample analysis suggest that there is a great deal of clustering of currency pairs amongst the top trading rules over time. Therefore the chapter looks at average White's statistics for single trading rule parameterisations across all currency pairs. This has not been done before.

The third chapter establishes a framework, to assess whether trading rule returns can be explained by systematic risk factors. The methodology applied in this chapter is inspired by the work proposed by Pojarliev and Levich (2008, 2010), who dissect the returns of active currency managers by applying a multiple OLS regression to currency fund returns. The chapter represents an extension of the work proposed by Pojarliev and Levich (2008, 2010), difference between their work and the chapter lies in the fact that the chapter looks to analyse systematic risk exposures of simple technical trading rules as opposed to the returns of active currency managers. Therein lies the first academic contribution. The second academic contribution lies in the test setup. While it is appropriate to run a series of independent multiple regressions for currency fund managers, which are following different investment strategies. Such test setup is not appropriate in the context of technical trading rules that use similar parameterisations such as SR1/LR5 or SR1/LR30, across a range of currency crosses. This is due to the fact that there is a high likelihood of commonalities between the SR1/LR5 trading rule for the USD/GBP and the USD/EUR cross. In order to account for these cross currency commonalities the proposed framework is based on a one step GMM model, which allows the calculation of the general sensitivity of the specific risk factor to the universe of trading rules that are calculated for each trading rule parameterisation. Finally, the chapter also it also looks at a wider set of risk factors than the work of Pojarliev and Levich (2008, 2010). Factors included are Trend, Momentum, Carry, Valuation, Risk Aversion and Volatility. This has also not been attempted before.

IV. Chapter 1:

Momentum Effects: G10 Currency Return Survivals

Abstract

The chapter analyses momentum effects in G10 currencies. For each of the currency crosses within the G10 universe the chapter models the “survival” probabilities of trading signals obtained from a wide set of dual crossover moving average combinations. The application of statistical tools that stem from survival time analysis sheds light on the subject of market efficiency within the currency market. Empirical momentum signals from shorter-term trading rules outlive respective benchmark signals, while longer-term moving average crossover signals have lower life expectancy than theory would suggest. Furthermore, a trading strategy constructed from a sub set of short-term moving average signals exhibits clear outperformance over a trading strategy that is generically composed from all moving average crossover signals. This outperformance persists over time.

A. Outline

1. Academic Background

The notion of efficient markets has surfaced in various forms and shapes throughout the twentieth century. The first formal definition of the concept, however, is given by Fama, who introduces three forms of market efficiency in his groundbreaking paper in 1970. Weak market efficiency suggests that all information that is contained in historical prices is fully reflected in current prices. Semi-strong market efficiency suggests that all publicly available information is fully reflected in current prices. And the strong market efficiency suggests that all publicly and privately available information is fully reflected in current prices. After this categorisation of market efficiency, it has taken fifteen years for academically meaningful papers, which challenge Fama's proposition, to appear. DeBondt and Thaler (1985, 1987), Jegadeesh and Titman (1993) and Rouwenhorst (1998) give significant evidence that stock prices do not move in a random fashion, as suggested by the Efficient Market Hypothesis. While DeBondt and Thaler (1985, 1987) report a significant reversal of previously outperforming stocks over a time horizon of three to five years, Jegadeesh and Titman (1993) find that over an intermediate horizon of three to twelve months, on average past winners continue to outperform past losers. Both studies are carried out on the US equity market, and have become the cornerstones of modern behavioural finance. Upon their findings a whole school of thought and a vast body of academic literature analysing various aspects of equity momentum has developed.

When it comes to the academic debate about market efficiency, the foreign exchange market had initially been neglected. This can be explained by two factors. Firstly, up until the early seventies, currencies were either on the gold standard, or in a state of hyperinflation, or in a fixed currency regime. Secondly, during the regime of floating currencies the degree of speculation persistently outweighs the degree of trade activity by many times. Hence, currency markets were regarded as being in a state of speculative efficiency. Consequently, early research focuses on the notion of speculative efficiency. Froot and Thaler (1990) point out that two schools of thought have developed within that. Friedman (1953) argues that speculators are in the market to make profits, hence their aim is to buy a currency cheap and sell it expensive. This should ensure that exchange rates reflect the fundamental or long run determinants of currency values. Nurske (1944) on the other hand points out that excess volatility

induced by currency speculation imposes large costs on producers and consumers, who as a consequence make less efficient capital allocation decisions. Therefore, speculation drives markets away from fundamentals. While the end of Breton Woods in 1973 would have given researchers an opportunity to investigate these initially very theoretical approaches further, only a handful of academic papers published before the eighties analyse currency market efficiency. Examples of that are Dooley and Shafer (1976, 1983) and Rogoff (1979). However, some of these papers such as Rogoff (1979) remain unpublished. Furthermore, Neely and Weller (2011) suggest that early papers that focus on trading strategies such as Dooley and Shafer (1976, 1984) had at the time of their publication been widely dismissed by academics. Nonetheless, erratic US dollar moves in the eighties brought the currency market efficiency debate into mainstream academia. Since then research on currency market efficiency has mushroomed. Researchers have generally focused on either analysing the causes of the forward discount bias or analysing technical trading rules and the source of their profitability.

While the forward discount bias is a very well documented phenomenon in modern finance, the analysis of technical trading rules is less well researched. The centre of gravity within the academic debate regarding the forward discount bias is the degree to which the carry derived from this effect is attributable to an unobservable time varying risk premium, or whether there is an element of market irrationality embedded in this phenomenon. Froot and Thaler (1990) review the literature up until the nineties. They suggest that evidence for both schools of thought can be found. Despite the relative ambiguity of the sources of the forward discount bias, the “carry” phenomenon has very quickly found its way into the finance industry as well as mainstream academic research. More recent papers such as Poljarliev and Levich (2008, 2010) use carry and other phenomena as distinctive style benchmarks, with which they assess the relative performance of active currency managers.

As pointed out earlier within the area of trading rule (“trend”) research there is a vast body of literature when it comes to equity markets, while technical trading rule research in the foreign exchange space is considerably less researched. The most noteworthy early studies are Dooley and Shafer (1976, 1984) and Logue and Sweeney (1977), both papers indicate very strong returns from applying simple filter rules. As suggested by Neely and Weller (2011), the main criticism of these studies is the fact that the observation time periods of the analysis are short and seem somewhat spurious to the academic establishment. Subsequent papers such as Sweeney (1986, 1988) and Levich

and Thomas (1993) have received wider academic attention. Subsequent research has developed in various branches, Neely and Weller (2011) suggest that traditionally three theories have been put forward to explain the apparent success of technical trading rules. The first line of arguments is based on the activity of central banks. Some of the rationale for that has been given earlier in this section. The other two theories have focused on the possibility of data snooping and systematic risk taking. More recent studies either use some of the “apparent inefficiencies” of foreign exchange markets as systematic risk factors, or seem to be more focussed on how to exploit diminishing trading rule returns. While conventional trading rule research is mostly focussed on time series analysis, some of the very recent academic papers also employ a cross sectional analysis of trading rule returns. The following sections give an overview of the different lines of research.

a) Central Bank Activity

When it comes to central bank activity, Taylor (1982) suggests that central bank activities, which are aimed to support currencies, also called "leaning against the wind" policies, were unprofitable during the 1970's and that half of the losses can be attributed to speculative positions aiming to benefit from central bank intervention. These early results are confirmed by Szakmary and Mathur (1997), who present strong evidence that market operations by central banks are indeed, key drivers of trading rule profitability. They show that trading against central bank intervention can yield significant excess returns. Their findings are based on a sample of five currencies² versus the US Dollar from 1977 to 1991. LeBaron (1999) confirms Szakmary and Mathur's (1997) results. His results suggest that the trading rule profits are highest during periods of central bank intervention. When removing the time periods where central banks are active in the currency market, the results are insignificant.

These observations stand in sharp contrast to the findings of Neely (2002), who analyses intraday data for five currency pairs³. The exact data range of each of the currency pairs varies slightly, but overall his study covers a time range from the early to mid-eighties to the mid to late nineties. At the outset of his analysis he documents three key points that define some of the common characteristics of central bank intervention.

² Deutsch Mark (DEM), Canadian Dollar (CAD), Japanese Yen (JPY), Swiss Franc (CHF), Pound Sterling (GBP)

³ AUD/USD, CHF/USD, DEM/USD, JPY/USD

Firstly, some central banks are more active in the currency market than others. However, if a central bank starts intervening, the likelihood of further intervention is relatively high, which means that intervention exhibits positive serial correlation. Moreover, intervention patterns cluster together. Finally, Neely (2002) also points out that the volume of intervention is generally very low compared to the general flow in these currencies. This observation already weakens the argument that trading rule profitability is merely a function of central bank intervention. His analysis is based on a 150 day moving average trading rule, which gives a buy signal if the current price is above the 150 day moving average and a sell signal if the current price is below the 150 day moving average. In a first step he compares the moving average trading results for a data sample that contains intervention dates and a data sample that does not contain intervention dates. The results of this test point are very similar to the results of LeBaron (1999). Nonetheless, Neely (2002) points out that many of the currency pairs analysed, do exhibit positive and statistically significant returns, even after removing intervention days. Neely (2002) also looks at the intraday return realisations from the moving average rule before and after central banks have intervened. The key finding of this analysis is that intervention does not generate returns itself. Currency intervention comes as a reaction to strong and very profitable short-term trends.

b) Data Snooping

Data Snooping is the second argument that is usually put forward to explain trading rule profits. Data snooping suggests the possibility that trading rules might be selected with a selection bias. Hence, certain rules are chosen that work well for one specific dataset, but might not work for any other set of data. For that reason early studies such as Dooley and Shafer (1984) and Sweeney (1986) focus on the most common and widely used trading rules in order to minimise this selection bias. Later studies such as Levich and Thomas (1993) utilise simulation techniques to establish an appropriate benchmark for their trading rules. In their paper they investigate a set of five currencies⁴ against the US Dollar, over a time period from 1976 to 1990. The study is based on daily closing settlement prices for currency futures contracts. Levich and Thomas (1993) test a set of nine trading rule signals, six of which are filter rules, and three of which are moving

⁴ Deutsch Mark (DEM), Canadian Dollar (CAD), Japanese Yen (JPY), Swiss Franc (CHF), Pound Sterling (GBP)

average rules. To evaluate the significance of these trading rules Levich and Thomas (1993) use a bootstrapping simulation technique. The bootstrap is a non-parametric simulation approach that re-samples an existing time series multiple times and hereby facilitates conducting standard statistical tests and inferences. The intuition behind this methodology is to evaluate the performance of the trading signals by applying them to a set of time series that have been created by resampling the original currency time series. This leads to a distribution of hypothetical trading rule returns against which the realised trading rule return is assessed. 25 of the tested filter rules and 14 of the 18 tested moving average rules offer results that suggest a statistically significant deviation from normality. The Levich and Thomas (1993) paper is insofar noteworthy as it introduces the idea of using resampling techniques, which has subsequently become somewhat of a benchmark methodology to assess the performance of trading rules. Nonetheless, their paper still focuses on a fairly narrow range of trading rules and does not eliminate the pre selection problem. Two lines of research have developed that look to eliminate the effects of the pre-selection bias, or data snooping. One line focuses on the development genetic programs and neuronal networks that naturally “grow” the best trading rules, while the other line looks to eliminate the effects of data snooping by incorporating them into a test statistic against which trading rules are tested.

Neely, Weller and Dittmar (1997) apply a genetic program that searches for an optimal trading rule. The paper is based on prices for six currency pairs⁵. The time sample spans from 1974 to 1995. It is split into three sub-periods, which constitute selection, training and testing period for the genetic code. One of the key findings of their study is the fact that different currency pairs produce higher trading returns than others. Furthermore, different currencies pairs also favour different sets of trading rules. Another main conclusion of the paper is the fact that all of the genetically grown trading rules show out of sample profitability. This holds even against bootstrapped benchmark simulations. Other more recent studies such as Evans, Pappas and Xhafa (2013) extend this line of research by applying artificial neural networks and genetic algorithms to intra day data for the GBP/USD, EUR/GBP and EUR/USD rates. Their results indicate that foreign exchange rates are not randomly distributed and that trading strategies built on their models produce very high, statistically significant returns after accounting for transaction cost. Qi and Wu (2006) apply a methodology previously introduced by White (2000) and Sullivan, Timmerman and White (1999) that allows the identification

⁵ GBP/USD, CHF/USD, DEM/USD, JPY/USD, DEM/JPY, GBP/CHF

best technical trading rules without the effect of data snooping to a series of foreign exchange rates. Their sample spans over a time period from April 1973 to December 1998. The results suggest that the best performing trading rules, according to the data snooping algorithm, are short-term channel breakout rules for the Japanese Yen and the Swiss Franc and short-term moving averages for the other currency pairs. Moreover, the results indicate that the performance of the best data snooping proven trading rules are in the range of 2.14% to 11.46% after accounting for transaction cost. However, on an out of sample basis the statistical significance and the profitability of trading rule returns have diminished considerably. These results are confirmed by the study of Kuang, Schoeder and Wang (2014), which undertakes a comprehensive examination of the profitability of technical trading rules emerging market exchange rates. While single trading rules indicate very strong profitability on an ex ante basis. On an ex post basis, once the data snooping bias is taken into account overall trading rule returns are negligible.

c) Systematic Risk Taking

The question as to whether systematic risk taking as opposed to true market inefficiency is the driver of technical trading rules returns, is the third key area analysed by academia. While some early studies such as Sweeney (1986), Taylor (1992) as well as Neely, Weller and Dittmar (1997), allow for risk adjustments. Kho (1996) is the first study that explicitly focuses on the aspect of systematic risk and trading rule returns. Moreover, his analysis allows for a time variation in the risk premium. Kho (1996) evaluates a set of moving average crossover rules with weekly data on foreign currency futures contracts from 1980 to 1991 for five different currencies.⁶ The results of the paper indicate that the profitability of the trading strategies tested is roughly on the same level as indicated by similar studies; however they also suggests that the returns have been obtained by systematic risk taking. Kho (1996) compares the results from the actual trading rules to results obtained from simulations that aim to replicate the historic evolution of time varying risk premia. Kho's analysis suggests that the model, which does not allow for time variation in the price of risk performs significantly worse than the actual trading rules. The other models that allow for time variations in the price of risk show similar results to the actual trading results, suggesting that the trading signals

⁶ British pound (GBP), Deutsche mark (DEM), Japanese yen (JPY), Swiss franc (CHF)

are correlated with the time-varying expected returns, and that the abnormal returns are close to zero. Other, papers such as Wang (2004) look at currency returns from a market microstructural perspective. His study incorporates the interaction between hedgers and speculators when designing tests of foreign exchange market efficiency. In that context Wang (2004) indicates that the aspect of hedging pressure has to be considered when analysing risk factors in the foreign exchange space. When doing so there is strong evidence that speculator profits are largely explained by risk premia.

d) Trading Rules as Risk Factors

Schulmeister (2006) argues that while traders do not follow technical signals, they monitor them frequently. By doing so, they are altering market behaviour in as far as traditional price discovery is violated, having implications on the link between trading rules and currency volatility, as well as systematic risk. The results of his study suggest that there is a pronounced feedback mechanism between trading rules and movements in the underlying exchange rates triggering a multiplier effect that is linked to technical trading rules, which translates small news flows into a market trend. Poljarliev and Levich (2008) contribute to this line of thought in as far as they establish a universe of four of currency benchmark strategies against which they compare various currency fund managers. Poljarliev and Levich (2008) highlight that the factors carry, trend, value and volatility explain 66% of returns of currency fund managers over their sample period. Poljarliev and Levich (2010) suggest that the volatility and correlation characteristics of currencies change if investors flock into one or another trading strategy. One of the examples that they give is the high correlation between the GBP/CHF cross and the NZD/JPY cross where there is no economic reason as to why these two currency pairs should be highly correlated. The only similarity that those two crosses share is the fact that GBP and NZD are traditionally high yielding currencies, while CHF and JPY are historically low yielding currencies. Poljarliev and Levich (2010) also indicate that the level of investor preference changes over time, their results show that crowdedness of carry strategies was very high in 2007 and 2008, while trend crowdedness was almost nonexistent in early 2008 with a strong pickup in the months after. Poti, Levich and Pattioni (2014) confirm the view that currency markets evolve over time. They suggest that currency predictability also changes over time. While the key finding of the paper is that the efficient market hypothesis does not hold, another notable aspect of their study is the fact that the trading rules, picked by their algorithm,

which they label “rational trading rules” tracked by popular proxy indices for carry and momentum. This strong relationship leads Poti, Levich and Pattioni (2014) to the conclusion that technical trading represents heuristics that allow portfolio managers to exploit mispricing relative to rational expectations.

e) Cross Sectional Analysis of Trading Rules

More recently the cross sectional approach by Jegadeesh and Titman (1993) has been adapted to foreign exchange research. Okunev and White (2003) and Chong and Ip (2009) make a contribution to this line of research. Their studies are based on a range of short-term and long-term moving average trading signals, which are periodically ranked according to the strength of the signal, from which long-short portfolios are implemented. In the case of Okunev and White (2003) this trading strategy yields excess returns over a specified the benchmark approximately 5%-6% after transaction cost per year. Chong and Ip (2009), who apply the strategy to emerging market currencies, indicate similar, trading cost adjusted results. More recent studies such as those by Burnside, Eichenbaum, and Rebelo (2011) and Menkhoff, Sarno, Schmeling and Schrimpf (2011) extend the cross sectional approach introduced by Okunev and White (2003). The study by Menkhoff, Sarno, Schmeling and Schrimpf (2011) replicates the traditional cross sectional momentum literature pioneered by Jegadeesh and Titman (1993), using foreign exchange data. Their sample consists of cross sectional data of 48 countries over a time period from January 1976 to January 2010, with markets being included in sample as they become available. In the spirit of Jegadeesh and Titman’s (1993) work, they create winner and loser portfolios on the basis of best and worst performing currency pairs over pre specified time periods. Their findings suggest that some combinations earn unconditional average excess returns of up to 10% per year. Moreover they confirm the results of the studies of Jegadeesh and Titman (1993, 2001), which find that momentum returns in equity markets go through different stages over time. At the point of the signal generation, returns are weak; later they become more pronounced and then fade away. Tajaddini and Crack (2012) also apply this cross sectional approach to emerging market currencies indicating that long-short momentum strategies deliver modest gains before accounting for transaction cost. After accounting for transaction cost, their results appear either negative or statistically insignificant.

The academic consensus views with regards to technical trading rule returns can be summarised as follows. Trading rule returns were high in the 70's and 80's even on a risk adjusted basis. These returns diminished considerably during the 90's and the period after 2000. Indicating that as markets have developed, the effect of increased competition has made them more efficient. However, some of the recent studies such as Poti, Levich and Pattioni (2014) suggest that in the most recent period currency predictability has decreased, and trading rule returns have started to increase again. When it comes to the argument of central bank intervention as source of trading rule profitability, early studies were clearly in favour of such argumentation, however Neely (2002) has presented strong evidence against that. His findings have not been challenged thereafter.

With regards to data snooping there are two lines of research. One focuses on the development genetic programs and neural networks, and the other looks to eliminate the effects of data snooping by incorporating them into a test statistic against which trading rules are tested. Both lines of research provide fair evidence against the argument of data snooping being the source of excess returns in technical trading rules. However, it has to be said that the second line of research, provides very weak out of sample evidence. The evidence for systematic risk taking is generally stronger than it is the case for the other lines of argument. However, Neely (2011) makes the point that if trading rule profitability was down to pure harvesting of risk premia, how are diminishing trading rules during the 90's and the 00's explainable. The new lines of research, that look at trading rule dynamics as source of risk also present strong evidence of such relationship and give grounds to Lo's (2004) Adaptive Market Hypothesis argument of an evolutionary path of foreign exchange markets. Finally, academic consensus regarding cross sectional trading rule analysis point towards diminishing trading rule returns in the most recent time period, with emerging market currencies being the main driver of returns. For an overview of the trading rule research in the foreign exchange space please refer to Appendix 1

2. Motivation of the Chapter and Main Contributions

The aim of this chapter is to analyse data dependencies and patterns in historic currency time series data, with the aim to implement trading rules that lead to abnormal currency returns that cannot be explained by any systematic risk taking. In that sense the paper challenges the notion of efficient markets, in particular the weak form of market efficiency, as outlined by Fama (1970). Although capital markets are generally regarded as weak and semi-strong efficient, currency markets, for reasons discussed earlier, seem to defy the market efficiency model persistently. The present chapter looks to extend previous research by not only analysing the returns achievable from implementing moving average rules, it also aims to show how long the signals of such moving average rules tend to persist. This, as shown in the second chapter, proves useful when designing trading rule enhancements. The present chapter will follow Jochum (2000) and Kos and Todorovic (2008) closely. Jochum (2000) bases his study on daily closing prices of the main equity indices in the United States, the United Kingdom and Switzerland. The sample period used spans from 1973 to 1997 and comprises 6523 data points. His study indicates that the returns of the three indices analysed exhibit microstructural pattern that cannot be explained by benchmark processes, namely that positive and negative momentum streams live longer than what theory would suggest. Kos and Todorovic (2008) confirm Jochum's (2000) findings. Their study also uses daily data, but for S&P sector indices and corresponding ETFs from 1998 to 2006. The results of the study suggest that various sectors show significant deviations from normality that can be exploited by simple trading rules.

Both these papers are based on the idea that under the notion of weak market efficiency, empirical equity returns should follow a random pattern. Hence, positive or negative returns of an empirical return time series should not systematically "outlive" positive or negative returns created from a random return time series. In that sense both papers can be seen as an extension of "Runs Tests" introduced by Fama (1965). The essence of "Runs Test" studies is to compare empirical return pattern to some pre-specified benchmark return pattern. While Fama (1965) compares the ratio of positive to negative returns with some theoretically derived value, Jochum (2000) and Kos and Todorovic (2008) utilise the Product Limit Estimator, which allows them to compare empirical survivorship curves to Monte Carlo simulated survivorship curves.

The Product Limit Estimator has been introduced by Kaplan and Meier (1958), it is a non-parametric measure that allows to estimate the probability of the length of survival of positive and negative return streams. Although Jochum (2000) and Kos and Todorovic (2008) offer an attractive alternative to the traditional Jegadeesh and Titman (1993, 2001) methodology, their methodology comes with some deficiencies in the implementation. Firstly, their papers analyse time series of returns. This has the disadvantage that all their momentum signals are very short lived, and quite difficult to interpret. Furthermore, there is no sub-sampling of the data and the results are highly dependent on the market environment.

A rampant bull market, such as the tech bubble in the late nineties, during a relatively short sample period, as is the case in Kos and Todorovic (2008), might have quite profound implications on the results. Another clear disadvantage of their methodology is that they utilise a Monte Carlo simulation, which is based on a standard asset price process such as ARMA (1,1). The main problem with this is the fact that one has to make assumptions about the distributional characteristics of the underlying data time series. Given the fact that the product limit estimator is non-parametric and the benchmark process is calculated based on a parametric model, the simulation process is prone to suffer from estimation errors. Furthermore, Kos and Todorovic (2008) evaluate the deviation of empirical survivorship curves from benchmark survivorship curves, by comparing average survival times. This is a very crude way of measuring differences between survivorship curves.

This present chapter applies the Jochum (2000) and Kos and Todorovic (2008) methodology to the currency space. This has not been done before. It also aims to improve previous work in various ways. First of all it analyses a long time period, which allows sub-sampling. In addition to that, the chapter defines a set of moving average pairs from which momentum signals are generated. Not only does this facilitate a wider breadth of the momentum analysis, but it also allows momentum signals to be longer and more interpretable. Furthermore, resampling is used as a simulation methodology to establish a set of benchmark survivorship curves. The advantage here is the fact that the resampling simulation itself can be constructed in a non-parametric framework hence it allows the to be free from any distributional assumptions. The disadvantage of this approach is the fact that a mere reshuffling of returns will break the volatility structure of a time series. This might raise questions on the appropriateness of the non-parametric resampling approach. In order to control for this aspect, the present

chapter also applies a resampling simulation that is based on an GARCH (1,1) process, which is based on the original historic time series of the various currencies. It also conducts a stationary bootstrap simulation proposed by Politis and Romano (1994) to control for potential autocorrelation in the underlying data. Finally, the chapter estimates the significance of the difference between empirical and benchmark survivorship curves by applying the Wilcoxon log-rank test, which is a standard test in lifetime statistics.

In summary, the present chapter extends previous work in various ways. Firstly, it applies the methodology to the foreign exchange space, which has not been done before. Secondly, while previous papers analyse simple return time series, the present chapter looks at trading signals derived from dual crossover moving average signals, which tend to live longer and provide more meaningful results. Furthermore, the chapter uses a longer data sample, which allows sub-sampling. Finally, most of the results in the present chapter are based on non-parametric simulation methods, mitigating the risk of parameter estimation errors.

B. Data and Methodology

1. Data, Return and Moving Average Calculations

a) Data Description

The dataset used for the empirical validation of the survivorship model contains daily New York closing mid-prices for G10 currencies. The sample spans from the 4th of January 1974 to the 31st of December 2009. After adjusting for non-trading days, the sample contains 9025 data points. Given the long history of the dataset, some adjustments have to be made. The EUR rate does not have a lifetime history that goes back to the mid-seventies. Hence, the sample is backfilled with the historic Deutschmark (DEM) rate, whereby the original EUR fixing rate of 1.95583 DEM per 1 EUR (as of 1 January 1999) is applied. Moreover the use of a foreign exchange data sample requires a brief discussion about the aspect of different exchange rate regimes as well as capital controls, which have been in place for many exchange rates over parts of the data sample. This discussion is undertaken at a later stage.

However, for the purpose of data description Figure 1-1 shows capital controls, as well as the different exchange rate regimes prevalent in different countries over the sample period.

The chapter aims to mimic the returns obtainable from a futures based trading strategy, hence the daily currency returns have to be interest rate adjusted. This is done, by using daily closing rates for 3-month T-bills of the respective countries. While it is appreciated that 3-month T-bills are only an approximation for the more appropriate, overnight rate. The decision to take the second best rate was made on the basis that a clean dataset, without the need for backfilling data points could be obtained.

FIGURE 1-1: FX CAPITAL CONTROLS AND EXCHANGE RATE REGIMES

Country	Controls on Foreign Exchange Capital	Managed Float	Free Float
United States	1963-1973	-	from 1974
United Kingdom	until 1979	-	from 1974
Japan	until 1980	-	from 1974
Germany (Later Europe)	until 1958	until 1979 participated in currency snake; until 1999 participated in EMS	from 1999
Switzerland	-	-	from 1974
Norway	until 1989	until 1978 participated in currency snake; until 1990 linked to a trade weighted currency basket; until 1992 linked to ECU	from 1992
Sweden	until 1989	until 1976 participated in currency snake; until 1991 linked to a trade weighted currency basket; until 1992 linked to ECU	from 1992
Canada	until 1951	-	from 1974
Australia	until 1983	until 1976 peg to effective exchange rate; until 1983 crawling peg to effective exchange rate	from 1984
New Zealand	until 1984	until 1979 peg to effective exchange rate; until 1983 crawling peg to effective exchange rate	from 1984

To verify that this approximation is leading to equivalent results as an interest rate adjustment based on a time series with overnight rates that are backfilled⁷ with 3 month

⁷ For the USD an overnight rate is available from 02.01.1975 onwards, before that a 3-month T-bill is used.

For the GBP an overnight rate is available from 02.01.1975 onwards, before that a 3 month T-bill is used.

For the JPY an overnight rate is available from 04.01.1982 onwards, before that a 3 month T-bill is used.

For the CHF an overnight rate is available from 02.01.1975 onwards, before that a 3-month T-bill is used.

For the NOK an overnight rate is available from 01.04.1997 onwards, before that a 3-month T-bill is used.

For the SEK an overnight rate is available from 01.04.1997 onwards, before that a 3-month T-bill is used.

For the CAD an overnight rate is available from 02.01.1975 onwards, before that a 3-month T-bill is used.

For the AUD an overnight rate is available from 01.04.1997 onwards, before that a 3-month T-bill is used.

For the NZD an overnight rate is available from 01.04.1997 onwards, before that a 3-month T-bill is used.

T-Bill rates, return calculations for the two time series adjustments are compared. The dataset is split into nine sub-samples, whereby the first eight sub-samples consist of exactly 1000 observations and the ninth sub-sample consists of 1025 observations. This split of sub-samples has been done completely agnostically of any underlying time period (and potential monetary regime). The reason for the almost equal split is the fact that each of the sub-samples will show similar levels of statistical confidence, given the equal amount of data analysed⁸. The dataset captures almost the entire regime of floating currencies since the mid-seventies until the end of 2009. It is designed to analyse the long-term behaviour of moving average rules. All calculations are carried out on all available cross spot exchange rates for each of the G10 currencies, whereby the currency crosses are obtained from the dollar crosses of each of the other G10 currencies. Both datasets are obtained from Factset, Datastream and Bloomberg.

b) Return Calculation

While the later parts of this chapter are based on simple return calculations, Equations 1 to 4 introduce both, simple and log returns. Moreover, descriptive statistics are calculated for both specifications. All exchange rates are expressed as units of domestic currency versus one unit of foreign currency. In order to obtain simple base currency $R_{B,t}$ returns, the return calculation as per Equation 1 is applied. S_t is hereby the currency spot price at time t .

$$(1) \quad R_{B,t} = \frac{S_t}{S_{t-1}} - 1,$$

Equation 2 shows the return calculation for $LR_{B,t}$, which are base currency returns in log terms.

$$(2) \quad LR_{B,t} = \ln(S_t) - \ln(S_{t-1})$$

⁸ This chapter uses a second dataset which contains daily New York closing mid-prices for G10 currencies, including bid/ask spreads for each of the currency crosses, whereby all the bid/ask spreads of the non-dollar crosses have been synthetically created from dollar crosses. It spans from the 27th of March 2002 to 31st of December 2009. This time period coincides with the last two sub-samples of the first dataset. The aim of using this dataset is to facilitate the analysis of the trading profitability of moving average rules, after accounting for Bid/Ask spreads. The results of this analysis are presented in APPENDIX 3

As mentioned earlier, the aim of this chapter is to replicate returns available from a futures trading strategy, which means that returns net of the interest rate differential have to be analysed. Throughout the sample period, various currencies have had significant interest rate differentials, which would potentially distort the results. In order to control for this, simple interest rate adjusted returns $R_{I,t}$ are calculated as per equation 3

$$(3) \quad R_{I,t} = \left[\left(\frac{1+r_f}{1+r} \right) * \left(\frac{S_t}{S_{t-1}} \right) \right] - 1,$$

The first term represents the daily interest rate differential between foreign (r_{ft}) and domestic (r_t) currencies. The second term of Equation 3 shows the return from currency appreciation. Equation 4 shows the same return calculation in log returns.

$$(4) \quad LR_{I,t} = \ln(1 + r_{ft}) - \ln(1 + r_t) + \ln(S_t) - \ln(S_{t-1})$$

Equations 3 and 4 are both based on the Money Market Basis convention (Actual/360). The adjusted return time series, obtained from both equations, result in approximate currency returns that can be earned by following a futures based investment strategy.

As mentioned earlier calculation of the interest rate differential is based upon the three month T-Bill rate. Using the 3-month T-bill rate represents a clean way of adjusting currency returns for the interest differential. This is due to the fact that for T-bill rates there are consistent time series available, which span over the sample period. Hence, there is no need for backfilling the data with other interest rate proxies, which have a longer time series available. However, as indicated earlier, using T-bill rates to calculate the interest rate differential is academically not fully appropriate, due to the fact that 3 month T-bills might have a slight term premium embedded in the rate, which might bias the interest rate adjustment. The appropriate measure in theory is the overnight rate, for which no clean time series is available. Hence there is a trade off between the quality of historic data and the bias due to a potential term premium. Hence, this section presents descriptive statistics for both adjustment factors and ascertains that both are equivalent. Tables 1-1, 1-2 and 1-3 show descriptive statistics for all cross currency pairs, whereby the currencies along the columns are base currencies and the currencies along the rows are foreign currencies. The first number of each of the currency pair blocks is the daily mean the second number is the daily standard deviation. Numbers 3, 4 and 5 are Skew,

Kurtosis and the p-value of the Jarque Bera test. The results of the Jarque Bera indicate a fair degree of non-normality in the data. All of the calculations are based on simple return calculations, the same analysis for log returns can be found in Appendix 2.

TABLE 1-1: DESCRIPTIVE STATISTICS (SIMPLE BASE CURRENCY RETURNS)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean Ret. (%)		0.006%	-0.011%	-0.004%	-0.007%	0.004%	0.009%	0.002%	0.007%	0.010%
Std. Dev. (%)		0.623%	0.689%	0.668%	0.755%	0.687%	0.712%	0.413%	0.746%	0.820%
Skew		0.183	-0.364	0.072	0.085	0.349	2.084	-0.153	3.779	4.325
Kurtosis		7.883	7.282	8.796	8.738	11.975	50.668	16.350	94.224	104.510
JB p-value		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
GBP										
Mean Ret. (%)	-0.002%		-0.015% **	-0.009% *	-0.012% **	-0.001%	0.004%	-0.001%	0.003%	0.006%
Std. Dev. (%)	0.622%		0.728%	0.503%	0.600%	0.553%	0.607%	0.648%	0.790%	0.836%
Skew	-0.054		-0.464	-0.511	-0.261	0.368	3.637	-0.077	2.624	3.737
Kurtosis	7.933		9.144	12.573	12.008	16.724	98.625	6.509	67.308	88.247
JB p-value	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
JPY										
Mean Ret. (%)	0.015% **	0.020% ***		0.009%	0.006%	0.017% **	0.022% ***	0.017% **	0.022% **	0.025%
Std. Dev. (%)	0.691%	0.731%		0.670%	0.698%	0.739%	0.784%	0.793%	0.946%	0.979%
Skew	0.494	0.642		0.478	0.262	0.859	2.290	0.415	2.601	2.899
Kurtosis	7.636	9.750		9.549	8.887	13.887	46.617	8.723	47.667	53.839
JB p-value	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
EUR										
Mean Ret. (%)	0.009%	0.012% **	-0.005%		-0.003%	0.008% **	0.014% ***	0.009%	0.014% *	0.017%
Std. Dev. (%)	0.668%	0.504%	0.669%		0.364%	0.419%	0.496%	0.687%	0.828%	0.869%
Skew	0.081	0.686	-0.314		0.221	1.851	7.051	0.035	2.610	3.429
Kurtosis	8.522	13.383	9.145		39.699	42.044	212.356	7.473	58.339	75.763
JB p-value	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
CHF										
Mean Ret. (%)	0.013% *	0.016% ***	-0.001%	0.004%		0.013% **	0.018% ***	0.014% *	0.018% **	0.021%
Std. Dev. (%)	0.755%	0.601%	0.697%	0.364%		0.543%	0.609%	0.783%	0.917%	0.950%
Skew	0.086	0.457	-0.099	0.199		1.143	3.976	0.070	2.100	2.761
Kurtosis	8.429	12.132	8.790	39.587		22.104	99.582	7.353	42.571	55.632
JB p-value	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%
NOK										
Mean Ret. (%)	0.001%	0.005%	-0.012%	-0.007%	-0.010% *		0.006%	0.002%	0.006%	0.009%
Std. Dev. (%)	0.686%	0.552%	0.734%	0.416%	0.540%		0.472%	0.685%	0.817%	0.866%
Skew	-0.130	-0.118	-0.601	-1.401	-0.834		4.452	-0.254	2.605	3.319
Kurtosis	11.539	15.620	12.650	37.480	20.004		163.479	9.944	63.440	76.518
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%
SEK										
Mean Ret. (%)	-0.004%	0.000%	-0.016% **	-0.011% **	-0.014% **	-0.004%		-0.003%	0.002%	0.004%
Std. Dev. (%)	0.704%	0.597%	0.772%	0.482%	0.597%	0.464%		0.696%	0.830%	0.887%
Skew	-1.251	-2.323	-1.487	-4.935	-2.677	-2.602		-1.475	1.614	2.356
Kurtosis	36.985	69.722	32.938	155.501	71.159	131.005		38.385	74.248	83.318
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%
CAD										
Mean Ret. (%)	0.000%	0.006%	-0.011%	-0.005%	-0.008%	0.003%	0.008%		0.006%	0.009%
Std. Dev. (%)	0.414%	0.649%	0.790%	0.687%	0.782%	0.687%	0.705%		0.697%	0.795%
Skew	0.346	0.185	-0.237	0.101	0.083	0.444	2.327		3.312	4.504
Kurtosis	16.995	6.610	8.571	7.843	7.673	10.873	53.627		84.396	112.281
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%
AUD										
Mean Ret. (%)	-0.002%	0.003%	-0.013%	-0.007%	-0.010%	0.000%	0.005%	-0.001%		0.005%
Std. Dev. (%)	0.729%	0.777%	0.927%	0.814%	0.903%	0.804%	0.825%	0.684%		0.701%
Skew	-2.398	-1.542	-1.680	-1.640	-1.296	-1.546	-0.076	-2.105		2.887
Kurtosis	56.402	39.235	31.075	34.556	26.219	37.679	58.088	52.625		179.449
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%
NZD										
Mean Ret. (%)	-0.003%	0.001%	-0.015%	-0.009%	-0.012%	-0.002%	0.003%	-0.003%	0.000%	
Std. Dev. (%)	0.798%	0.816%	0.956%	0.849%	0.930%	0.847%	0.876%	0.773%	0.696%	
Skew	-2.715	-2.321	-1.869	-2.163	-1.726	-2.013	-0.603	-2.840	0.482	
Kurtosis	60.387	50.877	33.028	44.074	32.730	45.238	61.242	64.039	165.191	
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	

The dataset consists of daily returns for individual currencies from the 4th of January 1974 to the 31st of December of 2009. The column labels denote base currency calculations and row labels denote foreign currency returns against the base currency. For each currency pair, mean return, standard deviation, as well as skew and kurtosis of daily returns are shown. A Jarque Bera test for normality is conducted. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

TABLE 1-2: DESCRIPTIVE STATISTICS (SIMPLE 3M T-BILL INTEREST RATE ADJ. RETURNS)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean Ret. (%)		0.000%	-0.003%	-0.003%	-0.001%	-0.002%	0.006%	-0.002%	0.000%	0.000%
Std. Dev. (%)		0.623%	0.689%	0.668%	0.755%	0.687%	0.712%	0.413%	0.746%	0.821%
Skew		0.179	-0.361	0.070	0.084	0.341	2.076	-0.146	3.768	4.310
Kurtosis		7.877	7.275	8.796	8.741	11.951	50.606	16.338	94.095	104.391
JB p-value		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
GBP										
Mean Ret. (%)	0.004%		-0.001%	-0.002%	0.000%	0.000%	0.007%	0.001%	0.002%	0.002%
Std. Dev. (%)	0.622%		0.728%	0.503%	0.600%	0.553%	0.607%	0.648%	0.790%	0.837%
Skew	-0.051		-0.468	-0.510	-0.260	0.361	3.626	-0.077	2.623	3.726
Kurtosis	7.927		9.141	12.545	11.994	16.678	98.437	6.504	67.401	88.247
JB p-value	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
JPY										
Mean Ret. (%)	0.008%	0.006%		0.003%	0.004%	0.005%	0.012%	0.006%	0.007%	0.007%
Std. Dev. (%)	0.691%	0.731%		0.671%	0.699%	0.739%	0.784%	0.793%	0.946%	0.979%
Skew	0.491	0.645		0.471	0.258	0.858	2.292	0.418	2.602	2.894
Kurtosis	7.627	9.749		9.534	8.880	13.864	46.593	8.722	47.705	53.841
JB p-value	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
EUR										
Mean Ret. (%)	0.007%	0.004%	0.001%		0.001%	0.002%	0.009% *	0.004%	0.005%	0.005%
Std. Dev. (%)	0.668%	0.504%	0.669%		0.364%	0.419%	0.496%	0.687%	0.828%	0.869%
Skew	0.082	0.684	-0.307		0.222	1.847	7.052	0.036	2.608	3.416
Kurtosis	8.520	13.349	9.132		39.654	41.940	212.280	7.475	58.328	75.684
JB p-value	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
CHF										
Mean Ret. (%)	0.007%	0.004%	0.001%	0.000%		0.002%	0.009%	0.004%	0.006%	0.005%
Std. Dev. (%)	0.755%	0.601%	0.698%	0.364%		0.544%	0.609%	0.783%	0.917%	0.950%
Skew	0.087	0.456	-0.096	0.198		1.138	3.973	0.070	2.099	2.753
Kurtosis	8.431	12.119	8.786	39.549		22.048	99.457	7.357	42.547	55.593
JB p-value	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%
NOK										
Mean Ret. (%)	0.006%	0.003%	0.001%	0.000%	0.001%		0.008% *	0.003%	0.004%	0.004%
Std. Dev. (%)	0.686%	0.552%	0.734%	0.416%	0.541%		0.472%	0.686%	0.817%	0.866%
Skew	-0.122	-0.112	-0.600	-1.398	-0.829		4.448	-0.248	2.606	3.314
Kurtosis	11.519	15.582	12.631	37.390	19.957		163.232	9.931	63.392	76.492
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%
SEK										
Mean Ret. (%)	-0.001%	-0.003%	-0.006%	-0.007%	-0.006%	-0.006%		-0.004%	-0.003%	-0.003%
Std. Dev. (%)	0.704%	0.597%	0.772%	0.482%	0.597%	0.464%		0.696%	0.830%	0.887%
Skew	-1.243	-2.313	-1.489	-4.937	-2.674	-2.599		-1.469	1.616	2.351
Kurtosis	36.943	69.583	32.921	155.441	71.068	130.806		38.386	74.265	83.294
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%
CAD										
Mean Ret. (%)	0.004%	0.003%	0.001%	0.001%	0.002%	0.002%	0.009%		0.003%	0.003%
Std. Dev. (%)	0.414%	0.649%	0.791%	0.687%	0.783%	0.687%	0.705%		0.697%	0.795%
Skew	0.339	0.185	-0.240	0.101	0.083	0.438	2.321		3.304	4.489
Kurtosis	16.983	6.605	8.567	7.847	7.679	10.856	53.625		84.363	112.202
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%
AUD										
Mean Ret. (%)	0.006%	0.004%	0.002%	0.001%	0.003%	0.002%	0.009%	0.002%		0.002%
Std. Dev. (%)	0.729%	0.778%	0.927%	0.814%	0.903%	0.804%	0.825%	0.684%		0.702%
Skew	-2.387	-1.540	-1.680	-1.638	-1.295	-1.548	-0.077	-2.096		2.879
Kurtosis	56.310	39.292	31.098	34.547	26.206	37.651	58.104	52.596		179.241
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%
NZD										
Mean Ret. (%)	0.007%	0.005%	0.003%	0.002%	0.004%	0.003%	0.011%	0.004%	0.003%	
Std. Dev. (%)	0.798%	0.816%	0.957%	0.849%	0.930%	0.847%	0.876%	0.773%	0.696%	
Skew	-2.700	-2.310	-1.863	-2.149	-1.718	-2.008	-0.598	-2.825	0.486	
Kurtosis	60.309	50.878	33.028	44.019	32.705	45.222	61.211	63.989	164.947	
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	

The dataset consists of daily returns for individual currencies from the 4th of January 1974 to the 31st of December of 2009. The column labels denote base currency calculations and row labels denote foreign currency returns against the base currency. For each currency pair, mean return, standard deviation, as well as skew and kurtosis of daily returns are shown. A Jarque Bera test for normality is conducted. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

TABLE 1-3: DESCRIPTIVE STATISTICS (SIMPLE O/N RATE INTEREST RATE ADJ. RETURNS)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean Ret. (%)		0.000%	-0.003%	-0.001%	0.002%	0.000%	0.007%	0.000%	0.001%	0.001%
Std. Dev. (%)		0.623%	0.689%	0.668%	0.755%	0.687%	0.712%	0.413%	0.746%	0.821%
Skew		0.179	-0.359	0.072	0.085	0.342	2.075	-0.154	3.767	4.310
Kurtosis		7.876	7.273	8.796	8.745	11.951	50.620	16.335	94.085	104.392
JB p-value		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
GBP										
Mean Ret. (%)	0.004%		-0.001%	0.000%	0.002%	0.001%	0.008%	0.003%	0.003%	0.003%
Std. Dev. (%)	0.622%		0.728%	0.503%	0.600%	0.553%	0.607%	0.648%	0.790%	0.837%
Skew	-0.051		-0.466	-0.507	-0.259	0.364	3.627	-0.077	2.625	3.726
Kurtosis	7.927		9.139	12.532	11.995	16.670	98.447	6.500	67.447	88.253
JB p-value	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
JPY										
Mean Ret. (%)	0.008%	0.006%		0.004%	0.007%	0.006%	0.013% *	0.007%	0.008%	0.008%
Std. Dev. (%)	0.691%	0.731%		0.671%	0.699%	0.739%	0.784%	0.793%	0.946%	0.979%
Skew	0.489	0.643		0.471	0.257	0.857	2.289	0.416	2.601	2.894
Kurtosis	7.624	9.746		9.537	8.874	13.866	46.597	8.717	47.698	53.844
JB p-value	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
EUR										
Mean Ret. (%)	0.005%	0.003%	0.000%		0.003%	0.002%	0.009% *	0.004%	0.005%	0.005%
Std. Dev. (%)	0.668%	0.504%	0.669%		0.364%	0.419%	0.496%	0.687%	0.828%	0.869%
Skew	0.081	0.681	-0.307		0.223	1.847	7.051	0.035	2.608	3.416
Kurtosis	8.520	13.334	9.135		39.697	41.936	212.299	7.474	58.331	75.692
JB p-value	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
CHF										
Mean Ret. (%)	0.004%	0.001%	-0.002%	-0.001%		0.001%	0.007%	0.003%	0.003%	0.003%
Std. Dev. (%)	0.755%	0.601%	0.698%	0.364%		0.543%	0.609%	0.783%	0.917%	0.950%
Skew	0.086	0.455	-0.094	0.196		1.138	3.970	0.070	2.098	2.750
Kurtosis	8.436	12.120	8.780	39.596		22.064	99.392	7.357	42.525	55.544
JB p-value	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%
NOK										
Mean Ret. (%)	0.004%	0.002%	-0.001%	0.000%	0.002%		0.008%	0.003%	0.003%	0.003%
Std. Dev. (%)	0.686%	0.552%	0.734%	0.416%	0.540%		0.472%	0.686%	0.817%	0.866%
Skew	-0.123	-0.115	-0.599	-1.397	-0.830		4.446	-0.249	2.606	3.314
Kurtosis	11.519	15.573	12.633	37.386	19.971		163.250	9.928	63.396	76.501
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%
SEK										
Mean Ret. (%)	-0.002%	-0.004%	-0.007%	-0.007%	-0.004%	-0.006%		-0.003%	-0.003%	-0.003%
Std. Dev. (%)	0.704%	0.597%	0.772%	0.482%	0.597%	0.464%		0.696%	0.830%	0.887%
Skew	-1.242	-2.313	-1.486	-4.936	-2.672	-2.597		-1.470	1.617	2.352
Kurtosis	36.954	69.587	32.921	155.454	71.018	130.817		38.388	74.274	83.299
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%
CAD										
Mean Ret. (%)	0.002%	0.002%	-0.001%	0.000%	0.003%	0.002%	0.008%		0.002%	0.002%
Std. Dev. (%)	0.414%	0.649%	0.791%	0.687%	0.783%	0.687%	0.705%		0.697%	0.795%
Skew	0.347	0.185	-0.237	0.102	0.082	0.439	2.322		3.306	4.489
Kurtosis	16.984	6.600	8.563	7.846	7.677	10.852	53.630		84.381	112.153
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%
AUD										
Mean Ret. (%)	0.005%	0.003%	0.001%	0.002%	0.005%	0.003%	0.010%	0.003%		0.002%
Std. Dev. (%)	0.729%	0.778%	0.927%	0.814%	0.903%	0.804%	0.825%	0.684%		0.702%
Skew	-2.386	-1.542	-1.680	-1.637	-1.295	-1.548	-0.078	-2.098		2.878
Kurtosis	56.304	39.319	31.093	34.547	26.196	37.652	58.109	52.611		179.234
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%
NZD										
Mean Ret. (%)	0.006%	0.004%	0.002%	0.003%	0.006%	0.004%	0.011%	0.004%	0.003%	
Std. Dev. (%)	0.798%	0.816%	0.957%	0.849%	0.930%	0.847%	0.876%	0.773%	0.696%	
Skew	-2.700	-2.310	-1.864	-2.149	-1.715	-2.008	-0.599	-2.826	0.487	
Kurtosis	60.308	50.880	33.029	44.024	32.675	45.226	61.214	63.957	164.947	
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	

The dataset consists of daily returns for individual currencies from the 4th of January 1974 to the 31st of December of 2009. The column labels denote base currency calculations and row labels denote foreign currency returns against the base currency. For each currency pair, mean return, standard deviation, as well as skew and kurtosis of daily returns are shown. A Jarque Bera test for normality is conducted. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

Table 1 bases its calculations on regular currency returns, which are not interest rate adjusted. The results of this analysis suggest that exchange rates do show statistically

significant long-term trends when they are not adjusted for the interest rate differential. For example GBP has depreciated against JPY over the observation time period by around -0.015% per day. When annualised with 250 trading days, this translates into roughly 3.75% per annum. This number is statistically significant. Tables 2 and 3, show the same analysis for interest rate adjusted returns. While Table 2 shows the interest rate adjusted returns using 3-month T-bill rates, Table 3 shows the same analysis on the basis of the overnight rate. Using the previous example again Tables 2 and 3 suggest that after the interest rate adjustment using either the 3 month T-bill rate or the overnight rate, the depreciation of GBP against the JPY has only been -0.001% per day, which is -0.25% when annualised. This is not statistically significant.

To ascertain that the interest rate adjustment with a 3 month T-rate is equivalent to the interest rate adjustment with overnight rates. This section conducts statistical tests for differences in the distributional characteristics of the currency time series under both interest rate adjustments. Moreover, the correlation between both interest rate adjusted currency return time series is calculated. Table 1-4 shows the output of this analysis. Currencies along the columns are base currencies and the currencies along the rows are foreign currencies. The first number of each of the currency pair blocks is the daily mean under the 3-month T-bill adjustment. The second number is the mean under the overnight rate adjustment. The following two numbers are p-values for various test specifications of mean equality tests for the two time series. The first of these four tests is a simple t-test. The second test is the Satterthwaite-Welch t-test, which allows for inequality in variances between the different time series. Given the fact that Tables 1-1, 1-2 and 1-3 indicate strong levels of non normality in the currency return time series, these two mean tests, which are parametric, hence assuming a normal distribution of the underlying data, are likely to be biased. Therefore the next two tests conducted are, nonparametric tests, which are not influenced by the distributional characteristics of the underlying time series. The Wilcoxon/Mann-Whitney test is a rank based test that analyses the distributional equality of two time series. The Wilcoxon/Mann-Whitney test adjusted for ties, corrects for observations that take the same value in both time series. The last number in each of the blocks shows the correlation between the two interest rate adjusted currency time series. The results in Table 1-4 make it evident that the interest rate adjustment using the 3-month T-bill is appropriate. None of the test results suggest a statistically significant difference in distributional characteristics. Moreover, the correlations across the interest rate adjusted time series are in the range of 0.999 to 1. Given these results, it is fair to conclude that both interest rate

adjustments are equivalent. The aspect of biases due to the fact that the 3 month T-bill rate might be subject to a term premium can be neglected. Hence, in the remainder of the study the 3 month T-bill rate is used for calculating interest rate adjusted returns.

TABLE 1-4: DIFFERENCES IN MEAN AND CORRELATION BETWEEN INTEREST RATE ADJUSTED RETURNS (3M T-BILL VS. OVERNIGHT RATE)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean (3M Tbill adjustment)		0.000%	-0.003%	-0.003%	-0.001%	-0.002%	0.006%	-0.002%	0.000%	0.000%
Mean (O/N Rate adjustment)		0.000%	-0.003%	-0.001%	0.002%	0.000%	0.007%	0.000%	0.001%	0.001%
p-value(t-test)		0.97	0.98	0.87	0.78	0.86	0.89	0.74	0.93	0.91
p-value(Satterthwaite-Weich t-test*)		0.97	0.98	0.87	0.78	0.86	0.89	0.74	0.93	0.91
p-value(Wilcoxon/Mann-Whitney)		0.97	0.98	0.84	0.73	0.82	0.85	0.59	0.88	0.85
p-value(Wilcoxon/Mann-Whitney (tie-adj.))		0.97	0.98	0.84	0.73	0.82	0.85	0.59	0.88	0.85
Correlation		0.99999	0.99999	1.00000	0.99994	1.00000	1.00000	0.99998	1.00000	1.00000
GBP										
Mean (3M Tbill adjustment)	0.004%		-0.001%	-0.002%	0.000%	0.000%	0.007%	0.001%	0.002%	0.002%
Mean (O/N Rate adjustment)	0.004%		-0.001%	0.000%	0.002%	0.001%	0.008%	0.003%	0.003%	0.003%
p-value(t-test)	0.97		0.99	0.86	0.75	0.86	0.91	0.86	0.96	0.94
p-value(Satterthwaite-Weich t-test*)	0.97		0.99	0.86	0.75	0.86	0.91	0.86	0.96	0.94
p-value(Wilcoxon/Mann-Whitney)	0.97		1.00	0.81	0.23	0.81	0.86	0.82	0.94	0.89
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.97		1.00	0.81	0.23	0.81	0.86	0.82	0.94	0.89
Correlation	0.99999		0.99999	0.99999	0.99990	0.99999	0.99999	0.99999	0.99999	1.00000
JPY										
Mean (3M Tbill adjustment)	0.008%	0.006%		0.003%	0.004%	0.005%	0.012%	0.006%	0.007%	0.007%
Mean (O/N Rate adjustment)	0.008%	0.006%		0.004%	0.007%	0.006%	0.013%	0.007%	0.008%	0.008%
p-value(t-test)	0.98	0.99		0.89	0.78	0.89	0.92	0.88	0.96	0.94
p-value(Satterthwaite-Weich t-test*)	0.98	0.99		0.89	0.78	0.89	0.92	0.88	0.96	0.94
p-value(Wilcoxon/Mann-Whitney)	0.98	1.00		0.86	0.71	0.84	0.87	0.83	0.95	0.92
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.98	1.00		0.86	0.71	0.84	0.87	0.83	0.95	0.92
Correlation	0.99999	0.99999		1.00000	0.99993	1.00000	1.00000	0.99999	1.00000	1.00000
EUR										
Mean (3M Tbill adjustment)	0.007%	0.004%	0.001%		0.001%	0.002%	0.009%	0.004%	0.005%	0.005%
Mean (O/N Rate adjustment)	0.005%	0.003%	0.000%		0.003%	0.002%	0.009%	0.004%	0.005%	0.005%
p-value(t-test)	0.87	0.86	0.89		0.78	0.97	0.97	0.97	0.95	0.98
p-value(Satterthwaite-Weich t-test*)	0.87	0.86	0.89		0.78	0.97	0.97	0.97	0.95	0.98
p-value(Wilcoxon/Mann-Whitney)	0.84	0.81	0.86		0.70	0.96	0.97	0.95	0.93	0.97
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.84	0.81	0.86		0.70	0.96	0.97	0.95	0.93	0.97
Correlation	1.00000	0.99999	1.00000		0.99975	1.00000	1.00000	0.99999	1.00000	1.00000
CHF										
Mean (3M Tbill adjustment)	0.007%	0.004%	0.001%	0.000%		0.002%	0.009%	0.004%	0.006%	0.005%
Mean (O/N Rate adjustment)	0.004%	0.001%	-0.002%	-0.001%		0.001%	0.007%	0.003%	0.003%	0.003%
p-value(t-test)	0.78	0.75	0.78	0.78		0.87	0.85	0.92	0.87	0.90
p-value(Satterthwaite-Weich t-test*)	0.78	0.75	0.78	0.78		0.87	0.85	0.92	0.87	0.90
p-value(Wilcoxon/Mann-Whitney)	0.73	0.23	0.71	0.70		0.82	0.77	0.90	0.82	0.86
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.73	0.23	0.71	0.70		0.82	0.77	0.90	0.82	0.86
Correlation	0.99994	0.99990	0.99993	0.99975		0.99989	0.99991	0.99994	0.99996	0.99996
NOK										
Mean (3M Tbill adjustment)	0.006%	0.003%	0.001%	0.000%	0.001%		0.008%	0.003%	0.004%	0.004%
Mean (O/N Rate adjustment)	0.004%	0.002%	-0.001%	0.000%	0.002%		0.008%	0.003%	0.003%	0.003%
p-value(t-test)	0.86	0.86	0.89	0.97	0.87		0.95	0.99	0.94	0.97
p-value(Satterthwaite-Weich t-test*)	0.86	0.86	0.89	0.97	0.87		0.95	0.99	0.94	0.97
p-value(Wilcoxon/Mann-Whitney)	0.82	0.81	0.84	0.96	0.82		0.94	0.97	0.92	0.96
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.82	0.81	0.84	0.96	0.82		0.94	0.97	0.92	0.96
Correlation	1.00000	0.99999	1.00000	1.00000	0.99989		1.00000	0.99999	1.00000	1.00000
SEK										
Mean (3M Tbill adjustment)	-0.001%	-0.003%	-0.006%	-0.007%	-0.006%	-0.006%		-0.004%	-0.003%	-0.003%
Mean (O/N Rate adjustment)	-0.002%	-0.004%	-0.007%	-0.007%	-0.004%	-0.006%		-0.003%	-0.003%	-0.003%
p-value(t-test)	0.89	0.91	0.92	0.97	0.84	0.95		0.95	0.97	1.00
p-value(Satterthwaite-Weich t-test*)	0.89	0.91	0.92	0.97	0.84	0.95		0.95	0.97	1.00
p-value(Wilcoxon/Mann-Whitney)	0.85	0.86	0.87	0.97	0.77	0.94		0.93	0.95	0.99
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.85	0.86	0.87	0.97	0.77	0.94		0.93	0.95	0.99
Correlation	1.00000	0.99999	1.00000	1.00000	0.99991	1.00000		1.00000	1.00000	1.00000
CAD										
Mean (3M Tbill adjustment)	0.004%	0.003%	0.001%	0.001%	0.002%	0.002%	0.009%		0.003%	0.003%
Mean (O/N Rate adjustment)	0.002%	0.002%	-0.001%	0.000%	0.003%	0.002%	0.008%		0.002%	0.002%
p-value(t-test)	0.74	0.86	0.88	0.97	0.92	0.99	0.96		0.92	0.96
p-value(Satterthwaite-Weich t-test*)	0.74	0.86	0.88	0.97	0.92	0.99	0.96		0.92	0.96
p-value(Wilcoxon/Mann-Whitney)	0.59	0.82	0.83	0.95	0.90	0.97	0.93		0.87	0.92
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.59	0.82	0.83	0.95	0.90	0.97	0.93		0.87	0.92
Correlation	0.99998	0.99999	0.99999	0.99999	0.99994	0.99999	1.00000		1.00000	1.00000
AUD										
Mean (3M Tbill adjustment)	0.006%	0.004%	0.002%	0.001%	0.003%	0.002%	0.009%	0.002%		0.002%
Mean (O/N Rate adjustment)	0.005%	0.003%	0.001%	0.002%	0.005%	0.003%	0.010%	0.003%		0.002%
p-value(t-test)	0.93	0.96	0.96	0.95	0.87	0.94	0.97	0.92		0.97
p-value(Satterthwaite-Weich t-test*)	0.93	0.96	0.96	0.95	0.87	0.94	0.97	0.92		0.97
p-value(Wilcoxon/Mann-Whitney)	0.88	0.94	0.95	0.93	0.82	0.92	0.95	0.87		0.94
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.88	0.94	0.95	0.93	0.82	0.92	0.95	0.87		0.94
Correlation	1.00000	0.99999	1.00000	1.00000	0.99996	1.00000	1.00000	0.99999		1.00000
NZD										
Mean (3M Tbill adjustment)	0.007%	0.005%	0.003%	0.002%	0.004%	0.003%	0.011%	0.004%	0.003%	
Mean (O/N Rate adjustment)	0.006%	0.004%	0.002%	0.003%	0.006%	0.004%	0.011%	0.004%	0.003%	
p-value(t-test)	0.91	0.94	0.94	0.98	0.90	0.97	1.00	0.95	0.97	
p-value(Satterthwaite-Weich t-test*)	0.91	0.94	0.94	0.98	0.90	0.97	1.00	0.95	0.97	
p-value(Wilcoxon/Mann-Whitney)	0.85	0.89	0.92	0.97	0.86	0.96	0.99	0.92	0.94	
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.85	0.89	0.92	0.97	0.86	0.96	0.99	0.92	0.94	
Correlation	1.00000	1.00000	1.00000	1.00000	0.99996	1.00000	1.00000	1.00000	1.00000	

The column labels denote base currency and row labels denote foreign currencies For each currency pair, mean returns for the different interest rate adjustments, and a series of equality tests as well as a correlation coefficient between time series are shown. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

c) *Moving Average Signal Calculations*

Having defined currency returns that mimic the returns obtainable from a futures-based trading strategy, the next step is to define momentum observations. In order to incorporate the interest rate adjustment in the survivorship analysis, the historic price time series is recalculated on the basis of interest rate adjusted returns, as given in Equation 2. Each of the historic currency price time series is rebased to 100 as of the 4th of January 1974.

$$(5) \quad \text{Positive Momentum} = \frac{1}{S} \sum_{i=0}^{S-1} S_{t-i} \geq \frac{1}{L} \sum_{i=0}^{L-1} S_{t-i}$$

$$(6) \quad \text{Negative Momentum} = \frac{1}{S} \sum_{i=0}^{S-1} S_{t-i} < \frac{1}{L} \sum_{i=0}^{L-1} S_{t-i}$$

If the short-term moving average is equal to or above the long-term moving average, then a positive momentum signal is observed. If the short-term moving average is below the long-term moving average, then a negative momentum signal is observed. There is no unified rule as to which moving average combination should be used.

FIGURE 1-2: MOVING AVERAGE COMBINATIONS

	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	1/5	1/10	1/15	1/20	1/25	1/30
SR 2	2/5	2/10	2/15	2/20	2/25	2/30
SR 3	3/5	3/10	3/15	3/20	3/25	3/30
SR 4	4/5	4/10	4/15	4/20	4/25	4/30
SR 5		5/10	5/15	5/20	5/25	5/30
SR 10			10/15	10/20	10/25	10/30
SR 15				15/20	15/25	15/30
SR 20					20/25	20/30
SR 25						25/30

The column labels denote long-term moving averages and row labels denote short-term moving averages. All short-term moving averages have to be shorter than any long-term moving average.

While Levich and Thomas (1991) apply rather short-term focused trading signals, practitioners such as Alexander Elder (2002) suggest moving average ranges starting from around 10 to 20 up to 50 days. The rationale behind the choice of this range is the


fact that 20 trading days effectively represent the time horizon of a month. The present chapter defines the range of short-term moving averages as 1 to 5 days as well as 10, 15, 20 and 25 days. Long-term moving averages are defined as 5, 10, 15, 20, 25 and 30 days. Any short-term moving average has to be shorter than any long-term moving average. This gives 39 different moving average combinations upon which the survivorship analysis is based. Figure 1-2 summarises all moving average combinations applied in this study. The column labels denote long-term moving averages and row labels denote short-term moving averages. All short-term moving averages have to be shorter than any long-term moving average.

2. Construction of Survivorship Curves

The basic idea behind creating survivorship curves is to model the probability of the persistence of some pre-specified signal within a given data sample. To illustrate the concept, Figure 1-3 shows hypothetical trading signals that have been created from a dual crossover moving average trading rule. The trading rule generates signals if the short-term moving average is above or equal to the long-term moving average. Whenever the short-term moving average is below the long-term moving average the previous trading signal ceases to exist and the trading rule gives an output of zero. This gives a series of trading signals of different lengths scattered along the empirical time series. Figure 1-3 gives a graphical description of the concept of duration data. The figure shows that on the positive side one momentum signal survives five days, three momentum signals survive three days, one signal survives two days, and two signals live for one day. On the negative side one signal survives for six days, one for five days, two signals survive for three days and two survive for two days and one signal survives for one day. The survivorship analysis aims to analyse the survival characteristics of the trading signals that have been created by the moving average crossover rules. This cannot be estimated at a single point in time because such observations do, as pointed out earlier, occur randomly within the sample. Survivorship and hazard curves, as laid out by Kaplan and Meier (1958), overcome the problem of analysing uncensored datasets. By constructing the Product Limit Estimator (PLE), Kaplan and Meier (1958) find a way of ordering data such that probabilities of survival can be calculated and inferences can be made. Originally, this methodology has been used in biomedical research to investigate the effectiveness of medical treatment on patient groups. However, over time, the methodology has found its use in analysing

economic problems, such as the analysis of unemployment rates or the estimation of credit default rates as suggested by Kiefer (1988).


FIGURE 1-3: GRAPHICAL DESCRIPTION OF DURATION DATA



On the positive side one momentum signal survives five days, three momentum signals survive three days, one signal survives two days, and two signals live for one day. On the negative side one signal survives for six days, one for five days, two signals survive for three days and two survive for two days and one signal survives for one day.

Despite the fact that this chapter uses only a small fraction of the actual breadth of possibilities offered by this methodology, it is beneficial to outline the statistical principals of survival time analysis, which is done by following the works of Kaplan and Meier (1958), Kiefer (1988), Lawless (2003) and Kleinbaum and Klein (2012) closely.

FIGURE 1-4: GRAPHICAL DESCRIPTION OF LIFETIME STATISTICS



A and B are the starting and ending dates of the survey, t_1 and t_2 are start-up observations, where birth and death of the start-up occurs within the survey time; t_3 is a right censored observation, and t_4 is a left censored observation.

For illustration purposes one imagines a survey that analyses the lifetime of start-up companies within an industry. The time span of the survey is assumed to be 120 months. Some of the companies within the survey will have started their business before

the start date of the survey and might have filed for bankruptcy within the timeframe of the survey. Other firms might start business operations sometimes within the survey period and they continue to exist even after the survey has been finished. Furthermore, the timing and the length of survival of various start-ups within the sample is not known. Figure 1-4 shows the described problem. Points A and B are starting and ending dates of the survey, t_1 and t_2 are start-up observations where birth and death of the start-up has occurred within the survey time t_3 is a right censored observation, and t_4 is a left censored observation. The general problem of this kind of test setup is the fact that right and left censored data cannot be estimated with perfect certainty. Moreover, uncensored observations within the dataset cannot easily be analysed with traditional statistical methods, hence making statistical inferences about the duration of such datasets difficult with traditional econometric tools. Survival time analysis utilises the concept of conditional probabilities. It periodically analyses the probability of whether an observation continues or ceases to exist post a specific date. From these results a string of conditional probabilities is created. In that sense it is not looking for the probability of a start-up filing for bankruptcy after an exact number of months, but it focuses on estimating the probability of a start-up filing for bankruptcy in each of the months in the survey period, given the start-up has survived the previous months beforehand. The advantage of this approach is the fact that the probability of a bankruptcy in month 10 for instance can be considered as the outcome of a sequence of simple conditional probabilities. To formalise the earlier descriptions, it can be assumed that within the timeframe of the survey, a start-up company can take two states, either it is operating or it goes bankrupt. Both states are mutually exclusive. Hence, the probability of failure or survival for a pre-specified time horizon can be written as follows:

$$(7) \quad F(t) = \Pr(T < t)$$

$$(8) \quad S(t) = \Pr(T \geq t)$$

$$(9) \quad S(t) = 1 - F(t)$$

Function $F(t)$ in Equation 7 is defined as the probability of T being smaller than a time t , whereby T is a random variable denoting the time of bankruptcy, as per the example. Hence, this is the probability for the time of bankruptcy to be before some pre-specified

time. Function $S(t)$ in Equation 8 on the other hand denotes the probability of T (time of bankruptcy) to be bigger than t . Hence, it is the probability of the bankruptcy occurring after some pre-specified date. Given the fact that $F(t)$ and $S(t)$ are mutually exclusive events, the link between both can be summarized in Equation 9. Taking the derivative of $F(t)$ and $S(t)$ produces the corresponding density functions for the two probabilities. Both are given in Equations 10 and 11. They can be seen as the rate of either bankruptcy, or survival per unit of time. Equation 12 shows the link between both density functions. This basic set of equations lays the foundation for any further analysis.

$$(10) \quad f(t) = \frac{dF(t)}{dt}$$

$$(11) \quad s(t) = \frac{dS(t)}{dt}$$

$$(12) \quad s(t) = \frac{d[1-F(t)]}{dt} = -f(t)$$

There are two other concepts in the subject of lifetime statistic that will help understanding the construction and the interpretation of the Kaplan Meier Product Limit Estimator. The first concept is the hazard function, shown in Equations 11 and 12. Given the linkage between failure and survival probabilities, there are obviously many ways to express the hazard function. For reasons of simplicity, this chapter will focus on the standard definition of the concept. For a more in-depth treatment please refer to Kiefer (1988), Lawless (2003) and Kleinbaum and Klein (2012)

$$(13) \quad \lambda(t) = \frac{f(t)}{s(t)}$$

$$(14) \quad \lambda(t)dt = Pr(t \leq T < t + dt | T \geq t)$$

Equation 13 shows the general definition of the hazard curve. Equation 14 gives the precise definition in terms of probabilities. A hazard curve denotes the conditional probability of an observation ceasing to exist within a pre-defined time horizon of t to $t+dt$, given that it has survived until t . Interpreting this measure in terms of the start-up example, it would for instance allow to calculate the conditional probability of a start-up

filing for bankruptcy in the time period between month 10 and month 11, given it has survived 10 months.

The next paragraphs will focus on the construction of the Product Limit Estimator (PLE) as defined by Kaplan and Meier in 1958. Furthermore, the calculation of the variance of the estimator is presented, as well as the Log Rank test is introduced. Both of these concepts represent the core methodology that is applied in this chapter. Survivorship analysis traditionally distinguishes two types of survival models, parametric and non-parametric models. The PLE is a non-parametric measure. Hence it does not rely on any assumption of distributional characteristics of the underlying data. This is particularly useful for the analysis of financial time series. To illustrate the dynamics of the PLE, the example of the start-up survey is used once again. For illustration, we assume that 80 start-ups have entered the survey of 120 months and bankruptcy filings for the companies in the sample happen in months 8, 31, 54 and 92.

FIGURE 1-5: RESULTS FROM THE HYPOTHETICAL START-UP SURVEY

Interval	Factor					
$t_{(i-1)}, t_i$	n_i	d_i	n_i'	l_i	s_i	$S(t)$
0-8	80	10	70	20	0.875	0.875
8-31	50	10	40	0	0.800	0.700
31-54	40	10	30	10	0.750	0.525
54-92	20	10	10	0	0.500	0.263
92-120	10	0	10	10	1.000	0.263

Column 1 shows the time intervals between bankruptcy filing. Columns 2 to 7 show various factors of the analysis; n_i represents the number of start-up firms in the sample immediately before bankruptcies occur in this time interval. n_i' represents the number of start-up firms in the sample immediately after bankruptcies in this time interval occur. d_i is the number of bankruptcies that occur during the interval. The column l_i gives the periodical number of firms that drop out of the survey for other reasons than bankruptcy. s_i is the periodic survival probability, and $s(t)$ is the cumulative survival probability.

Figure 1-5 summarises the hypothetical results for the survey of start-up companies. Column 1 shows the time intervals between bankruptcy filings. Columns 2 to 7 show various factors of the analysis, whereby n_i represents the number of start-up firms in the sample immediately before bankruptcies occur in this time interval. n_i' represents the number of start-up firms in the sample immediately after bankruptcies in this time interval occur. Hence, it is the number of observations that survive during the time interval. d_i is the number of bankruptcies that occur during the interval. The column l_i gives the periodical number of firms that drop out of the survey for other reasons than

bankruptcy. Given the fact that the present chapter is based on financial time series data, this element of the PLE is less important for the analysis in this chapter. The ratio between n_i' and n_i represents a conditional probability of survival for that time interval. It divides the number of observations that have survived over the interval by the number of observations that were at risk at the beginning of the interval.

Comparing this measure with the hazard rate in Equations 13 and 14 it is easy to understand the link between both measures. The conditional probability of survival is shown in Equation 15 and the link to the hazard curve is shown in Equation 16. Taking the product of the periodical survival probabilities, one can obtain the cumulative survival probability, which is in effect the PLE as given in Equation 17.

$$(15) \quad \hat{S}_i = \left(\frac{n_i'}{n_i} \right) = \frac{(n_i - d_i)}{n_i}$$

$$(16) \quad \hat{S}_i = (1 - \lambda_i)$$


$$(17) \quad \hat{S}(t) = \prod_{i=1}^k \left(\frac{n_i'}{n_i} \right) \text{ with } u_k = 1 \text{ and } n_i' = n_i - d_i$$

The PLE is obtained by setting the conditional probability of survival equal to the observed relative frequency of completion for any given interval. The PLE estimator will approach the true survival function, when a large enough sample is taken. In order to make inferences of the validity of the estimator, the variance of the estimator has to be calculated. The definition is given in Equation 18.

$$(18) \quad VAR[\hat{S}(t)] = \hat{S}^2(t) \sum_1^k \left[\frac{d_i}{n_i(n_i - d_i)} \right]$$

The PLE forms a step function, with steps in every time interval where a loss occurs. It is assumed that in between the time steps, the survival rate remains constant. Figure 1-6 shows the graphical description of the PLE for the example of the startup survey that has been introduced earlier. Figure 1-6 indicates that 87.5% survive up until 8 month, 70% survive up until 31 month, 52.5% up until 54 month and 26.3% survive 120 month.

FIGURE 1-6: GRAPHICAL REPRESENTATION OF THE PLE FOR THE START-UP SURVEY DATA EXAMPLE



The figure shows a graphical description of the PLE for the example of the start-up survey. The figure indicates that 87.5% survive up until 8 month, 70% survive up until 31 month, 52.5% up until 54 month and 26.3% survive 120 month.

Besides the calculation of the estimator itself, this chapter relies on the comparison of any empirical survival curve with a theoretical benchmark survival curve to assess the presence of momentum effects. In order to facilitate such comparison the chapter uses the log-rank test, to verify potential differences between the empirical model and benchmark processes. Lawless (2003) indicates that the log-rank test, developed by Mantel and Cox, is based on the premise that every observation point on the survivorship curve can be seen as a contest between the two survival samples. Equation 19 shows the basic methodology behind the test. It shows the test setup for two distinct samples ($z=1,2$).

$$(19) \quad \frac{\begin{bmatrix} d_{i1} & N_{i1}-d_{i1} \end{bmatrix} N_{i1}}{\begin{bmatrix} d_{i2} & N_{i2}-d_{i2} \end{bmatrix} N_{i2}} \frac{1}{d_i \quad N_i-d_i \quad N_i}$$

The first column contains the numbers of observations in each sample that were observed to fail at time i . The third column shows the number of observations that were at risk at time i . The column in the middle shows the number of observations that survive at time i . $0 = t_0 \leq t_1 < t_i < \dots < t_j$ are distinct times at which failure occurs in each of both samples.

$$(20) \quad V_1 = \sum_{i=1}^j w_i \left(d_{i1} - \frac{d_i N_{i1}}{N_i} \right)$$

The first step is given in Equation 20, which sums the periodical difference between the actual number of failures within sample one and the number of failures suggested by combining both samples. While d_i and N_i have been described earlier w_i has not been explained so far. It is a weighting parameter, which allows putting more or less weight on various survival observations.

This chapter will be using the Wilcoxon specification of the Cox Mantel test, where the weighting is based on the proportion of observations that are at risk at time i (for sample 1) relative to the total number of observations for the test sample. The reason for this specification is the fact that it will put equal emphasis on each of the lifetime observations. Under an equally weighted specification, for instance, longer survival observations will receive a weighting, which is disproportionate to the probability of occurrence. Hence, the test results of an equally weighted log-rank test would be heavily influenced by the results of the longer observations. These are by definition less reliable, due to their low frequency of occurrence. The weights of each time period are assigned as given by Equation 21.

$$(21) \quad w_i = \frac{N_{i1}}{N}$$

Equation 22 shows the variance of the sum of weighted differences between actual and expected failure rate.

$$(22) \quad Var(V_1) = \sum_{i=1}^j w_i^2 \frac{d_i(N_i - d_i)N_{i1}N_{i2}}{N_i^2(N_i - 1)}$$

Kleinbaum and Klein (2012) indicate that under the null hypothesis of no difference between survival curves, the tests statistic can be expressed as shown in equation 23,

this test statistic follows a χ^2 distribution with one degree of freedom. Given that the test setup of the log rank test is based on the χ^2 distribution, a one-sided hypothesis testing has to be applied.

$$(23) \quad \chi^2 \sim \frac{\left(\sum_{i=1}^j w_i \left(d_{i1} - \frac{d_i N_{i1}}{N_i} \right) \right)^2}{\sum_{i=1}^j w_i^2 \frac{d_i (N_i - d_i) N_{i1} N_{i2}}{N_i^2 (N_i - 1)}}$$

To assess the directionality of the test statistic, this chapter translates the χ^2 test statistic into a test statistic that follows a standard normal distribution for large samples. This transformation is shown in equation 24

$$(24) \quad Z \sim \frac{\sum_{i=1}^j w_i \left(d_{i1} - \frac{d_i N_{i1}}{N_i} \right)}{\sqrt{\sum_{i=1}^j w_i^2 \frac{d_i (N_i - d_i) N_{i1} N_{i2}}{N_i^2 (N_i - 1)}}}$$

The intuition behind this transformation can be described as follows. If Z follows a standard normal distribution then Z^2 follows a χ^2 distribution with 1 degree of freedom. For a further treatment of the subject see Lawless (2003) or Kleinbaum and Klein (2012).

C. Survivorship Analysis versus Runs Test

In the introduction it was pointed out that the presented survivorship analysis approach represents an extension of the concept of runs test, which was introduced by Fama in 1965. Given the fact that this is the main academic innovation of this chapter, it is beneficial not only to analyse the similarities between both approaches, but also to outline the main innovations that come with the survivorship analysis approach. The next section makes a short introduction to the concept of runs tests and it presents its main methodological aspects. Furthermore the section compares and contrasts both methodologies and it summarises the key differences between them.

Fama (1965) introduces the idea of runs tests as a novelty in finance, which was the case at that time. However, the concept itself was not new. It had previously been used in various other scientific fields, such as meteorology. Early work with regards to the subject can be found in Barton and David (1958, 1962). The general intuition behind the concept is to analyse the persistence of pre-defined mutually exclusive events. In Barton and David's case this analysis is mostly based on the analysis of rainfall or wind speed. Fama (1965) used this methodology to analyse persistence in equity returns. Dooley, Shafer (1976) also devote a chapter of their paper to the runs test for various exchange rates. The main reason for Fama's (1965) choice of the runs test is the fact that up until then the traditional way of analysing time series patterns of stock returns was based on autocorrelation tests. An autocorrelation based test setup is very sensitive to outliers in the underlying data. This is not the case when looking at returns as binary outcomes, which is what a runs test does. By defining a positive return as + and a negative return as -, one can count sequences of the same signs (which are called runs) and assess whether they are in line with what is expected. For the purpose of illustration, one could imagine a sequence of positive and negative returns as given in Equation 25. Such a sequence is then split into four runs, whereby two runs are positive and two runs are negative. Out of the positive return runs, there is a run of three and a run of two, on the side of negative returns there is a run of one and a run of two.

$$(25) \quad + + + - + + - -$$

Fama (1965) carries out three types of runs tests. The first test looks at all runs irrespective of sign. The second test looks at positive and negative runs separately. The third test calculates the length of positive and negative runs. For the purpose of comparing the runs tests to the survivorship analysis this chapter conducts the first and the third tests proposed by Fama (1965). The reason why the second test is not conducted is the fact that the results of the second test are embedded in the results of the third test. The first and the third tests will be carried out following the exact specifications proposed by Fama. The analysis is carried out on one day return observations. All tests in the following section are conducted on the basis of the full

data sample⁹, for which interest rate adjusted returns are calculated; it is carried out on all currency pairs.

All of the test specifications proposed by Fama (1965) are based on the assumption that stock returns follow a Markov type process. That means that the return realisations are independent from each other. It also means that transition probabilities between positive, negative and zero returns have to be assigned. These transition probabilities aim to reflect a split between positive negative and zero returns that reflect the true pattern of the time series. In order to facilitate this, Fama (1965) proposes to use the actual split of the historic time series as a representative estimate for these transition probabilities. Given the fact that the time series used in the present study is very long, Fama's assumption can be adopted. Therefore, considering the actual split of returns and assuming independence between observations, the number of total runs (for all signs) can be calculated as shown in Equation 26, whereby N represents the number of total observations and n_i represents the number of positive, negative and zero returns respectively.

$$(26) \quad m = \frac{[N(N+1) - \sum_{i=1}^3 n_i^2]}{N}$$

Applying survivorship curves to the first type of runs test suggested by Fama is not possible because the test looks at the absolute number of runs, without any element of direction or time. To establish a link between the methodology proposed in this chapter and Fama's concept, this chapter generates the expected number of runs as given in Equation 26 by carrying out a resampling simulation with 500 iterations. From the distribution of the empirically observed number of runs the mean is used as a fair estimate of an expected number of runs. This is then compared to the theoretical number derived by Fama (1965) and the empirically observed number of runs.

The calculation of the statistical significance of Fama's results follows closely the methodology proposed by Fama (1965)¹⁰, while the results of the simulation methodology are based on the Welch F-test, for differences in means. Figure 1-7 shows the percentage difference between results. Figure 1-7 is constructed in such way that it

⁹ 4th of January 1974 to the 31st of December 2009

¹⁰ Fama estimates the standard deviation of m as follows:

$$\sigma_m = \left(\frac{\sum_{i=1}^3 n_i^2 [\sum_{i=1}^3 n_i^2 + N(N+1)] - 2N \sum_{i=1}^3 n_i^3 - N^3}{N^2(n-1)} \right)^{1/2}.$$

He then calculates a standardised difference between the number of empirical runs R and the number of expected runs m as follows $K = \frac{(R+0.5)-m}{\sigma_m}$. From this formula the statistical significance is calculated.

displays the base currencies across the columns and the foreign currencies across the rows. Each of the boxes represents one currency pair and consists of three numbers. The first number on the upper left is the number of total runs observed in the empirical time series. The number on the upper right is the percentage difference between the observed number of runs and the number of runs calculated by Equation 26. The lower right number shows the percentage difference between the observed number of runs and the expected number of runs that comes out of the resampling simulation. The stars next to the upper and lower right numbers indicate statistical significance levels.

FIGURE 1-7: BASIC RUNS TEST ANALYSIS; NO ASSUMPTION ABOUT DIRECTION

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD		4383 -2.57% -2.53% **	4668 3.50% * 3.59% ***	4518 0.11% 0.18%	4605 2.16% 2.22% **	4547 0.75% 0.80%	4613 2.22% 2.24% **	4550 0.91% 0.90%	4584 1.86% 1.88% *	4615 2.55% 2.57%
GBP	4383 -2.57% -2.45% **		4321 -3.71% ** -3.62% ***	4435 -1.27% -1.22%	4192 -4.15% ** -4.07% ***	4601 1.97% 1.96% *	4613 2.31% 2.40% **	4411 -2.25% -2.20% **	4479 -0.74% -0.68%	4543 0.66% 0.80%
JPY	4668 3.50% * 3.57% ***	4321 -3.71% ** -3.68% ***		4323 -4.18% ** -3.96% ***	4431 -1.88% -1.75% *	4451 -1.29% -0.92%	4381 -2.83% -2.78% **	4646 3.33% * 3.37% ***	4566 1.74% 1.87% *	4608 2.49% 2.51%
EUR	4518 0.11% 0.18%	4435 -1.27% -1.30%	4323 -4.18% ** -3.90% ***		4647 3.13% * 3.24% ***	4834 7.38% *** 7.54% ***	4673 3.57% * 3.71% ***	4515 0.07% 0.15%	4601 2.19% 2.27% **	4652 3.35% 3.39% *
CHF	4605 2.16% 2.27% **	4192 -4.15% ** -4.15% ***	4431 -1.88% -1.78% *	4647 3.13% * 3.22% ***		4678 3.85% 3.92%	4732 4.98% *** 5.06% ***	4487 -0.46% -0.41%	4561 1.26% 1.31%	4628 2.78% *** 2.83%
NOK	4547 0.75% 0.72%	4601 1.97% 2.05% *	4451 -1.29% -0.89%	4834 7.38% *** 7.58% ***	4678 3.85% 3.93% ***		4668 3.09% * 3.84% ***	4575 1.34% 1.37%	4719 4.68% ** 4.73% ***	4730 4.95% 5.07% ***
SEK	4613 2.22% 2.26% **	4613 2.31% 2.34% **	4381 -2.83% -2.73% **	4673 3.57% * 3.69% ***	4732 4.98% *** 5.11% ***	4668 3.09% * 3.88% ***		4576 1.35% 1.49%	4711 4.67% ** 4.79% ***	4744 5.23% 5.52% ***
CAD	4550 0.91% 1.03%	4411 -2.25% -2.16% **	4646 3.33% * 3.32% ***	4515 0.07% 0.07%	4487 -0.46% -0.42%	4575 1.34% 1.40%	4576 1.35% 1.43%		4620 2.49% 2.56% **	4650 3.13% 3.21% *
AUD	4584 1.86% 1.92% *	4479 -0.74% -0.68%	4566 1.74% 1.83% *	4601 2.19% 2.19% **	4561 1.26% 1.29%	4719 4.68% ** 4.77% ***	4711 4.67% ** 4.66% ***	4620 2.49% 2.62% **		4860 7.47% 7.90% ***
NZD	4615 2.55% 2.59% **	4543 0.66% 0.73%	4608 2.49% 2.50% **	4652 3.35% * 3.35% ***	4628 2.78% 2.94% ***	4730 4.95% *** 5.08% ***	4744 5.23% 5.61% ***	4650 3.13% 3.25% ***	4860 7.47% 7.89% ***	

Base currencies are given across the columns and the foreign currencies across the rows. Each of the boxes represents one currency pair and consists of three numbers. The first number on the upper left is the number of runs observed in the empirical time series. The number on the upper right is the percentage difference between the actual number of runs and the number of runs calculated by Equation 26. The lower right number shows the percentage difference between the actual number of runs and the expected number of runs that comes out of the simulation analysis. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

The analysis shows that the currency matrices are symmetric. The empirical number of runs is the same for the GBPUSD and the USDGBP cross. The calculations suggest 4830 runs for the currency cross, irrespective of base currency. This observation can also be made for the differences between the actual and the theoretical number of runs. The GBPUSD as well as the USDGBP cross suggest that the theoretical number of runs is 2.57% less than the actual number of runs exhibited by the time series. Base currency effects can be found in the differences between the empirical and simulated curves. Here the GBPUSD cross suggests that the theoretical number of runs is 2.53% less than the empirical number of runs, while the USDGBP cross suggests that the difference is 2.45%. The differences between base currencies are very small. They can therefore be considered as immaterial.

The second observation that can be made is the fact that the differences between realised and expected number of runs is quite significant. For the USDGBP cross the difference is more than 2.5%. For other currency pairs such as the NZDAUD the difference is as high as 7%. This indicates that, depending on the currency pair, the empirical time series either outlives the theoretical benchmark, or it exhibits a much shorter survival pattern than theory would suggest. The most important observation, however, is the fact that the percentage differences between actual and theoretical as well as actual and simulated results are very similar. This indicates that both methodologies show comparable performance under the given test specifications. However, where both methodologies differ, is on the levels of statistical significance. Fama's runs test calculation seems to have wider confidence intervals than the simulation analysis. Hence, fewer of the currency pairs indicate a statistical significant difference between actual and hypothetical number of runs. Fama (1965) indicates that the proposed hypothesis test does suffer from sample size problems. He suggests that the best way to analyse the data is to look at the absolute differences between the empirical and theoretical number of runs. In both cases, for either Fama's proposed model, or the simulation methodology applied, these differences are very similar.

The other test proposed by Fama goes further than merely looking at the absolute number of runs. It allows calculating a theoretically derived, expected number of positive, negative and zero runs. Furthermore, it allows the calculation of the average life of positive and negative runs. The starting point of Fama's analysis is again the assumption of a Markov process with $P(+)$, $P(-)$ and $P(0)$, as probabilities assigned to positive, negative and zero price changes¹¹. From these probabilities the expected number of positive runs can be calculated using Equation 27.

$$(27) \quad \sum_{i=1}^{\infty} NP(+)^i [1 - P(+)]^2 = NP(+)[1 - P(+)]$$

To obtain the theoretically expected average life of runs, the total theoretical number of runs has to be disaggregated into time increments. That means one has to calculate an expected number of one day, two day, three day etc. runs. Equation 28 allows for the extraction of the expected proportion of positive runs that survive for i days, out of the total number of positive runs.

¹¹ As pointed out earlier probabilities are equivalent to the empirical split between positive, negative and zero returns, observed in the data sample.

$$(28) \quad \frac{NP(+)^i[1-P(+)]^2}{NP(+)[1-P(+)]} = P(+)^{i-1}[1 - P(+)]$$

Equation 28 gives the conditional probability of a run being of the length of i days, given the fact that it has been identified to be positive. The expected number of positive runs for the length of i can then be calculated as given in Equation 29. Whereby $R(+)$ is the total actual number of positive runs

$$(29) \quad \bar{R}_i(+) = R(+) P(+)^{i-1}[1-P(+)]$$

The same concept can be applied for negative and zero price changes, as given in Equations 30 (a, b, c) and 31 (a, b, c)

$$(30) \quad (a) \quad \sum_{i=1}^{\infty} NP(-)^i[1 - P(-)]^2 = NP(-)[1 - P(-)]$$

$$(b) \quad \frac{NP(-)^i[1-P(-)]^2}{NP(-)[1-P(-)]} = P(-)^{i-1}[1 - P(-)]$$

$$(c) \quad \bar{R}_i(-) = R(-) P(-)^{i-1}[1 - P(-)]$$

$$(31) \quad (a) \quad \sum_{i=1}^{\infty} NP(0)^i[1 - P(0)]^2 = NP(0)[1 - P(0)]$$

$$(b) \quad \frac{NP(0)^i[1-P(0)]^2}{NP(0)[1-P(0)]} = P(0)^{i-1}[1 - P(0)]$$

$$(c) \quad \bar{R}_i(0) = R(0) P(0)^{i-1}[1 - P(0)]$$

Figure 1-8 compares the average life of the empirical number of runs to the theoretical and simulated number of runs. The resampling simulation follows the same methodology as outlined earlier for Figure 1-7. The average life of a run is hereby a simple weighted average of the number of runs observed at every given day i . Figure 1-8 is split into two sub figures. The first figure shows the results of the positive runs and the second figure shows the results of the negative runs. The figure for zero runs is not displayed. The currency space is very liquid and continuously trading, hence there are no issues with respect to stale prices, therefore the number of zero runs in the sample is almost non-existent. Each of the figures is constructed in such a way that it displays the base currencies across the columns. Foreign currencies are shown across the rows. Each of the boxes represents one currency pair and consists of three numbers. The first number on the upper left is the average life of positive or negative runs of the empirical

time series. The number on the upper right is the percentage difference between the empirical average life of runs and the average life of runs calculated by Equations 29 and 30c. The lower right number shows the percentage difference between the empirical average life and the average life of runs that is derived from the resampling simulation. For this runs analysis Fama does not propose a test for statistical significance, his analysis merely looks at the differences between actual and theoretical value. This chapter follows Fama's approach.

FIGURE 1-8: AVERAGE LIFE; RUNS TEST FOR POSITIVE AND NEGATIVE RETURNS

Runs test for positive returns											
	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD	
USD		1.21 -8.4% -5.2%	0.93 -11.0% 1.3%	1.02 -10.7% -2.0%	0.92 -11.0% 2.1%	1.02 -9.7% 2.0%	1.00 -9.3% 4.8%	1.05 -9.4% 1.0%	1.49 -5.4% -33.2%	1.15 -7.7% -3.2%	
GBP	0.93 -13.2% -6.0%		0.97 -14.1% -11.6%	0.96 -12.9% -7.4%	0.81 -15.5% -10.0%	0.89 -10.9% 5.4%	0.89 -11.0% 4.1%	0.98 -11.9% -3.2%	1.03 -10.2% -0.8%	1.04 -9.7% -1.3%	
JPY	1.05 -9.0% 1.6%	1.22 -7.7% -5.2%		1.22 -9.7% -13.2%	1.06 -10.3% -3.7%	1.10 -9.0% -1.1%	1.07 -9.5% -0.8%	1.03 -7.8% 9.3%	1.12 -7.0% 5.2%	1.07 -7.4% 5.7%	
EUR	1.03 -10.8% -3.0%	1.21 -8.7% -8.8%	1.01 -13.0% -10.0%		0.93 -11.8% -2.1%	0.91 -7.8% 16.7%	0.96 -8.9% 8.0%	1.04 -10.2% -1.7%	1.03 -8.7% 5.0%	1.01 -8.3% 7.1%	
CHF	1.01 -9.1% 5.1%	1.56 -2.1% -16.3%	0.98 -11.2% -0.6%	1.06 -8.3% 3.1%		0.98 -8.2% 9.2%	0.96 -8.2% 10.6%	1.09 -9.5% -1.9%	1.07 -8.6% 2.0%	1.02 -8.2% 7.7%	
NOK	0.93 -10.7% 2.9%	1.02 -8.9% 4.0%	0.94 -11.9% -3.1%	0.81 -9.9% 11.5%	0.92 -10.8% 0.7%		0.91 -9.5% 8.2%	0.96 -9.5% 6.0%	0.99 -7.9% 10.0%	0.96 -7.6% 13.8%	
SEK	0.92 -10.7% 3.4%	1.03 -8.7% 4.5%	0.99 -12.2% -5.9%	0.88 -9.9% 8.4%	0.85 -10.4% 8.1%	0.89 -9.1% 11.1%		0.96 -9.8% 5.4%	1.06 -7.7% 4.9%	0.98 -7.2% 11.9%	
CAD	0.90 -10.9% 4.5%	1.12 -10.2% -6.7%	0.83 -11.1% 5.8%	0.99 -11.0% -0.8%	0.94 -11.6% -0.9%	0.93 -10.3% 5.6%	0.98 -10.5% 1.1%		0.99 -8.3% 9.1%	0.96 -8.6% 9.9%	
AUD	0.79 -11.5% 6.3%	0.98 -10.8% -0.2%	0.81 -11.7% 4.2%	0.88 -10.9% 4.6%	0.92 -11.6% -0.3%	0.81 -10.2% 10.7%	0.78 -10.3% 11.9%	0.85 -10.6% 7.6%		0.91 -8.0% 13.5%	
NZD	0.90 -11.9% -1.1%	0.93 -10.3% 3.7%	0.84 -11.4% 4.3%	0.84 -10.6% 8.2%	0.85 -10.9% 6.0%	0.79 -10.1% 12.1%	0.78 -10.1% 12.1%	0.88 -10.5% 6.5%	0.79 -9.2% 15.6%		

Runs test for negative returns											
	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD	
USD		0.93 -13.2% -6.5%	1.05 -9.0% 1.6%	1.03 -10.8% -3.1%	1.01 -9.1% 4.9%	0.93 -10.7% 2.8%	0.92 -10.7% 3.1%	0.90 -10.9% 4.7%	0.79 -11.5% 6.3%	0.90 -11.9% -1.2%	
GBP	1.21 -8.4% -6.0%		1.22 -7.7% -5.2%	1.21 -8.7% -8.9%	1.56 -2.1% -16.0%	1.02 -8.9% 4.2%	1.03 -8.7% 4.7%	1.12 -10.2% -6.7%	0.98 -10.8% -0.1%	0.93 -10.3% 3.2%	
JPY	0.93 -11.0% 1.1%	0.97 -14.1% -11.6%		1.01 -13.0% -10.0%	0.98 -11.2% -0.9%	0.94 -11.9% -2.9%	0.99 -12.2% -6.0%	0.83 -11.1% 6.1%	0.81 -11.1% 4.2%	0.84 -11.4% 4.4%	
EUR	1.02 -10.7% -2.3%	0.96 -12.9% -7.4%	1.22 -9.7% -13.2%		1.06 -8.3% 3.2%	0.81 -9.9% 11.7%	0.88 -9.9% 8.5%	0.99 -11.0% -1.1%	0.88 -10.9% 4.7%	0.84 -10.6% 8.3%	
CHF	0.92 -11.0% 2.3%	0.81 -15.5% -9.9%	1.06 -10.3% -3.8%	0.93 -11.8% -2.1%		0.92 -10.8% 0.6%	0.85 -10.4% 8.4%	0.94 -11.6% -0.8%	0.92 -11.6% -0.3%	0.85 -10.9% 6.2%	
NOK	1.02 -9.7% 2.1%	0.89 -10.9% 5.4%	1.10 -9.0% -1.0%	0.91 -7.8% 16.7%	0.98 -8.2% 9.0%		0.89 -9.1% 11.3%	0.93 -10.3% 5.5%	0.81 -10.2% 10.7%	0.79 -10.1% 12.2%	
SEK	1.00 -9.3% 4.7%	0.89 -11.0% 4.1%	1.07 -9.5% -0.6%	0.96 -8.9% 7.8%	0.96 -8.2% 10.8%	0.91 -9.5% 8.1%		0.98 -10.5% 1.1%	0.78 -10.3% 11.9%	0.78 -10.1% 12.3%	
CAD	1.05 -9.4% 0.9%	0.98 -11.9% -3.3%	1.03 -7.8% 9.2%	1.04 -10.2% -1.9%	1.09 -9.5% -2.2%	0.96 -9.5% 5.6%	0.96 -9.8% 5.2%		0.85 -10.6% 7.7%	0.88 -10.5% 6.6%	
AUD	1.49 -5.4% -32.8%	1.03 -10.2% -0.5%	1.12 -7.0% 4.9%	1.03 -8.7% 5.1%	1.07 -8.6% 1.6%	0.99 -7.9% 10.0%	1.06 -7.7% 5.4%	0.99 -8.3% 9.4%		0.79 -9.2% 15.5%	
NZD	1.15 -7.7% -3.0%	1.04 -9.7% -0.7%	1.07 -7.4% 5.6%	1.01 -8.3% 7.5%	1.02 -8.2% 7.5%	0.96 -7.6% 14.1%	0.98 -7.2% 11.9%	0.96 -8.6% 10.2%	0.91 -8.0% 13.4%		

Base currencies are given across the columns and the foreign currencies across the rows. The first figure shows the results of the positive runs and the second figure shows the results of the negative runs. Each of the boxes represents one currency pair and consists of three numbers. The first number on the upper left is the average life of the empirical time series. The number on the upper right is the percentage difference between the actual number of runs and the number of runs calculated by Equation 29 and 30c. The lower right number shows the percentage difference between the actual number of runs and the average life of runs that comes out of the simulation analysis. For this runs analysis Fama does not propose a test for statistical significance, his analysis looks at the differences between actual and theoretical value. This chapter follows Fama's approach.

The results in Figure 1-8 suggest that while the calculation of the total number of runs offers symmetrical results across various base currencies, the calculation for the average life of runs does have some base currency effects. For instance, in the case of positive runs, GBPUSD has an average life of 0.93 days versus USDGBP, which has an average life of 1.21 days. This is obviously the reverse when it comes to negative runs, as shown in the lower part of Figure 1-8. Furthermore, while Figure 1-7 suggested that there is very little difference between theoretical numbers of returns and expected number of returns, Figure 1-8 indicates that the theoretically calculated number average life of runs, as proposed by Fama (1965) systematically underestimates the average life of empirical runs. The average difference between the actual lifetime and the theoretical lifetime across all currencies is approximately 10%. The simulated average lifetime on the other hand does not show any systematic bias. For some currency pairs the actual average lifetime is longer than the simulated average lifetime. For other currency pairs it is shorter. The average difference between the simulated average lifetime and actual average lifetimes across all currencies is less than 3%.

The main conclusion from these results is that the theoretical construction of the Runs Test works appropriately for a total number of runs. However, when it comes to analysing the lifetime characteristics of runs, Fama's model fails to perform. This is different for the simulation approach; the simulated results for the total number of runs are equivalent to the theoretically obtained number. However, when it comes to the analysis of the average lifetime of runs the simulation approach gives numbers that are much closer to the empirical numbers than the numbers suggested by the runs test, without bias in either direction. In addition, there are various other aspects that make the proposed survival time methodology superior to the traditional runs test calculation. First, the original specification of the runs test bases its methodology on the stochastic characteristics of Markov type processes. It assumes independence between time increments (return observations) and it assigns probabilities of transition between states (between positive returns, negative returns and zero returns) for the respective time increments. This assumption limits the runs test specification to the analysis of single return observations only. Hence, the trading rules that are proposed in this chapter cannot be tested directly using Fama's framework. Any rule that gives signals ("runs") on the basis of moving average type filters does incorporate a degree of autocorrelation between time increments. Therefore it is not possible to obtain a theoretical number of "runs" for momentum signals based on Fama's methodology.

One might argue that while momentum signals cannot be tested directly, the actual return time series that has been generated by applying a trading signal could potentially be tested in the runs test setup. Figure 1-9 is based on the USDGBP currency cross. It shows the runs test results of the time series that have been created by applying positive momentum filters to the empirical time series.

FIGURE 1-9: USDGBP TRADING RULE; AVERAGE LIFE TEST FOR POSITIVE RETURNS

	LR 5		LR 10		LR 15		LR 20		LR 25		LR 30	
SR 1	1.1547	22.1% -1.8%	1.1583	19.1% -3.3%	1.1826	18.1% -6.5%	1.1851	17.4% -6.4%	1.1848	16.8% -6.4%	1.195	16.3% -7.8%
SR 2	1.1544	16.4% -2.0%	1.175	15.4% -3.5%	1.1807	15.2% -5.1%	1.1767	14.7% -4.3%	1.1974	14.7% -6.6%	1.2067	14.3% -7.5%
SR 3	1.0753	14.0% -0.3%	1.1683	14.0% -5.0%	1.1801	14.3% -7.1%	1.1772	13.9% -5.7%	1.1998	14.1% -7.7%	1.2063	13.8% -8.8%
SR 4	0.9564	12.3% 1.6%	1.1462	13.4% -5.1%	1.1704	13.9% -7.1%	1.1822	13.3% -7.0%	1.1968	13.6% -8.1%	1.2165	13.4% -10.6%
SR 5			1.1311	12.6% -4.7%	1.1865	13.7% -9.8%	1.1745	13.0% -7.0%	1.2147	13.5% -10.8%	1.2126	13.2% -11.1%
SR 10					1.1398	12.5% -9.7%	1.1513	12.3% -8.0%	1.2073	12.9% -10.9%	1.2063	12.7% -10.1%
SR 15							1.1257	12.0% -8.8%	1.1911	12.5% -12.0%	1.2119	12.3% -11.8%
SR 20									1.1498	12.0% -11.3%	1.1676	11.8% -9.0%
SR 25											1.139	11.7% -7.2%

The first number, on the upper left, is the average life of total runs observed in the empirical time series. The number on the upper right is the percentage difference between the actual average life of runs and the average life of runs calculated by Equation 29. The lower right number shows the percentage difference between the actual average life of runs and the expected average life of runs that comes out of the resampling simulation analysis. For this runs analysis Fama does not propose a test for statistical significance, his analysis looks at the differences between actual and theoretical value. This chapter follows Fama's approach.

Each of the boxes represents one moving average crossover combination and consists of three numbers. The first number, on the upper left, is the average life of total runs observed in the empirical time series. The number on the upper right is the percentage difference between the actual average life of runs and the average life of runs calculated by Equation 29. The lower right number shows the percentage difference between the actual average life of runs and the expected average life of runs that comes out of the resampling simulation analysis. These results suggest that the actual average life expectancy of the empirical time series is considerably longer than what is suggested by the theoretical runs test proposed by Fama (1965). In the case of the SR1/LR10 filter, the difference is almost 20%. Furthermore, these differences are similar across currency pairs and they are very strong. The results from the simulation analysis are much closer to the results that have been obtained from the empirical time series. In addition, they

also suggest that the empirical observation tends to live slightly shorter than suggested by theory when it comes to short-term moving averages. However longer-term moving averages such as the SR5/LR30 combination live considerably shorter than what is suggested by the resampling simulation. As discussed earlier, the runs test cannot be applied to momentum signals directly because of the correlation between signals. Furthermore, the results in Figure 1-9 suggest that even the application of the runs test to the return time series of trading rules is not informative. Hence, it is fair to conclude that a runs test specification as outlined by Fama (1965) is not suitable for more sophisticated trading signals. However, it is possible to obtain an expected number of “runs” by the means of a simulation. This is possible not only for time series that have been created by trading rules, but also for momentum signals directly. This is the first key difference between the traditional specification of the runs test and the proposed model. While Fama’s (1965) methodology is limited to the analysis of sequences of daily returns, the model presented in this study allows the analysis of sequences of signals for any kind of filter rule. Moreover, if an empirical return stream contains or requires a more sophisticated benchmark assumption, Fama’s (1965) runs test would not be able to capture this. The simulation approach proposed in this chapter, however, is very flexible and can easily be adapted to cater for more advanced benchmark assumptions. Finally, the specification of runs test analysis widely fails to allow for hypothesis testing. Here it has to be pointed out that Fama (1965) proposes a test design that allows for hypothesis testing of the total number of runs. However, he also highlights the statistical deficiencies of the proposed hypothesis testing methodology and suggests the analysis of percentage deviations (as given in Figure 1-7, Figure 1-8 and Figure 1-9) as the most appropriate way of analysing the validity of runs tests. Furthermore, Fama’s (1965) paper fails to make any suggestions about potential hypothesis testing of average survival times, which is another aspect where the proposed methodology offers a clear advantage. The log-rank test, introduced in an earlier section, is a highly accurate statistical tool that allows for hypothesis testing, applicable to either simple return streams, or more complicated signals.

Overall, it can be concluded that under the first test specification, which analyses the number of runs, both methodologies deliver similar returns. However, when looking at the estimates for the average life of runs, Fama’s calculations systematically overestimate the empirical results. The proposed methodology on the other hand gives results that do not show any systematic bias. In addition, the proposed methodology is

sufficiently flexible to test more complicated trading signals than Fama's (1965) runs test. It also allows for differing benchmark assumptions and it facilitates hypothesis testing.

D. Sensitivity of Survivorship Curves to External Factors.


Another key question that has to be discussed in the context of survivorship analysis is to which degree the survival time of moving average trading rules is impacted by factors such as underlying the interest rate differential of the underlying exchange rate pairs, the volatility skew and kurtosis of the underlying foreign exchange data. The next section discusses this aspect and conducts a time series as well as a cross sectional analysis of the set of moving averages introduced earlier.

Rationale for using factors such as the interest rate differential and the level of currency volatility can be found in studies of Poljarliev and Levich (2010) and Kho (1996). With regards to the interest rate differential, the so called "Carry" Strategy, which is in effect buying higher yielding currencies and selling lower yielding currencies at the same time, has become hugely popular amongst investors. This has resulted in the fact that strategies such as carry have become risk factors in their own right. Poljarliev and Levich (2010) indicate that high correlations between the GBP/CHF cross and the NZD/JPY are merely driven by investor preferences, as opposed to economic linkages. The only similarity that those two crosses share is the fact that GBP and NZD are traditionally high yielding currencies, while CHF and JPY are historically low yielding currencies. Hence, in time periods where carry becomes popular the correlation of currency crosses that combine high and low yielding currency pairs has gone up. While all return calculations of this study are adjusted for interest rate differentials, a sudden change in preferences away from carry might have a considerable impact on interest rate adjusted currency returns as well. Given the fact that carry comes with sharp reversals in times of market stress, one could expect a slightly negative relationship between survival time and trading rule returns. When it comes to currency volatility, Kho's (1996) analysis provides strong evidence of a systematic link between trading rule returns and market volatility. In his study he identifies the risk premium as the covariation of the return stream derived from the moving average trading strategy with a CAPM based benchmark, whereby the MSCI world equity index is used as a market proxy. In particular his results indicate that periods of higher or lower returns identified by technical trading rules largely correspond to those of higher or lower conditional

expected returns, due to high or low risk premia and volatility. Therefore, the majority of technical trading rule profits might well be a result of time varying risk premia. While the link between market risk and trading rule return provides only a mild approximation for a relationship between market risk and trading rule survival, it is definitely a relationship that is worth exploring. Given the very high level of adaptability of shorter term trading rules, one would expect that currency volatility has a neutral impact on survival time, while longer term trading rules, where there is lower adaptability, should be more negatively exposed. The rationale for choosing skew and kurtosis as external factors comes from the characteristics of the data, shown in Tables 1 to 3 which suggest a fair deviation from normality of many of the underlying currency data. As mentioned earlier, the analysis in this section is done on two dimensions, on a cross sectional basis as well as on a time series basis. The cross sectional analysis is structured as follows. Average interest differential, standard deviation; skew and kurtosis are calculated for all currency pairs. For each of the 39 moving average combinations the combined product limit estimator for positive and negative momentum signals is constructed and the average survival time, as shown in equation 39 in the next section, is calculated. In a second step for each trading rule parameterisation (i.e. SR1/LR5 or SR1/LR10) Spearman's rank correlation, between average survival time and average interest differential, standard deviation; skew and kurtosis are calculate across all exchange rates is calculated.


The rationale for using Spearman's rank correlation coefficient is due to the fact that it is a nonparametric measure. Hence, it does not depend on the distributional characteristic of the underlying data, which becomes important when, for instance, looking at relationship between survival time and kurtosis, which appears highly nonlinear, as shown in Figure 1-10. The Idea behind spearman's rank correlation is to analysis the strength of the monotonic relationship between two variables. Under monotonic relationship one can understand the relationship of the rank order of elements in different data sets. Hence, it facilitates the assessment as to whether higher or lower survival times of different trading rule parameterisations relate to higher or lower levels of external variables even if their relationship is non linear. Figure 1-11 shows the same relationship between survival time and kurtosis, however, on a ranked order, i.e. in the context of the monotonic relationship as described earlier. Spearman's sank correlation coefficient is defined between -1 and 1, whereby -1 indicates a perfect negative monotonic relationship and 1 a perfect positive relationship.

FIGURE 1-10: SR1/LR10 TRADING RULE: AVERAGE SURVIVAL TIME VS. KURTOSIS



The Figure shows a scatterplot of average survival time versus levels of kurtosis of the SR1/LR10 trading rule across all currency pairs. In The upper right corner the line of best fit is shown. The relationship is logarithmic, with a goodness of fit of 0.38.

FIGURE 1-11: SR1/LR10 TRADING RULE: AVERAGE SURVIVAL TIME VS. KURTOSIS



The Figure shows a scatterplot of the ranked values of average survival time versus levels of kurtosis of the SR1/LR10 trading rule across all currency pairs. The upper right corner the line of best fit is shown. The relationship is linear, with a goodness of fit of 0.4

Absolute ranges between 0 to 0.4 indicates a weak relationship, 0.4 to 0.6 a medium relationship and everything above 0.8 or below -0.8 indicate a strong relationship. Moreover, this section also analyses the statistical validity of the rank correlation coefficient.

FIGURE 1-12: CROSS SECTIONAL ANALYSIS OF AVERAGE SURVIVAL VS EXTERNAL FACTORS

CROSS SECTIONAL ANALYSIS: Spearman Rank Correlation of Average Survival Time vs.							
		LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	Average CCY Volatility	-0.15	-0.20	-0.21	-0.27 *	-0.17	-0.13
	Average Interest Rate Differential	0.22	-0.03	-0.06	0.04	-0.03	-0.02
	Average CCY Skewness	0.43 ***	0.31 **	0.28 *	0.33 **	0.26 *	0.23
	Average CCY Kurtosis	-0.67 ***	-0.63 ***	-0.60 ***	-0.50 ***	-0.54 ***	-0.50 ***
SR 2	Average CCY Volatility	-0.18	-0.25	-0.24	-0.15	-0.06	-0.16
	Average Interest Rate Differential	0.12	-0.07	-0.13	-0.09	-0.20	0.03
	Average CCY Skewness	0.47 ***	0.22	0.20	0.21	0.12	0.20
	Average CCY Kurtosis	-0.72 ***	-0.56 ***	-0.46 ***	-0.42 ***	-0.36 **	-0.31 **
SR 3	Average CCY Volatility	-0.05	-0.20	-0.21	-0.12	-0.08	-0.17
	Average Interest Rate Differential	0.02	-0.15	-0.16	-0.14	-0.11	0.00
	Average CCY Skewness	0.44 ***	0.19	0.13	0.18	0.15	0.19
	Average CCY Kurtosis	-0.70 ***	-0.52 ***	-0.38 **	-0.35 **	-0.34 **	-0.32 **
SR 4	Average CCY Volatility	-0.17	-0.18	-0.22	-0.19	-0.10	-0.18
	Average Interest Rate Differential	0.17	-0.19	-0.14	-0.08	-0.12	0.02
	Average CCY Skewness	0.48 ***	0.21	0.17	0.21	0.12	0.18
	Average CCY Kurtosis	-0.64 ***	-0.56 ***	-0.36 **	-0.29 **	-0.25	-0.21
SR 5	Average CCY Volatility		-0.22	-0.24	-0.17	-0.13	-0.16
	Average Interest Rate Differential		-0.14	-0.09	-0.14	-0.11	-0.01
	Average CCY Skewness		0.23	0.19	0.17	0.12	0.15
	Average CCY Kurtosis		-0.47 ***	-0.34 **	-0.22	-0.23	-0.20
SR 10	Average CCY Volatility			-0.24	-0.25 *	-0.20	-0.22
	Average Interest Rate Differential			-0.13	-0.01	0.06	0.06
	Average CCY Skewness			0.13	0.12	0.14	0.13
	Average CCY Kurtosis			-0.29 *	-0.14	-0.05	-0.12
SR 15	Average CCY Volatility				-0.31 **	-0.22	-0.23
	Average Interest Rate Differential				-0.04	-0.03	-0.05
	Average CCY Skewness				0.16	0.11	0.13
	Average CCY Kurtosis				-0.22	-0.18	-0.16
SR 20	Average CCY Volatility					-0.28 *	-0.14
	Average Interest Rate Differential					-0.04	-0.01
	Average CCY Skewness					0.18	0.20
	Average CCY Kurtosis					-0.24	-0.25 *
SR 25	Average CCY Volatility						-0.19
	Average Interest Rate Differential						0.03
	Average CCY Skewness						0.26 *
	Average CCY Kurtosis						-0.28 *

The Figure shows Spearman's Rank correlation between average survival time and average interest differential, standard deviation; skew and kurtosis are calculated for all currency pairs. The index of short term parameters of the moving average combinations are shown along the first column, the index of long term parameters is shown along the first row. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

Figure 1-12 shows Spearman's rank correlation between average survival time and average interest differential, standard deviation; skew and kurtosis are calculated for all currency pairs. The index of short term parameters of the moving average combinations

are shown along the first column, the index of long term parameters is shown along the first row. For each of the moving average combinations four rank correlation coefficients against the four variables discussed earlier are shown. Stars next to the rank correlation coefficients indicate the level of statistical significance. The results of this analysis suggest a mild negative cross sectional relationship between currency volatility and average survival time. This means that moving average signals of currency pairs that exhibit higher volatility tend to live shorter than moving average signals derived from lower volatility currency pairs. However, this relationship is only statistically significant in very few cases. It tends to become more significant for longer moving average pairs. The correlation between average survival time and the average interest rate differential is almost neutral and not statistically significant. When it comes to skew there is a mildly positive relationship. This indicates that momentum signals that are based on currency crosses with a positive skew have higher average survival time. While the relationship is only mildly positive, the level of statistical significance is reasonably high, namely, when looking at shorter term moving average combinations. Finally there is a strong and statistically significant negative relationship between average survival time and kurtosis. This indicates that momentum signals, which are based on currency pairs that exhibit high levels of kurtosis, tend to live shorter than others, which exhibit lower levels of kurtosis.

As mentioned earlier, the analysis presented in this section is two-dimensional. The first step looks at the cross sectional differences in survival times, the second step looks at the time series differences in survival times. This is done by calculating the product limit estimator for each of the 39 moving average combinations across all currency pairs and across all nine sub-samples. Spearman's rank correlation coefficient is then calculated across all nine sub-samples for each of the currency pairs and each of the moving average combinations. Taking the currency pair GBPUSD and parameterisation SR1/LR5 as an example, the average survival time average interest rate differential, standard deviation; skew and kurtosis are calculated for each of the nine sub samples. Spearman's rank correlation coefficient across the different sub samples is then calculated for each of the currency pairs and parameterisations. The set of Spearman's rank correlations associated with each trading rule parameterisation is averaged across all currency pairs and used as final output. This is shown in Figure 1-13. Figure 1-13 is structured similarly to Figure 1-12, whereby the index of short term parameters of the moving average combinations are shown along the first column, the index of long term parameters is shown along the first row. For each of the moving average combinations

four rank correlation coefficients against the four variables discussed earlier are shown. The results of this analysis indicate that there is no meaningful time series relationship in the data. While there is a slightly positive correlation between average survival time and levels of kurtosis, this relationship is not statistically significant.

FIGURE 1-13: TIME SERIES ANALYSIS OF AVERAGE SURVIVAL VS EXTERNAL FACTORS

TIME SERIES ANALYSIS: Spearman Rank Correlation of Average Survival Time vs.							
		LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	Average CCY Volatility	0.06	-0.07	-0.01	-0.03	0.00	0.06
	Average Interest Rate Differential	0.07	0.03	-0.02	0.02	-0.01	-0.03
	Average CCY Skewness	0.03	0.00	0.01	0.09	0.06	0.00
	Average CCY Kurtosis	-0.06	0.00	0.08	0.05	0.16	0.21
SR 2	Average CCY Volatility	0.05	-0.05	-0.02	-0.01	0.00	0.05
	Average Interest Rate Differential	0.07	-0.02	-0.02	0.00	-0.06	-0.05
	Average CCY Skewness	0.05	0.02	0.02	0.06	0.04	-0.04
	Average CCY Kurtosis	-0.03	0.08	0.15	0.19	0.22	0.30
SR 3	Average CCY Volatility	0.01	-0.06	0.01	0.00	0.02	0.01
	Average Interest Rate Differential	0.04	-0.04	-0.02	-0.03	-0.07	-0.05
	Average CCY Skewness	0.01	0.00	0.02	0.03	0.03	-0.03
	Average CCY Kurtosis	-0.07	0.11	0.13	0.18	0.24	0.30
SR 4	Average CCY Volatility	0.04	-0.08	-0.03	0.01	0.01	0.02
	Average Interest Rate Differential	0.11	-0.07	-0.02	-0.04	-0.04	-0.08
	Average CCY Skewness	0.03	0.04	0.02	0.04	0.03	-0.03
	Average CCY Kurtosis	-0.08	0.08	0.15	0.21	0.22	0.31
SR 5	Average CCY Volatility		-0.05	-0.05	0.00	0.00	0.02
	Average Interest Rate Differential		-0.05	-0.02	-0.07	-0.05	-0.06
	Average CCY Skewness		0.09	0.03	0.05	0.01	-0.03
	Average CCY Kurtosis		0.11	0.15	0.17	0.25	0.31
SR 10	Average CCY Volatility			-0.03	0.02	0.00	-0.03
	Average Interest Rate Differential			-0.02	-0.07	-0.06	-0.05
	Average CCY Skewness			0.10	0.06	0.00	-0.02
	Average CCY Kurtosis			0.12	0.14	0.24	0.26
SR 15	Average CCY Volatility				0.01	-0.02	0.04
	Average Interest Rate Differential				-0.05	-0.05	-0.08
	Average CCY Skewness				0.07	-0.02	-0.07
	Average CCY Kurtosis				0.15	0.21	0.27
SR 20	Average CCY Volatility					0.01	0.04
	Average Interest Rate Differential					-0.03	-0.06
	Average CCY Skewness					-0.01	-0.01
	Average CCY Kurtosis					0.19	0.24
SR 25	Average CCY Volatility						0.03
	Average Interest Rate Differential						-0.07
	Average CCY Skewness						0.01
	Average CCY Kurtosis						0.23

The Figure shows Spearman's Rank correlation between average survival time and average interest differential, standard deviation; skew and kurtosis are calculated for all currency pairs. The index of short term parameters of the moving average combinations are shown along the first column, the index of long term parameters is shown along the first row. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

This is shown in Figure 1-13. Figure 1-13 is structured similarly to Figure 1-12, whereby the index of short term parameters of the moving average combinations are

shown along the first column, the index of long term parameters is shown along the first row. For each of the moving average combinations four rank correlation coefficients against the four variables discussed earlier are shown. The results of this analysis indicate that there is no meaningful time series relationship in the data. While there is a slightly positive correlation between average survival time and levels of kurtosis, this relationship is not statistically significant. Given the fact that in this analysis Spearman's Rank correlation coefficient is calculated on only very few data points, i.e. 9 data points for each of the sub-samples defined earlier, one could expect lower levels of statistical significance. To test whether the lack of statistical significance stems from the low number of data points, the same analysis is conducted by splitting the sample period in 36 sub samples 35 of which consist of 250 data points and the last of 275 data points. The results of this analysis shown in Appendix 4 are similar to the results shown in Figure 1-13. They do not allow for any meaningful conclusions either.

In summary it can be stated that the analysis provides different conclusions, depending on the perspective from which the relationship between survival curves and external factors is analysed. When it comes to a cross sectional perspective, higher levels of currency volatility and kurtosis have in some cases a statistically significant negative impact on average survival time of momentum signals, while skew has a mildly positive impact on average survival times of momentum signals. However, from a time series perspective the analysis does not yield a clear conclusion. The fact that average survival time and survivorship curves are considerably impacted by skewness and kurtosis, the choice of a benchmark against which empirical curves are assessed becomes of high importance. This is discussed at length in the next section.

E. Resampling Techniques and Testing for Market Efficiency.

One central aim of the chapter is to evaluate whether empirical survivorship functions have unusual pattern, which cannot be explained by theoretical benchmark processes. Hence it becomes crucial to define benchmark processes that comprise a fair representation of the return generating process of the various exchange rate pairs. This comes with a certain amount of challenges.

1. Distributional characteristics in the underlying exchange rates

Firstly, Tables 1-1 to 1-3 indicate that exchange rate returns do not follow a normal distribution. Many exchange rate pairs exhibit a fair degree of skew and kurtosis. Moreover, the results of the previous section indicate that average survival times derived by the product limit estimator are impacted by distributional characteristics of the underlying foreign exchange data. Therefore, it is essential to create data simulations that incorporate the distributional characteristics of the underlying data. This can be done in various ways. While Monte Carlo simulations as proposed by Jochum (2000) and Kos and Todorovic (2008), can be conditioned to fit the specific distributional characteristics of the underlying data. One clear disadvantage of this implementation is the fact that one has to make assumptions about these distributional characteristics, which bears potential estimation errors. Moreover, as noted earlier the Product Limit Estimator is non-parametric. Hence, the use of a benchmark process that does not require any assumptions about distributional characteristics of the underlying data is a more appropriate simulation setup. Re-sampling is a non-parametric simulation approach that facilitates conducting standard statistical tests for any given dataset. By reshuffling the original dataset multiple times, the simulation will replicate a data distribution that is on average the same as the distribution of the original dataset.

When applying any re-sampling technique, it has to be decided whether it is preferable to conduct the simulation with replacement or without replacement. Sampling with replacement means that after drawing any random observation from the original sample, the observation is put back before drawing the next observation. This process is also called bootstrapping. Efron introduced this methodology in 1979. Since then it has found a very wide use in statistics and also in finance. Permutation, which is re-sampling without replacement, would result in the same set of numbers, however in a different order. While Karolyi and Kho (2004) point out that in the finance space the majority of studies employ re-sampling with replacement, general statistics literature gives only limited guidance as to which simulation methodology is preferable.

Both approaches have their advantages as well as disadvantages depending on the test setup. For example when it comes to the traditional cross sectional momentum literature Jegadeesh and Titman (2002) indicate that simulations with replacement bear a bias arising from the possibility that return observations are drawn from evaluation as well

as holding periods. In their view permutation is a more appropriate test setup for their cross sectional momentum analysis introduced in 1993.

FIGURE 1-14: TEST FOR AUTOCORRELATION AND HETEROSKEDASTICITY IN FOREIGN EXCHANGE RETURNS

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
F-Statistic (Breusch-Godfrey Serial Correlation LM Test)		7.12 ***	1.23	1.18	1.30	7.43 ***	6.07 ***	3.30 ***	7.01 ***	6.22 ***
F-Statistic (ARCH LM Test)		0.94	1.59	1.93 *	4.07 ***	7.89 ***	0.05	6.46 ***	0.01	0.03
GBP										
F-Statistic (Breusch-Godfrey Serial Correlation LM Test)	6.84 ***		9.37 ***	1.93 *	0.85	3.87 ***	1.52	2.77 **	1.46	3.69 ***
F-Statistic (ARCH LM Test)	0.79		4.77 ***	3.70 ***	5.23 ***	5.76 ***	0.02	2.79 **	0.14	0.05
JPY										
F-Statistic (Breusch-Godfrey Serial Correlation LM Test)	1.17	9.90 ***		4.34 ***	0.69	2.06 *	5.03 ***	0.75	4.31 ***	4.47 ***
F-Statistic (ARCH LM Test)	1.47	4.59 ***		4.43 ***	2.94 **	5.84 ***	0.06	2.56 **	0.02	0.02
EUR										
F-Statistic (Breusch-Godfrey Serial Correlation LM Test)	1.21	2.03 *	4.14 ***		30.65 ***	24.57 ***	4.77 ***	0.89	2.46 **	4.20 ***
F-Statistic (ARCH LM Test)	1.45	2.51 **	3.45 ***		1.01	0.87	0.02	3.47 ***	0.02	0.02
CHF										
F-Statistic (Breusch-Godfrey Serial Correlation LM Test)	1.30	0.85	0.68	30.70 ***		14.38 ***	3.08 ***	1.14	3.45 ***	4.73 ***
F-Statistic (ARCH LM Test)	3.47 ***	3.13 ***	1.91 *	1.25		1.87 *	0.05	4.36 ***	0.01	0.03
NOK										
F-Statistic (Breusch-Godfrey Serial Correlation LM Test)	7.81 ***	3.90 ***	2.18 *	25.58 ***	14.70 ***		10.99 ***	8.69 ***	8.14 ***	10.58 ***
F-Statistic (ARCH LM Test)	10.49 ***	8.25 ***	7.86 ***	1.36	3.03 ***		0.02	11.95 ***	0.01	0.02
SEK										
F-Statistic (Breusch-Godfrey Serial Correlation LM Test)	6.42 ***	1.63	5.14 ***	6.13 ***	3.88 ***	12.27 ***		6.36 ***	8.08 ***	10.92 ***
F-Statistic (ARCH LM Test)	0.07	0.03	0.09	0.02	0.06	0.03		0.07	0.03	0.03
CAD										
F-Statistic (Breusch-Godfrey Serial Correlation LM Test)	3.39 ***	2.87 **	0.82	0.86	1.19	8.53 ***	6.04 ***		21.88 ***	12.98 ***
F-Statistic (ARCH LM Test)	6.21 ***	2.70 **	2.64 **	3.84 ***	4.34 ***	9.23 ***	0.03		0.11	0.01
AUD										
F-Statistic (Breusch-Godfrey Serial Correlation LM Test)	7.94 ***	1.64	4.86 ***	2.77 **	3.81 ***	8.73 ***	8.60 ***	23.29 ***		27.10 ***
F-Statistic (ARCH LM Test)	0.01	0.25	0.04	0.05	0.03	0.04	0.05	0.18		0.04
NZD										
F-Statistic (Breusch-Godfrey Serial Correlation LM Test)	7.57 ***	4.26 ***	5.81 ***	4.90 ***	5.52 ***	11.70 ***	11.68 ***	14.90 ***	30.56 ***	
F-Statistic (ARCH LM Test)	0.06	0.09	0.04	0.04	0.06	0.04	0.04	0.02	0.03	

The table shows the results of the Breusch-Godfrey Serial Correlation LM test as well as the results of the ARCH LM test. The column labels denote base currency calculations and row labels denote foreign currencies returns against the base currency. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

One general counter argument in favour of resampling with replacement is that simple permutation bears the risk of a small-sample bias. This is due to the fact that permutation restricts simulated time series to be of the same length than the original time series and each single observation of the original time series is present in every simulated time series. Hence, few data observations, which might have a very low probability of occurrence in the empirical time series, will have a significant impact in the simulated time series. Other tradeoffs between resampling with and without replacement revolve around consistency and on rates of convergence. In that context Politis and Romano (1994) show that permutation consistently estimates the distribution of a statistic under very weak conditions. Horowitz (2001) indicates that, while conditions for consistency in resampling with replacement are not difficult to establish, permutation is useful when such conditions are not met. Generally, resampling with replacement is used to estimate confidence intervals, while permutation is more commonly used for hypothesis testing, where sample size is sufficiently large.

Hesterberg et al. (2003) emphasise that permutation as a tool for hypothesis testing becomes particularly useful when the underlying dataset is not normally distributed, as it enables the calculation of sampling distribution without making any assumptions about the shape or other parameters of the population distributions. However, the restriction that comes with permutation tests is the fact that single observations in the data have to be interchangeable. Hence data dependencies such as volatility clustering and autocorrelation within the underlying data cannot be fully reflected in this simulation method. Given the high degree of non normality in the underlying foreign exchange data, one of the data simulations implemented in this chapter uses permutation, as given the fact that it produces consistent estimates even under very weak assumptions with regards to distributional characteristics. Moreover, the sample size is considerable. Hence, any biases stemming from sample size are negligible. However, as shown in Figure 1-14 the underlying currency data show some signs of data dependencies. Most of the exchange rate returns show high levels of autocorrelation, and some degree of volatility clustering. Figure 1-14 shows the F-values for the Breusch-Godfrey LM test for serial correlation and Engle's ARCH LM test. Both tests are conducted to detect effects of serial correlation, as well as ARCH effects up to a lag length of ten days. The maximum lag length of ten days was chosen arbitrarily, yet it implies a time period of two weeks, which seems appropriate. In order to account for these data biases two further simulation methods are implemented. The first of which is a stationary bootstrap proposed by Politis and Romano (1994), which aims to keep the autocorrelation structure of foreign exchange returns intact, the second is a GARCH(1,1) based bootstrap, which aims to replicate the volatility structure of empirical data

2. Serial Correlation and Heteroskedasticity in the underlying exchange rates

As indicated in the previous section, foreign exchange returns exhibit a degree of serial correlation, which undermine the results of the permutation test to some degree. In order to adjust for that a stationary bootstrap simulation is undertaken, which embeds the serial correlation of the underlying data in the universe resampled time series. Politis and Romano (1994), introduce this variation of the block bootstrap, which is called the stationary bootstrap. While the block bootstrap resamples data blocks of fixed length,

the stationary bootstrap resamples data blocks of random length. The length size of data blocks; follows hereby a geometric distribution, with a mean block length of $1/q$, q being a smoothing parameter that has to be chosen. Both bootstrap methodologies are able to replicate mild data dependencies. However, the main difference between the regular block bootstrap and the variation of Politis and Romano (1994) is the fact that under their specification, the resampled time series will be stationary providing the original time series to be stationary. Under the specification of Politis and Romano (1994), larger values of q generate shorter block length, while lower levels of q increase the block length. This chapter follows Sullivan, Timmerman and White (1999) as well as Qi and Wu (2006) closely, who propose a value of 0.1 for q , which equates to an average block length of 10 observations.

Moreover, as highlighted earlier, the disadvantage of the simple permutation approach is the fact that a mere reshuffling of returns will break the volatility structure of a time series. This might raise questions about the appropriateness of the nonparametric resampling approach. In order to control for this, this chapter will also conduct a set of simulations that use a resampling methodology, which leaves the original volatility structure intact. This is done, by embedding a GARCH (1,1) process into the resampling simulation. The chapter hereby follows the works of Pascual, Romo and Ruiz (2005). The GARCH model has been developed independently by Bollerslev (1986) and Taylor (1986). It allows the conditional variance to be dependent on previous observations. The GARCH (1,1) specification is in this context very popular insofar as it captures volatility clustering to a sufficient degree, while being parsimonious in its estimation; it is given in Equations 32 and 33.

$$(32) \quad y_t = \sigma_t \varepsilon_t$$

$$(33) \quad \sigma_t^2 = \omega + \alpha y_{t-1}^2 + \beta \sigma_{t-1}^2, t = 1, \dots, T,$$

Whereby the process y_t can be defined by its variance σ_t and a white noise process ε_t . The variance σ_t is defined by parameters ω, α, β , whereby $\omega > 0, \alpha \geq 0, \beta \geq 0$ and $\alpha + \beta \leq 1$. The unknown parameters are estimated by using a quasi maximum likelihood estimation (QML) that is based on the original set of time series data for all the currency crosses, as shown in Appendix 5. After the estimation of the parameters $\hat{\omega}, \hat{\alpha}, \hat{\beta}$ the residuals can be estimated as per Equations 34 to 36

$$(34) \quad \hat{\varepsilon}_t = \frac{y_t}{\hat{\sigma}_t}, t = 1, \dots, T,$$

whereby the estimates for marginal variance, e.g. the starting point and conditional variance are given in Equations 35 and 36

$$(35) \quad \hat{\sigma}_1^2 = \frac{\hat{\omega}}{(1-\hat{\alpha}-\hat{\beta})}$$

$$(36) \quad \hat{\sigma}_t^2 = \hat{\omega} + \hat{\alpha}y_{t-1}^2 + \hat{\beta}\hat{\sigma}_{t-1}^2, t = 2, \dots, T,$$

As mentioned before the aim of the GARCH (1,1) bootstrap is to estimate replicates of the original time series, $Y_T^* = \{y_1^*, \dots, y_T^*\}$ that mimic its original volatility structure. This is done by recursive estimation of to Equations 37 and 38, whereby ε_t^* represent random draws from the empirical time series

$$(37) \quad y_t^* = \hat{\sigma}_t^* \varepsilon_t^*, t = 1, \dots, T,$$

$$(38) \quad \hat{\sigma}_t^{*2} = \hat{\omega} + \hat{\alpha}y_{t-1}^{*2} + \hat{\beta}\hat{\sigma}_{t-1}^{*2}$$

For each of the proposed bootstrap simulations, the log-rank test simulation is designed in such way that after every permutation/resampling a log-rank test between the empirical survivorship curve and the simulated survivorship curve is calculated. As noted earlier, for large datasets the log-rank test statistic expressed in Equation 24 follows a standard normal distribution. Hence repeatedly recalculating the test statistic will yield a distribution of the test statistic from which inferences can be made. The number of iterations is chosen to be 500; this number resembles a good trade-off between accuracy of results and calculation time.

F. Empirical Results

1. Survivorship curves and log-rank Test Results.

As pointed out in the previous section, the initial step when creating survivorship curves is to formalise a positive or negative trading signal upon which the survivorship curves are based. The definition of the trading signal can be chosen arbitrarily as long as it is applied to both the empirical series and the simulations. The filter rules used in this chapter and introduced in a previous section look at the difference between short and long-term return moving averages. If the short-term return moving average is above (or below) the long-term return moving average, then a positive (or negative) signal is obtained.

Figure 1-15 shows the empirical survival curve for positive momentum signals from the SR1/LR10 moving average combination for the USDGBP exchange rate. The results suggest that during the sample period, there have been 4445 observations where the prevailing price is above the ten day average price. Out of these 4445 observations, 3686 observations survive one further day. Hence, the Product Limit Estimator from day one to two is 82.92%. The probability of survival beyond three days is 60%. The survival probability diminishes to less than 10% after 14 days. The average survival time can be calculated as the sum of the periodical survival probabilities multiplied by the respective time increment. Daily returns are used as time increments. Hence, the average survival time represents the sum of all periodical PLE estimates. As given in Equation 39:

$$(39) \quad AVG \ SURVIVAL \ TIME = \sum_{i=1}^k \hat{S}(t)$$

The average survival time of the positive moving average curve amounts to 5.99 days. To assess whether a survival time of 6 days can be reasonably expected for this moving average combination, Figure 1-16 shows the survival rates that are obtained using a re-sampling simulation based on the same time series and the same filter rule.

FIGURE 1-15: EMPIRICAL PLE CURVE; POSITIVE RETURNS; USDGBP; SR1/LR10

	Ordered failure time	intact before t	ending at time t	contribution to KM estimator	KM estimator	Variance
j	t(j)	n _j	d _j	(n _j '/n _j)	S(t)	VAR(S(t))
1	2	4445	759	82.92%	82.92%	0.00003
2	3	3686	547	85.16%	70.62%	0.00002
3	4	3139	444	85.86%	60.63%	0.00002
4	5	2695	379	85.94%	52.10%	0.00002
5	6	2316	334	85.58%	44.59%	0.00001
6	7	1982	289	85.42%	38.09%	0.00001
7	8	1693	243	85.65%	32.62%	0.00001
8	9	1450	216	85.10%	27.76%	0.00001
9	10	1234	176	85.74%	23.80%	0.00001
10	11	1058	149	85.92%	20.45%	0.00001
11	12	909	125	86.25%	17.64%	0.00001
12	13	784	103	86.86%	15.32%	0.00000
13	14	681	93	86.34%	13.23%	0.00000
14	15	588	81	86.22%	11.41%	0.00000
15	16	507	66	86.98%	9.92%	0.00000
16	17	441	56	87.30%	8.66%	0.00000
17	18	385	50	87.01%	7.54%	0.00000
18	19	335	39	88.36%	6.66%	0.00000
19	20	296	34	88.51%	5.89%	0.00000
20	21	262	30	88.55%	5.22%	0.00000
21	22	232	25	89.22%	4.66%	0.00000
22	23	207	20	90.34%	4.21%	0.00000
23	24	187	19	89.84%	3.78%	0.00000
24	25	168	18	89.29%	3.37%	0.00000
25	26	150	17	88.67%	2.99%	0.00000
26	27	133	15	88.72%	2.65%	0.00000
27	28	118	14	88.14%	2.34%	0.00000
28	29	104	11	89.42%	2.09%	0.00000
29	30	93	8	91.40%	1.91%	0.00000
30	31	85	8	90.59%	1.73%	0.00000
31	32	77	7	90.91%	1.57%	0.00000
32	33	70	6	91.43%	1.44%	0.00000
33	34	64	6	90.63%	1.30%	0.00000
34	35	58	6	89.66%	1.17%	0.00000
35	36	52	5	90.38%	1.06%	0.00000
36	37	47	3	93.62%	0.99%	0.00000
37	38	44	3	93.18%	0.92%	0.00000
38	39	41	3	92.68%	0.85%	0.00000
39	40	38	3	92.11%	0.79%	0.00000
40	41	35	3	91.43%	0.72%	0.00000
41	42	32	3	90.63%	0.65%	0.00000
42	43	29	3	89.66%	0.58%	0.00000
43	44	26	3	88.46%	0.52%	0.00000
44	45	23	3	86.96%	0.45%	0.00000
45	46	20	3	85.00%	0.38%	0.00000
46	47	17	2	88.24%	0.34%	0.00000
47	48	15	2	86.67%	0.29%	0.00000
48	49	13	2	84.62%	0.25%	0.00000
49	50	11	2	81.82%	0.20%	0.00000
50	51	9	2	77.78%	0.16%	0.00000
51	52	7	2	71.43%	0.11%	0.00000
52	53	5	1	80.00%	0.09%	0.00000
53	54	4	1	75.00%	0.07%	0.00000
54	55	3	1	66.67%	0.05%	0.00000
55	56	2	1	50.00%	0.02%	0.00000
56	57	1	1	0.00%	0.00%	0.00000

The Figure shows the product limit estimator, the survival curve for positive momentum signals from the SR1/LR10 moving average combination for the USDGBP exchange rate. The first column labelled j is the count of observation intervals. The second column is the index of survival time t(i). The third column shows the number of observations intact before time t. The fourth column shows the observations that cease to exist at time t. the fifth and sixth column show the conditional and absolute survival probability at time t.

FIGURE 1-16 SIMULATED PLE POSITIVE RETURNS; USDGBP; SR1/LR10

Survival Function of POSITIVE Market Momentum (standard resampling)					Survival Function of POSITIVE Market Momentum (GARCH(1,1) resampling)				
	Ordered failure time	KM estimator	Variance	Significanc e Test		Ordered failure time	KM estimator	Variance	Significanc e Test
j	t(j)	S(t)	VAR(S(t))	T-Stat	j	t(j)	S(t)	VAR(S(t))	T-Stat
1	2	83.04% ***	0.004816	172.4284	1	2	82.73% ***	0.004744	174.4111
2	3	70.32% ***	0.007976	88.1661	2	3	69.95% ***	0.007223	96.84653
3	4	59.78% ***	0.010314	57.95755	3	4	59.57% ***	0.009253	64.3829
4	5	50.73% ***	0.012021	42.19906	4	5	50.71% ***	0.010308	49.18919
5	6	42.89% ***	0.012928	33.17545	5	6	43.06% ***	0.011053	38.95945
6	7	36.08% ***	0.013565	26.5955	6	7	36.43% ***	0.011395	31.97102
7	8	30.15% ***	0.013977	21.57043	7	8	30.70% ***	0.011099	27.65472
8	9	25.14% ***	0.014434	17.41832	8	9	25.81% ***	0.01066	24.21636
9	10	20.92% ***	0.014863	14.07602	9	10	21.67% ***	0.010698	20.25335
10	11	17.52% ***	0.014548	12.04209	10	11	18.28% ***	0.0105	17.40648
11	12	14.66% ***	0.014372	10.19654	11	12	15.43% ***	0.010489	14.71567
12	13	12.28% ***	0.013531	9.072796	12	13	13.08% ***	0.010387	12.59181
13	14	10.29% ***	0.012718	8.090354	13	14	11.08% ***	0.010141	10.92233
14	15	8.63% ***	0.012309	7.013936	14	15	9.33% ***	0.009711	9.603789
15	16	7.22% ***	0.011623	6.214947	15	16	7.84% ***	0.009365	8.370274
16	17	6.06% ***	0.010646	5.695002	16	17	6.58% ***	0.008871	7.412549
17	18	5.08% ***	0.010056	5.04795	17	18	5.54% ***	0.008456	6.552043
18	19	4.24% ***	0.009834	4.316327	18	19	4.68% ***	0.007678	6.093283
19	20	3.58% ***	0.00938	3.814482	19	20	3.94% ***	0.007055	5.583889
20	21	3.03% ***	0.008775	3.455387	20	21	3.33% ***	0.006747	4.935333
21	22	2.58% ***	0.008089	3.187164	21	22	2.82% ***	0.006403	4.396911
22	23	2.19% ***	0.007473	2.9271	22	23	2.38% ***	0.006045	3.931411
23	24	1.86% ***	0.006732	2.767654	23	24	2.00% ***	0.005636	3.545608
24	25	1.58% ***	0.006125	2.583668	24	25	1.69% ***	0.005121	3.303431
25	26	1.34% **	0.005568	2.39928	25	26	1.43% ***	0.004643	3.079462
26	27	1.12% **	0.005034	2.223252	26	27	1.21% ***	0.004178	2.897288
27	28	0.93% **	0.004603	2.019462	27	28	1.02% ***	0.003765	2.716116
28	29	0.78% *	0.00421	1.844579	28	29	0.85% ***	0.003277	2.601013
29	30	0.65% *	0.00384	1.695652	29	30	0.70% **	0.002875	2.44321
30	31	0.54%	0.003558	1.515739	30	31	0.59% **	0.002574	2.277532
31	32	0.45%	0.003258	1.375252	31	32	0.49% **	0.002316	2.106239
32	33	0.37%	0.002905	1.259453	32	33	0.41% **	0.002029	1.995771
33	34	0.30%	0.0026	1.135229	33	34	0.33% *	0.001761	1.868588
34	35	0.24%	0.002265	1.041379	34	35	0.27% *	0.001568	1.698507
35	36	0.19%	0.001934	0.971857	35	36	0.22%	0.001414	1.551202
36	37	0.15%	0.001639	0.896062	36	37	0.18%	0.001309	1.402753
37	38	0.11%	0.001339	0.822616	37	38	0.15%	0.001173	1.316792
38	39	0.08%	0.001032	0.801097	38	39	0.13%	0.001063	1.179422
39	40	0.06%	0.000852	0.727677	39	40	0.10%	0.000959	1.050682
40	41	0.04%	0.0007	0.590549	40	41	0.08%	0.000854	0.943681
41	42	0.03%	0.000578	0.43658	41	42	0.06%	0.000755	0.831087
42	43	0.02%	0.000506	0.36319	42	43	0.05%	0.000677	0.728908
43	44	0.01%	0.000436	0.316228	43	44	0.04%	0.00059	0.645663
44	45	0.01%	0.000363	0.316228	44	45	0.03%	0.000498	0.586262
45	46	0.01%	0.000291	0.316228	45	46	0.02%	0.000397	0.565763
46	47	0.01%	0.000218	0.316228	46	47	0.02%	0.000301	0.522095
47	48	0.00%	0.000145	0.316228	47	48	0.01%	0.000218	0.413294
48	49	0.00%	7.27E-05	0.316228	48	49	0.00%	0.000143	0.316228
					49	50	0.00%	7.14E-05	0.316228

The Figure shows the simulated product limit estimator, the simulated survival curve for positive momentum signals from the SR1/LR10 moving average combination for the USDGBP exchange rate. The curve on the left side has been obtained from re-sampling the USDGBP return time series 500 times based on standard permutation. The curve on the right side has been obtained from re-sampling the USDGBP return time series 500 times based on a GARCH (1,1) resampling. The first column in each of the two figures is labelled j and is the count of observation intervals. The second column is the index of survival time t(i). The third column shows the simulated Kaplan Meier estimator, with stars that denote the statistical significance. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%. The fourth column shows the variance of the estimator. The fifth column shows the test statistic of the Kaplan Meier estimator.


The figure is split into two parts, the column on the left is the benchmark survival curve that has been obtained from a resampling simulation using regular permutation; the column on the right shows the results that are obtained from the GARCH (1,1) simulation. From the results in Figures 1-15 and 1-16 it becomes evident that both simulated survival curves appear to be considerably shorter than the empirical curve.

The average survival time of the benchmark curve obtained by simple resampling is 5.27 days; the average survival time from the GARCH (1,1) simulation amounts to 5.37 days. Neither of the simulation methods can capture the extent of the actual average survival time, which is 5.99 days. The comparison between actual and theoretical survival time is based on the log-rank test, which facilitates the evaluation of differences between various survivorship curves. The test simulation is based on 500 iterations. For each of the iterations, a log-rank test between the empirical survival curve and the simulated survival curve is conducted. The final estimate for each log-rank test is the average value obtained from the series of 500 tests results. The analysis is carried out for all cross currency pairs across all ten base currencies. For each of the currency pairs the following moving average combination of Short Run SR (1, 2, 3, 4, 5, 10, 15, 20, 25) and Long Run LR (5, 10, 15, 20, 25, 30) are tested, whereby all Short Run moving averages have to be shorter than Long Run moving averages. This equates to 3510 moving average combinations that are tested. Given the vast amount of tests conducted, the results of the log-rank test will be presented in the form of heat maps. This makes the interpretation of the dataset more intuitive. Each of the following Figures will be based on the first dataset, as outlined in the methodology section.

2. Results for the Full Sample Period


Figures 1-17 to 1-22 show results of the log rank test, which have been obtained through the three simulation methods introduced earlier. The figures show positive and negative momentum signals separately. Each of the figures is split into eleven separate sub tables. The first ten tables show the outputs of the 10 base currencies, the eleventh shows the median value across all base currencies. The vertical axis of each heat map in the figures shows short-term moving averages and the horizontal axis long-term moving averages. The numbers in the figure show the results of the show the log-rank test results between the empirical and the simulated time series as z-values of the standard normal distribution.

FIGURE 1-17: LOG RANK TEST RESULTS, HEAT MAP FOR POSITIVE MOVING AVERAGE SIGNALS (SIMPLE RESAMPLING)




The vertical axis of each heat map in the Figure shows short-term moving averages and the horizontal axis long-term moving averages. The numbers in the Figure show the results of the log-rank test between the empirical and the simulated time series as z-values of the standard normal distribution. Statistical significance levels are shown two ways, implicitly via the heat map, and explicitly via stars next to the statistically significant numbers. Looking at the heat map the darkest red shade represents a significance level of less than 1%, the colour shades then change incrementally, as indicated in the legend of each of the Figures up to the darkest blue shade, which indicates a confidence level of 99% and more. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

FIGURE 1-18: LOG RANK TEST RESULTS, HEAT MAP FOR NEGATIVE MOVING AVERAGE SIGNALS (SIMPLE RESAMPLING)




The vertical axis of each heat map in the Figure shows short-term moving averages and the horizontal axis long-term moving averages. The numbers in the Figure show the results of the show the log-rank test results between the empirical and the simulated time series as z-values of the standard normal distribution. Statistical significance levels are shown two ways, implicitly via the heat map, and explicitly via stars next to the statistically significant numbers. Looking at the heat map the darkest red shade represents a significance level of less than 1%, the colour shades then change incrementally, as indicated in the legend of each of the Figures up to the darkest blue shade, which indicates a confidence level of 99% and more. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

FIGURE 1-19: LOG RANK TEST RESULTS, HEAT MAP FOR POSITIVE MOVING AVERAGE SIGNALS (GARCH(1,1) RESAMPLING)




The vertical axis of each heat map in the Figure shows short-term moving averages and the horizontal axis long-term moving averages. The numbers in the Figure show the results of the show the log-rank test results between the empirical and the simulated time series as z-values of the standard normal distribution. Statistical significance levels are shown two ways, implicitly via the heat map, and explicitly via stars next to the statistically significant numbers. Looking at the heat map the darkest red shade represents a significance level of less than 1%, the colour shades then change incrementally, as indicated in the legend of each of the Figures up to the darkest blue shade, which indicates a confidence level of 99% and more. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

FIGURE 1-20: LOG RANK TEST RESULTS, HEAT MAP FOR NEGATIVE MOVING AVERAGE SIGNALS (GARCH(1,1) RESAMPLING)




The vertical axis of each heat map in the Figure shows short-term moving averages and the horizontal axis long-term moving averages. The numbers in the Figure show the results of the show the log-rank test results between the empirical and the simulated time series as z-values of the standard normal distribution. Statistical significance levels are shown two ways, implicitly via the heat map, and explicitly via stars next to the statistically significant numbers. Looking at the heat map the darkest red shade represents a significance level of less than 1%, the colour shades then change incrementally, as indicated in the legend of each of the Figures up to the darkest blue shade, which indicates a confidence level of 99% and more. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

FIGURE 1-21: LOG RANK TEST RESULTS, HEAT MAP FOR POSITIVE MOVING AVERAGE SIGNALS (BLOCK RESAMPLING)



The vertical axis of each heat map in the Figure shows short-term moving averages and the horizontal axis long-term moving averages. The numbers in the Figure show the results of the log-rank test between the empirical and the simulated time series as z-values of the standard normal distribution. Statistical significance levels are shown two ways, implicitly via the heat map, and explicitly via stars next to the statistically significant numbers. Looking at the heat map the darkest red shade represents a significance level of less than 1%, the colour shades then change incrementally, as indicated in the legend of each of the Figures up to the darkest blue shade, which indicates a confidence level of 99% and more. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

FIGURE 1-22: LOG RANK TEST RESULTS, HEAT MAP FOR NEGATIVE MOVING AVERAGE SIGNALS (BLOCK RESAMPLING)



The vertical axis of each heat map in the Figure shows short-term moving averages and the horizontal axis long-term moving averages. The numbers in the Figure show the results of the log-rank test between the empirical and the simulated time series as z-values of the standard normal distribution. Statistical significance levels are shown two ways, implicitly via the heat map, and explicitly via stars next to the statistically significant numbers. Looking at the heat map the darkest red shade represents a significance level of less than 1%, the colour shades then change incrementally, as indicated in the legend of each of the Figures up to the darkest blue shade, which indicates a confidence level of 99% and more. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.


The darkest red shade represents a significance level of less than 1%, which means that the empirical observation is statistically significantly shorter than what is suggested by the benchmark process. The colour shades then change incrementally, as indicated in the legend of each of the figures up to the darkest blue shade, which indicates a confidence level of 99% and more. This result indicates that the empirical survivorship curve has statistically significantly longer life expectancy of momentum signals than suggested by the benchmark processes. Moreover, ***, ** and * indicate significance at the 1% and 5% and 10% levels, respectively. As shown in an earlier section the log rank follows a χ^2 distribution with one degree of freedom, which is then translated into a standard normal distribution via equation 24. Given the nature of the original test setup of the Log rank test, a one sided as opposed to a two sided test statistic has to be applied in the following figures.

The results in Figures 1-17 to 1-22 allow drawing various conclusions. First of all, under the simple, as well as the GARCH(1,1) resampling methodology, there is a fair number of log-rank tests, for positive as well as negative returns, that show statistically significant results. This is not only the case for smaller base currencies such as the NZD, AUD, NOK and SEK, but also for larger base currencies such as the USD and the JPY. This not the case when it comes to the results the log-rank tests of the stationary bootstrap simulation shown in Figures 1-21 and 1-22. While none of the results in these tables give any statistically significant results, it can still be argued that the reasonable level of statistical significance in Figures 1-17 to 1-20 presents valid evidence of systematic exposures. This is due to the fact that taking the median across base currencies reduces the strength of the results considerably. Moreover, given the fact that the technical trading rules applied in this chapter do focus on momentum, hence autocorrelation of currency returns, the results of Figures 1-21 and 1-22 should not come as a surprise either. Incorporating certain levels of autocorrelation into the benchmark assumption will automatically reduce the statistical significance of trading signals that aim to exploit the autocorrelation observed in the underlying data. Despite the fact that Figures 1-21 and 1-22 do not show any statistically significant results, the magnitude of the Log Rank test results of particularly the longer term moving averages are considerable. Looking at Figure 1-21 the SR3/LR30 for the USD is 1.13, which is very close to a statistically significant result.

Secondly, various currencies show differing levels of significance. Figures 1-17 and 1-20 indicate that positive momentum signals for the base currencies USD and JPY live

statistically shorter than what is suggested by the theoretical model. However, figures 1-17 to 1-20 also suggest that smaller currencies such as the NZD, AUD, NOK, SEK have positive and negative momentum signals that tend to live longer than suggested by benchmark simulations.

FIGURE 1-23: LOG RANK TEST RESULTS, HEAT MAP FOR MEDIAN VALUES OF POSITIVE AND NEGATIVE MOVING AVERAGE SIGNALS



The vertical axis of each heat map in the Figure shows short-term moving averages and the horizontal axis long-term moving averages. The numbers in the Figure show the results of the log-rank test between the empirical and the simulated time series as z-values of the standard normal distribution. Statistical significance levels are shown two ways, implicitly via the heat map, and explicitly via stars next to the statistically significant numbers. Looking at the heat map the darkest red shade represents a significance level of less than 1%, the colour shades then change incrementally, as indicated in the legend of each of the Figures up to the darkest blue shade, which indicates a confidence level of 99% and more. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

In many cases the statistical significance is relatively high. While the GARCH (1,1) simulation gives generally less strong results than regular permutation, it has to be pointed out that the differences between the two simulation methodologies are marginal. The results of the third simulation are generally less strong but they tend to follow similar pattern. As mentioned earlier, given the fact that taking the median generally tends to reduce the strength of results, the fair level of statistical significance of at least the first two simulation methods, presents evidence that shorter term moving averages tend to live longer than what theory suggests, while longer term moving averages tend to live shorter than indicated by theoretical models. Moreover, all of the heat maps follow the same pattern. To illustrate this Figure 1-23 shows the median Log Rank test results for positive as well as negative signals for all three simulations methods across all base currencies. While the results for the third simulation are less strong than the results for the first two simulations all three tables clearly indicate that SR (1, 2, 3, 4, 5) and LR (5, 10, 15) show fair levels of positive (or at least less negative) deviations from the benchmark simulation. Other, longer term moving average combinations indicate levels of negative deviations from the benchmark simulations.

The results of this analysis are somewhat mixed, the first two simulation methods provide some statistical evidence, while the third simulation does not. Nonetheless, one could argue that it is thorough to go through a series of benchmark assumptions against which momentum signals are compared. Yet incorporating autocorrelation in the benchmark assumption, as it is the case for the third simulation is very restrictive. This is due to the fact that the entire idea of momentum is based on the concept of autocorrelation. Despite this, the results of the third simulation follow the results of the first two simulations in terms of the directionality, pointing towards longer survival rates for shorter term moving averages and shorter survival rates for longer term moving averages.

3. Results for Sub-Sample Periods One to Nine

Given the fact that the data sample used in this chapter is very long, the analysis undertaken in the previous sections does incorporate a variety exchange rate regimes in the countries analysed. Figure 1-1 shows capital controls, as well as the different exchange rate regimes prevalent in different countries over the sample period. Over the observation time period exchange rates, such as the USD, the CAD, the CHF or the

EUR (DEM in its previous form) were not restricted by capital controls or pegs. With regards to the EUR it has to be noted that while the DEM was a part of the currency snake in the early years and the EMS (European Monetary System) thereafter, it was the driving factor of both of these institutions. Hence, the currency can be seen as unconstrained. Other exchange rates such as the GBP and the JPY are free floating over the sample period, however with capital controls in place. Then finally some of the smaller exchange rates such as the AUD, NZD, NOK, SEK have capital controls as well as pegs in place up until the mid 1980's for the AUD and the NZD and the mid 1990's for the NOK and the SEK. This begs the question to which extent some of the anomalies observed in the previous section are driven by the fact that many of the markets were closed and tightly controlled by central banks. While, Neely (2002) provides extensive evidence that in many cases trading rule returns precede the actual intervention by central banks. Hence exchange rate intervention comes as a reaction to strong and very profitable short-term trends within currency markets. In light of the fact that some foreign exchange markets were dominated by central bank activities over parts of the sample period, this section applies the same log rank analysis used in the previous section to sub-samples of the dataset. As indicated earlier the full dataset is split into nine sub-samples, whereby the first eight sub-samples consist of exactly 1000 observations and the ninth sub-sample consists of 1025 observations. Hence, the annual split of sub samples look as follows, SS1 (1975-1978), SS2 (1978-1982), SS3 (1982-1986), SS4 (1986-1990), SS5 (1990-1994), SS6 (1994-1998), SS7 (1998-2002), SS8 (2002-2006) and SS9 (2006-2009). This split of sub-samples has to be done somewhat independently of the different regimes as shown in Figure 1-1. The rationale for this is the fact that the statistical power of the survivorship analysis depends heavily on the number of observations in the sample. Hence, by almost equally splitting the dataset, each of the sub-samples will show similar levels of statistical confidence.


As indicated earlier the moving average combinations SR (1, 2, 3, 4, 5) and LR (5, 10,15) tend to show less negative deviations from the benchmark simulation. The question now arises whether this is persistent over time. Figure 1-24 aims to answer this question, the results in the 1- are calculated on the basis of the regular resampling simulation. The dark bars in the figure show the median log-rank test for positive momentum signals generated across all moving average pairs. Moreover, the light bars shows the difference between the median log-rank test result of all moving average pairs and the median log-rank test result of the short-term moving average pairs SR (1, 2, 3, 4, 5) and LR (5, 10,15). The results in Figure 1-24 suggest that the median log-rank

test result across all moving average pairs has increased from -0.5 in the first sub-sample to 0.2 in sub-sample nine. This indicates that over the observation period the empirical survival time has increased. This also indicates that in the first sub-sample the empirical survival time is shorter than suggested by the benchmark simulations, and in the last sub-sample it has become marginally longer. Moreover, the difference between short-term and long-term moving average results decrease over time. Figure 1-25 shows the same results for negative moving average combinations. None of the results in the graph are statistically significant. However, it has to be borne in mind that Figures 1-24 and 1-25 are based on median values across all currency pairs. Appendix 6 shows the same analysis for single base currencies and the result of the same analysis as shown in this section, however based on the GARCH (1,1) simulation.

The main outcome from this sub-sample analysis is the fact that the deviation from normality has diminished over time, to a point where there is no obvious non-normality across all moving average pairs. However, this move to normality has not been steady. Looking at Figures 1-24 and 1-25 it becomes evident that both the log rank median across all moving averages in dark grey, as well as the difference in medians between long and short term moving averages deteriorates steadily over the first three sub sample periods, covering the time from 1975 to 1986. By this time only the NOK and the SEK were pegged and under capital controls. From 1986 to 2002, both median values and differences in median are more pronounced again. This subsides then in the time period from 2002 to 2006. The results from the last sub-sample 2006 to 2009 suggest a small pickup in non-normality. This is the case across the board for all moving average combinations, with no distinction between long-term and short-term moving average combinations.


While one could argue, that some of the results in the previous section could be influenced by capital controls and central bank activity. Namely, positive deviations from normality for the NOK the SEK, the AUD and the NZD might be influenced by the restrictive monetary regimes in the early part of the sample period. Nonetheless, looking at Appendix 6 one has to bear in mind that for some of the exchange rates these deviations continue to persist in sub sample periods that are not subject to capital controls and exchange rate pegs. Moreover, when looking at the USD and the JPY, neither of the exchange rates have been pegged over the sample period, and capital controls for the JPY are abolished very early in 1980.

FIGURE 1-24: SUB-SAMPLE ANALYSIS; MEDIAN LOG RANK VALUES; POSITIVE SIGNALS



The dark bars show the median level of log-rank test results across all moving average signal combinations. The light bars shows the difference between the median log-rank test result of all moving average pairs and the median log-rank test result of the short-term moving average pairs SR (1, 2, 3, 4, 5) and LR (5, 10, 15).

FIGURE 1-25: SUB-SAMPLE ANALYSIS; MEDIAN LOG RANK VALUES; NEGATIVE SIGNALS




The dark bars show the median level of log-rank test results across all moving average signal combinations. The light bars shows the difference between the median log-rank test result of all moving average pairs and the median log-rank test result of the short-term moving average pairs SR (1, 2, 3, 4, 5) and LR (5, 10, 15).

Yet the results for both of the exchange rates suggest a strong negative deviation from normality over the sample period, which spans evenly with a deteriorating trajectory over many sub samples as indicated in Appendix 6. From these results the question arises whether the results of the survivorship analysis have some resemblance with results generated from actual trading rules. To facilitate this, the next section conducts a trading rule analysis.

G. Trading Rule Implementation

Given the fact that there is only an implicit relationship between the length of survival of a technical trading signal, and the return that such trading signal derives, this section aims to build a bridge between both. While it is crucial to design realistic trading rule setup, is also important that this setup is parsimonious, and easily understandable. The first step in making this trading rule implementation parsimonious is to combine all the moving average trading signals that are generated for each currency pair into one composite trading signal. The reason for this is the fact that generating trades based on single trading signals might incur high trading costs, which most likely make any trading strategy unprofitable.

FIGURE 1-26: EXPOSURE LEVEL OF THE COMPOSITE/BENCHMARK TRADING RULE FOR USDGBP



The dark line shows the exposure level of the composite (benchmark) trading rule of the USDGBP exchange rate over the first eight month of 2007.


Furthermore, a combined trading rule also incorporates the interaction between moving average combinations. For instance, while longer-term moving average combinations might still point towards a long position, shorter-term moving averages might indicate an increasing short bias. The composite moving average trading rule is constructed by taking all moving average combinations of SR (1, 2, 3, 4, 5, 10, 15, 20, 25) and LR (5, 10, 15, 20, 25, 30) and by summing up the number of positive trading signals, then deducting the number of negative trading signals at every point in time. Given that the chapter tests 39 moving average rules, a positive number of +39 would indicate maximum positive momentum, because it means that all moving average rules generate positive trading signals. A number of -39 would signal maximum negative momentum. The composite trading rule is then standardised between a range of +1 and -1. This is obtained by dividing the raw signal by the total number of moving averages. $+39/39 = +1$. Figure 1-26 shows the evolution of the exposure level of the composite trading strategy for the USDGBP cross over the first eight months of 2007. The trading strategy starts the year with a slightly negative stance, spikes to a positive level in early January and falls off to a negative stance thereafter. It switches to a positive exposure level in February, and falls off afterwards. Moreover the strategy oscillates between positive and negative exposure in the following months. During early January, late March, early June and late August the trading strategy reaches its most negative exposure level, while being most positively exposed to the USDGBP cross in late April and early May, as well as July and August. Given the fact that this trading strategy incorporates all trading rules tested, this trading strategy is labelled “benchmark” trading strategy.

The second trading strategy that is tested consists of a subset of the original universe of trading rules applied. As indicated in an earlier section the moving average combinations SR (1, 2, 3, 4, 5) and LR (5, 10, 15) tend to live longer or less short than what theory would suggest. To assess whether the information provided by the survival analysis does add any value in the context of trading rule profitability this subset of trading rules is used as the “enhanced” trading rule. It is constructed in a similar way as the “benchmark” trading rule.

Figure 1-27 shows the evolution of the exposure level of both strategies. The benchmark strategy and the enhanced strategy follow the same path when it comes to the directionality of the exposure level. Where they differ however is in the pace at which different exposure levels are reached, as well as the amount of time in which the strategy is fully long or short exposed. Considering the time period of Figures 1-26 and 1-27, the benchmark strategy has a 100% long and 100% short exposure 12.5% and

3.6% of the time respectively. The enhanced trading rule has a 100% long exposure 28% of the time and a 100% short exposure 18.5% of the time over the given period.

FIGURE 1-27: EXPOSURE LEVEL OF THE ENHANCED AND THE BENCHMARK TRADING RULE FOR USDGBP



The dark line shows the exposure level of the benchmark trading rule of the USDGBP exchange rate over the first eight month of 2007. The dotted line shows the exposure level of the enhanced trading rule over the same time period.

As mentioned earlier the aim of this section is to build a link between the results of the survival analysis and trading rule returns. To do so, this section translates the results from the survivorship analysis into a trading rule implementation that is realistic. The survivorship analysis presented in previous sections is based on interest rate adjusted currency spot returns. The reason for the interest rate adjustment is twofold. Firstly, it corrects for biases stemming from the interest rate differential between different currencies. However, the more important reason for the adjustment is the fact that interest rate adjusted currency returns are the actual returns that one obtains when trading in the currency market. Given the fact that the data used in the study are interest rate adjusted currency spot returns, which cannot be traded, one has to make an assumption about the design of an implementable trading strategy in order to assess the fair cost of implementation.


This can be done two ways, either one assumes that the trading strategies are implemented with currency forwards, or with currency futures. Both strategies are

almost identical, the only difference between both strategies are the fact that a futures based strategy requires daily margining, whilst not being subject to counterparty risk, whereas a forward based strategy does not require daily margining, yet it is subject to counterparty credit risk. For a performance perspective there might be slight differences between the two strategies depending on the correlation between the return on the margin cash and the underlying currency return. However, the aim of this trading rule implementation is to understand the number, as well as the nature of transactions that have to be entered into, in order to obtain the desired exposure level, an example of which is shown in Figures 1-26 and 1-27. Hence, nuances between trading futures or forward contracts can safely be ignored. For the purpose of illustrating the amount of transactions that go into the trading rule calculation, this section assumes that the strategy is implemented on the basis of futures contracts.

When designing a trading strategy it is always key to minimise the transaction cost involved in such strategy. Again this analysis applies to futures or forward contracts alike. The only difference here is the fact that forward contracts have more flexibility with regards to expiry dates than futures, this is due to the fact that futures contracts are more standardised. The strategy assumes that at the beginning of each quarter a 3-month contract with a fixed (end of quarter) expiry date is purchased. Any other futures (or forward) contract that is bought or sold within the quarter, as a result of changing trading signals, has the same expiry date. This means all of the other trades during the quarter net off, apart from a residual position, which reflects the amount of the trading signal on the last day of the quarter. Hence, only the open position at the end of the quarter has to be closed and rolled into a new three-month contract. Due to the fact that the currency space is very liquid, and no market disruptions are expected as a consequence of the futures roll. This chapter does not stagger the rolling of the contract over different maturities. This strategy minimises the amount of transactions, however, it might be not the cheapest strategy. This would depend on the liquidity of contracts. Some contracts are less liquid in longer maturities, hence more expensive. Moreover, the term structure of the futures curves might be such way that a strategy that rolls monthly or bi-monthly is more cost effective, despite the higher number of transactions entered into. This chapter does not have information about the term structure of futures (or forward) contracts. Hence it applies an interest rate adjustment to spot returns. This is done via 3 month t-bill rates. The reason for choosing a three month rolling cycle is due to the fact that most futures (or forward) contracts have the best liquidity point in the quarterly contracts. It is appreciated that a mere interest adjustment of spot returns

does not fully reflect all of the dynamics of the currency futures (or forward) curves, nonetheless the interest rate adjustment will correct for the biggest part of that. Moreover, the fact that the trading strategy is based on fairly fast moving trading signals that oscillate symmetrically between long and short exposure, one could argue that any futures curve movements in excess of the interest rate differential are likely to cancel out.

FIGURE 1-28: USDGBP BENCHMARK TRADING RULE, ASSUMED TRANSACTIONS



Schematic overview of transactions undertaken of the USDGBP benchmark trading rule from 15th of March 2007 to the 16th April 2007. On the 15/03/2007 the strategy is short 795 in notional, from that day to the next the trading signal signals a purchase of 256 notional. Therefore the strategy is short 538 notional on the 16/03/2007.

The dataset used in this chapter consist of daily mid-prices calculated on the close of the New York trading session. Given the fact that the currency market is a 24 hour

market¹², it is possible to initiate a trade. (or change in exposure), at the point when the signal has been obtained. Foreign exchange markets are the deepest and most liquid financial markets. Therefore the aspect of slippage, i.e. the price movement caused by undertaking a transaction, is negligible.

The implementation costs of the trading strategy are calculated on an adjustment transaction basis, with a quarterly roll adjustment as and when the assumed futures contracts expire. Figure 1-28 shows the actual exposure levels of the USDGBP trading strategy from the 15th of March 2007 to the 16th April 2007. For illustration purposes, to make the descriptions in Figure 1-28 more approachable, it is assumed that the notional of the trading strategy varies between 1000 and -1000. Figure 1-28 indicates that on the 15th of March 2007 the strategy is short 795 notional on the 16th of March the exposure level has to be adjusted to a short level of 539. Hence 256 worth notional has to be purchased on that day. On these 256 worth notional transaction cost are applied. The same is done for any other change in exposure during the time period when the three months futures contract is alive. Hence on the 19th of March 51 of notional has to be purchased, transaction costs are calculated on that purchase. On the 20th of March 308 worth of notional has to be purchased on the basis of which transaction cost have to be calculated, etc. At quarter ends the assumed futures positions have to be rolled into new quarterly contracts.

The notional of the quarter end roll in Figure 1-28 is a combination of the total exposure level on the last day of the quarter, which is 692, plus the adjustment that comes from the trading rule signal, which is -103. 103 is hereby the difference between notional exposure signalled by the trading strategy on the 30/03/2007 and the 02/04/2007. Hence the overall amount of the roll at quarter end is 589. When it comes to rolling existing positions these transaction cost are much lower than transaction cost that come as a result of changes in exposure levels. In practice the roll cost are mainly a function of the position of futures (forward) curves, which, adjusted for the interest rate differential, are assumed to cancel out, given the long short nature of the strategy. Hence at quarter end the part of the transaction that is attributed to the roll as opposed to the change in exposure level is assumed to be immaterial.

Given that the dataset of this study does not have any bid/ask spreads, all transactions are implemented on the interest rate adjusted spot price, mimicking futures (forward)

¹² The Australian currency market opens at the time of the New York close.

mid prices. In order to assess whether different trading strategies are profitable, this chapter will present the results of the trading rule implementation in form of breakeven transaction cost levels. The idea hereby is to compare actual transaction cost levels, consisting of bid/ask spreads, commissions, etc. with the breakeven transaction cost level. If the actual transaction cost level is lower than the breakeven transaction cost level, then the strategy is profitable. Breakeven transaction cost levels are defined as the level of per trade transaction cost that needs to be incurred, in order for the trading rule to yield a risk adjusted return of zero. This is done on an iterative basis, whereby the transaction costs are incrementally increased up until the point where the trading strategy becomes unprofitable. In that context the number of transactions, as well as the changes in exposure level are defined, as shown in Figure 1-28.


Given the fact that different trading rules consist of different currency pairs, or different parameterisations, and have different levels of exposure at different times, it is reasonable to assume that each of the moving average trading strategies that are analysed is likely to have different levels of risk. This suggests that calculating breakeven transaction cost levels by merely setting the trading rule return to zero is not sufficient to make a fair comparison between trading strategies. In order to make the breakeven transaction cost levels comparable across moving average combinations as well as base currencies, the calculation of the breakeven transaction cost levels focuses on the Sharpe Ratio, whereby the interest rate is assumed to be zero given the fact that the trading strategy aims to mimic a futures based trading strategy. Therefore, the Sharpe Ratio is defined as the unit of return received per unit of risk taken. This ensures that each of the breakeven transaction cost levels across trading strategies is obtained with exactly the same level of risk. The calculation process of the breakeven transaction cost level is conducted such way that assumed transaction cost are increased incrementally until the sharp ratio of each trading strategy reaches a level of 0.01¹³. As mentioned earlier the breakeven transaction cost analysis is based on an assumed futures based trading strategy, whereby changes in exposure levels and assumed transactions are calculated as shown in Figure 1-28. Moreover, it has to be pointed out that the analysis incorporates the element of turnover. Hence, if the composite trading signal indicates to turn over the position by 60% then the transaction cost are only applied to 60% of the portfolio. This ensures the comparability of different composite trading

¹³ The reason for using 0.01 as opposed to 0 is technical

rules. Therefore, a trading rule that reacts faster and has a higher turnover will have to deliver a higher return than a slower moving average trading rule.


Figure 1-29 shows breakeven transaction cost levels for the benchmark and the enhanced trading rules across all base currencies. For the full dataset, the median breakeven transaction cost level for the benchmark strategy is 7.8 basis points, while the median breakeven transaction cost level for the enhanced strategy is at 14.5 basis points. This amounts to a performance difference of 6.7 basis points per trade. Figure 1-30 shows the evolution of the median breakeven transaction cost levels over time. The results suggest that while the benchmark strategy sees an erosion of profitability over time (with the exception of the most recent sub-samples), the enhanced trading strategy maintains its levels of profitability. The range of breakeven transaction cost levels of the benchmark strategy is 15.4 basis points to 1.5 basis points, while the range of enhanced trading strategy breakeven transactions cost level spans from 18.3 to 11.2 basis points. The results for single base currencies are given in Appendix 6

FIGURE 1-29: BVTC; ENHANCED VS. BENCHMARK TRADING RULES (ACROSS CURRENCIES)



The Figure shows breakeven transaction cost levels across base currencies. The dark bars show the breakeven transaction cost levels for the benchmark long short trading strategy, which consists of all moving average signal combinations. The light bars show the breakeven transaction cost levels of the enhanced long short trading strategy, which consists of SR (1, 2, 3, 4, 5) and LR (5, 10, 15) moving average pairs. The dotted line shows the difference between both.

FIGURE 1-30: BVTC; ENHANCED VS. BENCHMARK TRADING RULES (ACROSS SUB-SAMPLES)




The Figure shows breakeven transaction cost levels across sub samples. The dark bars show the breakeven transaction cost levels for the benchmark long short trading strategy, which consists of all moving average signal combinations. The light bars show the breakeven transaction cost levels of the enhanced long short trading strategy, which consists of SR (1, 2, 3, 4, 5) and LR (5, 10,15) moving average pairs. The dotted line shows the difference between both.

H. Link between Survival Analysis and Trading Rule Results.

Given the fact that that survivorship analysis looks at the length of a trading signal, while trading rule analysis looks at the magnitude of return derived from trading signals, the link between the survivorship analysis and trading rule returns is indirect. The log-rank test results and the trading rule results, make it evident that it does exist. As indicate in an earlier section the basic idea behind creating survivorship curves is to model the probability of the persistence of some pre-specified signal within a given data sample. In the context of this chapter a dual crossover moving average trading rule generates the investigated signal. Any positive signal is generated if the short-term moving average is above or equal to the long-term moving average. If the short-term moving average moves below the long-term moving average the previous trading signal ceases to exist and a negative trading signal is initiated. If for example the SR1/LR10 trading rule for the USDGBP cross generates nine positive trading signals, one can deduct two things either that the magnitude of returns in few days when the trading rule is established are so high that the signal still persists even when the underlying returns

have turned negative, or that the returns remain marginally positive over the time period where the signal remains positive, or variations within the two extremes. What is known however is the fact that up until the point where a positive signal switches to neutral or negative, the return generation is on the margin positive. Moreover, the stop loss mechanism, in form of the long term moving average is ratcheted up as new (positive) returns enter the calculation. Therefore, the longer a signal lives the more opportunity exists for ratcheting up the stop loss level. Hence the total return derived from the strategy should be higher, the opposite may be the case for signals that live shorter than expected. Yet the increased adaptability of the model due to shorter survival might in turn add value. Hence, the higher the absolute level of deviation from normality in terms of survival of trading rules, the higher the opportunity for trading rules to add value.

FIGURE 1-31: ABS. LOG-RANK TEST RESULTS VS. BENCHMARK TRADING RULE RESULTS




The dark grey line shows the breakeven transaction cost levels of the benchmark trading strategy (the line refers to the right hand scale). The light grey line shows the absolute level of positive and negative log-rank test results (the line refers to the left hand scale).

As shown in earlier sections the survivorship analysis suggests that shorter-term moving averages outlive the benchmark simulations, while longer-term moving average combinations tend to have a shorter life expectancy than what is suggested by the benchmark model. A similar pattern materialises when looking at the returns of the enhanced versus the benchmark trading rule, the enhanced trading rule shows

considerably higher returns than the benchmark trading strategy. Moreover, Figure 1-31 shows that the same pattern also materialises over time. The dark grey line in Figure 1-31 shows the median breakeven transaction cost levels for all currency pairs over time. The light grey line shows the absolute value of the average median log-rank test result across positive and negative momentum signals. The reason for taking the absolute value of the average across median log-rank test values for positive and negative momentum signals is the fact that the trading strategy is a long-short strategy.

FIGURE 1-32: ABS. LOG-RANK TEST RESULTS VS. ENHANCED TRADING RULE RESULTS



The dark grey line shows the breakeven transaction cost levels of the enhanced trading strategy. (the line refers to the right hand scale). The light grey line shows the absolute level of positive and negative log-rank test results (the line refers to the left hand scale).

Hence, the trading rule switches continuously between positive and negative exposure depending on the trading signal. (For example, it is likely to be in long exposure if the underlying currency exhibits a positive trend, therefore participating in that positive trend. However, it also switches into short exposure when the trend turns and therefore also positively participating in a negative trend.) Hence, it should capture deviation from market normality either way regardless of whether it is a positive or negative deviation. Figure 1-31 clearly shows that there is a link between survivorship analysis

and the results of the benchmark trading strategy. The absolute levels of log-rank test results suggest that the deviation from normality has diminished over time with a small pickup in the most recent period. The same is true for the breakeven transaction cost levels. They have diminished over time and have seen some pickup in the most recent period. The correlation between the two lines in Figure 1-31 is 0.786. The results for single base currencies are given in Appendix 6 and are widely in line with what Figure 1-31 suggests.

Figure 1-32 shows the results of the same analysis comparing the results of the enhanced trading strategy with the absolute log-rank test results of the respective moving average combinations. The results clearly indicate that no such relationship, as shown in Figure 1-31, exists for shorter-term moving averages. The correlation between the two lines is -0.106. The results for different base currencies are given in Appendix 6. This suggests that while there is a link between the broad results of the survivorship analysis and the variation of the overall trading rule profitability over time, no such link can be established for shorter-term moving average trading rules. This allows for two observations. Firstly, some parts of the trading rule profitability are likely to be driven by deviations from market normality. The sub-sample analysis, however, suggests that these deviations diminish over time. As these deviations diminish, overall trading rule profitability diminishes as well. This suggests that the diminishing part of the trading rule profits can be attributed to diminishing market inefficiency. Secondly, there is a set of trading rules that maintains its level of profitability despite the fact that the survivorship analysis points towards diminishing deviation from market efficiency. This implies that some of the trading rules might have return drivers, other than market inefficiency.

I. Conclusions

The first chapter of this thesis introduces an alternative approach to detecting market inefficiency. The methodology is based on lifetime statistics, in particular survivorship analysis. The intuition behind the presented methodology finds its roots in the concept of Runs test. Runs tests aim to compare the probability of occurrence of positive or negative return streams within an empirical time series, with the probability of occurrence of similar return sequences theoretically derived, assuming independence between returns. The survivorship analysis follows partly the same intuition; it models

the survival probability of empirical return streams. It also allows the comparison between empirical survival probabilities and theoretically generated benchmark survival probabilities. However, it differs from the runs test when it comes to the assumption about the properties of the benchmark process. The runs test only allows for a benchmark specification that follows a Bernoulli type process, hence mere independence between returns. The survivorship analysis on the other hand, has the flexibility of using different benchmark processes. Furthermore, a runs test can only ever be applied to a mere return stream. When it comes to assessing signals that are generated from trading rules, the runs test specification breaks down. Survivorship analysis on the other hand, is able to facilitate the analysis of more complicated trading rules. Finally, when assessing market inefficiency by implementing runs test, the aspect of hypothesis testing becomes particularly problematic. This is not the case for survivorship analysis. The log-rank test, introduced earlier in this chapter, represents a reliable tool, to assess the statistical significance of results.

Further to these methodological advances, this chapter gives evidence of inefficiencies within the currency market. The results from the survivorship analysis suggest that various currencies have empirical exchange rate patterns that cannot be explained by any benchmark process. These anomalies either suggest that empirical momentum signals outlive benchmark signals, as is the case for moving average crossover signals that utilise a set of very short-term moving average combinations, or they suggest that momentum signals created from empirical curves, as is the case for some longer-term moving average crossover signals, have a lower life expectancy than theory would suggest. The results from a sub-sample analysis, however, indicates that most of the deviations from market efficiency deteriorate over time, up until the point where all of the momentum signals exhibit survival times that are statistically equivalent to what is suggested by benchmark processes.

The results obtained from implementing generic trading rules on the same set of moving average crossover signals as tested in the survivorship analysis, reinforce the validity of the survivorship methodology as a tool to detect market inefficiencies. The profitability of a generic trading rule that incorporates all moving average signals deteriorates continuously (as suggested by the survivorship analysis) to a point where the trading rule becomes unprofitable. While, a trading strategy that is constructed from a subset of moving average signals shows clear outperformance over a trading strategy that is generically composed from all moving average crossover signals. This outperformance

persists over time. Moreover, the overall profitability of such a short-term focused strategy remains within a reasonably high range over time. This result counters the results suggested by the survivorship analysis. It also suggests that the source of these returns might well be something other than market inefficiency.

V. Chapter 2:

Momentum Effects:

G10 Currency Return Survivals: Implications for Trading Rules

Abstract

The chapter models survival probabilities of positive and negative momentum signals that are obtained from a wide set of dual crossover moving average combinations for all G10 cross currency pairs. The results of this survival analysis are used to create trading rule enhancements that aim to outperform generic dual crossover moving average trading signals. The trading rule enhancements are assessed, by applying White's (1999) "data snooper". The results suggest that there is scope for trading rule enhancements to outperform generic trading rules. Moreover, results present strong evidence for Lo's (2004) Adaptive Market Hypothesis.

A. Outline

1. Academic Background

a) Trading Rules in Equity Markets

The urge to simplify complex structures and the desire to create rules of thumb that support a decision making process is deeply embedded in human nature. Therefore it is not unreasonable to assume that technical trading rules have existed since the existence of organised exchanges. However, the first documented evidence of trading rules comes from Charles Henry Dow, the founding editor of the Wall street journal. In his early work that is aimed to explain general business conditions, Charles Henry Dow lays the intellectual foundations of modern technical analysis. William Peter Hamilton (1922) and Robert Rhea (1932) later formalise Dow's theories. Brown, Goetzmann and Kumar (1998), who re-test Hamilton's formalisations, indicate that the algorithms proposed by Hamilton do have some degree of predictive power. Moreover, they conclude that Dow's theory is a momentum theory.

In the pre-market efficiency era technical trading rules were attractive instruments used by practitioners. Unfortunately there is comparatively little academic evidence that underpins the case for technical trading rules in that time. With the rise of the Efficient Market Hypothesis, the attitude of academics towards technical analysis deteriorated significantly. Malkiel, in his book "Random Walk Down Wall Street" (1981), makes the point that the methodology was widely regarded to be a flawed concept. Moreover, technical analysis was an easy subject to pick on. As a result, it must have been very difficult for researchers to get papers published that look at this field. An example of an early academic study that investigates technical trading rules is Levy (1967), who is able to generate superior trading rule returns on the basis of signals generated from strong historic price movements. Nonetheless, in his conclusion he makes the point that some of the returns must clearly come from extraordinary risk taking, which the study doesn't control for. Hence, Levy argues that despite the results being strong there is not sufficient evidence to refute the Random Walk hypothesis. Sweeney (1988) indicates that during the late seventies academic research focussed on market anomalies, and these anomalies are always put into context of the efficient market framework. Very regularly the underlying model assumptions are questioned and dismissed, in favour of the market efficiency framework. Bearing in mind the historic background, another

landmark paper in the context of early trading rules is Fama and Blume (1966). The study investigates a set of filter rules that are applied to daily returns of the constituents of the Dow Jones Industrials Index. Their sample spans depending on the single stocks from early 1956 or 1958 to late 1962 and it covers 1200 to 1700 daily return observations. The filter rules applied in the study are designed in such way that a long or short position is established if the up or down move in the price of the underlying stock exceeds a certain threshold. The total number of filters is 24 and the filter thresholds are ranging from 0.5 per cent to 50 per cent. The study covers all 30 Dow Jones Industrial stocks. Fama and Blume (1966) indicate that there are slight amounts of positive as well as negative dependences in the price changes. For three filter sizes (0.5%, 1.0%, 1.5%) the average returns per stock on long positions are greater than the average returns from a buy and hold strategy. A similar dynamic applies to the short side. The returns on both long and short positions fall dramatically as the filter size increases. The study also suggests that filters below 12% and above 25% produce negative average returns over the observation period, when adjusted for brokerage fees. Filters within that range produce positive returns. However, they are small compared to the buy and hold strategy. Therefore, after accounting for transaction costs the trading rule profitability is significantly diminished. Hence, Fama and Blume (1966) conclude that none of the investigated trading strategies provides any economic benefit, as transaction costs erode the profitability of the filter rules.

Fifteen years later, when the contra movement to Fama's Efficient Market Hypothesis started taking off with DeBondt and Thaler's (1985, 1987) studies of return reversals and Jegadeesh and Titman's (1993) groundbreaking momentum paper. Technical trading rule studies increasingly gained academic acceptance. Papers such as Sweeney (1988), which directly counters Fama and Blume's (1966) paper, received wide academic attention. Sweeney's (1988) main criticism of Fama and Blume's (1966) paper is that the authors do a mediocre job of finding potentially winning securities. Furthermore, their tests don't have sufficiently developed statistical confidence bounds for judging significance levels. Sweeney (1988) re-examines the results of the earlier paper of Fama and Blume (1966) and develops a statistical framework by which filter rule results can be tested for their significance. Sweeney's paper selects a subset of the Fama and Blume stocks, which exhibited the most promising results, hence the "winners" of their sample. These stocks are then investigated in the time period from 1970 to 1982. The dataset consists of daily closing prices. The paper focuses on long

only equity positions, the argument hereby is the fact that shorts don't perform well and only add layers of unnecessary transaction cost. In order to make trading rule results comparable with the buy and hold strategy, risk adjusted returns are used. The key finding of Sweeney's (1988) study is the fact that the winner stocks in Fama and Blumes (1966) paper seem to be winners in the following decade as well. Sweeney (1988) suggests that previous winners, which have been identified by a 0.5% filter rule, yield returns that are statistically significantly better than buy and hold returns. This even holds when floor trader transaction costs of 5 basis points are applied. The results clearly weaken if higher transaction costs are applied. Nonetheless, Sweeney (1988) suggests that over this time period, significant trading profits can be made, but only for investors that have a fairly low level of transaction cost such as floor traders, who are not dependent on intermediaries. For a cost structure of institutional investors, the trading rules are borderline profitable. Single retail investors, however, will not be able to benefit from systematically trading on the basis of these filter rules. He cautions the results as they depend on the assumption that the daily closing price is an unbiased estimate of the actual price at which transactions are undertaken.

Another landmark study within the early trading rule research is Brock, Lakonishok and LeBaron (1992). Their paper tests only two trading rules. The first rule is a dual moving average crossover signal. Hence, it gives a buy or sell signal if the short moving average is above or below the long moving average. They also apply a band around the moving averages. If the moving averages cross within that band, no signal is generated. This is a secondary signal that is aimed to smoothen the trading signal generation and mitigate the risk of trading signals getting whiplashed. With respect to this signal, the authors differentiate two implementations. The first implementation has a variable length in holding period after a signal has been established. Hence, as long as the short-term moving average is above the long-term moving average (plus band) the strategy will be long. Once the short-term moving average crosses the long-term moving average, the long position will be liquidated. The second implementation is fixed length, hence once the signal is established, the trading strategy will go long or short for an arbitrary period of 10 days. The second trading strategy is a break out strategy. That means that if the price exceeds the maximum price over a defined time period or falls below the minimum price over that same time period the trend is broken and a long or short signal is initiated. Brock, Lakonishok and LeBaron (1992) define three time ranges of 50, 150 and 200 days as appropriate for the trading signal generation. What makes this paper

somewhat different than the other papers that have been published before is the assessment of the trading rule profitability. Previous papers were merely comparing the returns from the filter rules with buy and hold returns, and conducting statistical significance tests on that basis. Brock, Lakonishok and LeBaron (1992) extend this standard statistical analysis via the use of simulation techniques, such as bootstrap analysis. They impose a set of four popular benchmark models: the random walk, the AR(1), the GARCH-M, and the Exponential GARCH. They use only one data series, the Dow Jones Industrial Average. However, their time window of daily closing values is very long, starting from 1897 to 1986. Conditional returns that have been generated by filter signals from the actual Dow Jones data exhibit consistently better returns than the unconditioned time series. In the case of the fixed length moving average signal, the difference between buy and sell signal is 0.08%, which is considerable when compared with a normal ten day upward drift of 0.17%. It is important to note that this difference cannot be explained by any difference in riskiness of strategy. Furthermore, under the variable length holding period, where the investor is continuously in the market, buy signals are followed by an average market increase of 12% per annum and sell signals are followed by a market decrease of 7% per annum. This is in sharp contrast to any research conducted up until that point.

The main criticism of Brock, Lakonishok and LeBaron (1992) is the fact that they use only a very narrow universe of trading rules. This is not only the case with their study, but all early trading rule papers investigate only a very small number of trading filters. The idea behind this is the fact that they wanted to avoid the look back bias. This occurs when a dataset is used not only for data inference, but also model selection. This effect is also called data snooping and might invalidate all trading rule result. Brock, Lakonishok and LeBaron (1992) acknowledge this problem, hence they introduce a resampling simulation that is based on various benchmark processes. This enhances the statistical validity of their results, but what it does not allow is to calculate a comprehensive test across moving averages. That means Brock, Lakonishok and LeBaron (1992), cannot assess whether certain rules that are chosen might merely work well for the specific dataset, or whether they truly represent a superior trading rule under any given environment. Brock, Lakonishok and LeBaron (1992) and other early researchers avoid this problem by using a very generic set of trading rules, which is accepted and widely used by the investment community, thus mitigating the data snooping risk. Nonetheless the aspect of data snooping remains a limiting factor to the

validity of early trading rule papers.

Sullivan, Timmerman and White (1999) address the problem of data snooping by applying a different methodology; they follow the methodology of White (2000) closely, which allows them to compare a specific trading rule with a “benchmark” that consists of a large set of trading rules. This new approach enables them to quantify any potential data snooping bias and fully adjust for its effects. In their test setup Sullivan, Timmerman and White (1999) follow Brock, Lakonishok, and LeBaron’s (1992) approach of using the very long return time series that is offered by the Dow Jones Industrial Average. However, they expand the very narrow set of 26 different trading rules used by Brock, Lakonishok, and LeBaron’s (1992) to a wider universe of 7846 filter rules.

The intuition behind White’s (2000) “Data Snooper” is the idea of evaluating the distribution of the performance of the optimal trading rule, considering the full set of models that led to the best performing trading rule. Hence, every single trading rule is iteratively tested against all other trading rules. Sullivan, Timmerman and White (1999) argue: “If the marginal trading rule does not lead to an improvement over the previously best performing trading rule, then the P-value for the null hypothesis that the best model does not outperform will increase, effectively accounting for the fact that the best trading rule has been selected from a larger set of rules. On the other hand, if the additional trading rule improves on the maximum performance statistic, then this can reduce the P-value since better performance increases the probability that the optimal model genuinely contains valuable economic information.” Their findings suggest that the 26 common trading rules that have been analysed by Brock, Lakonishok, and LeBaron’s (1992), stand up to the data snooping test. They validate this finding also on a sub-sample basis. Furthermore, they find that other data snooping proven trading rules perform even better. However, these results have to be considered with caution, given the fact that the data snooping proven trading rules that have supposedly performed best over the 100 year sample period completely break down when tested on S&P 500 futures returns on an out of sample basis, over the time period from 1987 to 1996. Further to this breakdown in performance of the top trading rules, the authors have also not been able to find any evidence of trading rule outperformance over the sample period. They put two arguments forward to explain this obvious breakdown. One might be the erratic return behaviour during the 1987 crash. They put caution around this argument, given that some trading strategies would have definitely benefitted from a

crash scenario. The second argument has to do with market efficiency where, as equity markets have become more efficient, the trading rule performance is diminished.

Other studies such as Lo, Mamaysky and Wang (2000) aim to automate investment decisions based on more complicated technical indicators, such as head and shoulder pattern. In their view the reason why technical analysis has widely been dismissed by academia is the fact that it requires a great deal of human judgement when interpreting historic price data. They undertake the automation of such judgement by the application of non-parametric kernel regressions, which are applied to a large number of U.S. stocks over a time period from 1962 to 1996. The idea behind their methodology is to predefine technical patterns in terms of their geometric properties, then construct a kernel estimator that allows identifying the local minima and maxima that have been set in the first step. The evaluation of the effectiveness of this methodology is then done by comparing the unconditional empirical distribution of daily stock returns to distributions that are conditional on the occurrence of technical pattern. The idea here is to evaluate whether technical pattern are informative or not. If they can be explained by the statistical model, both distributions should be closely related. Their results suggest that there is incremental value that can be created by the application of automated algorithms.

b) Trading Rules in Foreign Exchange Markets

While the early studies in technical analysis have focussed on equity markets, academic supporters in that time have persistently argued that technical analysis is universal and can be applied in one form or another to all asset markets. Hence, it shouldn't come as a surprise that traders in foreign exchange markets have increasingly started utilising technical trading rules to support their trading decisions. In the academia however, studies that have focused on the profitability of trading rules in currencies have witnessed similar resistance of the academic body, as it was the case within the equity space. The earliest noteworthy study is Sweeney (1986), which is the first publication that presents trading rule returns on a risk-adjusted basis, with a statistical framework that allows for significance test.

Sweeney (1986) bases his study on a set of daily data of the Dollar D-Mark exchange rate, as well as the overnight Fed funds rate and the Frankfurt interbank loan rate over a

time period from 1975 to 1980. After removing non-trading days the sample equates to 1289 trading days. In order to assess the performance between a set of technical filter rules, which he implements on a long only basis, and a buy and hold strategy, Sweeney (1986) uses the CAPM methodology. He transfers the CAPM setup into the currency space by defining the interest rate differential plus a constant risk premium as the market price for risk. The expected excess return of a trading filter over the market price of risk is then dependent on the beta and the ex-ante market premium for the days when the trading strategy is in the market. For the days when the strategy is out of the market the expected excess return of the strategy is zero. Sweeney (1986) also adjusts the return of the buy and hold strategy for the non-trading days on which no market risk premium is earned. The results of Sweeney's (1986) study suggest that even after correcting for risk by utilising the CAPM adjustment, the application of trading rules leads to significant excess returns. He suggests that these returns might well be a result of destabilizing speculation (i.e. intervention by central banks), time varying risk premia or market inefficiencies.

Indeed one of the main arguments brought forward when it comes to the profitability of technical trading rules is the fact that central banks, which are big players in the currency market, are not profit orientated. Taylor (1982) was one of the first to investigate central bank behaviour. His first observation is the fact that during the seventies central banks lost billions of dollars intervening in currency markets. Taylor (1992) points out that central banks follow a policy of "leaning against the wind", which might present itself in pegging the existing exchange rate to another currency when its equilibrium level changes. He also suggests that central banks can only support their currency for a limited amount of time but are eventually forced to allow the adjustment to take place, and when this happens they lose significant amounts of money. Taylor (1982) refers to Friedman's (1953) argument. If the aim of central bank intervention is to promote economic efficiency by reducing deviations from the equilibrium exchange rate, central banks should make profits from currency market intervention, because they are better informed than other market participants. However, central banks as a group have generally lost money in the process of intervention. Furthermore, he also points out that in the absence of speculators who were betting against central banks, the losses of currency intervention would have been only half the amount they actually have been. He also indicates that speculators' profit is likely to be a proportionate share of the central bank's loss. Hence following this argument, academia has posited central bank

intervention as one of the main deterrents for trading rule profitability. Szakmary and Mathur (1997) are the first to quantify the impact of central bank intervention. They indicate that central bank activity is the key driver of trading rule profitability. They base their results on a sample of daily returns of five currencies¹⁴ versus the US Dollar from 1977 to 1991. They arbitrarily test a set of various moving average rules. The median return of the moving average trading rule ranges between 5.4% and 9.8% depending on the currency pairs chosen. The exception hereby is the Canadian dollar, which exhibits negative trading rule returns. Based on a regression analysis they suggest that leaning against the wind intervention helps explaining that median moving average trading profits for various currencies are greater than zero.

The observations of the study of Szakmary and Mathur (1997) are confirmed by the work of LeBaron (1999). His results suggest that trading rule profits are highest during periods of central bank intervention. The study is based on weekly data from 1979 to 1992, for the Deutsch Mark and the Japanese Yen. The adjustment of the interest rate differential is done using the one-week eurorates. All return calculations are based on Wednesday London closing prices. In addition to that, the study utilises a time series of intervention values provided by the Federal Reserve. The trading rule tested is a simple moving average rule, whereby the prevailing price is compared to the average price of the previous thirty trading days. LeBaron (1999) tests the profitability of this trading rule using the full dataset. He also uses a dataset where weeks in which central bank intervention has occurred are removed. The key finding of his study is the when the impact of central bank activity is removed; the returns available from trading rules are significantly diminished. However, in his conclusion he makes the point that there might be a simultaneity problem. Interventions and profits may be driven by the same common factor and as a result of that, the causality lined out by him might be spurious. Neely (2002) picks up on this argument. His study analyses intraday data for five currency pairs¹⁵. The intraday data points for the range between 3 and 7 observations depending on currency. The exact data range of each of the currency pairs varies slightly, depending on the intervention activity of the various central banks. The data ranges are chosen to maximise the observation window. Overall the paper covers a time range from the early to mid-eighties to the mid to late nineties. The paper is based on a 150 day moving average trading rules. The results suggest that in the case of the US,

¹⁴ Deutsch Mark (DEM), Canadian Dollar (CAD), Japanese Yen (JPY), Swiss Franc (CHF), Pound Sterling (GBP)

¹⁵ AUD/USD, CHF/USD, DEM/USD, JPY/USD

German and Swiss interventions, trading rule returns precede the actual intervention. Trading rule returns in the Australian dollar on the other hand tend to lag intervention. Neely (2002), however, suggests that the trading rule returns for the Australian Dollar are unlikely to be a result of central bank intervention. The direction of trading signals and intervention make it implausible that the intervention is actually generating those returns. The key finding of Neely (2002) is that intervention does not generate returns itself. Currency intervention comes as a reaction to strong and very profitable short-term trends within currency markets. In that sense he confirms LeBaron's (1999) argument of simultaneity. Neely (2011) reemphasises in his survey of technical analysis in the foreign exchange market, the argument that central banks tend to intervene contrary to strong exchange rate trends and trading rules tend to profit from such trends. As a consequence, there is a positive correlation between intervention days and trading profits, however there is no causality between both.

The notion of time varying risk premia has also been used to explain the forward discount bias, which is the fact that the uncovered interest parity does not hold. Hence the prevailing spot rate does not converge to the expected spot rate implied by the interest rate differential. Cavaglia, Verschoor and Wolff (1994) point out that one of the most common rationales for the explanation of this carry effect is the fact that domestic and foreign assets are imperfect substitutes. Hence, any rational investor would demand a risk premium for holding foreign assets. None of the results of this area of academic research can be interpreted without any ambiguity.

Neely, Weller and Dittmar (1997) apply a genetic program that searches for an optimal trading rule. This approach aims to control for data snooping biases. The paper is based on prices for six currency pairs¹⁶, with a time sample spanning from 1974 to 1995. It is split into three sub-periods, which constitute selection, training and testing period for the genetic code. The key findings of their paper are that different currency pairs produce higher trading returns than others and that different currencies pairs also favour different sets of trading rules. Furthermore, all of the "genetically grown" trading rules show out of sample profitability when compared to bootstrapped benchmark simulations. Neely, Weller and Dittmar (1997) also investigate whether trading rule returns are a result of genuine market inefficiencies, or whether they are risk premiums received for taking systematic market risk. They use both a world market index and several national indices as benchmarks. Their results show only one value that suggests

¹⁶ GBP/USD, CHF/USD, DEM/USD, JPY/USD, DEM/JPY, GBP/CHF

a significant positive relationship between the trading rule results of a currency pair and a market index. Most of the results suggest no or even a negative relationship to equity market indices. Hence, excess returns observed are not a risk premium earned for taking systematic risk.

Kho (1996) also examines whether the results of various technical trading rules can be attributed to time varying risk premia. His study is based on daily data and his results indicate that periods of higher or lower returns identified by the technical rules largely correspond to those of higher or lower conditional expected returns, due to high or low risk premia and volatility. This suggests that there is an element of time varying risk premia.

Okunev and White (2003) point out that while higher frequency data might be subject to time varying risk premia; this is not the case for lower frequency data. Furthermore, trading strategies that are implemented on a long short basis tend to have zero covariance to markets. Chapter 3 of this thesis will have a closer look at the aspect of systematic risk taking.

A recent study by Qi and Wu (2006) picks up the argument of data snooping within currency trading rules. They apply the methodology of data snooper introduced by White (2000) to a universe of daily rates of seven¹⁷ currencies against the USD over a time period from April 1973 to December 1998. The results of their study are based on three data snooping proven test criteria: excess returns, Sharpe Ratio and Jensen's Alpha. The results suggest that the best performing trading rules, according to White's data snooper, are short-term channel breakout rules for the Japanese Yen and the Swiss Franc and short-term moving averages for the other currency pairs. Without accounting for transaction costs the mean excess returns over a buy and hold strategy are unanimously positive in the range of 4.02% to 12.81% per annum. After accounting for one-way transaction costs of 4bps the excess returns are still positive in the range of 2.14% to 11.46%. Incorporating transaction costs in the data snooping algorithm does not materially alter the results; the data snooper still favours short-term moving average rules for most of the currencies, with the exception for the Canadian dollar where it picks a trading range break rule and the Swiss Franc where it still favours a channel breakout rule. In a second step Qi and Wu (2006) split the dataset into two sub-samples. The first sub-sample spans from 1974 to 1985, the second sub-sample covers the time

¹⁷ CAD, DM, FRF, ITL, JPY, CHF and GBP

period from 1986 to 1998. While moving average rules still remain top performers amongst the universe of trading rules, the overall results suggest that the statistical significance and the profitability of trading rule returns has diminished considerably in the most recent time period. In the first sub-sample the data snooping proven trading rules for each of the currencies are statistically significant. In the second sub-sample this is only the case for the Japanese Yen, Canadian Dollar, Deutsch mark and Swiss Franc on a 10% confidence level. Trading rules for the other currencies are not statistically significant. Qi and Wu (2006) confirm the validity of these results by applying the data snooping test to cross exchange rates. Moreover, Qi and Wu (2006) apply an out of sample test to the trading rules identified in the first sub-sample. The returns generated on an out of sample basis are considerably less than the in-sample returns. Nonetheless, with the exception of the Italian Lira all of them are statistically significant on the 10% level. The results of the paper indicate that trading rules do offer statistically significant as well as economically important excess returns. However, these returns have diminished over time.

Neely (2011) makes the point that, within the context of currency trading rules, any explanation for trading rule profitability that is based on the arguments about foreign exchange intervention, systematic risk taking, or data mining, can be rejected confidently. Hence, it can be argued that trading rule profitability might stem from different sources such as behavioural traits of investors or market micro structure.

c) Other Explanations for Trading Rule Profitability

Recent studies such as Osler (2003) and Friesen, Weller and Dunham (2009) find the reasons for the strong profitability of trading rules and technical analysis in investor behaviour or more generally in the microstructure of the currency markets. Friesen, Weller and Dunham (2009) shed light on the aspect of investor behaviour from the perspective of behavioural biases such as the confirmation bias, which leads to autocorrelation in price pattern. The confirmation bias has extensively been documented in psychological literature. Within a financial context it suggests that investors, who acquire information and trade on the basis of that information tend to bias their interpretation of subsequent information in the direction of their original view. This would imply that, while market participants interpret large signals in a rational way, less informative signals are not interpreted rationally. Market participants' bias their

interpretation of less informative signals, which arrive more frequently, towards the most recently observed large signals. This behaviour generates price pattern such as “head and shoulders”. Most importantly it also constitutes a source of price momentum, which can be exploited by moving average trading rules. Friesen, Weller and Dunham (2009) use a jump-diffusion process in a discrete-time framework to model the process by which information is revealed. Hereby they assume that low-frequency signals are more informative than high frequency signals. Therefore, low frequency signals generate jumps while high frequency signals generate diffusion. They also assume that high frequency signals, which are subject to a cognitive bias, are risk neutral. Low frequency signals are assumed to be processed in a rational way, hence, risk averse. This simplification allows focusing exclusively on the expectations of biased traders in the price discovery process. The empirical analysis of their model, based on S&P 100 data, suggests that the pattern resulting from the model conform to a number of well-documented trading strategies. It also indicates that return autocorrelations are negative over very short horizons, positive over intermediate horizons, and become negative again over long horizons. These findings are very much in line with the well-documented empirical properties of US equity prices (see Jegadeesh and Titman (1993, 2001)). The paper also indicates the existence of negative weekly autocorrelations immediately after extreme information events with strong persistent momentum emerging several weeks after an extreme return.

Osler (2003) on the other hand side finds the rational for pattern in currency price movements in the microstructure of currency markets. The study analyses a dataset of almost 9700 stop-loss or take-profit orders placed by a large investment bank for three exchange rates¹⁸ from September 1999 to April 2000. The paper suggests that “support” and “resistance” levels can be key indicators for accelerated momentum or reversals, depending on whether they are broken or not. This is due to the distribution of the placement of stop-loss and take-profit orders by clients, which tends to cluster around round numbers. Furthermore, the most critical aspect of the clustering is the fact that it differs between stop-loss and take-profit orders. A take profit order is designed to lock in profits from a favourable move in the exchange rate; therefore it can be argued that it tends to reverse existing price trends. A stop-loss order on the other hand is designed to cut losses if the currency moves against the original view; hence, it is a factor that intensifies trends. Osler (2003) finds that while take-profit orders are mostly clustered

¹⁸ GBP/USD, DEM/USD, JPY/USD

around round numbers, stop loss orders have a pronounced tendency to be placed just beyond round numbers. Buy orders are often just above and sell orders are just below the round number. This would suggest that “support” and “resistance” levels, which tend to be round numbers, are key indicators for either a trend reversal if the spot price fails to cross the level, or trend acceleration when levels are crossed. Osler (2003) further investigated this thesis by applying a bootstrap simulation. Here results reaffirm the idea that there is a self-fulfilling dynamic between order placement and exchange-rate dynamics. Hence, it can be argued that placing orders around clusters is a rational action by market participants. Furthermore, technical analysis might be a fully rational method of exploiting the institutional features of foreign exchange markets.

These findings go hand in hand with the Adaptive Market Hypothesis, which is another recent concept that aims to explain the behaviour of asset markets. Neely (2011) makes a strong case for the Adaptive Market Hypothesis as the most appropriate framework to characterise modern capital markets. The Adaptive Market Hypothesis is introduced by Andrew Lo (2004), who challenges Fama’s Efficient Market Hypothesis in as far as its principles are based on the assumption of a steady state market environment. Lo (2004) argues that markets as well as economies follow evolutionary paths. As per his paper from 1999, Lo argues that, while the Efficient Market Hypothesis can be summarized by three P’s, prices, probabilities and preferences, the key weakness of the Efficient Market Hypothesis lies within the aspect of preferences and particularly behaviour of market participants. The asymmetry of risk aversion versus risk seeking in a scenario of potential gains and losses comes hereby to mind. He also presents evidence of Lo and Repin (2002), who find that even for highly experienced investors, physiological variables that are linked to the autonomic nervous system are highly correlated to market events and market variables. Hence, it can be argued that emotional responses are an important factor of processing financial risk. Therefore Lo’s (2004) main question with regards to the validity of the Efficient Market Hypothesis is whether market forces are sufficiently powerful to overcome any behavioural biases. While this question cannot be answered on academic grounds, recent years, during which the world has gone through a series of financial market crises, have provided plenty of anecdotal evidence that would suggest that market forces might not be sufficiently strong to overcome the behavioural aspect of market dynamics.

Under this paradigm, one can describe financial markets as an aggregation of behavioural biases combined with the market forces of supply and demand. Therefore

as market participants (humans) follow the concepts of “evolutionary psychology”, market behaviour is likely to follow similar concepts. Lo (2004) makes the point that capital market participants can be split into different groups, (“species”), such as retail investors, pension funds, market makers, hedge funds, which have very distinct preference. Therefore, a change in the market environment, might not only be induced by changing preferences of all market participants, but also by a change in the composition of investor groups (“species”) competing for specific assets. Moreover, he argues that within any given market environment there might be an asset such as US treasuries that is subject to competition by many investors within a group or across groups. Such an asset is likely to be highly efficiently priced due to the laws of demand and supply. Other assets that see less competition between market participants are likely to be less efficiently priced.

The practical implications of such market framework can be described as follows. Similarly, to the Efficient Market Hypothesis, the risk reward trade off will differ across assets; it will also change over time. Moreover, given the evolutionary aspect of the Adaptive Market Hypothesis, Lo’s (2004) concept also suggests that arbitrage opportunities continuously arise as the market environment changes. These arbitrage opportunities subsequently disappear, as market participants exploit them. Furthermore, the performance of certain investment strategies will be strong at times and weak at other times due to changing preferences. By the same token investors have to adapt to the changing market environment in order to achieve persistent levels of return. Therefore, they either search for new investment opportunities, or they innovate the way by which they have been able to generate returns within the existing investment opportunity set. Both of these propositions counter the traditional argument of the Efficient Market Hypothesis, which suggests a continuous trend towards more market efficiency as a consequence of persistent arbitrage.

2. Motivation of the Chapter and Main Contributions

This chapter is an extension of the first chapter, which analyses survival pattern within momentum trading signals. The methodology applied in the first chapter is based on the concept of “Runs Tests”, which has been used in fields such as meteorology. The essence of “Runs Test” studies is to compare empirical return pattern to some pre-specified benchmark return pattern. Fama (1965) picks this methodology up to analyse the persistence in equity returns. Jochum (2000) and then later Kos and Todorovic

(2008) extend the runs test framework, by utilising the Product Limit Estimator, which is a bio statistical tool that is widely used in medical research. This approach allows the analysis of return patterns of more complex trading rules, which had previously not been possible under the Fama (1965) specification. The objective of the first chapter is to introduce an alternative methodology for detecting market inefficiency. The key finding of the chapter is the fact that empirical momentum signals of very short-term moving average combinations outlive their theoretical benchmark signals, and that empirical momentum signals created of longer-term moving average combinations have lower lifetime expectancy than theory would suggest. Moreover, most of the deviations from market efficiency tend to deteriorate over time, to the point where all of the momentum signals exhibit survival times that are statistically equivalent to what is suggested by their respective benchmark processes. While trading strategy based on these findings is implemented in the first chapter, the aim of this chapter is not to search for a superior trading rule. Moreover, while the first chapter compares the empirical survival time of momentum trading signals with theoretical survival times, it does not make use of much of the other information that survival analysis can provide. This present chapter does. Lifetime statistics in form of the Product Limit Estimator also provides absolute as well as conditional survival probabilities for each stage of the life cycle of an investigated variable.

Therefore one can make statistically valid assessments about the likelihood of a momentum signal surviving two, three or four days or even more. Such information can be used to create trading rules that should be able to outperform generic trading rules. The key objective of this chapter is to assess whether trading rule enhancements that utilize information from lifetime analysis can generate returns that are superior to the returns generated from equivalent strategies that don't use such enhancements. The chapter uses White's (2000) data snooping methodology to assess whether the information provided by the survivorship analysis can be used to design trading rule enhancements that outperform generic trading rules. While Qi and Wu (2006) apply White's (2000) data snooping methodology within the context of foreign exchange markets, the academic contribution of this chapter can be summarised as follows. Firstly, the chapter analyses the performance of enhancements of moving average crossover trading rules, as opposed to picking the best trading rule out of a heterogeneous universe of trading rules. Secondly, the chapter undertakes an extended analysis of sub-samples, facilitating the analysis of persistence in performance of single trading rules. This has been done in previous studies as well, however, to a much lesser

extent. Finally, the chapter proposes to look at the results of White's data snooping test in a relative context as opposed to an absolute context. The results of the sub-sample analysis suggest that there is a great deal of clustering of currency pairs amongst the top trading rules over time. Therefore, the chapter looks at average White statistics for single trading rules across all currency pairs. This has not been done before.

B. Data and Methodology

1. Data, Return and Moving Average Calculations

The dataset used in this chapter is the same as the dataset used in the first chapter. It contains daily New York closing mid-prices for G10 currencies, as well as three month cash rates for corresponding countries. It spans from the 4th of January 1974 to the 31st of December 2009. After adjusting for non-trading days, the sample contains 9025 data points. The time series for the EUR rate is backfilled with the historic Deutschmark (DEM) rate, with the original EUR fixing rate of 1.95583 DEM per 1 EUR, as of 1 January 1999. The time period analysed is split into nine sub-samples, whereby the first eight sub-samples consist of exactly 1000 observations and the ninth sub-sample consists of 1025 observations. All calculations are based on returns that are adjusted for the interest rate differential. This is done to mimic the returns obtainable from a futures based trading strategy. The exchange rates are expressed in units of domestic currency versus one unit of foreign currency. Equation 1 calculates an interest adjusted return time series. The first term represents the daily interest rate differential between foreign (r_f) and domestic (r) currencies. The second term shows the return from currency movement. S_t is hereby the currency spot price at time t .

$$(1) \quad R_{I,t} = \left[\left(\frac{1+r_f}{1+r} \right) * \left(\frac{S_t}{S_{t-1}} \right) \right] - 1,$$

The interest rate calculations in Equation 1 are based on the Money Market Basis convention (Actual/360). The adjusted return time series, obtained from the equation, results in approximate currency returns that can be earned by following a futures based investment strategy. The calculations for the interest rate differential are based upon the three month T-Bill rate, for which clean time series across all countries in the G10 currency universe exists. While the 3 month T-bill rate is only the second best


adjustment factor after the overnight rate, the previous chapter has proved that both interest rate adjustments are equivalent.

This chapter also follows the previous chapter in its definition of trading rules. It uses a simple price moving average filter. The rationale for this very basic choice of trading signal is the fact that it is parsimonious. The aim of this chapter is to enhance trading rules by way of using the results of the survival analysis to identify periods, within the “life” of a trading signal, where it is optimal to increase exposure or to reduce exposure or to completely exit the position. If the survivorship analysis has the power to improve the profitability of a plain moving average trading signal, it is likely to improve the results of more sophisticated trading rules as well. Equations 2 and 3 describe the crossover signals used to calculate the trading filters.

$$(2) \quad \text{Positive Momentum} = \frac{1}{S} \sum_{i=0}^{S-1} S_{t-i} \geq \frac{1}{L} \sum_{i=0}^{L-1} S_{t-i}$$

$$(3) \quad \text{Negative Momentum} = \frac{1}{S} \sum_{i=0}^{S-1} S_{t-i} < \frac{1}{L} \sum_{i=0}^{L-1} S_{t-i}$$

The time periods for the short-term moving averages here denoted as (S) range between 1 to 5 days as well as 10, 15, 20 and 25 days. The time periods for the long-term moving averages (L) are defined as 5, 10, 15, 20, 25 and 30 days. Any short-term moving average has to be shorter than any long-term moving average. Equation 2 suggests that a positive momentum signal is established when the short-term moving average is equal to or above the long-term moving average. Equation 3 indicates that a negative moving average signal is established when the short-term moving average lies below the long-term moving average. When applied to the data, the moving average trading rule gives a series of positive or negative trading signals of different lengths, scattered along the empirical time series. Figure 2-1 presents a hypothetical trading signal and its survival distribution.



On the positive side, one momentum signal survives five days, three momentum signals survive three days, one signal survives two days, and two signals live for one day. On the negative side, one signal survives for six days, one for five days, two signals survive for three days and two survive for two days, one signal survives for one day.

2. Product Limit Estimator

Given the fact that these positive and negative signals occur randomly along the time series, it is not possible to make any statistical inferences from that data time series with traditional statistical methods. Kaplan and Meier (1958) construct the Product Limit Estimator (PLE), which allows ordering data such that survivorship probabilities can be calculated and inferences can be made. The idea behind the PLE is a stepwise ordering of any pre-specified signal according to the duration of its survival. Within the context of trading, any positive or negative signal is ordered according to its length of survival. This means that all signals that survive two periods are summed up. From this population of two period survivors, all signals that survive a further period are extracted and added together. This is done until the entire data sample is ordered according to the length of survival of positive and negative trading signals. Equation 4 shows the Kaplan Meier estimator, whereby n_i represents the number of observations in the sample that have survived until the time period i . d_i is the number of observations that cease to exist in period i


$$(4) \quad \hat{S}(t) = \prod_{i=1}^k \frac{(n_i - d_i)}{n_i}$$

The ratio between $(n_i - d_i)$ and n_i represents a conditional probability of survival for that time interval. It divides the number of observations that have survived over the interval by the number of observations that were at risk at the beginning of the interval. By multiplying these periodical survival probabilities, the PLE, which is the cumulative survival probability, can be obtained.

Figure 2-2 shows the absolute and the conditional survival probability of the 1 to 20 day moving average signal of the USDGBP exchange rate. As pointed out before, the PLE is the absolute survival probability of a signal. It is shown in Figure 2-2 as a dark grey downward sloping line. While Figure 2-2 shows the PLE as a smooth line, it is actually a step function with steps at points where the survival of an observation is assessed. This chapter assesses the survival of a signal on a daily basis. The PLE shown in Figure 2-2 can be interpreted as follows: 82.6% of the observations in the sample survive more than one day, 41.2% of the sample survive more than eleven days, 15.3% of the sample survive longer than twenty six days, and so on. Figure 2-2 also shows the evolution of

conditional survival probabilities in light grey. This measure also contains useful information. When looking at the life cycle of a trading signal over a certain time horizon, one can identify time periods within the life of such signal where the conditional periodic survival probability is relatively high. On the other hand side there are time sequences where the conditional periodic survival probability is relatively low. Figure 2-2 suggests that the conditional periodic survival probability gradually increases up to 25 days. Subsequently it decreases until 43 days. Thereafter it ceases to exist. Both survival measures, the unconditional as well as the conditional, contain potentially valuable information when it comes to the analysis of trading rules. The pattern of the unconditional survival probability might be useful to decide when to exit a position, or the relative survival probability might be informative when it comes to relative weighting of trading rules over time.

FIGURE 2-2: PRODUCT LIMIT ESTIMATOR AND CONDITIONAL SURVIVAL PROBABILITIES OF THE PLE



Absolute and conditional survival probability of the positive 1 to 20 day moving average signal of the USDGBP exchange rate.

3. White's Data Snooper

This study builds optimised trading rules utilising unconditional and conditional survival probabilities. The aim of this is to assess whether the information provided by these survival probabilities adds value when built into a trading rule. One of the key challenges hereby is to distinguish whether any potential value added stems from the

true superiority of a trading rule or whether it is a mere artefact of data mining. White (2000) presents a test methodology that allows distinguishing between them. He draws from the methodology developed by Diebold and Mariano (1995) and West (1996), which evaluates the superiority of any given statistical model relative to a benchmark. The innovation of the methodology lies in the definition of the benchmark, which comprises the distribution of all statistical models that have been included in the search for the best statistical model. Sullivan, Timmerman and White (1999) call this methodology the ‘Data Snooper’ and apply it to the analysis of trading rules. The Data Snooping test is based on $l \times 1$ performance statistic, with l representing the number of trading rules that are considered in the search for the best trading rule, comprising the benchmark population, and n is the number of prediction periods that run from R to T , hence $T = R + n - 1$.

$$(5) \quad \bar{f} = n^{-1} \sum_{t=R}^T \hat{f}_{t+1}$$

$\hat{f}_{t+1}$ is the performance measure for $t+1$. In its general form it can be described as shown in Equation 6, whereby Z can be generally described as a vector of dependent variables and predictor variables and $\hat{\beta}_t$ is a vector of estimated parameters.

$$(6) \quad \hat{f}_{t+1} = f(Z_t, \hat{\beta}_t)$$

When constructing trading rules, parameterisations are set, hence $\hat{\beta}_t$ does not have to be estimated. Furthermore, trading rules in different parameterisations generate returns directly, which can be used to measure performance. Each of the technical trading rules that are assessed is indexed by a subscript k . Consequently $f_{k,t+1}$ is given in Equation 7

$$(7) \quad f_{k,t+1} = \ln[1 + y_{t+1}S_k(x_t, \beta_k)] - \ln[1 + y_{t+1}S_0(x_t, \beta_0)] \quad k = 1, \dots, l$$

with

$$(8) \quad x_t = \{X_{t-i}\}_{i=0}^R$$

Whereby X_t is the interest rate adjusted price time series for different currency pairs as given in Equation 1 and y_{t+1} shown in Equation 9 are the returns from the time series X_t .

$$(9) \quad y_{t+1} = \frac{(X_{t+1} - X_t)}{X_t}$$

Reverting to Equation 7, Sullivan, Timmerman and White (1999) indicate that $S_k(\cdot)$ and $S_0(\cdot)$ are “signal” functions that transform the information embedded in the system parameters β_k and β_0 into market positions. The signal functions can take three values: 0 for neutral, 1 for long and -1 for short. In a general form, the null hypothesis of the framework is to test if the best trading rule delivers a performance that cannot be distinguished from the performance of the benchmark. A rejection of the null hypothesis, given in Equation 10, indicates that the best technical trading rule achieves performance that is superior to the benchmark.

$$(10) \quad H_0: \max_{k=1, \dots, l} \{E(f_k)\} \leq 0$$

The null hypothesis can be evaluated via the application of the stationary bootstrap, as introduced by Politis and Romano (1994). The stationary bootstrap is a version of the block bootstrap, whereby the length of the blocks that are patched together is random, following a geometric distribution. For the implementation of the stationary bootstrap, the present chapter follows Sullivan, Timmerman and White (1999) closely and Qi and Wu (2006) closely, who propose a block length of 10 observations. The mathematical proof of the appropriateness of the stationary bootstrap for the “Data Snooper” can be found in White (2000). The resampling procedure yields multiple observations of trading rule results of $\bar{f}_k$ which are defined as $\bar{f}_{k,i}^*$ whereby i is the index of the bootstrapped sample. The resampling procedure is based on 500 iterations. It facilitates the construction of the test statistic given in Equations 11 and 12.

$$(11) \quad \bar{V}_l: \max_{k=1, \dots, l} \{\sqrt{n}(\bar{f}_k)\}$$

$$(12) \quad \bar{V}_{l,i}^*: \max_{k=1, \dots, l} \{\sqrt{n}(\bar{f}_{k,i}^* - \bar{f}_k)\}, \quad i = 1, \dots, 500$$

The p-value for the null hypothesis of the white reality check is then obtained by comparing $\bar{V}_l$ to the quantiles of $\bar{V}_{l,i}^*$. This is done across all the l trading rules. Hence the reality check p-value incorporates the effects of data snooping from the search over the l rules. Sullivan, Timmerman and White (1999) implement the above test also on the basis of Sharpe Ratios, which is a more appropriate approach than merely looking at the trading strategy with the highest return, as it provides a risk adjusted performance evaluation framework. The Sharpe ratio measures the average excess return over the risk free rate per unit of risk taken. This chapter applies the Sharpe Ratio criteria. However, it simplifies the Sharpe ratio formula by assuming a zero interest rate as hurdle rate. This is appropriate given the fact that the trading strategies implemented in this chapter are based on interest rate adjusted returns, which aim to replicate the returns of a futures trading strategy. Furthermore Qi and Wu (2006), who apply the same methodology, argue that the natural benchmark of a currency speculator is one of not being invested in any currency and therefore not earning any interest rate. The Null hypothesis under this specification is given in Equation 13

$$(13) \quad H_0: \max_{k=1, \dots, l} \{g(E(h_k))\} \leq g(E(h_0))$$

h is a vector that consists of two components that are given in Equations 14 and 15 and g is given in Equation 16

$$(14) \quad h_{k,t+1}^1 = y_{t+1} S_k(x_t, \beta_k)$$

$$(15) \quad h_{k,t+1}^2 = (y_{t+1} S_k(x_t, \beta_k))^2$$

$$(16) \quad g(E(h_{k,t+1})) = \frac{E(h_{k,t+1}^1)}{\sqrt{E(h_{k,t+1}^2) - (E(h_{k,t+1}^1))^2}}$$

Equation 17 shows the sample statistic, whereby $\bar{h}_k$ and $\bar{h}_0$ are averages that are computed over the prediction sample for the k^{th} trading rules as well as the benchmark model. The calculation for $\bar{h}_k$ and $\bar{h}_0$ is given in Equation 18

$$(17) \quad \bar{f}_k = g(\bar{h}_k) - g(\bar{h}_0)$$

$$(18) \quad \bar{h}_k = n^{-1} \sum_{t=R}^T h_{k,t+1}, \quad k = 0, \dots, L$$

The evaluation of whether a trading rule is superior to other trading rules is then again done on the basis of the Bootstrapping procedure created by Politis and Romano (1994) whereby $\bar{f}_{k,i}^*$ is given by Equations 19 and 20.

$$(19) \quad \bar{f}_{k,i}^* = g(\bar{h}_{k,i}^*) - g(\bar{h}_{0,i}^*), \quad i = 1, \dots, 500$$

$$(20) \quad \bar{h}_{k,i}^* = n^{-1} \sum_{t=R}^T \bar{h}_{k,t+1,i}^*, \quad i = 1, \dots, 500$$

4. Trading Rules to be Evaluated and Link to Survival Curves

As shown in the previous section the data snooper evaluates the best trading rule relative to a universe of trading rules that go into the search for the best trading rule. As a consequence Sullivan, Timmerman and White (1999) specify a comprehensive and heterogeneous universe of 7846 trading rules that go into the search for their best trading rule. Later studies that also aim to identify the single best trading rule use similarly wide universes. The purpose of this chapter is somewhat different from other studies. The aim of the chapter is to use the information embedded in the analysis of survivorship curves and create enhanced trading rules on the basis of that information. The focus of the analysis is therefore a relative one as opposed to an absolute one. This aspect has to be borne in mind when constructing the test setup. Following the framework established in the first chapter, this chapter will use simple moving average crossover trading rules as a basis, which are then enhanced four ways. The four enhancements plus the original cross over trading rule represent the universe against which each of the trading rules is compared. After removing reverse currency quotations, each trading rule is compared to an overall universe of 8775 trading rules. Before describing the details of the trading rule enhancements, the chapter will touch on the link between survivorship analysis and trading rules, as well as the rationale for using the different enhancements. In the previous chapter the intellectual link between trading rule returns and the length of survival of trading signals was already introduced. In essence, the longer a trading signal survives, the more chance it has to generate returns. Moreover as time passes the stop loss mechanism, in form of the long term moving average is ratcheted up as new (positive) returns enter the calculation. Hence in

absolute terms it can be expected that signals that live on average longer than other signals, are likely to have higher returns as well. However, the distributional characteristic of the returns after a signal has been generated, is not known. Within the time period a trading signal is alive the return generation can be such way that either the magnitude of returns in few days when the trading rule is established are so high that the signal still persists even when the underlying returns have turned negative, or that the returns remain marginally positive over the time period where the signal remains positive, or variations within the two extremes.

However, what the survivorship analysis enables is to analyse the distributional characteristics of the trading rule signal, from which assumptions about the return generation of the strategy can be made. Figure 2-2 shows the product limit estimator, as well as the conditional survival probability of the SR1/LR10 trading signal of the USDGBP exchange rate. When looking at the conditional survival probability in Figure 2-2 it becomes evident that this conditional survival probability follows the shape of an inverted U. The light grey line, which indicates the conditional survival probability, indicates that the periodic (conditional) survival probability increases over time and then subsequently decreases. The periodic conditional survival probability of the trading signal given in the figure is 92.5% on day two. This means that assuming a momentum signal has been established, there is a 92.5% chance that it lives for two days. The conditional survival probability then increases to 93% for the third day. This means that there is a 93% probability that momentum signals that have survived for two days, will survive another day, etc. This pattern of gradually increasing and subsequently decreasing survival probability is common to all tested trading rules. The pattern described is a well-documented phenomenon within academic literature when it comes to the analysis the returns of a momentum strategy. Notably, Jegadeesh and Titman (1993, 2001) for equities and Menkhoff, Sarno, Schmeling and Schrimpf (2011) for currencies, find that momentum returns go through different stages over time. At the point of the signal generation, returns are weak; later they become more pronounced and then fade away. Bearing this dynamic in mind the first trading rule enhancement aims to exploit this well documented phenomenon by assuming that the returns generated from the dual crossover trading strategies follow the conditional survival pattern.

The second trading rule enhancement has a clearer link to the returns of the trading rule strategy than the first enhancement. This enhancement is linked to the product limit estimator, which is shown as the dark line in Figure 2-2. The product limit estimator

being the absolute, unconditional survival probability of a momentum signal is downward sloping. This indicates that the chance of a momentum signal surviving diminishes over time. Given the fact that the length of time a positive trading signal is generated defines the length of time the trading strategy is in place and has the potential to generate returns. Hence at the point when a trading signal is generated the expected value of a return generated on the first day of the trading strategy being implemented is higher than on the fifth or sixth day. This is due to the fact that the probability of survival of the trading signal (the potential to generate returns) is considerably lower on day six, than it is on day one. In order to reflect this, the strategy follows the product limit estimator.

- The first enhancement of the generic cross over trading rule is based on the historic conditional survival probability of the Product Limit Estimator. Given the fact that the conditional survival probability follows an inverse U shape, which is a widely documented pattern for returns generated from momentum signals, it is assumed that the evolution of returns in each cycle of trading signals follows the conditional survival. At the point of the signal generation, returns are weak; later they become more pronounced and then fade away. To exploit this pattern the first trading rule enhancement changes its exposure level according to the historically realised periodic conditional survival probability, while the generic trading rule remains fully invested as long as a momentum signal is alive. Using the example from 2- 2, this means that the enhanced strategy will be 92.5% invested on the day after the signal has been established; if the signal survives two days the exposure level goes up to 93% and so on.
- The second enhanced strategy is based on historic estimates of the unconditional periodic survival probability. Hence, it uses the historic realisations of the Product Limit Estimator to determine its exposure levels. Using the example of Figure 2-2, this means that the second enhanced strategy will be 76.5% invested on the day after the signal has been established; if the signal survives two days the exposure level goes down to 71.1% and so on. Trading rule returns that are generated directly after the trading signal has been established receive the highest weight, while later realisations have less weight.

The other two enhanced trading rules are variations of the first two trading rules. The disadvantage of the first two strategies is the fact that they have to rebalance exposures

every day. This adds to turnover, which can be very costly. The third and the fourth trading rule enhancements are done with the aim of reducing transaction costs, which come from continuous rebalancing of exposures. The design of both trading rule strategies is chosen such way that they replicate the essence of the first two strategies, while reducing the need to rebalance exposures. The exposure levels of the third and fourth trading rule enhancement vary between 90% and 110%. These levels are chosen arbitrarily. The intuition behind the choice of 110% vs. 100% vs. 90% exposure level is the fact that it represents a meaningful deviation from the unity exposure level of the standard trading rule, while not incurring too significant transaction costs when the rebalancing of exposures occurs.

- The third enhanced trading rule follows the intuition of the first enhanced trading rule, which assigns lower exposure to newly established trading signals. It increases that exposure over time and then reduces the exposure thereafter. In order to capture this dynamic, the trading filter assigns a 90% exposure to observations if the product limit estimator is above 0.8. Subsequently, during the observations where the product limit estimator is between 0.8 and 0.5 the exposure applied to the trading strategy will be 110%. As soon as the Product Limit Estimator falls below 0.5 the exposure level of the strategy goes to 100%. In practical terms this means that the third enhanced trading rule applies a lower exposure to returns observations following a newly established trading signal. It applies a higher exposure to the returns of trading signals that have been alive for a while and an equal weight to the returns of trading signals that have been established a longer time ago. Therefore the strategy mimics the dynamic of the exposure levels of the first trading rule enhancement.
- The fourth trading rule enhancement follows the intuition of the second trading rule enhancement. Mechanically it works similar to the third trading rule enhancement with the only difference that the exposure limits are assigned differently. It assigns an exposure of 110% when the PLE is above 0.8, an exposure of 100% if the PLE is between 0.8 and 0.5 and an exposure level of 90% if the PLE is below 0.5.

The key determinant as to whether the third and the fourth trading rule enhancements are able to outperform the first two trading rule enhancements is the difference between the incremental return generated versus incremental cost occurred from continuous

rebalancing. To illustrate that point one can take a ten day moving average signal as an example. Under the standard trading rule specification, there is no rebalancing of exposure levels over that time period. The first two trading rule enhancements would most likely require ten rebalancing transactions, while trading rule enhancements three and four only require three rebalancing transactions over that time period.

One could argue that each of these individual strategies has very different exposure levels embedded. Therefore the returns of these strategies are not comparable, because of the different levels of risk that come with the differences in exposures. For this reason each of the strategies are compared to each other on the basis of the Sharpe Ratio. The Sharpe Ratio criteria is applied in the data snooping test, which identify the best trading rules on an ex ante basis. It is also applied on an ex post basis when evaluating the trading rules. Therefore the general comparability of the different trading rule specifications is ensured.

To capture the dynamic of trading rule returns over time, White's data snooper is applied to a series of sub-sample. As indicated earlier the data sample spans from the 4th of January 1974 to the 31st of December 2009. After adjusting for non-trading days, the sample contains 9025 data points. This sample period is split into nine sub-samples, whereby the first eight sub-samples consist of exactly 1000 observations and the ninth sub-sample consists of 1025 observations. The data snooping test will be conducted for all but the first sub-sample. The first sub-sample is used to calculate the first survival curve, which is the basis for the enhanced trading rules of the second sub-sample. The survival curve of the second sub-sample is used as the basis of the enhanced trading rules of third sub-sample, and so on. This is done to avoid any circularity within the data snooping test. As pointed out earlier, the universe of the data snooping test is based on the set of regular moving average combinations as well as the four trading rule enhancements across all currency pairs. The fact that the data snooping test is done for a series of sub-samples is insightful as it allows the comparison of the performance of the trading rule variations as well as the persistency of best data snooping proven trading rules over time.

Moreover, the test results from the data snooping test can be ordered in different ways. One can look at the ten best trading rules out of a universe of 8775 trading rules. However, one can also look at trading rule returns in a relative context, by analysing the performance of the four trading rule enhancements against the generic trading rule across all currency pairs. To do so the chapter compares the distribution of percentiles

of White's P-Values for each individual enhanced trading rule specification (e.g. SR1/LR10, weighted according to unconditional probability) across all currency pairs. This distribution is then compared to the distribution of percentiles of the standard trading rule that has the same choice of parameters (e.g. SR1/LR10, standard trading rule). If the distribution of the enhanced trading rule is statistically significantly higher than the distribution of the standard trading rule, it can be concluded that the particular enhanced trading rule represents a genuinely better trading rule than the standard trading rule specification. This approach differs somewhat from Sullivan, Timmerman and White (1999) and Qi and Wu (2006) who merely look at the best few trading rules. Furthermore, the following sections will emphasise why such relative analysis is preferable to an absolute analysis with regard to foreign exchange markets.

5. Back-Test of Best Trading Rules

In a second step this chapter also back-tests the ten best "data snooping proven" trading rules. It has to be pointed out that the decision to search for the ten best trading rules is arbitrary. However, the rationale of choosing ten instead of five or twenty best trading rules lies in the fact that ten trading rules represent a good trade-off when it comes to the implementation of the back test. Using the top five trading rules is likely to make the back test results strong. However it imposes the risk of spuriousness. On the other hand, using the best twenty trading rules might not give any meaningfully strong results. When looking at trading rules on an absolute basis, 20 top trading rules would be probably preferable to 10 given the fact that the universe against which the back test is implemented is vast (8775 trading rules). However, when looking at trading rules in a relative context. This means assessing the parameterisations of the four trading rule enhancements (across all currencies) relative to the parameterisation of the generic trading rule. The universe against which the best trading rules are assessed becomes considerably smaller (156 trading rule enhancements and parameterisations). One could make a case for using a fixed percentage as opposed to a fixed number. Nonetheless, as shown later, the conclusions that are drawn from the results will stay the same regardless of whether the absolute number of best trading rules or the percentage number of best trading rules is used for the back test.

To avoid data circularities, the analysis spans from the second sub-sample to the last sub-sample. The data from the first sub-sample are used to create survivorship curves that go into the trading rule enhancement of the second sub-sample. The trading rule enhancements of the third sub-sample are based on survivorship data of the second sub-sample, and so on. Moreover, the back test is done in such way that the data snooping proven trading rules that have been identified in the second sub-sample are implemented in a trading rule starting from the third sub-sample to the end of the dataset. Hence, the ten best “data snooping proven” trading rules identified in the second sub-sample have a live track record that spans over the subsequent seven sub-samples, the enhanced trading rule in the third sub-sample will have a live track record in the following six sub-samples and so on.

a) Construction of Composite Trading Rules

Similar to the first chapter this back-test is carried out on the basis of a composite trading rule, which is constructed by summing up the number of positive trading signals and then deducting the number of negative trading signals from it. The reason for creating a composite trading rule is the fact that generating trades based on single trading signals might incur high trading costs, which most likely makes any trading strategy unprofitable. Furthermore, a combined trading rule also incorporates the interaction between moving average combinations. For instance, while longer-term moving average combinations might still point towards a long position, shorter-term moving averages might indicate an increasing short bias. The benchmark composite trading rule is equally weighting all moving average combinations of SR (1, 2, 3, 4, 5, 10, 15, 20, 25) and LR (5, 10, 15, 20, 25, 30), across all currencies. Given the fact that the standard composite trading rule consists of 39 moving average rules. The composite trading rule is then standardised by dividing the raw signal by the total number of moving averages. $+39/39 = +1$. This is done across all currency pairs, whereby all cross currency holdings are netted off, so that the end portfolio consists of long and short positions in nine currencies against the USD.

The trading rule back-test of the absolute specification of the data snooping proven trading rules will be conducted as follows: Once the ten best data snooping proven trading rules have been identified for a sub-sample, these ten trading rules are combined

in an equally weighted composite trading rule by summing up the number of positive trading signals and then deducting the number of negative trading signals from it. This composite trading rule is then compared to a benchmark composite trading rule. The components of the data snooping proven trading rule will vary depending on the sub-samples in which they have been identified.

When it comes to the results of the relative data snooping test, the back-test is conducted as follows: the ten best data snooping proven trading rules across the four enhanced specifications are combined in an equally weighted composite trading rule. This enhanced trading rule specification is then compared to an equally weighted benchmark composite trading rule, which consists of exactly the same trading rule parameterisation implemented on standard trading rules. So for instance if the enhanced composite trading rule consists of the parameterisation SR2/LR15 and SR3/LR20 of the first enhanced trading rule as well as the SR1/LR10 parameterisation of the second enhanced trading rule, the benchmark trading rule would consist of the SR1/LR10, SR2/LR15 and SR3/LR20 parameterisation of the standard trading rule.

b) Accounting for Implicit Transaction Cost

As mentioned in an earlier the aim of this section is to build a realistic link between the results of the survival analysis and trading rule returns. This is done by mimicking a futures based trading strategy, which is approximated by calculating interest rate adjusted returns as per equation 1. As mentioned earlier, the analysis in the dataset of the chapter consists of daily exchange rate levels, adjusted for the interest rate differential, calculated on the close of the New York trading session. Given the fact that the currency market is effectively a 24 hour market¹⁹, it is possible to trade at the point when the signal has been established. The previous chapter takes that approach. This is done under the assumption that foreign exchange markets are the deepest and most liquid financial markets. Therefore the aspect of slippage, i.e. the price movement caused by undertaking a transaction, is negligible. This assumption is valid, given the aim of the back test in the first chapter, which is merely linking the theoretical results from the survivorship analysis with trading data.

This chapter takes a different approach; this is because the aim of this chapter is different. The key objective of this chapter is a real live assessment of whether trading rule enhancements on the basis of survivorship models can generate returns that are

¹⁹ The Australian currency market opens at the time of the New York close.

superior to the returns generated from equivalent strategies that don't use such enhancements. Hence, neglecting the potential impact of slippage might be somewhat too onerous, when aiming to identify most optimal trading rules. In order to be conservative in the assessment, the chapter assumes that trading strategies gradually phases into positions. The assumption here is that once the composite trading rule alters its exposure level, the trading strategy can only alter the change in exposure gradually, without moving the market. In order to account for this opportunity cost of not being invested, which is in effect a form of "implicit" transaction cost, the trading strategy delays the increase or decrease in exposure levels by 24 hours. This assumption is very conservative. One would expect that real life trading strategies could be executed quicker than that without moving markets. However, for the purpose of a conservative analysis of trading rule profitability, it is an appropriate measure.

The time delay has a symmetric impact; it might have a negative or a positive impact on trading rule returns, depending on the direction of the currency movement. Therefore, besides the comparability of risk levels across trading rules, the breakeven transaction cost levels also incorporate all "explicit" transaction costs. Given the fact that differences in risk and turnover, as well as the implementation shortfall, are accounted for when constructing the transaction cost breakeven levels, the results of this analysis do not only allow for a wide comparison of different trading rules, they also give a representative indication of returns obtainable from a real life trading strategy. The transaction cost breakeven levels can be directly compared to historic estimates of bid-ask spreads and commissions. Similar to the previous chapter, the implementation cost of the trading strategy are calculated on an adjustment transaction basis, with a quarterly roll adjustment as and when futures contracts expire. For a graphical illustration of that please refer to Figure 1-33 in Chapter 1.

c) Calculation of Breakeven Transaction Cost Levels

Similar to the previous chapter, this chapter will also present the results of the trading rule implementation in form of breakeven transaction cost levels. Breakeven transaction cost levels are defined as the level of per trade transaction cost that needs to be incurred, in order for the trading rule to yield a risk adjusted return of zero. That implies that if the actual trading cost is lower than the transaction cost breakeven level, the strategy

will be profitable. Given the fact that different trading rules consist of different currency pairs, or different parameterisations, and have different levels of exposure at different times, it is reasonable to assume that each of the moving average trading strategies that are analysed is likely to have different levels of risk. This suggests that calculating breakeven transaction cost levels by merely setting the trading rule return to zero is not sufficient to make a fair comparison between trading strategies. In order to adjust for that, breakeven transaction cost levels are calculated against the ratio of trading rule return to trading rule risk. In order to achieve this, the transaction cost levels are incrementally increased until that ratio reaches a level of 0.01²⁰. Moreover, depending on its signal generation, each trading rule will have different levels of turnover. If the trading signal turns only 45% of positions over then the transaction cost are only applied to 45% of the portfolio. This means that each of the breakeven transaction cost levels is adjusted for the different levels of risk and turnover incurred by different trading rules and is therefore comparable across the universe of trading rules.

C. Empirical Evaluation

1. Best Absolute Data Snooping Proven Trading Rules

As pointed out earlier the aim of this section is to identify the best ten data snooping proven strategies for the nine different sub-samples. 2- 2-3 shows the results of White's data snooping test for all nine sub-samples. The first column in each sub-sample box specifies the enhancement, i.e. weighting according to conditional or unconditional survival probability. The second and third columns show the currency pair and the parameterisation of the trading rule. The last column of each sub-sample box shows White's P-value.

Looking at the results of the trading strategy, it becomes evident that most of the top ten trading rules across the sub-samples have P-values below the 10% or even 5% mark. This indicates that most of the trading rules shown below are statistically significantly better than the rest of the trading rule universe from which they have been selected. Nonetheless, the significance levels of the top trading rules are lower than the significance levels obtained in similar studies such as Sullivan, Timmerman and White (1999) and Qi and Wu (2006). The most likely explanation for this difference is the fact

²⁰ The reason for using 0.01 as opposed to 0 is technical

that the mentioned studies use a fairly heterogeneous universe of trading rules from moving average to channel breakout trading rules, to determine the best data snooping proven trading rule. Hence, one would expect that the heterogeneity of the universe of trading rules that is used in White (1999) and Qi and Wu (2006) results in higher P-Values of the best trading rules of that same universe. The present chapter, on the other hand, utilises a more homogeneous universe of trading rules.

FIGURE 2-3: BEST ABSOLUTE DATA SNOOPING PROVEN TRADING RULES FOR ALL SUB-SAMPLES

SUB SAMPLE 2				SUB SAMPLE 3			
ENHANCEMENT	CURRENCY	SR/LR	White P-VALUE	ENHANCEMENT	CURRENCY	SR/LR	White P-VALUE
Conditional Prob (rrb.)	USD/JPY	2/15	0.0480**	Conditional Prob (rrb.)	GBP/EUR	1/5	0.0445**
Unconditional Prob (rrb.)	USD/JPY	2/15	0.0545*	Unconditional Prob (rrb.)	GBP/EUR	1/5	0.0459**
Unconditional Prob (rrb.)	JPY/EUR	1/20	0.0573*	Standard	GBP/EUR	1/5	0.0477**
Unconditional Prob (rrb.)	JPY/EUR	3/25	0.0577*	Conditional Prob	GBP/EUR	1/5	0.0492**
Conditional Prob	USD/JPY	2/15	0.0577*	Unconditional Prob	GBP/EUR	1/10	0.0616*
Conditional Prob	JPY/EUR	1/25	0.0585*	Unconditional Prob (rrb.)	USD/EUR	3/15	0.0682*
Conditional Prob (rrb.)	JPY/EUR	1/25	0.0598*	Conditional Prob (rrb.)	USD/EUR	3/15	0.0716*
Conditional Prob	JPY/EUR	3/25	0.0614*	Conditional Prob	USD/EUR	3/15	0.0717*
Conditional Prob	JPY/EUR	1/15	0.0621*	Unconditional Prob (rrb.)	GBP/EUR	1/10	0.0766*
Standard	GBP/CHF	5/10	0.0622*	Unconditional Prob	GBP/EUR	1/5	0.0781*
SUB SAMPLE 4				SUB SAMPLE 5			
ENHANCEMENT	CURRENCY	SR/LR	White P-VALUE	ENHANCEMENT	CURRENCY	SR/LR	White P-VALUE
Unconditional Prob	USD/EUR	10/25	0.0634*	Unconditional Prob	EUR/CAD	1/5	0.1227
Conditional Prob	USD/GBP	10/20	0.0749*	Conditional Prob (rrb.)	EUR/CAD	1/5	0.1228
Unconditional Prob	JPY/EUR	2/30	0.0762*	Unconditional Prob (rrb.)	EUR/CAD	1/5	0.1265
Unconditional Prob (rrb.)	USD/EUR	4/25	0.0816*	Unconditional Prob	GBP/EUR	1/25	0.1274
Conditional Prob	USD/EUR	4/25	0.0819*	Standard	EUR/CAD	1/5	0.1282
Conditional Prob (rrb.)	USD/EUR	4/25	0.0823*	Unconditional Prob (rrb.)	EUR/CAD	1/10	0.1306
Conditional Prob (rrb.)	USD/GBP	10/20	0.0851*	Conditional Prob (rrb.)	EUR/CAD	1/10	0.1309
Conditional Prob	USD/GBP	1/15	0.0856*	Unconditional Prob	EUR/CAD	1/10	0.1336
Unconditional Prob (rrb.)	USD/GBP	10/20	0.0862*	Unconditional Prob (rrb.)	GBP/EUR	1/25	0.1356
Conditional Prob	USD/GBP	2/15	0.0862*	Unconditional Prob	JPY/EUR	1/5	0.1367
SUB SAMPLE 6				SUB SAMPLE 7			
ENHANCEMENT	CURRENCY	SR/LR	White P-VALUE	ENHANCEMENT	CURRENCY	SR/LR	White P-VALUE
Standard	JPY/NOK	5/30	0.0734*	Unconditional Prob	EUR/NZD	3/30	0.0823*
Standard	JPY/NOK	5/15	0.0736*	Unconditional Prob	USD/AUD	1/5	0.0839*
Standard	JPY/SEK	15/20	0.0823*	Unconditional Prob	EUR/CHF	1/5	0.0887*
Standard	JPY/NOK	10/20	0.0830*	Unconditional Prob	CHF/NOK	1/5	0.1023
Standard	JPY/NOK	4/30	0.0863*	Unconditional Prob	EUR/NOK	1/5	0.1148
Unconditional Prob (rrb.)	JPY/SEK	15/20	0.0880*	Unconditional Prob	EUR/SEK	1/5	0.1197
Standard	JPY/SEK	5/15	0.0893*	Conditional Prob	CHF/NOK	1/5	0.1316
Standard	GBP/EUR	20/25	0.0912*	Unconditional Prob	EUR/CHF	1/15	0.1323
Conditional Prob	JPY/SEK	15/20	0.0924*	Unconditional Prob	USD/NOK	1/5	0.1344
Conditional Prob (rrb.)	JPY/SEK	15/20	0.0925*	Unconditional Prob	USD/AUD	3/25	0.1351
SUB SAMPLE 8				SUB SAMPLE 9			
ENHANCEMENT	CURRENCY	SR/LR	White P-VALUE	ENHANCEMENT	CURRENCY	SR/LR	White P-VALUE
Unconditional Prob	GBP/NOK	1/20	0.0716*	Unconditional Prob	CHF/NOK	10/15	0.0555*
Unconditional Prob	GBP/NOK	3/20	0.0762*	Unconditional Prob	JPY/NZD	3/10	0.0872*
Unconditional Prob	GBP/NOK	1/15	0.1060	Unconditional Prob	USD/AUD	3/30	0.0909*
Conditional Prob	NOK/AUD	1/10	0.1080	Unconditional Prob (rrb.)	CHF/NOK	10/15	0.0957*
Unconditional Prob	GBP/NOK	1/25	0.1085	Conditional Prob	GBP/JPY	15/25	0.0976*
Conditional Prob	USD/GBP	5/10	0.1101	Standard	CHF/NOK	2/5	0.0985*
Unconditional Prob	GBP/NOK	4/20	0.1127	Unconditional Prob (rrb.)	GBP/JPY	15/25	0.0987*
Unconditional Prob (rrb.)	NOK/AUD	1/10	0.1134	Conditional Prob (rrb.)	GBP/JPY	15/25	0.0987*
Unconditional Prob	NOK/CAD	5/10	0.1164	Conditional Prob	CHF/NOK	10/15	0.1006
Unconditional Prob	CHF/AUD	2/10	0.1172	Unconditional Prob	USD/AUD	4/30	0.1023

The first column in each sub-sample box specifies the enhancement, i.e. weighting according to conditional or unconditional survival probability. The second and third columns show the currency pair and the parameterisation of the trading rule. The last column of each sub-sample box shows White's P-value. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

This is due to the fact that the motivation behind this chapter is somewhat different from the motivation behind the cited previous studies. While earlier studies aim to find a small number of the top trading rules out of a large universe of trading rules, this study, as pointed out before, aims to assess whether incrementally enhancing simple trading rules leads to trading rule results that are statistically superior to the results obtained by a standard trading rule. Therefore, a more homogeneous universe of trading rules is arguably preferable to a broader universe. Another aspect that becomes evident from Figure 2-3 is the fact that there seems to be somewhat of a clustering of trading rule parameters as well as currency pairs across the different sub-samples.

The best ten data snooping proven trading rules of the fifth sub-sample are either SR/LR 1/5, 1/10 or 1/25 trading rules. The seventh sub-sample shows similarly concentrated results. Moreover, the best data snooping proven trading rules of the third sub-sample consists almost exclusively of GBP/EUR or USD/EUR crosses. When it comes to the four enhancements of the trading rules, the picture is fairly mixed. Figure 2-3 indicates that in most of the earlier sub-samples there is an even split between the different trading rules. Sub-sample six on the other hand suggests a clear dominance of the standard trading rule over the enhanced trading rules. In the sub-samples thereafter, strategy number two, which weighs observations according to their unconditional probability, is the enhancement that features most often amongst the top ten trading rules. This gives rise to the question whether it is mere coincidence that strategy two features so well in sub-samples 7, 8 and 9, or whether this pattern is a broader trend that appears across the entire spectrum of sub-samples.

This begs the question to which extent the selection of data snooping proven trading rules is a function of the specificities of the different sub-samples, as opposed to mere coincidence. To assess whether there are any biases that come as a result of the choice of sub-samples, the next two 2s show the degree to which single trading rule parameterisations or single currency pairs are picked amongst the decile of best trading rules.

Figure 2-4 shows this analysis for trading rule parameterisations. For each of the sub-samples the top 10% (i.e. 877) trading rules that exhibit the highest White's P-value are used as the "adjusted" universe of trading rules. From this "adjusted" universe, the percentage occurrence of trading rules with different parameterisations is then calculated.

FIGURE 2-4: TRADING RULE PARAMETERISATION AMONGST TOP 10% OF TRADING RULES

SUB SAMPLE 2						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	2.1%	2.2%	3.3%	1.9%	2.1%	2.5%
SR 2	0.1%	2.9%	3.4%	3.4%	3.2%	2.5%
SR 3	1.0%	3.1%	3.5%	1.8%	3.3%	3.7%
SR 4	0.5%	2.3%	3.2%	2.3%	2.4%	3.8%
SR 5		1.6%	3.5%	3.2%	3.0%	3.1%
SR 10			1.4%	3.5%	2.6%	3.8%
SR 15				1.5%	2.1%	3.0%
SR 20					2.4%	1.9%
SR 25						3.1%

SUB SAMPLE 3						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	3.3%	4.1%	3.7%	3.3%	1.6%	1.3%
SR 2	3.3%	3.3%	3.3%	2.4%	2.2%	0.6%
SR 3	2.9%	4.7%	4.5%	1.7%	2.6%	2.6%
SR 4	3.3%	4.1%	1.9%	2.5%	3.7%	2.3%
SR 5		1.7%	1.9%	1.8%	4.8%	3.7%
SR 10			1.8%	2.2%	1.1%	1.8%
SR 15				1.5%	1.1%	1.8%
SR 20					2.5%	2.1%
SR 25						1.0%

SUB SAMPLE 4						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	0.0%	0.2%	1.0%	1.1%	2.1%	2.4%
SR 2	0.0%	1.7%	2.4%	1.8%	3.4%	4.6%
SR 3	0.0%	1.8%	2.2%	3.1%	3.1%	3.2%
SR 4	0.0%	1.3%	1.9%	4.1%	4.1%	4.7%
SR 5		1.9%	1.8%	5.1%	4.7%	4.7%
SR 10			1.7%	4.6%	5.5%	3.4%
SR 15				5.1%	4.2%	3.2%
SR 20					0.7%	1.4%
SR 25						1.7%

SUB SAMPLE 5						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	5.5%	4.1%	5.0%	6.3%	7.8%	5.0%
SR 2	7.1%	3.9%	4.7%	4.6%	3.7%	2.7%
SR 3	6.6%	2.7%	3.9%	3.0%	3.8%	2.4%
SR 4	4.9%	2.6%	4.5%	2.1%	1.4%	1.9%
SR 5		0.0%	0.0%	0.0%	0.0%	0.0%
SR 10			0.0%	0.0%	0.0%	0.0%
SR 15				0.0%	0.0%	0.0%
SR 20					0.0%	0.0%
SR 25						0.0%

SUB SAMPLE 6						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	2.2%	0.9%	2.4%	3.0%	2.9%	3.3%
SR 2	0.9%	0.0%	1.4%	2.1%	2.6%	3.3%
SR 3	0.0%	1.1%	2.1%	1.9%	2.3%	2.6%
SR 4	0.6%	0.9%	1.6%	3.3%	2.2%	2.4%
SR 5		1.0%	3.9%	3.7%	3.1%	3.5%
SR 10			2.4%	2.9%	2.4%	2.9%
SR 15				2.3%	5.1%	5.1%
SR 20					6.6%	5.1%
SR 25						4.1%

SUB SAMPLE 7						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	9.0%	7.1%	7.1%	3.9%	4.0%	3.2%
SR 2	2.5%	1.9%	2.3%	1.6%	1.4%	1.5%
SR 3	3.8%	1.8%	1.5%	1.9%	1.6%	0.9%
SR 4	4.6%	2.1%	1.1%	1.4%	2.1%	2.1%
SR 5		1.8%	0.7%	1.8%	3.5%	1.6%
SR 10			2.7%	2.1%	1.3%	1.4%
SR 15				1.3%	2.4%	1.0%
SR 20					0.8%	3.0%
SR 25						4.5%

SUB SAMPLE 8						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	3.0%	5.0%	6.1%	3.7%	3.2%	3.9%
SR 2	2.3%	5.3%	2.3%	2.3%	3.3%	2.2%
SR 3	1.4%	4.0%	2.2%	2.7%	2.6%	2.5%
SR 4	1.7%	4.9%	2.2%	3.0%	2.7%	2.4%
SR 5		5.3%	2.6%	3.8%	2.1%	2.1%
SR 10			1.7%	0.7%	1.1%	0.7%
SR 15				2.2%	0.2%	0.5%
SR 20					1.0%	1.1%
SR 25						0.3%


SUB SAMPLE 9						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	5.8%	1.8%	2.4%	2.2%	1.9%	3.4%
SR 2	2.4%	1.3%	1.1%	1.4%	1.3%	2.2%
SR 3	2.5%	0.8%	1.3%	0.6%	1.6%	2.2%
SR 4	3.1%	0.7%	0.7%	0.8%	1.5%	2.5%
SR 5		0.7%	1.5%	1.0%	2.1%	3.9%
SR 10			3.8%	3.7%	5.3%	5.8%
SR 15				3.3%	4.8%	5.9%
SR 20					6.4%	3.8%
SR 25						2.9%

The Figure shows the percentage occurrence of single trading rule parameterisations amongst the decile of trading rules with the highest White's P-Value. Each zero weight is bold and any weight below 2.56% is white. Values between 2.56% to 5% are light pink and values between 5% to 7.5% and then 7.5% to 10% are gradually darker shades of pink.

Given the fact that this study uses 39 different moving average combinations, it is to be expected that under a normal distribution of trading rule results approximately 23 (i.e. 2.56%) observations of each trading rule parameterisation show up amongst the top 10% of trading rules. An occurrence of more than 2.56% would indicate that a specific moving average combination tends to achieve higher White's P-values than others.

Figure 2-4 is coloured in such way that each zero weight is bold and any weight below 2.56% is white. Values between 2.56% to 5% are light pink; values between 5% and 7.5% and then 7.5% to 10% are gradually darker shades of pink. The results in Figure 2-4 would suggest that there is no sub-sample specific bias when it comes to different moving average combinations. The only sub-samples where there are signs of a slight bias would be sub-sample 4 and sub-sample 5. In sub-sample 4, none of the shorter-term moving averages appears in the top decile of trading rules. In sub-sample 5, none of the longer-term moving averages appears in the top decile of trading rules. Apart from these two sub-samples the occurrence of moving average combinations is fairly evenly split, with a slight overweight of the shorter-term moving averages.

FIGURE 2-5: OCCURRENCE OF TRADING RULE ENHANCEMENTS AMONGST TOP 10% OF TRADING RULES



The Figure shows the percentage split of the standard as well as the four enhanced trading rules amongst the decile of trading rules that have the highest White's P-Values.

Figure 2-5 shows a graph of the evolution of the five single trading strategies over time. It calculates the percentage split of the standard as well as the four enhanced trading rules, within the 10% of best data snooping proven trading rules, i.e. the decile of

trading rules that have the highest White's P-Value. Figure 2-5 suggests that during the first five sub-samples, there is a fairly even split across trading rule enhancements and all of the five trading rules are close to the 20% level, as indicated by the black dotted line. In sub-sample 6 the standard trading rule is the dominant trading rule, while trading rule enhancement three occurs only half as often as what would be expected under normality. However, thereafter, trading rule enhancement three, which weighs each trading signal according to its expected historic unconditional survival probability, occurs more than 30% of the time. This significant jump in the occurrence of the third enhanced trading rule suggests that there is somewhat of a shift within the behaviour of currencies over time. Therefore one could argue that Figure 2-5 does provide some indication that it might be beneficial to apply the White's "data snooper" not only to the best absolute but also to the "best relative" trading rules. Hence, looking at the standard, as well as the four enhanced trading rules relative to each other.

Moreover Figure 2-6, which shows the same analysis as given in Figure 2-4, but on the basis of currency pairs as opposed to parameterisations, provides further evidence that analysing different trading rules in a relative as opposed to an absolute perspective might lead to more stable overall results. Figure 2-6 also uses the top decile (i.e. 877) trading rules that exhibit the highest White's P-value are used as the "adjusted" universe of trading rules. The overall number of currency pairs analysed is 45. Therefore under the assumption of a normal distribution, approximately 19.5 observations (i.e. 2.2%) of each currency pair are expected to appear in the top decile of data snooping proven trading rules. The colour code of Figure 2-6 follows the same principle of the colour code in Figure 2-4. Any zero weight is bold, any weight below 2.2% is white and any value above 2.2% is shown in pink, whereby the shades of the pink get darker as the percentage increases. The results in Figure 2-6 suggest that there is a fairly high degree of concentration of single currency pairs amongst the top decile of trading rules. Taking sub-sample 6 as an example, approximately 33% or 290 of the 877 top trading rules are concentrated in the JPY/NOK and the JPY/SEK currency cross. This is about 7.5 times the amount of trading rules that would be expected under the assumption of normality. Moreover, this concentration fluctuates significantly across sub-samples. In the second sub-sample 13% of observations are concentrated in the GBP/CHF cross for any sub-sample thereafter the concentration does not exceed 2.1%, which is in line with what is expected normally. The pattern in Figure 2-6 would also suggest that there tends to be a higher concentration in currency pairs that include the four big liquid currencies USD, GBP, JPY and EUR. Given this high concentration of results when it comes to these

currency pairs, it is fair to conclude that the set of trading rules that are identified as the ten best trading rules is likely to be very time dependent. Hence, one could argue that this kind of test specification is not the best method of analysing the dataset.

FIGURE 2-6: CURRENCY PAIRS AMONGST TOP 10% OF TRADING RULES

SUB SAMPLE 2									
	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD	6.6%	12.6%	2.2%	6.2%	3.4%	0.0%	3.7%	2.9%	0.0%
GBP		0.0%	11.3%	13.0%	0.9%	0.0%	0.2%	0.0%	0.0%
JPY			16.3%	6.1%	3.0%	0.0%	4.7%	0.0%	0.0%
EUR				0.0%	0.2%	0.5%	1.6%	0.0%	0.0%
CHF					0.6%	0.0%	0.6%	0.0%	0.0%
NOK						0.1%	0.1%	0.0%	0.0%
SEK							0.7%	0.0%	0.0%
CAD								2.7%	0.0%
AUD									0.0%

SUB SAMPLE 3									
	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD	5.0%	3.0%	12.0%	6.0%	4.0%	2.0%	0.0%	6.0%	3.0%
GBP		2.0%	14.0%	1.0%	0.0%	0.0%	0.0%	1.0%	0.0%
JPY			8.0%	0.0%	0.0%	1.0%	3.0%	6.0%	0.0%
EUR				1.0%	0.0%	1.0%	7.0%	0.0%	0.0%
CHF					0.0%	0.0%	2.0%	0.0%	0.0%
NOK						0.0%	2.0%	1.0%	0.0%
SEK							1.0%	3.0%	0.0%
CAD								3.0%	0.0%
AUD									0.0%

SUB SAMPLE 4									
	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD	12.2%	10.2%	7.9%	4.2%	6.1%	7.9%	0.6%	0.0%	0.0%
GBP		0.0%	0.6%	0.0%	0.5%	0.1%	6.2%	0.1%	0.0%
JPY			8.6%	0.2%	9.7%	5.1%	6.5%	0.0%	0.0%
EUR				0.0%	1.0%	0.6%	7.5%	0.0%	0.0%
CHF					0.0%	0.0%	1.8%	0.5%	0.0%
NOK						0.0%	0.1%	0.0%	0.0%
SEK							1.7%	0.0%	0.1%
CAD								0.0%	0.0%
AUD									0.0%

SUB SAMPLE 5									
	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD	7.2%	0.0%	3.2%	0.9%	2.3%	0.6%	0.2%	0.0%	0.0%
GBP		9.4%	7.6%	0.1%	2.1%	0.0%	7.2%	8.9%	1.3%
JPY			7.3%	0.8%	2.9%	0.5%	0.5%	0.0%	0.0%
EUR				0.0%	0.0%	0.0%	10.4%	4.3%	0.8%
CHF					0.0%	0.0%	10.3%	1.4%	0.2%
NOK						0.0%	5.6%	2.6%	0.0%
SEK							0.7%	0.6%	0.3%
CAD								0.0%	0.0%
AUD									0.0%

SUB SAMPLE 6									
	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD	0.1%	11.4%	0.1%	0.6%	0.0%	0.9%	5.4%	0.0%	1.1%
GBP		10.2%	2.5%	2.1%	0.6%	0.6%	0.1%	2.4%	0.0%
JPY			9.4%	4.0%	16.7%	16.3%	0.1%	0.3%	0.6%
EUR				0.7%	0.9%	0.8%	0.0%	1.6%	0.0%
CHF					1.8%	0.1%	0.0%	1.5%	0.1%
NOK						0.1%	0.0%	2.4%	0.0%
SEK							0.1%	1.6%	1.7%
CAD								0.0%	0.0%
AUD									1.1%

SUB SAMPLE 7									
	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD	0.6%	5.6%	4.8%	2.3%	0.7%	1.7%	0.1%	12.8%	5.1%
GBP		1.2%	1.5%	0.3%	0.5%	0.1%	0.0%	1.6%	0.9%
JPY			2.3%	1.2%	2.9%	1.2%	1.3%	2.1%	3.8%
EUR				2.3%	3.3%	1.2%	4.9%	4.1%	1.7%
CHF					3.8%	1.6%	2.8%	7.6%	1.4%
NOK						2.9%	0.1%	0.9%	0.0%
SEK							0.9%	0.6%	0.0%
CAD								0.3%	0.0%
AUD									5.1%

SUB SAMPLE 8									
	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD	6.5%	0.5%	1.5%	1.4%	3.5%	12.1%	0.3%	1.3%	0.9%
GBP		1.0%	0.6%	0.2%	11.3%	4.1%	1.9%	1.7%	0.5%
JPY			1.8%	0.0%	0.9%	0.5%	0.0%	0.2%	0.2%
EUR				0.9%	7.7%	0.1%	1.4%	3.0%	0.6%
CHF					0.3%	0.1%	1.7%	4.0%	0.0%
NOK						1.4%	3.1%	3.9%	0.0%
SEK							4.7%	0.1%	1.7%
CAD								0.0%	0.0%
AUD									12.3%

SUB SAMPLE 9									
	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD	1.5%	3.1%	10.1%	1.0%	1.1%	1.1%	2.2%	7.5%	6.5%
GBP		6.4%	2.6%	1.6%	2.2%	0.1%	0.2%	5.5%	0.8%
JPY			1.5%	1.0%	0.9%	0.5%	1.0%	6.2%	4.7%
EUR				0.9%	9.8%	0.8%	1.3%	0.5%	0.9%
CHF					3.2%	1.7%	0.5%	0.1%	1.1%
NOK						0.3%	0.0%	0.0%	0.0%
SEK							0.0%	0.1%	0.6%
CAD								3.5%	1.8%
AUD									3.4%

The Figure shows the percentage occurrence of single currency pairs amongst the decile of trading rules with the highest White's P-Value. Any zero weight is bold, any weight below 2.2% is white and any value above 2.2% is shown in pink, whereby the shades of the pink get darker as the percentage increases.

This is partially confirmed when looking at Figure 2-7, which shows the results the back test of the search for the ten best trading rules from an absolute perspective. As mentioned earlier the back test of the absolute specification of the data snooping proven trading rules is conducted as follows. The ten best data snooping proven trading rules that have been identified for each sub-sample, as given in Figure 2-3, are combined in

an equally weighted composite trading rule. This enhanced composite trading rule is then compared to a benchmark composite trading rule. The components of the data snooping proven trading rule will vary depending on sub-samples. The benchmark composite trading rule is constructed in the same way as the enhanced composite trading rule. However, it is equally weighting all moving average combinations of SR (1, 2, 3, 4, 5, 10, 15, 20, 25) and LR (5, 10, 15, 20, 25, 30), across all currencies.

FIGURE 2-7: BVTC; BEST DATA SNOOPING PROVEN TRADING RULES (ABSOLUTE EVALUATION)

Sub Sample by Sub Sample breakeven transaction cost levels							
ABSOLUT	SS3	SS4	SS5	SS6	SS7	SS8	SS9
OPT SS2	11.98	14.73	-0.12	11.92	-1.06	1.83	0.25
OPT SS3		1.41	2.94	0.24	2.25	-1.28	1.81
OPT SS4			10.11	2.90	-5.26	4.47	-1.47
OPT SS5				-0.68	0.05	1.99	2.89
OPT SS6					-34.66	0.68	-6.05
OPT SS7						1.51	0.65
OPT SS8							-5.76
BM	5.63	2.88	-0.51	2.43	-2.95	-3.75	-4.07
RELATIVE	SS3	SS4	SS5	SS6	SS7	SS8	SS9
OPT SS2	6.35	11.85	0.40	9.49	1.89	5.58	4.32
OPT SS3		-1.47	3.45	-2.20	5.20	2.47	5.88
OPT SS4			10.62	0.47	-2.31	8.22	2.60
OPT SS5				-3.11	3.00	5.74	6.96
OPT SS6					-31.71	4.43	-1.98
OPT SS7						5.26	4.71
OPT SS8							-1.69

The Figure shows a backtest of the breakeven transaction cost levels of investment strategies with highest (absolute) White's P-value.. The first section shows the absolute breakeven transaction cost levels that are calculated sub-sample by sub-sample for all data snooping proven trading rules. It also shows the results of the benchmark trading rule in bold. The first line of the first section represents the back test of the top ten trading rules from the second sub-sample The second line of the first section gives the results of the optimised trading rule of sub-sample three, and so on. The second section shows the difference between both.

To avoid data circularities, the backtest analysis spans from the third sub-sample to the last sub-sample. The first two sub-samples are used to estimate the survivorship curves and to identify the best ten data snooping proven trading rules. The first section shows the absolute breakeven transaction cost levels that are calculated sub-sample by sub-sample for all the data snooping proven trading rules. It also shows the results of the benchmark trading rule in bold. The difference of the benchmark trading rule results of this chapter versus the first chapter stems from the fact that the trading strategy in this chapter only gradually phases into trading positions over a time period of 24 hours in

order to minimise market impact, while the trading strategy in the first chapter assumes positions to be established immediately, without considering aspects of market impact.

The first line of the first section represents the back test of the top ten trading rules from the second sub-sample. The second line of the first section gives the results of the optimised trading rule of sub-sample three, and so on. The second section shows the difference between both. While the optimised trading rules from sub-samples two to five deliver results that are fairly stable over time, the optimised trading rules in the sub-samples thereafter seem rather volatile. Namely the data snooping proven top trading rules identified in sub-sample six deliver very weak results in the subsequent sub-samples. Overall one can conclude that the results of the absolute data snooping analysis do not offer too much room for interpretation. The fact that there is a great deal of clustering of currency pairs amongst the highest ranking trading rules, and the fact that the second trading rule enhancement becomes the dominant, superior trading rule in the later sub-samples, gives scope to analyse the data obtained in the data snooper in a different way.


2. Best Relative Data Snooping Proven Trading Rules

The focus of the previous analysis is to find the ten trading rules that show the highest White's P-Values, without any initial conditioning of the universe of P-Values. Given the high level of parameterisation, currency and enhancement volatility within the ten best trading rules over time, a simple conditioning is implemented. This facilitates the analysis of the frequency of single currency pairs, or trading rule parameterisations within the top decile of trading rules over time. The key conclusion of this adjustment is the fact that any absolute analysis of data snooping proven trading rules within the foreign exchange space is subject to biases due to a great deal of clustering of currency pairs amongst the top trading rules over time. As a consequence of that this section goes further in the degree of conditioning of the universe of White's P-values. The aim of this is to obtain more granular results that help explaining the dynamics of trading rules over time. The analysis in this section divides the universe of White's P-Values in five ways, according to the P-Values obtained by the standard trading rule and the four enhancements. Within each sub-universe the P-Values are then split according to their parameterisations. As discussed in the previous section, the aim of this analysis is to smooth out erratic effects from currency variation. Hence, the average P-value across all

currency pairs is taken. Given the fact that the P-values have upper and lower boundaries, the mean and median averages of the distribution will be very close to each other. Therefore using the mean is preferable over the median as it represents a very parsimonious measure. The means of each of the parameterisations of the enhanced trading strategies are then compared to the mean P-Values of the standard trading rule. Figures 2-8 a, b, c show the differences in means between the average P-Values (across all currency pairs) of the enhanced trading rules and the standard trading rule, ordered by parameterisations.


When it comes to testing the statistical significance of the cross sectional differences in Whites P-values, a series of aspects have to be considered. An optimal test setup is non-parametric, with no assumption about the distribution of the underlying data. Moreover, the fact that the different trading rule enhancements are very similar in their construction opens the possibility that White's P-values derived from the different enhancements are not fully independent from the P-values derived from the standard trading rule. To allow for those two aspects, Wilcoxon's signed rank test is applied in figures 2-8 a, b, c. This test is designed to assess the distributional differences of two sets of data that are not independent from each other. The test is rank based. The underlying intuition can be described as follows; the differences between the distributions of P-values between the standard and the enhanced trading rule are ranked. If both trading rules are similar the ranking of the differences between the two trading rules should be evenly distributed. If this is not the case, then the tests indicates that the difference between trading rules is statistically significant. The colour codes in figures 2-8 a, b, c do not indicate levels of statistical significance. They merely visualise the percentiles in differences between enhanced and standard trading rules. The colour scheme in the Figures is organised in such way that the top decile of trading rule parameterisations have the darkest blue shade, while bottom decile has the darkest red shade.

FIGURE 2-8A: DIFFERENCES IN MEAN P-VALUES OF THE ENHANCED AND STANDARD TRADING RULE




The Figure shows the differences in means between average P-Values (across all currency pairs) of the enhanced standard trading rule across all currency pairs. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%. The top percentile has the darkest blue shade, while bottom percentile has the darkest red shade. The split of colour shades between the bottom and top percentile is given in the legend of each Figure. Each of the sub samples shows ten values in boxes, which represent the trading rule parameterisations with the highest difference in average P-Values relative to the standard trading rule.

FIGURE 2-8B: DIFFERENCES IN MEAN P-VALUES OF THE ENHANCED AND STANDARD TRADING RULE



The Figure shows the differences in means between average P-Values (across all currency pairs) of the enhanced standard trading rule across all currency pairs. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%. The top percentile has the darkest blue shade, while bottom percentile has the darkest red shade. The split of colour shades between the bottom and top percentile is given in the legend of each Figure. Each of the sub samples shows ten values in boxes, which represent the trading rule parameterisations with the highest difference in average P-Values relative to the standard trading rule.

FIGURE 2-8C: DIFFERENCES IN MEAN P-VALUES OF THE ENHANCED AND STANDARD TRADING RULE



The Figure shows the differences in means between average P-Values (across all currency pairs) of the enhanced standard trading rule across all currency pairs. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%. The top percentile has the darkest blue shade, while bottom percentile has the darkest red shade. The split of colour shades between the bottom and top percentile is given in the legend of each Figure. Each of the sub samples shows ten values in boxes, which represent the trading rule parameterisations with the highest difference in average P-Values relative to the standard trading rule.

The split of colour shades between the bottom and top deciles is given in the legend of each figure. The colour scheme is applied to each sub-sample individually. Therefore the colour associated with each absolute number may vary over time. Each of the sub samples shows ten values in boxes, which represent the trading rule parameterisations with the highest difference in average P-Values relative to the standard trading rule. These ten trading rules are later used as input for the trading rule backtest.

Looking through the results sub-sample by sub-sample, a series of observation can be made. Firstly, a high number of the differences are statistically significant. Secondly, the results suggest that the observation period can be split into two distinct regimes:

sub-samples two to six and sub-samples seven, eight and nine. The first regime is characterised by equal or lower average P-values of the enhanced trading rules versus the standard trading rule.

Moreover only, while there is a fair amount of statistically significant differences between the enhanced and the standard trading rules. In the first six sub samples there are comparatively few statistically significant positive differences in average P-values between enhanced and standard trading rules, while there are many negative differences that are statistically significant. The second regime on the other hand shows that enhanced trading rules exhibit much higher average P-Values than the values obtained in the standard trading rule, most of the differences are statistically significant. Moreover, during the first regime the top trading rules are scattered around all trading rule enhancements, while in the second regime the top trading rules can only be found in the second enhancement. Another observation that can be made is the fact that over time most of the statistically significant differences in average P-Values, may they be positive or negative, can be found at the second trading rule enhancement. Overall, the results of the relative trading rule analysis reconfirm the results from the absolute trading rule analysis. It also provides further insight into the dynamics of the enhanced trading rules. While conditioning the trading rule results of the absolute search for the best data proven merely shows that a disproportionately high amount of second trading rule enhancements appear amongst the top decile of trading rules during sub-samples seven, eight and nine, the relative analysis also suggests that the second trading rule enhancement is a relative underperformer before that. During the first regime the second trading rule enhancement shows persistently lower average P-values than the standard trading rule. In some sub-samples, such as the second, fifth or sixth the difference is as much as 10% and statistically significant. The results of the analysis of relative P-values give clear indication of two distinct regimes within the dataset. In the first regime there seems to be very little scope to add value via the application of the presented enhanced trading rules. However, in the second regime some of the incremental enhancements to simple moving average trading rules, presented in this chapter, produce superior results.

To verify the quality of this assessment, the chapter applies a trading rule backtest similar to the one applied for the absolute best trading rules. The trading rule backtest in the previous section is designed in such way that the best ten trading rules are compared to a composite trading rule that encompasses all 39 standard trading rules. This specification is not appropriate for the relative trading rule, as it would compare the

universe of all parameterisations of the standard trading rule to a small sub-set of parameterisations of enhanced trading rules. Hence, it would not be possible to distinguish the performance effect from the mismatch in parameterisation versus the performance effect that comes from the trading rule enhancement itself. Therefore the benchmark for the back-test of the relative trading rule consists of ten standard trading rules that have the same parameterisation as the top ten enhanced trading rules. As mentioned earlier Figures 2-8 a, b, c show for each of the sub-samples analysed values in boxes, which represent the trading rule parameterisations with the highest difference in average P-Values relative to the standard trading rule. For each of the sub-samples these ten trading rules are combined to a composite trading rule as described in an earlier section. For each of the sub-samples these composite trading rules are then compared to composite benchmark trading rules, which are composed of standard trading rules with exactly the same parameterisations as the top ten enhanced trading rules. If there is an overlap in parameterisations amongst the enhanced trading rules, then the benchmark trading rule counts the overlapping parameterisations as of then as they occur. Sub-sample 6 gives a good example of such overlap. Trading rule parameterisations SR2/LR5 and SR1/LR5 are counted three and two times in the benchmark trading rule, all the other five parameterisations are counted once. The results of this analysis are given in Figure 2-9. The Figure is similarly structured to Figure 2-7. It shows the back test of the top ten trading rules on a sub sample by sub sample basis. To avoid data circularities, the back-test analysis spans from the third sub-sample to the last sub-sample. The first two sub-samples are used to estimate the survivorship curves and to identify the best ten data snooping proven trading rules. The first section shows the absolute breakeven transaction cost levels that are calculated sub-sample by sub-sample for all the data snooping proven trading rules. It also shows the results of the benchmark trading rules in bold. The first line of the first section represents the back-test of the top ten trading rules from the second sub-sample. The second line of the first section shows the results of the respective benchmark. The third line gives the results of the optimised trading rule of sub-sample three, and so on. The second section shows the difference between both.

The results confirm the trend described previously. There are two distinct regimes in the dataset. The differences between breakeven transaction cost levels of the enhanced and the benchmark trading rules remain very low during the time period that is earlier described as the first regime. In the second regime, however, these differences are considerably higher. The first section of Figure 2-9 indicates that the difference between

standard and enhanced strategy has increased persistently over time. The second section of Figure 2-9 gives a more detailed insight in the dynamics of the trading rule. This figure indicates that the overall trading rule results are less volatile than the results from the absolute trading rule analysis given in Figure 2-7.

FIGURE 2-9: BVTC; BEST DATA SNOOPING PROVEN TRADING RULES (RELATIVE EVALUATION)

Sub Sample by Sub Sample breakeven transaction cost levels							
ABSOLUT	SS3	SS4	SS5	SS6	SS7	SS8	SS9
OPT SS2	6.23	1.94	-2.25	0.47	-4.79	2.32	-6.64
BM SS2	5.64	1.49	-2.59	0.72	-5.13	2.20	-7.21
OPT SS3		11.53	5.09	17.14	-11.77	-5.29	-0.64
BM SS3		12.18	5.35	19.61	-12.54	-6.06	0.53
OPT SS4			1.71	4.65	-5.44	-1.11	0.62
BM SS4			2.81	9.40	-7.83	-2.55	-0.81
OPT SS5				0.45	-1.15	0.91	-1.85
BM SS5				0.00	-1.60	0.32	-3.08
OPT SS6					-0.33	0.76	-1.58
BM SS6					-0.86	0.28	-2.82
OPT SS7						1.76	-1.79
BM SS7						1.33	-5.31
OPT SS8							-0.99
BM SS8							-4.64
RELATIVE	SS3	SS4	SS5	SS6	SS7	SS8	SS9
OPT SS2	0.60	0.45	0.34	-0.25	0.34	0.12	0.58
OPT SS3		-0.65	-0.26	-2.47	0.77	0.77	-1.17
OPT SS4			-1.10	-4.75	2.39	1.44	1.43
OPT SS5				0.45	0.45	0.58	1.22
OPT SS6					0.54	0.48	1.23
OPT SS7						0.43	3.52
OPT SS8							3.65

The Figure shows a backtest of the breakeven transaction cost levels of investment strategies with highest (relative) White's P-value, as shown in Figure 8a,b,c. The first section shows the absolute breakeven transaction cost levels that are calculated sub-sample by sub-sample for all the data snooping proven trading rules. The first line of the first section represents the backtest of the top ten trading rules from the second sub-sample. The second line of the first section shows the results of the respective benchmark (in bold). The third line gives the results of the optimised trading rule of sub-sample three, and so on. The second section shows the difference between both.

Nonetheless, the sub-sample per sub-sample breakeven transaction cost levels show a fair degree of volatility. The fact that the differences in White's P-values between standard and enhanced trading rules become highly significant in the second regime is only partially reflected in the analysis presented in Figure 2-9. As shown earlier, sub-samples 7, 8 and 9 indicate that the second enhanced trading rule exhibits average P-values that are significantly higher than the P-values of the standard trading rule. This

would suggest that the enhanced trading rules established in these sub-samples should perform very well. For two of the three sub-samples in which the outperformance of the second trading rule enhancement has been recorded, a trading rule back test is implemented. The results of the back tests shown in Figure 2-9 are the results of the optimised trading rule of sub-sample 7 and 8. The difference between enhanced and standard trading rules is positive in both cases. But they are only strongly positive for the ninth sub-sample. While one can clearly grasp a trend, this is by far not sufficient to verify a link between the differences in White's P-values and subsequent trading rule performance.

FIGURE 2-10: BVTC; FIRST AND SECOND ENHANCEMENT VERSUS BENCHMARK

Sub Sample by Sub Sample breakeven transaction cost levels							
ABSOLUT	SS3	SS4	SS5	SS6	SS7	SS8	SS9
OPT SS1	6.56	3.71	-0.20	2.03	-2.84	-3.48	-3.79
OPT SS2	3.36	-0.09	-1.01	-0.32	0.51	-1.05	-1.17
BM	5.63	2.88	-0.51	2.43	-2.95	-3.75	-4.07
RELATIVE	SS3	SS4	SS5	SS6	SS7	SS8	SS9
OPT SS1	0.93	0.83	0.31	-0.40	0.11	0.27	0.28
OPT SS2	-2.27	-2.97	-0.50	-2.75	3.46	2.70	2.90

Both the enhanced and the standard composite trading rule results shown in the Figure combine all 39 trading rule parameterisations, SR (1, 2, 3, 4, 5, 10, 15, 20, 25) and LR (5, 10, 15, 20, 25, 30), across all currencies. The first section shows the absolute breakeven transaction cost levels that are calculated sub-sample by sub-sample for the first and second trading rule enhancement. It also shows the results of the benchmark trading rule in bold. The second section shows the difference between the enhanced trading rules and the benchmark trading rule.

Given the fact that the second trading rule enhancement exhibits significantly lower average P-values in many of the sub-samples of the first regime, and much stronger results in the second regime, it is worthwhile to analyse the second trading rule enhancement further. The results of this analysis are shown in Figure 2-10. Both the enhanced and the standard composite trading rule results shown in the Figure combine all 39 trading rule parameterisations, SR (1, 2, 3, 4, 5, 10, 15, 20, 25) and LR (5, 10, 15, 20, 25, 30), across all currencies. For comparison reasons, Figure 2-10 shows the results for the first trading rule enhancement as well. The other two enhancements are ignored given the fact that they have not produced any meaningful results in White's P-Value analysis. As mentioned previously, to avoid data circularities, the backtest analysis spans from the third sub-sample to the last sub-sample. The first two sub-samples are used to estimate the survivorship curves and to identify the best ten data snooping

proven trading rules. The first section shows the absolute breakeven transaction cost levels that are calculated sub-sample by sub-sample. It also shows the results of the benchmark trading rule in bold. The second section shows the difference between both. The first trading rule enhancement shows small but persistently positive values. The exception is the sixth sub-sample, where the enhanced trading rule underperforms the standard trading rule. At a first glance the results of the back test of the first trading rule enhancement seem to be somewhat negligible. This is due to the fact that the results from the data snooper are not conclusive and that the magnitudes of the breakeven differentials are very small.

However, one has to put these results into the context of the overall level of transaction costs that are incurred when following such strategy. Neely (2011) estimates that since 2000, spot market participants have faced spreads of 2 basis points or less for transactions in the \$5 million to \$50 million range. Even conservatively assuming that forward or futures transactions for the same amounts cost 50% more than spot transactions, i.e. 3 basis points, one can still make a strong argument that applying the first trading rule enhancement would improve the end performance by almost 10% for sub-samples 8 and 9 and much more for the earlier sub-samples. Moreover, there seems to be very limited risk involved in applying the strategy, given its very persistent positive performance. Even in sub-sample 6, the time period when the first trading rule enhancement underperformed the benchmark trading rule, the underperformance was relatively mild.

When it comes to the second trading rule enhancement, Figure 2-10 provides clear evidence that the enhancement produces much weaker results during sub-samples three to six. Thereafter the enhancement produces stronger returns. Another observation that can be made from this analysis, is the fact that the benchmark trading strategy exhibits a persistent deterioration in profitability over time. Overall, it can be pointed out that the results from the back test reconfirm the observations made on the basis of the results of the relative analysis of whites P-values given in Figures 2-8 a, b, c and Figure 2-9. There is evidence of two distinct regimes within the data set. The back test also gives evidence that while overall trading rule profitability diminishes in the more recent time period, the first trading rule enhancement is not affected by this deterioration and the second trading rule enhancements performs better over time. The relative magnitude of that outperformance is considerable, when put into context of Neely's (2011) observations about transaction costs in foreign exchange markets. One could still make

the point that this strong relative outperformance of the second trading rule enhancement during the second regime might have something to do with the fact that the second trading rule has lower levels of total market exposure. Therefore it is likely to underperform during the first regime, when most of the benchmark trading rules exhibit positive breakeven transaction cost levels. It is also likely to outperform during the second regime when most of the benchmark trading rules show negative breakeven transaction cost levels. As pointed out in an earlier section the calculation of the breakeven transaction cost levels is done on a risk adjusted basis. Therefore, the results of the enhanced trading rule strategy, which is somewhat more defensive in its design than the benchmark trading rule, are fully comparable to the results of the benchmark trading rule. This observation is confirmed by the results of the first trading rule enhancement, which also exhibit lower exposure levels than the benchmark trading strategy. The results of the strategy are very much unrelated to the regime change exhibited in the data.

Both trading rule enhancements have some merit for practitioners. The persistency in performance of the first trading rule enhancement suggests that using conditional survival probabilities does incrementally improve the performance of more generic trading rules. Yet it adds very little incremental risk to the trading strategy. Hence, the use of this enhancement is likely to lead to small improvements in the Sharpe Ratio profiles of systematic trading rules.

The second trading rule enhancement has wider implications for practitioners given the higher magnitude of outperformance in recent sub-samples, but also given the higher level of volatility of that outperformance. The results in this section indicate a clear regime shift within foreign exchange markets. This regime shift has been in favour of shorter-term focussed trading rules. While simple trading strategies that were once profitable fail to deliver positive returns, the results presented in this section suggest that trading strategies, which are enhanced by applying unconditional survival probabilities to the exposure levels of returns, are able to add value versus a standard trading rule. However, access to the extra return derived from the second trading rule enhancement does come with a fair level of active risk taking. Hence, practitioners have to make a judgement about the future market environment in the foreign exchange space before following such trading rule enhancement. The fact that the profitability of the second trading rule enhancement has not yet deteriorated, would suggest that this

opportunity still exists. Moreover, the profound changes in foreign exchange markets would suggest that the opportunity might well persist for longer.

D. Conclusions

As pointed out earlier, Lo's (2004) Adaptive Market Hypothesis follows an evolutionary concept of market equilibrium as opposed to a steady state assumption. In that sense it incorporates a continuously changing market environment in which market inefficiencies arise, which are subsequently arbitrated away as market participant become aware of them. This suggests that the Adaptive Market Hypothesis is a more complex framework than Fama's (1969) Efficient Market Hypothesis. It captures the dynamics of market cycles in as far as it allows different assets to be subject to different levels of efficiency at the same time. Lo (2004) argues that asset prices are driven by the nature and preferences of market participants. He points out that there are different groups of investors, which have very distinct investment pattern and investment preferences. If one or many of these investment "species" find interest in one specific asset, the pricing of this asset becomes more efficient and vice versa. As a consequence of that, investment strategies can undergo different stages in which they show different levels of profitability. Therefore investors have to adapt to the changing market environment in order to achieve persistent levels of return.

One of the key conclusions that can be drawn from this study is the fact that the profitability of generic trading rules has continuously diminished over time. This would suggest that the excess returns that were available from undertaking systematic trading rules eroded as more investors participated in this kind of trading strategy. Indeed, since the late 1970s investment funds that manage client money by investing in trading strategies, so called "managed futures funds", grew exponentially. Initially these funds started in the commodity futures space. However, as client demand for these strategies grew, their investment universe expanded to currencies and other futures markets. Given the deep liquidity offered by the foreign exchange market, it should not come as a surprise that systematic trading within currencies currently represent one of the core areas of these trading funds. Other areas would be interest rates and equities. Barclay Hedge²¹, a database provider for systematic trading funds, estimates that in 1980 the

²¹ See: http://www.barclayhedge.com/research/indices/cta/mum/CTA_Fund_Industry.html

assets under management of systematic trading funds were in the range of US\$ 300m. This grew to US\$ 38bn by the end 2000. As of the end of 2009, when the data sample of this chapter ends, approximately US\$ 214bn were managed in systematic trading funds. Currently this number is close to US\$ 330bn. These numbers only represent the proportion of assets managed in explicit fiduciary mandates with clients. The actual amount of assets managed in systematic trading algorithms is likely to be much higher.

Since 1998 the Bank of International Settlements publishes a triennial survey of Foreign Exchange and Derivatives market activity. Besides this survey a series of working papers are published that shed light on the drivers in the change of trading volume. Since 2001 there has been a persistent increase in foreign exchange trading volumes. Galati and Melvin (2004) point out that since the beginning of the decade there has been a surge in foreign exchange trading. Carry and momentum trading strategies have predominantly driven this increase in activity. Galati and Melvin (2004) highlight the significant growth in the participation of Hedge Funds, in particular trend following strategies, which have considerably grown in numbers. However they also make the point that the landscape of Hedge Funds has changed considerably with time. Systematic trading funds that had entered the market were typically smaller than the trend following funds that had been there before. They also use algorithms that are much shorter-term in their nature than what has been used before. Galati and Heath (2007) reiterate the aspect of Hedge Fund participation in their review of the years from 2004 to 2007. Moreover they also point towards the aspect of algorithmic trading as one of the key sources of turnover within the foreign exchange markets.

King and Rime (2010) make particular reference to the concept of algorithmic trading in their analysis of foreign exchange volumes in the years from 2007 to 2010. With algorithmic trading King and Rime (2010) refer to systems that break up trades to optimise trade execution, or automated hedging by market makers or other forms of proprietary technical trading. Within the universe of algorithmic trading the concept of high frequency trading receives special attention in their paper. High frequency trading became an increasing contributor to the growth in trading volumes in foreign exchange markets since 2004 when the first electronic brokerage systems were launched. King and Rime (2010) estimate that high frequency trading takes up 25% of the volume of all spot transactions worldwide. From this it becomes evident that the concept of systematic trading has dramatically evolved over the past ten to fifteen years. The growth in systematic trading strategies has not only come from trend following funds,

which have seen spectacular growth over that time period, but also from a profound change in how foreign exchange markets operate. The advent of electronic brokerage has opened up infinite opportunities to exploit market dynamics by applying systematic trading rules. Therefore the decline in trading rule profitability, documented in this study and many other earlier studies is arguably a mere consequence of the evolution of foreign exchange markets.

Another conclusion of the results presented is the fact that White's "Data Snooping" framework does not work very well when it comes to the foreign exchange market. Namely the absolute search for the best data snooping proven trading rules out of the universe of 8775 trading rules that are investigated exhibits results with a great deal of variation over time. Namely, the periodic clustering of currency pairs amongst the top trading rules over time makes the interpretation of results challenging. Given the high level of "currency pair volatility" within the ten best trading rules over time, a relative analysis that looks across all currency pairs is preferable.

The results of the relative specification of the data snooping analysis allow for another key observation. Currency markets have gone through different regimes over the observation period. This aspect becomes evident when looking at the results of the second trading rule enhancement in Figures 2-8a b c. The second trading rule enhancement weighs the periodical strategy exposure according to the unconditional survival probability, thereby reducing the exposure level over time. During the first six sub-samples the second enhanced trading rule delivers White's P-Values that are significantly lower than the P-Values of the standard trading rule. However, in sub-samples seven to nine, this reverses. Also looking at the back test results in Figure 2-10 it can be seen that the trading rule added only very limited value in the early parts of the data sample. However, as general trading rule profitability deteriorated, the incremental value derived from applying this trading rule enhancement increases considerably.

This points strongly towards of Lo's (2004) Adaptive Market Hypothesis, which indicates that, in line with the evolutionary aspects of his theory, the market environment is subject to continuous change as new investor "species" enter the market pace, while existing investment opportunities cease to exist. Indeed with the increased participation of the money management industry in the foreign exchange space, a multitude of systematic trading strategies such as moving average crossovers, filter rules, channel breakout rules are used frequently. Many of them are trend following in their nature. Hence they profit from a directional trend in the underlying exchange rate.

Once this trend changes or reverses these strategies tend to make losses, which causes them to adjust their positioning. This dynamic has the ability of not only prolonging trends in exchange rates; it also gives a rationale for some of the other characteristics exhibited by exchange rates. Exchange rates can undergo very sharp reversals and they exhibit relatively high day-to-day volatility. One could argue that with the unprecedented growth in systematic trading strategies, the trending characteristic of currencies might have been exacerbated. This would suggest that the enhancements of simple trading rules or the creation of more sophisticated strategies that aim to reduce potential losses that could occur in a sudden reversal periods should add incremental value.

Given the high level of statistically significant differences in White's P-Values of the first and third trading rule enhancement, as shown in Figures 2-8 a, b, c, one can draw a clear conclusion that trading rule enhancements focussed on conditional survival probability do add incremental, yet persistent value versus a generic trading rule. The results of the backtest in Figure 2-10, allow for the same conclusion. The intuition of the trading rule enhancements presented is to apply weights to the currency exposure that change over time. The first enhancement weighs the exposure of the trading strategy according the historic conditional survival probability of trading rule signals, which tends to start at a medium level, then increases over time and falls off thereafter. As described earlier the pattern of the first trading rule enhancement assumes that the periodic survival probability of a trading signal is not constant and there is value to be added by applying higher levels of exposure to sections of the "life" of a trading rule signal where its relative survival probability is the highest. When compared to the overall level of transaction costs, the incremental value added is almost 10% for the two most recent sub-sample periods. The incremental value added in earlier sub-samples is somewhat higher. This result supports the findings of Jegadeesh and Titman (1993, 2001), who find that momentum returns in equity markets go through different stages over time. At the point of the signal generation, returns are weak; later they become more pronounced and then fade away. This pattern is also documented by other more recent studies such as Menkhoff, Sarno, Schmeling and Schrimpf (2011), who directly implement the Jegadeesh and Titman (1993, 2001) approach using foreign exchange data. However, the value added from the strategy deteriorates slightly over time. The line of argument that can be made within that context is that during the early parts of the data sample there was comparatively low competition amongst market participants. While the overall number of market participants was low, the universe of trading rules

applied must have been very narrow at the same time. Evidence of this can be found in early studies that investigate trading rules. Dooley and Shafer (1984) and Sweeney (1986) and later studies such as Levich and Thomas (1991) all focus on a narrow range of very basic, filter or moving average crossover trading rules. The aim of their choice of trading rules was to pick the most widely used set of trading rules in order to minimise a potential selection bias and to avoid data snooping. However this also meant that investors at that time were only using very generic trading rules to generate profits. However, over time the investor base has increased and competition has caused an erosion of trading rule profitability. Investors saw themselves forced to adapt. Hence they started exploring more sophisticated versions of trading rules. Algorithmic and high frequency trading systems were invented. Therefore, as market participants diversified their investment approach, some of the profitability of that enhancement has been arbitrated away.

One of the objectives of this study is to assess whether the methodology constitutes a market novelty that leads to superior trading rule returns. From the results presented in this chapter it is evident that there is an element of market novelty in the methodology applied. Therefore one can add incremental value by applying conditional and unconditional survival probabilities to moving average trading rules. While the former strategy, has only limited success in generating significant White's P-values, it delivers returns that are incrementally better than the returns achieved by a generic trading strategy. The incremental outperformance is persistent over time. The only exception is sub-sample 6 where the strategy underperforms. Strategy two, whose exposure is weighed according to the unconditional historic survival probability exhibits two distinct phases. The first phase comes with White's P-value's that are significantly worse than the White's P-values obtained by a standard trading rule. Likewise the performance of that strategy fails to stand out against any standard trading strategy. During the second phase, however, this strategy delivers high and significant White's P-values, as well as strong trading rule results. The fact that the strategy goes through two distinct regimes suggests that foreign exchange markets have changed over time. While traditional trading rules have gradually ceased to deliver positive returns, this trading rule enhancement is increasingly able to add incremental value.

This observation goes hand in hand with Lo's (2004) Adaptive Market Hypothesis. The foreign exchange market has evolved over the years. As a consequence, simple trading rule strategies that were once profitable fail to deliver positive returns in more recent

years as market participant arbitrage the excess returns that were once available away. As opportunities cease to exist, others arise. The results presented suggest that trading strategies, which are enhanced by applying conditional and unconditional survival probabilities to the exposure levels of returns, are able to add value versus a standard trading rule. The fact that the profitability of these enhanced trading rules has not yet deteriorated would suggest that this opportunity still exists. However, following Lo's logic, one would expect that this opportunity is likely to eventually fade away, while others are likely to arise.

Bearing this in mind areas of future research would be to look at broader universe of foreign exchange markets. In recent years emerging market currencies have become increasingly important. Nonetheless, while these markets have enjoyed a true pilgrimage of international investors, opportunities still seem to be plentiful there. This assumption is confirmed by studies such as Menkhoff, Sarno, Schmeling and Schrimpf (2011), who find that most of the momentum returns come from emerging market currencies, and Chong and Ip (2009) report 30% plus annualised returns by utilising a momentum based trading strategy in emerging market currencies. Another area of future research would be to extend the framework to more sophisticated trading rules. This chapter uses very basic trading rules as benchmarks, but those basic trading rules have mostly lost their power to generate positive returns. There are other, more sophisticated trading rules, that are still able to generate returns. The question as to whether the enhancements proposed in this chapter are also able to add incremental value to these strategies, might also be a subject of further research

VI. Chapter 3:

Momentum Effects:

Dissecting Generic G10 Trading Rule Returns

Abstract

This chapter builds on the work of Pojarliev and Levich (2008, 2010), who dissect the returns of active currency managers by applying a multiple ordinary least squares (OLS) regression to currency fund returns. Where the chapter differs is in the specification of the dependent variables, which are in the context of the present chapter a set of trading rule parameterisations that are applied to a broad range of currency pairs. The results of this chapter suggest that there is some alpha embedded in the returns of technical trading rules. Moreover, the chapter establishes a comparatively strong positive, statistically significant link between the risk factors Trend, Momentum, Risk Aversion. The results of the chapter clearly indicate that shorter-term moving averages exhibit less systematic exposure than longer term moving averages. Other factors such as Carry, Value and Volatility have a considerably less pronounced relationship; only few factor sensitivities are statistically significant. Moreover, the results also indicate that systematic risk exposures of trend following trading strategies change with small adjustments in the design of trading rules.

A. Outline

1. Academic Background

In recent years the use of technical trading rules has become an established way of managing money. Barclay Hedge, a database provider for systematic trading funds, estimates that currently US\$ 330bn is managed in systematic trading funds. Such funds apply complicated trading algorithms to a multitude of asset classes, in particular currencies. Whilst the performance that they have delivered historically is very strong, the question still remains whether the returns they generate are a true reflection of market inefficiencies or mere compensation for systematic risks that are taken on. Within academia, the main lines of argument that aim to explain the profitability of technical trading rules can be described as follows. Firstly, the activity of central banks is regarded to create inefficiencies. This is due to the fact that when intervening in foreign exchange markets, central banks are not aiming to maximise their profits. Moreover, data snooping is another common line of explanation. The idea behind that argument is the fact that when academic studies present evidence of highly profitable trading rules, the results could have come from choosing a particular trading rule that works very well at that time period, as opposed to true market inefficiency. Finally, the high returns from trading rules might well stem from taking on systematic market risk, and are therefore a risk premium. The first two arguments have received significant attention from the academic body. The key conclusions there are in favour of market inefficiencies. The argument of systematic risk taking, however, in the context of foreign exchange markets has mostly been explored within the context of the forward discount bias. The analysis of trading rule returns, on the other hand, have been somewhat neglected. This can partially be explained by the fact that the definition of a systematic risk factor against which trading rules are assessed is difficult to make in the context of foreign exchange markets. Most academic studies merely argue that a “long-short” implementation of a trading rule back tests will eliminate most of the systematic risk factors. Such an approach is valid; nonetheless it is not entirely satisfactory due to the fact that these strategies might still bear systematic risk, which is not accounted for in a traditional beta analysis. This chapter aims to demystify the sources of returns from generic trading rules.

a) Technical Trading Rules and Systematic Risk Taking

One of the early papers that look into trading rule returns and risk is Sweeney (1986) who uses the interest rate differential between different currencies plus a constant risk premium as the market price for risk. The results of his paper suggest that the application of trading rules leads to significant excess returns, which cannot be explained by systematic risk taking. Neely and Weller (2011) hereby point out that Sweeney's (1986) assumption of a constant risk premium is more appropriate for the equity space, which might show unconditional risk premia, but comes with problems for the foreign exchange space.

Neely, Weller and Dittmar (1997) are another example of an early study that analyses trading rules in foreign exchange markets, explicitly controlling for risk. They apply a genetic program that searches for an optimal trading rule, aiming to eliminate data snooping biases. The key findings are that different currency pairs produce higher trading returns than others and that different currencies pairs also favour different sets of trading rules. They indicate that their optimal trading rules show out of sample profitability when compared to bootstrap simulations. They also analyse to which degree return streams of these strategies are impacted by market risk, which they define as a series of broad equity market indices, one being the MSCI World index and several national indices. Their results show only one value that suggests a significant positive relationship between the trading rule results of a currency pair and a market index. Most of the results suggest no or even a negative relationship to equity market indices. They therefore conclude that excess returns derived from the trading strategies are not a risk premium earned for taking systematic risk.

Kho (1996) also examines whether the results of various technical trading rules can be explained by time varying risk premia. The focus of his study is somewhat different than Sweeney (1986) and Neely, Weller and Dittmar (1997), in as far as the previous studies are focussing on analysing and creating profitable trading rules, with systematic risk control as a second round criteria. Kho (1996), however, solely focuses on the aspect of systematic risk taking within the context of technical trading rules. His risk premium assumption is also derived from the CAPM literature, whereby the risk premium is defined as the co-variation of returns between the return stream generated by a moving average trading strategy and the MSCI world equity index, which is used a

market proxy. The results of his paper suggest that the application of trading rules leads to significant excess returns, which cannot be explained by systematic risk taking. He evaluates a set of moving average crossover rules using weekly data on foreign currency futures contracts from 1980 to 1991 for four different currencies against the USD.²² The author also compares the results from the actual trading rules to results obtained from simulations that aim to replicate the historic evolution of time varying risk premia. Kho (1996) confirms that the magnitude of trading rule results is in line with what has been suggested in previous papers. However, he points out that the profits might come from the existence of time varying risk premia. In particular his results indicate that periods of higher or lower returns identified by the technical trading rules largely correspond to those of higher or lower conditional expected returns, due to high or low risk premia and volatility. Therefore, the majority of technical trading rule profits might well be a result of time varying risk premia.

Other, more recent papers such as Wang (2004) look at currency returns from a market microstructural perspective. Wang (2004) incorporates the positioning of market participants such as hedgers and speculators when designing tests of foreign exchange market efficiency. The paper utilises weekly data of five²³ currency futures contracts against the USD listed on the CME. The sample spans from January 1993 to March 2000, the DEM sample ends in September 1999 due to a lack of liquidity in the futures contract following the introduction of the EUR. Besides price information, the sample also includes information about the market position of hedgers and speculators. Wang (2004) finds that speculator sentiment varies positively with future returns, while hedger sentiment varies negatively with future returns. Moreover, positive or negative extreme sentiment exhibits a higher correlation to price movements than moderate sentiment. The study also indicates that the aspect of hedging pressure has to be considered in the context of foreign exchange markets. Wang (2004) suggests that, while the relation between speculator sentiment and returns remain positive and statistically significant after accounting for market risk, which is defined by the paper as the value weighted CRSP Index, it becomes insignificant after accounting for hedging pressure. Hence the study gives evidence that speculator profits are largely explained by risk premia. The indication that hedgers lose to speculators in the foreign currency markets is in line with financial theory as well, given the fact that the losses can be interpreted as an insurance

²² British Pound (GBP), Deutsch Mark (DEM), Japanese Yen (JPY), Swiss Franc (CHF)

²³ British Pound (GBP), Canadian Dollar (CAD), Deutsch Mark (DEM), Japanese Yen (JPY), Swiss Franc (CHF)

premium paid by the hedgers to the speculators.

Both of these studies, Kho (1996) as well as Wang (2004), present a more elegant model for the risk premium to quantify systematic risk factors in trading rule returns than the model presented by Sweeny (1986). However, they assume equity indices as their market portfolio against which trading rule returns are compared. In the case of Kho (1996) this index is the MSCI world equity index, which is an index of global developed equity markets, in the case of Wang (2004) CRSP index is used, which is an all market capitalisation equity index covering almost 100% of the US equity market. The problem with this choice of equity indices as market portfolios for foreign exchange markets is the fact that it requires very strong assumptions about the degree to which foreign exchange markets are integrated. Moreover, even under the assumption of fully integrated markets, the use of an equity index as an appropriate market portfolio for foreign exchange risk remains questionable.

When looking at Kho's (1996) analysis it becomes evident that a comparison of currency market trading rule returns and equity benchmarks might not be fully appropriate. Kho's data suggests that the equity index exhibits a relatively high degree of negative skew and very high levels of kurtosis, leading to a rejection of the Bera-Jarque test, which suggests that returns do not follow a normal distribution. The returns of the trading rules are also not normally distributed. However, the rejection of the Bera-Jarque test is based on positive skewness and much lower levels of kurtosis. This raises the question whether both measures are comparable especially given the fact that, from an academic point of view, the return generating process is different for equity and currency markets. While the traditional frameworks of pricing systematic risks, such as the CAPM is undoubtedly valid from an intellectual perspective, the aspect of Roll's (1977) critique, which questions the appropriateness of the market portfolio, is particularly problematic in the context of foreign exchange markets. This aspect is a very likely explanation for the fact that academic literature in that particular line of research is rather sparse.

However, in recent years the notion of what is classified as a systematic risk factor has changed considerably. While beta has been associated with systematic risk exposure to a broad benchmark index that represents an asset class, more recent studies propose broader ways of looking at beta and proxies for systematic risk exposure. Anson (2008) is notable in this context. He suggests that there is no clear distinction between alpha and beta and that there is a continuum between both. Therefore, defining beta in a traditional sense does not fully capture all potential systematic risk exposure. He links

his thesis to various examples of beta within the equity, fixed income and currency space, such as bespoke, alternative, fundamental, cheap, active and bulk beta, which are in his view different ways of harvesting risk premia by obtaining exposure to systematic risk factors other than the traditional broad market exposure. Given this change in perception of what is systematic risk, research in the foreign exchange space started evolving as well. More recent studies about systematic risk and foreign exchange markets focus more on trading rule returns as being sources of currency risk in their own right.

b) Technical Trading Rules as Systematic Risk Taking

While, Lequeux and Acar (1998) establish a trading rule benchmark that replicates the risk and reward profile of the average actively managed currency fund, Schulmeister (2006) sheds some light on the systematic risk taking aspect of trading rules from the perspective of market microstructure. He links the behaviour of technical models and exchange rate dynamics. The principal idea of his line of argument is the fact that, while traders do not follow technical signals, they monitor them frequently. By doing so, they are altering market behaviour. This means that traditional price discovery under the efficient market hypothesis, which is driven by private information becoming public information, is somewhat violated, due to the additional aspect of market participants being aware of technical trading rules. This has implications on the link between trading rules and currency volatility, as well as the link between trading rules and systematic risk. His study is based on the analysis of the predictive power of aggregate trading signals. He analyses 1024 moving average and momentum models in the DEM/USD market between 1973 and 1999. He also conducts an out-of-sample test for the EUR/USD rate over the time period from 2000 to 2004. The results of his analysis suggest that when markets change direction, the majority of trading filters in his study tend to be on the same side, i.e. they are either long or short. For a trading rule to adapt to a trend, it usually takes 10 to 20 days. Schulmeister's (2006) results suggest that there is a pronounced feedback mechanism between trading rules and movements in the underlying exchange rates. He explains this by analysing exchange rate movements around time periods when the majority (97.5%) of trading rules change their position from long to short or vice versa. Therefore, according to his line of argument, there is a multiplier effect linked to technical trading rules, which translates a small news flow

into a market trend. Schulmeister (2006) also indicates that the majority of trading rules are profitable, in and out of sample, and that the profitability is exclusively due to persistence in exchange rate movements. Therefore one could argue that market participants expect the persistence in price movements to be sufficiently frequent to compensate for the potential loss that occurs due to a sudden reversal of a trend. Schulmeister (2006) notes that in a market that is purely rational, where market participants are utility maximizing, individual technical trading would not be profitable due to the laws of arbitrage. Moreover, he also argues that if someone assumes the imperfection of human knowledge and that decisions of market participants are only partly based on reason, then the occurrence of exploitable trends is likely. Therefore it is not far-fetched to assume that markets are not perfectly efficient. While Schulmeister's (2006) conclusion is perfectly congruent with the traditional way of looking at market efficiency, one could also look at his results from the perspective of Lo's (2004) Adaptive Market Hypothesis, which draws comparisons between markets and the ecology, both of which follow evolutionary paths. In his view, markets are a combination of behavioural biases and the market forces of supply and demand. Therefore as the behaviour of market participants (humans) follows the concepts of "evolutionary psychology", market behaviour is likely to follow similar concepts. Lo (2004) indicates that different investors follow different behavioural patterns. Therefore the market environment for a particular asset does not only change due to a general change in preferences by all market participants, but also by a change in the composition of investor groups competing for specific assets. This line of argument allows rationalising a multitude of characteristics of modern financial markets. However, within the context of Schulmeister's (2006) results, the argument that investors have to adapt to the changing market environment in order to achieve persistent levels of return is the most important one. One could argue that over time, as investors had to adapt to the changing market environment in foreign exchange markets, they have become aware of the reinforcing link between trading rules and market trends. Therefore, it has become rational for market participants to exploit that relationship. One could even go as far as arguing that exploiting such relationship is in fact harvesting a risk premium.

Brunnermeier, Petersen and Nagel (2008) contribute to this line of research by looking at the risk return profile of carry trades. Their research suggests that the returns derived from buying high yielding currencies, while selling lower yielding currencies can be

seen as compensation for “crash risk” associated with carry trades. Their findings suggest that strategies, which go long high-interest-rate and short low-interest-rate currencies are negatively skewed, which stems to a great extent from funding risk associated with carry trades. This is due to the fact that currency trades often have high levels of leverage behind them and, while the equity risk premium depends on investor preferences in an unlevered world, the carry premium depends on willingness of investors to take on leverage. This argument is reinforced by the observation that the carry trade tends to be loss making in time periods when traders have funding problems, as documented by a positive correlation of crashes in the carry trade with equity market volatility as well as changes in TED spreads. Moreover, they also document, that after controlling for other factors exchange rates with similar interest rate levels exhibit co-movements with each other. This indicates that carry trades do affect the behaviour of exchange rates.

Lustig, Roussanov and Verdelhan (2010, 2011) take this line of argument further and aim to formalise a country specific as well as a generic risk premium for currencies. The country specific risk premium looks at the relative differential between the interest rate in each of the respective countries analysed and the average interest rate across all other countries within the universe subject of the study. If the interest rate of the country exceeds the average interest rate of the basket, the risk premium is defined as the currency return of the country minus the return on the basket, and vice versa. Lustig, Roussanov and Verdelhan (2010, 2011) also define a global risk premium, which is the return on the highest minus the return on the lowest interest rate currencies. Their studies indicate that this measure offers a parsimonious explanation of currency risk premia. They also provide strong evidence that such risk premium has the power to explain the returns derived from carry-based strategies.

c) Technical Trading Rules and Multi Factor Models

Further examples of this more recent line of foreign exchange research, which treat anomalies such as momentum or carry as risk premium strategies in their own right, are Pojarliev and Levich (2008, 2010). Pojarliev and Levich (2008) contribute to this line of thought in as far as they establish a universe of four currency benchmark strategies against which they compare various currency fund managers. Pojarliev and Levich (2008) highlight that the factors Carry, Trend, Value and Volatility explain a significant

part of the returns of currency fund managers, with Carry and Trend being the most dominant factors. They indicate that over the entire sample period, spanning from 1996 to 2000, 66% of the variability in monthly returns of their manager universe can be explained by these four factors. In the time period after 2000 the explanatory power of the four factors rises to almost 77%. One of the notable differences in their approach versus other studies is their definition of beta benchmarks. Previous studies had either assigned the risk free rate or a zero return as benchmarks for currency strategies. This is due to the fact that currencies are deemed to be unpredictable, or the fact that they exhibit low correlations with equity benchmarks. Therefore, all returns greater than the interest rate or zero are assumed to be alpha. Pojarliev and Levich's (2008) framework introduces explicit benchmarks that rely on basic currency trading strategies, which can be harvested by investors with very low cost. They also suggest that the systematic risks that are associated with the currency beta indices are as follows: in the case of Carry, which is an investment strategy that buys high yielding currencies, while selling lower yielding currencies, the high interest rate currency may depreciate by more than the interest differential. For Momentum, the risks are sudden reversals of trends or patterns, trading based on false signals and excessive trading costs. The risks associated with investing in Value and volatility strategies are firstly the risk that convergence to fair value, in their case the Purchasing Power Parity (PPP), takes longer than expected, with potential further deviation from fair value. Moreover, there is also the risk of being long or short volatility in a falling or rising volatility environment. Pojarliev and Levich (2008) use simple regression analysis to evaluate their results. They regress an index that represents the universe of active currency managers on the four factors mentioned earlier. The results suggest that Trend is the most significant factor in their analysis. While the explanatory power of Trend has declined somewhat in recent years, the overall percentage of currency fund returns explained by the factor Trend is 65% throughout the 1990s and after 2000. They also indicate that the disappointing returns from currency managers are mainly the consequence of declining profitability of Momentum as a trading strategy.

Pojarliev and Levich (2010) expand the results of their 2008 paper by analysing the degree to which the aforementioned "benchmark" trading strategies are utilised by currency managers. In their search for currency trading strategies that crowded. Hence, are followed by a large number of active managers. They make reference to the aspect of changing volatility and correlation characteristics of currencies as a consequence investor preferences. Similar to the findings of Brunnermeier, Petersen and Nagel

(2008), one of the examples that they give is the high correlation between the GBP/CHF cross and the NZD/JPY cross where there is no economic reason as to why these two currency pairs should be highly correlated. The only similarity that those two crosses share is the fact that GBP and NZD are traditionally high yielding currencies, while CHF and JPY are historically low yielding currencies. If Carry becomes popular, investors will go long high yielding currencies, while funding these purchases with low yielding currencies. As a consequence, one would expect the correlation of currency crosses that combine high and low yielding currency pairs to go up. Indeed the correlation of these two currency pairs follows closely the cycles of popularity of carry trades. As formulated by Brunnermeier, Petersen and Nagel (2008), the returns derived from carry strategies compensate for “crash risks”, driven by funding and liquidity constraints of market participants. In the case of momentum trading rules, the systematic risk factor that is harvested might well be the compensation for increasing currency volatility due to the pronounced feedback mechanism between trading rules and movements in the underlying exchange rates as indicated by Schulmeister (2006). Pojarliev and Levich (2010) define crowdedness as the percentage of the funds with significant positive exposure to a given “benchmark” trading strategy less the percentage of the funds with significant negative exposure to that same strategy²⁴. Key findings of their study are a high degree of crowdedness of carry in 2007 and 2008. Their carry crowdedness measure peaks in April 2008 at 32%; it collapses subsequently as the global financial crisis reached its peak. Carry crowdedness sees a considerable increase throughout 2009 reaching also 32% in the latter part of that year, followed by a drop of Carry from the second half of 2010 onwards. Trend on the other hand side is very popular in the early parts of their sample, which spans from 2005 to 2010. The levels of Trend crowdedness range between 25% and 35%. By May 2008 Trend crowdedness declines to almost zero, a few months before the performance of the Trend factor starts picking up again throughout autumn of 2008. Subsequently the number of fund managers following the trend strategy picks up again. The measure for crowdedness reaches 21.6% in November 2009. It then declines again to almost zero by the middle of 2010. The results of Pojarliev and Levich (2010) link very well with Lo’s (2004) Adaptive Market Hypothesis, where different investment strategies can go through different stages of profitability. This is due to the fact that in Lo’s (2004) view there are different “species” of market participants, which have distinguishable

²⁴ The measure of crowdedness is calculated by analysing style betas of the universe of managers to the aforementioned “benchmark” trading strategies.

characteristics in terms of their investment pattern. Hence, market performance is not only driven by a change in overall market preferences, but also in a change of the composition of market participants competing in a specific asset class. Asset classes also go through cycles where competition of market participants varies. Therefore the risk/reward trade off will differ across assets and will also change over time. Indeed when Pojarliev and Levich (2010) identify Carry to be crowded in early 2008, the performance of carry strategies subsequently drops off. Conversely they also indicate that Momentum was unpopular with investors at around the same time, but delivers strong returns subsequently. Following these two seminal papers Pojarliev and Levich conduct a series of further studies which use generic technical trading rules to evaluate active currency managers and the quality of their alpha generation (2012, 2013). The test setup of these subsequent papers is very similar to the test setup of the 2008 and 2010 papers by Pojarliev and Levich, where systematic risk factors such as Trend, Carry, Value are created by applying simple technical trading rules, against which active managers are compared. Both of the later papers deal with the performance evaluation of active currency managers. The 2012 paper uses the style based regression analysis proposed in the 2008 paper, to evaluate the alpha generation accounting for systematic risk exposures of the “benchmark” trading strategies. This is done by analysing the goodness of fit of a multiple regression of between the returns of the active managers and the returns of the systematic risk factors. The 2013 paper focuses on the intercept of the multiple regression analysis, i.e. the regression alpha, which in the context of the paper represents the true alpha generation of active currency fund managers.

In conclusion, three lines of research that look at technical trading rules and systematic risk have been conducted. The first line of research looks at trading rules in the traditional CAPM context. The definition of an appropriate market portfolio proves hereby challenging, and the appropriateness of equity indices that are used in these studies remains questionable. The second line of research rationalises why certain generic trading rules are systematic risk factors in their own right. Schulmeister (2006) provides compelling evidence for the factor Trend, while Brunnermeier, Petersen and Nagel (2008) and Lustig, Roussanov and Verdelhan (2010, 2011) make a case for the factor Carry. The third line of research uses technical trading rules as specific risk factors. Pojarliev and Levich (2008, 2010) propose a multiple regression framework, whereby the returns of active currency managers are analysed. The motivation of this

chapter lies in a combination of all these three lines of research. Generic technical trading rules are hereby assessed against a series of systematic risk factors, widely accepted and used by the foreign exchange investment community.

2. Motivation of the Chapter and Main Contributions

The argument presented by earlier studies, such as Okunev and White (2003), that long-short trading strategies tend to have market covariance levels that are close to zero (hence are not subject to systematic risk factors) clearly does not hold. This chapter looks to shed light on whether the returns from systematic trading rules are due to pure market inefficiency, or whether there is an aspect of compensation for taking on systematic risk. This chapter proposes to analyse simple trading rule returns for systematic risk factors in a broader way, thereby addressing some of the criticisms of earlier studies. Assessing whether trading rule returns are a compensation for taking on risk is undoubtedly an important academic question. However, as pointed out by Neely and Weller (2011) it is heavily dependent on the construction of a convincing model for the risk premium.

In the spirit of Lo's (2004) Adaptive Market Hypothesis market participants have to continuously adapt to a changing market environment in the foreign exchange space. As a consequence they have learned to exploit the reinforcing link between trading rules and market trends as proposed by Schulmeister (2006), or they have learned to bear the risks associated with carry strategies as suggested by Brunnermeier, Petersen and Nagel (2008). Therefore, exploiting relationships such as trend and carry have in fact become a common way of harvesting risk premia, and indeed many of the newer studies in the field of foreign exchange markets such as Pojarliev and Levich (2008, 2010), treat anomalies such as trend or carry, and many others, as a risk premium strategy in their own right.

Given these developments in the foreign exchange space, this chapter looks to assess technical trading rule returns against a model of multiple risk factors such as Trend, Momentum, Carry, Valuation, Risk Aversion and Volatility. The chapter represents an extension of the work proposed by Pojarliev and Levich (2008, 2010), whereby a wider universe of factors is used. The main difference lies in the fact that the proposed chapter looks to analyse systematic risk exposures of simple technical trading rules as opposed to the returns of active currency managers. While it is appropriate to run a series of

independent multiple regressions for currency fund managers, which are following different investment strategies. Such test setup is not appropriate in the context of technical trading rules that use similar parameterisations such as SR1/LR5 or SR1/LR30, across a range of currency crosses. This is due to the fact that there is a high likelihood of commonalities between the SR1/LR5 trading rule for the USD/GBP and the USD/EUR cross. In order to account for these cross currency commonalities the proposed framework is based on a one step GMM model, which allows the calculation of the general sensitivity of the specific risk factor to the universe of trading rules that are calculated for each trading rule parameterisation. The results make it evident that factors such as Trend and Momentum and Risk Aversion have a relatively strong positive and statistically significant impact on trading rule returns. It should be noted, however, that this is less the case for shorter term moving averages, while as longer term moving averages exhibit more systematic exposure. Other factors such as Carry, Value and Volatility have a considerably less pronounced relationship to trading rule returns. Moreover, the findings of the paper also suggest a comparatively strong, positive and statistically significant link between the risk factors Trend, Momentum, Risk Aversion. These results make a strong case for the fact that at least a part of the returns from technical trading rules are driven by systematic factors. Paired with the finding that shorter term moving averages exhibit higher levels of alpha, the results in this chapter would suggest that shorter term moving averages are less affected by systematic risk factors than is the case for longer moving averages.

The idea of creating a universe of systematic risk factors for the foreign exchange market has become popular in the more recent academic literature, but creating a link to technical trading rules has not. Finally assessing the impact of different risk factors on trading rule parameterisations that spread across a series of currency crosses has also not been attempted before.

B. Data and Methodology

1. Data, Return and Moving Average Calculations

The chapter uses the same dataset as the previous chapters. The dataset contains daily New York end of day losing mid-prices for nine G10 currencies against the USD, as

well as three-months cash rates for corresponding countries. All other currency crosses are calculated from these nine pairs. The sample spans from the 4th of January 1974 to the 31st of December 2009. After adjusting for non-trading days, it contains 9025 data points. Given the lack of long-term history for the EUR, the time series for the EUR rate is backfilled with the historic Deutschmark (DEM) rate, with the original EUR fixing rate of 1.95583 DEM per 1 EUR, as of 1 January 1999. In addition to this dataset, the current chapter uses annual estimates of the Purchasing Power Parity rate of the nine currency crosses against the USD, which are published by the OECD. It also uses estimates of trading volume for the most important currency crosses, which are published in the Triennial Central Bank Survey of foreign exchange and derivatives market activity published by the Bank of International Settlement

All currency return calculations are based on returns that are adjusted for the interest rate differential. This is done to mimic the returns obtainable from a futures based trading strategy. The exchange rates are expressed in units of domestic currency versus one unit of foreign currency. Equation 1 calculates an interest adjusted return time series. The first term represents the daily interest rate differential between foreign (r_f) and domestic (r) currencies. The second term shows the return from currency appreciation. S_t is the currency spot price at time t .

$$(1) \quad R_{I,t} = \left[\left(\frac{1+r_f}{1+r} \right) * \left(\frac{S_t}{S_{t-1}} \right) \right] - 1,$$

The interest rate calculations in Equation 1 are based on the Money Market Basis convention (Actual/360). The adjusted return time series, obtained from the equation, results in approximate currency returns that can be earned by following a futures based investment strategy. The calculations for the interest rate differential are based upon the three-month T-Bill rate, for which clean time series across all countries in the G10 currency universe exists. While the three-month T-bill rate is only the second best adjustment factor after the overnight rate, the first chapter has proved that both interest rate adjustments are equivalent.

The chapter also follows the previous chapters in its definition of trading rules. It uses a simple price moving average filter. The rationale for this very basic choice of trading signal is the fact that it is very parsimonious. Equations 2 and 3 describe the crossover signals used to calculate the trading filters.

$$(2) \quad \text{Positive Momentum} = \frac{1}{S} \sum_{i=0}^{S-1} S_{t-i} \geq \frac{1}{L} \sum_{i=0}^{L-1} S_{t-i}$$

$$(3) \quad \text{Negative Momentum} = \frac{1}{S} \sum_{i=0}^{S-1} S_{t-i} < \frac{1}{L} \sum_{i=0}^{L-1} S_{t-i}$$

The time periods for the short-term moving averages, here denoted as (S), range between 1 to 5 days as well as 10, 15, 20 and 25 days. The time periods for the long-term moving averages (L) are defined as 5, 10, 15, 20, 25 and 30 days. Any short-term moving average has to be shorter than any long-term moving average. Equation 2 suggests that a positive momentum signal is established when the short-term moving average is equal to or above the long-term moving average. Equation 3 indicates that a negative moving average signal is established when the short-term moving average lies below the long-term moving average. These trading signals are then translated into a long-short trading rule. Under the assumption of zero transaction costs, the returns of such rule are defined in Equation 4

$$(4) \quad \text{TRR}_{I,t} = \text{SIG}_{t-1} R_{I,t},$$

$$\text{SIG}_{t-1} \begin{cases} +1, \text{positive momentum} \\ -1, \text{negative momentum} \end{cases}$$

Equation 4 indicates that the trading rule return at time t is the product of the periodic currency returns as given in Equation 1 and the trading signal established at the end of period t – 1. Signal SIG may take a value of +1 if a positive momentum signal is established, or a value of -1 in the case of a negative momentum signal. Transaction costs are assumed to be zero, given the fact that the chapter focuses on analysing the link between trading rule returns and a respective currency benchmark, as opposed to verifying the historic profitability of these trading rules.

After removing reverse currency pairs from the overall G10 universe a total of 45 currency crosses is analysed. This is done for a set of 39 trading rule parameterisations, as shown in Figure 3-1.

FIGURE 3-1: MOVING AVERAGE COMBINATIONS

	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	1/5	1/10	1/15	1/20	1/25	1/30
SR 2	2/5	2/10	2/15	2/20	2/25	2/30
SR 3	3/5	3/10	3/15	3/20	3/25	3/30
SR 4	4/5	4/10	4/15	4/20	4/25	4/30
SR 5		5/10	5/15	5/20	5/25	5/30
SR 10			10/15	10/20	10/25	10/30
SR 15				15/20	15/25	15/30
SR 20					20/25	20/30
SR 25						25/30

The column labels denote long-term moving averages and row labels denote short-term moving averages. All short-term moving averages have to be shorter than any long-term moving average. This equates to 39 different sets of moving average combinations.

2. Specification of Systematic Risk Factors

a) Factor 1: Trend

The risk factor Trend is one of the factors used by Pojarliev and Levich (2008, 2010), who propose the AFX Currency Management Index, constructed by Lequeux and Acar (1998). The aim of their paper is to create a parsimonious index representing the universe of trend following strategies, which can be implemented at relatively low cost. It consists of three moving average trading rules, SR1/LR32, SR1/LR61 and SR1/LR117, which are equally weighted for each currency pair. Whenever the long term moving average goes above or below the short term moving average, a long or short position is initiated for the duration of one day. The rationale behind choosing this particular set of moving average combinations is to match the investment time horizon of the average currency manager. The index is constructed by volume weighting the trading rule returns of the different currency pairs. Richard Levich publishes a daily time series of that index since January 1984. This chapter will use the time series of daily returns provided on the website of Levich²⁵. In order to undertake the analysis for the entire sample period, which starts in 1975. This chapter replicates the methodology and currency weighting proposed by Lequeux and Acar (1998). This time series is then used to backfill the data of the AFX index from the time period of 1975 until 1984. The correlation between the AFX index and the backfilled index is 0.9 for the time period

²⁵ http://people.stern.nyu.edu/rlevich/afx_index.html

from 1984 onwards where both indices are live and comparable. Hence, the backfilled data represent an appropriate proxy for the AFX Index.

b) Factor 2: Momentum

The second risk factor is Momentum. It finds its rationale in the work of Menkhoff, Sarno, Schmeling and Schrimpf (2011), which replicates the traditional cross sectional momentum literature pioneered by Jegadeesh and Titman (1993) using foreign exchange data, whereby currency pairs are ranked by their performance over a specified period. The best performers will receive a long position, while the worst performers receive a short position. Sarno, Schmeling and Schrimpf (2011) find similar pattern of returns of portfolios as reported by Jegadeesh and Titman (1993). They present strong evidence of under reaction for very short term holding periods, with strongly positive returns for medium term holding periods and a reversal of momentum returns for longer holding periods. While they do make the point that their returns are highly time varying and that a considerable part of the returns stem from exotic currency pairs that are outside of the G10 universe, it is still felt appropriate to include a variation of this strategy in the mix of systematic factors. Deutsche Bank assembled an index that replicates the returns available from applying a generic momentum strategy. The index is created in such way that G10 currencies are ranked by their 12-month return, from which the three top performing currencies are bought, while the three worst performing currencies are sold. The ranking is reassessed on a monthly basis. Deutsche Bank provides a daily time series of that index since June 1989. It is this index that is used here. Similar to the factor trend, this paper replicates the construction of the Deutsche bank Index. The constructed time series is then used to backfill the index data in the time period from 1975 to 1989. The backfilled time series exhibits again a high correlation in excess of 0.89 with the index during the time period from 1989 onwards, when both are live.

c) Factor 3: Carry

The third factor used in this chapter, is originally proposed by Lustig, Roussanov and Verdelhan (2010, 2011). It resembles the Carry factor in Pojarliev and Levich (2008,

2010), who use the Deutsche Bank G10 Harvest Index as the proxy for the returns of a carry strategy. This index, that is used here, is constructed from the G10 currency universe and it captures the return of being long the three highest yielding currencies within the universe, while being short the three lowest yielding currencies. The ranking of currencies is done on the basis of three months interest rates and the index is rebalanced quarterly. Deutsche Bank provides a daily time series of the index since September 2000; the time period before that is again backfilled by replicating the methodology. The correlation between the index and the replicated time series is 0.97 making both time series indistinguishable.

d) Factor 4: Value

Factor 4 is the factor Value; it is also one of the factors used in Pojarliev and Levich (2008, 2010). The rationale for using this factor is based on the idea that in the long term currencies are mean reverting, hence any currency that is deemed to be very cheap is more likely to appreciate than to depreciate. The relative degree of cheapness is determined by its valuation relative to the Purchasing Power Parity, which is based on the law of one price. The PPP relationship is long term in its nature and it is rather loose, hence it only exhibits strong explanatory power in time periods of extreme valuations. Similar to Pojarliev and Levich (2008, 2010) this chapter also uses the Deutsche Bank FX PPP Index. The Index ranks all G10 currency pairs according to their valuation relative to their PPP, which is provided by the OECD on an annual basis. The three most undervalued currency pairs are bought, while the three most overvalued currencies are sold. The index is rebalanced every three months. Deutsche Bank provides daily returns for this index from June 1989 onwards. In order to backfill the time period from 1975 to 1989, this chapter builds a time series that replicates the index. At the time when both time series are live, the correlation is again very high, with a value of 0.9.

e) Factor 5: Risk Aversion

The fifth risk factor is labelled the heuristic Risk Aversion factor. The rationale for this factor can be found in Lo's (2004) Adaptive Market Hypothesis, which advocates the notion that financial markets follow evolutionary paths. In that context Schulmeister

(2006) argues that investors use trading rule signals as heuristics in their assessment of markets, and that they consequently become risk factors in their own right. This argument can be extended to other heuristics used by the investment community, one of which is the notion of risk versus safe haven currencies. Market participants perceive the JPY and the CHF as safe haven currencies, while the EUR and the GBP are seen as currencies that are correlated with risk assets. Academic evidence for this assumption of JPY and CHF being safe haven currencies, while EUR and GBP being risk currencies can be found in publications such as Mueller, Stathopoulos, Vedolin (2012), who analyse correlation risk in foreign exchange markets. All of the currencies have a long time history of free float and ample market liquidity. Hence the inclusion of a strategy that deducts the average return of a long position in the JPY and the CHF versus the USD from the average return of a short position in the EUR and the GBP versus the USD, provides an appropriate proxy for risk aversion in foreign exchange markets.

f) Factor 6: Volatility

The last factor used in this chapter is the factor Volatility. This factor is also proposed by Pojarliev and Levich (2008, 2010), which is proxied with the Deutsche Bank Currency Volatility Index. This index follows in its construction the CBOE Volatility Index®, it is calculated as the weighted average of three-months implied volatility currency options in the most liquid currency pairs. The daily time series of this index spans back to September 2001, which is again not long enough to cover the sample period of this chapter, therefore the chapter will create a risk factor that mimics the dynamics of this index by estimating the stochastic volatility of the most liquid currency pairs by applying a GARCH (1,1) model and by combining these volatility measures into an index. The replicated index consists of the nine most liquid currency pairs USD/EUR, USD/JPY, USD/GBP, USD/AUD, USD/CHF, EUR/JPY, EUR/GBP and EUR/CHF. The weights of the currency pairs are determined by rebalancing the average trading volume from 2001 to 2010 as published by Triennial Central Bank Survey of foreign exchange and derivatives market activity published by the Bank of International Settlement. The correlation between the replicated index and the index provided by Deutsche Bank is very high, exceeding 0.99 over the time period that both time series are live, making both indices indistinguishable.

g) Correlations across Factors

The inclusion of single risk factors in a multi risk factor model depends, besides the economic meaningfulness of the risk factor, also on the correlation of the risk factor to other risk factors. If two risk factors are highly correlated with each other, the regression framework cannot distinguish which of the highly correlated factors drives the depended variable and which of the independent factor is merely coincident. Figure 3-2 shows a correlation matrix of the factors used in the analysis. The factors used in this chapter, while following different rationales, might well be constructed such way that some degree of correlation between them is inevitable. When looking at Trend and Momentum, both strategies look at a continuation of trends. While the factor Trend might be more short term in its nature, Momentum tends to be longer term. The correlation between both is 0.19, therefore well below the threshold of 0.8 – 0.9 that would raise concerns about multicollinearity. Momentum and Carry tend to be opposing forced, as suggested by Pojarliev and Levich (2010), the correlation between those two factors is therefore negative, but again below the threshold that would raise concerns.

FIGURE 3-2: CORRELATION ACROSS RISK FACTORS

TREND	MOMENTUM	CARRY	VALUE	RISK AVERSION	VOLATILITY
18.9%					
-5.1%	-25.5%				
2.9%	6.6%	-26.7%			
-1.2%	-7.8%	-31.5%	-15.3%		
3.4%	2.9%	-9.6%	1.7%	-4.2%	

The matrix in the figure shows the correlation across the risk factors used in the chapter.

The factors Value and Carry are also negatively correlated, this stems from the fact that high yielding currencies tend to be overvalued and the carry strategy tends to bet on a continuation of that overvaluation, while the Value strategy bets on a reversion to the mean. From its construction, the Risk Aversion factor has a fair level of negative

correlation to the factor Carry. This is due to the fact that the CHF and the JPY are perceived low yielders, while at least the GBP is traditionally a higher yielding currency. Yet the correlation between this heuristic Risk Aversion factor and the Carry factor proposed by Lustig, Roussanov and Verdelhan (2010, 2011) is in the range of -0.3 and can therefore be included in the analysis. The Volatility factor generally has a low correlation to all of the other risk factors.

3. Specification of the Multivariate Factor Model

As outlined earlier, this chapter follows the intuition of Pojarliev and Levich (2008, 2010), who dissect the returns of active currency managers by applying a multiple OLS regression to currency fund returns. The independent factors that are used are proxies for systematic risk factors within the foreign exchange market. While their test setup is inspired by the papers of Sharpe (1992) and Fung and Hsieh (1997, 2004), which apply an asset based style factor analysis to fund and hedge fund managers, the essence of their test has its roots in the Arbitrage Pricing Theory (APT).

Equation 5 describes a generic multi factor model, that allows for a factor that is common to all assets in the universe and a series of factors that affect individual assets only. $y_{i,t}$ is hereby the excess return of asset i at time t . $i = 1, \dots, N$ represents the universe of assets, which is in this chapter the universe of trading strategies. Moreover given the fact that the natural benchmark of a currency speculator is one of not being invested in any currency and therefore not earning any interest rate, as suggested by Qi and Wu (2006), excess returns as defined by $y_{i,t}$ are assumed to be the returns of the trading strategy. $t = 1, \dots, T$ represents the time period. j is the number of factors included in the model and $j = 1, \dots, F$ are in the case of this chapter the factors Trend, Momentum, Carry, Value, Risk Aversion and Volatility.

$$(5) \quad y_{i,t} = \beta_0 + \sum_{j=1}^F \beta_{i,j} d_{j,t} + \varepsilon_{i,t}$$

β_0 is a constant, $d_{j,t}$ are unexpected changes in exogenous risk factors, with $\beta_{i,j}$ being the systematic exposure of a specific asset i to common risk factors. $\varepsilon_{i,t}$ represents the idiosyncratic factor for the specific asset i . Appendix 1 shows some of the results of

such a regression framework. The table shows the multi-factor regressions for the SR1/LR5 parameterisation for all currency pairs analysed. While there is a relatively high degree of statistically significant relationships across all six factors, the test results of White's test for heteroskedasticity as well as the results of the Breusch-Godfrey test suggest that the residuals of these 45 individual regressions exhibit signs of heteroskedasticity as well as serial correlation. While these are aspects that can be adjusted, it is felt that such test setup is not fully appropriate for two reasons. Firstly the chapter looks to identify systematic factors that are common to trading rules across a set of currency pairs, which have equal parameterisations i.e. SR1/LR5. More important, however, is the fact that the residuals $\varepsilon_{i,t}$ from the OLS regression exhibit strong signs of cross-sectional heteroskedasticity and contemporaneous correlation. Appendix 2 shows the results of a test for cross sectional heteroskedasticity. Appendix 3 shows a correlation matrix of OLS residuals. This chapter aims to find a more general regression setup, which adjusts for both aspects, cross-sectional heteroskedasticity and contemporaneous correlation, as well as the problems of standard heteroskedasticity and autocorrelation in the residuals. This can be achieved by utilising a GMM framework proposed by Hansen (1982). The key assumption to achieve this is to verify that risk factors Trend, Momentum, Carry, Value, Risk Aversion and Volatility, which are used to analyse trading rule returns, are exogenous. From a purely academic point of view one could argue that, as illustrated earlier, under the assumption of Lo's (2004) Adaptive Market Hypothesis, which suggests that financial markets follow evolutionary paths, there are no exogenous variables in financial markets. However, when looking at the results in Appendix 4, which show the output of the Durbin-Wu-Hausman test, it is appropriate to assume that the instrumental variables, Trend, Momentum, Carry, Value, Risk Aversion and Volatility are exogenous. Hence, the assumption of exogeneity of risk factors holds in the context of this test framework. Moreover, the GMM model presented in this chapter allows for cross equation restriction on parameters. This means that for each of the earlier described risk factors a single beta factor that describes the general sensitivity of the specific risk factor to the universe of trading rules that are calculated for each trading rule parameterisation, can be calculated. As indicated earlier this chapter calculates 45 trading rules for each trading rule parameterisation, including all G10 currency pairs after adjusting for reverse crosses.

The approach maps an OLS regression into the GMM framework²⁶. This facilitates the construction of an asymptotic distribution that corrects for both serial correlation and general heteroskedasticity via a HAC estimator for the coefficient covariance function, which controls conditional heteroskedasticity and contemporaneous correlation of error terms under the assumption of general heteroskedasticity and serial correlation in the residuals. Under the assumption of exogeneity of factors the GMM specification of a one-step GMM yields results from a least squares estimator, if the orthogonality conditions as introduced in the following equations are satisfied. Equations 8 to 10 follow the notation of Cochrane (2004) closely. In very general terms an OLS regression looks to set the sensitivity parameter beta in such way that the variance of residuals is minimised, as given in Equation 6

$$(6) \quad \min_{\beta_j} E_T \left[(y_{i,t} - \beta_j' d_{j,t})^2 \right]$$

$\hat{\beta}_j$ is then derived from the orthogonality condition as given in Equation 7

$$(7) \quad g_T(\beta_j) = E_T[d_{j,t}(y_{i,t} - d_{j,t}'\beta_j)] = 0$$

As indicated earlier, due to the fact that the number of moments is equivalent to the number of parameters, the condition is identified and can be solved analytically as shown in Equation 8

$$(8) \quad \hat{\beta}_j = [E_T(d_{j,t}d_{j,t}')]^{-1} E_T(d_{j,t}y_{i,t})$$

This results in

$$(9) \quad d \equiv -E_T(d_{j,t}y_{i,t})$$

$$(10) \quad f(d_{j,t}, \beta_j) = d_{j,t}(y_{i,t} - d_{j,t}'\beta_j) = d_{j,t}\varepsilon_{i,t}$$

Then the standard GMM formulation becomes:

²⁶ I thank Dr. Elena Kalotychou (Cass Business School), for her help in validating the appropriateness of the GMM test framework

$$(11) \quad a_T g_T(\hat{\beta}_j) = 0$$

where a_T is a matrix that defines which linear combinations of $g_T(\hat{\beta}_j)$ are set to zero, whereby $g_T(\hat{\beta}_j)$ is defined as

$$(12) \quad g_T(\beta_j) = \frac{1}{T} \sum_{t=1}^T f(d_{j,t}, \beta_j)$$

The standard errors of the estimate $\hat{\beta}_j$ are then defined as

$$(13) \quad \sqrt{T}(\hat{\beta}_j - \beta_j) \rightarrow N[0, (ad)^{-1} aSa'(ad)^{-1}]$$

where d is given in Equation 9 with a and S defined in Equations 14 and 15

$$(14) \quad a \equiv \text{plim } a_T$$

$$(15) \quad S \equiv \sum_{v=-\infty}^{\infty} E[f(d_{j,t}, \beta_j), f(d_{j,t-v}, \beta_j)']$$

whereby $f(d_{j,t}, \beta_j)$ is defined as per Equation 10. Following the described substitutions the OLS standard errors can be defined as per Equation 16

$$(16) \quad \text{var}(\hat{\beta}_j) = \frac{1}{T} E_T(d_{j,t} d_{j,t}')^{-1} \left[\sum_{v=-\infty}^{\infty} E(\varepsilon_{i,t} d_{j,t} d_{j,t-v}' \varepsilon_{i,t-v}) \right] E_T(d_{j,t} d_{j,t}')^{-1}$$

Following this approach will result in an OLS estimate, which is robust with respect to both serial correlation and general heteroskedasticity while controlling for conditional heteroskedasticity and the contemporaneous correlation of error terms.

C. Empirical Evaluation

As indicated in an earlier section the aim of this study is to assess whether the returns derived from applying generic trading rules are due to market inefficiencies, or if they are a mere compensation for systematic risk taking. The presented GMM framework

facilitates this, due to the fact that it allows the construction of robust beta factors that describe the general sensitivity of each trading rule parameterisation, i.e. SR1/SR5 across all analysed currency pairs. The results of this analysis are shown in Figure 3-3. The table is structured such way that in the first column the various trading rule parameterisations are shown, while the next six columns show the GMM estimates (including constant) with their respective statistical significances. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%. In the last two columns the J-statistic and respective P-value for the GMM estimation are given. With respect to the last two columns it has to be pointed out that the results of the J-statistics are meaningless within the context of the conducted GMM estimation, as it is a statistical measure that looks at the degree of over-fitting. Over-fitting occurs if the number of moment conditions exceeds the number of parameters estimated. This becomes particularly relevant when the parameters are assumed to be endogenous. However the present paper assumes the parameters in the GMM framework to be exogenous. Hence, the number of moment conditions equals the number of parameters. This leads to an analytical solution, therefore the J-stats are very low, suggesting no signs of over-fitting of the model.

With respect to the remaining test results, there is clear evidence of statistical significance. In the case of the Trend, Momentum and Risk Aversion factors; this statistical significance is very high. The same is the case for the constant. Most of the factors exhibit a statistical significance at the 1% level and only very few parameters are not statistically significant. Amongst the other factors used, Volatility is the only factor that exhibits some statistical significance, while the factors Carry and Value do not show any statistically significant results. In terms of the signs of the beta factors the results are very intuitive. The constant is positive across all moving average combinations. The results in Figure 3-3 suggest that there is some alpha embedded in the technical trading strategies analysed in this chapter. When it comes to the Trend and Momentum factors all trading strategies exhibit a strong positive relationship, which makes intuitive sense given the nature of these factors. The trend index used is deliberately constructed such way that it replicates the returns that can be expected from an average trend following strategy. While the Momentum factor, as shown in Figure 3-2, exhibits comparatively low correlation to the Trend factor, the spirit behind the Momentum factor lies also in the continuation of trends, hence a positive relationship between the risk factor and trading rule returns is expected.

FIGURE 3-3: GMM ESTIMATION RESULTS

MOVAV	CONSTANT	TREND	MOMENTUM	CARRY	VALUE	RISKAV	VOLA	J-STAT	P-VALUE
SR1/LR5	0.34% ***	0.87%	1.86%	-2.46%	0.75%	4.50% **	-0.95%	0.0000001	(0.999)
SR1/LR10	0.25% ***	2.94% **	2.28%	-2.22%	0.59%	5.39% **	-0.55%	0.0000001	(0.999)
SR1/LR15	0.21% ***	3.83% ***	2.64%	-1.62%	0.55%	6.21% ***	-0.61%	0.0000001	(0.999)
SR1/LR20	0.18% ***	4.83% ***	3.16% *	-1.50%	0.58%	6.51% ***	-0.14%	0.0000001	(0.999)
SR1/LR25	0.17% ***	5.62% ***	4.12% **	-1.77%	0.74%	6.26% ***	0.23%	0.0000001	(0.999)
SR1/LR30	0.15% ***	6.44% ***	4.68% ***	-1.54%	0.82%	6.17% ***	0.66%	0.0000001	(0.999)
SR2/LR5	0.18% ***	1.45%	2.06%	-2.13%	0.50%	5.67% ***	-4.74% ***	0.0000001	(0.999)
SR2/LR10	0.13% ***	4.33% ***	2.80%	-1.96%	0.44%	6.00% ***	-3.06% **	0.0000001	(0.999)
SR2/LR15	0.11% ***	5.31% ***	3.26% *	-2.15%	0.46%	5.43% ***	-2.94% **	0.0000001	(0.999)
SR2/LR20	0.09% ***	6.24% ***	4.06% **	-2.28%	0.64%	5.53% **	-2.16%	0.0000001	(0.999)
SR2/LR25	0.08% ***	7.29% ***	4.64% ***	-2.20%	0.93%	5.52% **	-2.03%	0.0000001	(0.999)
SR2/LR30	0.08% ***	7.67% ***	5.25% ***	-1.61%	0.87%	5.51% **	-1.80%	0.0000001	(0.999)
SR3/LR5	0.12% ***	2.23% *	2.67% *	-1.93%	0.79%	5.35% ***	-2.35% *	0.0000001	(0.999)
SR3/LR10	0.09% ***	5.12% ***	3.19% *	-2.41%	0.47%	4.81% ***	-1.88%	0.0000001	(0.999)
SR3/LR15	0.08% ***	5.93% ***	3.53% **	-2.33%	0.61%	4.64% **	-1.90%	0.0000001	(0.999)
SR3/LR20	0.06% ***	6.87% ***	4.17% **	-1.91%	0.69%	4.96% **	-1.63%	0.0000001	(0.999)
SR3/LR25	0.06% ***	7.33% ***	4.98% ***	-1.43%	0.71%	5.03% **	-1.70%	0.0000001	(0.999)
SR3/LR30	0.05% ***	8.17% ***	5.12% ***	-1.03%	0.92%	5.10% **	-1.74%	0.0000001	(0.999)
SR4/LR5	0.09% ***	3.24% ***	3.22% **	-2.23%	1.18%	4.60% **	-1.13%	0.0000001	(0.999)
SR4/LR10	0.07% ***	5.79% ***	3.31% **	-1.97%	0.68%	4.31% **	-1.84%	0.0000001	(0.999)
SR4/LR15	0.06% ***	6.14% ***	3.77% **	-1.89%	0.87%	4.22% **	-1.89%	0.0000001	(0.999)
SR4/LR20	0.05% ***	7.04% ***	4.49% ***	-1.28%	0.67%	4.58% **	-1.74%	0.0000001	(0.999)
SR4/LR25	0.05% ***	7.60% ***	4.97% ***	-1.14%	0.79%	4.63% **	-1.71%	0.0000001	(0.999)
SR4/LR30	0.04% ***	8.10% ***	5.10% ***	-0.76%	1.10%	4.67% **	-1.78%	0.0000001	(0.999)
SR5/LR10	0.06% ***	5.45% ***	3.67% **	-1.84%	0.68%	4.06% **	-1.97%	0.0000001	(0.999)
SR5/LR15	0.05% ***	6.19% ***	4.05% **	-1.61%	0.86%	3.92% **	-1.93%	0.0000001	(0.999)
SR5/LR20	0.04% ***	7.29% ***	4.63% ***	-1.20%	0.60%	4.25% **	-1.76%	0.0000001	(0.999)
SR5/LR25	0.04% ***	7.76% ***	5.07% ***	-0.94%	0.82%	4.31% **	-2.11% *	0.0000001	(0.999)
SR5/LR30	0.04% ***	8.11% ***	5.34% ***	-0.40%	1.01%	4.40% **	-2.04% *	0.0000001	(0.999)
SR10/LR15	0.03% ***	6.48% ***	3.37% **	-1.05%	0.75%	2.82% *	-2.50% **	0.0000000	(0.999)
SR10/LR20	0.02% ***	7.65% ***	4.58% ***	-0.62%	0.82%	2.92% **	-3.04% ***	0.0000000	(0.999)
SR10/LR25	0.02% ***	7.79% ***	5.32% ***	-0.46%	0.69%	3.40% **	-2.33% ***	0.0000000	(0.999)
SR10/LR30	0.02% ***	8.12% ***	5.66% ***	0.14%	0.84%	3.75% **	-1.86% **	0.0000001	(0.999)
SR15/LR20	0.02% ***	7.62% ***	5.43% ***	-0.95%	0.31%	2.53% *	-1.99% **	0.0000000	(0.999)
SR15/LR25	0.02% ***	7.79% ***	5.94% ***	-0.15%	0.31%	3.37% *	-1.81% *	0.0000001	(0.999)
SR15/LR30	0.02% ***	8.05% ***	5.90% ***	0.59%	0.64%	3.55% *	-1.26%	0.0000001	(0.999)
SR20/LR25	0.02% ***	7.47% ***	6.07% ***	0.10%	0.36%	3.12%	-1.18%	0.0000001	(0.999)
SR20/LR30	0.01% ***	7.63% ***	6.07% ***	1.12%	0.59%	3.32%	-0.95%	0.0000001	(0.999)
SR25/LR30	0.01% ***	7.70% ***	5.89% ***	2.63%	0.47%	3.26%	-0.85%	0.0000001	(0.999)

The first column in the table shows the various trading rule parameterisations. The next six columns show the GMM estimates (including constant) with their respective statistical significances. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%. The last two columns show the J-statistic and respective P-value for the GMM estimation.

The Carry Factor, as mentioned earlier, does not yield any statistically significant results, yet it exhibits a negative sensitivity to most of the moving average trading rules, which confirms previous research such as Pojarliev and Levich (2010), who find that in time periods where trend is very popular, Carry tends to be less popular and vice versa. A rationale for the negative relationship between the Carry Factor and trading rule returns might be found in Brunnermeier, Petersen and Nagel (2008), who suggest that carry strategies are negatively skewed, which can largely be attributed to funding risk, due to the high leverage in currency trading strategies. In time periods where carry strategies sell off sharply, short term trading strategies adapt quickly to new trends and tend to exhibit positive returns. Hence they are not affected by the funding risk that is present in carry strategies. This link between the adaptability of trading strategies and funding risk is also reflected in the results of this chapter. Longer-term moving averages

such as the SR25/LR30, which clearly exhibit less adaptability than most of the other technical trading rules investigated, have a mildly positive sensitivity to the carry. In the case of the SR25/LR30 the sensitivity is 2.63%, however this is not statistically significant.

The Value Factor is mildly positive, yet not statistically significant. The sensitivities are in the range of 0.31% to 1.18%. Overall the sensitivities of this factor to trading rule results do not allow for any meaningful conclusion.

When it comes to Risk Aversion, the results are positive and fairly strong, suggesting that moving average trading rules mimic the returns of a safe haven strategy. This, again links into the adaptability argument whereby trading strategies tend to exhibit negative performance in more benign market time periods, due to the high level of turnover and transaction cost that come with it. However, in time periods of risk aversion, they tend to perform strongly due to their ability to adapt. One can compare this dynamic to currencies such as the Japanese Yen and the Swiss Frank, which tend to appreciate in times of stress, and depreciate in quiet time periods due to the lack of carry that these currencies bear.

The final factor, Volatility, also exhibits a mildly negative relationship to trading rule returns, with some statistically significant values. Most of the sensitivities are in the range of -1% to -2%. The statistically significant values cluster around the SR2/LR5 and SR2/LR15 as well as the SR10/LR15 and the SR10/LR30 moving average combinations. There are two trading rule parameterisations that exhibit a positive sensitivity, albeit very low and not statistically significant. The overall result is again not surprising, given the fact that short term focused trading strategies switch continuously between long and short exposure. Hence, in times of high volatility, when the variability between positive and negative returns increases, trading rule returns bear the risk of getting whiplashed. Therefore they exhibit negative sensitivity. This result counters the popular notion that trading rule returns are essentially long volatility strategies. This might be because the average currency trader might well associate market volatility with a spike in safe haven currencies as opposed to an increase in the variability of returns. The highly significant positive relationship to the Risk Aversion factor provides evidence for that assumption.

While Figure 3-3 provides some insight into the statistical relationships between technical trading rules and the risk factors proposed in this chapter, it is not very user

friendly, as it does not explicitly show the dynamics of factor sensitivities across trading rule parameterisations. For this reason Figure 3-4 reorganizes the results presented in Figure 3-3 in a visually more accessible way. Figure 3-4 is split into seven parts, which are equally structured. The columns represent the short term moving average parameters and the rows represent the long term moving average parameters. The first part shows the constant and the remaining six parts refer each to the sensitivities of the different trading rules to the proposed risk factors. The colour code is designed in such a way that it assigns different colour shades across different percentiles, which are calculated across all six parameters, excluding the constant. The rationale for excluding the constant from this analysis is the fact that it is not a sensitivity to a systematic risk factor. The colour coding is calculated across all factor sensitivities, giving a visual impression of the overall strength of sensitivities to single factors. The parameters that exhibit the highest positive values are shaded in dark blue, while the parameters with the highest negative values are shaded in dark red. The colour index is given in Figure 4. The rationale for creating a colour scheme across all parameters is that it provides a strong visual impression of the level of sensitivity between risk factors and technical trading rules.

As mentioned earlier, the constant is highly statistically significant across all moving average combinations. While Figure 3-3 did make this result fairly evident, what has not been fully shown in the previous analysis is the distribution of constant terms across trading rule parameterisations. Figure 3-4 shows that very short term moving average combinations exhibit a considerably higher alpha than longer term moving average combinations. In the case of the SR1/LR5 combination the constant term has a value of 0.34% while the constant of the SR5/LR30 combination only exhibits a value of 0.01%. Both of these numbers are per day alphas. However, as mentioned earlier, a part of that difference might well be explained by the fact that no transaction costs are assumed in this analysis, which, due to the high level of turnover of shorter term trading rules might well reduce the relative level of alpha. Nonetheless, the results do suggest that short-term trading rules deliver returns that cannot be explained by the systematic risk factors presented in this chapter. This finding is broadly in line with the observations of the first chapter, which finds that moving average crossover signals that utilise a set of very short-term moving average combinations outlive what is suggested by theory, while long-term moving average crossover signals' life expectancy is shorter than theory would suggest. When looking at the other results, the three key factors that exhibit high

statistical significance, and comparatively high positive sensitivity are the Trend, Momentum and Risk Aversion factors.

FIGURE 3-4: GMM ESTIMATION RESULTS, CONSTANT AND SENSITIVITIES OF TRADING RULE RETURNS TO RISK FACTORS

Coefficient(CONSTANT)							Coefficient (VALUE)						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30		LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	0.34%***	0.25%***	0.21%***	0.18%***	0.17%***	0.15%***	SR 1	0.75%	0.59%	0.55%	0.58%	0.74%	0.82%
SR 2	0.18%***	0.13%***	0.11%***	0.09%***	0.08%***	0.08%***	SR 2	0.5%	0.44%	0.46%	0.64%	0.93%	0.87%
SR 3	0.12%***	0.09%***	0.08%***	0.06%***	0.06%***	0.05%***	SR 3	0.79%	0.47%	0.61%	0.69%	0.71%	0.92%
SR 4	0.09%***	0.07%***	0.06%***	0.05%***	0.05%***	0.04%***	SR 4	1.18%	0.68%	0.87%	0.67%	0.79%	1.1%
SR 5		0.06%***	0.05%***	0.04%***	0.04%***	0.04%***	SR 5		0.68%	0.86%	0.6%	0.82%	1.01%
SR 10			0.03%***	0.02%***	0.02%***	0.02%***	SR 10			0.75%	0.82%	0.69%	0.84%
SR 15				0.02%***	0.02%***	0.02%***	SR 15				0.31%	0.31%	0.64%
SR 20					0.02%***	0.01%***	SR 20					0.36%	0.59%
SR 25						0.01%***	SR 25						0.47%

Coefficient (TREND)							Coefficient (RISKAV)						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30		LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	0.87%	2.94%**	3.83%***	4.83%***	5.62%***	6.44%***	SR 1	4.5%**	5.39%**	6.21%***	6.51%***	6.26%***	6.17%***
SR 2	1.45%	4.33%***	5.31%***	6.24%***	7.29%***	7.67%***	SR 2	5.67%***	6%**	5.43%***	5.53%**	5.52%**	5.51%**
SR 3	2.23%*	5.12%***	5.93%***	6.87%***	7.33%***	8.17%***	SR 3	5.35%***	4.81%***	4.64%**	4.96%**	5.03%**	5.1%**
SR 4	3.24%***	5.79%***	6.14%***	7.04%***	7.6%***	8.1%**	SR 4	4.6%**	4.31%**	4.22%**	4.58%**	4.63%**	4.67%**
SR 5		5.45%***	6.19%***	7.29%***	7.76%***	8.11%***	SR 5		4.06%**	3.92%**	4.25%**	4.31%**	4.4%**
SR 10			6.48%***	7.65%***	7.79%***	8.12%***	SR 10			2.82%*	2.92%**	3.4%**	3.75%**
SR 15				7.62%***	7.79%***	8.05%***	SR 15				2.53%*	3.37%*	3.55%*
SR 20					7.47%***	7.63%***	SR 20					3.12%	3.32%
SR 25						7.7%***	SR 25						3.26%

Coefficient (MOMENTUM)							Coefficient (VOLA)						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30		LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	1.86%	2.28%	2.64%	3.16%*	4.12%**	4.68%***	SR 1	-0.95%	-0.55%	-0.61%	-0.14%	0.23%	0.66%
SR 2	2.06%	2.8%	3.26%*	4.06%**	4.64%***	5.25%***	SR 2	-4.74%***	-3.06%**	-2.94%**	-2.16%	-2.03%	-1.8%
SR 3	2.67%*	3.19%*	3.53%**	4.17%**	4.98%***	5.12%***	SR 3	-2.35%*	-1.88%	-1.9%	-1.63%	-1.7%	-1.74%
SR 4	3.22%**	3.31%**	3.77%**	4.49%***	4.97%***	5.1%***	SR 4	-1.13%	-1.84%	-1.89%	-1.74%	-1.71%	-1.78%
SR 5		3.67%**	4.05%**	4.63%***	5.07%***	5.34%***	SR 5		-1.97%	-1.93%	-1.76%	-2.11%*	-2.04%*
SR 10			3.37%**	4.58%***	5.32%***	5.66%***	SR 10			-2.5%**	-3.04%***	-2.33%***	-1.86%**
SR 15				5.43%***	5.94%***	5.9%***	SR 15				-1.99%**	-1.81%*	-1.26%
SR 20					6.07%***	6.07%***	SR 20					-1.18%	-0.95%
SR 25						5.89%***	SR 25						-0.85%

Coefficient (CARRY)							Percentiles						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30		99.5%	99.0%	97.5%	95.0%	85.0%	50.0%
SR 1	-2.46%	-2.22%	-1.62%	-1.5%	-1.77%	-1.54%							
SR 2	-2.13%	-1.96%	-2.15%	-2.28%	-2.2%	-1.61%							
SR 3	-1.93%	-2.41%	-2.33%	-1.91%	-1.43%	-1.03%							
SR 4	-2.23%	-1.97%	-1.89%	-1.28%	-1.14%	-0.76%							
SR 5		-1.84%	-1.61%	-1.2%	-0.94%	-0.4%							
SR 10			-1.05%	-0.62%	-0.46%	0.14%							
SR 15				-0.95%	-0.15%	0.59%							
SR 20					0.1%	1.12%							
SR 25						2.63%							

Percentiles						
	99.5%	99.0%	97.5%	95.0%	85.0%	50.0%
	25.0%	10.0%	5.0%			

The Figure is split into seven sections, which are equally structured. The first section on the left shows the constant of the GMM regression. Each of the other sections exhibits the sensitivities of trading rules to risk factors, whereby the columns represent the short term moving average parameters and the rows represent the long term moving average parameters. The colour code is designed such way that it assigns different colour shades across different percentiles, legend of the colour index is given in the figure. The percentiles of the colour code are calculated using the overall universe of factor sensitivities, excluding the regression constant. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

Amongst these factors, Trend is the most pronounced factor with the highest level of sensitivity, while the Risk Aversion and Momentum factors exhibit similar levels of

sensitivity. In terms of the distribution of sensitivities across risk the Trend and Momentum factors follow a similar pattern whereby shorter-term moving averages show lower sensitivities to the risk factors. In the case of the SR1/LR5 trading rule combination, Trend has a sensitivity of +0.87% while Momentum has a sensitivity of +1.86%. This increases as the length of longer term and shorter term moving averages increase. In the case of Trend, the trading rule with the highest sensitivity is the SR3/LR30 trading rule, which equates to +8.17%. The highest sensitivity for the Momentum factor is +6.07%, which is the sensitivity for the SR20/LR25 and the SR25/LR30. While the patterns of statistical significance are very similar, there are some subtle differences between both factors. When it comes to Trend, the difference between short term and long term moving average tends to determine the level of sensitivity. This is illustrated very clearly in the LR30 column of the Trend section. Here the sensitivities increase from SR1 to SR3 and then fall off again. In the Momentum section, column LR30 sees a continuous increase in sensitivity from SR1 up to SR20 and then a fall thereafter. While the distribution of sensitivities to the Trend factor might well be explained by the construction of the Trend factor, which combines short term and long term moving averages, the distribution of sensitivities of the factor Momentum is somewhat more difficult to explain. It might well stem from the slower moving nature of trading rules that have longer SR moving averages as well as longer LR moving averages in their signal generation.

This means that they adapt only very slowly to a changing market environment and remain unaffected by short-term reversals in exchange rates. They only reverse when more powerful trend changes are happening. Therefore these trading rule parameterisations become more akin to the Momentum factor, which looks to go long past winners, while going short previous losers, aiming to capture a continuation of trends.

The Risk Aversion factor exhibits again a somewhat different pattern of sensitivities across trading rule parameterisations than the other two factors. The highest level of sensitivity can be found in the SR1 row. The sensitivity increases from the LR5 to the LR20 and decreases thereafter. The sensitivities of any other SR combinations are considerably lower. The highest sensitivity to this factor is +6.51%, which is borne by the SR1/LR20 combination. As argued by Schulmeister (2006), this might be because market participants use trading rules as heuristics in their assessment of markets. The SR1/LR20 is in that context a particularly followed moving average combination, given

the fact that in trading days this moving average combination equates to the one day / one month signal. The same market participants perceive the JPY and the CHF as safe haven currencies, which they tend to buy in times of economic and/or financial market stress. Therefore, the high sensitivity between this particular moving average combination and the Risk Aversion factor might well stem from psychological aspects.

While the Carry factor does not exhibit any statistically significant relationship the results shown in Carry section illustrate the dynamics mentioned earlier well. In time periods where carry strategies sell off sharply, short-term trading strategies adapt quickly to new trends and tend to exhibit a positive return. Hence, they are not affected by the funding risk that is present in carry strategies, as suggested by Brunnermeier, Petersen and Nagel (2008). The trading rule that exhibits the highest negative sensitivity of -2.46% is the SR1/LR5 trading rule, which is arguably the trading rule with the highest adaptability amongst all of the analysed trading rules, while the SR25/LR30 trading rule with a sensitivity of +2.63% is one of the slowest moving trading rules analysed. As mentioned earlier, the results of the factor Value are neither statistically significant, nor very conclusive. This has not changed when assessing the sensitivities to this risk factor as shown in Figure 3-4.

When it comes to the Volatility factor, the analysis in Figure 3-4 provides some valuable insight. While most of the results exhibit mildly negative sensitivities, some of them are statistically significant. However, the statistically significant relationships are scattered around two short term moving average levels SR2 and SR10. This might well be explained by the relative level of whiplash those strategies are exposed to, which as illustrated later, might well stem from the averaging process in the signal generation. In the case of the SR2/LR5 trading rule, the sensitivity is -4.74%, which is the most negative sensitivity to this factor. The SR1/LR5 moving average combination exhibits a negative sensitivity of -0.95%, which is amongst the least negative and not statistically significant. Yet the only difference between both trading rules is the fact that the short-term signal is generated by a two-day average in the case of SR2/LR5 versus a one-day observation in the case of SR1/LR5. In time periods of spiking volatility, price movements are very variable. Assuming that the exchange rate is 100 and has been static for the last few days, a fall of 20% would take the exchange rate to 80. 100% of that fall goes into the SR signal, as it is calculated over one day, while 20% of the fall goes into the LR signal, as it is calculated over five days. Hence, the SR1/LR5 is 80/96. A short signal is established and on the following day the trading signal continues to be

in a short mode, unless the price recovers from 80 to 96, which is again a 20% rise. In the case of the SR2/LR5 trading rule the triggers are somewhat different. 50% of the above-described fall goes into the SR, as it is calculated over two days, while 20% of the fall goes into the LR signal, as it is again calculated over five days. This means that the SR2/LR5 trading signal is 90/96. A short signal is established; the following day the trading signal continues to be in a short mode unless the price recovers to 120, which equates to a 50% move upwards. Assuming, that after a rise in the currency to 120, both trading strategies are long, if the price falls again the following day, under the SR1/LR5 trading rule a fall to 100 would again trigger a short position, while the SR2/LR5 rule would require a fall to 66 to be short again. Hence, the moving average calculation for the SR2/LR5 trading rule sets a very high hurdle rate for switching exposures, which means that the chances of the strategy being on the wrong side in volatile periods is high. This dynamic is considerably less pronounced for SR3/LR5 and even SR4/LR5, hence the considerably lower sensitivities. Following the same argument illustrated earlier, one can also see why the sensitivities in the SR1 row are becoming more and more positive as the LR moving averages increase, given that the hurdle rate of a switch falls as the longer term moving averages increase. This allows the strategy to be more adaptive to volatility changes.

D. Conclusions

This chapter builds on the work of Pojarliev and Levich (2008, 2010), who dissect the returns of active currency managers by applying a multiple OLS regression to currency fund returns. Where the chapter differs is in the specification of the depended variable, which is in the context of the present chapter a set of trading rule parameterisations that are applied to a broad range of currency pairs. Moreover the present chapter extends the work of Pojarliev and Levich (2008, 2010) in as far as it widens the universe of systematic risk factors.

First, the returns from moving average trading strategies deliver modest positive alpha, which is statistically significantly different from zero. The levels of statistical significance are generally high. This can be partially explained by the fact that there are no transaction costs factored into the analysis. Short-term moving averages tend to deliver higher levels of alpha than is the case for longer-term moving averages.

Second, the chapter establishes a comparatively strong, positive and statistically significant link between the risk factors Trend, Momentum, Risk Aversion. The results of the chapter clearly indicate that shorter-term moving averages exhibit less systematic exposure than longer term moving averages. The dynamics of systematic exposures are different for the three risk factors. When it comes to Trend, the difference between short term and long term moving average tends to determine the level of sensitivity. This is not the case for Momentum, where an increase in shorter term and longer term moving average combinations translates into higher sensitivity to the risk factor. For the Risk Aversion factor the highest level of sensitivity can be found in the SR1 row. Whereby the SR1/LR20, which coincides with the one day / one month signal bears the highest sensitivity to the risk factor, which might well be explained by psychological aspects.

Moreover, other factors such as Carry, Value and Volatility have a considerably less pronounced relationship; only a few factor sensitivities are statistically significant. Despite that, some of the dynamics between trading rules and these risk factors provide useful insights. When it comes to Carry, shorter-term moving averages exhibit higher negative sensitivity while longer term moving averages exhibit mildly positive sensitivity. This can be explained by differing levels of adaptability of the trading strategies, which means that shorter term trading rules run a higher chance of accumulating a positive return in time periods of carry crashes, than longer term trading rules. While the results of the factor Value are neither statistically significant, nor very conclusive, the factor Volatility exhibits some statistically significant sensitivity. However, some of these results might well be driven by the construction of the trading rules.

The overall conclusion that can be drawn from these results is that at least a part of the returns from technical trading rules are driven by systematic factors. While Trend, Momentum and Risk Aversion are the most dominant risk factors, very short-term moving averages are less exposed to these factors than longer term moving averages. Paired with the finding that shorter term moving averages exhibit higher levels of alpha, the results in this chapter would suggest that, shorter term moving averages are less affected by systematic risk factors than it is the case for longer moving averages. When looking at the returns of very short-term moving average trading rules such, as the SR1/LR5 day rule, it becomes evident that apart from the constant, which is the alpha contribution from this strategy, there is only one statistically significant risk factor, that being Risk Aversion, which impacts trading rule returns. Given that this factor is a

purely heuristic factor, the fact that it impacts very short term focussed trading rules is a valuable insight into market psychology. Hence, one might deduct that some of the trading rule returns are genuinely driven by human spirit and market inefficiency as opposed to risk taking. When it comes to SR1/LR30 day rule, Trend, Momentum, and Risk Aversion are statistically significant, while Volatility has a mildly positive sensitivity. Looking at the SR25/LR30 day rule only Trend and Momentum are statistically significant factors, while volatility has a negative sensitivity, carry does have a positive sensitivity to the risk factor. The fact that these three trading rule parameterisations, which are supposedly very homogeneous in terms of their design are subject to different risk factors allows for a second conclusion, in as far as trading rule returns are not only subject to systematic risk factors, but also that slight differences in parameterisation or design of trading rules exposes them to very different types of risk. The insight that systematic risk exposures of trend following trading strategies change with small adjustments in the design of trading rules is profound.

Areas of future research might be to extend the trading rule parameterisations, or to alter the framework of the analysis to more sophisticated trading rules. This chapter uses very basic trading rules as benchmarks, which may have become so popular that the systematic risk component has become very high within these trading rules. Looking at other, more sophisticated trading rules, might lead to very different results. Other areas of research might be to look at broader universe of foreign exchange markets. It would be insightful to understand to which degree trading rule returns composed of emerging market currencies are subject to the same universe of systematic risk factors as it is the case for trading rules that are based on G10 currencies. Recent studies such as Menkhoff, Sarno, Schmeling and Schrimpf (2011) find that most of the returns from momentum type trading rules come from emerging market currencies. This finding is confirmed by and Chong and Ip (2009), reporting 30% plus annualised returns from trading strategies that utilize a momentum based trading strategy in emerging market currencies. Hence, understanding whether trading rule returns in emerging markets are driven by risk factors or not, might shed some light on the question whether there is still a “free lunch” for currency traders in these markets.

VII. Overall Conclusion:

The first chapter of this thesis introduces a methodology that applies a variation of survivorship analysis. The aim of this is to compare the probability of occurrence of positive or negative return streams in an empirical time series with a theoretically derived probability. Empirical momentum signals either outlive benchmark signals, as is the case for moving average crossover signals that utilise a set of very short-term moving average combinations, or momentum signals created from empirical curves, have lower life expectancy than theory would suggest as is the case for some longer-term moving average crossover signals. The results of a sub-sample analysis suggest that most of the deviations from market efficiency deteriorate over time, up until the point where all of the momentum signals exhibit survival times that are statistically equivalent to what is suggested by benchmark processes. Moreover, when implementing trading rules on the same set of moving average crossover signals, it becomes evident that profitability of a generic trading rule that incorporates all moving average signals deteriorates continuously to a point where the trading rule becomes unprofitable. Furthermore, a trading strategy that is constructed from a sub-set of moving average signals, namely shorter-term moving average signals, shows clear outperformance over a trading strategy that is generically composed from all moving average crossover signals. This outperformance persists over time.

The second chapter extends the first chapter. It aims to search for a superior trading rule. Survivorship analysis provides a wide range of information about historic survival patterns of moving average trading rules, which can be used to establish the best exit points of a trading strategy. This chapter investigates a series of trading strategy enhancements. Similar to the results of the first chapter, the second chapter also indicates that the profitability of generic trading rules diminishes over time. Moreover the results also indicate that during the early years of the data sample, when the general trading rule profitability is high, the performance of trading rule enhancements is somewhat mixed. While the trading rule enhancement that weighs strategy exposures according to the conditional historic survival probability is able to add some value, the trading rule enhancement that weighs exposures according to the unconditional survival probability doesn't. This changes in the latter parts of the sample period where the enhancement that weighs strategy exposures according to the conditional historic survival probability fails to add value, while the enhancement that weighs strategy

exposures according to the unconditional historic survival probability performs strongly. The results of the second chapter indicate that trading rule returns exhibit two distinct regimes, which suggests that foreign exchange markets have changed over time.

The third chapter sheds light on whether the returns derived from applying generic technical trading rules embed some compensation for systematic risk taking. While factors such as Trend and Momentum and Risk Aversion have a relatively strong positive and statistically significant impact on trading rule returns. It should be noted, however, that the systematic exposure, albeit present, is less pronounced for shorter term moving averages, while as longer term moving averages exhibit more systematic exposure. Other factors such as Carry, Value and Volatility have a considerably less pronounced relationship to trading rule returns. Only few factor sensitivities are statistically significant. The results of this chapter make a strong case for the fact that at least a part of the returns from technical trading rules are driven by systematic factors. Paired with the finding that shorter term moving averages exhibit higher levels of alpha, the results in this chapter would suggest that shorter term moving averages are less affected by systematic risk factors than it is the case for longer moving averages. This becomes evident when analysing the results of very short term trading rules such as the SR1/LR5 rule are only influenced by the factor Risk Aversion, which is a purely heuristic factor.

One of the main conclusions that can be drawn from this thesis is that the profitability of generic trading rules has continuously diminished over time. The same is the case for deviations from market efficiency. As mentioned earlier, this is not the case for the returns of shorter-term moving averages, which remain generally higher even when more generic trading rules fail to perform. Within that context the results of the second chapter shed some light on the dynamics of technical trading rules. One of the key observations of this chapter is the fact that during the early years of the data sample, when general trading rule profitability is high, the scope of trading rule enhancements to outperform is somewhat limited. This, however, changes as the level of general trading rule profitability deteriorates.

Namely, one of the trading rule enhancements, which weights its exposures according to the historic unconditional survival probability of moving average crossover trading rule signals, shows strong results in the second half of the data sample. This allows for two conclusions.

Firstly, the results of the second chapter clearly point to a regime change in foreign exchange markets. Moreover, bearing in mind that the discussed trading rule enhancement reduces the exposure level of an established trading rule signal very quickly, the regime change has been in favour of technical trading rules that are shorter-term in their nature.

The results of the third chapter indicate that even shorter term focussed technical trading rule returns, with the exception of very short term focused trading rule returns, are still to a fair extent impacted by systematic risk factors. In this context the results of the second chapter, which suggest very weak returns of the trading rule enhancement that shorten exposure times, can be explained as follows. During the early part of the data sample, where pronounced trends persisted and deviations from market efficiency were high, as indicated in the first chapter, longer term trading rule returns might have been higher due to lower cost of rebalancing. As trends have faded away the aspect of adaptability has become more important. Hence, trading rules with shortened exposure perform better. Such explanation, however, does not explain the erosion of trading rule profitability in the wider universe of technical trading rules, which also have some embedded systematic risk factors, as shown in chapter three. For this reason another line of argument, which is in the spirit of Lo's (2004) Adaptive Market Hypothesis, can be presented.

The market environment is subject to continuous change as new investor "species" enter the market place, new investment opportunities become profitable while existing investment opportunities cease to exist. Therefore, different trading styles might be en vogue at different times. In the context of foreign exchange markets three observations can be put forward in favour of such line of argument.

First, competition amongst investors that follow systematic trading strategies within foreign exchange markets has increased considerably. This has also led to a change in the way of how investors compete in the foreign exchange market and how foreign exchange markets operate. Barclay Hedge²⁷, a database provider for systematic trading funds, estimates that in 1980 the assets under management of systematic trading funds were in the range of US\$ 300m. This grew to US\$ 38bn by the end 2000. As of the end of 2009, when the data sample of this chapter ends, approximately US\$ 214bn were managed in systematic trading funds. Currently this number is over US\$ 330bn. These

²⁷ See: http://www.barclayhedge.com/research/indices/cta/mum/CTA_Fund_Industry.html

numbers only represent the proportion of assets managed in explicit fiduciary mandates with clients. The actual amount of assets managed in systematic trading algorithms is likely to be much higher. Since 1998 the Bank of International Settlements publishes a triennial survey of Foreign Exchange and Derivatives market activity. Besides this survey, a series of working papers are published that shed light on the drivers in the change of trading volume. Galati and Melvin (2004) point out that since 2001 there has been a surge in foreign exchange trading. Galati and Melvin (2004) also highlight the significant growth in the participation of Hedge Funds, in particular trend following strategies, which have considerably grown in numbers. While reiterating the aspect of Hedge Fund participation, Galati and Heath (2007) also point towards the aspect of algorithmic trading as one of the key sources of turnover within the foreign exchange markets. King and Rime (2010) make particular reference to the concept of algorithmic trading in their analysis of foreign exchange volumes in the years from 2007 to 2010. With algorithmic trading, King and Rime (2010) refer to systems that break up trades to optimise trade execution, or automated hedging by market makers or other forms of proprietary technical trading. They estimate that high frequency trading takes up to 25% of the volume of all spot transactions worldwide. From this it becomes evident that the concept of systematic trading has dramatically evolved over the past ten to fifteen years. The growth in systematic trading strategies has come not only from trend following funds, which have seen spectacular growth over that time period, but also from a profound change in how foreign exchange markets operate.

Second, the wide acceptance of technical trading rules by the investment community has changed investment behaviour. This has evolved to the degree that technical trading rules have become systematic risk factors in their own right. Such thesis is supported by the findings of Osler (2003), which looks at the microstructural aspects of order books in foreign exchange markets. Osler (2003) argues that “support” and “resistance” levels can be key indicators for accelerated momentum or reversals, depending on whether they are broken or not. The key aspect behind this thesis is the distribution of the placement of stop-loss and take-profit orders by clients. Take-profit orders, designed to lock in profits, are mostly clustered around round numbers. Stop-loss orders, which are designed to cut losses, tend to be placed just beyond round numbers. A further investigation using bootstrap simulations reaffirms the idea that there is a self-fulfilling dynamic between order placement and exchange rate dynamics. Hence, technical analysis might be a fully rational method of exploiting the institutional features of currency markets. These arguments are supported by the findings of Schulmeister

(2006), who looks at the predictive power of aggregate trading signals. The results of his analysis suggest that when markets change direction, the majority of trading filters in his study tend to be on the same side, i.e. they are either long or short. His results indicate that there is a pronounced feedback mechanism between trading rules and movements in the underlying exchange rates. Therefore, technical trading rules act as a multiplier, translating small news flows into a market trend. In the spirit of Lo's (2004) one could argue that over time, as investors had to adapt to the changing market environment in the foreign exchange market, they have become aware of the self-enforcing link between trading rules and market trends. Indeed Poljarliev and Levich (2008) indicate that various technical trading rules or carry strategies have become so popular they have become risk factors in their own right. Poljarliev and Levich (2010) make reference to this in as far as suggesting that the volatility and correlation characteristics of currencies change as consequences of changing investor preferences. These results link very well with Lo's (2004) Adaptive Market Hypothesis, where different investment strategies can go through different stages of profitability. This is due to the fact that in Lo's (2004) view, there are different "species" of market participants, which have distinguishable characteristics in terms of their investment pattern. Hence, market performance is not only driven by a change in overall market preferences, but also in a change of the composition of market participants competing in a specific asset class. Hence, it should not come as a surprise that some aspects of markets, which have initially been viewed as market anomalies, have over time become systematic risk factors in their own right.

Finally, the lack of a short-term valuation framework for currencies exacerbates trending characteristics as well as the volatility of foreign exchange markets. As indicated in the introduction, the dividend yield of stocks gives a timely signal to investors whether a stock is cheap or expensive. Exchange rates don't have that same concise valuation framework. The valuation relationships such as the purchasing power parity or interest rate differentials are either loose relationships or the equating factor, which indicates cheapness or expensiveness is missing. While the dividend yield of a stock comes down as the stock price goes up, an appreciation of a currency is not automatically linked to a narrowing in the interest rate differential. This is due to the fact that interest rates are set by central bank policy, which does not necessarily change as a result of the valuation of an exchange rate. The profound implication of this missing anchoring device is the fact that exchange rates are prone to trend much more than other financial assets. This is particularly beneficial for short-term moving average

trading rules, given that they can adapt very quickly to trend changes, allowing these trading rules to benefit from very sharp reversals and relatively high day to day volatility.

The fact that currency markets have changed profoundly, the fact that market participants perceive systematic risk factors differently than they used to, and the fact currencies exhibit a high level of day to day volatility due to the lack of a short term valuation framework, makes the argument presented earlier compelling. While general trading rule profitability is higher in the earlier parts of the observation, some trading rule specifications deliver positive returns in the early part of the observation period, others perform better during the latter parts. This shift might partially explained by the deterioration of market inefficiencies, which were more pronounced during the first part of the sample period, and partially via some of the systematic risks embedded in trading rule returns. Yet the most important driver of this change in trading rule profitability might well be the change in the market environment itself. Investors adapt as the general market environment changes.

Bearing this in mind, areas of future research would be to extend the frameworks presented in this thesis to more sophisticated trading rules. The trading rules that are presented in this study are deliberately chosen to be very generic. The drawback of this is the fact that with time they have widely lost their power to generate positive returns. Understanding how more sophisticated trading rules, which still produce strong returns, would feature in the tests applied, could be a potentially insightful exercise. An alternative way of extending this thesis would be to look at broader universe of foreign exchange markets. In recent years emerging market currencies have become increasingly important. Nonetheless, while these markets have enjoyed a true pilgrimage of international investors, opportunities still seem to be plentiful there. This assumption is confirmed by studies such as Menkhoff, Sarno, Schmeling and Schrimpf (2011), who find that most of the momentum returns come from emerging market currencies, and Chong and Ip (2009) report that more than 30% of annualised returns can be generated by utilising a momentum based trading strategy in emerging market currencies.

VIII. REFERENCES:

Anson, M. (2008) "The Beta Continuum: From Classic Beta to Bulk Beta", *Journal of Portfolio Management*, 34 (2), pp. 53-64.

Baba Y., Engle R.F., Kraft D. and Kroner, K. (1990) "Multivariate simultaneous generalized ARCH", unpublished paper, University of California, San Diego.

Baillie, R. T. and Bollerslev, T. (2000) "The forward premium anomaly is not as bad as you think", *Journal of International Money and Finance*, 19, pp. 471-88.

Berndt, E. K., Hall, B. H., Hall, R. E. and Hausman, J. A. (1974) "Estimation and Inference in Nonlinear Structural Models", *Ann. Econ. and Soc. Measurement*, 3, pp. 653-65.

Bollerslev, T. (1986) "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics*, 31, pp. 307-27.

Bollerslev, T., Engle, R. F. and Wooldridge, J. M. (1988) "A Capital Asset Pricing Model with Time varying Covariances", *Journal of Political Economy*, 96, pp.116-131.

Brown, S. J., Kumar, A. and Goetzmann, W. N. (1998) "The Dow Theory: William Peter Hamilton's Track Record Re-Considered", Working paper.

Brunnermeier, M., Nagel S., and Pedersen l. (2008) "Carry trades and currency crashes.", *NBER Macroeconomics Annual*, 23, pp. 313-347.

Campbell, B. and Dufour, J. M. (1995) "Exact non-parametric orthogonality and random walk tests", *Review of Economics and Statistics*, 77, pp.1-16.

Campbell, B. and Dufour, J. M. (1997) "Exact non-parametric tests of orthogonality and random walk in the presence of a drift parameter", *International Economic Review*, 38, pp. 151-73.

Cavaglia, S. M. F. G., Verschoor W. F. C. and Wolff, C. C. P. (1994) "On the Biasedness of Forward Foreign Exchange Rates: Irrationality or Risk Premia? ", *The Journal of Business*, 67 (3), pp. 321-343.

Brock, W., Lakonishok, J. and Lebanon, B. (1992) "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns", *Journal of Finance*, 47, pp.1731-1764.

Chinn, M. and Meredith, G. (2004) "Monetary Policy and Long Horizon Uncovered Interest Parity", *IMF Staff Papers*, 51 (3), pp. 409-430.

Chong, T. T. and Ip, H. T. (2009) "Do momentum-based strategies work in emerging currency markets?", *Pacific-Basin Finance Journal*, 17, pp. 479-493

Cowles, A. (1934) "Can stock market forecasters forecast?", *Econometrica*, 1, pp. 309-324.

DeBondt, W. M. F., and Thaler, R. (1985) "Does the stock market overreact?", *Journal of Finance*, 40, pp. 793-805.

DeBondt, W. M. F., and Thaler, R. (1987) "Further evidence on investor overreaction and stock market seasonality?", *Journal of Finance*, 42, pp. 557-581.

De Grauwe, P. and Grimaldi, M. (2006) "Exchange rate puzzles: A tale of switching attractors", *European Economic Review*, 50, pp. 1–33.

Dooley, M. P. and Shafer, J. R. (1976) "Analysis of short-run exchange rate behavior: March 1973 to September 1975", *International Finance Discussion paper no. 123*, Federal Reserve Board.

Dooley, M.P. and Shafer, J. R. (1984) "Analysis of short-run exchange rate behavior: March 1973 to November 1981", in Bigman, D. and Taya, T. (ed.) *Floating Exchange Rates and the State of World Trade Payments*: Ballinger Publishing Company, Cambridge, MA.

Efron, B. (1979) "Bootstrap methods: Another look at the jackknife", *Annals of Statistics*, 7, pp. 1-26.

Engle, R. F. (1982) "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation", *Econometrica*, 50, pp. 987-1007.

Engle R., Kroner F.K. (1995) "Multivariate simultaneous generalized ARCH", *Econometric*

Theory, 11, pp.122-150.

Fama, E. F. (1965) "The Behavior of Stock-Market Prices", *Journal of Business*, 28, pp. 34-105.

Fama, E. F. (1970) "Efficient capital markets: A Review of Theory and Empirical Work", *Journal of Finance*, 25 (2), pp. 383-417.

Fama, E. F. (1984) "Forward and spot exchange rates", *Journal of Monetary Economics*, 14, pp. 319-38.

Fama, E. F. (1989) "Efficient capital markets II", *Journal of Finance*, 26 (5), pp. 1575-1617.

Friedman, M. (1953) "The Case for Flexible Exchange Rates", *Essays in Positive Economics*.

Friesen, G. C., Weller, P. A. and Dunham, L. M. (2009) "Price trends and patterns in technical analysis: A theoretical and empirical examination", *Journal of Banking and Finance*, 33, pp. 1089–1100.

Froot, K. A. and Frankel, J. A. (1989) "Forward Discount Bias: Is it an Exchange Risk Premium?", *Quarterly Journal of Economics*, 416, pp. 139-161.

Froot, K. A. and Thaler, R. H. (1990) "Anomalies: Foreign Exchange", *The Journal of Economic Perspectives*, 4 (3), pp. 179-192.

Froot, K. A., (1990) "Short Rates and Expected Asset Returns", NBER Working paper, 3247.

Fung, W., and Hsieh, D. (2004) “Hedge Fund Benchmarks: A Risk Based Approach.” *Financial Analysts Journal*, 60, 5, pp. 65-80.

Galati, G. and Heath, A. (2007) “What drives the growth in FX activity? Interpreting the 2007 triennial survey”, *BIS Quarterly Review*, December, pp. 63–72.

Galati, G. and Melvin, M. (2004) “Explaining the surge in FX turnover”, *BIS Quarterly Review*, December, pp. 67–74.

Hamilton, W. (1922) “The Stock Market Barometer: A Study of Its Forecast Value Based on Charles H. Dow's Theory of the Price Movement”, *Barron's*, New York, N.Y.

Hansen, L. P. (1982) “Large Sample Properties of Generalized Method of Moments Estimators.”, *Econometrica*, 50(4), pp. 1029–1054.

Hesterberg, T., Moore, D.S., Monaghan, S., Clipson, A., & Epstein, R. (2003) “Bootstrap methods and permutation tests” Chapter 14 in Moore, D.S. & McCabe, G.P. *Introduction to the practice of statistics*, 5th edition.

Gultekin, M. N., and Gultekin, B. N. (1983) "Stock Market Seasonality: International Evidence", *Journal of Financial Economics*, 12, pp. 469-481.

Hodrick, R. J. and Srivastava, S. (1986) “The Covariation of Risk Premiums and Expected Future Spot Rates”, *Journal of International Money and Finance*, 5, pp5-22.

Horowitz, J. L. (2001) “The Bootstrap.”, In: Heckman, J. J., Leamer, E. E., “*Handbook of econometrics*, Vol. 5., North Holland, Amsterdam, The Netherlands, pp. 3159-3228.

Jegadeesh, N. and Titman, S. (1993) "Returns to buying winners and selling losers: Implications for stock market efficiency", *Journal of Finance*, 48, pp. 65-91.

Jegadeesh, N. and Titman, S. (2001) "Profitability of Momentum Strategies: An Evaluation of Alternative Explanations", *Journal of Finance*, 56, pp. 699–720.

Jegadeesh, N., Titman, S. (2002) "Cross-sectional and time-series determinants of momentum returns.", *Review of Financial Studies* 15 (1), pp.143– 158.

Jensen, M. C. (1972) "Capital Markets: Theory and Evidence." *Bell Journal of Economics and Management Science*, 3, pp. 357-98.

Jochum, C. (2000) "Does Market Momentum survive longer than it should?", *Financial Markets and Portfolio Management*, 14 (1), pp. 12-23.

Kaplan, E. L. and Meier, P. (1958) "Non-parametric Estimation from Incomplete Observations", *Journal of the American Statistical Association*, 53, pp. 457-481.

Karolyi, G. A. and Kho, B. C. (2004), "Momentum strategies: some bootstrap tests", *Journal of Empirical Finance*, 11, pp. 509-536.

Kho, B. C. (1996) "Time-varying risk premia, volatility, and technical trading rule profits: Evidence from foreign currency futures markets", *Journal of Financial Economics*, 41, pp. 249-290.

Kiefer, N. M. (1988) "Economic Duration Data and Hazard Functions", *Journal of Economic Literature*, 26 (2), pp. 646-679.

King, M. R. and Rime, D. (2010) “The \$4 trillion question: what explains FX growth since the 2007 survey?”, *BIS Quarterly Review*, December, pp. 27–42.

Kleinbaum, D. G. and Klein, M. (2012) “Survival Analysis, a self learning text”, Third Edition, Springer Science, Business, Media, New York, N.Y.

Kos, H. and Todorovic, N. (2008) “Global Sector survivals: Momentum effects in sector indices underlying iShares”, *The Quarterly Review of Economics and Finance*, 48, pp. 520-540.

Lawless, J. F. (2003) “Statistical Models and Methods for Lifetime Data”, Second Edition, Wiley, Hoboken, New Jersey, N.Y.

LeBaron, B. (1999) “Technical trading rule profitability and foreign exchange intervention”, *Journal of International Economics*, 49, pp. 125–143.

Lequeux, P. and Acar, E. (1998) “A Dynamic Index for Managed Currencies Funds Using CME Currency Contracts.”, *European Journal of Finance*, 4, pp. 311-330.

Levich, R. M. and Thomas, L. R. (1991) "The significance of technical trading-rule profits in the foreign exchange market: A bootstrap approach", *NBER Working paper*, 3818.

Lintner, J. (1965) "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets", *The Review of Economics and Statistics*, 47, pp. 13-37.

Lo, A. W., Mamaysky, H. and Wang, J. (2000) “Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation”, *The Journal of Finance*, 55, pp. 1705–1770.

Lo, A. W. (2004) “The adaptive markets hypothesis: Market efficiency from an evolutionary perspective”, *Journal of Portfolio Management*, 30th Anniversary Issue, pp. 15–29.

Lo, A. W., D.Repin, (2002) “The Psychophysiology of Real-Time Financial Risk Processing.”, *Journal of Cognitive Neuroscience*, 14, 3, pp. 323–39.

Lustig, H., Roussanov, N. and Verdelhan, A. (2010), Countercyclical currency risk premia, NBER Working Paper 16427.

Lustig, H., Roussanov, N. and Verdelhan, A. (2011), “Common Risk Factors in Currency Returns,” *Review of Financial Studies*, 24 (11), pp. 3731–3777.

Logue, D. E. and Sweeney, R. J. (1977) “White noise in imperfect markets: The case of the Franc/Dollar exchange rate”, *Journal of Finance*, 32, pp. 761–768.

Maynard, A. (2006) “The Forward Premium Anomaly: Statistical Artefact or Economic Puzzle? New Evidence from Robust Tests”, *The Canadian Journal of Economic*, 39(4), pp. 1244-1281.

Menkhoff L., Sarno, L., Schmeling, M. and Schrimpf, A. (2012) “Currency Momentum Strategies”, *Journal of Financial Economics*, 106, pp. 660–684;

Merton, R. C. (1973) “An Intertemporal Capital Asset Pricing Model” *Econometrica*, Econometric Society, 41(5), pp. 867-87

Mueller, P., Stathopoulos, A. and Vedolin, A. (2012) “International correlation risk.”, Working Paper, London School of Economics.

Neely, C. J. , Weller, P. and Dittmar, R. (1997) “Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach”, *The Journal of Financial and Quantitative Analysis*, 32(4), pp. 405 -426.

Neely, C. J. (2002) “The temporal pattern of trading rule returns and exchange rate intervention: Intervention does not generate technical trading rule profits”, *Journal of International Economics*, 58, pp. 211-232.

Neely, C. J., Weller, P. and Ulrich, R. (2009) “The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market”, *Journal of Financial and Quantitative Analysis*, 44, pp. 467-488.

Neely, C. and Weller, P. (2011) “Technical Analysis in the Foreign Exchange Market”, Working Paper, 2011-001A, Federal Reserve Bank of St. Louis

Nurkse, R. (1944) “International Currency Experience”, Geneva: League of Nations.

Okunev, J. and White, D. (2003) “Do Momentum-Based Strategies Still Work in Foreign Currency Markets?”, *The Journal of Financial and Quantitative Analysis*, 38 (2), pp. 425 -447.

Osler, C. L. (2003) "Currency orders and exchange rate dynamics: An explanation for the predictive success of technical analysis", *The Journal of Finance*, 58, pp. 1791–1820.

Pascual, L., Romo, J. and Ruiz, E. (2005) "Forecasting Returns and Volatilities in GARCH Processes Using the Bootstrap", Working Paper, Universidad Carlos III de Madrid, Spain.

Politis, D. and Romano, J. (1994) "The stationary bootstrap", *Journal of the American Statistical Association*, 89, pp. 1303-1313.

Pojarliev, M. and Levich, R. M. (2008) "Do Professional Currency Managers Beat the Benchmark?", *Financial Analysts Journal*, Sept./Oct. , pp. 18-32.

Pojarliev, M. and Levich, R. M. (2010) "Detecting Crowded Trades in Currency Funds.", NBER Working Papers 15698, National Bureau of Economic Research, Inc.

Pojarliev, M. and Levich, R. M. (2012) "Hunting for Alpha Hunters in the Currency Jungle.", *Journal of Portfolio Management*, Apr. , pp. 3-7.

Pojarliev, M. and Levich, R. M. (2013) " Evaluating Absolute Return Managers.", *Financial Markets and Portfolio Management*, Feb. pp. 95-103.

Potì , V., Levich, R. M. and Pattison, P. (2014) "Predictability, Trading Rule Profitability and Learning in Currency Markets", *International Review of Financial Analysis*, In Press.

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

Rogoff, K. (1979) “Essays on Expectations and Exchange Rate Volatility”, Ph.D. Dissertation Massachusetts Institute of Technology, M.A.

Rhea, R. (1932) “The Dow Theory”, Barron's, New York, N.Y.

Roll, R. (1977) "A Critique of the Asset Pricing Theory's Tests: Part I: On Past and Potential Testability of the Theory", *Journal of Financial Economics*, 4, pp. 129-76.

Rouwenhorst, K. G. (1998) “International Momentum Strategies”, *Journal of Finance* 53, pp. 267-284.

Schulmeister, S. (2006) “The interaction between technical currency trading and exchange rate fluctuations.”, *Finance Research Letters*, Elsevier, 3(3), pp. 212-233.

Sharpe, W. F. (1964) “Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk”, *Journal of Finance*, 19, pp. 425-42.

Sharpe, W. (1992) “Asset Allocation: Management Style and Performance Measurement.”, *Journal of Portfolio Management*, 18, 2, pp.7-19.

Shleifer, A., and Summers, L. H. (1990) “The Noise Trader Approach to Finance.” *Journal of Economic Perspectives*, 4, pp. 19-33.

Sullivan, R., Timmermann, A. and White, H. (1999) “Data-snooping, technical trading rule performance, and the bootstrap”, *Journal of Finance*, 54, 1647–1691.

Sweeney, R. J. (1986) “Beating the foreign exchange market”, *Journal of Finance* 41, pp. 163–182.

Sweeney, R. J. (1991) “Technical strategies in the foreign exchange markets: an interim report”, *Recent Developments in International Banking and Finance*, 3, 243-276.

Szakmary, A. C. and Mathur, I. (1997) “Central Bank Intervention and Trading Rule Profits in Foreign Exchange Markets.”, *Journal of International Money and Finance*, 16, pp. 513-535.

Taylor, S. J. (1986) “Forecasting the Volatility of Currency Exchange Rates”, *International Journal of Forecasting*, 3, pp. 159-170.

Wang, C. (2004) “Futures trading activity and predictable foreign exchange market movements”, *Journal of Banking and Finance*, Elsevier, 28(5), pp.1023-1041.

White, H. (2000) “A reality check for data snooping”, *Econometrica*, 68, pp. 1097–1126.

IX. APPENDICES: Chapter 1

A. APPENDIX 1: Overview of FX Trading Rule Literature

Study	Title	Published	Key Conclusions	Time Series vs. Cross Section
Dooley and Shafer (1976)	Analysis of short-run exchange rate behavior: March 1973 to September 1975	Working paper, Federal reserve Board	Test of exchange rate dynamics, by applying runs tests and a series of simple technical trading rules. The results of their study suggest a rejection of the martingale model for exchange rates, given the high profitability of the trading rules tested. These results cannot be explained by interest rate differentials, the possibility of trading rule returns as a consequence of central bank intervention is mentioned but not tested, due to lack of data, no risk adjustment is made.	TS, trading rule profitability
Louge and Sweeney (1977)	'White-Noise' in Imperfect Markets: The Case of the Franc/Dollar Exchange Rate	Journal of Finance	Applies spectral analysis to the FRF/USD exchange rate and then tests a series of channel breakout rules. After adjusting for transaction cost 13 of the 14 trading rules tested outperform the buy and hold strategy.	TS, trading rule profitability
Taylor (1982)	Rewards available to currency futures speculators: Compensation for risk or evidence of inefficient pricing?	Economic Record	The study investigates central bank behaviour. "leaning against the wind", policies were not profitable during the 1970's The study indicates that central banks can only support their currency for a limited amount of time but are eventually forced to allow the adjustment to take place, and when this happens they lose significant amounts of money. Moreover, in the absence of speculators who were betting against central banks, the losses of currency intervention would have been only half the amount they actually have been. Hence intervention must have been one of the main deterrents for trading rule profitability.	TS, central bank behaviour
Dooley and Shafer (1984)	Analysis of short-run exchange rate behavior: March 1973 to November	Book: David Bigman and Teizo Taya, eds.: Floating Exchange Rates and the State of World Trade Payments	The study looks at trading filters. The results give evidence of substantial profits to all but the largest filters over the period 1973-81 for the deutschemark, yen and pound sterling.	TS, trading rule profitability
Sweeney (1986)	Beating the foreign exchange market	Journal of Finance	The study uses daily dollar/deutschemark data over the period 1975-80, he reported excess profits over buyandhold strategies of 4% per annum for a 0.5% filter. A risk adjustment is undertaken by applying of a constant risk premium (or discount).	TS, trading rule profitability, risk adjustment
Allen and Taylor (1990)	Charts, noise and fundamentals in the London foreign exchange market	The Economic Journal	The results of the paper suggest that chartist advice does not destabilise markets in the sense that chartists' expectations do not appear to overreact systematically to changes in the current exchange rate. The most that can be said, given the present evidence, is that chart advice may at most cause meanreverting or stationary deviations from the fundamentals. the study indicates that almost all foreign exchange professionals use technical analysis as a tool in decision making at least to some degree. Moreover, most foreign exchange professionals use some combination of technical analysis and fundamental analysis and the relative weight given to technical analysis as opposed to fundamental analysis rises as the trading or forecast horizon declines.	TS, trading rule profitability
Lebaron (1991)	Technical trading rule profitability and foreign exchange intervention	Journal of International Economics	The study suggests that simple rules used by traders have some predictive value over the future movement of foreign exchange prices. However, when the profitability of these trading rules is analysed in connection with central bank activity using intervention data from the Federal Reserve, excess returns of technical trading rules are considerably diminished. The results indicate that after removing periods in which the Federal Reserve is active, exchange rate predictability is dramatically reduced.	TS, central bank behaviour
Taylor (1992)	Rewards Available to Currency Futures Speculators: Compensation for Risk or Evidence of Inefficient Pricing?.	Economic Record	The study tests the profitability of technical trading rules allowing for a time-varying risk premium in the form of a first-order autoregressive process. For several parameterisations of the time varying risk premium, hundreds of time series are then simulated on which trading rule results are evaluated. It is found that there appears to be no reasonable constellation of parameters for the time-varying risk premium which would be needed to explain observed returns as a compensation for risk.	TS, trading rule profitability, risk adjustment
Levich and Thomas (1993)	The significance of technical trading-rule profits in the foreign exchange market: A bootstrap approach	Journal of International Money and Finance	The study uses futures prices for a number of foreign currencies to examine the profits earned by various moving average and filter rules over the period 1976-1990. To ascertain the validity of the results the study conducts bootstrapping simulations to assess the significance of their results. the results indicate that the profit levels generated by the trading rules cannot be reconciliated by the simulation	TS, trading rule profitability, risk adjustment
Taylor (1994)	Trading futures using a channel rule: A study of the predictive power of technical analysis with currency examples	Journal of Futures Markets	The study uses a Channel breakout trading rules, and evaluates the results on the basis of a series of econometric tests. The results suggest that channel breakout trading rules based on currency pairs appear to be profitable when compared to random walk assumptions.	TS, trading rule profitability
Osler and Chang (1995)	Head and shoulders: Not just a flaky pattern	Federal Reserve Bank of New York Staff Report	the study aims to evaluate the profitability of the head-and-shoulders pattern, applying bootstrap methodology the study finds evidence of significant profits for some of the exchange rates tested. Trading in all six currencies simultaneously would have yielded significant profits even after transactions costs, although these profits were lower than typically reported for the moving average and filter rules.	TS, trading rule profitability
Kho (1996)	Time-varying risk premia, volatility, and technical trading rule profits: Evidence from foreign currency futures markets	Journal of Financial Economics	The paper looks at systematic risk factors, by establishing a conditional CAPM framework, implemented through a bivariate GARCH-M model, between trading rule returns and the MSCI world equity index, which acts as a proxy for global market risk. The results of the study suggest that a substantial amount of the profitability of moving average rules in foreign currency markets can be explained by a time-varying risk premium.	TS, trading rule profitability, risk adjustment
Lee and Mathur (1996)	Trading rule profits in european currency spot cross-rates	Journal of Banking & Finance	The study applies moving average trading rules six European spot cross rates. The results suggest that MA trading rules are marginally profitable only for the JY/DM and the JY/SF cross rates, while trading rules are not profitable for the other four cross-rates. Bootstrapping and out of sample tests provide similar results. Examination of subsamples characterized by central bank intervention do not produce different results. Computation of Box Pierce statistics adjusted for heteroscedasticity show that daily returns for all six cross rates are serially uncorrelated.	TS, trading rule profitability, risk adjustment

Study	Title	Published	Key Conclusions	Time Series vs. Cross Section
Neely, Weller and Dittmar (1997)	Is technical analysis in the foreign exchange market profitable? A genetic programming approach	Journal of Financial and Quantitative Analysis	the study applies a genetic program that searches for an optimal trading rule. the evaluation of results is based on six currency pairs, which a sample spanning from 1974 to 1995 that is split into three sub-periods, which constitute selection, training and testing period for the genetic code. the key findings are that different currency pairs produce higher trading returns than others and that different currencies pairs also favour different sets of trading rules. Overall genetically grown trading rules show out of sample profitability, even when compared against bootstrapped benchmark simulations.	TS, trading rule profitability, risk adjustment
Szakmary and Mathur (1997)	Central bank intervention and trading rule profits in foreign exchange markets	Journal of International Money and Finance	The results of the study present strong evidence that market operations by central banks are indeed key drivers of trading rule profitability. They show that trading against central bank intervention can yield significant excess returns. These findings are based on a sample of five currencies versus the US Dollar from 1977 to 1991. The median return of the moving average trading rule ranges between 5.4% and 9.8% depending on the currency pairs chosen. Based on a regression analysis they suggest that leaning against the wind intervention helps explaining that median moving average trading profits for various currencies are greater than zero.	TS, central bank behaviour
Lequeux and Acar	A dynamic index for managed currencies funds using CME currency contracts	The European Journal of Finance	The goal of the paper is build a dynamic benchmark based on technical trading rules. That exhibits a high correlation to the average currency manager	TS, Performance evaluation
Chang and Osler (1999)	Methodical madness: Technical analysis and the irrationality of exchange-rate forecasts	Economic Journal	This paper identifies a widely used technical trading signal the head-and-shoulders pattern, as a potential source of departures from market efficiency. Forecasts based on this pattern are evaluated for daily dollar exchange rates over 1973 to 1994, using profitability and efficiency as evaluation parameters. When tested for statistical significance using a bootstrap technique, the results indicate that he strategy is profitable but not efficient, given the fact that simpler trading rules exhibit higher profitability.	TS, trading rule profitability
Lebaron (1999)	Technical trading rule profitability and foreign exchange intervention	Journal of International Economics	The study confirms the findings of previous research. The results suggest that the trading rule profits are highest during periods of central bank intervention. When removing the time periods where central banks are active in the currency market, the results are insignificant.	TS, central bank behaviour
Neely and Weller (1999)	Technical Trading Rules in the European Monetary System	Journal of International Money and Finance	The study analyses the performance of intraday technical trading rules which are constructed by using genetic programming. The tested trading rules generate significant excess returns for three of four EMS exchange rates over the out-of-sample period 1986–1996. Moreover the results cannot be duplicated by commonly used moving average rules and there is no evidence that the excess returns are compensation for bearing systematic risk. When realistic transaction costs and trading hours are taken into account there is no evidence of excess returns to the trading rules derived.	TS, trading rule profitability, risk adjustment
Lee, Gleason and Mathur (2001)	Trading rule profits in Latin American currency spot rates	International Review of Financial Analysis	The study applies applying the moving average and channel trading rules to Latin American currencies to see if opportunities for profitable trading exist. While not all of the Latin American currencies can be exploited through the use of trading rules, some appear amenable to technical analysis. Moving average rules are profitable the Brazilian real, the Mexican peso, the Peruvian new sol, and the Venezuelan bolivar, while channel trading rules rules are profitable for the Brazilian real, the Mexican peso, and the Venezuelan bolivar. Moreover the results indicate that some trading rules may be more suitable for certain types of currencies.	TS, trading rule profitability
Martin(2001)	Technical trading rules in the spot foreign exchange markets of developing countries	Journal of Multinational Financial Management	the results of the study indicates that technical trading rules generate profit opportunities in the spot foreign exchange markets of developing countries. Most of the technical trading rules generate statistically significant out-of-sample returns even after accounting for transaction cost. On a risk-adjusted basis performance measures indicate that trading rules do not outperform a simple short-selling strategy or risk-free strategy.	TS, trading rule profitability
Lebaron (2002)	Technical trading profitability in foreign exchange markets in the 1990's	Working paper, Brandeis University	The study gives evidence of changing profitability in trading rule returns during the 1990, previously good performance is no longer strong, evidence for regime shift	TS, trading rule profitability
Neely (2002)	The temporal pattern of trading rule returns and exchange rate intervention: Intervention does not generate technical trading rule profits	Journal of International Economics	The study analyses intraday data for five currency pairs, covering the time range from the early to mid-eighties to the mid to late nineties. The analysis is based on a 150 day moving average trading rule. First the study compares the moving average trading results for a data sample that contains intervention dates and a data sample that does not contain intervention dates, indicating similar results to previous studies. However when looking at intraday return realisations the results indicate that intervention does not generate returns itself. Currency intervention comes as a reaction to strong and very profitable short-term trends.	TS, central bank behaviour
Saacke (2002)	Technical analysis and the effectiveness of central bank intervention	Journal of International Money and Finance	the study provides evidence that central banks earn profits with interventions and that technical trading rules are unusually profitable on days on which interventions take place, these results are based on data of foreign exchange interventions of the Bundesbank and the Fed. The results indicate that intervention profits and trading rule profitability are measured over different horizons and after interventions, exchange rates tend to move contrary to central banks' intentions in the short run, but in agreement with their interventions in the long run.	TS, central bank behaviour
Gencay, Dacorogna, Olsen and Pictet (2003)	Real-time trading models and the statistical properties of foreign exchange rates	International Economic Review	the study compares the performance of a widely used commercial real-time trading model to asimple exponential moving average model, and the trading models are used as diagnostic tools to evaluate the statistical properties of foreign exchange rates. The trading models applied help to observe the data generating process in foreign exchange markets is a complex network of layers where each layer corresponds to a particular frequency. A successful characterization of such data generating processes should be estimated with models whose parameters are functions of intra and inter frequency dynamics.	TS, trading rule profitability
Neely and Weller (2003)	Intraday technical trading in the foreign exchange market	Journal of International Money and Finance	The study analyses the out-of-sample performance of intraday technical trading strategies selected using two methodologies, a genetic program and an optimized linear forecasting model. When realistic transaction costs and trading hours are taken into account, there is find no evidence of excess returns to the trading rules derived with either methodology.	TS, trading rule profitability, risk adjustment

Study	Title	Published	Key Conclusions	Time Series vs. Cross Section
Okunev and White (2003)	Do momentum-based strategies still work in foreign currency markets?	Journal of Financial and Quantitative Analysis	The study evaluates 354 moving average rules for eight currencies from January 1980 to June 2000. Every short-term moving average value ranges from one to twelve months, while the long-term moving average values range from two to 36 months. the implementation of the trading rule portfolio is done similar to the methodology proposed by Jegadeesh and Titman (1993, 2001) whereby they choose long-short portfolios based upon the strength of the moving average signal. For every short-term/long-term moving average combination they initiate a long position in the currency with the highest rank and short the currency with the lowest rank. After having corrected currency returns for the interest differential the trading strategy provides an excess returns over the benchmark of 5%-6% per year with low correlations between the trading rule returns and the benchmark currency basket.	CS, trading rule profitability
Osler (2003)	Currency orders and exchange rate dynamics: An explanation for the predictive success of technical analysis	The Journal of Finance	The study analyses a dataset of almost 9700 stop-loss or take-profit orders placed by a large investment bank for three exchange rates from September 1999 to April 2000. The paper suggests that "support" and "resistance" levels can be key indicators for accelerated momentum or reversals, depending on whether they are broken or not. This is due to the distribution of the placement of stop-loss and take-profit orders by clients, which tends to cluster around round numbers. While take-profit orders are mostly clustered around round numbers, stop loss orders have a pronounced tendency to be placed just beyond round numbers. Buy orders are often just above and sell orders are just below the round number. This would suggest that "support" and "resistance" levels, which tend to be round numbers, are key indicators for either a trend reversals if the spot price fails to cross the level, or trend acceleration when levels are crossed.	TS, trading rule profitability, market microstructure
Cheung and Chin (2004)	Currency traders and exchange rate dynamics: A survey of the US market	Journal of International Money and Finance	The survey indicates that in recent years electronically-brokered transactions have risen substantially, mostly at the expense of traditional brokers. Technical trading best characterizes about 30% of traders, with this proportion rising from five years ago. Moreover, economic fundamentals are perceived to be more important at longer horizons, while short-run deviations from the fundamentals are attributed to excess speculation. Speculation is generally viewed positively, as enhancing market efficiency and liquidity, even though it exacerbates volatility. Central bank intervention does not appear to have a substantial effect, although there is general agreement that it increases volatility.	TS, trading rule profitability
Olson (2004)	Have trading rule profits in the currency markets declined over time?	Journal of Banking and Finance	the study tests whether moving average trading rule profits have declined over the period from 1971 to 2000. The analysis is done using 18 exchange rate series. Trading rules are optimized for successive 5 year in sample periods from 1971 to 1995 and tested over subsequent 5 year out of sample periods. Results show that risk adjusted trading rule profits have declined over time from an average of over 3% in the late 1970s and early 1980s to about zero in the 1990s	TS, trading rule profitability
Wang (2004)	Futures trading activity and predictable foreign exchange market movements	Journal of Banking & Finance	The study looks at currency returns from a market microstructural perspective, incorporating the positioning of market participants such as hedgers and speculators when designing tests of foreign exchange market efficiency. The paper utilises weekly data of five currency futures contracts against the USD with a sample spanning January 1993 to March 2000 the Study finds that speculator sentiment varies positively with future returns, while hedger sentiment varies negatively with future returns. Moreover, positive or negative extreme sentiment exhibits a higher correlation to price movements than moderate sentiment. The study also indicates that the aspect of hedging pressure has to be considered in the context of foreign exchange markets.	TS, trading rule profitability, market microstructure
Osler (2005)	Stop-loss orders and price cascades in currency markets	Journal of International Money and Finance	This study gives evidence that currency stop-loss orders contribute to rapid, self-reinforcing price movements. The analysis based on high-frequency exchange rates offers three main results. Exchange rate trends are unusually rapid when rates reach exchange rate levels at which stop-loss order have been documented to cluster. The response to stop-loss orders is larger than the response to take-profit orders. The response to stop-loss orders lasts longer than the response to take-profit orders. In combination these results indicate that stop-loss orders propagate trends and are sometimes triggered in waves, contributing to price cascades.	TS, trading rule profitability, market microstructure
Dewachter and Lyrio (2006)	The cost of technical trading rules in the Forex market: A utility-based evaluation	Journal of International Money and Finance	the study analyses the opportunity cost for rational risk averse agents of using technical trading rules in the foreign exchange rate market. Opportunity cost are analysed as a cost related to the misallocation of wealth, which increases with the investor's level of risk aversion and as a cost related to the investor's erroneous belief regarding the sign of the expected excess return. The results of the study indicate that even for low levels of risk aversion the opportunity cost of using chartist rules tends to be prohibitively high.	TS, trading rule profitability
Menkhoff and Taylor (2006)	The obstinate passion of foreign exchange professionals: Technical analysis	Journal of Economic Literature	The survey gives an overview of the stylised facts of trading rule research. Technical analysis is an important and widely used method of analysis in the foreign exchange market and that applying certain technical trading rules over a sustained period may lead to significant positive excess returns. The study also analyses the four arguments that have been put forward to explain the continuing widespread use of technical analysis and its apparent profitability. Firstly that the foreign exchange market may be characterised by not-fully-rational behaviour. Secondly, that technical analysis may exploit the influence of central bank interventions. Moreover, that technical analysis may be an efficient form of information processing; and finally that it may provide information on non-fundamental influences on foreign exchange movements. Although all of these positions may be relevant to some degree, neither non-rationality nor official interventions seem to be widespread and persistent enough to explain the profitability of technical trading rules.	TS, trading rule profitability
Qi and Wu (2006)	Technical trading-rule profitability, data snooping, and reality check: Evidence from the foreign exchange market	Journal of Money, Credit, and Banking	The study applies the data snooper introduced by White (2000) to a universe of daily rates of seven currencies against the USD over a time period from April 1973 to December 1998. The results suggest that the best performing trading rules, according to White's data snooper, are short-term channel breakout rules for the Japanese Yen and the Swiss Franc and short-term moving averages for the other currency pairs. Without accounting for transaction costs the mean excess returns over a buy and hold strategy are unanimously positive in the range of 4.02% to 12.81% per annum. After accounting for one-way transaction costs of 4bps the excess returns are still positive in the range of 2.14% to 11.46%. The returns generated on an out of sample basis are considerably less than the in-sample returns. Nonetheless, with the exception of the Italian Lira all of them are statistically significant on the 10% level.	TS, trading rule profitability, risk adjustment

Study	Title	Published	Key Conclusions	Time Series vs. Cross Section
Schulmeister (2006)	The interaction between technical currency trading and exchange rate fluctuations	Finance Research Letters	The study links the behaviour of technical models and exchange rate dynamics. The basic idea behind this is the fact that traditional price discovery under the efficient market hypothesis is violated, due to the additional aspect of market participants being aware of technical trading rules. The study analyses 1024 moving average and momentum models in the DEM/USD market between 1973 and 1999. It also conducts an out-of-sample test for the EUR/USD rate over the time period from 2000 to 2004. The results of his analysis suggest that when markets change direction, the majority of trading filters in his study tend to be on the same side, suggesting that there is a pronounced feedback mechanism between trading rules and movements in the underlying exchange rates. The study also indicates that the majority of trading rules are profitable, in and out of sample, and that the profitability is exclusively due to persistence in exchange rate movements. Therefore one could argue that market participants expect the persistence in price movements to be sufficiently frequent to compensate for the potential loss that occurs due to a sudden reversal of a trend.	TS, trading rule profitability, market microstructure
Dueker and Neely (2006)	Can Markov switching models predict excess foreign exchange returns?	Journal of Banking & Finance	Study merges the literature on technical trading rules with the literature on Markov switching to develop economically useful trading rules. While the Markov models outperform standard technical rules modestly on an out of sample basis. A portfolio of Markov and standard technical rules outperforms either of the individual sets of trading rules on a risk-adjusted basis.	TS, trading rule profitability
Pukthuanthong-Le, Levich and Thomas(2007)	Do foreign exchange markets still trend?	The Journal of Portfolio Management	The study examines the major currency futures contracts which have been trading since the 1970s as well as more recent contracts on emerging market currencies. The main conclusion is that the era of easy profits from simple trend following strategies in major foreign currencies is over. The markets have adapted to the extent that profits from these simple trading strategies have vanished. Amongst the emerging market currencies there are more attractive profit opportunities.	CS, trading rule profitability
Brunnermeier, Nagel, Pedersen (2008)	Carry Trades and Currency Crashes	NBER Working Paper	This Study documents that carry traders are subject to crash risk: i.e. exchange rate movements between high-interest-rate and low-interest-rate currencies are negatively skewed. The authors argue that this negative skewness is due to sudden unwinding of carry trades, which tend to occur in periods in which risk appetite and funding liquidity decrease. Funding liquidity measures predict exchange rate movements, and controlling for liquidity helps explain the uncovered interest-rate puzzle. Carry-trade losses reduce future crash risk, but increase the price of crash risk. They also document excess co-movement among currencies with similar interest rate. Our findings are consistent with a model in which carry traders are subject to funding liquidity constraints.	TS, trading rules as systematic risk factors
Poljarliev and Levich (2008)	"Do Professional Currency Managers Beat the Benchmark?"	Financial Analyst Journal	The study establishes a universe of four of currency benchmark strategies against which they compare various currency fund managers, consisting of the factors carry, trend, value and volatility. The results indicate that over the entire sample period, spanning from 1996 to 2000, 66% of the variability in monthly returns of their manager universe can be explained by these four factors. In the time period after 2000 the explanatory power of the four factors rises to almost 77%, with carry being the most dominant factor. While the explanatory power of trend has declined somewhat in recent years, the overall percentage of currency fund returns explained by the factor trend is 65% throughout the 1990s and after 2000. They also indicate that the disappointing returns from currency managers are mainly the consequence of declining profitability of momentum as a risk factor.	TS, trading rules as systematic risk factors
Schulmeister (2008)	Components of the profitability of technical currency trading	Applied Financial Economics	The paper investigates the profitability of 1,024 moving average and momentum models and their components in the German mark (euro)/U.S. dollar market. The main results are as follows. First, each of these models would have been profitable over the entire sample period. Second, this profitability is exclusively due to the exploitation of exchange rate trends. Third, these results do not change substantially when trading is examined within subperiods. Fourth, the 25 best performing models in each in-sample period examined were profitable also out of sample in most cases. Fifth, the profitability of technical currency trading has been declining since the late 1980s.	TS, trading rule profitability, market microstructure
Chong and Ip (2009)	Do momentum-based strategies work in emerging currency markets?	Pacific-Basin Finance Journal	Chong and Ip (2009) extend Okunev and White's (2003) study to emerging market currencies. The study tests six developing country currencies with a sample spanning from January 1985 to December 2004. The results of the analysis suggests 30% plus annualised returns of the moving average trading strategy. They also find returns to be very steady throughout the observation period. After accounting for transaction costs of 5% per annum, the trading rule still delivers significant positive returns.	TS, trading rule profitability, risk adjustment
Harris and Yilmaz (2009)	A momentum trading strategy based on the low frequency component of the exchange rate	Journal of Banking & Finance	This study develops a momentum trading strategy based on the low frequency trend component of the spot exchange rate. This is done by kernel regression and the high pass filter of Hodrick and Prescott. The back tests of this strategy suggest that the results offer greater directional accuracy, higher returns and Sharpe ratios, lower maximum drawdown and less frequent trading than traditional moving average rules. This performance is also relatively robust across different time periods and choice of smoothing parameters as well as the distribution and bandwidth parameter.	TS, trading rule profitability
Neely, Weller and Ulrich (2009)	The adaptive markets hypothesis: Evidence from the foreign exchange market	Journal of Financial and Quantitative Analysis	the study analyses the intertemporal stability of excess returns to technical trading rules in the foreign exchange market by conducting true, out-of-sample tests on previously studied rules. The results suggest that the excess returns of the 1970s and 1980s were genuine and not just the result of data mining. But these profit opportunities had disappeared by the early 1990s for filter and moving average rules. Returns to less-studied rules also have declined but have probably not completely disappeared.	TS, trading rule profitability, market microstructure
Schulmeister (2009)	Aggregate trading behaviour of technical models and the yen/dollar exchange rate 1976-2007	Japan and the World Economy	The study investigates the profitability of 1,024 moving average and momentum models and their components in the yen/dollar market. It turns out that all models would have been profitable between 1976 and 1999. While the models produce more single losses than single profits, the size of the single profits is on average much higher than the size of single losses. Hence, the profitability of technical currency trading is exclusively due to the exploitation of persistent exchange rate trends. The results of the analysis hold over a series of over sub samples. However, the profitability of technical currency trading based on daily data has declined since the late 1980s and has disappeared over the out-of-sample period between 2000 and 2004.	TS, trading rule profitability, market microstructure

Study	Title	Published	Key Conclusions	Time Series vs. Cross Section
Wan and Kao (2009)	Evidence on the contrarian trading in foreign exchange markets	Economic Modelling	the study analyses the existence and price impacts of contrarian behavior in the foreign exchange markets. this is done utilizing a nonlinear behavioral model where the chartists and fundamentalists coexist, evidence obtained from two sample periods significantly supports the existence of contrarian trading in the British pound, the Japanese yen and the German mark markets. The contrarian trading can only partially offset the price impacts of trend-followers, therefore the price impact of the chartists as a whole is destabilizing. The ability that the contrarians can counterbalance the extrapolation of the trend-followers differs across markets. Traders in the BP market have the highest tendency to contrarian strategy, which in turn contributes to the least deviations of the BP exchange rates departing from its PPP fundamentals.	TS, trading rule profitability
Zwart, Markwat, Swinkels and Dijk (2009)	The economic value of fundamental and technical information in emerging currency markets	Journal of International Money and Finance	The study measures the economic value of information derived from macroeconomic variables and from technical trading rules for emerging markets currency investments. Basing its findings on the analysis of a sample of 21 emerging markets over the period 1997–2007, explicitly accounting for trading restrictions on foreign capital movements by using non-deliverable forward data. The study documents that both the use of fundamental and technical analysis improves the risk-adjusted performance of investment strategies when used in combination	TS, trading rule profitability
Lustig, Roussanov, Verdelhan (2010)	Countercyclical Currency Risk Premia	NBER Working Paper	The study builds novel currency investment strategy, the ‘dollar carry trade,’ which delivers large excess returns, uncorrelated with the returns on well-known carry trade strategies. Using a no-arbitrage model of exchange rates they show that these excess returns compensate U.S. investors for taking on aggregate risk by shorting the dollar in bad times, when the U.S. price of risk is high. The counter-cyclical variation in risk premia leads to strong return predictability: the average forward discount and U.S. industrial production growth rates forecast up to 25% of the dollar return variation at the one-year horizon. The estimated model implies that the variation in the exposure of U.S. investors to world-wide risk is the key driver of predictability.	TS, trading rules as systematic risk factors
Poljarliev and Levich (2010)	Detecting Crowded Trades in Currency Funds	NBER Working Papers	The paper focuses on crowdedness of styles of currency fund managers. The strategies used in the paper are carry, momentum, volatility and value. the study defines crowdedness as the percentage of the funds with significant positive exposure to a given “benchmark” trading strategy less the percentage of the funds with significant negative exposure to that same strategy. Key findings of their study are a high degree of crowdedness of carry in 2007 and 2008. Trend on the other hand side is very popular in the early parts of their sample, which spans from 2005 to 2010. By May 2008 trend crowdedness declines to almost zero, a few months before the performance of the trend factor starts picking up again throughout autumn of 2008. Subsequently the number of fund managers following the trend strategy picks up again. The measure for crowdedness reaches 21.6% in November 2009. It then declines again to almost zero by the middle of 2010.	TS, trading rules as systematic risk factors
Serban (2010)	Combining mean reversion and momentum trading strategies in foreign exchange markets	Journal of Banking and Finance	The study analyses momentum and mean reversion behaviour in foreign exchange markets, by implementing trading strategy that combines mean reversion and momentum in foreign exchange markets. The tested strategy, which was originally designed for equity markets, generates abnormal returns when applied to uncovered interest parity deviations for five countries. Quantitatively, the strategy performs better in foreign exchange markets than in equity markets, it also outperforms traditional foreign exchange trading strategies, such as carry trades and moving average rules.	TS, trading rule profitability
Burnside, Eichenbaum, and Rebelo (2011)	Carry Trade and Momentum in Currency Markets	NBER Working Paper	The study analyses two explanations for the profitability of the carry momentum strategies. The first is that investors are compensated for bearing risk, for which the study finds little evidence. The second is that the profitability results from a rare disaster problem. The study also indicates that a rare disaster event is not characterized by large losses to currency speculators. Instead, it features moderate losses and high values of the stochastic discount factor.	TS, trading rule profitability
Cialenco Protopapadakis (2011)	Do technical trading profits remain in the foreign exchange market? Evidence from 14 currencies	Journal of International Financial Markets	The paper examines the in and out of sample behavior of moving average filters for 14 developed country currencies using daily data with bid-ask spreads. The study records significant in sample returns in the early periods, while out of sample returns are lower and only occasionally significant. the results also suggest that a currency risk factor proposed in the literature is systematically related to these returns. Moreover the findings present no evidence that there is a link between falling transactions costs and trading profits.	TS, trading rule profitability
Poljarliev and Levich (2012)	Hunting for Alpha Hunters in the Currency Jungle	Journal of Portfolio Management	Editorial, that analyses how to look at active currency managers	TS, Performance evaluation
Neely and Weller (2012)	Technical Analysis in the Foreign Exchange Market”, Working Paper	Working Paper, Federal Reserve Bank of St. Louis	The adaptive markets hypothesis posits that trading strategies evolve as traders adapt their behavior to changing circumstances. This paper studies the evolution of trading strategies for a hypothetical trader who chooses portfolios from foreign exchange (forex) technical rules in major and emerging markets, the carry trade, and U.S. equities. The results show that a backtesting procedure to choose optimal portfolios improves upon the performance of nonadaptive rules. We also find that forex trading alone dramatically outperforms the S&P 500, with much larger Sharpe ratios over the whole sample, but there is little gain to coordinating forex and equity strategies, which explains why practitioners consider these tools separately. Forex trading returns dip significantly in the 1990s but recover by the end of the decade and have been markedly superior to an equity position since 1998. Overall, trading rule returns still exist in forex markets—with substantial stability in the types of rules—though they have migrated to emerging markets to a considerable degree.	TS, trading rule profitability
Tajaddini and Crack (2012)	Do momentum-based trading strategies work in emerging currency markets?	Journal of International Financial Markets	The study reports the profitability of emerging currency momentum strategies using a long time series and a good cross-sectional sample. Using a 1985–2009 sample period and six emerging currencies. the results indicate that that long-short momentum strategies gained about 1–3% per annum after actual transaction costs. These profits declined through time. Most strategies lose money after transaction costs during the last five years of our sample.	CS, trading rule profitability

Study	Title	Published	Key Conclusions	Time Series vs. Cross Section
Lustig, Roussanov, Verdelhan (2011)	Common Risk factors in Currency Markets	Review of Financial Studies	The study analyses the theycyclicalty and predictability of currency excess returns . The average excess returns on low interest rate currencies are 4.8 percent per annum smaller than those on high interest rate currencies after accounting for transaction costs. A single return-based factor, the return on the highest minus the return on the lowest interest rate currency portfolios, explains the cross-sectional variation in average currency excess returns from low to high interest rate currencies. In a simple affine pricing model, they show that the high-minus-low currency return measures that component of the stochastic discount factor innovations that is common across countries. To match the carry trade returns in the data, low interest rate currencies need to load more on this common innovation when the market price of global risk is high.	TS, trading rules as systematic risk factors
Menkhoff, Sarno, Schmeling and Schrimpf (2012)	Currency Momentum Strategies	Journal of Financial Economics	the study connects traditional cross sectional momentum literature of equity markets with foreign exchange markets. The study implements the Jegadeesh and Titman (1993) approach using foreign exchange data. The sample of the study consists of cross sectional data of 48 countries over a time period from January 1976 to January 2010. the results of the study suggests that some of the winner versus loser combinations earn unconditional average excess returns of up to 10% per year.	CS, trading rule profitability
Evans, Pappas and Xhafa (2013)	Utilizing artificial neural networks and genetic algorithms to build an algo-trading model for intra-day foreign exchange speculation	Mathematical and Computer Modelling	this study builds a prediction and decision making model based on Artificial Neural Networks and Genetic Algorithms. The dataset utilized ch comprises of 70 weeks of past currency rates of the GBP/USD, EUR/GBP and EUR/USD. The results of the study suggest that with a significance of more than 95% currency rates are not randomly distributed. The results of the proposed model achieve 72.5% prediction accuracy. Furthermore, implementing the optimal trading strategy, this model produces 23.3% annualised return after transaction cost.	TS, trading rule profitability
Poljarliev and Levich (2013)	Hunting for Alpha Hunters in the Currency Jungle	Financial Markets and Portfolio Management	The Study proposes a new performance metric that strips out beta returns associated with investment-style factors. This approach leads to a new statistic, the alpha ratio, which can dramatically impact the relative performance rankings of managers and provide a clearer signal of manager skill. One traditional measure of investment performance, the information ratio (IR), is defined as the active return (alpha) divided by the tracking error (the standard deviation of the active return). Calculating an IR is straightforward when the benchmark for performance is a buy-and-hold standard like the S&P 500. For absolute return managers, however, the typical benchmark is zero, meaning that any excess return is classified as alpha and deemed to represent the return from active management or skill. In this paper, we argue that this standard approach confuses beta returns and alpha returns. The former can be earned by following generic strategies that are easily implemented and often replicated by ETFs, while the later are associated with more original or complex strategies that more genuinely reflect unique skills or expertise.	TS, Performance evaluation
Neely and Weller (2013)	Lessons from the Evolution of Foreign Exchange Trading Strategies	NBER Working Paper	The study analyses the evolution of trading strategies for a hypothetical trader who chooses portfolios from foreign exchange technical rules in major and emerging markets, the carry trade, and U.S. equities. The results show that a backtesting procedure to choose optimal portfolios improves upon the performance of nonadaptive rules. The results also indicate that forex trading alone dramatically outperforms the S&P 500, with much larger Sharpe ratios over the whole sample with little gain to coordinating forex and equity strategies. Forex trading returns dip significantly in the 1990s but recover by the end of the decade and have been markedly superior to an equity position since 1998. Overall, trading rule returns still exist in forex markets though they have migrated to emerging markets to a considerable degree.	TS, trading rule profitability
Sager and Taylor (2014)	Generating currency trading rules from the term structure of forward foreign exchange premia	Journal of International Money and Finance	this study aims create a trading system that aims to outperform the random walk assumption by exploiting information embedded within the term structure of forward exchange rate premia and, whether such framework can be used to generate significant trading profits in combination with an acceptable degree of risk in a realistic investment portfolio context.	TS, trading rule profitability
Kuang, Schoder and Wang (2014)	Illusory profitability of technical analysis in emerging foreign exchange markets	International Journal of Forecasting	The study undertakes a comprehensive examination of the profitability of technical trading rules in ten emerging foreign exchange markets. Studying 25,988 trading strategies for emerging foreign exchange markets, the results suggest thaty the best rules can sometimes generate an annual mean excess return of more than 30%. Based on standard tests. Moreover the authors indicate that almost all of the trading rule returns reported vanish once the data snooping bias is taken into account.	TS, trading rule profitability, risk adjustment
Levich and Poti (2014)	Predictability and 'Good Deals'	International Journal of Forecasting	The study analyses the predictability of currency returns over the period 1971-2006. To assess the economic significance of currency predictability, the study calculates predictive regressions, with upper boundaries which are motivated by "no good-deal" restrictions that rule out unduly attractive investment opportunities. The results indicate that the Excess predictability is highest in the 1970s and tends to decrease over time, but it is still present in the final part of the sample period. Moreover, periods of high and low predictability tend to alternate.	TS, trading rule profitability
Poti, Levich and Pattioni (2014)	Predictability, Trading Rule Profitability and Learning in Currency Markets	International Review of Financial Analysis	This paper studies predictability of currency returns over time and the extent to which it is captured by trading rules commonly used in currency markets. the results indicate a close relation between these strategies and indices that track popular technical trading rules, namely moving average cross-over rules and the carry trade. this suggests that trading rules represent heuristics by which professional market participants exploit currency mispricing. The predictability is highest in the mid '90, subsequently decreases sharply, but increases again in the more recent time period, especially for the Euro and other emerging currencies. The key finding of the paper is that the efficient market hypothesis does not hold, and this is particularly evident in the early part of their data sample, which spans form 1988 to 2010. During the time period of the global financial crisis (2007-2010), there is also strong evidence of a deviation from the efficient market hypothesis. This stands in sharp contrast to the popular view of a gradual deterioration of trading rule profitability, which has gathered support amongst the academic community in recent years.	TS, trading rule profitability

B. APPENDIX 2: Descriptive Statistics Based on Log Returns

1. Descriptive Statistics (LOG Base Currency Returns)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean Ret. (%)		0.004%	-0.013% *	-0.006%	-0.010%	0.001%	0.006%	0.001%	0.004%	0.007%
Std. Dev. (%)		0.623%	0.690%	0.668%	0.755%	0.686%	0.707%	0.413%	0.737%	0.808%
Skew		0.119	-0.428	-0.005	-0.001	0.239	1.643	-0.249	3.022	3.438
Kurtosis		7.895	7.442	8.628	8.549	11.707	42.822	16.633	72.132	78.724
JB p-value		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
GBP										
Mean Ret. (%)	-0.004%		-0.017% **	-0.011% **	-0.014% **	-0.003%	0.002%	-0.004%	0.000%	0.002%
Std. Dev. (%)	0.623%		0.729%	0.503%	0.600%	0.552%	0.601%	0.648%	0.783%	0.825%
Skew	-0.119		-0.552	-0.597	-0.359	0.241	2.941	-0.131	2.026	2.955
Kurtosis	7.895		9.414	12.931	12.012	16.096	82.214	6.550	50.741	66.427
JB p-value	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
JPY										
Mean Ret. (%)	0.013% *	0.017% **		0.007%	0.003%	0.014% *	0.019% **	0.014% *	0.017% *	0.020% **
Std. Dev. (%)	0.690%	0.729%		0.669%	0.698%	0.736%	0.778%	0.791%	0.936%	0.967%
Skew	0.428	0.552		0.395	0.180	0.728	1.863	0.326	2.104	2.336
Kurtosis	7.442	9.414		9.318	8.805	13.197	38.734	8.617	38.043	41.704
JB p-value	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
EUR										
Mean Ret. (%)	0.006%	0.011% **	-0.007%		-0.004%	0.008% *	0.013% ***	0.007%	0.011%	0.013%
Std. Dev. (%)	0.668%	0.503%	0.669%		0.364%	0.418%	0.488%	0.687%	0.820%	0.858%
Skew	0.005	0.597	-0.395		0.011	1.622	5.945	-0.032	2.075	2.731
Kurtosis	8.628	12.931	9.318		39.435	39.554	180.464	7.634	44.364	57.236
JB p-value	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
CHF										
Mean Ret. (%)	0.010%	0.014% **	-0.003%	0.004%		0.011% **	0.016% ***	0.011%	0.014%	0.017% *
Std. Dev. (%)	0.755%	0.600%	0.698%	0.364%		0.542%	0.602%	0.782%	0.909%	0.939%
Skew	0.001	0.359	-0.180	-0.011		0.986	3.289	-0.006	1.660	2.191
Kurtosis	8.549	12.012	8.805	39.435		20.948	83.462	7.489	32.958	42.204
JB p-value	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%
NOK										
Mean Ret. (%)	-0.001%	0.003%	-0.014% *	-0.008% *	-0.011% **		0.005%	-0.001%	0.003%	0.005%
Std. Dev. (%)	0.686%	0.552%	0.736%	0.418%	0.542%		0.467%	0.686%	0.810%	0.856%
Skew	-0.239	-0.241	-0.728	-1.622	-0.986		3.499	-0.347	2.022	2.601
Kurtosis	11.707	16.096	13.197	39.554	20.948		144.991	10.362	48.292	58.227
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%
SEK										
Mean Ret. (%)	-0.006%	-0.002%	-0.019% **	-0.013% ***	-0.016% ***	-0.005%		-0.006%	-0.002%	0.000%
Std. Dev. (%)	0.707%	0.601%	0.778%	0.488%	0.602%	0.467%		0.700%	0.827%	0.880%
Skew	-1.643	-2.941	-1.863	-5.945	-3.289	-3.499		-1.875	0.810	1.434
Kurtosis	42.822	82.214	38.734	180.464	83.462	144.991		44.905	63.166	68.954
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%
CAD										
Mean Ret. (%)	-0.001%	0.004%	-0.014% *	-0.007%	-0.011%	0.001%	0.006%		0.004%	0.006%
Std. Dev. (%)	0.413%	0.648%	0.791%	0.687%	0.782%	0.686%	0.700%		0.690%	0.783%
Skew	0.249	0.131	-0.326	0.032	0.006	0.347	1.875		2.655	3.585
Kurtosis	16.633	6.550	8.617	7.634	7.489	10.362	44.905		65.967	84.125
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%
AUD										
Mean Ret. (%)	-0.004%	0.000%	-0.017% *	-0.011%	-0.014%	-0.003%	0.002%	-0.004%		0.002%
Std. Dev. (%)	0.737%	0.783%	0.936%	0.820%	0.909%	0.810%	0.827%	0.690%		0.697%
Skew	-3.022	-2.026	-2.104	-2.075	-1.660	-2.022	-0.810	-2.655		1.189
Kurtosis	72.132	50.741	38.043	44.364	32.958	48.292	63.166	65.967		161.919
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%
NZD										
Mean Ret. (%)	-0.007%	-0.002%	-0.020% **	-0.013%	-0.017% *	-0.005%	0.000%	-0.006%	-0.002%	
Std. Dev. (%)	0.808%	0.825%	0.967%	0.858%	0.939%	0.856%	0.880%	0.783%	0.697%	
Skew	-3.438	-2.955	-2.336	-2.731	-2.191	-2.601	-1.434	-3.585	-1.189	
Kurtosis	78.724	66.427	41.704	57.236	42.204	58.227	68.954	84.125	161.919	
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	

The dataset consists of daily returns for individual currencies from the 4th of January 1974 to the 31st of December of 2009. The column labels denote base currency calculations and row labels denote foreign currencies returns against the base currency. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

2. Descriptive Statistics (LOG 3M T-bill Interest Rate Adj. Currency Returns)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean Ret. (%)		-0.002%	-0.006%	-0.005%	-0.004%	-0.004%	0.003%	-0.003%	-0.003%	-0.004%
Std. Dev. (%)		0.623%	0.690%	0.668%	0.755%	0.686%	0.707%	0.413%	0.737%	0.808%
Skew		0.115	-0.425	-0.007	-0.002	0.231	1.635	-0.242	3.011	3.423
Kurtosis		7.889	7.434	8.627	8.552	11.685	42.770	16.621	72.022	78.628
JB p-value		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
GBP										
Mean Ret. (%)	0.002%		-0.004%	-0.003%	-0.002%	-0.002%	0.005%	-0.001%	-0.001%	-0.002%
Std. Dev. (%)	0.623%		0.729%	0.504%	0.601%	0.552%	0.602%	0.648%	0.783%	0.825%
Skew	-0.115		-0.555	-0.596	-0.358	0.235	2.930	-0.131	2.025	2.944
Kurtosis	7.889		9.412	12.900	11.999	16.054	82.052	6.545	50.813	66.428
JB p-value	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
JPY										
Mean Ret. (%)	0.006%	0.004%		0.001%	0.001%	0.002%	0.009%	0.003%	0.003%	0.002%
Std. Dev. (%)	0.690%	0.729%		0.670%	0.698%	0.736%	0.778%	0.792%	0.936%	0.967%
Skew	0.425	0.555		0.389	0.177	0.727	1.865	0.329	2.105	2.331
Kurtosis	7.434	9.412		9.305	8.799	13.176	38.714	8.615	38.072	41.705
JB p-value	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
EUR										
Mean Ret. (%)	0.005%	0.003%	-0.001%		0.001%	0.001%	0.008%	0.002%	0.002%	0.001%
Std. Dev. (%)	0.668%	0.504%	0.670%		0.364%	0.418%	0.488%	0.687%	0.820%	0.858%
Skew	0.007	0.596	-0.389		0.012	1.619	5.947	-0.032	2.073	2.717
Kurtosis	8.627	12.900	9.305		39.393	39.457	180.397	7.637	44.353	57.170
JB p-value	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
CHF										
Mean Ret. (%)	0.004%	0.002%	-0.001%	-0.001%		0.000%	0.007%	0.001%	0.001%	0.001%
Std. Dev. (%)	0.755%	0.601%	0.698%	0.364%		0.542%	0.602%	0.783%	0.909%	0.939%
Skew	0.002	0.358	-0.177	-0.012		0.981	3.286	-0.006	1.659	2.182
Kurtosis	8.552	11.999	8.799	39.393		20.897	83.356	7.494	32.940	42.173
JB p-value	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%
NOK										
Mean Ret. (%)	0.004%	0.002%	-0.002%	-0.001%	0.000%		0.007%	0.001%	0.001%	0.000%
Std. Dev. (%)	0.686%	0.552%	0.736%	0.418%	0.542%		0.467%	0.686%	0.810%	0.856%
Skew	-0.231	-0.235	-0.727	-1.619	-0.981		3.495	-0.341	2.024	2.597
Kurtosis	11.685	16.054	13.176	39.457	20.897		144.769	10.348	48.255	58.205
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%
SEK										
Mean Ret. (%)	-0.003%	-0.005%	-0.009%	-0.008%	-0.007%	-0.007%		-0.006%	-0.006%	-0.007%
Std. Dev. (%)	0.707%	0.602%	0.778%	0.488%	0.602%	0.467%		0.700%	0.827%	0.880%
Skew	-1.635	-2.930	-1.865	-5.947	-3.286	-3.495		-1.869	0.812	1.429
Kurtosis	42.770	82.052	38.714	180.397	83.356	144.769		44.904	63.182	68.927
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%
CAD										
Mean Ret. (%)	0.003%	0.001%	-0.003%	-0.002%	-0.001%	-0.001%	0.006%		0.000%	0.000%
Std. Dev. (%)	0.413%	0.648%	0.792%	0.687%	0.783%	0.686%	0.700%		0.690%	0.783%
Skew	0.242	0.131	-0.329	0.032	0.006	0.341	1.869		2.646	3.570
Kurtosis	16.621	6.545	8.615	7.637	7.494	10.348	44.904		65.936	84.062
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%
AUD										
Mean Ret. (%)	0.003%	0.001%	-0.003%	-0.002%	-0.001%	-0.001%	0.006%	0.000%		-0.001%
Std. Dev. (%)	0.737%	0.783%	0.936%	0.820%	0.909%	0.810%	0.827%	0.690%		0.697%
Skew	-3.011	-2.025	-2.105	-2.073	-1.659	-2.024	-0.812	-2.646		1.183
Kurtosis	72.022	50.813	38.072	44.353	32.940	48.255	63.182	65.936		161.708
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%
NZD										
Mean Ret. (%)	0.004%	0.002%	-0.002%	-0.001%	-0.001%	0.000%	0.007%	0.000%	0.001%	
Std. Dev. (%)	0.808%	0.825%	0.967%	0.858%	0.939%	0.856%	0.880%	0.783%	0.697%	
Skew	-3.423	-2.944	-2.331	-2.717	-2.182	-2.597	-1.429	-3.570	-1.183	
Kurtosis	78.628	66.428	41.705	57.170	42.173	58.205	68.927	84.062	161.708	
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	

The dataset consists of daily returns for individual currencies from the 4th of January 1974 to the 31st of December of 2009. The column labels denote base currency calculations and row labels denote foreign currencies returns against the base currency. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

3. Descriptive Statistics (LOG O/N Rate Interest Rate Adj. Currency Returns)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean Ret. (%)		-0.002%	-0.005%	-0.003%	-0.001%	-0.002%	0.005%	-0.001%	-0.002%	-0.002%
Std. Dev. (%)		0.623%	0.690%	0.668%	0.755%	0.687%	0.707%	0.413%	0.737%	0.808%
Skew		0.115	-0.423	-0.005	-0.001	0.232	1.634	-0.250	3.010	3.423
Kurtosis		7.888	7.432	8.627	8.556	11.685	42.783	16.620	72.015	78.628
JB p-value		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
GBP										
Mean Ret. (%)	0.002%		-0.004%	-0.002%	0.001%	0.000%	0.006%	0.000%	0.000%	-0.001%
Std. Dev. (%)	0.623%		0.729%	0.504%	0.600%	0.552%	0.602%	0.648%	0.783%	0.825%
Skew	-0.115		-0.553	-0.593	-0.357	0.238	2.931	-0.131	2.026	2.945
Kurtosis	7.888		9.409	12.886	12.000	16.045	82.058	6.541	50.848	66.431
JB p-value	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
JPY										
Mean Ret. (%)	0.005%	0.004%		0.002%	0.004%	0.003%	0.010%	0.004%	0.003%	0.003%
Std. Dev. (%)	0.690%	0.729%		0.670%	0.698%	0.736%	0.778%	0.792%	0.936%	0.967%
Skew	0.423	0.553		0.389	0.175	0.726	1.862	0.326	2.104	2.331
Kurtosis	7.432	9.409		9.308	8.794	13.178	38.716	8.610	38.066	41.707
JB p-value	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
EUR										
Mean Ret. (%)	0.003%	0.002%	-0.002%		0.002%	0.001%	0.008%	0.002%	0.001%	0.001%
Std. Dev. (%)	0.668%	0.504%	0.670%		0.364%	0.418%	0.488%	0.687%	0.820%	0.858%
Skew	0.005	0.593	-0.389		0.013	1.618	5.946	-0.032	2.073	2.717
Kurtosis	8.627	12.886	9.308		39.438	39.453	180.413	7.636	44.355	57.176
JB p-value	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
CHF										
Mean Ret. (%)	0.001%	-0.001%	-0.004%	-0.002%		-0.001%	0.006%	0.000%	-0.001%	-0.001%
Std. Dev. (%)	0.755%	0.600%	0.698%	0.364%		0.542%	0.602%	0.783%	0.909%	0.939%
Skew	0.001	0.357	-0.175	-0.013		0.981	3.284	-0.005	1.658	2.180
Kurtosis	8.556	12.000	8.794	39.438		20.912	83.298	7.494	32.925	42.134
JB p-value	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%
NOK										
Mean Ret. (%)	0.002%	0.000%	-0.003%	-0.001%	0.001%		0.007%	0.001%	0.000%	0.000%
Std. Dev. (%)	0.687%	0.552%	0.736%	0.418%	0.542%		0.467%	0.686%	0.810%	0.856%
Skew	-0.232	-0.238	-0.726	-1.618	-0.981		3.494	-0.342	2.023	2.597
Kurtosis	11.685	16.045	13.178	39.453	20.912		144.783	10.344	48.257	58.212
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%
SEK										
Mean Ret. (%)	-0.005%	-0.006%	-0.010%	-0.008%	-0.006%	-0.007%		-0.006%	-0.007%	-0.007%
Std. Dev. (%)	0.707%	0.602%	0.778%	0.488%	0.602%	0.467%		0.700%	0.827%	0.880%
Skew	-1.634	-2.931	-1.862	-5.946	-3.284	-3.494		-1.870	0.813	1.429
Kurtosis	42.783	82.058	38.716	180.413	83.298	144.783		44.907	63.188	68.931
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%
CAD										
Mean Ret. (%)	0.001%	0.000%	-0.004%	-0.002%	0.000%	-0.001%	0.006%		-0.001%	-0.001%
Std. Dev. (%)	0.413%	0.648%	0.792%	0.687%	0.783%	0.686%	0.700%		0.690%	0.783%
Skew	0.250	0.131	-0.326	0.032	0.005	0.342	1.870		2.648	3.570
Kurtosis	16.620	6.541	8.610	7.636	7.494	10.344	44.907		65.952	84.022
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%
AUD										
Mean Ret. (%)	0.002%	0.000%	-0.003%	-0.001%	0.001%	0.000%	0.007%	0.001%		0.000%
Std. Dev. (%)	0.737%	0.783%	0.936%	0.820%	0.909%	0.810%	0.827%	0.690%		0.697%
Skew	-3.010	-2.026	-2.104	-2.073	-1.658	-2.023	-0.813	-2.648		1.182
Kurtosis	72.015	50.848	38.066	44.355	32.925	48.257	63.188	65.952		161.704
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%
NZD										
Mean Ret. (%)	0.002%	0.001%	-0.003%	-0.001%	0.001%	0.000%	0.007%	0.001%	0.000%	
Std. Dev. (%)	0.808%	0.825%	0.967%	0.858%	0.939%	0.856%	0.880%	0.783%	0.697%	
Skew	-3.423	-2.945	-2.331	-2.717	-2.180	-2.597	-1.429	-3.570	-1.182	
Kurtosis	78.628	66.431	41.707	57.176	42.134	58.212	68.931	84.022	161.704	
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	

The dataset consists of daily returns for individual currencies from the 4th of January 1974 to the 31st of December of 2009. The column labels denote base currency calculations and row labels denote foreign currencies returns against the base currency. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

C. APPENDIX 3: Trading Rule Analysis including Bid/Ask Spreads

1. Dataset 2: Descriptive Statistics (Simple Base Currency Returns)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean Ret. (%)		-0.004%	-0.016% **	-0.023% ***	-0.022% ***	-0.017% **	-0.015% *	-0.018% ***	-0.022% **	-0.020% **
Std. Dev. (%)		0.644%	0.669%	0.628%	0.686%	0.824%	0.808%	0.671%	0.943%	0.947%
Skew		0.192	-0.548	-0.064	-0.156	-0.039	-0.220	-0.038	1.347	0.617
Kurtosis		7.906	6.816	6.811	6.765	8.064	7.202	8.178	18.902	9.178
JB p-value		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
GBP										
Mean Ret. (%)	0.008%		-0.008%	-0.017% ***	-0.016% ***	-0.013% *	-0.010%	-0.012% *	-0.017% **	-0.016% **
Std. Dev. (%)	0.644%		0.855%	0.489%	0.601%	0.649%	0.631%	0.664%	0.761%	0.764%
Skew	-0.058		-0.781	-0.262	-0.516	-0.380	-0.282	0.002	1.037	0.530
Kurtosis	8.036		12.292	7.779	9.670	8.259	8.617	5.169	15.941	6.981
JB p-value	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
JPY										
Mean Ret. (%)	0.020% ***	0.016% *		-0.004%	-0.004%	0.002%	0.004%	0.002%	-0.001%	0.000%
Std. Dev. (%)	0.672%	0.862%		0.751%	0.688%	0.959%	0.948%	0.952%	1.189%	1.174%
Skew	0.662	1.061		0.880	0.411	0.749	0.517	0.526	2.275	1.428
Kurtosis	7.313	13.268		11.808	8.828	11.116	9.826	10.285	30.503	16.810
JB p-value	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
EUR										
Mean Ret. (%)	0.027% ***	0.020% ***	0.009%		0.001%	0.005%	0.007% *	0.006%	0.001%	0.002%
Std. Dev. (%)	0.629%	0.490%	0.747%		0.295%	0.473%	0.424%	0.630%	0.742%	0.761%
Skew	0.176	0.360	-0.663		-0.142	0.148	0.195	0.018	1.753	0.876
Kurtosis	7.053	7.873	10.799		12.708	8.234	7.912	5.528	22.990	8.828
JB p-value	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
CHF										
Mean Ret. (%)	0.026% ***	0.020% ***	0.008%	0.000%		0.005%	0.007%	0.006%	0.002%	0.003%
Std. Dev. (%)	0.688%	0.603%	0.686%	0.295%		0.586%	0.557%	0.745%	0.891%	0.897%
Skew	0.277	0.670	-0.258	0.245		0.575	0.220	0.105	1.804	0.976
Kurtosis	7.076	10.096	8.370	12.500		8.288	7.856	6.341	23.103	10.451
JB p-value	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%
NOK										
Mean Ret. (%)	0.024% ***	0.017% ***	0.007%	-0.003%	-0.001%		0.004%	0.003%	-0.002%	-0.001%
Std. Dev. (%)	0.825%	0.651%	0.953%	0.473%	0.586%		0.481%	0.732%	0.792%	0.816%
Skew	0.219	0.523	-0.487	-0.044	-0.454		-0.034	0.120	2.041	1.244
Kurtosis	8.567	8.861	10.062	8.448	8.017		8.096	5.106	29.318	15.673
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%
SEK										
Mean Ret. (%)	0.021% ***	0.014% **	0.005%	-0.006%	-0.004%	-0.002%		0.000%	-0.005%	-0.004%
Std. Dev. (%)	0.810%	0.632%	0.944%	0.423%	0.556%	0.481%		0.703%	0.762%	0.786%
Skew	0.374	0.427	-0.276	-0.108	-0.106	0.135		0.159	1.444	0.665
Kurtosis	7.675	9.000	9.572	7.863	7.835	7.976		4.687	17.954	7.446
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%
CAD										
Mean Ret. (%)	0.023% ***	0.017% **	0.007%	-0.002%	-0.001%	0.002%	0.005%		-0.003%	-0.001%
Std. Dev. (%)	0.672%	0.664%	0.948%	0.630%	0.745%	0.732%	0.702%		0.743%	0.792%
Skew	0.185	0.081	-0.271	0.069	0.015	-0.031	-0.083		0.900	0.583
Kurtosis	8.450	5.191	10.070	5.669	6.409	5.020	4.625		19.132	7.709
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%
AUD										
Mean Ret. (%)	0.031% ***	0.023% ***	0.015%	0.004%	0.006%	0.008%	0.011%	0.009%		0.002%
Std. Dev. (%)	0.933%	0.756%	1.163%	0.734%	0.879%	0.781%	0.755%	0.739%		0.500%
Skew	-0.915	-0.730	-1.487	-1.354	-1.334	-1.529	-1.124	-0.519		0.045
Kurtosis	16.587	14.724	23.922	20.068	19.461	22.879	15.696	18.420		6.555
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%
NZD										
Mean Ret. (%)	0.029% ***	0.022% ***	0.014%	0.003%	0.005%	0.007%	0.010%	0.008%	0.000%	
Std. Dev. (%)	0.943%	0.762%	1.158%	0.757%	0.890%	0.809%	0.782%	0.789%	0.500%	
Skew	-0.398	-0.401	-0.978	-0.721	-0.758	-0.953	-0.527	-0.435	0.039	
Kurtosis	8.851	6.720	14.061	8.160	9.412	12.766	7.070	7.292	6.673	
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	

The dataset spans from the 27th of March 2002 to 31st of December 2009. The column labels denote base currency calculations and row labels denote foreign currencies returns against the base currency. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

2. Dataset 2: Descriptive Statistics (3M T-bill Adj. Currency Returns)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean Ret. (%)		-0.009%	-0.010%	-0.024% ***	-0.019% ***	-0.022% ***	-0.016% *	-0.019% ***	-0.028% ***	-0.030% ***
Std. Dev. (%)		0.644%	0.669%	0.628%	0.686%	0.824%	0.808%	0.671%	0.942%	0.947%
Skew		0.192	-0.555	-0.074	-0.163	-0.049	-0.224	-0.040	1.339	0.613
Kurtosis		7.904	6.822	6.808	6.767	8.061	7.203	8.180	18.883	9.172
JB p-value		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
GBP										
Mean Ret. (%)	0.013% **		0.002%	-0.014% ***	-0.009%	-0.013% **	-0.006%	-0.009%	-0.019% **	-0.021% ***
Std. Dev. (%)	0.644%		0.855%	0.489%	0.601%	0.649%	0.631%	0.664%	0.761%	0.764%
Skew	-0.058		-0.789	-0.277	-0.531	-0.393	-0.291	0.000	1.031	0.526
Kurtosis	8.035		12.297	7.784	9.681	8.275	8.619	5.169	15.931	6.981
JB p-value	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
JPY										
Mean Ret. (%)	0.015% **	0.006%		-0.011%	-0.006%	-0.008%	-0.002%	-0.004%	-0.013%	-0.015%
Std. Dev. (%)	0.672%	0.862%		0.751%	0.688%	0.959%	0.948%	0.952%	1.189%	1.174%
Skew	0.669	1.069		0.878	0.411	0.746	0.521	0.532	2.276	1.429
Kurtosis	7.321	13.275		11.793	8.822	11.099	9.827	10.289	30.504	16.808
JB p-value	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
EUR										
Mean Ret. (%)	0.028% ***	0.017% ***	0.016% **		0.006% *	0.002%	0.008% *	0.007%	-0.004%	-0.006%
Std. Dev. (%)	0.629%	0.490%	0.747%		0.295%	0.473%	0.424%	0.630%	0.742%	0.761%
Skew	0.185	0.376	-0.660		-0.142	0.145	0.202	0.024	1.759	0.880
Kurtosis	7.053	7.882	10.787		12.689	8.234	7.915	5.528	23.020	8.835
JB p-value	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
CHF										
Mean Ret. (%)	0.023% ***	0.012% **	0.011%	-0.005%		-0.003%	0.004%	0.002%	-0.008%	-0.010%
Std. Dev. (%)	0.688%	0.603%	0.686%	0.295%		0.588%	0.557%	0.745%	0.891%	0.897%
Skew	0.284	0.684	-0.257	0.245		0.573	0.229	0.112	1.810	0.980
Kurtosis	7.081	10.111	8.366	12.482		8.282	7.858	6.343	23.132	10.457
JB p-value	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%
NOK										
Mean Ret. (%)	0.029% ***	0.017% ***	0.018% *	0.001%	0.006%		0.008%	0.007%	-0.004%	-0.006%
Std. Dev. (%)	0.825%	0.651%	0.953%	0.473%	0.586%		0.481%	0.733%	0.792%	0.816%
Skew	0.229	0.536	-0.483	-0.041	-0.452		-0.027	0.127	2.048	1.248
Kurtosis	8.568	8.882	10.050	8.449	8.012		8.097	5.108	29.354	15.692
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%
SEK										
Mean Ret. (%)	0.022% ***	0.010%	0.011%	-0.006%	0.000%	-0.005%		0.000%	-0.011%	-0.013%
Std. Dev. (%)	0.810%	0.632%	0.944%	0.423%	0.556%	0.481%		0.703%	0.762%	0.786%
Skew	0.378	0.436	-0.280	-0.115	-0.115	0.128		0.162	1.444	0.665
Kurtosis	7.678	9.004	9.573	7.865	7.835	7.975		4.689	17.958	7.446
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%
CAD										
Mean Ret. (%)	0.024% ***	0.013% **	0.013%	-0.003%	0.003%	-0.001%	0.005%		-0.008%	-0.010%
Std. Dev. (%)	0.672%	0.664%	0.949%	0.630%	0.745%	0.732%	0.702%		0.743%	0.792%
Skew	0.187	0.083	-0.277	0.063	0.008	-0.038	-0.085		0.893	0.580
Kurtosis	8.453	5.191	10.071	5.667	6.408	5.019	4.626		19.121	7.709
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%
AUD										
Mean Ret. (%)	0.037% ***	0.025% ***	0.027% **	0.009%	0.015% *	0.010%	0.016% **	0.014% *		-0.001%
Std. Dev. (%)	0.933%	0.757%	1.163%	0.734%	0.879%	0.781%	0.755%	0.739%		0.500%
Skew	-0.907	-0.724	-1.488	-1.361	-1.340	-1.536	-1.124	-0.512		0.045
Kurtosis	16.575	14.716	23.924	20.091	19.483	22.908	15.700	18.418		6.561
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%
NZD										
Mean Ret. (%)	0.039% ***	0.027% ***	0.029% **	0.011%	0.018% **	0.012%	0.019% **	0.016% **	0.003%	
Std. Dev. (%)	0.943%	0.762%	1.158%	0.757%	0.890%	0.809%	0.782%	0.789%	0.500%	
Skew	-0.394	-0.397	-0.979	-0.724	-0.762	-0.957	-0.527	-0.432	0.039	
Kurtosis	8.847	6.721	14.060	8.165	9.416	12.781	7.070	7.294	6.679	
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	

The dataset spans from the 27th of March 2002 to 31st of December 2009. The column labels denote base currency calculations and row labels denote foreign currencies returns against the base currency. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

3. Dataset 2: Descriptive Statistics (O/N Rate Adj. Currency Returns)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean Ret. (%)		-0.008%	-0.009%	-0.023% ***	-0.017% **	-0.020% **	-0.015% *	-0.019% ***	-0.029% ***	-0.030% ***
Std. Dev. (%)		0.644%	0.669%	0.628%	0.686%	0.824%	0.808%	0.671%	0.942%	0.947%
Skew		0.191	-0.554	-0.072	-0.162	-0.047	-0.227	-0.042	1.341	0.615
Kurtosis		7.906	6.823	6.809	6.767	8.065	7.207	8.178	18.891	9.172
JB p-value		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
GBP										
Mean Ret. (%)	0.012% *		0.002%	-0.013% ***	-0.008%	-0.012% *	-0.007%	-0.009%	-0.020% ***	-0.022% ***
Std. Dev. (%)	0.644%		0.855%	0.489%	0.601%	0.649%	0.631%	0.663%	0.761%	0.764%
Skew	-0.057		-0.788	-0.271	-0.524	-0.390	-0.292	0.001	1.033	0.528
Kurtosis	8.037		12.298	7.781	9.671	8.270	8.620	5.167	15.937	6.979
JB p-value	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
JPY										
Mean Ret. (%)	0.014% **	0.005%		-0.010%	-0.006%	-0.008%	-0.003%	-0.005%	-0.015%	-0.016%
Std. Dev. (%)	0.672%	0.862%		0.751%	0.688%	0.959%	0.948%	0.952%	1.189%	1.174%
Skew	0.668	1.068		0.882	0.414	0.748	0.519	0.531	2.278	1.430
Kurtosis	7.322	13.276		11.807	8.816	11.106	9.824	10.285	30.516	16.815
JB p-value	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
EUR										
Mean Ret. (%)	0.027% ***	0.016% ***	0.016% **		0.006% *	0.002%	0.007%	0.006%	-0.006%	-0.007%
Std. Dev. (%)	0.629%	0.490%	0.747%		0.295%	0.473%	0.424%	0.630%	0.742%	0.761%
Skew	0.183	0.369	-0.665		-0.142	0.143	0.193	0.020	1.754	0.877
Kurtosis	7.053	7.877	10.798		12.679	8.233	7.912	5.527	22.986	8.825
JB p-value	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
CHF										
Mean Ret. (%)	0.022% ***	0.011% *	0.010%	-0.005%		-0.002%	0.002%	0.001%	-0.010%	-0.011%
Std. Dev. (%)	0.688%	0.603%	0.686%	0.295%		0.588%	0.557%	0.745%	0.891%	0.897%
Skew	0.282	0.677	-0.260	0.244		0.572	0.220	0.108	1.806	0.977
Kurtosis	7.080	10.099	8.359	12.472		8.288	7.853	6.342	23.110	10.446
JB p-value	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%	0.000%
NOK										
Mean Ret. (%)	0.027% ***	0.016% **	0.017% *	0.000%	0.006%		0.006%	0.005%	-0.007%	-0.008%
Std. Dev. (%)	0.825%	0.651%	0.953%	0.473%	0.586%		0.481%	0.733%	0.792%	0.816%
Skew	0.227	0.533	-0.486	-0.039	-0.451		-0.032	0.125	2.045	1.246
Kurtosis	8.571	8.875	10.055	8.449	8.018		8.091	5.108	29.332	15.683
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%	0.000%
SEK										
Mean Ret. (%)	0.022% ***	0.011% *	0.012%	-0.005%	0.001%	-0.004%		0.000%	-0.012%	-0.013% *
Std. Dev. (%)	0.810%	0.632%	0.944%	0.423%	0.556%	0.481%		0.703%	0.762%	0.786%
Skew	0.380	0.437	-0.277	-0.106	-0.106	0.133		0.162	1.448	0.667
Kurtosis	7.682	9.006	9.571	7.862	7.832	7.972		4.689	17.970	7.449
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%	0.000%
CAD										
Mean Ret. (%)	0.023% ***	0.013% **	0.014%	-0.002%	0.004%	0.000%	0.005%		-0.010%	-0.011%
Std. Dev. (%)	0.672%	0.664%	0.948%	0.630%	0.745%	0.732%	0.702%		0.743%	0.792%
Skew	0.188	0.082	-0.276	0.067	0.012	-0.036	-0.086		0.896	0.581
Kurtosis	8.452	5.188	10.069	5.667	6.409	5.020	4.626		19.136	7.707
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%	0.000%
AUD										
Mean Ret. (%)	0.038% ***	0.026% ***	0.029% **	0.011%	0.018% **	0.013%	0.018% **	0.015% **		0.000%
Std. Dev. (%)	0.933%	0.756%	1.163%	0.734%	0.879%	0.781%	0.755%	0.739%		0.500%
Skew	-0.909	-0.726	-1.490	-1.355	-1.336	-1.533	-1.127	-0.514		0.044
Kurtosis	16.581	14.722	23.931	20.066	19.467	22.891	15.710	18.431		6.563
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%		0.000%
NZD										
Mean Ret. (%)	0.039% ***	0.028% ***	0.030% ***	0.013% *	0.019% **	0.014% *	0.019% **	0.017% **	0.003%	
Std. Dev. (%)	0.943%	0.762%	1.158%	0.757%	0.890%	0.809%	0.782%	0.788%	0.500%	
Skew	-0.396	-0.399	-0.981	-0.721	-0.759	-0.956	-0.529	-0.434	0.040	
Kurtosis	8.846	6.719	14.065	8.157	9.407	12.774	7.073	7.292	6.682	
JB p-value	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	


The dataset spans from the 27th of March 2002 to 31st of December 2009. This time period coincides with the last two sub-samples of the first dataset. The column labels denote base currency calculations and row labels denote foreign currencies returns against the base currency. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

4. Dataset 2: Differences in IR adjusted returns (3M T-bill vs. Overnight rate)

	USD	GBP	JPY	EUR	CHF	NOK	SEK	CAD	AUD	NZD
USD										
Mean (3M Tbill adjustment)		-0.009%	-0.010%	-0.024%	-0.019%	-0.022%	-0.016%	-0.019%	-0.028%	-0.030%
Mean (O/N Rate adjustment)		-0.008%	-0.009%	-0.023%	-0.017%	-0.020%	-0.015%	-0.019%	-0.029%	-0.030%
p-value(t-test)		0.98	0.96	0.94	0.94	0.94	0.99	0.98	0.98	1.00
p-value(Satterthwaite-Welch t-test*)		0.98	0.96	0.94	0.94	0.94	0.99	0.98	0.98	1.00
p-value(Wilcoxon/Mann-Whitney)		0.96	0.93	0.91	0.92	0.91	0.98	0.97	0.95	0.99
p-value(Wilcoxon/Mann-Whitney (tie-adj.))		0.96	0.93	0.91	0.92	0.91	0.98	0.97	0.95	0.99
Correlation		1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000
GBP										
Mean (3M Tbill adjustment)	0.013%		0.002%	-0.014%	-0.009%	-0.013%	-0.006%	-0.009%	-0.019%	-0.021%
Mean (O/N Rate adjustment)	0.012%		0.002%	-0.013%	-0.008%	-0.012%	-0.007%	-0.009%	-0.020%	-0.022%
p-value(t-test)	0.98		0.98	0.95	0.96	0.95	0.99	0.99	0.95	0.98
p-value(Satterthwaite-Welch t-test*)	0.98		0.98	0.95	0.96	0.95	0.99	0.99	0.95	0.98
p-value(Wilcoxon/Mann-Whitney)	0.96		0.96	0.93	0.94	0.92	0.99	0.99	0.92	0.96
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.96		0.96	0.93	0.94	0.92	0.99	0.99	0.92	0.96
Correlation	1.00000		1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000
JPY										
Mean (3M Tbill adjustment)	0.015%	0.006%		-0.011%	-0.006%	-0.008%	-0.002%	-0.004%	-0.013%	-0.015%
Mean (O/N Rate adjustment)	0.014%	0.005%		-0.010%	-0.006%	-0.008%	-0.003%	-0.005%	-0.015%	-0.016%
p-value(t-test)	0.96	0.98		0.99	0.98	0.98	0.98	0.98	0.96	0.98
p-value(Satterthwaite-Welch t-test*)	0.96	0.98		0.99	0.98	0.98	0.98	0.98	0.96	0.98
p-value(Wilcoxon/Mann-Whitney)	0.93	0.96		0.98	0.98	0.96	0.96	0.97	0.91	0.95
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.93	0.96		0.98	0.98	0.96	0.96	0.97	0.91	0.95
Correlation	1.00000	1.00000		1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000
EUR										
Mean (3M Tbill adjustment)	0.028%	0.017%	0.016%		0.006%	0.002%	0.008%	0.007%	-0.004%	-0.006%
Mean (O/N Rate adjustment)	0.027%	0.016%	0.016%		0.006%	0.002%	0.007%	0.006%	-0.006%	-0.007%
p-value(t-test)	0.94	0.95	0.99		1.00	0.98	0.93	0.96	0.92	0.95
p-value(Satterthwaite-Welch t-test*)	0.94	0.95	0.99		1.00	0.98	0.93	0.96	0.92	0.95
p-value(Wilcoxon/Mann-Whitney)	0.91	0.93	0.98		1.00	0.95	0.91	0.95	0.87	0.93
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.91	0.93	0.98		1.00	0.95	0.91	0.95	0.87	0.93
Correlation	1.00000	1.00000	1.00000		1.00000	1.00000	0.99999	1.00000	1.00000	1.00000
CHF										
Mean (3M Tbill adjustment)	0.023%	0.012%	0.011%	-0.005%		-0.003%	0.004%	0.002%	-0.008%	-0.010%
Mean (O/N Rate adjustment)	0.022%	0.011%	0.010%	-0.005%		-0.002%	0.002%	0.001%	-0.010%	-0.011%
p-value(t-test)	0.94	0.96	0.98	1.00		0.98	0.95	0.96	0.93	0.96
p-value(Satterthwaite-Welch t-test*)	0.94	0.96	0.98	1.00		0.98	0.95	0.96	0.93	0.96
p-value(Wilcoxon/Mann-Whitney)	0.92	0.94	0.98	1.00		0.96	0.93	0.95	0.89	0.93
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.92	0.94	0.98	1.00		0.96	0.93	0.95	0.89	0.93
Correlation	1.00000	1.00000	1.00000	1.00000		1.00000	1.00000	1.00000	1.00000	1.00000
NOK										
Mean (3M Tbill adjustment)	0.029%	0.017%	0.018%	0.001%	0.006%		0.008%	0.007%	-0.004%	-0.006%
Mean (O/N Rate adjustment)	0.027%	0.016%	0.017%	0.000%	0.006%		0.006%	0.005%	-0.007%	-0.008%
p-value(t-test)	0.94	0.95	0.98	0.98	0.98		0.92	0.95	0.91	0.94
p-value(Satterthwaite-Welch t-test*)	0.94	0.95	0.98	0.98	0.98		0.92	0.95	0.91	0.94
p-value(Wilcoxon/Mann-Whitney)	0.91	0.92	0.96	0.95	0.96		0.89	0.92	0.87	0.91
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.91	0.92	0.96	0.95	0.96		0.89	0.92	0.87	0.91
Correlation	1.00000	1.00000	1.00000	1.00000	1.00000		1.00000	1.00000	1.00000	1.00000
SEK										
Mean (3M Tbill adjustment)	0.022%	0.010%	0.011%	-0.006%	0.000%	-0.005%		0.000%	-0.011%	-0.013%
Mean (O/N Rate adjustment)	0.022%	0.011%	0.012%	-0.005%	0.001%	-0.004%		0.000%	-0.012%	-0.013%
p-value(t-test)	0.99	0.99	0.98	0.93	0.95	0.92		1.00	0.96	0.99
p-value(Satterthwaite-Welch t-test*)	0.99	0.99	0.98	0.93	0.95	0.92		1.00	0.96	0.99
p-value(Wilcoxon/Mann-Whitney)	0.98	0.99	0.96	0.91	0.93	0.89		0.99	0.93	0.97
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.98	0.99	0.96	0.91	0.93	0.89		0.99	0.93	0.97
Correlation	1.00000	1.00000	1.00000	0.99999	1.00000	1.00000		1.00000	1.00000	1.00000
CAD										
Mean (3M Tbill adjustment)	0.024%	0.013%	0.013%	-0.003%	0.003%	-0.001%	0.005%		-0.008%	-0.010%
Mean (O/N Rate adjustment)	0.023%	0.013%	0.014%	-0.002%	0.004%	0.000%	0.005%		-0.010%	-0.011%
p-value(t-test)	0.98	0.99	0.98	0.96	0.96	0.95	1.00		0.96	0.99
p-value(Satterthwaite-Welch t-test*)	0.98	0.99	0.98	0.96	0.96	0.95	1.00		0.96	0.99
p-value(Wilcoxon/Mann-Whitney)	0.97	0.99	0.97	0.95	0.95	0.92	0.99		0.93	0.97
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.97	0.99	0.97	0.95	0.95	0.92	0.99		0.93	0.97
Correlation	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000		1.00000	1.00000
AUD										
Mean (3M Tbill adjustment)	0.037%	0.025%	0.027%	0.009%	0.015%	0.010%	0.016%	0.014%		-0.001%
Mean (O/N Rate adjustment)	0.038%	0.026%	0.029%	0.011%	0.018%	0.013%	0.018%	0.015%		0.000%
p-value(t-test)	0.98	0.95	0.96	0.92	0.93	0.91	0.96	0.96		0.96
p-value(Satterthwaite-Welch t-test*)	0.98	0.95	0.96	0.92	0.93	0.91	0.96	0.96		0.96
p-value(Wilcoxon/Mann-Whitney)	0.95	0.92	0.91	0.88	0.89	0.87	0.93	0.93		0.94
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.95	0.92	0.91	0.88	0.89	0.87	0.93	0.93		0.94
Correlation	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000		0.99999
NZD										
Mean (3M Tbill adjustment)	0.039%	0.027%	0.029%	0.011%	0.018%	0.012%	0.019%	0.016%	0.003%	
Mean (O/N Rate adjustment)	0.039%	0.028%	0.030%	0.013%	0.019%	0.014%	0.019%	0.017%	0.003%	
p-value(t-test)	1.00	0.98	0.98	0.95	0.96	0.94	0.99	0.99		
p-value(Satterthwaite-Welch t-test*)	1.00	0.98	0.98	0.95	0.96	0.94	0.99	0.99		
p-value(Wilcoxon/Mann-Whitney)	0.99	0.96	0.95	0.93	0.93	0.91	0.97	0.97		
p-value(Wilcoxon/Mann-Whitney (tie-adj.))	0.99	0.96	0.95	0.93	0.93	0.91	0.97	0.97		
Correlation	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	0.99999	

The dataset spans from the 27th of March 2002 to 31st of December 2009. This time period coincides with the last two sub-samples of the first dataset. The column labels denote base currency calculations and row labels denote foreign currencies returns against the base currency. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

5. Dataset 2: BVTG; Enhanced vs. BM Trading Rules (Across Currencies)



The rationale for incorporating the bid-ask spread into the breakeven transaction cost analysis can be explained as follows. In a real trading scenario, the impact of bid-ask spreads is not symmetric. This is due to the fact that there is a continuous compounding element embedded in the evolution of bid/ask spreads. The results presented in the main body of this chapter assume that an incremental amount is deducted each time a transaction is made. Nonetheless, the bid/ask spread remains volatile and it will be widest in periods of stress in the financial system. Therefore, whether a trading strategy is very profitable or very unprofitable in the periods where the bid/ask spread is the widest will have a profound impact on the results of the trading strategy. The Figure above shows the median transaction cost breakeven levels for the benchmark trading rule as well as the enhanced trading rule. The analysis has been carried out for sub-sample 8 and sub-sample 9 using the second dataset, which incorporates bid/ask spreads. The results for the benchmark strategy suggest that it either destroys value as it is the case for sub-sample 8 or fails to add value, as it is the case for sub-sample 9. The results of the enhanced trading strategy on the other hand suggest that short-term trading rules will create value even in periods where a generic trading rule fails to perform.

D. APPENDIX 4: Sensitivity of Survival times based on 36 Sub-Samples

1. Time Series Analysis of Average Survival vs External Factors

TIME SERIES ANALYSIS 36 Sub-Samples: Spearman Rank Correlation of Average Survival Time vs.							
		LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	Average CCY Volatility	0.02	0.03	0.03	0.06	0.05	0.06
	Average Interest Rate Differential	0.05	-0.03	-0.04	-0.02	-0.01	-0.02
	Average CCY Skewness	-0.01	-0.04	-0.03	-0.01	0.01	0.00
	Average CCY Kurtosis	-0.06	0.03	0.04	0.08	0.10	0.12
SR 2	Average CCY Volatility	0.03	0.03	0.03	0.05	0.05	0.04
	Average Interest Rate Differential	0.01	-0.06	-0.05	-0.03	-0.04	-0.02
	Average CCY Skewness	-0.03	-0.04	-0.02	0.00	0.01	0.00
	Average CCY Kurtosis	-0.01	0.06	0.07	0.11	0.13	0.14
SR 3	Average CCY Volatility	0.01	0.03	0.02	0.04	0.05	0.04
	Average Interest Rate Differential	0.00	-0.05	-0.02	-0.03	-0.03	-0.03
	Average CCY Skewness	-0.03	-0.04	0.00	0.01	0.01	0.00
	Average CCY Kurtosis	-0.03	0.07	0.09	0.10	0.13	0.16
SR 4	Average CCY Volatility	0.04	-0.01	0.01	0.06	0.05	0.03
	Average Interest Rate Differential	0.04	-0.05	-0.02	-0.04	-0.03	-0.03
	Average CCY Skewness	-0.03	-0.02	0.01	0.02	0.03	0.01
	Average CCY Kurtosis	-0.06	0.07	0.08	0.10	0.10	0.15
SR 5	Average CCY Volatility		0.01	0.01	0.05	0.04	0.03
	Average Interest Rate Differential		-0.05	-0.02	-0.03	-0.02	-0.03
	Average CCY Skewness		-0.03	0.00	0.02	0.03	0.02
	Average CCY Kurtosis		0.08	0.07	0.09	0.09	0.14
SR 10	Average CCY Volatility			-0.01	0.04	0.01	0.00
	Average Interest Rate Differential			0.00	-0.02	0.00	-0.02
	Average CCY Skewness			0.00	0.04	0.05	0.02
	Average CCY Kurtosis			0.08	0.06	0.08	0.10
SR 15	Average CCY Volatility				0.03	0.03	0.04
	Average Interest Rate Differential				0.00	0.01	-0.01
	Average CCY Skewness				0.04	0.04	0.01
	Average CCY Kurtosis				0.05	0.05	0.05
SR 20	Average CCY Volatility					0.03	0.06
	Average Interest Rate Differential					0.01	0.00
	Average CCY Skewness					0.04	0.04
	Average CCY Kurtosis					0.03	0.05
SR 25	Average CCY Volatility						0.06
	Average Interest Rate Differential						0.00
	Average CCY Skewness						0.04
	Average CCY Kurtosis						0.05

The Figure shows Spearman's Rank correlation between average survival time and average interest differential, standard deviation; skew and kurtosis are calculated for all currency pairs. The index of short term parameters of the moving average combinations are shown along the first column, the index of long term parameters is shown along the first row. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%.

E. APPENDIX 5: GARCH (1,1) Parameters for Resampling Simulation

1. GARCH (1,1) Parameters


Parameters of Garch estimation															
	USD			GBP			JPY			EUR			CHF		
	a0	a1	b	a0	a1	b	a0	a1	b	a0	a1	b	a0	a1	b
USD				4.9E-07	0.064	0.924	6.1E-07	0.053	0.936	8.8E-07	0.067	0.916	7.7E-07	0.060	0.928
p-value				0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
GBP	4.9E-07	0.063	0.925				8.1E-07	0.067	0.918	7.4E-07	0.079	0.893	9.1E-07	0.075	0.901
p-value	0.00	0.00	0.00				0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
JPY	6.3E-07	0.053	0.935	8.2E-07	0.067	0.918				7.0E-07	0.081	0.905	1.1E-06	0.095	0.885
p-value	0.00	0.00	0.00	0.00	0.00	0.00				0.00	0.00	0.00	0.00	0.00	0.00
EUR	8.9E-07	0.065	0.917	7.5E-07	0.079	0.893	7.0E-07	0.080	0.906				2.8E-07	0.088	0.896
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00				0.00	0.00	0.00
CHF	7.7E-07	0.059	0.929	9.1E-07	0.073	0.903	1.1E-06	0.096	0.885	2.5E-07	0.086	0.901			
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
NOK	6.9E-07	0.070	0.917	4.9E-07	0.067	0.919	4.8E-07	0.062	0.931	1.7E-07	0.100	0.903	4.9E-07	0.092	0.897
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SEK	5.3E-06	0.135	0.768	4.1E-06	0.131	0.769	1.5E-06	0.099	0.888	1.5E-06	0.205	0.791	3.6E-06	0.254	0.691
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
CAD	1.2E-07	0.094	0.903	1.1E-06	0.070	0.906	8.4E-07	0.062	0.925	1.7E-06	0.073	0.894	1.7E-06	0.067	0.907
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
AUD	4.1E-06	0.099	0.824	1.5E-06	0.134	0.858	4.1E-06	0.081	0.869	3.9E-06	0.085	0.861	5.7E-06	0.090	0.841
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
NZD	1.1E-06	0.181	0.854	2.1E-06	0.234	0.796	7.5E-06	0.108	0.814	7.0E-06	0.132	0.788	6.4E-06	0.114	0.824
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

	NOK			SEK			CAD			AUD			NZD		
	a0	a1	b	a0	a1	b	a0	a1	b	a0	a1	b	a0	a1	b
USD	7.0E-07	0.070	0.917	7.3E-06	0.159	0.715	1.2E-07	0.093	0.904	5.5E-06	0.094	0.807	1.1E-06	0.223	0.840
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
GBP	4.7E-07	0.066	0.921	5.2E-06	0.150	0.732	1.1E-06	0.069	0.908	1.3E-06	0.138	0.862	2.2E-06	0.285	0.775
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
JPY	4.6E-07	0.061	0.933	2.3E-06	0.128	0.853	8.7E-07	0.062	0.925	5.4E-06	0.081	0.857	9.8E-06	0.112	0.791
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
EUR	1.6E-07	0.099	0.906	1.7E-06	0.241	0.777	1.7E-06	0.072	0.895	4.7E-06	0.081	0.856	8.7E-06	0.139	0.768
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
CHF	4.6E-07	0.089	0.901	4.3E-06	0.306	0.651	1.7E-06	0.066	0.908	7.1E-06	0.085	0.833	7.8E-06	0.115	0.813
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
NOK				2.4E-06	0.250	0.723	1.1E-06	0.071	0.909	7.6E-06	0.110	0.785	1.2E-05	0.117	0.740
p-value				0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SEK	1.8E-06	0.223	0.756				1.7E-05	0.171	0.486	1.4E-05	0.100	0.707	2.5E-05	0.125	0.567
p-value	0.00	0.00	0.00				0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
CAD	1.1E-06	0.070	0.910	2.7E-05	0.204	0.271				4.3E-06	0.311	0.668	6.8E-06	0.143	0.776
p-value	0.00	0.00	0.00	0.00	0.00	0.00				0.00	0.00	0.00	0.00	0.00	0.00
AUD	6.4E-06	0.113	0.795	1.6E-05	0.118	0.657	4.2E-06	0.263	0.693				2.2E-06	0.373	0.726
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00				0.00	0.00	0.00
NZD	9.7E-06	0.120	0.756	3.8E-05	0.170	0.350	5.3E-06	0.133	0.797	2.4E-06	0.389	0.714			
p-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			

The figure shows the GARCH(1,1) parameters for individual currency pairs based on historic time series data from the 4th of January of 2074 to 31st of June of 2010. P-values are given to assess the statistical significance.


F. APPENDIX 6: Trading Rule Results across Base Currencies

1. Median Log Rank Values; Positive Signals; Simple Resampling




The dark bars in the Figure show the median level of log-rank test results across all moving average signal combinations. The light bars in the Figure show the difference between short-term and long-term moving average combination

2. Median Log Rank Values; Negative Signals; Simple Resampling




The dark bars in the Figure show the median level of log-rank test results across all moving average signal combinations. The light bars in the Figure show the difference between short-term and long-term moving average combination

3. Median Log Rank Values; Positive Signals; GARCH (1,1) Resampling




The dark bars in the Figure show the median level of log-rank test results across all moving average signal combinations. The light bars in the Figure show the difference between short-term and long-term moving average combination

4. Median Log Rank Values; Negative Signals; GARCH (1,1) Resampling




The dark bars in the Figure show the median level of log-rank test results across all moving average signal combinations. The light bars in the Figure show the difference between short-term and long-term moving average combination

5. Breakeven Transaction Cost; Enhanced vs. Benchmark Trading Rules




Breakeven transaction cost levels for the benchmark strategy are shown in dark grey, breakeven transaction cost levels for the enhanced strategy are shown in light grey. The dotted line represents the difference between the two

6. Abs. log-rank test Results vs. Benchmark Trading Rule Results



The dark grey line shows the breakeven transaction cost levels of the benchmark trading rule. The light grey line shows the absolute level of positive and negative log-rank test results.

7. Abs. log-rank test Results vs. Enhanced Trading Rule Results



The dark grey line shows the breakeven transaction cost levels of the benchmark trading rule. The light grey line shows the absolute level of positive and negative log-rank test results.

X. APPENDICES: Chapter 3

A. APPENDIX 1: OLS Regression on Systematic Risk Factors

1. Output of the SR1/LR5 Regression

CCY	CONSTANT	TREND	MOMENTUM	CARRY	VALUE	RISKAV	VOLA	R2	Sig(White)	Sig(Breusch-Godfrey, 10 order)
'1/2	0.32% ***	4.22% ***	0.33%	0.04%	1.11%	4.15% ***	-5.84% ***	0.78%	***	***
'1/3	0.36% ***	4.78% ***	-0.23%	-1.03%	-1.27%	7.42% ***	-7.62% ***	1.54%	***	***
'1/4	0.35% ***	5.72% ***	-0.20%	-2.19% *	-1.14%	2.79% ***	-5.10% ***	0.53%	***	***
'1/5	0.40% ***	4.21% **	-1.26%	-1.93%	-1.63%	1.19%	-7.41% ***	0.37%	***	***
'1/6	0.36% ***	-0.07%	-0.67%	-0.27%	-3.45% ***	4.73% ***	-6.62% ***	0.74%	***	***
'1/7	0.35% ***	-1.09%	4.83% ***	-0.91%	-1.41%	3.08% ***	-5.10% ***	0.57%	***	***
'1/8	0.21% ***	-1.33%	0.94%	0.26%	-0.21%	2.33% ***	1.57% **	0.32%	***	***
'1/9	0.33% ***	-4.27% **	0.18%	-6.34% ***	0.88%	6.05% ***	-1.82%	1.29%	***	***
'1/10	0.35% ***	-1.95%	3.89% ***	0.22%	8.57% ***	6.11% ***	1.51%	1.39%	***	***
'2/3	0.38% ***	1.08%	1.28%	2.86% **	2.70% **	10.62% ***	0.03%	2.03%	***	***
'2/4	0.26% ***	6.25% ***	-0.31%	0.27%	2.76% ***	3.27% ***	2.26% **	1.02%	***	***
'2/5	0.30% ***	3.70% ***	-1.51% *	-0.20%	0.96%	4.65% ***	-0.93%	0.67%	***	***
'2/6	0.27% ***	2.29% *	1.09%	-1.17%	-0.29%	3.88% ***	1.00%	0.62%	***	***
'2/7	0.28% ***	1.11%	5.98% ***	-0.28%	0.58%	3.19% ***	1.33%	0.90%	***	***
'2/8	0.34% ***	2.14%	0.71%	-0.19%	2.61% **	3.16% ***	-2.11% *	0.38%	***	***
'2/9	0.38% ***	1.69%	1.04%	-5.33% ***	0.94%	5.98% ***	-3.89% ***	1.08%	***	***
'2/10	0.38% ***	-0.68%	4.46% ***	-1.69%	8.76% ***	4.43% ***	1.79%	1.31%	***	***
'3/4	0.35% ***	3.10% **	-0.97%	-0.39%	-2.13% **	8.69% ***	0.65%	1.65%	***	***
'3/5	0.36% ***	1.07%	-1.70% *	0.52%	-1.56%	7.19% ***	-1.59%	0.98%	***	***
'3/6	0.37% ***	-0.16%	-0.24%	1.42%	-3.63% ***	10.01% ***	-1.53%	1.64%	***	***
'3/7	0.37% ***	-1.31%	6.56% ***	0.02%	-2.78% **	9.49% ***	2.29%	1.83%	***	***
'3/8	0.41% ***	3.60% **	1.98% *	-1.15%	-2.13% *	9.22% ***	-3.25% **	1.49%	***	***
'3/9	0.46% ***	-1.14%	-0.45%	-5.96% ***	-1.26%	11.16% ***	-1.97%	1.76%	***	***
'3/10	0.46% ***	-1.49%	4.23% ***	0.03%	7.93% ***	8.51% ***	-0.38%	1.35%	***	***
'4/5	0.17% ***	3.03% ***	-1.14% **	0.19%	-0.37%	1.59% ***	3.14% ***	0.53%	***	***
'4/6	0.20% ***	2.81% ***	-0.42%	-2.21% ***	-1.93% ***	2.18% ***	2.75% ***	0.69%	***	***
'4/7	0.21% ***	0.00%	5.66% ***	-0.88%	-1.03%	1.31% **	1.80% *	0.88%	***	***
'4/8	0.37% ***	4.38% ***	-0.15%	-1.30%	-1.38%	2.22% ***	-2.13% *	0.20%	***	***
'4/9	0.40% ***	4.57% **	-0.94%	-7.37% ***	-0.53%	5.16% ***	-2.13%	0.98%	***	***
'4/10	0.41% ***	2.65%	2.08% *	-4.18% ***	8.04% ***	3.54% ***	1.75%	1.13%	***	***
'5/6	0.26% ***	-0.18%	-1.41% *	-1.37%	-2.62% ***	3.92% ***	0.06%	0.56%	***	***
'5/7	0.28% ***	-1.77%	3.98% ***	-1.05%	-1.72% *	2.47% ***	0.17%	0.42%	***	***
'5/8	0.42% ***	2.68%	-0.98%	-1.11%	-2.59% **	2.21% **	-3.87% ***	0.16%	***	***
'5/9	0.45% ***	1.61%	-1.75%	-7.12% ***	-1.21%	4.86% ***	-3.20% *	0.65%	***	***
'5/10	0.46% ***	0.16%	1.71%	-3.40% **	6.28% ***	4.86% ***	-2.03%	0.78%	***	***
'6/7	0.19% ***	-2.46% **	4.86% ***	-2.47% ***	-1.67% **	2.41% ***	1.22%	1.10%	***	***
'6/8	0.36% ***	-0.70%	1.35%	-2.10% *	-3.02% ***	2.98% ***	-1.70%	0.29%	***	***
'6/9	0.40% ***	0.06%	-0.13%	-10.08% ***	-1.88%	4.49% ***	-2.16%	1.18%	***	***
'6/10	0.41% ***	-1.36%	5.78% ***	-5.01% ***	5.77% ***	2.71% ***	2.23%	1.12%	***	***
'7/8	0.36% ***	-1.17%	7.13% ***	-1.00%	-2.20% *	1.96% **	-1.72%	0.67%	***	***
'7/9	0.40% ***	-2.10%	8.39% ***	-8.45% ***	-2.07%	3.73% ***	-0.99%	1.56%	***	***
'7/10	0.40% ***	-2.14%	11.55% ***	-3.50% **	6.32% ***	2.19% **	2.54%	1.68%	***	***
'8/9	0.33% ***	0.51%	0.73%	-8.19% ***	2.39% **	4.11% ***	1.03%	1.52%	***	***
'8/10	0.36% ***	0.16%	5.03% ***	-4.84% ***	9.32% ***	1.75% **	4.05% ***	1.74%	***	***
'9/10	0.28% ***	-3.24% **	2.38% **	-11.77% ***	4.82% ***	0.53%	-1.06%	1.94%	***	***

The figure shows the results from an OLS regression. The first column shows the various currency pairs, whereby 1=USD, 2=GBP, 3=JPY, 4=EUR, 5=CHF, 6=NOK, 7=SEK, 8=CAD, 9=AUD, 10=NZD. The next seven columns show the OLS estimates (including constant) with their respective statistical significances. Three stars indicate a significance level of 1%, two stars of 5% and one star of 10%. In the last three columns the R^2 and the significance levels of the White test for heteroskedasticity and the Breusch-Godfrey test for autocorrelation are shown.

B. APPENDIX 2: Test for Cross-Sectional Heteroskedasticity

1. Test results for equality of variance across currency pairs, ordered by trading rule parameterisations

	Bartlett TEST					
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	***	***	***	***	***	***
SR 2	***	***	***	***	***	***
SR 3	***	***	***	***	***	***
SR 4	***	***	***	***	***	***
SR 5		***	***	***	***	***
SR 10			***	***	***	***
SR 15				***	***	***
SR 20					***	***
SR 25						***


	Levene TEST					
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	***	***	***	***	***	***
SR 2	***	***	***	***	***	***
SR 3	***	***	***	***	***	***
SR 4	***	***	***	***	***	***
SR 5		***	***	***	***	***
SR 10			***	***	***	***
SR 15				***	***	***
SR 20					***	***
SR 25						***

	Brown-Forsythe TEST					
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	***	***	***	***	***	***
SR 2	***	***	***	***	***	***
SR 3	***	***	***	***	***	***
SR 4	***	***	***	***	***	***
SR 5		***	***	***	***	***
SR 10			***	***	***	***
SR 15				***	***	***
SR 20					***	***
SR 25						***

The Figure shows the results of the significance levels of equality of variance test for each of the trading rule parameterisations across all currencies. The first part of the Figure shows the test results of the Bartlett test, the second and third part of the figure shows the test results of the Bartlett, Levene and Brown-Forsythe test. “*” indicates statistically significant cross sectional heteroskedasticity at the 10% level, “**” indicates statistically significant cross sectional heteroskedasticity at the 5% level, “***” indicates statistically significant cross sectional heteroskedasticity at the 1% level.

C. APPENDIX 3: Correlation of OLS Residuals

1. Correlation matrix of the residuals of the SR1/LR5 OLS regression, given in Appendix 1



The Figure shows the correlation matrix of the residuals of the SR1/LR5 regression analysis given in Appendix 1. The darkest blue shade represents a correlation of less than 0.95, the colour shades then change incrementally, as indicated in the legend. The darkest red shade indicates a correlation level of 0.95 and more.

D. APPENDIX 4: Test of Endogeneity of Systematic Factors

1. Median P-value of the Durbin-Wu-Hausman test across all analysed currency pairs

Coefficient (TREND)							Coefficient (VALUE)						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30		LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	45.84%	48.61%	47.81%	37.24%	25.2%	19.95%	SR 1	9.23%*	9.88%*	13.5%	16.08%	17.19%	17.33%
SR 2	48.66%	44.77%	41.42%	30.15%	25.29%	20.98%	SR 2	10.7%	13.76%	17.22%	15.41%	14.83%	15.23%
SR 3	42.31%	36.32%	38.81%	29.75%	25.99%	24.39%	SR 3	11.92%	16.94%	16.71%	18.75%	19.57%	21.34%
SR 4	33.3%	35.27%	33.01%	27.98%	25.62%	26.58%	SR 4	17.6%	18.32%	19.08%	24.04%	22.39%	24.72%
SR 5		32.73%	29.52%	31.04%	27.26%	24.16%	SR 5		17.43%	23.63%	26.43%	26.89%	28.1%
SR 10			38.16%	26.48%	21.93%	20.97%	SR 10			40.92%	35.26%	31.48%	33.46%
SR 15				22.7%	17.84%	19.45%	SR 15				34.43%	34.7%	35.15%
SR 20					15.31%	16.14%	SR 20					38.47%	39.54%
SR 25						14.81%	SR 25						42.75%

Coefficient (MOMENTUM)							Coefficient (RISKAV)						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30		LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	15.43%	10.83%	11.52%	10.95%	10.48%	10.21%	SR 1	55.33%	52.63%	54.47%	54.75%	52.29%	56.84%
SR 2	25.97%	15.28%	14.93%	8.65%*	9.7%*	8.41%*	SR 2	48.41%	49.81%	54.39%	55.15%	58.04%	59.51%
SR 3	24.14%	14.81%	12.11%	9.46%*	10.42%	9.02%*	SR 3	49.19%	53.48%	62.04%	56.74%	57.84%	59.11%
SR 4	22.2%	13.82%	15.71%	12.16%	9.96%*	9.39%*	SR 4	53.25%	52.97%	61.29%	57.07%	59.65%	58.35%
SR 5		17.63%	15.12%	12.49%	11.6%	10.05%	SR 5		56.95%	58.88%	57.06%	55.16%	55.72%
SR 10			17.58%	15.26%	13.22%	12.55%	SR 10			50.4%	52.79%	57.98%	56.3%
SR 15				10.27%	13.34%	12.4%	SR 15				61.31%	54.69%	52.35%
SR 20					15.95%	18.63%	SR 20					49.28%	51.23%
SR 25						23.12%	SR 25						46.16%

Coefficient (CARRY)							Coefficient (VOLA)						
	LR 5	LR 10	LR 15	LR 20	LR 25	LR 30		LR 5	LR 10	LR 15	LR 20	LR 25	LR 30
SR 1	37.32%	39.25%	29.7%	22.44%	16.8%	15.38%	SR 1	10.22%	8.04%*	8.9%*	10.06%	10.95%	10.38%
SR 2	40.66%	34.13%	37.43%	28.06%	26.22%	22.71%	SR 2	11.66%	10.85%	11.18%	8.72%*	8.66%*	9.99%*
SR 3	32.22%	37.68%	36.23%	31.69%	27.49%	25.67%	SR 3	11.58%	10.16%	10.91%	10.55%	12.39%	12.26%
SR 4	28.13%	37.34%	35.95%	28.87%	28.75%	24.58%	SR 4	16.85%	11.51%	12.25%	14.55%	13.35%	14.92%
SR 5		37.84%	32.07%	31.14%	28.01%	23.66%	SR 5		11.35%	16.83%	16.27%	16.28%	14.26%
SR 10			38.87%	35.33%	31.81%	26.2%	SR 10			27.48%	20.5%	18.4%	22.11%
SR 15				32.68%	27.38%	22.46%	SR 15				17.3%	20.7%	20.87%
SR 20					23.6%	21.4%	SR 20					21.5%	24.81%
SR 25						19.6%	SR 25						34.69%

The Figure shows the median P-values of the Durbin-Wu-Hausman test across all currency pairs for each given trading rule parameterisation. "*" indicates statistically significant endogeneity at the 10% level, "***" indicates statistically significant endogeneity at the 5% level, "****" indicates statistically significant endogeneity at the 1% level.